Intrinsic Robustness of the Price of Anarchy*

Tim Roughgarden†

July 3, 2013

Abstract

The price of anarchy, defined as the ratio of the worst-case objective function value of a Nash equilibrium of a game and that of an optimal outcome, quantifies the inefficiency of selfish behavior. Remarkably good bounds on this measure are known for a wide range of application domains. However, such bounds are meaningful only if a game's participants successfully reach a Nash equilibrium. This drawback motivates inefficiency bounds that apply more generally to weaker notions of equilibria, such as mixed Nash equilibria and correlated equilibria, and to sequences of outcomes generated by natural experimentation strategies, such as successive best responses and simultaneous regret-minimization.

We establish a general and fundamental connection between the price of anarchy and its seemingly more general relatives. First, we identify a "canonical sufficient condition" for an upper bound on the price of anarchy of pure Nash equilibria, which we call a *smoothness argument*. Second, we prove an "extension theorem": every bound on the price of anarchy that is derived via a smoothness argument *extends automatically*, with no quantitative degradation in the bound, to mixed Nash equilibria, correlated equilibria, and the average objective function value of every outcome sequence generated by no-regret learners. Smoothness arguments also have automatic implications for the inefficiency of approximate equilibria, for bicriteria bounds, and, under additional assumptions, for polynomial-length best-response sequences. Third, we prove that in congestion games, smoothness arguments are "complete" in a proof-theoretic sense: despite their automatic generality, they are guaranteed to produce optimal worst-case upper bounds on the price of anarchy.

*A preliminary version of this paper appeared in the Proceedings of the 41st Annual Symposium on Theory of Computing, May-June 2009. A condensed overview appeared in *Communications of the ACM*, July 2012.

†Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. This research was supported in part by NSF grants CCF-0448664 and CCF-1016885, an AFOSR MURI grant, an ONR Young Investigator Award, an ONR PECASE Award, and an Alfred P. Sloan Fellowship. Email: tim@cs.stanford.edu.

1 Introduction

Self-interested behavior by autonomous decision-makers generally leads to an inefficient result — an outcome that could be improved upon given dictatorial control over everyone's actions. Imposing such control can be costly or infeasible in many systems, with large networks furnishing obvious examples. This fact motivates the search for conditions under which decentralized optimization by competing individuals is guaranteed to produce a near-optimal outcome.

A rigorous guarantee of this type requires a formal behavioral model, to define "the outcome of self-interested behavior." The majority of previous research studies *pure-strategy Nash equilibria*, defined as follows. Each player i selects a strategy s_i from a set S_i, like a path in a network. The cost $C_i(\mathbf{s})$ incurred by a player i in a game is a function of the entire vector $\mathbf{s}$ of players' chosen strategies, which is called a *strategy profile* or *outcome*. By definition, a strategy profile $\mathbf{s}$ of a game is a *pure Nash equilibrium* if no player can decrease its cost via a unilateral deviation:

$$C_i(\mathbf{s}) \leq C_i(s_i', \mathbf{s}_{-i}) \tag{1}$$

for every i and $s_i' \in S_i$, where $\mathbf{s}_{-i}$ denotes the strategies chosen by the players other than i in $\mathbf{s}$. These concepts can be defined equally well via payoff-maximization rather than cost-minimization; see also Examples 2.6 and 2.7.

The *price of anarchy (POA)* [74] measures the suboptimality caused by self-interested behavior. Given a game, a notion of an "equilibrium" (such as pure Nash equilibria), and a nonnegative objective function (such as the sum of players' costs), the POA of the game is defined as the ratio between the largest cost of an equilibrium and the cost of an optimal outcome. An upper bound on the POA has an attractive worst-case flavor: it applies to every equilibrium and obviates the need to predict a single outcome of selfish behavior. Many researchers have proved remarkably good bounds on the POA in a wide range of models; see the surveys in Nisan et al. [87, Chapters 17–21] and the references therein.

1.1 The Need for More Robust Bounds

A good bound on the price of anarchy of a game is not enough to conclude that self-interested behavior is relatively benign. Such a bound is meaningful only if a game's participants successfully reach an equilibrium. For pure Nash equilibria, however, there are a number of reasons why this might not occur: the players might fail to coordinate on one of multiple equilibria; they might be playing a game in which computing a pure Nash equilibrium is a computationally intractable problem [47], or, even more fundamentally, a game in which pure Nash equilibria do not exist. These critiques motivate worst-case performance bounds that apply to as wide a range of outcomes as possible, and under minimal assumptions about how players play and coordinate in a game.

This paper presents a general theory of "robust" bounds on the price of anarchy, meaning bounds that apply to equilibrium concepts that are much more permissive than pure Nash equilibria, including those shown in Figure 1. We formally define these concepts — mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria — in Section 3.1.

Enlarging the set of equilibria weakens the behavioral and technical assumptions necessary to justify equilibrium analysis. First, while there are games with no pure Nash equilibria — "Rock-Paper-Scissors" being a simple example — every (finite) game has at least one mixed Nash equilibrium [86]. As a result, the "non-existence critique" for pure Nash equilibria does not apply to any of the more general equilibrium concepts in Figure 1. Second, while computing a mixed

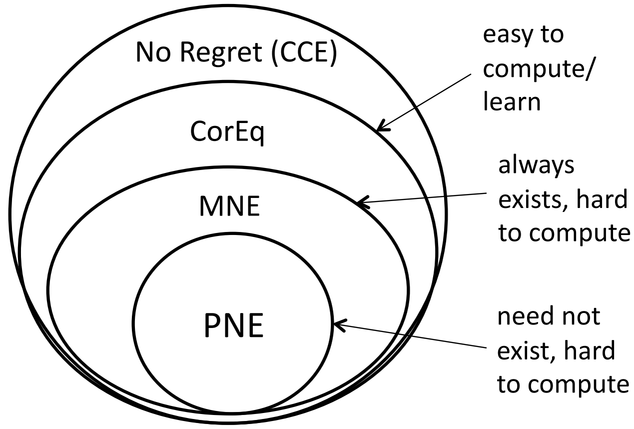

Figure 1: Generalizations of pure Nash equilibria. "PNE" stands for pure Nash equilibria; "MNE" for mixed Nash equilibria; "CorEq" for correlated equilibria; and "No Regret (CCE)" for coarse correlated equilibria, which correspond to the empirical distributions that arise from repeated joint play by no-regret learners. See Section 3.1 for formal definitions.

Nash equilibrium is a computationally intractable problem in general [30, 42, 45], computing a correlated equilibrium is not [58]. Thus, the "intractability critique" for pure and mixed Nash equilibria does not apply to the two largest sets of Figure 1. More importantly, these two sets are "easily learnable": when a game is played repeatedly over time, there are natural classes of learning dynamics — processes by which each player chooses its strategy for the next time step, as a function only of its own past play and payoffs — such that the empirical distribution of joint play converges to these sets (see Blum and Mansour [21]).

1.2 Overview of Results

The primary goal of this paper is the formulation and proof of the following general result:

> *In many fundamental game-theoretic models, worst-case bounds on the POA apply even when players have not converged to a (Nash) equilibrium.*

Our contributions can be divided into three parts. First, we identify a sufficient condition for an upper bound on the POA of pure Nash equilibria of a game for the welfare objective function. This condition encodes a canonical proof template for deriving such bounds. We call such proofs "smoothness arguments." Many of the POA upper bounds in the literature can be recast as instantiations of this canonical method.

Second, we prove an "extension theorem": every bound on the price of anarchy that is derived via a smoothness argument extends automatically, with no quantitative degradation in the bound, to all of the more general equilibrium concepts pictured in Figure 1. We also show that smoothness arguments have automatic implications for the inefficiency of approximate equilibria, for bicriteria bounds, and, under additional assumptions, for polynomial-length best-response sequences.

Third, we prove that congestion games, with cost functions restricted to some arbitrary set, are "tight" in the following sense: smoothness arguments, despite their automatic generality, are guaranteed to produce optimal worst-case upper bounds on the POA, even for the set of pure Nash equilibria. POA bounds for these classes of games are "intrinsically robust," in that the worst-case POA is the same for each of the equilibrium concepts of Figure 1. This result also provides an understanding of the worst-case POA of congestion games that is as complete as that given for nonatomic congestion games by Roughgarden and Tardos [99] and Correa et al. [40], in the form of tight bounds and a characterization of worst-case examples for all classes of cost functions.

1.3 Organization of Paper

Section 2 provides formal and intuitive definitions of smooth games, along with several examples and non-examples. Section 3 states and proves the main extension theorem, that every smoothness argument automatically applies to all of the equilibrium concepts shown in Figure 1. Section 4 derives consequences of smoothness arguments for approximate equilibria, bicriteria bounds, and best-response sequences. Section 5 proves that smoothness arguments always give optimal POA bounds in congestion games. Section 6 describes related work, including advances following the conference version of this paper [95]. Section 7 concludes.

2 Smooth Games

Section 2.1 formally defines smooth games. Section 2.2 interprets this definition in terms of POA bounds that make minimal use of the Nash equilibrium hypothesis. Section 2.3 shows how three well-known POA bounds from disparate problem domains can be interpreted as smoothness arguments. Section 2.4 defines tight classes of games, in which smoothness arguments yield optimal POA bounds. Section 2.5 explains why not all POA bounds are equivalent to smoothness arguments.

2.1 Definitions

By a *cost-minimization game*, we mean a game — players, strategies, and cost functions — together with the joint cost objective function $C(\mathbf{s}) = \sum_{i=1}^{k} C_i(\mathbf{s})$. Essentially, a "smooth game" is a cost-minimization game that admits a POA bound of a canonical type — a *smoothness argument*. We give the formal definition and then explain how to interpret it.

Definition 2.1 (Smooth Game) A cost-minimization game is (λ, μ)-*smooth* if for every two outcomes $\mathbf{s}$ and $\mathbf{s}^*$,

$$\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}). \tag{2}$$

There is an analogous definition of smooth games for maximization objectives; see Examples 2.6 and 2.7. Roughly, smoothness controls the cost of a set of "one-dimensional perturbations" of an outcome, as a function of both the initial outcome $\mathbf{s}$ and the perturbations $\mathbf{s}^*$.

We claim that if a game is (λ, μ)-smooth, with $\lambda > 0$ and $\mu < 1$, then each of its pure Nash equilibria $\mathbf{s}$ has cost at most $\lambda/(1-\mu)$ times that of an optimal solution $\mathbf{s}^*$. In proof, we derive

$$
\begin{aligned}
C(\mathbf{s}) &= \sum_{i=1}^{k} C_i(\mathbf{s}) && (3)\\
&\leq \sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) && (4)\\
&\leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}), && (5)
\end{aligned}
$$

where (3) follows from the definition of the objective function; inequality (4) follows from the Nash equilibrium condition (1), applied once to each player i with the hypothetical deviation s_i^*; and inequality (5) follows from the defining condition (2) of a smooth game. Rearranging terms yields the claimed bound.

Definition 2.1 is sufficient for the last line of this three-line proof (3)–(5), but it insists on more than what is needed: it demands that the inequality (2) holds for *every* outcome $\mathbf{s}$, and *not only for Nash equilibria*. This is the basic reason why smoothness arguments imply worst-case bounds beyond the set of pure Nash equilibria.

We define the *robust POA* as the best upper bound on the POA that is provable via a smoothness argument.

Definition 2.2 (Robust POA) The *robust price of anarchy* of a cost-minimization game is

$$
\inf \left\{ \frac{\lambda}{1-\mu} \; : \; (\lambda, \mu) \text{ such that the game is } (\lambda, \mu)\text{-smooth} \right\},
$$

with μ always constrained to be less than 1.

Remark 2.3 (Relaxations of Smoothness) There are two ways to weaken Definition 2.1 that preserve all of the consequences proved in this paper. First, the assumption that the objective function satisfies $C(\mathbf{s}) = \sum_{i=1}^{k} C_i(\mathbf{s})$ can be replaced by the inequality $C(\mathbf{s}) \leq \sum_{i=1}^{k} C_i(\mathbf{s})$; we exploit this fact in Examples 2.6 and 2.7 below. Second, in Definition 2.1, the inequality (2) only needs to hold for some optimal solution $\mathbf{s}^*$ and all outcomes $\mathbf{s}$, rather than for all pairs $\mathbf{s}, \mathbf{s}^*$ of outcomes. See Example 2.7 and Section 6 for applications of this relaxation.

Remark 2.4 ((Non-)Existence of Pure Nash Equilibria) Games can be smooth with nontrivial values of λ and μ despite possessing no pure Nash equilibria. Examples of such games include valid utility games [105] (see Example 2.6) and weighted versions of the congestion games studied in Section 5 [2, 11, 14, 32, 60, 88]. The derivation in (3)–(5) proves that *if* a (λ, μ)-smooth game has at least one pure Nash equilibrium, then its POA for such equilibria is at most $\lambda/(1-\mu)$. (We leave the POA undefined if no equilibria exist.)

Our smoothness framework provides an explanation for the arguably mystifying fact that meaningful POA bounds for valid utility and weighted congestion games do not seem to require a universal equilibrium existence result. All of the known upper bounds on the POA of pure Nash equilibria in these games follow from smoothness arguments. As such, these POA bounds are not fundamentally about pure Nash equilibria, but rather the more permissive equilibrium concepts shown in Figure 1, for which existence is guaranteed. Of course, when some of the mixed Nash equilibria happen to be pure, such a POA bound applies to them as a special case.

2.2 Intuition

Smoothness arguments are a class of upper bound proofs for the POA of pure Nash equilibria that are confined to use the equilibrium hypothesis in a minimal way. To explain, recall the canonical three-line proof (3)–(5). The first inequality (4) uses the Nash equilibrium hypothesis, but only to justify why each player i selects its equilibrium strategy s_i rather than its strategy s_i^* in the optimal outcome. If we care only about the POA of pure Nash equilibria, then we are free to establish an upper bound using any argument that we please. For example, such an argument could invoke the Nash equilibrium hypothesis again to generate further inequalities of the form $C_i(\mathbf{s}) \leq C_i(\hat{s}_i, \mathbf{s}_{-i})$, with the hypothetical deviations $\hat{s}_i$ chosen as a function of the particular Nash equilibrium $\mathbf{s}$. Using a smoothness argument — that is, proving inequality (5) for *all* outcomes $\mathbf{s}$ — is tantamount to discarding the Nash equilibrium hypothesis after deriving the first inequality (4) using only the hypothetical deviations suggested by the optimal outcome $\mathbf{s}^*$.

2.3 Examples

Concern about the range of applicability of a definition grows as its interesting consequences accumulate. To alleviate such fears and add some concreteness to the discussion, we next show how three well-known POA analyses can be recast as smoothness arguments; more are discussed in Sections 5 and 6.

The first example is congestion games with affine cost functions. The POA in these games was first studied by Awerbuch et al. [11] and Christodoulou and Koutsoupias [33]. Section 5 treats congestion games with general cost functions in detail. The second example concerns Vetta's well-studied utility games [105]. This example illustrates how smoothness arguments work in payoff-maximization games, and also with a "one-sided" variant of sum objective functions (cf., Remark 2.3). The third example recasts as a smoothness argument the analysis in Christodoulou et al. [36] of simultaneous second-price auctions, exploiting the second relaxation mentioned in Remark 2.3.

2.3.1 Cost-Minimization Games

Example 2.5 (Congestion Games with Affine Cost Functions [11, 33]) A congestion game is a cost-minimization game defined by a ground set E of resources, a set of k players with strategy sets $S_1, \ldots, S_k \subseteq 2^E$, and a cost function $c_e : \mathcal{Z}^+ \to \mathcal{R}$ for each resource $e \in E$. Congestion games were defined by Rosenthal [91]. In this paper, we always assume that cost functions are nonnegative and nondecreasing. For this example, we make the much stronger assumption that every cost function is *affine*, meaning that $c_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$ for every resource $e \in E$. A canonical example is routing games, where E is the edge set of a network, and the strategies of a player correspond to paths between its source and sink vertices. Given a strategy profile $\mathbf{s} = (s_1, \ldots, s_k)$, with $s_i \in S_i$ for each i, we say that $x_e = |\{i : e \in s_i\}|$ is the *load* induced on e by $\mathbf{s}$, defined as the number of players that use it in $\mathbf{s}$. The cost to player i is defined as $C_i(\mathbf{s}) = \sum_{e \in s_i} c_e(x_e)$, where $\mathbf{x}$ is the vector of loads induced by $\mathbf{s}$. A reversal of sums shows that $C(\mathbf{s}) = \sum_{i=1}^k C_i(\mathbf{s}) = \sum_{e \in E} c_e(x_e) x_e$.

We claim that every congestion game with affine cost functions is $(\frac{5}{3}, \frac{1}{3})$-smooth, and hence has robust POA at most $\frac{5}{2}$. The basic reason for this was identified by Christodoulou and Koutsoupias [33, Lemma 1], who noted that

$$y(z+1) \leq \tfrac{5}{3}y^2 + \tfrac{1}{3}z^2$$

for all nonnegative integers y, z.[1] Thus, for all $a, b \geq 0$ and nonnegative integers y, z,

$$ay(z+1) + by \leq \tfrac{5}{3}\left(ay^2 + by\right) + \tfrac{1}{3}\left(az^2 + bz\right). \tag{6}$$

To establish smoothness, consider a pair $\mathbf{s}, \mathbf{s}^*$ of outcomes of a congestion game with affine cost functions, with induced loads $\mathbf{x}, \mathbf{x}^*$. Since the number of players using resource e in the outcome $(s_i^*, \mathbf{s}_{-i})$ is at most one more than that in $\mathbf{s}$, and this resource contributes to precisely x_e^* terms of the form $C_i(s_i^*, \mathbf{s}_{-i})$, we have

$$\begin{aligned} \sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) &\leq \sum_{e \in E} (a_e(x_e + 1) + b_e) x_e^* \\ &\leq \sum_{e \in E} \frac{5}{3}(a_e x_e^* + b_e) x_e^* + \sum_{e \in E} \frac{1}{3}(a_e x_e + b_e) x_e \qquad (7) \\ &= \frac{5}{3} C(\mathbf{s}^*) + \frac{1}{3} C(\mathbf{s}), \end{aligned}$$

where (7) follows from (6), with x_e^* and x_e playing the roles of y and z, respectively. The canonical three-line argument (3)–(5) then implies an upper bound of 5/2 on the POA of pure Nash equilibria in every congestion game with affine cost functions. This fact was first proved independently in [11] and [33], along with matching worst-case lower bounds. Our extension theorem (Theorem 3.2) implies that the bound of 5/2 extends to the other three sets of outcomes shown in Figure 1. These extensions were originally established in two different papers [20, 32] subsequent to the original POA bound [11, 33].

2.3.2 Payoff-Maximization Games

The next two examples are naturally phrased as *payoff-maximization* games, where each player has a payoff function $\Pi_i(\mathbf{s})$ that it strives to maximize. We use W to denote the objective function of a payoff-maximization game. We call such a game (λ, μ)-*smooth* if

$$\sum_{i=1}^{k} \Pi_i(s_i^*, \mathbf{s}_{-i}) \geq \lambda \cdot W(\mathbf{s}^*) - \mu \cdot W(\mathbf{s})$$

for every pair $\mathbf{s}, \mathbf{s}^*$ of outcomes. A derivation similar to (3)–(5) shows that, in a (λ, μ)-smooth payoff-maximization game, the objective function value of every pure Nash equilibrium is at least a $\lambda/(1+\mu)$ fraction of the maximum possible. We define the *robust POA* of a payoff-maximization game as the supremum of $\lambda/(1+\mu)$ over all legitimate smoothness parameters (λ, μ).

Example 2.6 (Valid Utility Games [105]) Our second example concerns *valid utility games* [105]. Such a game is defined by a ground set E, a nonnegative submodular function V defined on subsets of E, and a strategy set $S_i \subseteq 2^E$ and a payoff function Π_i for each player $i = 1, 2, \ldots, k$.[2] For example, the set E could denote a set of locations where facilities can be built, and a strategy $s_i \subseteq E$ could denote the locations at which player i chooses to build facilities. For an outcome $\mathbf{s}$, let $U(\mathbf{s}) \subseteq E$ denote the union $\cup_{i=1}^{k} s_i$ of players' strategies in $\mathbf{s}$. The objective function value of an

[1] The statement of this lemma in [32, 33] contains a typo, but it is applied correctly in both works.

[2] A set function $V : 2^E \to \mathcal{R}$ is *submodular* if $V(X \cap Y) + V(X \cup Y) \leq V(X) + V(Y)$ for every $X, Y \subseteq E$.

outcome $\mathbf{s}$ is defined as $W(\mathbf{s}) = V(U(\mathbf{s}))$. Furthermore, two conditions hold, by definition, in a valid utility game: (i) for each player i, $\Pi_i(\mathbf{s}) \geq W(\mathbf{s}) - W(\emptyset, \mathbf{s}_{-i})$ for every outcome $\mathbf{s}$; and (ii) $\sum_{i=1}^{k} \Pi_i(\mathbf{s}) \leq W(\mathbf{s})$ for every outcome $\mathbf{s}$. The second condition is precisely the relaxation of the sum objective function discussed in Remark 2.3, and thus the applications of smoothness arguments apply in such games. One concrete example of such a game is competitive facility location with price-taking markets and profit-maximizing firms [105]; further applications are explored by Goemans et al. [59].

We claim that every valid utility game with a nondecreasing objective function V is $(1,1)$-smooth, and hence has robust POA at least $1/2$. The proof is essentially a few key inequalities from [105, Theorem 3.2], as follows. Let $\mathbf{s}, \mathbf{s}^*$ denote arbitrary outcomes of such a game. Let $Z_i \subseteq E$ denote the union of all of the players' strategies in $\mathbf{s}$, together with the strategies employed by players $1, 2, \ldots, i$ in $\mathbf{s}^*$. Then

$$\begin{aligned} \sum_{i=1}^{k} \Pi_i(s_i^*, \mathbf{s}_{-i}) &\geq \sum_{i=1}^{k} [V(U(s_i^*, \mathbf{s}_{-i})) - V(U(\emptyset, \mathbf{s}_{-i}))] && (8) \\ &\geq \sum_{i=1}^{k} [V(Z_i) - V(Z_{i-1})] && (9) \\ &\geq W(\mathbf{s}^*) - W(\mathbf{s}), && (10) \end{aligned}$$

where inequality (8) follows from condition (i) of valid utility games; inequality (9) follows from the submodularity of V, with $X = U(s_i^*, \mathbf{s}_{-i})$ and $Y = Z_{i-1}$; and inequality (10) follows from the assumption that V is nondecreasing. This smoothness argument implies a lower bound of 1/2 on the POA of pure Nash equilibria in every valid utility game with a nondecreasing objective function — a result first proved in [105], along with an extension to mixed Nash equilibria and a matching worst-case upper bound. Our extension theorem (Theorem 3.2) shows that this lower bound applies more generally to all of the equilibrium concepts depicted in Figure 1, a fact first established by Blum et al. [20].

Example 2.7 (Simultaneous Second-Price Auctions [36]) There is a set $\{1, 2, \ldots, m\}$ of goods for sale. Each player $i \in \{1, 2, \ldots, k\}$ has a nonnegative valuation $v_i(T)$, or willingness to pay, for each subset T of goods. We assume that every valuation function is submodular. A strategy for a player i consists of a nonnegative bid b_{ij} for each good j such that the sum of the bids $\sum_{j \in T} b_{ij}$ for each subset T of goods is at most the bidder's valuation $v_i(T)$ for it. Each good is allocated independently, to the highest bidder for it, at a price equal to the second-highest bid for the good.

For a bid profile $\mathbf{b}$, let $X_i(\mathbf{b}) \subseteq \{1, 2, \ldots, m\}$ denote the goods that i wins — those on which it is the highest bidder. Define $p_i(\mathbf{b}) = \sum_{j \in X_i(\mathbf{b})} b_{(2)j}$ as the total payment of bidder i, where $b_{(2)j}$ denotes the second-highest bid for the good j. Finally, the payoff $\Pi_i(\mathbf{b}) = v_i(X_i(\mathbf{b})) - p_i(\mathbf{b})$ of bidder i with the bid profile $\mathbf{b}$ is simply its net gain from the auctions. We consider the welfare objective function — the sum of all payoffs, including the revenue of the seller — and denote it by $W(\mathbf{b}) = \sum_{i=1}^{k} v_i(X_i(\mathbf{b}))$. As in Example 2.6, the sum of players' payoffs $\sum_{i=1}^{k} \Pi_i(\mathbf{b})$ is always bounded above by the objective function, and all of the implications of smoothness arguments apply.

Christodoulou et al. [36] show that this game satisfies the following relaxation of (1,1)-smoothness:

there is an optimal bid profile $\mathbf{b}^*$ such that the inequality

$$\sum_{i=1}^{k} \Pi_i(b_i^*, \mathbf{b}_{-i}) \geq W(\mathbf{b}^*) - W(\mathbf{b}) \tag{11}$$

holds for every bid profile $\mathbf{b}$.[3] As mentioned in Remark 2.3, this relaxed property is sufficient for all of the applications of smoothness arguments discussed in this paper. This smoothness argument implies a lower bound of 1/2 on the POA of pure Nash equilibria, which is tight in the worst case [36]. Our extension theorem (Theorem 3.2) shows that this lower bound applies more generally to all of the equilibria depicted in Figure 1.[4]

Here is the bid profile $\mathbf{b}^*$ we use to prove the relaxed smoothness condition. Consider a welfare-maximizing allocation of the goods to the players, in which the goods T_i^* are allocated to bidder i. Consider a bidder i, and assume by relabeling that T_i^* contains the goods $1, 2, \ldots, d$ for some $d \in \{0, 1, \ldots, m\}$. Set $b_{ij}^* = v_i(\{1, 2, \ldots, j\}) - v_i(\{1, 2, \ldots, j-1\})$ for $j = 1, 2, \ldots, d$ and $b_{ij}^* = 0$ for $j > d$. Submodularity of v_i implies that $\sum_{j \in T} b_{ij}^* \leq v_i(T)$ for every bundle $T \subseteq \{1, 2, \ldots, m\}$, and equality holds when $T = T_i^*$. The bids $\mathbf{b}^*$ recover the welfare-maximizing allocation $(T_1^*, \ldots, T_k^*)$.

To verify (11), consider the payoff of a bidder i in the outcome $(b_i^*, \mathbf{b}_{-i})$. On a good $j \in T_i^*$, the bidder either wins at a price of $\max_{\ell \neq i} b_{\ell j}$ or, if $\max_{\ell \neq i} b_\ell \geq b_{ij}^*$, loses and pays nothing. Let T denote the goods of T_i^* that i wins in $(b_i^*, \mathbf{b}_{-i})$. Bidder i's payoff can then be bounded below as follows:

$$\begin{aligned}
\Pi_i(b_i^*, \mathbf{b}_{-i}) &= v_i(T) - \sum_{j \in T} \max_{\ell \neq i} b_{\ell j} \\
&\geq \sum_{j \in T} \left(b_{ij}^* - \max_{\ell \neq i} b_{\ell j} \right) \\
&\geq \sum_{j \in T_i^*} \left(b_{ij}^* - \max_{\ell \neq i} b_{\ell j} \right) \\
&\geq v_i(T_i^*) - \sum_{j \in S_i^*} \max_{\ell=1}^{k} b_{\ell j}.
\end{aligned} \tag{12}$$

Let T_i denote the goods allocated to bidder i in the bid profile $\mathbf{b}$. Summing inequality (12) over all of the bidders and using the fact that highest bidders win, we have

$$\sum_{i=1}^{k} \Pi_i(b_i^*, \mathbf{b}_{-i}) \geq \sum_{i=1}^{k} v_i(T_i^*) - \sum_{j=1}^{m} \max_{\ell} b_{\ell j} = W(\mathbf{b}^*) - \sum_{i=1}^{k} \sum_{j \in T_i} b_{ij}.$$

Bids are constrained to satisfy $\sum_{j \in T_i} b_{ij} \leq v_i(T_i)$ for each bidder i, so the final term of the inequality is at most $W(\mathbf{b})$. This completes the verification of (11).

[3] To see that such games are not always (1,1)-smooth in the stronger sense of Definition 2.1, consider an example with two bidders and one good, with $v_1(\{1\}) = 1$, $v_2(\{1\}) = 2$, $b_{11}^* = 0$, $b_{21}^* = \frac{1}{3}$, $b_{11} = 1$, and $b_{21} = \frac{2}{3}$.

[4] Christodoulou et al. [36] did not discuss correlated or coarse correlated equilibria, but they did prove guarantees for the Bayes-Nash equilibria of the incomplete information version of this game, which is not considered here. See Section 6 for a discussion of smooth games of incomplete information and recent extension theorems for Bayes-Nash equilibria.

2.4 Tight Classes of Games

Smoothness arguments are a restricted form of POA bound that enjoy automatic extensions to, among other things, all of the equilibrium concepts shown in Figure 1. What is the cost of restricting ourselves to this class of proofs? For example, if we care only about the performance of the pure Nash equilibria of a game, can we prove better bounds by departing from the smoothness paradigm?

Examples 2.5–2.7 share a remarkable property: smoothness arguments, despite their restricted form and automatic generality, provide a tight bound on the POA, *even for pure Nash equilibria.* Thus, in these classes of games, the worst-case POA is *exactly the same* for each of the equilibrium concepts of Figure 1.

To define this property precisely, let $\mathcal{G}$ denote a set of cost-minimization games, each with a nonnegative objective function. Let $\mathcal{A}(\mathcal{G})$ denote the parameter values (λ, μ) such that every game of $\mathcal{G}$ is (λ, μ)-smooth. Let $\widehat{\mathcal{G}} \subseteq \mathcal{G}$ denote the games with at least one pure Nash equilibrium, and $\rho_{pure}(G)$ the POA of pure Nash equilibria in a game $G \in \widehat{\mathcal{G}}$. The canonical three-line proof (3)–(5) shows that for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{G})$ and every $G \in \widehat{\mathcal{G}}$, $\rho_{pure}(G) \leq \lambda/(1-\mu)$. We call a set of games *tight* if equality holds for suitable choices of $(\lambda, \mu) \in \mathcal{A}(\mathcal{G})$ and $G \in \widehat{\mathcal{G}}$.

Definition 2.8 (Tight Class of Games) A set $\mathcal{G}$ of cost-minimization games is *tight* if

$$\sup_{G \in \widehat{\mathcal{G}}} \rho_{pure}(G) = \inf_{(\lambda,\mu) \in \mathcal{A}(\mathcal{G})} \frac{\lambda}{1-\mu}. \tag{13}$$

The right-hand side of (13) is the best worst-case upper bound provable via a smoothness argument, and it applies to all of the sets shown in Figure 1. The left-hand side of (13) is the actual worst-case POA of pure Nash equilibria in $\mathcal{G}$ — corresponding to the smallest set in Figure 1 — among games with at least one pure Nash equilibrium. That the left-hand side is trivially upper bounded by the right-hand side is reminiscent of "weak duality." Tight classes of games are characterized by the min-max condition (13), which can be loosely interpreted as a "strong duality-type" result. In a tight class of games, *every* valid upper bound on the worst-case POA of pure Nash equilibria is superseded by a suitable smoothness argument. Thus, every such bound — whether or not it is proved using a smoothness argument — is "intrinsically robust," in that it applies to all of the sets of outcomes in Figure 1.

Prior work [11, 33, 36, 105] effectively showed that the classes of games presented in Examples 2.5–2.7 are tight in the sense of Definition 2.8. The main result in Section 5 is that, for *every* set $\mathcal{C}$ of allowable cost functions, congestion games with cost functions in $\mathcal{C}$ constitute a tight class.

2.5 Non-Examples

Not all POA bounds are equivalent to smoothness arguments, and not all interesting classes of games are tight. Here are two concrete examples.

Example 2.9 (Network Formation Games) Consider a network formation game in which links are formed unilaterally; Fabrikant et al. [46] is one well-known example. That is, the players are the vertices of an undirected graph, each player selects a subset of the other players to connect to via an edge, and an edge is then formed if and only if at least one of its endpoints wants to be connected to the other. A typical player objective function is the sum of two cost terms, one that is increasing in the number of incident edges and one that is increasing in the vertex's distances from the other vertices of the network. Thus a player wants to be central in the network without investing undue resources in maintaining local relationships.

In many such models, all players incur infinite cost when the formed network has more than one connected component. Since an arbitrary "hybrid" outcome $(s_i^*, \mathbf{s}_{-i})$ might well correspond to a disconnected network, even when $\mathbf{s}^*$ is an optimal outcome and $\mathbf{s}$ is a pure Nash equilibrium, such network formation games are not smooth for any finite values of λ and μ. Nonetheless, non-trivial bounds on the POA of pure Nash equilibria are known for such games; see [104] for a survey and [3, 4] for some of the most recent results. These bounds fail to qualify as smoothness proofs because the Nash equilibrium hypothesis is invoked for a hypothetical deviation s_i^* that is a function of the other players' equilibrium strategies $\mathbf{s}_{-i}$.

More generally, in most network formation models the worst-case POA of coarse correlated equilibria is strictly worse than that of pure Nash equilibria, and hence no lossless extension theorem like Theorem 3.2 can apply. Thus, these classes of network formation games are not tight.

Example 2.10 (Symmetric Congestion Games with Singleton Strategies) A more subtle example is provided by symmetric congestion games with *singleton* strategies — equivalently, networks of parallel links — and affine cost functions. The worst-case POA of pure Nash equilibria in such games is precisely 4/3 [6, 52, 78]. The proofs of the POA upper bound use the Nash equilibrium hypothesis in non-obvious ways (cf., Section 2.2). For example, one proof follows from Anshelevich et al. [6, Theorem 3.4] and relies on a characterization of the Nash equilibria of these games as the minimizers of a potential function [51, 69]. Other proofs of this upper bound [52, 78] rely on inequalities beyond the canonical ones in (4), which hold for pure Nash equilibria but not for arbitrary outcomes. None of these proofs cannot be recast as smoothness arguments.

More generally, there is *no* smoothness proof that yields an upper bound of 4/3 on the POA. The reason is that for mixed-strategy Nash equilibria, the worst-case POA in congestion games with singleton strategies and affine cost functions is strictly larger than 4/3 [74]; see also Example 3.1. We conclude that such games do not form a tight class.

3 Extension Theorems

This section states and proves the extension theorems discussed in Section 1.2: every POA bound for pure Nash equilibria that follows from a smoothness argument extends automatically to the more general equilibrium concepts in Figure 1, and to the corresponding outcome sequences in games played over time. Further consequences of smoothness arguments are discussed in Section 4. We work with cost-minimization games, though analogous results hold for smooth payoff-maximization games, such as those in Examples 2.6 and 2.7.

3.1 One-Shot Games

We begin with implications of Definition 2.1 for randomized equilibrium concepts in one-shot games; the next section treats outcome sequences generated by repeated play.

A set $(\sigma_1, \ldots, \sigma_k)$ of independent probability distributions over strategy sets — one per player of a cost-minimization game — is a *mixed Nash equilibrium* of the game if no player can decrease its expected cost under the product distribution $\sigma = \sigma_1 \times \cdots \times \sigma_k$ via a unilateral deviation:

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s}_{-i}\sim\sigma_{-i}}[C_i(s_i', \mathbf{s}_{-i})]$$

for every i and $s_i' \in S_i$, where σ_{-i} is the product distribution of all σ_j's other than σ_i. (By linearity, it suffices to consider only pure-strategy unilateral deviations.) Obviously, every pure Nash equilibrium is a mixed Nash equilibrium and not conversely; indeed, many games have no pure Nash equilibria, but every finite game has at least one mixed Nash equilibrium [86].

A *correlated equilibrium* [9] of a cost-minimization game G is a (joint) probability distribution σ over the outcomes of G with the property that

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})|s_i] \leq \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(s_i', \mathbf{s}_{-i})|s_i] \tag{14}$$

for every i and $s_i, s_i' \in S_i$. A classical interpretation of a correlated equilibrium is in terms of a mediator, who draws an outcome $\mathbf{s}$ from the publicly known distribution σ and privately "recommends" strategy s_i to each player i. The equilibrium condition requires that following a recommended strategy always minimizes the expected cost of a player, conditioned on the recommendation. Mixed Nash equilibria correspond to the correlated equilibria that are also product distributions. Correlated equilibria have been widely studied as strategies for a benevolent mediator, and also because of their relative tractability. The set of correlated equilibria is explicitly described by a small set of linear inequalities, so computing (and even optimizing over) correlated equilibria can be done in time polynomial in the size of the game [58]. They are also relatively "easy to learn," as discussed in the next section.

Finally, a *coarse correlated equilibrium* [84] of a cost-minimization game is a probability distribution σ over outcomes that satisfies

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(s_i', \mathbf{s}_{-i})] \tag{15}$$

for every i and $s_i' \in S_i$. The set of all such distributions is sometimes called the *Hannan set*, after Hannan [61]. While a correlated equilibrium (14) protects against deviations by a player aware of its recommended strategy, a coarse correlated equilibrium (15) is only constrained by player deviations that are independent of the sampled outcome. Since every correlated equilibrium is also a coarse correlated equilibrium, coarse correlated equilibria can only be easier to compute and learn, and are thus a still more plausible prediction for the realized play of a game.

Example 3.1 (Equilibrium Concepts) All of the inclusions in Figure 1 are generally strict. To see this and to illustrate the different equilibrium concepts, consider a congestion game (Example 2.5) with four players, a set $E = \{0, 1, 2, 3, 4, 5\}$ of six resources each with cost function $c(x) = x$, and singleton strategies, meaning $S_i = E$ for each player i. The pure Nash equilibria of this game are the $\binom{6}{4}$ outcomes in which each player chooses a distinct resource. Every player suffers only unit cost in such an equilibrium. One mixed Nash equilibrium that is obviously not pure has each player independently choosing a resource uniformly at random. Every player suffers expected cost $3/2$ in this equilibrium. The uniform distribution over all outcomes in which there is one resource with two players and two resources with one player each is a (non-product) correlated equilibrium, since both sides of (14) read $\frac{3}{2}$ for every i, s_i, and s_i'. The uniform distribution over the subset of these outcomes in which the set of chosen resources is either $\{0, 2, 4\}$ or $\{1, 3, 5\}$ is a coarse correlated equilibrium, since both sides of (15) read $\frac{3}{2}$ for every i and s_i'. It is not a correlated equilibrium, since a player i that is recommended the resource s_i can reduce its conditional expected cost to 1 by choosing the deviation s_i' to the successive resource (modulo 6).

We now give our extension theorem for equilibrium concepts in one-shot games: every POA bound proved via a smoothness argument extends automatically to the set of coarse correlated equilibria. With the "correct" definitions in hand, the proof writes itself.

Theorem 3.2 (Extension Theorem — Static Version) *For every cost-minimization game G with robust POA $\rho(G)$, every coarse correlated equilibrium σ of G, and every outcome $\mathbf{s}^*$ of G,*

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})] \le \rho(G) \cdot C(\mathbf{s}^*).$$

Proof: Let G be a (λ, μ)-smooth cost-minimization game, σ a coarse correlated equilibrium, and $\mathbf{s}^*$ an outcome of G. We can write

$$\begin{aligned}
\mathbf{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})] &= \mathbf{E}_{\mathbf{s}\sim\sigma}\left[\sum_{i=1}^{k} C_i(\mathbf{s})\right] && (16)\\
&= \sum_{i=1}^{k} \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] && (17)\\
&\le \sum_{i=1}^{k} \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(s_i^*, \mathbf{s}_{-i})] && (18)\\
&= \mathbf{E}_{\mathbf{s}\sim\sigma}\left[\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i})\right] && (19)\\
&\le \mathbf{E}_{\mathbf{s}\sim\sigma}[\lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})] && (20)\\
&= \lambda \cdot C(\mathbf{s}^*) + \mu \cdot \mathbf{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})], && (21)
\end{aligned}$$

where equality (16) follows from the definition of the objective function, equalities (17), (19), and (21) follow from linearity of expectation, inequality (18) follows from the definition (15) of a coarse correlated equilibrium (applied once per player i, with the hypothetical deviation s_i^*), and inequality (20) follows from the assumption that the game is (λ, μ)-smooth. Rearranging terms completes the proof. ■

3.2 Repeated Play and No-Regret Sequences

The extension theorem (Theorem 3.2) applies equally well to certain outcome sequences generated by repeated play. To illustrate this point, consider a sequence $\mathbf{s}^1, \mathbf{s}^2, \ldots, \mathbf{s}^T$ of outcomes of a (λ, μ)-smooth game and a minimum-cost outcome $\mathbf{s}^*$ of the game. For each i and t, define

$$\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, \mathbf{s}_{-i}^t) \tag{22}$$

as the hypothetical improvement in player i's cost at time t had it used the strategy s_i^* in place of s_i^t. When $\mathbf{s}^t$ is a Nash equilibrium, $\delta_i(\mathbf{s}^t)$ is nonnegative; for an arbitrary outcome $\mathbf{s}^t$, $\delta_i(\mathbf{s}^t)$ can be positive or negative. We can mimic the derivation in (3)–(5) to obtain

$$C(\mathbf{s}^t) \le \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{\sum_{i=1}^{k} \delta_i(\mathbf{s}^t)}{1-\mu} \tag{23}$$

for each t.

This section concerns outcome sequences in which every player i experiences *vanishing average (external) regret*, meaning that its cost over time is asymptotically competitive with that of every time-invariant strategy:

$$\frac{1}{T}\sum_{t=1}^{T} C_i(\mathbf{s}^t) \le \frac{1}{T}\left[\min_{s_i'} \sum_{t=1}^{T} C_i(s_i', \mathbf{s}_{-i}^t)\right] + o(1), \tag{24}$$

where the $o(1)$ term denotes some function that goes to 0 as $T \to \infty$. The condition (24) is a time-averaged analog of the Nash equilibrium condition (1), but it does not preclude highly oscillatory behavior over large time horizons. For example, repeatedly cycling through all of the outcomes in the support of the coarse correlated equilibrium in Example 3.1 yields arbitrarily long outcome sequences in which every player has zero regret. The most significant motivation for considering outcome sequences in which every player has vanishing average regret is that there are several simple "off-the-shelf" online learning algorithms with good convergence rates that are guaranteed to generate such sequences. See, for example, Cesa-Bianchi and Lugosi [26].

For such a sequence, we can proceed as follows. Averaging (23) over the T time steps and reversing the order of the resulting double summation yields

$$\frac{1}{T}\sum_{t=1}^{T} C(\mathbf{s}^t) \leq \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{1}{1-\mu}\sum_{i=1}^{k}\left(\frac{1}{T}\sum_{t=1}^{T}\delta_i(\mathbf{s}^t)\right). \tag{25}$$

Recalling from (22) that $\delta_i(\mathbf{s}^t)$ is the additional cost incurred by player i at time t due to playing strategy s_i^t instead of the (time-invariant) strategy s_i^*, the no-regret guarantee (24) implies that $[\sum_{t=1}^{T}\delta_i(\mathbf{s}^t)]/T$ is bounded above by a term that goes to 0 with T. Since this holds for every player i, inequality (25) implies that the average cost of outcomes in the sequence is no more than the robust POA times the cost of an optimal outcome, plus an error term that approaches zero as $T \to \infty$.

Theorem 3.3 (Extension Theorem — Repeated Version) *For every cost-minimization game G with robust POA $\rho(G)$, every outcome sequence $\mathbf{s}^1, \ldots, \mathbf{s}^T$ that satisfies* (24) *for every player, and every outcome $\mathbf{s}^*$ of G,*

$$\frac{1}{T}\sum_{t=1}^{T} C(\mathbf{s}^t) \leq [\rho(G) + o(1)] \cdot C(\mathbf{s}^*)$$

as $T \to \infty$.

Blum et al. [20] were the first to consider bounds of this type, calling them "the price of total anarchy."

We reiterate that the approximation bound in Theorem 3.3 is significantly more compelling, and assumes much less from both the game and its participants, than one that applies only to Nash equilibria. Nash equilibria can be intractable or impossible to find while, as mentioned, simple online learning algorithms guarantee vanishing average regret for every player. Of course, the guarantee bound in Theorem 3.3 makes no reference to which learning algorithms, if any, the players' use to play the game — the bound applies whenever repeated joint play has low regret, whatever the reason.

Remark 3.4 (Mixed-Strategy No-Regret Sequences) For simplicity, the condition (24) and Theorem 3.3 are stated for sequences of pure outcomes. These are easily extended to mixed outcomes: for every cost-minimization game G with robust POA $\rho(G)$, every sequence $\sigma^1, \ldots, \sigma^T$ of (not necessarily product) probability distributions over outcomes that satisfies

$$\mathbf{E}_{\mathbf{s}^t \sim \sigma^t}\left[\sum_{t=1}^{T} C_i(\mathbf{s}^t)\right] \leq \mathbf{E}_{\mathbf{s}^t_{-i} \sim \sigma^t_{-i}}\left[\sum_{t=1}^{T} C_i(s_i', \mathbf{s}^t_{-i})\right] + o(T)$$

for every player i, we have

$$\frac{1}{T}\sum_{t=1}^{T}\mathbf{E}_{\mathbf{s}^t\sim\sigma^t}[C(\mathbf{s}^t)] \le [\rho(G)+o(1)]\cdot C(\mathbf{s}^*)$$

as $T\to\infty$.

Remark 3.5 (Equivalence of Theorems 3.2 and 3.3) Theorems 3.2 and 3.3 are essentially equivalent, in that either one can be derived from the other. The reason is that the empirical distributions of sequences in which every player has vanishing average regret approximate, arbitrarily closely as $T\to\infty$, the set of coarse correlated equilibria.

Remark 3.6 (Correlated Equilibria and Internal Regret) Correlated equilibria correspond, in the sense of Remark 3.5, to outcome sequences in which every player has nonpositive "internal" regret; see Blum and Mansour [21]. There are also several computationally efficient online learning algorithms that guarantee a player vanishing average internal regret in an arbitrary game [50, 65].

4 Additional Consequences of Smoothness

Smooth games enjoy robustness properties beyond those guaranteed by the main extension theorems (Theorems 3.2 and 3.3). Section 4.1 proves that approximate equilibria of smooth games approximately obey the robust POA guarantee. Section 4.2 establishes bicriteria bounds for smooth games, where the performance of equilibria is compared to that of an optimal outcome with a different number of players. Section 4.3 considers smooth potential games and shows that many forms of best-response dynamics rapidly converge to outcomes in which the robust POA guarantee approximately holds.

4.1 Approximate Equilibria

Every POA bound proved via a smoothness argument applies automatically to approximate equilibria, with the bound degrading gracefully as a function of the approximation parameter. For instance, define an ϵ-*Nash equilibrium* of a cost-minimization game as a strategy profile $\mathbf{s}$ in which no player can decrease its cost by more than a $(1+\epsilon)$ factor via a unilateral deviation:

$$C_i(\mathbf{s}) \le (1+\epsilon)\cdot C_i(s_i',\mathbf{s}_{-i}) \tag{26}$$

for every i and $s_i'\in S_i$. Approximate versions of the other three equilibrium concepts studied in Section 3.1 can be defined in the same way.

Mimicking the derivation (3)–(5) for exact Nash equilibria, replacing in (4) the exact equilibrium condition (1) by the approximate one (26), shows that every ϵ-Nash equilibrium of a (λ,μ)-smooth cost-minimization game with $\epsilon<\frac{1}{\mu}-1$ has cost at most $\frac{(1+\epsilon)\lambda}{1-\mu(1+\epsilon)}$ times that of an optimal solution. Just as the bound for exact pure Nash equilibria extends to coarse correlated equilibria (Theorem 3.2), this bound for ϵ-Nash equilibria extends to ϵ-coarse correlated equilibria.

Theorem 4.1 (Extension Theorem for Approximate Equilibria) *For every (λ, μ)-smooth cost-minimization game G, every $\epsilon < \frac{1}{\mu} - 1$, every ϵ-coarse correlated equilibrium σ of G, and every outcome $\mathbf{s}^*$ of G,*

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C(\mathbf{s})] \leq \frac{(1+\epsilon)\lambda}{1-\mu(1+\epsilon)} \cdot C(\mathbf{s}^*). \tag{27}$$

Example 4.2 (Congestion Geams with Affine Cost Functions) Since every congestion game with affine cost functions is $(\frac{5}{3}, \frac{1}{3})$-smooth (Example 2.5), Theorem 4.1 implies that every ϵ-coarse correlated equilibrium of such a game with $\epsilon < 2$ has expected cost at most $\frac{5-5\epsilon}{2-\epsilon}$ times that of an optimal outcome.

Remark 4.3 (Optimal Smoothness Parameters Can Depend on ϵ) Theorem 4.1 applies to every choice (λ, μ) of smoothness parameters for a class of games, not just the choice that minimizes the robust POA $\frac{\lambda}{1-\mu}$. The smoothness parameters that minimize the POA bound $\frac{(1+\epsilon)\lambda}{1-\mu(1+\epsilon)}$ for ϵ-equilibria for some $\epsilon > 0$ need not be the optimal parameters for the $\epsilon = 0$ case. For example, in congestion games with affine cost functions, the standard smoothness parameters $(\frac{5}{3}, \frac{1}{3})$ give no POA bounds whatsoever when $\epsilon \geq 2$. Christodoulou et al. [35] show how to choose smoothness parameters $(\lambda(\epsilon), \mu(\epsilon))$ as a function of ϵ to obtain tight POA bounds for the ϵ-approximate equilibria of such games, for all $\epsilon \geq 0$.

Remark 4.4 (Payoff-Maximization Games) Analogous results hold in smooth payoff-maximization games. Here, for $\epsilon \in [0, 1)$, we define an ϵ-coarse correlated equilibria as a probability distribution σ over outcomes that satisfies $\mathbf{E}_{\mathbf{s}\sim\sigma}[\Pi_i(\mathbf{s})] \geq (1-\epsilon)\mathbf{E}_{\mathbf{s}\sim\sigma}[\Pi_i(s'_i, \mathbf{s}_{-i})]$ for every player i and unilateral deviation $s'_i \in S_i$. For every (λ, μ)-smooth payoff-maximization game G, every $\epsilon \in [0, 1)$, every ϵ-coarse correlated equilibrium σ of G, and every outcome $\mathbf{s}^*$ of G, the expected objective function value under σ is at least $\frac{(1-\epsilon)\lambda}{1+\mu(1-\epsilon)}$ times that of $\mathbf{s}^*$. For instance, ϵ-coarse correlated equilibria of valid utility games with a nondecreasing objective function (Example 2.6) and simultaneous second-price auctions with submodular bidder valuations (Example 2.7) have expected welfare at least $\frac{1-\epsilon}{2-\epsilon}$ times that of an optimal outcome.

4.2 Bicriteria Bounds

This section derives "bicriteria" or "resource augmentation" bounds for smooth games, where the objective function value of the worst equilibrium is compared to the optimal outcome with a different number of players.

4.2.1 Cost-Minimization Games

This section and the next consider sets of games $\mathcal{G}$ that are closed under player deletions and player duplications, meaning that applying either of these operations to a game $\mathcal{G}$ yields another game of $\mathcal{G}$. Congestion games (Example 2.5), several natural classes of utility games (Example 2.6), and simultaneous second-price auction games (Example 2.7) are all closed in this sense.

Bicriteria bounds follow from a strengthened version of Definition 2.1 that accommodates duplicated players. Below, we write C^G and C_i^G for the objective function and player cost functions of a cost-minimization game G.

Definition 4.5 (Smooth Closed Sets of Cost-Minimization Games) Let $\mathcal{G}$ be a set of cost-minimization games that is closed under player deletions and duplications. The set $\mathcal{G}$ is (λ, μ)-*smooth* if for every outcome $\mathbf{s}$ of a game $G \in \mathcal{G}$, and every outcome $\mathbf{s}^*$ of a game $\widehat{G}$ obtained from G by duplicating each player i n_i times,

$$\sum_{i=1}^{k} \sum_{j=1}^{n_i} C_i^G(s^*_{(i,j)}, \mathbf{s}_{-i}) \leq \lambda \cdot C^{\widehat{G}}(\mathbf{s}^*) + \mu \cdot C^G(\mathbf{s}), \tag{28}$$

where $s^*_{(i,j)}$ denotes the strategy chosen by the jth copy of player i in $\mathbf{s}^*$.

For example, consider congestion games with affine cost functions. The derivation in Example 2.5, based on the inequality (6), shows that this (closed) set of games is $(\frac{5}{3}, \frac{1}{3})$-smooth in the sense of Definition 4.5. More generally, the results of Section 5 imply that whenever Definition 2.1 holds for congestion games with cost functions restricted to a set $\mathcal{C}$, Definition 4.5 holds as well, with the same values of λ and μ.

Theorem 4.6 (Bicriteria Bound for Smooth Cost-Minimization Games) *Let $\mathcal{G}$ be a (λ, μ)-smooth set of cost-minimization games that is closed under player deletions and duplications, and ℓ a positive integer. For every pure Nash equilibrium $\mathbf{s}$ of a game $G \in \mathcal{G}$ and every outcome $\mathbf{s}^*$ of the game $\widehat{G}$ in which each player of G is duplicated ℓ times,*

$$C^G(\mathbf{s}) \leq \frac{\lambda}{\ell - \mu} \cdot C^{\widehat{G}}(\mathbf{s}^*).$$

Proof: Write $s^*_{(i,j)}$ for the strategy used by the jth copy of player i in $\mathbf{s}^*$. Applying our usual assumption about the objective function, the Nash equilibrium condition, and smoothness yields

$$\begin{aligned} \ell \cdot C^G(\mathbf{s}) &\leq \sum_{j=1}^{\ell} \sum_{i=1}^{k} C_i^G(\mathbf{s}) \\ &\leq \sum_{j=1}^{\ell} \sum_{i=1}^{k} C_i^G(s^*_{(i,j)}, \mathbf{s}_{-i}) \\ &\leq \lambda \cdot C^{\widehat{G}}(\mathbf{s}^*) + \mu \cdot C^G(\mathbf{s}); \end{aligned}$$

rearranging as usual completes the proof. ■

Example 4.7 (Congestion Games with Affine Cost Functions) Theorem 4.6 implies that the Nash equilibria of a game in a (λ, μ)-smooth closed set cost no more than that of an optimal outcome after every player has been duplicated at least $\lambda + \mu$ times. For example, in congestion games with affine cost functions, the cost of every Nash equilibrium is bounded above by that of an optimal outcome with two copies of every player.

4.2.2 Payoff-Maximization Games

In payoff-maximization games, bicriteria bounds award additional players to the Nash equilibrium, rather than to the optimal outcome. The analog of Definition 4.5 is the following.

Definition 4.8 (Smooth Closed Sets of Payoff-Maximization Games) Let $\mathcal{G}$ be a set of payoff-maximization games that is closed under player deletions and duplications. The set $\mathcal{G}$ is (λ, μ)-*smooth* if for every outcome $\mathbf{s}$ of a game $G \in \mathcal{G}$, and every outcome $\mathbf{s}^*$ of a game $\widehat{G}$ with a subset A of the players of G,

$$\sum_{i \in A} \Pi_i^G(s_i^*, \mathbf{s}_{-i}) \geq \lambda \cdot W^{\widehat{G}}(\mathbf{s}^*) - \mu \cdot W^G(\mathbf{s}). \tag{29}$$

For example, the derivation (8)–(10) shows that the condition in (29) is satisfied, with $\lambda = \mu = 1$, by every valid utility game with a nondecreasing objective function. The derivation in Example 2.7 shows that the set of simultaneous second-price auction games with submodular bidder valuations satisfies a relaxed version of Definition 4.8, with $\lambda = \mu = 1$, in which the inequality (29) holds for all outcomes $\mathbf{s}$ of G and for a judiciously chosen welfare-maximizing outcome $\mathbf{s}^*$ of $\widehat{G}$. The following bicriteria bound also holds under this weaker condition.

Theorem 4.9 (Bicriteria Bound for Smooth Payoff-Maximization Games) *Let $\mathcal{G}$ be a (λ, μ)-smooth set of payoff-maximization games that is closed under player deletions and duplications, and ℓ a positive integer. For every outcome $\mathbf{s}^*$ of a game $\widehat{G} \in \mathcal{G}$ and every pure Nash equilibrium $\mathbf{s}$ of the game G in which each player of $\widehat{G}$ is duplicated ℓ times,*

$$W^G(\mathbf{s}) \geq \frac{\ell\lambda}{1 + \ell\mu} \cdot W^{\widehat{G}}(\mathbf{s}^*) = \frac{\lambda}{\frac{1}{\ell} + \mu} \cdot W^{\widehat{G}}(\mathbf{s}^*).$$

Proof: Write $\mathbf{s}_{-(i,j)}$ for the strategies of $\mathbf{s}$ other than that chosen by the jth copy of player i, and $\Pi_{i,j}^G$ for the payoff function of the jth copy of player i in G. Applying our usual assumption about the objective function, the Nash equilibrium condition, and smoothness yields

$$\begin{aligned} W^G(\mathbf{s}) &\geq \sum_{j=1}^{\ell} \sum_{i=1}^{k} \Pi_{i,j}^G(\mathbf{s}) \\ &\geq \sum_{j=1}^{\ell} \sum_{i=1}^{k} \Pi_{i,j}^G(s_i^*, \mathbf{s}_{-(i,j)}) \\ &\geq \ell \left(\lambda \cdot W^{\widehat{G}}(\mathbf{s}^*) - \mu \cdot W^G(\mathbf{s}) \right); \end{aligned}$$

rearranging completes the proof. ■

Remark 4.10 (Impossibility of Recovering the Optimal Welfare) No number ℓ of duplicate players is sufficient to guarantee that Nash equilibria in the modified payoff-maximization game have performance as good as that of an optimal outcome in the original game (cf., Theorem 4.6). To see this, fix a positive number ℓ and consider the following ($(1,1)$-smooth) valid utility game. In the original game $\widehat{G}$, there are $k \geq 2$ players and k resources. Player 1 can only use resource 1; player $i > 1$ can use resource 1 or resource i. If x players choose resource 1, then they each receive payoff $k\ell/x$. If x players choose resource i, then they each receive payoff $1/x$. Define the objective function to be the sum of players' payoffs. In the optimal solution of $\widehat{G}$, player i chooses resource i and the welfare is $k\ell + k - 1$. In the game G in which every player is duplicated ℓ times, there is a Nash equilibrium in which every player chooses the resource 1, yielding welfare only $k\ell$.

4.2.3 Further Discussion

We conclude this section with three comments on Theorems 4.6 and 4.9. First, analogous to bounds for approximate equilibria (Remark 4.3), the smoothness parameters that optimize the POA ($\frac{\lambda}{1-\mu}$ or $\frac{\lambda}{1+\mu}$) need not be the same ones that optimize bicriteria bounds ($\frac{\lambda}{\ell-\mu}$ or $\frac{\ell\lambda}{1+\ell\mu}$).

Second, if the POA of a cost-minimization game is ρ and the cost of an optimal outcome increases at least linearly with the number ℓ of player duplicates, as in congestion games with nondecreasing cost functions, then there is a trivial bicriteria bound of ρ/ℓ. For smooth games, this bound of $\frac{\lambda}{\ell-\ell\mu}$ is significantly weaker than that in Theorem 4.6. For payoff-maximization games, it is not obvious that duplicating players improves the approximation bound for worst-case Nash equilibria at all. For smooth games, Theorem 4.9 guarantees such an improvement.

Finally, following the proofs of Theorems 3.2 and 3.3 shows that the approximation guarantees of Theorems 4.6 and 4.9 also extend to all coarse correlated equilibria of and no-regret sequences in the game G.

4.3 Short Best-Response Sequences

Our extension theorem for outcome sequences generated by no-regret learners (Theorem 3.3) shows that good approximation bounds apply to fundamental classes of learning dynamics in smooth games, even when such dynamics fail to converge to a Nash equilibrium. This section investigates another well-studied learning procedure, *best-response dynamics*.

Best-response dynamics (BRD) is a natural myopic model of how players might search for a pure Nash equilibrium: if the current outcome $\mathbf{s}$ is not a pure Nash equilibrium, then some player i that can benefit from a unilateral deviation switches to a strategy that minimizes its cost $C_i(s_i', \mathbf{s}_{-i})$. BRD cannot converge in games without pure Nash equilibria, and might fail to converge even in games that do have such equilibria. These facts motivate the search for a general extension theorem, analogous to Theorem 3.3, for BRD. Sadly, Mirrokni and Vetta [80] showed that no such general extension theorem exists: there are ($(1,1)$-smooth) valid utility games in the sense of Example 2.6, with an arbitrarily large number k of players, such that BRD only visits outcomes with welfare $1/k$ times that of an optimal solution.

We next prove guarantees on the performance of BRD in smooth games under two additional assumptions. First, we restrict attention to *potential games* [81], meaning games that admit a *potential function* Φ, which is a function on the game outcomes that satisfies

$$\Phi(\mathbf{s}) - \Phi(s_i', \mathbf{s}_{-i}) = C_i(\mathbf{s}) - C_i(s_i', \mathbf{s}_{-i}) \tag{30}$$

for every outcome $\mathbf{s}$, player i, and deviation $s_i' \in S_i$. That is, a potential function tracks the change in a unilateral deviator's cost.

BRD converges to a pure Nash equilibrium in every finite potential game, since equation (30) guarantees that every iteration strictly improves the potential function.

Example 4.11 (Congestion Games [91]) Every congestion game is a potential game, with the potential function

$$\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{x_e} c_e(i),$$

where x_e denotes the number of players using resource e in $\mathbf{s}$.

Example 4.12 (Basic Utility Games [105]) Valid utility games do not always possess pure Nash equilibria [105], so they are not always potential games. Vetta [105] defined a *basic* utility game as one for which the welfare function W is a potential function. Concrete examples of basic utility games include competitive facility location [105] and certain market-sharing games [59].

Since BRD eventually converges to a pure Nash equilibrium in a finite potential game, every POA bound — robust or otherwise — applies to sufficiently long BRD sequences. However, the number of iterations required for convergence can be exponential in the number of players, no matter how the deviating player is chosen in each iteration, even in congestion games with affine cost functions [1, 47]. Similar lower bounds hold for reaching an approximate Nash equilibrium [100]. For potential games, the technically and conceptually interesting goal is to prove approximation bounds for BRD that apply after a relatively small number of iterations, long before convergence to an (approximate) Nash equilibrium is guaranteed.

Our second assumption restricts how the deviating player in each iteration of BRD is chosen. Without some such restriction, BRD can require an exponential number of iterations to reach a state with a non-trivial approximation guarantee, even in congestion games with affine cost functions [10, Theorem 3.4]. Several different assumptions imply convergence to outcomes that approximately obey the robust POA bound within a polynomial number of iterations. Roughly, as long as the deviating player is likely to have at least an approximately average incentive to deviate, relative to the other players, then BRD rapidly reaches near-optimal outcomes in smooth potential games. We next treat two concrete restrictions of BRD in detail.

Define *maximum-gain BRD* as the specialization of BRD in which, in each iteration, the player with the most-improving unilateral deviation is chosen.

Theorem 4.13 (Maximum-Gain BRD in Smooth Potential Games) *Let G be a (λ, μ)-smooth cost-minimization game with k players and a positive potential function Φ that satisfies $\Phi(\mathbf{s}) \leq C(\mathbf{s})$ for every outcome $\mathbf{s}$. Let $\mathbf{s}^0, \ldots, \mathbf{s}^T$ be a sequence generated by maximum-gain BRD, $\mathbf{s}^*$ a minimum-cost outcome, and $\epsilon > 0$ a parameter. Then all but*

$$O\left(\frac{k}{\epsilon(1-\mu)} \log \frac{\Phi(\mathbf{s}^0)}{\Phi_{min}}\right) \tag{31}$$

states $\mathbf{s}^t$ satisfy

$$C(\mathbf{s}^t) \leq \left(\frac{\lambda}{1-\mu} + \epsilon\right) \cdot C(\mathbf{s}^*),$$

where Φ_{min} is the minimum potential function value of an outcome of G.

Proof: Define $\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, \mathbf{s}_{-i}^t)$ and $\Delta(\mathbf{s}^t) = \sum_{i=1}^k \delta_i(\mathbf{s}^t)$. Call a state $\mathbf{s}^t$ *bad* if

$$\Delta(\mathbf{s}^t) \geq \epsilon(1-\mu)C(\mathbf{s}^t).$$

Inequality (23) reduces the theorem to proving that only $O(\frac{k}{\epsilon(1-\mu)} \log \frac{\Phi(\mathbf{s}^0)}{\Phi_{min}})$ states $\mathbf{s}^t$ are bad.

The potential function Φ strictly decreases in each iteration of BRD. For a nonnegative integer j, we define the *jth phase* of the sequence $\mathbf{s}^0, \ldots, \mathbf{s}^T$ as the (contiguous) subsequence of states in which $\Phi(\mathbf{s}^t) \in (2^{-(j+1)} \cdot \Phi(\mathbf{s}^0), 2^{-j} \cdot \Phi(\mathbf{s}^0)]$. There are only $\approx \log_2 \frac{\Phi(\mathbf{s}^0)}{\Phi_{min}}$ phases.

We complete the proof by showing that each phase contains $O(\frac{k}{\epsilon(1-\mu)})$ bad states. In a bad state $\mathbf{s}^t$, since Φ underestimates C, $\Delta(\mathbf{s}^t) \geq \epsilon(1-\mu)C(\mathbf{s}^t) \geq \epsilon(1-\mu)\Phi(\mathbf{s}^t)$. If a player i chooses a best response to the outcome $\mathbf{s}^t$, its cost decreases by at least $\delta_i(\mathbf{s}^t)$. Thus, in a bad state $\mathbf{s}^t$, the cost of the player chosen by maximum-gain BRD decreases by at least $\frac{\epsilon(1-\mu)}{k}\Phi(\mathbf{s}^t)$. Since Φ is a potential function, $\Phi(\mathbf{s}^{t+1}) \leq (1-\frac{\epsilon(1-\mu)}{k})\Phi(\mathbf{s}^t)$ whenever $\mathbf{s}^t$ is a bad state. This implies that there are $O(\frac{k}{\epsilon(1-\mu)})$ bad states in each phase, completing the proof. ■

Define *random BRD* as the specialization of BRD in which the deviating player in each iteration is chosen independently and uniformly at random.

Theorem 4.14 (Random BRD in Smooth Potential Games) *Let G be a (λ,μ)-smooth cost-minimization game with a positive potential function Φ that satisfies $\Phi(\mathbf{s}) \leq C(\mathbf{s})$ for every outcome $\mathbf{s}$. Let $\mathbf{s}^0,\ldots,\mathbf{s}^T$ be a sequence generated by random BRD, $\mathbf{s}^*$ a minimum-cost outcome, and $\epsilon > 0$ a parameter. Let Ψ denote $\frac{\Phi(\mathbf{s}^0)}{\Phi_{min}}$, where Φ_{min} is the minimum potential function value of an outcome of G. Then with probability at least $1-\frac{1}{k}$ over the generated sequence, all but*

$$O\left(\frac{k(\log k + \log\log\Psi)}{\epsilon(1-\mu)}\log\Psi\right) \tag{32}$$

states $\mathbf{s}^t$ satisfy

$$C(\mathbf{s}^t) \leq \left(\frac{\lambda}{1-\mu}+\epsilon\right)\cdot C(\mathbf{s}^*).$$

Proof: Using the notation of the proof of Theorem 4.13, in each iteration t of random BRD we choose a (random) player $i(t)$ to deviate that satisfies $\mathbf{E}\left[\delta_{i(t)}(\mathbf{s}^t)|\mathbf{s}^t\right] \geq \Delta(\mathbf{s}^t)/k$. Thus, in a bad state $\mathbf{s}^t$,

$$\mathbf{E}\left[\Phi(\mathbf{s}^{t+1})|\mathbf{s}^t\right] \leq \left(1-\frac{\epsilon(1-\mu)}{k}\right)\Phi(\mathbf{s}^t). \tag{33}$$

Define phases as in the proof of Theorem 4.13. The potential function value $\Phi(\mathbf{s}^t)$ is decreasing in t with probability 1, and there are only $\approx \log_2\Psi$ phases. We prove that, for each j, with probability at least $1-\frac{1}{k\cdot\log_2\Psi}$, phase j contains $O(\frac{k(\log k+\log\log\Psi)}{\epsilon(1-\mu)})$ bad states. The theorem then follows from a union bound over the phases.

The argument is the same for every phase, so we focus on the outcomes in phase 0. Let X_t denote $\Phi(\mathbf{s}^t)$ if $\mathbf{s}^t$ belongs to the 0th phase, and 0 otherwise. Define $Y_t = X_t - X_{t-1}$ for $t \geq 1$, which lies in $[0, X_t]$ with probability 1. Also, if $\mathbf{s}^t$ is a bad state in phase 0, then inequality (33) implies that $\mathbf{E}\left[Y_t|X_t\right] \geq 2\alpha X_t \geq \alpha\Phi(\mathbf{s}^0)$, where $\alpha = \frac{\epsilon(1-\mu)}{2k}$.

The probability that phase 0 contains more than m bad states is bounded above by the probability that the sum of m independent nonnegative random variables, each bounded by $\Phi(\mathbf{s}^0)$ and with expectation $\alpha\Phi(\mathbf{s}^0)$, is at most $\Phi(\mathbf{s}^0)/2$. Standard Chernoff bounds (see, e.g., Motwani and Raghavan [83]) imply that the latter probability is at most $\frac{1}{k\cdot\log_2\Psi}$ provided $m \geq \frac{c}{\alpha}(\log k + \log\log\Psi)$ for a sufficiently large constant c. The proof is complete. ■

Remark 4.15 (Discussion of Theorems 4.13 and 4.14) The convergence bounds (31) and (32) in Theorems 4.13 and 4.14 are polynomial in the number of players k and the number of bits used to specify the potential function. By contrast, the number of iterations needed for BRD to converge to an (approximate) Nash equilibrium of a potential game can be exponential in both of these quantities [1, 47, 100].

The condition that $\Phi(\mathbf{s}) \leq C(\mathbf{s})$ for every outcome $\mathbf{s}$ in Theorems 4.13 and 4.14 holds in congestion games (Example 4.11) since cost functions are assumed to be nondecreasing. If this condition is relaxed to $\Phi(\mathbf{s}) \leq M \cdot C(\mathbf{s})$ for some constant M, then the upper bounds on the number of bad states increase by a factor of M. For payoff-maximization games, the potential function restriction changes to $\Phi(\mathbf{s}) \geq W(\mathbf{s})$. This condition is satisfied with equality in the basic utility games of Example 4.12.

Theorems 4.13 and 4.14 provide upper bounds N on the total number of "bad states" — states that fail to approximately obey the robust POA bound. These results do not imply that every state beyond the first N are good, however — since an iteration of BRD can strictly increase the overall cost, a good state can be followed by a bad one. When the potential function gives a two-sided approximation of the cost function, however, the first good state can only be succeeded by "approximately good" states. Precisely, if, for some $M \geq 1$, $\phi(\mathbf{s}) \leq C(\mathbf{s}) \leq M \cdot \phi(\mathbf{s})$ for every outcome $\mathbf{s}$, then an arbitrary number of BRD iterations can only increase the cost by a factor of M. For instance, in congestion games with cost functions that are polynomials with nonnegative coefficients and degree at most d, $M = d + 1$.

Remark 4.16 (Guarantees for ϵ-BRD) Here is a third approach for ensuring that the improvement of the deviating player chosen by BRD is related to that of an average player — and thus, along the lines of Theorems 4.13 and 4.14, that there can only be a polynomial number of bad states. In every iteration of *ϵ-BRD*, a deviating player is chosen such that a best response will decrease its cost by at least a $(1 + \epsilon)$ factor. ϵ-BRD can only terminate at an ϵ-approximate Nash equilibrium (see Section 4.1). Theorem 4.1 applies upon convergence, but convergence can take an exponential number of iterations [100]. Awerbuch et al. [10] effectively proved that, under a mild additional Lipschitz condition on a game's cost functions, and assuming that every player is given the opportunity to move at least once every polynomial number of iterations, ϵ-BRD generates only a polynomial number of bad states in smooth potential games.

5 Congestion Games Are Tight

Example 2.5 introduced congestion games, and considered the special case of affine resource cost functions. The worst-case POA in such games is 5/2. More generally, the POA in congestion games depends on the set of allowable cost functions. For example, with cost functions that are polynomials with degree at most d and nonnegative coefficients, the worst-case POA in congestion games is exponential in d, but independent of the network size and the number of players [2, 11, 33, 88]. This dependence motivates parameterizing POA bounds for congestion games via the class $\mathcal{C}$ of allowable resource cost functions. We do not expect the worst-case POA in congestion games to be expressible in closed form for every set $\mathcal{C}$, and instead seek a relatively simple characterization of this value, as a function of the set $\mathcal{C}$.

This section proves that, for every fixed set of nonnegative and nondecreasing cost functions $\mathcal{C}$, the set $\mathcal{G}(\mathcal{C})$ of congestion games with cost functions in $\mathcal{C}$ form a tight class of games. Recalling Definition 2.8 from Section 2.4, and the fact that every congestion game has at least one pure Nash equilibrium (see Rosenthal [91] or Example 4.11), this means that

$$\sup_{G \in \mathcal{G}(\mathcal{C})} \rho(G) = \inf_{(\lambda,\mu) \in \mathcal{A}(\mathcal{G}(\mathcal{C}))} \frac{\lambda}{1 - \mu}, \tag{34}$$

where $\mathcal{A}(\mathcal{G}(\mathcal{C}))$ is the set of parameter values (λ, μ) with $\mu < 1$ for which every game of $\mathcal{G}(\mathcal{C})$ is (λ, μ)-smooth. Combining this result with Theorem 3.2 shows that, for every set $\mathcal{C}$, the worst-case POA of games in $\mathcal{G}(\mathcal{C})$ is the same for each of the equilibrium concepts shown in Figure 1.

We proceed as follows. Section 5.1 uses the additive structure of congestion games to simplify the search for optimal smoothness parameters. Sections 5.2 and 5.3 form the heart of our argument. The former section shows that optimal smoothness parameters can generally be characterized as the unique intersection of two smoothness constraints. The latter section extracts the ingredients of a worst-case congestion game from these two constraints, and establishes tightness for finite sets of positive cost functions subject to a bounded load on every resource. Section 5.4 extends this tightness result to arbitrary sets of cost functions. Section 5.5 explains how our proof provides two characterizations of the worst-case POA in congestion games, as a function of the set $\mathcal{C}$ of allowable cost functions — one stemming from each side of (34).

5.1 Simplifying the Smoothness Constraints

We begin by simplifying the right-hand side of equation (34). We exploit the fact that, in a congestion game, the objective function and players' cost functions are additive over the resources E. This reduces the search for parameters (λ, μ) that satisfy condition (2) of Definition 2.1 — which imposes one constraint for every congestion game with cost functions in $\mathcal{C}$, and every pair $\mathbf{s}, \mathbf{s}^*$ of outcomes in that game — to a much simpler one.

Let $\mathcal{C}$ be a non-empty set of cost functions. As always, we assume that every function $c \in \mathcal{C}$ is nonnegative and nondecreasing. For convenience, we disallow the irrelevant all-zero cost function. Let $\mathcal{A}(\mathcal{C})$ denote the set of parameters (λ, μ) with $\mu < 1$ that satisfy

$$c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x \tag{35}$$

for every cost function $c \in \mathcal{C}$, non-negative integer x, and positive integer x^*. When $\mathcal{C}$ is the set of affine cost functions, the condition (35) specializes to the one (6) used in Example 2.5.

We define the value $\gamma(\mathcal{C})$ as the best POA bound that can be proved for games in $\mathcal{G}(\mathcal{C})$ via the condition (35). That is, define $\gamma(\mathcal{C})$ as

$$\inf \left\{ \frac{\lambda}{1-\mu} \, : \, (\lambda, \mu) \in \mathcal{A}(\mathcal{C}) \right\}, \tag{36}$$

with $\gamma(\mathcal{C}) = +\infty$ if $\mathcal{A}(\mathcal{C})$ is empty.

We have not constrained μ to be nonnegative. When $c(1) > 0$ for every cost function $c \in \mathcal{C}$, however, all points of $\mathcal{A}(\mathcal{C})$ have this property.

Lemma 5.1 (Nonnegativity of μ) *For every non-empty set $\mathcal{C}$ of strictly positive functions and every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$, $\mu \geq 0$.*

Proof: Taking $x = n$ and $x^* = 1$ in (35), with an arbitrary cost function $c \in \mathcal{C}$, shows that

$$\mu \geq \frac{c(n+1) - \lambda c(1)}{c(n)n} \geq \frac{1-\lambda}{n}, \tag{37}$$

where we have used that c is nondecreasing. Since n can be arbitrarily large, inequality (37) implies that $\mu \geq 0$ for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$. ■

We next show that the set $\mathcal{A}(\mathcal{C})$ can only be smaller than $\mathcal{A}(\mathcal{G}(\mathcal{C}))$ — that is, every game of $\mathcal{G}(\mathcal{C})$ is (λ, μ)-smooth for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$ — and hence the value $\gamma(\mathcal{C})$ is an upper bound on the worst-case robust POA of games in $\mathcal{G}(\mathcal{C})$. The proof is similar to the derivation in Example 2.5, and we record it for future reference.

Proposition 5.2 ($\gamma(\mathcal{C})$ Is an Upper Bound on the Robust POA) *For every non-empty set $\mathcal{C}$ of cost functions, the robust POA of every game of $\mathcal{G}(\mathcal{C})$ is at most $\gamma(\mathcal{C})$.*

Proof: We can assume that $\gamma(\mathcal{C})$ is finite. If $\mathcal{C}$ contains a cost function c that is not strictly positive — say $c(z) = 0$ and $c(z+1) > 0$ for some $z \geq 1$ — then taking $x = x^* = z$ in (35) shows that $\mathcal{A}(\mathcal{C}) = \emptyset$ and hence $\gamma(\mathcal{C}) = +\infty$. We can therefore assume that all functions of $\mathcal{C}$ are strictly positive.

We show that every game $G \in \mathcal{G}(\mathcal{C})$ is (λ, μ)-smooth for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$. Fix G and $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$; by Lemma 5.1, $\mu \geq 0$. For every outcome pair $\mathbf{s}, \mathbf{s}^*$ of G with induced load vectors $\mathbf{x}, \mathbf{x}^*$, we have

$$\begin{aligned}
\sum_{i=1}^{k} C_i(s_i^*, \mathbf{s}_{-i}) &\leq \sum_{e \in E \,:\, x_e^* \geq 1} c_e(x_e + 1) x_e^* && (38)\\
&\leq \sum_{e \in E \,:\, x_e^* \geq 1} [\lambda c_e(x_e^*) x_e^* + \mu c_e(x_e) x_e] && (39)\\
&\leq \sum_{e \in E} [\lambda c_e(x_e^*) x_e^* + \mu c_e(x_e) x_e] && (40)\\
&= \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}),
\end{aligned}$$

where in (38) we use that exactly x_e^* players ponder a deviation s_i^* that contains the resource e, which in turn is used by at most x_e players in $\mathbf{s}_{-i}$; inequality (39) follows from the definition of $\mathcal{A}(\mathcal{C})$; and inequality (40) follows from the nonnegativity of μ. ∎

5.2 Characterization of the Optimal Smoothness Parameters

The next step is to characterize the right-hand side of (34), meaning the parameters that minimize the objective function $\lambda/(1-\mu)$ over the feasible region $\mathcal{A}(\mathcal{C})$ defined in the preceding section. This optimization problem has several nice properties. First, there are only two decision variables — λ and μ — so $\mathcal{A}(\mathcal{C})$ is contained in the plane. Second, while there are an infinite number of constraints (35) that define this feasible region, each is linear in λ and μ. Thus, $\mathcal{A}(\mathcal{C})$ is the intersection of halfplanes. Third, the objective function $\lambda/(1-\mu)$ is decreasing in both decision variables. Thus, ignoring some edge cases that can be handled separately, the choice of (λ, μ) that minimizes the objective function lies on the "southwestern boundary" of $\mathcal{A}(\mathcal{C})$, and can be characterized as the unique point of $\mathcal{A}(\mathcal{C})$ that satisfies with equality a particular pair of constraints of the form (35).

Precisely, we assume in this section that the set $\mathcal{C}$ is finite, that every cost function is strictly positive, and that there is an upper bound n on the maximum load of a resource. Let $\mathcal{A}(\mathcal{C}, n)$ denote the set of parameters (λ, μ) with $\mu < 1$ that satisfy (35) for every $c \in \mathcal{C}$, $x \in \{0, 1, \ldots, n\}$, and $x^* \in \{1, 2, \ldots, n\}$. Geometrically, $\mathcal{A}(\mathcal{C}, n)$ is the intersection of a finite number of halfplanes, one open ($\mu < 1$), and each of the others a closed halfplane containing everything "northeast" of a line with negative slope. In contrast to Lemma 5.1, the set $\mathcal{A}(\mathcal{C}, n)$ can contain points (λ, μ) with

$\mu < 0$. See also Example 5.4 and Figure 2 for an illustration. The set $\mathcal{A}(\mathcal{C}, n)$ is non-empty because it includes the point $(\max_{c \in \mathcal{C}} \frac{c(n+1)}{c(1)}, 0)$. Define

$$\gamma(\mathcal{C}, n) = \inf \left\{ \frac{\lambda}{1-\mu} \ : \ (\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n) \right\}. \tag{41}$$

Because $\mathcal{A}(\mathcal{C}, n)$ is generally not bounded, this infimum need not be attained.

The following technical but key lemma asserts "first-order conditions" for optimal smoothness parameters: if the value $\gamma(\mathcal{C}, n)$ is attained by some point of $\mathcal{A}(\mathcal{C}, n)$, then this point is the intersection of two equality constraints of the form (35) that have particular properties. See also Example 5.4. These two constraints encode the cost functions, optimal resource loads, and equilibrium resource loads used in the lower-bound construction of Section 5.3.

Lemma 5.3 (Optimality Conditions) *Let $\mathcal{C}$ be a finite set of strictly positive cost functions and n a positive integer. Suppose there exist $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C}, n)$ such that*

$$\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n).$$

Then there exist $c_1, c_2 \in \mathcal{C}$, $x_1, x_2 \in \{0, 1, \ldots, n\}$, $x_1^*, x_2^* \in \{1, 2, \ldots, n\}$, and $\eta \in [0, 1]$ such that

$$c_j(x_j + 1)x_j^* = \hat{\lambda} \cdot c_j(x_j^*)x_j^* + \hat{\mu} \cdot c_j(x_j)x_j \tag{42}$$

for $j = 1, 2$; and

$$\eta \cdot c_1(x_1 + 1)x_1^* + (1-\eta) \cdot c_2(x_2 + 1)x_2^* = \eta \cdot c_1(x_1)x_1 + (1-\eta) \cdot c_2(x_2)x_2. \tag{43}$$

Proof: Write

$$\mathcal{H}_{c,x,x^*} = \{(\lambda, \mu) \ : \ c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x\}$$

for the halfplane corresponding to $c \in \mathcal{C}$, $x \in \{0, 1, \ldots, n\}$, and $x^* \in \{1, 2, \ldots, n\}$. Write $\partial\mathcal{H}_{c,x,x^*}$ for its boundary, meaning the points (λ, μ) that satisfy the inequality with equality. Define

$$\beta_{c,x,x^*} = \frac{c(x)x}{c(x+1)x^*}, \tag{44}$$

which is well defined because $x^* \geq 1$ and cost functions are strictly positive. If $x \geq 1$, then we can uniquely express λ in terms of μ along the line $\partial\mathcal{H}_{c,x,x^*}$ and derive

$$\frac{\lambda}{1-\mu} = \frac{c(x+1)}{c(x^*)} \frac{1 - (\beta_{c,x,x^*})\mu}{1-\mu} \tag{45}$$

for points of $\partial\mathcal{H}_{c,x,x^*}$. The only non-redundant constraint with $x = 0$ has the form $\mathcal{H}_{c,0,1} = \{(\lambda, \mu) :$ $\lambda \geq 1\}$. In this case, $\beta_{c,0,1} = 0$ and $\frac{\lambda}{1-\mu} = \frac{1}{1-\mu}$ for points in $\partial\mathcal{H}_{c,0,1}$. In any case, as λ increases and μ decreases along the line $\partial\mathcal{H}_{c,x,x^*}$, the value $\frac{\lambda}{1-\mu}$ is strictly decreasing if $\beta_{c,x,x^*} < 1$, strictly increasing if $\beta_{c,x,x^*} > 1$, and constant if $\beta_{c,x,x^*} = 1$. We accordingly call a line $\partial\mathcal{H}_{c,x,x^*}$ *decreasing*, *increasing*, or *constant*.

By assumption, there is a point $(\hat{\lambda}, \hat{\mu})$ that attains the infimum in (41). The optimality of $(\hat{\lambda}, \hat{\mu})$ has several implications. Since $\lambda/(1-\mu)$ is strictly decreasing in both λ and μ, $(\hat{\lambda}, \hat{\mu})$ inhabits the

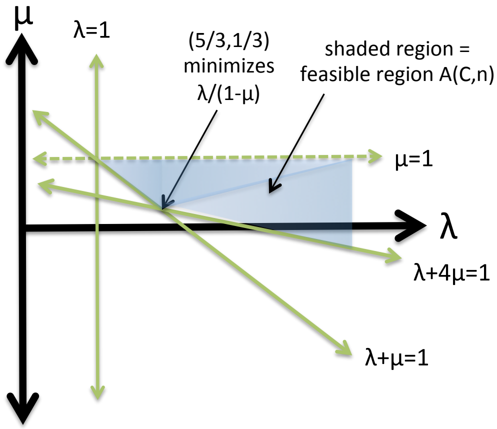

Figure 2: Example 5.4. The halfplanes that define the parameter $\gamma(\mathcal{C}, n)$.

boundary of $\mathcal{A}(C, n)$. In particular, it belongs to at least one line segment of the form $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$, and these choices of c, x, x^* satisfy equation (42). In the lucky event that $(\hat{\lambda}, \hat{\mu})$ is contained in a constant line $\partial\mathcal{H}_{c,x,x^*}$ — and thus $\beta_{c,x,x^*} = 1$ and $c(x+1)x^* = c(x)x$ — we can take $c_1 = c_2 = c$, $x_1 = x_2 = x$, $x_1^* = x_2^* = x^*$, and an arbitrary value of $\eta \in [0, 1]$ to satisfy (43).

If $(\hat{\lambda}, \hat{\mu})$ does not belong to a constant line, then it is an endpoint of the line segment $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$ — the endpoint with minimum μ-value if $\partial\mathcal{H}_{c,x,x^*}$ is decreasing, or with maximum μ-value if $\partial\mathcal{H}_{c,x,x^*}$ is increasing. Hence, $(\hat{\lambda}, \hat{\mu})$ is also an endpoint of a *second* boundary segment $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c',y,y^*}$, with $\partial\mathcal{H}_{c',y,y^*}$ increasing (decreasing) if $\partial\mathcal{H}_{c,x,x^*}$ is decreasing (increasing).

Relabel c, c', x, x^*, y, y^* so that $(\hat{\lambda}, \hat{\mu})$ is the endpoint of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c_1,x_1,x_1^*}$ with minimum μ-value and the endpoint of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c_2,x_2,x_2^*}$ with maximum μ-value. Equation (42) is satisfied. Since $\beta_{c_1,x_1,x_1^*} < 1$ and $\beta_{c_2,x_2,x_2^*} > 1$, $c_1(x_1 + 1)x_1^* > c_1(x_1)x_1$ while $c_2(x_2 + 1)x_2^* < c_2(x_2)x_2$. Choosing a suitable $\eta \in [0, 1]$ then satisfies equation (43). ■

Example 5.4 ($\mathcal{C} = \{c(x) = x\}$ and $n = 2$) Consider the special case in which $n = 2$ and $\mathcal{C}$ contains only the identity function $c(x) = x$. Not counting the constraint that $\mu < 1$, there are six constraints in the definition (36) of $\gamma(\mathcal{C}, n)$, corresponding to the two and three permitted values of x^* and x, respectively. Three of these are redundant, leaving the feasible choices of (λ, μ) constrained by the inequalities $\lambda \geq 1$, corresponding to $x = 0$, $x^* = 1$; $\lambda + \mu \geq 2$, corresponding to $x = x^* = 1$; and $\lambda + 4\mu \geq 3$, corresponding to $x = 2$ and $x^* = 1$. See Figure 2. Since $\beta_{c,1,1} < 1 < \beta_{c,1,2}$, the value $\gamma(\mathcal{C}, n)$ is attained at the intersection of the two corresponding lines, with $(\hat{\lambda}, \hat{\mu}) = (\frac{5}{3}, \frac{1}{3})$. Hence, $\gamma(\mathcal{C}, n) = \frac{5}{2}$.

The next lemma complements Lemma 5.3 by addressing cases where the infimum in (41) is not attained by any point of $\mathcal{A}(\mathcal{C}, n)$.[5]

[5]This can occur if, for example, $\mathcal{C}$ contains only a rapidly growing function like the factorial function.

Lemma 5.5 (Necessary Condition for $\gamma(\mathcal{C}, n)$ To Not Be Attained) *Let $\mathcal{C}$ be a finite set of strictly positive cost functions and n a positive integer. Suppose no point $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n)$ satisfies $\frac{\lambda}{1-\mu} = \gamma(\mathcal{C}, n)$. Then there exists $c \in \mathcal{C}$ such that*

$$\gamma(\mathcal{C}, n) = \frac{c(n)n}{c(1)} \tag{46}$$

and

$$c(n)n < c(n+1). \tag{47}$$

Proof: We use the notation and terminology from the proof of Lemma 5.3. The key point is to show that the infimum in (41) is not attained only when $\mathcal{A}(\mathcal{C}, n)$ has an unbounded boundary face $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$ for which $\partial\mathcal{H}_{c,x,x^*}$ is decreasing, meaning that the value β_{c,x,x^*} defined in (44) is less than 1.

Since $\mathcal{C}$ is a finite set of positive cost functions, $\mathcal{A}(\mathcal{C}, n)$ is non-empty and $\gamma(\mathcal{C}, n)$ is finite. Consider a sequence $\{(\lambda_k, \mu_k)\}$ in $\mathcal{A}(\mathcal{C}, n)$ with $\frac{\lambda_k}{1-\mu_k} \downarrow \gamma(\mathcal{C}, n)$. Since $\frac{\lambda}{1-\mu}$ is decreasing in both arguments, we can assume that every point (λ_k, μ_k) lies on the boundary of $\mathcal{A}(\mathcal{C}, n)$. Instantiating (35) with $x = 0$ and $x^* = 1$ proves that $\lambda_k \geq 1$ for every k. Since $\lim_{\mu \uparrow 1} \frac{1}{1-\mu} = +\infty$, we can assume that the μ_k's are bounded above by some $b < 1$. Since $\{(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n) : \mu \leq b\}$ is closed, $\frac{\lambda}{1-\mu}$ is continuous on this set, $\frac{\lambda_k}{1-\mu_k} \downarrow \gamma(C, n)$, and $\gamma(\mathcal{C}, n)$ is not attained, the sequence $\{(\lambda_k, \mu_k)\}$ has no limit point.

Every halfplane boundary $\partial\mathcal{H}_{c,x,x^*}$ has compact intersection with $\{(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n) : \mu \leq b\}$, except for a single constraint with a boundary with the least negative slope $-\frac{c(x^*)x^*}{c(x)x}$. Since cost functions are positive and nondecreasing, this constraint boundary necessarily has the form $\partial\mathcal{H}_{c,n,1}$ for some $c \in \mathcal{C}$. We complete the proof by showing that this cost function c satisfies (46) and (47).

Since $\{(\lambda_k, \mu_k)\}$ has no limit point and there are only finitely many constraints, this sequence is eventually contained in $\partial\mathcal{H}_{c,n,1} \cap \mathcal{A}(\mathcal{C}, n)$, with $\lambda_k \to +\infty$ and $\mu_k \to -\infty$ as $k \to +\infty$. Using equation (45), we have

$$\gamma(\mathcal{C}, n) = \lim_{k \to \infty} \frac{\lambda_k}{1 - \mu_k} = \beta_{c,n,1} \cdot \frac{c(n+1)}{c(1)} = \frac{c(n)n}{c(1)},$$

which verifies condition (46). Moreover, since $\frac{\lambda_k}{1-\mu_k}$ is strictly decreasing in k, equation (45) implies that $\beta_{c,n,1} < 1$ or, equivalently, $c(n)n < c(n+1)$. This verifies condition (47) and completes the proof. ■

5.3 Lower Bound Construction: The Finite Case

We now present the main lower bound construction. In this section, we continue to assume that $\mathcal{C}$ is a finite set of strictly positive cost functions, and that there is an upper bound on the maximum load of a resource. The next section treats the general case.

Constructing a congestion game with POA equal to $\gamma(\mathcal{C}, n)$ is tantamount to constructing a game in which the inequalities in (38)–(40) hold with equality. The plan is to construct a congestion game in which each player has two strategies, one that uses a small number of resources, and a disjoint strategy that uses a large number of resources. In the optimal outcome, all players use their small strategies and incur low cost. This outcome is also a pure Nash equilibrium. In the suboptimal pure Nash equilibrium, all players use their large strategies, thereby "flooding" all resources and incurring a large cost. How can this suboptimal outcome persist as a Nash equilibrium? If a player deviates unilaterally, it enjoys the benefit of fewer resources in its strategy, but the load on each of the new resources is one more than that on the previously used resources. We show that, implemented optimally, this construction produces a congestion game and a pure Nash equilibrium of it with cost $\frac{\hat{\lambda}}{1-\hat{\mu}}$ times that of the optimal outcome, where $(\hat{\lambda}, \hat{\mu})$ are the optimal smoothness parameters identified in the previous section. See also Example 5.7.

Theorem 5.6 (Main Construction) *Let $\mathcal{C}$ be a non-empty finite set of strictly positive cost functions and n a positive integer. There exist congestion games with cost functions in $\mathcal{C}$ and (pure) POA arbitrarily close to $\gamma(\mathcal{C}, n)$.*

Proof: We first dispense with the case in which the value $\gamma(\mathcal{C}, n)$ is not attained by any point $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n)$. Let $c \in \mathcal{C}$ be the cost function guaranteed by Lemma 5.5, which satisfies conditions (46) and (47). Define a congestion game as follows. Let $E = \{e_1, \ldots, e_{n+1}\}$ and introduce $n+1$ players, where player i's two strategies are $\{e_i\}$ and $E \setminus \{e_i\}$. If players choose their singleton strategies, then the resulting outcome has cost $(n+1) \cdot c(1)$. If players choose their non-singleton strategies, then every resource is used by all by one player and the cost is $(n+1) \cdot c(n)n$. Condition (47) implies that the latter outcome is a Nash equilibrium. By condition (46), the POA of this game is at least $c(n)n/c(1) = \gamma(\mathcal{C}, n)$, as required.

For the rest of the proof, we assume that there is a point $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C}, n)$ with $\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n)$. Choose c_1, c_2, x_1, x_2, x_1^*, x_2^*, η as in Lemma 5.3. Define a congestion game as follows. The ground set $E_1 \cup E_2$ should be thought of as two disjoint "cycles," where each cycle has $k = \max\{x_1 + x_1^*, x_2 + x_2^*\}$ resources that are labeled from 1 to k. Resources from E_1 and E_2 are each given the cost function $\eta \cdot c_1(x)$ and $(1-\eta) \cdot c_2(x)$, respectively. There are also k players, each with two strategies. Player i's first strategy P_i uses precisely x_j consecutive resources of E_j (for $j = 1, 2$), starting with the ith resource of each cycle and wrapping around to the beginning, if necessary. Player i's second strategy Q_i uses x_j^* consecutive resources of E_j (for $j = 1, 2$), ending with the $(i-1)$th resource of each cycle and wrapping around from the end, if necessary. We have chosen k large enough that, for each i, the strategies P_i and Q_i are disjoint.

Let $\mathbf{y}$ and $\mathbf{y}^*$ denote the outcomes in which each player selects the strategy P_i and Q_i, respectively. By symmetry, $y_e = x_1$ and $y_e^* = x_1^*$ for resources $e \in E_1$, while $y_e = x_2$ and $y_e^* = x_2^*$ for resources $e \in E_2$. Thus, for example, the value x_1 serves both as the cardinality of every set $P_i \cap E_1$, and as the load y_e of every resource $e \in E_1$ in the outcome $\mathbf{y}$.

To verify that $\mathbf{y}$ is a pure Nash equilibrium, fix a player i and derive

$$\begin{aligned} C_i(\mathbf{y}) &= \sum_{e \in P_i \cap E_1} \eta \cdot c_1(y_e) + \sum_{e \in P_1 \cap E_2} (1-\eta) \cdot c_2(y_e) \\ &= \eta \cdot c_1(x_1)x_1 + (1-\eta) \cdot c_2(x_2)x_2 \\ &= \eta \cdot c_1(x_1+1)x_1^* + (1-\eta) \cdot c_2(x_2+1)x_2^* \qquad (48) \\ &= \sum_{e \in Q_i \cap E_1} \eta \cdot c_1(y_e+1) + \sum_{e \in Q_i \cap E_2} (1-\eta) \cdot c_2(y_e+1) \\ &= C_i(y_i^*, \mathbf{y}_{-i}), \qquad (49) \end{aligned}$$

where equation (48) follows from requirement (43) in Lemma 5.3, and equation (49) follows from the disjointness of P_i and Q_i.

Moreover, using (48) as a launching pad, we can derive

$$
\begin{aligned}
C(\mathbf{y}) &= \sum_{i=1}^{k} C_i(\mathbf{y}) \\
&= k \cdot [\eta \cdot c_1(x_1+1)x_1^* + (1-\eta) \cdot c_2(x_2+1)x_2^*] \\
&= k\eta \cdot \left(\hat{\lambda} \cdot c_1(x_1^*)x_1^* + \hat{\mu} \cdot c_1(x_1)x_1\right) + k(1-\eta) \cdot \left(\hat{\lambda} \cdot c_2(x_2^*)x_2^* + \hat{\mu} \cdot c_2(x_2)x_2\right) \qquad (50) \\
&= \hat{\lambda} \cdot k \cdot (\eta \cdot c_1(x_1^*)x_1^* + (1-\eta) \cdot c_2(x_2^*)x_2^*) + \hat{\mu} \cdot k \cdot (\eta \cdot c_1(x_1)x_1 + (1-\eta) \cdot c_2(x_2)x_2) \\
&= \hat{\lambda} \cdot C(\mathbf{y}^*) + \hat{\mu} \cdot C(\mathbf{y}),
\end{aligned}
$$

where (50) follows from condition (42) in Lemma 5.3. Rearranging gives a lower bound of

$$\frac{C(\mathbf{y})}{C(\mathbf{y}^*)} = \frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n)$$

on the POA of this congestion game.

This construction uses resource cost functions $\eta \cdot c_1(x)$ and $(1-\eta) \cdot c_2(x)$ that are scalar multiples of cost functions c_1, c_2 that lie in the given set $\mathcal{C}$. The construction can be extended to use only the original cost functions c_1, c_2 via standard scaling and replication tricks, as in [93]. In more detail, η and $(1-\eta)$ can be approximated by nonnegative rational numbers so that the outcome $\mathbf{y}$ remains a Nash equilibrium and the POA goes down by an arbitrarily small amount. Then, scaling all cost functions so that η and $(1-\eta)$ are integers does not change the POA of the constructed game. Finally, replacing each resource with cost function $m \cdot c(x)$ by a set of m resources, each with cost function $c(x)$, and modifying players' strategy sets accordingly, does not change the POA. This completes the proof. ■

Example 5.7 ($\mathcal{C} = \{c(x) = x\}$ and $n = 2$, Continued) The proof of Theorem 5.6, specialized to Example 5.4, regenerates a construction by Christodoulou and Koutsoupias [33, Theorem 2] that gives a matching lower bound on the POA of pure Nash equilibria. Awerbuch et al. [11, Theorem 3.5] independently gave a slightly different construction that provides the same lower bound. First, recall the relevant two halfplanes from Example 5.4, with $x = x^* = 1$ and with $x = 2$, $x^* = 1$. The unique value of η that satisfies condition (43) of Lemma 5.3 is $\frac{1}{2}$. Define a congestion game with three players $0, 1, 2$ and six resources $u_0, u_1, u_2, v_0, v_1, v_2$, all with the cost function $c(x) = x/2$. (Using the function $c(x) = x$ instead yields an equivalent example.) Player i has two strategies, $\{u_i, v_i\}$ and $\{u_{i+1}, v_{i+1}, v_{i+2}\}$, where all arithmetic is modulo 3. See also Figure 3. If all players use their smaller strategies, each incurs cost 1. If all players use their larger strategies, each incurs cost $\frac{1}{2} + 2 \cdot \frac{2}{2} = \frac{5}{2}$. Since a unilateral deviation would also yield cost $\frac{2}{2} + \frac{3}{2} = \frac{5}{2}$ to the deviator, this outcome is a pure Nash equilibrium. The POA in this game is at least 5/2.

5.4 Lower Bound Construction: The General Case

We now generalize Theorem 5.6 to arbitrary sets of nondecreasing cost functions and arbitrarily many players.

Theorem 5.8 (Tightness of Congestion Games) *For every non-empty set $\mathcal{C}$ of cost functions, the set of congestion games with cost functions in $\mathcal{C}$ is tight.*

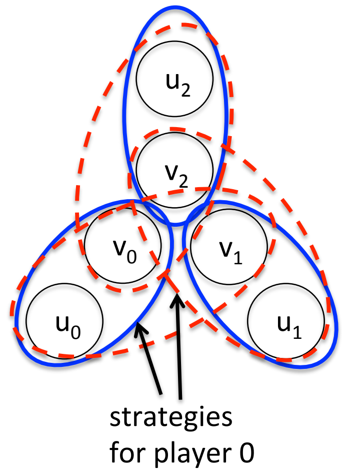

Figure 3: Example 5.7. There are six resources, each with cost function $c(x) = x/2$. Solid and dashed ovals denote the strategies used in the optimal and bad Nash equilibrium outcomes, respectively. The two strategies of a player contain disjoint sets of resources.

Proof: We first observe that if $\mathcal{C}$ contains a cost function that is not strictly positive, then there is a congestion game with cost functions in $\mathcal{C}$ and infinite POA. Suppose $c \in \mathcal{C}$ satisfies $c(z) = 0$ and $c(z+1) > 0$, with $z \geq 1$. Mimic the main construction in the proof of Theorem 5.6 with $c_1 = c_2 = c$, $\eta = \frac{1}{2}$, $x_1 = x_1^* = x_2^* = z$, and $x_2 = z + 1$. The outcome $\mathbf{y}^*$ has cost $kc(z)z = 0$. The outcome $\mathbf{y}$ has cost $\frac{k}{2}c(z)z + \frac{k}{2}c(z+1)(z+1) > 0$. Since

$$C_i(y_i^*, \mathbf{y}_{-i}) = \tfrac{1}{2}(zc(z+1) + zc(z+2)) \geq \tfrac{1}{2}(z+1)c(z+1) = C_i(\mathbf{y}),$$

the outcome $\mathbf{y}$ is a Nash equilibrium. The POA of this congestion game is infinite.

For the rest of the proof, we can assume that $\mathcal{C}$ contains only strictly positive cost functions. For now, we assume that $\mathcal{C}$ is countable. Order the cost functions of $\mathcal{C}$ and let $\mathcal{C}_n$ denote the first n functions. Theorem 5.6 applies to each (non-empty) set $\mathcal{A}(\mathcal{C}_n, n)$.

We now proceed in several steps. The first four steps address cases in which the worst-case POA of congestion games with cost functions in $\mathcal{C}$ is unbounded (and $\gamma(\mathcal{C}) = +\infty$). The final two steps handle sets of cost functions for which this worst-case POA is finite.

1. Suppose that, for infinitely many n, $\gamma(\mathcal{C}_n, n)$ is not attained by a point of $\mathcal{A}(\mathcal{C}_n, n)$. Then, Theorem 5.6 and condition (46) imply that there are congestion games with cost functions in $\mathcal{C}$ and arbitrarily large POA. We can therefore assume that, for all sufficiently large n, there is a pair $(\lambda_n, \mu_n) \in \mathcal{A}(\mathcal{C}_n, n)$ with $\frac{\lambda_n}{1-\mu_n} = \gamma(\mathcal{C}_n, n)$.

2. Instantiating constraint (35) with an arbitrary cost function $c \in \mathcal{C}_n$, $x = n$, and $x^* = 1$ shows that

$$\lambda \geq \frac{c(n+1) - \mu c(n)n}{c(1)} \geq 1 - \mu n$$

for all $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}_n, n)$. Thus, $\frac{\lambda_n}{1-\mu_n} > \frac{n}{2}$ whenever $\mu_n < -1$. By Theorem 5.6, we can assume henceforth that $\mu_n \geq -1$ for all sufficiently large n.

3. If λ_n grows without bound and $\mu_n \geq -1$ for all sufficiently large n, then Theorem 5.6 provides a sequence of congestion games with arbitrarily large POA. We can therefore assume that, for some constant M, $\lambda_n \leq M$ for all sufficiently large n.

4. Recall that $\lambda \geq 1$ for every point in $\mathcal{A}(\mathcal{C}, n)$. Thus, for all sufficiently large n, (λ_n, μ_n) lies in the compact set $[1, M] \times [-1, 1]$. Consider a convergent subsequence, with limit (λ^*, μ^*). If $\mu^* = 1$, then there is an infinite sequence of values (λ_n, μ_n) with $\lambda_n \geq 1$ and $\mu_n \to 1$. Applying Theorem 5.6 to this sequence yields a sequence of congestion games with arbitrarily large POA.

5. If $\mu^* < 1$, then, by continuity of the function $\lambda/(1-\mu)$ on $[1, M] \times [-1, \mu^*]$, Theorem 5.6 yields a sequence of congestion games with POA arbitrarily close to $\frac{\lambda^*}{1-\mu^*}$.

6. To complete the proof for countable sets of cost functions, we claim that $(\lambda^*, \mu^*) \in \mathcal{A}(\mathcal{C})$ when $\mu^* < 1$, and hence $\gamma(\mathcal{C}) \leq \frac{\lambda^*}{1-\mu^*}$. For if not, there is a cost function $c \in \mathcal{C}$, a nonnegative integer x, and a positive integer x^* with $c(x+1)x^* > \lambda^* c(x^*)x^* + \mu^* c(x)x$. By continuity, $c(x+1)x^* > \lambda_n c(x^*)x^* + \mu_n c(x)x$ for all sufficiently large n. But $c \in \mathcal{C}_n$ and $x, x^* \leq n$ eventually, contradicting the fact that $(\lambda_n, \mu_n) \in \mathcal{A}(\mathcal{C}, n)$ for every n.

Finally, we reduce general sets of (strictly positive) cost functions to the case of countable sets via a density argument. For a rational number $r \in (0, 1)$, we say that a cost function $\tilde{c}$ is a *rational r-approximation* of a strictly positive cost function c if can be constructed by the following inductive process: $\tilde{c}(1)$ is a rational number in the open interval $(c(1), (1+r)c(1))$; and for $x = 2, 3, 4, \ldots$, $\tilde{c}(x)$ is a rational number that is at least $\tilde{c}(x-1)$, strictly greater than $c(x)$, and strictly less than $\frac{\tilde{c}(x-1)}{c(x-1)} c(x)$. The purpose of rational r-approximations is to approximate cost functions in a manner that preserves certain pure Nash equilibria. The next four properties follow from the definition of such approximations.

1. For every strictly positive and nondecreasing cost function c and rational number $r > 0$, a rational r-approximation of c can be constructed by induction.

2. Every rational r-approximation of a strictly positive and nondecreasing cost function is again strictly positive and nondecreasing.

3. Every rational r-approximation has a countable domain and range.

4. For every rational r-approximation $\tilde{c}$ of a cost function c and positive integer x,

$$1 < \frac{\tilde{c}(x+1)}{c(x+1)} < \frac{\tilde{c}(x)}{c(x)} < 1 + r.$$

Fix an arbitrary set $\mathcal{C}$ of strictly positive cost functions, and a parameter $\delta > 0$. Let $\mathcal{C}_\delta$ denote the set of rational r-approximations of cost functions in $\mathcal{C}$ for all rational $r \in (0, \delta)$. By the third property of rational r-approximations, the set $\mathcal{C}_\delta$ is countable. Thus, there is a sequence $\widetilde{G}_1, \widetilde{G}_2, \ldots$ of congestion games with cost functions in $\mathcal{C}_\delta$ and POA arbitrarily close to $\gamma(\mathcal{C}_\delta)$. We can assume that each game $\widetilde{G}_i$ was produced by one of the two constructions in the proof of Theorem 5.6.

Consider the congestion game $\widetilde{G}_i$ and let $\mathbf{y}$ and $\mathbf{y}^*$ denote its bad Nash equilibrium and optimal outcome, respectively, as constructed in the proof of Theorem 5.6. Obtain the game G_i from $\widetilde{G}_i$

by replacing each resource cost function $\tilde{c}$ by a cost function $c \in \mathcal{C}$ for which $\tilde{c}$ is a rational r-approximation with $r < \delta$. The fourth property of rational r-approximations implies that: the cost of $\mathbf{y}$ in G_i is at least $\frac{1}{1+\delta}$ times that in $\widetilde{G}_i$; the cost of $\mathbf{y}^*$ in G_i is at most that in $\widetilde{G}_i$; the equilibrium condition that holds for $\mathbf{y}$ in $\widetilde{G}_i$ — condition (47) or (49), depending on the construction — continues to hold, as an inequality, for $\mathbf{y}$ in G_i. Thus, $\mathbf{y}$ is a Nash equilibrium of G_i and the POA of the game is at least $\frac{1}{1+\delta}$ times that of $\widetilde{G}_i$. This shows that there are congestion games with cost functions in $\mathcal{C}$ and POA arbitrarily close to $\frac{\gamma(\mathcal{C}_\delta)}{1+\delta}$, where $\delta > 0$ is an arbitrarily small rational number.

We complete the proof by arguing that $\gamma(\mathcal{C}_\delta) \geq \gamma(\mathcal{C})$ for every rational number $\delta > 0$. Choose $c \in \mathcal{C}$, a nonnegative integer x, and a positive integer x^*. By continuity, a pair (λ, μ) satisfies the corresponding constraint (35) only if it satisfies the analogous constraints for the same integers x, x^* and for every rational r-approximation $\tilde{c}$ of c for every rational number $r \in (0, \delta)$. That is, $\mathcal{A}(\mathcal{C}_\delta) \subseteq A(\mathcal{C})$ and hence $\gamma(C_\delta) \geq \gamma(\mathcal{C})$. ■

5.5 Universal Proofs and Universal Worst-Case Examples

Our proof of Theorem 5.8 gives two characterizations of the worst-case POA of congestion games with cost function restricted to a set $\mathcal{C}$. These give, in particular, the first bounds on the POA in congestion games with non-polynomial cost functions.

First, Theorem 5.8 shows that the worst-case POA in congestion games with cost functions in $\mathcal{C}$ is precisely $\gamma(\mathcal{C})$, the best POA upper bound provable using the smoothness arguments generated by Proposition 5.2. Of course, computing the exact value of $\gamma(\mathcal{C})$ is not trivial, even for simple sets $\mathcal{C}$. Aland et al. [2] and Olver [88] provide a (complex) closed-form expression for $\gamma(\mathcal{C})$ when $\mathcal{C}$ is a set of polynomials with nonnegative coefficients. Similar computations should be possible for some other simple sets $\mathcal{C}$. More broadly, the proof of Lemma 5.3 indicates how to compute good lower bounds on $\gamma(\mathcal{C})$ when there is a particular set $\mathcal{C}$ of interest. Computing the exact value of $\gamma(\mathcal{C}, n)$ for a finite set $\mathcal{C}$ reduces to computing the upper envelope of $O(n^2|\mathcal{C}|)$ lines. Taking n sufficiently large and a sufficiently representative finite subset of cost functions should typically permit close approximation of $\gamma(\mathcal{C})$, the worst-case POA of congestion games with cost functions in $\mathcal{C}$.

Second, Theorem 5.8 provides a characterization of the worst-case POA in congestion games via a class of "universal" worst-case examples. The proof of Theorem 5.8 only requires congestion games that comprise a "cycle" or the product of two cycles of the same size (a "double-cycle"), where every resource on the same cycle has the same cost function. These congestion games can be realized as network congestion games using a bidirected cycle network, provided zero-cost edges are also permitted; see Gairing [53, Figure 5.2]. Thus, the worst-case POA of congestion games with cost functions in $\mathcal{C}$ equals the worst-case POA of such double-cycles. This observation is analogous to a simpler such characterization in *nonatomic* congestion games —- in which there is a continuum of players, each of negligible size — where under modest assumptions on $\mathcal{C}$, the worst-case POA is always achieved in two-node two-link networks [39, 99].

Remark 5.9 (Necessity of Double-Cycles) The proof of Theorem 5.8 shows that double-cycle congestion games are universal worst-case examples for the POA. Could there be a simpler set of universal worst-case examples? In general, the answer is "no." The basic reason is that, in light of Theorem 5.8, a congestion game is a worst-case example *only if* all of the inequalities in Proposition 5.2 hold with equality.

For example, consider a finite set $\mathcal{C}$ of cost functions and a value of n so that, in the language of the proof of Lemma 5.3, $\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n)$ for some $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C}, n)$. If the halfplanes defining $\mathcal{A}(\mathcal{C}, n)$ are in general position, then $(\hat{\lambda}, \hat{\mu})$ satisfies at most two of them with equality. Call these hyperplanes $\mathcal{H}_{c_1,x_1,x_1^*}$ and $\mathcal{H}_{c_2,x_2,x_2^*}$. Inequality (39) holds with equality only if, for every resource e on which the worst pure Nash equilibrium or optimal outcome incurs non-zero cost, there is an $i \in \{1, 2\}$ such that $c_e = c_i$ and the equilibrium and optimal loads on e are x_i and x_i^*, respectively. If the value β_{c_i,x_i,x_i^*} in (44) equals 1 for one of $i = 1, 2$, then there is a single-cycle congestion game that uses only the cost function c_i and realizes the worst-case POA. The construction is similar to the first one in the proof of Theorem 5.6.

In general, however, $\beta_{c_i,x_i,x_i^*} \neq 1$ for $i = 1, 2$ (cf., Example 5.4). Here, a single-cycle construction — with cost function c_i, equilibrium load x_i, and optimal load x_i^* on every resource — does not work. If $\beta_{c_i,x_i,x_i^*} < 1$, then this construction does not satisfy inequality (38) with equality. If $\beta_{c_i,x_i,x_i^*} > 1$, then it does not satisfy inequality (38) at all, and hence fails to produce the bad Nash equilibrium. Thus, two groups of resources with distinct cost functions are generally necessary to construct congestion games that realize the worst-case POA.

6 Related Work

Sections 6.1 and 6.2 discuss related work published prior and subsequent to the conference version of this paper [95], respectively.

6.1 Previous Work

The Price of Anarchy. The price of anarchy was first studied by Koutsoupias and Papadimitriou [74] for the makespan minimization objective in scheduling games. This is not a sum objective function, and the worst-case POA in this model was immediately recognized to be different for different equilibrium concepts [20, 41, 73, 74]. See Vöcking [106] for a survey of the literature on this model, and Heydenreich et al. [67] for results on the POA in other scheduling models.

The POA with a sum objective was first studied by Roughgarden and Tardos [98] in the nonatomic selfish routing games of Wardrop [107], which are discussed further below. Most work to date on the POA concerns sum objective functions. Extensively studied problem domains include congestion games and their variants (see Roughgarden [94] and below), network design and utility games (see Tardos and Wexler [104]), and auction games (see Section 6.2).

The POA in Congestion Games. The first general results on the POA of pure Nash equilibria for the congestion games of Rosenthal [91] and their weighted variants [81] are by Awerbuch et al. [11] and Christodoulou and Koutsoupias [33], who independently gave tight bounds for games with affine cost functions and qualitatively similar upper and lower bounds for games with polynomial cost functions with nonnegative coefficients. Subclasses of congestion games with affine cost functions were studied earlier by Lücking et al. [78] and Suri et al. [101]. Matching upper and lower bounds for both unweighted and weighted congestion games with polynomial cost functions with nonnegative coefficients were given by Aland et al. [2] and, subsequently but independently, Olver [88]. The lower bound construction in Section 5 generalizes those in [2, 88].

The POA of Mixed Nash, Correlated, and Coarse Correlated Equilibria. The importance of POA bounds that apply beyond Nash equilibria was first articulated by Mirrokni and Vetta [80]. Most of the previous works that established more general POA bounds relied on arguments that can be recast as smoothness proofs (as in Examples 2.5–2.7), so our extension theorems immediately imply, and often strengthen, these previously proved bounds.

In more detail, the authors in [2, 11, 105] extend most of their upper bounds on the worst-case POA of pure Nash equilibria in congestion or valid utility games to mixed Nash equilibria. Christodoulou and Koutsoupias [32] show that the worst-case POA of correlated equilibria is the same as for pure Nash equilibria in unweighted and weighted congestion games with affine cost functions. Blum et al. [20] rework and generalize several bounds on the worst-case POA of pure Nash equilibria to show that the same bounds hold for the average objective function value of no-regret sequences. Their applications include valid utility games [105] and the (suboptimal) bounds of [11, 33] for unweighted congestion games with polynomial cost functions, and also a constant-sum location game and a fairness objective, which falls outside of our framework.

Another line of work identifies games in which POA bounds for pure Nash equilibria *cannot* be extended to more general equilibrium concepts. Recall from Example 3.1 that all of the inclusions in Figure 1 can be strict, even in simple games such as congestion games with affine cost functions and symmetric singleton strategies (i.e., networks of parallel links). More generally, there are specific such games in which the worst-case cost of an equilibrium is different for each of the four equilibrium concepts shown in Figure 1; see Ashlagi et al. [8], Blum et al. [20], Bradonjić et al. [22], and Kleinberg et al. [71] for some concrete examples. These examples do not contradict our tightness result for congestion games (Theorem 5.8), which only precludes such separations for *worst-case* congestion games for a given set of cost functions.

There are also interesting classes of games for which, in contrast to congestion games, the worst-case POA is different for different equilibrium concepts. In addition to the two examples in Section 2.5 and the makespan minimization games mentioned above, where there is already a gap between the worst-case POA of pure and mixed Nash equilibria, recent work by Roughgarden and Schoppmann [97] and Bhawalkar et al. [15] identified natural classes of games with a sum objective in which the worst-case POA is the same for pure Nash, mixed Nash, and correlated equilibria, but is strictly larger for coarse correlated equilibria.

Nonatomic Congestion Games. Previous work that bounds the POA in *nonatomic* congestion games, where there is a fixed mass of infinitesimal players, can be viewed as a precursor to and special case of our smoothness framework. These games were introduced by Roughgarden and Tardos [99], motivated by the nonatomic routing games of Wardrop [107] and Beckmann et al. [13] and the congestion games of Rosenthal [91]. Tight POA bounds for almost all classes of cost functions are also in [99]. The relationship between POA bounds for these games and smoothness arguments is clearest in the work of Correa et al. [40].

The POA in nonatomic congestion games can be bounded above using the same ideas as in Example 2.5 and Section 5. In this sense, the optimal POA analyses for nonatomic and atomic congestion games follow the same template. We next highlight the main differences.

For a set $\mathcal{C}$ of allowable cost functions, define the set $\mathcal{A}(\mathcal{C})$ as the parameters (λ, μ) with $\mu < 1$ that satisfy

$$c(x)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x \tag{51}$$

for every cost function $c \in \mathcal{C}$ and nonnegative real numbers x^*, x. The difference between (51) and its atomic analog (35) is that the "+1" on the left-hand side has disappeared. This reflects the negligible size of every player. Also, x and x^* are no longer constrained to be integral. Analogous to Proposition 5.2, for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$, the POA of every nonatomic congestion game with cost functions in $\mathcal{C}$ is at most $\frac{\lambda}{1-\mu}$.

A remarkable property of nonatomic congestion games is that, assuming only that the set $\mathcal{C}$ contains at least one function that is positive at zero, there is a pair $(\hat{\lambda}, \hat{\mu}) \in \mathcal{A}(\mathcal{C})$ of optimal smoothness parameters with $\hat{\lambda} = 1$ [39, 99]. For instance, if $\mathcal{C}$ is the set of affine cost functions, then the optimal smoothness parameters are $\lambda = 1$ and $\mu = \frac{1}{4}$, leading to a POA upper bound of $\frac{4}{3}$. For this reason, early work on the POA in nonatomic congestion games did not suggest the two-parameter smoothness framework developed in this paper.

A second special property of nonatomic congestion games is that equilibria are essentially unique. Precisely, it follows from Blum et al. [19] that all coarse correlated equilibria of a nonatomic congestion game have the same cost. This renders our extension theorem (Theorem 3.2) superfluous for such games — bounds on the POA of pure Nash equilibria automatically extend to the other equilibrium sets in Figure 1, simply because all of the sets are essentially identical.

Further Precursors of Smoothness. Several previous papers gave special cases of our two-parameter smoothness definition, in each case for a specific model and without any general applications to robust POA guarantees: Perakis [90] for a nonatomic selfish routing model with non-separable cost functions; Christodoulou and Koutsoupias [32] for congestion games with affine cost functions; Aland et al. [2] for congestion games with polynomial cost functions and nonnegative coefficients; Harks [62] for splittable congestion games; and Gairing et al. [56], Dumrauf and Gairing [43], Harks and Végh [64], and Farzad et al. [49] for different variants of selfish routing games. Awerbuch et al. [10] defined a "β-nice" condition for potential games that is essentially a single-parameter instantiation of our smoothness definition, and gave general guarantees for the performance of BRD in such games (cf., Section 4.3). The parameter β in [10] corresponds to the value $\frac{\lambda}{1-\mu}$ in our framework.

The POA of Approximate Nash Equilibria. The POA of approximate Nash equilibria was first studied by Roughgarden and Tardos [98] and Weitz [108] in nonatomic selfish routing games. Vetta [105, Theorem 8.1] extended his POA analysis of valid utility games to approximate Nash equilibria (cf., Remark 4.4). Christodoulou et al. [35] compute precisely the POA of approximate Nash equilibria in both nonatomic and atomic congestion games with cost functions that are polynomials with nonnegative coefficients.

Bicriteria Bounds. Bicriteria bounds like our Theorems 4.6 and 4.9 are sometimes called "resource augmentation bounds" or "pseudo-approximations." Roughgarden and Tardos [98] proved the first such bound for Nash equilibria: in every nonatomic selfish routing game, the cost of an equilibrium is bounded above by that of an optimal solution that routes twice as much traffic. This bound is tight for general edge cost functions, but Chakrabarty [27] and Correa et al. [40] showed how to improve it for restricted sets of cost functions. The only previous results on bicriteria bounds in atomic congestion games were negative [98].

Guarantees for Best-Response Dynamics. Mirrokni and Vetta [80] were the first to study the performance of best-response dynamics (BRD), and they gave both positive and negative results for valid utility games (Example 2.6). Goemans et al. [60] refined some of the results in [80] and also proved that, in weighted congestion games, random BRD quickly reaches states that approximately obey the robust POA bounds proved in [2, 88]. Awerbuch et al. [10] identified general conditions under which several variants of BRD quickly reach states that approximately satisfy the corresponding robust POA bounds. Our development in Section 4.3 subsumes several of the positive results in [10, 60, 80]. A notable exception is the upper bounds in [60] on the "price of sinking" in weighted congestion games, which generally do not possess the potential function required by Theorems 4.13 and 4.14 (see [63]). For recent work that derives quantitative trade-offs between the number of BRD iterations in congestion games with affine cost functions and the approximation factor of the state reached, see Fanelli et al. [48], Bilò et al. [17], and the references therein.

Guarantees with Irrational Players. Still another way to weaken the rationality assumptions implicit in equilibrium analysis is to allow some players to behave irrationally. Ideally, the POA of a game degrades gracefully with the fraction of irrational players.

Karakostas and Viglas [70] were the first to give POA bounds with irrational players, in nonatomic selfish routing networks with some "malicious" traffic. Babaioff et al. [12] and Gairing [54] analyzed other models of malicious players in routing games, while Moscibroda et al. [82] studied similar issues in inoculation games. Blum et al. [20] and Roth [92] proposed a more general model in which irrational players play arbitrary strategies, rather than optimizing a specific payoff function meant to model malice, and gave POA bounds with such players for classes of congestion games (Example 2.5), valid games (Example 2.6), and location games.

The smoothness framework of this paper can be used to derive some of these POA bounds with irrational players. For example, consider a valid utility game G with a nondecreasing objective function, rational players R, and irrational players I. Let $\mathbf{s}^I$ denote an arbitrary strategy profile for the players of I and $\mathbf{s}^R$ an "induced Nash equilibrium," meaning that no player of R can increase its payoff via a unilateral deviation from $(\mathbf{s}^R, \mathbf{s}^I)$. Blum et al. [20] proved that the welfare $W^G(\mathbf{s}^R, \mathbf{s}^I)$ of such an induced Nash equilibrium of G is at least half of the welfare $W^{\widehat{G}}(\mathbf{s}^*)$ of an optimal solution $\mathbf{s}^*$ in the game $\widehat{G}$ induced by the rational players R. This guarantee also follows directly from Definition 4.8, which by (8)–(10) G satisfies with $\lambda = \mu = 1$. Precisely, we have

$$W^G(\mathbf{s}^R, \mathbf{s}^B) \geq \sum_{i \in R} \Pi_i^G(\mathbf{s}^R, \mathbf{s}^B) \geq \sum_{i \in R} \Pi_i^G(s_i^*, \mathbf{s}_{-i}^R, \mathbf{s}^B) \geq W^{\widehat{G}}(\mathbf{s}^*) - W^G(\mathbf{s}^R, \mathbf{s}^B),$$

which implies the claimed bound. Following the proof of Theorem 3.3 extends this guarantee to the average payoff of rational players over repeated play, with irrational players playing arbitrarily at each time step, provided each rational player has vanishing average regret. A second example of such a guarantee is provided by Lucier [75], for a natural class of smooth combinatorial auction games.

Analogous guarantees apply to cost-minimization games that satisfy Definition 4.5, although the resulting upper bound on cost includes a term that depends on the irrational players' strategies. This significant performance degradation with irrational players is consistent with the negative results in Roth [92] for congestion games with affine cost functions.

6.2 Subsequent Work

Further Applications. Recent applications of the smoothness framework presented here were given by Bhawalkar et al. [14] and Kollias and Roughgarden [72] to congestion games with weighted players and arbitrary cost functions; Bhawalkar and Roughgarden [16] to simultaneous second-price auctions (Example 2.7) with bidders with subadditive valuations; Caragiannis et al. [24, 25] and Lucier and Paes Leme [77] to the extensively-studied "generalized second price" sponsored search auctions; Chen et al. [29] to games with altruistic players; Cole et al. [38] and Cohen et al. [37] to coordination mechanisms [34] in scheduling; Hoeksma and Uetz [68] and Anshelevich et al. [7] to classes of scheduling games; and Marden and Roughgarden [79] to subclasses of basic utility games (Example 2.6).

Relaxing the Smoothness Condition. The fact that smoothness arguments do not always provide optimal POA bounds motivates weakening the condition in Definition 2.1. As discussed in Remark 2.3, a costless relaxation is to require the inequality (2) for some optimal outcome $\mathbf{s}^*$ and all outcomes $\mathbf{s}$, rather than for all pairs $\mathbf{s}, \mathbf{s}^*$ of outcomes. Several papers noted this relaxation and gave applications of it [68, 77, 85, 97]; the work of Christodoulou et al. [36] (Example 2.7) is a precursor to this definition. Lucier and Paes Leme [77] called this relaxed condition "semi-smoothness."

Two recent works propose relaxations of Definition 2.1 that result in weaker extension theorems. Roughgarden and Schoppmann [97] proposed a "local smoothness" framework for games in which players' strategy sets and cost functions are convex. Local smoothness intuitively requires the inequality (2) only for pairs of strategy profiles that are arbitrarily close to each other. They were motivated by congestion games in which each player can split its weight fractionally over multiple strategies, and showed that local smoothness arguments suffice always give optimal POA bounds while standard smoothness arguments do not [97]. Bhawalkar et al. [15] recently showed that local smoothness arguments give optimal POA bounds for a class of opinion formation games introduced by Bindel et al. [18], while standard smoothness arguments do not. The drawback of local smoothness arguments is that the corresponding POA bounds extend automatically to the correlated equilibria of a game but not, in general, to the coarse correlated equilibria [97]. Since correlated equilibria can be learned efficiently with relatively simple online algorithms [50, 65], this weaker extension theorem continues to significantly relax the assumptions needed to justify Nash equilibrium analysis.

To explain the second relaxation, recall that the outcome $\mathbf{s}^*$ that supplies hypothetical unilateral deviations in Definition 2.1 cannot depend on the outcome $\mathbf{s}$. Syrgkanis and Tardos [103] proved that, if the hypothetical deviation s_i^* for player i depends on its strategy s_i in $\mathbf{s}$ but not on the other strategies $\mathbf{s}_{-i}$, then every consequent POA bound extends to all correlated equilibria. This relaxation is used in [103] to prove composition theorems stating that, under suitable assumptions on bidders' valuations, analyzing the robust POA of many single-item auctions, executed simultaneously or sequentially, reduces to analyzing the robust POA of a single such auction. All of the standard single-item auctions, including first- and second-price auctions, admit good smoothness bounds [66, 103].

The POA in Games of Incomplete Information. This paper, like most previous work on the price of anarchy, has focused on full-information games, where players' payoff functions are common knowledge. It is also important to understand the POA of Bayes-Nash equilibria in games of incomplete information, where players' payoff functions are drawn from a prior distribution that is common knowledge. Auctions, where each player has a private willingness-to-pay for each bundle of goods, provide numerous motivating examples. Early examples of POA analyses of Bayes-Nash equilibria include Gairing et al. [55], Georgiou et al. [57], and Gairing [54] for subclasses of congestion games; Christodoulou et al. [36] for simultaneous single-item auctions; Lucier and Borodin [76] for greedy combinatorial auctions; and Paes Leme and Tardos [89] for sponsored search auctions.

Several recent works proved extension theorems for games of incomplete information. Of these, the extension theorem implicit in Lucier and Paes Leme [77] and explicit in Caragiannis et al. [25] imposes the most stringent hypothesis and has the strongest conclusion. It gives conditions under which POA bounds for pure Nash equilibria of full-information games extend to all Bayes-Nash equilibria of the corresponding incomplete-information games, even when players' payoff functions are correlated. These conditions are met in the important case of sponsored search auctions [25, 77]. For most of the other games of incomplete information in which the POA has been studied, there cannot be an extension theorem for arbitrary correlated prior distributions [16, 96]. Roughgarden [96] and Syrgkanis [102] gave significantly weaker conditions under which full-information pure POA bounds extend to the Bayes-Nash equilibria of every corresponding game of incomplete information with a *product* prior distribution. These extension theorems are general enough to capture most previous results on the POA of Bayes-Nash equilibria.

Limits of Smoothness. Does the robust POA of a game apply to any outcome distributions outside the set of coarse correlated equilibria? Nadav and Roughgarden [85] gave a precise answer to this question. Inspecting the proofs of Theorems 3.2 and 3.3 shows that these extension theorems continue to hold when the coarse correlated equilibrium and no-regret conditions, respectively, are satisfied only on average over the players (rather than by every player). The conditions cannot be weakened further: the robust POA with respect to an optimal solution $\mathbf{s}^*$ of a game is precisely the ratio between the worst-case expected cost of such a "no-average-regret" solution and the cost of an optimal outcome [85].

7 Conclusions and Future Directions

Pure-strategy Nash equilibria — where each player deterministically picks a single strategy — are easier to reason about than their more general cousins like mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria. On the other hand, inefficiency guarantees for more general classes of equilibria are crucial for several reasons: pure Nash equilibria do not always exist; they can be intractable to compute; and even when efficiently computable by a centralized algorithm, they can elude natural learning dynamics.

This paper presented an extension theorem, which extends, in "black-box" fashion, price of anarchy bounds for pure Nash equilibria to the more general equilibrium concepts listed above. Such an extension theorem can hold only under some conditions, and the key idea is to restrict *the method of proof used* to bound the price of anarchy of pure Nash equilibria. We defined smooth games to formalize a canonical method of proof, in which the Nash equilibrium hypothesis is used in only a minimal way, and proved an extension theorem for smooth games. Many of the games in which the price of anarchy has been studied are smooth games in our sense.

For the fundamental model of congestion games with arbitrarily restricted cost functions, we showed that this canonical proof method is guaranteed to produce an optimal upper bound on the worst-case POA. In this sense, POA bounds for congestion games are "intrinsically robust."

There remain many opportunities for interesting work on smooth games. One important research direction is to discover further natural game-theoretic models for which smoothness arguments or variants thereof give good POA bounds. Given the diversity of examples thus far, we expect there will be many more applications.

There are also several interesting open questions concerning the basic smoothness framework and potential refinements.

1. Section 4.3 showed that, in potential games, several variants of best-response dynamics quickly converge to outcomes that approximately obey the robust POA bound. Are more general results possible? A construction of Mirrokni and Vetta [80] shows that the potential game hypothesis cannot be dropped entirely, but the positive results in Goemans et al. [60] for weighted congestion games and Lucier [75] for combinatorial auction games suggest that it can be weakened significantly.

2. Which classes of games are tight in the sense of Definition 2.8? The known tightness results, in Section 5 and [14, 79], employ domain-dependent constructions. Is it enough to have "sufficiently rich" strategy sets?

3. The price of anarchy concerns, by definition, the worst-case equilibrium of a game. Many other measures of inefficiency of game-theoretic equilibria have been studied, and some of these should also admit a "canonical upper bound argument." For example, the price of stability [6] is the ratio between the objective function values of the best Nash equilibrium and an optimal outcome. The hypothesis that an outcome $\mathbf{s}$ is the best Nash equilibrium seems hard to apply in a generic way. However, most analyses of the price of stability concern potential games (see Section 4.3) and analyze the global potential function optimizer, which may or may not be the best Nash equilibrium. The analyses by Christodoulou and Koutsoupias [32] and Caragiannis et al. [23] on the worst-case cost of potential function minimizers in congestion games with affine cost functions suggests that a more general theory should be possible. Very recent progress by Christodoulou and Gairing [31] provides further support for this possibility.

4. A strong Nash equilibrium is an outcome in which no coalition of players can collectively change strategies to make them all better off. Can some of the existing analyses of the corresponding "strong price of anarchy" [5, 28, 44] be unified via smoothness-type arguments, ideally with implications for larger sets of outcomes?

References

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *Journal of the ACM*, 55(6), 2008.

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

[3] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 89–98, 2006.

[4] N. Alon, E. D. Demaine, M. T. Hajiaghayi, and F. T. Leighton. Basic network creation games. *SIAM Journal on Discrete Mathematics*, 2013. To appear.

[5] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. *Games and Economic Behavior*, 65(2):289–317, 2009.

[6] E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

[7] E. Anshelevich, J. Postl, and T. Wexler. Assignment games with conflicts: Price of total anarchy and convergence results via semi-smoothness. arXiv:1304.5149, 2013.

[8] I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. *Journal of Artificial Intelligence Research*, 33:575–613, 2008.

[9] R. J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[10] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrkoni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pages 264–273, 2008.

[11] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

[12] M. Babaioff, R. D. Kleinberg, and C. H. Papadimitriou. Congestion games with malicious players. In *Proceedings of the 8th ACM Conference on Electronic Commerce (EC)*, pages 103–112, 2007.

[13] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation.* Yale University Press, 1956.

[14] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, volume 2, pages 17–28, 2010.

[15] K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 41–50, 2013.

[16] K. Bhawalkar and T. Roughgarden. Welfare guarantees for combinatorial auctions with item bidding. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 700–709, 2011.

[17] V. Bilò, A. Fanelli, M. Flammini, and L. Moscardelli. Performance of one-round walks in linear congestion games. *Theory of Computing Systems*, 49(1):24–45, 2011.

[18] D. Bindel, J. M. Kleinberg, and S. Oren. How bad is forming your own opinion? In *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 57–66, 2011.

[19] A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. *Theory of Computing*, 6:179–199, 2010.

[20] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pages 373–382, 2008.

[21] A. Blum and Y. Mansour. Learning, regret minimization, and equilibria. In Nisan et al. [87], chapter 4, pages 79–101.

[22] M. Bradonjić, G. Ercal-Ozkaya, A. Meyerson, and A. Roytman. On the price of mediation. In *Proceedings of the 10th ACM Conference on Electronic Commerce (EC)*, pages 315–324, 2009.

[23] I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. *Algorithmica*, 61(3):606–637, 2011.

[24] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, and M. Kyropoulou. On the efficiency of equilibria in generalized second price auctions. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 81–90, 2011.

[25] I. Caragiannis, C. Kaklamanis, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. Paes Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. arXiv:1201.6429, 2012.

[26] N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

[27] D. Chakrabarty. Improved bicriteria results for the selfish routing problem. Unpublished manuscript, 2004.

[28] S. Chein and A. Sinclair. Strong and Pareto price of anarchy in congestion games. In *Proceedings of the 36th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 279–291, 2009.

[29] P. Chen, D. de Keijzer, D. Kempe, and G. Schaefer. The robust price of anarchy of altruistic games. In *Proceedings of the 7th International Workshop on Internet and Network Economies (WINE)*, pages 383–390, 2011.

[30] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.

[31] G. Christodoulou and M. Gairing. Price of stability in polynomial congestion games. In *Proceedings of the 40th International Colloqium on Automata, Languages and Programming (ICALP)*, 2013. To appear.

[32] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proceedings of the 13th Annual European Symposium on Algorithms (ESA)*, pages 59–70, 2005.

[33] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

[34] G. Christodoulou, E. Koutsoupias, and A. Nanavati. Coordination mechanisms. *Theoretical Computer Science*, 410(36):3327–3336, 2009.

[35] G. Christodoulou, E. Koutsoupias, and P. G. Spirakis. On the performance of approximate equilibria in congestion games. *Algorithmica*, 61(1):116–140, 2011.

[36] G. Christodoulou, A. Kovács, and M. Schapira. Bayesian combinatorial auctions. In *Proceedings of the 35th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 820–832, 2008.

[37] J. Cohen, C. Dürr, and N. K. Thang. Smooth inequalities and equilibrium inefficiency in scheduling games. In *Proceedings of the 8th International Workshop on Internet and Network Economies (WINE)*, pages 350–363, 2012.

[38] R. Cole, J. R. Correa, V. Gkatzelis, V. S. Mirrokni, and N. Olver. Decentralized utilitarian mechanisms for scheduling games. *Games and Economic Behavior*, 2013. To appear.

[39] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[40] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. A geometric approach to the price of anarchy in nonatomic congestion games. *Games and Economic Behavior*, 64(2):457–469, 2008.

[41] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Transactions on Algorithms*, 3(1), 2007. Article 4.

[42] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[43] D. Dumrauf and M. Gairing. Price of anarchy for polynomial wardrop games. In *Proceedings of the 6th International Workshop on Internet and Network Economies (WINE)*, pages 319–330, 2006.

[44] A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. *Games and Economic Behavior*, 67(1):51–68, 2009.

[45] K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010.

[46] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. J. Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[47] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pages 604–612, 2004.

[48] A. Fanelli, L. Moscardelli, and A. Skopalik. On the impact of fair best response dynamics. In *Proceedings of the 37th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pages 360–371, 2012.

[49] B. Farzad, N. Olver, and A. Vetta. A priority-based model of routing. *Chicago Journal of Theoretical Computer Science*, 2008. Article 1.

[50] D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.

[51] D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. *Theory of Computing Systems*, 47(1):113–136, 2010.

[52] D. Fotakis. Stackelberg strategies for atomic congestion games. *Theory of Computing Systems*, 47(1):218–249, 2010.

[53] M. Gairing. *Selfish Routing in Networks*. PhD thesis, Universität Paderborn, 2006.

[54] M. Gairing. Malicious Bayesian congestion games. In *Proceedings of the 6th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 119–132, 2008.

[55] M. Gairing, B. Monien, and K. Tiemann. Selfish routing with incomplete information. *Theory of Computing Systems*, 42(1):91–130, 2008.

[56] M. Gairing, B. Monien, and K. Tiemann. Routing (un-) splittable flow in games with player-specific affine latency functions. *ACM Transactions on Algorithms*, 7(3), 2011. Article 31.

[57] C. Georgiou, T. Pavlides, and A. Philippou. Selfish routing in the presence of network uncertainty. *Parallel Processing Letters*, 19(1):141–157, 2009.

[58] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

[59] M. X. Goemans, E. L. Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. *IEEE Journal on Selected Areas in Communications*, 24(5):1020–1033, 2006.

[60] M. X. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[61] J. Hannan. Approximation to Bayes risk in repeated play. *Ccontributions to the Theory of Games*, 3:97–139, 1957.

[62] T. Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems*, 48(4):781–802, 2011.

[63] T. Harks, M. Klimm, and R. H. Möhring. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems*, 49(1):46–70, 2011.

[64] T. Harks and L. A. Végh. Nonadaptive selfish routing with online demands. In *Proceedings of the 4th Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN)*, pages 27–45, 2007.

[65] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibria. *Econometrica*, 68(5):1127–1150, 2000.

[66] A. Hassidim, H. Kaplan, M. Mansour, and N. Nisan. Non-price equilibria in markets of discrete goods. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 295–296, 2011.

[67] B. Heydenreich, R. Müller, and M. Uetz. Games and mechanism design in machine scheduling - an introduction. *Production and Operations Management*, 16(4):437–454, 2007.

[68] R. Hoeksma and M. Uetz. The price of anarchy for minsum related machine scheduling. In *Proceedings of the 9th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 261–273, 2011.

[69] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and Economic Behavior*, 21(1/2):85–101, 1997.

[70] G. Karakostas and A. Viglas. Equilibria for networks with malicious users. *Mathematical Programming, Series A*, 110(3):591–613, 2007.

[71] R. D. Kleinberg, G. Piliouras, and É. Tardos. Multiplicative updates outperform generic no-regret learning in congestion games. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 533–542, 2009.

[72] K. Kollias and T. Roughgarden. Restoring pure equilibria to weighted congestion games. In *Proceedings of the 38th International Colloqium on Automata, Languages and Programming (ICALP)*, pages 539–551, 2011.

[73] E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. Approximate equilibria and ball fusion. *Theory of Computing Systems*, 36(6):683–693, 2003.

[74] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.

[75] B. Lucier. Beyond equilibria: Mechanisms for repeated combinatorial auctions. In *Proceedings of the First Conference on Innovations in Computer Science*, pages 166–177, 2010.

[76] B. Lucier and A. Borodin. Price of anarchy for greedy auctions. In *Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 537–553, 2010.

[77] B. Lucier and R. Paes Leme. GSP auctions with correlated types. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC)*, pages 71–80, 2011.

[78] T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *Theoretical Computer Science*, 406(3):187–206, 2008.

[79] J. R. Marden and T. Roughgarden. Generalized efficiency bounds in distributed resource allocation. In *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, pages 2233–2238, 2010.

[80] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 183–194, 2004.

[81] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[82] T. Moscibroda, S. Schmid, and R. Wattenhofer. When selfish meets evil: Byzantine players in a virus inoculation game. In *Proceedings of the 25th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 35–44, 2006.

[83] R. Motwani and P. Raghavan. *Randomized Algorithms*. Cambridge University Press, 1996.

[84] H. Moulin and J. P. Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7(3/4):201–221, 1978.

[85] U. Nadav and T. Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *Proceedings of the 6th International Workshop on Internet and Network Economies (WINE)*, pages 319–326, 2010.

[86] J. F. Nash. Equilibrium points in N-person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.

[87] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.

[88] N. Olver. The price of anarchy and a priority-based model of routing. M.S. thesis, McGill University, 2006.

[89] R. Paes Leme and É. Tardos. Pure and Bayes-Nash price of anarchy for generalized second price auction. In *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 735–744, 2010.

[90] Georgia Perakis. The "price of anarchy" under nonlinear and asymmetric costs. *Mathematics of Operations Research*, 32(3):614–628, 2007.

[91] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[92] A. Roth. The price of malice in linear congestion games. In *Proceedings of the 8th International Workshop on Internet and Network Economies (WINE)*, pages 118–125, 2008.

[93] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[94] T. Roughgarden. Routing games. In Nisan et al. [87], chapter 18, pages 461–486.

[95] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.

[96] T. Roughgarden. The price of anarchy in games of incomplete information. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC)*, pages 862–879, 2012.

[97] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in atomic splittable congestion games. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 255–267, 2011.

[98] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.

[99] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49(2):389–403, 2004.

[100] A. Skopalik and B. Vöcking. Inapproximability of pure Nash equilibria. In *Proceedings of the 40th ACM Symposium on Theory of Computing (STOC)*, pages 355–364, 2008.

[101] S. Suri, C. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.

[102] V. Syrgkanis. Bayesian games and the smoothness framework. arXiv.cs.GT:1203.5155v1, 2012.

[103] V. Syrgkanis and É. Tardos. Composable and efficient mechanisms. In *Proceedings of the 45th ACM Symposium on Theory of Computing (STOC)*, pages 211–220, 2013.

[104] É. Tardos and T. Wexler. Network formation games and the potential function method. In Nisan et al. [87], chapter 19, pages 487–516.

[105] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–425, 2002.

[106] B. Vöcking. Selfish load balancing. In Nisan et al. [87], chapter 20, pages 517–542.

[107] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.

[108] D. Weitz. The price of anarchy. Manuscript, 2001.