Intrinsic Robustness of the Price of Anarchy

Tim Roughgarden*
Department of Computer Science
Stanford University
353 Serra Mall, Stanford, CA 94305
tim@cs.stanford.edu

ABSTRACT

The price of anarchy (POA) is a worst-case measure of the inefficiency of selfish behavior, defined as the ratio of the objective function value of a worst Nash equilibrium of a game and that of an optimal outcome. This measure implicitly assumes that players successfully reach some Nash equilibrium. This drawback motivates the search for inefficiency bounds that apply more generally to weaker notions of equilibria, such as mixed Nash and correlated equilibria; or to sequences of outcomes generated by natural experimentation strategies, such as successive best responses or simultaneous regret-minimization.

We prove a general and fundamental connection between the price of anarchy and its seemingly stronger relatives in classes of games with a sum objective. First, we identify a "canonical sufficient condition" for an upper bound of the POA for pure Nash equilibria, which we call a *smoothness argument*. Second, we show that every bound 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 regret-minimizing players (or "price of total anarchy"). Smoothness arguments also have automatic implications for the inefficiency of approximate and Bayesian-Nash equilibria and, under mild additional assumptions, for bicriteria bounds and for polynomial-length best-response sequences.

We also identify classes of games — most notably, congestion games with cost functions restricted to an arbitrary fixed set — that are *tight*, in the sense that smoothness arguments are guaranteed to produce an optimal worst-case upper bound on the POA, even for the smallest set of interest (pure Nash equilibria). Byproducts of our proof of this result include the first tight bounds on the POA in congestion games with non-polynomial cost functions, and the first structural characterization of atomic congestion games that are universal worst-case examples for the POA.

*Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

STOC'09, May 31–June 2, 2009, Bethesda, Maryland, USA.

Categories and Subject Descriptors

F.0 [**Theory of Computation**]: General

General Terms

Algorithms, Economics, Theory

Keywords

Price of anarchy; congestion games; Nash equilibria; regret-minimization

1. INTRODUCTION

Self-interested behavior by autonomous decision-makers with competing objectives 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 large systems (the Internet being an obvious example), motivating the search for conditions under which decentralized optimization by competing individuals is guaranteed to produce a near-optimal outcome.

Any guarantee of this type requires a formal behavioral model, in order to define "the outcome of selfish behavior". The majority of the literature studies *pure-strategy Nash equilibria*, defined as follows. Each of k players selects a strategy s_i from a set S_i (e.g., a path in a network), with the cost $C_i(\mathbf{s})$ of each player i depending on all of the selected strategies $\mathbf{s}$. A strategy profile $\mathbf{s}$ is then a pure-strategy Nash equilibrium if no player can decrease its cost via a unilateral deviation: $C_i(\mathbf{s}) \leq C_i(s'_i, s_{-i})$ for every i and $s'_i \in S_i$, where s_{-i} denotes the strategies chosen by the players other than i. (These concepts can be defined equally well via payoff-maximization rather than cost-minimization.)

The *price of anarchy (POA)* is a standard measure of the suboptimality introduced by self-interested behavior. Given a game, a notion of an "equilibrium" (such as pure Nash equilibria), and an 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 (minimum-cost) outcome. A POA guarantee has an attractive worst-case flavor: it applies to every possible equilibrium and obviates the need to predict a single outcome of selfish behavior. Many POA guarantees for pure Nash equilibria and mixed-strategy Nash equilibria (in which players

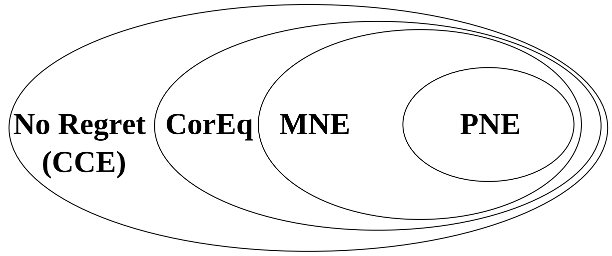

Figure 1: Some of the generalizations of pure Nash equilibria that are amenable to POA bounds. "PNE" stands for pure Nash equilibria; "MNE" for mixed Nash equilibria; "CorEq" for correlated equilibria; and "No Regret (CCE)" for coarse correlated equilibria, which are the empirical distributions of joint play in which every player has no asymptotic external regret.

randomize to minimize their expected costs) are known, in a number of different models; see [25, Chapters 17–21] and the references therein.

A bound on the POA implicitly assumes that players successfully reach *some* equilibrium. For pure Nash equilibria, however, there are a number of reasons why this might not occur: perhaps the players fail to coordinate on one of multiple equilibria; or they are playing a game in which computing a pure Nash equilibrium is PLS-complete [15]; or, even more fundamentally, a game in which pure Nash equilibria simply 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.

For example, Figure 1 shows a hierarchy of sets of probability distributions over outcomes that generalize pure Nash equilibria: mixed-strategy Nash equilibria, correlated equilibria, and coarse correlated equilibria (see Section 2 for formal definitions). The three largest sets in Figure 1 are guaranteed to be non-empty in every finite game [24], and some of them significantly relax the behavioral assumptions necessary to justify Nash equilibrium analysis. For instance, the largest set comprises the empirical distributions of joint play by players that incur no external regret, and this joint play need not "converge" in any sense. All of the inclusions shown in Figure 1 are generally strict, and there are a number of specific games where, for example, correlated equilibria can have much larger expected cost than (mixed or pure) Nash equilibria [2, 8, 9].

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

> *For many fundamental classes of games and for sum objective functions, the worst-case POA is identical for all of the equilibrium concepts in Figure 1.*

In this sense, POA bounds for pure Nash equilibria in such classes of games are "intrinsically robust".

Our techniques also imply performance guarantees for approximate Nash equilibria and Bayes-Nash equilibria in these classes of games as well as, under mild extra conditions, for polynomial-length best-response sequences and "bicriteria bounds".

1.1 Overview

Our contributions can be divided into three parts.

(A) We identify a sufficient condition for an upper bound on the POA of pure Nash equilibria of a game, which encodes a "canonical proof template" for deriving such bounds. Most of the POA upper bounds in the literature for games with a sum objective can be recast as instantiations of this canonical method.

(B) We show that, for every game with a sum objective, every POA bound proved using this canonical technique *extends automatically*, without any quantitative degradation, to all of the sets in Figure 1 (among other applications).

(C) We identify classes of games — most notably, congestion games with cost functions restricted to an arbitrary fixed set — that are *tight*, meaning that our canonical upper bound technique is *guaranteed* to produce an optimal worst-case upper bound on the POA, even for the smallest set of interest (pure Nash equilibria).

We now provide some details and examples to illustrate these three points. By a *cost-minimization game*, we mean a game (defined as above) together with the total cost objective function $C(\mathbf{s}) = \sum_{i=1}^{k} C_i(\mathbf{s})$. Our canonical method of upper bounding the POA of a cost-minimization game, as discussed in step (A) above, is by what we call a "smoothness" argument.

Definition 1.1 (Smooth Games) A cost-minimization game is (λ, μ)-*smooth* if for every two outcomes $\mathbf{s}$ and $\mathbf{s}^*$,

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

If a game is (λ, μ)-smooth (with $\lambda \geq 0$ 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 $\mathbf{s}^*$. To justify this, use the definition of the cost objective, the Nash equilibrium condition, and smoothness to derive

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

rearranging terms yields the claimed bound.

We define the *robust POA* as the best (i.e., least) upper bound on the POA that is provable via a smoothness argument.

Definition 1.2 (Robust POA) The *robust price of anarchy* of a cost-minimization game is

$$\inf \left\{ \frac{\lambda}{1-\mu} \, : \, (\lambda, \mu) \text{ s.t. the game is } (\lambda, \mu)\text{-smooth} \right\}.$$

The strength of Definition 1.1 is that inequality (1) is required to hold for *every* outcome $\mathbf{s}$, and *not only for Nash equilibria*; this is the reason that an upper bound on the robust POA has immediate implications beyond pure Nash equilibria. To appreciate this and highlight a special case of our second contribution (B), consider a sequence $\mathbf{s}^1, \mathbf{s}^2, \ldots, \mathbf{s}^T$ of outcomes of a (λ, μ)-smooth game. (We consider only pure strategies here; see Section 2 for the general mixed-strategy case.) Define $\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, s_{-i}^t)$ for each i and t, where $\mathbf{s}^*$ again denotes a minimum-cost outcome. We can mimic the derivation in (2) to obtain

$$C(\mathbf{s}^t) \le \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{\sum_{i=1}^k \delta_i(\mathbf{s}^t)}{1-\mu} \tag{3}$$

for each t. Suppose that every player i experiences vanishing average external regret, in that

$$\sum_t C_i(\mathbf{s}^t) \le \left[\min_{s_i'} \sum_t C_i(s_i', s_{-i}^t)\right] + o(T).$$

Averaging (3) over the T time steps and reversing the order of the resulting double summation yields

$$\frac{1}{T}\sum_{t=1}^{T} C(\mathbf{s}^t) \le \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{1}{1-\mu}\sum_{i=1}^{k}\left(\frac{1}{T}\sum_{t=1}^{T}\delta_i(\mathbf{s}^t)\right). \tag{4}$$

Recalling that $\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, s_{-i}^t)$ is the additional cost of player i at time t beyond that incurred by playing the (fixed) strategy s_i^*, the no-regret guarantee implies that $[\sum_t \delta_i(\mathbf{s}^t)]/T$ goes to 0 with T. Since this holds for every player i, inequality (4) implies that the average cost of outcomes in the sequence is no more than the robust POA times the minimum-possible cost, plus an $o(1)$ factor that goes to zero as $T \to \infty$.

1.2 Two Concrete Examples

Nervousness 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 single out two well-known POA analyses that can be recast as smoothness arguments. (See the full version of this paper for a much longer list.) These examples demonstrate that Definition 1.1 is non-vacuous and, at the very least, can be used to derive a host of known results in a simple and unified way. (New results will have to wait until Section 3.)

The first example is a special class of congestion games; Section 3 studies the general case in detail. The second example, which concerns Vetta's well-studied utility games [33], illustrates how smoothness arguments can be defined and used in payoff-maximization games, and also with a "one-sided" variant of sum objective functions.

Example 1.3 (Congestion Games) A congestion game is a cost-minimization game defined by a ground set E, a set of k players with strategy sets $S_1, \ldots, S_k \subseteq 2^E$, and a non-negative, non-decreasing cost function $c_e : \mathcal{Z}^+ \to \mathcal{R}^+$ for each element $e \in E$ [28]. A canonical example is atomic selfish routing (or "network congestion") games, where E is the edge set of a network, and S_i is the set of paths from a source vertex to a sink vertex. Given a strategy profile $\mathbf{x} = (x_1, \ldots, x_k)$, with $x_i \in S_i$ for each i, let $x_e = |\{i : e \in x_i\}|$ denote the *load* on e, defined as the number of players that use it. The cost to player i is defined as $C_i(\mathbf{x}) = \sum_{e \in x_i} c_e(x_e)$. For this example, we assume that every cost function is *affine*, meaning that $c_e(x_e) = a_e x_e + b_e$ with $a_e, b_e \ge 0$ for every $e \in E$.

We claim that every congestion game with affine cost functions is $(5/3, 1/3)$-smooth. The basic reason for this was identified by Christodoulou and Koutsoupias [13, Lemma 1], who noted that

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

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

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

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

$$\begin{aligned}
\sum_{i=1}^{k} C_i(x_i^*, x_{-i}) &\le \sum_{e \in E}(a_e(x_e+1) + b_e)x_e^* \\
&\le \sum_{e \in E}\frac{5}{3}(a_e x_e^* + b_e)x_e^* + \sum_{e \in E}\frac{1}{3}(a_e x_e + b_e)x_e \\
&= \frac{5}{3}C(\mathbf{x}^*) + \frac{1}{3}C(\mathbf{x}),
\end{aligned}$$

as desired. This smoothness guarantee implies an upper bound of 5/2 for the POA of pure Nash equilibria; this was first proved in [6, 13], where matching lower bounds were also supplied. The framework introduced in this paper implies that the bound automatically extends to, in particular, the three other sets of outcomes shown in Figure 1; these extensions were originally established in two different papers [8, 12] subsequent to the original POA bound [6, 13].

Example 1.4 (Valid Utility Games) Our second example concerns a class of games called *valid utility games* [33]. These games 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 Π to denote the objective function of a payoff-maximization game. We call such a game (λ, μ)-*smooth* if

$$\sum_i \Pi_i(s_i^*, s_{-i}) \ge \lambda \cdot \Pi(\mathbf{s}^*) - \mu \cdot \Pi(\mathbf{s})$$

for every pair $\mathbf{s}, \mathbf{s}^*$ of outcomes. A derivation similar to (2) 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. (Similarly to (3) and (4), the same bound applies to, for example, no-regret sequences.) We define the *robust POA* of a payoff-maximization game as the supremum of $\lambda/(1+\mu)$ over all legitimate smoothness parameters (λ, μ).

A valid utility game is defined by a ground set E, a non-negative 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 $i = 1, 2, \ldots, k$. For an outcome $\mathbf{s}$, let $U(\mathbf{s}) \subseteq E$ denote the union of players' strategies in $\mathbf{s}$. The objective function value of an outcome $\mathbf{s}$ is defined as $\Pi(\mathbf{s}) = V(U(\mathbf{s}))$. Furthermore, the definition requires that two conditions hold: (1) for each player i, $\Pi_i(\mathbf{s}) \ge V(U(\mathbf{s})) - V(U(\emptyset, s_{-i}))$ for every outcome $\mathbf{s}$; and (2) $\sum_{i=1}^k \Pi_i(\mathbf{s}) \le \Pi(\mathbf{s})$ for every outcome

[1] The statement of this lemma in [12, 13] contains a typo, but it is applied correctly in both works.

$\mathbf{s}$. One concrete example of such a game is competitive facility location with price-taking markets and profit-maximizing firms [33].

Valid utility games do *not* have a sum objective function in the sense we have defined them, in that the inequality in condition (2) is generally strict. But inspection of the (payoff-maximization analog of) the derivation in (2) shows that this one-sided inequality is sufficient to conclude a POA bound of $\lambda/(1+\mu)$ from an (λ, μ)-smoothness guarantee. (In a cost-minimization context, all of our derivations remain valid for objective functions C satisfying $C(\mathbf{s}) \le \sum_{i=1}^{k} C_i(\mathbf{s})$ for every outcome $\mathbf{s}$.)

We claim that every valid utility game with a nondecreasing objective function V is $(1,1)$-smooth. The proof is essentially a few key inequalities from [33, Theorem 3.2], as follows. Let $\mathbf{s}, \mathbf{s}^*$ denote arbitrary outcomes of a valid utility game with a nondecreasing objective function. Let $U_i \subseteq E$ denote the union of all of the players' strategies in $\mathbf{s}$, together with the strategies employed by players $1, 2, \ldots, i$ in $\mathbf{s}^*$. Applying condition (1), the submodularity of V, and the nondecreasing property of V yields

$$\begin{aligned}\sum_{i=1}^{k} \Pi_i(s_i^*, s_{-i}) &\ge \sum_{i=1}^{k} [V(U(s_i^*, s_{-i})) - V(U(\emptyset, s_{-i}))] \\ &\ge \sum_{i=1}^{k} [V(U_i) - V(U_{i-1})] \\ &\ge \Pi(\mathbf{s}^*) - \Pi(\mathbf{s}),\end{aligned}$$

as desired. This smoothness argument implies a lower bound of 1/2 on the POA of pure Nash equilibria (first proved in [33], along with a matching upper bound) and more generally of no-regret sequences (recently established in [8]).

1.3 Tight Classes of Games

The best-possible POA upper bound for a set of allowable outcomes increases as the set grows bigger. *Or so one would think.* Examples 1.3 and 1.4 share a remarkable property: smoothness arguments, despite their automatic generality, provide a tight bound on the POA, *even for pure Nash equilibria.* Our third contribution (C) studies whether this property is coincidental or fundamental.

Let $\mathcal{G}$ denote a set of games, and assume that a nonnegative objective function has been defined on the outcomes of these games. 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 $G \in \widehat{\mathcal{G}}$. Our work thus far shows, as a very special case, that for every $(\lambda, \mu) \in \mathcal{A}(\mathcal{G})$ and every $G \in \widehat{\mathcal{G}}$, $\rho_{pure}(G) \le \lambda/(1-\mu)$. We call a class of games *tight* if equality holds for suitable choices of $(\lambda, \mu) \in \mathcal{A}(\mathcal{G})$ and $G \in \widehat{\mathcal{G}}$.

Definition 1.5 (Tight Class of Games) A class $\mathcal{G}$ of games is *tight* if

$$\sup_{G \in \widehat{\mathcal{G}}} \rho_{pure}(G) = \inf_{(\lambda,\mu) \in \mathcal{A}(\mathcal{G})} \frac{\lambda}{1-\mu}. \qquad (5)$$

The right-hand side of (5) is the best upper bound provable via a worst-case smoothness argument. The left-hand side of (5) is the actual worst-case POA of pure Nash equilibria in $\mathcal{G}$, among games with at least one pure Nash equilibrium. We reiterate that the left-hand side is trivially upper bounded by the right-hand side, in the spirit of "weak duality". Tight classes of games are characterized by the min-max condition (5), which can be very loosely interpreted as a "strong duality-type" result.

For instance, Example 1.3 shows that, if $\mathcal{G}$ is the set of congestion games with affine cost functions, then the right-hand side of (5) is at most 5/2. Constructions of Awerbuch et al. [6] and Christodoulou and Koutsoupias [13] show that the left-hand side is at least 5/2 for this class of games. Thus, congestion games with affine cost functions form a tight class.

The final main result of this paper shows that this result is no fluke: *for every fixed set $\mathcal{C}$ of allowable cost functions, the class of congestion games with cost functions in $\mathcal{C}$ is tight.* Byproducts of our proof of this result include the first tight bounds on the POA in congestion games with non-polynomial cost functions, and the first structural characterization of (atomic) congestion games that are universal worst-case examples for the POA.

1.4 Related Work

The price of anarchy was first studied in [22] for makespan minimization in scheduling games; note this is *not* a sum objective and our framework does not imply anything interesting about it. Indeed, the worst-case POA in this model was immediately recognized to be different for pure and mixed Nash equilibria [14, 21, 22], and later for correlated equilibria and regret-minimizing players [8]. For other gaps between the sets in Figure 1, see e.g. [2, 9].

The POA with a sum objective was first studied in [31] for "nonatomic" games, with infinitely many players of negligible size. The first general results on the POA of pure Nash equilibria for atomic congestion games and their weighted variants are in [6, 13], who gave tight bounds for games with affine cost functions and reasonably close upper and lower bounds for games with polynomial cost functions with nonnegative coefficients; matching upper and lower bounds for the latter class were later given independently in [1] and [26]. While most of the proof techniques in [1, 26] apply only to games with polynomial cost functions, this paper greatly generalizes some of the core ideas in [1], resulting in a characterization of the worst-case POA as general as that previously proved in Roughgarden [30] for nonatomic congestion games.

Many results extending a POA bound for pure Nash equilibria to a larger set of outcomes are scattered across the literature. The underlying bound on the POA of pure Nash equilibria can be formulated as a smoothness argument in almost all of these cases, so our general analysis immediately implies (and in most cases strengthens) these previous results. Specifically, the authors in [1, 6, 13, 33] each observe that their upper bounds on the worst-case POA of pure Nash equilibria carry over easily to mixed Nash equilibria. In [12] the worst-case POA of correlated equilibria is shown to be the same as for pure Nash equilibria in unweighted and weighted congestion games with affine cost functions. Blum et al. [7, 8] 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 earned by no-regret players, a type of bound that they dub the "price of total anarchy". They accomplish this for the optimal bounds for nonatomic congestion games in [32], the optimal bounds for valid utility games in [33], and the (suboptimal) bounds of [6, 13] for unweighted congestion games with polynomial cost functions. Blum et al. [8] also give a result of this type for a constant-sum location game and a fairness objective, which falls outside of our framework. Finally, Goemans et al. [19] and Awerbuch et al. [5] prove that, under various conditions, best-response dynamics quickly reaches a state that approximately obeys the worst-case POA bound for pure Nash equilibria. Our results imply those in [5, 19] except for the upper bounds in [19] on the "price of sinking" in non-potential games (see [23]), such as weighted congestion games with non-affine cost functions; it is unclear if smoothness can be used to derive directly the latter results.

2. SOME CONSEQUENCES OF ROBUST POA BOUNDS

This section outlines some of the many promised consequences of an upper bound on the robust POA of a game with a sum objective. There are also easy consequences for the worst-case inefficiency of approximate Nash equilibria and for Bayes-Nash equilibria; these applications are detailed in the full version of this extended abstract. We work with cost-minimization games, though similar results hold for smooth payoff-maximization games (as in Example 1.4).

2.1 Nash, Correlated, and Coarse Correlated Equilibria

We begin with the most direct consequences of a smoothness argument: bounds on the four sets in Figure 1. While the stronger results in this section directly imply the weaker ones, we prefer to develop the theory one step at a time, defining the necessary concepts along the way.

Nash and Correlated Equilibria.

We already noted in (2) that an upper bound on the robust POA of a cost-minimization game easily implies the same bound on the POA of its pure Nash equilibria.

Proposition 2.1 (POA of Pure Nash Equilibria) *For every cost-minimization game G with robust POA $\rho(G)$, for every pure Nash equilibrium $\mathbf{s}$ and outcome $\mathbf{s}^*$ of G,*

$$C(\mathbf{s}) \leq \rho(G) \cdot C(\mathbf{s}^*).$$

To extend Proposition 2.1 to more general equilibrium concepts, we prove that the smoothness condition (1) for pairs of outcomes automatically extends to (not necessarily product) probability distributions over outcomes.

Lemma 2.2 (Smoothness for Distributions) *Let G be a (λ, μ)-smooth cost-minimization game, σ a probability distribution over the outcomes of G, and s^* an outcome of G. Then*

$$\sum_{i=1}^{k} \mathbf{E}_{s_{-i} \sim \sigma_{-i}}[C_i(s_i^*, s_{-i})] \leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot \mathbf{E}_{\mathbf{s} \sim \sigma}[C(\mathbf{s})],$$

where σ_{-i} denotes the marginal distribution of σ for players other than i.

PROOF. Using linearity of expectation and the smoothness guarantee, we obtain

$$\begin{aligned}
\sum_{i=1}^{k} \mathbf{E}_{s_{-i} \sim \sigma_{-i}}[C_i(s_i^*, s_{-i})] &= \sum_{i=1}^{k} \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s_i^*, s_{-i})] \\
&= \mathbf{E}_{\mathbf{s} \sim \sigma}\left[\sum_{i=1}^{k} C_i(s_i^*, s_{-i})\right] \\
&\leq \mathbf{E}_{\mathbf{s} \sim \sigma}\left[\lambda \cdot C(s^*) + \mu \cdot C(\mathbf{s})\right] \\
&= \lambda \cdot C(\mathbf{s}^*) + \mu \cdot \mathbf{E}_{\mathbf{s} \sim \sigma}[C(\mathbf{s})].
\end{aligned}$$

□

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-strategy 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:

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

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-strategy Nash equilibrium and not conversely; indeed, many games have no pure Nash equilibrium, but every finite game has a mixed-strategy Nash equilibrium [24].

Substituting expected cost $\mathbf{E}_{\mathbf{s} \sim \sigma}[C(\mathbf{s})]$ in place of cost $C(\mathbf{s})$ and Lemma 2.2 in place of Definition 1.1 in the derivation (2) shows that the robust POA of a cost-minimization game upper bounds its POA with respect to mixed Nash equilibria.

Proposition 2.3 (POA of Mixed Nash Equilibria) *For every cost-minimization game G with robust POA $\rho(G)$, every mixed-strategy Nash equilibrium $\sigma_1, \ldots, \sigma_k$ of G, and every outcome $\mathbf{s}^*$ of G,*

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

A *correlated equilibrium* [3] of a cost-minimization game G is a joint probability distribution σ over the outcomes of G with the property that

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s}) | s_i] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s_i', s_{-i}) | s_i] \qquad (6)$$

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). One easily verifies that mixed-strategy Nash equilibria are precisely the correlated equilibria that are also product distributions. Correlated equilibria have been widely studied as strategies for a benevolent mediator, and also because of strong connections to Bayesian rationality [4] and to learning (e.g. [16, 20]), and for their relative computational tractability [18, 27].

Since Lemma 2.2 applies to arbitrary joint distributions, the proof of Proposition 2.3 carries over without change to bound the POA for correlated equilibria.

Proposition 2.4 (POA of Correlated Equilibria) *For every cost-minimization game G with robust POA $\rho(G)$, every correlated equilibrium σ of G, and every outcome $\mathbf{s}^*$ of G,*

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

Regret-Minimization, Coarse Correlated Equilibria, and the Price of Total Anarchy.

A *no-regret sequence* $\sigma^1, \ldots, \sigma^T$ of (not necessarily product) probability distributions over outcomes is defined by the property that the total expected cost of each player is at most $o(T)$ more than that of the best fixed strategy in hindsight: for all i and $s_i' \in S_i$,

$$\mathbf{E}\left[\sum_{t=1}^{T} C_i(\mathbf{s}^t)\right] \le \mathbf{E}\left[\sum_{t=1}^{T} C_i(s_i', s_{-i}^t)\right] + o(T),$$

where $\mathbf{s}^t \sim \sigma^t$ and $s_{-i}^t \sim \sigma_{-i}^t$ for every t. Players that employ no-regret algorithms — and many fast and simple such algorithms exist, see e.g. [10] — are guaranteed to generate a no-regret sequence, and such sequences can also arise for other reasons. No-regret sequences can simulate the expected costs of every correlated equilibrium and are strictly more general; and even the empirical distribution of such a sequence need not converge as $T \to \infty$.

Blum et al. [8] define the *price of total anarchy* as the worst-case ratio between the expected average cost of a no-regret sequence and the cost of an optimal outcome. Substituting expected costs $\mathbf{E}_{\mathbf{s}^t\sim\sigma^t}[C(\mathbf{s}^t)]$ in place of costs $C(\mathbf{s}^t)$ and Lemma 2.2 in place of Definition 1.1 in the derivations (3) and (4) shows that the robust POA of a cost-minimization game upper bounds its price of total anarchy.

Proposition 2.5 (Price of Total Anarchy) *For every cost-minimization game G with robust POA $\rho(G)$, every no-regret sequence $\sigma^1, \ldots, \sigma^T$, and every outcome $\mathbf{s}^*$ of G,*

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

as $T \to \infty$.

In particular, the robust POA of a game applies to all of its *coarse correlated equilibria*, meaning the probability distributions σ over outcomes that satisfy

$$\mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(\mathbf{s})] \le \mathbf{E}_{\mathbf{s}\sim\sigma}[C_i(s_i', s_{-i})] \tag{7}$$

for every i and $s_i, s_i' \in S_i$. The set of all such distributions is also sometimes called the *Hannan set*; see Young [34], for example. While a correlated equilibrium (6) protects against deviations by players aware of their recommended strategy, a coarse correlated equilibrium (7) is only constrained by player deviations that are committed to in advance of the sampled outcome.

2.2 Short Best-Response Sequences

Best-response dynamics (BRD) is a natural myopic model of how players search for a pure Nash equilibrium: if the current outcome $\mathbf{s}$ is not a pure Nash equilibrium, then choose some player i who can decrease its cost and switch the player to a strategy that minimizes $C_i(s_i', s_{-i})$. BRD always converges in a cost-minimization game that admits a *potential function* Φ, which by definition satisfies

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

for every outcome $\mathbf{s}$, player i, and deviation s_i' [23]. For convenience, we define a *lower* potential function as one that only underestimates cost: $\Phi(\mathbf{s}) \le C(\mathbf{s})$ for every $\mathbf{s}$. It will be evident that the result below holds, with a worse convergence bound, for general potential functions. Congestion games (see Section 3) are prominent examples of games with (lower) potentials; for others, see e.g. [5].

Since every round of BRD decreases the potential function, it eventually converges to a pure Nash equilibrium. An upper bound on the POA obviously applies to BRD at termination, but convergence requires exponential time in general [15]. We next observe that in smooth games with a lower potential function, two forms of BRD are guaranteed to reach quickly outcomes that meet the robust POA bound. (The full version of this extended abstract describes additional results along these lines.) The key to these results is the following lemma, which is inspired by arguments in [5, 11]. Roughly, it states that as long as the player chosen in each round of BRD is likely to have at least an average incentive to deviate, relative to the other players, then BRD rapidly reaches outcomes that obey the robust POA bound.

Proposition 2.6 (Bounds for BRD) *Let $\mathbf{s}^0, \ldots, \mathbf{s}^T$ be a best-response sequence of a cost-minimization game G with robust POA ρ, $\mathbf{s}^*$ a minimum-cost outcome of G, and $\epsilon, \alpha > 0$ parameters. Define $\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, s_{-i}^t)$ and $\Delta(\mathbf{s}^t) = \sum_{i=1}^{k} \delta_i(\mathbf{s}^t)$.*

If G has a nonnegative, integral lower potential Φ and the (possibly randomly chosen) player $i(t)$ at each time t satisfies

$$\mathbf{E}\left[\delta_{i(t)}(\mathbf{s}^t) | \mathbf{s}^t\right] \ge \Delta(\mathbf{s}^t)/\alpha,$$

then with high probability, all but $O(\frac{\alpha}{\epsilon}\log\Phi(\mathbf{s}^0))$ of the outcomes $\mathbf{s}^t$ satisfy $C(s^t) \le \frac{\rho}{1-\epsilon} \cdot C(s^*)$.

Proof. Let G be (λ, μ)-smooth. We treat $1 - \mu$ as a constant in asymptotic quantities. Call a state $\mathbf{s}^t$ *bad* if $\Delta(\mathbf{s}^t) \ge \epsilon \cdot (1-\mu) C(\mathbf{s}^t)$. Inequality (3) reduces the proof to showing that only $O(\frac{\alpha}{\epsilon}\log\Phi(\mathbf{s}^0))$ states $\mathbf{s}^t$ are bad (with high probability). Since Φ underestimates C, $\Delta(\mathbf{s}^t) \ge \epsilon \cdot (1-\mu)\Phi(\mathbf{s}^t)$ in a bad state $\mathbf{s}^t$. By assumption, $\mathbf{E}\left[\delta_{i(t)}(\mathbf{s}^t)|\mathbf{s}^t\right] \ge \frac{\epsilon(1-\mu)}{\alpha} \cdot \Phi(\mathbf{s}^t)$. Since $\delta_i(\mathbf{s}^t)$ lower bounds the improvement of a best response by i from $\mathbf{s}^t$ and Φ is a potential function, $\mathbf{E}\left[\Phi(\mathbf{s}^{t+1})|\Phi(\mathbf{s}^t)\right] \le (1 - \frac{\epsilon(1-\mu)}{\alpha})\Phi(\mathbf{s}^t)$ whenever $\mathbf{s}^t$ is a bad state. Since $\Phi(s^t)$ is decreasing in t with probability 1, standard large deviation arguments complete the proof. ◻

We note two immediate corollaries. By *maximum-gain BRD* we mean the variant in which, in each round, the player with the maximum available improvement is chosen; in *random BRD*, a player is chosen each round uniformly at random. Since $\delta_i(\mathbf{s}^t)$ denotes the benefit of switching to the particular strategy s_i^*, it lower bounds the cost decrease that i can achieve in $\mathbf{s}$. Thus maximum-gain BRD always picks a player i with $\delta_i(\mathbf{s}^t) \ge \Delta(\mathbf{s}^t)/k$, and random BRD satisfies $\mathbf{E}\left[\delta_i(\mathbf{s}^t)|\mathbf{s}^t\right] \ge \Delta(\mathbf{s}^t)/k$. By Proposition 2.6, both variants generate best-response sequences in which the approximation factor of all but $O(\frac{k}{\epsilon}\log\Phi(\mathbf{s}^0))$ outcomes are within a $1/(1-\epsilon)$ factor of the robust POA bound. (For random BRD, this guarantee holds with high probability.)

2.3 Bicriteria Bounds

We next consider cost-minimization games that, informally, remain well defined and smooth after duplicating one or more players. By a "bicriteria bound", we mean a bound on the cost of an equilibrium relative to that of an optimal outcome with additional players. A canonical application is in routing, where a comparison to an optimal solution with extra traffic translates to a comparison between additional network capacity and centralized control [29, §3.5].

Abusing notation, we write $\mathbf{s}^1 + \mathbf{s}^2$ to denote the superposition of two outcomes $\mathbf{s}^1$ and $\mathbf{s}^2$ of an underlying cost-minimization game G. We assume that cost functions are defined more generally for such superpositions. We assume that cost is superadditive, in that $C(\mathbf{s}^1+\mathbf{s}^2) \geq C(\mathbf{s}^1)+C(\mathbf{s}^2)$ for every pair $\mathbf{s}^1, \mathbf{s}^2$ of outcomes, and similarly that players' cost functions are superadditive.

Proposition 2.7 (Bicriteria Bounds) *For every cost-minimization game G that is (λ, μ)-smooth and superadditive with duplicated players, for every pure Nash equilibrium $\mathbf{s}$ and outcomes $\mathbf{s}^1, \ldots, \mathbf{s}^\ell$ of G,*

$$C(\mathbf{s}) \leq \frac{\lambda}{\ell - \mu} \cdot C(\mathbf{s}^1 + \cdots + \mathbf{s}^\ell).$$

PROOF. Suppose G is (λ, μ)-smooth. Applying the Nash equilibrium condition and adding, followed by superadditivity of players' cost functions, followed by smoothness, we can derive

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

rearranging as usual completes the proof. □

For example, consider congestion games with affine cost functions (Example 1.3), which satisfy all of the extra assumptions needed in Proposition 2.7. Since $\lambda/(2 - \mu) = 1$ for such games, the proposition implies that the cost of a pure Nash equilibrium of such a game is no more than the cost of an optimal solution with two copies of every player.

We note that a weaker bound of $\lambda/\ell(1 - \mu)$ — or more generally, the POA divided by ℓ — is easy to prove without using smoothness. This weaker bound is not sufficient, for example, to prove the preceding fact about congestion games with affine cost functions.

3. CONGESTION GAMES ARE TIGHT

This section proves that, for every fixed set of cost functions $\mathcal{C}$, congestion games with cost functions in $\mathcal{C}$ form a tight class of games in the sense of Definition 1.5. (Several variants of congestion games are also tight, as we detail in the full version.) In addition to showing that smoothness arguments always give optimal POA bounds in congestion games, this result implies the first POA bounds of any sort for congestion games with non-polynomial cost functions, and the first structural characterization of universal worst-case examples for the POA in congestion games.

Recall from Example 1.3 the definition of and notation for congestion games. Here we consider arbitrary nonnegative and nondecreasing cost functions c_e. To streamline the exposition, we assume that cost functions are strictly positive; more general results follow from extra case analysis in the proof of Theorem 3.4, below.

The worst-case POA in congestion games depends on the "degree of nonlinearity" of the allowable cost functions and is $+\infty$ unless these functions are restricted in some way. For example, for polynomial cost functions with nonnegative coefficients and degree at most d, the worst-case POA in congestion games grows exponentially with d [1, 6, 13, 26].

An ideal analysis of the POA in congestion games would solve every possible special case simultaneously, in the form of a characterization, for every set $\mathcal{C}$ of allowable cost functions, of the worst-case POA among congestion games with cost functions in $\mathcal{C}$. Such a characterization was previously known in *nonatomic* congestion games, where players are infinitesimally small: roughly, the largest-possible POA, among all nonatomic congestion games with cost functions lying in a prescribed set, is always attained in a game with only two resources and singleton strategies (i.e., a two-link network), in which one of the resources has a constant cost function. (See [30, 32] for precise statements.) This section provides the first such characterization for (atomic) congestion games. Since we use smoothness arguments to establish our POA upper bounds, they apply more generally to, for example, all of the equilibrium concepts shown in Figure 1.

Theorem 3.1 *For every non-empty set $\mathcal{C}$ of nondecreasing, positive cost functions, the set of congestion games with cost functions in $\mathcal{C}$ is tight.*

We prove Theorem 3.1 in several steps, and then conclude by pointing out important implications of the proof. Fix a set $\mathcal{C}$ of positive, nondecreasing cost functions and let $\mathcal{G}(\mathcal{C})$ denote the set of congestion games with cost functions in $\mathcal{C}$. Let $L = \{(\lambda, \mu) : \mu < 1\}$ denote the legal values for parameters λ and μ. We begin with a generic upper bound on the robust POA. Define $\gamma(\mathcal{C})$ as

$$\inf \left\{ \frac{\lambda}{1 - \mu} : (\lambda, \mu) \text{ s.t. } c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x \right\}, \tag{8}$$

where $(\lambda, \mu) \in L$ and the inequality constraints range over all integers $x \geq 0$, $x^* \geq 1$, and functions $c \in \mathcal{C}$. (We define $\gamma(\mathcal{C}) = +\infty$ if no pairs $(\lambda, \mu) \in L$ meet all of the constraints.) Intuitively, $\gamma(\mathcal{C})$ quantifies the best POA bound achievable via a local "elementwise" smoothness argument. It 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 1.3 and we record it for future reference.

Proposition 3.2 *For every set $\mathcal{C}$, 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. Consider a game in $\mathcal{G}(\mathcal{C})$ and a pair $\mathbf{x}, \mathbf{x}^*$ of its outcomes. We show that the game is (λ, μ)-smooth for every (λ, μ) in the feasible set in (8). Indeed, for every such pair we have

$$\begin{aligned}
\sum_{i=1}^{k} C_i(x_i^*, x_{-i}) &\leq \sum_{e \in E} c_e(x_e + 1)x_e^* && (9) \\
&\leq \sum_{e \in E} [\lambda c_e(x_e^*)x_e^* + \mu c_e(x_e)x_e] && (10) \\
&= \lambda \cdot C(x^*) + \mu \cdot C(x),
\end{aligned}$$

where in the first inequality we use that exactly x_e^* players ponder a deviation to a strategy x_i^* that contains e, which in turn is used by at most x_e players in x_{-i}. $\square$

In the language of Definition 1.5, the proof of Proposition 3.2 shows that every pair (λ, μ) meeting the defining constraints in (8) also belongs to the set $\mathcal{A}(\mathcal{G}(\mathcal{C}))$; as a consequence, the number $\gamma(\mathcal{C})$ upper bounds the right-hand side of (5).

The hard part of Theorem 3.1 is to show a matching lower bound on the left-hand side of (5). We now exhibit, for every set $\mathcal{C}$, a game of $\mathcal{G}(\mathcal{C})$ with (pure) POA arbitrarily close to $\gamma(\mathcal{C})$. For the time being, assume that $\mathcal{C}$ is finite and that the allowable number of players is bounded by a parameter n. Let $\mathcal{A}(\mathcal{C}, n)$ denote the pairs (λ, μ) that meet the constraints in (8) for all values of x, x^* that are at most n. The set $\mathcal{A}(\mathcal{C}, n)$ is non-empty: since cost functions are strictly positive and we are assuming that $\mathcal{C}$ is finite, for any fixed $\mu < 1$ we can take λ sufficiently large so that all of the inequalities are satisfied. Geometrically, $\mathcal{A}(\mathcal{C}, n)$ is the intersection of the halfplane L with a finite number of halfplanes, each containing everything "northeast" of a line with negative slope (Figure 2). Let $\gamma(\mathcal{C}, n)$ denote the infimum of $\lambda/(1-\mu)$ for $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n)$.

Our goal is to exhibit a congestion game of $\mathcal{G}(\mathcal{C})$ with POA equal to $\gamma(\mathcal{C}, n)$. In such a game, the inequalities in (9) and (10) must hold (essentially) with equality. This is impossible for values of $(\lambda, \mu) \in \mathcal{A}(\mathcal{C}, n)$ other than ones that determine $\gamma(\mathcal{C}, n)$; the following technical lemma identifies special properties that such optimal values possess. We give the lemma (Lemma 3.3), then our main construction (Theorem 3.4), and finally an example to illustrate the proof techniques (Example 3.5).

Lemma 3.3 *Fix finite $\mathcal{C}$ and n and suppose there exist $(\hat\lambda, \hat\mu)$ 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

$$c_j(x_j+1)x_j^* = \hat\lambda \cdot c_j(x_j^*)x_j^* + \hat\mu \cdot c_j(x_j)x_j \tag{11}$$

for $j = 1, 2$; and

$$\eta \cdot c_1(x_1+1)x_1^* + (1-\eta)\cdot c_2(x_2+1)x_2^* = \eta \cdot c_1(x_1)x_1 + (1-\eta)\cdot c_2(x_2)x_2. \tag{12}$$

Proof. We begin with preliminary observations about the candidate residences for $(\hat\lambda, \hat\mu)$. Write

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

for the halfplane corresponding to c, x, x^*, and $\partial\mathcal{H}_{c,x,x^*}$ for its boundary (points (λ, μ) that satisfy the inequality with equality). Fix c, x, x^* and define

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

this is well defined as $x^* \ge 1$ and cost functions are strictly positive. If $x \ge 1$, then the line $\partial\mathcal{H}_{c,x,x^*}$ is in general position; we can then uniquely express λ in terms of μ and derive

$$\frac{\lambda}{1-\mu} = \frac{c(x+1)}{c(x^*)}\frac{1-(\beta_{c,x,x^*})\mu}{1-\mu} \tag{13}$$

for points of the polyhedral face (i.e., line segment) $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$. Using (13), we make several assertions.

(a) If $\beta_{c,x,x^*} = 1$, then every point of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$ has the same value of $\lambda/(1-\mu)$.

(b) If $\beta_{c,x,x^*} < 1$ and $x \ge 1$, then the unique minimizer of $\lambda/(1-\mu)$ among points of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$ is the one with the minimum value of μ, and hence the maximum value of λ (i.e., the rightmost endpoint). If $\mathcal{H}_{c,x,x^*}$ determines the rightmost (infinite) face of the boundary, then the value of $\lambda/(1-\mu)$ is strictly decreasing as $\lambda \to \infty$ over the entire face.

(c) If $\beta_{c,x,x^*} > 1$, then the unique minimizer of $\lambda/(1-\mu)$ among points of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$ is the one with the maximum value of μ, and hence the minimum value of λ (i.e., the leftmost endpoint).

(d) All halfplanes of the form $\mathcal{H}_{c,0,x^*}$ are redundant, except possibly for one of the form $\mathcal{H}_{c,0,1} = \{(\lambda, \mu) : \lambda \ge 1\}$. The unique minimizer of $\lambda/(1-\mu)$ among points of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,0,1}$ is the one with $\lambda = 1$ and minimum value of μ (i.e., the bottommost endpoint).

By assumption, there is a point $(\hat\lambda, \hat\mu)$ that attains the infimum in (8). Since $\lambda/(1-\mu)$ is strictly decreasing in both λ and μ, $(\hat\lambda, \hat\mu)$ inhabits the closed, "southwestern" boundary of $\mathcal{A}(c, n)$. In particular, it belongs to some line $\partial\mathcal{H}_{c,x,x^*}$, where c, x, x^* satisfy equation (11). In the lucky event that $\beta_{c,x,x^*} = 1$, we can take $c_1 = c_2 = c$, $x_1 = x_2 = x$, and $x_1^* = x_2^* = x^*$ (and any value of $\eta \in [0,1]$) to satisfy (12) and complete the proof.

When $\beta_{c,x,x^*} \ne 1$, we claim that $(\hat\lambda, \hat\mu) \in \partial\mathcal{H}_{c',y,y^*}$ for a *second* halfplane boundary $\partial\mathcal{H}_{c',y,y^*}$. This follows from our assertions (b)–(d) and our assumption that $(\hat\lambda, \hat\mu)$ exists: if $\beta_{c,x,x^*} < 1$ ($\beta_{c,x,x^*} > 1$), $(\hat\lambda, \hat\mu)$ must be the rightmost (leftmost) endpoint of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c,x,x^*}$, and hence also the leftmost (rightmost) endpoint of $\mathcal{A}(\mathcal{C}, n) \cap \partial\mathcal{H}_{c',y,y^*}$ for some halfplane $\mathcal{H}_{c',y,y^*}$ with $\beta_{c',y,y^*} > 1$ ($\beta_{c',y,y^*} < 1$).

Relabel c, c', x, x^*, y, y^* so that $(\hat\lambda, \hat\mu)$ is the right endpoint of $\mathcal{H}_{c_1,x_1,x_1^*}$ and the left endpoint of $\mathcal{H}_{c_2,x_2,x_2^*}$. Equation (11) 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 (12) as well, completing the proof. $\square$

Theorem 3.4 *For every set $\mathcal{C}$ of cost functions, there exist congestion games with cost functions in $\mathcal{C}$ and (pure) POA arbitrarily close to $\gamma(\mathcal{C})$.*

Proof. By a limiting argument, we only need to prove that, for arbitrarily large finite sets $\mathcal{C}$ and natural numbers n, there is a congestion game with cost functions in $\mathcal{C}$ and POA equal to $\gamma(\mathcal{C}, n)$. By using standard scaling and replication tricks (as in [30, Lemma 4.8]), we also have the luxury of deploying positive scalar multiples of functions in $\mathcal{C}$.

Fix such a $\mathcal{C}$ and n, and suppose there are parameters $(\hat\lambda, \hat\mu)$ satisfying $\hat\lambda/(1-\hat\mu) = \gamma(\mathcal{C}, n)$. Choose c_1, c_2, x_1, x_2, x_1^*, x_2^*, η as in Lemma 3.3. The ground set $E_1 \cup E_2$ should be thought of as two disjoint "cycles", where each cycle has $k = \max\{x_1 + x_1^*, x_2 + x_2^*\}$ elements that are labeled from 1 to k. Elements 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 elements of E_j (for $j = 1, 2$), starting with the ith element of each cycle (wrapping around to the beginning, if necessary). Player i's second strategy Q_i uses x_j^* consecutive elements of E_j (for $j = 1, 2$), ending with the $(i-1)$th element of each cycle (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 elements $e \in E_1$, while $y_e = x_2$ and $y_e^* = x_2^*$ for elements $e \in E_2$. 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 element $e \in E_1$ (in the outcome $\mathbf{y}$).

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

$$\begin{aligned}
C_i(\mathbf{y}) &= \sum_{e \in P_i \cap E_1} \eta \cdot c_1(y_e) + \sum_{e \in P_1 \cap E_2} (1-\eta) \cdot c_2(y_e) \\
&= \eta \cdot c_1(x_1)x_1 + (1-\eta) \cdot c_2(x_2)x_2 \\
&= \eta \cdot c_1(x_1+1)x_1^* + (1-\eta) \cdot c_2(x_2+1)x_2^* \qquad (14) \\
&= \sum_{e \in Q_i \cap E_1} \eta \cdot c_1(y_e+1) + \sum_{e \in Q_i \cap E_2} (1-\eta) \cdot c_2(y_e+1) \\
&= C_i(y_i^*, y_{-i}), \qquad (15)
\end{aligned}$$

where equation (14) follows from requirement (12) in Lemma 3.3, and equation (15) follows from the disjointness of P_i and Q_i, Moreover, using (14) as a launching pad, we can derive

$$\begin{aligned}
C(\mathbf{y}) &= \sum_{i=1}^{k} C_i(\mathbf{y}) \\
&= k \cdot [\eta \cdot c_1(x_1+1)x_1^* + (1-\eta) \cdot c_2(x_2+1)x_2^*] \\
&= k\eta \cdot \left(\hat{\lambda} \cdot c_1(x_1^*)x_1^* + \hat{\mu} \cdot c_1(x_1)x_1\right) + \\
&\quad k(1-\eta) \cdot \left(\hat{\lambda} \cdot c_2(x_2^*)x_2^* + \hat{\mu} \cdot c_2(x_2)x_2\right) \qquad (16) \\
&= \hat{\lambda} \cdot k \cdot (\eta \cdot c_1(x_1^*)x_1^* + (1-\eta) \cdot c_2(x_2^*)x_2^*) + \\
&\quad \hat{\mu} \cdot k \cdot (\eta \cdot c_1(x_1)x_1 + (1-\eta) \cdot c_2(x_2)x_2) \\
&= \hat{\lambda} \cdot C(\mathbf{y}^*) + \hat{\mu} \cdot C(\mathbf{y}),
\end{aligned}$$

where (16) follows from condition (11) in Lemma 3.3. Rearranging gives a lower bound of $C(\mathbf{y})/C(\mathbf{y}^*) = \hat{\lambda}/(1-\hat{\mu}) = \gamma(\mathcal{C}, n)$ on the POA of this congestion game, completing the proof in the (common) case where the infimum in (8) is attained by a pair $(\hat{\lambda}, \hat{\mu})$.

Finally, we consider the remaining case in which $\gamma(\mathcal{C})$ is not attained by any pair (λ, μ). (This can in fact occur, for example if $\mathcal{C}$ contains only a very fast-growing function like the factorial function.) Assertion (b) in the proof of Lemma 3.3 shows that this case arises only when the rightmost (infinite) face of the boundary of $\mathcal{A}(\mathcal{C}, n)$ has β-value less than 1, in which case the infimum is approached by the value of $\lambda/(1-\mu)$ for points on this face as $\lambda \to \infty$ and $\mu \to -\infty$. This face corresponds to the line $\partial\mathcal{H}_{c,x,x^*}$ with least negative slope $-c(x^*)x^*/c(x)x$, which in turn must satisfy $x = n$ and $x^* = 1$. Using (13), in this case we have

$$\gamma(\mathcal{C}, n) = \beta_{c,n,1} \cdot \frac{c(n+1)}{c(1)} = \frac{c(n)n}{c(1)}.$$

Also, $\beta_{c,n,1} < 1$ implies that $c(n+1) > c(n)n$.

Now, define $E = \{e_1, \ldots, e_{n+1}\}$ 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, the resulting outcome has cost $(n+1) \cdot c(1)$. If players choose their non-singleton strategies, the cost is $(n+1) \cdot c(n)n$. Since $c(n+1) > c(n)n$, the latter outcome is a Nash equilibrium. The POA of this game is therefore at least $c(n)n/c(1) = \gamma(\mathcal{C}, n)$, and thus the proof is complete. □

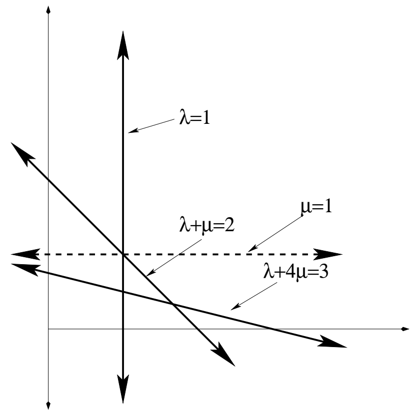

Figure 2: Example 3.5. The halfplanes that define the parameter $\gamma(\mathcal{C}, n)$.

Example 3.5 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 (8) of $\gamma(\mathcal{C}, n)$, corresponding to the two and three permitted values of x^* and x, respectively. Four of these are redundant, leaving the feasible choices of (λ, μ) constrained by the inequalities $\lambda + \mu \geq 2$ (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 $(\hat{\lambda}, \hat{\mu}) = (\frac{5}{3}, \frac{1}{3})$ and $\gamma(\mathcal{C}, n) = \frac{5}{2}$.

The proof of Theorem 3.4, specialized to this example, regenerates an construction from [13] that gives a matching lower bound on the POA of pure Nash equilibria. First observe that, in the notation of Lemma 3.3, the (unique) value of η corresponding to these two halfplanes 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 $c(x) = x$ 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. 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 switching strategies would also yield cost $\frac{2}{2} + \frac{3}{2} = \frac{5}{2}$, this outcome is a pure Nash equilibrium and shows that the POA in the game is at least 5/2.

Remark 3.6 (POA Bounds for All Cost Functions) Proposition 3.2 and Theorem 3.4 give the first characterization (namely, $\gamma(\mathcal{C})$) of the worst-case POA in congestion games with cost functions in an arbitrary set $\mathcal{C}$. Of course, precisely computing the value of $\gamma(\mathcal{C})$ is not trivial, even for simple sets $\mathcal{C}$. Arguments in [1, 26] imply 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 $\mathcal{C}$. Also, numerical work should produce good bounds on $\gamma(\mathcal{C})$ for many sets $\mathcal{C}$ of interest. In particular, computing the exact value of $\gamma(\mathcal{C}, n)$ reduces to computing the upper envelope of $O(n^2|\mathcal{C}|)$ lines.

Remark 3.7 (Worst-Case Congestion Games) The proof of Theorem 3.4 shows that congestion games comprising two parallel cycles are universal worst-case examples for the POA, no matter what the allowable set of cost functions. Such games can be realized as an atomic selfish routing (i.e., network congestion) game using a bidirected cycle (cf. [17, Figure 5.2]). This corollary is an analog of a simpler such sufficient condition for *nonatomic* congestion games where, under modest assumptions on $\mathcal{C}$, the worst-case POA is always achieved in two-node two-link networks [30].

Informally speaking, the tightness of congestion games implies that structural complexity comparable to this double-cycle structure is generally necessary to attain the worst-case POA among all congestion games with cost functions in a given set $\mathcal{C}$. (For very special sets $\mathcal{C}$, as in the last paragraph of the proof of Theorem 3.4, a single cycle suffices.) The reason is that, in light of Theorem 3.1, a congestion game is a worst-case example *only if* all of the inequalities in the derivation (2) hold with equality. This generally requires using two different types of cost function/equilibrium load/optimal load combinations (cf., Lemma 3.3).

4. REFERENCES

[1] S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *STACS '06*, pages 218–229.

[2] I. Ashlagi, D. Monderer, and M. Tennenholtz. On the value of correlation. *Journal of Artificial Intelligence*, 2009. To appear.

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

[4] R. J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1–18, 1987.

[5] B. Awerbuch, Y. Azar, A. Epstein, V. S. Mirrkoni, and A. Skopalik. Fast convergence to nearly optimal solutions in potential games. In *EC '08*, pages 264–273.

[6] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *STOC '05*, pages 57–66.

[7] A. Blum, E. Even-Dar, and K. Ligett. Routing without regret: On convergence to Nash equilibria of regret-minimizing algorithms in routing games. In *PODC '06*, pages 45–52.

[8] A. Blum, M. Hajiaghayi, K. Ligett, and A. Roth. Regret minimization and the price of total anarchy. In *STOC '08*, pages 373–382.

[9] M. Bradonjić, G. Ercal-Ozkaya, and A. Meyerson. On the price of mediation. Manuscript, 2008.

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

[11] S. Chien and A. Sinclair. Convergence to approximate Nash equilibria in congestion games. In *SODA '07*, pages 169–178.

[12] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *ESA '05*, pages 59–70.

[13] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05*, pages 67–73.

[14] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. In *SODA '02*, pages 413–420.

[15] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *STOC '04*, pages 604–612.

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

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

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

[19] M. X. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *FOCS '05*, pages 142–151.

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

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

[22] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *STACS '99*, pages 404–413.

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

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

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

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

[27] C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3), 2008. Article 14.

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

[29] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, 2002.

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

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

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

[33] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS '02*, pages 416–425.

[34] H. P. Young. *Strategic Learning and Its Limits*. Oxford University Press, 2004.