

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, 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, correlated equilibria, or to sequences of outcomes generated by natural experimentation strategies, such as simultaneous regret-minimization.

We prove 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 no-regret sequence of joint repeated play. Third, we prove that in routing games, smoothness arguments are “complete” in a proof-theoretic sense: despite their automatic generality, they are guaranteed to produce an optimal worst-case upper bound on the price of anarchy.

1. INTRODUCTION

Every student of game theory learns early and often that *equilibria are inefficient* — self-interested behavior by autonomous decision-makers generally leads to an outcome inferior to the one that a hypothetical benevolent dictator would choose. Such inefficiency is ubiquitous in real-world situations and arises for many different reasons: congestion externalities, network effects, mis-coordination, and so on. It can also be costly or infeasible to eliminate in many situ-

The original version of this article was published in the Proceedings of the 41st Annual ACM Symposium on Theory of Computing, May 2009.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 2012 ACM 0001-0782/08/0X00 ...\$5.00.

ations, with large networks being one obvious example. The past ten years have provided an encouraging counterpoint to this widespread equilibrium inefficiency: in a number of interesting application domains, decentralized optimization by competing individuals *provably approximates* the optimal outcome.

A rigorous guarantee of this type requires a formal behavioral model, in order 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 an *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:

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

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 Example 2.5.

The *price of anarchy (POA)* measures the suboptimality caused 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 outcome. An upper bound on the POA 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 researchers have proved remarkably good bounds on the POA in a wide range of models; see [17, 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: 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 a computationally intractable problem [9]; 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

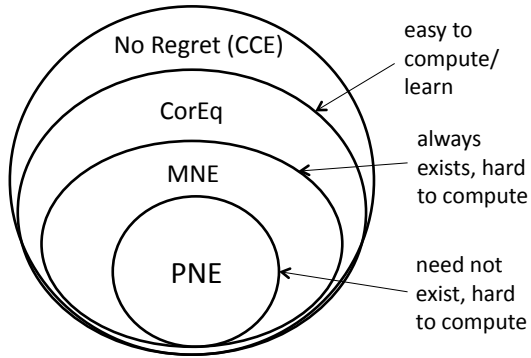


Figure 1: Generalizations of pure Nash equilibria. “PNE” stands for pure Nash equilibria; “MNE” for mixed Nash equilibria; “CorEq” for correlated equilibria; and “No Regret (CCE)” for coarse correlated equilibria, which are the empirical distributions corresponding to repeated joint play in which every player has no (external) regret.

a game.

This article presents a general theory of “robust” bounds on the price of anarchy. We focus on the hierarchy of fundamental equilibrium concepts shown in Figure 1; the full version [22] discusses additional generalizations of pure Nash equilibria, including approximate equilibria and outcome sequences generated by best-response dynamics. We formally define the equilibrium concepts of Figure 1 — mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria — in Section 3.1, but mention next some of their important properties.

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 — “Matching Pennies” being a simple example — every (finite) game has at least one mixed Nash equilibrium [16]. As a result, the “non-existence critique” for pure Nash equilibria does not apply to any of the more general concepts in Figure 1. Second, while computing a mixed Nash equilibrium is in general a computationally intractable problem [5, 8], computing a correlated equilibrium is not (see, e.g., [17, Chapter 2]). 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 a player chooses its strategy for the next time step, as a function only of its own payoffs and the history of play — that are guaranteed to converge quickly to these sets of equilibria (see [17, Chapter 4]).

1.2 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. We call such proofs “smoothness arguments”.

Many of the POA upper bounds in the literature can be recast as instantiations of this canonical method.

- (B) 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.
- (C) We prove that routing 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. Thus, in these classes of games, the worst-case POA is the same for each of the equilibrium concepts of Figure 1.

2. SMOOTH GAMES

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

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

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

$$C(\mathbf{s}) = \sum_{i=1}^k C_i(\mathbf{s}) \quad (3)$$

$$\leq \sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i}) \quad (4)$$

$$\leq \lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s}), \quad (5)$$

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 (Variations on Smoothness) Examining the three-line proof (3)–(5) reveals that the assumptions can be weakened in two ways. First, the assumption that the objective function satisfies $C(\mathbf{s}) = \sum_{i=1}^k C_i(\mathbf{s})$ can be replaced by the assumption $C(\mathbf{s}) \leq \sum_{i=1}^k C_i(\mathbf{s})$; we exploit this in Example 2.5 below. Second, in Definition 2.1, the inequality (2) only needs to hold for all outcomes \mathbf{s} and some optimal solution \mathbf{s}^* , rather than for all pairs \mathbf{s}, \mathbf{s}^* of outcomes. This relaxation is useful in some applications [4, 23].

Finally, there is an analogous definition of smooth games for maximization objectives; see Example 2.5.

2.2 Intuition

Smoothness arguments should be interpreted as 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 invoke the Nash equilibrium hypothesis again to prove the second inequality (5), or more generally to establish an upper bound using any argument that we please. Using a smoothness argument—that is, proving inequality (5) for *all* outcomes \mathbf{s} —is tantamount to discarding the Nash equilibrium hypothesis after it is used to justify the first inequality (4).

2.3 Two Examples

Concern about the range of applicability of a definition grows as its interesting consequences accumulate. Given that smoothness arguments enable the extension theorem discussed in Section 1.2, how many games can be (λ, μ) -smooth with interesting values of λ, μ ? 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. More generally, many but not all of the known price of anarchy bounds follow from smoothness proofs; see the full version [22] for a detailed discussion.

The first example is a special class of congestion games; Section 4 studies the general case in detail. The second example concerns Vetta’s well-studied utility games [25], and also illustrates how smoothness arguments can be defined and used in payoff-maximization games, and also with a “one-sided” variant of sum objective functions (cf., Remark 2.3).

The most common reason that a price of anarchy bound fails to qualify as a smoothness proof is that the Nash equilibrium hypothesis is invoked for a hypothetical deviation s_i^* that is a function of the other players’ equilibrium actions \mathbf{s}_{-i} . In most of these cases, it is also known that the worst-case POA of mixed Nash equilibria is strictly worse than that of pure Nash equilibria, and hence no lossless extension theorem exists.

Example 2.4 (Atomic Congestion Games) A congestion game is a cost-minimization game defined by a ground set E or resources, a set of k players with strategy sets $S_1, \dots, S_k \subseteq 2^E$, and a cost function $c_e : \mathcal{Z}^+ \rightarrow \mathcal{R}$ for each resource $e \in E$ [20]. In this article, we always assume that cost functions are non-negative and non-decreasing. 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 $\mathbf{s} = (s_1, \dots, s_k)$, 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} . For this example, we assume 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 $e \in E$.

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 [7, Lemma 1], who noted that

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

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

$$ay(z+1) + by \leq \frac{5}{3}(ay^2 + by) + \frac{1}{3}(az^2 + bz). \quad (6)$$

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

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

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 [2] and [7], where matching lower bounds were also supplied. Our extension theorem (Theorem 3.1) 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 [3, 6] subsequent to the original POA bound [2, 7].

Example 2.5 (Valid Utility Games) Our final example concerns a class of games called *valid utility games* [25]. 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

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

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) .

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 $i = 1, 2, \dots, k$. 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 outcome \mathbf{s} is defined as $\Pi(\mathbf{s}) = V(U(\mathbf{s}))$. Furthermore, the definition requires that two conditions hold: (i) for each player i , $\Pi_i(\mathbf{s}) \geq V(U(\mathbf{s})) - V(U(\emptyset, \mathbf{s}_{-i}))$ for every outcome \mathbf{s} ; and (ii) $\sum_{i=1}^k \Pi_i(\mathbf{s}) \leq \Pi(\mathbf{s})$ for every outcome \mathbf{s} . One concrete example of such a game is competitive facility location with price-taking markets and profit-maximizing firms [25].

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 [25, 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 $1, 2, \dots, i$ in \mathbf{s}^* . Applying condition (i), the submodularity of V , and the nondecreasing property of V yields

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

as desired. This smoothness argument implies a lower bound of $1/2$ on the POA of pure Nash equilibria in every valid utility game with a nondecreasing objective function — a result first proved in [25], along with a matching upper bound. Our extension theorem shows that this lower bound applies more generally to all of the equilibria depicted in Figure 1, a fact first established in [3].

3. AN EXTENSION THEOREM

This section states and proves the extension theorem discussed in Section 1.2: every POA bound on pure Nash equilibria derived 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. Several less direct consequences of smoothness arguments are discussed in the full version [22]. We work with cost-minimization games, though similar results hold for smooth payoff-maximization games (cf., Example 2.5).

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

3.1 Static Equilibrium Concepts

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 $(\sigma_1, \dots, \sigma_k)$ 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 $\sigma = \sigma_1 \times \dots \times \sigma_k$ 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 equilibrium, but every finite game has a mixed Nash equilibrium [16].

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

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s}) | s_i] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s'_i, \mathbf{s}_{-i}) | s_i] \quad (8)$$

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 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 their relative tractability. Because the set of correlated equilibria is explicitly described by a small set of linear inequalities, computing (and even optimizing over) correlated equilibria can be done in time polynomial in the size of the game (see, e.g., [17, Chapter 2]). They are also relatively “easy to learn,” as discussed in the next section.

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

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s})] \leq \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s'_i, \mathbf{s}_{-i})] \quad (9)$$

for every i and $s'_i \in S_i$. While a correlated equilibrium (8) protects against deviations by players aware of their recommended strategy, a coarse correlated equilibrium (9) 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 are “even easier” to compute and learn, and are thus a still more plausible prediction for the realized play of a game.

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 3.1 (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 ,

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

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

$$\mathbf{E}_{\mathbf{s} \sim \sigma}[C(\mathbf{s})] = \mathbf{E}_{\mathbf{s} \sim \sigma} \left[\sum_{i=1}^k C_i(\mathbf{s}) \right] \quad (10)$$

$$= \sum_{i=1}^k \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(\mathbf{s})] \quad (11)$$

$$\leq \sum_{i=1}^k \mathbf{E}_{\mathbf{s} \sim \sigma}[C_i(s_i^*, \mathbf{s}_{-i})] \quad (12)$$

$$= \mathbf{E}_{\mathbf{s} \sim \sigma} \left[\sum_{i=1}^k C_i(s_i^*, \mathbf{s}_{-i}) \right] \quad (13)$$

$$\leq \mathbf{E}_{\mathbf{s} \sim \sigma}[\lambda \cdot C(\mathbf{s}^*) + \mu \cdot C(\mathbf{s})] \quad (14)$$

$$= \lambda \cdot C(\mathbf{s}^*) + \mu \cdot \mathbf{E}_{\mathbf{s} \sim \sigma}[C(\mathbf{s})], \quad (15)$$

where equality (10) follows from the definition of the objective function, equalities (11), (13), and (15) follow from linearity of expectation, inequality (12) follows from the definition (9) of a coarse correlated equilibrium (applied once per player, with the hypothetical deviation s_i^*), and inequality (14) follows from the assumption that the game is (λ, μ) -smooth. Rearranging terms completes the proof. \square

3.2 Repeated Play and No-Regret Sequences

The extension theorem (Theorem 3.1) applies equally well to certain outcome sequences generated by repeated play, because of a well-known correspondence between such sequences and static equilibrium concepts. To illustrate this point, consider a sequence $\mathbf{s}^1, \mathbf{s}^2, \dots, \mathbf{s}^T$ of outcomes of a (λ, μ) -smooth game and a minimum-cost outcome \mathbf{s}^* of the game. For each i and t , define

$$\delta_i(\mathbf{s}^t) = C_i(\mathbf{s}^t) - C_i(s_i^*, \mathbf{s}_{-i}^t) \quad (16)$$

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, $\delta_i(\mathbf{s}^t)$ cannot be positive; 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

$$C(\mathbf{s}^t) \leq \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{\sum_{i=1}^k \delta_i(\mathbf{s}^t)}{1-\mu} \quad (17)$$

for each t .

Suppose that every player i experiences *vanishing average (external) regret*, meaning that its cost over time is competitive with that of every time-invariant strategy:

$$\sum_{t=1}^T C_i(\mathbf{s}^t) \leq \left[\min_{s_i'} \sum_{t=1}^T C_i(s_i', \mathbf{s}_{-i}^t) \right] + o(T). \quad (18)$$

Repeating the same pure Nash equilibrium over and over again yields a degenerate example, but in general such sequences can exhibit highly oscillatory behavior over arbitrarily large time horizons (see e.g. [3, 13]).

Averaging (17) over the T time steps and reversing the order of the resulting double summation yields

$$\frac{1}{T} \sum_{t=1}^T C(\mathbf{s}^t) \leq \frac{\lambda}{1-\mu} \cdot C(\mathbf{s}^*) + \frac{1}{1-\mu} \sum_{i=1}^k \left(\frac{1}{T} \sum_{t=1}^T \delta_i(\mathbf{s}^t) \right). \quad (19)$$

Recalling from (16) 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 (18) 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 (19) implies that the average cost of outcomes in the sequence is no more than the robust POA times the minimum-possible cost, plus an error term that approaches zero as $T \rightarrow \infty$.

Theorem 3.2 (Extension Theorem — Repeated Version)

For every cost-minimization game G with robust POA $\rho(G)$, every outcome sequence $\sigma^1, \dots, \sigma^T$ that satisfies (18) 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 $T \rightarrow \infty$.

Blum et al. [3] were the first to consider bounds of this type, calling them “the price of total anarchy.”

We reiterate the type of bound in Theorem 3.2 is significantly more compelling, and assumes much less from both the game and its participants, than one that applies only to Nash equilibria. While Nash equilibria can be intractable or impossible to find, there are several computationally efficient “off-the-shelf” learning algorithms with good convergence rates that are guaranteed, in any game, to generate outcome sequences with vanishing average regret (see e.g. [17, Chapter 4]). Of course, the guarantee in Theorem 3.2 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 3.3 (Theorems 3.1 and 3.2)

Theorems 3.1 and 3.2 are essentially equivalent, in that either one can be derived from the other. The reason is that the set of coarse correlated equilibria of a game is precisely the closure of the empirical distributions of (arbitrarily long) sequences in which every player has nonpositive average regret.

Remark 3.4 (Correlated Equilibria and Swap Regret)

There is a more stringent notion of regret — *swap regret* — under which there is an analogous correspondence between the correlated equilibria of a game and the outcome sequences in which every player has nonpositive (swap) regret. There are also computationally efficient “off the shelf” learning algorithms that guarantee each player vanishing average swap regret in an arbitrary game [10, 12].

4. CONGESTION GAMES ARE TIGHT

The worst-case POA for a set of allowable outcomes can only increase as the set grows bigger. This section proves that, in congestion games with restricted cost functions, the worst-case POA is *exactly the same* for each of the equilibrium concepts of Figure 1. We prove this by showing that smoothness arguments, despite their automatic generality, provide a tight bound on the POA, *even for pure Nash equilibria*.

More precisely, let \mathcal{G} denote a set of cost-minimization 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 $\rho_{\text{pure}}(G)$ 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 $G \in \widehat{\mathcal{G}}$, $\rho_{\text{pure}}(G) \leq \lambda/(1 - \mu)$. 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 4.1 (Tight Class of Games) A set \mathcal{G} of games is *tight* if

$$\sup_{G \in \widehat{\mathcal{G}}} \rho_{\text{pure}}(G) = \inf_{(\lambda, \mu) \in \mathcal{A}(\mathcal{G})} \frac{\lambda}{1 - \mu}. \quad (20)$$

The right-hand side of (20) 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 (20) 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 (20), 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.

Recall from Example 2.4 the definition of and notation for congestion games. Here we consider arbitrary nonnegative and nondecreasing cost functions c_e . The worst-case POA in congestion games depends on the “degree of nonlinearity” of the allowable cost functions. For example, for polynomial cost functions with nonnegative coefficients and degree at most d , the worst-case POA in congestion games is finite but exponential in d [1, 2, 7, 18].

Example 2.4 shows that, if \mathcal{G} is the set of congestion games with affine cost functions, then the right-hand side of (20) is at most $5/2$. Constructions in [2, 7] show that the left-hand side of (20) is at least $5/2$ for this class of games. Thus, congestion games with affine cost functions form a tight class. Our final result shows that this fact is no fluke.

Theorem 4.2 *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.*

In addition to showing that smoothness arguments always give optimal POA bounds in congestion games, this result and its proof imply 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.

The proof of Theorem 4.2 is technical and we provide only a high-level outline; the complete proof can be found in the full version [22]. For the following discussion, fix a set \mathcal{C} of cost functions. The first step is to use the fact that, in

See Nadav and Roughgarden [15] for a formal treatment of the duality between equilibrium concepts and POA bounds.

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 the search for parameters (λ, μ) that satisfy

$$c(x+1)x^* \leq \lambda \cdot c(x^*)x^* + \mu \cdot c(x)x \quad (21)$$

for every cost function $c \in \mathcal{C}$, non-negative integer x , and positive integer x^* . This condition is the same as (6) in Example 2.4 for the case where \mathcal{C} is the set of affine cost functions.

The second step of the proof is to understand the optimization problem of minimizing the objective function $\lambda/(1 - \mu)$ over the “feasible region” $\mathcal{A}(\mathcal{C})$, where $\mathcal{A}(\mathcal{C})$ denotes the set of values (λ, μ) that meet the condition (21) above. This optimization problem is almost the same as the right-hand side of (20), and it 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 is an infinite number of constraints (21), 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 (21).

The third and most technical part of the proof is to show a matching lower bound on the left-hand side of (20). The intuition behind the construction is to arrange 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 one player deviates unilaterally, it enjoys the benefit of fewer resources in its strategy, but each of these new resources now has load one more than that of each of the resources it was using previously. Implemented optimally, this construction produces a congestion game and a pure Nash equilibrium of it with cost a $\lambda/(1 - \mu)$ factor larger than that of the optimal outcome, where (λ, μ) are the optimal smoothness parameters identified in the second step of the proof.

Remark 4.3 (POA Bounds for All Cost Functions)

Theorem 4.2 gives the first solution to the worst-case POA in congestion games with cost functions in an arbitrary set \mathcal{C} . Of course, precisely computing the exact value of the worst-case POA is not trivial, even for simple sets \mathcal{C} . Arguments in [1, 18] imply a (complex) closed-form expression for the worst-case POA when \mathcal{C} is a set of polynomials with non-negative coefficients. Similar computations should be possible for some other simple sets \mathcal{C} . More broadly, the second and third steps of the proof of Theorem 4.2 indicate how to numerically produce good upper and lower bounds, respectively, on the worst-case POA when there is a particular set \mathcal{C} of interest.

Remark 4.4 (Worst-Case Congestion Games) The details of the construction in the third step of the proof of Theorem 4.2 show that routing games on a bidirected cycle are universal worst-case examples for the POA, no matter what the allowable set of cost functions. This corollary is an analog of a simpler sufficient condition for *nonatomic* congestion games — where there is a continuum of players, each of negligible size — in which, under modest assumptions on \mathcal{C} , the worst-case POA is always achieved in two-node two-link networks [21].

5. FURTHER RELATED WORK

The price of anarchy was first studied in [14] for makespan minimization 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 [3, 14]. See [17, Chapter 20] for a survey of the literature on this model.

The POA with a sum objective was first studied in [24] for nonatomic selfish routing games. The first general results on the POA of pure Nash equilibria for (atomic) congestion games and their weighted variants are in [2, 7], 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 [18].

Many previous works recognized the possibility of and motivation for more general POA bounds. 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 extension theorem immediately implies, and often strengthens, these previously proved robust bounds. Specifically, the authors in [1, 2, 7, 25] each observe that their upper bounds on the worst-case POA of pure Nash equilibria carry over easily to mixed Nash equilibria. In [6] 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. [3] 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 [25] and the (suboptimal) bounds of [2, 7] 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.

Versions of our two-parameter smoothness definition are implicit in a few previous papers, in each case for a specific model and without any general applications to robust POA guarantees: Perakis [19] for a nonatomic routing model with non-separable cost functions; Christodoulou and Koutsoupias [6] for congestion games with affine cost functions; and Harks [11] for splittable congestion games.

6. CONCLUSIONS

Pure-strategy Nash equilibria — where each player deterministically picks a single strategy — are often 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, even when they are guaranteed to exist; and even when efficiently computable by a centralized algorithm, they can elude natural learning dynamics.

This article presented an extension theorem, which automatically 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 only exist 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 proved 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”.

7. ACKNOWLEDGMENTS

This research was supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

8. REFERENCES

- [1] 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.
- [2] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.
- [3] 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.
- [4] I. Caragiannis, C. Kaklamani, P. Kanellopoulos, M. Kyropoulou, B. Lucier, R. P. Leme, and É. Tardos. On the efficiency of equilibria in generalized second price auctions. Submitted, 2012. Earlier versions appeared in *FOCS '10* and *EC '11*.
- [5] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of two-player Nash equilibria. *Journal of the ACM*, 56(3), 2009.
- [6] 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)*, volume 3669 of *Lecture Notes in Computer Science*, pages 59–70, 2005.
- [7] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of*

Since the conference version of this article, the definition of smooth games has been refined and extended in several ways, and new smoothness arguments have been discovered for a number of interesting models. See the full version [22] for details and references.

- the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.
- [8] 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.
- [9] 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.
- [10] D. Foster and R. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.
- [11] T. Harks. Stackelberg strategies and collusion in network games with splittable flow. In E. Bampis and M. Skutella, editors, *Proceedings of the 6th International Workshop on Approximation and Online Algorithms (WAOA'08)*, volume 5426 of *LNCS*, pages 133–146, 2008.
- [12] S. Hart and A. Mas-Colell. A simple adaptive procedure leading to correlated equilibria. *Econometrica*, 68(5):1127–1150, 2000.
- [13] 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.
- [14] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.
- [15] U. Nadav and T. Roughgarden. The limits of smoothness: A primal-dual framework for price of anarchy bounds. In *The 6th International Workshop on Internet and Network Economies (WINE)*, pages 319–326, 2010.
- [16] J. F. Nash. Equilibrium points in N -person games. *Proceedings of the National Academy of Science*, 36(1):48–49, 1950.
- [17] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [18] N. Olver. The price of anarchy and a priority-based model of routing. M.S. thesis, McGill University, 2006.
- [19] G. Perakis. The “price of anarchy” under nonlinear and asymmetric costs. *Mathematics of Operations Research*, 32(3):614–628, 2007.
- [20] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [21] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [22] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *41st ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
- [23] T. Roughgarden. The price of anarchy in games of incomplete information. In *13th ACM Conference on Electronic Commerce (EC)*, 2012.
- [24] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [25] 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.