Local Smoothness and the Price of Anarchy in Splittable Congestion Games*

Tim Roughgarden[†] Florian Schoppmann[‡]

January 26, 2012

Abstract

Congestion games have player costs that are additive over a set of resources that have anonymous cost functions, with pure strategies corresponding to certain subsets of resources. In a *splittable* congestion game, each player can choose a convex combination of subsets of resources. Congestion games have been used in several engineering, economic, and biological models. For example, splittable congestion games naturally model users of a communication network that each route traffic (possibly over multiple paths) from an origin to a destination in order to minimize the total delay of the user's traffic. The *price of anarchy* is a quantitative measure of the inefficiency of game-theoretic equilibria, such as the inefficiency caused by negative externalities in congestion games.

We characterize the worst-case price of anarchy of splittable congestion games. Our approximation guarantee is parameterized by the set of allowable resource cost functions, and degrades with the "degree of nonlinearity" of these cost functions. We prove that our guarantee is the best possible for every set of cost functions that satisfies mild technical conditions. We prove our guarantee using a novel "local smoothness" proof framework, and as a consequence the guarantee applies not only to the Nash equilibria of splittable congestion games, but also to all correlated equilibria.

1 Introduction

Congestion games play a central role in the theory of worst-case approximation guarantees for game-theoretic equilibria. They are expressive enough to capture a number of otherwise unrelated applications — including routing, network design, oligopoly models, and the migration of species [2, 13, 14, 20, 21] — yet structured enough to permit interesting theoretical guarantees. In the standard model introduced by Rosenthal [20], there is a ground set of resources, and each player selects a subset of them (e.g., a path in a network). Each resource has a univariate cost function that depends on the load induced by the players that use it, and each player strives to minimize the sum of the resources' costs in its chosen strategy (given the strategies chosen by the other players). Because of

*An extended abstract of the paper appeared in the Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, January 23-25, 2011.

†Stanford University, Department of Computer Science. Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an ONR PECASE Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship. Email: `tim@cs.stanford.edu`.

‡Stanford University, Department of Computer Science. Supported by a fellowship within the Postdoc Programme of the German Academic Exchange Service (DAAD). Email: `fschopp@uni-paderborn.de`.

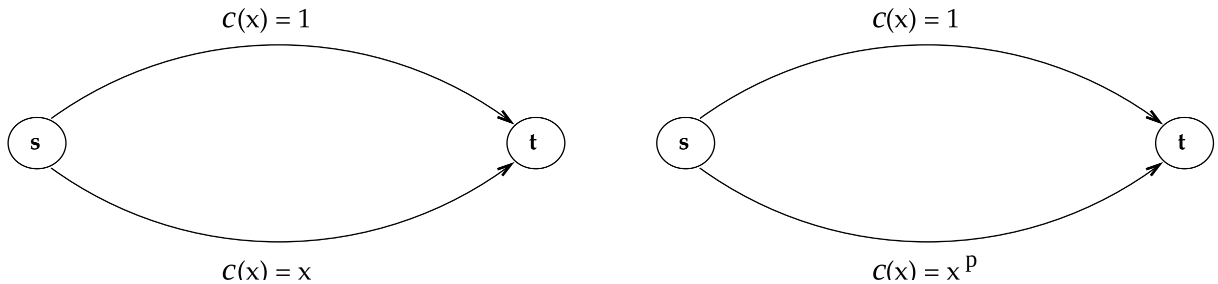

(a) Pigou's Example (b) A Nonlinear Variant

Figure 1: The price of anarchy grows with the "degree of nonlinearity" of the resource cost functions.

congestion externalities — that is, because each player ignores the extra cost its action imposes on the other players — Nash equilibria of congestion games typically do not minimize the joint cost of the players.

We study the *splittable* variant of congestion games, where each player has a *weight* w_i and a list of available strategies (each a subset of resources), and each player chooses how to split fractionally its weight over its strategies.[1] The splittable model is more appropriate than the traditional "unsplittable" model in some applications, such as multipath routing in networks. Indeed, in the computer networking literature, the splittable model was studied a decade prior to the unsplittable model (beginning with [16]).

The goal of this paper is to quantify the inefficiency of Nash equilibria in splittable congestion games. To measure inefficiency, we use the *price of anarchy (POA)* [12]: the worst-case ratio between the sum of players' costs in a Nash equilibrium and in a minimum-cost outcome. To develop intuition for the POA in congestion games, we informally review a simple example, essentially due to Pigou [18]. Consider the two-vertex, two-edge network shown in Figure 1(a). Resources correspond to edges, and strategies correspond to s-t paths. Assume that there is a very large number of players, each with negligible weight, with the total weight of all players summing to 1. Each edge is labeled with a cost function, describing the cost incurred by traffic on that edge, as a function of the sum of the weights of the players on that edge. With negligible-size players, the lower edge is a dominant strategy for every player. Thus, there is a Nash equilibrium in which the average player cost is 1. On the other hand, in an outcome where the players are split equally between the two edges, the average player cost is only $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 1 = \frac{3}{4}$. For these reasons, the POA of this game is at least $\frac{4}{3}$.

Now suppose we replace the previously linear cost function $c(x) = x$ on the lower edge with the highly nonlinear one $c(x) = x^p$ for p large (Figure 1(b)). There is still a Nash equilibrium with average cost 1. In the outcome with minimum average player cost, there is a small ϵ fraction of the players on the upper edge, and the average cost is $\epsilon + (1-\epsilon)^{p+1}$. Since this approaches 0 as ϵ tends to 0 and p tends to infinity, the POA grows without bound as p grows large.

The first point of the previous example is that Nash equilibria are suboptimal even in extremely simple splittable congestion games. Of course, there might be examples (with linear cost functions,

[1] Deterministically spreading weight over multiple strategies is *not* equivalent to probabilistically selecting a single strategy, except in the trivial case of load-independent resource cost functions.

say) with POA even larger than that in Figure 1(a) due to more complicated strategy sets or to non-negligible player weights. The second point of the example above is that the worst-case inefficiency of Nash equilibria seems to grow with the "degree of nonlinearity" of the resource cost functions. Thus, we expect an optimal upper bound on the worst-case POA of splittable congestion games to be parameterized by the set of allowable resource cost functions.

1.1 Our Results

In this paper, we resolve the worst-case price of anarchy in splittable congestion games. Prior to this work, no tight bounds on the POA in splittable congestion games were known, even for the simplest non-trivial special case of affine cost functions. In contrast, tight bounds for essentially all classes of cost functions were proved for the nonatomic model (where these is a continuum of players, like in Figure 1) and for the standard (unsplittable) model some years ago [1, 7, 24, 25]. Our bounds imply that the worst-case POA in splittable congestion games is reasonably close to 1 provided the cost functions are "not too nonlinear". The degree of nonlinearity that can be tolerated (to obey a target upper bound on the POA) is qualitatively smaller than in nonatomic congestion games, but is qualitatively larger than in standard (unsplittable) congestion games. Thus, with respect to the worst-case POA measure, *allowing non-negligible-sized players to choose fractional strategies substantially reduces inefficiency.*

Technically, we make two distinct contributions. On the upper-bound side, we define the framework of "local smoothness", which provides a sufficient condition for a game to have a bounded POA. This framework refines the smoothness paradigm introduced in [24] for games with convex strategy sets, intuitively by requiring certain inequalities only for *nearby* pairs of outcomes, rather than for *all* pairs of outcomes as in [24]. The smoothness paradigm in [24] provably cannot establish tight bounds on the POA in splittable congestion games, whereas local smoothness arguments can (as we show). Further, we prove that every POA bound derived via local smoothness applies automatically, without any quantitative degradation, to every correlated equilibrium of the game (and hence also to every mixed Nash equilibrium). Extending POA bounds to more general equilibrium concepts is important because it weakens the rationality assumptions under which the bounds are valid. An upper bound that applies only to pure Nash equilibria presumes that players reach one. A bound that applies more generally to correlated equilibria does not require players to converge to anything: if a game is played repeatedly and each player has vanishing per-step "swap regret" [9, 11], then the bound applies to their time-averaged cost.[2]

Our second contribution is a general lower bound. For a set $\mathcal{L}$ of allowable resource cost functions, we denote by $\gamma(\mathcal{L})$ the smallest upper bound on the POA that is provable via a local smoothness argument. We prove that for *every* set $\mathcal{L}$ that satisfies mild technical conditions, the worst-case POA in splittable congestion games with cost functions in $\mathcal{L}$ is *exactly* $\gamma(\mathcal{L})$. Thus, the worst-case POA of pure Nash equilibria, mixed Nash equilibria, and correlated equilibria coincide in such games.

The technical challenge in proving our lower bound stems from its generality: we need to exhibit a worst-case splittable congestion game for a set $\mathcal{L}$ of cost functions without knowing anything about $\mathcal{L}$! Our high-level approach is to exhibit an example for which all of the inequalities used in the upper bound proof are tight, in the spirit of "complementary slackness" arguments in linear programming. This goal translates to a labyrinth of restrictions on a candidate worst-case splittable congestion

[2]The blunter "smoothness framework" in [24] yields upper bounds that apply even more generally to the coarse correlated equilibria of the game (see Young [26] for a definition); this is not always the case for local smoothness proofs (Example 3.3).

game — on the allowable cost functions, on the resource loads in equilibrium and optimal outcomes, and on the relative use of a resource by different players in an equilibrium. Nevertheless, we show that all of these conditions can be met simultaneously and thus there are splittable congestion games with POA arbitrarily close to our upper bound of $\gamma(\mathcal{L})$.

Table 1 illustrates our exact bounds for the special case of bounded-degree polynomials with non-negative coefficients. (The necessary calculations are not immediately obvious and are given in Section 6.) The worst-case price of anarchy in splittable congestion games is generally strictly larger than that in nonatomic congestion games (with a continuum of players) and strictly less than that in standard (unsplittable) congestion games.

Table 1: Price of anarchy with polynomial cost functions.

Degree	**Atomic splittable**	Atomic unsplittable (weighted) [1]	Nonatomic [25]
1	**1.500**	2.618	1.333
2	**2.549**	9.909	1.626
3	**5.063**	47.82	1.896
4	**11.09**	277.0	2.151
5	**26.32**	1,858	2.394
6	**66.88**	14,099	2.630
7	**180.3**	118,926	2.858
8	**512.0**	1,101,126	3.081
d	$(\frac{1+\sqrt{d+1}}{2})^{d+1}$	$\Theta(\frac{d}{\log d})^{d+1}$	$\Theta(\frac{d}{\log d})$

1.2 Related Work

See [22, §4.8] for general references on splittable congestion games; here, we focus only on the prior research most relevant to the present work.

Splittable congestion games seem more difficult to reason about than other congestion game models. For example, it was only shown recently that Nash equilibria need not be unique in such games [3]. Splittable congestion games also exhibit counterintuitive behavior, like the fact that fusing two players into one — seemingly, increasing the amount of cooperation in the game — can increase the cost of a game's Nash equilibrium [4]. Finally, two independent proofs claimed that the worst-case price of anarchy in splittable congestion games is never worse than that in nonatomic congestion games [8, 23]. Cominetti et al. [6] showed, however, that these proofs are valid only in symmetric games — where all players have the same weight and the same set of strategies — and adapted an example in [4] to refute the general claims.

The first upper bounds on the POA in general splittable congestion games were given in [6]. These bounds are derived using a special case of our local smoothness framework in which one of our two parameters (λ) is fixed at 1. This restricted approach yields finite upper bounds on the worst-case POA only for cost functions that are polynomials with degree at most 3 and non-negative coefficients. Later Harks [10] used what we are calling local smoothness arguments to derive significantly better upper bounds. The upper bounds in [10] are equivalent to ours (as in, e.g., Table 1), but are given in a more complicated form that renders intractable the problems of constructing matching lower bounds and giving exact closed-form expressions for polynomial cost functions. Also, prior to our work, there were no known upper bounds on the POA for any equilibrium concept more general than pure Nash equilibria.

The only lower bounds known prior to the present work follow from the counterexamples in [6]. As an example comparison, for bounded-degree polynomial cost functions, our new lower bounds are exponentially larger (in the degree d) than those in [6].

1.3 Paper Organization

Section 2 formally defines splittable congestion games, the equilibrium concepts that we study, and the price of anarchy. Section 3 defines "local smoothness proofs" for games with convex strategy sets, shows that such proofs yield upper bounds on the price of anarchy of correlated equilibria, and that these upper bounds do not generally apply to all coarse correlated equilibria. Section 4 instantiates this general framework for the special case of splittable congestion games, thereby deriving a generic POA upper bound that is parameterized by the set of allowable resource cost functions. Section 5 constructs families of splittable congestion games and pure Nash equilibria in them to show that the POA upper bound in Section 4 is tight for every set of cost functions that satisfies mild technical conditions. Section 6 supplies the calculations necessary to derive closed-form expressions for the worst-case POA in splittable congestion games with resources cost functions that are polynomials with nonnegative coefficients (cf., Table 1). The Appendix simplifies and strengthens the lower bound construction of Section 5 for specific classes of allowable resource cost functions, such as monomials.

2 The Model

Splittable Congestion Games In an (atomic) splittable congestion game, a set E of *resources* has to be shared between $n \in \mathbb{N}$ players. Each resource $e \in E$ exhibits a load-dependent cost, defined by its *cost function* $\ell_e : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$. Each player $i \in [n] := \{1, \ldots, n\}$ has a set $\mathcal{P}_i \subseteq 2^E \setminus \emptyset$ of *basic strategies* available. A *fractional strategy* of player i is any distribution of its *weight* $w_i \in \mathbb{R}_{>0}$ among the basic strategies available to it, i.e., the player i's set of (fractional) strategies is $S_i := \{\vec{x}^i \in \mathbb{R}^{\mathcal{P}_i} \mid \sum_{p \in \mathcal{P}_i} x_p^i = w_i\}$. A *strategy profile* is a vector $\vec{x} = (\vec{x}^i)_{i \in [n]}$ of all players' strategies. We sometimes call a single fractional strategy a *pure* strategy.

Resource Cost Functions Following standard terminology, we say a cost function ℓ is *semi-convex* if $x \cdot \ell(x)$ is convex. For a non-decreasing function ℓ, this assumption is weaker than convexity, and is almost always satisfied in concrete applications of congestion games. In this work, we always assume that cost functions are non-decreasing, continuously differentiable, and semi-convex (the latter two conditions are necessary for a nice characterization of Nash equilibria, see below). We say that a set of cost functions $\mathcal{L}$ is *scale-invariant* if $\ell \in \mathcal{L}$ implies that $\sigma \cdot \ell(\tau \cdot x) \in \mathcal{L}$ for every $\sigma, \tau > 0$. Scale-invariance means that the set of allowable functions is invariant under changes in the units of measurement.

Load Given a strategy profile $\vec{x}$ and a resource $e \in E$, we define $x_e^i := \sum_{p \in \mathcal{P}_i : e \in p} x_p^i$ as the *load* player i puts on resource e and $x_e := \sum_{i \in [n]} x_e^i$ as the total load on e. We also use the abbreviating notation $\vec{x}_e := (x_e^i)_{i \in [n]}$.

Cost and Equilibria Given a strategy profile $\vec{x}$, the *cost of player* i is defined as $c_i(\vec{x}) := \sum_{e\in E} x_e^i \cdot \ell_e(x_e)$. We are interested in *equilibria* of the game, i.e., states where no player can reduce its (expected) cost by unilaterally deviating. To make this notion precise, we consider the following hierarchy of equilibrium concepts (see, e.g., [26] for more details and context). A *(pure) Nash equilibrium* — the most restricted concept — is a strategy profile $\vec{x}$ such that for every player i and every fractional strategy $\vec{y}^i$ it holds that $c_i(\vec{x}) \le c_i(\vec{y}^i, \vec{x}^{-i})$, where $\vec{x}^{-i}$ denotes the strategies chosen by the players other than i in $\vec{x}$. It is known that (pure) Nash equilibria always exist in splittable congestion games [16, 19]. In general, an equilibrium is a probability distribution P over the set of strategy profiles (we can regard a pure Nash equilibrium as a point mass). We say that P is a *mixed Nash equilibrium* if all players' strategies are stochastically independent and for all players i and all fractional strategies $\vec{y}^i$ of player i, we have

$$\mathrm{E}_{\vec{x}\sim P}[c_i(\vec{x})] \le \mathrm{E}_{\vec{x}\sim P}[c_i(\vec{y}^i, \vec{x}^{-i})]\,. \tag{2.1}$$

We say that P is a *correlated equilibrium* if for all players i and all functions $\delta : S_i \to S_i$ it holds that

$$\mathrm{E}_{\vec{x}\sim P}[c_i(\vec{x})] \le \mathrm{E}_{\vec{x}\sim P}[c_i(\delta(\vec{x}^i), \vec{x}^{-i})].$$

Finally, P is a *coarse correlated equilibrium* if (2.1) holds for all players i and all strategies $\vec{y}^i$. To summarize, we have the following chain of inclusions for the sets of equilibria in a game: (pure) Nash $\subseteq$ mixed Nash $\subseteq$ correlated $\subseteq$ coarse correlated.

Since cost functions are differentiable and semi-convex, a necessary and sufficient condition for a Nash equilibrium is that for every player i, its *marginal costs* for all used basic strategies are equal and at most that of every unused basic strategy. That is, for all players $i \in [n]$, and all $p, p' \in \mathcal{P}_i$ with $x_p^i > 0$ it must hold that

$$\sum_{e\in p} \ell_e^i(\vec{x}_e) \le \sum_{e\in p'} \ell_e^i(\vec{x}_e)\,,$$

where $\ell_e^i(\vec{x}_e) := \ell_e(x_e) + x_e^i \cdot \ell_e'(x_e)$. This condition can alternatively be stated as a variational inequality: For every player $i \in [n]$ and for every strategy $\vec{y}^i$ it holds that $\sum_{e\in E} \ell_e^i(\vec{x}_e)\cdot(y_e^i - x_e^i) \ge 0$. The overall measure for the quality of a strategy profile $\vec{x}$ is its *social cost*

$$\mathrm{SC}(\vec{x}) := \sum_{i\in[n]} c_i(\vec{x}).$$

By a reversal of sums, we can also write $\mathrm{SC}(\vec{x}) = \sum_{e\in E} x_e \cdot \ell_e(x_e)$.

Price of Anarchy The *price of anarchy* of an equilibrium concept in a game is the largest ratio between the (expected) social cost of an equilibrium and that of a minimum-cost strategy profile.

3 Local Smoothness

This section presents a "local" refinement of the smoothness framework in [24]. This refinement can lead to better upper bounds on the price of anarchy for games with convex strategy sets, and in particular permits optimal upper bounds for splittable congestion games. Bounds proved using local smoothness extend automatically to the correlated equilibria of a game; but in contrast to standard smoothness bounds, they do not extend to the coarse correlated equilibria of a game.

By a *cost-minimization game*, we mean a finite set of players, a strategy set S_i for each player i, and a cost function c_i for each player that maps outcomes (i.e., strategy profiles) to the nonnegative reals. Roughgarden [24] generalized several previous works [1, 5, 10, 17] by defining a (λ, μ)-*smooth* cost-minimization game as one that satisfies

$$\sum_{i=1}^{n} c_i(\vec{y}^i, \vec{x}^{-i}) \leq \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \tag{3.1}$$

for every pair $\vec{x}, \vec{y}$ of outcomes. Every coarse correlated equilibrium of a (λ, μ)-smooth game has expected cost at most $\lambda/(1-\mu)$ times the cost of an optimal outcome [24].

For the rest of this section, we consider cost-minimization games for which each strategy set S_i is a convex subset of some Euclidean space $\mathbb{R}^{m_i}$ and each cost function c_i is continuously differentiable with a bounded derivative.[3] The rough intuition behind local smoothness is to require the constraint (3.1) only for outcomes $\vec{y}$ that are "arbitrarily close to" $\vec{x}$. Since dropping constraints increases the set of feasible values for λ and μ, this idea has the potential to yield improved upper bounds.[4] Formally, we implement this idea as follows.

Definition 3.1 (Locally Smooth Games) A cost-minimization game is *locally* (λ, μ)-*smooth* with respect to the outcome $\vec{y}$ if for every outcome $\vec{x}$,

$$\begin{aligned} &\sum_{i=1}^{n} \left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i) \right] \\ &\leq \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \,. \end{aligned} \tag{3.2}$$

In Definition 3.1, $\nabla_i c_i := (\partial c_i / \partial x^i_1, \ldots, \partial c_i / \partial x^i_{m_i})$ denotes the gradient of c_i with respect to $\vec{x}^i$.

We next prove that if a game is locally (λ, μ)-smooth with respect to an optimal outcome with $\mu < 1$, then the expected cost of every correlated equilibrium (and hence every pure and mixed Nash equilibrium) is at most $\lambda/(1-\mu)$ times that of an optimal outcome.

Theorem 3.2 (Local Smoothness Bounds All Correlated Equilibria) *Let P be a correlated equilibrium of a cost-minimization game. If the game is locally (λ, μ)-smooth with respect to the outcome $\vec{y}$, then $\mathrm{E}_{\vec{x} \sim P}[\mathrm{SC}(\vec{x})] \leq \frac{\lambda}{1-\mu} \cdot \mathrm{SC}(\vec{y})$.*

Proof: The key claim is that

$$\mathrm{E}_{\vec{x} \sim P} \left[\nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i) \right] \geq 0$$

for every player i. Assuming the claim is true, we can complete the proof by using (3.2) and the

[3] The splittable congestion games that we consider, which have continuously differentiable cost functions and compact convex strategy sets, satisfy these assumptions.

[4] To see why standard smoothness arguments cannot prove optimal upper bounds on the POA of splittable congestion games, note that the strategy sets in a splittable game contain those of its unsplittable counterpart. Thus, for a fixed set of cost functions, the requirement (3.1) is only more constraining in splittable games, and the best-provable upper bound can only be larger. But, as Table 1 shows, the worst-case POA in splittable games is generally *smaller* than that in the corresponding class of unsplittable games.

linearity of expectation (twice) to derive

$$\begin{aligned}
&\mathrm{E}_{\vec{x}\sim P}[\mathrm{SC}(\vec{x})] \\
&\leq \sum_{i=1}^{n} \mathrm{E}_{\vec{x}\sim P}\left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T(\vec{y}^i - \vec{x}^i)\right] \qquad (3.3)\\
&\leq \mathrm{E}_{\vec{x}\sim P}\left[\lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x})\right]
\end{aligned}$$

and then rearrange the terms.

To prove the key claim, suppose for contradiction that $\mathrm{E}_{\vec{x}\sim P}\left[\nabla_i c_i(\vec{x})^T(\vec{y}^i - \vec{x}^i)\right] < 0$. For brevity, define $\vec{x}^{\epsilon} := ((1-\epsilon)\cdot \vec{x}^i + \epsilon \cdot \vec{y}^i, \vec{x}^{-i})$. Since strategy sets are convex, $\vec{x}^{\epsilon}$ is a well-defined strategy for every ϵ between 0 and 1. In the limit as ϵ goes to zero, $\mathrm{E}_{\vec{x}\sim P}[\frac{1}{\epsilon}(c_i(\vec{x}^{\epsilon}) - c_i(\vec{x}))]$ tends to $\mathrm{E}_{\vec{x}\sim P}[\nabla_i c_i(\vec{x})^T(\vec{y}^i - \vec{x}^i)]$, which is strictly negative by assumption.[5] Thus, there is a sufficiently small $\epsilon > 0$ such that $\mathrm{E}_{\vec{x}\sim P}[c_i(\vec{x}^{\epsilon})] < \mathrm{E}_{\vec{x}\sim P}[c_i(\vec{x})]$, which contradicts the assumption that P is a correlated equilibrium.[6] ■

Example 3.3 (Local Smoothness Does Not Bound All Coarse Correlated Equilibria) Consider the cost-minimization game defined by $N = \{1,2\}$, $S_1 = S_2 = [0,1]$, and $c_1(\vec{x}) = c_2(\vec{x}) = (x_1 - x_2)^2 + \varepsilon$, where $\varepsilon > 0$. Note that the game is an identical-interest game and has continuously differentiable convex cost functions. The sole purpose of ε is to ensure that the optimum has strictly positive social cost. Let P be a randomized strategy profile that chooses $(0,\alpha)$ and $(1, 1-\alpha)$ with equal probability, where $\alpha > 0$. Elementary calculations verify that this is a coarse correlated equilibrium with expected social cost $2\alpha^2 + 2\varepsilon$ when $\alpha \leq 1/4$. Further calculations show that for every profile $\vec{x}$ and every optimal profile $\vec{y}$ (i.e., $y_1 = y_2$) it holds that $\sum_{i=1}^{2} \nabla_i c_i(\vec{x})(y_i - x_i) = -2(x_1 - x_2)^2 = -\mathrm{SC}(\vec{x}) + \mathrm{SC}(\vec{y})$. Consequently, the game is $(1,0)$-smooth with respect to $\vec{y}$. The corresponding approximation factor of $\lambda/(1-\mu) = 1$ obviously does not apply to all coarse correlated equilibria.

4 A Locally Smooth Upper Bound

We now instantiate the local smoothness framework of Section 3 for splittable congestion games. We first need a simple observation. Define $\kappa(x,y)$ as $y^2/4$ if $x \geq y/2$ and $x(y-x)$ otherwise.

Lemma 4.1 *Let $n \in \mathbb{N}$ and $x, y \geq 0$. For all $\vec{x}, \vec{y} \in \mathbb{R}^n_{\geq 0}$ with $\sum_i x_i = x$ and $\sum_i y_i = y$ it holds that $\sum_i \left(y_i \cdot x_i - x_i^2\right) \leq \kappa(x,y)$.*

Proof: Let $x_{\max} = \max_i x_i$ and note that

$$\begin{aligned}
\sum_i \left(y_i \cdot x_i - x_i^2\right) &\leq y \cdot x_{\max} - x_{\max}^2 \\
&= \frac{y^2}{4} - \left(\frac{y}{2} - x_{\max}\right)^2 .
\end{aligned}$$

[5] This can be formally justified using the dominated convergence theorem: Since cost functions are continuously differentiable with bounded derivatives, there is some $M < \infty$ so that for all $\vec{x} \in S$ we have $|\frac{1}{\epsilon}(c_i(\vec{x}^{\epsilon}) - c_i(\vec{x}))| < M$. Hence, $\lim_{\epsilon \searrow 0} \int \frac{1}{\epsilon}(c_i(\vec{x}^{\epsilon}) - c_i(\vec{x}))\, dP(\vec{x}) = \int \nabla_i c_i(\vec{x})^T(\vec{y}^i - \vec{x}^i)\, dP(\vec{x})$.

[6] A similar trick was used by Neyman [15] to prove a rather different result, that games with convex strategy sets and a concave potential function have a unique correlated equilibrium.

■

Next is a simple univariate condition on cost functions that implies local smoothness of the corresponding class splittable congestion games.

Proposition 4.2 Let $\mathcal{L}$ be a class of allowable cost functions. If

$$y \cdot \ell(x) + \kappa(x,y) \cdot \ell'(x) \leq \lambda \cdot y \cdot \ell(y) + \mu \cdot x \cdot \ell(x) \tag{4.1}$$

for every $\ell \in \mathcal{L}$ and $x, y \geq 0$, then every splittable congestion game with cost functions in $\mathcal{L}$ is locally (λ, μ)-smooth with respect to every outcome.

Proof: In a splittable congestion game, we have for all strategy profiles $\vec{x}, \vec{y}$,

$$\begin{aligned}
& \sum_{i=1}^{n} \left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i) \right] && (4.2) \\
&= \sum_{i \in [n]} \sum_{e \in E} \left[x_e^i \cdot \ell_e(x_e) + y_e^i \cdot \ell_e^i(\vec{x}_e) - x_e^i \cdot \ell_e^i(\vec{x}_e) \right] \\
&= \sum_{e \in E} \left[y_e \cdot \ell_e(x_e) + \ell_e'(x_e) \cdot \sum_{i \in [n]} \left(y_e^i \cdot x_e^i - (x_e^i)^2 \right) \right] \\
&\leq \sum_{e \in E} \left[y_e \cdot \ell_e(x_e) + \kappa(x_e, y_e) \cdot \ell_e'(x_e) \right] && (4.3) \\
&\leq \sum_{e \in E} \left[\lambda \cdot y_e \cdot \ell_e(y_e) + \mu \cdot x_e \cdot \ell_e(x_e) \right] && (4.4) \\
&= \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \,.
\end{aligned}$$

Inequality (4.3) is due to Theorem 4.1, and inequality (4.4) follows from the hypothesis (4.1). ■

We now define the quantity $\gamma(\mathcal{L})$ as, intuitively, the best upper bound on the POA that is provable using Proposition 4.2. Formally, we first define

$$g_{\ell,x,y}(\mu) := \frac{y \cdot \ell(x) + \kappa(x,y) \cdot \ell'(x) - \mu \cdot x \cdot \ell(x)}{y \cdot \ell(y) \cdot (1 - \mu)}$$

for every cost function ℓ and values $x \geq 0, y > 0$ with $\ell(y) > 0$ (in short, we say an *admissible* triple ℓ, x, y). The constraint (4.1) is equivalent to $g_{\ell,x,y}(\mu) \leq \frac{\lambda}{1-\mu}$. Given a set of cost functions $\mathcal{L}$, we then define

$$\gamma(\mathcal{L}) := \inf_{\mu \in [0,1)} \sup_{\substack{\ell \in \mathcal{L} \\ x \geq 0, y > 0, \ell(y) > 0}} g_{\ell,x,y}(\mu) \,,$$

with the interpretation that $\sup \emptyset = 1$. Theorem 3.2, Proposition 4.2, and the definition of $\gamma(\mathcal{L})$ imply the following generic upper bound.[7]

Corollary 4.3 *For every set $\mathcal{L}$ of cost functions and every splittable congestion game with cost functions in $\mathcal{L}$, the price of anarchy of correlated equilibria is at most $\gamma(\mathcal{L})$.*

[7] We can ignore triples ℓ, x, y in which $y = 0$ or $\ell(y) = 0$ for the following reason. If $y = 0$ then inequality (4.1) is guaranteed by $\mu \geq 0$ alone. If $\ell(y) = 0$ and ξ denotes $\max \ell^{-1}(\{0\})$, then if (4.1) holds for all $y > \xi$ it also holds for $y = \xi$ (by continuity), and hence for all $y \in [0, \xi]$ (since the left-hand side of (4.1) is monotonic in y).

5 A Matching Lower Bound for All Scale-Invariant Classes of Cost Functions

We now show that for every scale-invariant set of cost functions $\mathcal{L}$, the worst-case price of anarchy of pure Nash equilibria in splittable congestion games with cost functions in $\mathcal{L}$ is *exactly* $\gamma(\mathcal{L})$.

We first need two important technical lemmas that show that $\gamma(\mathcal{L})$ can be arbitrarily approximated by the intersection of two curves $g_{\ell_1,x_1,y_1}(\mu)$ and $g_{\ell_2,x_2,y_2}(\mu)$, where one is non-increasing and one is non-decreasing.

5.1 Approximating $\gamma(\mathcal{L})$ by Two Curves $g_{\ell,x,y}$

Define $\Gamma_{\mathcal{L}} : \mathbb{R}_{<1} \to \mathbb{R} \cup \{\infty\}$ as the inner part of the infimum in the definition of $\gamma(\mathcal{L})$, i.e.,

$$\Gamma_{\mathcal{L}}(\mu) := \sup_{\substack{\ell \in \mathcal{L} \\ x \geq 0, y > 0, \ell(y) > 0}} g_{\ell,x,y}(\mu)\,.$$

Figure 3 in Section 6 provides plots of the functions $\Gamma_{\mathcal{L}}$ when $\mathcal{L}$ contains only linear functions and constants.

We saw in the Introduction that $\gamma(\mathcal{L}) < \infty$ only for restricted sets of cost functions $\mathcal{L}$. If that is the case, the following Theorem 5.1 shows that the infimum of $\Gamma_{\mathcal{L}}$ is always attained on a closed interval (hence, the infimum is in fact also a minimum). Define

$$h_{\ell,x,y} := (y - x) \cdot \ell(x) + \kappa(x, y) \cdot \ell'(x)\,.$$

A simple calculation shows that the derivative of $g_{\ell,x,y}(\mu)$ with respect to μ always has the same sign as $h_{\ell,x,y}$.

Lemma 5.1 *Let $\mathcal{L}$ be a set of cost functions such that $\gamma(\mathcal{L}) < \infty$. Then, $\Gamma_{\mathcal{L}}^{-1}(\{\gamma(\mathcal{L})\})$ is a non-empty closed interval $[s, t] \subseteq [0, 1)$. If $\mathcal{L}$ contains an unbounded function, then $s > 0$.*

Proof: For convenience, fix $\Gamma = \Gamma_{\mathcal{L}}$. The proof proceeds in several steps.

1. For all $\ell \in \mathcal{L}$ and any $\mu < 0$, we have $g_{\ell,x,1}(\mu) \xrightarrow{x\to\infty} \infty$. If ℓ is unbounded, then also $g_{\ell,x,1}(0) \xrightarrow{x\to\infty} \infty$.

2. For all $\ell \in \mathcal{L}$ and $y > x > 0$, we have $h_{\ell,x,y} > 0$ and therefore $g_{\ell,x,y}(\mu) \xrightarrow{\mu\to 1} \infty$. Hence, $\Gamma(\mu) \xrightarrow{\mu\to 1} \infty$.

3. Let $(\mu_i)_i \subset [0, 1)$ be an arbitrary sequence with $\lim_{n\to\infty} \Gamma(\mu_n) = \gamma(\mathcal{L})$. Since $(\mu_i)_i$ is bounded, we may assume w.l.o.g. (Bolzano-Weierstraß) that (μ_i) converges to some μ, where due to the previous limit argument we have $\mu \in [0, 1)$. By definition, we have $\Gamma(\mu) \geq \gamma(\mathcal{L})$. Now all $g_{\ell,x,y}$ are continuous, so $\Gamma(\mu) < \infty$ and for all $\varepsilon > 0$ there is a $\delta > 0$ so that for all $z \in (\mu - \delta, \mu + \delta)$ it holds that $\Gamma(z) \geq \Gamma(\mu) - \varepsilon$. Hence, $\Gamma(\mu) \leq \gamma(\mathcal{L})$, i.e., even equality holds. Moreover, the set $\Gamma^{-1}(\{\gamma(\mathcal{L})\})$ is non-empty and compact.

4. Suppose that $\mu_1, \mu_2 \in [0, 1)$ with $\mu_1 < \mu_2$ and $\Gamma(\mu_1) = \Gamma(\mu_2) = \gamma(\mathcal{L})$. Since all $g_{\ell,x,y}$ are monotone, it must hold that also $\Gamma(z) = \gamma(\mathcal{L})$ for all $z \in [\mu_1, \mu_2]$. Consequently, $\Gamma^{-1}(\{\gamma(\mathcal{L})\})$ is a closed interval $[s, t] \subseteq [0, 1)$, and $s > 0$ if $\mathcal{L}$ contains unbounded functions.

■

Lemma 5.2 *Let $\mathcal{L}$ be a set of cost functions. For every $\widehat{\gamma} < \gamma(\mathcal{L})$, there are $\mu < 1$ and admissible triples ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2 so that*

$$g_{\ell_1,x_1,y_1}(\mu) = g_{\ell_2,x_2,y_2}(\mu) > \widehat{\gamma} \quad \textit{and}$$
$$\operatorname{sgn}(h_{\ell_1,x_1,y_1}) = -\operatorname{sgn}(h_{\ell_2,x_2,y_2})\,.$$

Proof: Suppose first that $\gamma(\mathcal{L}) < \infty$. Due to Theorem 5.1, $\Gamma_{\mathcal{L}}$ attains its minimum $\gamma(\mathcal{L})$ on some closed interval $[s,t]$. Now note that for every $\xi \in (s,t)$ there are ℓ, x, y so that $g_{\ell,x,y}(\xi) > \widehat{\gamma}$. Consequently, there are ℓ_1, x_1, y_1 so that one of the following holds.

- $g_{\ell_1,x_1,y_1}(s) > \widehat{\gamma}$ and $h_{\ell_1,x_1,y_1} \geq 0$ (i.e., g_{ℓ_1,x_1,y_1} is non-decreasing)

  Due to continuity, for every $\delta > 0$ there are ℓ_2, x_2, y_2 so that $g_{\ell_2,x_2,y_2}(s-\delta) > \gamma(\mathcal{L})$ and $h_{\ell_2,x_2,y_2} < 0$. The claim follows by choosing δ small enough so that $g_{\ell_1,x_1,y_1}(s-\delta) \geq \widehat{\gamma}$.

- $g_{\ell_1,x_1,y_1}(t) > \widehat{\gamma}$ and $h_{\ell_1,x_1,y_1} \leq 0$ (i.e., g_{ℓ_1,x_1,y_1} is non-increasing)

  The argument for the first case can be adapted correspondingly.

In the remainder of the proof, suppose now that $\gamma(\mathcal{L}) = \infty$. Let

$$\mu^* = \sup\{\mu < 1 \mid \exists \text{ admissible triple } \ell, x, y \text{ with } g_{\ell,x,y} \geq \gamma(\mathcal{L}) \text{ and } h_{\ell,x,y} < 0\}\,.$$

Then $\mu^* \in [0,1]$ because for any $\mu < 0$ we could choose an arbitrary ℓ, a large enough x, and a small enough y so that $g_{\ell,x,y}(\mu) \geq \frac{-\mu \cdot x \cdot \ell(x)}{(1-\mu)\cdot y \cdot \ell(y)} \geq \gamma(\mathcal{L})$ and $h_{\ell,x,y} = (y-x)\cdot\ell(x) + \frac{y^2}{4}\cdot\ell'(x) < 0$. Now, the definition of μ^* ensures the following:

- There are ℓ_1, x_1, y_1 with $h_{\ell_1,x_1,y_1} \geq 0$ and $\lim_{\mu\to\mu^*} g_{\ell_1,x_1,y_1}(\mu) > \widehat{\gamma}$.

  If $\mu^* = 1$, this follows by the definition of the functions $g_{\ell,x,y}$. Otherwise, it follows from $\Gamma_{\mathcal{L}}(\mu^*) = \infty$ but the fact that $g_{\ell,x,y}(\mu^*) \leq \gamma(\mathcal{L})$ for all ℓ, x, y with $h_{\ell,x,y} < 0$.

- For every $\mu < \mu^*$, there are ℓ_2, x_2, y_2 with $h_{\ell_2,x_2,y_2} < 0$ and $g_{\ell_2,x_2,y_2}(\mu) > \widehat{\gamma}$.

The claim follows by choosing $\mu < \mu^*$ large enough. ■

5.2 The Construction

An exact lower bound requires the inequalities (3.3), (4.3), and (4.4) from the proofs of Theorem 3.2 and Proposition 4.2 to be asymptotically tight.

Motivated by Theorem 5.2, the plan for our construction is as follows. We construct a family of instances that contain only two groups of resources, one with cost function ℓ_1 and one with ℓ_2. Each instance needs to possess a Nash equilibrium $\vec{x}$ so that the load on all resources of group $i \in \{1,2\}$ is x_i, yet there must be some other strategy profile $\vec{y}$ where the load is only y_i on each resource of group i. Suppose now that $g_{\ell_i,x_i,y_i}(\mu) = \frac{\lambda}{1-\mu}$. By definition of $h_{\ell_i,x_i,y}$, we then have

$$x_i \cdot \ell_i(x_i) = \lambda \cdot y_i \cdot \ell_i(y_i) + \mu \cdot x_i \cdot \ell_i(x_i) - h_{\ell_i,x_i,y_i}\,.$$

We hence need $\operatorname{sgn}(h_{\ell_1,x_1,y_1}) = -\operatorname{sgn}(h_{\ell_2,x_2,y_2})$ and the number of resources in groups 1 and 2 to be chosen so that the sum, over all resources, of the h_{ℓ_i,x_i,y_i}-terms vanishes. Then $\frac{\mathrm{SC}(\vec{x})}{\mathrm{SC}(\vec{y})} = \frac{\lambda}{1-\mu}$ as needed.

Tightness of inequality (4.3) alone gives another constraint:

Remark 5.3 Suppose both $\vec{x}$ and $\vec{y}$ in Theorem 4.1 are sorted in descending order (without loss of generality). As $n \to \infty$, the inequality is asymptotically tight when $x_1 = \min\{\frac{y}{2}, x\}$, $x_2 = \cdots = x_n = o(1)$, and $y_1 = y$, $y_2 = \cdots = y_n = 0$. The left-hand side of the inequality is maximized when $x_2 = \cdots = x_n = \frac{x-x_1}{n-1}$ (see Cominetti et al. [6, Theorem 3.1]).

In view of this observation, for each resource e of group i there needs to be one player who in the Nash equilibrium puts load $\min\{\frac{y_i}{2}, x_i\}$ on resource e, and all other players only put an infinitesimal load on e. In the following, we show that all of the above conditions can indeed be met simultaneously.

Theorem 5.4 *Let $\lambda \in \mathbb{R}$, $\mu < 1$. Moreover, let ℓ_1, ℓ_2 be cost functions and $x_1, x_2 \geq 0$ and $y_1, y_2 > 0$. If $x_2 < \frac{y_2}{2}$ or $\ell_2'(x_2) = 0$, define $r := \ell_2(x_2) + x_2 \cdot \ell_2'(x_2)$. Otherwise, define $r := \ell_2(x_2) + \frac{y_2}{2} \cdot \ell_2'(x_2)$. Suppose that*

$$\ell_1(x_1) = \ell_2(x_2) = 1\,,$$
$$g_{\ell_1,x_1,y_1}(\mu) = g_{\ell_2,x_2,y_2}(\mu) = \tfrac{\lambda}{1-\mu}\,,\ \text{and}$$
$$h_{\ell_2,x_2,y_2} = -h_{\ell_1,x_1,y_1} \cdot r \geq 0\,.$$

Then, there is an infinite family of splittable congestion games with cost functions in $\{\sigma_1\ell_1, \ell_2 : \sigma_1 \geq 1\}$ and with limiting price of anarchy at least $\frac{\lambda}{1-\mu}$.

Proof: We construct a family of instances determined by two scaling parameters $n, p_2 \in \mathbb{N}$. All other variables (described in Theorem 2) are functions of n, p_2. For convenience, we also define $h_i := h_{\ell_i,x_i,y_i}$ for $i \in \{1,2\}$, and we use the notation $\overline{1} := 2$ and $\overline{2} := 1$. Theorem 2 illustrates our construction.

Before we start, note that the theorem's assumptions imply that $h_1 \leq 0$ and hence $x_1 > \frac{y_1}{2}$.

Table 2: Symbols used in description of lower-bound construction

Symbol	Meaning (load refers to load in Nash equilibrium)
n	number of players per group
p_i	size of "optimal" strategies in group i
q_i	size of "non-optimal" strategies in group i
t_i	number of "non-optimal" strategies for each player in group i
α_i	load each player from group i puts on its "optimal" strategy
β_i	load each player from group i puts on its "non-optimal" strategies
γ_i	load each player from group $\overline{i}$ puts on each "optimal" strategy of group i
σ_i	scaling factor for cost functions in group i

Resources There are two groups of resources, with group $i \in \{1,2\}$ consisting of $n \cdot p_i$ resources that we denote by $(i,0), \ldots, (i, n \cdot p_i - 1)$. A good intuition is to think of two circles. Resources in group i have the cost function $\sigma_i \cdot \ell_i$, where σ_1 will be determined later and $\sigma_2 := 1$.

Players and Strategies There will be two groups of atomic players, with group $i \in \{1, 2\}$ consisting of n players denoted by $(i, 0), \ldots, (i, n-1)$. Each player (i, j) has one "optimal" strategy $\mathcal{P}_{i,j,0}$ available. If $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, player (i, j) has also $t_i := \frac{p_i \cdot (n-1)}{q_i}$ "non-optimal" strategies $\mathcal{P}_{i,j,1}, \ldots, \mathcal{P}_{i,j,t_i}$ available. Finally, players from group 2 can also choose all "optimal" strategies in group 1, i.e., $\mathcal{P}_{1,0,0}, \ldots, \mathcal{P}_{1,n-1,0}$. Formally:

$$\begin{aligned} \mathcal{P}_{i,j,0} &:= \{(i, j \cdot p_i), \ldots, (i, (j+1) \cdot p_i - 1)\}\,, \text{ and} \\ \mathcal{P}_{i,j,k} &:= \{(i, (j+1) \cdot p_i + (k-1) \cdot q_i), \ldots, \\ &\qquad (i, (j+1) \cdot p_i + k \cdot q_i - 1)\} \text{ for } k \geq 1\,. \end{aligned}$$

The weight of each player in group i will be $w_i := \alpha_i + t_i \cdot \beta_i + n \cdot \gamma_{\bar{i}}$, where $\gamma_1 := \frac{-h_1}{n}$ and $\gamma_2 := 0$ (by construction, players from group 1 cannot use any resources in group 2).

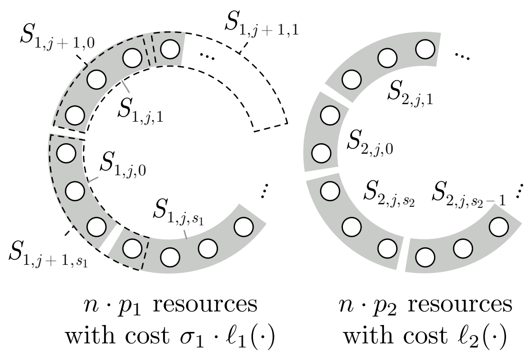

Figure 2: Illustration of construction with $p_1 = 3$, $q_1 = 4$ and $p_2 = 2$, $q_2 = 3$

The Equilibrium Consider the strategy profile where each atomic player (i, j) uses strategy $\mathcal{P}_{i,j,0}$ with load α_i and each of the strategies $\mathcal{P}_{i,j,1}, \ldots, \mathcal{P}_{i,j,t_i-1}$ with load β_i (if $x_i < \frac{y_i}{2}$ or $\ell_i'(x_i) = 0$ then β_i is necessarily 0). Moreover, each player in group 2 also uses each of the n "optimal" strategies in group 1 with load γ_1. In the following, we show that we can choose values for the variables from Theorem 2 so that the following six conditions are satisfied:

1. The load on each resource of group i is exactly x_i. That is,
$$\begin{aligned} \alpha_i + (n-1) \cdot \beta_i + n \cdot \gamma_i = x_i\,, &\quad \text{i.e.,} \\ \beta_i = \frac{x_i - \alpha_i - n \cdot \gamma_i}{n-1}\,. & \end{aligned} \tag{5.1}$$

2. Each player is faced with equal marginal costs for all its strategies. The first condition for players in group 2 is
$$\begin{aligned} & p_1 \cdot \sigma_1 \cdot \big(\ell_1(x_1) + \gamma_1 \cdot \ell_1'(x_1)\big) \\ & = p_2 \cdot \sigma_2 \cdot \big(\ell_2(x_2) + \alpha_2 \cdot \ell_2'(x_2)\big)\,. \end{aligned} \tag{5.2}$$

   Moreover, if $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, then also (for both groups)
$$\begin{aligned} & p_i \cdot \big(\ell_i(x_i) + \alpha_i \cdot \ell_i'(x_i)\big) \\ & = q_i \cdot \big(\ell_i(x_i) + \beta_i \cdot \ell_i'(x_i)\big)\,. \end{aligned} \tag{5.3}$$

3. If $\ell_i'(x_i) > 0$, then for each resource in group i there is one player who puts load $\min\{\frac{y_i}{2}, x_i\} \pm o(1)$ on it whereas all other players put load $o(1)$ on it.

   If $i = 2$ and $x_2 \le \frac{y_2}{2}$, there is nothing to show because $\alpha_2 = x_2$. Otherwise, plugging $\ell_i'(x_i) = \frac{4(x_i - y_i + h_i)}{y_i^2}$ and (5.1) into (5.3) yields

$$\alpha_i = \left[\frac{y_i^2 \cdot \left(\frac{q_i}{p_i} - 1\right)}{4 \cdot (x_i - y_i + h_i)} + \frac{q_i \cdot (x_i - n \cdot \gamma_i)}{(n-1) \cdot p_i} \right] \cdot \left[1 + \frac{q_i}{(n-1) \cdot p_i} \right]^{-1} .$$

   Now $\alpha_i \xrightarrow{n, p_2 \to \infty} \frac{y_i}{2}$ and $\beta_i \xrightarrow{n, p_2 \to \infty} 0$ are fulfilled if

$$\frac{q_i}{p_i} \xrightarrow{p_2 \to \infty} \frac{2x_i - y_i + 2h_i}{y_i}, \quad \text{i.e.,}$$
$$q_i := \left\lfloor p_i \cdot \frac{2x_i - y_i + 2h_i}{y_i} \right\rfloor . \tag{5.4}$$

4. In the "optimal" strategy profile, where each player from group i only uses its "optimal" strategy in group i, the load on any resource in group i is $y_i + o(1)$.

   We need $w_i \xrightarrow{n, p_2 \to \infty} y_i$. First note that if $x_2 \ge \frac{y_2}{2}$ and $\ell_2'(x_2) > 0$, then

$$\begin{aligned} n \cdot \gamma_1 = -h_1 = \frac{h_2}{r} &= \frac{h_2 \cdot y_2}{y_2 + \frac{y_2^2}{2} \cdot \ell_2'(x_2)} \\ &= \frac{y_2}{2} \cdot \frac{2h_2}{2x_2 - y_2 + 2h_2} . \end{aligned} \tag{5.5}$$

   If, on the other hand, $x_2 \le \frac{y_2}{2}$ or $\ell_2'(x_2) = 0$, then

$$\begin{aligned} n \cdot \gamma_1 = -h_1 &= \frac{h_2}{r} \\ &= \frac{(y_2 - x_2) \cdot (\ell_2(x_2) + x_2 \cdot \ell_2'(x_2))}{\ell_2(x_2) + x_2 \cdot \ell_2'(x_2)} \\ &= y_2 - x_2 . \end{aligned}$$

   Now, consider $i \in \{1, 2\}$.

   - If $x_i \ge \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, then we have due to (5.4) and (5.5) that

$$\begin{aligned} w_i &= \alpha_i + t_i \cdot \beta_i + n \cdot \gamma_{\bar{i}} \\ &= \alpha_i + \frac{p_i}{q_i} \cdot (x_i - \alpha_i - n \cdot \gamma_i) + n \cdot \gamma_{\bar{i}} \\ &\xrightarrow{n, p_2 \to \infty} \frac{y_i}{2} \cdot \left(1 + \frac{2x_i - y_i - 2n \cdot \gamma_i}{2x_i - y_i + 2h_i}\right) + n \cdot \gamma_{\bar{i}} \\ &= y_i . \end{aligned}$$

- If $\ell'_1(x_1) = 0$, then $w_1 = \alpha_1 = x_1 - n \cdot \gamma_1 = x_1 + h_1 = x_1 + (y_1 - x_1) = y_1$.
- If $x_2 \le \frac{y_2}{2}$ or $\ell'_2(x_2) = 0$, then $w_2 = \alpha_2 + n \cdot \gamma_1 = x_2 + (y_2 - x_2) = y_2$.

5. The social cost is $(\frac{\lambda}{1-\mu} - o(1))$ times as in the "optimal" strategy profile.

   The social cost contributed by any resource in group i in the equilibrium is

$$\begin{aligned} &\sigma_i \cdot x_i \cdot \ell_i(x_i) \\ &= \sigma_i \cdot (\lambda \cdot y_i \cdot \ell_i(y_i) + \mu \cdot x_i \cdot \ell_i(x_i) - h_i) , \end{aligned}$$

   where the equality is due to $g_{\ell_i,x_i,y_i}(\mu) = \frac{\lambda}{1-\mu}$ and the definition of h_i.

   Let $\Phi(n,p_2) := \sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot y_i \cdot \ell_i(y_i)$ be a lower bound on the social cost in the "optimal" strategy profile, which we denote by $\mathrm{SC}^{\mathrm{opt}}(n,p_2)$. Note that $\Phi(n,p_2) \ge \mathrm{SC}^{\mathrm{opt}}(n,p_2)$ because, in the "optimal" strategy profile, the actual load on each resource of group i is larger than y_i. Moreover, let $\Lambda(n,p_2)$ be the sum of the $(\sigma_i \cdot h_i)$-terms, over all resources. That is, $\Lambda(n,p_2) := \sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot h_i$.

   We want that

$$\frac{\mathrm{SC}^{\mathrm{eq}}(n,p_2)}{\Phi(n,p_2)} \xrightarrow{n,p_2 \to \infty} \frac{\lambda}{1-\mu} .$$

   Since $\frac{\Phi(n,p_2)}{\mathrm{SC}^{\mathrm{opt}}(n,p_2)} \xrightarrow{n,p_2 \to \infty} 1$ (as $w_i \xrightarrow{n,p_2 \to \infty} y_i$), we then also have

$$\frac{\mathrm{SC}^{\mathrm{eq}}(n,p_2)}{\mathrm{SC}^{\mathrm{opt}}(n,p_2)} \xrightarrow{n,p_2 \to \infty} \frac{\lambda}{1-\mu}$$

   as needed.

   W.l.o.g., we may assume that $\Lambda(n,p_2) \le 0$ because we are done otherwise. We are also done if $\Lambda(n,p_2) \xrightarrow{n,p_2 \to \infty} 0$, which motivates us to set $p_1 := \lceil p_2 \cdot r \rceil$. We have

$$\begin{aligned} &\Phi(n,p_2) \\ &\ge n \cdot p_2 \cdot (r \cdot \sigma_1 \cdot y_1 \cdot \ell_1(y_1) + \sigma_2 \cdot y_2 \cdot \ell_2(y_2)) \end{aligned}$$

   and

$$|\Lambda(n,p_2)| \le n \cdot p_2 \cdot h_2 \cdot |\sigma_2 - \sigma_1| .$$

   Consequently, $\frac{\Lambda(n,p_2)}{\Phi(n,p_2)} \xrightarrow{n,p_2 \to \infty} 0$ provided that $\sigma_1 \xrightarrow{n,p_2 \to \infty} 1$. (Recall that always $\sigma_2 = 1$.)

   Let $\mathrm{SC}^{\mathrm{eq}}(n,p_2) := \sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot x_i \cdot \ell_i(x_i)$ be the social cost in the equilibrium. By definition,

$$\begin{aligned} &\mathrm{SC}^{\mathrm{eq}}(n,p_2) \\ &= \lambda \cdot \Phi(n,p_2) + \mu \cdot \mathrm{SC}^{\mathrm{eq}}(n,p_2) + \Lambda(n,p_2) , \end{aligned}$$

   i.e.,

$$\begin{aligned} \frac{\mathrm{SC}^{\mathrm{eq}}(n,p_2)}{\Phi(n,p_2)} &= \frac{\lambda}{1-\mu} + \frac{\Lambda(n,p_2)}{\Phi(n,p_2) \cdot (1-\mu)} \\ &\xrightarrow{n,p_2 \to \infty} \frac{\lambda}{1-\mu} . \end{aligned}$$

6. All parameters are feasible, i.e.,

$$n, p_i, q_i, t_i \in \mathbb{N}, \quad \alpha_i, \beta_i, \gamma_i \geq 0, \quad \sigma_i > 0\,.$$

We argue that all six conditions can indeed be satisfied simultaneously: Suppose the scaling parameters $n, p_2 \in \mathbb{N}$ are given. This determines p_1, q_i, α_i, and β_i (in this order) according to conditions 5, 3, and 1, respectively, and then σ_1 according to (5.2) of condition 2. Now, conditions 1–3 imply also condition 4, as shown above.

By the theorem's assumptions and our definitions, we have $\sigma_i \xrightarrow{n,p_2\to\infty} 1$ so that also condition 5 is met. Finally, we need to show condition 6: Since we may assume without loss of generality that $p_i \cdot (n-1)$ is a multiple of q_i (because p_i and q_i do not depend on n), we have $t_i \in \mathbb{N}$. The only other non-obvious part of condition 6 is that $\beta_1 \geq 0$. This holds because $x_1 - \alpha_1 \xrightarrow{n,p_2\to\infty} x_1 - \frac{y_1}{2}$ and $n \cdot \gamma_1 = -h_1 = x_1 - y_1 - \frac{y_1^2 \cdot \ell_1'(x_1)}{4} < x_1 - \frac{y_1}{2}$. This verifies the construction and completes the proof. ■

We now have everything to formally state the main result of this section:

Corollary 5.5 *Let $\mathcal{L}$ be a scale-invariant set of cost functions. Then, the worst-case price of anarchy in atomic splittable congestion games with cost functions in $\mathcal{L}$ is exactly $\gamma(\mathcal{L})$.*

Proof: The upper bound is due to Theorem 4.3. For the lower bound, it suffices to show that for any two triples ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2 "output" by Theorem 5.2 we can find triples $\widehat{\ell}_1, \widehat{x}_1, \widehat{y}_1$ and $\widehat{\ell}_2, \widehat{x}_2, \widehat{y}_2$ with that we can feed the lower-bound construction of Theorem 5.4 and that induce the same functions $g_{\ell,x,y}$.

We start with a simple observation. Let ℓ be a cost function and $\sigma, \tau > 0$. Define $\widehat{\ell}(x) := \sigma \cdot \ell(\tau \cdot x)$, which belongs to $\mathcal{L}$ by scale-invariance. Then, $\widehat{\ell}'(x) = \sigma \cdot (\ell(\tau \cdot x))' = \sigma \cdot \tau \cdot \ell'(\tau \cdot x)$. Consequently, $g_{\widehat{\ell},x,y} = g_{\ell,\tau\cdot x,\tau\cdot y}$ and $\tau \cdot h_{\widehat{\ell},x,y} = \sigma \cdot h_{\ell,\tau\cdot x,\tau\cdot y}$.

We can assume that $\ell_i(x_i) > 0$ because otherwise $g_{\ell_i,x_i,y_i} = 0$. This cannot happen provided we use $\widehat{\gamma} > 1$ in Theorem 5.2. Consequently, we can set $\widehat{\ell}_2(x) := \frac{1}{\ell_2(x_2)} \cdot \ell_2(x)$, $\widehat{x}_2 = x_2$, $\widehat{y}_2 = y_2$ and $\widehat{\ell}_1(x) := \frac{1}{\ell_1(x_1)} \cdot \ell_1(\tau \cdot x)$, $\widehat{x}_1 = \frac{x_1}{\tau}$, $\widehat{y}_1 = \frac{y_1}{\tau}$. Here, we have the freedom to choose τ as needed. ■

Remark 5.6 Since each player's basic strategies in the lower-bound construction are disjoint, the construction can be transformed easily into a (directed) network congestion game — orient both circles, give each player its own source and sink vertices (outside the circles), and paths corresponding to its basic strategies in the construction above.

Remark 5.7 A "lucky case" in the lower-bound construction occurs when $h_{\ell_i,x_i,y_i} = 0$, i.e., when g_{ℓ_i,x_i,y_i} is a constant. Then, only one circle of resources is needed and the scale-invariance hypothesis can be dropped. Theorem A.1 gives an example. When $h_{\ell_i,x_i,y_i} = 0$, a price of anarchy of $\gamma(\mathcal{L})$ can even arise with singleton strategies (this, however, only with scale-invariance). Section A.2 details the construction. The lucky case occurs, e.g., when $\mathcal{L}$ contains only monomials (see Section 6).

6 Polynomials

This section determines the exact price of anarchy — that is, evaluates the parameter $\gamma(\mathcal{L})$ — when the cost functions are polynomials of degree at most $d \in \mathbb{N}$ and with non-negative coefficients. We therefore introduce the following notation: For $d \in \mathbb{N}$, let $\mathcal{P}_d$ denote this set of cost functions (with degree bound d). Moreover, we write X^d to denote the monomial function $x \mapsto x^d$ (where X^0 is the constant function 1), and we let $\mathcal{M}_d := \{X^d, X^{d-1}, \dots, X^0\}$ be the set of all monomials of degree at most d. We define Ψ_d as the unique positive real x with $x^d + \frac{d \cdot x^{d-1}}{4} = x^{d+1}$, i.e., $\Psi_d := \frac{1}{2}(1+\sqrt{d+1})$. To save work, we define $g^*_{\ell,x,y}$, $h^*_{\ell,x,y}$, $\Gamma^*_{\mathcal{L}}$, and $\gamma^*(\mathcal{L})$ as in Theorem 4 and Theorem 5.1, respectively, but replace $\kappa(x,y)$ by $\frac{y^2}{4}$. We start with three lemmas to simplify $\gamma^*(\mathcal{P}_d)$. In the end, it will turn out that $\gamma(\mathcal{P}_d) = \gamma^*(\mathcal{P}_d)$.

Lemma 6.1 *Let $\mu \in (0,1)$ and $d \geq 1$. Define $g : \mathbb{R}_{\geq 0} \to \mathbb{R}$, $g(x) := x^d + \frac{d \cdot x^{d-1}}{4} - \mu \cdot x^{d+1}$. Then, it holds that g has exactly one global maximum, at*

$$\xi = \frac{d + \sqrt{d^2 + d \cdot \mu \cdot (d^2 - 1)}}{2\mu \cdot (d+1)} .$$

Moreover, ξ is the only local extremum on $\mathbb{R}_{>0}$.

Proof: We first show that $x = 0$ is not a global maximum. If $d = 1$, then $g(\frac{1}{2\mu}) = \frac{1}{2} + \frac{1}{4\mu} > \frac{1}{4} = g(0)$. If $d > 1$, then $g(\Psi_d) = (1-\mu) \cdot \Psi_d > 0 = g(0)$. Now since $\lim_{x \to \infty} g(x) = -\infty$, g is continuous, and we know that g attains values strictly larger than $g(0)$ somewhere on $\mathbb{R}_{>0}$, it suffices to show that there is a unique local extremum on $\mathbb{R}_{>0}$. For $x > 0$, we have as necessary first-order condition for a local extremum that

$$g'(x) = dx^{d-2}\left(x + \frac{d-1}{4}\right) - \mu(d+1)x^d = 0 . \tag{6.1}$$

Indeed, ξ is the unique real value for x that satisfies (6.1). ■

Remark 6.2 Note that Theorem 6.1 gives a closed form for the function $\mu \mapsto \sup_{x \geq 0} g^*_{X^d,x,1}(\mu)$, which is the envelope function of all $g^*_{X^d,x,y}$. In Theorem 3, this curve is shown as a thick line for the case $d = 1$. In detail, the envelope function is here

$$\mu \mapsto \frac{1+\mu}{4 \cdot \mu \cdot (1-\mu)} .$$

Lemma 6.3 *Let $d \in \mathbb{N}$. Then,*

$$\gamma^*(\mathcal{P}_d) = \gamma^*(\mathcal{M}_d) = \min_{\mu \in (0,1)} \max_{\substack{\ell \in \mathcal{M}_d \\ x \geq 0}} g^*_{\ell,x,1}(\mu) .$$

Proof: We can rewrite

$$\gamma(\mathcal{P}_d) = \inf_{(\lambda,\mu) \in \mathbb{R} \times (0,1)} \left\{ \frac{\lambda}{1-\mu} \,\middle|\, \forall \ell \in \mathcal{P}_d, x \geq 0, y > 0 : \lambda \geq \frac{y \cdot \ell(x) + \frac{y^2 \cdot \ell'(x)}{4} - \mu \cdot x \cdot \ell(x)}{y \cdot \ell(y)} \right\} . \tag{6.2}$$

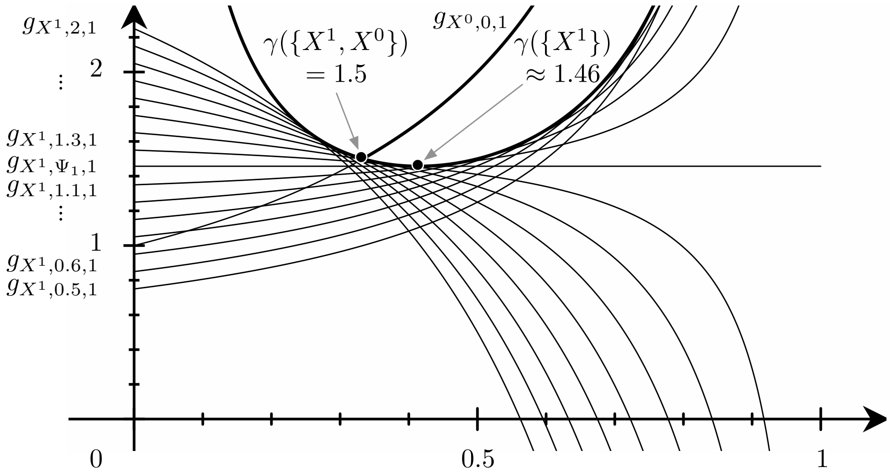

Figure 3: Functions $g_{\ell,x,y}$ when ℓ is the identity or a constant function.

Now observe that the condition

$$\forall \ell \in \mathcal{P}_d, x \geq 0, y > 0 : \lambda \geq \frac{y \cdot \ell(x) + \frac{y^2 \cdot \ell'(x)}{4} - \mu \cdot x \cdot \ell(x)}{y \cdot \ell(y)} \tag{6.3}$$

is fulfilled if and only if the inequality holds for all monomials. Moreover, when ℓ is constant, the inequality boils down to $\lambda \geq 1 - \mu \cdot \frac{x}{y}$. Consequently, (6.3) is equivalent to

$$\forall r \in [d], x \geq 0, y > 0 : \lambda \geq \frac{y \cdot x^r + \frac{y^2 \cdot r \cdot x^{r-1}}{4} - \mu \cdot x^{r+1}}{y^{r+1}} \quad \text{and} \quad \lambda \geq 1\,. \tag{6.4}$$

Now the first condition in (6.4) is equivalent to

$$\forall r \in [d], x \geq 0 : \lambda \geq x^r + \frac{r \cdot x^{r-1}}{4} - \mu \cdot x^{r+1}\,. \tag{6.5}$$

Consequently,

$$\begin{aligned}
\gamma^*(\mathcal{P}_d) &= \inf_{(\lambda,\mu)\in\mathbb{R}\times(0,1)} \left\{ \frac{\lambda}{1-\mu} \,\middle|\, \forall r \in [d], x \geq 0 : \right. \\
&\qquad\qquad \left. \frac{\lambda}{1-\mu} \geq g_{X^r,x,1}(\mu) \text{ and } \frac{\lambda}{1-\mu} \geq g_{X^0,0,1}(\mu) \right\} \\
&= \gamma^*(\mathcal{M}_d) \\
&= \min_{\mu\in(0,1)} \max_{\substack{\ell\in\mathcal{M}_d \\ x\in\mathbb{R}_{\geq 0}}} g^*_{\ell,x,1}(\mu)\,.
\end{aligned}$$

Here the last equality can be seen as follows: Theorem 6.1 implies that $\gamma(\mathcal{P}_d) < \infty$ and that for each $\mu \in (0,1)$ the supremum $\Gamma^*_{\mathcal{P}_d}(\mu)$ is attained by some $g^*_{\ell,x,1}(\mu)$. Moreover, Theorem 5.1 implies that the infimum $\gamma^*(\mathcal{P}_d)$ is indeed attained by some $\Gamma_{\mathcal{P}_d}(\mu)$ with $\mu \in (0,1)$. ■

Lemma 6.4 *Let $d \in \mathbb{N}$. Then:*

1. $\gamma^*(\{X^d\}) = \Psi_d^{d+1}$.
2. $\gamma^*(\{X^1, X^0\}) = \frac{3}{2}$. *If $d \geq 2$, then* $\gamma^*(\{X^d, X^0\}) = \gamma^*(\{X^d\}) = \Psi_d^{d+1}$.
3. *If $\mathcal{L}$ is one of $\{X^d\}$ or $\{X^d, X^0\}$, then* $\gamma(\mathcal{L}) = \gamma^*(\mathcal{L})$.
4. $\gamma(\mathcal{P}_d) = \gamma(\{X^d, X^0\})$.

Proof: For $x > 0$ define

$$\mu_x := \frac{d \cdot (4x + d - 1)}{(d+1) \cdot 4x^2} .$$

Now, for any ξ it holds that ξ fulfills the necessary first order condition (6.1) for local extrema of $x \mapsto g^*_{X^d,x,1}(\mu_\xi)$. By Theorem 6.1, we get that ξ is even a global maximum on $\mathbb{R}_{\geq 0}$. Hence, $g^*_{X^d,\xi,1}(\mu_\xi) = \max_{x \in \mathbb{R}_{\geq 0}}\{g^*_{X^d,x,1}(\mu_\xi)\}$.

1. Fix $\xi := \Psi_d$. Note that $\Psi_d^2 = \Psi_d + \frac{d}{4}$ and hence

   $$\mu_\xi = \frac{d \cdot (4\Psi_d + d - 1)}{(d+1) \cdot (4\Psi_d + d)} \in (0, 1) .$$

   So far, we have shown that $\gamma^*(\{X^d\}) \leq g^*_{X^d,\xi,1}(\mu_\xi) = \Psi_d^{d+1}$. Since $h^*_{X^d,\xi,1} = 0$, it holds that $g^*_{X^d,\xi,1}$ is a constant, so $\Gamma^*_{\{X^d\}}(\mu) \geq \Psi_d^{d+1}$ for all $\mu \in (0, 1)$ and we have indeed $\gamma^*(\{X^d\}) = \Psi_d^{d+1}$.

2. Consider first the case $d = 1$: Fix $\xi := \frac{3}{2}$ and note that $\mu_\xi = \frac{1}{3} \in (0, 1)$. We have that $g^*_{X^0,0,1}(\frac{1}{3}) = \frac{3}{2} = g^*_{X^d,\xi,1}(\frac{1}{3})$. Consequently, also $\gamma^*(\{X^d, X^0\}) = \frac{3}{2}$.
   Otherwise, if $d \geq 2$, choose again $\xi := \Psi_d$. It holds that

   $$\begin{aligned} g^*_{X^d,\xi,1}(\mu_\xi) = \Psi_d^{d+1} &= \left(\frac{1 + \sqrt{d+1}}{2}\right)^{d+1} \\ &> \frac{2 \cdot (d+1)}{d + 1 + \sqrt{d+1}} \cdot \left(\frac{1 + \sqrt{d+1}}{2}\right)^2 \\ &= \frac{1}{1 - \mu_\xi} = g^*_{X^0,0,1}(\mu_\xi) . \end{aligned}$$

   As for (1.), we hence have $\gamma^*(\{X^d, X^0\}) = \Psi_d^{d+1}$.

3. For $x < \frac{y}{2}$, we have that $\kappa(x, y) = x \cdot (y - x) = \frac{y^2}{4} - \left(\frac{y}{2} - x\right)^2 \leq \frac{y^2}{4}$. Therefore, for every admissible triple ℓ, x, y we have $g_{\ell,x,y} \leq g^*_{\ell,x,y}$ (pointwise), where equality holds if $\frac{x}{y} \geq \frac{1}{2}$. Hence, if $\xi \geq \frac{1}{2}$, then also $g_{X^d,\xi,1}(\mu_\xi) = \max_{x \in \mathbb{R}_{\geq 0}}\{g_{X^d,x,1}(\mu_\xi)\}$.

4. The derivative of $g_{X^r,\xi,1}(\mu)$ with respect to r is

$$\begin{aligned}&\frac{\partial}{\partial r}\frac{\xi^r+\frac{r\cdot\xi^{r-1}}{4}-\mu\cdot\xi^{r+1}}{1-\mu}\\ &=\frac{\xi^{r-1}}{4(1-\mu)}+\ln(\xi)\cdot g_{X^r,\xi,1}(\mu)\,,\end{aligned}$$

which is positive if $\xi > 1$ and $g_{X^r,\xi,1}(\mu) \geq 0$. Consequently, if $\xi > 1$, then

$$g_{X^d,\xi,1}(\mu_\xi) = \max_{\substack{r\in[d]\\ x\in\mathbb{R}_{\geq 0}}}\{g_{X^r,x,1}(\mu_\xi)\}\,.$$

■

Theorem 5.5, Theorem 6.3, and Theorem 6.4 immediately imply:

Corollary 6.5 *Suppose that $\mathcal{L}$ is one of the following sets of cost functions.*

1. *If $\mathcal{L}$ is the set of linear functions, then $\gamma(\mathcal{L}) = \Psi_1^2 \approx 1.457$,*
2. *If $\mathcal{L} = \mathcal{P}_1$, then $\gamma(\mathcal{L}) = \frac{3}{2} > \Psi_1^2$,*
3. *If $\mathcal{L} = \mathcal{P}_d$ and $d \in \mathbb{N}_{\geq 2}$, then $\gamma(\mathcal{L}) = \Psi_d^{d+1}$.*

Acknowledgments We thank Kshipra Bhawalkar, Martin Gairing, Uri Nadav, and Tobias Harks for helpful discussions, and an anonymous SODA reviewer for very insightful remarks.

References

[1] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. Exact price of anarchy for polynomial congestion games. *SIAM Journal on Computing*, 40(5):1211–1233, September 2011.

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

[3] Umang Bhaskar, Lisa Fleischer, Darrell Hoy, and Chien-Chung Huang. Equilibria of atomic flow games are not unique. In *Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'09)*, pages 748–757, 2009.

[4] Stefano Catoni and Stefano Pallottino. Traffic equilibrium paradoxes. *Transportation Science*, 25(3):240–244, 1991.

[5] George Christodoulou and Elias Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC'05)*, pages 67–73. ACM Press, 2005.

[6] Roberto Cominetti, José R. Correa, and Nicolás E. Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.

[7] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, November 2004.

[8] José R. Correa, Andreas S. Schulz, and Nicolás E. Stier-Moses. On the inefficiency of equilibria in congestion games. In Michael Jünger and Volker Kaibel, editors, *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO'05)*, volume 3509 of *LNCS*, pages 167–181, 2005.

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

[10] Tobias Harks. Stackelberg strategies and collusion in network games with splittable flow. In Evripidis Bampis and Martin Skutella, editors, *Proceedings of the 6th International Workshop on Approximation and Online Algorithms (WAOA'08)*, volume 5426 of *LNCS*, pages 133–146, 2008.

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

[12] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In Christoph Meinel and Sophie Tison, editors, *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, volume 1563 of *LNCS*, pages 404–413. Springer Verlag, 1999.

[13] I. Milchtaich. Congestion models of competition. *American Naturalist*, 147(5):760–783, 1996.

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

[15] A. Neyman. Correlated equilibrium and potential games. *International Journal of Game Theory*, 26(2):223–227, 1997.

[16] Ariel Orda, Raphael Rom, and Nahum Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

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

[18] A. C. Pigou. *The Economics of Welfare.* Macmillan, 1920.

[19] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.

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

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

[22] Tim Roughgarden. *Selfish Routing and the Price of Anarchy.* MIT Press, 2005.

[23] Tim Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 1184–1185, 2005.

[24] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC'09)*, pages 513–522, 2009.

[25] Tim Roughgarden and Éva Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 47(2):389–403, 2004.

[26] H. Peyton Young. *Strategic learning and its limits.* Oxford University Press, 2004.

Appendix

A Details to Theorem 5.7 from Section 5 (Lower Bound)

A.1 Construction Example

As an example for our construction, consider the lucky case mentioned in Theorem 5.7 where only a single group of resources is needed. The admissible triple ℓ, x, y with $\ell(z) = z^3$, $x = \frac{3}{2}$, $y = 1$ constitutes a lucky case. It is easy to verify that

$$\begin{aligned} h_{\ell,x,y} &= (y-x)\cdot\ell(x) + \frac{y^2}{4}\cdot\ell'(x) \\ &= x^2\cdot\left((y-x)\cdot x + \frac{3}{4}\right) = 0\,. \end{aligned}$$

Let λ, μ such that $g_{\ell,x,y}(\mu) = \frac{\lambda}{1-\mu}$.

The family of instances is as follows. There are n players and n resources, each with cost function ℓ. The players' "optimal" strategies have size $p = 1$, whereas their "non-optimal" strategies have size $q = 2$. Each player thus has $t = \frac{n-1}{2}$ "non-optimal" strategies. We consider the strategy profile where every player puts load $\alpha = [\frac{1}{2} + \frac{3}{n-1}]\cdot[1 + \frac{2}{n-1}]^{-1} = \frac{n+5}{2\cdot(n+1)}$ on its "optimal" and $\beta = \frac{x-\alpha}{n-1} = \frac{1}{n+1}$ on all its "non-optimal" strategies. Then:

1. The load on each resource is exactly

$$\alpha + (n-1)\cdot\beta = x\,.$$

2. Each player is faced with equal marginal costs for all its strategies, because

$$\begin{aligned} p\cdot\big(\ell(x) + \alpha\cdot\ell'(x)\big) &= x^2\cdot(x + 3\cdot\alpha) \\ &= x^2\cdot\frac{6n+18}{2\cdot(n+1)} \\ &= x^2\cdot 2\cdot(x + 3\cdot\beta) \\ &= q\cdot\big(\ell(x) + \beta\cdot\ell'(x)\big)\ . \end{aligned}$$

3. For each resource, there is one player who puts load $\alpha = \frac{1}{2} \pm o(1)$ on it whereas all other players put load $\beta = o(1)$ on it.

4. In the “optimal” strategy profile, where each player only uses its “optimal” strategy, the load on any resource is $1 + o(1)$, because each player has weight

$$\alpha + t \cdot \beta = \alpha + \frac{n-1}{2 \cdot (n+1)} \xrightarrow{n \to \infty} 1 .$$

5. The social cost is $(\frac{\lambda}{1-\mu} - o(1))$ times as in the “optimal” strategy profile. This holds because each resource contributes cost

$$x \cdot \ell(x) = \lambda \cdot y \cdot \ell(y) + \mu \cdot x \cdot \ell(x) ,$$

where the equality is due to $g_{\ell,x,y}(\mu) = \frac{\lambda}{1-\mu}$ and the definition of $h_{\ell,x,y}$.

Together with the upper bound from Theorem 6, this construction alone immediately shows that the price of anarchy for atomic splittable congestion games with polynomial cost functions of degree at most 3 is *exactly* $(\frac{3}{2})^4 = 5.0625$.

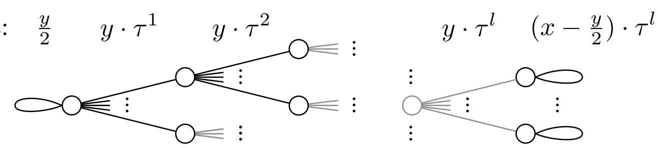

Figure 4: Illustration of construction with singleton strategies

A.2 Construction with Singleton Strategies

In the lucky case mentioned in Theorem 5.7, a price of anarchy of $\gamma(\mathcal{L})$ can even arise with singleton strategies. The lucky case occurs, e.g., when $\mathcal{L}$ contains only monomials (see Theorem 6).

Theorem A.1 *Let $\lambda \in \mathbb{R}$, $\mu < 1$. Moreover, let $\mathcal{L}$ be a scale-invariant set of cost functions, $\ell \in \mathcal{L}$, and $x \geq y > 0$. Suppose that*

$$g_{\ell,x,y}(\mu) = \frac{\lambda}{1-\mu} \text{ and } h_{\ell,x,y} = 0 .$$

Then, there is an infinite family of atomic splittable congestion games with singleton strategies, with cost functions in $\mathcal{L}$, and with limiting price of anarchy at least $\frac{\lambda}{1-\mu}$.

Proof: We define a family of singleton congestion games, represented by full k-ary trees of height l. To simplify our presentation, assume that the root node and each leaf node have self-loops. Then, each edge corresponds to a player, and each node in the tree corresponds to a resource. The strategies of a players are its (at most two) incident nodes. An illustration of the construction is shown in Theorem 4.

Let $\sigma, \tau > 0$ be values to be determined later (dependent on k and l). The cost function for resources on level j is $\ell_j(z) := \frac{1}{\sigma^j} \cdot \ell(\frac{z}{\tau^j})$. Note that the root resource has cost function $\ell_0 = \ell$. We say a player is in level $j \in [n]$ if its edge is between resource levels $j-1$ and j. The weight of each player in level j is $y \cdot \tau^j$. The player who only has the root resource as strategy has weight $\frac{y}{2}$, and the players who only have a leaf resource as strategy have weight $(x - \frac{y}{2}) \cdot \tau^l$.

We first show that we can choose σ and τ such that the profile in which each player splits its weight equally (i.e., each player on level j puts load $\frac{y}{2} \cdot \tau^j$ on both of its strategies) is a Nash equilibrium. Let $\tau := \frac{2x-y}{y \cdot k}$, so that the equilibrium load on each resource of level $j \in [l]_0$ is $\frac{y}{2} \cdot \tau^j + k \cdot \frac{y}{2} \cdot \tau^{j+1} = x \cdot \tau^j$. We need that each player faces equal marginal costs on each of its strategies, i.e., for players on all levels $j \in [l]$ that

$$\begin{aligned} &\ell_{j-1}(x \cdot \tau^{j-1}) + \left(\frac{y}{2} \cdot \tau^j\right) \cdot \ell'_{j-1}(x \cdot \tau^{j-1}) \\ &= \ell_j(x \cdot \tau^j) + \left(\frac{y}{2} \cdot \tau^j\right) \cdot \ell'_j(x \cdot \tau^j) . \end{aligned}$$

By plugging in that $\ell_j(z) = \frac{1}{\sigma^j} \cdot \ell(\frac{z}{\tau^j})$ and $\ell'_j(z) = \frac{1}{\sigma^j \cdot \tau^j} \cdot \ell'(\frac{z}{\tau^j})$, this is equivalent to

$$\ell(x) + \frac{y}{2} \cdot \tau \cdot \ell'(x) = \frac{1}{\sigma} \cdot \left[\ell(x) + \frac{y}{2} \cdot \ell'(x)\right] ,$$

i.e.,

$$\sigma = \frac{\ell(x) + \frac{y}{2} \cdot \ell'(x)}{\ell(x) + \frac{y}{2} \cdot \tau \cdot \ell'(x)} \xrightarrow{k\to\infty} 1 + \frac{y}{2} \cdot \frac{\ell'(x)}{\ell(x)} = \frac{2x-y}{y} ,$$

where the last equality follows from $h_{\ell,x,y} = 0$. Consequently, $k \cdot \tau \cdot \frac{1}{\sigma} \xrightarrow{k\to\infty} 1$, and the social cost contributed by all resources on level $j \in [l]_0$ is $k^j \cdot x \cdot \tau^j \cdot \frac{\ell(x)}{\sigma^j} \xrightarrow{k\to\infty} x \cdot \ell(x)$.

Now consider the profile where each player uses only the strategy farther away from the root. The social cost contributed by all resources on level $j \in [l-1]$ is $y \cdot \ell(y)$. The root resource on level 0 contributes $\frac{y}{2} \cdot \ell(\frac{y}{2})$, and level l contributes $k^l \cdot (x + \frac{y}{2}) \cdot \tau^l \cdot \frac{\ell(x+\frac{y}{2})}{\sigma^l} \xrightarrow{k\to\infty} (x + \frac{y}{2}) \cdot \ell(x + \frac{y}{2})$, which is constant in l.

Consequently, as both $k, l \to \infty$, we have that the ratio of social cost in the Nash equilibrium divided by the social cost in the other profile goes to $\frac{x \cdot \ell(x)}{y \cdot \ell(y)} = g_{\ell,x,y}(\mu)$. ■