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

Tim Roughgarden[†] Florian Schoppmann[‡]

April 29, 2013

Abstract

Congestion games are multi-player games in which players' costs 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. We characterize the worst-case price of anarchy — a quantitative measure of the inefficiency of equilibria — in 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, 18, 19, 24, 28] — yet structured enough to permit interesting theoretical guarantees. In the standard model introduced by Rosenthal [24], 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 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

*An extended abstract of this 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 grant CCF-1016885, an ONR Young Investigator Award, an ONR PECASE Award, and an AFOSR MURI grant. 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`.

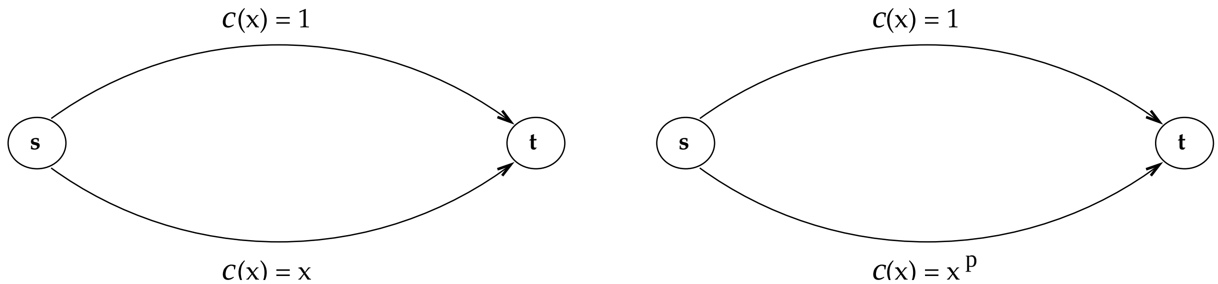

(a) Pigou's Example (b) A Nonlinear Variant

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

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 [22]. The splittable model also arises naturally when studying coalitions of players in nonatomic congestion games, where there is a continuum of players [7, 8, 14, 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)* [17]: 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 [23]. 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, 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.

[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.

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. By contrast, tight bounds for essentially all classes of cost functions were proved some years ago for both nonatomic congestion games (with a continuum of players, as in Figure 1) and standard (unsplittable) congestion games [1, 4, 9, 27, 29]. 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 [27] 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 [27]. While the smoothness paradigm in [27] provably cannot establish tight bounds on the POA in splittable congestion games, we show that local smoothness arguments can. Further, we prove the following "extension theorem": every POA bound derived via local smoothness applies automatically, without any quantitative degradation, to every correlated equilibrium, and hence also to every mixed Nash equilibrium, of the game.

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 time-averaged "swap regret" [11, 15], 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 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

[2]The blunter "smoothness framework" in [27] yields upper bounds that apply even more generally to the coarse correlated equilibria [12, 20] of the game; this is not always the case for local smoothness proofs (Example 3.3).

Table 1: The price of anarchy in the special case of polynomial cost functions with nonnegative coefficients. For splittable congestion games, the lower bounds are contributed by the present work. The upper bound of $\frac{3}{2}$ for affine cost functions was first proved by Cominetti et al. [8]. For higher-degree polynomials, we give the first closed-form POA upper bounds, essentially matching the numerical upper bounds computed by Harks [14].

Degree	**Atomic splittable**	Atomic unsplittable (weighted) [1]	Nonatomic [29]
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	$(\frac{1+\sqrt{d+1}}{2})^{d+1}$	$\Theta(\frac{d}{\log d})^{d+1}$	$\Theta(\frac{d}{\log d})$

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.

1.2 Related Work

We next describe the prior research that is most relevant to the present work. See [25, §4.8] for the history of and many more references on splittable congestion games.

Splittable congestion games seem more difficult to reason about than other congestion game models. For example, while the existence of pure Nash equilibria in such games was established early on via fixed-point arguments [13, 22], Bhaskar et al. [3] showed only recently that such equilibria need not be unique. 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 [7]. 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 [10, 26]. Cominetti et al. [8] 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 [7] to refute the general claims.

The first upper bounds on the POA in general splittable congestion games were given by Cominetti et al. [8]. These bounds are derived using a special case of our local smoothness framework in which one of our two parameters (λ in Definition 3.1) 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 nonnegative coefficients — bounds of $\frac{3}{2}$, 2.564, and 7.826 for affine, quadratic, and cubic cost functions, respectively. Harks [14] showed that allowing the parameter λ to vary yields significantly better POA bounds. The generic upper bound framework in [14] is equivalent to ours, though it produces bounds with a more complicated form. The simplified form derived here permits the first closed-form expressions for the POA for polynomial cost functions with nonnegative coefficients and, more importantly, enables the construction of matching lower bounds for all classes of allowable cost functions that satisfy mild technical conditions.

Prior to our work, there were no upper bounds on the POA of splittable congestion games for any equilibrium concept more general than pure Nash equilibria.

The best lower bounds on the POA that were known previously follow from counterexamples in Cominetti et al. [8]. For polynomials with nonnegative coefficients, these lower bounds grow linearly with the maximum degree d; for example, they are $1.343, 1.67, 1.981, 2.287$ for $d = 1, 2, 3, 4$, respectively. Our tight lower bounds are exponentially larger in the degree d.

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 resource cost functions that are polynomials with nonnegative coefficients (cf., Table 1). Section 7 concludes. 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 $e \in E$ has 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 a distribution of its *weight* $w_i \in \mathbb{R}_{>0}$ among the basic strategies available to it, i.e., player i's set of (fractional) strategies is $S_i := \{\vec{x}^i \in \mathbb{R}_{\geq 0}^{\mathcal{P}_i} \mid \sum_{p \in \mathcal{P}_i} x_p^i = w_i\}$. A *strategy profile* is a vector $\vec{x} = (\vec{x}^i)_{i \in [n]}$ of all players' strategies. We sometimes call a fractional strategy that uses only one basic 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 enable a useful characterization of Nash equilibria; see (2), below. We say that a set of cost functions $\mathcal{L}$ is *non-trivial* if it contains at least one function that is not everywhere zero, and *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 $c_i(\vec{x}) := \sum_{e \in E} x_e^i \cdot \ell_e(x_e)$. 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 $\mathrm{SC}(\vec{x}) = \sum_{e \in E} x_e \cdot \ell_e(x_e)$.

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., [31] for more details and context). A *(pure) Nash equilibrium* — the most restrictive concept — is a strategy profile $\vec{x}$ such that for every player i and every fractional strategy $\vec{y}^i$ it holds that $c_i(\vec{x}) \leq c_i(\vec{y}^i, \vec{x}^{-i})$, where $\vec{x}^{-i}$ denotes the strategies chosen by the players other than i in $\vec{x}$. Pure Nash equilibria always exist in splittable congestion games [13, 22].

A *mixed Nash equilibrium* is a profile of mixed strategies — stochastically independent probability distributions $P_1, \ldots, P_n$ over $S_1, \ldots, S_n$ — such that

$$\mathrm{E}_{\vec{x} \sim P}[c_i(\vec{x})] \leq \mathrm{E}_{\vec{x} \sim P}[c_i(\vec{y}^i, \vec{x}^{-i})] \tag{1}$$

for all players i and all fractional strategies $\vec{y}^i \in S_i$, where P denotes the product distribution over strategy profiles induced by $P_1, \ldots, P_n$. Pure Nash equilibria are the mixed Nash equilibria in which no player randomizes.

A (not necessarily product) distribution P over the set of strategy profiles is a *correlated equilibrium* if for all players i and all functions $\delta : S_i \to S_i$ it holds that

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

Mixed Nash equilibria correspond to the correlated equilibria that are product distributions.

Finally, such a distribution P is a *coarse correlated equilibrium* if (1) holds for all players i and all strategies $\vec{y}^i \in S_i$. Every correlated equilibrium is a coarse correlated equilibrium, and the converse is false in general (e.g., Example 3.3).

Characterization of Nash Equilibria Since cost functions are differentiable and semi-convex, a necessary and sufficient condition for a strategy profile to be a (pure) Nash equilibrium is that for every player i, the *marginal cost* of every used basic strategy is the same and at most that of every unused basic strategy. That is,

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

for all players $i \in [n]$ and all $p, p' \in \mathcal{P}_i$ with $x_p^i > 0$, where $\ell_e^i(\vec{x}_e)$ denotes $\ell_e(x_e) + x_e^i \cdot \ell'_e(x_e)$. This condition can alternatively be stated as a variational inequality:

$$\sum_{e \in E} \ell_e^i(\vec{x}_e) \cdot (y_e^i - x_e^i) \geq 0 \tag{2}$$

for every player $i \in [n]$ and every strategy $\vec{y}^i$. See Harks [14, Lemma 1], for example, for formal proofs of these characterizations.

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 [27]. 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 always extend to the coarse correlated equilibria of a game.

For context and comparison, we next review the standard definition of smooth games [27].[3] 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. A cost-minimization game is (λ, μ)-*smooth* if

$$\sum_{i=1}^{n} c_i(\vec{y}^i, \vec{x}^{-i}) \leq \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \tag{3}$$

for every pair $\vec{x}, \vec{y}$ of outcomes. The main extension theorem in [27] states that 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.

For the rest of this section, we consider cost-minimization games for which every strategy set S_i is a convex compact subset of some Euclidean space $\mathbb{R}^{m_i}$ and every cost function c_i is continuously differentiable. The splittable congestion games that we consider satisfy these assumptions. The rough intuition behind local smoothness is to require the constraint (3) 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 on the POA.[4] Formally, we implement this idea as follows.

[3] There are several precursors to and recent variations on this definition; see [27] for a detailed discussion.

[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) 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.

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 $\vec{x}$,

$$\sum_{i=1}^{n} \left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i) \right] \leq \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \,. \tag{4}$$

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 $\vec{y}$ with $\mu < 1$, 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 (4) and the linearity of expectation (twice) to derive

$$\mathrm{E}_{\vec{x} \sim P}[\mathrm{SC}(\vec{x})] \leq \sum_{i=1}^{n} \mathrm{E}_{\vec{x} \sim P} \left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i) \right] \leq \mathrm{E}_{\vec{x} \sim P} \left[\lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x}) \right] \tag{5}$$

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 some player i. For brevity, define the deviation function $\delta_\epsilon : S_i \to S_i$ by $\delta_\epsilon(\vec{x}^i) := (1-\epsilon) \cdot \vec{x}^i + \epsilon \cdot \vec{y}^i$. Intuitively, we are considering the hypothetical deviation by player i that always replaces its strategy $\vec{x}^i$ by one that is "a little closer" to $\vec{y}^i$. Since strategy sets are convex, $\delta_\epsilon(\vec{x}^i)$ is a well-defined strategy for every ϵ between 0 and 1. In the limit as ϵ goes to zero, $\mathrm{E}_{\vec{x} \sim P}[\frac{1}{\epsilon}(c_i(\delta_\epsilon(\vec{x}^i), \vec{x}^{-i}) - c_i(\vec{x}))]$ 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 $\mathrm{E}_{\vec{x} \sim P}[c_i(\delta_\epsilon(\vec{x}^i), \vec{x}^{-i})] < \mathrm{E}_{\vec{x} \sim P}[c_i(\vec{x})]$, 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 $\varepsilon > 0$ is an arbitrarily small constant. This identical-interest game has positive, continuously differentiable, convex cost functions and convex compact strategy sets. Let P be the uniform distribution over the strategy profiles $(0, \alpha)$ and $(1, 1-\alpha)$, where $\alpha \in (0, \frac{1}{4}]$. Elementary calculations verify that this is a coarse correlated equilibrium with expected social cost $2\alpha^2 + 2\varepsilon$. Further calculations show that for every strategy profile $\vec{x}$ and every optimal strategy profile $\vec{y}$ (i.e.,

[5] This can be formally justified using the dominated convergence theorem: Since the strategy sets are compact and the cost functions are continuously differentiable, there is a constant $M < \infty$ such that $|\frac{1}{\epsilon}(c_i(\delta_\epsilon(\vec{x}^i), \vec{x}^{-i}) - c_i(\vec{x}))| < M$ for every strategy profile $\vec{x}$. Hence, $\lim_{\epsilon \searrow 0} \int \frac{1}{\epsilon}(c_i(\delta_\epsilon(\vec{x}^i), \vec{x}^{-i}) - c_i(\vec{x}))\, dP(\vec{x}) = \int \nabla_i c_i(\vec{x})^T (\vec{y}^i - \vec{x}^i)\, dP(\vec{x})$.

[6] A similar trick was used by Neyman [21] to prove a rather different result, that every game with convex compact strategy sets and a strictly concave potential function has a unique correlated equilibrium.

$y_1 = y_2$) 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 locally $(1,0)$-smooth with respect to every optimal strategy profile. The corresponding approximation factor of $\lambda/(1-\mu) = 1$ obviously does not apply to the coarse correlated equilibria P.

Remark 3.4 (Smoothness Versus Local Smoothness) Here is one reason why standard smoothness arguments extend to coarse correlated equilibria but local smoothness arguments do not. In the definition (3) of (λ, μ)-smoothness, the outcome $\vec{y}$ is used to propose hypothetical deviations $\vec{y}^1, \ldots, \vec{y}^n$ for the players. These proposed deviations are *independent* of the strategy profile $\vec{x}$, and for this reason the resulting approximation bound of $\frac{\lambda}{1-\mu}$ extends to all coarse correlated equilibria. In Definition 3.1 and the proof of Theorem 3.2, however, the outcome $\vec{y}$ induces the hypothetical deviations $(1-\epsilon)\vec{x}^1 + \epsilon\vec{y}^1, \ldots, (1-\epsilon)\vec{x}^n + \epsilon\vec{y}^n$, which *do* depend on $\vec{x}$. Fortunately, the proposed deviation $(1-\epsilon)\vec{x}^i + \epsilon\vec{y}^i$ for player i depends only $\vec{x}^i$ and not on $\vec{x}^{-i}$, and for this reason the resulting approximation bound of $\frac{\lambda}{1-\mu}$ extends to all 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 every $\vec{x}, \vec{y} \in \mathbb{R}^n_{\geq 0}$ with $\sum_{i=1}^n x_i = x$ and $\sum_{i=1}^n y_i = y$, $\sum_{i=1}^n \left(y_i \cdot x_i - x_i^2\right) \leq \kappa(x,y)$.*

Proof: Denote $x_{\max} = \max_{i=1}^n x_i$. We have

$$\sum_{i=1}^{n} \left(y_i \cdot x_i - x_i^2\right) \leq \sum_{i=1}^{n} (y_i \cdot x_i) - x_{\max}^2 \leq y \cdot x_{\max} - x_{\max}^2 = \frac{y^2}{4} - \left(\frac{y}{2} - x_{\max}\right)^2 \leq \frac{y^2}{4}.$$

For the case where $x < y/2$, observe that $z \mapsto (y \cdot z - z^2)$ is increasing on $[0, y/2]$. Consequently, $y \cdot x_{\max} - x_{\max}^2 \leq y \cdot x - x^2 = x(y-x)$, as required. ■

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

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

$$y \cdot \ell(x) + \kappa(x,y) \cdot \ell'(x) \leq \lambda \cdot y \cdot \ell(y) + \mu \cdot x \cdot \ell(x) \tag{6}$$

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: Consider a splittable congestion game with cost functions in $\mathcal{L}$ and two strategy profiles $\vec{x}$ and $\vec{y}$. Recall that $\ell^i_e(\vec{x}_e)$ denotes the marginal cost $\ell_e(x_e) + x^i_e \cdot \ell'_e(x_e)$. We have

$$
\begin{aligned}
\sum_{i=1}^{n} \left[c_i(\vec{x}) + \nabla_i c_i(\vec{x})^T(\vec{y}^i - \vec{x}^i)\right] &= \sum_{i\in[n]}\sum_{e\in E} \left[x^i_e \cdot \ell_e(x_e) + y^i_e \cdot \ell^i_e(\vec{x}_e) - x^i_e \cdot \ell^i_e(\vec{x}_e)\right] \\
&= \sum_{e\in E}\left[y_e \cdot \ell_e(x_e) + \ell'_e(x_e) \cdot \sum_{i\in[n]} \left(y^i_e \cdot x^i_e - (x^i_e)^2\right)\right] \\
&\leq \sum_{e\in E} \left[y_e \cdot \ell_e(x_e) + \kappa(x_e, y_e) \cdot \ell'_e(x_e)\right] \qquad (7) \\
&\leq \sum_{e\in E} \left[\lambda \cdot y_e \cdot \ell_e(y_e) + \mu \cdot x_e \cdot \ell_e(x_e)\right] \qquad (8) \\
&= \lambda \cdot \mathrm{SC}(\vec{y}) + \mu \cdot \mathrm{SC}(\vec{x})\,,
\end{aligned}
$$

where inequalities (7) and (8) follow from Lemma 4.1 and assumption (6), respectively. ■

We now define the quantity $\gamma(\mathcal{L})$ as, intuitively, the best upper bound on the POA that is provable using Theorem 3.2 and Proposition 4.2. Formally, we first define $g_{\ell,x,y} : \mathbb{R}_{<1} \to \mathbb{R} \cup \{\infty\}$ by

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

for every *admissible triple* ℓ, x, y, meaning a cost function $\ell \in \mathcal{L}$ and values $x \geq 0, y > 0$ with $\ell(y) > 0$. If $\mu < 1$, then for every admissible triple ℓ, x, y, the constraint (6) is equivalent to

$$g_{\ell,x,y}(\mu) \leq \frac{\lambda}{1 - \mu}\,; \qquad (9)$$

that is, $g_{\ell,x,y}(\mu)$ is a lower bound on the best POA bound that can be proved using Proposition 4.2 and a given value of $\mu < 1$.

Non-admissible triples ℓ, x, y can be ignored in Proposition 4.2. First, if ℓ is the zero function, inequality (6) reduces to $0 \leq 0$ irrespective of λ and μ. Second, if ℓ is not the zero function, then define $\xi := \max\{y \geq 0 \mid y \cdot \ell(y) = 0\}$. This maximum is guaranteed to exist because $y \mapsto y \cdot \ell(y)$ is continuous. Now if (6) holds for all $y > \xi$, then it also holds for $y = \xi$ (since both sides of (6) are continuous in y), and hence for all $y \in [0, \xi]$ (since the left-hand side of (6) is nondecreasing in y).

The upshot is that, for $\mu < 1$, the requirement of Proposition 4.2 — that is, the conjunction of all constraints (6) over all triples $\ell \in \mathcal{L}$, $x, y \geq 0$ — is equivalent to

$$\sup_{\substack{\ell\in\mathcal{L} \\ x\geq 0, y>0, \ell(y)>0}} g_{\ell,x,y}(\mu) \leq \frac{\lambda}{1 - \mu}\,. \qquad (10)$$

Put differently, for a fixed value of $\mu < 1$, the value of λ that minimizes $\frac{\lambda}{1-\mu}$ subject to condition (6) for all admissible triples is $(1 - \mu)$ times the left-hand side of (10).

Given a non-trivial set of cost functions $\mathcal{L}$, the best POA bound provable using Theorem 3.2 and Proposition 4.2 is the infimum of $\frac{\lambda}{1-\mu}$ over all choices of (λ, μ) with $\mu < 1$ that meet condition (6) for all admissible triples. Since condition (6) reduces to $0 \leq \mu \cdot x \cdot \ell(x)$ if $y = 0$, any finite POA bound also requires $\mu \geq 0$. The left-hand side of (10) is the best POA bound for a given choice of μ, and the definition of $\gamma(\mathcal{L})$ simply minimizes this POA bound over the choices for μ:

$$\gamma(\mathcal{L}) := \inf_{\mu \in [0,1)} \sup_{\substack{\ell \in \mathcal{L} \\ x \geq 0, y > 0, \ell(y) > 0}} g_{\ell,x,y}(\mu) \,. \tag{11}$$

The definition of $\gamma(\mathcal{L})$, Proposition 4.2, and Theorem 3.2 immediately imply the following.

Corollary 4.3 *For every non-trivial 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 $\gamma(\mathcal{L})$.*

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

In this section, we show that for every non-trivial 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* $\gamma(\mathcal{L})$. Before giving the main construction in Section 5.2, we prove in Section 5.1 that $\gamma(\mathcal{L})$ can "usually" be approximated arbitrarily well by the intersection of a non-decreasing curve $g_{\ell_1,x_1,y_1}(\mu)$ and a non-increasing curve $g_{\ell_2,x_2,y_2}(\mu)$. These two curves encode the cost functions and resource loads that we use in the construction of a worst-case congestion game. The "unusual" cases, in which $\gamma(\mathcal{L})$ must be $+\infty$, are handled directly in Section 5.2.

5.1 Approximating $\gamma(\mathcal{L})$ by Two Curves

Define $\Gamma_{\mathcal{L}} : [0,1) \to \mathbb{R} \cup \{\infty\}$ as the inner part of the infimum in the definition (11) of $\gamma(\mathcal{L})$:

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

This is the optimal POA bound that can be proved using local smoothness (Theorem 3.2 and Proposition 4.2) with the given value of μ. Figure 4 in Section 6 provides plots of the functions $g_{\ell,x,y}$ and $\Gamma_{\mathcal{L}}$ when $\mathcal{L}$ contains only linear and constant functions. In general, the function $\Gamma_{\mathcal{L}}$ is non-increasing on $(0, \mu]$ and non-decreasing on $[\mu, 1)$ for some μ, and unbounded as μ approaches 0 or 1.

Given an admissible triple ℓ, x, y, define the scalar $h_{\ell,x,y}$ by

$$h_{\ell,x,y} := (y - x) \cdot \ell(x) + \kappa(x,y) \cdot \ell'(x) \,. \tag{12}$$

A simple calculation shows that, for every admissible triple ℓ, x, y and $\mu < 1$, $h_{\ell,x,y}$ and $g_{\ell,x,y}(\mu)$ have the same sign. Specifically, $g_{\ell,x,y}(\mu)$ has the form $\frac{a - \mu \cdot b}{c \cdot (1-\mu)}$, with $a, b \geq 0$, $c > 0$, the derivative of which is $\frac{a-b}{c \cdot (1-\mu)^2}$. Hence,

$$\frac{\partial g_{\ell,x,y}(\mu)}{\partial \mu} = \frac{h_{\ell,x,y}}{y \cdot \ell(y) \cdot (1-\mu)^2} \,. \tag{13}$$

Thus, the sign of $h_{\ell,x,y}$ indicates whether the function $g_{\ell,x,y}$ is strictly increasing, strictly decreasing, or constant in μ. The values $h_{\ell,x,y}$ arise as "error terms" in the construction in Section 5.2, and must be carefully managed to produce a worst-case example.

Lemma 5.1 (Two Curves Lemma) *Let $\mathcal{L}$ be a set of non-trivial cost functions. Suppose there is an admissible triple ℓ, x, y with $h_{\ell,x,y} < 0$. Then, 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*

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

Proof: The easy case is when there is an admissible triple ℓ, x, y such that $g_{\ell,x,y}$ is a constant function larger than $\widehat{\gamma}$. In this case, $h_{\ell,x,y} = 0$, and we can use this triple for both ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2 to satisfy the requirements of the lemma. Relatively simple tight lower-bound constructions are possible in this special case, as we show later. In the rest of this proof, we assume that no such triple exists.

Define

$$\mu^* := \inf\{\mu \in [0,1) \mid \exists \text{ admissible triple } \ell, x, y \text{ with } g_{\ell,x,y}(\mu) \geq \widehat{\gamma} \text{ and } g_{\ell,x,y} \text{ is strictly increasing } \} .$$

This infimum is taken over a non-empty set and hence $\mu^* < 1$. To see this, choose $\ell \in \mathcal{L}$ and $y > x > 0$ such that $\ell(x) > 0$. Note that $h_{\ell,x,y} > 0$. Then $g_{\ell,x,y}(\mu)$ has the form $\frac{a+b-\mu \cdot c}{1-\mu}$ where $0 < a \leq 1$, $b \geq 0$, and $0 < c < a$. Therefore, $\lim_{\mu \nearrow 1} g_{\ell,x,y}(\mu) = \infty$. This shows that the condition in the definition of μ^* is met for values of μ that are sufficiently close to 1.

The key claim is that there is a value $\widehat{\mu} < 1$ and admissible triples ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2 so that g_{ℓ_1,x_1,y_1} is strictly increasing, g_{ℓ_2,x_2,y_2} is strictly decreasing, and $g_{\ell_2,x_2,y_2}(\widehat{\mu}) \geq g_{\ell_1,x_1,y_1}(\widehat{\mu}) \geq \widehat{\gamma}$. Then, since both functions are unbounded at $\mu = 1$, they must intersect at a point (μ, γ) with $\widehat{\mu} \leq \mu < 1$ and $\gamma \geq \widehat{\gamma}$, which completes the proof.

To prove the key claim, we distinguish two cases.

(1) There is a strictly increasing function g_{ℓ_1,x_1,y_1} with $g_{\ell_1,x_1,y_1}(\mu^*) > \widehat{\gamma}$.

Since g_{ℓ_1,x_1,y_1} is a continuous function, there is a value $\widehat{\mu} < \mu^*$ so that also $g_{\ell_1,x_1,y_1}(\widehat{\mu}) > \widehat{\gamma}$. We must have $\mu^* = 0$ in this case, as otherwise we could have found a smaller value for μ^*.

Next, by the assumption of the lemma, there is an admissible triple ℓ, x, y with $h_{\ell,x,y} < 0$, which implies $0 < y < x$. Define $\xi := \max\{y \geq 0 \mid y \cdot \ell(y) = 0\}$. Note that $g_{\ell,x,y}(\widehat{\mu}) \geq \frac{-\widehat{\mu} \cdot x \cdot \ell(x)}{(1-\widehat{\mu}) \cdot y \cdot \ell(y)} \xrightarrow{y \searrow \xi} \infty$, since $\widehat{\mu} < 0$. Denote $\ell_2 = \ell$, $x_2 = x_2$, and let y_2 be such that $g_{\ell_2,x_2,y_2}(\widehat{\mu}) \geq g_{\ell_1,x_1,y_1}(\widehat{\mu})$.

(2) For every strictly increasing function $g_{\ell,x,y}$, $g_{\ell,x,y}(\mu^*) \leq \widehat{\gamma}$.

Since $\Gamma_{\mathcal{L}}(\mu^*) \geq \gamma(\mathcal{L}) > \widehat{\gamma}$, in this case there must be a strictly decreasing function g_{ℓ_2,x_2,y_2} with $g_{\ell_2,x_2,y_2}(\mu^*) > \widehat{\gamma}$. Since g_{ℓ_2,x_2,y_2} is continuous, we can choose δ so that $\mu^* + \delta < 1$ and $g_{\ell_2,x_2,y_2}(\mu^* + \delta) > \widehat{\gamma}$. Moreover, by the definition of μ^*, there is a strictly increasing function g_{ℓ_1,x_1,y_1} with $g_{\ell_1,x_1,y_1}(\mu^* + \delta) \geq \widehat{\gamma}$. Since $g_{\ell_1,x_1,y_1}(\mu^*) \leq \widehat{\gamma}$ by assumption, continuity and monotonicity imply that there is a value $\widehat{\mu} \in [\mu^*, \mu^* + \delta]$ with $g_{\ell_2,x_2,y_2}(\widehat{\mu}) \geq g_{\ell_1,x_1,y_1}(\widehat{\mu}) \geq \widehat{\gamma}$.

■

Remark 5.2 The requirement in Lemma 5.1 that there is an admissible triple ℓ, x, y with $h_{\ell,x,y} < 0$ is not without loss of generality. For instance, suppose that $\mathcal{L}$ contains only a function ℓ that satisfies $\ell(x) = 0$ for $x \in [0, 2]$ and $\ell'(x) \geq x \cdot \ell(x) > 0$ for all $x > 2$. Every admissible triple satisfies $y > 2$. Definition (12) implies that $h_{\ell,x,y} \leq 0$ only if $y \leq x$. For all such admissible triples, $h_{\ell,x,y} = (y - x) \cdot \ell(x) + \frac{y^2}{4} \cdot \ell'(x) > y \cdot \ell(x) > 0$.

5.2 The Construction

5.2.1 Guiding Necessary Conditions

To construct a family of examples with POA approaching the upper bound proved in Theorem 3.2 and Proposition 4.2, it is necessary that all of the inequalities in the upper bound — inequalities (5), (7), and (8) — hold with equality in the limit.

The plan for our construction is as follows. We first apply Lemma 5.1 to obtain two admissible triples ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2. We then construct a family of instances that each contain two groups of resources, one with cost functions ℓ_1 and one with cost functions ℓ_2. Each instance will possess a Nash equilibrium $\vec{u}$ in which players are indifferent between all of their basic strategies and the load on all resources of group $i \in \{1, 2\}$ is x_i, and yet there is another strategy profile $\vec{v}$ in which the load approaches y_i on each resource of group i. Suppose now that $g_{\ell_i,x_i,y_i}(\mu) = \frac{\lambda}{1-\mu}$ for $i = 1, 2$. By the definition of h_{ℓ_i,x_i,y_i}, we have

$$x_i \cdot \ell_i(x_i) = \lambda \cdot y_i \cdot \ell_i(y_i) + \mu \cdot x_i \cdot \ell_i(x_i) - h_{\ell_i,x_i,y_i} \,. \tag{14}$$

This indicates that we need $\operatorname{sgn}(h_{\ell_1,x_1,y_1}) = -\operatorname{sgn}(h_{\ell_2,x_2,y_2})$ and to choose the number of resources in groups 1 and 2 so that in the sum of the above equations, over all resources, the h_{ℓ_i,x_i,y_i}-terms vanish. Then $\frac{\mathrm{SC}(\vec{u})}{\mathrm{SC}(\vec{v})} = \frac{\lambda}{1-\mu}$ as needed.

So far, our construction idea provides tightness for the variational inequality (5) and for the (λ, μ)-smoothness inequality (8). To see how to make inequality (7) tight as well, we extend an observation of Cominetti et al. [8, Theorem 3.1]. Consider Lemma 4.1, which distills inequality (7). As $n \to \infty$, Lemma 4.1 is asymptotically tight when $x_1 = \min\{\frac{y}{2}, x\}$, $x_2 = \cdots = x_n$, and $y_1 = y$, $y_2 = \cdots = y_n = 0$. To see this, note that if $x \geq \frac{y}{2}$, then $x_1 = \frac{y}{2}$, $x_2 = \cdots = x_n = \frac{2x-y}{2n-2}$, and thus $\sum_i (y_i \cdot x_i - x_i^2) = \frac{y^2}{4} - \frac{(2x-y)^2}{4n-4}$. If $x < \frac{y}{2}$, then $x_1 = x$, $x_2 = \cdots = x_n = 0$, and thus $\sum_i (y_i \cdot x_i - x_i^2) = x(y - x)$.

To take advantage of this observation in our construction, we ensure that for each resource of group i, one player contributes load $\min\{\frac{y_i}{2}, x_i\}$ to the resource in the Nash equilibrium, while all other players contribute only infinitesimal amounts.

5.2.2 The Main Construction

The following theorem is the main construction of worst-case examples. The edge case in which Lemma 5.1 does not apply is treated separately in the following section.

Theorem 5.3 (Main Construction) *Let $\lambda, \mu \in \mathbb{R}$ with $\mu < 1$. Let ℓ_1, ℓ_2 be cost functions and $x_1, x_2 \geq 0$ and $y_1, y_2 > 0$. Define ω by $\ell_2(x_2) + \frac{y_2}{2} \cdot \ell_2'(x_2)$ if $x_2 \geq y_2/2$ and $\ell_2'(x_2) > 0$, and by $\ell_2(x_2) + x_2 \cdot \ell_2'(x_2)$ otherwise. Suppose that all of the following conditions hold:*

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

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 of the other variables, described in Table 2, are functions of n and p_2. For convenience, we also denote $h_i := h_{\ell_i,x_i,y_i}$ for $i \in \{1,2\}$, and we use the notation $\overline{1} := 2$ and $\overline{2} := 1$.

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

Symbol	Meaning (load refers to load in Nash equilibrium)	Definition (references to paragraph "The Equilibrium")
n	number of players per group	*free scaling parameter*
p_i	size of "optimal" strategies in group i	$p_1 := \lceil p_2 \cdot \omega \rceil$ p_2: *free scaling parameter*
q_i	size of "non-optimal" strategies in group i	$q_i := \left\lfloor p_i \cdot \frac{2x_i - y_i + 2h_i}{y_i} \right\rfloor$
t_i	number of "non-optimal" strategies for each player in group i	$t_i := \frac{p_i \cdot (n-1)}{q_i}$
α_i	load each player from group i puts on its "optimal" strategy	*see* (18) *in condition (3.)*
β_i	load each player from group i puts on its "non-optimal" strategies	$\beta_i := \frac{x_i - \alpha_i - n \cdot \gamma_i}{n-1}$
γ_i	load each player from group $\bar{i}$ puts on each "optimal" strategy of group i	$\gamma_1 := \frac{-h_1}{n}$ $\gamma_2 := 0$
w_i	weight of players in group i	$w_i := \alpha_i + t_i \cdot \beta_i + n \cdot \gamma_{\bar{i}}$
σ_i	scaling factor for cost functions in group i	σ_1: *see* (16) *in condition (2.)* $\sigma_2 := 1$

Resources There are two groups of resources, with group $i \in \{1,2\}$ consisting of $n \cdot p_i$ resources that we denote by $(i,0),\dots,(i,n\cdot p_i - 1)$. A good intuition is to think of two cycles; see also Figure 2, which illustrates our construction. 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 players, with group $i \in \{1,2\}$ consisting of n players denoted by $(i,0),\dots,(i,n-1)$. Each player (i,j) has one "optimal" strategy $\mathcal{P}_{i,j,0}$, which comprises p_i resources. Different players' optimal strategies are disjoint, so they partition the resources of a group. If $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, then player (i,j) has also $t_i := \frac{p_i \cdot (n-1)}{q_i}$ "non-optimal" strategies $\mathcal{P}_{i,j,1},\dots,\mathcal{P}_{i,j,t_i}$, each comprising q_i resources. These non-optimal strategies are mutually disjoint, and also disjoint from the player's optimal strategy. Finally, players from group 2 can also use the "optimal" strategies for group 1, i.e., $\mathcal{P}_{1,0,0},\dots,\mathcal{P}_{1,n-1,0}$. Formally:

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

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

The Equilibrium Define the strategy profile $\vec{u}$ as follows. Each player (i,j) uses strategy $\mathcal{P}_{i,j,0}$ with load α_i and each of the strategies $\mathcal{P}_{i,j,1},\dots,\mathcal{P}_{i,j,t_i-1}$ with load β_i. If $x_i < \frac{y_i}{2}$ or $\ell_i'(x_i) = 0$,

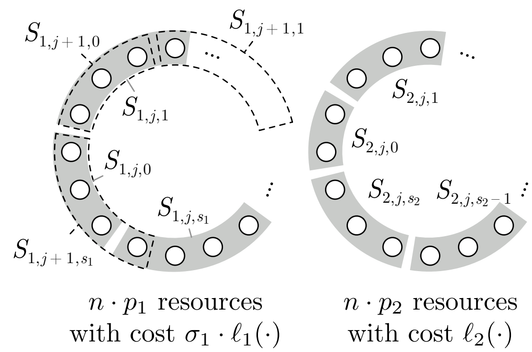

Figure 2: Illustration of construction with $p_1 = 3$, $q_1 = 4$ and $p_2 = 2$, $q_2 = 3$

then β_i is necessarily 0. In addition, each player in group 2 uses each of the n "optimal" strategies in group 1 with load γ_1.

Define the strategy profile $\vec{v}$ as that in which every player uses only its "optimal" strategy.

We next state six conditions that formalize the high-level plan outlined in the previous section. After their statements, we explain how to choose values for the parameters in Table 2 so that all of the conditions are satisfied simultaneously.

1. In the profile $\vec{u}$, the load on each resource of group i is exactly x_i. That is,

$$\alpha_i + (n-1) \cdot \beta_i + n \cdot \gamma_i = x_i \,; \quad \text{equivalently,}$$
$$\beta_i = \frac{x_i - \alpha_i - n \cdot \gamma_i}{n-1} \,. \tag{15}$$

2. In the profile $\vec{u}$, each player is faced with equal marginal costs for all its strategies, and hence the profile is a Nash equilibrium. The first condition for players in group 2 is

$$p_1 \cdot \sigma_1 \cdot \big(\ell_1(x_1) + \gamma_1 \cdot \ell_1'(x_1)\big) = p_2 \cdot \sigma_2 \cdot \big(\ell_2(x_2) + \alpha_2 \cdot \ell_2'(x_2)\big) \ . \tag{16}$$

Second, for $i = 1, 2$, if $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, then

$$p_i \cdot \big(\ell_i(x_i) + \alpha_i \cdot \ell_i'(x_i)\big) = q_i \cdot \big(\ell_i(x_i) + \beta_i \cdot \ell_i'(x_i)\big) \ . \tag{17}$$

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

If $i = 2$ and $x_2 \leq \frac{y_2}{2}$, there is nothing to show because $\alpha_2 = x_2$. (For $i = 1$, the assumption that $h_1 \leq 0$ implies that $x_1 > y_1 > \frac{y_1}{2}$.) Otherwise, $\frac{y_i}{2} \leq x_i$ and, recalling the assumption that $\ell_i(x_i) = 1$, we can plug in $\ell_i'(x_i) = \frac{4(x_i - y_i + h_i)}{y_i^2}$ and (15) into (17) to obtain

$$\alpha_i = \left[\frac{y_i^2 \cdot \left(\frac{q_i}{p_i} - 1\right)}{4 \cdot (x_i - y_i + h_i)} + \frac{q_i \cdot (x_i - n \cdot \gamma_i)}{(n-1) \cdot p_i}\right] \cdot \left[1 + \frac{q_i}{(n-1) \cdot p_i}\right]^{-1} \,. \tag{18}$$

The desired limits $\alpha_i \xrightarrow{n,p_2\to\infty} \frac{y_i}{2}$ and $\beta_i \xrightarrow{n,p_2\to\infty} 0$ hold provided

$$\frac{q_i}{p_i} \xrightarrow{p_2\to\infty} \frac{2x_i - y_i + 2h_i}{y_i}\,, \quad \text{which holds if we set}$$

$$q_i := \left\lfloor p_i \cdot \frac{2x_i - y_i + 2h_i}{y_i} \right\rfloor . \tag{19}$$

4. In the strategy profile $\vec{v}$, the load on every resource in group i is $y_i + o(1)$. That is, $w_i \xrightarrow{n,p_2\to\infty} y_i$.

   We first make some preliminary calculations. If $x_2 \geq \frac{y_2}{2}$ and $\ell_2'(x_2) > 0$, then

$$\begin{aligned} n \cdot \gamma_1 = -h_1 = \frac{h_2}{\omega} &= \frac{h_2 \cdot y_2}{y_2 + \frac{y_2^2}{2} \cdot \ell_2'(x_2)} \\ &= \frac{y_2}{2} \cdot \frac{2h_2}{2x_2 - y_2 + 2h_2}\,. \end{aligned} \tag{20}$$

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

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

   Now, consider $i \in \{1, 2\}$. Recall that our assumption that $h_1 \leq 0$ implies that $x_1 \geq y_1$.

   - 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 $x_2 \leq \frac{y_2}{2}$ or $\ell_2'(x_2) = 0$, then $w_2 = \alpha_2 + n \cdot \gamma_1 = x_2 + (y_2 - x_2) = y_2$.
   - Otherwise, $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$. Using equations (15) and (19), and also equation (20) for the $i = 2$ case, we have

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

5. The social cost of the Nash equilibrium $\vec{u}$ is $(\frac{\lambda}{1-\mu} - o(1))$ times that of the profile $\vec{v}$.

   Using condition 1, write $\mathrm{SC}(\vec{u}) = \sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot x_i \cdot \ell_i(x_i)$. The assumption that $g_{\ell_i,x_i,y_i}(\mu) = \frac{\lambda}{1-\mu}$ for $i = 1, 2$ means, as in (14), that

$$\mathrm{SC}(\vec{u}) = \lambda \cdot \Phi + \mu \cdot \mathrm{SC}(\vec{u}) + \Delta\,,$$

where $\Phi = \sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot y_i \cdot \ell_i(y_i)$ and $\Delta = -\sum_{i=1,2} n \cdot p_i \cdot \sigma_i \cdot h_i$. That is,

$$\frac{\mathrm{SC}(\vec{u})}{\Phi} = \frac{\lambda}{1-\mu} + \frac{\Delta}{\Phi \cdot (1-\mu)} .$$

Assuming condition 4, we have $\Phi \xrightarrow{n,p_2\to\infty} \mathrm{SC}(\vec{v})$. Thus, the present condition follows provided $\frac{\Delta}{\Phi} \xrightarrow{n,p_2\to\infty} 0$. Recalling that $h_2 = -\omega \cdot h_1$, if we set $p_1 \approx p_2 \cdot \omega$, then

$$|\Delta| \leq n \cdot p_2 \cdot h_2 \cdot |\sigma_2 - \sigma_1| .$$

Consequently, $\frac{\Delta}{\Phi} \xrightarrow{n,p_2\to\infty} 0$ provided $\sigma_1 \xrightarrow{n,p_2\to\infty} 1$. (Recall that always $\sigma_2 = 1$.) We check that this is indeed the case below.

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 now argue that all six conditions can indeed be satisfied simultaneously. Choose values for the scaling parameters $n, p_2 \in \mathbb{N}$. Set $\gamma_1 = -\frac{h_1}{n}$ and $\gamma_2 = 0$. Next set p_1 according to condition 5 (as $\approx p_2 \cdot \omega$), q_i according to (19) in condition 3, t_i as $\approx p_i(n-1)/q_i$, and α_i, β_i to satisfy the simultaneous equations (15) and (17). (If $x_i < \frac{y_i}{2}$ or $\ell_i'(x_i) = 0$, then equation (17) is replaced by the equation $\beta_i = 0$.) Set $\sigma_2 = 1$ and σ_1 according to (16) of condition 2. Now, conditions 1–3 imply also condition 4, as shown above. Condition 5 reduces to showing that $\sigma_1 \xrightarrow{n,p_2\to\infty} 1$. After solving for σ_1 in (16), this follows since $\gamma_1 \xrightarrow{n,p_2\to\infty} 0$ and $\frac{p_2}{p_1} \xrightarrow{n,p_2\to\infty} \frac{1}{\omega}$ by definition, $\alpha_2 \xrightarrow{n,p_2\to\infty} \min\{\frac{y_2}{2}, x_2\}$ by condition 3, and using the definition of ω. Finally, consider the non-negativity constraints in condition 6. These hold for γ_1, γ_2 by definition and for α_1, α_2 by condition 3. For β_i, we can assume that $x_i \geq \frac{y_i}{2}$ and $\ell_i'(x_i) > 0$, as otherwise $\beta_i = 0$. Since $\gamma_2 = 0$, equation (15) and condition 3 imply that $\beta_2 \geq 0$. For $i = 1$, we have $x_1 - \alpha_1 \xrightarrow{n,p_2\to\infty} x_1 - \frac{y_1}{2}$ and $n \cdot \gamma_1 = -h_1 = x_1 - y_1 - \frac{y_1^2 \cdot \ell_1'(x_1)}{4} < x_1 - \frac{y_1}{2}$; inspecting (15) shows that $\beta_1 \geq 0$. This verifies the construction and completes the proof. ■

Remark 5.4 (Network Congestion Games) Since each player's basic strategies in this construction are disjoint, these congestion games can be represented as (directed) network congestion games: orient both cycles, give each player its own source and sink vertices (outside the cycles), and paths corresponding to its basic strategies.

5.2.3 An Edge Case

Before combining our results into a generally applicable lower bound, we need to give a related construction for the sets of cost functions $\mathcal{L}$ with no triples ℓ, x, y such that $h_{\ell,x,y} < 0$. The next lemma shows that, in this case, there is a family of games that admit strategy profiles with a per-resource cost approaching zero and Nash equilibria with positive per-resource cost (bounded away zero). Thus, the worst-case POA is $+\infty$ with respect to such sets of cost functions. This special case does not require scale-invariance.

Lemma 5.5 *Let ℓ be a cost function so that $h_{\ell,x,y} > 0$ for every admissible triple ℓ, x, y. There is a sequence of congestion games using only the cost function ℓ and with infinite limiting price of anarchy.*

Proof: Clearly, ℓ is not the zero function. Moreover, $\xi := \max\{x \mid \ell(x) = 0\} > 0$ (and this is well defined). To see this, suppose for contradiction that $\ell(y) > 0$ for all $y > 0$. Then, for fixed $x > 0$ and arbitrary $y > 0$, we have $h_{\ell,x,y} = (y-x)\cdot\ell(x)+\kappa(x,y)\cdot\ell'(x) \xrightarrow{y\searrow 0} -x\cdot\ell(x) < 0$, a contradiction.

We give a sequence of instances similar to but simpler than the lower-bound construction in Theorem 5.3. There is only one group of resources and players. As in the previous construction, we leave open several parameters to enable limiting arguments:

- The number of players and resources is an odd number n.
- The load on each resource in the Nash equilibrium is denoted by $\widehat{x}$ and will approach $\frac{3\xi}{2}$.
- The load each player puts on its "optimal" strategy in the Nash equilibrium is α and will approach $\frac{\xi}{2}$.

All other parameters are defined as follows.

- The size of the "optimal" strategy of each player is $p = 1$.
- The size of the "non-optimal" strategy of each player is $q = 2$.
- Each player has $t = \frac{p\cdot(n-1)}{q}$ non-optimal strategies.
- The load each player puts on each of its "non-optimal" strategies is $\beta = \frac{\widehat{x}-\alpha}{n-1}$.
- The load on each resource in the optimum is equal to the weight of each player, which is $w = \alpha + t\cdot\beta$.

For a given choice of $\widehat{x}$ and α, the corresponding strategy profile is a Nash equilibrium if the variational inequality (2) — corresponding to condition (17) in Theorem 5.3 — holds with equality:

$$\ell(\widehat{x}) + \alpha\cdot\ell'(\widehat{x}) = 2\cdot\left(\ell(\widehat{x}) + \frac{\widehat{x}-\alpha}{n-1}\cdot\ell'(\widehat{x})\right), \quad \text{i.e.,} \quad \ell'(\widehat{x}) = \frac{\ell(\widehat{x})}{\alpha} + \frac{2\cdot(\widehat{x}-\alpha)}{\alpha\cdot(n-1)}\cdot\ell'(\widehat{x}). \qquad (21)$$

Every triple ℓ, x, y with $x \ge y > \xi$ is admissible and, by assumption, satisfies $h_{\ell,x,y} = (y-x)\cdot\ell(x) + \frac{y^2}{4}\cdot\ell'(x) > 0$. Due to continuity of $h_{\ell,x,y}$ in y, the previous inequality also holds (not necessarily strictly) for $y = \xi$; that is, $\ell'(x) \ge \frac{4}{\xi^2}\cdot(x-\xi)\cdot\ell(x)$. Hence, for every $x > \frac{3\xi}{2}$ we have $\ell'(x) > \frac{2}{\xi}\cdot\ell(x)$.

By the previous observation, for every $\delta > 0$ we can choose $\widehat{x} \in [\frac{3\xi}{2}, \frac{3\xi}{2}+\delta)$ so that $\ell'(\widehat{x}) > \frac{2}{\xi}\cdot\ell(\widehat{x})$. Thus, we can choose $n \in \mathbb{N}$ large enough so that

$$\ell'(\widehat{x}) > \frac{2\cdot\ell(\widehat{x})}{\xi} + \frac{4\cdot(\widehat{x}-\frac{\xi}{2})}{\xi\cdot(n-1)}\cdot\ell'(\widehat{x}).$$

Since the right-hand side of (21) is continuous and monotonically decreasing in α, and unbounded for $\alpha \searrow 0$, we can find $\alpha \in (0, \frac{\xi}{2})$ so that (21) holds with equality.

Recall that the weight of each player is

$$w = \alpha + t\cdot\beta = \alpha + \frac{p\cdot(n-1)}{q}\cdot\frac{\widehat{x}-\alpha}{n-1} = \frac{\widehat{x}+\alpha}{2} < \frac{2\widehat{x}+\xi}{4} < \xi + \frac{\delta}{2}.$$

Consequently, we can find a sequence of games so that the load on each resource in some Nash equilibrium approaches $\frac{3\xi}{2}$, while the load on each resource in a different strategy profile approaches ξ. Since $\ell(\frac{3\xi}{2}) > 0$, $\ell(\xi) = 0$, and cost functions are continuous, the POA grows without bound as $\delta \to 0$ and $n \to \infty$. ■

5.2.4 Putting It All Together

We can now prove the main result of this section.

Corollary 5.6 (Tight Lower Bound) *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 Corollary 4.3. For the lower bound, the special case in which $\mathcal{L}$ does not admit any triples ℓ, x, y with $h_{\ell,x,y} < 0$ is addressed by Lemma 5.5. In the rest of the proof, we assume that there is an admissible triple ℓ, x, y with $h_{\ell,x,y} < 0$.

We show that, for any two triples ℓ_1, x_1, y_1 and ℓ_2, x_2, y_2 produced by Lemma 5.1, there are triples $\widehat{\ell}_1, \widehat{x}_1, \widehat{y}_1$ and $\widehat{\ell}_2, \widehat{x}_2, \widehat{y}_2$ that can be used in the lower-bound construction of Theorem 5.3 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 Lemma 5.1. Now set $\widehat{\ell}_2(x) := \frac{1}{\ell_2(x_2)} \cdot \ell_2(x)$, $\widehat{x}_2 = x_2$, $\widehat{y}_2 = y_2$. Define ω as in Theorem 5.3 in terms of $\widehat{\ell}_2, \widehat{x}_2, \widehat{y}_2$. Let

$$\tau := \frac{-h_{\ell_1,x_1,y_1} \cdot \omega}{\ell_1(x_1) \cdot h_{\widehat{\ell}_2,\widehat{x}_2,\widehat{y}_2}} .$$

Let $\widehat{\ell}_1(x) := \frac{1}{\ell_1(x_1)} \cdot \ell_1(\tau \cdot x)$, $\widehat{x}_1 = \frac{x_1}{\tau}$, $\widehat{y}_1 = \frac{y_1}{\tau}$. Then

$$h_{\widehat{\ell}_2,\widehat{x}_2,\widehat{y}_2} = \frac{-h_{\ell_1,x_1,y_1} \cdot \omega}{\ell_1(x_1) \cdot \tau} = -h_{\widehat{\ell}_1,\widehat{x}_1,\widehat{y}_1} \cdot \omega ,$$

as needed. ■

5.2.5 Example: Cubic Cost Functions

We give an example of our lower-bound construction when $\mathcal{L}$ consists of the cubic monomials $\{ax^3 : a \geq 0\}$. Monomial cost functions are a "lucky case" where, in Theorem 5.3, we can take $h_{\ell_i,x_i,y_i} = 0$. In such cases, similarly to the construction in Lemma 5.5, only one cycle of resources is needed and the scale-invariance hypothesis can be dropped.

Consider the admissible triple ℓ, x, y with $\ell(z) = z^3$, $x = \frac{3}{2}$, $y = 1$. It is easy to verify that

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

the function $g_{\ell,x,y}$ is identically equal to $(\frac{3}{2})^4 = 5.0625$. Choose $\lambda, \mu \in (0, 1)$ 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 each of 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 that 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 in Section 6, this construction shows that the price of anarchy for splittable congestion games with polynomial cost functions of degree at most 3 is *exactly* $(\frac{3}{2})^4 = 5.0625$.

5.2.6 Construction with Singleton Strategies

Continuing with the "lucky case" of the previous section (including monomial cost functions), we reimpose the scale-invariance assumption and give a tight lower-bound construction that uses only singleton strategies.

Theorem 5.7 *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*

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

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

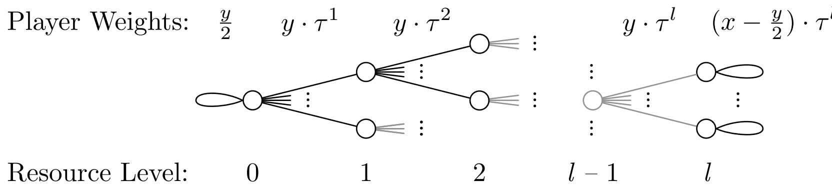

Figure 3: Illustration of construction with singleton strategies

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 player are its (at most two) incident nodes. Figure 3 illustrates the construction.

Let $\sigma, \tau > 0$ be values to be determined later (dependent on k and l). The cost function for resources at 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 a strategy has weight $\frac{y}{2}$, and the players who only have a leaf resource as a 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

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

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 the k^j resources at 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 further away from the root. Reasoning as above, the social cost contributed by the k^j resources at level $j \in [l-1]$ approaches $y \cdot \ell(y)$. 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 a constant independent of l.

Consequently, as $l \to \infty$ and $k \to \infty$ suitably quickly in l,the ratio of the social cost in the Nash equilibrium and that of the other profile approaches $\frac{x \cdot \ell(x)}{y \cdot \ell(y)} = g_{\ell,x,y}(\mu) = \frac{\lambda}{1-\mu}$. ■

6 Polynomial Cost Functions

This section gives a closed-form expression for the exact price of anarchy — that is, analytically evaluates the parameter $\gamma(\mathcal{L})$ — when the cost functions are polynomials with degree at most $d \in \mathbb{N}$ and non-negative coefficients. For $d \in \mathbb{N}$, let $\mathcal{P}_d$ denote this set of cost functions. Also, we write X^d to denote the monomial function $x \mapsto x^d$, and we let $\mathcal{M}_d := \{X^d, X^{d-1}, \ldots, X^0\}$ be the set of all monomials of degree at most d. We define Ψ_d as the unique positive real x with $x^d + \frac{d \cdot x^{d-1}}{4} = x^{d+1}$, that is, as $\Psi_d := \frac{1}{2}(1 + \sqrt{d+1})$. To save work, we let $g^*_{\ell,x,y}$ denote $g_{\ell,x,y}$, as defined in Section 4, except with $\kappa(x,y)$ replaced by $\frac{y^2}{4}$. We similarly define $\gamma^*(\mathcal{L})$ (cf., (11)). $h^*_{\ell,x,y}$ (cf., (12)), and $\Gamma^*_{\mathcal{L}}$. We start with three lemmas to simplify $\gamma^*(\mathcal{P}_d)$. In the end, it will turn out that $\gamma(\mathcal{P}_d) = \gamma^*(\mathcal{P}_d)$. The point of the next lemma is to give a closed-form formula for the function $\mu \mapsto \sup_{x \geq 0} g^*_{X^d,x,1}(\mu)$.

Lemma 6.1 *Let $\mu \in (0,1)$ and $d \geq 1$. Define $g : \mathbb{R}_{\geq 0} \to \mathbb{R}$ by $g(x) := x^d + \frac{d \cdot x^{d-1}}{4} - \mu \cdot x^{d+1}$. Then, 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 $g(\Psi_d) = (1 - \mu) \cdot \Psi_d > 0 = g(0)$. 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 $x > 0$, the necessary first-order condition for a local extremum is

$$g'(x) = dx^{d-2}\left(x + \frac{d-1}{4}\right) - \mu(d+1)x^d = 0\,. \tag{22}$$

Indeed, ξ is the unique positive value for x that satisfies (22). ■

The next lemma shows that we can restrict attention to monomial cost functions and admissible triples ℓ, x, y in which $y = 1$.

Lemma 6.2 *Let $d \in \mathbb{N}$. Then,*

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

Proof: We can rewrite

$$\gamma^*(\mathcal{P}_d) = \inf_{(\lambda,\mu) \in \mathbb{R} \times (0,1)} \left\{ \frac{\lambda}{1-\mu} \;\middle|\; \forall \ell \in \mathcal{P}_d, x \geq 0, y > 0 : \lambda \geq \frac{y \cdot \ell(x) + \frac{y^2 \cdot \ell'(x)}{4} - \mu \cdot x \cdot \ell(x)}{y \cdot \ell(y)} \right\}. \tag{23}$$

The defining condition in (23) holds for a given (λ, μ) if and only if it holds with ℓ restricted to $\mathcal{M}_d$. This implies the first equality in the lemma statement. Moreover, when ℓ is constant (and non-zero),

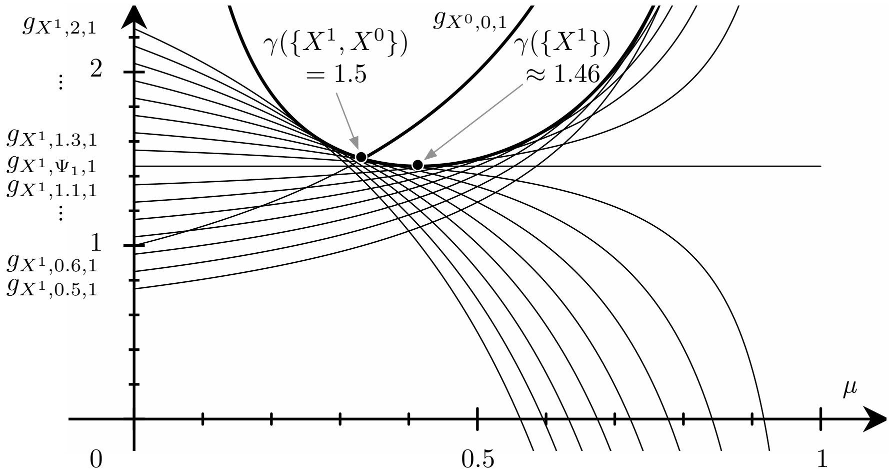

Figure 4: The functions $g_{\ell,x,y}$ when ℓ is the identity or a constant function, and the corresponding upper-envelope function (the thick line). Precisely, the envelope function here turns out to be $\mu \mapsto \frac{1+\mu}{4\cdot\mu\cdot(1-\mu)}$.

the inequality boils down to $\lambda \geq 1 - \mu \cdot \frac{x}{y}$ for all $x \geq 0$ and $y > 0$. Consequently, this defining condition is equivalent to

$$\forall r \in [d], x \geq 0, y > 0 : \lambda \geq \frac{y \cdot x^r + \frac{y^2 \cdot r \cdot x^{r-1}}{4} - \mu \cdot x^{r+1}}{y^{r+1}} \quad \text{and} \quad \lambda \geq 1\,. \tag{24}$$

In (24), the values $\frac{x}{y}$ and 1 yield the same inequality as the values x and y. We can therefore fix $y = 1$ without loss of generality. Consequently,

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

■

Lemma 6.3 *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

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

By construction, every ξ fulfills the necessary first-order condition (22) for local extrema of the function $x \mapsto g^*_{X^d,x,1}(\mu_\xi)$. By Lemma 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 $\gamma^*(\{X^d\}) \leq g^*_{X^d,\xi,1}(\mu_\xi) = \Psi_d^{d+1}$, with the equality holding by the definition of Ψ_d. Since $h^*_{X^d,\xi,1} = 0$, $g^*_{X^d,\xi,1}$ is a constant function and $\Gamma^*_{\{X^d\}}(\mu) \geq \Psi_d^{d+1}$ for every $\mu \in (0,1)$. Thus, $\gamma^*(\{X^d\}) = \Psi_d^{d+1}$.

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 $g^*_{X^0,0,1}(\frac{1}{3}) = \frac{3}{2} = g^*_{X^d,\xi,1}(\frac{1}{3})$. Because $g^*_{X^0,0,1}$ and $g^*_{X^d,\xi,1}$ are increasing and decreasing functions, respectively, $\gamma^*(\{X^d, X^0\}) = \frac{3}{2}$.
   Otherwise, if $d \geq 2$, choose $\xi := \Psi_d$ as in the first step. 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 in step 1, we have $\gamma^*(\{X^d, X^0\}) = \Psi_d^{d+1}$.

3. For $x < \frac{y}{2}$, we have $\kappa(x,y) \leq \frac{y^2}{4}$. Therefore, for every admissible triple ℓ, x, y we have $g_{\ell,x,y} \leq g^*_{\ell,x,y}$ pointwise, with equality holding whenever $\frac{x}{y} \geq \frac{1}{2}$. Hence, when $\xi \geq \frac{1}{2}$, we have $g_{X^d,\xi,1}(\mu_\xi) = \max_{x \in \mathbb{R}_{\geq 0}}\{g_{X^d,x,1}(\mu_\xi)\}$. Since the arguments above use values of ξ larger than $\frac{1}{2}$, they extend to the computation of γ.

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 $\xi > 1$, as it is in all computations above, then

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

■

Corollary 5.6, Lemma 6.2, and Lemma 6.3 immediately imply:

Corollary 6.4 *The following exact bounds on the worst-case price of anarchy in splittable congestion games with cost functions in $\mathcal{L}$ hold.*

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 $\gamma(\mathcal{L}) = \Psi_d^{d+1} = (\frac{1+\sqrt{d+1}}{2})^{d+1}$.*

7 Future Directions

We conclude with three proposals for further work. First, it would be interesting to discover more applications of the local smoothness framework defined in Section 3. One such application was given recently by Bhawalkar et al. [5], who used the framework to obtain tight bounds on the POA in a family of opinion formation games. In these games, each player i has an intrinsic opinion $s_i \in [0,1]$ and expresses a (possibly different) opinion $z_i \in [0,1]$. A player is interested both in how similar its expressed opinion is to its intrinsic one, and how its expressed opinion compares to those expressed by other players. Formally, the cost to player i in the strategy profile $\vec{z}$ has the form $g_i(z_i - s_i) + \sum_{j \neq i} f_{ij}(z_i - z_j)$, where g_i and f_{ij} are given cost functions. Bindel et al. [6] were the first to study the POA in such games, and they give exact worst-case bounds when $g_i(x) = x^2$ and $f_{ij}(x) = w_{ij}x^2$, where w_{ij} is a player pair-specific weight. Bhawalkar et al. [5] used the local smoothness framework to obtain tight POA bounds for all convex cost functions.

Second, while the present work obtains tight POA bounds for the correlated equilibria of splittable congestion games, the analogous question for coarse correlated equilibria remains open. We showed that local smoothness bounds do not extend to coarse correlated equilibria in general (Example 3.3), but we have not found an analogous example in a splittable congestion game. Very recently, von Falkenhausen and Roughgarden [30] showed that, in splittable congestion games with affine cost functions, every coarse correlated equilibrium is a mixture of Nash equilibria and hence the POA bound of $\frac{3}{2}$ applies. With nonlinear cost functions, however, there are splittable congestion games that possess coarse correlated equilibria that are costlier than all of their correlated equilibria [30]. The examples in [30] do not prove that the worst-case POA for coarse correlated equilibria is larger than that for correlated equilibria, however.

Finally, it would be interesting to resolve the worst-case POA in splittable congestion games in which every player has the same set of basic strategies. In symmetric games, where every player also has the same weight, the worst-case POA is identical to that in nonatomic congestion games [8]. With identical basic strategies but different player weights, it remains open to improve over the upper bounds of [8, 14] and the present work for general splittable congestion games, or over the lower bounds of [29] for nonatomic congestion games.

Acknowledgments We thank Kshipra Bhawalkar, Martin Gairing, Tobias Harks, and Uri Nadav for helpful discussions, and the anonymous SODA and JET reviewers for a number of useful remarks.

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] 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] U. Bhaskar, L. Fleischer, D. Hoy, and C. 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] K. Bhawalkar, M. Gairing, and T. Roughgarden. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. In *Proceedings of the 18th Annual European Symposium on Algorithms (ESA)*, pages 17–28, 2010.

[5] K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *45th ACM Symposium on Theory of Computing (STOC)*, 2013. To appear.

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

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

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

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

[10] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games. In *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO'05)*, volume 3509 of *LNCS*, pages 167–181, 2005.

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

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

[13] P. T. Harker. Multiple equilibrium behaviors on networks. *Transportation Science*, 22(1):39–46, 1988.

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

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

[16] A. Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 89–98, 2006.

[17] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th International Symposium on Theoretical Aspects of Computer Science (STACS'99)*, volume 1563 of *LNCS*, pages 404–413, 1999.

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

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

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

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

[22] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

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

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

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

[26] T. Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 1184–1185, 2005.

[27] T. 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.

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

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

[30] P. von Falkenhausen and T. Roughgarden. Convergence of regret-minimizing algorithms in atomic splittable routing games. Submitted, 2013.

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