

Existence and Efficiency of Equilibria for Cost-Sharing in Generalized Weighted Congestion Games

MARTIN GAIRING, University of Liverpool, UK
 KOSTAS KOLLIAS, Google Research, USA
 GRAMMATEIA KOTSIALOU, King's College London, UK

This work studies the impact of cost-sharing methods on the existence and efficiency of (pure) Nash equilibria in weighted congestion games. We also study *generalized weighted congestion games*, where each player may control multiple commodities. Our results are fairly general, we only require that our cost-sharing method and our set of cost functions satisfy certain natural conditions. For general weighted congestion games, we study the existence of pure Nash equilibria in the induced games and we exhibit a separation from the standard single-commodity per player model by proving that the *Shapley value* is the only cost-sharing method that guarantees existence of pure Nash equilibria. With respect to efficiency, we present general tight bounds on the price of anarchy, which are robust and apply to general equilibrium concepts. Our analysis provides a tight bound on the price of anarchy, which depends only on the used cost-sharing method and the set of allowable cost functions. Interestingly, the same bound applies to weighted congestion games and generalized weighted congestion games. We then turn to the price of stability and prove an upper bound for the Shapley value cost-sharing method, which holds for general sets of cost functions and which is tight in special cases of interest, such as bounded degree polynomials. Also for bounded degree polynomials, we provide a somewhat surprising result, showing that a slight deviation from the Shapley value has a huge impact on the price of stability. In fact, for this case, the price of stability becomes as bad as the price of anarchy. Again, our bounds on the price of stability are independent on whether players are single or multi-commodity.

CCS Concepts: • **Theory of Computation** → **Algorithmic game theory; Network Games; Quality of equilibria; Solution concepts in game theory;**

Additional Key Words and Phrases: existence of equilibria; price of anarchy; price of stability; congestion games; selfish routing; Shapley value; multi-commodity players;

ACM Reference Format:

Martin Gairing, Kostas Kollias, and Grammateia Kotsialou. 2020. Existence and Efficiency of Equilibria for Cost-Sharing in Generalized Weighted Congestion Games. *ACM Transactions on Economics and Computation* 0, 0, Article TEAC-D-18-00044R1 (2020), 28 pages. <https://doi.org/0000001.0000001>

1 INTRODUCTION

The class of weighted congestion games [27, 29] encapsulates a large collection of important applications in the study of solution concepts and inefficiencies induced by strategic behavior in large systems. The applications that fall within this framework involve a set of players who place demands on a set of resources and one of the most prominent such applications is *selfish routing* in

This work was supported by EPSRC grants EP/J019399/1 and EP/L011018/1 and combines results from the two conference papers [12] and [13].

Authors' addresses: Martin Gairing, University of Liverpool, UK, gairing@liverpool.ac.uk; Kostas Kollias, Google Research, USA, kostaskollias@google.com; Grammateia Kotsialou, King's College London, UK, grammateia.kotsialou@kcl.ac.uk.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

XXXX-XXXX/2020/0-ARTTEAC-D-18-00044R1 \$15.00

<https://doi.org/0000001.0000001>

a telecommunication or traffic network [3, 8, 34]. When total demand on a resource is increased, the resource becomes more scarce and the quality of service experienced by its users degrades. More specifically, in a weighted congestion game, there is a set of players N and a set of resources E . Each player $i \in N$ has a positive weight w_i and she gets to select the subset of the resources that she will use. The possible subsets of resources are given in her set of possible strategies \mathcal{P}_i . Once players make their decisions, each resource $e \in E$ generates a joint cost $C_e(f_e)$, where f_e is the total weight of the users of e and C_e is the cost function of e . The joint cost of resource e is paid by her users S_e , i.e., $\sum_{i \in S_e} \chi_{ie} = C_e(f_e)$, where χ_{ie} is the *cost share* of player i on resource e .

In many applications of interest, there is motivation to focus on a more general model where each player controls multiple *commodities* with their own weight and set of available strategies. In the selfish routing example, players might control multiple flows in the graph and each commodity has its own source, destination and weight. Settings such as routing road traffic controlled by competing ride-sharing applications or routing Internet traffic controlled by competing service providers lie within this framework. In this more general model, f_e^i is the combined flow placed on resource e by player i , which might be the sum of multiple commodities. The total flow on the resource is then $f_e = \sum_{i \in N} f_e^i$ and, as previously, the joint cost must be covered by the players who assign some flow on e , i.e., $\sum_{i: f_e^i > 0} \chi_{ie} = C_e(f_e)$.

The way the cost shares χ_{ie} are calculated is given by the *cost-sharing method* used in the game. A cost-sharing method determines the cost-shares of the players on a resource, given the joint cost that each subset of them generates, i.e., the cost shares are functions of the state on that resource alone. In most applications of interest, it is important that the cost-sharing method possesses this *locality property*, since we expect the system designer's control method to scale well with the size of the system and to behave well as resources are dynamically added to or removed from the system. Altering our cost-sharing method of choice changes the individual player costs. Given that our candidate outcomes are expected to be game-theoretic equilibrium solutions, this modification of the individual player costs also changes the possible outcomes players can reach. The *price of anarchy* (PoA) and the *price of stability* (PoS) measure the performance of a cost-sharing method by comparing the worst and best equilibrium, respectively, to the optimal solution, and taking the worst-case ratio over all instances.

Certain examples of cost-sharing methods include *proportional sharing* (PS) and the *Shapley value* (SV). In PS, the cost share of a player is proportional to her flow, i.e., $\chi_{ie} = (f_e^i / f_e) \cdot C_e(f_e)$, while the SV of a player on a resource e is her average marginal cost increase over a uniform ordering of the players in S_e . Other than different PoA and PoS values, different cost-sharing methods also possess different equilibrium existence properties. The *pure Nash equilibrium* (PNE) is the most widely accepted solution concept in such games. In a PNE, no player can improve her cost with a unilateral deviation to another strategy. In a *mixed Nash equilibrium* (MNE) players randomize over strategies and no player can improve her expected cost by picking a different distribution over strategies. By Nash's famous theorem, a MNE is guaranteed to exist in every weighted congestion game. However, existence of a PNE is not guaranteed for some cost-sharing methods. As examples, in routing games with single-commodity players, PS does not guarantee existence of a PNE (see e.g. [11, 16] and [19] for a characterization), while the SV does for players with a single commodity [25]. Gopalakrishnan et al. [17] showed that only the class of *generalized weighted Shapley values* (see Section 2 for a definition) guarantees the existence of a PNE in games with single-commodity players.

As a metric that is worst-case by nature, the PoA of a cost-sharing method that does not always induce a PNE must be measured with respect to more general concepts, such as the MNE, which are guaranteed to exist. Luckily, PoA upper bounds are typically *robust* [31] which means they

apply to MNE and even more general classes (such as correlated and coarse-correlated equilibria). On the other hand, the motivation behind the study of the PoS assumes a PNE will exist, hence the PoS is more meaningful when the cost-sharing method does guarantee a PNE.

1.1 Assumptions

For most of our results, we naturally assume the following (see Section 2 for more details):

- (1) Every cost function in the game is continuous, nondecreasing, and convex.
- (2) All players are charged only according to their contribution to the joint resource cost.
- (3) The cost share of a player on a resource is a convex function of her weight on it.

Assumption 1 is standard in congestion-type settings. For example, linear cost functions have obvious applications in many network models, as do queuing delay functions, while higher degree polynomials (such as quartic) have been proposed as realistic models of road traffic [36]. Assumption 2 states that the cost share of a player only depends on how much the player impacts the joint cost function of a resource and there is no discrimination between players in any other way. Assumption 3 states that the curvature of the cost shares function should be consistent with the convexity of the resource cost function (given by Assumption 1), i.e., the share of a player on a resource is a convex function of her weight. Without this assumption, we land in a setting where increasing the weight of a player i will increase i 's cost share at a slower rate than that of another player j whose weight remains constant, something we view as unnatural.

1.2 Contribution

Our contributions lie on two aspects: existence and efficiency of equilibria. In particular, we focus on the existence of pure Nash equilibria, on the price of anarchy and the price of stability.

PNE existence: It is known that generalized weighted Shapley values are the only methods that guarantee existence of a PNE in games with single-commodity players [17]. In this work we prove that the SV continues to guarantee the existence of a PNE even in games with multi-commodity players and cost functions as stated in Assumption 1 (Theorem 3.1). We then exhibit a separation from the single-commodity case by proving that a subclass of weighted Shapley values do not guarantee existence of a PNE when we have multi-commodity players (Theorem 3.2). Given that the class of weighted Shapley values are the unique cost-sharing methods that guarantee pure Nash equilibria in the single-commodity player model [17], our results suggest that the Shapley value is essentially the unique anonymous cost-sharing method that guarantees pure Nash equilibria in the multi-commodity player model.

General PoA bounds: On the PoA side, we present *tight* bounds for general classes of allowable cost functions and for general cost-sharing methods (satisfying Assumptions 1, 2, 3), i.e., we parameterize the PoA by (i) the set of allowable cost functions (which changes depending on the application under consideration) and (ii) the cost-sharing method (Theorem 4.2). We note that our upper bound is robust and applies to general equilibrium concepts that are guaranteed to exist for all cost-sharing methods. Our upper bound applies to games with multi-commodity players, while our lower bound uses an instance with single-commodity players, which proves that the PoA coincides for the two models.

PoS bounds for Shapley Values: Studying the PoS is best motivated in settings where a trusted mediator or some other authority can place the players in an initial configuration from which they will not be willing to deviate. For this reason, the PoS is a very interesting concept, especially for games possessing a PNE. Hence, we focus on cost-sharing methods which always induce games with a PNE (at least in the single-commodity model). For SV cost sharing, we prove an upper

bound on the PoS which holds for all sets of cost functions that satisfy Assumption 1 (Theorem 5.1). We show that for the interesting case of polynomials of bounded degree d , this upper bound is $d + 1$ (Corollary 5.3), which is asymptotically tight and always very close to the exact value given by Christodoulou and Gairing [7].

Moreover, we show that this linear dependence on the maximum degree d is very fragile. To do so, we consider a parameterized class of weighted Shapley values, where players with larger weight get an advantage or disadvantage, which is determined by a single parameter γ . When $\gamma = 0$ this recovers the SV. For all other values $\gamma \neq 0$, we show (Theorem 5.8 and Theorem 5.10) that the PoS is very close and, for $\gamma > 0$, even matches the upper bound on the *price of anarchy* by Gkatzelis et al. [15]. In other words, for this case the PoS and the PoA coincide, which we found very surprising due to the fact that the upper bound in [15] applies even to general cost-sharing methods. We note that these weighted Shapley values are the only cost-sharing methods that guarantee existence of a PNE in games with single-commodity players and satisfy Assumption 2 [15, 17]. Similarly to our PoA bounds, our PoS upper bound applies to games with multi-commodity players, while the lower bounds construct instances with single-commodity players, hence obtaining the desired generality.

1.3 Related work and comparison to our results

In this section, we discuss the related work (which is mostly for the single-commodity case) on PNE existence, price of anarchy and price of stability, and compare it with our findings.

PNE existence: In weighted congestion games with single-commodity players, the most common method for sharing the total cost of a resource among her users is the proportional sharing, which, unfortunately, lacks the desirable property of guaranteeing existence of pure Nash equilibria (see examples by Fotakis et al. [11], and Goemans et al. [16]). However, PNE exist, if we allow only affine [11] or only exponential resource cost functions [19]; Harks and Klimm [19] show that those are the only classes of costs functions that guarantee this property. Harks et al. [20] characterize the existence of potential functions in weighted congestion games under proportional sharing. Using Shapley values instead of proportional sharing, Kollias and Roughgarden [25] show that also such games are potential games. A characterization for a much wider class of cost-sharing methods was given by Gopalakrishnan et al. [17]. More specifically, [17] shows that the only cost-sharing method that guarantees pure Nash equilibria in games is a generalization of weighted Shapley values.

A number of previous works [11, 23, 30, 37] study settings that share similarities to routing games with multi-commodity players. Rosenthal [30] studies weighted congestion games where each player may split her integer flow size among different sub-flows of integer size. Focusing on proportional sharing, Rosenthal [30] proves that there exist such games with no PNE. Tran-Thanh et al. [37] identify special cases where PNE exist in Rosenthal's model. Our approach differs from the work in [30] and [37], as we allow multi-commodity players, i.e., we do not restrict to players who control unit flows with the same start vertex. Fotakis et al. [11] focus on coalitions of atomic players in routing games (equivalent to multi-commodity players) and mostly on the objective of minimizing the maximum cost. For the sum of costs objective (which we consider in this paper), they prove that the game always admits a PNE under PS and quadratic edge cost functions. With respect to existence of PNE, we provide comprehensive results for more general methods and more general classes of cost functions.

PoA: The notion of price of anarchy was proposed by Koutsoupias and Papadimitriou in [26], and most of the work on efficiency of equilibria for weighted congestion games has focused on proportional sharing and single-commodity players. More specifically, for linear cost functions, Awerbuch et al. [3] and, Christodoulou and Koutsoupias [8] obtain tight bounds, while Gairing

and Schoppmann [14] prove PoA bounds for multiple classes of singleton congestion games. For the interesting class of polynomial cost functions, Aland et al. [1] provide the exact PoA value. Roughgarden [31] generalized the technique used in [1, 3, 8] to hold for general cost functions, defined it as the *smoothness* framework, and showed that PoA bounds obtained in this way are robust, i.e., apply to general equilibrium concepts. Bhawalkar et al. [4] characterize the PoA in weighted atomic congestion games as a function of the allowable resource cost functions.

Using Shapley values in weighted congestion games (as an alternative to proportional sharing), is an approach introduced by Kollias and Roughgarden [25], who also give tight bounds on the PoA. Gkatzelis et al. [15] show that, among all cost-sharing methods that guarantee PNE existence in weighted congestion games, the Shapley value minimizes the worst-case PoA. With respect to proportional sharing, they show that it is near optimal in general, for convex cost functions and single-commodity players. Klimm and Schmand [24], and Roughgarden and Schrijvers [33], also discuss the optimality of the SV for the extended model with non-anonymous costs given by set functions. Network cost-sharing games with fixed resource costs were studied by Chen et al. [6] while variants of it were studied by von Falkenhausen and Harks [38] and Christodoulou and Sgouritsa [9]. Note that our model is different to the ones with fixed resource costs as the latter allow concave resource cost functions (which we disallow). To see this, consider a resource with any fixed resource cost, i.e., $c_e(n_e) = 2$ for any $n_e > 0$ and $c_e(n_e) = 0$ for $n_e = 0$ (any empty resource generates zero cost), which gives a concave function. In our setting but for different cost-sharing methods, Harks and Miller [21] give efficiency bounds that depend on the cost functions and the cost-sharing methods. The methods they study are (a) *average cost-sharing*, (b) *marginal cost pricing* and (c) *incremental cost-sharing*. Similar to our generalized setting of congestion games with multi-commodity players, is the setting of Hayrapetyan et al. [23]. They show that there exist games where merging atomic players into a coalition may degrade the quality of the induced PNE when PS is used. In a small contrast, we focus on worst-case metrics and show that the PoA and PoS of multi-commodity games is no worse than the PoA and PoS of the single-commodity case, for general cost-sharing methods.

In addition to extending to the case of multi-commodity players, our PoA results greatly generalize the work on cost-sharing methods for weighted congestion games and give a recipe for tight bounds in a large array of applications. Prior to our work, only a handful of cost-sharing methods have been tightly analyzed. Our results facilitate the better design of such systems, beyond the optimality criteria considered by Gkatzelis et al. [15]. For example, the SV has the drawback that it cannot be computed efficiently, while PS (on top of not always inducing a PNE) might have equilibria that are hard to compute. In cases where existence and efficient computation of a PNE is considered important, the designer might opt for a different cost-sharing method, such as a priority based one (that fixes a weight-dependent ordering of the players and charges them the marginal increase they cause to the joint cost by entering the resource in this order), which has polynomial time computable shares and equilibria. Our results show how the inefficiency of equilibria is quantified for all such possible choices, to help evaluate the tradeoffs between different options. Our work closely parallels the work on network cost-sharing games by Kollias and Roughgarden [25], which provides tight bounds for general cost-sharing methods.

PoS. The term price of stability was introduced by Anshelevich et al. in [2] for the network cost-sharing games, which was further studied for weighted players and various cost-sharing methods by Chen and Roughgarden [5], Chen et al. [6], and Kollias and Roughgarden [25]. With respect to congestion games, results on the PoS are only known for polynomial *unweighted* games and single-commodity players, for which Christodoulou and Gairing [7] provide exact bounds.

Our PoS upper bound is the first for weighted congestion games that applies to any class of convex costs. Roughgarden and Schoppmann present SV PoS bounds in a more general setting with non-anonymous but submodular cost functions [33]. In a similar vein, Klimm and Schmand [24] present tight PoS bounds on the SV in games with non-anonymous costs, by allowing any cost function and parameterizing by the number of players in the game, i.e., they show that for the set of all cost functions the PoS of the SV is $\Theta(n \log n)$ and for the set of supermodular cost functions it becomes n , where n is the number of players. These upper bounds apply to our games as well, however we adopt a slightly different approach. We allow an infinite number of players for our bounds to hold and parameterize by the set of possible cost functions, to capture the PoS of different applications. For example, for polynomials of degree at most d , we show that the PoS is at most $d + 1$. Observe that for unweighted games PS and SV are identical. Thus, the lower bound by Christodoulou and Gairing [7], which approaches $d + 1$, also applies to our setting, showing that our bound for polynomials is asymptotically tight.

Our lower bounds on the PoS for the parameterized class of weighted Shapley values build on the corresponding lower bounds on the PoA by Gkatzelis et al. [15]. Our construction matches these bounds by ensuring that the instance possesses a unique Nash equilibrium. Together with our upper bound this shows an interesting contrast: For the special case of SV the PoS is exponentially better than the PoA, but as soon as we give some weight dependent priorities to the players, the PoA and the PoS essentially coincide.

2 PRELIMINARIES

In this section we present the notation and preliminaries for our model in terms of a weighted congestion game with multi-commodity players. In such a game, there is a set Q of k commodities which are partitioned into $n \leq k$ non-empty and disjoint subsets Q_1, Q_2, \dots, Q_n . Each set of commodities Q_i , for $i = 1, 2, \dots, n$, is controlled by an independent player. Denote $N = \{1, 2, \dots, n\}$ the set of players. The players in N share access to a set of resources E . Each resource $e \in E$ has a cost function $C_e : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$.

Strategies. Each commodity $q \in Q$ has a set of possible strategies $\mathcal{P}^q \subseteq 2^E$. Associated with each commodity q is a weight w_q , which has to be allocated to a strategy in \mathcal{P}^q . For a player i , a strategy $P_i = (P_q)_{q \in Q_i}$, defines the strategy for each commodity q player i controls. An outcome $P = (P_1, P_2, \dots, P_n)$ is a tuple of strategies of the n players.

Resource load. For an outcome P , the flow $f_e^i(P)$ of a player i on resource e equals the sum of the weights of all her commodities using e , i.e., $f_e^i(P) = \sum_{q \in Q_i, e \in P_q} w_q$. For any $T \subseteq N$, denote $f_e^T(P) = (f_e^i(P))_{i \in T}$ the vector of flows on e restricted to players in T . The total flow on a resource e is given by $f_e(P) = \sum_{i \in N} f_e^i(P)$, while the set of users of e (players who assign positive flow on e) for an outcome P is given by $S_e(P)$.

Cost shares. The cost-sharing method of the game determines how the flow-dependent joint cost of a resource $C_e(f_e(P))$ is divided among its users. For any $T \subseteq N$, let $f_e^T(P)$ be the vector of the flows that each player in T assigns to e . Then the cost share of player i is defined as a function of the player's identity, the vector of flows assigned to e and the resource's cost function, i.e., $\chi_{ie}(P) = \xi(i, f_e^T(P), C_e)$. Note that there is no resource dependent cost-sharing methods and that the cost shares are computed only based on the players in T . For a resource e , the cost shares are such that $\sum_{i \in S_e(P)} \chi_{ie}(P) = C_e(f_e(P))$. The total cost of a player is $X_i(P) = \sum_{e \in P_i} \chi_{ie}(P)$ and the social cost of the game is given by the sum of the player costs,

$$SC(P) = \sum_{i \in N} \sum_{e \in E} \chi_{ie}(P) = \sum_{i \in N} \sum_{e \in E} \xi(i, f_e^N(P), C_e) = \sum_{e \in E} C_e(f_e(P)). \quad (1)$$

On the *cost-sharing method* and the set of allowable cost functions, we make three natural assumptions:

- (1) The cost functions of the game are drawn from a given set C of allowable cost functions, such that every $C \in C$ must be continuous, increasing and convex. We also make the mild technical assumption that C is closed under dilation, i.e., that if $C(x) \in C$, then also $C(a \cdot x) \in C$ for $a > 0$. Without loss of generality, every C is also closed under scaling, i.e., if $C(x) \in C$, then also $a \cdot C(x) \in C$ for $a > 0$ (this is given by simple scaling and replication arguments).
- (2) The second assumption states that the cost-sharing method only charges players based on how they contribute to the joint cost. We formally state this assumption in two parts:
 - (a) Suppose resource e with players i_1, i_2, \dots, i_k and resource e' with players i'_1, i'_2, \dots, i'_k are such that, for every set of indices I , we have:

$$C_e \left(\sum_{j \in I} f_e^{i_j} \right) = C_{e'} \left(\sum_{j \in I} f_{e'}^{i'_j} \right).$$

Then it should be the case that $\chi_{i_j, e} = \chi_{i'_j, e'}$ for every index j .

- (b) Suppose we scale the joint cost on a resource e by a positive factor α , i.e., $\tilde{C}_e(f_e(P)) = \alpha \cdot C_e(f_e(P))$. The new cost shares of the players would be a scaled by factor α version of their initial cost shares, i.e., $\tilde{\chi}_{ie} = \alpha \cdot \chi_{ie}$. This is necessary for the cost-sharing method to be consistent, in the sense that, a resource with a cost function $\alpha \cdot C_e(f_e(P))$ should induce the same costs to the players as α resources with a cost function $C_e(f_e(P))$ each.
- (3) Last, we assume that the cost share of a player on a resource is a continuous, increasing and convex function of her flow. This is something to expect from a reasonable cost-sharing method, given that the joint cost on the resource is a continuous increasing convex function of the resource's total flow $f_e(P)$.

Pure Nash equilibrium. We now proceed with the definition of our solution concept. The *pure Nash equilibrium* (PNE) condition on an outcome P states that for every player i it must be the case that

$$\sum_{e \in E} \chi_{ie}(P) \leq \sum_{e \in E} \chi_{ie}(P'_i, P_{-i}), \text{ for any other strategy } P'_i. \quad (2)$$

Price of Anarchy and Price of Stability. Let \mathcal{Z} be the set of outcomes and \mathcal{Z}^N be the set of pure Nash equilibria outcomes of the game. Then the *price of anarchy* (PoA) and the *price of stability* (PoS) are defined as follows:

$$PoA = \frac{\max_{P \in \mathcal{Z}^N} SC(P)}{\min_{P \in \mathcal{Z}} SC(P)} \quad \text{and} \quad PoS = \frac{\min_{P \in \mathcal{Z}^N} SC(P)}{\min_{P \in \mathcal{Z}} SC(P)}. \quad (3)$$

The PoA and PoS for a class of games are defined as the largest such ratios among all games in the class.

Weighted Shapley values. The *weighted Shapley value* defines how the cost $C_e(\cdot)$ of resource e is distributed among the players using it. Given an ordering π of N , let $F_e^{<i, \pi}(P)$ be the sum of flows of players preceding i in π . Then the marginal cost increase caused by a player i 's flow is $C_e(F_e^{<i, \pi}(P) + f_e^i(P)) - C_e(F_e^{<i, \pi}(P))$. For a given distribution Π over orderings, the cost share of player i on e is $E_{\pi \sim \Pi} [C_e(F_e^{<i, \pi}(P) + f_e^i(P)) - C_e(F_e^{<i, \pi}(P))]$. For the weighted Shapley value, the distribution over orderings is given by a sampling parameter $\lambda_e^i(P)$ for each player i . The last flow in the ordering is picked proportional to the sampling parameters $\lambda_e^i(P)$. This process is then repeated iteratively for the remaining players.

As in [15], we study a parameterized class of weighted Shapley values defined by a parameter γ . For this class $\lambda_e^i(P) = f_e^i(P)^\gamma$ for all players i and resources e . For $\gamma = 0$, this reduces to the (unweighted) Shapley value, where we have a uniform distribution over orderings.

3 EXISTENCE OF PURE NASH EQUILIBRIA

In this section we focus on the existence of pure Nash equilibria. Recall that for the single-commodity case, [17] characterized a generalization of weighted Shapley values as the only cost-sharing methods to guarantee existence of pure Nash equilibria. As soon as we extend the model to congestion games with multi-commodity players, our results in this section show a separation to the result of [17]. We prove that the only cost-sharing method that admits pure Nash equilibria restricts to Shapley values, where all players have identical sampling weights. More specifically, while Shapley values still ensures PNE existence (see Theorem 3.1), already for a parameterized class of weighted Shapley values PNE might not exist (see Theorem 3.2). We close the section with a discussion on computing the cost shares on a per-commodity basis (see Theorem 3.4).

THEOREM 3.1. *Using the Shapley value to share player costs in congestion games with multi-commodity players yields a potential game.*

PROOF. Consider any ordering π of the players in N and let $f_e^{\leq i, \pi}(P)$ denote the vector that we get after truncating $f_e^N(P)$ by removing all entries for players that succeed i in π . We prove that the following function $\Phi()$ is a potential function of the game.

$$\Phi(P) = \sum_{e \in E} \sum_{i \in N} \xi(i, f_e^{\leq i, \pi}(P), C_e).$$

The potential property we wish to prove is that, when player i unilaterally changes the outcome of the game from P to (P'_i, P_{-i}) , we get:

$$\Phi(P) - \Phi(P'_i, P_{-i}) = \sum_{e \in E} \xi(i, f_e^N(P), C_e) - \sum_{e \in E} \xi(i, f_e^N(P'_i, P_{-i}), C_e).$$

It is clearly sufficient to prove this property per resource. Specifically, if the potential on each resource e is defined as:

$$\Phi_e(P) = \sum_{i \in N} \xi(i, f_e^{\leq i, \pi}(P), C_e), \quad (4)$$

then we need to prove:

$$\Phi_e(P) - \Phi_e(P'_i, P_{-i}) = \xi(i, f_e^N(P), C_e) - \xi(i, f_e^N(P'_i, P_{-i}), C_e). \quad (5)$$

Hart and Mas-Colell [22] proved that (4) is independent of the ordering π in which player flows are considered. We can then assume that π is such that player i goes last in the ordering and write $\Phi_e(P)$ as:

$$\Phi_e(P) = \sum_{j \in N, j \neq i} \xi(j, f_e^{\leq j, \pi}(P), C_e) + \xi(i, f_e^N(P), C_e).$$

Observe that using this last expression for $\Phi_e()$, we get:

$$\begin{aligned} \Phi_e(P) - \Phi_e(P'_i, P_{-i}) &= \sum_{j \in N, j \neq i} \xi(j, f_e^{\leq j, \pi}(P), C_e) + \xi(i, f_e^N(P), C_e) \\ &\quad - \sum_{j \in N, j \neq i} \xi(j, f_e^{\leq j, \pi}(P'_i, P_{-i}), C_e) - \xi(i, f_e^N(P'_i, P_{-i}), C_e). \end{aligned} \quad (6)$$

Note that, since i is the last one in the ordering and the only who changes strategy, we get $f_e^{\leq j, \pi}(P) = f_e^{\leq j, \pi}(P'_i, P_{-i})$. Using this fact in (6) recovers (5) and completes the proof. \square

One might expect that, similarly to standard congestion games, the same potential function argument would apply to weighted Shapley values as well. However, as we prove next, this is not the case.

THEOREM 3.2. *There is a congestion game with multi-commodity players that admits no PNE for any weighted Shapley value defined by sampling weights of the form $f_e^i(P)^\gamma$ with $\gamma \neq 0$.*

PROOF. We prove this theorem by showing two examples admitting no PNE, for $\gamma > 0$ and $\gamma < 0$. We start with the $\gamma > 0$ case. Consider two players, 1 and 2, who compete for two parallel (meaning each commodity must pick exactly one of them) resources e, e' with identical cost functions $C_e(x) = C_{e'}(x) = x^{1+\delta}$ with $\delta > 0$ and $\frac{\gamma}{\delta}$ a large positive number (note that for $\delta = 0$, we have linear cost functions where in this case we have an equilibrium. As soon as we deviate from linearity, we use convexity to construct an example with no equilibrium). Player 1 controls a unit commodity $p \in Q_1$. Player 2 controls two commodities $q, q' \in Q_2$, with $w_{q'} = 1$ and $w_q = k$, for k a very large number. Recall, that the sampling weight of a player i on a resource e is given by $\lambda_e^i = (f_e^i)^\gamma$. This means that smaller weights are favored when constructing the weighted Shapley ordering.

We continue by proving the following lemma which we use in the instance afterwards.

LEMMA 3.3. *For any positive δ and ϵ , there exists some k sufficiently large such that*

$$(1 + \delta) \cdot k^\delta + k^{1+\delta} \leq (k + 1)^{1+\delta} \leq (1 + \epsilon) \cdot (1 + \delta) \cdot k^\delta + k^{1+\delta}.$$

PROOF. Given any $\delta > 0$, we have that

$$\lim_{k \rightarrow \infty} \frac{(k + 1)^{1+\delta} - k^{1+\delta}}{(1 + \delta) \cdot k^\delta} = \lim_{k \rightarrow \infty} \frac{k^{1+\delta} \left(\left(1 + \frac{1}{k}\right)^{1+\delta} - 1 \right)}{(1 + \delta) \cdot k^\delta} = \lim_{k \rightarrow \infty} \frac{\left(1 + \frac{1}{k}\right)^{1+\delta} - 1^{1+\delta}}{(1 + \delta) \cdot \frac{1}{k}} = \frac{(x^{1+\delta})'|_{x=1}}{1 + \delta} = 1.$$

By definition of limit, we get that for every $\epsilon > 0$, there exists a M such that

$$\left| \frac{(k + 1)^{1+\delta} - k^{1+\delta}}{(1 + \delta) \cdot k^\delta} - 1 \right| < \epsilon \quad (7)$$

whenever $k > M$. Then, observe that both our inequalities follow from (7). \square

Suppose, without loss of generality, that player 1 places commodity p on resource e . Then the best response of player 2 is to place the large commodity q alone on e' and the small commodity q' on e . To see this, note that player 2's cost on this outcome will be

$$\frac{1}{2} \cdot (2^{1+\delta} - 1) + \frac{1}{2} + k^{1+\delta} = 2^\delta + k^{1+\delta}. \quad (8)$$

Then we show that any other outcome results in a larger cost for player 2. First, let her large commodity be assigned with commodity p of player 1 on resource e . This would result in a cost of at least

$$\frac{1}{1 + k^\gamma} \cdot k^{1+\delta} + \frac{k^\gamma}{1 + k^\gamma} \cdot ((k + 1)^{1+\delta} - 1) + 1. \quad (9)$$

By Lemma 3.3, we get that

$$(9) \geq \frac{k^{1+\delta} + k^\gamma \cdot ((1 + \delta) \cdot k^\delta + k^{1+\delta} - 1) + 1 + k^\gamma}{1 + k^\gamma} = k^{1+\delta} + \frac{(1 + \delta) \cdot k^{\gamma+\delta} + 1}{1 + k^\gamma}. \quad (10)$$

Note that for large k , the second term is larger than 2^δ . Therefore (10) > (8). If now both commodities of player 2 were assigned on e' , this would result in a cost of

$$(k+1)^{1+\delta} \stackrel{\text{Lemma 3.3}}{\geq} \left(k^{1+\delta} + (1+\delta) \cdot k^\delta \right) > k^{1+\delta} + 2^\delta = (8).$$

We focus now on player 1. Her cost on resource e , given player 2's best strategy, is

$$2^\delta, \tag{11}$$

which is a fixed number larger than 1, since $\delta > 0$. We show that player 1 would prefer to assign her commodity p on resource e' , together with the large commodity q of player 2. In this outcome, player 1's cost would be

$$\frac{k^\gamma}{1+k^\gamma} \cdot 1 + \frac{1}{1+k^\gamma} \cdot \left((k+1)^{1+\delta} - k^{1+\delta} \right) \stackrel{\text{Lemma 3.3}}{\leq} \frac{k^\gamma}{1+k^\gamma} + \frac{(1+\epsilon) \cdot (1+\delta) \cdot k^\delta}{1+k^\gamma}. \tag{12}$$

which approaches 1 for large enough k , since $\gamma > \delta$. Thus, (12) < (11) which proves the desirable deviation. Therefore there is no equilibrium.

We now switch to the case with $\gamma < 0$. Consider players $i = 1, 2, \dots, k$, who compete for two parallel resources e_1, e_2 with identical cost functions $C_{e_1}(x) = C_{e_2}(x) = x^3$. Player k controls two commodities $p, q \in Q_k$ with weights $w_p = k$ and $w_q = 1$. Each player $i < k$ controls only one commodity $r_i \in Q_i$ with $w_{r_i} = 1$. The sampling weight of a player i on a resource e is given by $\lambda_e^i = (f_e^i)^\gamma$, for $\gamma < 0$. Assume player k assigns commodity p to resource e_1 and commodity q to e_2 . In this case, note that it is a dominant strategy for players $1, \dots, k-1$ to use resource e_2 . For player k , this gives a cost share of

$$\frac{1}{k} \cdot C_{e_2}(f_{e_2}) + C_{e_1}(f_{e_1}) = k^2 + k^3. \tag{13}$$

On the other hand, if she assigns commodity p to e_2 and q to e_1 , by Lemma 5.6(b), her cost can be arbitrarily close to $1 + k^3$, for large k , which is strictly smaller than her previous cost (13). Therefore there is no PNE, which completes the proof of Theorem 3.2. \square

Alternative weighted Shapley value cost-sharing

One might consider a different way of generalizing weighted Shapley values to multi-commodity congestion games: Based on the vector of commodity flows, compute the weighted Shapley value for each commodity and charge a player the sum of the weighted Shapley values of the commodities controlled by her. When all commodities have the same weight, then this cost-sharing method coincides with proportional cost-sharing (i.e., every player pays a cost-share that is proportional to her flow on any given resource). Below we use one such instance with unit commodities to prove that all these methods do not guarantee pure Nash equilibrium existence.

Our instance is based on an example in [11], where Fotakis et al. prove that network unweighted congestion games with linear resource cost functions and equal cardinality coalitions do not have the *finite improvement property*, therefore they admit no potential function. Their example translates to a restricted setting of our model where each player controls an equal number of unit commodities. We strengthen their result by proving non-existence of pure Nash equilibria for congestion games with multi-commodity players and cubic resource cost functions (we construct even a network congestion game with no pure Nash equilibrium).

A similar example has already been given by Rosenthal [30]. However, Rosenthal's example uses concave cost functions, which we disallow in our setting. In contrast, our example only uses convex functions.

Table 1. Players' costs in example of Theorem 3.4.

| | P_1 | P_3 |
|-------|-----------------------------------|--|
| P_2 | $20 + 8$ $20 + 8$ | $5 + (13 - 9 \cdot \epsilon)$ $(16 - 9 \cdot \epsilon) + (13 - 9 \cdot \epsilon)$ |
| P_4 | $20 + 7$ $20 + (2 - \epsilon)$ | $5 + 13$ $(17 - 4 \cdot \epsilon) + (5 - 4 \cdot \epsilon)$ |




Fig. 1. The network congestion game in Corollary 3.5.

THEOREM 3.4. *There is a congestion game with multi-commodity players and cubic cost functions that admits no pure Nash equilibrium under weighted Shapley value applied on commodity weights.*

PROOF. We prove the theorem by constructing an instance with unit commodities such that best-response dynamics from any initial configuration cycles. This cycle is based on an example in [11], where Fotakis et al. prove that network unweighted congestion games with linear delays and equal cardinality coalitions do not have the *finite improvement property*, therefore they admit no potential function. Their model translates to a setting of our model where each player controls an equal number of unit commodities. We strengthen their result by proving non-existence of a pure Nash equilibria for (network) congestion games with multi-commodity players and cubic cost functions.

Consider two players, 1 and 2, who control two commodities each, $p, q \in Q_1$ and $p', q' \in Q_2$ with $w_p = w_q = w_{p'} = w_{q'} = 1$. There is a set of resources $E = \{1, 2, \dots, 6\}$ with cost functions $C_1(x) = C_5(x) = C_6(x) = 5 \cdot x^3$, $C_2(x) = x^3$, $C_4(x) = 2 \cdot x^3$ and $C_3(x) = (1 - \epsilon) \cdot x^3$, for small ϵ . The strategy sets for each commodity are: $\mathcal{P}^p = \{P_1\}$, $\mathcal{P}^q = \{P_2, P_4\}$, $\mathcal{P}^{p'} = \{P_1, P_3\}$ and $\mathcal{P}^{q'} = \{P_2\}$, where $P_1 = \{1\}$, $P_2 = \{2, 3\}$, $P_3 = \{3, 4, 5\}$ and $P_4 = \{4, 6\}$. Note that each player has a fixed strategy for the first commodity and two possible strategies for the second. Therefore there are four possible states. We model this instance as a bimatrix game in Table 1, where player 1 is the row player and 2 the column player. The costs are described as the sum of commodities' costs for each player. We claim that for any $\epsilon \in (\frac{1}{18}, \frac{3}{2})$, there is no pure Nash equilibrium. \square

We strengthen this result by showing non-existence of pure Nash equilibria even for *network* congestion games under this setting. The proof uses the network in Figure 1.

COROLLARY 3.5. *There is a network congestion game with multi-commodity players and cubic cost functions that admits no pure Nash equilibrium under weighted Shapley value applied on commodity weights.*

Table 2. Players' costs in example of Corollary 3.5.

| | P_1 | P_3 |
|-------|---|---|
| P_2 | $20 + 8$ $20 + 8$ | $5 + (13 - 9\epsilon)$ $(16 - 9\epsilon) + (13 - 9\epsilon)$ |
| P_4 | $20 + 7$ $20 + (2 - \epsilon)$ | $5 + 13$ $(17 - 4\epsilon) + (5 - 4\epsilon)$ |
| P_3 | $20 + (11 - 4\epsilon)$ $20 + (5 - 4\epsilon)$ | $5 + (37 - 9\epsilon)$ $(37 - 9\epsilon) + (10 - 9\epsilon)$ |

PROOF. Consider the same set of players, commodities, resources and the associated cost functions, with those described in the proof of Theorem 3.4. In addition, for each commodity z there is a source-destination pair (s_z, t_z) with at least one path between them. In total, we have the following seven paths in terms of edges: $P_1 = \{(s_p, a), (a, b), (b, t_p)\}$, $P_2 = \{(s_{p'}, a), (a, b), (b, t_{p'})\}$, $P_3 = \{(s_{p'} f), (f, g), (g, d), (d, e), (e, t_{p'})\}$, $P_4 = \{(s_{q'}, c), (c, d), (d, e), (e, t_{q'})\}$, $P_5 = \{(s_q, c), (c, d), (d, e), (e, t_q)\}$, $P_6 = \{(s_q, f), (f, g), (g, d), (d, e), (e, t_q)\}$, $P_7 = \{(s_q, f), (f, g), (g, t_q)\}$ and we see that the strategy sets for each commodity are: $\mathcal{P}^p = \{P_1\}$, $\mathcal{P}^q = \{P_5, P_6, P_7\}$, $\mathcal{P}^{p'} = \{P_2, P_3\}$ and $\mathcal{P}^{q'} = \{P_4\}$. A graphical interpretation of this network is given in Figure 1. Note that commodity $q \in Q_1$ has an additional strategy choice compared to the example of Theorem 3.4, therefore the total states are increased to six. We now model this instance as a bimatrix game, given by Table 2, where player 1 is the row player and 2 the column player. Observe, that the additional strategy P_7 for commodity p is strictly dominated by P_2 and P_4 , which implies the corollary. \square

4 TIGHT POA BOUNDS FOR GENERAL COST-SHARING METHODS

In this section, we focus on the price of anarchy of congestion games with multi-commodity players and prove general tight price of anarchy bounds, which are robust. More specifically, we prove an upper bound that holds for the multi-commodity case and afterwards complement this with a lower bound for the single-commodity case. This implies that the price of anarchy is the same both in the single- and multi-commodity model.

Upper bound

We first generalize the (λ, μ) -smoothness framework of [31] to accommodate any cost-sharing method and set of possible cost functions. For any set S , define f^S to be the vector whose components represent the flows of the players in S , and define as F^S the sum of components of f^S , that is $F^S = \sum_{i \in S} (f^S)_i$, $\forall S$. Suppose we identify positive parameters λ and $\mu < 1$ such that for every cost function in our allowable set $C \in \mathcal{C}$, and every pair of sets of players T and T^* , we get

$$\sum_{i \in T^*} \xi(i, f^{T \cup \{i\}}, C) \leq \lambda \cdot C(F^{T^*}) + \mu \cdot C(F^T). \quad (14)$$

Then, for P a PNE and P^* the optimal solution, we would get

$$SC(P) \stackrel{(1)}{=} \sum_{i \in N} \sum_{e \in E} \xi(i, f_e^N(P), C_e) \stackrel{(2)}{\leq} \sum_{i \in N} \sum_{e \in E} \xi(i, f_e^N(P_i^*, P_{-i}), C_e). \quad (15)$$

To simplify, let T_e and T_e^* be the sets of players who assign positive flow on resource e on an equilibrium outcome P and the optimal P^* , accordingly. Thus, $T_e = S_e(P)$ and $T_e^* = S_e(P^*)$. Then

$$\begin{aligned}
 (15) &= \sum_{e \in E} \sum_{i \in N} \xi(i, f_e^{T_e \cup \{i\}}(P_i^*, P_{-i}), C_e) \\
 &\stackrel{(14)}{\leq} \sum_{e \in E} \lambda \cdot C_e(F^{T_e^*}) + \mu \cdot C_e(F^{T_e}) \\
 &= \sum_{e \in E} \lambda \cdot C_e \left(\sum_{i \in T_e^*} (f^{T_e^*})_i \right) + \mu \cdot C_e \left(\sum_{i \in T_e} (f^{T_e})_i \right) \\
 &= \sum_{e \in E} \lambda \cdot C_e(f_e(P^*)) + \mu \cdot C_e(f_e(P)) \\
 &\stackrel{(1)}{=} \lambda \cdot SC(P^*) + \mu \cdot SC(P). \tag{16}
 \end{aligned}$$

Rearranging (16) yields a $\lambda/(1 - \mu)$ upper bound on the PoA. The same bound can be easily shown to apply to MNE and more general concepts (correlated and coarse correlated equilibria), though we omit the details (see, e.g., [31] for more). We then get the following lemma.

LEMMA 4.1. *Consider the following optimization program with variables λ, μ .*

$$\text{Minimize } \frac{\lambda}{1 - \mu} \tag{17}$$

$$\text{Subject To } \mu \leq 1 \tag{18}$$

$$\sum_{i \in T^*} \xi(i, f^{T \cup \{i\}}, C) \leq \lambda \cdot C(F^{T^*}) + \mu \cdot C(F^T), \tag{19}$$

where constraint (19) needs to hold for any function $C \in \mathcal{C}$, pair of sets $T, T^* \subseteq N$ and any positive flow vector f^S . Every feasible solution yields a $\lambda/(1 - \mu)$ upper bound on the PoA of the cost-sharing method given by $\xi(i, f^S, C)$ and the set of cost functions \mathcal{C} .

The upper bound holds for any cost-sharing method and set of allowable cost functions.

Bound tightness

We now proceed to show that the optimal solution to Program (17)-(19) gives a tight upper bound when our assumptions described in Section 2 hold. Observe that, since the PoA bound given by Lemma 4.1 is given by the optimal point of the mathematical program, the PoA is either infinite or a constant that does not depend on the number of players.

THEOREM 4.2. *Let (λ^*, μ^*) be the optimal point of Program (17)-(19). In congestion games with multi-commodity players, the PoA of the cost-sharing method given by $\xi(i, f^S, C)$ and the set of cost functions \mathcal{C} is precisely $\lambda^*/(1 - \mu^*)$ and equal to the PoA of the single-commodity case.*

PROOF. The proof proceeds in two parts. We first argue about strong duality of the mathematical program that provides the upper bound and then present the lower bound construction.

Strong duality of the mathematical program. Define $\zeta(y, x, C)$ for $y, x > 0$ as

$$\zeta(y, x, C) = \max_{T^*: FT^*=y, T:FT=x} \sum_{i \in T^*} \xi(i, f^{T \cup \{i\}}, C). \tag{20}$$

With this definition, we can rewrite Program (17)-(19) as

$$\text{Minimize } \frac{\lambda}{1 - \mu} \quad (21)$$

$$\text{Subject To } \mu \leq 1 \quad (22)$$

$$\zeta(y, x, C) \leq \lambda \cdot C(y) + \mu \cdot C(x), \quad \forall C \in \mathcal{C} \text{ and } x, y \in \mathbb{R}_{>0}. \quad (23)$$

Observe that for every constraint, we can scale the weights of the players by a factor a , dilate the cost function by a factor $1/a$ and scale it by an arbitrary factor, and keep the constraint intact (by Assumption 2). This suggests we can assume that every constraint has $y = 1$ and $C(1) = 1$. Then we rewrite Program (21)-(23) as

$$\text{Minimize } \frac{\lambda}{1 - \mu} \quad (24)$$

$$\text{Subject To } \mu \leq 1 \quad (25)$$

$$\zeta(1, x, C) \leq \lambda + \mu \cdot C(x), \quad \forall C \in \mathcal{C} \text{ and } x \in \mathbb{R}_{>0}. \quad (26)$$

The Lagrangian dual of Program (24)-(26) is

$$\text{Maximize } \inf_{\lambda, \mu} \left\{ \frac{\lambda}{1 - \mu} + \sum_{C \in \mathcal{C}, x > 0} z_{Cx} \cdot (\zeta(1, x, C) - \lambda - \mu \cdot C(x)) + z_{\mu} \cdot (\mu - 1) \right\} \quad (27)$$

$$\text{Subject To } z_{Cx}, z_{\mu} \geq 0. \quad (28)$$

Our primal is a semi-infinite program with an objective that is continuous, differentiable, and convex in the feasible region, and with linear constraints. We get that strong duality holds (see also [35, 39] for a detailed treatment of strong duality in this setting). We first treat the case when the optimal value of the primal is finite and is given by point (λ^*, μ^*) . Before concluding our proof we will explain how to deal with the case when the primal is infinite or infeasible. The KKT conditions yield for the optimal $\lambda^*, \mu^*, z_{Cx}^*$:

$$\frac{1}{1 - \mu^*} = \sum_{C \in \mathcal{C}, x > 0} z_{Cx}^* \quad (29)$$

$$\frac{\lambda^*}{(1 - \mu^*)^2} = \sum_{C \in \mathcal{C}, x > 0} z_{Cx}^* \cdot C(x). \quad (30)$$

Calling $\eta_{Cx} = z_{Cx}^* / \sum_{C \in \mathcal{C}, x > 0} z_{Cx}^*$ and dividing (30) with (29) we get

$$\frac{\lambda^*}{1 - \mu^*} = \sum_{C \in \mathcal{C}, x > 0} \eta_{Cx} \cdot C(x). \quad (31)$$

By (31) and the fact that all constraints for which $z_{Cx}^* > 0$ are tight (by complementary slackness), we get

$$\sum_{C \in \mathcal{C}, x > 0} \eta_{Cx} \cdot \zeta(1, x, C) = \sum_{C \in \mathcal{C}, x > 0} \eta_{Cx} \cdot C(x). \quad (32)$$

We now proceed to our matching lower bound which holds even for the single-commodity per player model.

Lower bound construction. The construction starts off with a single-commodity player i , who has flow 1 and, in the PNE, uses a single resource e_i by herself. The cost function of resource e_i is an arbitrary function from C such that $C_{e_i}(1) \neq 0$ (it is easy to see that such a function exists, since C is closed under dilation, unless all functions are 0, which is a trivial case) scaled so that $C_{e_i}(1) = \sum_{(C,x) \in \mathcal{T}} \eta_{Cx} \cdot \zeta(1, x, C)$. The other option of player i is to use a set of resources, one for each $(C, x) \in \mathcal{T}$ with cost functions $\eta_{Cx} \cdot C(\cdot)$. The resource corresponding to each (C, x) is used in the PNE by a player set that is equivalent to the T that maximizes the expression in (20) for the corresponding C, x . We now prove that player i does not gain by deviating to her alternative strategy. The key point is that due to convexity of the cost shares (Assumption 3), the worst case T^* in definition (20) will always be a single player. Then we can see that the cost share of i on each (C, x) resource in her potential deviation will be $\eta_{Cx} \cdot \zeta(1, x, C)$. It then follows that she is indifferent between her two strategies. Note that the PNE cost of i is $\sum_{(C,x) \in \mathcal{T}} \eta_{Cx} \cdot \zeta(1, x, C)$, which by (31) and (32) is equal to $\lambda^*/(1 - \mu^*)$. Also note that if player i could use her alternative strategy by herself, her cost would be 1.

We now make the following observation which allows us to complete the lower bound construction: Focus on the players and resources of the previous paragraph. Suppose we scale the weight of player i , as well as the weights of the users of the resources in her alternative strategy by the same factor $a > 0$. Then, suppose we dilate the cost functions of all these resources (the one used by i in the PNE and the ones in her alternative strategy) by a factor $1/a$ so that the costs generated by the players go back to the values they had in the previous paragraph. Finally, suppose we scale the cost functions by an arbitrary factor $b > 0$. We observe that the fact that i has no incentive to deviate is preserved (by Assumption 2, page 7) and the ratio of PNE cost versus alternative cost for i remains the same, i.e., $\lambda^*/(1 - \mu^*)$. This suggests that for every player generated by our construction so far in the PNE, we can repeat these steps by looking at her weight and PNE cost and appropriately constructing her alternative strategy and the users therein. After repeating this construction for a large number of layers $M \rightarrow \infty$, we complete the instance by creating a single resource for each of the players in the final layer. The cost functions of these resources are arbitrary nonzero functions from C scaled and dilated so that each one of these players is indifferent between her PNE strategy and using the newly constructed resource.

Consider the outcome that has all players play their alternative strategies and not the ones they use in the PNE. Every player other than the ones in the final layer would have a cost $\lambda^*/(1 - \mu^*)$ smaller, as we argued above. We can now see that, by (32), the cost of every player in the PNE is the same as that of the players in the resources of her alternative strategy. This means the cost across levels of our construction is identical and the final layer is negligible, since $M \rightarrow \infty$. This proves that the cost of the PNE is $\lambda^*/(1 - \mu^*)$ times larger than the outcome that has all players play their alternative strategies, which gives the tight lower bound.

Note on case with primal infeasibility. Recall that during our analysis we assumed that the primal program (24)-(26) had a finite optimal solution. Now suppose the program is either infeasible or $\mu = 1$, which means the minimizer yields an infinite value. This implies that, if we set μ arbitrarily close to 1, then there exists some $C \in C$, such that, for any arbitrarily large λ , there exists $x > 0$ such that $\zeta(1, x, C) > \lambda + \mu \cdot C(x)$. We can rewrite this last expression as $\zeta(1, x, C)/C(x) > \mu + \lambda/C(x)$, which shows we have values C and x such that $\zeta(1, x, C)$ is arbitrarily close to $C(x)$ or larger (since μ is arbitrarily close to 1). We can then replace λ with λ' such that the constraint becomes tight. It is not hard to see that these facts give properties parallel to (31) and (32) by setting $\eta_{Cx} = 1$ for our C and x and every other such variable to 0. Then our lower bound construction goes through for this arbitrarily large $\lambda'/(1 - \mu)$, which shows we can construct a lower bound with as high PoA as desired. This completes the proof of Theorem 4.2. \square

5 POS FOR WEIGHTED SHAPLEY VALUES

In this section we study the price of stability for a class of weighted Shapley values, where the sampling parameter of each player i is defined by $\lambda_i(P) = (f_e^i(P))^\gamma$ for any γ and outcome P .

PoS upper bound for Shapley values ($\gamma = 0$)

We start with an upper bound on the PoS for the case that $\gamma = 0$, i.e., for the *Shapley value* (SV) cost-sharing method, which is tight for the interesting subclass of polynomial cost functions.

THEOREM 5.1. *In congestion games with multi-commodity players, the PoS under the Shapley value cost-sharing method is at most*

$$\max_{C \in \mathcal{C}, x > 0} \frac{C(x)}{\int_0^x \frac{C(x')}{x'} dx'},$$

where \mathcal{C} the set of allowable cost functions.

PROOF. We begin with the potential function of the game:

$$\Phi(P) = \sum_{e \in E} \sum_{i \in N} \xi(i, f_e^{\leq i, \pi}(P), C_e)$$

and we prove the following lemma which is the main tool for proving our upper bound on the PoS. Briefly, the lemma states the following. For any instance with N players and any strategy profile, we can construct a new instance with $N + 1$ players by splitting one player in half into two new players. Then this can only reduce the potential value of the game. More precisely, we do this by splitting in half the flow of each commodity controlled by a player i on a resource creating two new commodities, which we assign to the new two players, say i' and i'' .

LEMMA 5.2. *Consider an outcome P of the game and assume that on a resource e , we substitute the total flow of a player i with the flows of two other players i', i'' such that $f_e^{i'}(\hat{P}) = f_e^{i''}(\hat{P}) = \frac{f_e^i(P)}{2}$. Then we claim that*

$$\Phi_e(P) \geq \Phi'_e(\hat{P}),$$

where $\Phi'_e(\hat{P})$ is the potential value of resource e after the substitution.

PROOF. First, rename the flows such that the substituted one $f_e^i(P)$ to have the highest index. Assign indices i' and i'' to the new ones, with $i < i' < i''$ in ordering π . Then, for any resource e , the new potential value equals to

$$\Phi'_e(\hat{P}) = \sum_{j=1}^{i-1} \xi(j, f_e^{\leq j, \pi}(P), C_e) + \xi(i', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P})), C_e) + \xi(i'', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P}), f_e^{i''}(\hat{P})), C_e).$$

Note that the contribution to the potential value of the flows before player's i flow is the same as before the substitution. Therefore it is enough to show that

$$\xi(i, f_e^N(P), C_e) \geq \xi(i', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P})), C_e) + \xi(i'', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P}), f_e^{i''}(\hat{P})), C_e). \quad (33)$$

To simplify, in what follows call $\chi = \xi(i, f_e^N(P), C_e)$, $\chi' = \xi(i', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P})), C_e)$ and $\chi'' = \xi(i'', (f_e^{\leq i, \pi}(P), f_e^{i'}(\hat{P}), f_e^{i''}(\hat{P})), C_e)$. Define as $S_e^i(\pi)$ the set of players preceding player i in π . Then, for every ordering π and permutation τ^i of set $S_e^i(\pi) \cup \{i\}$, define as $F_e^{\leq i, \pi, \tau^i}(P)$ the sum of players' flows who precede i in both π and τ^i . Let now $|S_e(P)| = r$. By definition of SV, we get

$$\chi = \frac{1}{r!} \sum_{\tau^i} \left(C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^i(P) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) \right) \right), \quad (34)$$

$$\chi' = \frac{1}{r!} \sum_{\tau^i} \left(C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i'}(\hat{P}) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) \right) \right). \quad (35)$$

For χ'' , since the position of $f_e^{i'}(\hat{P})$ in the ordering is unspecified, we give an upper bound for this value as follows. For any permutation τ , let $A(\tau)$ be the marginal cost increase caused by $f_e^{i''}(\hat{P})$ when she precedes $f_e^{i'}(\hat{P})$ in π , and $B(\tau)$ when she succeeds. That is

$$\begin{aligned} A(\tau) &= C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i''}(\hat{P}) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) \right), \\ B(\tau) &= C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i'}(\hat{P}) + f_e^{i''}(\hat{P}) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i'}(\hat{P}) \right). \end{aligned} \quad (36)$$

Let now p equal the probability of $f_e^{i'}(\hat{P})$ preceding $f_e^{i''}(\hat{P})$. Then SV definition gives

$$\chi'' = (1-p) \cdot \frac{1}{r!} \cdot \sum_{\tau^i} A(\tau) + p \cdot \frac{1}{r!} \cdot \sum_{\tau^i} B(\tau). \quad (37)$$

Due to convexity, $A(\tau) \leq B(\tau)$. Therefore, by substituting $A(\tau)$ with $B(\tau)$ in definition (37), we get the following upper bound for χ'' ,

$$\chi'' \leq \frac{1}{r!} \sum_{\tau^i} B(\tau). \quad (38)$$

Towards proving inequality (33), we have

$$\begin{aligned} \chi' + \chi'' &\stackrel{(35),(38)}{\leq} \frac{1}{r!} \sum_{\tau^i} C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i''}(\hat{P}) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) \right) + \\ &\quad + C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i'}(\hat{P}) + f_e^{i''}(\hat{P}) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^{i'}(\hat{P}) \right). \end{aligned}$$

Since $f_e^{i'}(\hat{P}) = f_e^{i''}(\hat{P}) = \frac{f_e^i(P)}{2}$, we get

$$\chi' + \chi'' \leq \frac{1}{r!} \sum_{\tau^i} \left(C_e \left(F_e^{<i,\pi,\tau^i}(P) + f_e^i(P) \right) - C_e \left(F_e^{<i,\pi,\tau^i}(P) \right) \right) \stackrel{(34)}{=} \chi,$$

as desired. This completes the proof of Lemma 5.2. \square

We now continue with the PoS upper bound. By repeatedly applying Lemma 5.2, we can break the total flow on each resource in identical flows of infinitesimal size without increasing the value of the potential. This implies that

$$\Phi_e(P) \geq \int_0^{f_e(P)} \frac{C_e(x)}{x} dx. \quad (39)$$

Now call P^* the optimal outcome and $P = \arg \min_{P'} \Phi(P')$ the minimizer of the potential function, which is, by definition, also a PNE. Then

$$\begin{aligned} SC(P^*) &\stackrel{??}{\geq} \Phi(P^*) \stackrel{\text{Def. } P}{\geq} \Phi(P) \stackrel{(39)}{\geq} \sum_{e \in E} \int_0^{f_e(P)} \frac{C_e(x)}{x} dx \\ &= \frac{\sum_{e \in E} \int_0^{f_e(P)} \frac{C_e(x)}{x} dx}{\sum_{e \in E} C_e(f_e(P))} \cdot SC(P) \geq \min_{e \in E} \frac{\int_0^{f_e(P)} \frac{C_e(x)}{x} dx}{C_e(f_e(P))} \cdot SC(P). \end{aligned} \quad (40)$$

Rearranging (40) yields the upper bound $PoS \leq \max_{C \in \mathcal{C}, x > 0} \frac{C(x)}{\int_0^x \frac{C(x')}{x'} dx'}$, which completes the proof of Theorem 5.1. \square

COROLLARY 5.3. *For polynomials with non-negative coefficients and degree at most d , the PoS of the SV is at most $d + 1$.*

REMARK 5.4. *From Corollary 5.3 and the corresponding lower bound from [7] it follows that the PoS of the SV is $\Theta(d)$.*

PoS for weighted Shapley values ($\gamma \neq 0$)

In the remainder of this section, we show that this linear dependence on the maximum degree d of the polynomial cost functions is very fragile. More precisely, for all values $\gamma \neq 0$, we show an exponential (in d) lower bound on the PoA, which asymptotically matches the corresponding upper bound by Gkatzelis et al. [15]. For $\gamma > 0$, our bound exactly matches the upper bound in [15], which holds for the weighted Shapley value in general. Our lower bound constructions hold for the single-commodity model and modify the corresponding instances from [15], making sure that they have a unique Nash equilibrium.

THEOREM 5.5. *In congestion games with single-commodity players and polynomial cost functions with non-negative coefficients and maximum degree d , the PoS for the class of weighted Shapley values with sampling parameters $\lambda_e^i(P) = (f_e^i(P))^\gamma$ is at least*

- (a) $(2^{\frac{1}{d}} - 1)^{-d}$, for all $\gamma > 0$, and
- (b) d^d , for all $\gamma < 0$.

In the following we prove Theorem 5.5. To do so, we first show a technical lemma, which will be crucial for proving our lower bounds on the PoS. We then introduce an instance in Example 5.7 and show in Theorem 5.8 that it gives the lower bound for $\gamma > 0$. Afterwards, Example 5.9 and Theorem 5.10 provide the corresponding lower bound for $\gamma < 0$.

LEMMA 5.6. *Consider a resource e , a player i with flow f_e^i and a set T of k players with total weight F_e^T , where $f_e^t = \frac{F_e^T}{k}$ for $t \in T$. Assume that the set of players $\{i\} \cup T'$, for some $T' \subseteq T$, is using a resource where players' cost shares are computed by weighted Shapley values with sampling weights $\lambda_z(f_e^z) = (f_e^z)^\gamma$, for each player z . Let $j \in T'$. Then (a) for $\gamma > 0$*

$$\lim_{k \rightarrow \infty} \chi_{ie} = C_e \left(|T'| \cdot \frac{F_e^T}{k} + f_e^i \right) - C_e \left(|T'| \cdot f_e^j \right), \quad (41)$$

$$\lim_{k \rightarrow \infty} \chi_{je} = \frac{1}{|T'|} \cdot C_e \left(|T'| \cdot \frac{F_e^T}{k} \right), \quad \text{if } |T'| \neq \emptyset, \quad (42)$$

and, (b) for $\gamma < 0$

$$\lim_{k \rightarrow \infty} \chi_{ie} = C_e(f_e^i), \quad (43)$$

$$\lim_{k \rightarrow \infty} \chi_{je} = \frac{1}{|T'|} \cdot \left(C_e \left(|T'| \cdot \frac{F_e^T}{k} + f_e^i \right) - C_e(f_e^i) \right), \quad \text{if } |T'| \neq \emptyset. \quad (44)$$

PROOF. First, notice that if $|T'| = \emptyset$, equations (41) and (43) of the lemma follow immediately, since only player i uses the resource. We now prove the lemma for the case where $|T'| \neq \emptyset$.

(a) Assume that the sampling weights are given by $\lambda_z(f^z) = (f^z)^\gamma$ for $\gamma > 0$ and let $k \rightarrow \infty$. We show that player i is the last player who enters the resource with probability $1 - o(1)$, i.e., she is among the first $\delta \cdot (|T'| + 1)$ players being drawn by the sampling procedure, for any arbitrary small $\delta > 0$. Consider now the probability p that player i is not among the first $\delta \cdot (|T'| + 1)$ drawn players. Then p is upper bounded by the probability that player i is not drawn among everyone in set $\{i\} \cup T'$ for $\delta \cdot (|T'| + 1)$ times,

$$p \leq \left(1 - \frac{f_e^i}{f_e^i + |T'| \cdot \left(\frac{F_e^T}{k} \right)^\gamma} \right)^{\delta \cdot (|T'| + 1)}. \quad (45)$$

Let $\beta = \frac{|T'|}{|T|} = \frac{|T'|}{k}$ where $\beta \in (0, 1]$. By substituting in (45), we have

$$\left(1 - \frac{f_e^i}{f_e^i + \beta \cdot k \cdot \left(\frac{F_e^T}{k} \right)^\gamma} \right)^{\delta \cdot (\beta \cdot k + 1)} = \left(1 - \frac{f_e^i}{f_e^i + \beta \cdot k^{1-\gamma} \cdot (F_e^T)^\gamma} \right)^{\delta \cdot (\beta \cdot k + 1)}. \quad (46)$$

For $\gamma \geq 1$, we have

$$(46) \leq \left(1 - \frac{f_e^i}{f_e^i + (F_e^T)^\gamma} \right)^{\delta \cdot (\beta \cdot k + 1)}$$

which goes to 0 as $k \rightarrow \infty$. Since $(1 - x) \leq e^{-x}$, we have for all $0 < \gamma < 1$,

$$\begin{aligned} (46) &\leq \exp \left(-\delta \cdot \beta \cdot k \cdot \frac{f_e^i}{f_e^i + \beta \cdot k^{1-\gamma} \cdot (F_e^T)^\gamma} \right) \\ &= \exp \left(-\delta \cdot \beta \cdot \frac{f_e^i}{\frac{f_e^i}{k} + \beta \cdot \left(\frac{F_e^T}{k} \right)^\gamma} \right), \end{aligned}$$

which also goes to 0 as $k \rightarrow \infty$. Therefore probability p is upper bounded by an arbitrarily small ϵ .

According to the weighted Shapley value method, the cost share of a player equals the expected marginal contribution she causes to the resource cost. Thus her cost share is affected by any player who enters resource before her. For any small $\delta > 0$, player i is with probability $1 - o(1)$ among the last $\delta \cdot (|T'| + 1)$ players who enter the resource. Therefore her cost share is affected by the players who enter the resource earlier. Thus player i incurs a cost of

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= C_e(|T'| \cdot f_e^j + f_e^i) - C_e(|T'| \cdot f_e^j) \\ &= C_e \left(|T'| \cdot \frac{F_e^T}{k} + f_e^i \right) - C_e \left(|T'| \cdot \frac{F_e^T}{k} \right). \end{aligned}$$

In contrast, the cost of each player $j \in T'$ is not affected by the flow of player i . Therefore for each $j \in T'$, we have

$$\lim_{k \rightarrow \infty} \chi_{je} = \frac{1}{|T'|} \cdot C_e(|T'| \cdot f_e^j) = \frac{1}{|T'|} \cdot C_e\left(|T'| \cdot \frac{F_e^T}{k}\right)$$

which completes the first part of the proof.

(b) Assume now that the sampling weights are given by $\lambda(f^z) = (f^z)^\gamma$ for $\gamma < 0$ and let $k \rightarrow \infty$. We prove that, for any $\delta > 0$, player i is, with probability $1 - o(1)$, among the first $\delta \cdot (|T'| + 1)$ players entering the resource. This probability equals the probability that player i is not drawn for the first $(1 - \delta) \cdot (|T'| + 1)$ sampling rounds. The following formula gives the probability that a player i is not drawn for the first q sampling rounds:

$$\prod_{r=1}^q \left(1 - \frac{f_e^i}{f_e^i + (|T'| - (r - 1)) \cdot \left(\frac{F_e^T}{k}\right)^\gamma} \right). \quad (47)$$

The probability of a player i being drawn increases with the number of sampling rounds. This implies that the probability of player i not being drawn becomes the smallest in the last sampling round. That is, when $r = q = (1 - \delta) \cdot (|T'| + 1)$, we get the smallest term of (47). Thus we can lower bound (47) by

$$\left(1 - \frac{f_e^i}{f_e^i + (\delta \cdot |T'| + 1) \cdot \left(\frac{F_e^T}{k}\right)^\gamma} \right)^{(1-\delta) \cdot (|T'|+1)}. \quad (48)$$

Define $\beta = \frac{|T'|}{|T'|} = \frac{|T'|}{k}$ where $\beta \in (0, 1]$. By substituting in (48), we have

$$\left(1 - \frac{f_e^i}{f_e^i + (\delta \cdot \beta \cdot k + 1) \cdot \left(\frac{F_e^T}{k}\right)^\gamma} \right)^{(1-\delta) \cdot (\beta \cdot k + 1)}$$

which is lower bounded by

$$\left(1 - \frac{f_e^i}{f_e^i + \delta \cdot \beta \cdot k^{1-\gamma} \cdot (F_e^T)^\gamma} \right)^{2 \cdot \beta \cdot k} \quad (49)$$

since $\beta \geq \frac{1}{k}$ (due to $|T'| \neq \emptyset$ in this case). By letting $k \rightarrow \infty$, we get that

$$\lim_{k \rightarrow \infty} (49) \geq \lim_{k \rightarrow \infty} \left(1 - \frac{f_e^i}{\delta \cdot \beta \cdot k^{1-\gamma} \cdot (F_e^T)^\gamma} \right)^{2 \cdot \beta \cdot k^{1-\gamma} \cdot k^\gamma}. \quad (50)$$

To simplify, let $x = \frac{f_e^i}{\delta \cdot (F_e^T)^\gamma}$. Then (50) equals to

$$\lim_{k \rightarrow \infty} \left(\left(1 - \frac{x}{\beta \cdot k^{1-\gamma}} \right)^{\beta \cdot k^{1-\gamma} - 1} \cdot \left(1 - \frac{x}{\beta \cdot k^{1-\gamma}} \right) \right)^{2 \cdot k^\gamma}. \quad (51)$$

Since $(1 - \frac{x}{n})^{n-1} \geq e^{-x}$ for any positive x and n , we lower-bound (51) by

$$\lim_{k \rightarrow \infty} \left(e^{-x} \cdot \left(1 - \frac{x}{\beta \cdot k^{1-\gamma}} \right) \right)^{2 \cdot k^\gamma}$$

which equals to 1 for any $\gamma < 0$.

As we mention in part (a) of the proof, the cost share of a player under the weighted Shapley value method equals to the marginal contribution she causes to the resource cost. As we proved, player i is the first player who enters the resource with probability $1 - o(1)$. Therefore her cost share is not affected by any player who enters after her. Even if some of the players in T' are introduced before player i in the resource, they will negligible affect player i 's cost. This is due to the fact that $f_e^j = \frac{F_e^T}{k} \rightarrow 0$ for any player $j \in T'$, since $k \rightarrow \infty$. Therefore the cost share of player i is given by

$$\chi_{ie} = C_e(f_e^i)$$

while the cost share of any player $j \in T'$ is given by

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{je} &= \frac{1}{|T'|} \cdot \left(C_e \left(f_e^i + |T'| \cdot f_e^j \right) - C_e \left(f_e^i \right) \right) \\ &= \frac{1}{|T'|} \cdot \left(C_e \left(f_e^i + |T'| \cdot \frac{F_e^T}{k} \right) - C_e \left(f_e^i \right) \right), \end{aligned}$$

since they enter resource after player i , which completes the proof of the Lemma 5.6. \square

Lower bound for positive γ

In this section, we give the proof for part (a) of Theorem 5.5. First, we describe the example that gives our lower bound, for $\gamma > 0$, which is stated in Theorem 5.8 afterwards.

Example 5.7. Consider a complete k -ary tree $G = (E, N)$ with l levels, where the root is positioned at level 0. Each vertex of the tree corresponds to a resource and each edge to a player. For $0 \leq j \leq l-1$, let $E_j \subset E$ be the set of resources (nodes) at level j and, for $0 \leq j \leq l-2$, let $N_j \subset N$ be the set of players (edges) between levels j and $j+1$ of the tree. The strategies of every player are the endpoints of the edge she is associated with, i.e. a player associated with edge $(j, j-1)$ has strategies $\{j\}, \{j-1\}$. For $0 \leq j \leq l-2$, the flow f^i of a player $i \in N_j$ is given by ϕ^{l-j-2} , where

$$\phi = k \cdot \left(2^{\frac{1}{d}} - 1 \right). \quad (52)$$

Define $\alpha = \frac{k}{\phi^d}$ and let $0 < \epsilon < 1$. The cost function of resources $e \in E_j$, is given by

$$C_e(x) = \left(1 + \epsilon^{l-j} \right) \cdot \alpha^{l-j-2} \cdot x^d, \quad \text{for } 0 \leq j \leq l-2,$$

$$\text{and } C_e(x) = (1 + \epsilon) \cdot k^{d-1} \cdot x^d, \quad \text{for } j = l-1.$$

THEOREM 5.8. *In congestion games with single-commodity players and polynomial cost functions with non-negative coefficients and maximum degree d , the PoS for the class of weighted Shapley values with sampling parameters $\lambda_i(f^i) = (f^i)^\gamma$, with $\gamma > 0$, is at least*

$$\left(2^{\frac{1}{d}} - 1 \right)^{-d}.$$

PROOF. Consider the instance in Example 5.7. Then, let P be the outcome where all players use the resource closer to the root, and P^* be the outcome where all players use the resource further from the root. We prove by induction that P is a unique Nash equilibrium with a total cost equal to $(2^{\frac{1}{d}} - 1)^{-d}$ times the total cost of outcome P^* .

Consider a player $i \in N_{l-2}$, i.e. a player connected to a leaf resource and assume she uses the leaf resource, $e \in E_{l-1}$. Note that she is the only player who can use this resource. Then player i 's cost share equals to

$$\chi_{ie} = (1 + \epsilon) \cdot k^{d-1}. \quad (53)$$

For $k' \leq k$, consider k' players from set N_{l-2} (including player i) and assume they use the resource e closer to the root, $e \in E_{l-2}$. Choose the (worst) case where a player $b \in N_{l-3}$ also uses resource e . We show that even in this case, player i prefers the resource closer to the root. Using (42) of Lemma 5.6, player i 's cost share is given by

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= \frac{1}{k'} \cdot (1 + \epsilon^2) \cdot \alpha \cdot (k' \cdot \phi)^d \\ &\stackrel{k' \leq k}{\leq} (1 + \epsilon^2) \cdot k^{d-1}, \end{aligned} \quad (54)$$

where (54) is strictly smaller than (53), since $0 < \epsilon < 1$.

For $l-1 \geq j' \geq j+1$, assume that each player of set $N_{j'}$ uses the strategy closer to the root, $e \in E_{j'}$. Then consider a player $i \in N_j$ who uses the resource e further from the root, $e \in E_{j+1}$. By assumption, this resource is also used by k players from set N_{j+1} . Then, using (41) of Lemma 5.6, player i 's cost share is given by

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= (1 + \epsilon^{l-j-1}) \cdot \alpha^{l-j-3} \cdot \left((k \cdot \phi^{l-j-3} + \phi^{l-j-2})^d - (k \cdot \phi^{l-j-3})^d \right) \\ &= (1 + \epsilon^{l-j-1}) \cdot \alpha^{l-j-3} \cdot k^d \cdot \phi^{d \cdot (l-j-3)} \cdot \left(\left(1 + \frac{\phi}{k} \right)^d - 1^d \right). \\ &= (1 + \epsilon^{l-j-1}) \cdot (\alpha \cdot \phi^d)^{l-j-3} \cdot k^d. \end{aligned}$$

By substituting α , we have

$$\lim_{k \rightarrow \infty} \chi_{ie} = (1 + \epsilon^{l-j-1}) \cdot (k^{l-j-3+d}). \quad (55)$$

For $k' \leq k$, consider k' players from set N_j (including player i) and assume they use the resource e closer to the root, $e \in E_j$. Choose the (worst) case where one player $b \in N_{j-1}$ also uses resource e . Then we show that player i still prefers to use resource $e \in E_j$. By (42) of Lemma 5.6, i 's cost share equals to

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= \frac{1}{k'} \cdot (1 + \epsilon^{l-j}) \cdot \alpha^{l-j-2} \cdot (k' \cdot \phi^{l-j-2})^d \\ &= (1 + \epsilon^{l-j}) \cdot \alpha^{l-j-2} \cdot (k')^{d-1} \cdot \phi^{d \cdot (l-j-2)} \\ &\stackrel{k' \leq k}{\leq} (1 + \epsilon^{l-j}) \cdot (\alpha \cdot \phi^d)^{l-j-2} \cdot k^{d-1}. \end{aligned} \quad (56)$$

By substituting α , we have

$$\lim_{k \rightarrow \infty} \chi_{ie} \leq (1 + \epsilon^{l-j}) \cdot k^{l-j-3+d},$$

which is strictly smaller than (55), since $0 < \epsilon < 1$.

PoS. Let P and P^* be the outcomes where each player chooses the strategy closer and further from the root respectively. In this section, we compute the total costs of outcomes P and P^* , and present a lower bound to the PoS.

For outcome P , since every player chooses the resource closer to the root, no player uses any resource $e \in E_{l-1}$, therefore the cost at the leaves of the tree is zero. For $0 \leq j \leq l-2$, we proved that each player using a resource $e \in E_j$ incurs a cost of $(1 + \epsilon^{l-j}) \cdot k^{l-j-3+d}$. Since each of the k^j resources in E_j is used by k players, the total cost at level j equals to

$$\begin{aligned} \sum_{e \in E_j} C_e(f_e(P)) &= k^{j+1} \cdot (1 + \epsilon^{l-j}) \cdot k^{l-j-3+d} \\ &= (1 + \epsilon^{l-j}) \cdot (k^{l-2+d}). \end{aligned}$$

Summing up for the l levels, we get that the social cost in outcome P equals to

$$SC(P) = \sum_{j=0}^{l-1} \sum_{e \in E_j} C_e(f_e(P)) = k^{l-2+d} \cdot \left(l-1 + \sum_{i=2}^l \epsilon^i \right). \quad (57)$$

For outcome P^* , we now compute the cost of each player on each level. Since every player chooses to use the resource further from the root, the cost of resource $e \in E_0$ is zero. For $1 \leq j \leq l-2$, each player using a resource $e \in E_j$ incurs a cost of

$$(1 + \epsilon^{l-j}) \cdot \alpha^{l-j-2} \cdot \phi^{(l-j-1) \cdot d} = (1 + \epsilon^{l-j}) \cdot (\alpha \cdot \phi^d)^{l-j-2} \cdot \phi^d = (1 + \epsilon^{l-j}) \cdot k^{l-j-2} \cdot \phi^d.$$

In this case, each of the k^j resources in E_j is used by only one player, therefore the total cost at level j equals to

$$\sum_{e \in E_j} C_e(f_e(P)) = k^j \cdot (1 + \epsilon^{l-j}) \cdot (k^{l-j-2}) = (1 + \epsilon^{l-j}) \cdot k^{l-2} \cdot \phi^d. \quad (58)$$

Last, the cost of each player using a leaf resource, $e \in E_{l-1}$, equals to $(1 + \epsilon) \cdot k^{d-1}$. Since there are k^{l-1} leaf resources, the total cost at level $l-1$ equals to

$$\sum_{e \in E_{l-1}} C_e(f_e(P)) = k^{l-1} \cdot (1 + \epsilon) \cdot k^{d-1} = (1 + \epsilon) \cdot (k^{l-2+d}). \quad (59)$$

Summing up for the l levels, we have that the social cost in outcome P^* equals to

$$\begin{aligned} SC(P^*) &= \sum_{j=0}^{l-1} \sum_{e \in E_j} C_e(f_e(P)) = \sum_{j=1}^{l-2} \sum_{e \in E_j} C_e(f_e(P)) + \sum_{e \in E_{l-1}} C_e(f_e(P)) \\ &\stackrel{(58),(59)}{=} k^{l-2} \cdot \phi^d \cdot \left(l-2 + \sum_{i=2}^l \epsilon^i \right) + (1 + \epsilon) \cdot k^{l-2+d} \\ &\stackrel{(52)}{=} k^{l-2+d} \cdot \left(\left(2^{\frac{1}{d}} - 1 \right)^d \cdot \left(l-2 + \sum_{i=2}^l \epsilon^i \right) + 1 + \epsilon \right). \quad (60) \end{aligned}$$

Using (57) and (60), the PoS is lower bounded by

$$\frac{SC(P)}{SC(P^*)} = \frac{l-1 + \sum_{i=2}^l \epsilon^i}{\left(2^{\frac{1}{d}} - 1 \right)^d \cdot \left(l-2 + \sum_{i=2}^{l-1} \epsilon^i \right) + 1 + \epsilon}$$

As $l \rightarrow \infty$ and ϵ is arbitrary small, this ratio goes to $\left(2^{\frac{1}{d}} - 1 \right)^{-d}$, as desired. \square

Lower bound for negative γ .

In this section, we give the proof for part (b) of Theorem 5.5. First, we describe the example that gives our lower bound, for $\gamma < 0$, which is stated in Theorem 5.10 afterwards.

Example 5.9. Consider a complete k -ary tree $G = (E, N)$ with l levels, where the root is positioned at level 0. Each vertex of the tree corresponds to a resource and each edge to a single-commodity player. For $0 \leq j \leq l-1$, let $E_j \subset E$ be the set of resources (nodes) at level j , and, for $0 \leq j \leq l-2$, let $N_j \subset N$ be the set of players (edges) between levels j and $j+1$ of the tree. The strategies of every player are the endpoints of the edge she is associated with, i.e. a player associated with edge $(j, j-1)$ has strategies $\{\{j\}, \{j-1\}\}$. The flow of a player $i \in N_j$ is given by ϕ^j , where

$$\phi = \frac{1}{k \cdot d}. \quad (61)$$

Define $\alpha = k^{d-1} \cdot d^d$. The cost function of resources $e \in E_j$, is given by

$$C_e(x) = (1 + \epsilon) \cdot x^d, \quad \text{for } j = 0,$$

$$\text{and } C_e(x) = (1 + \epsilon^{j+1}) \cdot \alpha^{j-1} \cdot x^d, \quad \text{for } 1 \leq j \leq l-1.$$

where ϵ is an arbitrarily small parameter.

THEOREM 5.10. *In congestion games with single-commodity players and polynomial cost functions with non-negative coefficients and maximum degree d , the PoS for the class of weighted Shapley values with sampling parameters $\lambda_i(f^i) = (f^i)^\gamma$, with $\gamma < 0$, is at least d^d .*

PROOF. Choose the instance described in Example 5.9. Then, let P be the outcome where all players use the resource further from the root, and P^* be the outcome where all players use the resource closer to the root. We prove by induction that P is a unique Nash equilibrium with a total cost equal to d^d times the total cost of outcome P^* .

For $k' \leq k$, consider k' players from set N_0 including a player i and assume they use the resource on the root, $e \in E_0$. Then player i 's cost share equals to

$$\lim_{k \rightarrow \infty} \chi_{ie} = \frac{1}{k'} \cdot C_e(f_e(P)) = \frac{1}{k'} \cdot (1 + \epsilon) \cdot (k')^d \stackrel{k' \geq 1}{\geq} 1 + \epsilon. \quad (62)$$

Assume now that player i deviates to resource e further from the root, $e \in E_1$, and she uses this resource together with k' players from set N_1 , for $k' \leq k$. By (43) of Lemma 5.6, player i 's cost share is given by

$$\lim_{k \rightarrow \infty} \chi_{ie} = (1 + \epsilon^2) \cdot 1^d = 1 + \epsilon^2$$

for large enough k , which is strictly lower than (62).

For $0 \leq j' \leq j-2$, assume that each player of set $N_{j'}$ uses the strategy further from the root, $e \in E_{j'+1}$. For $k' \leq k$, consider k' players from set N_{j-1} including a player i , and assume they use the resource e closer to the root, $e \in E_{j-1}$. By assumption, this resource is also used by one player $b \in N_{j-2}$. We show that player i prefers to deviate. Using (44) of Lemma 5.6, player i incurs a cost of

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= \frac{1}{k'} \cdot (1 + \epsilon^j) \cdot \alpha^{j-2} \cdot \left((k' \cdot \phi^{j-1} + \phi^{j-2})^d - \phi^{(j-2) \cdot d} \right) \\ &\stackrel{k' \geq 1}{\geq} (1 + \epsilon^j) \cdot \alpha^{j-2} \cdot \left((\phi^{j-1} + \phi^{j-2})^d - \phi^{(j-2) \cdot d} \right) \\ &= (1 + \epsilon^j) \cdot \alpha^{j-2} \cdot \phi^{(j-2) \cdot d} \cdot \left((\phi + 1)^d - 1 \right). \end{aligned} \quad (63)$$

Define function $g(x) = x^d$. Since g is convex and $\phi \rightarrow 0$, we have that

$$(\phi + 1)^d - 1 \geq \phi \cdot d.$$

Using the above inequality, we get

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &\geq (63) \geq (1 + \epsilon^j) \cdot (\alpha \cdot \phi^d)^{j-2} \cdot \phi \cdot d \\ &= (1 + \epsilon^j) \cdot \left(\frac{1}{k}\right)^{j-2} \cdot \frac{1}{k \cdot d} \cdot d \\ &= (1 + \epsilon^j) \cdot \frac{1}{k^{j-1}}. \end{aligned} \quad (64)$$

Assume now that player i deviates to resource e further from the root, $e \in E_j$. Choose the case where k' players from set N_j also use this resource e , for a $k' \leq k$. Then by (43) of Lemma, player i 's cost share is given by

$$\begin{aligned} \lim_{k \rightarrow \infty} \chi_{ie} &= (1 + \epsilon^{j+1}) \cdot \alpha^{j-1} \cdot (\phi^{j-1})^d \\ &= (1 + \epsilon^{j+1}) \cdot \frac{1}{k^{j-1}}, \end{aligned}$$

which is strictly lower than (64).

PoS. First, we compute the total cost of the outcome P . Since every player chooses the resource further from the root, the cost of resource $e \in E_0$ (root) is zero. As we proved, each player using a resource $e \in E_j$ incurs a cost of $(1 + \epsilon^{j+1}) \cdot \frac{1}{k^{j-1}}$, for $1 \leq j \leq l-1$. Each of the k^j resources in E_j is used by only one player. Therefore the total cost at level j equals to

$$\sum_{e \in E_j} C_e(f_e(P)) = k^j \cdot (1 + \epsilon^{j+1}) \cdot \frac{1}{k^{j-1}} = (1 + \epsilon^{j+1}) \cdot k.$$

Computing the sum for the l levels, we have that the cost of outcome P equals

$$SC(P) = \sum_{j=0}^{l-1} \sum_{e \in E_j} C_e(f_e(P)) = k \cdot \left(l - 1 + \sum_{i=2}^l \epsilon^i \right). \quad (65)$$

Now let P^* be the outcome where each player chooses the resource closer to the root. The cost of level $l-1$ (leaves) of the tree is zero. The cost at level 0 is $(1 + \epsilon) \cdot k^d$. For $1 \leq j \leq l-2$, each resource $e \in E_j$ incurs a cost of

$$\begin{aligned} (1 + \epsilon^{j+1}) \cdot \alpha^{j-1} \cdot (k \cdot \phi^j)^d &= (1 + \epsilon^{j+1}) \cdot \alpha^{j-1} \cdot k^d \cdot (\phi^{j-1} \cdot \phi)^d \\ &= (1 + \epsilon^{j+1}) \cdot (\alpha \cdot w^d)^{j-1} \cdot (k \cdot \phi)^d \\ &= (1 + \epsilon^{j+1}) \cdot \frac{1}{k^{j-1} \cdot d^d}. \end{aligned}$$

Since there are k^j resources in E_j , the total cost of level j of the tree is given by

$$\sum_{e \in E_j} C_e(f_e(P)) = k^j \cdot (1 + \epsilon^{j+1}) \cdot \frac{1}{k^{j-1} \cdot d^d} = (1 + \epsilon^{j+1}) \cdot \frac{k}{d^d}.$$

Summing up for all levels, we get that the social cost in outcome P^* equals to

$$\begin{aligned} SC(P^*) &= \sum_{j=0}^{l-1} \sum_{e \in E_j} C_e(f_e(P)) = \sum_{e \in E_0} C_e(f_e(P)) + \sum_{j=1}^{l-1} \sum_{e \in E_j} C_e(f_e(P)) \\ &= (1 + \epsilon) \cdot k^d + \left(l - 1 + \sum_{i=2}^l \epsilon^i \right) \cdot \frac{k}{d^d}. \end{aligned} \quad (66)$$

Using (65) and (66), the PoS is lower bounded by the following ratio

$$\begin{aligned} \frac{SC(P)}{SC(P^*)} &= \frac{k \cdot \left(l - 1 + \sum_{i=2}^l \epsilon^i \right)}{(1 + \epsilon) \cdot k^d + \left(l - 1 + \sum_{i=2}^l \epsilon^i \right) \cdot \frac{k}{d^d}} \\ &= \frac{l - 1 + \sum_{i=2}^l \epsilon^i}{(1 + \epsilon) \cdot k^{d-1} + \left(l - 1 + \sum_{i=2}^l \epsilon^i \right) \cdot \frac{1}{d^d}}. \end{aligned}$$

With $l \rightarrow \infty$ and $\epsilon \rightarrow 0$, this ratio goes to d^d for any k , which completes the proof of Theorem 5.10. \square

6 CONCLUSION

In this work we conducted a complete study of cost-sharing methods in weighted congestion games. We characterized the price of anarchy and price of stability for general classes of cost functions and cost methods and extended our results to generalizations of congestion games where each player can control multiple commodities. The main conceptual messages coming out of our work are that:

- these generalized congestion games are tight and robust with respect to the price of anarchy in a very general framework of cost functions and cost-sharing methods.
- the Shapley value maintains strong properties in weighted and generalized congestion games as it guarantees the existence of pure Nash equilibria and (almost) minimizes the price of stability.
- weighted variants of the Shapley value do not fare as well, since they either fail to maintain pure Nash equilibrium existence (when moving to multiple commodities per player) or have a much higher price of stability than the unweighted Shapley value (both for standard and generalized weighted congestion games).

Our work provides a complete understanding in the case of generalized weighted congestion games and it is interesting to explore the landscape in the related framework of *splittable congestion games* [10, 18, 23, 28, 32]. In splittable congestion games with multi-commodity players, the weight of a commodity can be split among its strategies. A result from Orda et al. [28] implies the existence of pure Nash equilibria in the multi-commodity splittable model, if the cost share of a player on a resource is a convex function of her flow on the resource. This result [28] is based on the Kakutani Fixed Point theorem. This immediately gives us existence of pure Nash equilibria for the standard Shapley value. Furthermore, we have observed that simple examples can show this is not the case for variants of the Shapley value that impose predefined orderings on the players. In this setting, it is interesting to focus on open questions which include:

- How do the price of anarchy and price of stability behave for different cost-sharing methods and classes of cost functions?
- For which classes of cost-sharing methods is existence of a pure Nash equilibrium guaranteed?

REFERENCES

- [1] Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. 2011. Exact price of anarchy for polynomial congestion games. *SIAM J. Comput.* 40, 5 (2011), 1211–1233.
- [2] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. 2008. The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38, 4 (2008), 1602–1623.
- [3] Baruch Awerbuch, Yossi Azar, and Amir Epstein. 2013. The price of routing unsplittable flow. *SIAM J. Comput.* 42, 1 (2013), 160–177.
- [4] Kshipra Bhawalkar, Martin Gairing, and Tim Roughgarden. 2014. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation* 2, 4 (2014), 14:1–14:23.
- [5] Ho-Lin Chen and Tim Roughgarden. 2009. Network design with weighted players. *Theory of Computing Systems* 45, 2 (2009), 302–324.
- [6] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. 2010. Designing network protocols for good equilibria. *SIAM J. Comput.* 39, 5 (2010), 1799–1832.
- [7] George Christodoulou and Martin Gairing. 2016. Price of stability in polynomial congestion games. *ACM Transactions on Economics and Computation* 4, 2 (2016), 10:1–10:17.
- [8] George Christodoulou and Elias Koutsoupias. 2005. The price of anarchy of finite congestion games. In *Proceedings of the ACM Symposium on Theory of Computing*. 67–73.
- [9] George Christodoulou and Alkmini Sgouritsa. 2016. Designing networks with good equilibria under uncertainty. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 72–89.
- [10] Roberto Cominetti, José R. Correa, and Nicolás E. Stier Moses. 2009. The impact of oligopolistic competition in networks. *Operations Research* 57, 6 (2009), 1421–1437.
- [11] Dimitris Fotakis, Spyros Kontogiannis, and Paul Spirakis. 2008. Atomic congestion games among coalitions. *ACM Transactions on Algorithms* 4, 4 (2008), 52:1–52:27.
- [12] Martin Gairing, Konstantinos Kollias, and Grammateia Kotsialou. 2015. Tight bounds for cost-sharing in weighted congestion games. In *Proceedings of the International Colloquium in Automata, Languages, and Programming (ICALP)*. Springer, 626–637.
- [13] Martin Gairing, Konstantinos Kollias, and Grammateia Kotsialou. 2017. Cost-sharing in generalised selfish routing. In *Proceedings of the International Conference on Algorithms and Complexity (CIAC)*. 272–284.
- [14] Martin Gairing and Florian Schoppmann. 2007. Total latency in singleton congestion games. In *Proceedings of the International Workshop of Internet and Network Economics (WINE)*. 381–387.
- [15] Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden. 2016. Optimal cost-sharing in general resource selection games. *Operations Research* 64, 6 (2016), 1230–1238.
- [16] Michel Goemans, Vahab Mirrokni, and Adrian Vetta. 2005. Sink equilibria and convergence. In *Proceedings of the Annual IEEE Symposium on Foundations of Computer Science (FOCS)*. 142–151.
- [17] Ragavendran Gopalakrishnan, Jason R Marden, and Adam Wierman. 2014. Potential games are necessary to ensure pure Nash equilibria in cost sharing games. *Mathematics of Operations Research* 39, 4 (2014), 1252–1296.
- [18] Tobias Harks. 2011. Stackelberg strategies and collusion in network games with splittable flow. *Theory of Computing Systems* 48, 4 (2011), 781–802.
- [19] Tobias Harks and Max Klimm. 2012. On the existence of pure Nash equilibria in weighted congestion games. *Mathematics of Operations Research* 37, 3 (2012), 419–436.
- [20] Tobias Harks, Max Klimm, and Rolf H Möhring. 2011. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems* 49, 1 (2011), 46–70.
- [21] Tobias Harks and Konstantin Miller. 2011. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations Research* 59, 6 (2011), 1491–1503.

- [22] Sergiu Hart and Andreu Mas-Colell. 1989. Potential, value, and consistency. *Econometrica: Journal of the Econometric Society* 57, 3 (1989), 589–614.
- [23] Ara Hayrapetyan, Éva Tardos, and Tom Wexler. 2006. The effect of collusion in congestion games. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*. 89–98.
- [24] Max Klimm and Daniel Schmand. 2015. Sharing Non-Anonymous Costs of Multiple Resources Optimally. In *Proceedings of the International Conference on Algorithms and Complexity (CIAC)*. 274–287.
- [25] Konstantinos Kollias and Tim Roughgarden. 2015. Restoring pure equilibria to weighted congestion games. *ACM Transactions on Economics and Computation* 3, 4 (2015), 21:1–21:24.
- [26] Elias Koutsoupias and Christos Papadimitriou. 2009. Worst-case equilibria. *Computer Science Review* 3, 2 (2009), 65–69.
- [27] Dov Monderer and Lloyd S Shapley. 1996. Potential games. *Games and Economic Behavior* 14, 1 (1996), 124–143.
- [28] Ariel Orda, Raphael Rom, and Nahum Shimkin. 1993. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking* 1, 5 (1993), 510–521.
- [29] Robert W Rosenthal. 1973. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory* 2, 1 (1973), 65–67.
- [30] Robert W Rosenthal. 1973. The network equilibrium problem in integers. *Networks* 3, 1 (1973), 53–59.
- [31] Tim Roughgarden. 2009. Intrinsic robustness of the price of anarchy. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*. 513–522.
- [32] Tim Roughgarden and Florian Schoppmann. 2015. Local smoothness and the price of anarchy in splittable congestion games. *Journal in Economic Theory* 156 (2015), 317–342.
- [33] Tim Roughgarden and Okke Schrijvers. 2014. Network cost-sharing without anonymity. In *Proceedings of the Symposium on Algorithmic Game Theory (SAGT)*. 134–145.
- [34] Tim Roughgarden and Éva Tardos. 2002. How bad is selfish routing? *J. ACM* 49, 2 (2002), 236–259.
- [35] Alexander Shapiro. 2005. On duality theory of convex semi-infinite programming. *Optimization* 54, 6 (2005), 535–543.
- [36] Yosef Sheffi. 1985. *Urban transportation networks: equilibrium analysis with mathematical programming methods*. Prentice Hall.
- [37] Long Tran-Thanh, Maria Polukarov, Archie C. Chapman, Alex Rogers, and Nicholas R. Jennings. 2011. On the existence of pure strategy Nash equilibria in integer-splittable weighted congestion games. In *Proceedings of the Symposium on Algorithmic Game Theory (SAGT)*. 236–253.
- [38] Philipp von Falkenhausen and Tobias Harks. 2011. Optimal cost sharing protocols for scheduling games. In *Proceedings of the ACM Conference on Electronic Commerce (ACM-EC)*. 285–294.
- [39] Soon-Yi Wu and S-C Fang. 1999. Solving convex programs with infinitely many linear constraints by a relaxed cutting plane method. *Computers & Mathematics with Applications* 38, 3 (1999), 23–33.