Tight Bounds for Cost-Sharing in Weighted Congestion Games*

Martin Gairing[1], Konstantinos Kollias[2], and Grammateia Kotsialou[1]

[1] University of Liverpool, Liverpool, Merseyside L69 3BX, UK
[2] Stanford University, Stanford, CA 94305, USA

Abstract. This work studies the price of anarchy and the price of stability of cost-sharing methods in weighted congestion games. We require that our cost-sharing method and our set of cost functions satisfy certain natural conditions and we present general tight price of anarchy bounds, which are robust and apply to general equilibrium concepts. 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 close the paper with a somehow 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.

1 Introduction

The class of weighted congestion games [16, 17] encapsulates a large collection of important applications in the study of the 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. As an example, one of the most prominent such applications is *selfish routing* in a telecommunications or traffic network [3, 8, 20]. When more total demand is placed on a resource, the resource becomes scarcer, and the quality of service experienced by its users degrades. More specifically, in weighted congestion games there is a set of players N and a set of resources E. Each player $i \in N$ has a weight $w_i > 0$ and she gets to select the subset of the resources that she will use. The possible subsets she can pick 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 $f_e \cdot c_e(f_e)$, 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 a resource is covered by the set of players S_e using e, i.e., $\sum_{i \in S_e} \chi_{ie} = f_e \cdot c_e(f_e)$, where χ_{ie} is the *cost share* of player i on resource 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

* This work was supported by EPSRC grants EP/J019399/1 and EP/L011018/1.

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 weight, i.e., $\chi_{ie} = w_i \cdot c_e(f_e)$, 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, PS does not guarantee equilibrium existence (see [12] for a characterization), while the SV does. In [11], it is shown that only the class of *generalized weighted Shapley values* (see Section 2 for a definition) guarantees the existence of a PNE in such games.

As a metric that is worst-case by nature, the POA of a 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* [18] which means they apply to MNE and even more general classes (such as correlated and coarse-correlated equilibria). On the other hand, the motivations behind the study of the POS assume a PNE will exist, hence the POS is more meaningful when the method does guarantee a PNE.

1.1 Our Contributions

In this work we make two main contributions, one with respect to the POA and one with respect to the POS.

General POA bounds: On the POA side, we present *tight* bounds for general classes of allowable cost functions and for general cost-sharing methods, 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. To obtain our tight bounds we make use of the following natural assumptions, which we explain in more detail in Section 2:

1. Every cost function in the game is continuous, nondecreasing, and convex.
2. Cost-sharing is consistent when player sets generate costs in the same way.
3. The cost share of a player on a resource is a convex function of her weight.

We now briefly discuss these assumptions. Assumption 1 is standard in congestion-type settings. For example, linear cost functions have obvious applications in many network models, as do queueing delay functions, while higher degree polynomials (such as quartic) have been proposed as realistic models of road traffic [22]. Assumption 2 asks that the cost-sharing method only looks at how players generate costs and does not discriminate between them in any other way. Assumption 3 asks that the curvature of the cost shares is consistent, i.e., given Assumption 1, that the share of a player on a resource is a convex function of her weight (otherwise, we would get that the share of the player increases in a slower than convex way but the total cost of the constant weight players increases in a convex way, which we view as unfair). We note that our upper bounds are robust and apply to general equilibrium concepts that are guaranteed to exist for all cost-sharing methods.

SV based POS bounds: Studying the POS is most well-motivated in settings where a trusted mediator or some other authority can place the players in an initial configuration and they will not be willing to deviate from it. 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. For SV cost sharing, we prove an upper bound on the POS which holds for all sets of cost functions that satisfy Assumption 1. We show that for the interesting case of polynomials of bounded degree d, this upper bound is $d + 1$, which is asymptotically tight and always very close to the lower bound in [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 that the POS is very close and for $\gamma > 0$ even matches the upper bound on the *price of anarchy* in [10]. In other words, for this case the POS and the POA coincide, which we found very surprising, in particular because the upper bound in [10] even applies to general cost-sharing methods. We note that these weighted Shapley values are the only cost-sharing methods that guarantee existence of a PNE and satisfy Assumption 2 [10, 11].

1.2 Related Work and Comparison to Previous Results

POA. The POA was proposed in [15]. Most work on the inefficiency of equilibria for weighted congestion games has focused on PS. Tight bounds for the case of linear cost functions have been obtained in [3, 8]. The case of bounded degree polynomials was resolved in [1] and subsequent work [4, 9, 18] concluded the study of PS. In particular, [18] formalized the *smoothness* framework which shows how robust POA bounds (i.e., POA bounds that apply to general equilibrium concepts) are obtained.

Further cost-sharing methods have been considered in [10, 14]. Here, [14] provides tight bounds for the SV in games with convex cost functions, while [10] proves that the SV is optimal among the methods that guarantee the existence of a PNE and that PS is near-optimal in general, for games with polynomial cost functions. The authors also show tight bounds on the *marginal contribution method* (which charges a player the increase her presence causes to the joint cost) in games with polynomial costs. Optimality of the SV in closely related settings has also been discussed in [13, 19].

POS. The term price of stability was introduced in [2] for the network cost-sharing game, which was further studied in [5, 6, 14] for weighted players and various cost-sharing methods. With respect to congestion games, results on the POS are only known for polynomial *unweighted* games, for which [7] provides exact bounds. Work in [13, 19] studies the POS of the Shapley value in related settings.

Comparison to previous work. 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 in [10]. For example, the SV has the drawback that it can't 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 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 in [14], which provides tight bounds for general cost-sharing methods.

Our POS upper bound is the first for weighted congestion games that applies to any class of convex costs. The work in [19] presents SV POS bounds in a more general setting with non-anonymous but submodular cost functions. In a similar vein, [13] presents 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 Shapley value 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 $d + 1$, even when $n \to \infty$. Observe that for unweighted games PS and SV are identical.

Thus, the lower bound in [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 in [10]. 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 our model in more detail. We write $N = \{1, 2, \ldots, n\}$ for the *players* and $E = \{1, 2, \ldots, m\}$ for the *resources*. Each player $i \in N$ has a positive weight w_i and a *strategy set* $\mathcal{P}_i$, each element of which is a subset of the resources, i.e., $\mathcal{P}_i \subseteq 2^E$. We write $P = (P_1, \ldots, P_n)$ for an *outcome*, with $P_i \in \mathcal{P}_i$ the *strategy* of player i. Let $S_e(P) = \{i : e \in P_i\}$ be the set of users of e and $f_e(P) = \sum_{i \in S_e(P)} w_i$ be the *total weight* on e. The *joint cost* on e is $C_e(f_e(P)) = f_e(P) \cdot c_e(f_e(P))$, with c_e a function that is drawn from a given set of allowable functions $\mathcal{C}$. We write $\chi_{ie}(P)$ for the *cost share* of player i on resource e. These shares are such that $\sum_{i \in S_e(P)} \chi_{ie}(P) = C_e(f_e(P))$. We make the following assumptions on the *cost-sharing method* and the set of allowable cost functions:

(1) Every function that can appear in the game is continuous, nondecreasing, and convex. We also make the mild technical assumption that $\mathcal{C}$ is closed under dilation, i.e., that if $c(x) \in \mathcal{C}$, then also $c(a \cdot x) \in \mathcal{C}$ for $a > 0$. We note that without loss of generality, every $\mathcal{C}$ is also closed under scaling, i.e., if $c(x) \in \mathcal{C}$, then also $a \cdot c(x) \in \mathcal{C}$ for $a > 0$ (this is given by simple scaling and replication arguments).

(2) Given a player set S and a cost function c, suppose we alter the players (e.g., change their weights or identities) and the cost function c in a manner such that the cost generated by every subset of S on c remains unchanged. (For example suppose we initially have two players with weights 1 and 2 and cost function $c(x) = x^2$ and we modify them so that the weights are now 2 and 4 and the cost function is now $c(x) = x^2/4$.) Our second assumption states that the cost shares of the players will remain the same. In effect, we ask that the cost-sharing method only charges players based on how they contribute to the joint cost. We also assume, without loss of generality, that if the costs of all subsets of $\hat{S} = \{1, 2, \ldots, k\}$ are scaled versions of those corresponding to $S = \{1, 2, \ldots, k\}$, then the cost shares are also simply scaled by the same factor (again this is given by simple scaling and replication arguments).

(3) Since each cost share is a function of the cost function and player set on a resource, we will also write $\xi_c(i, S)$ for the share of a player i when the cost function is c and the player set is S. This means that the cost share of player i on resource e can be written both as $\chi_{ie}(P)$ and as $\xi_{c_e}(i, S_e(P))$. Our third assumption now states that the expression $\xi_{c_e}(i, S_e(P))$ is a continuous, nondecreasing, and convex function of the weight of player i. This is something to expect from a reasonable cost-sharing method, given that the joint cost on the resource is a continuous nondecreasing convex function of the weight of player i.

The *pure Nash equilibrium* (PNE) condition on an outcome P states that for every player i it must be the case that:

$$\sum_{e \in P_i} \chi_{ie}(P) \leq \sum_{e \in P_i'} \chi_{ie}(P_i', P_{-i}), \text{ for every } P_i' \in \mathcal{P}_i. \tag{1}$$

The social cost in the game will be the sum of the player costs, i.e.,

$$C(P) = \sum_{i \in N} \sum_{e \in P_i} \chi_{ie}(P) = \sum_{i \in N} \sum_{e \in P_i} \xi_{c_e}(i, S_e(P)) = \sum_{e \in E} f_e(P) \cdot c_e(f_e(P)). \tag{2}$$

Let $\mathcal{P}$ be the set of outcomes and $\mathcal{P}^N$ be the set of PNE outcomes of the game. Then the *price of anarchy* (POA) is defined as $POA = \frac{\max_{P \in \mathcal{P}^N} C(P)}{\min_{P \in \mathcal{P}} C(P)}$, and the *price of stability* (POS) is defined as $POS = \frac{\min_{P \in \mathcal{P}^N} C(P)}{\min_{P \in \mathcal{P}} C(P)}$. 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 partitioned among the set of players S_e using e. Given an ordering π of the players in S_e, the marginal cost increase by players $i \in S_e$ is $C(f_i^\pi + w_i) - C(f_i^\pi)$, where f_i^π is the total weight of players preceding i in the ordering. For a given distribution Π over orderings, the cost share of player i is $E_{\pi \sim \Pi}[C(f_i^\pi + w_i) - C(f_i^\pi)]$. For the weighted Shapley value, the distribution over orderings is given by a sampling parameter λ_i for each player i. The last player of the ordering is picked proportional to the sampling parameter λ_i. This process is then repeated iteratively for the remaining players.

As in [10], we study a *parameterized class of weighted Shapley values* defined by a parameter γ. For this class $\lambda_i = w_i^\gamma$ for all players i. For $\gamma = 0$ this reduces to the (normal) *Shapley value* (SV), where we have a uniform distribution over orderings.

3 Tight POA Bounds for General Cost-Sharing Methods

We first generalize the (λ, μ)-smoothness framework of [18] to accommodate any cost-sharing method and set of possible cost functions. 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

$$\sum_{i \in T^*} \xi_c(i, T \cup \{i\}) \leq \lambda \cdot w_{T^*} \cdot c(w_{T^*}) + \mu \cdot w_T \cdot c(w_T), \tag{3}$$

where $w_S = \sum_{i \in S} w_i$ for any set of players S. Then, for P a PNE and P^* the optimal solution, we would get

$$
\begin{aligned}
C(P) &\overset{(2)}{=} \sum_{i \in N} \sum_{e \in P_i} \xi_{c_e}(i, S_e(P)) \overset{(1)}{\leq} \sum_{i \in N} \sum_{e \in P_i^*} \xi_{c_e}(i, S_e(P) \cup \{i\}) \\
&= \sum_{e \in E} \sum_{i \in S_e(P^*)} \xi_{c_e}(i, S_e(P) \cup \{i\}) \\
&\overset{(3)}{\leq} \sum_{e \in E} \lambda \cdot w_{S_e(P^*)} c_e(w_{S_e(P^*)}) + \mu \cdot w_{S_e(P)} \cdot c_e(w_{S_e(P)}) \\
&\overset{(2)}{=} \lambda \cdot C(P^*) + \mu \cdot C(P). \qquad (4)
\end{aligned}
$$

Rearranging (4) 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., [18] for more). We then get the following lemma.

Lemma 1. *Consider the following optimization program with variables λ, μ.*

$$
\begin{aligned}
&\textit{Minimize} \quad \tfrac{\lambda}{1-\mu} && (5)\\
&\textit{Subject to} \quad \mu \leq 1 && (6)\\
&\textstyle\sum_{i \in T^*} \xi_c(i, T \cup \{i\}) \leq \lambda \cdot w_{T^*} \cdot c(w_{T^*}) + \mu \cdot w_T \cdot c(w_T), \forall c, T, T^* && (7)
\end{aligned}
$$

Every feasible solution yields a $\lambda/(1-\mu)$ upper bound on the POA of the cost sharing method given by $\xi_c(i, S)$ and the set of cost functions $\mathcal{C}$.

The upper bound holds for any cost-sharing method and set of cost functions. We now proceed to show that the optimal solution to Program (5)-(7) gives a tight upper bound when our assumptions described in Section 2 hold.

Theorem 1. *Let (λ^*, μ^*) be the optimal point of Program (5)-(7). The POA of the cost-sharing method given by $\xi_c(i, S)$ and the set of cost functions $\mathcal{C}$ is precisely $\lambda^*/(1-\mu^*)$.*

Proof. First define $\zeta_c(y, x)$ for $y, x > 0$ as

$$
\zeta_c(y, x) = \max_{T^*: w_{T^*} = y, T: w_T = x} \sum_{i \in T^*} \xi_c(i, T \cup \{i\}). \qquad (8)
$$

With this definition, we can rewrite Program (5)-(7) as

$$
\begin{aligned}
&\text{Minimize} \quad \tfrac{\lambda}{1-\mu} && (9)\\
&\text{Subject to} \quad \mu \leq 1 && (10)\\
&\zeta_c(y, x) \leq \lambda \cdot y \cdot c(y) + \mu \cdot x \cdot c(x), \ \forall c \in \mathcal{C}, x, y && (11)
\end{aligned}
$$

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 the cost function 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 (9)-(11) as

$$\begin{aligned} \text{Minimize} \quad & \tfrac{\lambda}{1-\mu} && (12)\\ \text{Subject to} \quad & \mu \le 1 && (13)\\ & \zeta_c(1,x) \le \lambda + \mu \cdot x \cdot c(x),\ \forall c \in \mathcal{C}, x && (14) \end{aligned}$$

The Lagrangian dual of Program (12)-(14) is

$$\text{Minimize} \quad \tfrac{\lambda}{1-\mu} + \sum_{c\in\mathcal{C},x>0} z_{cx} \cdot (\zeta_c(1,x) - \lambda - \mu \cdot x \cdot c(x)) + z_\mu \cdot (\mu - 1) \quad (15)$$

$$\text{Subject to} \quad z_{cx}, z_\mu \ge 0 \quad (16)$$

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 [21, 23] 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 $\lambda^*, \mu^*, z^*_{cx}$:

$$\frac{1}{1-\mu^*} = \sum_{c\in\mathcal{C},x>0} z^*_{cx} \quad (17)$$

$$\frac{\lambda^*}{(1-\mu^*)^2} = \sum_{c\in\mathcal{C},x>0} z^*_{cx} \cdot x \cdot c(x) \quad (18)$$

Calling $\eta_{cx} = z^*_{cx} / \sum_{c\in\mathcal{C},x>0} z^*_{cx}$ and dividing (18) with (17) we get

$$\frac{\lambda^*}{1-\mu^*} = \sum_{c\in\mathcal{C},x>0} \eta_{cx} \cdot x \cdot c(x). \quad (19)$$

By (19) and the fact that all constraints for which $z^*_{cx} > 0$ are tight (by complementary slackness), we get

$$\sum_{c\in\mathcal{C},x>0} \eta_{cx} \cdot \zeta_c(1,x) = \sum_{c\in\mathcal{C},x>0} \eta_{cx} \cdot x \cdot c(x). \quad (20)$$

Lower bound construction. Let $\mathcal{T} = \{(c,x) : z^*_{cx} > 0\}$. The construction starts off with a single player i, who has weight 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 $\mathcal{C}$ such that $c_{e_i}(1) \neq 0$ (it is easy to see that such a function exists, since $\mathcal{C}$ is closed under dilation, unless all function are 0, which is a trivial case) scaled so that $c_{e_i}(1) = \sum_{(c,x)\in\mathcal{T}} \eta_{cx} \cdot \zeta_c(1,x)$. 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 (8) 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 (8) 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 $\eta_{cx} \cdot \zeta_c(1, x)$. It then follows that she is indifferent between her two strategies. Note that the PNE cost of i is $\sum_{(c,x)\in\mathcal{T}} \eta_{cx} \cdot \zeta_c(1, x)$, which by (19) and (20) 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) 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 $M \to \infty$, 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 $\mathcal{C}$ 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. Evey 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 (20), 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 $M \to \infty$. 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 (12)-(14) 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 $c \in \mathcal{C}$, such that, for any arbitrarily large λ, there exists $x > 0$ such that $\zeta_c(1, x) > \lambda + \mu \cdot x \cdot c(x)$. We can rewrite this last expression as $\zeta_c(1, x)/(x \cdot c(x)) > \mu + \lambda/(x \cdot c(x))$, which shows we have c, x values such that $\zeta_c(1, x)$ is arbitrarily close to $x \cdot c(x)$ 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 (19) and (20) by setting $\eta_{cx} = 1$ for our c, 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. □

4 Shapley Value POS

In this section we study the POS for a class of weighted Shapley values, where the sampling parameter of each player i is defined by $\lambda_i = w_i^\gamma$ for some γ.

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. For the SV, existence of a PNE has been shown in [14] with the help of the following potential function, which is defined for an arbitrary ordering of the players:

$$\Phi(P) = \sum_{e \in E} \Phi_e(P) = \sum_{e \in E} \sum_{i \in S_e(P)} \xi_{c_e}(i, \{j : j \leq i, j \in S_e(P)\}). \quad (21)$$

We first prove the following lemma which is the main tool for proving our upper bound on the POS.

Lemma 2. *Suppose we are given an outcome of the game and suppose we substitute any given player i with two players who have weight $w_i/2$ each and who use the exact same resources as i. Then the value of the potential function will be at most the same as before the substitution.*

Proof. First rename the players so that the substituted player i has the highest index. Assign indices i' and i'' to the new players, with $i'' > i' > i$. On every resource e that is used by these players, the potential decreases by $\xi_{c_e}(i, S_e(P))$, while it increases by $\xi_{c_e}(i', S_e(P) \cup \{i'\} \setminus \{i\}) + \xi_{c_e}(i'', S_e(P) \cup \{i', i''\} \setminus \{i\})$. Hence, it suffices to show that

$$\xi_{c_e}(i', S_e(P) \cup \{i'\} \setminus \{i\}) + \xi_{c_e}(i'', S_e(P) \cup \{i', i''\} \setminus \{i\}) \leq \xi_{c_e}(i, S_e(P)). \quad (22)$$

For simplicity, in what follows call $\xi = \xi_{c_e}(i, S_e(P))$, $\xi' = \xi_{c_e}(i', S_e(P) \cup \{i'\} \setminus \{i\})$, and $\xi'' = \xi_{c_e}(i'', S_e(P) \cup \{i', i''\} \setminus \{i\})$. Consider every ordering π of the players in $S_e(P) \setminus \{i\}$ and every possible point in the ordering where a new player can be placed. If we assume that player is i and we average all possible joint cost jumps i can cause (by definition of the SV) we get ξ. Similarly, with i', we get ξ'. If we repeat the same thought process for i'', we are not getting ξ'', since the position of i' in the ordering is unspecified. However, we get a value that is larger than ξ'' if we always place i' right before i''. Call this larger value $\hat{\xi}''$. Observe that if we take every ordering π of $S_e(P) \setminus \{i\}$ and in every possible position, we place first i' and then i'' and we take the average of the combined joint cost jump that they cause, we will be getting $\xi' + \hat{\xi}''$, which, as we explained is at least $\xi' + \xi''$. Now note that this combined jump of the two players will also be the jump that i would cause in that particular position (since $w_{i'} + w_{i''} = w_i$), which means $\xi' + \hat{\xi}'' = \xi$, which in turn gives $\xi \geq \xi' + \xi''$ and completes the proof. □

By repeatedly applying Lemma 2, we can break the total weight on each resource in players of infinitesimal size and the value of the potential will not increase. This suggests:

$$\Phi_e(P) \geq \int_0^{f_e(P)} c_e(x)dx. \tag{23}$$

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. We get:

$$\begin{aligned} C(P^*) &\overset{(21)}{\geq} \Phi(P^*) \overset{\text{Def.}P}{\geq} \Phi(P) \overset{(23)}{\geq} \sum_{e \in E} \int_0^{f_e(P)} c_e(x)dx \\ &= \frac{\sum_{e \in E} \int_0^{f_e(P)} c_e(x)dx}{\sum_{e \in E} f_e(P) \cdot c_e(f_e(P))} \cdot C(P) \geq \min_{e \in E} \frac{\int_0^{f_e(P)} c_e(x)dx}{f_e(P) \cdot c_e(f_e(P))} \cdot C(P). \end{aligned}$$

Rearranging yields the following theorem.

Theorem 2. *The POS of the SV with $\mathcal{C}$ the set of allowable cost functions is at most* $\max_{c \in \mathcal{C}, x > 0} \frac{x \cdot c(x)}{\int_0^x c(x')dx'}$.

Corollary 1. *For polynomials with non-negative coefficients and degree at most d, the POS of the SV is at most $d+1$, which asymptotically matches the lower bound of [7] for unweighted games.*

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 which matches the corresponding lower bound on the POA in [10]. Our bound for $\gamma > 0$ even matches the upper bound on the POA [10], which holds for the weighted Shapley value in general. Our constructions modify the corresponding instances in [10], making sure that they have a unique Nash equilibrium. Due to page restrictions we defer the proof to our full version.

Theorem 3. *For polynomial cost functions with non-negative coefficients and maximum degree d, the POS for the class of weighted Shapley values with sampling parameters $\lambda_i = w_i^\gamma$ is at least*

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

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact price of anarchy for polynomial congestion games. SIAM Journal on Computing 40(5), 1211–1233 (2011)

2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM Journal on Computing 38(4), 1602–1623 (2008)
3. Awerbuch, B., Azar, Y., Epstein, A.: The price of routing unsplittable flow. In: Proceedings of STOC, pp. 57–66. ACM (2005)
4. Bhawalkar, K., Gairing, M., Roughgarden, T.: Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. ACM Transactions on Economics and Computation 2(4), 14 (2014)
5. Chen, H.L., Roughgarden, T.: Network design with weighted players. Theory of Computing Systems 45(2), 302–324 (2009)
6. Chen, H.L., Roughgarden, T., Valiant, G.: Designing network protocols for good equilibria. SIAM Journal on Computing 39(5), 1799–1832 (2010)
7. Christodoulou, G., Gairing, M.: Price of stability in polynomial congestion games. In: Proceedings of ICALP, pp. 496–507. Springer (2013)
8. Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: Proceedings of STOC, pp. 67–73. ACM (2005)
9. Gairing, M., Schoppmann, F.: Total latency in singleton congestion games. In: Proceedings of WINE, pp. 381–387. Springer (2007)
10. Gkatzelis, V., Kollias, K., Roughgarden, T.: Optimal cost-sharing in weighted congestion games. In: Proceedings of WINE, pp. 72—88. Springer (2014)
11. Gopalakrishnan, R., Marden, J.R., Wierman, A.: Potential games are necessary to ensure pure Nash equilibria in cost sharing games. Mathematics of Operations Research (2014)
12. Harks, T., Klimm, M.: On the existence of pure Nash equilibria in weighted congestion games. Mathematics of Operations Research 37(3), 419–436 (2012)
13. Klimm, M., Schmand, D.: Sharing non-anonymous costs of multiple resources optimally. arXiv preprint arXiv:1412.4456 (2014); to appear in CIAC 2015
14. Kollias, K., Roughgarden, T.: Restoring pure equilibria to weighted congestion games. In: Proceedings of ICALP, pp. 539–551. Springer (2011)
15. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. Computer Science Review 3(2), 65–69 (2009)
16. Monderer, D., Shapley, L.S.: Potential games. Games and Economic Behavior 14(1), 124–143 (1996)
17. Rosenthal, R.W.: A class of games possessing pure-strategy Nash equilibria. International Journal of Game Theory 2(1), 65–67 (1973)
18. Roughgarden, T.: Intrinsic robustness of the price of anarchy. In: Proceedings of STOC, pp. 513–522. ACM (2009)
19. Roughgarden, T., Schrijvers, O.: Network cost-sharing without anonymity. In: Proceedings of SAGT, pp. 134–145. Springer (2014)
20. Roughgarden, T., Tardos, É.: How bad is selfish routing? Journal of the ACM (JACM) 49(2), 236–259 (2002)
21. Shapiro, A.: On duality theory of convex semi-infinite programming. Optimization 54(6), 535–543 (2005)
22. Sheffi, Y.: Urban transportation networks: equilibrium analysis with mathematical programming methods. Prentice-Hall (1985)
23. Wu, S.Y., Fang, S.C.: Solving convex programs with infinitely many linear constraints by a relaxed cutting plane method. Computers & Mathematics with Applications 38(3), 23–33 (1999)