

Optimal Cost-Sharing in General Resource Selection Games

Vasilis Gkatzelis

Department of Computer Science, Stanford University, Stanford, CA 94305, gkatz@cs.stanford.edu

Konstantinos Kollias

Department of Management Science and Engineering, Stanford University, Stanford, CA 94305, kkollias@stanford.edu

Tim Roughgarden

Department of Computer Science, Stanford University, Stanford, CA 94305, tim@cs.stanford.edu

Resource selection games provide a model for a diverse collection of applications where a set of resources is matched to a set of demands. Examples include routing in traffic and in telecommunication networks, service of requests on multiple parallel queues, acquisition of services or goods with demand-dependent prices, etc. In reality, demands are often submitted by selfish entities (players) and congestion on the resources results in negative externalities for their users. We consider a policy maker that can set a priori rules to minimize the inefficiency induced by selfish players. For example, these rules may assume the form of scheduling policies or pricing decisions. We explore the space of such rules abstracted as cost-sharing methods. We prescribe desirable properties that the cost-sharing method should possess and prove that, in this natural design space, the cost-sharing method induced by the Shapley value minimizes the worst-case inefficiency of equilibria.

Key words: resource selection, cost-sharing, Shapley value, price of anarchy, network routing

History: Submitted October 2014; revision submitted August 2015; accepted April 2016.

1. Introduction

Resource selection games offer an abstraction for many interesting applications sharing a common theme: there is a set of selfish demands (players) that choose a set of resources to use. The presence of multiple players on the same resource causes an undesirable congestion effect which increases the cost each one of them suffers. The applications that lie within this framework are diverse: routing information packets in a telecommunications network or vehicles in a road network (Roughgarden and Tardos (2002), Awerbuch et al. (2005)), constructing the infrastructure of a network (Anshelevich et al. (2008), Chen et al. (2010), von Falkenhausen and Harks (2013)), scheduling tasks for

processing on parallel servers (Haviv and Roughgarden (2007)), and pricing congestion-dependent services (Cominetti et al. (2009)). A common characteristic of these applications is that the joint cost of a resource depends only on the total demand on that resource. Resource selection games can be considered as generalizations of congestion games (Rosenthal (1973a)).

Enforcing a socially optimal allocation of the various resources (links, channels, roads, processors, etc.) to the players is typically not feasible. We consider players that are not subject to centralized control and pick the resources that allow them to optimize their own objectives. In the presence of self-interested players, the obvious goal of a policy maker or system designer is to carefully design the rules of the system to incentivize players to reach a socially desirable outcome. More concretely, we assume that players reach an equilibrium of the induced game – a state such that no unilateral deviation is beneficial to a player – and the goal is to provide a guarantee on how well this equilibrium approximates a socially optimal allocation.

The way a policy maker or system designer can influence the game may vary, depending on the application specifics, but a general abstraction that captures many such game-theoretic control concepts is that of cost-sharing methods. In resource selection games the total demand that is allocated to a specific resource generates a joint cost. This cost can be monetary (e.g., a total payment requested for specific services) or not (e.g., aggregate queueing delay). The decision of the policy maker then reduces to picking the cost-sharing method that determines which fraction of this generated cost each of the users is responsible for.

Example: Routing of Packets in a Telecommunications Network. Consider the application of routing packets from many different senders in a network. In typical cases, such as the Internet and wireless networks, data from a sender to a receiver is sent on a fixed route. All packets that traverse a given link in the network are initially queued and then forwarded by the local router. Each sender has a given rate of injecting packets in the network and selects the route that minimizes the total delay of her packets over all links (queues) they traverse. Consider a given link in the network and the corresponding queue, which for the sake of this example can be assumed to be

M/M/1. The arrival rate at the queue is given by the sum of rates of the users of the link (i.e., all senders that include this link in their routes). This total arrival rate can be used to calculate the M/M/1 average delay of a packet on this link. The total delay is then given as the product of the average delay and the number of packets. A cost-sharing method in this setting would correspond to a scheduling policy that assigns priorities to different senders. Rearranging the order of packets in the queue, without introducing any idle time, does not impact the average delay and, hence, each policy yields a different method of sharing the same joint cost (Shenker (1995), Mosk-Aoyama and Roughgarden (2009)). We note that the converse does not hold and that in this application not every cost-sharing method can be implemented by a scheduling policy (see Coffman Jr. and Mitrani (1980) for a characterization).

Example: Economics of the Transportation of Goods. In the spirit of Cominetti et al. (2009), consider the transportation of goods by logistics and freight companies. Shipments go through transportation hubs such as airports, harbors, and train stations where operational costs vary depending on the total volume of shipments. Freight companies are then charged fees which cover these operational costs. The objective of each freight company is to identify the least costly route between the origin and destination of their shipment. In such situations where the joint cost is monetary, the connection with cost-sharing methods is clear and unrestricted. The policy maker designs the method that distributes the operational costs to the companies using each transportation hub.

We assume that the policy maker seeks the cost-sharing method that maximizes the efficiency of the resulting outcome. The design space is vast and we next identify some important properties that the cost-sharing method should possess (Chen et al. (2010), von Falkenhausen and Harks (2013)) and that give a crisper image of the available alternatives. We give an informal definition of these properties here and present them formally in Section 2.

1. *Budget-balance*: the joint cost on each resource is covered precisely by its users.
2. *Stability*: the induced resource selection game is guaranteed to possess an equilibrium.
3. *Locality*: cost-sharing on a resource is independent of the system's state beyond that resource.

Property 1 requests that the cost shares exactly cover the joint cost. We also consider the impact of overcharging the players in Section 5. Property 2 requires that the cost-sharing method guarantees the existence of a pure Nash equilibrium, which is a highly desired property. There are, of course, more general equilibrium concepts like the *mixed Nash equilibrium* that are guaranteed to exist in every finite game, but randomized solution concepts suffer from well-known drawbacks in interpretation and implementation (see, e.g., (Osborne and Rubinstein 1994, §3.2)). Again, we discuss the case of dropping this condition in our conclusions. Finally, Property 3 is important for resource selection games since all the motivating applications concern large systems where knowledge regarding the state of the system beyond the resource at hand is either infeasible or very costly. Also, in such systems, resources are added and removed constantly so introducing complicated dependencies among them is bound to cause scalability issues to the system. We call a cost-sharing method *admissible* if it possesses all three properties.

1.1. Our Contributions and Paper Structure

We study resource selection games from the policy maker’s perspective and, given any set of allowable convex and increasing joint resource cost functions, we characterize the optimal admissible cost-sharing method. Our main result states that, among all such methods, the (unweighted) Shapley value (Shapley (1953)) is the one that minimizes the worst-case *price of anarchy*, i.e., the worst-case ratio between the total player cost in an equilibrium to the total player cost in the socially optimal allocation of demands to resources.

In Section 2 we present our model and some examples. To illustrate our main ideas, we present a case study on resource selection games with polynomial cost functions in Section 3. There, we examine the performance of a class of parameterized Shapley values such that a single parameter controls the relative advantage or disadvantage that is given to players with larger demand. We characterize the price of anarchy as a function of this parameter and we prove that it is minimized when the parameter is zero (which corresponds to the Shapley value). In Section 4 we strengthen this result by showing that the Shapley value remains optimal among all admissible cost-sharing

methods for all convex cost functions. We conclude our paper and discuss extensions to cost-sharing methods that are not admissible in Section 5. Omitted proofs are included in the electronic companion section of the paper.

1.2. Related Work

The performance of cost-sharing methods in resource selection games has recently received a lot of attention, leading to a sequence of results including the work of Chen et al. (2010) and von Falkenhausen and Harks (2013), which is the most relevant to our results. In Chen et al. (2010), the authors examine the whole space of admissible cost-sharing methods as well, but in networks where all players have equal demand and the cost functions of the resources are constant, i.e., in a setting with positive externalities. Their results exhibit optimality of the Shapley value in directed networks and optimality of (non-Shapley) simple priority protocols for undirected networks. The work of von Falkenhausen and Harks (2013) studies the inefficiency of cost-sharing methods in a model very similar to the one we adopt here. The differences from our approach are the following: first, rather than considering arbitrary strategy sets like we do, von Falkenhausen and Harks (2013) focus on games where the players' strategies are singletons or matroids. Also, von Falkenhausen and Harks (2013) keep the set of allowable cost-functions unrestricted and prove that the price of anarchy of any admissible cost-sharing method is unbounded. Apart from admissible cost-sharing methods they also consider methods that violate the locality property in different ways and they characterize the optimal method in each case.

Cost-sharing methods for resource selection games were also the focus of Harks and Miller (2011), Marden and Wierman (2013), and Gopalakrishnan et al. (2014). In Harks and Miller (2011), the authors study the performance of several cost-sharing methods in a slightly modified setting, where each player declares a different demand for each resource. Marden and Wierman (2013) study various cost-sharing methods in a utility maximization model for the players, while Gopalakrishnan et al. (2014) characterized the space of admissible cost-sharing methods as the set of generalized weighted Shapley values; a characterization that we will use in this work.

An interesting family of resource selection games is the class of weighted congestion games. There has been a long line of work on these games focusing on the proportional sharing method, according to which the players share their joint cost in proportion to the size of their demands (Rosenthal (1973b), Milchtaich (1996), Monderer and Shapley (1996), Awerbuch et al. (2005), Gairing and Schoppmann (2007), Bhawalkar et al. (2014), Harks and Klimm (2012), Harks et al. (2011)). Proportional sharing is not an admissible cost-sharing method because it does not always guarantee the existence of an equilibrium, except for special cases such as when the per unit of load cost functions are quadratic or exponential (Fotakis and Spirakis (2008), Harks and Klimm (2012)). Kollias and Roughgarden (2015) were the first to propose using Shapley value based methods in congestion games in order to restore stability.

The impact of cost-sharing methods on the quality of equilibria has also been studied in other models: Moulin and Shenker (2001) focused on participation games, while Moulin (2008) and Mosk-Aoyama and Roughgarden (2009) studied queueing games. Also, very closely related in spirit is previous work on coordination mechanisms, beginning with Christodoulou et al. (2009) and subsequently in Immorlica et al. (2009), Azar et al. (2015), Caragiannis (2013), Abed and Huang (2012), Kollias (2013), Cole et al. (2015), Christodoulou et al. (2014), and Bhattacharya et al. (2014). Most work on coordination mechanisms concerns scheduling games and how the price of anarchy varies with the choice of local machine policies (i.e., the order in which to process jobs assigned to the same machine). Some of the recent work comparing the price of anarchy of different auction formats, such as Lucier and Borodin (2010), Bhawalkar and Roughgarden (2012), and Syrgkanis and Tardos (2013) also has a similar flavor.

2. The Model

In this section, we provide more details regarding our model. We begin with a set of definitions (Section 2.1) and then we expand on how this model applies to our motivating applications.

2.1. Definitions

DEFINITION 1 (RESOURCE SELECTION GAME). A *resource selection game* is defined by a tuple $(N, R, \mathcal{A}, w, C, \Xi)$. N is a finite set of players, R is a finite set of resources, and $\mathcal{A} = \times_{i \in N} \mathcal{A}_i$ is the

set of strategy profiles, with $\mathcal{A}_i \subseteq 2^R$ being the strategy set of player i . The vector $w = (w_i)_{i \in N}$ contains the demand w_i of each player i , and $C = (C_r)_{r \in R}$ is a vector with $C_r : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ being the cost function of resource r . Finally, Ξ is a cost-sharing method (defined below).

We write $l(S) = \sum_{i \in S} w_i$ to denote the total demand of the players in $S \subseteq N$. Given a strategy vector $A \in \mathcal{A}$, we write $S_r(A) = \{i \in N : r \in A_i\}$ for the set of users of resource r . For ease of notation we also use $l_r(A)$ for $l(S_r(A))$, the total demand on resource r . Function C_r outputs the joint cost on the corresponding resource r as $C_r(l_r(A))$.

EXAMPLE 1. Consider a game with two players $N = \{1, 2\}$ having demands $w_1 = w_2 = 1$, and two resources $R = \{a, b\}$ with cost functions $C_a(x) = x^2$ and $C_b(x) = 2 \cdot x^3$. The action sets are $\mathcal{A}_1 = \mathcal{A}_2 = \{\{a\}, \{b\}\}$, which means that each of the players will use exactly one of the resources. If $A_1 = A_2 = \{a\}$ (both players decide to use resource a), then the total demand on a is $l_a(A) = 2$ and the joint cost on a is $C_a(l_a(A)) = C_a(2) = 4$. Now, if $A_1 = A_2 = \{b\}$ (both players use resource b) and the total demand on b is $l_b(A) = 2$ leading to a joint cost of $C_b(l_b(A)) = C_b(2) = 16$. Finally, if the players decide to use different resources, then $C_a(l_a(A)) = C(1) = 1$ and $C_b(l_b(A)) = C_b(1) = 2$.

A *cost-sharing method* Ξ defines a cost share function $\xi_{i,r} : \mathcal{A} \rightarrow \mathbb{R}_{\geq 0}$ for all $i \in N$ and $r \in R$. The total cost of a player i , when the strategy profile is A , is $\xi_i(A) = \sum_{r \in A_i} \xi_{i,r}(A)$. We proceed with the definition of a pure Nash equilibrium.

DEFINITION 2 (PURE NASH EQUILIBRIUM). The strategy vector A is a *pure Nash equilibrium* (PNE) if for every player i and every $A'_i \in \mathcal{A}_i$

$$\xi_i(A) \leq \xi_i(A'_i, A_{-i}). \quad (1)$$

EXAMPLE 2. We revisit the game from Example 1 and suppose the cost-sharing method simply dictates that the users of each resource share the joint cost of that resource equally. For example, consider the $A_1 = A_2 = \{a\}$ case. The joint cost is 4 and hence for both players we get $\xi_1(A) = \xi_2(A) = 2$. If any of them were to deviate to strategy $\{b\}$, the incurred cost would be $C_b(1) = 2$, which shows that the equilibrium condition (1) holds. The same can be shown for the outcome $A_1 =$

$\{a\}$ and $A_2 = \{b\}$. Player 1 has cost $\xi_1(A) = 1$ while, by deviating to the outcome $A' = (\{b\}, \{b\})$, she would have a cost $\xi_1(A') = C_b(2)/2 = 8$. Similarly Player 2 suffers cost $\xi_2(A) = 2$, while by deviating to outcome $A'' = (\{a\}, \{a\})$, she would incur a cost $\xi_2(A'') = C_a(2)/2 = 2$. Note that the PNE $(\{a\}, \{a\})$ has total cost 4, while the PNE $(\{a\}, \{b\})$ has total cost 2 (which is also the optimal total cost).

We now formally define the properties that make cost-sharing method Ξ admissible.

DEFINITION 3 (COST-SHARING METHOD PROPERTIES). An admissible method Ξ satisfies:

1. *Budget-balance*: for every resource r , $C_r(A) = \sum_{i \in S_r(A)} \xi_{i,r}(A)$, and $\xi_{i,r}(A) = 0$ if $i \notin S_r(A)$.
2. *Stability*: every induced game $(N, R, \mathcal{A}, w, C, \Xi)$ always possesses at least one PNE.
3. *Locality*: for every $r \in R$ and any two profiles $A, A' \in \mathcal{A}$ with $S_r(A) = S_r(A')$, $\xi_{i,r}(A) = \xi_{i,r}(A')$

The work of Gopalakrishnan et al. (2014) shows that the set of methods satisfying the above properties is precisely a class that generalizes the weighted variants of the Shapley value, called *generalized weighted Shapley values*. We now define the class of weighted Shapley values, a subclass of generalized weighted Shapley values. As we explain at the end of this section (and the electronic companion), it is sufficient to focus on this more restricted class when seeking to optimize the worst-case PoA among admissible cost-sharing methods.

DEFINITION 4 (SHAPLEY VALUE). Consider a resource r and the set $S_r(A)$ of its users. For a given ordering π of the players in $S_r(A)$, let $S_r^{\pi,i}(A)$ denote the players preceding i in π . Then, the quantity $C_r(l(S_r^{\pi,i}(A)) + w_i) - C_r(l(S_r^{\pi,i}(A)))$ is the marginal increase in the joint cost caused by player i when only the players preceding her in π and herself are using the resource. The Shapley value of a Player $i \in S_r(A)$ is the expected value of this increase with respect to the uniform distribution over all orderings π :

$$\xi_{i,r}^{\text{SV}}(A) = \mathbb{E}_{\pi \sim \text{uniform}} [C_r(l(S_r^{\pi,i}(A)) + w_i) - C_r(l(S_r^{\pi,i}(A)))] . \quad (2)$$

Weighted variants of the Shapley value use non-uniform distributions over orderings. To be more precise, these variants assign a positive sampling weight λ_i to each player i and they iteratively

pick the player that goes last in the ordering among the still unassigned players in each iteration, with probabilities proportional to these sampling weights.

DEFINITION 5 (WEIGHTED SHAPLEY VALUE). A weighted Shapley value is defined by a vector λ of sampling weights, one for each player. Given a resource r and its users $S_r(A)$, consider the following process of generating an ordering π of the players in $S_r(A)$. The player that goes last in the ordering is picked with probabilities proportional to the sampling weights, i.e., player i has probability $\lambda_i / \sum_{j \in S_r(A)} \lambda_j$ of being the last player. The process is repeated with the penultimate player being selected among the remaining ones in the same fashion, and so on. The weighted Shapley value of player i is the expected increase she causes to the joint cost with respect to the distribution Δ_λ on orderings induced by λ :

$$\xi_{i,r}^w = \mathbb{E}_{\pi \sim \Delta_\lambda} [C_r(l(S_r^{\pi,i}(A)) + w_i) - C_r(l(S_r^{\pi,i}(A)))] . \quad (3)$$

A weighted Shapley value with all sampling weights equal is equivalent to the Shapley value.

An important observation is that changing the cost-sharing method being used leads to a different game, and hence to different equilibrium outcomes in the induced game. As a result, the efficiency of this game crucially depends on the choice of the cost-sharing method. As a measure of the inefficiency of an outcome A , we will be using the total cost $Q(A) = \sum_{r \in R} C_r(l_r(A))$. In order to quantify the quality of different cost-sharing methods we use the price of anarchy metric.

DEFINITION 6 (PRICE OF ANARCHY). Given a collection of cost functions \mathcal{C} and a cost-sharing method Ξ , the price of anarchy $\text{PoA}(\mathcal{C}, \Xi)$ is the worst-case ratio of equilibrium cost to optimal cost over all resource selection games $\mathcal{G} = (N, R, \mathcal{A}, w, C, \Xi)$ with $C_r \in \mathcal{C}$ for all $r \in R$:

$$\text{PoA}(\mathcal{C}, \Xi) = \sup_{\mathcal{G}} \frac{\max_{A \in \text{PNE}(\mathcal{G})} Q(A)}{\min_{A^* \in \mathcal{A}} Q(A^*)} . \quad (4)$$

Depending on the particular application of interest, the set of cost functions \mathcal{C} can be restricted in order to model the corresponding types of resources being used. Moreover, appropriately restricting the set \mathcal{C} is important since the setting where any cost function can appear in the game is technically

uninteresting as it leads to infinite PoA for all cost-sharing methods (see von Falkenhausen and Harks (2013)). Therefore, to obtain meaningful PoA bounds, we parameterize the PoA by the set of allowable resource cost functions. In particular, we focus on resources whose cost functions are positive, increasing, convex, and chosen from a given set \mathcal{C} . On the other hand, we restrict neither the number of resources and players nor the strategy sets of the players, which allows us to evaluate the quality of systems whose structure may change over time, i.e., new resources may be added to the system and new players with different strategy sets may join over time.

The main result of Gopalakrishnan et al. (2014) states that for any given set of cost functions \mathcal{C} that are *local*, i.e., such that the cost of a resource depends only on the players using it, the only admissible cost-sharing methods are the ones that belong to the class of generalized weighted Shapley values. The cost functions that we focus on in this work depend only on the weights of these users, hence we can focus on this class of cost-sharing methods without loss of generality.

A generalized weighted Shapley value is allowed to partition the players into groups $1, 2, \dots, k$, so that the players in group i always appear before the players in group $j > i$ in the random Shapley ordering. The relative order of the players *within* a given group is then determined using their sampling weights, similarly to the weighted Shapley value. As far as optimizing the PoA is concerned though, it suffices to focus on the more restricted class of weighted Shapley values (as opposed to the class of generalized weighted Shapley values). We explain this fact in Section 6 of the electronic companion.

2.2. Motivating Applications Revisited

Here we present some instances with specific cost functions and policies. We show how these policies are mapped to cost-sharing methods and present numerical examples.

Packet Routing. Consider the packet routing application from Section 1. Focus on a single resource r and suppose the resource cost function is the aggregate delay function of an M/M/1 queue with capacity 4, i.e., $C_r(x) = x/(4-x)$ for $x < 4$. Let players 1 and 2 be the sole users of r when the action profile is A and let their demands be $w_1 = 1$ and $w_2 = 2$ respectively. We then get

$C_r(l_r(A)) = 3$ for the joint cost of the two players. We present three alternative scheduling policies: (1) the First In First Out (FIFO) protocol, (2) a priority protocol that always forwards the packets of the smallest demand first, and (3) a protocol that randomly selects a player at each step and forwards her next packet. We now discuss the cost-sharing interpretations of these policies.

FIFO is known to correspond to proportional sharing, according to which the cost-share of a player is proportional to the demand that the player has placed on the resource (Shenker (1995)). Hence, in our example, Player 1 receives a $1/3$ fraction of the total cost, while Player 2 receives the remaining $2/3$ fraction, leading to $\xi_{1,r}(A) = 1$ and $\xi_{2,r}(A) = 2$. Remember that proportional sharing is not admissible since there exist instances for which it violates the stability property (see Goemans et al. (2005)).

The priority protocol always prefers packets of Player 1 whenever packets of both players are in the queue. This suggests that the delay experienced by Player 1 is the same as if she were the only user of r , i.e., $\xi_{1,r}(A) = 1/(4-1) = 1/3$. The aggregate delay does not change by this rearrangement of packets, so Player 2 suffers the remaining $\xi_{2,r}(A) = 3 - 1/3 = 8/3$ cost.

The method that randomly selects the next player whose packet gets forwarded can be interpreted as the Shapley value, which is a randomization over all priority protocols, i.e., over all the orderings of the players. From Definition 4 we get that Player 1 suffers a cost equal to:

$$\xi_{1,r}(A) = \frac{1}{2} \cdot \frac{1}{4-1} + \frac{1}{2} \cdot \left(\frac{3}{4-3} - \frac{2}{4-2} \right) = \frac{7}{6}. \quad (5)$$

Similarly we get the cost share of Player 2 as:

$$\xi_{2,r}(A) = \frac{1}{2} \cdot \frac{2}{4-2} + \frac{1}{2} \cdot \left(\frac{3}{4-3} - \frac{1}{4-1} \right) = \frac{11}{6}. \quad (6)$$

Transportation of goods. We now focus on the second application from Section 1, the economics of the transportation of goods. Let the joint cost on a hub r be cubic, i.e., $C_r(x) = x^3$. Suppose again that the sole users of r in strategy profile A are players 1 and 2 with demands $w_1 = 1$ and $w_2 = 2$. Here the costs are monetary and the cost-sharing rule can be arbitrary. The Shapley value outputs the following cost shares. For Player 1:

$$\xi_{1,r}(A) = \frac{1}{2} \cdot 1^3 + \frac{1}{2} \cdot (3^3 - 2^3) = 10. \quad (7)$$

For Player 2:

$$\xi_{2,r}(A) = \frac{1}{2} \cdot 2^3 + \frac{1}{2} \cdot (3^3 - 1^3) = 17. \quad (8)$$

An order-based cost-sharing method that assumes the players join the resource, one at a time, in increasing demand order, and charges them the marginal increase they cause on the joint cost, assigns the cost shares $\xi_{1,r}(A) = 1$ and $\xi_{2,r}(A) = 26$.

Equilibrium considerations. Switching from one cost-sharing method to another impacts the set of equilibria that the game possesses. Consider a game $(N, R, \mathcal{A}, w, C, \Xi)$: the player set is $N = \{1, 2\}$ and the demands are $w_1 = 1$ and $w_2 = 2$. The resources are $R = \{a, b\}$ with cost functions $C_a(x) = x^2$ and $C_b(x) = 3 \cdot x$. Strategies are singletons and both players are called to pick one of the two resources. Suppose Ξ is the Shapley value. Then, when both players select resource a , we get a PNE with total cost 9. If we switch to an order-based protocol that gives priority to smaller demands, this strategy profile is not a PNE any more. The unique PNE has Player 1 use resource a and Player 2 use resource b for a total cost of 7. We can see that 7 is also the optimal total cost, which suggests that the PoA of the Shapley value in this instance is $9/7$, while the PoA of the order-based protocol is 1. Hence, this shortest-first protocol performs better in this instance. However, we can verify that the opposite holds in the following game. Let $N = \{1, 2\}$ with demands $w_1 = 1$ and $w_2 = 2$. Let $R = \{a, b, c\}$ with cost functions $C_a(x) = x^2$, $C_b(x) = C_c(x) = x^2/2$. Player 1 must choose one of resources b and c , while Player 2 must choose one of a and b and in the unique equilibrium of the Shapley value cost-sharing method, 1 picks c and 2 picks b . The shortest-first order-based method however also admits the equilibrium with 1 to b and 2 to a which has a higher total cost. In this instance, the PoA of the Shapley value is 1 and the PoA of the order-based protocol is $9/5$.

As hinted at above, no cost-sharing method is optimal across all instances. This motivates our focus on worst-case guarantees on the price of anarchy.

3. Polynomial Cost Functions and Parameterized Weighted Shapley Values

As discussed in Section 2, the price of anarchy is determined as a function of two parameters: the set of resource cost functions and the cost-sharing method. This section studies a natural and




Figure 1 A plot of the PoA when the set of cost functions is the set of polynomials with degree at most 2, and the cost-sharing method is the weighted Shapley value with sampling weights $\lambda_i = w_i^\gamma$.

interesting special case where the set of resource cost functions is the set of polynomials with nonnegative coefficients and degrees between 1 and d , and the cost-sharing method is a weighted Shapley value induced by sampling weights of the form $\lambda_i = w_i^\gamma$, for some $\gamma \in \mathbb{R}$. We think of the set of allowable cost functions (parameterized by d in this section’s case study) as determined by the nature of the application, while the cost-sharing method (parameterized by γ in this section’s case study) is set by the policy maker. Since resource cost functions are convex, players benefit from being positioned earlier in the random ordering generated by a weighted Shapley value, and the parameter γ controls the relative advantage (for $\gamma > 0$) or disadvantage (for $\gamma < 0$) that is given to smaller demands when generating the player ordering on a resource. Letting $\gamma \rightarrow +\infty$ (resp. $-\infty$) would yield a priority-based protocol that favors smaller (resp. larger) demands, while setting $\gamma = 0$ recovers the Shapley value, since all orderings are equally likely.

Different values for d and γ result in a different price of anarchy. Figure 1 plots how the PoA changes if we fix $d = 2$ and we let γ vary, which exhibits that even the simple case with $d = 2$ is interesting. In this plot, $\gamma = 0$ (the Shapley value) achieves a PoA equal to 2.618. Every other γ has PoA at least 4. In fact, any positive γ has PoA precisely 5.828, which is also the case for $\gamma \rightarrow -\infty$. The plot is derived as a special case of our statements and proofs in Section 3.1 and Section 3.2.

Degree	$\gamma < 0$	Shapley	$\gamma > 0$
2	≥ 4	2.618	5.828
3	≥ 27	12.626	56.948
4	≥ 256	101.58	780.279
d	$\geq d^d$	$(0.9 \cdot d)^d$	$(1.4 \cdot d)^d$

Table 1 PoA values for different values of the maximum degree d .

3.1. PoA Lower Bounds

Let Ξ^γ be a weighted Shapley value induced by sampling weights of the form $\lambda_i = w_i^\gamma$. Also, let \mathcal{C}_d denote the set of polynomials with nonnegative coefficients and degrees between 1 and d . In this section, we prove the following theorem.

THEOREM 1 (Optimal parameterized weighted Shapley value for polynomials). *For cost functions in \mathcal{C}_d , the Shapley value is optimal among all the cost-sharing methods Ξ^γ , i.e.,*

$$\arg \min_{\gamma} \text{PoA}(\mathcal{C}_d, \Xi^\gamma) = 0$$

As we explained earlier, the $\gamma = 0$ case corresponds to the Shapley value (since all orderings of the players on a resource have the same likelihood). In Kollias and Roughgarden (2015), it is shown that the PoA of the Shapley value, when the cost functions are drawn from \mathcal{C}_d , is equal to $\chi_d^d \approx (0.9 \cdot d)^d$, where χ_d is the largest root of $3 \cdot x^d - 1 - (x + 1)^d$. Our lower bounds show that any $\gamma \neq 0$ results in a PoA larger than χ_d^d .

LEMMA 1 (Lower bound for polynomials and $\gamma > 0$). *For positive values of γ :*

$$\text{PoA}(\mathcal{C}_d, \Xi^\gamma) \geq (2^{1/d} - 1)^{-d} \approx (1.4 \cdot d)^d.$$

As we prove in Theorem 2, this lower bound is in fact tight.

LEMMA 2 (Lower bound for polynomials and $\gamma < 0$). *For negative values of γ :*

$$\text{PoA}(\mathcal{C}_d, \Xi^\gamma) \geq d^d.$$

We prove an even stronger, but more complicated, lower bound for $\gamma < 0$ in Lemma 3 in the electronic companion. From Lemma 1 and Lemma 2, we get Theorem 1, which is the main result of this section and an early indication of the general optimality of the Shapley value among admissible cost-sharing methods.

3.2. PoA Upper Bounds

We first present an upper bound that applies to all weighted Shapley values (not only the ones parameterized by γ), for games with polynomial cost functions of maximum degree d .

THEOREM 2 (Upper bound for polynomials). *For any admissible cost-sharing method Ξ :*

$$PoA(\mathcal{C}_d, \Xi) \leq (2^{1/d} - 1)^{-d} \approx (1.4 \cdot d)^d.$$

Proof. Consider a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi)$. Let A be a PNE of the game and A^* the optimal outcome. For the total cost of A , we get:

$$Q(A) = \sum_{r \in R} C_r(l_r(A)) = \sum_{r \in R} \sum_{i \in N} \xi_{i,r}(A) = \sum_{i \in N} \sum_{r \in A_i} \xi_{i,r}(A) \leq \sum_{i \in N} \sum_{r \in A_i^*} \xi_{i,r}(A_i^*, A_{-i}). \quad (9)$$

The inequality follows from the equilibrium condition (1). Note that the cost share of any player on any resource, when the cost-sharing method is a weighted Shapley value and the resource costs are convex, is upper bounded by the increase that would be caused on the joint resource cost, if that player were to be the last in the ordering. This means that for every $r \in A_i^*$, we get:

$$\xi_{i,r}(A_i^*, A_{-i}) \leq C_r(l_r(A_i^*, A_{-i})) - C_r(l_r(A_i^*, A_{-i}) - w_i) \leq C_r(l_r(A) + w_i) - C_r(l_r(A)). \quad (10)$$

Combining (9) with (10), we get:

$$Q(A) \leq \sum_{i \in N} \sum_{r \in A_i^*} C_r(l_r(A) + w_i) - C_r(l_r(A)) \quad (11)$$

$$= \sum_{r \in R} \sum_{i \in S_r(A^*)} C_r(l_r(A) + w_i) - C_r(l_r(A)) \quad (12)$$

$$\leq \sum_{r \in R} C_r(l_r(A) + l_r(A^*)) - C_r(l_r(A)). \quad (13)$$

The last inequality follows by convexity of the expression as a function of w_i . We now claim that the following is true, for any $x, y \in \mathbb{R}$, and $d \geq 1$:

$$(x + y)^d - x^d \leq \hat{\lambda} \cdot y^d + \hat{\mu} \cdot x^d, \quad (14)$$

with

$$\hat{\lambda} = 2^{(d-1)/d} \cdot (2^{1/d} - 1)^{-(d-1)} \quad \text{and} \quad \hat{\mu} = 2^{(d-1)/d} - 1. \quad (15)$$

We can verify this as follows. Note that, without loss of generality, we can set $y = 1$ (equivalent to dividing both sides of (14) with y^d and renaming x/y to x). We can then see that the value of x that maximizes $(x + 1)^d - \hat{\mu} \cdot x^d$ is $x = 1/(2^{1/d} - 1)$, for which inequality (14) is tight. Also, note that the expressions for $\hat{\lambda}$ and $\hat{\mu}$ are increasing as functions of d , which implies that the given values for degree d , satisfy (14) for smaller degrees as well. This means we can combine (13), (14), and (15), to get:

$$Q(A) \leq \sum_{r \in R} \hat{\lambda} \cdot C_r(l_r(A^*)) + \hat{\mu} \cdot C_r(l_r(A)) = \hat{\lambda} \cdot Q(A^*) + \hat{\mu} \cdot Q(A). \quad (16)$$

Rearranging, we get $Q(A)/Q(A^*) \leq \hat{\lambda}/(\hat{\mu} + 1) = (2^{1/d} - 1)^{-d}$. \square

We prove a stronger, though more complicated, upper bound for $\gamma < 0$ in Lemma 4 in the electronic companion.

4. Convex Cost Functions and Admissible Cost-Sharing Methods

This section builds up to our main result. First we extend the results of Section 3 to the class of all weighted Shapley values (not just those of the form $\lambda_i = w_i^\gamma$) and polynomial cost functions and, subsequently, we prove optimality of the Shapley value in games that draw resource costs from arbitrary sets of allowable cost functions that satisfy the following conditions: cost functions are positive, increasing, convex, assign 0 cost to 0 demand, and the cost function set is closed under dilation (and scaling).

Our main result shows that the case study of Section 3 on polynomial costs and parameterized Shapley values actually captures the crucial properties and the whole difficulty of characterizing

the optimal admissible cost-sharing method in the general case. Our analysis shows that, in essence, every admissible cost-sharing method has to decide whether to treat all players roughly the same, to offer benefits to large demands, or to offer benefits to small demands. We then present lower bounds that either match the Shapley value upper bound from Kollias and Roughgarden (2015), or generalize the lower bounds of Section 3 for the special case of polynomials. Finally, we generalize these lower bounds even further, by showing that they hold for all classes of convex cost functions and do not depend on the special structure of polynomial costs.

A weighted Shapley value has the power to determine a player's sampling weight as a function of her identity and demand. Hence, we suppose the description of a weighted Shapley value is given to us as a collection of functions, one for each player identity, that map a demand to a sampling weight. We make the following mild technical assumptions: (a) the sampling weight functions are continuous, i.e., slightly perturbing a player's demand, also slightly perturbs her sampling weight and (b) the limit of the sampling weight functions at 0 exists.

Consider a very large population of players N and for every $i \in N$, call $\lambda_i(\cdot)$ the function that outputs the sampling weight of player i , given her demand. There are three cases: (i) the limit $\lim_{w \rightarrow 0} \lambda_i(w)$ is a positive constant for at least $|N|/3$ identities i (not necessarily the same positive constant for all), (ii) the same limit is $+\infty$ for at least $|N|/3$ identities, or (iii) the same limit is 0 for at least $|N|/3$ identities. Fix a very large set N and admissible cost-sharing method Ξ . If Ξ falls under case (i) we call it *balanced with respect to N* , if it falls under case (ii) we call it *small-demand-punishing with respect to N* , and if it falls under case (iii) we call it *small-demand-favoring with respect to N* . We then get the following:

THEOREM 3 (Optimal admissible cost-sharing method for polynomials). *Let Ξ be any admissible cost-sharing method:*

1. *If Ξ is balanced with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq \chi_d^d \approx (0.9 \cdot d)^d$.*
2. *If Ξ is small-demand-punishing with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq d^d$.*
3. *If Ξ is small-demand-favoring with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq (2^{1/d} - 1)^{-d} \approx (1.4 \cdot d)^d$.*

It follows that the PoA of χ_d^d achieved by the Shapley value is optimal.

We now generalize this result to get our main theorem:

THEOREM 4 (Optimal admissible cost-sharing method for convex functions). *Let Ξ^{SV} denote the Shapley value and Ξ any admissible cost-sharing method. Also, let \mathcal{C} denote any given set of positive, increasing, and convex cost functions which assign 0 cost to 0 demand, which has the property that if $C(x) \in \mathcal{C}$, then $a \cdot C(b \cdot x) \in \mathcal{C}$ for any positive a, b . Then, $\text{PoA}(\mathcal{C}, \Xi^{\text{SV}}) \leq \text{PoA}(\mathcal{C}, \Xi)$.*

5. Discussion and Open Questions

This paper studies the interactions among self-interested agents who place demands on shared resources which incur a congestion-dependent cost. In this setting, and when the cost functions are convex, we make the case that distributing the generated cost among the participants using the Shapley value leads to the best price of anarchy among all cost-sharing methods possessing certain desirable properties: budget-balance, stability, and locality. It is worth noting that the upper bounds on the price of anarchy make use of the (λ, μ) -smoothness framework of Roughgarden (2009), which implies that these bounds are *robust*, i.e., they hold even for equilibrium concepts more general than the pure Nash equilibrium (e.g., the mixed Nash equilibrium) which are guaranteed to exist. As a result, the optimality of the Shapley value holds much more generally.

Our prescribed properties of budget-balance, stability, and locality are important in our motivating applications. However, understanding the design space beyond such admissible cost-sharing methods is an interesting question. For example, various ways in which the locality property can be relaxed are discussed in von Falkenhausen and Harks (2013). The parameterization of the PoA there is different, the authors do not restrict the set of allowable cost functions but the number of players. It is an interesting open question to study this problem from our perspective and identify the PoA one may achieve if the specific game instance is available to the policy maker.

Violating the stability requirement implies dropping the existence of a pure Nash equilibrium and it may, hence, induce unappealing gaming behavior. Nevertheless, we briefly addressed this question in a recently announced extended abstract (Gkatzelis et al. (2014)). In particular, for polynomial

cost functions of maximum degree d , we prove that the proportional sharing method, which is known to violate the stability property, achieves a robust price of anarchy that is essentially optimal among all cost-sharing methods that satisfy the other two properties. This PoA of approximately $(1.4 \cdot d / \ln d)^d$ compares favorably to the $(0.9 \cdot d)^d$ of the Shapley value. However, proportional sharing has been shown to perform poorly with slight modifications of the model of the players' behavior. This is exhibited by metrics such as the *price of sinking* (Goemans et al. (2005)). On the other hand, methods that induce a potential game (such as the weighted Shapley values) do not suffer from such problems (e.g., the price of sinking is at most the PoA for these methods).

Maybe the most interesting extension would be to relax the budget-balance requirement. Gopalakrishnan et al. (2014) show that, as long as the stability property is enforced, the cost-sharing method will necessarily distribute a total cost according to some weighted Shapley value, although the cost being distributed need not be the same as the generated cost anymore. The cost-sharing method within this class that is arguably the most natural and well studied one is the marginal contribution, according to which each player suffers a cost equal to her marginal contribution in the total cost. For convex cost functions, like the ones that we considered in this paper, this method charges the players a total cost that is higher than the cost that their use of the resources generated. In Gkatzelis et al. (2014) we also show that, for polynomial cost functions of maximum degree d , the marginal contribution method leads to PoA worse than that of any weighted Shapley value. Moreover, there are further indications that the Shapley value remains optimal in this relaxed setting, since our lower bounds apply to most reasonable approaches for defining a cost-sharing method that overcharges players, such as when at least one of the following is true:

1. The modified cost functions remain convex and dominate the actual cost functions.
2. A player that uses a resource by herself is not overcharged.
3. The cost for larger demand is inflated more than the cost for smaller demand.
4. The cost inflation is higher, as the number of players on a resource increases.

Finally, other interesting open problems are to identify the optimal cost-sharing methods in games that are symmetric, i.e., all the players have the same strategy set and in games where the policy maker has a budget which she can use to subsidize player costs.

References

- Abed, Fidaa, Chien-Chung Huang. 2012. Preemptive coordination mechanisms for unrelated machines. *European Symposium on Algorithms*. Springer, 12–23.
- Anshelevich, Elliot, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, Tim Roughgarden. 2008. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing* **38**(4) 1602–1623.
- Awerbuch, Baruch, Yossi Azar, Amir Epstein. 2005. The price of routing unsplittable flow. *ACM Symposium on Theory of Computing*. ACM, 57–66.
- Azar, Yossi, Lisa Fleischer, Kamal Jain, Vahab S. Mirrokni, Zoya Svitkina. 2015. Optimal coordination mechanisms for unrelated machine scheduling. *Operations Research* **63**(3) 489–500.
- Bhattacharya, Sayan, Sungjin Im, Janardhan Kulkarni, Kamesh Munagala. 2014. Coordination mechanisms from (almost) all scheduling policies. *Proceedings of the 5th conference on Innovations in theoretical computer science*. ACM, 121–134.
- Bhawalkar, Kshipra, Martin Gairing, Tim Roughgarden. 2014. Weighted congestion games: Price of anarchy, universal worst-case examples, and tightness. *ACM Transactions on Economics and Computation* **2**(4) 14.
- Bhawalkar, Kshipra, Tim Roughgarden. 2012. Simultaneous single-item auctions. *Internet and Network Economics*. Springer, 337–349.
- Caragiannis, Ioannis. 2013. Efficient coordination mechanisms for unrelated machine scheduling. *Algorithmica* **66**(3) 512–540.
- Chen, Ho-Lin, Tim Roughgarden, Gregory Valiant. 2010. Designing network protocols for good equilibria. *SIAM Journal on Computing* **39**(5) 1799–1832.
- Christodoulou, George, Elias Koutsoupias, Akash Nanavati. 2009. Coordination mechanisms. *Theoretical Computer Science* **410**(36) 3327–3336.

- Christodoulou, Giorgos, Kurt Mehlhorn, Evangelia Pyrga. 2014. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica* **69**(3) 619–640.
- Coffman Jr., Edward G., Isi Mitrani. 1980. A characterization of waiting time performance realizable by single-server queues. *Operations Research* **28**(3) 810–821.
- Cole, Richard, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, Neil Olver. 2015. Decentralized utilitarian mechanisms for scheduling games. *Games and Economic Behavior* **92** 306–326.
- Cominetti, Roberto, José R. Correa, Nicolás E. Stier Moses. 2009. The impact of oligopolistic competition in networks. *Operations Research* **57**(6) 1421–1437.
- Fotakis, Dimitris, Paul G. Spirakis. 2008. Cost-balancing tolls for atomic network congestion games. *Internet Mathematics* **5**(4) 343–363.
- Gairing, Martin, Florian Schoppmann. 2007. Total latency in singleton congestion games. *Internet and Network Economics*. Springer, 381–387.
- Gkatzelis, Vasilis, Konstantinos Kollias, Tim Roughgarden. 2014. Optimal cost-sharing in weighted congestion games. *Conference on Web and Internet Economics*. Springer, 72–88.
- Goemans, Michel, Vahab S. Mirrokni, Adrian Vetta. 2005. Sink equilibria and convergence. *IEEE Symposium on Foundations of Computer Science*. IEEE, 142–151.
- Gopalakrishnan, Ragavendran, Jason R. Marden, Adam Wierman. 2014. Potential games are necessary to ensure pure nash equilibria in cost sharing games. *Mathematics of Operations Research* .
- Harks, Tobias, Max Klimm. 2012. On the existence of pure nash equilibria in weighted congestion games. *Mathematics of Operations Research* **37**(3) 419–436.
- Harks, Tobias, Max Klimm, Rolf H. Möhring. 2011. Characterizing the existence of potential functions in weighted congestion games. *Theory of Computing Systems* **49**(1) 46–70.
- Harks, Tobias, Konstantin Miller. 2011. The worst-case efficiency of cost sharing methods in resource allocation games. *Operations Research* **59**(6) 1491–1503.
- Haviv, Moshe, Tim Roughgarden. 2007. The price of anarchy in an exponential multi-server. *Operations Research Letters* **35**(4) 421–426.

- Immorlica, Nicole, Li Erran Li, Vahab S. Mirrokni, Andreas S. Schulz. 2009. Coordination mechanisms for selfish scheduling. *Theoretical Computer Science* **410**(17) 1589–1598.
- Kollias, Konstantinos. 2013. Nonpreemptive coordination mechanisms for identical machines. *Theory of Computing Systems* **53**(3) 424–440.
- Kollias, Konstantinos, Tim Roughgarden. 2015. Restoring pure equilibria to weighted congestion games. *ACM Transactions on Economics and Computation* **3**(4) 13–46.
- Lucier, Brendan, Allan Borodin. 2010. Price of anarchy for greedy auctions. *Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 537–553.
- Marden, Jason R., Adam Wierman. 2013. Distributed welfare games. *Operations Research* **61**(1) 155–168.
- Milchtaich, Igal. 1996. Congestion games with player-specific payoff functions. *Games and Economic Behavior* **13**(1) 111–124.
- Monderer, Dov, Lloyd S. Shapley. 1996. Potential games. *Games and Economic Behavior* **14**(1) 124–143.
- Mosk-Aoyama, Damon, Tim Roughgarden. 2009. Worst-case efficiency analysis of queueing disciplines. *International Colloquium on Automata, Languages and Programming*. Springer, 546–557.
- Moulin, Hervé. 2008. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory* **36**(3) 379–405.
- Moulin, Hervé, Scott J. Shenker. 2001. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory* **18**(3) 511–533.
- Osborne, Martin J., Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT press.
- Rosenthal, Robert W. 1973a. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* **2**(1) 65–67.
- Rosenthal, Robert W. 1973b. The network equilibrium problem in integers. *Networks* **3**(1) 53–59.
- Roughgarden, Tim. 2009. Intrinsic robustness of the price of anarchy. *ACM Symposium on Theory of Computing*. ACM, 513–522.
- Roughgarden, Tim, Éva Tardos. 2002. How bad is selfish routing? *Journal of the ACM (JACM)* **49**(2) 236–259.

Shapley, Lloyd S. 1953. *Additive and non-additive set functions*. Princeton University.

Shenker, Scott J. 1995. Making greed work in networks: A game-theoretic analysis of switch service disciplines. *IEEE/ACM Transactions on Networking (TON)* **3**(6) 819–831.

Syrçkanis, Vasilis, Éva Tardos. 2013. Composable and efficient mechanisms. *ACM Symposium on Theory of Computing*. ACM, 211–220.

von Falkenhausen, Philipp, Tobias Harks. 2013. Optimal cost sharing protocols for scheduling games. *Mathematics of Operations Research* **38**(1) 184–208.

Proofs of Statements

6. Reduction from Generalized Weighted to Weighted Shapley Values

CLAIM 1. *For any generalized weighted Shapley value Ξ , there exists a weighted Shapley value Ξ' that approximates every player's cost shares in Ξ with arbitrary precision.*

Proof. A generalized weighted Shapley value Ξ defines a partition of the game's player set into groups $1, 2, \dots, k$ and sampling weights for the players. To get the sampling weights of the weighted Shapley value Ξ' that approximates Ξ we multiply the sampling weights of Ξ for the players in group i by M^i , where M is an arbitrarily large positive number. As $M \rightarrow \infty$, the orderings not allowed by Ξ have realization probability arbitrarily close to 0 in Ξ' , while the relative sampling probabilities among players in the same group remain intact. \square

CLAIM 2. *If $PoA(\mathcal{C}, \Xi') \geq \beta$ for a given set of cost functions \mathcal{C} and any weighted Shapley value Ξ' , then for any generalized weighted Shapley value Ξ and any $\epsilon > 0$, $PoA(\mathcal{C}, \Xi) \geq \beta - \epsilon$.*

Proof. Consider a generalized weighted Shapley value Ξ and let Ξ' be a weighted Shapley value that approximates it with arbitrary precision, as in Claim 1. Also, consider the game that yields the highest PoA when Ξ' is used and let A be the worst PNE of that game. Using this game we construct a modified game for Ξ that yields approximately the same PoA.

If one switches from Ξ' to Ξ , then A might not remain a PNE in the initial game since the cost shares of the players may change, even though this change would be arbitrarily small. In order to ensure that A remains a PNE we modify the game as follows: for each player i , we create a new resource r^i which we add to all of that player's possible strategies, except for A_i . The cost of this resource is chosen so that $C_{r^i}(w_i) = \delta$, where δ is equal to the largest change in any player's cost due to the switch from Ξ' to Ξ . Note that, by Claim 1, this δ can be made arbitrarily small. Therefore, in the modified game, outcome A remains a PNE, while the cost of the PNE remains unchanged and the cost of the optimal solution may increase only by a negligible amount. Since $PoA(\mathcal{C}, \Xi') \leq \beta$, this implies that $PoA(\mathcal{C}, \Xi) \leq \beta - \epsilon$, where $\epsilon > 0$ depends on the choice of δ and can be arbitrarily small. \square

7. Proof of Lemma 1.

LEMMA 1. *For positive values of γ :*

$$PoA(\mathcal{C}_d, \Xi^\gamma) \geq (2^{1/d} - 1)^{-d} \approx (1.4 \cdot d)^d.$$

Proof. Define $\rho = (2^{1/d} - 1)^{-1}$ and let T be a set of ρ/ϵ players with demand ϵ each, where $\epsilon > 0$ is an arbitrarily small parameter. Consider a player i with demand $w_i = 1$ and suppose she uses a resource r with cost function $C_r(x) = x^d$ with the players in T , while the cost-sharing rule is our Ξ^γ . We now argue that, as we let $\epsilon \rightarrow 0$, the cost share of i in r becomes $(\rho + 1)^d - \rho^d$. Consider the probability p that i is not among the last $\delta \cdot |T|$ players of the random ordering generated by our sampling weights (i.e., i is not among the first $\delta \cdot |T|$ players sampled), for some $\delta < 1$. This probability is upper bounded by the probability that i is not drawn, using our sampling weights, among everyone in $T \cup \{i\}$, $\delta \cdot |T| = \delta \cdot \rho/\epsilon$ times. Note that the sampling weight of i is 1 and the total sampling weight of the players in T is $\rho \cdot \epsilon^{\gamma-1}$. Hence, if $\gamma \geq 1$, we get:

$$p \leq \left(1 - \frac{1}{1 + \rho \cdot \epsilon^{\gamma-1}}\right)^{\delta \cdot \rho/\epsilon} \leq \left(1 - \frac{1}{1 + \rho}\right)^{\delta \cdot \rho/\epsilon}, \quad (17)$$

which goes to 0 as $\epsilon \rightarrow 0$. Similarly, if $\gamma < 1$, we get:

$$p \leq \left(1 - \frac{1}{1 + \rho \cdot \epsilon^{\gamma-1}}\right)^{\delta \cdot \rho/\epsilon} \leq \exp\left(-\delta \cdot \frac{\rho}{\epsilon} \cdot \frac{1}{1 + \rho \cdot \epsilon^{\gamma-1}}\right) = \exp\left(-\delta \cdot \rho \cdot \frac{\epsilon^{-\gamma}}{\epsilon^{1-\gamma} + \rho}\right), \quad (18)$$

which always goes to 0 as $\epsilon \rightarrow 0$, for any arbitrarily small $\delta > 0$. Then, by letting $\delta \rightarrow 0$, our claim that the cost share of i is $(\rho + 1)^d - \rho^d$ follows by (3). Similarly, it follows that if a player with demand w shares a resource with cost function $a \cdot x^d$ with ρ/ϵ players with demand $w \cdot \epsilon$ each, her cost share will be $a \cdot w^d \cdot ((\rho + 1)^d - \rho^d)$ (since scaling the cost function and the player demands does not change the fractions of the cost that are assigned to the players), which, for our choice of ρ is equal to $a \cdot w^d \cdot \rho^d$.

Using facts from the previous paragraph as building blocks, we construct a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi^\gamma)$, such that the total cost in the worst equilibrium is ρ^d times the optimal. Suppose our resources are organized in a tree graph $G = (R, E)$, where each vertex corresponds

to a resource. There is a one-to-one mapping between the set of edges of the tree, E , and the set of players of the game, N . The player i , that corresponds to edge $e = (r, r')$, has strategy set $\mathcal{A}_i = \{\{r\}, \{r'\}\}$, i.e., she must choose one of the two endpoints of her designated edge. Tree G has branching factor ρ/ϵ and l levels, with the root positioned at level 1 and the leaves positioned at level l .

Player demands. The demand of every player (edge) between resources (vertices) at levels j and $j + 1$ of the tree is ϵ^{j-1} .

Cost functions. The cost function of any resource (vertex) at level $j = 1, 2, \dots, l - 1$, is:

$$C^j(x) = \left(\frac{1}{\rho \cdot \epsilon^{d-1}} \right)^{j-1} \cdot x^d. \quad (19)$$

The cost functions of any resource (vertex) at level l is equal to:

$$C^l(x) = \frac{\rho^{d-l+1}}{\epsilon^{(d-1) \cdot (l-1)}} \cdot x^d. \quad (20)$$

PNE. Let A be the outcome that has all players play the resource closer to the root. We claim that this outcome is a PNE. The cost of every player, using a resource at level $j < l$, in A , is $(\rho/\epsilon)^{d-j}$. If one of the players that are adjacent to the leaves were to switch to her other strategy (play the leaf resource), she would incur a cost equal to $(\rho/\epsilon)^{d-l+1}$, which is the same as the one she has in A . Consider any other player and her potential deviation from the resource at level j , to the resource at level $j + 1$. By the analysis in the first paragraph of this proof (she would be a player with demand ϵ^{j-1} sharing a resource with ρ/ϵ players with demand ϵ^j), her cost would be $(\rho \cdot \epsilon^{d-1})^j \cdot \epsilon^{d \cdot (j-1)} \cdot \rho^d = (\rho/\epsilon)^{d-j}$, which is her current cost in A . This proves that the equilibrium condition (1) holds for all players in A .

PoA. As we have shown, every player using a resource at level j has cost $(\rho/\epsilon)^{d-j}$ in A . There are $(\rho/\epsilon)^j$ such players, which implies the total cost of A is $Q(A) = (l-1) \cdot (\rho/\epsilon)^d$, since there are $l-1$ levels of nonempty resources, and every level has the same total cost, $(\rho/\epsilon)^d$. Now, let A^* be the outcome that has all players play the resource further from the root. In this outcome, every player using a resource at level $j = 2, \dots, l-1$, has cost $\rho^{-j+1}/\epsilon^{d-j+1}$. There are $(\rho/\epsilon)^{j-1}$ such players, hence, the total cost at level j is $(1/\epsilon)^d$. Similarly, we get that the total cost at level l is $(\rho/\epsilon)^d$. In total, $Q(A^*) = (l-2) \cdot (1/\epsilon)^d + (\rho/\epsilon)^d$. We can then see that, $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = \rho^d$. \square

8. Proof of Lemma 2.

LEMMA 2. *For negative values of γ :*

$$PoA(\mathcal{C}_d, \Xi^\gamma) \geq d^d$$

Proof. We construct a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi^\gamma)$, such that the total cost in the worst equilibrium is d^d times the optimal. Suppose our resources are organized in a tree graph $G = (R, E)$, where each vertex corresponds to a resource. There is a one-to-one mapping between the set of edges of the tree, E , and the set of players of the game, N . The player i , that corresponds to edge $e = (r, r')$, has strategy set $\mathcal{A}_i = \{\{r\}, \{r'\}\}$, i.e., she must choose one of the two endpoints of her designated edge. Tree G has branching factor $1/(d \cdot \epsilon)$, with $\epsilon > 0$ an arbitrarily small parameter, and l levels. The root is positioned at level 1 and the leaves are positioned at level l .

Player demands. The demand of every player (edge) between resources (vertices) at levels j and $j + 1$ of the tree is ϵ^{j-1} .

Cost functions. The cost function of any resource (vertex) at level $j = 2, 3, \dots, l$, is:

$$C^j(x) = \left(\frac{d}{\epsilon^{d-1}} \right)^{j-2} \cdot x^d. \quad (21)$$

The cost function of the root is:

$$C^1(x) = x^d. \quad (22)$$

PNE. Let A be the outcome that has all players play the resource further from the root. We prove that this outcome is a PNE. The cost of every player that has played a resource at level j is $(d \cdot \epsilon)^{j-2}$. If one of the players that are adjacent to the root were to switch to her other strategy (play the root), she would incur a cost equal to 1, which is the same as the one she has in A . Consider any other player and her potential deviation from the resource at level j , to the resource at level $j - 1$. Since our construction considers ϵ arbitrarily close to 0, the deviating player will go last with probability 1 in the Shapley ordering (since $\gamma < 0$ for our Ξ^γ) and her cost will, by (3), be equal to $(d/\epsilon^{d-1})^{j-3} \cdot ((\epsilon^{j-3} + \epsilon^{j-2})^d - \epsilon^{(j-3) \cdot d}) = (d \cdot \epsilon)^{j-2}$, which is equal to her current cost in A . Hence, the equilibrium condition (1) holds for all players.

PoA. As we have shown, every player playing a resource at level j has cost $(d \cdot \epsilon)^{j-2}$ in A . There are $1/(d \cdot \epsilon)^{j-1}$ such players, hence, the total cost at level j is $1/(d \cdot \epsilon)$. Then, it follows that $Q(A) = (l-1)/(d \cdot \epsilon)$. Now let A^* be the outcome that has all players play the resource closer to the root. Then the joint cost at the root is $1/(d \cdot \epsilon)^d$. The joint cost of every other resource at level j is $(d \cdot \epsilon)^{j-2}/d^d$, and the number of resources at level j is $1/(d \cdot \epsilon)^{j-1}$. Hence, we get in total, $Q(A^*) = (l-2)/(d^{d+1} \cdot \epsilon) + 1/(d \cdot \epsilon)^d$. We can then see that $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = d^d$. \square

9. Statement and Proof of Lemma 3.

LEMMA 3 (**Improved lower bound for polynomials and $\gamma < 0$**). *Consider Ξ^γ with $\gamma < 0$. We write,*

$$D_{\gamma,d}(x) = \sup_{T:l(T)=1} \sum_{i \in T} \frac{w_i^\gamma}{x^\gamma + w_i^\gamma} \cdot ((x + w_i)^d - x^d) + \frac{x^\gamma}{x^\gamma + w_i^\gamma} \cdot w_i^d, \quad (23)$$

for the worst-case sum of cost-shares of players with total demand 1, when each one of them uses a resource with cost function x^d with another player with demand x . Then, $PoA(C_d, \Xi^\gamma) \geq \psi_{\gamma,d}^d$, where $\psi_{\gamma,d}$ is the largest root of $D_{\gamma,d}(x) - x^d$.

Proof. Let T be a set of players that satisfies $l(T) = 1$ and:

$$\sum_{i \in T} \frac{w_i^\gamma}{\psi_{\gamma,d}^\gamma + w_i^\gamma} \cdot ((\psi_{\gamma,d} + w_i)^d - \psi_{\gamma,d}^d) + \frac{\psi_{\gamma,d}^\gamma}{\psi_{\gamma,d}^\gamma + w_i^\gamma} \cdot w_i^d = D_{\gamma,d}(\psi_{\gamma,d}), \quad (24)$$

i.e, it achieves the worst-case sum of deviations into a resource with cost function x^d , when the resource is used by a player with demand $\psi_{\gamma,d}$. We construct a game $(N, R, \mathcal{A}, w, C_d, \Xi^\gamma)$ with the following specifics. The resources in R are organized in a tree graph $G = (R, E)$, such that each vertex corresponds to a resource. There is a one-to-one mapping between the edges in E and the players in N , with the interpretation that the player i that is mapped to edge (r, r') has strategy set $\mathcal{A}_i = \{\{r\}, \{r'\}\}$. The tree has l levels (level 1 is the root and level l are the leaves) and the branching parameter is equal to $|T|$, with the exception of the root that only has one child.

Player demands. Index the players in T as $t_1, t_2, \dots, t_{|T|}$. The demand of the player (edge) between the root and its i -th child is w_{t_i} . We define the remaining demands recursively. Let w be the demand of a player between a resource at level j and a resource r at level $j+1$, with $j < l-1$. Then the demand of the player between r and the i -th child of r is $w \cdot w_{t_i} / \psi_{\gamma,d}$.

Cost functions. The cost function at the root is x^d . The same holds for the root's child. We define the remaining cost functions recursively. Suppose a non-leaf resource r has cost function $C_r(x)$. Then, its i -th child, r^i , has cost function:

$$C_{r^i}(x) = \frac{((w_i + 1)^d - 1) \cdot w_i^\gamma + 1}{w_i^\gamma + 1} \cdot \frac{1}{\psi_{\gamma,d}^d} \cdot C_r(x). \quad (25)$$

PNE. Let A be the outcome that has all players play the resource further from the root. We claim that this outcome is a PNE. The equilibrium condition (1) trivially holds for the player that plays the child of the root (instead of the root). Consider any other player, who plays the i -th child, r^i , of a resource r , instead of r . This player has demand $w \cdot w_i / \psi_{\gamma,d}$, with w the demand of the player that uses resource r in A . Let $C_r(x) = a \cdot x^d$. Then, by (25), it follows that:

$$C_{r^i}(x) = \frac{((w_i + 1)^d - 1) \cdot w_i^\gamma + 1}{w_i^\gamma + 1} \cdot \frac{1}{\psi_{\gamma,d}^d} \cdot a \cdot x^d. \quad (26)$$

The cost of the player under consideration in A is $C_{r^i}(w \cdot w_i / \psi_{\gamma,d})$. By (3), we get that her cost, if she were to deviate to resource r , would be:

$$\frac{(w \cdot w_i / \psi_{\gamma,d})^\gamma \cdot (C_r(w + w \cdot w_i / \psi_{\gamma,d}) - C_r(w)) + w^\gamma \cdot C_r(w \cdot w_i / \psi_{\gamma,d})}{(w \cdot w_i / \psi_{\gamma,d})^\gamma + w^\gamma} = C_{r^i}(w \cdot w_i / \psi_{\gamma,d}), \quad (27)$$

which proves that the equilibrium condition (1) holds for all players.

PoA. Consider a resource r with cost function $a \cdot x^d$, that is neither the root nor a leaf. By (27), we get that the players using the children of r have a cost equal to their potential deviations to r . Let w be the demand of the user of r . By a straightforward extension of (25), we get that the total costs of the players using the children of r is $a \cdot (w / \psi_{\gamma,d})^d \cdot D_{\gamma,d}(\psi_{\gamma,d})$, which due to the property of $\psi_{\gamma,d}$ becomes $a \cdot (w / \psi_{\gamma,d})^d \cdot \psi_{\gamma,d}^d = a \cdot w^d$, which is the same as the cost of the user of r in A . Let A^* be the outcome that has all players play the resource closer to the root. Then, the total cost of the players in r (which are the same as the players that use the children of r in A) would be $a \cdot (w / \psi_{\gamma,d})^d$. Hence, for every set of players that use sibling resources, we get that their total cost in A is $\psi_{\gamma,d}^d$ times their total cost in A^* . The only player that remains to be examined is the player adjacent to the root, who has the same cost, 1, in both outcomes. Note however, that,

as we showed earlier in this paragraph, the cost on a resource r in A is equal to the sum of the costs on its children, which implies the cost across levels is the same. We, then, get $Q(A) = l - 1$ and $Q(A^*) = 1 + (l - 2)/\psi_{\gamma,d}^d$, which means $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = \psi_{\gamma,d}^d$. \square

10. Statement and Proof of Lemma 4.

LEMMA 4 (**Improved upper bound for polynomials and $\gamma < 0$**). *Consider Ξ^γ with $\gamma < 0$.*

We write,

$$D_{\gamma,k}(x) = \sup_{T:l(T)=1} \sum_{i \in T} \frac{w_i^\gamma}{x^\gamma + w_i^\gamma} \cdot ((x + w_i)^k - x^k) + \frac{x^\gamma}{x^\gamma + w_i^\gamma} \cdot w_i^k, \quad (28)$$

for the worst-case sum of cost-shares of players with total demand 1, when each one of them uses a resource with cost function x^k with another player with demand x . Then, $PoA(C_d, \Xi^\gamma) \leq \omega_{\gamma,d}$, where $\omega_{\gamma,d}$ is the smallest number, among the ones that satisfy $D_{\gamma,k}(x) \leq x^k$ for all k , for which there exists some $\hat{\mu} \in (0, 1)$, such that, $\arg \max_{x>0} \max_{k \in [1,d]} D_{\gamma,k}(x) - \hat{\mu} \cdot x^k = \omega_{\gamma,d}$.

Proof. Consider a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi)$ and suppose A is a PNE of the game and A^* the optimal outcome. We get

$$Q(A) = \sum_{r \in R} \sum_{i \in N} \xi_{i,r}(A) = \sum_{i \in N} \sum_{r \in A_i} \xi_{i,r}(A) \leq \sum_{i \in N} \sum_{r \in A_i^*} \xi_{i,r}(A_i^*, A_{-i}) = \sum_{r \in R} \sum_{i \in S_r(A^*)} \xi_{i,r}(A_i^*, A_{-i}). \quad (29)$$

At this point we wish to identify appropriate $\hat{\lambda}$ and $\hat{\mu}$, such that the last expression in (29) is upper bounded by $\hat{\lambda} \cdot Q(A^*) + \hat{\mu} \cdot Q(A)$, which would yield a $\hat{\lambda}/(1 - \hat{\mu})$ upper bound on the PoA. To do so, we ask the stronger fact that $\hat{\lambda}$ and $\hat{\mu}$ satisfy this for every resource (i.e., every degree $k \in [1, d]$) and every possible set of users of the resource and deviating players (i.e., users of the resource in A^*). That is, we ask that the sum of deviation costs, of the players in an arbitrary set of players T^* , into a resource with cost function x^k , which is being used by an arbitrary set of players T , is always upper bounded by $\hat{\lambda}$ times $l(T^*)^k$ plus $\hat{\mu}$ times $l(T)^k$. We observe that the worst-case for the deviating players (hence, the worst-case for our upper bound) is when T is a single player, because $\gamma < 0$ (there is higher probability that the demand already in the resource is earlier than the deviating player in the Shapley ordering). In short, if $\hat{\lambda}$ and $\hat{\mu}$ satisfy, for every $k \in [1, d]$, $x, y > 0$, and every T^* such that $l(T^*) = x$:

$$\sum_{i \in T^*} \frac{w_i^\gamma}{w_i^\gamma + x^\gamma} \cdot ((w_i + x)^k - x^k) + \frac{x^\gamma}{w_i^\gamma + x^\gamma} \cdot w_i^k \leq \hat{\lambda} \cdot y^k + \hat{\mu} \cdot x^k, \quad (30)$$

then, combining with (29), we get a $\hat{\lambda}/(1 - \hat{\mu})$ upper bound on the PoA. We can see that scaling the player demands does not impact the constraint (30). Hence, by considering x and 1, as opposed to x and y , and by considering only the worst-case T^* , we can rewrite (30) as:

$$D_{\gamma,k}(x) \leq \hat{\lambda} + \hat{\mu} \cdot x^k. \quad (31)$$

We choose $\hat{\mu}$ to be the value that lets $\omega_{\gamma,d}$ be used as the worst-case value for x in (31). Then we choose $\hat{\lambda}$ such as to make this worst-case for (31) tight. Then, from (31), we get:

$$D_{\gamma,k}(\omega_{\gamma,d}) = \hat{\lambda} + \hat{\mu} \cdot \omega_{\gamma,d}^k \Rightarrow \frac{\hat{\lambda}}{1 - \hat{\mu}} = \frac{D_{\gamma,k}(\omega_{\gamma,d}) - \hat{\mu} \cdot \omega_{\gamma,d}^k}{1 - \hat{\mu}} \leq \omega_{\gamma,d}^k \leq \omega_{\gamma,d}^d, \quad (32)$$

where the final inequality follows from the properties of $\omega_{\gamma,d}$. \square

Typically, and for the vast majority of γ, d pairs, this upper bound matches the lower bound from Lemma 3. It is possible that there exist rare pathological cases, for which the two bounds do not match, however, even finding one such γ, d pair is a difficult task.

11. Proof of Theorem 3.

THEOREM 3. *Let Ξ be any admissible cost-sharing method and N an arbitrarily large set of players:*

1. *If Ξ is balanced with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq \chi_d^d \approx (0.9 \cdot d)^d$.*
2. *If Ξ is small-demand-punishing with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq d^d$.*
3. *If Ξ is small-demand-favoring with respect to N , then $PoA(\mathcal{C}_d, \Xi) \geq (2^{1/d} - 1)^{-d} \approx (1.4 \cdot d)^d$.*

It then follows that the PoA of χ_d^d achieved by the Shapley value is optimal.

We prove the theorem via the following sequence of lemmas.

LEMMA 5. *Let Ξ be a weighted Shapley value that is balanced with respect to N , an arbitrarily large set of players. With \mathcal{C}_d the set of polynomials with nonnegative coefficients and degrees between 1 and d , we get that $PoA(\mathcal{C}_d, \Xi) \geq \chi_d^d$.*

Proof. We will construct a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi)$ such that the worst equilibrium cost is χ_d^d times the optimal cost. Focus on the subset of players whose designated sampling weight function has a positive and finite limit at 0. In this subset, find a sequence of l players, i_1, i_2, \dots, i_l , such

that, if the limit at 0 of the sampling weight function for player i is λ_i^0 , then $\lambda_{i_1}^0 \geq \lambda_{i_2}^0 \geq \dots \geq \lambda_{i_l}^0$. Such a sequence trivially exists. Now, suppose ϵ is a very small number, such that the value of the sampling weight function of player i for any demand at most ϵ is arbitrarily close to λ_i^0 . By continuity, such an ϵ exists. The resources are arranged on a line as r_1, r_2, \dots, r_{l+1} , and there is an extra resource r_0 , which does not belong on the line. Each Player i_j of the sequence we identified at the start of the proof has strategy set $\mathcal{A}_{i_j} = \{\{r_j\}, \{r_{j+1}\}\}$, i.e., she must pick one of r_j and r_{j+1} . All other players in N have only one possible strategy, which is to play r_0 .

Player demands. The demand of Player i_j , for $j = 1, 2, \dots, l$, is ϵ/χ_d^j . The demand of any other player in N is 0.

Cost functions. The cost function of r_1 is x^d and the cost function of r_j , for $j = 2, \dots, l$, is $\chi_d^{d \cdot (j-1)} \cdot x^d$. The cost function of r_0 is x^d .

PNE. Let A be the outcome where Player i_j plays resource r_{j+1} , for $j = 1, 2, \dots, l$. The equilibrium condition trivially holds for all players using r_0 (since they have no other option) and for Player i_1 (since r_1 and r_2 have the same cost). Every other player has cost ϵ^d in A . The cost of her potential deviation is at least her Shapley value on the other resource (by the choice of the sequence of players i_1, i_2, \dots, i_l in the beginning), which is also ϵ^d .

PoA. Let A^* be the outcome that has Player i_j play resource j , for $j = 1, 2, \dots, l$. We observe that $Q(A) = l \cdot (\epsilon/\chi_d)^d$ and $Q(A^*) = (\epsilon/\chi_d)^d + (l-1) \cdot (\epsilon/\chi_d^2)^d$. Then $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = \chi_d^d$. Note that we can let l grow to infinity, since Ξ is balanced with respect to N and N is arbitrarily large. \square

LEMMA 6. *Let Ξ be a weighted Shapley value that is small-demand-punishing with respect to N , an arbitrarily large set of players. With \mathcal{C}_d the set of polynomials with nonnegative coefficients and degrees between 1 and d , we get that $PoA(\mathcal{C}_d, \Xi) \geq d^d$.*

Proof. We construct a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi)$ such that the cost of the worst equilibrium is d^d times the cost of the optimal outcome. The resources in R are organized in a tree graph $G = (R \setminus \{r_0\}, E)$, where each vertex corresponds to a resource, with the exception of a single resource r_0 ,

which does not belong to the tree. The tree has l levels and the branching factor is $1/(d \cdot \epsilon)$, where ϵ is an arbitrarily small parameter, with the exception of the root, which has only one child. Each edge $(r, r') \in E$ corresponds to a distinct Player $i \in N$, who has strategy set $\mathcal{A}_i = \{\{r\}, \{r'\}\}$. Every player from N that is not mapped to an edge, has playing r_0 as her only available strategy.

Player placement and demands. Focus only on players such that their designated sampling weight functions approach infinity as their demands approach 0. Since Ξ is small-demand-punishing with respect to N and N is arbitrarily large, there is an infinite supply of such players. Fill the edges of G with such players. As already stated, all other players will have r_0 as their only strategy. All players placed in r_0 have demand 0. A player between levels j and $j+1$ of the tree has demand ϵ^{j-1} .

Cost functions. The cost function of any resource at level $j = 2, 3, \dots, l$, is $(d/\epsilon^{d-1})^{j-2} \cdot x^d$. The cost functions of the root is x^d and the cost function of r_0 is also x^d .

PNE. Let A be the outcome that has all players play the resource further from the root. We prove that this outcome is a PNE. The cost of every player that has played a resource at level j is $(d \cdot \epsilon)^{j-2}$. If one of the players that are adjacent to the root were to switch to her other strategy (play the root), she would incur a cost equal to 1, which is the same as the one she has in A . Consider any other player and her potential deviation from the resource at level j , to the resource at level $j-1$. Since our construction considers ϵ arbitrarily close to 0, the deviating player will go last with probability 1 in the Shapley ordering (since $\gamma < 0$ for our Ξ^γ) and her cost will, by (3), be equal to $(d/\epsilon^{d-1})^{j-3} \cdot ((\epsilon^{j-3} + \epsilon^{j-2})^d - \epsilon^{(j-3) \cdot d}) = (d \cdot \epsilon)^{j-2}$, which is equal to her current cost in A . Hence, the equilibrium condition (1) holds for all players.

PoA. As we have shown, every player playing a resource at level j has cost $(d \cdot \epsilon)^{j-2}$ in A . There are $1/(d \cdot \epsilon)^{j-1}$ such players, hence, the total cost at level j is $1/(d \cdot \epsilon)$. Then, it follows that $Q(A) = (l-1)/(d \cdot \epsilon)$. Now let A^* be the outcome that has all players play the resource closer to the root. Then the joint cost at the root is $1/(d \cdot \epsilon)^d$. The joint cost of every other resource at level j is $(d \cdot \epsilon)^{j-2}/d^d$, and the number of resources at level j is $1/(d \cdot \epsilon)^{j-1}$. Hence, we get in total, $Q(A^*) = (l-2)/(d^{d+1} \cdot \epsilon) + 1/(d \cdot \epsilon)^d$. We can then see that $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = d^d$. \square

LEMMA 7. Let Ξ be a weighted Shapley value that is small-demand-favoring with respect to N , an arbitrarily large set of players. With \mathcal{C}_d the set of polynomials with nonnegative coefficients and degrees between 1 and d , we get that $PoA(\mathcal{C}_d, \Xi) \geq (2^{1/d} - 1)^{-d}$.

Proof. Let $\rho = (2^{1/d} - 1)^{-1}$. We construct a game $(N, R, \mathcal{A}, w, \mathcal{C}_d, \Xi)$ such that the cost of the worst equilibrium is ρ^d times the cost of the optimal outcome. The resources are organized in a tree graph $G = (R \setminus \{r_0\}, E)$, which has a vertex for each resource, with the exception of a single resource r_0 , which does not belong to the tree. The tree has l levels and the branching factor is ρ/ϵ , where ϵ is an arbitrarily small parameter. Each edge $(r, r') \in E$ corresponds to a distinct Player $i \in N$, who has strategy set $\mathcal{A}_i = \{\{r\}, \{r'\}\}$. Every player from N that is not mapped to an edge, has playing r_0 as her only available strategy.

Player placement and demands. Focus only on players such that their designated sampling weight functions goes to 0 as their demands approach 0. Since Ξ is small-demand-favoring with respect to N and N is arbitrarily large, there is an infinite supply of such players. Fill the edges of G with such players. All other players will have r_0 as their only strategy. All players placed in r_0 have demand 0. A player between levels j and $j + 1$ of the tree has demand ϵ^{j-1} .

Cost functions. The cost function at the root is x^d . For levels $2, 3, \dots, l - 1$, we define the cost functions recursively. Suppose the cost function of a resource r at level $j < l - 2$ is $C_r(x)$ and consider the player i that is between the resource and its child r' . Let A be the outcome that has all players on the tree play the resource closer to the root. Suppose the fraction of the joint cost of r in A that is covered by player i is α_r^i . Then the cost function of resource r' is $C_{r'}(x) = \alpha_r^i \cdot C_r(x)/\epsilon^d$. Finally, if the cost function of a resource r at level $l - 1$ is $C_r(x)$, and the cost fraction that is assigned to the player i between r and its child r' in A is α_r^i , then the cost function of r' is $C_{r'}(x) = \alpha_r^i \cdot (\rho/\epsilon)^d \cdot C_r(x)$.

PNE. Consider outcome A as in the previous paragraph. We will show that it is a PNE. Consider a player i adjacent to a leaf. Let r be the resource that the player is using in A . Then her cost is $\alpha_r^i \cdot C_r(w_i \cdot \rho/\epsilon)$. If she were to deviate to the leaf resource, her cost would be $\alpha_r^i \cdot (\rho/\epsilon)^d \cdot C_r(w_i)$.

We can easily verify that the two quantities are equal. Now consider any other player i that is using a resource r at level $j < l - 1$. The cost of i in A is $a_r^i \cdot C_r(w_i \cdot \rho/\epsilon)$. If i deviates to the child of r , her cost will be $a_r^i \cdot (C_r(w_i \cdot (1 + \rho)) - C_r(w_i \cdot \rho))/\epsilon^d$ (since, as we explained in Lemma 1, with such sampling weights, the player's cost share is as if she goes last in the Shapley ordering with probability 1), which, by the fact that for our choice of ρ , $(1 + \rho)^d - \rho^d = \rho^d$, is equal to her current cost in A . Hence, the equilibrium condition holds for all players.

PoA. The total cost of A is $Q(A) = (l - 1) \cdot (\rho/\epsilon)^d$, since there are $l - 1$ levels of nonempty resources, and every level has the same total cost, $(\rho/\epsilon)^d$, by the fact that $\sum_{i \in S_r(A)} \alpha_r^i = 1$, for every r . Now, let A^* be the outcome that has all players play the resource further from the root. In this outcome, the total cost at level j is $(1/\epsilon)^d$. Similarly, we get that the total cost at level l is $(\rho/\epsilon)^d$. In total, $Q(A^*) = (l - 2) \cdot (1/\epsilon)^d + (\rho/\epsilon)^d$. We can then see that, $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = \rho^d$. \square

12. Proof of Theorem 4.

THEOREM 4. *Let Ξ^{SV} denote the Shapley value and Ξ any admissible cost-sharing method. Also, let \mathcal{C} denote any given set of positive, increasing, and convex cost functions which assign 0 cost to 0 demand, which has the property that if $C(x) \in \mathcal{C}$, then $a \cdot C(b \cdot x) \in \mathcal{C}$ for any positive a, b . Then, $PoA(\mathcal{C}, \Xi^{\text{SV}}) \leq PoA(\mathcal{C}, \Xi)$.*

In this section, we will be using the following fact from Kollias and Roughgarden (2015):

PROPOSITION 1 (Kollias and Roughgarden (2015)). *Let $PoA(\mathcal{C}, \Xi^{\text{SV}})$ denote the PoA of the Shapley value for a given set of positive, increasing, and convex cost functions, which is closed under dilation. Then, there exist $C_1, C_2 \in \mathcal{C}$, $x_1, x_2 > 1$, and $\eta \in [0, 1]$, such that*

$$\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2) = \eta \cdot \hat{\xi}_{C_1}^{\text{SV}}(1, x_1) + (1 - \eta) \cdot \hat{\xi}_{C_2}^{\text{SV}}(1, x_2) = PoA(\mathcal{C}, \Xi^{\text{SV}}), \quad (33)$$

where $\hat{\xi}_C^{\text{SV}}(1, x)$ denotes the Shapley value of a player with demand 1 that shares a resource with cost function C with a player with demand x .

We now prove Theorem 4 via the following lemmas.

LEMMA 8. *Let Ξ be a weighted Shapley value that is balanced with respect to N , an arbitrarily large set of players. Then $PoA(\mathcal{C}, \Xi) \geq PoA(\mathcal{C}, \Xi^{SV})$.*

Proof. Suppose C_1, C_2, x_1, x_2, η , are as in Proposition 1 for \mathcal{C} and the Shapley value Ξ^{SV} . We describe a game $(N, R, \mathcal{A}, w, \mathcal{C}, \Xi)$ and we present a PNE of the game such that the cost is $\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2)$ times the optimal. To simplify the presentation of this proof, we henceforth denote the expression $\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2)$ by ζ . We describe the game and the equilibrium strategies simultaneously. We first construct a resource r_0 which has some arbitrary cost function $C_{r_0}(x)$ from \mathcal{C} . In what follows in our construction, we only consider players i for whom the limit of the sampling weight function at 0, λ_i^0 , is a positive constant. Since N is arbitrarily large and Ξ is balanced with respect to N , there is an infinite supply of such players. Every other player ID that is in N and is not used in what follows, is given r_0 as her only possible strategy and is supposed to have demand 0.

Construction: players, resources, and strategies. We resume our construction with a single resource and a player that is using it by herself. The cost function of the resource is $C_1(x)$ and the demand of the player is ϵ , an arbitrarily small parameter the precise value of which we set later. The ID of the player is the one that maximizes λ_i^0 among the players we are willing to consider (i.e., the ones that have positive constant λ_i^0). We then construct l levels using the following recursive step. Consider any player that uses a resource which lies at a level earlier than $l - 1$. Suppose the player has demand w and cost $\tilde{\xi}$ on the resource. The other alternative of the player would be to use two resources which we now construct and place at the next level. The first one has cost function $\eta \cdot (\tilde{\xi}/\zeta) \cdot C_1(x/w)$ and is used by a player with demand $w \cdot x_1$. The second one has cost function $(1 - \eta) \cdot (\tilde{\xi}/\zeta) \cdot C_2(x/w)$ and is used by a player with demand $w \cdot x_2$. Again, for the new players we place on our construction, we pick the ones that have the maximum λ_i^0 among the remaining available players. Finally, consider the players at level $l - 1$. Focus on one of them and suppose she has demand w and cost $\tilde{\xi}$ on her currently selected resource. Her other alternative is a single resource at level l , which we construct now and has cost function $\tilde{\xi} \cdot C_1(x/w)$. We are

now ready to fix the value of ϵ . We ask that it is small enough, so that the largest demand in the game δ is such that for every $\lambda_i(\cdot)$ of the players in N , the value of $\lambda_i(\delta)$ is within $\tilde{\delta}$ of λ_i^0 , for some parameter $\tilde{\delta}$. As we let $\tilde{\delta} \rightarrow 0$, all sampling weights become equal to the corresponding λ_i^0 .

PNE. Call our constructed outcome A . We can see that it is a PNE. Players in r_0 have no other alternative, players at level $l - 1$ are clearly indifferent between their two alternatives. Consider any other player with demand w and cost $\tilde{\xi}$ in A . Since we have ensured the sampling weights are nonincreasing as we move down the levels of our construction, the cost of the player's possible deviation would be at least her Shapley value on the two resources, which, using (33), becomes $(\tilde{\xi}/\zeta) \cdot (\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2)) = \tilde{\xi}$. Then it follows that the equilibrium condition holds for all players.

PoA. Players in r_0 have no contribution to the total cost since they have 0 demand and, hence, 0 cost. We first make the observation that every resource at level at most $l - 2$ incurs the exact same cost as the total of its children. Then it follows that the cost is the same across levels in A . Call A^* the outcome that has all players play the opposite strategy from the one in A . This means every player plays the children of her resource in A . We can see that every player, other than the ones from level $l - 1$ in A , will have a cost that is ζ times smaller in A^* . The players from level $l - 1$ move to level l where they have the same cost as before. However, due to the fact that the cost across levels is the same in A , these player become negligible as $l \rightarrow +\infty$ and we get the lower bound of ζ . \square

LEMMA 9. *Let Ξ be a weighted Shapley value that is small-demand-favoring with respect to N , an arbitrarily large set of players. Then $PoA(\mathcal{C}, \Xi) \geq PoA(\mathcal{C}, \Xi^{SV})$.*

Proof. Suppose C_1, C_2, x_1, x_2, η , are as in Proposition 1 for \mathcal{C} and the Shapley value Ξ^{SV} . We describe a game $(N, R, \mathcal{A}, w, \mathcal{C}, \Xi)$ and we present a PNE of the game such that the cost is $\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2) = \zeta$ times the optimal. We describe the game and the PNE strategies simultaneously. We first construct a resource r_0 which has some arbitrary cost function $C_{r_0}(x)$ from \mathcal{C} . In what follows in our construction, we only consider players i for whom the limit of the

sampling weight function at 0, λ_i^0 , is 0. Since N is arbitrarily large and Ξ is small-demand-favoring with respect to N , there is an infinite supply of such players. Every other player ID that is in N and is not used in what follows, is given r_0 as her only possible strategy and is supposed to have demand 0.

Construction: players, resources, and strategies. We resume our construction with a single resource and a player that is using it by herself. The cost function of the resource is $C_1(x)$ and the demand of the player is 1. We then construct l levels using the following recursive step. Consider any player that uses a resource which lies at a level earlier than $l - 1$. Suppose the player has demand w and cost $\tilde{\xi}$ on the resource. The other alternative of the player would be to use two resources which we now construct and place at the next level. The first one has cost function $\eta \cdot (\tilde{\xi}/\zeta) \cdot C_1(x/w)$ and is used by x_1/ϵ players with demand $w \cdot \epsilon$ each, where ϵ is an arbitrarily small parameter. The second one has cost function $(1 - \eta) \cdot (\tilde{\xi}/\zeta) \cdot C_2(x/w)$ and is used by x_2/ϵ players with demand $w \cdot \epsilon$ each. Finally, consider the players at level $l - 1$. Focus on one of them and suppose she has demand w and cost $\tilde{\xi}$ on her currently selected resource. Her other alternative is a single resource at level l , which we construct now and has cost function $\tilde{\xi} \cdot C_1(x/w)$.

PNE. Call our constructed outcome A . We can see that it is a PNE. Players in r_0 have no other alternative, players at level $l - 1$ are clearly indifferent between their two alternatives. Consider any other player with demand w and cost $\tilde{\xi}$ in A . Consider the possible deviation of the player to the two children resources that have a large number of players with total demands $w \cdot x_1$ and $w \cdot x_2$ respectively. As we explained in the proof of Lemma 1, with such sampling weights that favor smaller demands, the player with demand w would be the last in the Shapley ordering with probability 1 when she shares a resource with a large number of players with demand ϵ , as $\epsilon \rightarrow 0$. Then, her cost share would be at least her Shapley value on the same resources, versus single players with demands $w \cdot x_1$ and $w \cdot x_2$ respectively. Then, we can see (with the same analysis as in Lemma 8) that the equilibrium condition holds for all players.

PoA. We first make the observation that every resource at level at most $l - 2$ incurs the exact same cost as the total of its children. Then it follows that the cost is the same across levels in A . Call A^* the outcome that has all players play the opposite strategy from the one in A . This means every player plays the children of her resource in A . We can see that every player, other than the ones from level $l - 1$ in A , will have a cost that is ζ times smaller in A^* . The players from level $l - 1$ move to level l where they have the same cost as before. However, due to the fact that the cost across levels is the same in A , these player become negligible as $l \rightarrow +\infty$ and we get the lower bound of ζ . \square

LEMMA 10. *Let Ξ be a weighted Shapley value that is small-demand-punishing with respect to N , an arbitrarily large set of players. Then $PoA(\mathcal{C}, \Xi) \geq PoA(\mathcal{C}, \Xi^{SV})$.*

Proof. Suppose C_1, C_2, x_1, x_2, η , are as in Proposition 1 for \mathcal{C} and the Shapley value Ξ^{SV} . We describe a game $(N, R, \mathcal{A}, w, \mathcal{C}, \Xi)$ and we present a PNE of the game such that the ratio between the PNE cost and the optimal cost is $\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2) = \zeta$. Our construction begins with a resource r_0 which has some arbitrary cost function $C_{r_0}(x)$ from \mathcal{C} . In what follows, we only consider players i for whom the limit l_i^0 of the corresponding sampling weight function at 0, is $+\infty$. Since N is arbitrarily large and Ξ is small-demand-punishing with respect to N , there is an infinite supply of such players. Every other player ID that is in N and is not used in what follows, is assigned demand 0 and is given r_0 as her only possible strategy.

Construction: players, resources, and strategies. With the exception of r_0 , our resources are organized in $l + 1$ levels. Level $j = 1, 2, \dots, l$ has precisely $2^{l-j}/\epsilon^{j-1}$ resources, where ϵ is a very small parameter. Level 0 has the same number of resources as level 1. In the outcome, A , that will serve as our PNE, every resource at level $j \geq 1$ has exactly one distinct user. The player demands, player alternative strategies, and resource cost functions are determined as follows. We begin from level l , where the demand of every players is ϵ^{l-1} and the cost function of every resource is the same arbitrary $C(x)$ from \mathcal{C} . Now the construction proceeds recursively up to level 1, with the following specifics. At the current level, j , partition the players into groups with cardinality $1/\epsilon$, in a way

such that all players in the same group have the same demand w and the same cost $\tilde{\xi}$ in A . This is trivially possible at level l . We will later see that this is the case for all other levels, up to level 2, as well. For each group, select two resources from level $j - 1$ and let them be the only alternative strategy of all the players in the group under consideration from level j . Note that our construction can (by the number of resources/players in each level), and will, select a different pair of resources at level $j - 1$ for each group of the partition of level j . Order the pair of resources selected for the group under consideration and assign cost function $\eta \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_1(x \cdot \epsilon/w)$ to the first and cost function $(1 - \eta) \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_2(x \cdot \epsilon/w)$ to the second. Also, let the player that is using the first have demand $x_1 \cdot w/\epsilon$ and the player that is using the second have demand $x_2 \cdot w/\epsilon$. We can now see that at level j we will have 2^{l-j} distinct demand values, with each one appearing $(1/\epsilon)^{j-1}$ times. Moreover, players with the same demand are always using resources with the same cost functions. Hence, we can always proceed with our partitioning up to level 2. Now, it only remains to specify the construction at level 0 and the alternative strategies of the players at level 1. There are no players at level 0. For every resource at level 1, there is exactly one resource at level 0 that has the same cost function and is the only available alternative strategy of the corresponding player from level 1.

PNE. We now show that A is in fact a PNE. Players in r_0 have no alternative strategy and players at level 1 are trivially indifferent between their action in A and their alternative. Consider any other player who has demand w and cost $\tilde{\xi}$ in A . Her other alternative would be to pick the two corresponding resources of the previous level that have cost functions $\eta \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_1(x \cdot \epsilon/w)$ and $(1 - \eta) \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_2(x \cdot \epsilon/w)$ and are occupied by players with demands $x_1 \cdot w/\epsilon$ and $x_2 \cdot w/\epsilon$ respectively. As we let $\epsilon \rightarrow 0$, it follows that the deviating player will be last in the Shapley ordering of both resources with probability 1. Then the cost $\hat{\xi}$ she would incur by deviating would be:

$$\hat{\xi} = \eta \cdot \frac{\tilde{\xi}}{\zeta \cdot \epsilon} (C_1(x_1 + \epsilon) - C_1(x_1)) + (1 - \eta) \cdot \frac{\tilde{\xi}}{\zeta \cdot \epsilon} (C_2(x_2 + \epsilon) - C_2(x_2)) \quad (34)$$

$$= \frac{\tilde{\xi}}{\zeta} \cdot \left(\eta \cdot \frac{C_1(x_1 + \epsilon) - C_1(x_1)}{(x_1 + \epsilon) - x_1} + (1 - \eta) \cdot \frac{C_2(x_2 + \epsilon) - C_2(x_2)}{(x_2 + \epsilon) - x_2} \right) \quad (35)$$

$$= \frac{\tilde{\xi}}{\zeta} \cdot (\eta \cdot C_1'(x_1) + (1 - \eta) \cdot C_2'(x_2)) \quad (36)$$

$$\geq \frac{\tilde{\xi}}{\zeta} \cdot \left(\eta \cdot \left(\frac{1}{2} \cdot \frac{C_1(x_1 + 1) - C_1(x_1)}{(x_1 + 1) - x_1} + \frac{1}{2} \cdot \frac{C_1(0 + 1) - C_1(0)}{(0 + 1) - 0} \right) + (1 - \eta) \cdot \left(\frac{1}{2} \cdot \frac{C_2(x_2 + 1) - C_2(x_2)}{(x_2 + 1) - x_2} + \frac{1}{2} \cdot \frac{C_2(0 + 1) - C_2(0)}{(0 + 1) - 0} \right) \right) \quad (37)$$

$$= \frac{\tilde{\xi}}{\zeta} \cdot (\eta \cdot C_1(x_1) + (1 - \eta) \cdot C_2(x_2)) = \tilde{\xi}, \quad (38)$$

where (36) follows by the fact that $\epsilon \rightarrow 0$, (37) follows by convexity of C_1, C_2 , and (38) follows by (33). This completes the proof that A is a PNE.

PoA. Let A^* be the outcome that has all players that are out of r_0 play the opposite strategy from the one they picked in A . The players in r_0 have 0 cost in both outcomes. Consider some player at some level larger than 1 who has demand w and cost $\tilde{\xi}$ in A . In A^* , this same player is sharing two resources with cost functions $\eta \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_1(x \cdot \epsilon/w)$ and $(1 - \eta) \cdot (\tilde{\xi}/(\zeta \cdot \epsilon)) \cdot C_2(x \cdot \epsilon/w)$ with another $1/\epsilon - 1$ players with demand w . It follows that their aggregate cost in A^* is $\tilde{\xi}/(\zeta \cdot \epsilon)$, while their aggregate cost in A is clearly $\tilde{\xi}/\epsilon$. The players at level 1 have the same cost in both outcomes, however as $l \rightarrow +\infty$, their contribution to the total costs $Q(A)$ and $Q(A^*)$ becomes negligible, since the cost across levels remains the same (by the fact that any group of players given by the partition described at the start of the proof has the same cost in A as the cost that is induced at the pair of resources given to that group as their alternative strategy). Hence we get $\lim_{l \rightarrow +\infty} Q(A)/Q(A^*) = \zeta$. \square