

# Designing cost-sharing methods for Bayesian Games

George Christodoulou<sup>1</sup>, Stefano Leonardi<sup>\*2</sup>, and Alkmini Sgouritsa<sup>1</sup>

<sup>1</sup> Computer Science Department, University of Liverpool

<sup>2</sup> Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza University of Rome

**Abstract.** We study the design of cost-sharing protocols for two fundamental resource allocation problems, the *Set Cover* and the *Steiner Tree Problem*, under environments of incomplete information (Bayesian model). Our objective is to design protocols where the worst-case Bayesian Nash equilibria, have low cost, i.e. *the Bayesian Price of Anarchy (PoA)* is minimized. Although budget balance is a very natural requirement, it puts considerable restrictions on the design space, resulting in high PoA. We propose an alternative, relaxed requirement called *budget balance in the equilibrium (BBiE)*. We show an interesting connection between algorithms for *Oblivious Stochastic* optimization problems and cost-sharing design with low PoA. We exploit this connection for both problems and we enforce approximate solutions of the stochastic problem, as Bayesian Nash equilibria, with the same guarantees on the PoA. More interestingly, we show how to obtain the same bounds on the PoA, by using *anonymous* posted prices which are desirable because they are easy to implement and, as we show, induce *dominant strategies* for the players.

**Keywords:** Price of Anarchy, Bayesian Games, Network Design

## 1 Introduction

A *cost-sharing game*, is an abstract setting that describes interactions of selfish players in environments where the cost of the produced solution needs to be shared among the participants. A *cost-sharing protocol* prescribes how the incurred cost is split among the users. This defines a game that is played by the participants, who try to select outcomes that incur low personal costs. Chen, Roughgarden and Valiant [6] initiated the *design* aspect, seeking for protocols that induce approximately efficient equilibria, *with low Price of Anarchy (PoA)* [27]. Similarly, we study the design of cost-sharing protocols, for two well-studied and very general resource allocation problems with numerous applications, the *Set Cover* and the *Steiner tree (multicast)* problem.

---

\* Partially supported by the ERC Advanced Grant 788893 AMDROMA "Algorithmic and Mechanism Design Research in Online Markets" and MIUR PRIN project ALGADIMAR "Algorithms, Games, and Digital Markets".

**Set Cover Game.** In the (weighted) set cover problem, there is a universe of  $n$  elements,  $U = \{1, \dots, n\}$ , and a family of subsets of  $U$ ,  $\mathcal{F} = \{F_1, \dots, F_m\}$ , with weights/costs  $c_{F_1}, \dots, c_{F_m}$ . A subset of elements,  $X \subseteq U$ , needs to be covered by the  $F_i$ 's so that the total cost is *minimized*. We are interested in a game theoretic version, where there are  $|X|$  players and  $|U|$  possible *types*;  $X$  corresponds to the set of players and each player's type associates her with a specific element of  $U$ . Multiple players may have the same type. A player's action is to choose a subset from  $\mathcal{F}$  that covers her element, and pay some cost-share for using it. A cost-sharing method prescribes how the subsets' costs are split among players.

**Multicast Game.** In a multicast game, there is a rooted (connected) undirected graph  $G = (V, E, t)$ , where each edge  $e$  carries a nonnegative weight  $c_e$  and  $t$  is a designated root. There are  $k$  players and  $|V|$  possible types; each player's type associates him with a specific vertex of  $V$  which needs to establish connectivity with  $t$ . The players' strategies are all the paths that connect their terminal with  $t$ . A cost-sharing method defines the cost-shares of the players.

**Cost-Sharing under Uncertainty.** There are two different possible sources of uncertainty that may need to be considered in the scenarios discussed above. Firstly, the designer needs to specify the cost-sharing protocol, having only partial information about the types of the players. Moreover, the players themselves, when they select their actions, they may have incomplete knowledge about the types of the other players. We approach the former using a stochastic model similar to [10], and the latter, by modeling the situation as a *Bayesian game*, introduced by [22] as an elegant way of modeling selfishness in partial-information settings. In a Bayesian game, players do not know the private information (types) of the other players, but only have *beliefs*, expressed by probability distributions over the different possible realizations of this private information.

The order of events is as follows; first, the designer specifies the cost-sharing methods, using the product probability distribution over the players' types, then the players interact in the induced Bayesian game, and end up in a Bayesian Nash Equilibrium. We are interested in the design of protocols, where *all* equilibria have low cost i.e., the (Bayesian) PoA of the induced game is *low*.

**Budget-Balance in the Equilibrium (BBiE).** One of the axioms that [6] required in their design space, that every cost-sharing protocol should satisfy, is *budget balance* i.e., that the players' cost-shares cover *exactly* the cost of *any* solution. Although budget balance is a very natural requirement, it puts considerable restrictions on the design space. However, since we expect that the players will end up in a Nash equilibrium, it is not clear why one should be interested to impose budget balance in non-equilibrium states; the players are going to deviate from such states anyways. We propose an alternative, relaxed requirement that we call *budget balance in the equilibrium (BBiE)*. A BBiE cost-sharing protocol satisfies budget balance in *all equilibria*; for any non-equilibrium profile we do not impose this requirement. This natural relaxation, enlarges the design space but maintains the desired property of balancing the cost in the equilibrium. More importantly, this amplification of the design space, allows us to design protocols that dramatically outperform the best possible PoA bounds

obtained by budget-balanced protocols. Indeed, the inefficiency of equilibria in budget balance protocols can be very high. For the set cover game, a lower bound of  $\Omega(n)$  exists, for the complete information setting [6], where  $n$  is the number of elements; we extend this lower bound for the Bayesian setting. We further show a lower bound of  $\Omega(\sqrt{n})$ , for the multicast Bayesian game, where  $n$  is the number of vertices of the underlying graph.

We demonstrate that, by designing BBiE protocols, we can enforce better solutions, that dramatically improve the PoA, and still retain the budget-balance in the resulting equilibria. We improve the PoA from  $\Omega(n)$  to  $O(\log n)$ , if  $m = \text{poly}(n)$  and to  $O\left(\frac{\log m}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ , which is  $O(n/\log n)$  in the worst case (recall that  $m$  is the number of all existing subsets of elements). For the multicast game, we improve the PoA from  $\Omega(\sqrt{n})$  to  $O(1)$ .

**Posted Prices.** It is a very common practice, especially in large markets and double auctions, for sellers to use posted prices on their items. More closely to cost-sharing games is the model proposed by Kelly [25] regarding *bandwidth allocation*. Kelly’s mechanism processes players’ willingness to pay for the bandwidth and posts a price for the whole bandwidth. Then each player pays a price proportional to the portion of bandwidth she wants to use. This can be seen as pricing an infinitesimal quantity of the bandwidth and the players, acting as price-takers, choose some number of quantities to buy. It turns out that it is in the best interest of the players to buy the whole bandwidth.

The use of posted prices to serve as cost-sharing mechanism, is highly desirable, but not always possible to achieve; a price is posted for each resource and then the players behave as price takers, picking up the cheapest possible resources that satisfy their requirements. Such a mechanism is desirable because it is extremely easy to implement and also induces *dominant strategies*. We stress that our main results can be implemented by *anonymous* posted prices.

## 1.1 Results and Discussion

We study the design of cost-sharing protocols for two fundamental resource allocation problems, the *Set Cover* and the *Steiner tree problem (multicast)*. We are interested in environments of incomplete information where both the designer and the players have partial information, described by prior probability distributions over types. We design cost-sharing protocols that are BBiE and our objective is to design protocols where the worst-case Bayesian Nash equilibria have low cost, i.e. the *Bayesian PoA* is minimized.

We show an interesting connection between algorithms for *Oblivious Stochastic* optimization problems and cost-sharing design with low PoA. We exploit this for both problems and we are able to enforce approximate solutions of the stochastic problem, as Bayesian Nash equilibria, with the same guarantees on the PoA. Although this connection is quite simple, it results in significant improvement on the PoA comparing to budget-balanced protocols. More precisely, we map each player to a *single* specific strategy and charge very high costs for any alternative strategy. In this way, their mapped strategy becomes a (strongly)

*dominant strategy*. For the set cover game, we enforce the oblivious solution given by [20]. They apriori map each player  $i$  to some subset  $F_i \in \mathcal{F}$ ; then, if  $i$  is sampled,  $F_i$  should be in the induced solution. Regarding the multicast game, the algorithm of Garg et al. [17], for the online Steiner tree problem, provides an oblivious solution.

**Budget-Balanced Protocols** (Sect. 3). First, we provide lower bounds for the PoA of budget-balanced protocols. It is not hard to see that there exists a set cover game that reduces to the lower bound of Chen, Roughgarden and Valiant [6] for the multicast directed network games, resulting in  $\text{PoA} = \Omega(n)$  in the complete information case (see Appendix B.1 for the reduction). For the stochastic or Bayesian setting, where players are i.i.d., we show that the same lower bound holds, but a further analysis is needed. Regarding the multicast game, the PoA is  $O(1)$  for the complete information case [6] and the stochastic case [10, 17]. However, we show that for the Bayesian setting there is a lower bound of  $\Omega(\sqrt{n})$ .

**BBiE Protocols** (Sect. 4). For the Bayesian (and stochastic) set cover game there exists an *ex-post*<sup>3</sup> BBiE protocol (determined in polynomial time) with  $\text{PoA} = O(\log n)$ , if  $m = \text{poly}(n)$ , and  $O\left(\frac{\log m}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ . An *ex-post* BBiE protocol also exists for the Bayesian multicast game resulting in constant PoA.

**Posted Prices** (Sect. 5). For the Bayesian (and the stochastic) settings, *ex-post* BBiE cannot be obtained by anonymous prices. Hence, we examine prices that are *ex-ante* BBiE. In Appendix A, we discuss limitations of other concepts, such as BBiE with “high” probability or bounded possible excess and deficit. In Sect. 5 we present anonymous prices with the same upper bounds as the BBiE protocols, for the unweighted set cover and for the multicast games, respectively. Regarding the weighted set cover game, we can only provide *semi-anonymous* prices with the same bounds; by semi-anonymous we mean that the prices for each player do not depend on her identity, but only on her type. We leave the case of anonymous prices as an open question. We remark that in all cases, posted prices induce *dominant strategies* for the players. At last, for the poly-time determinable prices, we give tight lower bounds.

We stress here that oblivious solutions may not be sufficient to guarantee low PoA for anonymous posted prices, in contrast to the BBiE protocols. The reason is that it is not clear anymore how to enforce players to choose desirable strategies, since *anonymous* posted prices are available to anyone. The anonymous prices we set for the unweighted set cover and the multicast games exist due to the specific properties of the oblivious solution.

**Prior-independent mechanisms.** Clearly, the above BBiE protocols and posted prices depend on the prior distribution. Prior-independent mechanisms are also of high interest and in Sect. 6 we discuss their limitations.

In Appendix C we study the complete information setting and we give results regarding BBiE protocols and *anonymous* posted prices. Due to lack of space, we give all the missing proofs in Appendix D.

<sup>3</sup> In *ex-post* budget-balance we require budget-balance in every realization of the game. If the *expected* excess and deficit are zero, the budget balance is called *ex-ante*.

## 1.2 Related Work

There is a vast amount of research in cost-sharing games and so, we only mention some of the most related. Moulin and Shenker [29] studied cost-sharing games under mechanism design context. In similar context, other papers considered (group)strategy proof and efficient mechanisms and relaxed the budget-balanced constraint; Devanur, Mihail and Vazirani [12] and Immorlica, Mahdian and Mirrokni [24] studied the set cover game under this context showing positive and negative bounds on the fraction of the cost that is covered.

Regarding the network design games, there is a long line of works mainly focusing on fair cost allocation originated by [2]. Anshelevich et al. [2] showed a tight  $\Theta(\log k)$  bound on the PoS for directed networks, while for undirected networks the exact value of PoS still remains an open problem. For multicast games, Li [28] proved an upper bound of  $O(\log k / \log \log k)$ , while for broadcast games, a constant upper bound is known due to Bilò, Flammini and Moscardelli [4]. Chen, Roughgarden and Valiant [6] were the first to study the design aspects for this game, identifying the best protocol with respect to the PoA and PoS in various cases, followed by [10, 13, 18]. The Bayesian Price of anarchy was first studied in auctions by [8]; see also [30] for routing games, and [32] for the PoS of Shapley protocol in cost-sharing games.

Close in spirit to our work is the notion of Coordination Mechanisms [7] which provide a way to improve the PoA in cases of incomplete information. Similar to our context, the designer has to decide in advance game-specific policies, without knowing the exact input. Such mechanisms have been used for scheduling and simple routing games, see [1, 3, 9] and the papers cited therein.

Posted prices have been used for pricing in large markets. Kelso and Crawford [26] and Gul and Stacchetti [21] proved the existence of prices, for gross substitute valuations, that clear the market efficiently. Pricing bundles for combinatorial Walrasian equilibria was introduced by Feldman, Gravin and Lucier [15], who showed that half of the social welfare can be achieved. In a follow-up work [16], they considered Bayesian combinatorial auctions and they could guarantee half of the optimum welfare, by using anonymous posted prices.

The underlying problems that we consider here, the set cover and the minimum Steiner tree problems, are very well studied and known to be in NP-complete. The best known approximations are  $O(\log(k))$  [11] (by using a simple greedy algorithm) and 1.39 [5]; in fact, for the set cover problem, Feige [14] showed that no improvement by a constant factor is likely. Research has been done regarding the stochastic model, Grandoni et al. [20] showed a roughly  $O(\log nm)$  tight bound for the set cover problem and Garg et al. [17] gave bounds on the approximation of the stochastic online Steiner tree problem. A slightly different distribution is the independent activations; [31] and [10] demonstrated constant approximation algorithms for the universal TSP problem and the multicast game, respectively.

## 2 Model

**Cost-Sharing Protocol.** In the cost-sharing games, we consider that there are  $k$  players who are interested in a set of resources,  $R = \{r_1, \dots, r_m\}$ . Each resource  $r$  carries a cost  $c_r$ . Whenever a subset of players uses a resource  $r$ , they are charged some cost-share, defined by a cost-sharing (resource-specific) method  $\xi$ . A cost-sharing protocol  $\Xi$  decides a cost-sharing method for each resource. In accordance with previous works, [6, 10, 13], the following are some natural properties that  $\Xi$  needs to satisfy:

- *Stability*: The induced game has always a *pure* Nash equilibrium.
- *Separability*: The cost shares of each resource  $r$  are completely determined by the set of players that choose it.
- *BBiE*: In any pure (Bayes) Nash equilibrium profile, the cost shares of the players choosing  $r$  should cover exactly the cost of  $r$ .

For the rest of the paper, by  $k$  we denote the number of players and by  $n$  the number of different types of the players, i.e. in the set cover game,  $|U| = n$ , and in the multicast game,  $|V| = n$ .

**Information Models.** We study several information models, from the point of view of the designer and of other players, regarding the knowledge of players' type. A player's type is some resource: in the set cover game, it is some element from  $U$  that needs to be covered, and in the multicast game, it is some vertex of  $G$ , on which the player's terminal lies, and requires connectivity with the root  $t$ . The parameters of the game is known to both the protocol designer and the participants. To be more specific, the tuple  $(U, \mathcal{F}, c)$  in the set cover game and the underlying (weighted) graph in the multicast game are commonly known.

The information models that we consider are the following:

- *Complete Information*: The types of the players are common knowledge, i.e. they are known to all players and to the designer.
- *Stochastic/A priori*: The players' types are drawn from some product distribution  $D$  defined over the type set ( $U$  for set cover and  $V$  for multicast). The actual types are unknown to the designer, who is only aware of  $D$ . However, the players decide their strategies by knowing other players' types.
- *Bayesian*: The players' types are drawn from some product distribution  $D$  defined over the type sets. Both the designer and the players know only  $D$ . The players now decide their strategies by knowing only  $D$  and not the actual types. A natural assumption is that every player knows her own type.

We assume that the players' types are distributed i.i.d. ( $D = \pi^k$ ) and the type of each player is drawn independently from some probability distribution  $\pi : R \rightarrow [0, 1]$ , with  $\sum_{r \in R} \pi(r) = 1$ ;  $R$  is either  $U$  in the set cover or  $V$  in the multicast. For simplicity we write  $\pi_r$  instead of  $\pi(r)$ .

**Price of Anarchy (PoA).** Let  $opt(\mathbf{t})$  be the optimum solution given the players' types  $\mathbf{t}$ , and  $NE(\mathbf{t})$  and  $BNE$  be the set of pure Nash equilibria and pure Bayesian Nash equilibria, respectively. We denote the cost of any solution

$A$  as  $c(A)$ . Then, the *Price of Anarchy* (PoA) for the complete information, stochastic and Bayesian settings is defined, respectively, as:

$$PoA = \max_{\mathbf{t}} \max_{\mathbf{s} \in NE(\mathbf{t})} \frac{c(\mathbf{s})}{c(\text{opt}(\mathbf{t}))} ; \quad PoA = \max_D \frac{\mathbb{E}_{\mathbf{t} \sim D} [\max_{\mathbf{s} \in NE(\mathbf{t})} c(\mathbf{s})]}{\mathbb{E}_{\mathbf{t} \sim D} [c(\text{opt}(\mathbf{t}))]} ;$$

$$PoA = \max_{D, \mathbf{s} \in BNE} \frac{\mathbb{E}_{\mathbf{t} \sim D, \mathbf{s}(\mathbf{t})} [c(\mathbf{s}(\mathbf{t}))]}{\mathbb{E}_{\mathbf{t} \sim D} [c(\text{opt}(\mathbf{t}))]} .$$


### 3 Lower Bounds for Budget-Balanced Protocols

**Theorem 1.** *The Bayesian or stochastic PoA of any budget-balanced protocol, for the unweighted set cover game, is  $\Omega(n)$ .*

*Proof.* Consider  $n$  players and  $n$  elements/types  $U = (1, \dots, n)$  and the family of sets  $\mathcal{F} = \{F_1 = \{1\}, F_2 = \{2\}, \dots, F_n = \{n\}, F_{all} = U\}$  with unit costs. Suppose that  $\pi$  is the uniform distribution over  $U$ . Then the probability that element  $i$  is drawn as the type of at least one player is  $q_i = 1 - (1 - \frac{1}{n})^n \geq 1 - \frac{1}{e}$ . By using any budget-balanced protocol, it is a (Bayes) Nash equilibrium if each player of type  $i$  chose set  $F_i$ . Her cost-share does not exceed 1, while by deviating to  $F_{all}$  her cost-share becomes 1. The expected cost of that equilibrium is  $nq_i = \Omega(n)$ , whereas the optimum solution (all players choose the set  $F_{all}$ ) has cost 1.  $\square$

**Theorem 2.** *The Bayesian PoA of any budget-balanced protocol, for the multi-cast game, is  $\Omega(\sqrt{n})$ .*

*Proof.* Consider the graph of Fig. 1. We set  $p = 1 - (1 - \frac{1}{\sqrt{n}})^{\frac{1}{n}}$ , such that the probability that vertex  $v_i$  is drawn as the type of at least one player is  $q_i = 1 - (1 - p)^n = \frac{1}{\sqrt{n}}$ . We claim that, for any budget-balanced protocol, it is a Bayes-Nash equilibrium if any player with type  $v_i$  uses the direct edges  $(v_i, t)$ . Indeed, if player  $i$  uses any other path  $(v_i, v, v_j, t)$  her cost-share will be at



**Fig. 1.** Lower bound on the PoA of any budget-balanced protocol.

least  $\frac{2}{\sqrt{n}} + (1 - q_j) = 1 + \frac{1}{\sqrt{n}}$ , which is greater than her current cost-share of at most 1. The expected social cost and optimum are:  $\mathbb{E}[SC] = \sum_i q_i = \sqrt{n}$  and  $\mathbb{E}[Opt] \leq \sum_i q_i \cdot \frac{1}{\sqrt{n}} + 1 = n \frac{1}{n} + 1 = 2$ . So, the Bayes PoA is at least  $\frac{1}{2}\sqrt{n}$ .  $\square$

## 4 BBiE Protocols

In this section we drop the requirement of budget balance and instead we consider a more general class of cost-sharing protocols  $\mathcal{C}$ , where the requirement is to preserve the budget balance in the equilibrium. For the rest of the paper, by  $h$  we denote a very high value with respect to the parameters of the game.  $h$  should be larger than the total cost-share of any player by using any budget-balanced protocol. It is safe to assume that  $h > \sum_{r \in R} c_r$ . For the set cover game it is sufficient that  $h > \max_j c_{F_j}$ . To show our results we will use known oblivious algorithms of the corresponding optimization problems and we will enforce their solution by applying appropriate cost-sharing protocols (or posted prices in Sect. 5); e.g. choices, not consistent with this solution, are highly expensive.

The types of the players correspond to the input components of the problem, and the set of the resources are the domain of players action space. An oblivious algorithm assigns an action for each input component, based on the prior distribution, and *independently of the realization* of all other input components. Take as an example, the multicast game, where the actions of an input (source) corresponds to the paths connecting the source to the root. An oblivious solution, maps each vertex to some path that connects it to the root, and is used in *any realization* of the input that contains this source.

**Theorem 3.** *Let  $G$  be any cost-sharing game and  $\Pi$  the underlying optimization resource allocation problem. Given any oblivious algorithm of  $\Pi$  with approximation ratio  $\rho$ , there exists a cost-sharing protocol  $\Xi \in \mathcal{C}$  for  $G$  with  $PoA = O(\rho)$ .*

The following corollaries hold for both the Bayesian and the stochastic setting.

**Set Cover Game.** Grandoni et al. [20] studied the stochastic problem, and they showed two mapping algorithms for the oblivious set cover problem (one for the *unweighted* problem which is *length-oblivious* and one for the *weighted* problem which is *length-oblivious*), which are almost  $O(\log mn)$ -competitive. For completeness, we give the two algorithms of [20] in the Appendix B.2. Theorem 3 implies the following theorem.

**Theorem 4.** *In the unweighted and weighted set cover game, there exist length-oblivious protocol  $\Xi_1 \in \mathcal{C}$  and length-aware protocol  $\Xi_2 \in \mathcal{C}$ , respectively, both computed in polynomial time, and with  $PoA$  of  $O(\log n)$ , if  $m = \text{poly}(n)$ , and  $O\left(\frac{\log m}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ .*

**Multicast Game.** Garg et al. [17] showed a constant approximation on the online Steiner tree problem. The idea is the following: sample a set  $S$  from the distribution  $\pi^k$  over the vertices and construct a minimum Steiner tree (or a constant approximation). Then connect each other vertex with its nearest vertex from  $S$  via shortest path. That way we end up with a spanning tree  $T$  (standard derandomization techniques can apply [10, 31, 33]).  $T$  defines a single path from each vertex to the root and this is an oblivious strategy for each players' type. By using Theorem 3 and any constant approximation of the minimum Steiner tree (the best known is by [5]), the following theorem holds.

**Theorem 5.** *In the multicast game, there exists  $\Xi \in \mathcal{C}$  with  $PoA = O(1)$ .*



## 5 Posted Prices

In this section, we show how to set *anonymous* or *semi-anonymous* prices for the resources. Strict BBiE cannot be obtained by using anonymous posted prices, as Example 13 (Appendix A) illustrates. Instead, we require *ex-ante* budget-balance, meaning budget-balance on expectation. In Appendix A we discuss other natural-variations of budget balance and show their limitations. The following proofs are for Bayes-Nash equilibria, but they also hold for the stochastic setting. For the rest of the section we define, for any subset of types  $A$ ,  $k_A$  to be the expected number of players having type in  $A$  and  $k_A^1$  to be the expected number of players having type in  $A$ , given there exists at least one such player:

$$k_A = \mathbb{E}_{\mathbf{t}}[|i : t_i \in A|] = k \sum_{i \in A} \pi_i ;$$

$$k_A^1 = \mathbb{E}_{\mathbf{t}}[|i : t_i \in A| \text{ given } |i : t_i \in A| \geq 1] = \frac{k \sum_{i \in A} \pi_i}{1 - (1 - \sum_{i \in A} \pi_i)^k} . \quad (1)$$

**Set Cover Game.** To determine anonymous prices for the unweighted set cover game, we first state Lemma 6 to be used in stability arguments.

**Lemma 6.** *For any  $a > b > 0$  and integer  $k \geq 2$ ,  $\frac{a}{1-(1-a)^k} > \frac{b}{1-(1-b)^k}$ .*

**Theorem 7.** *In the unweighted set cover game, there exist length-oblivious and anonymous prices (computed in polynomial time) with PoA  $O(\log n)$ , if  $m = \text{poly}(n)$ , and  $O\left(\frac{\log m}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ .*

*Proof.* In order to set the prices, we run the greedy algorithm of [11] and at each step we set the price for the selected set. Algorithm 1 describes this procedure.

---

### ALGORITHM 1: Bayesian posted prices.

---

**Input:**  $(U, \mathcal{F})$ .  
**while**  $U \neq \emptyset$  **do**  
    let  $F \leftarrow$  set in  $\mathcal{F}$  maximizing  $\sum_{i \in F \cap U} \pi_i$ ;  
    set the price for  $F$  to  $\frac{1}{k_{F \cap U}^1}$ ; Let  $U \leftarrow U \setminus F$ .  
**end**  
Set the price of all other sets to  $h$ .

---

We first argue that there exists a unique Bayes-Nash equilibrium, where each player  $i$  chooses the set picked earlier by Algorithm 1 and covers her. For that it is sufficient to show that for any two sets  $A$  and  $B$ , such that  $\sum_{i \in A} \pi_i > \sum_{i \in B} \pi_i$ ,  $k_A^1 > k_B^1$ . From (1), we need to show that  $\frac{k \sum_{i \in A} \pi_i}{1 - (1 - \sum_{i \in A} \pi_i)^k} > \frac{k \sum_{i \in B} \pi_i}{1 - (1 - \sum_{i \in B} \pi_i)^k}$ , which is true due to Lemma 6, by setting  $a = \sum_{i \in A} \pi_i$  and  $b = \sum_{i \in B} \pi_i$ ; note that for  $k = 1$ , there exists only one player and this is a trivial case.

Next notice that, given that a set  $F$  is chosen by some player, the expected number of players paying for it is  $k_F^1$ , resulting in ex-ante BBiE.

Regarding the PoA, Grandoni et al. [20] analyzed the performance of Algorithm 1, for the stochastic set cover problem. They didn't consider any prices, instead they mapped each player to the set first picked by their algorithm and covers her (see Algorithm 2 in Appendix B.2). They cared about the coverage of the drawn elements and used the mapped sets to form a set cover. Their cover though coincide with the solution of the equilibrium and therefore their results immediately provide bounds on the PoA, and the theorem follows.  $\square$

**Theorem 8.** *In the weighted set cover game, there exist length-aware and semi-anonymous prices (computed in polynomial time) with PoA  $O(\log n)$ , if  $m = \text{poly}(n)$ , and  $O\left(\frac{\log nm}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ .*

**Theorem 9.** *For  $m = \text{poly}(n)$  and  $k = \Omega(n)$ , there are no anonymous prices for the unweighted set cover, or semi-anonymous prices for the weighted set cover, computed in poly-time, with  $\text{PoA} = o(\log n)$ , unless  $NP \subseteq DTIME(n^{O(\log \log n)})$ .*

**Theorem 10.** *For  $m \gg n$  and  $k = \Omega(n)$ , there are no anonymous prices for the unweighted set cover, or semi-anonymous prices for the weighted set cover, with  $\text{PoA} = o\left(\frac{\log m}{\log \log m - \log \log n}\right)$ .*

**Multicast Game.** We construct a spanning tree  $T$  in the same way as in Sect. 4 and then we use it in order to set the posted prices.

**Theorem 11.** *In the multicast game, there exist prices (computed in polynomial time) with  $\text{PoA} = O(1)$ .*

*Proof.* For each edge  $e \in E(T)$ , let  $V(e)$  be the set of vertices that are disconnected from the root  $t$  in  $T \setminus \{e\}$ . We set the price for each  $e \in E(T)$  as  $c_e/k_{V(e)}^1$ . For each  $e \notin E(T)$ , the price is set to  $h$ . In the equilibrium each player chooses the path that connects her terminal with  $t$  via  $T$ . The constant PoA follows by [17] and the approximation of [5]. The expected total prices for  $e \in E(T)$  is  $k_{V(e)}^1 c_e / k_{V(e)}^1 = c_e$ , if  $e$  is used, and 0 otherwise, resulting in ex-ante BBiE.  $\square$

## 6 Prior-Independent Mechanisms

The design of prior-independent mechanisms is a more difficult task, as the objective now is to identify a single mechanism that always has good performance, under any distributional assumption. In this section, we show limitations of prior-independent mechanisms even for the restricted class of i.i.d. prior distributions.

**BBiE Protocols.** Satisfying BBiE with prior-independent protocols highly restricts the class of cost-sharing protocols and seems hard for natural classes of distribution, e.g. i.i.d., to find ex-post BBiE protocols with low PoA. Regarding the weighted set cover game with i.i.d. distributions, we can construct a lower bound of  $\Omega(\sqrt{n})$  for all prior-independent mechanisms, which are ex-post BBiE.

**Theorem 12.** *In the weighted set cover game, any prior-independent, ex-post BBiE protocol  $\Xi \in \mathcal{C}$  has  $\text{PoA} = \Omega(\sqrt{n})$ .*

*Proof.* Consider  $n$  players,  $n+1$  elements/types  $U = \{0, 1, \dots, n\}$  and the family of sets  $\mathcal{F} = \{F_0, F_1, \dots, F_n, F_{all}\}$ , with  $F_j = \{j\}$ ,  $c_{F_j} = 1$  for all  $j$ , and  $F_{all} = \{1, \dots, n\}$ ,  $c_{F_{all}} = \sqrt{n}$ . Note that 0 is covered only by  $F_0$ , serving as dummy set.

Given a BBiE, prior-independent protocol  $\Xi$ , suppose that there exists some  $F_j$ ,  $j \neq 0$ , where  $\Xi$  is not budget-balanced, i.e. there exists a set of players  $S$ , such that if only  $S$  chooses  $F_j$ , the sum of their cost-shares are different from 1. Consider the prior distribution  $D_1 = \pi^n$  with  $\pi(0) = \pi(j) = 1/2$  and  $\pi(j') = 0$  for any  $j' \notin \{0, j\}$ . With positive probability,  $1/2^n$ , all player of  $S$  have type  $j$  and all other players have type 0. If all players of  $S$  choose  $F_j$  in any *pure* Bayes-Nash equilibrium, ex-post BBiE is violated. So, there exists a player choosing  $F_{all}$  (and this happens with probability  $1/2$ ) which results in  $\text{PoA} = \Omega(\sqrt{n})$ .

Suppose now that  $\Xi$  is budget-balanced for any  $F_j$ , where  $j \neq 0$ . Let  $I$  be the set of players such that whenever  $i \in I$  is the only player choosing  $F_{all}$ ,  $\Xi$  doesn't charge  $\sqrt{n}$  to  $i$ . Consider the prior distribution  $D_2 = \pi^n$  with  $\pi(0) = 1/2$  and  $\pi(j) = 1/(2^n n)$  for all other  $j$ . With positive probability,  $1/(2^n n)$ , player  $i$ 's type is some  $j \neq 0$  and all other players' type is 0. If for any type  $j \neq 0$  player  $i$  chooses  $F_{all}$  in any Bayes-Nash equilibrium, ex-post BBiE is violated.

We claim that the strategy profile, where any player  $i$  with type  $t_i$  chooses  $F_{t_i}$  is a Bayes-Nash equilibrium. For any player  $i \in I$  there is no other valid strategy. For each player  $i \notin I$ , whenever  $t_i \neq 0$ , player  $i$  always pays at most 1 (due to budget balanced in  $F_{t_i}$ ), whereas if she deviates to  $F_{all}$  she pays  $\sqrt{n}$ .

Each element  $j \neq 0$  is a type of a player with probability  $1 - (1 - \frac{1}{2^n})^n \geq 1 - \frac{2}{e}$ , giving an expected cost of  $\Omega(n)$  in the equilibrium. The expected optimum is at most  $1 + \sqrt{n}$  by using only  $F_0$  and  $F_{all}$  and so  $\text{PoA} = \Omega(\sqrt{n})$ .  $\square$

**Posted Prices.** Setting posted prices in the adversarial model cannot guarantee any budget-balance in equilibrium, even ex-ante. Consider the set cover game (similar example exists for the multicast game) with  $n$  players,  $n$  elements and two subsets of unit costs, one containing element 1 and the other containing the rest. Suppose now that we post a price  $q$  for the first subset. If  $q \leq 1/\sqrt{n}$ , for the uniform prior distribution, the expected number of players with type 1, given that there exists at least one, is  $\frac{n \cdot 1/n}{1 - (1 - 1/n)^n} \leq \frac{e}{e-1}$ . The expected cost shares for the first set are  $O(1/\sqrt{n})$ , meaning that its cost is undercovered by a factor of  $\Omega(\sqrt{n})$ . If  $q > 1/\sqrt{n}$ , consider the prior  $D = \pi^n$ , where  $\pi(1) = 1$  and  $\pi(j) = 0$  for all  $j \neq 1$ . All players choose the first set and their total shares are  $n \cdot 1/\sqrt{n} = \sqrt{n}$  which exceeds the set's cost by a factor of  $\sqrt{n}$ . So, there is no way to avoid an over/under-charge of a resource by a factor better than  $\Theta(\sqrt{n})$ .

## References

1. Abed, F., Correa, J.R., Huang, C.: Optimal coordination mechanisms for multi-job scheduling games. In: ESA (2014)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J.M., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM J. Comput. (2008)

3. Bhattacharya, S., Kulkarni, J., Mirrokni, V.S.: Coordination mechanisms for selfish routing over time on a tree. In: ICALP (2014)
4. Bilò, V., Flammini, M., Moscardelli, L.: The price of stability for undirected broadcast network design with fair cost allocation is constant. In: FOCS (2013)
5. Byrka, J., Grandoni, F., Rothvoß, T., Sanità, L.: An improved lp-based approximation for steiner tree. In: STOC (2010)
6. Chen, H., Roughgarden, T., Valiant, G.: Designing network protocols for good equilibria. *SIAM J. Comput.* (2010)
7. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination mechanisms. *Theor. Comput. Sci.* (2009)
8. Christodoulou, G., Kovács, A., Schapira, M.: Bayesian combinatorial auctions. In: ICALP (2008)
9. Christodoulou, G., Mehlhorn, K., Pyrga, E.: Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica* (2014)
10. Christodoulou, G., Sgouritsa, A.: Designing networks with good equilibria under uncertainty. In: SODA (2016)
11. Chvatal, V.: A greedy heuristic for the set-covering problem. *Mathematics of Operations Research* (1979)
12. Devanur, N.R., Mihail, M., Vazirani, V.V.: Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems* (2005)
13. von Falkenhausen, P., Harks, T.: Optimal cost sharing for resource selection games. *Mathematics of Operations Research* (2013)
14. Feige, U.: A threshold of  $\ln n$  for approximating set cover. *J. ACM* (1998)
15. Feldman, M., Gravin, N., Lucier, B.: Combinatorial walrasian equilibrium. In: STOC (2013)
16. Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: SODA (2015)
17. Garg, N., Gupta, A., Leonardi, S., Sankowski, P.: Stochastic analyses for online combinatorial optimization problems. In: SODA (2008)
18. Gkatzelis, V., Kollias, K., Roughgarden, T.: Optimal cost-sharing in weighted congestion games. In: WINE (2014)
19. Gopalakrishnan, R., Marden, J.R., Wierman, A.: Potential games are necessary to ensure pure nash equilibria in cost sharing games. In: EC (2013)
20. Grandoni, F., Gupta, A., Leonardi, S., Miettinen, P., Sankowski, P., Singh, M.: Set covering with our eyes closed. *SIAM J. Comput.* (2013)
21. Gul, F., Stacchetti, E.: Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* (1999)
22. Harsanyi, J.C.: Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science* (1968)
23. Imase, M., Waxman, B.M.: Dynamic Steiner tree problem. *SIAM J. Discrete Math.* (1991)
24. Immorlica, N., Mahdian, M., Mirrokni, V.S.: Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms* (2008)
25. Kelly, F.: Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.* (1997)
26. Kelso, A.S., Crawford, V.P.: Job matching, coalition formation, and gross substitutes. *Econometrica* (1982)
27. Koutsoupias, E., Papadimitriou, C.H.: Worst-case equilibria. In: STACS (1999)
28. Li, J.: An  $o(\log(n)/\log(\log(n)))$  upper bound on the price of stability for undirected shapley network design games. *Inf. Process. Lett.* (2009)
29. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory* (2001)

30. Roughgarden, T.: The price of anarchy in games of incomplete information. *ACM Trans. Economics and Comput.* (2015)
31. Shmoys, D.B., Talwar, K.: A constant approximation algorithm for the a priori-traveling salesman problem. In: *IPCO* (2008)
32. Syrgkanis, V.: Price of stability in games of incomplete information. *CoRR* (2015)
33. Williamson, D.P., van Zuylen, A.: A simpler and better derandomization of an approximation algorithm for single source rent-or-buy. *Oper. Res. Lett.* (2007)

## APPENDIX

### A Posted Prices and BBiE

In Sect. 5 we showed results for posted prices that satisfy *ex-ante* BBiE. Here, we argue that other natural variations of budget balance cannot be very promising: a) BBiE with “high” probability, b) bounded possible excess and deficit. Example 13 indicates that *any* anonymous posted prices may result in BBiE with probability at most  $O(1/\sqrt{k})$ . The same example serves to demonstrate that *no* posted prices can guarantee good bounds on possible excess and deficit, i.e. for *any* posted prices, there are cases where the total shares for some resource are either at least  $\sqrt{k}$  or at most  $1/\sqrt{k}$  of the resource’s cost.

*Example 13.* Consider  $k$  i.i.d. players whose type is the uniform distribution over two elements  $e_1, e_2$  in set cover or two vertices  $v_1, v_2$  in multicast. In the set cover game, there are only two subset of unit cost,  $\mathcal{F} = \{\{e_1\}, \{e_2\}\}$ . In the multicast game there are only two edges,  $(v_1, t)$  and  $(v_2, t)$ , of unit cost. The question that arises in both cases is how to set a price on a resource  $r$  of unit cost, when each player may use it with probability  $1/2$ . Let  $q$  be the price for resource  $r$ . If  $1/q$  is not an integer in  $\{1, \dots, k\}$ , then budget-balance appears with zero probability. So, suppose that  $1/q = k' \in \{1, \dots, k\}$ , then budget-balance appears only when  $k'$  players use resource  $r$  and this happens with probability,  $\Pr[\# \text{ players} = k'] = \binom{k}{k'} \left(\frac{1}{2}\right)^{k'} \left(1 - \frac{1}{2}\right)^{k-k'} \leq \binom{k}{\lfloor k/2 \rfloor} 1/2^k < 1/\sqrt{k}$ . Furthermore, for any price  $q$  for resource  $r$ , if  $q \geq 1/\sqrt{k}$  then, in the case that all players use  $r$ , the total shares sum up to at least  $k \cdot 1/\sqrt{k} = \sqrt{k}$ . On the other hand, if  $q < 1/\sqrt{k}$  then, in the case that only one player uses  $r$ , her share is at most  $1/\sqrt{k}$ . This means that we cannot guarantee good bounds on any possible excess and deficit.

### B Set Cover

#### B.1 Set Cover Reduction

The following multicast (directed) network cost-sharing game models the set cover game. Consider a directed bipartite graph with  $U$  and  $\mathcal{F}$  as the two sets of vertices. For each pair of vertices  $i \in U$  and  $F_j \in \mathcal{F}$ , we add a directed edge  $(i, F_j)$  if and only if  $i \in F_j$ ; the cost of such an edge is set to 0. Each vertex  $i$  of  $U$  is associated with the terminal  $t_i$  of some player. We further add an extra vertex  $t$  as the common destination and we add a directed edge  $(F_j, t)$  for every  $F_j \in \mathcal{F}$ , with cost  $c_{F_j}$ .

All strategies are two length paths and for each player/element  $i \in U$  the space of their strategies are all the paths where their middle vertex is a set that  $i$  belongs to. The cost-sharing protocol for each  $(F_j, t)$  edge determines exactly the cost-shares for the players that choose  $F_j$ . From the lower bound of Chen et al. [6] (Proposition 4.12), the following theorem can be trivially derived.

**Theorem 14.** *The PoA of the (unweighted) set cover cost-sharing game for the complete information setting is  $\Omega(n)$*

## B.2 Set Cover Algorithms

For completeness, we give here the two algorithms of [20], Algorithms 2 and 3, respectively. These algorithm are applied in the case of uniform distribution and then a reduction applies to generalize the results.

---

**ALGORITHM 2:** Algorithm 1 of [20]: Mapping for unweighted set cover.

---

**Input:**  $(U, \mathcal{F})$ .  
**while**  $U \neq \emptyset$  **do**  
     let  $F \leftarrow$  set in  $\mathcal{F}$  maximizing  $|F \cap U|$ ;  
     for each  $u \in F \cap U$ , map  $u$  to  $F$ ;  
      $U \leftarrow U \setminus F$ .  
**end**

---



---

**ALGORITHM 3:** Algorithm 2 of [20]: Mapping for weighted set cover.

---

**Input:**  $(U, \mathcal{F}, c), \mathbb{E}[c(opt)]$ .  
**while**  $U \neq \emptyset$  **do**  
     let  $F \leftarrow$  set in  $\mathcal{F}$  minimizing  $\frac{c_F}{|F \cap U|}$ ;  
     **if**  $\frac{c_F}{|F \cap U|} > \frac{64\mathbb{E}[c(opt)]}{|U|}$ ;  
         **then**  
             let  $F \leftarrow$  set in  $\mathcal{F}$  minimizing  $c_F$ ;  
         **end**  
     for each  $u \in S \cap U$ , map  $u$  to  $F$ ;  
      $U \leftarrow U \setminus F$  and  $\mathcal{F} \leftarrow$  all sets covering at least one element remaining in  $U$ .  
**end**

---

## C Complete Information

In the complete information setting, the input now is known and therefore for any feasible solution we can consider oblivious strategies. If  $\mathcal{F}' \subseteq \mathcal{F}$  is any feasible solution with cost  $c(\mathcal{F}')$  and  $\mathcal{F}^* \subseteq \mathcal{F}$  is the optimum solution with cost  $c(\mathcal{F}^*)$ , then by Theorem 3 we can enforce the solution  $\mathcal{F}'$  and get a PoA of  $c(\mathcal{F}')/c(\mathcal{F}^*)$ . In the following we consider  $\mathcal{F}'$  as either the optimum solution or an approximation. Note that, while trying to bound the PoA, computational issues are not of primary concern.

**Set Cover Game.** By considering  $\mathcal{F}'$ , as either the optimum solution or its  $O(\log k)$ -approximation of the greedy algorithm, where  $k$  is the number of players to be covered, and by using Theorem 3 we get the following corollary.

**Corollary 15.** *For the (weighted) set cover game there exists a protocol  $\Xi \in \mathcal{C}$ , that can be defined in exponential time, with  $PoA = 1$  and a protocol  $\Xi \in \mathcal{C}$ , that can be defined in polynomial time, with  $PoA = O(\log k)$ .*

Next we show that there exist posted prices that can be computed in polynomial time with  $PoA = O(\log k)$ . We show that, under the restriction of setting the prices in polynomial time, this bound is tight. Then we drop the constrain of defining the prices in polynomial time, and we define posted prices with  $PoA = 1$ .

**Theorem 16.** *There exist posted prices, that can be set in polynomial time, for the (weighted) set cover game, with  $PoA = O(\log k)$ .*

*Proof.* Let  $k_e$  be the number of players with type  $e \in U$  and  $X \subseteq U$  be the set of elements that needs to be covered, i.e.  $k_e = 0$  if and only if  $e \notin X$ . Clearly  $\sum_{e \in X} k_e$  equals  $k$ , the total number of players. To set the prices we run Algorithm 4:

---

**ALGORITHM 4:** Complete information posted prices.

---

**Input:**  $(X, \mathcal{F}, c)$ .  
**while**  $X \neq \emptyset$  **do**  
    let  $F \leftarrow$  set in  $\mathcal{F}$  minimizing  $\frac{c(F)}{\sum_{e \in F \cap X} k_e}$ ;  
    set the price for  $F$  to  $\frac{c(F)}{\sum_{e \in F \cap X} k_e}$ ;  
     $X \leftarrow X \setminus F$ .  
**end**  
Set the price of all other sets to  $h$ .

---

There exists a unique Nash equilibrium, where each player  $i$  chooses the set picked earlier by Algorithm 4 and covers her, let it be  $F^i$ . This is because any other set  $F$  that covers  $i$  and picked after  $F^i$  should have at least the same value  $\frac{c(F)}{\sum_{e \in F \cap X} k_e}$  with  $\frac{c(F^i)}{\sum_{e \in F^i \cap X} k_e}$  at the time that  $F^i$  was picked. After processing  $F^i$ ,  $F \cap X$  is diminished by at least 1, and so the price of  $F$  should be strictly higher than the price of  $F^i$ .

In order to show the bound on the  $PoA$ , we consider a reduction from the set cover game  $G = (U, X, \mathcal{F}, c, \mathbf{t})$ , where  $\mathbf{t}$  are players' types, to a set cover problem  $\Pi = (U', \mathcal{F}', c')$ , such that the  $PoA$  equals the approximation ratio of  $\Pi$ .  $U'$  is derived by  $X$ , after replacing each of its elements,  $e$ , by  $k_e$  copies  $e_1, \dots, e_{k_e}$ . For each set  $F \in \mathcal{F}$ , we construct a set  $F' \in \mathcal{F}'$  of the same cost ( $c_F = c'_{F'}$ ), first by erasing all elements belonging to  $U \setminus X$ , and then by replacing each of its remaining elements,  $e$ , by  $k_e$  copies  $e_1, \dots, e_{k_e}$ . We assume that each player  $i \in \{1, \dots, k_e\}$  of type  $e$  is associated with element  $e_i$  of the constructed set cover problem.

It is easy to see that the greedy algorithm on  $\Pi$  chooses the sets in the same order with Algorithm 4. Therefore, the approximation ratio for  $\Pi$  equals the



PoA of  $G$ . Notice that  $|U'| = k$  and since the approximation ratio of the greedy is  $O(\log k)$ , the bound on the PoA follows. Note further that the sum of the prices for each such set used in the Nash equilibrium equals the cost of the set, as desired.  $\square$

**Theorem 17.** *There are no posted prices computed in polynomial time, for the weighted set cover game, with  $PoA = o(\log k)$ , unless  $NP \subseteq DTIME(n^{O(\log \log n)})$ .*

*Proof.* On the contrary, suppose there exist posted prices computed in polynomial time with  $PoA = o(\log k)$ . This would imply a polynomial time algorithm for the set cover problem with approximation ratio of  $o(\log k)$ , but this is a contradiction due to [14].  $\square$

**Theorem 18.** *There exist posted prices (set in exponential time), for the (weighted) set-cover game, with  $PoA = 1$ .*

*Proof.* As in the proof of Theorem 16, let  $k_e$  be the number of players with type  $e \in U$  and  $X \subseteq U$  be the set of elements/players that need to be covered. Moreover, let  $\mathcal{F}^* \subseteq \mathcal{F}$  be the optimum solution, found in exponential time. To set the prices we run Algorithm 4 but for input  $(X, \mathcal{F}^*)$ .

By using similar arguments as in the proof of Theorem 16, in the (unique) Nash equilibrium, each player chooses the set picked earlier by Algorithm 4 and covers her. The prices for each set used in the Nash equilibrium should equal the cost of the set. The difference here is that the Nash equilibrium uses only the sets of  $\mathcal{F}^*$ , resulting in  $PoA = 1$ .  $\square$

**Multicast Game.** Similarly with the set cover game, we can easily get the following corollary by using Theorem 3, for the multicast game. For the second part, we use the 1.39-approximation algorithm of [5].

**Corollary 19.** *For the multicast game there exists a cost-sharing protocol  $\Xi \in \mathcal{C}$ , that can be defined in exponential time, with  $PoA = 1$  and a cost-sharing protocol  $\Xi \in \mathcal{C}$ , that can be defined in polynomial time, with  $PoA \leq 1.39$ .*

We next use posted prices and show that the PoA is constant for the case of multicast game. By using the 1.39-approximation algorithm of [5], the PoA is constant even if we require the prices to be computed in polynomial time.

**Theorem 20.** *For the multicast game, there exist posted prices, computed in polynomial time, with  $PoA \leq 1.39$  and posted prices, computed in exponential time, with  $PoA = 1$ .*

*Proof.* Let  $V'$  be the set of players' terminal and  $T$  be the solution that approximates the minimum Steiner tree on the requested vertices  $S \cup \{t\}$  derived by the 1.39-approximation algorithm of [5]. If we drop the requirement of computing the prices in polynomial time,  $T$  is the minimum Steiner tree. For each edge  $e \in E(T)$ , let  $k_e$  be the number of players that are disconnected from  $t$  in  $T \setminus \{e\}$ . We set the price for  $e \in E(T)$  as  $c_e/k_e$ . For each  $e \notin E(T)$ , we set the

price to be  $h$ . In the Nash equilibrium each player will choose the unique path that connects her source with  $t$  in  $T$ , since any other path has high cost at least  $h$ . Obviously, the players covers exactly the cost of each used edge and the cost of the Nash equilibrium equals the cost of  $T$ .  $\square$

## D Missing Proofs

**Proof of Theorem 3:** Suppose that  $R_i$  is the set of the resources allocated by the oblivious algorithm to the the input component that serves as the type of some player  $i$ . Even though it is not quite correct, we will say that  $R_i$  are the resources allocated to player  $i$ . Let  $S_r$  be the set of players to whom resource  $r$  is allocated.

Then  $\Xi$  assigns the following cost-share to any player  $i$  for choosing any resource  $r$ , when the set of players choosing  $r$  is  $S$ ,

$$\xi_r(i, S) = \begin{cases} c_r/|S| & \text{if } i \in S_r \\ h & \text{if } i \notin S_r \\ 0 & \text{otherwise} \end{cases}$$

Note that  $\Xi$  assigns equal shares restricted to  $S_r$  and a high value  $h$  for other players. In fact, instead of equal shares we could use any budget-balanced protocol restricted to  $S_r$ , for instance any generalized weighted Shapley protocol (for definition see [19]).

Note that any player  $i$  using a resource  $r \notin R_i$  should pay  $h$ . By the definition of  $h$ , this is *strictly* more than  $\sum_{r' \in R_i} c_{r'}$ , which is the maximum she may pay if she deviates to  $R_i$ . Therefore, the only Nash equilibria are for each player  $i$  to choose some subset of  $R_i$ . This results to a PoA which is at most the same with the approximation ratio of the optimization problem, so  $\text{PoA} = O(\rho)$ . Moreover, by the construction of  $\Xi$ , BBiE holds.  $\square$

**Proof of Lemma 6:** We prove the lemma by mathematical induction on  $k$ . For  $k = 2$ ,

$$\frac{a}{1 - (1 - a)^2} = \frac{a}{2a - a^2} = \frac{1}{2 - a} > \frac{1}{2 - b} = \frac{b}{1 - (1 - b)^2}.$$

Suppose that the statement holds for  $k - 1$ , i.e.  $\frac{a}{1 - (1 - a)^{k-1}} > \frac{b}{1 - (1 - b)^{k-1}}$ . We show the equivalent inequality  $\frac{1 - (1 - a)^k}{a} < \frac{1 - (1 - b)^k}{b}$ ,

$$\begin{aligned} \frac{1 - (1 - a)^k}{a} &= \frac{1 - (1 - a)^{k-1}(1 - a)}{a} = \frac{1 - (1 - a)^{k-1} + a(1 - a)^{k-1}}{a} \\ &= \frac{1 - (1 - a)^{k-1}}{a} + (1 - a)^{k-1} < \frac{1 - (1 - b)^{k-1}}{b} + (1 - b)^{k-1} = \frac{1 - (1 - b)^k}{b}. \end{aligned}$$

$\square$

**Proof of Theorem 8:** By using the mapping of Algorithm 3 (Algorithm 2 of [20]), let  $S(F)$  be the set of elements mapped to set  $F$ . For each set  $F \in \mathcal{F}$ , set the price to be  $\frac{c_F}{k_{S(F)}^1}$ , for any player whose type is in  $S(F)$ , and let the price be  $h$  for any other player. Such way, in the Bayes-Nash equilibrium, each player chooses the set, to which is mapped by Algorithm 3. Grandoni et al. [20] showed that the expected cost of that mapping is  $O(\log n)$ , if  $m = \text{poly}(n)$ , and  $O\left(\frac{\log m}{\log \log m - \log \log n}\right)$ , if  $m \gg n$ , away from the expected cost of the optimum solution. Those also serve as upper bounds on the PoA of the induced game. Finally note that those prices are ex-ante budget-balanced in the equilibrium, for the same reasons as the unweighted case.  $\square$

**Proof of Theorem 9:** For any set cover problem  $\Pi = (U, \mathcal{F}, c)$ , we consider the game  $G$  with  $|U|$  players where each one is associated with a different element. Consider the stochastic or Bayesian game, where  $k \gg n$ , and  $k$  is sufficiently large such that the probability that each element is the type of some player converges to 1. Then, we apply the prices on the stochastic  $G$ . It is easy to see that we can compute a Nash equilibrium in polynomial time,  $O(nm)$ ; players choose, among the sets that covers them, some with minimum price. All the chosen sets define a set cover for  $U$ . If there exist posted prices computed in polynomial time with  $\text{PoA} = o(\log nm) = o(\log n)$ , this would imply a polynomial time algorithm for the set-cover problem with approximation ratio  $o(\log n)$ . However, by [14], no polynomial time algorithm for the set cover problem can approximate the optimal solution by  $o(\log k)$ , unless  $NP \subseteq DTIME(n^{O(\log \log n)})$ , which results to a contradiction.  $\square$

**Proof of Theorem 10:** On the contrary, suppose that such prices exist. Then, they would determine a mapping from the elements to the sets, meaning each element is covered by a specific set. The expected cost of the sets that cover the sampled elements would be  $o\left(\frac{\log m}{\log \log m - \log \log n}\right)$  away from the expected cost of the optimum solution. This contradicts the lower bound given by [20] (Theorem 4.2).  $\square$