Designing Cost-Sharing Methods for Bayesian Games

George Christodoulou[1,⋆,⋆⋆], Stefano Leonardi[2,⋆⋆⋆], and Alkmini Sgouritsa[1]

[1] Computer Science Department, University of Liverpool, UK
{gchristo,salkmini}@liv.ac.uk
[2] Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza University of Rome, Italy
leonardi@dis.uniroma1.it

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*

⋆ This work was supported by EP/M008118/1, EP/K01000X/1 and LT140046.

⋆⋆ Part of this work was done while all authors were visiting the Economics and Computation program of the Simons Institute for the Theory of Computing.

⋆⋆⋆ The work of S. Leonardi was partially supported by Google Focused award "Algorithms for Large-scale Data analysis" and EU FET project no. 317532 "Multiplex".

(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.

Set Cover Game. In the (weighted) set cover problem, there is a universe of n elements, $U = \{1, \ldots, n\}$, and a family of subsets of U, $\mathcal{F} = \{F_1, \ldots, F_m\}$, with weights/costs $c_{F_1}, \ldots, c_{F_m}$. 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 chose 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 $|V| = n$ 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 above scenarios. Firstly, the designer needs to specify the cost-sharing protocol, having only partial information about the players' types. Moreover, the players themselves, when they select their actions, may have incomplete knowledge about the types of the other players. We approach the former by using a stochastic model similar to [10], and the latter, as a *Bayesian game*, introduced by [22], which is an elegant way of modelling selfishness in partial-information settings. In a Bayesian game, players do not know the private types of the other players, but only have *beliefs*, expressed by probability distributions over the possible realizations of the types.

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, by restricting to budget-balanced protocols, a lower bound of $\Omega(n)$ exists, for the complete information set cover game [6]; we extend this lower bound for the Bayesian setting. We further show a lower bound of $\Omega(\sqrt{n})$, for the multicast Bayesian game. We demonstrate that, by designing BBiE protocols, we can enforce better solutions, that dramatically improve the PoA. For the set cover game, we improve the PoA to $O(n/\log n)$ (or $O(\log n)$ if $m = poly(n)$). Regarding the multicast game, we improve the PoA to O(1).

Posted Prices. It is a very common practice, especially in large markets and double auctions, for sellers to use posted prices. More closely to cost-sharing games is the model proposed by Kelly [25] regarding *bandwidth allocation*. Kelly's mechanism processes players' willingness to pay and posts a price for the whole bandwidth. Then each player pays a price proportional to the bandwidth she uses. This can be seen as pricing an infinitesimal quantity of 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*. 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. Our objective is to design cost-sharing protocols that are *BBiE* and the worst-case 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. For the multicast game, the algorithm of [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 PoA= $\Omega(n)$ in the complete information case; we refer the reader to the full version of the paper for this 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. We refer the reader to the full version of the paper for this reduction. 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 $\Omega(\sqrt{n})$ (see Table 1.1 for a summary of the main results).

BBiE Protocols (Sect. 4). For the Bayesian (and stochastic) set cover game there exists an *ex-post*[3] BBiE protocol (determined in polynomial time) with PoA $O(\log n)$, if $m = poly(n)$, 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 stochastic) settings, ex-post BBiE cannot be obtained by anonymous prices. Hence, we examine prices that are *ex-ante* BBiE. In the full version of the paper, 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. We stress that oblivious solutions may not be sufficient to guarantee low PoA for anonymous posted prices, in contrast to the BBiE protocols. This is because it is not clear anymore how to enforce players to choose desirable strategies, since *anonymous* prices are available to anyone. The reason that they exist here is due to the specific properties of the oblivious solution.

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.

Table 1. PoA of budget-balanced and BBiE protocols.

	BB protocols		BBiE protocols/posted prices	
	Set cover	Undirected	Set cover	Undirected
Full information	$\Theta(n)$ [6]	$O(1)$ [6]	1	1
Bayesian	$\Omega(n)$	$\Omega(\sqrt{n})$	$O(n/\log n)$	$O(1)$

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.

[3] 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.

In the full version of the paper we further study the complete information setting (see Table 1.1). Due to lack of space, we refer the reader to the full version of the paper for all the missing proofs.

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 Anshelevich et al. [2]. They 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 well studied NP-complete problems. 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 or 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, $R = \{r_1, \ldots, r_m\}$. 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 ξ. A cost-sharing protocol Ξ decides a cost-sharing method for each resource. In accordance with previous works, [6], [10], [13], the following are some natural properties that Ξ 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 $\pi : R \to [0, 1]$, 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 π_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:

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

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

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 π is the uniform distribution over U. Then the probability that element i is drawn as the type of at least one player is $q_i = 1 - \left(1 - \frac{1}{n}\right)^n \geq 1 - \frac{1}{e}$. 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. □

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

Proof. Consider the graph of Fig. 1. We set $p = 1 - \left(1 - \frac{1}{\sqrt{n}}\right)^{\frac{1}{n}}$, 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 least

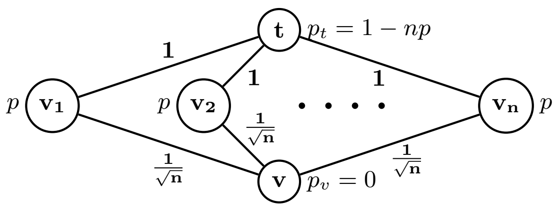

Fig. 1. Lower bound on the PoA of any budget-balanced protocol.

$\frac{2}{\sqrt{n}} + (1 - q_j) = 1 + \frac{1}{\sqrt{n}}$, 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 $\frac{1}{2}\sqrt{n}$. □

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 Π the underlying optimization resource allocation problem. Given any oblivious algorithm of Π with approximation ratio ρ, 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 *unweighted* problem which is *length-oblivious*), which are almost $O(\log mn)$-competitive. Theorem 3 implies the following corollary.

Corollary 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 $m = poly(n)$, 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 π^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 corollary holds.

Corollary 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. Ex-post BBiE cannot be obtained by using anonymous posted prices. Instead, we require *ex-ante* BBiE. For the rest of the section we define 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:

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

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

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 $k \geq 2$, $\frac{a}{1-(1-a)^k} > \frac{b}{1-(1-b)^k}$.*

Proposition 7. *In the* unweighted *set cover game, there exist* length-oblivious *and* anonymous *prices (computed in polynomial time) with PoA $O(\log n)$, if $m = poly(n)$, 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 $\frac{1}{k^1_{F \cap U}}$; 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 $\frac{k \sum_{i \in A} \pi_i}{1-\left(1-\sum_{i \in A} \pi_i\right)^k} > \frac{k \sum_{i \in B} \pi_i}{1-\left(1-\sum_{i \in B} \pi_i\right)^k}$, 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. As for the PoA, Grandoni et al. [20] analyzed the performance of Algorithm 1, for the stochastic problem. They didn't consider any prices, instead they mapped each player to the first set considered by the algorithm and they used the mapping in order to form a set cover. Their cover though coincide with the equilibrium solution and therefore their results immediately provide bounds on the PoA. □

Proposition 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* $m = poly(n)$*, and* $O\left(\frac{\log nm}{\log\log m - \log\log n}\right)$*, if* $m \gg n$*.*

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

Multicast Game. We construct a spanning tree T in the same way as in Sect. 4 and we use it to set the posted prices (computed in polynomial time).

Proposition 10. *In the multicast game, there exist prices with 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 $c_e/k^1_{V(e)}$. 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 $k^1_{V(e)} c_e / k^1_{V(e)} = c_e$, if e is used, and 0 otherwise, resulting in ex-ante BBiE. □

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.

Proposition 11. *In the* weighted *set cover game, any prior-independent, ex-post BBiE protocol* $\Xi \in \mathcal{C}$ *has PoA=* $\Omega(\sqrt{n})$*.*

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

Given a BBiE, prior-independent protocol Ξ, suppose that there exists some F_j, $j \neq 0$, where Ξ 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 PoA= $\Omega(\sqrt{n})$.

Suppose now that Ξ 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}, Ξ doesn't charge $\sqrt{n}$ to i. Consider the prior distribution $D_2 = \pi^n$ with $\pi(0) = 1/2$ and $\pi(j) = 1/2n$ 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 $1 - \left(1 - \frac{1}{2n}\right)^n \geq 1 - \frac{2}{e}$, 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 PoA= $\Omega(\sqrt{n})$. □

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)