

Designing Cost-Sharing Methods for Bayesian Games

George Christodoulou · Stefano
Leonardi · Alkmini Sgouritsa

Received: date / Accepted: date

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, Cost-Sharing Games

An extended abstract of the current paper appears in [27].

G. Christodoulou
Computer Science Department, University of Liverpool, Liverpool, UK
E-mail: gchristo@liv.ac.uk

S. Leonardi
Department of Computer, Control and Management Engineering Antonio Ruberti, Sapienza
University of Rome, Rome, Italy
E-mail: leonardi@diag.uniroma1.it

A. Sgouritsa
Computer Science Department, University of Liverpool, Liverpool, UK
E-mail: salkmini@liv.ac.uk

Acknowledgements 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 G. Christodoulou was supported by EPSRC grants EP/M008118/1, EP/K01000X/1 and Royal Society Leverhulme Trust Senior Research Fellowship LT140046. 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”.

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 [23] initiated the *design* aspect, seeking for protocols that induce approximately efficient equilibria, *with low Price of Anarchy (PoA)* [55]. 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, $U = \{1, \dots, n\}$, 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*; each player's type associates her with a specific element of U and X corresponds to the set of players' types. 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 $|V| = n$ 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 [29], and the latter, as a *Bayesian game*, introduced by [48], 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 [23] 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 anyway. 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 [23]; 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 $m = \text{poly}(n)$). 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 [51] 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 [45]. 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 [42], 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 [23] for the multicast directed network games, resulting in $\text{PoA} = \Omega(n)$ in the complete information case; (see Appendix A for the reduction). For the stochastic or Bayesian setting, where players are *i.i.d.*, we show that the same lower bound holds. Regarding the multicast game, the PoA is $O(1)$ for the complete information case [23] and the stochastic case [29], [42]. However, we show that for the Bayesian setting there is a lower bound of $\Omega(\sqrt{n})$ (see Table 1 for a summary).

BBiE Protocols (Sect. 4). For the Bayesian (and stochastic) set cover game there exists an *ex-post*¹ BBiE protocol (determined in polynomial time) 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 $m \gg n$. An *ex-post* BBiE protocol also exists for the Bayesian multicast game resulting in constant PoA (see Table 2 for a summary).

¹ 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.

Posted Prices (Sect. 5). For the Bayesian (and stochastic) settings, ex-post BBiE cannot be obtained by anonymous prices. We first discuss limitations of other concepts, such as BBiE with “high” probability or bounded possible excess and deficit. Then, we examine prices that result in *ex-ante* BBiE. 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 are only able to 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. Finally, for the poly-time determinable prices, we give tight lower bounds.

	Budget-balanced protocols	
	Set cover	Multicast
Complete information	$\Theta(n)$ [23]	$O(1)$ [23]
Bayesian	$\Omega(n)$	$\Omega(\sqrt{n})$

Table 1 PoA of budget-balanced protocols

	BBiE protocols/posted prices	
	Set cover	Multicast
Complete information	1	1
Bayesian	$O(n/\log n)$	$O(1)$

Table 2 PoA of BBiE protocols

Prior-Independent Mechanisms (Sect. 6). 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.

Complete Information Setting (Sect. 7). We further study the complete information setting (see Tables 1 and 2). By using either BBiE protocols or *anonymous* posted prices, we enforce the strategy profile of the optimum solution. Note that, while trying to bound the PoA, computational issues are not of primary concern. However, if we stick to protocols that can be determined in polynomial time, we can upper bound the PoA of set cover and multicast games by $O(\log k)$ and 1.39, respectively, where k is the number of players. Moreover, we argue that there are *no* anonymous prices, computed in polynomial time, for the set cover game, with $\text{PoA} = o(\log k)$.

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 [57] studied cost-sharing games under mechanism design context; they characterized the budget-balanced and group strategyproof mechanisms and identify the one with minimum welfare loss. In similar context, other papers considered (group)strategy proof and efficient mechanisms and relaxed the budget-balanced constraint; Devanur, Mihail and Vazirani [34] and Immorlica, Mahdian and Mirrokni [50] studied the set cover game under this context showing positive and negative bounds on the fraction of the cost that is covered.

Regarding network design games, there is a long line of works mainly focusing on fair cost allocation (Shapley cost-sharing mechanism), originated by [5]. Anshelevich et al. [5] showed a tight $\Theta(\log k)$ bound on the PoS for directed networks, while for undirected networks several variants have been studied [11, 12, 22, 24, 35] but the exact value of PoS still remains an open problem. For multicast games, Li [56] proved an upper bound of $O(\log k / \log \log k)$, while for broadcast games, Fiat et al. [39] proved an $O(\log \log k)$ upper bound which was improved to constant due to Bilò, Flammini and Moscardelli [13]. The PoA of some special equilibria has been also studied in [17, 21].

Chen, Roughgarden and Valiant [23] 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 [62] for parallel links, [40, 43, 54] for weighted congestion games, [29, 59, 47] for network games, [41] for routing games and [52] for resource allocation. The Bayesian Price of anarchy was first studied in auctions by [26]; see also [58] for routing games, and [61] for the PoS of Shapley protocol in cost-sharing games.

Close in spirit to our work is the notion of Coordination Mechanisms [25] 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 problems under the objective of minimising the makespan [2, 7, 16, 49, 53] or minimising the sum of players' costs [1, 9, 33], as well as for simple routing games [28, 10].

Posted prices have been used for pricing in large markets. Kelso and Crawford [4] and Gul and Stacchetti [46] 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 [37], who showed that half of the social welfare can be achieved. In a follow-up work [38], they considered Bayesian combinatorial auctions and they could guarantee half of the optimum welfare, by using anonymous posted prices. Dynamic pricing schemes has been used by Cohen, Eden, Fiat and Jez [31] in several online settings to induce the same performance as the best online algorithm, and by Cohen-Addad, Eden, Feldman and Fiat [32] in matching markets in order to achieve the optimal social welfare, for any tie breaking

rule. For maximizing the revenue with posted price mechanisms see [6, 8, 14, 18–20, 3].

We further discuss some related work to the underlying problems that we consider here, the set cover and the minimum Steiner tree problems. Both problems are very well studied and known to be in NP-complete. The best known approximations are $O(\log(k))$ [30] (by using a simple greedy algorithm) and 1.39 [15]; in fact, for the set cover problem, Feige [36] showed that no improvement by a constant factor is likely. Research has been done regarding the stochastic model, Grandoni et al. [45] showed a roughly $O(\log nm)$ tight bound for the set cover problem and Garg et al. [42] gave bounds on the approximation of the stochastic online Steiner tree problem. A slightly different distribution is the independent activations. Shmoys and Talwar [60] demonstrated randomized and deterministic algorithms with constant approximations for the universal TSP problem, and Christodoulou and Sgouritsa [29] studied the multicast game presenting an ordered protocol with constant PoA.

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 ξ . A cost-sharing protocol Ξ decides a cost-sharing method for each resource. In accordance with previous works, [23, 29, 62], the following are some natural properties that Ξ needs to satisfy:

- *Stability:* The induced game has always a *pure* (Bayes) 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 game or V in the multicast game. 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/strategy profile \mathbf{s} as $c(\mathbf{s})$. Then, the *Price of Anarchy* (PoA) for the complete information, stochastic and Bayesian settings is defined, respectively, as:

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

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

3 Lower Bounds for Budget-Balanced Protocols

In this section, we show the lower bounds of budget-balanced protocols, for the Bayesian setting.

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. \square

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) .




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

Indeed, if player i uses any other path (v_i, v, v_j, t) her cost-share will be at 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 respectively

$$\mathbb{E}[SC] = \sum_i q_i = \sqrt{n} ; \quad \mathbb{E}[Opt] \leq \sum_i q_i \cdot \frac{1}{\sqrt{n}} + 1 = n \frac{1}{n} + 1 = 2 .$$

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 sufficient that $h > \max_j c_{F_j}$ for the set cover game and $h > \sum_{e \in E} c_e$ for the multicast game. 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.

In an optimization problem, 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. We associate the types of the players to the input components of the problem, and then we would like to enforce the players to follow the action decided by the oblivious algorithm for their type.

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 $PoA = O(\rho)$.*

Proof Suppose that R_i is the set of the resources allocated by the oblivious algorithm to 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 Ξ 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 Ξ 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 [44]).

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 $PoA = O(\rho)$. Moreover, by the construction of Ξ , BBiE holds. \square

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

Set Cover Game. Grandoni et al. [45] 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-aware*), which are almost $O(\log mn)$ -competitive.

For completeness, we give the two algorithms of [45] in Algorithms 1 and 2. These algorithm are applied in the case of uniform distribution and then a reduction applies to generalize the results.

ALGORITHM 1: Algorithm 1 of [45]: 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 2: Algorithm 2 of [45]: 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

Theorem 3 implies the following corollary by using the results of [45].

Corollary 1 *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. [42] 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 [29], [60], [63]). 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 [15]), the following corollary holds.

Corollary 2 *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, as the following example illustrates. Instead, we require *ex-ante* BBiE. Furthermore, the same example serves to demonstrate that other natural variations of budget balance cannot be very promising: a) BBiE with “high” probability, b) bounded possible excess and deficit. Example 1 indicates that *any* anonymous posted prices may result in BBiE with probability at most $O(1/\sqrt{k})$ and 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 1 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.

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:

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

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

Set Cover Game. To determine anonymous prices for the unweighted set cover game, we first state Lemma 1 to be used in stability arguments.

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

Proof 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 $\frac{1-(1-a)^k}{a} < \frac{1-(1-b)^k}{b}$,

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

□

Theorem 4 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 [30] and at each step we set the price for the selected set. Algorithm 3 describes this procedure.

ALGORITHM 3: 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 3 and covers her. For that it is sufficient to show that for any two sets A and B , such that $\sum_{i \in A} \pi_i > \sum_{i \in B} \pi_i$, $k_A^1 > k_B^1$; so, if some player is covered by both A and B , the price set for A should be less than the price set for B and the player prefers A that is picked by Algorithm 3 before B .

By using (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 for $k \geq 2$, due to Lemma 1 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 in the game and this is a trivial case to solve.

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. [45] analyzed the performance of Algorithm 3, 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 coincides with the equilibrium solution and therefore their results immediately provide bounds on the PoA.

Theorem 5 *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 $m \gg n$.*

Proof By using the mapping of Algorithm 2 (Algorithm 2 of [45]), 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 she is mapped by Algorithm 2. Grandoni et al. [45] 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 satisfy ex-ante BBiE, for the same reasons as the unweighted case. \square

We complement our results by providing tight lower bounds for poly-time determinable prices.

Proposition 1 *For $m = \text{poly}(n)$, there may not exist anonymous posted prices for the unweighted set cover, or semi-anonymous posted prices for the weighted set cover, computed in poly-time, with PoA = $o(\log n)$, unless $NP \subseteq DTIME(n^{O(\log \log n)})$.*

Proof 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 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 [36], 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 in a contradiction. \square

Proposition 2 *For $m \gg n$, there may not exist anonymous prices for unweighted set cover, or semi-anonymous prices for weighted set cover, with PoA = $o\left(\frac{\log m}{\log \log m - \log \log n}\right)$.*

Proof 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 [45] (Theorem 4.2). \square

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).

Theorem 6 *In the multicast game, there exist anonymous posted 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 $c_e/k_{V(e)}^1$. 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 [42] and the approximation of [15]. 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.

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 7 *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 $U = \{0, 1, \dots, n\}$ 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 a 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 at least one player from S such that, whenever her type is j , she chooses F_{all} (and this happens with probability $1/2$) which results in $\text{PoA} = \Omega(\sqrt{n})$.

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. Therefore, for any player $i \in I$, whenever her type is j , she chooses F_j .

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 balance in F_{t_i}), whereas if she deviates to F_{all} she pays \sqrt{n} .

Each element $j \neq 0$ is a type of at least one player with probability $1 - (1 - \frac{1}{2n})^n \geq 1 - \frac{2}{e}$, 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 $\text{PoA} = \Omega(\sqrt{n})$.

Posted Prices. Setting prior-independent posted prices cannot guarantee any BBiE, 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 $\Theta(\sqrt{n})$.

7 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 3 *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 8 *There exist posted prices satisfying BBE, 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 Π . 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 Π chooses the sets in the same order with Algorithm 4. Therefore, the approximation ratio for Π 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, that results in BBiE as desired. \square

Proposition 3 *There may not exist posted prices, that are computed in polynomial time, for the 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 [36]. \square

Theorem 9 *There exist posted prices satisfying BBiE, computed in exponential time, for the (weighted) set-cover game, with PoA = 1.*

Proof As in the proof of Theorem 8, 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}^*) . We set the prices for the rest of the sets $\mathcal{F} \setminus \mathcal{F}^*$ equal to h .

By using similar arguments as in the proof of Theorem 8, 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 equal the cost of the set resulting in BBiE. 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 [15].

Corollary 4 *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 ≤ 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 [15], the PoA is constant even if we require the prices to be computed in polynomial time.

Theorem 10 *For the multicast game, there exist posted prices, computed in polynomial time, with PoA ≤ 1.39 and posted prices, computed in exponential time, with PoA = 1.*

Proof Let $S \subseteq V$ 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 [15]. 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 of at least h . Obviously, the players cover exactly the cost of each used edge and the cost of the Nash equilibrium equals the cost of T . \square

References

1. Fidaa Abed, José R. Correa, and Chien-Chung Huang. Optimal coordination mechanisms for multi-job scheduling games. In *Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings*, pages 13–24, 2014.
2. Fidaa Abed and Chien-Chung Huang. Preemptive coordination mechanisms for unrelated machines. In *Algorithms - ESA 2012 - 20th Annual European Symposium, Ljubljana, Slovenia, September 10-12, 2012. Proceedings*, pages 12–23, 2012.
3. Marek Adamczyk, Allan Borodin, Diodato Ferraioli, Bart de Keijzer, and Stefano Leonardi. Sequential posted price mechanisms with correlated valuations. In *Web and Internet Economics - 11th International Conference, WINE 2015, Amsterdam, The Netherlands, December 9-12, 2015, Proceedings*, pages 1–15, 2015.
4. Vincent P. Crawford Alexander S. Kelso. Job matching, coalition formation, and gross substitutes. *Econometrica*, 50(6):1483–1504, 1982.
5. Elliot Anshelevich, Anirban Dasgupta, Jon M. Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
6. Elliot Anshelevich, Koushik Kar, and Shreyas Sekar. Envy-free pricing in large markets: Approximating revenue and welfare. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 52–64, 2015.
7. Yossi Azar, Kamal Jain, and Vahab Mirrokni. (Almost) optimal coordination mechanisms for unrelated machine scheduling. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '08*, pages 323–332, Philadelphia, PA, USA, 2008. Society for Industrial and Applied Mathematics.
8. Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. In *Innovations in Computer Science - ICS 2010, Tsinghua University, Beijing, China, January 7-9, 2011. Proceedings*, pages 166–178, 2011.
9. Sayan Bhattacharya, Sungjin Im, Janardhan Kulkarni, and Kamesh Munagala. Coordination mechanisms from (almost) all scheduling policies. In *Innovations in Theoretical Computer Science, ITCS'14, Princeton, NJ, USA, January 12-14, 2014*, pages 121–134, 2014.
10. Sayan Bhattacharya, Janardhan Kulkarni, and Vahab S. Mirrokni. Coordination mechanisms for selfish routing over time on a tree. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I*, pages 186–197, 2014.
11. Vittorio Bilò and Roberta Bove. Bounds on the price of stability of undirected network design games with three players. *Journal of Interconnection Networks*, 12(1-2):1–17, 2011.
12. Vittorio Bilò, Ioannis Caragiannis, Angelo Fanelli, and Gianpiero Monaco. Improved lower bounds on the price of stability of undirected network design games. *Theory Comput. Syst.*, 52(4):668–686, 2013.

13. Vittorio Bilò, Michele Flammini, and Luca Moscardelli. The price of stability for undirected broadcast network design with fair cost allocation is constant. In *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, 26-29 October, 2013, Berkeley, CA, USA*, pages 638–647, 2013.
14. Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *Proceedings 9th ACM Conference on Electronic Commerce (EC-2008), Chicago, IL, USA, June 8-12, 2008*, page 49, 2008.
15. Jaroslaw Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. An improved lp-based approximation for steiner tree. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 583–592, 2010.
16. Ioannis Caragiannis. Efficient coordination mechanisms for unrelated machines scheduling. In *SODA*, pages 815–824, 2009.
17. Moses Charikar, Howard J. Karloff, Claire Mathieu, Joseph Naor, and Michael E. Saks. Online multicast with egalitarian cost sharing. In *SPAA*, pages 70–76, 2008.
18. Shuchi Chawla, Jason D. Hartline, and Robert D. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings 8th ACM Conference on Electronic Commerce (EC-2007), San Diego, California, USA, June 11-15, 2007*, pages 243–251, 2007.
19. Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320, 2010.
20. Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
21. Chandra Chekuri, Julia Chuzhoy, Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. Non-cooperative multicast and facility location games. *IEEE Journal on Selected Areas in Communications*, 25(6):1193–1206, 2007.
22. Ho-Lin Chen and Tim Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.
23. Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing network protocols for good equilibria. *SIAM J. Comput.*, 39(5):1799–1832, 2010.
24. George Christodoulou, Christine Chung, Katrina Ligett, Evangelia Pyrga, and Rob van Stee. On the price of stability for undirected network design. In *Approximation and Online Algorithms, 7th International Workshop, WAOA 2009, Copenhagen, Denmark, September 10-11, 2009. Revised Papers*, pages 86–97, 2009.
25. George Christodoulou, Elias Koutsoupias, and Akash Nanavati. Coordination mechanisms. *Theor. Comput. Sci.*, 410(36):3327–3336, 2009.
26. George Christodoulou, Annamária Kovács, and Michael Schapira. Bayesian combinatorial auctions. *J. ACM*, 63(2):11:1–11:19, 2016.
27. George Christodoulou, Stefano Leonardi, and Alkmini Sgouritsa. Designing cost-sharing methods for bayesian games. In *Algorithmic Game Theory - 9th International Symposium, SAGT 2016, Liverpool, UK, September 19-21, 2016. Proceedings*, pages 327–339, 2016.
28. George Christodoulou, Kurt Mehlhorn, and Evangelia Pyrga. Improving the price of anarchy for selfish routing via coordination mechanisms. *Algorithmica*, 69:619–640, 2014.
29. George Christodoulou and Alkmini Sgouritsa. Designing networks with good equilibria under uncertainty. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 72–89, 2016.
30. V. Chvatal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
31. Ilan Reuven Cohen, Alon Eden, Amos Fiat, and Lukasz Jez. Pricing online decisions: Beyond auctions. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 73–91, 2015.
32. Vincent Cohen-Addad, Alon Eden, Michal Feldman, and Amos Fiat. The invisible hand of dynamic market pricing. *CoRR*, abs/1511.05646, 2015.

33. Richard Cole, José R. Correa, Vasilis Gkatzelis, Vahab S. Mirrokni, and Neil Olver. Inner product spaces for minsum coordination mechanisms. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 539–548, 2011.
34. Nikhil R. Devanur, Milena Mihail, and Vijay V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location games. *Decision Support Systems*, 39(1):11–22, 2005.
35. Yann Disser, Andreas Emil Feldmann, Max Klimm, and Matús Mihalák. Improving the H_k -bound on the price of stability in undirected shapley network design games. In *Algorithms and Complexity, 8th International Conference, CIAC 2013, Barcelona, Spain, May 22-24, 2013. Proceedings*, pages 158–169, 2013.
36. Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
37. Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial walrasian equilibrium. In *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 61–70, 2013.
38. Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135, 2015.
39. Amos Fiat, Haim Kaplan, Meital Levy, Svetlana Olonetsky, and Ronen Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, pages 608–618, 2006.
40. Martin Gairing, Konstantinos Kollias, and Grammateia Kotsialou. Tight bounds for cost-sharing in weighted congestion games. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part II*, pages 626–637, 2015.
41. Martin Gairing, Konstantinos Kollias, and Grammateia Kotsialou. Cost-sharing in generalised selfish routing. In *Algorithms and Complexity - 10th International Conference, CIAC 2017, Athens, Greece, May 24-26, 2017, Proceedings*, pages 272–284, 2017.
42. Naveen Garg, Anupam Gupta, Stefano Leonardi, and Piotr Sankowski. Stochastic analyses for online combinatorial optimization problems. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2008, San Francisco, California, USA, January 20-22, 2008*, pages 942–951, 2008.
43. Vasilis Gkatzelis, Konstantinos Kollias, and Tim Roughgarden. Optimal cost-sharing in weighted congestion games. In *Web and Internet Economics - 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, pages 72–88, 2014.
44. Ragavendran Gopalakrishnan, Jason R. Marden, and Adam Wierman. Potential games are necessary to ensure pure nash equilibria in cost sharing games. In *EC*, pages 563–564, 2013.
45. Fabrizio Grandoni, Anupam Gupta, Stefano Leonardi, Pauli Miettinen, Piotr Sankowski, and Mohit Singh. Set covering with our eyes closed. *SIAM J. Comput.*, 42(3):808–830, 2013.
46. Faruk Gul and Ennio Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95 – 124, 1999.
47. Tobias Harks, Anja Huber, and Manuel Surek. A characterization of undirected graphs admitting optimal cost shares. *CoRR*, abs/1704.01983, 2017.
48. John C. Harsanyi. Games with incomplete information played by "bayesian" players, i-iii. part ii. bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.
49. Nicole Immorlica, Li Li, Vahab S. Mirrokni, and Andreas Schulz. Coordination mechanisms for selfish scheduling. In *Internet and Network Economics, First International Workshop (WINE)*, pages 55–69, 2005.
50. Nicole Immorlica, Mohammad Mahdian, and Vahab S. Mirrokni. Limitations of cross-monotonic cost-sharing schemes. *ACM Transactions on Algorithms*, 4(2), 2008.
51. Frank Kelly. Charging and rate control for elastic traffic. *Eur. Trans. Telecomm.*, 8(1):3337, January 1997.

52. Max Klimm and Daniel Schmand. Sharing non-anonymous costs of multiple resources optimally. In *Algorithms and Complexity - 9th International Conference, CIAC 2015, Paris, France, May 20-22, 2015. Proceedings*, pages 274–287, 2015.
53. Konstantinos Kollias. Non-preemptive coordination mechanisms for identical machine scheduling games. In *SIROCCO*, pages 197–208, 2008.
54. Konstantinos Kollias and Tim Roughgarden. Restoring pure equilibria to weighted congestion games. *ACM Trans. Economics and Comput.*, 3(4):21:1–21:24, 2015.
55. Elias Koutsoupias and Christos H. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
56. Jian Li. An $o(\log(n)/\log(\log(n)))$ upper bound on the price of stability for undirected shapley network design games. *Inf. Process. Lett.*, 109(15):876–878, 2009.
57. Hervé Moulin and Scott Shenker. Strategyproof sharing of submodular costs: budget balance versus efficiency. *Economic Theory*, 18(3):511–533, 2001.
58. Tim Roughgarden. The price of anarchy in games of incomplete information. *ACM Trans. Economics and Comput.*, 3(1):6, 2015.
59. Tim Roughgarden and Okke Schrijvers. Network cost-sharing without anonymity. *ACM Trans. Economics and Comput.*, 4(2):8:1–8:24, 2016.
60. David B. Shmoys and Kunal Talwar. A constant approximation algorithm for the a priori traveling salesman problem. In *Integer Programming and Combinatorial Optimization, 13th International Conference, IPCO 2008, Bertinoro, Italy, May 26-28, 2008, Proceedings*, pages 331–343, 2008.
61. Vasilis Syrgkanis. Price of stability in games of incomplete information. *CoRR*, abs/1503.03739, 2015.
62. Philipp von Falkenhausen and Tobias Harks. Optimal cost sharing for resource selection games. *Math. Oper. Res.*, 38(1):184–208, 2013.
63. David P. Williamson and Anke van Zuylen. A simpler and better derandomization of an approximation algorithm for single source rent-or-buy. *Oper. Res. Lett.*, 35(6):707–712, 2007.

APPENDIX

A 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. [23] (Proposition 4.12), the following corollary can be trivially derived.

Corollary 5 *The PoA of the (unweighted) set cover cost-sharing game for the complete information setting is $\Omega(n)$*