Worst-Case Mechanism Design via Bayesian Analysis

Bei, Xiaohui; Chen, Ning; Gravin, Nick; Lu, Pinyan

2017

Bei, X., Chen, N., Gravin, N., & Lu, P. (2017). Worst-Case Mechanism Design via Bayesian Analysis. *SIAM Journal on Computing*, 46(4), 1428-1448.

<https://hdl.handle.net/10356/86533>

<https://doi.org/10.1137/16M1067275>

© 2017 Society for Industrial and Applied Mathematics (SIAM). This paper was published in *SIAM Journal on Computing* and is made available as an electronic reprint (preprint) with permission of Society for Industrial and Applied Mathematics (SIAM). The published version is available at: [<http://dx.doi.org/10.1137/16M1067275>]. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper is prohibited and is subject to penalties under law.

Downloaded on 26 Apr 2023 21:02:01 SGT

WORST-CASE MECHANISM DESIGN VIA BAYESIAN ANALYSIS*

XIAOHUI BEI[†], NING CHEN[†], NICK GRAVIN[‡], AND PINYAN LU[§]

Abstract. Budget feasible mechanism design is the study of procurement combinatorial auctions in which the sellers have private costs to produce items, and the buyer (auctioneer) aims to maximize her valuation function on a subset of purchased items under the budget constraint on the total payment. One of the most important questions in the field is “which valuation domains admit truthful budget feasible mechanisms with ‘small’ approximations to the social optimum?” Singer [*Proceedings of the 51st FOCS*, IEEE Press, Piscataway, NJ, 2010, pp. 765–774] showed that submodular functions have a constant approximation mechanism. Dobzinski, Papadimitriou, and Singer [*Proceedings of the 12th ACM Conference on Electronic Commerce*, ACM, New York, 2011, pp. 273–282] gave an $O(\log^2 n)$ approximation mechanism for subadditive functions and remarked that “A fundamental question is whether, regardless of computational constraints, a constant-factor budget feasible mechanism exists for subadditive functions.” In this paper, we give an affirmative answer to this question. To this end we relax the prior-free mechanism design framework to the Bayesian mechanism design framework (these are two standard approaches from computer science and economics, respectively). Then we convert our results in the Bayesian setting back to the prior-free framework by employing Yao’s minimax principle. Along the way, we obtain the following results: (i) a polynomial time constant approximation for XOS valuations (a.k.a. fractionally subadditive valuations, a superset of submodular functions), (ii) a polynomial time $O(\log n / \log \log n)$ -approximation for general subadditive valuations, (iii) a constant approximation for general subadditive functions in the Bayesian framework—we allow correlation in the distribution of sellers’ costs and provide a universally truthful mechanism, (iv) the existence of a prior-free constant approximation mechanism via Yao’s minimax principle.

Key words. mechanism design, budget feasible, prior-free, Bayesian

AMS subject classifications. 68W25, 91A80

DOI. 10.1137/16M1067275

1. Introduction. Consider a procurement auction problem where there is a buyer who wants to purchase resources from a set of agents A . Each agent $i \in A$ is able to supply a resource at an incurred cost $c(i)$. The buyer has a budget B on the compensation that could be distributed among the agents, and a function $v(\cdot)$ that describes the value of the buyer for each subset of A . This defines a natural optimization problem: find a subset $S \subseteq A$ that maximizes $v(S)$ subject to $\sum_{i \in S} c(i) \leq B$. The budgeted optimization problem has been considered in a variety of domains with respect to different valuation functions, e.g., additive (a.k.a. knapsack) and submodular [48].

The suppliers, as self-interested agents, may want to get as much compensation as possible. In particular, agent i may conceal his true incurred cost $c(i)$ (which is known only to himself) and claim any amount $b(i)$ instead. Thus we face an additional challenge of dealing with selfish and strategic behavior of the agents. To cope with this problem we adopt a mechanism design approach: given submitted bids $b(i)$ from

*Received by the editors March 23, 2016; accepted for publication (in revised form) April 24, 2017; published electronically August 17, 2017. Results (i)–(iii) were presented at STOC 12 [8]; result (iv) is new.

<http://www.siam.org/journals/sicomp/46-4/M106727.html>

[†]Nanyang Technological University, Singapore, Singapore 637371 (xhbei@ntu.edu.sg, ningc@ntu.edu.sg).

[‡]Shanghai University of Finance and Economics, Shanghai, China (nikolai@mail.shufe.edu.cn).

[§]Corresponding author. Shanghai University of Finance and Economics, Shanghai, China (lu.pinyan@mail.shufe.edu.cn).

all agents, a *mechanism* determines the winning set $S \subseteq A$ and the payment $p(i)$ to each winner $i \in S$. A mechanism is called *truthful* (a.k.a. incentive compatible) if it is in the best interest of every agent to bid his true cost, i.e., claim $b(i) = c(i)$. Truthfulness is one of the central solution concepts in mechanism design, and indeed might be desirable in the context of procurement auctions.

Our mechanism design problem has an important and practical ingredient: the budget; i.e., the total payment of a mechanism should not exceed B . In single-parameter domains where the private information of every individual is a single value (which is true in our case), a monotone allocation rule with associated threshold payments provides a sufficient and necessary condition for truthfulness [3, 39]. However, it may not necessarily generate a budget feasible solution. A number of well-known truthful designs, such as the seminal Vickrey–Clarke–Groves (VCG) mechanism [19, 31, 49], do not apply anymore, and new ideas have to be developed.

Another unavoidable issue caused by the budget constraint is that, unlike when using the VCG mechanism, which always generates a socially optimal solution, we cannot hope to have a solution that is both socially optimal and budget feasible. Indeed, for many simple valuations that exhibit complements [46], the approximation to the optimum of any budget feasible mechanism can be arbitrarily bad. The following simple example illustrates the problem with non-complement-free valuations. Consider the objective in which there is a crucial item i^* and where the valuation $v(S)$ is $|S|$ if $i^* \in S$ and 0 otherwise. It is easy to see that no approximation guarantee can be achieved for this valuation, as the seller of the critical item i^* can extract the whole budget (publicly known) from the auctioneer. In this sense the complement-free valuations are the most general class of functions for which we can hope to obtain budget feasible mechanisms. This impossibility result naturally raises the following question: which valuation domains admit truthful budget feasible mechanisms that are close to the socially optimal solution? The answer to this question depends on the properties of the valuation functions under consideration. Since valuations with complements are not well aligned with the budget constraint, the research in this area has always been focused on complement-free valuations. In particular, given the following hierarchy of valuation classes [36],

$$\text{additive} \subset \text{gross substitutes} \subset \text{submodular} \subset \text{XOS} \subset \text{subadditive},$$

which ones admit constant approximation budget feasible mechanisms?

Singer [46] initiated the study of approximate budget feasible mechanism design and gave constant approximation mechanisms for additive and submodular valuations. In subsequent work, Dobzinski, Papadimitriou, and Singer [25] considered subadditive functions and showed an $O(\log^2 n)$ approximation. They also remarked in [25]:

A fundamental question is whether, regardless of computational constraints, a constant-factor budget feasible mechanism exists for subadditive functions.

In the present paper, we give a positive answer to this question, albeit we don't obtain an explicit mechanism, but rather show its existence.

1.1. Our results and techniques. Our final result is the following.

THEOREM 1. *There exists a truthful budget feasible mechanism for subadditive functions with a constant approximation ratio.*

In Theorem 1, the approximation guarantee must hold for every instance of the costs, i.e., in the worst case over all possible inputs. Such prior-free worst-case analysis

was used in all previous works on budget feasible mechanism design [17, 25, 29, 46] and is the standard framework in this particular area and broadly in theoretical computer science. However, to obtain our main result, as an intermediate step we cast the problem into the Bayesian framework.

Bayesian mechanism design. The Bayesian analysis framework [39] is a standard approach in economics and game theory. In the Bayesian analysis it is assumed that the agents' private information ($c(i)$'s in our model) is drawn from a fixed distribution and the performance is measured ex ante, i.e., in expectation over this distribution. Bayesian analysis offers a realistic framework that often provides more flexibility than analysis in the worst case. Bayesian mechanism design, in particular, has received significant attention in the computer science community in the past few years; see, e.g., [9, 10, 13, 14, 15, 20, 26, 32, 33, 34, 35]. The central result of this paper is the following theorem.

THEOREM 2. *There exists a constant approximation truthful budget feasible Bayesian mechanism for subadditive functions for any given distribution of the costs.¹*

Theorem 2 is a weaker version of Theorem 1, as any approximation result in the worst case implies at least as good approximation guarantees in the Bayesian framework. Nevertheless, this result serves as a useful and necessary step towards understanding the problem and obtaining our final result in the prior-free framework. Furthermore, it should be noted that our result does not rely on the common assumptions of Bayesian analysis in the following aspects.

- *Notion of truthfulness.* In most of the previous work on Bayesian mechanism design with a focus on social welfare maximization, e.g., [9, 14, 32, 33], the considered solution concept is *Bayesian truthfulness*. That is, truth-telling is in expectation a dominant strategy when other agents' profiles are drawn from the given prior distribution. Our mechanism guarantees *universal truthfulness*, meaning that truth-telling is the best that an agent can do for any instance of other agents' reported costs and for any coin flips of the mechanism. Stronger than Bayesian truthfulness, universal truthfulness has already been considered in Bayesian mechanism design prior to our work, e.g., in [13], however, with a different focus on profit maximization.
- *Distributional assumptions.* The vast majority of previous related work considers only independent distributions, e.g., [9, 18, 32, 33, 34]. Our results apply to distributions that allow correlations of the costs. Correlation among different pieces of private information is a natural phenomenon arising in practice and has been considered in optimal auction design [23, 37, 42, 43, 44]. In our setting, correlation among private costs of different agents appears to be well motivated. For example, if the price on the crude oil goes up, the production costs go up as well simultaneously for every agent.

Back to prior-free mechanism design. After taking a detour to the Bayesian framework, we return to the prior-free framework and show that our Bayesian result implies the existence of a constant approximation truthful prior-free mechanism. It is worth noting that this result comes from a rather general observation: In an arbitrary problem setting, suppose that we have a mechanism with a certain approximation guarantee in the Bayesian framework which both (i) is universally truthful and (ii) allows for arbitrary correlated distributions. Then, using *Yao's minimax*

¹The result holds for any distribution with a finite support. For distributions with infinite support, we require a certain technical integrability assumption for any subset of individual costs.

principle, one can always render guarantees in the Bayesian framework into the same approximation guarantees in the worst-case framework.

Techniques. In the design of budget feasible mechanisms, the major approach used in previous works [17, 25, 46] is based on a simple idea of selecting winning agents sequentially and greedily. The main technical challenge in these methods is to define a good selection rule which ensures that (i) the final selected set of winners has a value close to the optimum and (ii) respective threshold payments are budget feasible. Our mechanisms, from a high level structural point of view, use a different set of ideas and methods.

First, we employ the idea of random sampling, a standard tool from the mechanism design literature on profit maximization [30]. In our setting random sampling works as follows. We sample uniformly at random a test set T and find the (approximately) optimal budget feasible solution on T . This solution gives us a close estimate on the value of the optimal solution for the remaining agents with high probability. We use set T only for “evaluation” purposes and, therefore, eliminate all incentive issues with the agents within. Using the estimate from T as a threshold, we then select the final set of winners from the remaining agents $A \setminus T$.

Second, we consider the optimization problem $\max_{S \subset A \setminus T} v(S)$ s.t. $c(S) \leq B$ and observe several remarkable incentive and approximation properties of the corresponding Lagrangian function $L(X, t) = v(X) - t \cdot c(X)$. The estimated value from a random sample T helps us to calibrate the parameter t in $L(X, t)$. For the class of fractionally subadditive (a.k.a. XOS) functions this allows us to show the following.

THEOREM 3. *There is a budget feasible truthful mechanism for XOS functions with a constant approximation ratio.*

To obtain this result we rely on the characteristic property of XOS functions, i.e., that any XOS function is a maximum over a collection of linear functions. We reduce further the set of potential winners to $S^* = \operatorname{argmax}_{X \subset A \setminus T} L(X, t)$ and substitute $v(\cdot)$ with a linear supporting function f for the set S^* . We then select the final winning set from S^* by running a truthful budget feasible mechanism for the additive valuation f .

Third, for subadditive valuations, we make use of the extra knowledge about prior distribution in the Bayesian setting. In particular, we generate a cost vector according to our prior distribution and use it as threshold payments for the prospective winners. Our choice of the threshold vector ensures certain symmetry between random thresholds and private costs, which in combination with subadditivity implies good approximation guarantees.² It is worth noting that the mechanism for subadditive valuations bears a resemblance to the prior-free mechanism for XOS valuations. In particular, in the Bayesian setting we still sample the test set T and estimate the optimum for $A \setminus T$ from this sample rather than from the prior distribution. Interestingly and perhaps surprisingly, while the latter approach works well when the private costs are drawn independently, it fails when costs are correlated (see relevant examples in section 4).

1.2. Computational aspects. In this paper, we mainly focus on the approximation ratio caused by the inherent nature of the truthfulness requirement, and computational issues are not our main concern. Due to the impossibility results for

²Results on the Bayesian price of anarchy in simultaneous auctions with subadditive bidders [28] use unilateral deviations where every single bidder, competing with a random price vector, submits a vector of random bids drawn from the same distribution as the random prices. This bidding strategy was inspired by our random-thresholds mechanism.

fractionally subadditive valuations with value queries by Singer [46], even without the truthfulness requirement, we are forced to focus on more powerful oracles when studying the computational complexity. Following are some claims regarding the computational complexity for our results:

- Theorem 3 actually implies a polynomial time mechanism if we assume XOS and demand oracles. As an example, the matching valuations, which are XOS but not submodular, have polynomial time algorithms with XOS and demand oracles.
- Theorem 2 also can be implemented in polynomial time if, in addition to the XOS and demand oracles, we have conditional sampling access to the distribution of cost vectors.
- Theorem 1 is a purely existential result and may not imply any computationally efficient mechanisms, as the existence of the mechanism is based on the minimax principle rather than an explicit construction.

We note that one can design a brute-force $O(\log n)$ approximation for the subadditive valuation v by running an XOS mechanism for the XOS approximation $\hat{v}$ of subadditive v , i.e., such an XOS function $\hat{v}$ that $\hat{v}(S) \leq v(S) \leq O(\log n)\hat{v}(S)$ for any $S \subset A$. However, our following result shows that the $\log(n)$ barrier is not the right bound for the polynomial time mechanisms for subadditive valuations.

THEOREM 4. *There is a budget feasible truthful mechanism for subadditive valuations with an $O(\frac{\log n}{\log \log n})$ approximation ratio that runs in polynomial time with the access to demand oracle.*

1.3. Related work. Our work falls into the field of algorithmic mechanism design, which is a fascinating area initiated by the seminal work of Nisan and Ronen [40]. There are many mechanism design models (see, e.g., [41] for a survey).

As mentioned earlier, the study of approximate mechanism design with a budget constraint was originated by Singer [46], and constant approximation mechanisms were given for additive and submodular functions. The approximation ratios were later improved in [17]. Dobzinski, Papadimitriou, and Singer [25] considered subadditive functions and showed an $O(\log^2 n)$ approximation mechanism. A better approximation for a budget feasible mechanism is known when one has an additional “large markets” assumption [2]. Ghosh and Roth [29] considered a budget feasible mechanism design model for selling privacy, where there are externalities for each agent’s cost. All these models considered prior-free worst-case analysis.

For Bayesian mechanism design, Hartline and Lucier [33] first proposed a Bayesian reduction in single-parameter settings that converts any approximation algorithm to a Bayesian truthful mechanism that approximately preserves social welfare. The black-box reduction results were later improved to multiparameter settings in [9] and [32] independently. Chawla, Malec, and Malekian [14] considered budget-constrained agents and gave Bayesian truthful mechanisms in various settings. A number of other Bayesian mechanism design works considered profit maximization, e.g., [10, 13, 15, 20, 21, 34]. Ours is the first to consider Bayesian analysis in budget feasible mechanisms with a focus on the valuation (social welfare) maximization. In a subsequent work, Balkanski and Hartline [7] studied Bayesian budget feasible posted pricing mechanisms that approximate the value obtained by the Bayesian optimal mechanism.

2. Preliminaries. In a marketplace, there are n agents (or items), denoted by A , and a single buyer, who runs the auction. Each agent $i \in A$ incurs a privately known cost $c(i) \geq 0$ when delivering a service or selling his item to the buyer. We

denote by $c = (c(i))_{i \in A}$ the cost vector of the agents. For any given subset $S \subseteq A$, the buyer derives publicly known value $v(S)$. We assume that $v(\emptyset) = 0$ and that the valuation function is monotone; i.e., $v(S) \leq v(T)$ for any $S \subset T \subseteq A$. A centralized authority (the buyer) wants to pick a subset of agents with maximum possible value, given a budget B to compensate agents' incurred costs, i.e., $\max_{S \subseteq A} v(S)$ s.t. $c(S) = \sum_{i \in S} c(i) \leq B$. We denote the optimal solution to the latter optimization problem by $\text{opt}(A)$ (or $\text{opt}(c)$) and its value by $v(\text{opt}(A))$.

We consider XOS and subadditive valuations in the paper; both are rather general classes and contain a number of well studied functions as special cases, e.g., additive, gross substitutes, and submodular:

- Subadditive (a.k.a. complement-free): $v(S) + v(T) \geq v(S \cup T)$ for any $S, T \subseteq A$.
- XOS (a.k.a. fractionally subadditive, as defined in [27]): there is a set of linear functions $f_1, \dots, f_m$ such that

$$v(S) = \max \{f_1(S), f_2(S), \dots, f_m(S)\} \quad \text{for each } S \subseteq A.$$

Note that the number of functions m can be exponential in $n = |A|$. The class of XOS functions is equivalent to the class of fractionally subadditive functions [27], which is a proper subclass of subadditive valuations.

The representation of a subadditive or XOS valuation function has size exponential in n . Instead we assume that we are given a query access to a *demand oracle*, which, for any given price vector $p(1), \dots, p(n)$, returns a subset

$$T \in \operatorname{argmax}_{S \subseteq A} \left(v(S) - \sum_{i \in S} p(i) \right)$$

in $O(1)$ time. It was shown in [46] that a weaker value query oracle is not sufficient to obtain a constant approximation polynomial time mechanism even for XOS valuation, and thus later work [23] adopted stronger computational model with demand oracles. We also assume that we have access to an *XOS oracle* (as defined in [24]), which allows one to find a supporting linear function $f_{i(S)}$ for any specified set S in $O(1)$ time; i.e., for a given set S the oracle returns a linear function f_S such that $v(S) = f_S(S)$ and $v(T) \geq f_S(T)$ for any $T \subseteq A$.

Agents, as self-interested entities, have their own objective as well; each agent i may not tell his true privately known cost $c(i)$ but, instead, submit a *bid* $b(i)$ strategically. Upon receiving $b(i)$ from each agent, a mechanism determines an *allocation* $S \subseteq A$ of the winners and a *payment* $p(i)$ to each $i \in A$. We assume that the mechanism has no positive transfer (i.e., $p(i) = 0$ if $i \notin S$) and is individually rational (i.e., $p(i) \geq b(i)$ if $i \in S$).

In a mechanism, each agent bids strategically to maximize her utility, which is $p(i) - c(i)$ if i is a winner and 0 otherwise. We say a mechanism is *truthful* if it is in the best interest of each agent to report her true cost, i.e., $b(i) = c(i)$. For randomized mechanisms, we consider universal truthfulness in this paper: a randomized mechanism is called *universally truthful* if it is a distribution over deterministic truthful mechanisms.

Our model is in the single parameter domain. Thus, by the well-known characterization [3, 39], any truthful mechanism can be equivalently described as a monotone allocation rule with the corresponding threshold payments. We, therefore, may explicitly specify only the allocation rule of our mechanism, and usually omit description of the payments.

A mechanism is said to be *budget feasible* if its total payment is within the budget constraint, i.e., $\sum_i p(i) \leq B$. Our goal in this paper is to design truthful and budget feasible mechanisms for XOS and subadditive valuations in two frameworks: prior-free and Bayesian.

We first establish the following technical lemma about random sampling, which is used in many places in our analysis.

LEMMA 2.1. *Consider any subadditive function $v(\cdot)$. For a given subset $S \subseteq A$ and a positive integer k we assume that $v(S) \geq k \cdot v(i)$ for any $i \in S$. Further, suppose that S is divided uniformly at random into two groups T_1 and T_2 . Then, with probability of at least $\frac{1}{2}$, we have $v(T_1) \geq \frac{k-1}{4k}v(S)$ and $v(T_2) \geq \frac{k-1}{4k}v(S)$.*

Proof. We first claim that there are disjoint subsets S_1 and S_2 with $S_1 \cup S_2 = S$ such that $v(S_1) \geq \frac{k-1}{2k}v(S)$ and $v(S_2) \geq \frac{k-1}{2k}v(S)$. This can be seen by the following recursive process: Initially let $S_1 = \emptyset$ and $S_2 = S$; we move items from S_2 to S_1 one by one in an arbitrary order, until the point when $v(S_1) \geq \frac{k-1}{2k}v(S)$. Consider the S_1, S_2 at the end of the process; we claim that at this point, we also have $v(S_2) \geq \frac{k-1}{2k}v(S)$. Note that $v(S) \leq v(S_1) + v(S_2)$. Let i be the last item moved from S_2 to S_1 ; therefore, $v(S_1 \setminus \{i\}) < \frac{k-1}{2k}v(S)$, which implies that $v(S_2 \cup \{i\}) > \frac{k+1}{2k}v(S)$. Thus, $v(S_2) + v(i) \geq v(S_2 \cup \{i\}) > \frac{k+1}{2k}v(S)$. As $v(i) \leq \frac{1}{k}v(S)$, we know that $v(S_2) > \frac{k-1}{2k}v(S)$.

Consider sets $X_1 = S_1 \cap T_1$, $Y_1 = S_1 \cap T_2$, $X_2 = S_2 \cap T_1$, and $Y_2 = S_2 \cap T_2$. Due to subadditivity we have $\frac{k-1}{2k}v(S) \leq v(S_1) \leq v(X_1) + v(Y_1)$; hence, either $v(X_1) \geq \frac{k-1}{4k}v(S)$ or $v(Y_1) \geq \frac{k-1}{4k}v(S)$. Similarly, we have that either $v(X_2) \geq \frac{k-1}{4k}v(S)$ or $v(Y_2) \geq \frac{k-1}{4k}v(S)$. With probability $\frac{1}{2}$ the most valuable parts of S_1 's partition and S_2 's partition get into different sets T_1 and T_2 , respectively. Thus the lemma follows. $\square$

3. Prior-free mechanism design. In this section, we first consider designing budget feasible mechanisms for XOS functions in the prior-free setting. Here we evaluate a mechanism according to its *approximation ratio*, which is defined as $\max_c \frac{v(\text{opt}(c))}{\mathcal{M}(c)}$, where $\mathcal{M}(c)$ is the (expected) value of a mechanism $\mathcal{M}$ on instance $c = (c(i))_{i \in A}$ and $v(\text{opt}(c))$ is its optimal value. Without loss of generality, we assume that $c(i) \leq B$ for any $i \in A$.

We recall that XOS function $v(\cdot)$ is given by

$$v(S) = \max \{f_1(S), f_2(S), \dots, f_m(S)\} \quad \text{for any } S \subseteq A,$$

where each $f_j(\cdot)$ is a nonnegative additive function, i.e., $f_j(S) = \sum_{i \in S} f_j(i)$.

In our mechanism, we use randomized mechanism ADDITIVE-MECHANISM for additive valuation functions as an auxiliary procedure, where ADDITIVE-MECHANISM is a universally truthful mechanism and has an approximation factor of at most 3 (see, e.g., Theorem B.2 of [17]).

Algorithm 1 XOS-RANDOM-SAMPLING.

1. Pick each item i.i.d. with probability $\frac{1}{2}$ into group T .
 2. Compute an optimal solution $\text{opt}(T)$ for items in T given budget B .
 3. Set a threshold $t = \frac{v(\text{opt}(T))}{8B}$.
 4. Find a set $S^* \in \arg\max_{S \subseteq A \setminus T} \{v(S) - t \cdot c(S)\}$.
 5. Find additive function f s.t. $f(S^*) = v(S^*)$, $f(X) \leq v(X) \forall X \subseteq A$.
 6. Run ADDITIVE-MECHANISM for $f(\cdot)$ on the set S^* and budget B .
 7. Output the winning set of ADDITIVE-MECHANISM.
-

In the above mechanism, we first sample in expectation half of the items to form a testing group T , and compute an optimal solution for T with a budget constraint

B. By Lemma 2.1, we know that $v(\text{opt}(A)) \geq v(\text{opt}(T)) \geq \frac{k-1}{4k}v(\text{opt}(A))$ and $v(\text{opt}(A \setminus T)) \geq \frac{k-1}{4k}v(\text{opt}(A))$ with a probability of at least $\frac{1}{2}$. That is, we are able to learn the rough value of the optimal solution by random sampling and still have a constant fraction of the optimal solution within the remaining items. We then set a threshold t for the proper value-per-cost conversion rate. We find a subset $S^* \subseteq A \setminus T$ with the largest difference between its value and cost, multiplied by the threshold t . (In the computation of S^* , if there are multiple choices, ties are broken in a fixed order.) Finally, we find a supporting linear function of the XOS valuation $v(\cdot)$ for the set S^* and run the truthful mechanism for this linear function on the set S^* .

XOS-RANDOM-SAMPLING extensively uses demand queries, and as such it bears resemblance to the $O(\log^2 n)$ approximation mechanism of Dobzinski, Papadimitriou, and Singer [25] for subadditive valuations. However, their mechanism does not use random sampling to estimate the value of the optimal solution, but rather builds on the ideas of a cost-sharing scheme, which we also employ in our mechanism SA-RANDOM-SAMPLING with sublogarithmic approximation ratio.

Our mechanism is designated for XOS functions. It is also used as an auxiliary procedure for more general subadditive functions in the subsequent sections. We note that the mechanism can be implemented in polynomial time if we are given access to XOS and demand query oracles.³

Note that in step 4, the function $v(S) - t \cdot c(S)$ that we maximize is simply the Lagrangian function

$$v(S) - x \cdot c(S) + x \cdot B$$

($x \cdot B$ is a fixed constant) of the original optimization problem $\max_S v(S)$ subject to $c(S) \leq B$. While we do not know the actual value of the variable x in the Lagrangian, a carefully chosen parameter t in the sampling step with a high probability ensures that $\max_S \{v(S) - t \cdot c(S) + t \cdot B\}$ gives a constant approximation of the optimum $\max_S \{v(S) - x \cdot c(S) + x \cdot B\}$ of the Lagrangian, which is precisely the target value $v(\text{opt}(A))$.

The linearity of the Lagrangian, together with the subadditivity of the valuations, is important in deriving the following properties.

CLAIM 3.1 (Claim 3.1 in [25]). *Given the threshold t , subset S^* , and additive function f defined in the XOS-RANDOM-SAMPLING, for any $S \subseteq S^*$, $f(S) - t \cdot c(S) \geq 0$.*

Proof. Suppose for the sake of contradiction that there exists a subset $S \subseteq S^*$ s.t. $f(S) - t \cdot c(S) < 0$. Let $S' = S^* \setminus S$. Since f is an additive function, we have $c(S') + c(S) = c(S^*)$ and $f(S') + f(S) = f(S' \cup S) = f(S^*) = v(S^*)$. Thus,

$$\begin{aligned} v(S') - t \cdot c(S') &\geq f(S') - t \cdot c(S') \\ &= v(S^*) - t \cdot c(S^*) - (f(S) - t \cdot c(S)) \\ &> v(S^*) - t \cdot c(S^*), \end{aligned}$$

which contradicts the definition of S^* . $\square$

The following claim says that any agent in S^* cannot manipulate the selection of

³In fact, in step 2 of the procedure, we can use any approximate solution. Step 4 requires a single query to the demand oracle. In step 5 we need to use XOS oracle [24]. For some XOS valuations like in the matching setting (the value of a subset of edges is equal to the size of the largest matching induced by them), XOS and demand oracles can be implemented in polynomial time.

the set S^* by bidding a smaller cost. This fact is critical for the monotonicity, and thus the truthfulness, of the mechanism.

CLAIM 3.2. *If any item $j \in S^*$ reports a smaller cost $b(j) < c(j)$, then set S^* remains the same.*

Proof. Let b be the bid vector, where j reports $b(j)$ and other bids remain unchanged. First we notice that, for any set S with $j \in S$, $(v(S) - t \cdot b(S)) - (v(S) - t \cdot c(S)) = t(c(j) - b(j))$ is a fixed positive value. Hence,

$$\begin{aligned} v(S^*) - t \cdot b(S^*) &= v(S^*) - t \cdot c(S^*) + t(c(j) - b(j)) \\ &\geq v(S) - t \cdot c(S) + t(c(j) - b(j)) \\ &= v(S) - t \cdot b(S). \end{aligned}$$

Further, for any set S with $j \notin S$, we have

$$\begin{aligned} v(S^*) - t \cdot b(S^*) &> v(S^*) - t \cdot c(S^*) \\ &\geq v(S) - t \cdot c(S) \\ &= v(S) - t \cdot b(S). \end{aligned}$$

Therefore, we conclude that $S^* = \operatorname{argmax}_{S \subseteq A \setminus T} (v(S) - t \cdot b(S))$, and by the fixed tie-breaking rule, S^* is selected as well. $\square$

Our main mechanism for XOS functions is a randomized mixture of the mechanism XOS-RANDOM-SAMPLING and one that always picks an item from $\operatorname{argmax}_i v(i)$.

Algorithm 2 XOS-MECHANISM-MAIN.

With probability $\frac{1}{2}$, run XOS-RANDOM-SAMPLING.

With probability $\frac{1}{2}$, pick the most-valuable item; pay B to the agent.

THEOREM 3.1. *The mechanism XOS-MECHANISM-MAIN is budget feasible and truthful and provides a constant approximation ratio for XOS valuation functions.*

In the remainder of this section, we prove this theorem. We split the proof into the following three lemmas.

LEMMA 3.2. *XOS-MECHANISM-MAIN is universally truthful.*

Our mechanism, at a high level point of view, has a flavor similar to that of the mechanism composition introduced in [1]. In particular, we may consider steps 1–4 of Algorithm 1 as one mechanism of choosing candidate winners, and steps 5–7 as another mechanism restricted to the set of the surviving agents; then the whole mechanism is a composition of the two. It was shown in [1] that if the first mechanism is composable (i.e., truthful plus the property that any winner cannot manipulate the winner set without losing) and the second mechanism is truthful, then the composition mechanism is truthful. In our mechanism XOS-MECHANISM-MAIN, composability of steps 1–4 follows from Claim 3.2, and truthfulness of steps 5–7 is by the corresponding property of ADDITIVE-MECHANISM. Therefore, XOS-MECHANISM-MAIN is truthful.

LEMMA 3.3. *XOS-MECHANISM-MAIN is budget feasible.*

In the mechanism XOS-RANDOM-SAMPLING, the payment to each winner is the maximum amount that the agent can bid and still win. This amount is the minimum of the threshold bids in each of the intermediate steps. In particular, the payment

is upper bounded by the threshold of the mechanism ADDITIVE-MECHANISM in step 6. As ADDITIVE-MECHANISM is budget feasible [17], our mechanism XOS-RANDOM-SAMPLING is budget feasible as well. Finally, it is also budget feasible to pay our entire budget to the agent with the most valuable item.

LEMMA 3.4. XOS-MECHANISM-MAIN has a constant approximation ratio.

Proof. Let $\text{opt} = \text{opt}(A)$ denote the optimal winning set within the budget B , and let $k = \min_{i \in \text{opt}} \frac{v(\text{opt})}{v(i)}$. Thus $v(\text{opt}) \geq k \cdot v(i)$ for each $i \in \text{opt}$. By Lemma 2.1, we have $v(\text{opt} \cap T) \geq \frac{k-1}{4k} v(\text{opt})$ with a probability of at least $\frac{1}{2}$. Thus, we have $v(\text{opt}(T)) \geq v(\text{opt} \cap T) \geq \frac{k-1}{4k} v(\text{opt})$ with a probability of at least $\frac{1}{2}$. (The first inequality is because $\text{opt} \cap T$ is a particular solution and $\text{opt}(T)$ is an optimal solution for set T with the budget constraint.)

We let $\text{opt}^* = \text{opt}_f(S^*)$ be the optimal solution with respect to the item set S^* , additive value-function f , and budget B . In the following we show that $f(\text{opt}^*)$ is a good approximation of the actual optimum $v(\text{opt})$. Consider the following two cases:

- $c(S^*) > B$. Then we can find a subset $S' \subseteq S^*$ such that $\frac{B}{2} \leq c(S') \leq B$. This S' can be obtained from S^* by taking elements of S and adding them to S' (S' is empty initially, items are added in the order of decreasing costs). By Claim 3.1, we know that $f(S') \geq t \cdot c(S') \geq \frac{v(\text{opt}(T))}{8B} \cdot \frac{B}{2} \geq \frac{v(\text{opt}(T))}{16}$. Then by the fact that opt^* is an optimal solution and S' is a particular solution with budget constraint B , we have $f(\text{opt}^*) \geq f(S') \geq \frac{v(\text{opt}(T))}{16} \geq \frac{k-1}{64k} v(\text{opt})$ with a probability of at least $\frac{1}{2}$.
- $c(S^*) \leq B$. Then $\text{opt}^* = S^*$. Let $S' = \text{opt} \setminus T$; thus, $c(S') \leq c(\text{opt}) \leq B$. By Lemma 2.1, we have $v(S') \geq \frac{k-1}{4k} v(\text{opt})$ with a probability of at least $\frac{1}{2}$. Recall that $S^* = \arg\max_{S \subseteq A \setminus T} (v(S) - t \cdot c(S))$. Then with a probability of at least $\frac{1}{2}$, we have

$$\begin{aligned} f(\text{opt}^*) &= f(S^*) = v(S^*) \\ &\geq v(S^*) - t \cdot c(S^*) \\ &\geq v(S') - t \cdot c(S') \\ &\geq \frac{k-1}{4k} v(\text{opt}) - \frac{v(\text{opt}(T))}{8B} \cdot B \\ &\geq \frac{k-1}{4k} v(\text{opt}) - \frac{v(\text{opt})}{8} \\ &= \frac{k-2}{8k} v(\text{opt}). \end{aligned}$$

In either case, we get

$$f(\text{opt}^*) \geq \min \left\{ \frac{k-1}{64k} v(\text{opt}), \frac{k-2}{8k} v(\text{opt}) \right\} \geq \frac{k-2}{64k} v(\text{opt})$$

with a probability of at least $\frac{1}{2}$. At the end we output the result of ADDITIVE-MECHANISM(f, S^*, B) in the last step of XOS-RANDOM-SAMPLING. We recall that ADDITIVE-MECHANISM has an approximation factor of at most 3 with respect to the optimal solution $f(\text{opt}^*)$. Thus the solution given by XOS-RANDOM-SAMPLING is at least $\frac{1}{3} \cdot f(\text{opt}^*) \geq \frac{1}{3} \cdot \frac{1}{2} \cdot \frac{k-2}{64k} v(\text{opt}) = \frac{k-2}{384k} v(\text{opt})$.

On the other hand, since $k = \min_{i \in \text{opt}} \frac{v(\text{opt})}{v(i)}$, the solution given by picking the largest item satisfies $\max_i v(i) \geq \frac{1}{k} v(\text{opt})$. Therefore, our main mechanism XOS-

MECHANISM-MAIN receives a set with the value of at least

$$\left(\frac{1}{2} \cdot \frac{k-2}{384k} + \frac{1}{2} \cdot \frac{1}{k}\right) v(\text{opt}) = \frac{k+382}{768k} v(\text{opt}) \geq \frac{1}{768} v(\text{opt}).$$

This completes the proof of the lemma. $\square$

Subadditive valuations via XOS approximation. For fractionally subadditive valuations Theorem 3.1 tells us that there is a constant approximation mechanism. We now ask the same question of whether there exists a constant approximation truthful mechanism for arbitrary subadditive function $v(\cdot)$. A straightforward approach would be to approximate subadditive valuation $v(\cdot)$ by an XOS function $\tilde{v}(\cdot)$. In other words, we want to find $\tilde{v}(\cdot)$ such that $\frac{v(S)}{\alpha} \leq \tilde{v}(S) \leq v(S)$ for some $\alpha \geq 1$ and any $S \subseteq A$. Then we may run XOS-MECHANISM-MAIN as if the real valuation was $\tilde{v}(\cdot)$ and obtain $O(\alpha)$ approximation to the subadditive valuation $v(\cdot)$. It is well known [27] that any subadditive function admits a $O(\log n)$ XOS approximation. In general, the approximation gap can be as large as $\Theta(\log n)$ [11, 22]. However, for some special cases one can improve on the $O(\log n)$ bound. From the classic Bondareva–Shapley theorem [12, 45] the best possible XOS approximation can be described in terms of the integrality gap of the corresponding fractional cover linear program for $v(\cdot)$. For some problems the corresponding gaps are bounded by a constant (e.g., facility location [41]).

3.1. Sublogarithmic approximation for subadditive valuations. Here, we give another mechanism for subadditive functions. Unlike the brute-force XOS approximation, this mechanism runs in polynomial time and has an $o(\log n)$ approximation ratio. It improves upon the previously best known upper bound $O(\log^2 n)$ of Dobzinski, Papadimitriou, and Singer [25]. Similar to [25], we employ the classic cost sharing scheme to compensate the sellers.

As a subroutine in our mechanism we use a constant factor approximation (non-truthful) *algorithm* for subadditive maximization under a knapsack constraint. Badaniyuru, Dobzinski, and Oren [6] gave a $2 + \epsilon$ approximation algorithm, which we will use as a black box in the description of our *truthful mechanism*.

Algorithm 3 SA-RANDOM-SAMPLING.

1. Pick each item i.i.d. with probability $\frac{1}{2}$ into group T .
2. Let v be the value of β -approximation algorithm $\mathcal{A}$ to

$$\text{optimization problem: } \max_{X \subseteq T} v(X) \quad \text{s.t. } c(X) \leq B.$$

3. **for** $k = 1$ **to** $|A \setminus T|$ **do**
 Let $X_k = \{i \in A \setminus T \mid c(i) \leq \frac{B}{k}\}$.

Let $\hat{X}_k$ be the output of $\mathcal{A}$ for the problem: $\max_{X \subseteq X_k} v(X)$ s.t. $|X| \leq k$.

if $v(\hat{X}_k) \geq \frac{\log \log n}{80 \log n} \cdot v$, **then**

return $\hat{X}_k$ as the winners; pay $\frac{B}{k}$ to everyone in $\hat{X}_k$.

end if

end for

4. **return** $\emptyset$.
-

Similarly to XOS-RANDOM-SAMPLING, the computed value v with high probability is within a constant factor of the optimum when the value of a single most valuable item is negligible compared to the optimal solution. That is, in the latter case we are able to learn the rough value of the optimal solution by random sampling. Next we try to find a big value subset $\hat{X}_k \subset A \setminus T$ in which every item will accept a compensation equal to the budget share of $\frac{B}{k}$. This part of our mechanism can be seen as an instance of the classic cost sharing scheme. Finally, we use v as a benchmark to determine whether $\hat{X}_k$ is good enough to be the final winning set. The final mechanism for subadditive functions is as follows.

Algorithm 4 SA-MECHANISM-MAIN-2.

With probability $\frac{1}{2}$, run SA-RANDOM-SAMPLING.

With probability $\frac{1}{2}$, pick the most-valuable item; pay B to the agent.

THEOREM 3.5. SA-MECHANISM-MAIN-2 runs in polynomial time given a demand oracle and is a truthful budget feasible mechanism for subadditive functions with an approximation ratio of $O(\frac{\log n}{\log \log n})$.

Proof. Let $S = A \setminus T$. We assume that $\mathcal{A}$ is a polynomial time algorithm given access to the demand query oracle. Hence, SA-MECHANISM-MAIN-2 works in polynomial time as well. Since $|\hat{X}_k| = k$, the payment share of $\frac{B}{k}$ is always affordable within the budget. Therefore, SA-RANDOM-SAMPLING and consequently SA-MECHANISM-MAIN-2 are budget feasible. To show universal truthfulness we only need to argue about SA-RANDOM-SAMPLING. As the mechanism is incentive compatible for items in T , it suffices to consider only items in $A \setminus T$. As we further use the standard cost sharing scheme, the truthfulness follows [38].

Approximation ratio. It remains to estimate the approximation ratio. Let $\text{opt} = \text{opt}(A)$ denote the optimal solution for the whole set. If there exists an item $i \in A$ such that $v(i) \geq \frac{1}{2}v(\text{opt})$, then picking the largest item will generate a value of at least $\frac{1}{2}v(\text{opt})$, and we are done. In the following, we assume that $v(i) < \frac{1}{2}v(\text{opt})$ for all $i \in A$. Then by Lemma 2.1,

$$\Pr \left[v(\text{opt}(T)) \geq \frac{v(\text{opt})}{8}, v(\text{opt}(S)) \geq \frac{v(\text{opt})}{8} \right] \geq 0.5.$$

Hence, with probability at least $\frac{1}{4}$ we have

$$(1) \quad v(\text{opt}(S)) \geq v(\text{opt}(T)) \geq \frac{1}{8}v(\text{opt}).$$

Therefore, it suffices to show that the approximation ratio is $O(\frac{\log n}{\log \log n})$ for any fixed S and T for which (1) holds true. Then, as $\mathcal{A}$ is a β approximation to $v(\text{opt}(T))$, we have $v \geq \frac{1}{\beta}v(\text{opt}(T)) \geq \frac{1}{8\beta}v(\text{opt})$. If the algorithm outputs $\hat{X}_k$ as winning set, then its value is at least $v \cdot \Omega(\frac{\log \log n}{\log n})$. Thus it remains to prove that the mechanism always outputs a nonempty set, assuming that (1) holds true. In the following proof we assume the contrary.

Let $\text{opt}(S) = S^* = \{1, 2, 3, \dots, m\}$, where $c_1 \geq c_2 \geq \dots \geq c_m$. We divide agents of S^* into disjoint groups $Z_1, \dots, Z_{r+1}$ by going over the elements of S^* in decreasing order of the costs. We start by forming Z_1 with $|Z_1| = \lfloor \frac{B}{c_1} \rfloor$ first items; to form each new group Z_i ($i \geq 2$) we start from the smallest yet unassigned agent $j(i)$ (we let

$j(1) = 1$) and pick next $|Z_i| = \lfloor \frac{B}{c_{j(i)}} \rfloor$ elements (or maybe a smaller number if we run out of m elements).

If there exists a set Z_i such that $v(Z_i) \geq \frac{\log \log n}{10 \log n} \cdot v$, then the mechanism does not output an empty set. Indeed, the mechanism would buy $|Z_i|$ items at the price $\frac{B}{|Z_i|}$, because $\mathcal{A}$ is a β -approximation algorithm and its output would pass the threshold of $\frac{\log \log n}{10 \beta \log n} \cdot v$. Therefore, we may assume that $v(Z_i) < \frac{\log \log n}{10 \log n} \cdot v$ for $1 \leq i \leq r+1$. On the other hand, by subadditivity, we have

$$\sum_{i=1}^{r+1} v(Z_i) \geq v(S^*) = v(\text{opt}(S)) \geq v(\text{opt}(T)) \geq v.$$

We conclude that $(r+1) \cdot \frac{\log \log n}{10 \log n} \cdot v > v$, which implies

$$r > \frac{10 \log n}{\log \log n} - 1 \geq \frac{5 \log n}{\log \log n} \geq \frac{5 \log m}{\log \log m}.$$

We recall that S^* is a budget feasible solution. Thus $\sum_{j=1}^m c_j \leq B$. Furthermore, $c_{j(1)} > \frac{B}{|Z_1|+1}$, $c_{j(2)} > \frac{B}{|Z_2|+1}$, $\dots$, $c_{j(r)} > \frac{B}{|Z_r|+1}$. We have

$$\begin{aligned} B &\geq \sum_{j=1}^m c_j \geq c_1 + |Z_1|c_{j(2)} + \dots + |Z_r|c_{j(r+1)} \\ &> \frac{B}{|Z_1|+1} + \frac{|Z_1|B}{|Z_2|+1} + \dots + \frac{|Z_{r-1}|B}{|Z_r|+1}. \end{aligned}$$

Hence,

$$\begin{aligned} 1 &\geq \frac{1}{|Z_1|+1} + \frac{|Z_1|}{|Z_2|+1} + \dots + \frac{|Z_{r-1}|}{|Z_r|+1} \\ &\geq \frac{1}{2|Z_1|} + \frac{|Z_1|}{2|Z_2|} + \dots + \frac{|Z_{r-1}|}{2|Z_r|} \geq \frac{1}{2} \cdot r \sqrt[r]{\frac{1}{|Z_1|} \frac{|Z_1|}{|Z_2|} \dots \frac{|Z_{r-1}|}{|Z_r|}}. \end{aligned}$$

Hence, $2 \geq r \sqrt[r]{\frac{1}{|Z_r|}}$ or, equivalently, $|Z_r| \geq (\frac{r}{2})^r$. Finally, we observe that $m \geq |Z_r| \geq (\frac{r}{2})^r$ and $r \geq \frac{5 \log m}{\log \log m}$. Therefore, $\log m \geq r \cdot \log \frac{r}{2} \geq \frac{5 \log m}{\log \log m} \cdot (\log \log m - \log \log \log m + \log \frac{5}{2})$, and we come to a contradiction. This concludes the proof. $\square$

4. Bayesian mechanism design. In this section, we study budget feasible mechanisms for subadditive valuations from a standard economics viewpoint, where the costs of all agents $(c(i))_{i \in A}$ are drawn from a prior known distribution $\mathcal{D}$. More specifically, the mechanism designer and all participants know in advance the $\mathcal{D}$ from which the real cost vector $(c(i))_{i \in A}$ is drawn. However, each $c(i)$ is private information of agent i . We allow dependencies on the agents' costs in $\mathcal{D}$.⁴

Every agent submits a bid $b(i)$ as before and seeks to maximize his own utility. We again consider universally truthful mechanisms; i.e., for every sequence of

⁴We need some mild technical restriction on $\mathcal{D}$ in order to sample certain conditional random variables. We assume that the density function $\rho(\cdot)$ of $\mathcal{D}$ is integrable over each subset $S \subseteq A$ of its variables for any choice of the rest parameters; i.e., $\rho(c_{A \setminus S}) = \int_{\Omega} \rho(c) dx_S$ is bounded. This condition is reminiscent of the integrability of marginal density functions (see, e.g., p. 331 of [47]), though in our case it is slightly stronger. Every finitely supported or product distribution satisfies the required condition.

coin flips of the mechanism and each cost vector, truth-telling should be a dominant strategy for every agent. The performance of a mechanism $\mathcal{M}$ is measured by $\mathbf{E}[\mathcal{M}] = \mathbf{E}_{c \sim \mathcal{D}}[\mathcal{M}(c)]$. We compare a mechanism with the optimal expected value $\mathbf{E}[\text{opt}] = \mathbf{E}_{c \sim \mathcal{D}}[v(\text{opt}(c))]$; we say that mechanism $\mathcal{M}$ is a (Bayesian) α -approximation if $\frac{\mathbf{E}[\text{opt}]}{\mathbf{E}[\mathcal{M}]} \leq \alpha$. Before describing the general mechanism, we first consider a simpler problem that helps to understand where Bayesian assumption might be helpful.

A promise problem. Consider a problem in which all possible cost vectors c are guaranteed to be budget feasible; i.e., $c(A) \leq B$ for any cost vector c in the support of $\mathcal{D}$. In this version of the problem $\text{opt} = A$ with the value $v(A)$. In the auction scenario agents may want to get higher compensation than their true costs, which may disallow the auctioneer to purchase the entire set A . We use the following simple mechanism to handle this issue.

Algorithm 5 RANDOM-THRESHOLDS.

Draw a random vector $d \sim \mathcal{D}$.

Let $X = \{i \in A \mid \text{s.t. } c(i) \leq d(i)\}$ be the winners.

We note that RANDOM-THRESHOLDS is budget feasible, because, for each fixed realization of d , the vector of the mechanism's threshold payments is d , and $d(A) \leq B$. We denote by $X(c, d)$ the winners of the RANDOM-THRESHOLDS mechanism for the cost vector c and the random draw d . We observe that $\Pr[X(c, d) \text{ wins}] = \Pr[X(d, c) \text{ wins}]$ for any two fixed c and d , because vectors c and d are i.i.d. Therefore, we may conclude that

$$\mathbf{E}_{c \sim \mathcal{D}, d \sim \mathcal{D}}[v(X(c, d))] = \frac{1}{2} \mathbf{E}_{c \sim \mathcal{D}, d \sim \mathcal{D}}[v(X(c, d)) + v(X(d, c))] \geq \frac{v(A)}{2},$$

because $X(c, d) \cup X(d, c) = A$ and, by subadditivity, $v(X(c, d)) + v(X(d, c)) \geq v(A)$. Thus RANDOM-THRESHOLDS is a 2 approximation to the opt .

For the remainder of this section let $\text{opt}_v(c, S)$ denote the winning set in an optimal solution when the valuation function is $v(\cdot)$, the cost vector is c , and the agent set is S (the parameters are omitted if they are clear from the context); let $v(\text{opt}_v(c, S))$ denote the value of $\text{opt}_v(c, S)$. Our mechanism for subadditive functions in the Bayesian setting is as follows.

Algorithm 6 SA-BAYESIAN-MECHANISM.

- With probability $\frac{1}{2}$ select the most-valuable item; pay B .
 - With probability $\frac{1}{2}$, run the following:
 1. Pick each item i.i.d. with probability $\frac{1}{2}$ into group T .
 2. Compute an optimal solution $\text{opt}(c, T)$ for items in T given budget B .
 3. Set a threshold $t = \frac{v(\text{opt}(c, T))}{8B}$.
 4. For items in $A \setminus T$ find a set $S^* \in \arg\max_{S \subseteq A \setminus T} \{v(S) - t \cdot c(S)\}$.
 5. Sample cost vector $d \sim \mathcal{D}$ conditioned on
 - (a) $d(i) = c(i)$ for each $i \in T$, and
 - (b) $S^* \in \arg\max_{S \subseteq A \setminus T} \{v(S) - t \cdot d(S)\}$.
 6. Let $Y = \{i \in S^* \mid \text{s.t. } c(i) \leq d(i)\}$.
 7. If $d(Y) \leq B$, let all $i \in Y$ be the winners.
 8. Else, in a fixed order select the winners $X \subset Y$ s.t. $B \geq d(X) \geq \frac{B}{2}$.
-

In the mechanism, steps 1–3 are the same as in XOS-RANDOM-SAMPLING, where we randomly sample a test group T and generate a threshold value t . In steps 4–8,

we consider a specific subset $S^* \subseteq A \setminus T$ and select winners only from it. Step 5 helps guide us on the threshold payments of the winners (see more discussions below).

A few remarks about the mechanism are in order:

- It is tempting to remove the random sampling part as, given $\mathcal{D}$, one may consider a “prior sampling” approach: Generate some virtual instances according to $\mathcal{D}$ and compute a threshold t based on them; then apply this threshold to all agents in A . Interestingly, the prior sampling approach works well in our mechanism when, e.g., all $c(i)$ ’s are independent, but it does not work for the case when variables are dependent.

For instance, consider an additive valuation $v(\cdot)$ with $v(S) = |S|$, budget $B = 2^k$ for a large k , and a set of $N = 2^k$ agents with the following discrete distribution over costs ($c = \ell$ means that every $c(i) = \ell$):

$$\Pr[c = 1] = \frac{1}{2^{k+1}}, \Pr[c = 2] = \frac{1}{2^k}, \dots, \Pr[c = 2^k] = \frac{1}{2},$$

$$\Pr[c = 2^{k+1}] = \frac{1}{2^{k+1}}.$$

Note that

$$v(\text{opt}(c = 1)) = 2^k, v(\text{opt}(c = 2)) = 2^{k-1}, \dots, v(\text{opt}(c = 2^k)) = 1,$$

$$v(\text{opt}(c = 2^{k+1})) = 0.$$

Then the expected optimal value is $\mathbf{E}[\text{opt}] = \frac{k+1}{2}$, and it is equally spread over all possible costs except the last one, $c = 2^{k+1}$. Roughly speaking, on a given instance c , any prior estimate on $v(\text{opt}(c))$ that gives a constant approximation applies to only a constant number of distinct costs (the contribution of these cases to $\mathbf{E}[\text{opt}]$ is negligible). Hence for almost all other possible costs, we get a meaningless estimate for $\text{opt}(c)$. Therefore, the prior sampling will lead to a bad approximation ratio.

- Why do we generate another cost vector d in step 5? Recall that our target winner set is S^* , whose value $v(S^*)$ in expectation gives a constant approximation of $\mathbf{E}[\text{opt}]$. However, we are faced with the problems of selecting a winning set in S^* with a sufficiently large value and distributing the budget among the winners. These two problems together are closely related to cooperative game theory and the notion of an approximate core. For subadditive functions, a constant approximate core may not exist [41] (e.g., set cover gives a logarithmic lower bound [11]). Thus we might not be able to pick a winning set with a constant approximation and set threshold payments in accordance with the valuation function. The question then is: Is there any other guidance we can take to bound budget feasible threshold payments and give a constant approximation?

Our solution is to use another random vector d to serve as such a guidance. Conditions in steps 5a and 5b guarantee that cost vectors c and d are distributed identically and can be switched in expectation while preserving some important parameters such as t and S^* . We set $d(i)$ as an upper bound on the payment of each agent $i \in S^*$, which guarantees that we are always within the budget constraint.

THEOREM 4.1. *SA-BAYESIAN-MECHANISM is a universally truthful budget feasible mechanism for subadditive functions and gives in expectation a constant approximation.*

Proof. Budget feasibility follows simply from the description of the mechanism and the fact that threshold payments are upper bounded by the random vector d .

For universal truthfulness, we note that in the mechanism, the sampled vector d comes from a distribution that depends on actual bid vector c . To see why our mechanism takes a distribution over deterministic truthful mechanisms, we can draw up front all possible samples d for (i) all possible cost vectors on T and (ii) all possible choices $S \subseteq A \setminus T$ of S^* . Note that the selection rule of S^* is monotone, and, similarly to Claim 3.2, each agent in S^* cannot manipulate (i) our choice of S^* and (ii) the choice of d , as long as he stays in S^* . Therefore, the composition of the first selection rule, where we choose S^* (step 4), with the next monotone rule, where we pick winners in S^* (steps 7–8), is again a monotone rule. Hence, the mechanism is universally truthful.

Approximation analysis. We assume that no single item can have cost more than B in the cost vector c . Thus the first part of SA-BAYESIAN-MECHANISM guarantees that in expectation we obtain one half of the largest item value. We denote by $X(c, d, T)$ the set of winners in the second part of SA-BAYESIAN-MECHANISM (steps 1–8), when the cost vector is c , the sampled test set is T , and the random cost vector is d . Similarly to the promise version of the problem, for every random sample T and any fixed cost vectors c, d we have

$$\Pr[X(c, d, T) \text{ wins}] = \Pr[X(d, c, T) \text{ wins}].$$

Therefore, we may substitute $v(X(c, d, T))$ with $\frac{1}{2}(v(X(c, d, T)) + v(X(d, c, T)))$ in our calculations of the expected performance of SA-BAYESIAN-MECHANISM. To simplify notation we denote the expression $\frac{1}{2}(v(X(c, d, T)) + v(X(d, c, T)))$ by $f(c, d, T)$. Before writing estimates on $f(c, d, T)$, we need the following preliminary claim.

In the mechanism we compute a threshold $t(c, T) = \frac{v(\text{opt}(c, T))}{8B}$ and a set $S^*(c, T)$, both of which depend on the sampled set T and cost vector c . We also know that the vector d is such that $t(d, T) = t(c, T)$ and $S^*(d, T) = S^*(c, T)$.

CLAIM 4.1. For a given subset $S \subset S^*(c, T)$

1. $v(S) \geq t(c, T) \cdot c(S)$,
2. $v(S) \geq t(c, T) \cdot d(S)$.

Proof. The argument is similar to that of Claim 3.1 and is identical for c and d cost vectors. Thus we prove only the first part of Claim 4.1. We suppose for the sake of contradiction that there exists a subset $S \subseteq S^*$ such that $v(S) - t \cdot c(S) < 0$. Let $S' = S^* \setminus S$. We have $c(S') + c(S) = c(S^*)$ and also $v(S') + v(S) \geq v(S^*)$, since v is a subadditive function. Therefore,

$$\begin{aligned} v(S') - t \cdot c(S') &\geq v(S^*) - v(S) - t \cdot c(S') \\ &= v(S^*) - t \cdot c(S^*) - (v(S) - t \cdot c(S)) \\ &> v(S^*) - t \cdot c(S^*), \end{aligned}$$

which contradicts the definition of S^* in step 5b. $\square$

Now we are ready to give some estimates on $f(c, d, T)$, as follows.

LEMMA 4.2. Let T be a fixed set and c, d be fixed cost vectors from SA-BAYESIAN-MECHANISM. Then

$$f(c, d, T) \geq \min \left(\frac{1}{32} v(\text{opt}(c, T)), \frac{1}{2} v(S^*) \right).$$

Proof. We first observe that

$$Y(c, d, T) \cup Y(d, c, T) = \{i \in S^* \mid \text{s.t. } c(i) \leq d(i)\} \cup \{i \in S^* \mid \text{s.t. } d(i) \leq c(i)\} = S^*.$$

Therefore, if both $d(Y(c, d, T)) \leq B$ and $c(Y(d, c, T)) \leq B$, then according to step 7, $X(c, d, T) = Y(c, d, T)$ and $X(d, c, T) = Y(d, c, T)$. We get the desired bound

$$f(c, d, T) \geq \frac{1}{2}(v(Y(c, d, T)) + v(Y(d, c, T))) \geq \frac{1}{2}v(S^*).$$

Now if $d(Y(c, d, T)) > B$ or $c(Y(d, c, T)) > B$, then according to step 8, either $d(X(c, d, T)) \geq \frac{B}{2}$ or $c(X(d, c, T)) \geq \frac{B}{2}$. In either case by Claim 4.1,

$$\begin{aligned} f(c, d, T) &= \frac{1}{2}(v(X(c, d, T)) + v(X(d, c, T))) \\ &\geq \frac{t}{2} \cdot d(X(c, d, T)) + \frac{t}{2} \cdot c(X(d, c, T)) \\ &\geq \frac{t}{2} \cdot \frac{B}{2} = \frac{v(\text{opt}(c, T))}{32}. \end{aligned}$$

We conclude the proof by combining the above two lower bounds on $f(c, d, T)$. $\square$

Proof of Theorem 4.1 (continued). The rest of the proof of Theorem 4.1 proceeds analogously to the proof of Lemma 3.4 for XOS valuations. From Lemma 2.1 we conclude that

$$\Pr_{T \sim 2^A} \left[v(\text{opt}(T)) \geq \frac{k-1}{4k} v(\text{opt}), v(\text{opt} \setminus T) \geq \frac{k-1}{4k} v(\text{opt}) \right] \geq \frac{1}{2}.$$

We also observe that if $v(\text{opt}(T)) \geq \frac{k-1}{4k} v(\text{opt})$ and $v(\text{opt} \setminus T) \geq \frac{k-1}{4k} v(\text{opt})$, then

$$\begin{aligned} v(S^*) &\geq v(S^*) - t \cdot c(S^*) \geq v(\text{opt} \setminus T) - t \cdot c(\text{opt} \setminus T) \\ &\geq \frac{k-1}{4k} v(\text{opt}) - t \cdot B \\ &\geq \frac{k-1}{4k} v(\text{opt}) - \frac{v(\text{opt}(T))}{8B} \cdot B \\ &\geq \frac{k-1}{4k} v(\text{opt}) - \frac{v(\text{opt})}{8} = \frac{k-2}{8k} v(\text{opt}). \end{aligned}$$

Therefore,

$$\Pr_{T \sim 2^A} \left[v(\text{opt}(T)) \geq \frac{k-1}{4k} v(\text{opt}), v(S^*) \geq \frac{k-2}{8k} v(\text{opt}) \right] \geq \frac{1}{2}.$$

By Lemma 4.2 we further conclude that

$$\Pr_{T \sim 2^A} \left[f(c, d, T) \geq \min \left(\frac{1}{32} \frac{k-1}{4k} v(\text{opt}), \frac{k-2}{16k} v(\text{opt}) \right) \right] \geq \frac{1}{2}.$$

Now we can estimate the expected value of the second part of SA-BAYESIAN-MECHANISM as

$$\mathbf{E}_{T, c, d} [f(c, d, T)] \geq \frac{1}{2} \cdot \min \left(\frac{k-1}{128k}, \frac{k-2}{16k} \right) v(\text{opt}) \geq \frac{k-2}{256k} v(\text{opt}).$$

Combining the estimates for the two parts of SA-BAYESIAN-MECHANISM together, we can lower-bound the performance of SA-BAYESIAN-MECHANISM as

$$\frac{1}{2} \cdot \frac{k-2}{256k} v(\text{opt}) + \frac{1}{2} \cdot \frac{1}{k} v(\text{opt}) = \frac{k+254}{512k} v(\text{opt}) \geq \frac{1}{512} v(\text{opt}).$$

This concludes the proof of Theorem 4.1. $\square$

5. Subadditive valuations: From Bayesian to prior-free. In this section, we return to the prior-free setting. Building on the constant approximation mechanism for the Bayesian setting in the previous section, we are finally able to show the existence of a constant approximation budget feasible mechanism for subadditive functions.

Our result is based on the following rather general observation: suppose that, for any given distribution $\mathcal{D}$ within the Bayesian framework, we can achieve a certain approximation guarantee on the performance of a truthful mechanism to the optimum solution (or any other benchmark). Further, suppose that we do not need to rely on Bayesian analysis in the following aspects:

Truthfulness. We are looking for *universal truthfulness*, i.e., require our mechanism to be incentive compatible for any realization of coin flips of the mechanism.

Distribution. We allow for arbitrary finitely supported distributions with interdependent bids.

With the above conditions, we can return to the prior-free worst-case framework and, based on the results from the Bayesian framework, show the existence of an incentive compatible mechanism with the same approximation to the optimum for any bid vector.

We illustrate below how it works in the context of budget feasible mechanism design. Recall that the result for the Bayesian framework states that, for any given distribution $\mathcal{D}$ of possibly interdependent sellers' bids b , there is a universally truthful budget feasible mechanism $\mathcal{M}_{\mathcal{D}}$ that derives in expectation over the distribution $b \sim \mathcal{D}$ at least a constant fraction δ of the expected optimal value $\text{opt}(b)$. We may render this result into a similar result in a worst-case prior-free framework.

THEOREM 5.1. *For any given subadditive valuation $v(\cdot)$ and budget B there exists a budget feasible incentive compatible mechanism with constant approximation δ to the optimum.*

Proof. Let $\mathcal{A}$ be the space of all universally truthful budget feasible mechanisms. Note that $\mathcal{A}$ forms a convex set, since for any two mechanisms $A_1, A_2 \in \mathcal{A}$ one may define another universally truthful budget feasible mechanism $A = \lambda \cdot A_1 + (1 - \lambda) \cdot A_2$, which with probability λ runs A_1 and with probability $1 - \lambda$ runs A_2 .

Without loss of generality, we may consider only those situations where every agent could submit finitely many different numbers as a bid (say, all integer multiples of $\frac{B}{2^n}$ not exceeding $B + 1$). Note that there are only finitely many possible allocation rules over the finite space of feasible bids. Therefore, we may also assume that there are only finitely many deterministic truthful mechanisms in $\mathcal{A}$.

We now recall our constant ($\lambda_0 > 0$) approximation results from Theorem 4.1 in the Bayesian framework: for any given distribution $\mathcal{D}$ of cost vectors there is a mechanism $A \in \mathcal{A}$ such that

$$\mathbf{E}_{c \sim \mathcal{D}}[v(A(c))] \geq \delta \cdot \mathbf{E}_{c \sim \mathcal{D}}[v(\text{opt}(c))].$$

Let us consider a two-player game with the first player deciding on a feasible cost vector c and the second player choosing a deterministic truthful budget feasible mechanism $A \in \mathcal{A}$. Note that each player has a finite number of pure strategies in this game. We define each entry of the payoff matrix (i.e., the amount that player 1 pays to player 2) for the pair (c, A) as $v(A(c)) - \delta \cdot v(\text{opt}(c))$ in the normal form of our game. Next we apply Yao's minimax principle:

$$\min_{\mathcal{D}} \max_{A \in \mathcal{A}} \mathbf{E}_{c \sim \mathcal{D}}[v(A(c)) - \delta \cdot v(\text{opt}(c))] = \max_{\mathcal{D}} \min_c \mathbf{E}_{A \sim \mathcal{D}}[v(A(c)) - \delta \cdot v(\text{opt}(c))].$$

From our results for the Bayesian setting the left-hand side is nonnegative. Hence there is a distribution $\mathfrak{D}$ of deterministic truthful budget feasible mechanisms in the right-hand side such that, for any cost vector c ,

$$\mathbf{E}_{A \sim \mathfrak{D}}[v(A(c))] \geq \delta \cdot v(\text{opt}(c)).$$

This concludes the proof, as we can take a randomized universally truthful budget feasible mechanism that simply runs $A \sim \mathfrak{D}$ and achieves a δ factor of approximation. $\square$

6. Conclusions. Our work on budget feasible mechanism design bridges prior-free and Bayesian analysis frameworks. First, using the random sampling technique, we give a prior-free constant approximation mechanism for XOS valuations, which also implies a prior-free $\log n$ mechanism for subadditive functions. Next we turn to the Bayesian framework. We present a Bayesian constant approximation mechanism for subadditive valuation functions. Finally, we convert the Bayesian mechanism to a prior-free mechanism while preserving the same approximation ratio, which is eventually a constant approximation mechanism for subadditive functions in the prior-free framework. All our mechanisms are universally truthful.

Our mechanisms continue to work for the extension when the valuation functions are nonmonotone; i.e., $v(S)$ is not necessarily less than $v(T)$ for any $S \subset T \subseteq A$. For instance, the cut function studied in [23] is nonmonotone. For such functions, we can define $\hat{v}(S) = \max_{T \subseteq S} v(T)$ for any $S \subseteq A$. It is easy to see that $\hat{v}(\cdot)$ can be computed easily given a demand oracle, is monotone, and inherits the classification of $v(\cdot)$. Further, any solution maximizing $v(\cdot)$ is also an optimal solution of $\hat{v}(\cdot)$. Hence, we can apply our mechanisms to $\hat{v}(\cdot)$ directly and obtain the same approximations.

Random sampling appears to be a powerful approach and has been used successfully in other domains of mechanism design, e.g., digital goods auctions [30], the secretary problem [4, 5], social welfare maximization [22], and mechanism design without money [16]. It is intriguing to find applications of random sampling in other mechanism design problems.

For XOS valuations our mechanism can be implemented in polynomial time given access to demand and XOS query oracles. We also give a complementary computationally efficient $o(\log n)$ -approximation mechanism for subadditive functions using a cost-sharing payment scheme. However, our constant approximation mechanism for subadditive valuations is not computationally efficient. Thus it is natural to ask whether there are truthful designs with the same approximations that can be implemented in polynomial time. Further, all of our mechanisms are randomized; it is intriguing to consider the approximability of deterministic mechanisms. We leave these questions for future work.

Acknowledgments. We thank the anonymous reviewers at STOC 2011 for their helpful comments. We are also grateful to Jason Hartline for many valuable suggestions (in particular, for the intuitive connection to Lagrangian multipliers).

REFERENCES

- [1] G. AGGARWAL AND J. D. HARTLINE, *Knapsack auctions*, in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2006, pp. 1083–1092.
- [2] N. ANARI, G. GOEL, AND A. NIKZAD, *Mechanism design for crowdsourcing: An optimal $1-1/e$ competitive budget-feasible mechanism for large markets*, in Proceedings of the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE Press, Piscataway, NJ, 2014, pp. 266–275.

- [3] A. ARCHER AND É. TARDOS, *Truthful mechanisms for one-parameter agents*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science, IEEE Press, Piscataway, NJ, 2001, pp. 482–491.
- [4] M. BABAIOFF, M. DINITZ, A. GUPTA, N. IMMORLICA, AND K. TALWAR, *Secretary problems: Weights and discounts*, in Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2009, pp. 1245–1254, <https://doi.org/10.1137/1.9781611973068.135>.
- [5] M. BABAIOFF, N. IMMORLICA, AND R. KLEINBERG, *Matroids, secretary problems, and online mechanisms*, in Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2007, pp. 434–443.
- [6] A. BADANIDIYURU, S. DOBZINSKI, AND S. OREN, *Optimization with demand oracles*, in Proceedings of the 13th Annual ACM Conference on Electronic Commerce, ACM, New York, 2012, pp. 110–127.
- [7] E. BALKANSKI AND J. D. HARTLINE, *Bayesian budget feasibility with posted pricing*, in Proceedings of the 25th International Conference on World Wide Web, 2016, pp. 189–203.
- [8] X. BEI, N. CHEN, N. GRAVIN, AND P. LU, *Budget feasible mechanism design: From prior-free to Bayesian*, in Proceedings of the 44th Annual ACM Symposium on Theory of Computing, ACM, New York, 2012, pp. 449–458.
- [9] X. BEI AND Z. HUANG, *Bayesian incentive compatibility via fractional assignments*, in Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2011, pp. 720–733, <https://doi.org/10.1137/1.9781611973082.57>.
- [10] S. BHATTACHARYA, G. GOEL, S. GOLLAPUDI, AND K. MUNAGALA, *Budget constrained auctions with heterogeneous items*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, ACM, New York, 2010, pp. 379–388.
- [11] K. BHAWALKAR AND T. ROUGHGARDEN, *Welfare guarantees for combinatorial auctions with item bidding*, in Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2011, pp. 700–709, <https://doi.org/10.1137/1.9781611973082.55>.
- [12] O. N. BONDAREVA, *Some applications of linear programming methods to the theory of cooperative games*, Probl. Kibernet., 10 (1963), pp. 119–139.
- [13] S. CHAWLA, J. D. HARTLINE, D. L. MALEC, AND B. SIVAN, *Multi-parameter mechanism design and sequential posted pricing*, in Proceedings of the 44th ACM Symposium on Theory of Computing, ACM, New York, 2010, pp. 311–320.
- [14] S. CHAWLA, D. L. MALEC, AND A. MALEKIAN, *Bayesian mechanism design for budget-constrained agents*, in Proceedings of the 12th ACM Conference on Electronic Commerce, ACM, New York, 2011, pp. 253–262.
- [15] S. CHAWLA, D. L. MALEC, AND B. SIVAN, *The power of randomness in Bayesian optimal mechanism design*, in Proceedings of the 11th ACM Conference on Electronic Commerce, ACM, New York, 2010, pp. 149–158.
- [16] N. CHEN, N. GRAVIN, AND P. LU, *Mechanism design without money via stable matching*, preprint <https://arxiv.org/abs/1104.2872>, 2011.
- [17] N. CHEN, N. GRAVIN, AND P. LU, *On the approximability of budget feasible mechanisms*, in Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, 2011, pp. 685–699, <https://doi.org/10.1137/1.9781611973082.54>.
- [18] X. CHEN, G. HU, P. LU, AND L. WANG, *On the approximation ratio of k -lookahead auction*, in Internet and Network Economics, Springer, New York, 2011, pp. 61–71.
- [19] E. H. CLARKE, *Multipart pricing of public goods*, Public Choice, 11 (1971), pp. 17–33.
- [20] N. DEVANUR, J. HARTLINE, A. KARLIN, AND T. NGUYEN, *Prior-independent multi-parameter mechanism design*, in Internet and Network Economics, Springer, New York, 2011, pp. 122–133.
- [21] P. DHANGWATNOTAI, T. ROUGHGARDEN, AND Q. YAN, *Revenue maximization with a single sample*, Games Econom. Behav., 91 (2015), pp. 318–333.
- [22] S. DOBZINSKI, *Two randomized mechanisms for combinatorial auctions*, in Approximation, Randomization, and Combinatorial Optimization, Algorithms and Techniques, Springer, New York, 2007, pp. 89–103.
- [23] S. DOBZINSKI, H. FU, AND R. D. KLEINBERG, *Optimal auctions with correlated bidders are easy*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, ACM, New York, 2011, pp. 129–138.
- [24] S. DOBZINSKI, N. NISAN, AND M. SCHAPIRA, *Approximation algorithms for combinatorial auctions with complement-free bidders*, Math. Oper. Res., 35 (2010), pp. 1–13.
- [25] S. DOBZINSKI, C. H. PAPADIMITRIOU, AND Y. SINGER, *Mechanisms for complement-free procurement*, in Proceedings of the 12th ACM Conference on Electronic Commerce, ACM, New York, 2011, pp. 273–282.

- [26] S. DUGHMI, T. ROUGHGARDEN, AND Q. YAN, *From convex optimization to randomized mechanisms: Toward optimal combinatorial auctions*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, ACM, New York, 2011, pp. 149–158.
- [27] U. FEIGE, *On maximizing welfare when utility functions are subadditive*, SIAM J. Comput., 39 (2009), pp. 122–142, <https://doi.org/10.1137/070680977>.
- [28] M. FELDMAN, H. FU, N. GRAVIN, AND B. LUCIER, *Simultaneous auctions are (almost) efficient*, in Proceedings of the 45th Annual ACM Symposium on Theory of Computing, ACM, New York, 2013, pp. 201–210.
- [29] A. GHOSH AND A. ROTH, *Selling privacy at auction*, Games Econom. Behav., 91 (2015), pp. 334–346.
- [30] A. V. GOLDBERG, J. D. HARTLINE, A. R. KARLIN, M. SAKS, AND A. WRIGHT, *Competitive auctions*, Games Econom. Behav., 55 (2006), pp. 242–269.
- [31] T. GROVES, *Incentives in teams*, Econometrica, 41 (1973), pp. 617–631.
- [32] J. D. HARTLINE, R. KLEINBERG, AND A. MALEKIAN, *Bayesian incentive compatibility via matchings*, Games Econom. Behav., 92 (2015), pp. 401–429.
- [33] J. D. HARTLINE AND B. LUCIER, *Bayesian algorithmic mechanism design*, in Proceedings of the 42nd ACM Symposium on Theory of Computing, ACM, New York, 2010, pp. 301–310.
- [34] J. D. HARTLINE AND T. ROUGHGARDEN, *Optimal mechanism design and money burning*, in Proceedings of the 40th Annual ACM Symposium on Theory of Computing, ACM, New York, 2008, pp. 75–84.
- [35] J. D. HARTLINE AND T. ROUGHGARDEN, *Simple versus optimal mechanisms*, in Proceedings of the 10th ACM Conference on Electronic Commerce, ACM, New York, 2009, pp. 225–234.
- [36] B. LEHMANN, D. LEHMANN, AND N. NISAN, *Combinatorial auctions with decreasing marginal utilities*, in Proceedings of the 3rd ACM Conference on Electronic Commerce, ACM, New York, 2001, pp. 18–28.
- [37] P. R. MILGROM, *Putting Auction Theory to Work*, Cambridge University Press, Cambridge, UK, 2004.
- [38] H. MOULIN, *Incremental cost sharing: Characterization by coalition strategy-proofness*, Social Choice and Welfare, 16 (1999), pp. 279–320.
- [39] R. B. MYERSON, *Optimal auction design*, Math. Oper. Res., 6 (1981), pp. 58–73.
- [40] N. NISAN AND A. RONEN, *Algorithmic mechanism design*, in Proceedings of the 31st Annual ACM Symposium on Theory of Computing, ACM, New York, 1999, pp. 129–140.
- [41] N. NISAN, T. ROUGHGARDEN, E. TARDOS, AND V. V. VAZIRANI, *Algorithmic Game Theory*, Cambridge University Press, Cambridge, UK, 2007.
- [42] C. H. PAPADIMITRIOU AND G. PIERRAKOS, *On optimal single-item auctions*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing, ACM, New York, 2011, pp. 119–128.
- [43] A. RONEN, *On approximating optimal auctions*, in Proceedings of the 3rd ACM Conference on Electronic Commerce, ACM, New York, 2001, pp. 11–17.
- [44] A. RONEN AND A. SABERI, *Optimal auctions are hard*, in Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science, IEEE Press, Piscataway, NJ, 2002, pp. 396–405.
- [45] L. S. SHAPLEY, *On balanced sets and cores*, Naval Res. Logistics Quart., 14 (1967), pp. 453–460.
- [46] Y. SINGER, *Budget feasible mechanisms*, in Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, IEEE Press, Piscataway, NJ, 2010, pp. 765–774.
- [47] C. J. STONE, *A Course in Probability and Statistics*, Duxbury Press, Pacific Grove, CA, 1996.
- [48] M. SVIRIDENKO, *A note on maximizing a submodular set function subject to a knapsack constraint*, Oper. Res. Lett., 32 (2004), pp. 41–43.
- [49] W. VICKREY, *Counterspeculation, auctions, and competitive sealed tenders*, J. Finance, 16 (1961), pp. 8–37.