

Budget Feasible Mechanism Design: From Prior-Free to Bayesian

Xiaohui Bei^{*}

Ning Chen[†]

Nick Gravin[†]

Pinyan Lu[‡]

ABSTRACT

Budget feasible mechanism design studies procurement combinatorial auctions in which the sellers have private costs to produce items, and the buyer (auctioneer) aims to maximize a social valuation function on subsets of 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 (compared to the social optimum)?” Singer [35] showed that additive and submodular functions have a constant approximation mechanism. Recently, Dobzinski, Papadimitriou, and Singer [20] gave an $O(\log^2 n)$ approximation mechanism for subadditive functions; further, they 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 address this question from two viewpoints: prior-free worst case analysis and Bayesian analysis, which are two standard approaches from computer science and economics, respectively.

- For the prior-free framework, we use a linear program (LP) that describes the fractional cover of the valuation function; the LP is also connected to the concept of approximate core in cooperative game theory. We provide a mechanism for subadditive functions whose approximation is $O(\mathcal{I})$, via the worst case integrality gap $\mathcal{I}$ of this LP. This implies an $O(\log n)$ -approximation for subadditive valuations, $O(1)$ -approximation for XOS valuations, as well as for valuations having a constant

integrality gap. XOS valuations are an important class of functions and lie between the submodular and the subadditive classes of valuations. We further give another polynomial time $O(\frac{\log n}{\log \log n})$ sub-logarithmic approximation mechanism for subadditive functions. Both of our mechanisms improve the best known approximation ratio $O(\log^2 n)$.

- For the Bayesian framework, we provide a constant approximation mechanism for all subadditive functions, using the above prior-free mechanism for XOS valuations as a subroutine. Our mechanism allows correlations in the distribution of private information and is universally truthful.

Categories and Subject Descriptors

F.0 [Theory of Computation]: General; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Theory, Economics

Keywords

Mechanism design, budget feasible, prior-free, Bayesian, submodular, subadditive, approximation

1. INTRODUCTION

Consider a procurement combinatorial auction problem where there is a buyer who wants to purchase resources from a set of agents A . Each $i \in A$ is able to supply a resource at an incurred cost $c(i)$. The buyer has a sharp budget B that gives an upper bound on the compensation that is distributed among agents, and a function $v(\cdot)$ describing the valuation that the buyer obtains 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 [37].

Agents, as self-interested entities, may want to get as many subsidies as possible. In particular, an agent can hide his true incurred cost $c(i)$ (which is known only to himself) and claim ‘any’ amount $b(i)$ instead. We adopt the approach of mechanism design to manage self-interested, but strategic, behaviors of the agents: Given submitted bids $b(i)$ from all agents, a *mechanism* decides a winning set $S \subseteq A$ and a payment $p(i)$ to each winner $i \in S$. A mechanism is called

^{*}Institute for Interdisciplinary Information Sciences, Tsinghua University, China. Email: bxh08@mails.tsinghua.edu.cn.

[†]Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore.

Email: ningc@ntu.edu.sg, ngravin@pmail.ntu.edu.sg.

[‡]Microsoft Research Asia. Email: pinyanl@microsoft.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC’12, May 19–22, 2012, New York, New York, USA.

Copyright 2012 ACM 978-1-4503-1245-5/12/05 ...\$10.00.

truthful (a.k.a. incentive compatible) if for every agent it is a dominant strategy to bid his true cost¹, i.e., $b(i) = c(i)$. Truthfulness is one of the central solution concepts in mechanism design. It ensures that every participant will behave truthfully to his best interest.

Our mechanism design problem has an important and practical ingredient: the budget, i.e., the total payment of a mechanism should be upper bounded by B . The budget constraint introduces a new dimension to mechanism design and restricts the space of truthful mechanisms. For example, in single-parameter domains where the private information of every individual is a single value (which is the case in our model), a monotone allocation rule with associated threshold payments provides a sufficient and necessary condition for truthfulness [31]. However, it may not necessarily generate a budget feasible solution. Thus, a number of well known truthful designs (e.g., the seminal VCG mechanism [38, 16, 24]) do not apply, and new ideas have to be developed.

Another significant challenge due to the budget constraint is that, unlike 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 even if we are given unlimited computational power. Indeed, in a simple setting like path procurement with 0 or 1 valuation [35], any budget feasible mechanism may have an arbitrarily bad solution. Therefore, the question that one may ask is “under which valuation domains do there exist truthful budget feasible mechanisms that admit ‘small’ approximations (compared to the socially optimal solution)?”

The answer to this question crucially depends on the properties and classifications of the valuation function under consideration. In particular, given the following function hierarchy [29]:

$$\begin{aligned} \text{additive} &\subset \text{gross substitutes} \subset \text{submodular} \\ &\subset \text{XOS} \subset \text{subadditive}, \end{aligned}$$

which one admits a positive answer? Singer [35] initiated the study of approximate budget feasible mechanism design and gave constant approximation mechanisms for additive and submodular functions. In subsequent work, Dobzinski, Papadimitriou, and Singer [20] considered subadditive functions and showed an $O(\log^2 n)$ approximation. Further, it was remarked in [20] that:

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

— Dobzinski, Papadimitriou, Singer

In the present paper we attempt to answer this question.

1.1 Our Results and Techniques

We address this question from two viewpoints: prior-free worst case analysis and Bayesian analysis. The former is the standard framework used in computer science: in our model, the private cost $c(i)$ of every agent is assumed to be arbitrary without any prior knowledge. All previous research on budget feasible mechanism design, e.g. [35, 14, 20, 22]

¹The focus of our work is to consider strategic behaviors of the agents rather than the buyer. We thus assume that the information related to the buyer, i.e., budget B and valuation function $v(\cdot)$, is public knowledge.

falls into this framework. The latter Bayesian analysis [31] is a classic economic approach that assumes the private information of the agents is drawn from a given prior-known distribution. Bayesian mechanism design has received a lot of attention in the computer science community in recent years, see, e.g., [27, 28, 26, 7, 11, 18, 13, 25, 5, 12, 17].

Prior-free mechanism design.

Consider the following linear program (LP), where $\alpha(\cdot)$'s are variables.

$$\begin{aligned} \min \quad & \sum_{S \subseteq A} \alpha(S) \cdot v(S) \\ \text{s.t.} \quad & \alpha(S) \geq 0, \quad \forall S \subseteq A \\ & \sum_{S: i \in S} \alpha(S) \geq 1, \quad \forall i \in A \end{aligned}$$

Constraints of this LP describe a fractional set cover of A , where each set S receives weight $\alpha(S)$ and we require that all agents in A are covered. An important observation about this LP is that for any monotone subadditive function $v(\cdot)$, the value of the optimal integral solution is precisely $v(A)$.

The above LP has a strong connection to the core of cost sharing games (considering $v(\cdot)$ instead as a cost function), which is a central notion in cooperative game theory [33]. Roughly speaking, the core of a game is a stable cooperation among all agents to share $v(A)$ where no subset of agents can benefit by breaking away from the grand coalition. It is well known that the cores of many cost sharing games are empty. This motivates the notion of α -approximate core, which requires all the agents to share only an α -fraction of $v(A)$. The classic Bondareva-Shapley Theorem [9, 34] says that for subadditive functions, the largest value α for which the α -approximate core is nonempty is equal to the integrality gap of the LP. Further, the integrality gap of the LP equals one (i.e., $v(A)$ is also an optimal fractional solution) if and only if the valuation function is XOS, which is also equivalent to the non-emptiness of the core.

Given an instance of our problem with an agent set A , we may consider the above LP and its integrality gap for every subinstance $A' \subseteq A$. We denote $\mathcal{I}$ as the largest integrality gap among all subinstances $A' \subseteq A$. In other words, the gap $\mathcal{I}$ characterizes the worst case scenario between the optimal integral and fractional solutions of the problem. We have the following result.

Theorem 1. *There is a budget feasible truthful mechanism for subadditive functions with approximation ratio $O(\mathcal{I})$. In particular, for XOS functions, the mechanism has a constant approximation ratio.*

Our results show an interesting connection between the budget feasible mechanism design and the integrality gap of the above linear program, as well as the existence of an α -approximate core. Note that the tight bound of the integrality gap is known to be $\Theta(\log n)$ [19, 8]; thus, our mechanism in the worst case has an approximation ratio of $O(\log n)$. (But for some special functions whose integrality gaps are bounded by constants, e.g., facility location [33], our mechanism gives a constant approximation.) Further, the mechanism may have exponential running time, though for some special XOS functions like matching it can be implemented in polynomial time. To remedy these issues, we further give a polynomial time sub-logarithmic approxima-

tion mechanism. Both of our mechanisms improve the best known approximation ratio $O(\log^2 n)$ of [20].

Theorem 2. *There is a polynomial time budget feasible truthful mechanism for subadditive functions with an approximation ratio $O(\frac{\log n}{\log \log n})$, where n is the number of agents.*

Bayesian mechanism design.

As a standard game theoretic model for incomplete information, Bayesian mechanism design assumes that agents' private information (i.e., $c(i)$ in our model) is drawn from a known distribution. In contrast to prior-free worst case analysis, if we have prior knowledge of the distribution, we can obtain more positive results in the form of constant approximation truthful mechanisms. In the Bayesian setting, we are able to answer the above question posed in [20] affirmatively.

Theorem 3. *There is a constant approximation budget feasible truthful mechanism for subadditive functions for any distribution under a mild assumption².*

It should be noted that our result does not completely rely on Bayesian analysis in the following aspects.

- **Truthfulness.** In most of the previous works in Bayesian mechanism design regarding social welfare maximization, e.g., [26, 5, 25, 12], the considered solution concept is Bayesian truthfulness, i.e., truth-telling is in expectation an equilibrium strategy when other agents' profiles are drawn from the known distribution. Our mechanism guarantees *universal truthfulness*, meaning that truth-telling is a dominant strategy of each agent for any coin flips of the mechanism and any instance of the costs. Thus universal truthfulness is a stronger solution concept than Bayesian truthfulness. Universal truthfulness in Bayesian mechanism design has also been used in, e.g. [11], but their focus was on profit maximization.
- **Distribution.** Regarding prior knowledge of the distribution, most of the previous related works consider independent distributions, e.g., [27, 26, 5, 25]. Our mechanism applies to general distributions that allow correlations on costs. Dependency on private information is a natural phenomenon arising in practice and it has been considered for, e.g., auctions [30]. In our model where costs are private information, correlations appear to be very common. For example, if the price of crude oil goes up, the costs of producing the items for *all* agents may go up as well.

Techniques.

In the design of budget feasible mechanisms, the major approach used in previous works [35, 14, 20] is based on a simple idea of adding agents one by one greedily and carefully ensuring that the budget constraint is not violated. Our mechanisms, from a high level structural point of view, use another simple, but powerful, approach: random sampling. We add agents into a test set T with probability

²Technically, we require that the distribution has integrable marginal densities for any subset of variables (e.g., jointly independent distributions trivially satisfy this condition). The formal definition is referred to Section 4.

half for each agent and compute an (approximately) optimal budget feasible solution on T . We use the agents in T only for the purpose of 'evaluation' and none of them will be a winner. The computed solution on T gives a close estimate of the optimal solution for the whole set with a high probability. We then compute a real winning set from the remaining agents using the evaluation from T as a threshold.

In the Bayesian setting, random sampling is often deemed to be unnecessary, because, when we have knowledge of the distribution, it is tempting to use a 'prior sampling' approach to generate random virtual instances and based on them to compute a threshold. While this works well when the private cost $c(i)$ of every agent is drawn independently, interestingly (and surprisingly), it fails when costs $c(i)$'s are correlated in the distribution. We therefore still use random sampling to compute a threshold based on the sampled test set; the collected information from random sampling correctly reflects the structure of the private costs (with a high probability) even for correlated distributions. To derive a constant approximation budget feasible mechanism for subadditive functions, we first generate a cost vector sampled from the known distribution and use it as a guidance for the payments to the winners. Then we adopt our (constant approximation) mechanism for prior-free XOS functions by feeding to this mechanism another valuation $\tilde{v}(\cdot)$, which we define as the solution of the above LP computed for the various subsets of A .

Random sampling appears to be a powerful approach and has been used successfully in other domains of mechanism design, e.g., digital goods auctions [23], secretary problem [3, 2], social welfare maximization [19], and mechanism design without money [15]. It is intriguing to find applications of random sampling in other mechanism design problems.

1.2 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 [32]. There are many mechanism design models; see, e.g., [33], for a survey.

As mentioned earlier, the study of approximate mechanism design with a budget constraint was originated by Singer [35] and constant approximation mechanisms were given for additive and submodular functions. The approximation ratios were later improved in [14]. Dobzinski, Papadimitriou, and Singer [20] considered subadditive functions and showed an $O(\log^2 n)$ approximation mechanism. Ghosh and Roth [22] 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 [26] 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 blackbox reduction results were later improved to multi-parameter settings in [5] and [25] independently. Chawla et al. [12] 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., [27, 7, 13, 18, 11, 17]. Ours is the first to consider Bayesian analysis in budget feasible mechanisms with a focus on the valuation (social welfare) maximization.

2. PRELIMINARIES

In a marketplace, there are n agents (or items), denoted by A . Each agent $i \in A$ has a privately known incurred cost $c(i) \geq 0$. We denote by $c = (c(i))_{i \in A}$ the cost vector of the agents. For any given subset $S \subseteq A$ of agents, there is a publicly known valuation $v(S)$, meaning the social welfare derived from S . We assume that $v(\emptyset) = 0$ and the valuation function is monotone, i.e., $v(S) \leq v(T)$ for any $S \subseteq T \subseteq A$. A centralized authority wants to pick a subset of agents with maximum possible valuation given a sharp budget B to cover their incurred costs, i.e., $\max_{S \subseteq A} v(S)$ given $c(S) = \sum_{i \in S} c(i) \leq B$. We denote the optimal solution of this optimization question by $\text{opt}(A)$ (or $\text{opt}(c)$) and its valuation by $v(\text{opt}(A))$.

We will consider XOS and subadditive functions 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): there is a set of linear functions $f_1, \dots, f_m$ such that

$$v(S) = \max \{f_1(S), f_2(S), \dots, f_m(S)\}$$

for any $S \subseteq A$. Note that the number of functions m can be exponential in $n = |A|$.

An equivalent definition [21] is $v(S) \leq \sum_{T \subseteq A} x(T) \cdot v(T)$ whenever $\sum_{T \subseteq A: i \in T} x(T) \geq 1$ for any $i \in S$, where $0 \leq x(T) \leq 1$. That is, if every element in S is fractionally covered, then the sum of the values of all subsets weighted by the corresponding fractions is at least as large as $v(S)$.

Note that the representation of a valuation function usually requires exponential size in n . Thus, we assume that we are given an access to a *demand oracle*, which, for any given price vector $p(1), \dots, p(n)$, returns us a subset $T \in \arg\max_{S \subseteq A} (v(S) - \sum_{i \in S} p(i))$; every such query is assumed to take unit time. The demand oracle is used in [20] as well, and it was shown that a weaker value query oracle is not sufficient [35].

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. We use mechanism design and its solution concept truthfulness to manage strategic behaviors of the agents. Upon receiving $b(i)$ from each agent, a mechanism decides 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, agents bid strategically to maximize their utilities, which is $p(i) - c(i)$ if i is a winner and 0 otherwise. We say a mechanism is *truthful* if it is of the best interest for each agent to report his 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 takes a distribution over deterministic truthful mechanisms.

Note that our model is in a single parameter domain, as each agent has only one private parameter which is his cost.

Thus, by the well known characterization of single parameter truthful mechanisms [31], designing a monotone allocation, plus the corresponding threshold payment rule, is sufficient to derive a truthful mechanism. We therefore do not specify the payments to the winners in our mechanisms explicitly.

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 the present paper is to design truthful and budget feasible mechanisms for XOS and subadditive functions in two frameworks: prior-free and Bayesian.

We first establish the following technical lemma, which is useful in the analysis of our mechanisms.

LEMMA 2.1. *Consider any subadditive function $v(\cdot)$. For any 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 $v(T_2) \geq \frac{k-1}{4k} v(S)$.*

3. PRIOR-FREE MECHANISM DESIGN

In this section we consider designing budget feasible mechanisms for XOS and subadditive functions in the prior-free setting. That is, the mechanism designer has no prior knowledge of the private information $c(i)$ of every agent, which can be an arbitrary cost, and the performance of a designed mechanism is analyzed in the worst case framework. That is, 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. (We assume without loss of generality that $c(i) \leq B$ for any $i \in A$, since such an agent will never win in any budget feasible truthful mechanism.)

3.1 Constant Approximation for XOS

We will first consider XOS functions. Given an XOS function $v(\cdot)$, by its definition, we assume that

$$v(S) = \max \{f_1(S), f_2(S), \dots, f_m(S)\}$$

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 a random 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, [14]).

XOS-RANDOM-SAMPLE

1. Pick each item independently at random 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. Consider items in $A \setminus T$ and find a set $S^* \in \arg\max_{S \subseteq A \setminus T} \{v(S) - t \cdot c(S)\}$.
5. Find an additive function f with $f(S^*) = v(S^*)$ in the XOS representation of $v(\cdot)$.
6. Run ADDITIVE-MECHANISM for function $f(\cdot)$ with respect to set S^* and budget B .
7. Output the result of ADDITIVE-MECHANISM.

In the above mechanism, we first sample in expectation half items to form a testing group T , and compute an optimal solution for T given 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 keep a nearly optimal solution formed with the remaining items. We then use the information from random sampling to compute a proper threshold t for the rest of items. Specifically, 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, we break ties by any given fixed order). Finally, we use the property of XOS functions to find a linear representation of $v(S^*)$ and run a truthful mechanism for linear functions with respect to S^* .

The mechanism is designated for XOS functions; it is also used crucially as an auxiliary procedure for the more general subadditive functions in the subsequent sections. Note that the runtime of the mechanism for general XOS functions is exponential³.

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

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

(note that $x \cdot B$ 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 targeted value $v(\text{opt}(A))$.

The linearity of the Lagrangian, together with the subadditivity of the valuations, is important in order to derive the following properties. (The threshold t , subset S^* , and additive function f are defined in the XOS-RANDOM-SAMPLE.)

CLAIM 3.1. *For any $S \subseteq S^*$, $f(S) - t \cdot c(S) \geq 0$.*

PROOF. Suppose by a contradiction that there exists a subset $S \subseteq S^*$ such that $f(S) - t \cdot c(S) < 0$. 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 item in S^* cannot manipulate the selection of the set S^* if bidding a smaller cost. This fact is critical for the monotonicity, and thus, the truthfulness of the mechanism.

³In fact, in the second step of the mechanism, we can use any constant approximation solution (e.g., algorithm SA-ALG-MAX established in Section 3.3), which suffices for our purpose. Further, Step (4) can be done by simply asking a demand query. Hence, the mechanism can be implemented in polynomial time, if we have access to an oracle that, for any subset X of items, gives a linear function f with $f(X) = v(X)$ and $f(S) \leq v(S)$ for each $S \subseteq X$. For some classic XOS problems like matching (the value of a subset of edges is the size of the largest matching induced by them), our mechanism can be implemented in polynomial time.

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

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

XOS-MECHANISM-MAIN

- With half probability, run XOS-RANDOM-SAMPLE.
- With half probability, pick a most-valuable item as the only winner and pay him B .

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 remaining of this section, we will prove the theorem. It follows from the following three lemmas.

LEMMA 3.1. *XOS-MECHANISM-MAIN is universally truthful.*

Our mechanism, at a high level point of view, has a similar flavor to the mechanism composition introduced in [1]. In particular, we may consider Steps (1-4) as one mechanism of choosing candidate winners and Steps (5-7) as the other restricted on the survived 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 one is truthful, then the composite 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 property of ADDITIVE-MECHANISM. Therefore, the mechanism is truthful.

LEMMA 3.2. *XOS-MECHANISM-MAIN is budget feasible.*

In the mechanism XOS-RANDOM-SAMPLE, 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 (e.g., Step (4) and (6)). In particular, the payment is upper bounded by the threshold of the mechanism ADDITIVE-MECHANISM in Step (6). As ADDITIVE-MECHANISM is budget feasible [14], our mechanism XOS-RANDOM-SAMPLE is budget feasible as well. Finally, picking the largest item and paying it the whole budget are clearly a budget feasible mechanism.

LEMMA 3.3. XOS-MECHANISM-MAIN has a constant approximation ratio.

PROOF. Let $\text{opt} = \text{opt}(A)$ denote the optimal winning set given 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 social optimum $v(\text{opt})$. Consider the following two cases:

- $c(S^*) > B$. With such assumption, we can always find a subset $S' \subseteq S^*$, such that $\frac{B}{2} \leq c(S') \leq B$. By Claim 3.1, we know $f(S') \geq t \cdot c(S') \geq \frac{v(\text{opt}(T))}{8B} \cdot \frac{B}{2} \geq \frac{v(\text{opt}(T))}{16}$. 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 $S^* = \text{argmax}_{S \subseteq A \setminus T} (v(S) - t \cdot c(S))$. 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-SAMPLE. 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-SAMPLE 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 $\max_i v(i) \geq \frac{1}{k} v(\text{opt})$. Combining the two mechanisms together, our main mechanism XOS-MECHANISM-MAIN has a performance 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. $\square$

3.2 Integrality-Gap Approximations for Sub-additive

Next we use our result for XOS functions to design a truthful mechanism for subadditive functions. Let $S_1, \dots, S_N$ be a permutation of all possible subsets of A , where $N = |2^A|$ is the size of the power set 2^A . We consider the following linear program for each subset $S \subseteq A$, where each subset S_j is associated with a variable α_j .

$$\begin{aligned} LP(S) : \quad & \min \sum_{j=1}^N \alpha_j \cdot v(S_j) & (\diamond) \\ \text{s.t.} \quad & \alpha_j \geq 0, \quad 1 \leq j \leq N \\ & \sum_{j: i \in S_j} \alpha_j \geq 1, \quad \forall i \in S \end{aligned}$$

In the above linear program, the minimum is taken over all possible non-negative values of $\alpha = (\alpha_1, \dots, \alpha_N)$. If we consider each α_j as the fraction of the cover by subset S_j , the last constraint implies that all items in S are fractionally covered. Hence, $LP(S)$ describes a linear program for the set cover of S . For any subadditive function $v(\cdot)$, it can be seen that the value of the optimal integral solution to the above $LP(S)$ is always $v(S)$. Indeed, one has $S \subseteq \bigcup_{j: \alpha_j \geq 1} S_j$ and $\sum_j \alpha_j \cdot v(S_j) \geq \sum_{j: \alpha_j \geq 1} v(S_j) \geq v(\bigcup_{j: \alpha_j \geq 1} S_j) \geq v(S)$.

Let $\tilde{v}(S)$ be the value of the optimal fractional solution of $LP(S)$, and $\mathcal{I}(S) = \frac{v(S)}{\tilde{v}(S)}$ be the integrality gap of $LP(S)$. Let $\mathcal{I} = \max_{S \subseteq A} \mathcal{I}(S)$; the integrality gap $\mathcal{I}$ gives a worst case upper bound on the integrality gap of all subsets. Hence, we have $\frac{v(S)}{\mathcal{I}} \leq \tilde{v}(S) \leq v(S)$ for any $S \subseteq A$. The classic Bondareva-Shapley Theorem [9, 34] says that the integrality gap $\mathcal{I}(S)$ is one (i.e., $v(S)$ is also an optimal fractional solution to the LP) if and only if $v(\cdot)$ is an XOS function.

LEMMA 3.4. $\tilde{v}(\cdot)$ is an XOS function.

We are now ready to present our mechanism for subadditive functions.

SA-MECHANISM-MAIN

1. For each subset $S \subseteq A$, compute $\tilde{v}(S)$.
2. Run XOS-MECHANISM-MAIN for the instance with respect to the XOS function $\tilde{v}(\cdot)$.
3. Output the result of XOS-MECHANISM-MAIN.

THEOREM 3.2. The mechanism SA-MECHANISM-MAIN is budget feasible and truthful, and provides an approximation ratio of $O(\mathcal{I})$ for subadditive functions, where $\mathcal{I}$ is the largest integrality gap of $LP(S)$ for all $S \subseteq A$.

PROOF. Note that the valuation $v(\cdot)$ is public knowledge and utilities of agents do not depend on $v(\cdot)$; thus computing $\tilde{v}(\cdot)$ and running XOS-MECHANISM-MAIN with respect to $\tilde{v}(\cdot)$ do not affect truthfulness. The claim then follows from Theorem 3.1 and the fact that $\frac{v(S)}{\mathcal{I}} \leq \tilde{v}(S) \leq v(S)$ for any $S \subseteq A$ (i.e., by using $\tilde{v}(\cdot)$ instead of $v(\cdot)$ we lose at most factor of $\mathcal{I}$ in the approximation ratio). $\square$

In general, the approximation ratio of the mechanism can be as large as $\Theta(\log n)$ [19, 8]. But for those instances when the integrality gap of $(\diamond)$ is bounded by a constant (e.g., facility location [33]), our mechanism gives a constant approximation.

3.3 Sub-Logarithmic Approximations for Sub-additive

In this section, we give another mechanism for subadditive functions based on the ideas of random sampling and cost sharing. In contrast to the previous section, the mechanism runs in polynomial time and has an $o(\log n)$ approximation ratio, improving the previously best known ratio $O(\log^2 n)$ [20]. Our mechanism relies on a constant factor approximation *algorithm* for subadditive function maximization under a knapsack constraint, which may have its own interest.

3.3.1 Subadditive Maximization with Budget

We first give an algorithm that approximates $\max_{S \subseteq A} v(S)$ given that $c(S) \leq B$. That is, we ignore for a while strategic behaviors of the agents and consider a pure optimization problem. Dobzinski et al. [20] considered the same question and gave a 4-approximation algorithm for the unweighted case (i.e., the restriction is on the size of a selected subset). Our algorithm extends their result to the weighted case and runs in polynomial time if we are given a demand oracle⁴.

SA-ALG-MAX

- Let $v^* = \max_{i \in A} v(i)$ and $\mathcal{V} = \{v^*, 2v^*, \dots, nv^*\}$
- For each $v \in \mathcal{V}$
 - Set $p(i) = \frac{v}{2B} \cdot c(i)$ for each $i \in A$, and find $T \in \arg\max_{S \subseteq A} (v(S) - \sum_{i \in S} p(i))$.
 - Let $S_v = \emptyset$.
 - If $v(T) < \frac{v}{2}$, then continue to next v .
 - Else, in decreasing order of $c(i)$ put items from T into S_v while preserving the budget constraint.
- Output: S_v with the largest value $v(S_v)$ for all $v \in \mathcal{V}$.

LEMMA 3.5. SA-ALG-MAX is an 8-approximation algorithm for subadditive maximization given a demand oracle.

Note that we can actually modify the algorithm to get a $(4 + \epsilon)$ -approximation with running time polynomial in n and $\frac{1}{\epsilon}$. To do so one may simply replace $\mathcal{V}$ by a larger set $\{\epsilon v^*, 2\epsilon v^*, \dots, \lceil \frac{n}{\epsilon} \rceil \epsilon v^*\}$. Both algorithms suffice for our purpose; for the rest of the paper, for simplicity we will use the 8-approximation algorithm to avoid the extra parameter ϵ in the description.

We will use SA-ALG-MAX as a subroutine to build a mechanism SA-RANDOM-SAMPLE for subadditive functions in the next subsection. When there are different sets maximizing $v(S) - \sum_{i \in S} p(i)$, we require that the demand query oracle always returns a fixed set T . This property is important for the truthfulness of our mechanism. To implement this, we set a fixed order on all the items $i_1 \prec i_2 \prec \dots \prec i_n$. We first compute

$$T_1 \in \arg\max_{S \subseteq A} (v(S) - \sum_{i \in S} p(i))$$

and

$$T_2 \in \arg\max_{S \subseteq A \setminus \{i_1\}} (v(S) - \sum_{i \in S} p(i)).$$

⁴Independent to our work, Badanidiyuru et al. [4] gave a $2 + \epsilon$ approximation algorithm to the same weighted problem.

If $v(T_1) - \sum_{i \in T_1} p(i) = v(T_2) - \sum_{i \in T_2} p(i)$, we know that there is a subset without i_1 that gives us the maximum; thus, we ignore i_1 for further consideration. If $v(T_1) - \sum_{i \in T_1} p(i) > v(T_2) - \sum_{i \in T_2} p(i)$, we know that i_1 should be in any optimal solution; hence, we keep i_1 and proceed with the process iteratively for $i_2, i_3, \dots, i_n$. This process clearly gives a fixed outcome that maximizes $v(S) - \sum_{i \in S} p(i)$.

3.3.2 Mechanism

Let us first consider the following mechanism based on random sampling and cost sharing.

SA-RANDOM-SAMPLE

1. Pick each item independently at random with a probability of $\frac{1}{2}$ into group T .
2. Run SA-ALG-MAX for items in group T , and let v be the value of the returned subset.
3. For $k = 1$ to $|A \setminus T|$
 - Run SA-ALG-MAX on the set $\{i \in A \setminus T \mid c(i) \leq \frac{B}{k}\}$ where each item has cost $\frac{B}{k}$, denote the output by X .
 - If $v(X) \geq \frac{\log \log n}{80 \log n} \cdot v$
 - Output X as the winning set and pay $\frac{B}{k}$ to each item in X .
 - Halt.
4. Output $\emptyset$.

In the above mechanism, we again first sample in expectation half of the items to form a testing group T , and then use SA-ALG-MAX to compute an approximate solution for the items in T given the budget constraint B . As can be seen in the analysis of the mechanism, the computed value v is in expectation within a constant factor of the optimal value of the whole set A . That is, we are able to learn the rough value of the optimal solution by random sampling. Next we consider the remaining items $A \setminus T$ and try to find a subset X with a relatively big value in which every item is willing to “share” the budget B at a fixed share $\frac{B}{k}$. (This part of our mechanism can be viewed as a reversion of the classic cost sharing mechanism.) Finally, we use the information v from random sampling as a benchmark to determine whether X should be a winning set or not.

The final mechanism for subadditive functions is described as follows.

SA-MECHANISM-MAIN-2

- With half probability, run SA-RANDOM-SAMPLE.
- With half probability, pick a most-valuable item as the only winner and pay him B .

THEOREM 3.3. 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 $O(\frac{\log n}{\log \log n})$.

4. BAYESIAN MECHANISM DESIGN

In this section, we study budget feasible mechanisms for subadditive functions 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 $\mathcal{D}$ in advance from which the real cost vector $(c(i))_{i \in A}$ is drawn. However, each $c(i)$ is the private information of agent i . Distribution $\mathcal{D}$ is given on the probability space Ω with the corresponding density function $\rho(\cdot)$ on $\mathbb{R}^{|A|}$. We allow dependencies on the agents' costs in $\mathcal{D}$ and consider the distributions that have integrable marginal densities for any subset of variables⁵; this includes, e.g., independent distributions as special cases.

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 coin flips of the mechanism and each cost vector, truth-telling is 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 mechanism $\mathcal{M}$ is a (Bayesian) α -approximation if $\frac{\mathbf{E}[\text{opt}]}{\mathbf{E}[\mathcal{M}]} \leq \alpha$.

In 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 is as follows.

SA-BAYESIAN-MECHANISM

- With a probability of $\frac{1}{2}$, let a most-valuable item be the only winner and pay him B .
- With a probability of $\frac{1}{2}$, run the following:
 1. Pick each item independently at random with a probability of $\frac{1}{2}$ 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 a 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. If $d(S^*) < B$, let all $i \in S^*$ with $c(i) \leq d(i)$ be the winners.
 7. If $d(S^*) \geq B$,
 - run XOS-MECHANISM-MAIN w.r.t. valuation $\tilde{v}(\cdot)$, set S^* , cost $c(\cdot)$, and budget B .
 - Output the result of XOS-MECHANISM-MAIN.

In the mechanism, Steps (1-3) are the same as XOS-RANDOM-SAMPLE where we randomly sample a test group T and generate a threshold value t . In Steps (4-7), we consider a specific subset $S^* \subseteq A \setminus T$ and select winners only

⁵We need some mild technical restriction on $\mathcal{D}$ in order to sample a conditioned random variable. 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 integrability of marginal density functions (see, e.g., page 331 of [36]), though in our case we require a slightly stronger condition.

inside of it. Step (5) helps to give us a guidance on the threshold payments of the winners (see more discussions below). Step (7) runs XOS-MECHANISM-MAIN on the function $\tilde{v}(\cdot)$ (defined as the optimal value of the LP ($\diamond$)), which is XOS according to Lemma 3.4.

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 only applies to 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 approximate core. For subadditive functions, a constant approximate core may not exist [33] (e.g., set cover gives a logarithmic lower bound [8]). 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), from a high level point of view, guarantee that the vector d is not too ‘far’ from c for the agents in S^* , in the sense that both vectors are derived from the same distribution. Thus, cost vectors c and d are distributed symmetrically and can be switched while preserving

some important parameters such as t and S^* in expectation.) If $d(S^*) \leq B$ (Step (6)), then we set $d(i)$ as an upper bound on the payment of each agent $i \in S^*$; this guarantees that we are always within the budget constraint. If $d(S^*) > B$, setting $d(i)$ as an upper bound is not sufficient to ensure budget feasibility; then we adopt our approach for XOS functions with inputs subset S^* and XOS valuation $\tilde{v}(\cdot)$ defined by $(\diamond)$.

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

Budget feasibility follows simply from the description of the mechanism and the fact that XOS-MECHANISM-MAIN is budget feasible.

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 describe all possible samples d for (i) all possible cost vectors on T and (ii) all possible choices $S \subseteq A \setminus T$ of S^* ; then we tell all flipped d 's to the agents before looking at the costs of T . (Practically, we can provide all our randomness as a black box accessible by all agents.) Note that the selection rule of S^* is monotone, and, similarly to Claim 3.2, each agent in S^* cannot manipulate (i) the composition of S^* given c and T , and (ii) the choice of d , as long as he stays in S^* . Therefore, composing the first part choosing S^* (Step (4)) with the next monotone rule picking winners in S^* (Steps (6-7)), we again get a monotone winner selection rule. Hence, the mechanism is universally truthful.

Next we give a sketch of the idea of proving the constant approximation. Details of the proof are deferred to the full version of the paper.

Approximation analysis (sketch). We sketch the proof idea of the approximation ratio of the mechanism. First, similar to our analysis in Section 3.1, the optimal solution $v(\text{opt}(c, T))$ obtained from random sampling in expectation gives a constant approximation to the optimal solution $\mathbf{E}[\text{opt}]$. Further, we observe the following facts (which are reminiscent of Claim 3.1):

$$\tilde{v}(S) - t \cdot c(S) \geq 0 \text{ and } \tilde{v}(S) - t \cdot d(S) \geq 0, \quad \forall S \subseteq S^*$$

where the second inequality is based on the conditional distribution we choose for d .

If $c(S^*) \geq B$ and $d(S^*) \geq B$ (i.e., the mechanism runs Step (7)), we can pick a subset $S^0 \subseteq S^*$ with $B \geq c(S^0) \geq \frac{B}{2}$. By Theorem 3.1, XOS-MECHANISM-MAIN gives a constant approximation to the optimum of $\tilde{v}(\cdot)$ on S^* . (This is the reason why in Step (7) of the mechanism, we run the whole XOS-MECHANISM-MAIN on the input instance $\tilde{v}(\cdot)$ and S^* .) Hence,

$$\tilde{v}(\text{opt}_{\tilde{v}}(c, S^*)) \geq \tilde{v}(S^0) \geq t \cdot c(S^0) \geq t \cdot \frac{B}{2} \geq \frac{v(\text{opt}(c, T))}{16},$$

where the first inequality follows from the fact that $S^0 \subseteq S^*$ is a budget feasible set. Thus, the optimum of $\tilde{v}(\cdot)$ on S^* is within a constant factor of $v(\text{opt}(c, T))$, as well as the benchmark $\mathbf{E}[\text{opt}]$.

If $c(S^*) < B$ and $d(S^*) \geq B$, we have $\tilde{v}(S^*) \geq t \cdot d(S^*) \geq \frac{v(\text{opt}(c, T))}{8}$. Further, we notice that S^* is budget feasible with respect to cost vector c ; thus, XOS-MECHANISM-MAIN gives

a constant approximation to $\tilde{v}(S^*)$, which in turn is within a constant factor of $v(\text{opt}(c, T))$ and $\mathbf{E}[\text{opt}]$.

We observe that the vectors d and c are restricted to the agents in S^* and conditioned on

$$S^* \in \operatorname{argmax}_{S \subseteq A \setminus T} \{v(S) - t \cdot d(S)\}$$

and

$$S^* \in \operatorname{argmax}_{S \subseteq A \setminus T} \{v(S) - t \cdot c(S)\}$$

have exactly the same distributions. Therefore, due to such a symmetry between d and c , in a run of our mechanism in expectation we will have the outcome T, t, S^* and a pair of vectors (c, d) as often as the outcome T, t, S^* and the pair (d, c) . This implies that in the case when $d(S^*) < B$ and $c(S^*) < B$, we get on average a value of at least $\frac{1}{2}v(S^*)$, since (i) the winning sets on the two instances where c (resp., d) is the private cost and d (resp., c) is the sampled cost altogether cover S^* , and (ii) v is a subadditive function. By the choice of threshold t , we also know that

$$\begin{aligned} v(S^*) &\geq v(S^*) - t \cdot c(S^*) \\ &\geq v(\text{opt}(c, A \setminus T)) - t \cdot c(\text{opt}(c, A \setminus T)) \\ &\geq v(\text{opt}(c, A \setminus T)) - t \cdot B. \end{aligned}$$

Thus, our mechanism gives a constant approximation to $v(\text{opt}(c))$ with some constant probability.

The last case is when $c(S^*) \geq B$ and $d(S^*) < B$. Again due to the symmetry between c and d , intuitively, we can treat this case as the above one when c and d are switched; thus we also get a constant approximation of $\mathbf{E}[\text{opt}]$. (The formal argument, however, due to multiple randomness used in the mechanism, is much more complicated.)

Therefore, the mechanism SA-BAYESIAN-MECHANISM on average has a constant approximation of the expected socially optimal value $\mathbf{E}[\text{opt}]$. $\square$

5. CONCLUSIONS

Our work considers budget feasible mechanism design in two analysis frameworks: prior-free and Bayesian. For XOS functions, we give a prior-free constant approximation mechanism. For subadditive functions, we present two prior-free mechanisms with integrality-gap and sub-logarithmic approximations, respectively, as well as a Bayesian constant approximation mechanism. All our mechanisms are universally truthful.

Our mechanisms continue to work for the extension when the valuation functions are non-monotone, 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 [20] is non-monotone. 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.

We give a constant approximation mechanism for subadditive functions in the Bayesian framework where the costs are drawn from a given *known* distribution. Considering the gap between Bayesian and prior-free, a natural question is under which frameworks a constant approximation mechanism still exists. An interesting step along this direction is the prior-independent setting, where the costs are

still drawn from an underlying distribution, but the mechanism designer does not have the prior knowledge of it. Our mechanism SA-BAYESIAN-MECHANISM can be adopted to the framework where all the costs are identically and independently distributed. (Specifically, by random sampling we are able to learn the underlying unknown distribution with sufficient precision.) However, we do not know how to handle independent but not necessarily identical distributions in the prior-independent framework, as well as the most general prior-free setting in the worst-case analysis.

Indeed, whether subadditive functions admit a prior-free constant approximation mechanism still remains an open problem. Our results show a separation between XOS and subadditive functions. Another angle to have such a distinction between the two classes is from exponential concentration: In the case of XOS the valuation of a randomly selected subset obeys an exponential concentration around its expected value⁶, whereas in the case of general subadditive valuations it does not (see [39] for a counterexample). Such a difference on exponential concentration may suggest a possible distinction between XOS and subadditive functions in terms of their approximability in (budget feasible) mechanism design.

For those mechanisms with exponential runtime, it is natural to ask if 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 as future work.

Acknowledgements

This work was supported in part by the National Basic Research Program of China Grant 2011CBA00300, 2011CBA00301, the National Natural Science Foundation of China Grant 61033001, 61061130540, 61073174.

6. REFERENCES

- [1] G. Aggarwal, J. Hartline, *Knapsack Auctions*, SODA 2006, 1083-1092.
- [2] M. Babaioff, M. Dinitz, A. Gupta, N. Immorlica, K. Talwar: *Secretary Problems: Weights and Discounts*, SODA 2009, 1245-1254.
- [3] M. Babaioff, N. Immorlica, R. Kleinberg, *Matroids, Secretary Problems, and Online Mechanisms*, SODA 2007, 434-443.
- [4] A. Badanidiyuru, S. Dobzinski, S. Oren, *Optimization with Demand Oracles*, arXiv.org/abs/1107.2869.
- [5] X. Bei, Z. Huang, *Bayesian Incentive Compatibility via Fractional Assignments*, SODA 2011, 720-733.
- [6] A. Berger, V. Bonifaci, F. Grandoni, G. Schafer, *Budgeted Matching and Budgeted Matroid Intersection via the Gasoline Puzzle*, Math. Program., V.128, 355-372, 2011.
- [7] S. Bhattacharya, G. Goel, S. Gollapudi, K. Munagala, *Budget Constrained Auctions with Heterogeneous Items*, STOC 2010, 379-388.
- [8] K. Bhawalkar, T. Roughgarden, *Welfare Guarantees for Combinatorial Auctions with Item Bidding*, SODA 2011, 700-709.
- [9] O. Bondareva, *Some Applications of Linear Programming to Cooperative Games*, Problemy Kibernetiki, V.10, 119-139, 1963.
- [10] S. Boucheron, G. Lugosi, P. Massart. *Concentration Inequalities Using the Entropy Method*, Annals of Probability, V.31, 1583-1614, 2003.
- [11] S. Chawla, J. Hartline, D. Malec, and B. Sivan, *Multi-Parameter Mechanism Design and Sequential Posted Pricing*, STOC 2010, 311-320.
- [12] S. Chawla, D. Malec, A. Malekian, *Bayesian Mechanism Design for Budget-Constrained Agents*, EC 2011, 253-262.
- [13] S. Chawla, D. Malec, B. Sivan, *The Power of Randomness in Bayesian Optimal Mechanism Design*, EC 2010, 149-158.
- [14] N. Chen, N. Gravin, P. Lu, *On the Approximability of Budget Feasible Mechanisms*, SODA 2011, 685-699.
- [15] N. Chen, N. Gravin, P. Lu, *Mechanism Design without Money via Stable Matching*, arxiv.org/abs/1104.2872.
- [16] E. Clarke, *Multipart Pricing of Public Goods*, Public Choice, V.11, 17-33, 1971.
- [17] N. Devanur, J. Hartline, A. Karlin, T. Nguyen, *A Prior-Independent Mechanism for Profit Maximization in Unit-demand Combinatorial Auctions*, WINE 2011.
- [18] P. Dhangwatnotai, T. Roughgarden, Q. Yan, *Revenue Maximization with a Single Sample*, EC 2010, 129-138.
- [19] S. Dobzinski, *Two Randomized Mechanisms for Combinatorial Auctions*, APPROX 2007, 89-103.
- [20] S. Dobzinski, C. Papadimitriou, Y. Singer, *Mechanisms for Complement Free Procurement*, EC 2011, 273-282.
- [21] U. Feige, *On Maximizing Welfare When Utility Functions Are Subadditive*, SIAM J. Computing, V.39, 122-142, 2009.
- [22] A. Ghosh, A. Roth, *Selling Privacy at Auction*, EC 2011, 199-208.
- [23] A. Goldberg, J. Hartline, A. Karlin, M. Saks, A. Wright, *Competitive Auctions*, Games and Economic Behavior, V.55(2), 242-269, 2006.
- [24] T. Groves, *Incentives in Teams*, Econometrica, V.41, 617-631, 1973.
- [25] J. Hartline, R. Kleinberg, A. Malekian, *Bayesian Incentive Compatibility via Matchings*, SODA 2011, 734-747.
- [26] J. Hartline, B. Lucier, *Bayesian Algorithmic Mechanism Design*, STOC 2010, 301-310.
- [27] J. Hartline, T. Roughgarden, *Optimal Mechanism Design and Money Burning*, STOC 2008, 75-84.
- [28] J. Hartline, T. Roughgarden, *Simple versus Optimal Mechanisms*, EC 2009, 225-234.
- [29] B. Lehmann, D. Lehmann, N. Nisan, *Combinatorial Auctions with Decreasing Marginal Utilities*, EC 2001, 18-28.
- [30] P. Milgrom, *Putting Auction Theory to Work*, Cambridge University Press, 2004.
- [31] R. Myerson, *Optimal Auction Design*, Mathematics of Operations Research, V.6(1), 1981.
- [32] N. Nisan, A. Ronen, *Algorithmic Mechanism Design*, STOC 1999, 129-140.
- [33] N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, *Algorithmic Game Theory*, Cambridge University Press, 2007.
- [34] L. Shapley, *On Balanced Sets and Cores*, Naval Research Logistics Quarterly, V.14, 453-460, 1967.
- [35] Y. Singer, *Budget Feasible Mechanisms*, FOCS 2010, 765-774.
- [36] C. Stone, *A Course in Probability and Statistics*, Duxbury Press, 1995.
- [37] M. Sviridenko, *A Note on Maximizing a Submodular Set Function Subject to Knapsack Constraint*, Operations Research Letters, V.32, 41-43, 2004.
- [38] W. Vickrey, *Counterspeculation, Auctions and Competitive Sealed Tenders*, Journal of Finance, V.16, 8-37, 1961.
- [39] J. Vondrak, *A Note on Concentration of Submodular Functions*, arxiv.org/abs/1005.2791.

⁶Note that this fact can be used to improve the approximation ratios of the mechanisms of XOS functions substantially (but still up to a constant factor).