Bayesian Optimal No-deficit Mechanism Design

Shuchi Chawla[1] Jason Hartline[2] Uday Rajan[3]
R. Ravi[4]

October 26, 2004
CMU-CS-04-153

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

One of the most fundamental problems in mechanism design is that of designing the auction that gives the optimal profit to the auctioneer. For the case that the probability distribution on the valuations of the bidders is known and independent, Myerson [13] reduces the problem to that of computing the *efficient allocation* on the bidders' *virtual valuations*. We review this technique for any single parameter agent problem, i.e., where there is an arbitrary cost function on the outcome that the mechanism produces. Further, we consider the problem of merging the worst case *no-deficit* condition with this average case *Bayesian expected profit maximization* problem. When restricting our attention to ex post incentive compatible mechanisms for this problem, we find that the Myerson mechanism is the optimal no-deficit mechanism for *supermodular* costs, that Myerson merged with a simple *thresholding mechanism* is optimal for *all-or-nothing* costs, and that neither mechanism is optimal for general *submodular* costs. Addressing computational side of the problem, for supermodular costs the Myerson mechanism is hard to compute. For all-or-nothing costs we show that the optimal thresholding mechanism is NP-hard to compute. Finally, for submodular costs if we relax the ex post incentive compatibility constraint, we can specify a Bayesian incentive compatible mechanism that achieves the same expected profit as Myerson, but never has a loss.

[1]Computer Science Department, Carnegie Mellon University, E-mail: `shuchi@cs.cmu.edu`
[2]Microsoft Research, Silicon Valley Campus, Mountain View, CA, E-mail: `hartline@microsoft.com`
[3]University of Michigan Business School, E-mail: `urajan@umich.edu`
[4]Tepper School of Business, Carnegie Mellon University, E-mail: `ravi@cmu.edu`

Keywords: Auctions, Mechanism Design, Profit Maximization, Bayes-Nash Equilibrium

1 Introduction

Suppose a seller is able to provide a service at total cost C to any number of users. Suppose further that they have done market research to determine the probability distribution from which their users valuations for receiving the service are drawn. What selling mechanism then should the seller run to to obtain the highest possible profit? In a seminal paper [13], Myerson essentially answers this question. If the seller aims to maximize their expected profit, they must first compute each user's *virtual valuation*, and then sell the item to all users with non-negative virtual valuation if the sum of their virtual valuations is above the cost C. It turns out that in many natural scenarios, this mechanism has a deficit on some possible instances of the users' values. A seller that is averse to such a loss might prefer a different mechanism.

In this paper we consider the general problem of Bayesian optimal mechanism design for arbitrary single-parameter agent problems (see e.g., [10, 2, 1]) when the seller requires the mechanism to never produce a deficit. In a single-parameter agent problem each agent has a publicly known partitioning of possible outcomes into two sets, the *reject* set and the *accept* set. It is assumed that agent i has valuation zero for any outcome in the reject set and private valuation v_i for any outcome in the accept set. For auction-like problems, agent i's accept set is simply the set of allocations where agent i is allocated their desired good and the reject set is the set of allocations where i is not allocated their desired good. The truthtelling strategy for agent i would be to report to the mechanism (a.k.a. bid) their true private value, v_i.

We follow a standard approach to profit maximization and assume that the agents' private valuations come from a known probability distribution. Our goal then is to design the seller optimal mechanism given knowledge of this distribution. We also make the assumption that this *Bayesian prior distribution* is independent but not necessarily identically distributed.

Motivating Problems

This paper considers a number of motivating problems, all of which fit in this single-parameter agent framework. Consider the following examples:

Fixed cost excludible good. In this case, the example from the first paragraph, the seller must pay a fixed cost C if any items are sold and zero otherwise. A motivating example of such a good is a digital good with production cost C. This is a special case of the general *multicast pricing* problem considered in [5, 6, 11].

Fixed cost non-excludible good. In this case there is a fixed cost, C, for providing the good or service to all users and no cost for serving nobody. However, the mechanism is not allowed to serve some users and not others (i.e., the cost for such allocations is infinite). We will sometimes refer to this as the *all-or-none* case. The classic example of a fixed cost non-excludible good is the bridge building problem where if the bridge is built than anyone can use it.

Submodular costs. In this case the additional cost in providing the good to any users is a decreasing function of the set of users already being provided. The excludible and non-excludible fixed cost problem and the multicast pricing problem are special cases of the general submodular cost problem.

Combinatorial auction (single-parameter). In this case each agent desires a subset of a set of items. The cost function is such that allocations to agents with disjoint subsets have cost zero and all other allocations have infinite cost. See e.g., [10, 1].

Supermodular costs. In this case the additional cost in providing the good to any users is an increasing function of the set of users already being provided. The single-parameter combinatorial auction problem is a special case of a supermodular cost function.

Mechanism Design Solution Concepts

The fundamental difference between mechanism design and algorithm design is that the inputs to a mechanism are the private values of selfish agents that will attempt to submit bids that result in outcomes that maximize their own *utility*. We adopt the following solution concepts for designing mechanisms for such selfish agents.

Ex post incentive compatibility. Otherwise known as *truthful* or *strategyproof* mechanisms, ex post incentive compatible mechanisms (via the revelation principle) are such that each agent, independent of the acts of any other agent, has a dominant strategy of stating their true valuation as their bid.

Bayesian incentive compatibility. Bayesian incentive compatible mechanisms are those where each agent has an optimal strategy of bidding their true valuation given that the other agents values come from a prior distribution and that all other agents bid their true values. Note that such a truthtelling strategy may not be optimal ex post, i.e., once the bids of other agents are known.

Overview of Results

The major focus of this paper, besides describing the Bayesian optimal mechanism, is computing it efficiently, i.e., in polynomial time. Myerson's optimal mechanism solves the single-parameter agent optimal mechanism design problem for any cost function given that the seller only wants to maximize their expected profit and spurious deficits are acceptable. For submodular costs, via a general algorithm due to Iwata et al. [9], it is possible to compute this optimal mechanism. However, for the single parameter combinatorial auction (and, thus general supermodular costs) this computational problem is NP-hard [10]. Of course the usual questions arise here as to whether it is possible to approximate the optimal mechanism via a polynomial time computation. For this problem, we show that Myerson's reduction from the efficient mechanism to the optimal mechanism via virtual valuations respects approximations. Given an incentive compatible mechanism that approximates efficiency, the Myerson approach can be used to obtain an incentive compatible mechanism that gives the same approximation factor against the optimal mechanism.

For the problem of designing the ex post incentive compatible optimal no-deficit mechanism we consider both the form in which the optimal mechanism takes as well as the problem of computing it. Like above, the answer to these questions depends on types of cost functions we are considering. We show that for supermodular costs functions, the Myerson mechanism is indeed no-deficit. Of course, by the above discussion such a mechanism is hard to compute. For the submodular case, and even the special case of a fixed cost excludible good, we show that Myerson is not no-deficit. We then consider the most natural way to try to obtain a no-deficit mechanism that achieves good expected profit: merging the Myerson mechanism which has optimal expected profit with a thresholding mechanism, e.g. Moulin and Shenker's [12] cost sharing mechanism, which has no-deficit. We show that even for the fixed cost excludible good problem when bidders are independent and identically distributed, this mechanism is not optimal. We further show the somewhat surprising result that even though in this case the problem is completely symmetrical, the optimal no-deficit mechanism is not. None-the-less, as these thresholding mechanisms are intuitively easy to understand, we ask two questions, first, when are thresholding mechanisms optimal, and second can we compute them. We show that these mechanisms are indeed optimal for all-or-nothing costs; yet computing the optimal thresholding mechanism on this special case is NP-hard.

We then consider relaxing the our solution concept from ex post incentive compatibility to Bayesian incentive compatibility. We show that while the ex post incentive compatibility payment rule of Myerson is not no-deficit on some realizations of the agents' valuations, there is a bayesian incentive compatible payment rule for Myerson's mechanism that obtains the same expected profit as the original Myerson payment rule and guarantees that there is never a deficit. We leave the problem of computing this payment rule as an open question.

Related Work

This work is based heavily on the result of Myerson [13] and generalizations observed by Bulow and Roberts [3] on optimal mechanism design. Cornelli re-derives these results for the special case of a fixed cost excludible good and considers the related problem of designing optimal non-direct revelation mechanisms (where the set of allowable bids is a subset of possible valuations of the bidders) [4]. Mehta and Vazirani consider the related computational question of how to compute the optimal "take it or leave it" offers for each agent (prior to seeing any agent bids), given a Bayesian prior distribution on agent valuations and a cost functions given by the multicast pricing special case of submodular costs [11].

Another branch of related work is that of worst-case profit maximizing mechanism design. For the trivial cost function, Goldberg et al. give an approximately optimal worst case auction [8, 7]. Fiat et al. consider the fixed cost excludible good problem and more general multicast pricing problem. They give approximately optimal mechanisms under certain assumptions [6].

Organization

The rest of the paper is organized as follows, in Section 2 we give formal definitions for the concepts considered above. In Section 3 we review the Myerson mechanism, the derivation of virtual valuations, and the proof that Myerson is optimal in the general context of any single parameter agent problem. In Section 4 we prove that Myerson is no-deficit for supermodular cost functions. Since it is NP-hard to compute, we also consider approximating the optimal Myerson mechanism for (single-parameter) combinatorial auction problem. For the submodular case, in Section 5 we consider the use of Myerson combined with thresholding mechanisms. We show that such a hybrid mechanism is optimal for all-or-nothing cost, but NP-hard to compute. For other submodular costs we show that this technique does not give the optimal mechanism. We consider the relaxed solution concept of bayesian incentive compatibility in Section 6 and show that the optimal no-deficit b.i.c. mechanism obtains the same expected profit as Myerson. We conclude in Section 7 with a discussion and open questions.

2 Notation and preliminaries

Let $\mathcal{B}$ be the set of bidders and $c(A)$ denote the cost of an allocation $A \subseteq \mathcal{B}$. We say that cost is submodular, if for all allocations A_1 and A_2, $c(A_1) + c(A_2) \geq c(A_1 \cup A_2) + c(A_1 \cap A_2)$. Likewise, cost is supermodular, if for all allocations A_1 and A_2, $c(A_1) + c(A_2) \leq c(A_1 \cup A_2) + c(A_1 \cap A_2)$.

Each agent has a true type t_i distributed according to the distribution F_i and corresponding density function f_i. Without loss of generality, we assume that t_i is in the range $[0, T]$ for all i. Given a bid b_i of agent i, we define the Virtual Valuation of this bidder to be $v_i(b_i) = b_i - \frac{1-f_i(b_i)}{F_i(b_i)}$. We restrict our attention to distributions F_i for which v_i is an increasing function of b_i. This is a standard assumption and is called *Monotone Hazard Rate*. This assumption essentially implies the uniqueness of the pricing strategy that maximizes the profit of the auctioneer. Under this assumption, there is a one-to-one correspondence between $v_i(b_i)$ and b_i. So we will use b and v interchangeably in the rest of the paper.

For a mechanism M, the allocation served at bid vector b is denoted by $M(b)$. When M is a randomized mechanism, $M(b)$ is a random variable. We use v to denote a vector in the virtual valuation space. For a virtual valuation vector v and allocation A, we define the *virtual surplus* of this allocation at vector v to be $\mathcal{S}_v(A) = \sum_{i \in A} v_i - c(A)$.

The price charged to bidder i by a mechanism M at bid vector b is denoted by $p_i^M(b)$. Each bidder aims to maximize her profit and therefore reports a bid $b_i = \text{argmax}\,[t_i - p_i^M(b_{-i}, b_i)]$. A mechanism is called truthful, if for all bidders i and all vectors b_{-i}, we have $b_i = t_i$. The following lemma is a well-known characterization of truthful mechanisms.

Lemma 1 *For any truthful (randomized or deterministic) mechanism M, let $q_i^M(b)$ be the probability that bidder i gets served when the bid vector b is reported. Then, $p_i^M(b) = b_i q_i^M(b) - \int_{t=0}^{t=b_i} q_i^M(t, b_{-i}) dt$.*

The above lemma suggests that for describing a mechanism, it is sufficient to specify the (distribution over) allocation(s) for every possible bid vector. If a mechanism is deterministic, the price $p_i^M(b)$ is the minimum bid that i must bid in order to be served. Whenever it is clear from the context, we drop the superscript M.

3 The Extended Meyerson mechanism is optimal

In this section, we charaterize the optimal mechanism for Bayesian auctions. When there are only two bidders and a single item with cost c, Meyerson [13] gives a simple characterization of the optimal mechanism. M^* serves the bidder with the higher virtual valuation, as long as this virtual valuation is larger than the cost c of the item. Below we review how this mechanism can be generalized to a larger number of bidders and general cost functions.

For a mechanism M and bid vector b, let $q_A(b)$ denote the probability of serving allocation A at bid vector b. Then, the probability of serving bidder i is given by $q_i(b) = \sum_{A: i \in A} q_A(b)$. Recall that by Lemma 1, the price charged to agent i is given by $p_i(b) = b_i q_i(b) - \int_{t=0}^{t=b_i} q_i(t, b_{-i}) dt$. The expected revenue of the auctioneer is given by $\sum_i \int_b p_i(b) f(b) db$. Here f is the joint density of the bid vector, $f = \prod_i f_i$. We can rewrite the revenue collected from bidder i as follows:

$$
\begin{aligned}
\int_b p_i(b)f(b)db &= \int_{b_{-i}} \left[\int_{b_i} b_i q_i(b) f_i(b_i) db_i - \int_{b_i} \left(\int_{t=0}^{t=b_i} q_i(t, b_{-i})dt\right) f_i(b_i)db_i\right] f(b_{-i})db_{-i} \\
&= \int_{b_{-i}} \left[\int_{b_i} b_i q_i(b) f_i(b_i) db_i - \int_{t} \left(\int_{b_i=t}^{b_i=\infty} f_i(b_i)db_i\right) q_i(t, b_{-i})dt\right] f(b_{-i})db_{-i} \\
&= \int_{b_{-i}} \left[\int_{b_i} b_i q_i(b) f_i(b_i) db_i - \int_{t} (1 - F_i(b_i))\, q_i(t, b_{-i})dt\right] f(b_{-i})db_{-i} \\
&= \int_{b_{-i}} \left[\int_{b_i} \left(b_i - \frac{1 - F_i(b_i)}{f_i(b_i)}\right) q_i(b) f_i(b_i) db_i\right] f(b_{-i})db_{-i} \\
&= \int_b v_i q_i(b) f(b) db
\end{aligned}
$$

The net profit of the auctioneer is given by

$$
\begin{aligned}
\int_b \left[\sum_i v_i q_i(b) - \sum_A c(A) q_A(b)\right] f(b)db &= \int_b \left[\sum_i \sum_{A:i\in A} v_i q_A(b) - \sum_A c(A)q_A(b)\right] f(b)db \\
&= \int_b \sum_A q_A(b) \left[\sum_{i\in A} v_i - c(A)\right] f(b)db \\
&= \int_b \sum_A q_A(b)\mathcal{S}_b(A)\, f(b)db
\end{aligned}
$$

This gives us the following lemma.

Lemma 2 *The expected profit of any truthful deterministic or randomized mechanism M is exactly equal to its expected virtual surplus —* $\int_b (\sum_i p_i(b) - c(M(b)))f(b)db = \int_b \mathcal{S}_b(M(b))f(b)db$.

Therefore, in order to maximize profit, the mechanism should deterministically select the allocation A that maximizes the virtual surplus $\mathcal{S}_v(A) = \sum_{i\in A} v_i - c(A)$. We call this mechanism the *Extended Myerson* mechanism[1].

Theorem 3 *The Extended Myerson mechanism given by $EM(b) = argmax_A \mathcal{S}_b(A)$ is the optimal truthful mechanism for the Bayesian auction problem.*

3.1 The discrete-valued case

Although all the definitions given above assume that the buyers' bids are continuous variables, it is easy to formulate similar expressions when bids are discrete-valued. We give analogues for the discrete case below, and leave the proofs to the reader.

For the ith bidder, let $x_{i,j}$ denote the jth value that b_i can take. Let the corresponding probability be given by $f_{i,j}$, and let $F_{i,j} = \sum_{k=0}^{k=j} f_{i,k}$ denote the cumulative probability. The jth virtual valuation of bidder i is given by $v_{i,j} = x_{i,j} - \frac{1-F(i,j)}{f_{i,j}}(x_{i,j+1} - x_{i,j})$.

[1] We assume for convenience, that the mechanism breaks ties in a consistent fashion. That is, we assume there is an implicit ordering over allocations, and whenever the virtual surplus maximizing allocation is not unique, the mechanism picks the allocation that precedes all other maximizers in the ordering. This is important for truthfulness.

The price $p_i(b)$ charged by a truthful mechanism M, when bidder i reports $b_i = x_{i,j}$, is given by $p_i(b) = b_i q_i(b) - \sum_{k=0}^{k=j-1} q_i(x_{i,k}, b_{-i})(x_{i,k+1} - x_{i,k})$. The virtual surplus of an allocation and the Extended Myerson mechanism are defined as before.

3.2 An example and the no-deficit constraint

Consider the following example. There are three bidders with independent but identically distributed bids — 9 with probability 0.6, 12 with probability 0.2 and 16 with probability 0.2. The virtual valuations of the bidders are given by $v(9) = 7$, $v(12) = 8$ and $v(16) = 16$. Let the cost of any allocation containing one bidder be 15, the cost of an allocation with two bidders be 24, and the cost of serving all of them be 30.4. This is a submodular cost function.

The allocations served by the Extended Myerson mechanism and the corresponding virtual surpluses for all possible bid vectors are given in Figure 1.

Note that when all the bidders bid 16, they all get served and the price charged to each is 9. The seller earns 27, and therefore makes a loss of 3.4. On the other hand, the expected profit of the allocation is 1.4016, which is larger than that of any other truthful mechanism.

$v_3 = 7$

	7	8	16
16	(2);1	(2);1	(123);8.6
8			(1);1
7			(1);1

$v_3 = 8$

	7	8	16
16	(2);1	(123);1.6	(123);9.6
8			(123);1.6
7			(1);1

$v_3 = 16$

	7	8	16
16	(123);8.6	(123);9.6	(123);17.6
8	(3);1	(123);1.6	(123);9.6
7	(3);1	(3);1	(123);8.6

Figure 1: An example in which EM does not satisfy no-deficit. Each table gives the allocation and virtual surplus of EM for a particular setting of v_3. The horizontal-axis represents v_1 and vertical-axis represents v_2.

Definition 1 The no-deficit constraint. *A mechanism M is said to satisfy the no-deficit constraint, if for all bid vectors b, the profit of the seller $\sum_i p_i^M(b) - c(M(b))$ is positive.*

4 Extended Myerson satisfies no-deficit for supermodular functions

In this section we prove that the Extended Myerson mechanism always satisfies the no-deficit constraint if the cost function is supermodular. We start with a few properties of the EM mechanism.

Lemma 4 *If any bidder served by EM unilaterally increases her bid, then the allocation of the mechanism stays the same. That is, given any two bid vectors b and b' with $b_j = b'_j$ for all $j \neq i$, and $b_i < b'_i$, if $i \in EM(b)$, then $EM(b') = EM(b)$.*

Proof: For an allocation A, let $\Delta(A) = \mathcal{S}_{b'}(A) - \mathcal{S}_b(A)$. Then, for any allocation A containing i, $\Delta(A) = b'_i - b_i > 0$, whereas, for any other allocation, $\Delta(A) = 0$. If $i \in EM(b)$, then $\Delta(EM(b)) \geq \Delta(A)$ for any allocation A. Also, we have $\mathcal{S}_b(EM(b)) \geq \mathcal{S}_b(A)$ for all A, by definition. Therefore, $\mathcal{S}_{b'}(EM(b)) \geq \mathcal{S}_{b'}(A)$ for all A, and $EM(b') = EM(b)$. ■

We get the following corollary:

Corollary 5 *The payment for bidder i given allocation A with $i \in A$ and other bids b_{-i} is the minimum bid b_i such that $EM(b) = A$.*

Proof: Note that the payment for bidder i is less than the minimum bid b_i with $M(b) = A$, because $i \in A$. Suppose the payment is less than b_i, and is given by b'_i. Then, $A_2 = M(b'_i, b_{-i})$ contains bidder i, but $A_1 \neq A_2$. This contradicts the lemma above. ■

Now we are ready to prove the main theorem:

Theorem 6 *Extended Myerson satisfies no-deficit for supermodular costs.*

Proof: Note that for all bid vectors b with $EM(b) = A$, and for all $i \in A$, we have $\mathcal{S}_b(A) \geq \mathcal{S}_b(A \backslash \{i\})$. So we have $b_i \geq v_i(b_i) \geq c(A) - c(A \setminus \{i\})$.

Now let $\min_i(A)$ be the minimum bid b_i of bidder i, with $i \in A$, such that for some bid vector b_{-i}, A is served, that is, $M(b_i, b_{-i}) = A$. Then Corollary 5 implies that the payment of bidder i at any vector b with $M(b) = A$ is given by $p_i(b) \geq \min_i(A)$, which is larger than $c(A) - c(A \setminus \{i\})$ by our observation above.

Now, taking a sum over all i, we get that the total payment collected is at least $\sum_i \min_i(A) \geq \sum_i [c(A) - c(A \setminus \{i\})]$. The net profit obtained is at least $\sum_i [c(A) - c(A \setminus \{i\})] - c(A)$. Note that submodularity implies $c(A) - c(A \setminus \{i\}) \geq c(B) - c(B \setminus \{i\})$, for any set $B \subset A$ with $i \in B$. We order the bidders in A in an arbitrary order – $n_1, \cdots, n_k$. Then, the net profit is given by $\sum_i [c(A) - c(A \setminus \{i\})] - c(A) \geq \sum_i [c(\cup_{j \leq i}\{n_j\}) - c(\cup_{j<i}\{n_j\})] - c(A) = c(A) - c(A) = 0$. ■

4.1 Computation of the optimal mechanism

Next we consider the problem of computing the Extended Meyerson mechanism for supermodular costs. In particular, we consider the problem of determining the winning allocation, given the bid vector, bid distributions and the cost function. Note that the input to this problem is in itself exponential in size. We therefore assume that the algorithm is provided with the following in order to find the optimal mechanism:

- A distribution oracle. Given i and b_i, it returns $F_i(b_i)$.

- A virtual value oracle. Given i and b_i, it returns v_i.
- A cost oracle. Given A, it returns $c(A)$.

In order to solve this problem, given a bid vector b, the mechanism designer must determine the set A that maximizes the surplus $\mathcal{S}_b(A) = \sum_{i \in A} v_i - c(A)$. Note that because costs are supermodular, the function $\mathcal{S}_b(A)$ is a submodular function. The NP-hardness of this problem now follows from the NP-hardness of maximization over submodular functions.

As an example, we reduce the problem of finding the maximum independent to this problem. In particular, let the bidders correspond to vertices. The cost of an allocation A is 0 if the corresponding set of vertices is an independent set (note that this can be computed in $O(n^2)$ time), and ∞ otherwise. It is easy to check that this cost is increasing and supermodular.

Bids of each bidder take on values 0 or 1. Each bid vector corresponds to a set of vertices containing exactly those vertices that have bid 1. The allocation that maximizes the virtual surplus corresponds to the maximum independent set in the graph. The reduction is approximation-preserving. Therefore, it is hard to approximate Problem 1 better than $\Omega(n^{1-\epsilon})$ unless P=NP.

Theorem 7 *It is NP-hard to approximate the expected profit of EM for a supermodular cost function better than an $\Omega(n^{1-\epsilon})$ factor.*

Given the above result, we consider a special case of this problem, namely the single-parameter combinatorial auction problem. We first note a relationship between profit maximization in a Bayesian setting and computing the social optimum in a worst-case setting. Note that for any bid vector b, the socially optimal solution is one that maximizes $\mathrm{SOPT}_b(A) = \sum_{i \in A} b_i - c(A)$. On the other hand, the profit maximizing solution aims to maximize $\mathcal{S}_b(A) = \sum_{i \in A} v_i(b_i) - c(A)$. If there exists an algorithm to compute an optimal solution to the first problem, then we can use the same algorithm to compute the optimal solution to the second problem, by simply giving the virtual valuation vector to the algorithm as input. A truthful mechanism in the first case would transform into a truthful mechanism in the second, because virtual valuations are increasing functions of bids.

Note that it is NP-hard to find the socially optimal solution for combinatorial auctions. However, for the special setting of single-parameter agents, there are randomized auctions [1] that approximate the social optimum well, and are truthful in expectation, or with high probability.

We first show that any truthful deterministic mechanism that approximates the social optimum gives an approximation to profit maximization in the Bayesian setting. Next we show how to modify the randomized auction of [1] to get a $(1+\epsilon)$-approximation to profit for the Bayesian setting via an auction that is truthful in expectation.

Theorem 8 *Given a polynomial-time truthful deterministic mechanism M that approximates the social optimum of a combinatorial auction in a worst-case setting, there exists a polynomial-time truthful mechanism M' that approximates the expected profit in a Bayesian setting.*

Proof: The mechanism M' can be described as follows. Given the bid vector b and costs c, we first determine the virtual valuations vector v. We then run the mechanism M on the virtual valuations v along with the cost vector c. Let A be the allocation determined by M and p be the vector of prices to be charged to bidders. M' returns the allocation A along with price vector p', where $p'_i = v_i^{-1}(p_i)$.[2]

Noting that $\mathrm{SOPT}_v(A) = \mathcal{S}_b(A)$, it is clear that the virtual surplus of the mechanism M' is the same as the profit of M, and therefore the two mechanisms obtain the same approximation. It only remains to show that M' is truthful.

Note that the truthfulness of M implies that M is monotone in all its input bids, viz., the vector v. Since the allocation of M' is the same as that of M, and v_i is increasing in b_i, M' is also monotone in the bids b. Furthermore, the price charged by M is the minimum v_i for which $i \in A$. Correspondingly, M' charges the bidder the minimum bid $(v_i^{-1}(p_i))$ for which the bidder $i \in A$. Therefore, M' is truthful. ■

Finally, we use the same technique as in the above proof to adapt the randomized mechanism of [1] to our setting. We obtain a $(1+\epsilon)$ approximation to profit. The resulting mechanism is monotone in the input bids, because the mechanism of [1] is monotone. However, because the mechanism is randomized, we cannot simply take the virtual inverse of the price charged by M to compute the price charged by M'. The price charged in this case is given by $p_i(b) = b_i q_i(b) - \int_{t=0}^{t=b_i} q_i(t, b_{-i})dt$, where $q_i(b)$ is the probability that bidder i gets served when the bid vector is b. Aarcher et. al. [1] compute this price by using a sampling technique. We use the same technique but replace the bid values used in their method, by the virtual inverse of the values. This gives us an estimate of the price $p_i(b)$.

Theorem 9 *There exists a polynomial time mechanism that is truthful in expectation and obtains a $(1+\epsilon)$-approximation to the single-parameter combinatorial auction in a Bayesian setting.*

5 Submodular costs, threshold mechanisms and all-or-none costs

In this section we consider submodular cost functions. As shown in Section 3.2, in this case, the Extended Meyerson mechanism does not always satisfy the no-deficit constraint. We now consider a simple family of mechanisms that always satisfies the no-deficit constraint. A threshold mechanism is characterized by values $\tau_i(A)$ for $i \in A$ signifying the minimum bid of bidder i such that for some bids of other bidders, allocation A is served. In particular, the mechanism characterized by these values is given by the following:

$$M(b) = \mathrm{argmax}_{A:\forall i \in A, b_i \geq \tau_i(A)} \mathcal{S}_b(A)$$

Note that if the threshold mechanism M satisfies the condition in Corollary 5, and the constraint that for all allocations A, $\sum_{i \in A} \tau_i(A) \geq c(A)$, then it satisfies the no-deficit constraint.

A natural question to ask is whether threshold mechanisms are optimal in the class of all truthful mechanisms satisfying no-deficit. Unfortunately, this is not the case, even when the cost function is symmetric and submodular, and all the bids are identically distributed, as we show through an example below.

Consider the example of section 3.2 again. Recall that in this example, EM makes a loss. The optimal no-deficit mechanism for this function is given in Figure 2, and has strictly larger profit than any threshold mechanism, or any symmetric mechanism.

[2]In the case of discrete-valued variables, we define $v_i^{-1}(p_i)$ to be the smallest bid b_i for which $v_i(b_i) \geq p_i$.

$v_3 = 7$

	7	8	16
16	(2);1	(2);1	(123);8.6
8			(1);1
7			(1);1

$v_3 = 8$

	7	8	16
16	(2);1	(123);1.6	(123);9.6
8			(123);1.6
7			(1);1

$v_3 = 16$

	7	8	16
16	(123);8.6	(123);9.6	(123);17.6
8	(3);1	(123);1.6	(13);9.6
7	(3);1	(3);1	(13);8.6

Figure 2: An example in which the optimal no-deficit mechanism is not a threshold mechanism. The expected revenue of the optimal no-deficit mechanism is 1.3744. The expected revenue of the optimal no-deficit threshold mechanism is 1.3728.

Although threshold mechanisms are not optimal for arbitrary submodular cost functions, we now show that they are indeed optimal for a special class of cost functions, that we call *all-or-none* costs. An all-or-none cost function is one in which the only allocations served are the empty allocation or the one containing all bidders. That is, for all allocations A with $A \neq \emptyset$ and $A \neq \mathcal{B}$, the cost $c(A) = \infty$.

Lemma 10 *Let M be any truthful mechanism for an all-or-none cost function c. Then, $p_i^M(b)$ is non-increasing in bids b_j with $j \neq i$.*

Proof: Suppose that there are bidders i and j such that $p_i^M(b)$ not non-increasing in b_j. Then, there are bid vectors b and b' with $b'_k = b_k$ for all $k \neq j$ and $b'_j > b_j$, such that $p_i(b') > p_i(b)$. Note that M is truthful, and so p_i does not depend on b_i. So we choose $b_i = b'_i = \frac{p_i(b')+p_i(b)}{2}$. Now, i is served at b but not at b'. However, since c is an all-or-none cost function, we have $M(b) = \mathcal{B}$ and $M(b') = \emptyset$. This means that the profit of bidder j is zero at b'. Then it is immediate, that when the bid vector is b', it is beneficial for bidder j to lie and report b_j instead of b'_j, and thereby obtain strictly positive profit. This contradicts the truthfulness of M. ∎

Theorem 11 *There exists a threshold mechanism that is optimal for any all-or-none cost function among the class of all truthful mechanisms satisfying no-deficit.*

Proof: Let M be any optimal truthful mechanism satisfying no-deficit. We will define a threshold mechanism M' with profit at least as large as the profit of M, thereby proving the theorem.

Let β be the bid vector with $b_i = T$, the highest bid, for every i. Let $\tau_i(\mathcal{B}) = p_i^M(\beta)$ for all i. Then, $\sum_i \tau_i(\mathcal{B}) = \sum_i p_i^M(\beta) \geq c(\mathcal{B})$, because M satisfies no-deficit. Consider the threshold mechanism M' given by thresholds τ_i.

For any bid vector b with $M(b) = \mathcal{B}$, we must have $\mathcal{S}_b(\mathcal{B}) > 0$. Otherwise, we can increase the profit of M by serving $\mathcal{B}$ only when $\mathcal{S}_b(\mathcal{B}) > 0$ and serving the empty set otherwise. Note also, that for all b with $M(b) = \mathcal{B}$ and all i, we have $b_i \geq p_i^M(b) \geq p_i^M(\beta) = \tau_i(\mathcal{B})$. Here the second inequality follows from Lemma 10. These two conditions along with the definition of M' imply that $M'(b) = \mathcal{B}$.

This means that for all b with $M(b) = \mathcal{B}$, we have $M'(b) = \mathcal{B}$. We get $\mathcal{S}_b(M(b)) \leq \mathcal{S}_b(M'(b))$, for all vectors b. Lemma 2 now implies that M' has a larger expected profit than M. ■

5.1 Computational hardness of the optimal threshold mechanism

Although threshold mechanisms are not always optimal, their simplicity is appealing and may make them practically useful. In this section we investigate the complexity of computing the optimal threshold mechanism. In particular, given bid distributions and a cost function, we consider the decision problem of determining whether there is a threshold mechanism with total expected profit greater than some given value P.

We show that even for a very simple input, in which every bidder has only two possible bids, and the cost function is an all-or-none function, it is NP-hard to compute the optimal threshold mechanism (which is also the optimal mechanism satisfying no-deficit in this case). As before, we assume that the mechanism has oracle access to bid distributions, virtual valuations and the cost function. We reduce the knapsack problem to the problem of computing the optimal mechanism.

Theorem 12 *The problem of computing the optimal no-deficit mechanism is NP-hard.*

Proof of Theorem 12: We reduce the knapsack problem to the problem of computing the optimal mechanism. The input to the knapsack problem is a set of items $1, \cdots, n$, with sizes $s_1, \cdots, s_n$, and values $u_1, \cdots, v_n$, and two positive real numbers S and U. The goal is to determine whether there is a subset A of items with small total size $\sum_{i \in A} s_i \leq S$ and large total value $\sum_{i \in A} u_i \geq U$.

The reduction is as follows. We assume that $S < \sum_i s_i$, otherwise the knapsack problem is trivial to solve. We rescale the sizes s_i and S such that for all i, $S - s_i > u_i$, and for all i and j, $S > u_i + u_j$, and $s_i s_j > u_i u_j$.

We construct a bidder for every item. The ith bidder has the following bid distribution — $l_i = S - u_i - s_i + \alpha$ with probability $f_i = \frac{s_i + u_i}{S + u_i}$ and $h_i = S - u_i + \alpha$ with probability $1 - f_i = \frac{S - s_i}{S + u_i}$. The virtual valuations of bidder i are given by $v_i(h_i) = h_i$ and $v_i = v_i(l_i) = \alpha + \frac{u_i}{s_i + u_i}(S - u_i - 2s_i) > \alpha - \frac{s_i u_i}{s_i + u_i}$. We pick $\alpha > \frac{s_i u_i}{s_i + u_i}$ for all i so that $v_i(l_i) > 0$. Let the cost of serving all users be $C = \sum_i h_i - S$, and for serving any other non-empty subset be ∞.

Now, it is easy to check that for all i, $v_i + \sum_{j \neq i} h_j > C$, but for all i and j, $v_i + v_j + \sum_{k \neq i,j} h_k < C$. Therefore, any mechanism should serve all bidders at bid vectors in which at most one bidder bids a low value and all the others bid high values. Let b^i denote the bid vector $\{l_i, h_{-i}\}$. The optimal mechanism serves all bidders at all these bid vectors b^i, along with $\{h_i\}$. Note however, the payments made by bidders when everyone bids h_i are l_i, and $\sum_i l_i = \sum_i h_i - \sum_i s_i < \sum_i h_i - S = C$. Therefore, the optimal mechanism does not satisfy no-deficit.

The optimal no-deficit mechanism serves some subset of the vectors b^i. Let this subset be A. Then, in order to satisfy no deficit, it must be the case that $\sum_{i \in A} l_i + \sum_{i \notin A} h_i \geq C$. This is satisfied if and only if $\sum_{i \in A} s_i \leq S$. The virtual surplus at vector b^i is given by $\Pi_j(1 - f_j)\frac{f_i}{1 - f_i}(v_i + \sum_{j \neq i} h_j - C) =$

$\Pi_j(1-f_j)u_i$. Define t such that the virtual surplus at bid vector $\{h_i\}$ is $\Pi_j(1-f_j)t$. Then, the total expected profit of the auctioneer is $[\Pi_j(1-f_j)](\sum_{i\in A} u_i + t)$. This profit exceeds $[\Pi_j(1-f_j)](U+t)$ if and only if $\sum_{i\in A} u_i \geq U$.

Therefore, the reduction is complete and a solution to the problem of computing the optimal no-deficit mechanism gives a solution to the knapsack problem. ■

6 Optimality via Bayesian Incentive Compatibility

In this section, we consider weakening the solution concept to Bayesian Incentive Compatibility. A mechanism M is said to be Bayesian incentive compatible (b.i.c.), if for all bidders i, and for all types t_i, the *expected* profit of the bidder is maximized when the bidder reveals her true value. That is, $\mathrm{argmax}_{b_i} E_{b_j, j\neq i}[q_i^M(b)(t_i - p_i^M(b))] = t_i$. As before, $q_i^M(b)$ denotes the probability that bidder i gets served by M when the bid vector is b.

The no-deficit constraint in the Bayesian case is the same as before – for any given set of bids, the net profit of the auctioneer should be non-negative.

Given that the optimal incentive-compatible mechanism does not always satisfy the no-deficit constraint, a natural question to ask is whether there exists an optimal b.i.c. mechanism that satisfies the constraint. As we show below, this is always the case, regardless of the cost function. Although the proof is constructive, unfortunately, it does not give us a means of computing the optimal no-deficit b.i.c. mechanism in polynomial time.

Theorem 13 *For all submodular cost functions, there exists an optimal b.i.c. mechanism that satisfies the no-deficit constraint.*

Note that for any incentive compatible mechanism M, we can define another mechanism M^i, in which the prices charged to bidder i change, however, in expectation, $E_{b_j, j\neq i}[q_i^M(b)(t_i - p_i^M(b))]$ remains the same for all types t_i and bids b_i. This can be done, for instance, by decreasing the price charged to bidder i for some setting of the bids of other bidders, b_j for $j \neq i$, and matching it with a corresponding increase in $p_i^M(b)$ for some other setting of the bids b_j, $j \neq i$.

If the mechanism M is b.i.c., then so is M^i, because the expected profit of any bidder in the two mechanisms is exactly the same. The expected profit of the auctioneer, the expectation being taken over bids of all bidders, also stays the same. However, the profit for any particular setting of bids changes. Note that the outcome (allocation) of the mechanism stays the same. In particular, it is the same as the allocation of EM. The only quantity that changes is the price charged to bidders that are served.

We use this idea to rebalance the profit of the auctioneer, such that the expected profit stays the same, whereas now, for all settings of the bids, the actual profit is non-negative.

We allow the new mechanism M' to charge prices $p_i^M(b)$ that are greater than b_i or smaller than zero. That is, the mechanism is not necessarily ex-post IR, but it is ex-ante IR. Likewise, the auctioneer may sometimes make a payment to one of the bidders in order for the mechanism to be b.i.c., however, she always makes non-negative profit by charging the other bidders a high price.

Let EM be the optimal truthful mechanism, as before. Let $\rho(b) = \sum_{i \in M(b)} p_i^M(b) - c(M(b))$ be the profit of the auctioneer at bid vector b, and $\mathcal{R}(b) = \rho(b)f(b)$ be its contribution to the total expected profit. First we describe a procedure that transfers an ϵ amount of expected profit from one bid vector to another, while still maintaining bayesian incentive compatibility.

Procedure *Transfer*:

Input: A b.i.c. mechanism M; two bid vectors b^1 and b^2; a bidder i with $b_i^1 = b_i^2$ and $i \in M(b^1) \cap M(b^2)$; a small positive number $\epsilon > 0$.

Output: A b.i.c. mechanism M' with the same allocation rule as M, and with $\mathcal{R}^{M'}(b^1) = \mathcal{R}^M(b^1) + \epsilon$ and $\mathcal{R}^{M'}(b^2) = \mathcal{R}^M(b^2) - \epsilon$.

1. $p_i^{M'}(b^1) = p_i^M(b^1) + \frac{\epsilon}{f(b^1)}$; $p_i^{M'}(b^2) = p_i^M(b^2) - \frac{\epsilon}{f(b^2)}$.
2. All other prices in M' are the same as that in M.

It is immediate that $\mathcal{R}(b^1)$ increases by an ϵ amount and $\mathcal{R}(b^2)$ decreases by the same amount. Note that because we only change the price of the ith bidder, and $b_i^1 = b_i^2$, the expected profit of bidder i remains the same. Therefore, the mechanism M' output by the above procedure is also b.i.c.

Now, we apply the above procedure iteratively, until the mechanism satisfies no-deficit. The algorithm is given below.

The rebalancing algorithm:

Repeat the following steps until there are no bid vectors with $\mathcal{R}(b) < 0$.

1. Let b^1 be a bid vector for which $\mathcal{R}(b^1) < 0$, and b^2 be another bid vector for which $\mathcal{R}(b^2) > 0$, with $M(b^1) = M(b^2)$.
2. Let $\epsilon = \min\{-\mathcal{R}(b^1), \mathcal{R}(b^2)\}$.
3. If there exists a bidder $i \in M(b^1)$ with $b_i^1 = b_i^2$, run the procedure *Transfer* with parameters b^1, b^2, i, and ϵ.
4. Otherwise, we define a new bid vector b^3 and transfer profit from b^2 to b^1 via b^3, as follows.
   (a) Suppose there exist bidders i and j in $M(b^1)$, with $b_i^1 < b_i^2$ and $b_j^1 > b_j^2$. Let b^3 be defined as $b_k^3 = \max\{b_k^1, b_k^2\}$ for all bidders k.
   (b) Suppose $b_k^1 < b_k^2$ for all bidders k. Define b^3 as $b_j^3 = b_j^1$ and $b_k^3 = b_k^2$ for all $k \neq j$.
   (c) Finally, suppose $b_k^1 > b_k^2$ for all bidders k. Define b^3 as $b_i^3 = b_i^2$ and $b_k^3 = b_k^1$ for all $k \neq i$.
5. Run the procedure Transfer with parameters b^1, b^3, j, and ϵ, and then again with parameters b^3, b^2, i, and ϵ.

Next we argue the correctness of the rebalancing algorithm.

6.1 Correctness of the rebalancing algorithm

In order to verify the correctness of the rebalancing algorithm, we begin by showing that if M does not satisfy the no-deficit condition, then we can always find bid vectors b^1 and b^2 satisfying the conditions in step 1. If M does not satisfy no-deficit, then, by definition, there must exist a bid vector b^1 with $\mathcal{R}(b^1) < 0$. Furthermore, for every allocation A, the expected profit of M over bids b for which $M(b) = A$ is positive. In particular, it is given by $\sum_{b:M(b)=A}\{\sum_{i\in A} v_i(b) - c(A)\}$, which is strictly positive by definition. Let $A = M(b^1)$. Then, there exists a bid vector b^2 with $M(b^2) = A$ and $\mathcal{R}(b^2) > 0$. The pair b^1 and b^2 satisfy the conditions of step 1.

If there exists a bidder $i \in M(b^1)$ with $b^1_i = b^2_i$, then the pair b^1 and b^2 satisfies all the conditions required for the transfer procedure, and we are done.

Otherwise, note that the bidder b^3 defined in step 4 satisfies $EM(b^1) \subseteq EM(b^3)$, or $EM(b^2) \subseteq EM(b^3)$. This follows from Lemma 4, because we either have $b^3_k \geq b^1_k$ for all $k \in EM(b^1)$, or $b^3_k \geq b^2_k$ for all $k \in EM(b^2)$. Again in this case, the pairs b^1 and b^3, and the pairs b^2 and b^3 satisfy the conditions of the transfer procedure.

Now note that the profit of b^3 stays the same as before, because in one transfer step, it decreases by an ϵ amount and in the second, it increases by the same amount ϵ.

It only remains to show that the algorithm converges to a mechanism that satisfies no-deficit. In order to do this, we consider a non-negative potential function that is non-zero only when the mechanism does not satisfy no-deficit, and show that at every step of the iteration, the potential function decreases. The potential function is given by $\vartheta^M = \sum_b \vartheta^M(b) = -\sum_{b:\mathcal{R}(b)<0} \mathcal{R}(b)$.

We only change the prices at bid vectors b^1, b^2 and b^3, and therefore, it is sufficient to consider the value of the potential at these bid vectors. As we argued before, $\mathcal{R}(b_3)$ stays the same before and after the iteration. The contribution of b^2 to the potential function is zero before and after the iteration, because $\mathcal{R}(b^2) \geq 0$ at both times. Finally, the contribution of b^1 is $-\mathcal{R}(b^1)$ before the iteration but decreases by an ϵ amount during the iteration. Therefore, there is a net decrease in the potential after each iteration, and the algorithm converges to a no-deficit mechanism.

Furthermore, note that at every step, the profit of at least one bid vector goes from a non-zero value to zero. Since there are only a finite number of bid vectors, the process terminates in a finite number of steps.

This concludes the proof of Theorem 13.

7 Conclusions

Acknowledgements

We are grateful to Amitabh Sinha for several helpful discussions.

References

[1] A. Archer, C. Papadimitriou, K. Tawar, and E. Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *Proc. 14th Symp. on Discrete*

Alg. ACM/SIAM, 2003.

[2] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

[3] J. Bulow and J. Roberts. The Simple Economics of Optimal Auctions. *The Journal of Political Economy*, 97:1060–90, 1989.

[4] F. Cornelli. Optimal Selling Procedures with Fixed Costs. *Journal of Economic Theory*, 71:1–30, 1996.

[5] J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the Cost of Multicast Transmissions. In *Proc. of 32nd Symposium Theory of Computing*, pages 218–226. ACM Press, New York, 2000.

[6] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing.* ACM Press, New York, 2002.

[7] A. V. Goldberg and J. D. Hartline. Competitiveness via Concensus. In *Proc. 14th Symp. on Discrete Alg.* ACM/SIAM, 2003.

[8] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Alg.*, pages 735–744. ACM/SIAM, 2001.

[9] S. Iwata, S. Fujishige, and L. Fleischer. A Combinatorial, Strongly Polynomial-Time Algorithm for Minimizing Submodular Functions. *Journal of the ACM*, 48(4):761–777, 2001.

[10] D. Lehmann, L. I. O'Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.

[11] A. Mehta and V. Vazirani. Profit-Maximizing Multicast Pricing Via Approximate Fixed Points. In *Proc. of 4th ACM Conference on Electronic Commerce.* ACM Press, New York, 2003.

[12] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511–533, 2001.

[13] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.