Bayesian Optimal No-Deficit Mechanism Design*

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

[1] Microsoft Research, Silicon Valley Campus, Mountain View, CA
shuchi@cs.wisc.edu
[2] Microsoft Research, Silicon Valley Campus, Mountain View, CA
hartline@microsoft.com
[3] Ross School of Business, University of Michigan, Ann Arbor, MI
urajan@bus.umich.edu
[4] Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA
ravi@cmu.edu

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 distributions on the valuations of the bidders are known and independent, Myerson [15] reduces the problem to that of maximizing the common welfare by considering the *virtual valuations* in place of the bidders' actual valuations. The Myerson auction maximizes the seller's profit over the class of all mechanisms that are truthful and individually rational for all the bidders; however, the mechanism does not satisfy ex post individual rationality for the seller. In other words, there are examples in which for certain sets of bidder valuations, the mechanism incurs a loss.

We consider the problem of merging the worst case *no-deficit* (or ex post seller individual rationality) 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 the computational side of the problem, we note that for supermodular costs the Myerson mechanism is NP-hard to compute. Furthermore, we show that for all-or-nothing costs the optimal thresholding mechanism is NP-hard to compute. Finally, we consider relaxing the ex post incentive compatibility constraint and show that there is a Bayesian incentive compatible mechanism that achieves the same expected profit as Myerson, but never incurs a loss.

1 Introduction

Suppose a seller is able to provide a service at total cost C to any number of users. Suppose further that the seller has done market research to determine the

* This work was done at Carnegie Mellon University and sponsored in part by the ALADDIN project.

P. Spirakis et al. (Eds.): WINE 2006, LNCS 4286, pp. 136–148, 2006.

probability distribution from which the user valuations for receiving the service are drawn. What selling mechanism should the seller then run to obtain the highest possible profit? In a seminal paper [15], 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. The Myerson mechanism is optimal for expected revenue, in the class of all mechanisms that induce the users to participate and reveal their preferences truthfully. However, 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.

We consider the general problem of Bayesian optimal mechanism design for arbitrary single-parameter agent problems (see e.g., [12,2,1]) when the seller requires the mechanism to never produce a deficit. Here, the seller must pay a cost that is a function of the outcome that the seller chooses. A deficit would arise if the total payments of the agents does not cover the cost of the outcome produced. 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 (or service) 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 their true private valuation, v_i.

We follow the standard economics 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 assume that the agents valuations are independent but not necessarily identically distributed.[1]

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. 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 and zero marginal cost. This is a special case of the general *multicast pricing* problem considered in [5,6,13].

Fixed cost non-excludible good. 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 then anyone can use it.

[1] It is NP-hard to compute the optimal auction when valuations are correlated [16].

Submodular costs. The additional (marginal) 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. Goods with concave production costs or increasing returns to scale fall in this category.

Combinatorial auction (single-parameter). 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., [12,1].

Supermodular costs. 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.

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 no-deficit mechanism, is to study the complexity of computing it. 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. [10], 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 [12]. 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, it is relatively easy to see 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 that 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 (Section 4). 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 [14] 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, such a mechanism is not optimal (Section 5). We further show the somewhat surprising result that even though in this case the problem is completely symmetrical, the optimal deterministic 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 our solution concept from ex post incentive compatibility to Bayesian incentive compatibility. We show that while the ex post incentive compatible 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 results of Myerson [15] on optimal mechanism design and generalizations observed by Bulow and Roberts [3]. 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 [13] consider the related computational question of how to compute the optimal "take it or leave it" offers for each agent prior to seeing any bids, for the aforementioned multicast pricing special case of submodular costs.

Another branch of related work is that of worst-case profit maximizing mechanism design. There is much work in this area. (See, for example, [9].) As an example, 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].

2 Notation and Preliminaries

Let $S = \{1, \ldots, n\}$ denote a set of n agents. We represent the outcome of the mechanism as an *allocation* $A \subset S$ of *accepted* agents. We assume there is a general cost function $c(A)$ over allocations. As noted in the introduction, this allows us to represent any (binary) single parameter agent problem.

A cost function is said to be *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, it is said to be *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 i has a valuation v_i for being accepted. We assume that v_i is drawn independently from distribution F_i and corresponding density function f_i. The joint distribution, $\mathbf{F}$, is the product $F_1 \times \cdots \times F_n$. Without loss of generality, we assume that v_i is in the range $[0, h]$ for all i. We define the *virtual valuation* of agent i to be $\phi_i(v_i) = v_i - \frac{1-F_i(v_i)}{f_i(v_i)}$. Where v_i is implicit, we will refer to ϕ_i as agent i's virtual valuation. We restrict our attention to distributions F_i for which ϕ_i is an increasing function of v_i. This is implied by the *monotone hazard rate* assumption which is standard in mechanism design.

Assume for mechanism $\mathcal{M}$ that the agents submit bids $\mathbf{b} = (b_1, \ldots, b_n)$. We denote the allocation served by $\mathcal{M}(\mathbf{b})$. When $\mathcal{M}$ is a randomized mechanism, $\mathcal{M}(\mathbf{b})$ is a random variable. For valuations $\mathbf{v}$ and allocation A, we define the *surplus* of this allocation to be $\mathcal{S}_{\mathbf{v}}(A) = \sum_{i \in A} v_i - c(A)$. The *virtual surplus* we denote by $\hat{\mathcal{S}}_{\mathbf{v}}(A) = \mathcal{S}_{\phi(\mathbf{v})}(A) = \sum_{i \in A} \phi_i(v_i) - c(A)$. For ex post IC mechanisms, we have $\mathbf{b} = \mathbf{v}$, so we sometimes use $\hat{\mathcal{S}}_{\mathbf{b}}$ to denote the virtual surplus.

Let $p_i(b_i)$ denote the payment charged by mechanism $\mathcal{M}$ to agent i when he bids b_i. Define $q_i(b_i)$ as the probability that agent i is allocated when bidding b_i. Notice that this payment and probability are dependent on the randomization in the other bids, $\mathbf{b}_{-i}$, and the randomization in the mechanism, $\mathcal{M}$. A mechanism is incentive compatible if this agent's utility is maximized when bidding their true valuation. I.e., $v_i \in \operatorname{argmax}_b[v_i q_i(\mathbf{b}_{-i}, b) - p_i(\mathbf{b}_{-i}, b)]$. A mechanism is (ex post) *incentive compatible* (IC) if this holds for all values of the other agents bids, $\mathbf{b}_{-i}$, and *Bayesian incentive compatible* (BIC) if it holds when the other agents bid their true values, so that $\mathbf{b}_{-i}$ is drawn from the prior distribution $\mathbf{F}_{-i} = F_1 \times \cdots \times F_{i-1} \times F_{i+1} \times \cdots \times F_n$. It is well known [11] that the allocation rule and the expected payment of each agent satisfies the following conditions[2].

Lemma 1. *For any ex post incentive compatible mechanism $\mathcal{M}$, for $\mathbf{b}_{-i}$ fixed, $q_i(b_i)$ is non-decreasing in b_i, and $p_i(b_i) = b_i q_i(\mathbf{b}) - \int_{b=0}^{b=b_i} q_i(\mathbf{b}_{-i}, b)db$.*

Lemma 2. *For any Bayesian incentive compatible mechanism $\mathcal{M}$, when $\mathbf{b}_{-i}$ are drawn from $\mathbf{F}_{-i}$ then $q_i(b_i) = \mathbf{E}_{\mathbf{b}_{-i}} q_i(\mathbf{b})$ is non-decreasing in b_i, and $p_i(b_i) = b_i q_i(b_i) - \int_{b=0}^{b=b_i} q_i(b)db$.*

The above lemmas imply that for describing an incentive compatible mechanism, it is sufficient to specify an allocation rule that is monotone in the bids of each

[2] In general, the expressions for p_i may contain a constant $p_i(0)$ term, but because we are interested in profit maximization, we assume that this term is zero.

agent. Notice that for a deterministic ex post IC mechanism, the price $p_i(b_i)$ is the minimum bid that i must bid in order to be served.

In addition to incentive compatibility, we also require our mechanisms to satisfy the *no-deficit* condition defined below.

No-Deficit Condition. A mechanism $\mathcal{M}$ is said to satisfy the no-deficit condition if and only if for all bid vectors $\mathbf{b}$, the profit of the mechanism is non-negative: $\sum_i p_i(\mathbf{b}) - c(\mathcal{M}(\mathbf{b})) \geq 0$.

3 The Myerson Mechanism

Notice that the Vickrey-Clarke-Groves (VCG) mechanism applied to our single parameter setting is the mechanism that chooses the allocation that maximizes the surplus (defined above). It is easy to see that this allocation rule is monotone and thus there exist prices that incentivize agents to bid their true values. Myerson reduced the problem of Bayesian profit maximization to that of maximizing surplus via the concept of virtual valuations. He shows that the Bayesian optimal mechanism is the one that maximizes the virtual surplus. His theorem generalizes directly to our single parameter agent setting as follows.

Lemma 3. *[15] The expected profit of any truthful mechanism is exactly equal to its expected virtual surplus.*

Theorem 1. *[15] Given agents with valuations drawn from distribution* $\mathbf{F} = F_1 \times \cdots \times F_n$ *with each* F_i *satisfying the monotone hazard rate condition, the ex post IC mechanism with the maximum expected profit selects the outcome to maximize the virtual surplus, i.e.,* $\mathcal{M}(\mathbf{b}) = \operatorname{argmax}_A \hat{\mathcal{S}}_{\mathbf{b}}(A)$. *The expected profit of this mechanism is given by* $\mathbf{E}_{\mathbf{b}} \hat{\mathcal{S}}_{\mathbf{b}}(\mathcal{M}(\mathbf{b}))$.

One view of this theorem is that to maximize profit, first compute virtual valuations assuming that the agents bid their valuations, and then run the VCG mechanism on these virtual valuations. Payments can be determined by the payments of VCG in this setting by applying each agent's inverse virtual valuation function to their VCG payment. We refer to the mechanism that maximizes the virtual surplus as the *Myerson* mechanism.

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 analogs for the discrete case below. These descriptions are standard and we leave the proofs to the reader.

For the ith bidder, let $x_{i,j}$ denote the jth value that v_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 $\phi_{i,j} = x_{i,j} - \frac{1-F_{i,j}}{f_{i,j}}(x_{i,j+1} - x_{i,j})$.

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

3.2 Some Examples and the No-Deficit Constraint

Consider the following example. There are three bidders, each having a value drawn uniformly from the interval $[0, 1]$. The cost of serving any non-empty subset of them is $C = 2$. Then, the virtual valuation function of the ith bidder is given by $\phi_i(v_i) = 2v_i - 1$. Consider the case when all the three valuations are 1. Then the total virtual valuation is 3. Therefore, the Myerson mechanism serves all three of the bidders. The payment of bidder i is given by the minimum virtual valuation at which bidder i gets served. This is $C - \sum_{j \neq i} \phi_j(v_j) = 2 - 2 = 0$. Therefore, the payment of bidder i is ${\phi_i}^{-1}(0) = 0.5$. The revenue of the mechanism at the bid vector $(1, 1, 1)$ is therefore 1.5, whereas the cost of serving the three bidders is 2. The mechanism incurs a loss.

A slightly different example shows that the ratio between the worst-case loss of the Myerson mechanism and its expected profit can be unbounded. Consider an example with n identically distributed bidders, each with bid distribution uniform over $[1, 2]$. The cost of serving any subset of the bidders is $C = 2n - 2$. The reader is encouraged to verify that the worst-case loss of the Extended Myerson mechanism in this case is $n - 2$ (for the bid vector $(2, \cdots, 2)$), whereas the expected profit of the mechanism is less than 2.

4 The No-Deficit Constraint for Supermodular Functions

In this section we prove that the Myerson mechanism always satisfies the no-deficit constraint if the cost function is supermodular. We start with a few properties of the Myerson mechanism that will be useful in our analysis.

4.1 Strong Monotonicity of Allocations in the Myerson Mechanism

We show that if any bidder served by Myerson unilaterally increases her bid, then the allocation of the mechanism stays the same. Note that if a bidder being served by the mechanism raises her bid, truthfulness (and thus, monotonicity) implies that the bidder continues being served. The next lemma however says something stronger—when the bidder raises her bid, no other bidder gets added or removed from the set being served.

Lemma 4. *Given any two bid vectors* $\mathbf{b}$ *and* $\mathbf{b}'$ *with* $b_j = {b_j}'$ *for all* $j \neq i$*, and* $b_i < {b_i}'$*, if* $i \in \text{Myerson}(\mathbf{b})$*, then* $\text{Myerson}(\mathbf{b}') = \text{Myerson}(\mathbf{b})$.

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

Corollary 1. *The payment for bidder i given allocation A with $i \in A$ and other bids $\mathbf{b}_{-i}$ is the minimum bid b_i such that* Myerson($\mathbf{b}$) $= A$.

Proof. Note that the payment for bidder i is less than the minimum bid b_i with Myerson(b) $= A$, because $i \in A$. Suppose the payment is $p_i(b_i) = b_i{}'$ with $b_i{}' < b_i$. Then the allocation $A_2 = \text{Myerson}(b_i{}', \mathbf{b}_{-i})$ contains i. This contradicts the lemma above, because when i increases her bid from $b_i{}'$ to b_i, the allocation changes from A_2 to $A \neq A_2$. □

4.2 Myerson Satisfies No-Deficit

Now we are ready to prove the main theorem of this section:

Theorem 2. *Myerson satisfies no-deficit for supermodular costs.*

Proof. Note that for all bid vectors $\mathbf{b}$ with Myerson($\mathbf{b}$) $= A$, and for all $i \in A$, we have $\hat{\mathcal{S}}_{\mathbf{b}}(A) \geq \hat{\mathcal{S}}_{\mathbf{b}}(A \setminus \{i\})$. Then by the definition of $\hat{\mathcal{S}}_{\mathbf{b}}(A)$ and using the monotone hazard rate condition, we have $b_i \geq \phi_i(b_i) \geq c(A) - c(A \setminus \{i\})$.

Now let $\min_i(A)$ be the minimum bid b_i of bidder i, with $i \in A$, such that for some bid vector $\mathbf{b}_{-i}$, A is served, that is,

$$\min_i(A) = \min\{b_i : \exists \mathbf{b}_{-i} \text{ with } \text{Myerson}(b_i, \mathbf{b}_{-i}) = A\}.$$

Then Corollary 1 implies that the payment of bidder i at any vector $\mathbf{b}$ with Myerson($\mathbf{b}$) $= A$ is given by $p_i(\mathbf{b}) \geq \min_i(A)$, 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 supermodularity implies $c(A) - c(A \setminus \{i\}) \geq c(B) - c(B \setminus \{i\})$, for any set $B \subset A$ with $i \in B$. Without loss of generality, let $|A| = k$ and $A = \{1, \ldots, k\}$. Then, the net profit is at least

$$\begin{aligned}\sum\nolimits_i [c(A) - c(A \setminus \{i\})] - c(A) &\geq \sum\nolimits_i [c(\{1, \ldots, i\}) - c(\{1, \ldots, i-1\})] - c(A) \\ &= c(A) - c(\emptyset) - c(A) = 0.\end{aligned}$$

□

4.3 Computation of the Optimal Mechanism

Next we consider the problem of computing the Myerson 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. Supermodularity of the cost function implies that in general the optimal allocation is NP-hard to approximate better than an $\Omega(n^{1-\epsilon})$ factor. However, in special cases, given an approximate truthful mechanism for welfare mechanism for the same cost function, we can design an approximate truthful mechanism for profit maximization in the Bayesian setting. We obtain the following results.

Theorem 3. *Given a polynomial-time truthful deterministic mechanism $\mathcal{M}$ that α-approximates the social optimum in a worst-case setting, there exists a polynomial-time truthful mechanism $\mathcal{M}'$ that α-approximates the expected profit in a Bayesian setting.*

Theorem 4. *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.*[3]

We note that Theorem 4 is a non-trivial extension of Theorem 3 to auctions satisfying truthfulness in expectation, as the inverse virtual valuation function is not generally linear and thus it affects the expected utility of the agents. See the full paper for details.

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 Myerson mechanism does not always satisfy the no-deficit constraint. Intuitively, when the Myerson mechanism serves a large set A of bidders, the marginal cost of serving a bidder $i \in A$, and therefore the price charged to i, is very small. A simple way of dealing with these low costs is to supplement the Myerson mechanism with reserve prices or *thresholds* for each bidder, below which the bidder is not served. Precisely, let τ denote a budget-balanced cost-sharing method, and $\tau_i(A)$ denote the cost-share assigned to bidder i in coalition A. Then, if a mechanism serves the set A only if the bids of all bidders in A are above their respective thresholds, then the mechanism obtains prices at least $\tau_i(A)$ from each bidder $i \in A$, and therefore, meets the cost of serving the set. Furthermore, if the mechanism picks a set A with the maximum virtual surplus over all sets satisfying the thresholds, then it also achieves good expected profit. We call such a mechanism a *threshold mechanism*. Note that the price charged to a bidder still depends on other bidders' bids and not just the threshold (and can therefore change as others' bids change, even when the allocation stays the same); the threshold only ensures that this price is never too low.

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. See an example in the full paper for details.

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 $A \neq \mathcal{B}$, we have $c(A) = \infty$.

[3] This mechanism is based on a social welfare maximizing mechanism due to Archer et al. [1] and assumes that multiple units of each item are available.

Lemma 5. *Let $\mathcal{M}$ be any truthful mechanism for an all-or-none cost function c. Then, $p_i(\mathbf{b})$ is non-increasing in bids b_j with $j \neq i$.*

Proof. Suppose that there are bidders i and j such that $p_i(\mathbf{b})$ not non-increasing in b_j. That is, there are bid vectors $\mathbf{b}$ and $bids'$ with $b_k{}' = b_k$ for all $k \neq j$ and $b_j{}' > b_j$, such that $p_i(\mathbf{b}') > p_i(\mathbf{b})$. Note that $\mathcal{M}$ is truthful, and so p_i does not depend on b_i. So we choose $b_i = b_i{}' = \frac{p_i(\mathbf{b}')+p_i(\mathbf{b})}{2}$. Now, i is served at $\mathbf{b}$ but not at $\mathbf{b}'$. However, since c is an all-or-none cost function, we have $\mathcal{M}(\mathbf{b}) = S$ (all bidders) and $\mathcal{M}(\mathbf{b}') = \emptyset$. This means that the allocation given by $\mathcal{M}$ is not a non-increasing function of b_j. Lemma 1 then implies a contradiction to the truthfulness of $\mathcal{M}$. □

Theorem 5. *For any all-or-none cost function, there exists a threshold mechanism that is optimal among the class of all truthful no-deficit mechanisms.*

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

Let $\bar{\mathbf{b}}$ be the bid vector with $\bar{b}_i = h$, the highest bid, for every i. Let $\tau_i(S) = p_i(\bar{\mathbf{b}})$ for all i. Then, $\sum_i \tau_i(S) = \sum_i p_i(\bar{\mathbf{b}}) \geq c(S)$, because $\mathcal{M}$ satisfies no-deficit. Consider the threshold mechanism $\mathcal{M}'$ given by thresholds τ_i.

For any bid vector $\mathbf{b}$ with $\mathcal{M}(\mathbf{b}) = S$, we must have $\hat{\mathcal{S}}_{\mathbf{b}}(S) > 0$. Otherwise, the mechanism $\mathcal{M}''$ given by $\mathcal{M}''(\mathbf{b}) = \mathcal{B}$ if $\mathcal{M}(\mathbf{b}) = S$ and $\hat{\mathcal{S}}_{\mathbf{b}}(S) > 0$ achieves a higher profit than $\mathcal{M}$ and also satisfies the no-deficit condition. Note also, that for all $\mathbf{b}$ with $\mathcal{M}(\mathbf{b}) = S$ and all i, we have $b_i \geq p_i(\mathbf{b}) \geq p_i(\bar{\mathbf{b}}) = \tau_i(S)$. Here the second inequality follows from Lemma 5. These two conditions along with the definition of $\mathcal{M}'$ imply that $\mathcal{M}'(\mathbf{b}) = S$.

This means that for all $\mathbf{b}$ with $\mathcal{M}(\mathbf{b}) = S$, we have $\mathcal{M}'(\mathbf{b}) = S$. Furthermore, for all $\mathbf{b}$ with $\mathcal{M}(\mathbf{b}) = \emptyset$, we have $\hat{\mathcal{S}}_{\mathbf{b}}(\mathcal{M}'(\mathbf{b})) \geq 0 = \hat{\mathcal{S}}_{\mathbf{b}}(\mathcal{M}(\mathbf{b}))$. Therefore, we get $\hat{\mathcal{S}}_{\mathbf{b}}(\mathcal{M}(\mathbf{b})) \leq \hat{\mathcal{S}}_{\mathbf{b}}(\mathcal{M}'(\mathbf{b}))$, for all vectors $\mathbf{b}$. Lemma 3 now implies that $\mathcal{M}'$ has a larger expected profit than $\mathcal{M}$. □

5.1 The Hardness of Computing the Optimal 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. Via a reduction from the knapsack problem, 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). See the full paper for details.

Theorem 6. *Computing the optimal no-deficit mechanism is NP-hard.*

6 Bayesian Incentive Compatible Mechanisms

We now consider relaxing ex post incentive compatibility to consider Bayesian incentive compatible (BIC) mechanisms. See the full paper for proofs of the following results.

Theorem 7. *The optimal BIC no-deficit mechanism gets the same expected profit as Myerson.*

The allocation procedure of this optimal BIC mechanism is precisely the allocation procedure of Myerson; the payment rule, however, is different. A BIC no-deficit payment rule can be derived by shifting payment from inputs where there is a deficit to ones where their is a surplus. These shifts can be done based on the joint density function, $\mathbf{F}$, so as to keep the expected payment of an agent, given their valuation, the same.

Although the proof of this theorem is constructive, it does not give a polynomial time procedure for computing the prices in general. Interestingly, when there are only two agents, there is a much simpler way of achieving optimality.

Lemma 6. *The optimal BIC no-deficit mechanism for two agents is to charge each agent the expected payment they must make conditioned on being allocated.*

Unfortunately, as the next lemma shows, this simple technique does not extend to more than two bidders.

Lemma 7. *The BIC mechanism for three or more agents that charges each agent the expected payment they must make conditioned on being allocated does not always satisfy the no-deficit constraint.*

Proof. Consider the following counter-example: there are three identical agents; each (independently) has a value of 2 with probability 0.9 and 11 with probability 0.1. The corresponding virtual valuations are 1 and 11 respectively. The costs of serving any one, any two, or all three of the agents are 10.99, 20 and 20 respectively. When all three bidders bid 11, Myerson serves all of them at a price of 2 each, incurring a deficit of 14. When two of the bidders bid 11, they are all served and each is charged a price equal to her bid. When only one bidder bids 11, the bidder is served at a price of 11. The expected payment of a bidder when bidding 11 can be computed to be 10.91. So when all the bidders bid 11, their combined expected payments are sufficient to cover the total cost of 20. On the other hand, the expected payment of a bidder on bidding 2 and losing is 0. Therefore, when the three bidders bid 11, 2, and 2, the sum of their expected payments is $10.91 < 10.99$, which is insufficient to cover the cost of serving the highest bidder. □

The counter-example in the above proof shows that other natural approaches fail as well and implies that the proof of Theorem 7 is necessarily not simple.

7 Conclusions

In this work we have explored the issue of merging the worst case no-deficit condition with the average case Bayesian optimality objective. We have found that for many interesting classes of problems it is not easy to describe the optimal solution nor is there a known algorithm for computing it. Particular questions of interest are:

1. Is there a concise description of the Bayesian optimal no-deficit ex post incentive compatible mechanism? In particular this question is interesting for submodular and general cost functions.
2. Is there a concise description of the payment rule of the Bayesian optimal no-deficit Bayesian incentive compatible mechanism? (Recall that the allocation rule is the same as Myerson's.)
3. Is there an algorithm that computes the Bayesian optimal no-deficit Bayesian incentive compatible mechanism for submodular costs? It is possible to compute the allocation so the open question is to compute the payments.
4. The BIC no-deficit mechanism constructed in our proof of Theorem 7 is only *ex interim* individually rational for the agents (i.e., they may have negative utility). This is standard for no-deficit mechanism design in economics. It is an open question as to whether there is a no-deficit mechanism that is also ex post individually rational.

Acknowledgments

We are grateful to Ilya Segal, Amitabh Sinha, and Rakesh Vohra for many 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. J. Hartline. Stanford cs364b: Topics in algorithmic game theory. Course Notes, 2006.
10. 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.
11. Vijay Krishna. *Auction theory*. Academic Press, San Diego, CA, USA, 2002.
12. 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.
13. A. Mehta, S. Shenker, 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.
14. H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance Versus Efficiency. *Economic Theory*, 18:511–533, 2001.
15. R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
16. A. Ronen. On Approximating Optimal Auctions. In *Proc. of Third ACM Conference on Electronic Commerce*. ACM Press, New York, 2001.