Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers

Saeed Alaei
Dept. of Computer Science
University of Maryland
College Park, MD 20742
Email: saeed@cs.umd.edu

Abstract— For Bayesian combinatorial auctions, we present a general framework for approximately reducing the mechanism design problem for multiple buyers to the mechanism design problem for each individual buyer. Our framework can be applied to any setting which roughly satisfies the following assumptions: (i) the buyer's types must be distributed independently (not necessarily identically), (ii) the objective function must be linearly separable over the set of buyers, and (iii) the supply constraints must be the only constraints involving more than one buyer. Our framework is general in the sense that it makes no explicit assumption about any of the following: (i) the buyer's valuations (e.g., submodular, additive, etc), (ii) The distribution of types for each buyer, and (iii) the other constraints involving individual buyers (e.g., budget constraints, etc).

We present two generic n-buyer mechanisms that use 1-buyer mechanisms as black boxes. Assuming that we have an α-approximate 1-buyer mechanism for each buyer[1] and assuming that no buyer ever needs more than $\frac{1}{k}$ of all copies of each item for some integer $k \geq 1$, then our generic n-buyer mechanisms are $\gamma_k \cdot \alpha$-approximation of the optimal n-buyer mechanism, in which γ_k is a constant which is at least $1 - \frac{1}{\sqrt{k+3}}$. Observe that γ_k is at least $\frac{1}{2}$ (for $k = 1$) and approaches 1 as k increases. As a byproduct of our construction, we improve a generalization of prophet inequalities. Furthermore, as applications of our main theorem, we improve several results from the literature.

1. Introduction

The main challenge of stochastic optimization arises from the fact that all instances in the support of the distribution are relevant for the objective and this support is exponentially big in the size of problem. This paper addresses this challenge by giving a general decomposition technique for assignment problems on independently distributed inputs where the objective is linearly separable over the inputs. The main challenge faced by such a decomposition approach is that the feasibility constraint of an assignment problem introduces correlation in the outcome of the optimal solution. In mechanism design problems, such constraints are usually the supply constraints. For example a revenue maximizing seller with unlimited supply who is facing independent buyers can decompose the problem over the buyers and optimize for each buyer independently. However, in the presence of supply constraints, a direct decomposition is not possible. Our decomposition technique can be roughly described as the following. (i) Construct a mechanism that satisfies the supply constraints only in expectation (ex-ante). The optimization problem for constructing such a mechanism can be fully decomposed over the set of buyers. (ii) Convert the mechanism from the previous step to another mechanism that satisfies the supply constraint at every instance.

We restrict our discussion to Bayesian combinatorial auctions. We are looking for mechanisms for allocating a set of heterogenous items with limited supply to a set of buyers in order to maximize the expected value of a certain objective function that is linearly separable over the buyers (e.g., welfare, revenue, etc). The buyers' types are distributed independently according to a publicly know prior. We defer the formal statement of our assumptions to section 2.

There are two challenges in designing mechanisms for multiple buyers:

(I) The decisions made by the mechanism for different buyers should be coordinated because of supply constraints.
(II) The decisions made by the mechanism for each buyer has to be optimal (or approximately optimal).

The first difficulty is due to the fact that the desired mechanism has to be optimized over the joint type space of all buyers, the size of which grows exponentially in the number of buyers. The second difficulty is usually due to the incentive compatibility (IC) constraints, specially in multi-dimensional settings where these constraints cannot be encoded compactly. In this paper, we mostly address the first difficulty by providing a frame-

Part of this work was done when the author was visiting Microsoft Research, New England. This work was partially supported by the NSF grant CCF-0728839

[1]Note that we can use different 1-buyer mechanisms to accommodate different classes of buyers.

work for approximately decomposing and reducing the mechanism design problem for multiple buyer to the mechanism design problem for individual buyers.

Our framework can be summarized as follows. We start by relaxing the supply constraints (i.e., the constraints specifying that the total number of allocated copies of each item cannot exceed the supply of that item) to hold only in expectation. In other words, we consider the set of mechanisms with the property that for each item only the ex-ante expected number of allocated copies does not exceed the supply. We show that the optimal mechanism for the relaxed problem can be constructed by independently running n single buyer mechanisms, where each single buyer mechanism is optimal for the corresponding buyer among all single buyer mechanisms that are restricted not to allocate each item with an ex-ante expected probability exceeding a certain threshold specific to each item. In particular, if we can construct and α-approximate single buyer mechanism for each buyer, then we can combine them to get an α-approximate mechanism for all buyers. We then present two methods for converting this mechanism to a mechanism that meets the supply constraints at every instance, while losing only a small constant factor in the approximation. In the first method, we serve buyers sequentially by running, for each buyer, the corresponding single buyer mechanism from the previous step. However, we sometimes randomly preclude some of the items from being offered to some buyers in order to ensure that buyers that are served later also get a chance of being offered with those items. We do this in such a way that would ensure that the ex-ante expected probability of preclusion is equalized over all buyers, and therefore simultaneously minimized for all buyers. In the second method, we run all of the single buyer mechanisms simultaneously and then modify the outcomes by deallocating some copies of the over-allocated items at random while adjusting the payments respectively. We do this in such a way that would ensure the ex-ante expected probability of deallocation for each item is equalized among all copies of that item and therefore simultaneously minimized for all buyers.

We also introduce a toy problem, *the magician's problem*, along with a near optimal solution for it. The solution of this problem is used as the main ingredient for converting mechanisms for the relaxed problem to mechanisms for the original problem. It also yields improved generalized prophet inequalities through a direct reduction.

As applications of our general framework, we construct improved mechanisms for several settings from the literature[2]. For each setting we only construct a single buyer mechanism that satisfies the requirements of our framework, and then our generic construction can be applied to construct a mechanism for multiple buyers, using the single buyer mechanism as the building block.

Some of the proofs and other results are omitted from the current version of the paper due to space constraints, but are available in the full version.

1.1. Related Work

In single dimensional settings, most of the related works form the CS literature are either focused on approximating the VCG mechanism for welfare maximization, or approximating the Myerson's mechanism [16] for revenue maximization (e.g., [5], [2], [4], [14], [9], [6], [17]). Most of them consider mechanisms that have simple implementation and are computationally efficient. For welfare maximization in single dimensional settings, [13] gives a blackbox reduction from mechanism design to algorithmic design.

In multidimensional setting, for welfare maximization, [12] presents a blackbox reduction from mechanism design to algorithm design which subsumes the earlier work of [13]. For revenue maximization, [7] presents several sequential posted pricing mechanisms for various settings with different types of matroid feasibility constraints. These mechanisms have simple implementation and approximate the revenue of the optimal mechanism. For a special form of combinatorial auctions with hard budget constraints, [3] presents an all pay BIC mechanism and a sequential posted pricing mechanism. [8] also considers various settings with hard budget constraints.

Prophet inequalities have been extensively studied in the past (e.g. [15]). The best known bound for the generalization to sum of k choices was $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ by [11] which we improve to $1 - \frac{1}{\sqrt{k+3}}$. Note that the current bound is not only asymptotically better than the previous bound, but is also tight for $k = 1$, where as the previous bound would be useful only for large k.

2. Model & Overview of Approach

Model: We consider mechanisms for selling m indivisible heterogenous items to n buyers where there are k_j copies of each item $j \in [m]$. All the relevant private information of each buyer $i \in [n]$ is represented by her type $t_i \in T_i$ where T_i is the type space of buyer i. The profile of types $t = (t_1, \cdots, t_n)$ is distributed according to a publicly known prior $\mathcal{D}$. We are restricted to mechanisms from a given space of mechanisms $\mathbb{M}$. For a mechanism $\mathcal{M}$, we use $X_{ij}^{\mathcal{M}}(t)$ and $P_i^{\mathcal{M}}(t)$ to denote the random variables[3] respectively

[2]Only one application is included in the conference version of the paper.

[3]Note that these random variables are often correlated. Furthermore, if $\mathcal{M}$ is a deterministic mechanism then for any given t these variables take deterministic values as a function of t.

for allocation of item j to buyer i and payment of buyer i, when the profile of types is t. We are looking for mechanisms that maximize the expected value of a given objective function $W(t,x,p)$ where t, x, and p respectively represent the types, the allocations, and the payments of all buyers. Formally, we are looking for a mechanism $\mathcal{M} \in \mathbb{M}$ that (approximately) maximizes $E_{t\sim\mathcal{D}}[W(t, X^{\mathcal{M}}(t), P^{\mathcal{M}}(t))]$.

Assumptions: We make the following assumptions:

(A1) The buyers' types must be distributed independently, i.e., $\mathcal{D} = \mathcal{D}_1 \times \cdots \times \mathcal{D}_n$ must be a product distribution.

(A2) The objective function must be linearly separable over the buyers, i.e., $W(t,x,p) = \sum_i W_i(t_i, x_i, p_i)$ where t_i, x_i, and p_i respectively represent the type, the allocations, and the payment of buyer i.

(A3) No buyer must ever need more than one copy of each item, i.e., $X^{\mathcal{M}}_{ij}(t) \in \{0,1\}$ for all t.[4]

(A4) $\mathbb{M}$ must be restricted to (Bayesian) incentive compatible mechanisms. By direct revelation principle this assumption is without loss of generality.

(A5) $\mathbb{M}$ must be a convex space. In other words, any convex combination of any two mechanisms from $\mathbb{M}$ must itself be a mechanism in $\mathbb{M}$. A convex combination of two mechanisms $\mathcal{M}, \mathcal{M}' \in \mathbb{M}$ is another mechanism $\mathcal{M}''$ which simply runs $\mathcal{M}$ with probability α and runs $\mathcal{M}'$ with probability $1-\alpha$, for some $\alpha \in [0,1]$. In particular, if $\mathbb{M}$ is restricted to deterministic mechanisms, then it is not convex. [5].

(A6) The set of constraints that specify $\mathbb{M}$ must be decomposable to supply constraints and single buyer constraints. Note that IC constraints, budget constraints, etc., are all single buyer constraints. We define this assumption formally as follows. For any mechanism $\mathcal{M}$, let $[[\mathcal{M}]]_i$ be the single buyer mechanism perceived by buyer i, as if the other buyers are part of the mechanism. Let $\mathbb{M}_i = \{[[\mathcal{M}]]_i | \mathcal{M} \in \mathbb{M}\}$ be the space of single buyer mechanisms perceived by buyer i resulting from mechanisms in $\mathbb{M}$. We require that for any mechanism $\mathcal{M}$, if $\mathcal{M}$ satisfies the supply constraints and also $[[\mathcal{M}]]_i \in \mathbb{M}_i$ (for all $i \in [n]$), then it must be that $\mathcal{M} \in \mathbb{M}$.

We shall clarify the last assumption by giving an example. Suppose $\mathbb{M}$ is the space of all truthful buyer specific item pricing mechanisms, then $\mathbb{M}$ satisfies the last assumption. On the other hand, if $\mathbb{M}$ is the space of mechanisms that offer the same set of prices to every buyer, then it does not satisfy the decomposability assumption.

Formally, the problem we are looking at is to find a mechanism $\mathcal{M}$ that is a solution to the following program:

$$\text{maximize:} \qquad \sum_i E_{t\sim\mathcal{D}}[W_i(t_i, X^{\mathcal{M}}_i(t), P^{\mathcal{M}}_i(t))] \tag{P}$$

subject to:

$$\forall t \in T, \forall j \in [m]: \qquad \sum_i X^{\mathcal{M}}_{ij}(t) \le k_j \tag{S}$$

$$\forall i \in [n]: \qquad [[\mathcal{M}]]_i \in \mathbb{M}_i \tag{M}$$

Summary of Approach: We now present an overview of our general framework for constructing approximately optimal mechanisms for the above program. We start by relaxing the supply constraints to hold only in expectation. We show that the optimal mechanism for the relaxed problem can be constructed by combining n independent single buyer mechanisms. We then present two approaches for converting the mechanism constructed in the previous step to a mechanism for the original problem. Each step is explained in more details next.

The problem is initially relaxed by requiring the supply constraints to hold only in expectation. In other words, the constraints (S) are replaced with the following constraints:

$$\forall j \in [m]: \qquad E_{t\sim\mathcal{D}}[\sum_i X^{\mathcal{M}}_{ij}(t)] \le k_j \tag{S'}$$

We show that an optimal mechanism for the relaxed problem can be constructed by combining n independent single buyer mechanisms. We shall first present the following definition.

Definition 1 (Primary Mechanism/Primary Benchmark)**.** *A primary mechanism for buyer i is a single buyer mechanism $\mathcal{M}_i$ that allows specifying an upper bound on the expected probability of allocating each item. For every $\bar{q}_i \in [0,1]^m$, $\mathcal{M}_i(\bar{q}_i)$ is a single buyer mechanism in $\mathbb{M}_i$ for which the ex-ante expected probability of allocating a copy of item j to buyer i is at most $\bar{q}_{ij}$. The optimal primary mechanism is the one that has the highest expected objective value.*

A primary benchmark for buyer i is a function $R_i : [0,1]^m \to \mathbb{R}_+$ that returns an upper bound on the expected objective value of the optimal primary mechanism for buyer i. If R_i is the optimal primary benchmark for buyer i then $R_i(\bar{q}_i)$ is exactly equal to the expected objective value of the optimal primary mechanism subject to $\bar{q}_i$.

[4] This assumption is not necessary and can be lifted as explained in the full paper.

[5] As an example of a randomized space of mechanisms without this property, consider the space of mechanisms where the expected payment of every type must be either less than \$2 or more than \$4

We show that an optimal mechanism for the relaxed problem can be constructed from n independent optimal primary mechanisms. Let $\mathcal{M}^*$ be any optimal mechanism for the relaxed problem and let $q^*_{ij} = E_{t\sim\mathcal{D}}[X^{\mathcal{M}^*}_{ij}(t)]$ be the expected probability that $\mathcal{M}^*$ allocates a copy of item j to buyer i. Let $\mathcal{M}_i$ denote the optimal primary for each buyer i. In section 4, we prove that the mechanism that runs $\mathcal{M}_i(q^*_i)$ independently for buyer i has the same expected objective value as $\mathcal{M}^*$. Therefore, if we can construct the optimal primary mechanism for each buyer, then we can construct an optimal mechanism for the relaxed problem by simply using $\mathcal{M}_i(q^*_i)$ independently for each buyer i, assuming that we know how to compute the q^*_{ij}. We will show that q^*_{ij} is the optimal assignment for the following program in which R_i is the optimal primary benchmark for buyer i:

$$\begin{aligned} &\text{maximize:} && \sum_i R_i(\bar{q}_i) && (CP_R)\\ &\forall j\in[m]: && \sum_i \bar{q}_{ij} \le k_j \\ &\forall i\in[n], \forall j\in[m]: && \bar{q}_{ij}\in[0,1] \end{aligned}$$

In particular, in section 4, we prove that the optimal primary benchmarks $R_i(\cdot)$ are always concave, and therefore the above program is a convex program. Consequently, this program can be efficiently solved to compute q^*_{ij}. Note that usually each function R_i is itself the optimal objective value of a linear or convex program, and does not have a closed form; in that case, all the corresponding linear/convex programs can be merged into one. So far, we have explained how the problem of constructing an optimal mechanism for the relaxed problem can be reduced to the problem of constructing the optimal primary mechanisms/primary benchmarks. Next, we explain how to convert it to a mechanism for the original problem.

We now present two approaches for converting a mechanism for the relaxed problem to a mechanism for the original problem, while losing only a small constant fraction of the expected objective value. Let $\mathcal{M}$ be the mechanism constructed in the previous step, which uses $\mathcal{M}_i(\bar{q}_i)$ independently for each buyer i. Since $\mathcal{M}$ satisfies the supply constraints only in expectation, it most likely violates those constraints in some instances. We propose two separate approaches for dealing with this issue, each one yielding a generic mechanism for the original (non-relaxed) problem. The following is a high level description of these two generic mechanisms:

1) **Pre-Rounding:** This mechanism serves buyers sequentially in an arbitrary order. For each buyer i, it runs $\mathcal{M}_i(\bar{q}'_i)$ in which $\bar{q}'_i$ is the same as $\bar{q}_i$ except that some of its entries are set to 0 as explained next. The outcome of $M_i(\bar{q}'_i)$ is taken as the final outcome for buyer i. Setting $\bar{q}'_{ij}$ to 0 effectively precludes $\mathcal{M}_i(\bar{q}'_i)$ from allocating a copy of item j to buyer i. The supply constraints are enforced by setting $\bar{q}'_{ij}$ to 0 for any item j that is sold out prior to serving buyer i. Moreover, for each item, the mechanism tries to minimize simultaneously for all buyers the expected probability of preclusion by equalizing this expected probability for all buyers. Effectively, the mechanism sometimes precludes some items from being offered to earlier buyers in order to make sure that later buyers get the same chance of being offered with those items. Note that, for any given pair of buyer and item, we only care about the expected probability of preclusion where the expectation is taken over the types of other buyers. In particular, an item might be precluded from the current buyer with probability 1 if certain scenarios of outcomes have been realized for buyers served prior to the current buyer. We show that if there are at least k copies of each item then the expected probability of preclusion of each item for each buyer is no more than $\frac{1}{\sqrt{k+3}}$.
2) **Post-Rounding:** This mechanism runs $M_i(\bar{q}_i)$ for each buyer i independently and then modifies the outcomes by deallocating some of the items at random to ensure that the supply constraints are met at every instance. This is done in such a way that would minimize the expected probability of deallocation observed by each buyer by equalizing this probability over all copies of each item. The payments are also scaled down accordingly by the same probability. Note that, for any given pair of buyer and item, we only care about the expected probability of deallocation, where the expectation is taken over the types of other buyers. In particular, a buyer who faces a small expected probability of deallocation could still face a deallocation probability of 1 for some items when certain profiles of types are reported by other buyers. We show that if there are at least k copies of each item, then the expected probability of deallocation is no more than $\frac{1}{\sqrt{k+3}}$ for each copy.

In section 4 we explain the above mechanisms in more details and present some technical assumptions that are sufficient to ensure that they retain at least a $1-\frac{1}{\sqrt{k+3}}$ fraction of the expected objective value of $\mathcal{M}$.

Throughout the above discussion, we have assumed that we can construct the optimal primary mechanisms and the optimal primary benchmarks. However, it is likely that we can only construct an approximation of them. Suppose for each buyer i, we only have an α-approximate primary mechanism and a corresponding concave primary benchmark R_i (i.e., the expected objective value of $\mathcal{M}_i(\bar{q}_i)$ is at least $R_i(\bar{q}_i)$ for every $\bar{q}_i \in [0,1]^m$). Then we can still used $\mathcal{M}_i$ and R_i in the above construction, but the final approximation factor will be multiplied by α.

Main Result: The main result of this paper can be summarized in the following informal theorem.

Theorem 1 (Market Expansion)**.** *Suppose for each buyer $i \in [n]$, we have an α-approximate primary mechanism $\mathcal{M}_i$ and a corresponding concave primary benchmark R_i. Then, with some further assumptions (explained later), we can efficiently construct a mechanism $\mathcal{M} \in \mathbb{M}$, using the primary mechanisms as building blocks, such that the expected objective value of $\mathcal{M}$ is at least $\gamma_k \cdot \alpha$-fraction of the expected objective value of the optimal mechanism, where $k = \min_j k_j$ and γ_k is a constant which is at least $1 - \frac{1}{\sqrt{k+3}}$.*

In order to explain our construction in more details, we shall first describe the magician's problem and its solution, which is used in equalizing the expected probabilities of preclusion/deallocation over all buyers.

3. The Magician's Problem

In this section, we present an abstract online stochastic toy problem and a near-optimal solution for it. The solution to this problem is the main ingredient for combining single buyer mechanisms to construct mechanisms for multiple buyers. The solution to this problem is used to prove a generalized prophet inequality. Furthermore, it has applications in online stochastic optimization[6]

Definition 2 (The Magician's Problem)**.** *A magician is presented with a series of boxes one by one. There is a prize hidden in one of the boxes. He has k magic wands that can be used to open the boxes. On each box is written a probability. If a wand is used on a box, it opens, but with at most the written probability the wand breaks. Let q_i denote this probability for the i^{th} box. The magician wishes to maximize the probability of obtaining the prize, but unfortunately the sequence of boxes, the written probabilities, and the box in which the prize is hidden are arranged by a villain, and the magician has no prior information about them (not even the number of boxes). However, it is given that $\sum_i q_i \leq k$, and that the villain cannot make any changes once the process has started.*

The magician could fail to open a box either because he might choose to skip the box or because he might run out of magic wands before coming to the box. Note that once the magician fixes his strategy, the best strategy for the villain is to put the prize in the box that has the lowest ex-ante expected probability of being opened, based on the magician's strategy. Therefore, in order for the magician to obtain the prize with a probability of at least γ, he has to devise a strategy that guarantees an ex-ante expected probability of at least γ for opening each box. Note that the nature of the prize or even whether there is actually a prize does not affect the problem. It is easy to show the following strategy ensures an ex-ante expected probability of at least $\frac{1}{4}$ for opening each box: For each box randomize and use a wand with probability $\frac{1}{2}$. But can we do better? Next we present an algorithm that takes a parameter γ and tries to ensure a minimum ex-ante expected probability of γ for opening each box. In Theorem 2, we show that for any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$ this algorithm indeed guarantees that the ex-ante expected probability of opening each box is at least γ.

Algorithm 1 (γ-Conservative Magician)**.** *The magician constructs a strategy table y_i^j using the dynamic programs given below. y_i^j specifies the probability with which the magician should choose to open the i^{th} box if j wands have been broken prior to seeing the i^{th} box. So if $y_i^j = 0$ or $y_i^j = 1$, then the magician makes a deterministic decision, otherwise he should randomize and open the i^{th} box with probability y_i^j. We use Y_i as the indicator random variable which is 1 iff the magician chooses to open the i^{th} box. The strategy table can be computed using the following dynamic programs (note that γ is a parameter that is given):*

$$y_i^j = \begin{cases} 1 & i\geq 1, 0\leq j<\theta_i \\ (\gamma-\phi_i^{\theta_i -1})/(\phi_i^{\theta_i}-\phi_i^{\theta_i -1}) & i\geq 1, j=\theta_i \\ 0 & \textit{otherwise.} \end{cases} \qquad (DP.y)$$

$$\theta_i = \min\{j | \phi_i^j \geq \gamma\} \qquad (DP.\theta)$$

$$\phi_i^j = \begin{cases} 1 & i=1, j\geq 0 \\ y_{i-1}^j q_{i-1} \phi_{i-1}^{j-1} + (1-y_{i-1}^j q_{i-1})\phi_{i-1}^j & i\geq 2, j\geq 0 \\ 0 & \textit{otherwise.} \end{cases} \qquad (DP.\phi)$$

Note that computing y_i^j only requires the knowledge of $q_1, \cdots, q_{i-1}$, so computing y_i^j and making a decision about the i^{th} box can be done even before seeing the i^{th} box itself.

Interpretation of γ-Conservative Magician(Alg. 1): The main idea of the algorithm is the following: After seeing the first $i-1$ boxes and prior to the arrival of the i^{th} box, the magician computes a threshold θ_i as follows. θ_i is the smallest integer such that the ex-ante expected probability of having lost at most θ_i wands, on the first $i-1$ boxes, is at least γ. In other words, if S_i is the random variable that represents the number of magic wands broken prior to seeing the i^{th} box, then θ_i

[6]Refer to the full version for more details.

is chosen to be the smallest integer such that $Pr[S_i \leq \theta_i] \geq \gamma$. Observe that if the magician always opens the i^{th} box when the number of wands broken so far is no more than θ_i, and otherwise discards the box, then he can guarantee an ex-ante probability of at least γ for opening the i^{th} box. Furthermore, if $Pr[S_i \leq \theta_i] > \gamma$, i.e. if the inequality is strict, then in the event of having broken exactly θ_i wands prior to the i^{th} box, the magician randomizes and opens the i^{th} box with a probability strictly less than 1, which is just enough to ensure that the total ex-ante expected probability of opening the i^{th} box is at least γ. It can be verified that ϕ_i^j, as defined by the dynamic program, is a lower bound on $Pr[S_i \leq j]$. In fact, if $q_1, \cdots, q_{i-1}$ are the exact probabilities of breaking a wand for each of the first $i-1$ boxes, then $Pr[S_i \leq j] = \phi_i^j$. In order to prove that the above strategy ensures that each box is opened with an ex-ante expected probability of at least γ, we need to show that $y_i^j = 0$ for all $j \geq k$ and all i. i.e., we need to show that the strategy table of the magician does not instruct him to open a box if he has broken all of his k wands. In Theorem 2 we present a sufficient condition on γ that ensures $y_i^j = 0$ for all $j \geq k$ and all i.

Theorem 2 (γ-Conservative Magician)**.** *For any $\gamma \leq 1 - \frac{1}{\sqrt{k+3}}$, a γ-conservative magician guarantees that each box is opened with an ex-ante expected probability at least γ. Furthermore, if q_i are the exact probabilities of breaking a wand, then the γ-conservative magician opens each box with an ex-ante expected probability exactly γ*[7]

Definition 3 (γ_k)**.** *We define γ_k to be the largest probability such that for any instance of the magician's problem with k' wands, where $k' \geq k$, a γ_k-conservative magician with k' wands can guarantee that each box is opened with an ex-ante expected probability at least γ_k. By Theorem 2, we know that γ_k must be at least $1 - \frac{1}{\sqrt{k+3}}$ because for any $k' \geq k$ obviously $1 - \frac{1}{\sqrt{k+3}} \leq 1 - \frac{1}{\sqrt{k'+3}}$.*

Observe that γ_k is a non-decreasing function in k which is at least $\frac{1}{2}$ (when $k = 1$) and approaches 1 as k gets larger. The next theorem shows that the lower bound of $1 - \frac{1}{\sqrt{k+3}}$ on γ_k is almost tight.

Theorem 3 (Optimal Magician)**.** *For any $\epsilon > 0$, it is not possible to guarantee an ex-ante expected probability of at least $1 - \frac{k^k}{e^k k!} + \epsilon$ for opening each box(i.e., no magician can guarantee it). Note that $1 - \frac{k^k}{e^k k!} \approx 1 - \frac{1}{\sqrt{2\pi k}}$.*

Next, we prove a generalization of prophet inequalities by a direct reduction to the magician's problem.

Definition 4 (Sum of k-Choices)**.** *A sequence of n non-negative random variables $V_1, \cdots, V_n$ are presented to a gambler one by one in an arbitrary order. The gambler knows n and the distribution of each random variable in advance but not the order in which they are presented. Upon being presented with the random variable V_i, the gambler observes the actual draw of V_i and he has to decide whether to keep it or to discard. This decision cannot be changed later. The gambler must select k of the random draws from the sequence. His objective is to maximize the sum of the selected draws. The prophet knows all the actual draws in advance, so he chooses the k highest draws. We assume that the order in which the random variables are presented to the gambler is fixed in advance and does not change during the process.*

It was shown in [11] that there is a strategy for the gambler that guarantees at least $1 - O(\frac{\sqrt{\ln k}}{\sqrt{k}})$ fraction of the payoff of the prophet, in expectation, by using a non-decreasing sequence of k stopping rules (thresholds) [8]. Next, we construct a gambler that obtains at least γ_k fraction of the prophet's payoff, in expectation, by using a γ_k-conservative magician as a black box. Note that $\gamma_k \geq 1 - \frac{1}{\sqrt{k+3}}$. This gambler uses only a single threshold. However, he may skip some of the random variables at random.

Theorem 4 (Prophet Inequalities for Sum of k Choices)**.** *The following strategy ensures that the gambler obtains at least γ_k fraction of payoff of the prophet in expectation.* [9]

- *Find the threshold τ such that $\sum_i Pr[V_i > \tau] = k$. This can be done by doing a binary search on τ.*
- *Use a γ_k-conservative magician with k magic wands. Upon seeing each V_i, create a box and write $q_i = Pr[V_i > \tau]$ on the box and present it to the magician. If the magician chooses to open the box and also $V_i > \tau$, then select V_i and break the magician's wand, otherwise skip V_i.*

Proof: First, we compute an upper bound on the expected payoff of the prophet. Let q_i be the probability that the prophet chooses V_i (i.e. the probability that V_i is among the k highest draws). Now let $u_i(q_i)$ denote the maximum possible contribution of the random variable V_i to the expected payoff of the prophet

[7] In particular the fact that the probability of breaking a wand for the i^{th} box is exactly q_i conditioned on any sequence of prior events implies that for each box the event of breaking a wand has to be independent of the sequence of past events and independent of other boxes.

[8] A gambler with stopping rules $\tau_1, \cdots, \tau_k$ works as follows. Upon seeing V_i, he selects it iff $V_i \geq \tau_{j+1}$ where j is the number of random draws selected so far.

[9] To simplify the exposition we assume that the distribution of each of V_i does not have any point mass. Our theorem holds with slight modifications if we allow point masses.

if V_i is selected with probability q_i. Note that $u_i(q_i)$ is equal to the expected value of V_i conditioned on being above the $1-q_i$ quantile. Let $F_i(\cdot)$ and $f_i(\cdot)$ denote the CDF and PDF of V_i. $u_i(q_i)$ can be defined as $u_i(q_i) = \int_{F_i^{-1}(1-q_i)}^{\infty} v f_i(v) dv$. By changing the integration variable and applying the chain rule we get $u_i(q_i) = \int_0^{q_i} F_i^{-1}(1-q)dq$. Observe that $\frac{d}{dq_i}u_i(q_i) = F_i^{-1}(1-q_i)$ is a non-increasing function so $u_i(q_i)$ is a concave function. Furthermore, $\sum_i q_i \leq k$ because the prophet cannot choose more than k random draws. So the optimal objective value of the following convex program is an upper bound on the payoff of the prophet:

$$\begin{aligned} \text{maximize:} \quad & \sum_i u_i(q_i) && (U) \\ & \sum_i q_i \leq k && (\tau) \\ \forall i \in [n]: \quad & q_i \geq 0 && (\mu_i) \end{aligned}$$

Now let $L(q,\tau,\mu) = -\sum_i u_i(q_i) + \tau(\sum_i q_i - k) - \sum_i \mu_i q_i$ be the Lagrangian. By KKT stationarity condition, at the optimal assignment, it must be that $\frac{\partial}{\partial q_i}L(q,\tau,\mu) = 0$. On the other hand, $\frac{\partial}{\partial q_i}L(q,\tau,\mu) = -F_i^{-1}(1-q_i) + \tau - \mu_i$. Assuming that $q_i > 0$, then by complementary slackness $\mu_i = 0$, which then implies that $q_i = 1 - F_i(\tau)$, so $q_i = Pr[V_i > \tau]$. Furthermore, it is easy to show that the first constraint must be tight, which implies that $\sum_i Pr[V_i > \tau] = k$. Observe that the contribution of each V_i to the objective value of the convex program is exactly $E[V_i|V_i > \tau]Pr[V_i > \tau]$. Now, by using a γ_k-conservative magician we can ensure that each box is opened with probability at least γ_k which implies the contribution of each V_i to the expected payoff of the gambler is $E[V_i|V_i > \tau]Pr[V_i > \tau]\gamma_k$ which proves that the expected payoff of the gambler is at least γ_k fraction of optimal objective value of the convex program, which was itself and upper bound on the expected payoff of the prophet. ∎

4. The Two Generic Mechanisms

In this section, we present the details of the approach that was outlined in section 2. The model and assumptions were explained in that section. We start by proving that an (approximately) optimal mechanism for the relaxed problem can always be constructed from n (approximately) optimal primary mechanisms.

Theorem 5. *Suppose for each buyer i, we have an α-approximate primary mechanism $\mathcal{M}_i$ and a matching concave primary benchmark R_i, as defined in Def. 1. Consider the mechanism $\mathcal{M}$ which simply uses $\mathcal{M}_i(\bar{q}_i)$ independently for each buyer i, where $\bar{q}_i$ is the optimal assignment of the following convex program. Then $\mathcal{M}$ is a feasible mechanism for the relaxed problem and its expected object value is at least an α-fraction of the optimal objective value of the convex program, which is itself an upper bound on the expected objective value of the optimal mechanism.*

$$\begin{aligned} \textit{maximize:} \quad & \sum_i R_i(\bar{q}_i) && (CP_R) \\ \forall j \in [m]: \quad & \sum_i \bar{q}_{ij} \leq k_j \\ \forall i \in [n], \forall j \in [m]: \quad & \bar{q}_{ij} \in [0,1] \end{aligned}$$

Proof: Let $\mathcal{M}^*$ be any optimal mechanism for the relaxed problem. For each buyer i, we construct a single buyer mechanism $\mathcal{M}_i^*$ as follows. $\mathcal{M}_i^*$ creates $n-1$ dummy buyers whose types are randomly drawn from $\mathcal{D}_{-i}$. It then runs $\mathcal{M}^*$ on buyers i and the $n-1$ other dummy buyers. Note that buyer i cannot tell the different between $\mathcal{M}_i^*$ and the original $\mathcal{M}^*$ because buyers types are distributed independently. Observe that the contribution of buyer i to the expected objective value of $\mathcal{M}^*$ is the same as her contribution to the expected objective value of $\mathcal{M}_i^*$. So the mechanism that runs $\mathcal{M}_1^*, \cdots, \mathcal{M}_n^*$ independently, has the same expected objective value and the same expected probabilities of allocation as $\mathcal{M}^*$. Now let $q_{ij}^* = E_{t\sim\mathcal{D}}[X_{ij}^{\mathcal{M}^*}(t)]$ be the expected probability that $\mathcal{M}^*$ allocates a copy of item j to buyer i. Observe that q_{ij}^* is a feasible assignment for the convex program. Furthermore, the expected object value of $\mathcal{M}^*$ is equal to the sum of the expected objective values of $\mathcal{M}_1^*, \cdots, \mathcal{M}_n^*$ which is upper bounded by $\sum_i R_i(q_i^*)$. So the optimal objective value of the convex program may only be higher than the expected objective value of $\mathcal{M}^*$. Now observe that the expected objective value of the mechanism $\mathcal{M}$ is at least $\sum_i \alpha \cdot R_i(\bar{q}_i)$ where $\bar{q}_i$ is the optimal assignment for the convex program. So the expected objective value of $\mathcal{M}$ is at least α-fraction of the expected objective value of $\mathcal{M}^*$. ∎

Note that in Theorem 5, R_i are concave functions by definition. However, we shall show that the optimal primary benchmarks are also concave.

Theorem 6. *The optimal primary benchmarks are always concave.*

Proof: We prove this for an arbitrary buyer i. Let $\mathcal{M}_i$ and R_i denote the optimal primary mechanism and the optimal primary benchmark for buyer i. To show that R_i is concave, it is enough to show that for any $q, q' \in [0,1]^m$ and any $\alpha \in [0,1]$, the following inequality holds: $R_i(\alpha q + (1-\alpha)q') \geq \alpha R_i(q) + (1-\alpha)R_i(q')$. Consider the single buyer mechanism $\mathcal{M}''$ that works as follows: $\mathcal{M}''$ uses $\mathcal{M}_i(q)$ with probability α and uses $\mathcal{M}_i(q')$ with probability $1-\alpha$. Because we assumed $\mathbb{M}_i$ is a convex space, $\mathcal{M}'' \in \mathbb{M}_i$. Observe that by linearly of expectation, the expected probabilities of allocation for $\mathcal{M}''$ is no more than $\alpha q + (1-\alpha)q'$ and the expected objective value of $\mathcal{M}''$ is $\alpha R_i(q) + (1-\alpha)R_i(q')$. So the expected objective value of the optimal primary mechanism, subject to the upper bound of $\alpha q + (1-\alpha)q'$ on the expected probabilities of allocation, may only be higher. That implies $R_i(\alpha q + (1-\alpha)q') \geq \alpha R_i(q) + (1-\alpha)R_i(q')$ which proves our claim. ■

Next, we present a detailed description of the two generic mechanisms that were outlined in section 2. Throughout the rest of this section, we assume that for each buyer i, we have an α-approximate primary mechanism $\mathcal{M}_i$ and a corresponding concave primary benchmark R_i. First, we present the pre-rounding mechanism.

Mechanism 1 (γ-Pre-Rounding)**.**

(I) Solve the convex program of (CP_R) and let $\bar{q}_{ij}$ denote an optimal assignment for it.

(II) For each item $j \in 1 \cdots m$: create an instance of γ-conservative magician (see Alg. 1) with k_j magic wands. We will use these magicians through the rest of the mechanism. Note that γ is a parameter that is given.

(III) For each buyer $i \in 1 \cdots n$:

a) For each $j \in 1 \cdots m$: write $\bar{q}_{ij}$ on a box and present it to the j^{th} magician. Let Y_{ij} denote the indicator random variable which is 1 iff the magician opens the box. Set $\bar{q}'_{ij} \leftarrow \bar{q}_{ij} Y_{ij}$.

b) Run the mechanism $\mathcal{M}_i(\bar{q}'_i)$ on buyer i and use its outcome as the final outcome for buyer i. Furthermore, let $X_{i1}, \ldots, X_{im}$ denote the indicator random variables for the allocation of $\mathcal{M}_i(\bar{q}'_i)$.

c) For each $j \in 1 \cdots m$: if $X_{ij} = 1$, then break the wand of the j^{th} magician.

In order for Mech. 1 to retain at least a γ-fraction of the the expected objective value of each $\mathcal{M}_i(\bar{q}_i)$, we have to make further technical assumptions. We show that it is enough to assume that each R_i has a budget-balanced and cross monotonic cost sharing scheme. Next we define this formally.

Definition 5 (Budget Balanced Cross Monotonic Cost Sharing Scheme)**.** *For any subset of items $S \subset [m]$ and any vector $\bar{q} \in [0,1]^m$, let $\bar{q}[S]$ denote the vector whose j^{th} component is $\bar{q}_j$ if $j \in S$ and is 0 otherwise. A primary benchmark function R_i has a budget balanced cross monotonic cost sharing scheme iff there exist a cost share function $\xi : [m] \times \{0,1\}^m \times [0,1]^m \to \mathbb{R}_+$ with the following two properties. ξ must be budget balances which means for all $\bar{q}_i \in [0,1]^m$ and $S \subset [m]$, $\sum_{j \in S} \xi(j, S, \bar{q}_i) = R_i(\bar{q}_i[S])$. Also ξ must be cross monotonic which means for all $\bar{q}_i \in [0,1]^m$, $j \in [m]$ and $S, T \subset [m]$, $\xi(j, S, \bar{q}_i) \geq \xi(j, S \cup T, \bar{q}_i)$.*

The above requirement on R_i can be interpreted in the following manner. Roughly it means that the contribution of each item j to the expected objective value of the primary mechanism for buyer i should not decrease when items other than j are being precluded from buyer i. The following are some examples of environments where this assumption holds. (i) When $R_i(\bar{q}_i[S])$ is a submodular function of S. (ii) For welfare objective, assuming that the buyers' valuations are submodular. (iii) For revenue maximization when the buyers' valuations are submodular and $\mathbb{M}$ is restricted to mechanisms that can be interpreted as buyer specific item pricing.

Theorem 7 (γ-Pre-Rounding)**.** *Suppose for each buyer i we have an α-approximate primary mechanism $\mathcal{M}_i$ and a corresponding concave primary benchmark R_i that has a budget balanced cross monotonic cost sharing scheme. Then, for any $\gamma \in [0, \gamma_k]$, the γ-pre-rounding mechanism (Mech. 1) is a feasible mechanism in $\mathbb{M}$. Furthermore, it is a $\gamma \cdot \alpha$-approximation of the optimal mechanism in $\mathbb{M}$. Note that the resulting mechanism is dominant strategy incentive compatible (DSIC).*

Remark 1. *The γ-pre-rounding mechanism assumes no control and no prior information about the order in which buyers are visited. The order specified in the mechanism is arbitrary and could be replaced by any other ordering which may be unknown in advance. In particular, this mechanism can be adopted to online settings where buyers are served in an unknown order.*

Next, we present the post-rounding mechanism.

Mechanism 2 (γ-Post-Rounding)**.**

(I) Solve the convex program of (CP_R) and let $\bar{q}_{ij}$ denote an optimal assignment for it.

(II) For each buyer $i \in 1 \cdots n$: run the corresponding primary mechanism $\mathcal{M}_i(\bar{q}_i)$ on buyer i and let $X_{i1}, \ldots, X_{im}$ and P_i denote the random variables for the allocation and the payment of $\mathcal{M}_i$. Furthermore, let $\hat{q}_{ij}$ be the actual marginal probability of allocating item j to buyer i by $\mathcal{M}_i(\bar{q}_i)$. Note that $\hat{q}_{ij} \leq \bar{q}_{ij}$.

(III) For each item type $j \in 1 \cdots m$:

a) Create a new instance of the γ-conservative magician (see Alg. 1) with k_j magic wands. This is the j^{th} magician.

b) For each $i \in 1 \cdots n$: create a box corresponding to X_{ij} and write $\hat{q}_{ij}$ on the box and present it to the j^{th} magician. Let Y_{ij} denote the indicator random variable which is 1 iff the magician chooses to open the box. Set $X'_{ij} \leftarrow X_{ij} Y_{ij}$. If $X'_{ij} = 1$ then break the *magician's wand.*

(IV) For each buyer $i \in 1 \cdots n$: charge buyer i a payment of $P'_i \leftarrow \gamma P_i$ and for each $j \in 1 \cdots m$, allocate a copy of item j to buyer i iff $X'_{ij} = 1$.

In order for Mech. 2 to be applicable, we need to make further technical assumptions. The following assumptions are sufficient to ensure that Mech. 2 is truthful and retains at least a γ-fraction of the expected objective value of each $\mathcal{M}_i(\bar{q}_i)$.

- We must be able to compute the actual expected probabilities of allocation for each $\mathcal{M}_i(\bar{q}_i)$. Note that $\bar{q}_i$ is only an upper bound on these probabilities.
- The objective function $W_i(t_i, x_i, p_i)$, must be submodular in x_i and linear in payment.
- The space of mechanism $\mathbb{M}$ should only be restricted to Bayesian incentive compatible mechanisms and may not be restricted to any stronger concept of truthfulness.
- The valuations of each buyer should be in the form of a weighted rank function of some matroid.

Note that the above assumptions are not the only possible set of assumptions that are sufficient for the applicability of Mech. 2. Next we define the last assumption formally.

Definition 6 (Valuations as Matroid Weighted Rank Functions)**.** *Valuations of a buyer for bundles of items can be represented as a weighted rank function of a matroid if there is a matroid whose ground set is the set of items, such that for any bundle S of items:*

- *If S is an independent set of the matroid, then the valuation of the buyer for S is just the sum of her valuations for each item in S.*
- *If S is not an independent set, then the valuation of the buyer for S is equal to her valuation of an independent subset $S' \subset S$ with the maximum valuation.*

In particular, additive valuations with capacities, unit demand valuations, etc. can be represented as matroid weighted rank functions [10].

Theorem 8 (γ-Post-Rounding)**.** *Suppose for each buyer i we have an α-approximate primary mechanism $\mathcal{M}_i$ and a corresponding concave primary benchmark R_i. Suppose all of the assumptions mentioned earlier hold. Then, for any $\gamma \in [0, \gamma_k]$, the γ-post-rounding mechanism (Mech. 2) is a feasible mechanism in $\mathbb{M}$. Furthermore, it is a $\gamma \cdot \alpha$-approximation of the optimal mechanism in $\mathbb{M}$. Note that the resulting mechanism is only Bayesian incentive compatible (BIC).*

[10]note that budget constraints are not part of the valuations.

5. AN EXAMPLE PRIMARY MECHANISM

In this section, we present an example of a primary mechanism for the following setting. We consider a buyer with correlated and additive valuations with a capacity and a hard budget constraint. We assume that the size of the type space of the buyer is polynomially bounded. Since a primary mechanism only interacts with a single buyer, we shall drop the subscript i. For each possible type t of the buyer, let v_{tj} denote her valuation for item j. Also let $f(t)$ denote the probability that the buyer's type is t. Furthermore, suppose that the buyer has a total budget of B and is interested in at most C items. We assume that the only private information of the buyer is her type and everything else is publicly known. Note that this is exactly the setting considered in [3]. They present a $\frac{1}{4}$-approximate BIC mechanism for maximizing revenue. Next, we present a truthful optimal primary mechanism for maximizing revenue, which can be converted to a γ_k-approximate BIC mechanism for multiple buyers, using the γ-post-rounding. Remember that γ_k is at least $\frac{1}{2}$ and approaches 1 as k increases.

Consider the following linear program in which x_{tj} is the variable corresponding to the probability of allocating item j when the buyer has reported type t and p_t is the variable for the payment for type t. The optimal objective value of this LP is an upper bound on the revenue of the optimal primary mechanism when restricted to allocate each item j with probability at most $\bar{q}_j$:

$$
\begin{array}{lcr}
\text{maximize:} & \displaystyle\sum_t f(t) p_t & (LP_{rev}) \\
\forall j \in [m]: & \displaystyle\sum_t f(t) x_{tj} \le \bar{q}_j & \\
\forall t \in T: & \displaystyle\sum_j x_{tj} \le C & \\
\forall t, t' \in T: & \displaystyle\sum_j v_{tj} x_{tj} - p_t \ge \sum_j v_{tj} x_{t'j} - p_{t'} & \\
\forall t \in T, \forall j \in [m]: & x_{tj} \in [0,1] & \\
\forall t \in T: & p_t \in [0, B] &
\end{array}
$$

We construct the optimal primary mechanism as follows.

Mechanism 3.

- *Define the optimal primary benchmark $R(\bar{q})$ to be the optimal objective value of (LP_{rev}) as a function of $\bar{q} = (\bar{q}_1, \cdots, \bar{q}_m)$.*
- *Given $\bar{q}_1, \cdots, \bar{q}_m$, solve the linear program of (LP_{rev}) to compute x_{tj} and p_t.*
- *If the buyer reports her type as t then charge her a payment of p_t and allocate each item j*

with probability x_{tj} as explained next. Use the dependent randomize rounding algorithm of [10] to round each x_{tj} to either 0 or 1 such that if X_j is the result of rounding the x_{tj} then $E[X_j] = x_{tj}$ and such that $\sum_j X_j \leq C$. Then, for each j allocate a copy of item j to the buyer iff $X_j = 1$.

Theorem 9. *The primary mechanism Mech. 3 is a truthful optimal primary mechanism for maximizing revenue and satisfies all the requirements of γ-post-rounding.*

Proof: The proof of truthfulness and optimality of Mech. 3 trivially follows from the (LP_{rev}). So, we only focus on proving that this mechanism satisfies the requirements of Theorem 8. First, observe that the benchmark function, $R(\bar{q})$, is concave (this follows by applying Lem. 1). Second, observe that the valuations of the buyer can be represented as a weighted rank function of a uniform matroid of rank C. Third, notice that given $\bar{q}_1, \cdots, \bar{q}_m$, we can compute the exact marginal probabilities of allocation, i.e. $\hat{q}_1, \cdots, \hat{q}_m$ as follows: $\hat{q}_j = \sum_t f(t) x_{tj}$. So the mechanism Mech. 3 and its associated benchmark satisfy the requirements of Theorem 8 for γ-post-rounding. ∎

Lemma 1. *Consider any convex program of the following form, in which $u(\cdot)$ is a concave function, $g_j(\cdot)$ are convex functions, and $\mathbb{X}$ is a convex set. Let $R(\bar{q})$ denote the optimal objective value of this program as a function of $\bar{q} = (\bar{q}_1, \cdots, \bar{q}_m)$. Then $R(\bar{q})$ is concave.*

$$\begin{array}{lcr} \textit{maximize:} & u(x) & (CP_u) \\ \forall j: & g_j(x) \leq \bar{q}_j & \\ & x \in \mathbb{X} & \end{array}$$

Remark 2. *Observe that if we replace the objective function of (LP_{rev}) with $\sum_{t,j} f(t) v_{tj} x_{tj}$ we get a truthful optimal primary mechanism for maximizing welfare instead of revenue, which can be converted to a γ_k-approximate BIC mechanism for multiple buyers and for maximizing welfare.*

6. Conclusion

In this paper, for Bayesian combinatorial auctions, we presented an approximate reduction from n-buyer mechanisms to 1-buyer mechanisms. This shows that the inherent difficulty of designing Bayesian mechanisms in multidimensional settings does not stem from the problem of making coordinated decisions for all buyer, but instead it stems from the difficulty of aligning the behavior of the mechanism with the incentives of each individual buyer, even in the absence of other buyers.

7. Acknowledgment

I thank Jason Hartline for many helpful suggestions and comments. I also thank Azarakhsh Malekian for a fruitful discussion that lead to the proof of the sand theorem (included in the full version).

References

[1] M. Armstrong, "Multiproduct nonlinear pricing," *Econometrica*, vol. 64, no. 1, pp. 51–75, January 1996. [Online]. Available: http://ideas.repec.org/a/ecm/emetrp/v64y1996i1p51-75.html

[2] M. Babaioff, R. Lavi, and E. Pavlov, "Single-value combinatorial auctions and implementation in undominated strategies," in *SODA*, 2006, pp. 1054–1063.

[3] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala, "Budget constrained auctions with heterogeneous items," in *STOC*, 2010, pp. 379–388.

[4] L. Blumrosen and T. Holenstein, "Posted prices vs. negotiations: an asymptotic analysis," in *ACM Conference on Electronic Commerce*, 2008, p. 49.

[5] J. Bulow and J. Roberts, "The simple economics of optimal auctions," *Journal of Political Economy*, vol. 97, no. 5, pp. 1060–90, October 1989. [Online]. Available: http://ideas.repec.org/a/ucp/jpolec/v97y1989i5p1060-90.html

[6] T. Chakraborty, E. Even-Dar, S. Guha, Y. Mansour, and S. Muthukrishnan, "Approximation schemes for sequential posted pricing in multi-unit auctions," in *WINE*, 2010.

[7] S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan, "Multi-parameter mechanism design and sequential posted pricing," in *STOC*, 2010, pp. 311–320.

[8] S. Chawla, D. Malec, and A. Malekian, "Bayesian mechanism design for budget-constrained agents," in *ACM Conference on Electronic Commerce*, 2011.

[9] P. Dhangwatnotai, T. Roughgarden, and Q. Yan, "Revenue maximization with a single sample," in *ACM Conference on Electronic Commerce*, 2010, pp. 129–138.

[10] R. Gandhi, S. Khuller, S. P. 0002, and A. Srinivasan, "Dependent rounding in bipartite graphs," in *FOCS*, 2002, pp. 323–332.

[11] M. T. Hajiaghayi, R. D. Kleinberg, and T. Sandholm, "Automated online mechanism design and prophet inequalities," in *AAAI*, 2007, pp. 58–65.

[12] J. D. Hartline, R. Kleinberg, and A. Malekian, "Bayesian incentive compatibility and matchings," in *SODA*, 2011.

[13] J. D. Hartline and B. Lucier, "Bayesian algorithmic mechanism design," in *STOC*, 2010, pp. 301–310.

[14] J. D. Hartline and T. Roughgarden, "Simple versus optimal mechanisms," in *ACM Conference on Electronic Commerce*, 2009, pp. 225–234.

[15] T. P. Hill and R. P. Kertz, "A survey of prophet inequalities in optimal stopping theory," in *Contemporary Mathematics*, 1992.

[16] R. B. Myerson, "Optimal auction design," 1981.

[17] Q. Yan, "Mechanism design via correlation gap," in *SODA*, 2011.