THEORY OF COMPUTING, Volume 8 (2012), pp. 429–460
www.theoryofcomputing.org

SPECIAL ISSUE IN HONOR OF RAJEEV MOTWANI

Budget-Constrained Auctions with Heterogeneous Items*

Sayan Bhattacharya† Gagan Goel Sreenivas Gollapudi
Kamesh Munagala‡

Received: August 18, 2010; published: September 4, 2012.

Abstract: We present the first approximation algorithms for designing revenue-optimal incentive-compatible mechanisms in the following setting: There are multiple (heterogeneous) items, and bidders have arbitrary demand and budget constraints (and additive valuations). Furthermore, the type of a bidder (which specifies her valuations for each item) is private knowledge, and the types of different bidders are drawn from publicly known mutually independent distributions. Our mechanisms are surprisingly simple.

First, we assume that the type of each bidder is drawn from a discrete distribution with polynomially bounded support size. This restriction on the type-distribution, however, allows the random variables corresponding to a bidder's valuations for different items to be arbitrarily correlated. In this model, we describe a sequential all-pay mechanism that is truthful in expectation and Bayesian incentive compatible. The outcome of our all-pay

*A preliminary version of this paper appeared in STOC 2010.

†Supported by NSF grants CCF-0745761 and CCF-1008065.

‡Supported by an Alfred P. Sloan Research Fellowship, a gift from Cisco, and by NSF grants CCF-0745761, CCF-1008065, and IIS-0964560.

ACM Classification: G.1.6, J.4

AMS Classification: 68W25

Key words and phrases: mechanism design, approximation algorithms

DOI: 10.4086/toc.2012.v008a020

mechanism can be computed in polynomial time, and its revenue is a 4-approximation to the revenue of the optimal truthful-in-expectation Bayesian incentive-compatible mechanism.

Next, we assume that the valuations of each bidder for different items are drawn from mutually independent discrete distributions satisfying the monotone hazard-rate condition. In this model, we present a sequential posted-price mechanism that is universally truthful and incentive compatible in dominant strategies. The outcome of the mechanism is computable in polynomial time, and its revenue is a $O(1)$-approximation to the revenue of the optimal truthful-in-expectation Bayesian incentive-compatible mechanism. If the monotone hazard-rate condition is removed, then we show a logarithmic approximation, and we complete the picture by proving that no sequential posted-price scheme can achieve a sub-logarithmic approximation. Finally, if the distributions are regular, and if the space of mechanisms is restricted to sequential posted-price schemes, then we show that there is a $O(1)$-approximation within this space. Our results are based on formulating novel LP relaxations for these problems, and developing generic rounding schemes from first principles.

1 Introduction

In several scenarios, such as the Google TV ad mechanism [31] and the FCC spectrum mechanisms [11], where mechanisms have been applied in the recent past, bidders are constrained by the amount of money they can spend. This leads to the study of mechanisms with budget-constrained bidders, which is the focus of this paper. The key difficulty with budgets is that the utility of a bidder is equal to her valuation minus price *if and only if* the price is below the budget constraint, and the utility is $-\infty$ whenever the price exceeds her budget. As a consequence, well-known mechanisms such as the VCG mechanism [17, 23, 34] are no longer directly applicable. Before proceeding further, we formally define our model.

1.1 Our model

There are n bidders and m heterogeneous items. The *type* of a bidder, which is her private knowledge, is an m-tuple representing her valuations for each item. Every bidder i has two publicly known constraints: A *demand constraint* d_i on the maximum number of items she is willing to buy, and a *budget constraint* B_i on the maximum total price she can afford to pay. In a *mechanism*, the bidders report their types to the auctioneer, and based on the reported types, the auctioneer computes an allocation of the items and payments. The utility of bidder i is defined as follows: Suppose that she gets a subset A of items where $|A| \leq d_i$, and pays a total price P_i. Let v_{ij} denote the valuation of bidder i for item j. If $P_i \leq B_i$, then the utility of bidder i is equal to $\sum_{j \in A} v_{ij} - P_i$. In contrast, if $P_i > B_i$, then her utility is $-\infty$. The revenue of the auctioneer is given by $\sum_i P_i$. The mechanism should be *incentive compatible* in that no bidder gains in utility by misreporting her type, and *individually rational*, meaning that a bidder gets nonnegative utility if she reports her true type.

There are two well-established ways of proceeding from here. In the adversarial setting, no assumptions are made on the types of the bidders, while in the Bayesian setting, it is assumed that the bidders' private types are drawn from mutually independent (but not necessarily identical) publicly known prior distributions. We take the latter Bayesian approach, and our goal is to design an incentive-compatible and individually-rational mechanism that (approximately) maximizes the expected revenue of the auctioneer. This line of research was pioneered by Myerson [30].

1.2 Preliminaries

We first distinguish between four kinds of mechanisms that satisfy the incentive-compatibility (and individual-rationality) constraints.

Dominant-strategy incentive-compatible (DSIC) Fix any bidder i. Suppose that her true type is $\mathbf{t}_i$.

- In a *universally truthful DSIC mechanism*, the utility of bidder i is maximized (at a nonnegative value) when she reveals her true type $\mathbf{t}_i$, regardless of the types reported by other bidders and the random choices made by the mechanism.
- In a *truthful-in-expectation DSIC mechanism*, the expected utility of bidder i (the expectation is over the random choices made by the mechanism) is maximized (at a nonnegative value) when she reveals her true type $\mathbf{t}_i$, regardless of the types reported by other bidders.

Bayesian incentive-compatible (BIC) Fix any bidder i and suppose that her true type is $\mathbf{t}_i$.

- In a *universally truthful BIC mechanism*, the expected utility of bidder i (the expectation is over the prior distributions of the types of the other bidders) is maximized (at a nonnegative value) when she reveals her true type $\mathbf{t}_i$, regardless of the random choices made by the mechanism.
- In a *truthful-in-expectation BIC mechanism*, the expected utility of bidder i (the expectation is over the prior distributions of the types of the other bidders and the random choices made by the mechanism) is maximized (at a nonnegative value) when she reveals her true type $\mathbf{t}_i$.

Note that if a mechanism is universally truthful DSIC, then the same mechanism satisfies all of the other three notions of incentive compatibility. On the other hand, the union of all the universally truthful DSIC, truthful-in-expectation DSIC, and universally truthful BIC mechanisms is a subset of the set of all truthful-in-expectation BIC mechanisms. Hence, among the four notions described above, universally truthful DSIC (respectively, truthful-in-expectation BIC) is the strongest (respectively, weakest) notion of incentive compatibility.

A standard assumption in Economics is that a bidder's valuation for an item is drawn from a distribution satisfying some useful properties. Specifically, we will be interested in the following classes of distributions.

Definition 1.1. Suppose that the valuation of a bidder i for item j is a discrete random variable v_{ij} with integral support $\{1,\ldots,L_{ij}\}$. The distribution of v_{ij} is *regular* if and only if

$$r - \frac{\Pr[v_{ij} > r]}{\Pr[v_{ij} = r]}$$

is a non-decreasing function of $r \in \{1,\ldots,L_{ij}\}$.

Definition 1.2. Suppose that the valuation of a bidder i for item j is a discrete random variable v_{ij} with integral support $\{1,\ldots,L_{ij}\}$. The distribution of v_{ij} is *monotone hazard rate* if and only if

$$\frac{\Pr[v_{ij} > r]}{\Pr[v_{ij} = r]}$$

is a non-increasing function of $r \in \{1,\ldots,L_{ij}\}$.

Note that if a distribution is monotone hazard rate, then it is regular. Examples of monotone hazard-rate distributions include geometric distributions and uniform distributions. In contrast, if we have $\Pr[v_{ij} \geq r] = 1/r$ for $r = 1,2,\ldots,L_{ij}$, then the distribution of v_{ij} is regular but not monotone hazard rate. Finally, if we have a bimodal distribution, such as $\Pr[v_{ij} = 1] = \Pr[v_{ij} = 3] = 4/9$ and $\Pr[v_{ij} = 2] = 1/9$, then it is easy to check that the random variable v_{ij} does not follow a regular distribution.

1.3 Our results

Recall that the type of a bidder is an m-tuple representing her valuations for each item, and the types of different bidders follow mutually independent public distributions. From the type-distribution of a bidder i, we can determine the distribution of the random variable v_{ij}, which denotes the valuation of bidder i for item j. Our results depend on whether the random variables $\{v_{ij}\}_j$ (denoting the valuations of the same bidder for different items) are correlated or mutually independent. All our results can be viewed as presenting simple characterizations of approximately revenue-optimal mechanisms in these contexts.

1.3.1 Our result in Section 2

We consider the following scenario in Section 2: For every bidder i, the random variables $\{v_{ij}\}_j$ (denoting the valuations of bidder i for different items) can be arbitrarily correlated. However, the type of bidder i is drawn from a discrete distribution with polynomially bounded support size.

Truthful-in-expectation BIC mechanism We present a simple *all-pay* mechanism whose outcome is computable in time polynomial in the input size. The mechanism charges each bidder a fixed price that depends only on her revealed type, while the allocation made to the bidder depends on the reported types of other bidders and the random choices made by the mechanism. The resulting scheme is truthful-in-expectation BIC, and we show that its revenue is a 4-approximation to the revenue of the optimal truthful-in-expectation BIC mechanism (Theorem 2.2).

Inapproximability of the optimal universally truthful DSIC mechanism An all-pay mechanism is unrealistic in several situations, since a bidder is forced to participate even if she obtains a negative utility when the mechanism concludes. A natural question to ask is whether we can compute a universally truthful DSIC mechanism with good revenue properties. We can show that this problem generalizes the problem of unlimited supply unit-demand profit-maximizing envy-free pricing [24], as described below.

Consider a single bidder and m items, and suppose that the bidder's type is drawn from a uniform distribution over the discrete set $\mathcal{T} = \{\mathbf{t}^{(1)},\ldots,\mathbf{t}^{(n)}\}$ of size n, and that the bidder has unit demand and infinite budget. In this setting, the optimal universally truthful DSIC mechanism will post a price for each item, and it will allow the bidder to pick the item that gives her maximum (nonnegative) utility. This is exactly equivalent to the following instance of the unlimited supply unit-demand envy-free pricing problem [24]: We have m items and n agents, and the valuation profile of each agent $i \in \{1, \dots, n\}$ (which gives her valuations for each item) is specified by the type $\mathbf{t}^{(i)}$.

For the unlimited supply unit-demand envy-free pricing problem, the best known polynomial-time algorithm gives a logarithmic approximation, and there is strong evidence that a better polynomial-time approximation is not possible [9]. Consequently, it is highly unlikely that we will be able to design a polynomial-time computable and universally truthful DSIC mechanism with good revenue properties. In Section 3, we impose further restrictions on the type-distributions to circumvent this negative result.

1.3.2 Our results in Section 3

We make the following assumption throughout Section 3: For every bidder i, the random variables $\{v_{ij}\}_j$ (denoting the valuations of bidder i for different items) are mutually independent. However, in contrast to Section 2, we no longer require that the type-distribution of a bidder should have polynomial support.

The assumption mentioned above has been used previously in the literature. The work of Chawla et al. [12] considers the special case of a single bidder with unit-demand and infinite budget, whose valuations for different items are drawn from mutually independent public distributions. For this problem, Chawla et al. give a constant factor approximation to the revenue maximizing universally truthful DSIC mechanism, by deriving an elegant connection to Myerson's mechanism [30]. Independently of our work, Chawla et al. [13] extend the previous result to a setting with multiple bidders. However, this result also crucially requires the unit-demand assumption.

Monotone hazard rates and universally truthful DSIC mechanism In Section 3.2, we further assume that for every bidder i and item j, the random variable v_{ij} (denoting the valuation of bidder i for item j) is drawn from a distribution that satisfies the monotone hazard-rate (Definition 1.2) condition. We give a universally truthful DSIC mechanism, whose outcome is computable in polynomial time, and whose revenue is a constant-factor approximation to the revenue of the optimal truthful-in-expectation BIC mechanism. Our mechanism is a *Sequential Posted Price* scheme. Any sequential posted-price scheme has the following simple structure: The auctioneer considers the bidders sequentially in arbitrary order, and each bidder is offered a subset of the available items, so that each item in the subset has to be purchased at a pre-computed price, and the bidder herself picks the items she wants to buy under these prices. Hence, we get a constant-factor gap between the revenues of the optimal truthful-in-expectation BIC mechanism and the optimal universally truthful DSIC mechanism (Theorem 3.14), which is in sharp contrast with the corresponding negative result [10] when the valuations of a bidder for different items are drawn from correlated distributions.

Regular distributions and adaptive posted-price mechanisms In Section 3.3, we show that the monotone hazard-rate condition is indeed necessary if we want to design a sequential posted-price mechanism with good revenue properties. Suppose that the monotone hazard-rate condition is slightly relaxed, and we consider the scenario where there is a single bidder and her valuations for different items are drawn from mutually independent regular distributions (Definition 1.1). In this case, the optimal universally truthful DSIC mechanism will have a logarithmic gap against the revenue of the optimal sequential posted-price scheme. We prove that this gap is tight by showing the existence of a sequential posted-price mechanism achieving this approximation ratio (Theorem 3.16). On a positive note, we prove that for regular distributions, if the space of feasible mechanisms is restricted to those that consider the bidders in some adaptive order and post prices for the items that may depend on the outcomes so far, then there is a $O(1)$-approximation within this space that considers the bidders in an arbitrary but fixed order, and pre-computes the posted prices (Theorem 3.21).

1.4 Related work

The Bayesian setting is widely studied in the economics literature [5, 11, 14, 15, 27, 29, 32, 33, 35]. In this setting, the optimal mechanism can always be computed by encoding the incentive-compatibility constraints in an integer program and maximizing expected revenue. However, the number of variables (and constraints) in this IP is exponential in the number of bidders, as there are variables for the allocations and prices for each scenario of revealed types. Therefore, the key difficulty in the Bayesian mechanism design case is *computational*: *Can the optimal (or approximately optimal) mechanism be efficiently computed and implemented?*

Much of the literature in economics considers the case where the auctioneer has one item (or multiple copies of one item). In the absence of budget constraints, Myerson [30] presents the characterization of any BIC mechanism in terms of expected allocation made to a bidder: This allocation must be monotone in the revealed valuation of the bidder. This yields a linear-time computable optimal revenue-maximizing mechanism that is both BIC and DSIC. The key issue with budget constraints is that the allocations need to be thresholded in order for the prices to be below the budgets [14, 27, 32]. However, even in this case, the optimal BIC mechanism follows from a polymatroid characterization that can be solved by the Ellipsoid algorithm and an all-pay condition [7, 32]. By *all-pay*, we mean that the bidder pays a fixed amount given his revealed type, regardless of the allocation made. This also yields a DSIC mechanism that is an $O(1)$-approximation to the optimal BIC revenue [6], but the result holds only for homogeneous items.

An alternative line of work deals with the *adversarial* setting, where no distributional assumption is made on the bidders' private valuations. In this setting, the budget-constrained mechanism problem is notorious, mainly because standard mechanism concepts such as VCG, efficiency, and competitive equilibria do not directly carry over [31]. Most previous results deal with the case of multiple units of a homogeneous good. In this setting, based on the random partitioning framework of Goldberg et al. [22, 21], Borgs et al. [8] presented an incentive-compatible mechanism whose revenue is asymptotically within a constant factor of the optimal revenue (see also [1]). Furthermore, Borgs et al. showed that no incentive-compatible mechanism can (approximately) maximize social welfare [8]. Consequently, the focus has been on weaker notions than efficiency, such as Pareto-optimality, where no pair of agents (including the auctioneer) can simultaneously improve their utilities by trading with each other. Dobzinski et al. [18] present an ascending price mechanism based on the clinching mechanism of Ausubel [3], which they show to be the only Pareto-optimal mechanism in the public budget setting. This result was extended to the private budget setting by Bhattacharya et al. [6]; see [25] for a related result.

Finally, several researchers have considered restricted scenarios. The first type of restriction is on the type of mechanism; examples include mechanisms that are sequential by item and second price within each item [5, 20], and ascending price mechanisms [4, 11]. The goal here is to analyze the improvement in revenue (or social welfare) by optimal sequencing, or to study incentive compatibility of commonly used ascending price mechanisms. However, analyzing the performance of sequential or ascending price mechanisms is difficult in general, and there is little known in terms of optimal mechanisms (or even approximately optimal mechanisms) in these settings. A different type of restriction is on the valuations of each bidder. The most well-known instance is keyword search mechanisms [19], where the items are ad slots, and the bidder valuations are typically single-dimensional, with the valuations for each slot being scaled down by a publicly known value. However, these mechanisms have typically been considered without budget constraints, and the goal here has been to analyze mechanisms used in practice, such as the generalized second-price mechanism [19].

1.5 Our techniques

If, for every bidder, the valuations for different items are drawn from correlated distributions (Section 2), then the optimal truthful-in-expectation BIC revenue can be bounded from above by a linear program (LP1) that requires the incentive-compatibility, individual-rationality, supply, and demand constraints to hold *only in expectation*. We construct a truthful-in-expectation BIC all-pay mechanism (Figure 1) that basically implements a rounding scheme on the optimal solution to LP1, losing a constant factor in revenue (Theorem 2.2). This approach is based on the techniques used in [6].

As mentioned before, Chawla et al. [12] consider the Bayesian unit-demand pricing problem. There are m heterogeneous items, a single bidder with unit demand, and her valuations (v_j for item $j \in [1,\dots,m]$) are drawn from independent distributions. They present an elegant pricing scheme that is a constant approximation to the optimal revenue by upper bounding it using the revenue of Myerson's mechanism in the following setting: There is a single item, m bidders, and the valuation of each bidder j follows the same distribution as that of v_j. However, this technique cannot be applied if the unit demand assumption is removed.

In contrast, our approach in Section 3 (where the valuations of a bidder for different items are drawn from mutually independent distributions) does *not* require the unit-demand assumption, and is based on a novel LP relaxation (LPREV) for the problem (Lemma 3.3). Unlike the LP relaxation of Section 2, and perhaps surprisingly, LPREV does not encode any incentive-compatibility constraints, and our universally truthful DSIC mechanism (Figure 2), which competes against this LP, is in fact a constant approximation to the optimal truthful-in-expectation BIC revenue. One limitation of our approach is that we have to (necessarily) assume that the bidders' valuations are drawn from montone hazard rate distributions (Definition 1.2). In the process of proving our main result (Theorem 3.14), we describe a crucial property of monotone hazard rate distributions (Lemma 3.9) that can be used to extend the type of results shown in [26]. For example, in multi-item settings with only demand constraints, *posted-price schemes generate revenue that is a constant factor of the optimal social welfare*, assuming monotone hazard rate distributions (Corollary 3.15). The LP formulations also generalize the stochastic matching setting in Chen et al. [16].

2 Truthful-in-expectation BIC mechanism

In this section, we consider the problem of approximating the optimal Bayesian incentive-compatible mechanism. We show an all-pay mechanism that is a 4-approximation to optimal revenue. Our solution techniques are inspired by the LP relaxation and rounding scheme in [6] for multi-unit mechanisms.

2.1 Notations

There is a set $\mathcal{I}$ of n bidders, and a set $\mathcal{J}$ of m indivisible items. A *type* $\mathbf{t}$ is an m-tuple $\langle \mathbf{t}(1), \mathbf{t}(2), \ldots, \mathbf{t}(m) \rangle$. If a bidder $i \in \mathcal{I}$ has type $\mathbf{t}_i$, then her valuation for item $j \in \mathcal{J}$ is given by $v_{ij} = \mathbf{t}_i(j)$. Every bidder $i \in \mathcal{I}$ has a *demand* $d_i \geq 1$, which upper bounds the number of items that can be allocated to her, and a *budget* B_i, which specifies the maximum total payment she can make. Both her demand and budget are publicly known. Suppose that she gets a subset $A \subseteq \mathcal{J}$ of items, where $|A| \leq d_i$, and is charged a price of P. Her utility $u_i(A,P)$ is given by the following expression.

$$u_i(A,P) = \begin{cases} -\infty & \text{if } P > B_i\,, \\ \sum_{j \in A} v_{ij} - P & \text{if } P \leq B_i\,. \end{cases}$$

The type of a bidder $i \in \mathcal{I}$ is private knowledge. Furthermore, it is drawn from a *discrete* probability distribution $f_i(\cdot)$ with support $\mathcal{T}_i \subseteq \mathbf{R}^m$. For all $\mathbf{t}_i \in \mathcal{T}_i$, we have $f_i(\mathbf{t}_i) = \Pr[\text{type of bidder } i = \mathbf{t}_i]$, and $\sum_{\mathbf{t}_i \in \mathcal{T}_i} f_i(\mathbf{t}_i) = 1$. The distributions $f_1(\cdot), \ldots, f_n(\cdot)$ are mutually independent and publicly known. The notation $f_{ij}(\cdot)$ denotes the marginal distribution of the valuation of bidder $i \in \mathcal{I}$ for item $j \in \mathcal{J}$:

$$f_{ij}(v) \;=\; \Pr[v_{ij} = v] \;=\; \sum_{\mathbf{t}_i \in \mathcal{T}_i : \mathbf{t}_i(j) = v} f_i(\mathbf{t}_i)\,.$$

The distribution $f_{ij}(\cdot)$ has support $\mathcal{T}_{ij}$, that is, $\mathcal{T}_{ij} = \{v \in \mathbf{R} : f_{ij}(v) > 0\}$. Note that a bidder's valuations for different items can be correlated. Specifically, we may have

$$\Pr[v_{ij} = v \mid v_{ij'} = v'] \;\neq\; f_{ij}(v) \quad \text{for some bidder } i, \text{ items } j \neq j, \text{ and } v \in \mathcal{T}_{ij}, v' \in \mathcal{T}_{ij'}\,.$$

In a mechanism, the bidders first report their types. Based on these reported types, the auctioneer computes an allocation of items and payments. The price charged to a bidder i should not exceed her budget B_i, and the number of items allocated to bidder i should be at most d_i.

2.2 The problem and the LP-relaxation

We want to find an mechanism that is incentive compatible (and individually rational) in the following sense: Fix any bidder i and suppose that her (private) type is $\mathbf{t}_i$. Her expected utility, where the expectation is over the distributions of types of other bidders and the random choices made by the mechanism, is maximized (at a nonnegative value) if she reveals her true type $\mathbf{t}_i$. We are interested in a mechanism that (approximately) maximizes the expected revenue, and can be computed in time polynomial in the input size, i. e., in n, m, and $\max_{i \in \mathcal{I}}\{|\mathcal{T}_i|\}$. Throughout the rest of Section 2, we will make the following assumption.

Assumption 2.1. *The distributions $f_1(\cdot), \ldots, f_n(\cdot)$ have polynomial supports.*

Linear programming relaxation For any feasible mechanism, let $x_{ij}(\mathbf{t}_i)$ denote the probability that bidder i obtains item j if her reported type is $\mathbf{t}_i$. Let $P_i(\mathbf{t}_i)$ denote the expected price paid by bidder i when she reports type $\mathbf{t}_i$. We have the following LP.

$$\text{Maximize} \quad \sum_{i\in\mathcal{I}} \sum_{\mathbf{t}_i\in\mathcal{T}_i} f_i(\mathbf{t}_i)P_i(\mathbf{t}_i) \qquad \text{(LP1)}$$

$$\begin{array}{rcll}
\sum_{i\in\mathcal{I}}\sum_{\mathbf{t}_i\in\mathcal{T}_i} f_i(\mathbf{t}_i)x_{ij}(\mathbf{t}_i) & \leq & 1 & \forall j\in\mathcal{J} \\
\sum_{j\in\mathcal{J}} x_{ij}(\mathbf{t}_i) & \leq & d_i & \forall i\in\mathcal{I}, \mathbf{t}_i\in\mathcal{T}_i \\
\sum_{j\in\mathcal{J}} \mathbf{t}_i(j)x_{ij}(\mathbf{t}_i) - P_i(\mathbf{t}_i) & \geq & \sum_{j\in\mathcal{J}} \mathbf{t}_i(j)x_{ij}(\mathbf{t}'_i) - P_i(\mathbf{t}'_i) & \forall i\in\mathcal{I}, \mathbf{t}_i, \mathbf{t}'_i\in\mathcal{T}_i \\
\sum_{j\in\mathcal{J}} \mathbf{t}_i(j)x_{ij}(\mathbf{t}_i) - P_i(\mathbf{t}_i) & \geq & 0 & \forall i\in\mathcal{I}, \mathbf{t}_i\in\mathcal{T}_i \\
x_{ij}(\mathbf{t}_i) & \in & [0,1] & \forall i\in\mathcal{I}, j\in\mathcal{J}, t\in\mathcal{T}_i \\
P_i(\mathbf{t}_i) & \in & [0,B_i] & \forall i\in\mathcal{I}, \mathbf{t}_i\in\mathcal{T}_i
\end{array}$$

The optimal truthful-in-expectation BIC mechanism is feasible for the above constraints. The first constraint encodes the fact that, in expectation, each item is assigned at most once. Now fix any bidder $i\in\mathcal{I}$, and suppose that her true type is $\mathbf{t}_i\in\mathcal{T}_i$. The second constraint encodes the demand: Bidder i can get at most d_i items in expectation. The third constraint encodes Bayesian incentive-compatibility: The expected utility of bidder i, when she reports any false type $\mathbf{t}'_i$, cannot be greater than her expected utility when she reports her true type $\mathbf{t}_i$. The fourth constraint encodes individual rationality: If bidder i reports her true type $\mathbf{t}_i$, then her expected utility is nonnegative. Therefore, the LP1 value is an upper bound on the expected revenue.

Remark Fix any bidder $i\in\mathcal{I}$. Note that the first constraint in LP1 holds in expectation over *both* her own type and the types of other bidders. In contrast, the next three constraints are enforced for every possible type $\mathbf{t}_i\in\mathcal{T}_i$, and in these constraints the expectations are taken *only* over the types of other bidders.

2.3 The all-pay mechanism

Suppose that the optimal solution to LP1 assigns a value of $x^*_{ij}(\mathbf{t}_i)$ (respectively, $P^*_i(\mathbf{t}_i)$) to each variable $x_{ij}(\mathbf{t}_i)$ (respectively, $P_i(\mathbf{t}_i)$). We design an all-pay mechanism (Figure 1). The key observation is that the mechanism satisfies the following property: Fix any bidder $i\in\mathcal{I}$ and her reported type $\mathbf{t}^*_i\in\mathcal{T}_i$. Furthermore, suppose that every other bidder $i'\in\mathcal{I}\setminus\{i\}$ reveals her true type. In this case, the probability that bidder i gets any item $j\in\mathcal{J}$, over the randomness introduced by the mechanism and the distributions of types of other bidders, is equal to $x^*_{ij}(\mathbf{t}^*_i)/4$; bidder i is charged a fixed payment of $P^*_{ij}(\mathbf{t}^*_i)/4$. Since both the expected allocation and the payment are scaled down by *exactly* the same factor [2, 28], this scheme preserves the Bayesian incentive-compatibility and individual-rationality conditions enforced by the constraints in LP1. Finally, note that the auctioneer's expected revenue is given by the expression

$$\sum_{i\in\mathcal{I}}\sum_{\mathbf{t}_i\in\mathcal{T}_i} \frac{f_i(\mathbf{t}_i)\cdot P^*_i(\mathbf{t}_i)}{4},$$

All-Pay Mechanism

1. Collect the reported types of all bidders. Denote the reported type of bidder i by $\mathbf{t}_i^*$.
2. Find the optimal solution to LP1, and denote the variable values by $\{x_{ij}^*(\mathbf{t}_i), P_i^*(\mathbf{t}_i)\}$.
3. Choose an arbitrary but fixed ordering of all bidders and denote it by $1, 2, \ldots, n$.
4. For all bidders $i = 1, 2, \ldots, n$, items $j \in \mathcal{J}$ and types $\mathbf{t}_i \in \mathcal{T}_i$, let

   $$\tilde{x}_{ij}(\mathbf{t}_i) = x_{ij}^*(\mathbf{t}_i)/2\,, \quad X_{ij} = \textstyle\sum_{\mathbf{t}_i \in \mathcal{T}_i} f_i(\mathbf{t}_i)\tilde{x}_{ij}(\mathbf{t}_i)\,, \quad Z_{ij} = \prod_{i'=1}^{i-1}(1 - X_{i'j})\,.$$

5. FOR all bidders i and types $\mathbf{t}_i \in \mathcal{T}_i$,
   Partition the set of items $\mathcal{J}$ into d_i disjoint groups such that
   in each group $\mathcal{G}(i, \mathbf{t}_i, k)$, we have $\sum_{j \in \mathcal{G}(i,\mathbf{t}_i,k)} \tilde{x}_{ij}(\mathbf{t}_i) \leq 1$.
6. FOR $i = 1, 2, \ldots n$,
   Initialize $W_i, S_i \leftarrow \emptyset$;
   FOR $k = 1$ to d_i
   Pick a single item $j \in \mathcal{G}(i, \mathbf{t}_i^*, k)$ with probability $\tilde{x}_{ij}(\mathbf{t}_i^*)$;
   $S_i \leftarrow S_i \cup \{j\}$;
   $Q_i \leftarrow S_i \setminus \cup_{i'=1}^{i-1} S_{i'}$;
   FOR ALL items $j \in Q_i$
   $W_i \leftarrow W_i \cup \{j\}$ with probability $1/(2Z_{ij})$;
   Bidder i gets the (random) set W_i, and pays a (fixed) price $P_i^*(\mathbf{t}_i^*)/4$.

Figure 1: BIC Mechanism for correlated valuations

which is $1/4$ times the optimal objective value of LP1. Hence, the mechanism gives a 4-approximation to optimal revenue.

In Step 1 of the All-Pay Mechanism (Figure 1), we ask the bidders to reveal their types. The reported type of bidder $i \in \mathcal{I}$ is denoted by $\mathbf{t}_i^*$. In Step 2, we solve (LP1). In Step 3, we order the bidders as $1, 2, \ldots, n$. In Step 4, we scale down the values of the allocation variables by a factor of 2, and introduce the notations $\tilde{x}_{ij}(\mathbf{t}_i)$, X_{ij} and Z_{ij}. Applying the second constraint in LP1, we get $\sum_{j \in \mathcal{J}} \tilde{x}_{ij}(\mathbf{t}_i) \leq d_i/2$, for all bidders $i \in \mathcal{I}$ and types $\mathbf{t}_i \in \mathcal{T}_i$. In Step 5, we exploit this property while partitioning the set of items $\mathcal{J}$ into d_i disjoint groups. Specifically, we employ the following greedy strategy.

- Initialize $V \leftarrow \mathcal{J}$ and $k \leftarrow 1$.
- REPEAT UNTIL $V = \emptyset$
  - Pack maximal number of items from V into the group $\mathcal{G}(i, \mathbf{t}_i, k)$, so that we have

    $$1/2 \leq \sum_{j \in \mathcal{G}(i,\mathbf{t}_i,k)} \tilde{x}_{ij}(\mathbf{t}_i) \leq 1\,.$$

  - Set $V \leftarrow V \setminus \mathcal{G}(i, \mathbf{t}_i, k)$ and $k \leftarrow k + 1$.

In Step 6, we visit the bidders according to the ordering $1,2,\dots,n$. While visiting bidder i, the notation S_i denotes the *tentative* allocation to bidder i, whereas her *actual* allocation is denoted by W_i. Each item $j \in \mathcal{J}$ is included in S_i with probability $\tilde{x}_{ij}(\mathbf{t}_i^*)$. Since the set S_i contains at most one item from each group $\mathcal{G}(i,\mathbf{t}_i^*,k)$ and there are d_i such groups, we get $|S_i| \le d_i$, and hence at most d_i items can be allocated to bidder i. The notation Q_i denotes a subset of S_i, consisting of all the items in S_i that were *not* included in the tentative allocation of any bidder $i' < i$. Each item $j \in Q_i$ is included in W_i with probability $1/(2Z_{ij})$. Bidder i gets the (random) set of items W_i and pays a (fixed) price $P_i^*(\mathbf{t}_i^*)/4$. Note that the subsets $Q_1,\dots,Q_m$ are mutually disjoint by definition, and furthermore we have $W_i \subseteq Q_i$ for all bidders $i = 1,\dots,n$. Hence, the subsets of items $W_1,\dots,W_n$ allocated to the different bidders are also mutually disjoint.

Theorem 2.2. *The All-Pay Mechanism (Figure 1) is truthful-in-expectation BIC, and its revenue is a 4-approximation to the revenue of the optimal truthful-in-expectation BIC mechanism.*

Proof. Applying the first constraint in LP1, we get $\sum_i X_{ij} \le 1/2$ for all items j. The inequality holds since we scaled down the LP variables by a factor of 2. Consequently, we have

$$Z_{ij} \;=\; \prod_{i'<i}(1-X_{ij}) \;\ge\; 1-\sum_{i'<i} X_{ij} \;\ge\; \frac{1}{2}\,.$$

This implies $1/(2Z_{ij}) \le 1$.

Fix a bidder i, her reported type $\mathbf{t}_i^*$, an item j, and suppose that every other bidder $i' \neq i$ reveals her true type. In other words, the reported type of any bidder $i' \neq i$ follows the distribution $f_{i'}(\cdot)$. Hence, for all bidders $i' \neq i$, we have

$$\Pr[j \notin S_{i'}] = 1 - \Pr[j \in S_{i'}] = 1 - \sum_{\mathbf{t}_{i'} \in \mathcal{T}_{i'}} f_{i'}(\mathbf{t}_{i'})\tilde{x}_{i'j}(\mathbf{t}_{i'}) = 1 - X_{i'j}\,.$$

We now show that bidder i gets item j with probability $x_{ij}^*(\mathbf{t}_i^*)/4$:

$$\begin{aligned}
\Pr[j \in W_i] &= \Pr[j \in W_i \mid j \in Q_i] \cdot \Pr[j \in Q_i] \\
&= \Pr[j \in W_i \mid j \in Q_i] \cdot \Pr\Big[j \in S_i \setminus \bigcup_{i'=1}^{i-1} S_{i'}\Big] \\
&= \Pr[j \in W_i \mid j \in Q_i] \cdot \Pr[j \in S_i] \cdot \prod_{i'=1}^{i-1} \Pr[j \notin S_{i'}] \qquad (2.1)\\
&= \frac{1}{2Z_{ij}} \cdot \tilde{x}_{ij}(\mathbf{t}_i^*) \cdot \prod_{i'=1}^{i-1}(1-X_{i'j}) \\
&= \frac{1}{2Z_{ij}} \cdot \tilde{x}_{ij}(\mathbf{t}_i^*) \cdot Z_{ij} \\
&= \frac{x_{ij}^*(\mathbf{t}_i^*)}{4}\,.
\end{aligned}$$

Equality (2.1) holds since the distributions of types of all the bidders are mutually independent, and we have assumed that every bidder $i' \neq i$ reveals her true type.

By linearity of expectation, the expected welfare of items allocated to bidder i is given by the expression

$$\frac{1}{4} \sum_{j \in \mathcal{J}} \mathbf{t}_i(j) \cdot x^*_{ij}(\mathbf{t}^*_i),$$

when her true type is $\mathbf{t}_i$. Since bidder i pays a fixed price $P^*_i(\mathbf{t}^*_i)/4$, both the expected welfare and expected price are scaled down by a factor exactly 4 relative to the LP values *regardless of her revealed type*. This preserves the incentive-compatibility constraints in the LP, and makes the scheme be incentive compatible and satisfy individual rationality in expectation over the types of other bidders and the randomness introduced by the mechanism. The theorem follows. □

Remark We note that if the objective function is replaced with

$$\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{\mathbf{t}_i \in \mathcal{T}_i} f_i(\mathbf{t}_i) \mathbf{t}_i(j) x_{ij}(\mathbf{t}_i),$$

then the resulting scheme gives a 4-approximation to the optimal expected social welfare.

3 Universally truthful DSIC mechanisms

We described an all-pay mechanism (Figure 1) in Section 2 that gives constant approximation to optimal revenue. This mechanism, however, is truthful-in-expectation BIC; for example, a bidder has to pay a fixed price even if she does not get any item in a random allocation. In contrast, it is more desirable in practice to implement mechanisms that are universally truthful DSIC and have good revenue properties. Unfortunately, we cannot achieve this goal unless we make additional assumptions on the input, for the following reason: The setting considered in Section 2 allows a bidder's valuations for different items to be drawn from correlated distributions, and it is computationally hard to approximate the revenue-optimal universally truthful DSIC mechanism [9] under such settings (see Section 1.3.1). We now formally state the assumptions that will be used throughout the rest of Section 3 to circumvent this hardness. We use the same notations as in Section 2.1.

Assumption 3.1. *We assume that the marginal distributions $\{f_{ij}(\cdot)\}_{i \in \mathcal{I}, j \in \mathcal{J}}$ are mutually independent (see Section 2.1), and any distribution $f_{ij}(\cdot)$ has a support $\mathcal{T}_{ij} = \{1, \dots, L_{ij}\}$, where L_{ij} is a positive integer. In other words, the valuation of bidder i for item j is a positive integer-valued random variable $v_{ij} \in [1, L_{ij}]$ with probability mass function $f_{ij}(\cdot)$, and the random variables $\{v_{ij}\}_{i \in \mathcal{I}, j \in \mathcal{J}}$ are mutually independent.*

Note that, in contrast to Section 2.2, we no longer require the distributions $f_1(\cdot), \dots, f_n(\cdot)$ to have polynomial supports. Furthermore, the assumption that the distributions $f_{ij}(\cdot)$ are defined at positive integer values is made without any loss of generality, since we can discretize continuous distributions in powers of $(1 + \varepsilon)$ and apply the same arguments. This also holds when the values taken by the random variables are not polynomially bounded.

Sequential posted-price mechanisms In this section, our main results are sequential posted-price mechanisms that give good approximations to optimal revenue. Such a mechanism considers the bidders sequentially in arbitrary order, and for each bidder, posts a certain price for each item, and lets the bidder select the subset of items she wants to buy under these prices. Note that any sequential posted-price mechanism is universally truthful DSIC by definition.

Our results We first consider the scenario where the random variables v_{ij} satisfy the monotone hazard-rate condition (Section 3.2). In this case, we present a sequential posted-price mechanism, and show that it is a $O(1)$-approximation to the optimal truthful-in-expectation BIC mechanism (Theorem 3.14). One of the interesting aspects of this result is that although our mechanism is universally truthful DSIC (which is the strongest notion of incentive compatibility), it still approximates the revenue of the optimal truthful-in-expectation BIC mechanism.

In Section 3.3, we relax the monotone hazard rate assumption, and consider the scenario where the random variables v_{ij} are arbitrary integer-valued variables over the domain $[1,L]$, that is, $L_{ij} = L$ for all $i \in \mathcal{I}, j \in \mathcal{J}$. In this case, we present a sequential posted-price mechanism that gives $O(\log L)$-approximation to the revenue of the optimal universally truthful DSIC mechanism, and we also prove that no other sequential posted-price mechanism can achieve an asymptotically better approximation ratio (Theorem 3.16).

3.1 Linear programming relaxation

We start with a simple definition.

Definition 3.2. For all $i \in \mathcal{I}, j \in \mathcal{J}$: Define $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4)$, let its probability mass function be $g_{ij}(\cdot)$, and let $\mathcal{R}_{ij} = \{1,\dots,|\mathcal{R}_{ij}|\}$ be the range of values the random variable $\mathcal{V}_{ij}$ can possibly take. Since the random variable v_{ij} takes integer values in the range $\{1,\dots,L_{ij}\}$, we get

$$\Pr[\mathcal{V}_{ij} = r] = g_{ij}(r) = \begin{cases} f_{ij}(r) & \text{if } r < B_i/4\,, \\ \sum_{v=r}^{L_{ij}} f_{ij}(v) & \text{if } r = B_i/4\,, \\ 0 & \text{otherwise.} \end{cases}$$

We now describe an LP-relaxation in order to upper bound the revenue of the optimal truthful-in-expectation BIC mechanism.

$$\text{Maximize} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} r\cdot g_{ij}(r)\cdot x_{ij}(r) \qquad (\text{LPR}\textsc{ev})$$

$$\begin{aligned}
\textstyle\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) &\le n_i && \forall i\in\mathcal{I} && (1)\\
\textstyle\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} r\cdot g_{ij}(r)\cdot x_{ij}(r) &\le B_i && \forall i\in\mathcal{I} && (2)\\
\textstyle\sum_{i\in\mathcal{I}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) &\le 1 && \forall j\in\mathcal{J} && (3)\\
x_{ij}(r) &\in [0,1] && \forall i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij} && (4)
\end{aligned}$$

Lemma 3.3. *The optimal value of the linear program (*LPREV*) is at least* $1/4$ *times the revenue of the optimal truthful-in-expectation BIC mechanism. Furthermore, in the LP solution,* $x_{ij}(r)$ *is a monotonically non-decreasing function of* r*, for all bidders* $i \in \mathcal{I}$ *and items* $j \in \mathcal{J}$*.*

Proof. Without any loss of generality, we assume that whenever a bidder is allocated a subset of items, the total price she has to pay is distributed amongst the individual items obtained. It is easy to ensure that the expected price on a single item is never greater than the valuation for the item or the overall budget. Given that $\mathcal{V}_{ij} = r \in \mathcal{R}_{ij}$, let $x_{ij}(r)$ denote the probability that item $j \in \mathcal{J}$ is allocated to bidder $i \in \mathcal{I}$, and let $p_{ij}(r)$ denote the expected price paid conditioned on obtaining the item. Since $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4)$, it is easy to see that

$$p_{ij}(r) \leq \min(v_{ij}, B_i) \leq 4 \cdot \min(v_{ij}, B_i/4) = 4 \cdot r\,.$$

Hence, the revenue of the optimal truthful-in-expectation BIC mechanism can be relaxed as:

$$\text{Maximize} \quad \sum_{i\in\mathcal{I}} \sum_{j\in\mathcal{J}} \sum_{r\in\mathcal{R}_{ij}} p_{ij}(r) \cdot g_{ij}(r) \cdot x_{ij}(r)$$

$$\begin{array}{rclll}
\sum_{j\in\mathcal{J}} \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & n_i & \forall i \in \mathcal{I} & (1)\\
\sum_{j\in\mathcal{J}} \sum_{r\in\mathcal{R}_{ij}} p_{ij}(r) \cdot g_{ij}(r) \cdot x_{ij}(r) & \leq & B_i & \forall i \in \mathcal{I} & (2)\\
\sum_{i\in\mathcal{I}} \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & 1 & \forall j \in \mathcal{J} & (3)\\
x_{ij}(r) & \in & [0,1] & \forall i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij} & (4)\\
p_{ij}(r) & \in & [0,4r] & \forall i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij} & (5)
\end{array}$$

The above program is nonlinear. Scale $p_{ij}(r)$ down by a factor of 4 so that $p_{ij}(r) \leq r$. This preserves the constraints and loses a factor of 4 in the objective. Now, for all i, j, if $p_{ij}(r) < r$, increase $p_{ij}(r)$ and decrease $x_{ij}(r)$ while preserving their product until $p_{ij}(r)$ becomes equal to r. This yields the constraints of (LPREV), and preserves the value of the objective function; but now, the objective becomes $\sum_{i,j} \sum_r r g_{ij}(r) x_{ij}(r)$. This shows that (LPREV) is a 4-approximation to the revenue of the optimal truthful-in-expectation BIC mechanism.

To show that the objective is maximized when the $x_{ij}(r)$ are monotonically non-decreasing in r, for any (i, j), preserve $\sum_r r g_{ij}(r) x_{ij}(r)$ by increasing $x_{ij}(r_2)$ and decreasing $x_{ij}(r_1)$ for $r_1 < r_2$. In this process, $\sum_r g_{ij}(r) x_{ij}(r)$ can never increase (see Corollary 3.4), preserving all the constraints, which implies the monotonicity. □

Corollary 3.4. *Fix any bidder* $i \in \mathcal{I}$*, item* $j \in \mathcal{J}$*, and* $r_1, r_2 \in \mathcal{R}_{ij}$ *such that* $r_1 < r_2$*. Furthermore, let* $\phi_{ij}(r)$ *be any positive non-decreasing function of* $r \in \mathcal{R}_{ij}$*. If we have* $0 < x_{ij}(r_1), x_{ij}(r_2) < 1$*, and we preserve* $\sum_{r\in\mathcal{R}_{ij}} \phi_{ij}(r) g_{ij}(r) x_{ij}(r)$ *while increasing* $x_{ij}(r_2)$ *and decreasing* $x_{ij}(r_1)$*, then the expression* $\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) x_{ij}(r)$ *can never increase.*

Proof. Suppose that we increase $x_{ij}(r_2)$ by an amount $\delta_2 > 0$, and we decrease $x_{ij}(r_1)$ by an amount $\delta_1 < 0$. Since $x_{ij}(r)$ remains unchanged for every $r \notin \{r_1, r_2\}$, and the sum $\sum_{r\in\mathcal{R}_{ij}} \phi_{ij}(r) g_{ij}(r) x_{ij}(r)$ is preserved, we must have

$$\phi_{ij}(r_1) \cdot g_{ij}(r_1) \cdot \delta_1 = \phi_{ij}(r_2) \cdot g_{ij}(r_2) \cdot \delta_2\,.$$

Since $\phi_{ij}(r)$ is a positive non-decreasing function of r, we have $\phi_{ij}(r_1) \leq \phi_{ij}(r_2)$. It follows that $g_{ij}(r_1)\delta_1 \geq g_{ij}(r_2)\delta_2$. Hence, we get

$$g_{ij}(r_1) \cdot (x_{ij}(r_1) - \delta_1) + g_{ij}(r_2) \cdot (x_{ij}(r_2) + \delta_2) \leq g_{ij}(r_1) \cdot x_{ij}(r_1) + g_{ij}(r_2) \cdot x_{ij}(r_2)\,.$$

Since the value of $x_{ij}(r)$ remains unchanged for all $r \notin \{r_1, r_2\}$, the expression $\sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) x_{ij}(r)$ can never increase. □

Remark Due to the presence of budget constraints, (LPREV) bounds only the expected revenue of any truthful-in-expectation BIC mechanism and *not* the expected social welfare, which can be larger by an unbounded amount. However, if the budget constraints are removed, the resulting LP also bounds the optimal social welfare.

3.2 Monotone hazard rates

We will present a constant-factor approximation to the optimal truthful-in-expectation BIC mechanism via sequential posted-price schemes, assuming that the random variables v_{ij} satisfy the monotone hazard-rate condition (Definition 1.2). Formally, we will make the following assumption throughout Section 3.2.

Assumption 3.5. *For all $i \in \mathcal{I}$ and $j \in \mathcal{J}$, the random variable v_{ij} (which denotes the valuation of bidder i for item j) satisfies the monotone hazard rate (MHR) condition.*

Claim 3.6. *If X is a positive-integer-valued random variable satisfying the monotone hazard-rate condition, then the random variable* $\min(X,a)$ *also satisfies the MHR condition, for any integer $a \geq 1$.*

Proof. Suppose that the random variable X is drawn from an MHR distribution with support $\{1,\ldots,k\}$. If $a \geq k$, then clearly the random variable $\min(X,a)$ also satisfies the MHR condition, since it has a support of $\{1,\ldots,k\}$ and we have

$$\Pr[\min(X,a) = r] = \Pr[X = r] \quad \text{for all } r \in \{1,\ldots,k\}\,.$$

Hence, we will assume that $a < k$ throughout the rest of the proof. In this case, the random variable $\min(X,a)$ follows a distribution with support $\{1,\ldots,a\}$, and we have

$$\Pr[\min(X,a) = r] = \Pr[X = r] \quad \text{for all } r \in \{1,\ldots,a-1\}\,, \tag{3.1}$$

$$\Pr[\min(X,a) > r] = \Pr[X > r] \quad \text{for all } r \in \{1,\ldots,a-1\}\,, \tag{3.2}$$

$$\Pr[\min(X,a) = a] = \sum_{r=a}^{k} \Pr[X = r]\,, \tag{3.3}$$

$$\Pr[\min(X,a) > a] = 0\,. \tag{3.4}$$

From equations (3.1), (3.2), (3.3) and (3.4), it follows that

$$\frac{\Pr[\min(X,a) > r]}{\Pr[\min(X,a) = r]} \geq \frac{\Pr[\min(X,a) > r']}{\Pr[\min(X,a) = r']} \quad \text{for all } r, r' \in \{1,\ldots,a\} \text{ such that } r < r'\,. \tag{3.5}$$

Hence, the random variable $\min(X,a)$ satisfies the MHR condition. □

Since the random variables v_{ij} satisfy the MHR condition (Assumption 3.5), we get the following corollary.

Corollary 3.7. *For all $i \in \mathcal{I}$ and $j \in \mathcal{J}$, the random variable $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4)$ satisfies the MHR condition.*

The next definition is due to Myerson [30]. First, recall that the random variable $\mathcal{V}_{ij}$ has a probability mass function $g_{ij}(\cdot)$ with support $\mathcal{R}_{ij}$ (Definition 3.2).

Definition 3.8. For all $r \in \mathcal{R}_{ij}$, let $G_{ij}(r) = \Pr[\mathcal{V}_{ij} > r]$. Define the *virtual valuation* as

$$\varphi_{ij}(r) = r - \frac{G_{ij}(r)}{g_{ij}(r)} \quad \text{for all } i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij}\,.$$

The distribution $g_{ij}(\cdot)$ is said to be *regular* if and only if we have

$$\varphi_{ij}(r) \le \varphi_{ij}(r') \quad \text{for all } r, r' \in \mathcal{R}_{ij} \text{ such that } r < r'\,.$$

Clearly, monotone hazard-rate distributions are regular. We now present a crucial lemma for monotone hazard-rate distributions.

Lemma 3.9. *If the random variable $\mathcal{V}_{ij}$ satisfies the MHR condition, then we have*

$$\Pr\left[\varphi_{ij}(\mathcal{V}_{ij}) > \frac{\mathcal{V}_{ij}}{2}\right] \ge \frac{1}{e^2}\,.$$

Proof. Recall that for all $r \in \mathcal{R}_{ij}$, we have $g_{ij}(r) = \Pr[\mathcal{V}_{ij} = r]$ and $G_{ij}(r) = \Pr[\mathcal{V}_{ij} > r]$. Let the support of the distribution $g_{ij}(\cdot)$ be given by $\mathcal{R}_{ij} = \{1, \ldots, k\}$, where k is a positive integer. Before proceeding any further, note that if $k = 1$, then $G_{ij}(1) = 0$, $\varphi_{ij}(1) = 1$, and the lemma holds since $\Pr[\varphi_{ij}(\mathcal{V}_{ij}) > \mathcal{V}_{ij}/2] = 1 \ge 1/e^2$. Thus, throughout the rest of the proof, we assume that $k \ge 2$.

Define

$$h_{ij}(r) = \frac{g_{ij}(r)}{G_{ij}(r)} \quad \text{for all } r \in \{1, \ldots, k\}\,.$$

It is easy to see that $\varphi_{ij}(r) = r - 1/h_{ij}(r)$. Hence, we have $\varphi_{ij}(r) \ge r/2$ if and only if $h_{ij}(r) \ge 2/r$. As the random variable $\mathcal{V}_{ij}$ satisfies the MHR condition, we infer that $h_{ij}(r)$ is a non-decreasing function of r. Finally, we note that $2/r$ is a non-increasing function of r, and that $h_{ij}(k) = \infty > 2/k$. These observations lead us to conclude the following:

> There is an integer $k^* \in \{1, \ldots, k\}$ such that $\varphi_{ij}(r) > r/2$ if and only if $r \ge k^*$, and $h_{ij}(r) \le h_{ij}(k^* - 1) \le 2/(k^* - 1)$ for all $r \in \{1, \ldots, k^* - 1\}$.

Therefore, we have

$$\Pr[\varphi_{ij}(\mathcal{V}_{ij}) > \mathcal{V}_{ij}/2] = \Pr[\mathcal{V}_{ij} \ge k^*]\,.$$

If $k^* = 1$, then $\Pr[\mathcal{V}_{ij} \ge k^*] = 1 > 1/e^2$, and the lemma holds. Consequently, throughout the rest of the proof, we assume that $k^* \ge 2$.

We construct a continuous probability distribution with probability density function $\hat{g}_{ij}(\cdot)$ and support $[0,k]$ such that

$$g_{ij}(r) = \int_{t=(r-1)}^{r} \hat{g}_{ij}(t)\,dt \quad \text{for all} \quad r \in \{1,\dots,k\}.$$

Next, for every real number $0 \le t \le k$, we define

$$\hat{G}_{ij}(t) = 1 - \int_{q=0}^{t} \hat{g}_{ij}(t)\,dt \quad \text{and} \quad \hat{h}_{ij}(t) = \frac{\hat{g}_{ij}(t)}{\hat{G}_{ij}(t)}.$$

Since $\frac{d}{dt}(\hat{G}_{ij}(t)) = -\hat{g}_{ij}(t)$, it follows that

$$\int \hat{h}_{ij}(t)\,dt = -\log(\hat{G}_{ij}(t)).$$

Hence, we have

$$\exp\Big(-\int_{t=0}^{(k^*-1)} \hat{h}_{ij}(t)\,dt\Big) = \frac{\hat{G}_{ij}(k^*-1)}{\hat{G}_{ij}(0)} = \hat{G}_{ij}(k^*-1). \tag{3.6}$$

Note that $\hat{G}_{ij}(t)$ is a non-increasing function of t. Thus, for all $r \in \{1,\dots,(k^*-1)\}$, we have

$$-\int_{t=(r-1)}^{r} \hat{h}_{ij}(t)\,dt \ge -\frac{1}{\hat{G}_{ij}(r)} \int_{t=(r-1)}^{r} \hat{g}_{ij}(t)\,dt = -\frac{g_{ij}(r)}{G_{ij}(r)} = -h_{ij}(r) \ge -\frac{2}{(k^*-1)}.$$

Summing over $r = 1,\dots,(k^*-1)$, we get

$$-\int_{t=0}^{(k^*-1)} \hat{h}_{ij}(t)\,dt = -\sum_{r=1}^{(k^*-1)} \int_{t=(r-1)}^{r} \hat{h}_{ij}(t)\,dt \ge -\sum_{r=1}^{(k^*-1)} \frac{2}{(k^*-1)} \ge -2. \tag{3.7}$$

To conclude the proof of the lemma, we deduce that

$$\begin{aligned}
\Pr[\varphi_{ij}(\mathcal{V}_{ij}) > \mathcal{V}_{ij}/2] &= \Pr[\mathcal{V}_{ij} \ge k^*] \\
&= 1 - \sum_{r=1}^{(k^*-1)} g_{ij}(r) \\
&= 1 - \int_{t=0}^{(k^*-1)} \hat{g}_{ij}(t)\,dt \\
&= \hat{G}_{ij}(k^*-1) \\
&= e^{-\int_{t=0}^{(k^*-1)} \hat{h}_{ij}(t)\,dt} && \text{(by equation (3.6))} \\
&\ge e^{-2} && \text{(by equation (3.7)).}
\end{aligned}$$

□

Note that the bound established is Lemma 3.9 is tight for exponential distributions. Suppose that X is a continuous random variable drawn from an exponential distribution with rate parameter μ, and let $\varphi(X)$ denote the virtual valuation function of X. We get

$$\Pr\left[\varphi(X) > \frac{X}{2}\right] = \Pr\left[X - \frac{e^{-\mu X}}{\mu \cdot e^{-\mu X}} > X/2\right] = \Pr\left[X > \frac{2}{\mu}\right] = \frac{1}{e^2}.$$

Recall Definition 3.2 and Definition 3.8. The next observation follows from Corollary 3.7 and Lemma 3.9.

Observation 3.10. For all bidders $i \in \mathcal{I}$ and items $j \in \mathcal{J}$, there exists an integer $v_{ij}^* \in \mathcal{R}_{ij}$ such that

$$\Pr[\mathcal{V}_{ij} \geq v_{ij}^*] \geq e^{-2} \quad \text{and} \quad \varphi_{ij}(\mathcal{V}_{ij}) > \mathcal{V}_{ij}/2 \quad \text{whenever} \quad \mathcal{V}_{ij} \geq v_{ij}^*.$$

3.2.1 Incorporating virtual valuations

First, we state a characterization of virtual valuations that was proved by Myerson [30]. We will invoke Lemma 3.11 multiple times throughout the rest of the paper.

Lemma 3.11. *For all bidders $i \in \mathcal{I}$ and items $j \in \mathcal{J}$, suppose that the random variable $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4)$ follows a discrete distribution $g_{ij}(\cdot)$ with support $\mathcal{R}_{ij}$, and $0 \leq x_{ij}(r) \leq 1$ for all $r \in \mathcal{R}_{ij}$. We have*

$$\sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x_{ij}(r) = \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \left(r \cdot x_{ij}(r) - \sum_{s \in \mathcal{R}_{ij} : s < r} x_{ij}(s) \right).$$

Proof. Recall that $\varphi_{ij}(r) = r - G_{ij}(r)/g_{ij}(r)$, where $G_{ij}(r) = \sum_{s \in \mathcal{R}_{ij} : s > r} g_{ij}(s)$. Hence, we get

$$\begin{aligned}
\sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x_{ij}(r) &= \sum_{r \in \mathcal{R}_{ij}} \left(r \cdot g_{ij}(r) - \sum_{s \in \mathcal{R}_{ij} : s > r} g_{ij}(s) \right) \cdot x_{ij}(r) \\
&= \sum_{r \in \mathcal{R}_{ij}} r \cdot g_{ij}(r) \cdot x_{ij}(r) - \sum_{r \in \mathcal{R}_{ij}} \sum_{s \in \mathcal{R}_{ij} : s > r} g_{ij}(s) \cdot x_{ij}(r) \\
&= \sum_{r \in \mathcal{R}_{ij}} r \cdot g_{ij}(r) \cdot x_{ij}(r) - \sum_{r \in \mathcal{R}_{ij}} \sum_{s \in \mathcal{R}_{ij} : s < r} g_{ij}(r) \cdot x_{ij}(s) \\
&= \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \left(r \cdot x_{ij}(r) - \sum_{s \in \mathcal{R}_{ij} : s < r} x_{ij}(s) \right).
\end{aligned}$$

This concludes the proof. □

Now consider the following linear program obtained from (LPREV).

$$\text{Maximize} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot\varphi_{ij}(r)\cdot x_{ij}(r) \qquad \text{(LP2)}$$

$$\begin{aligned}
\textstyle\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) &\le n_i & \forall i\in\mathcal{I} && (1)\\
\textstyle\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot\varphi_{ij}(r)\cdot x_{ij}(r) &\le B_i & \forall i\in\mathcal{I} && (2)\\
\textstyle\sum_{i\in\mathcal{I}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) &\le 1 & \forall j\in\mathcal{J} && (3)\\
x_{ij}(r) &\in [0,1] & \forall i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij} && (4)
\end{aligned}$$

Lemma 3.12. *The value of* (LP2) *is at least* $1/(2e^2)$ *times the value of* (LPREV).

Proof. Let $\{x^*_{ij}(r)\}$ denote the values taken by the $x_{ij}(r)$ variables in the optimal solution to (LPREV). For all $i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij}$, define:

$$\tilde{x}_{ij}(r) = \begin{cases} 0 & \text{if } r < v^*_{ij}\,, \\ x^*_{ij}(r) & \text{if } r \ge v^*_{ij}\,. \end{cases} \tag{3.8}$$

The optimal objective of (LPREV) is given by

$$\sum_{i,j} \mathbf{E}_{\mathcal{V}_{ij}\sim g_{ij}(\cdot)}[\mathcal{V}_{ij}\cdot x^*_{ij}(\mathcal{V}_{ij})]\,.$$

Lemma 3.3 implies that the function $\mathcal{V}_{ij}\cdot x^*_{ij}(\mathcal{V}_{ij})$ is monotonically non-decreasing in $\mathcal{V}_{ij}$. Hence, we get

$$\begin{aligned}
\mathbf{E}_{\mathcal{V}_{ij}}[\mathcal{V}_{ij}\cdot x^*_{ij}(\mathcal{V}_{ij})] &\le \mathbf{E}_{\mathcal{V}_{ij}}[\mathcal{V}_{ij}\cdot x^*_{ij}(\mathcal{V}_{ij}) \mid \mathcal{V}_{ij}\ge v^*_{ij}] \\
&\le \mathbf{E}_{\mathcal{V}_{ij}}[2\cdot\varphi_{ij}(\mathcal{V}_{ij})\cdot x^*_{ij}(\mathcal{V}_{ij}) \mid \mathcal{V}_{ij}\ge v^*_{ij}] && \text{(Observation 3.10)}\\
&= \frac{2\cdot\mathbf{E}_{\mathcal{V}_{ij}}[\varphi_{ij}(\mathcal{V}_{ij})\cdot\tilde{x}_{ij}(\mathcal{V}_{ij})]}{\Pr[\mathcal{V}_{ij}\ge v^*_{ij}]} && \text{(by equation (3.8))}\\
&\le 2e^2\cdot\mathbf{E}_{\mathcal{V}_{ij}}[\varphi_{ij}(\mathcal{V}_{ij})\cdot\tilde{x}_{ij}(\mathcal{V}_{ij})] && \text{(Observation 3.10)}\\
&= 2e^2\cdot\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot\varphi_{ij}(r)\cdot\tilde{x}_{ij}(r)\,.
\end{aligned}$$

Summing over all bidders $i\in\mathcal{I}$ and items $j\in\mathcal{J}$, we infer:

$$\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} r\cdot g_{ij}(r)\cdot x^*_{ij}(r) \le 2e^2\cdot\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot\varphi_{ij}(r)\cdot\tilde{x}_{ij}(r)\,. \tag{3.9}$$

Furthermore, we note the following inequalities.

$$0 \le \tilde{x}_{ij}(r) \le x^*_{ij}(r) \quad \text{for all } i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij}\,, \tag{3.10}$$

$$\varphi_{ij}(r)\cdot\tilde{x}_{ij}(r) \le r\cdot x^*_{ij}(r) \quad \text{for all } i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij}\,. \tag{3.11}$$

Equation (3.10) holds by definition, and equation (3.11) holds since $\varphi_{ij}(r) \le r$ (Definition 3.8). Recall that the $x^*_{ij}(r)$ values constitute an optimal solution to (LPREV). Hence, equations (3.9), (3.10), and (3.11) imply that the $\tilde{x}_{ij}(r)$ values constitute a feasible solution to (LP2), having an objective that is within a factor of $2e^2$ of the optimal objective of (LPREV). The lemma follows. □

Lemma 3.13. *Let $x^*_{ij}(r)$ denote the values assigned to the variables in the optimal solution to* (LP2). *For all $i \in \mathcal{I}, j \in \mathcal{J}$, we can express the corresponding $\{x^*_{ij}(r)\}_{r\in\mathcal{R}_{ij}}$ values as a convex combination of two solutions. The first (respectively, second) solution has variable values $\gamma^*_{ij}(r)$ (respectively,$\lambda^*_{ij}(r)$), and an integer $r^*_{ij} \in \mathcal{R}_{ij}$ (respectively, $s^*_{ij} = r^*_{ij} + 1 \in \mathcal{R}_{ij}$) so that*

$$\gamma^*_{ij}(r) = \begin{cases} 0 & \text{if } r < r^*_{ij}, r \in \mathcal{R}_{ij}, \\ 1 & \text{if } r \ge r^*_{ij}, r \in \mathcal{R}_{ij}, \end{cases} \quad \text{and} \quad \lambda^*_{ij}(r) = \begin{cases} 0 & \text{if } r < s^*_{ij}, r \in \mathcal{R}_{ij}, \\ 1 & \text{if } r \ge s^*_{ij}, r \in \mathcal{R}_{ij}. \end{cases}$$

Suppose that in the convex combination, the first solution has weight $0 \le w_{ij} \le 1$ and the second solution has weight $1 - w_{ij}$. In this case, we have $x^*_{ij}(r) = w_{ij} \cdot \gamma^*_{ij}(r) + (1 - w_{ij}) \cdot \lambda^*_{ij}(r)$, and furthermore

$$\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x^*_{ij}(r) = w_{ij} \cdot r^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \ge r^*_{ij}] + (1 - w_{ij}) \cdot s^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \ge s^*_{ij}]. \tag{3.12}$$

Proof. Since the variable values $\{x^*_{ij}(r)\}$ constitute an optimal solution to (LP2), we have $x^*_{ij}(r) = 0$ whenever $\varphi_{ij}(r) \le 0$. Furthermore, the virtual valuation $\varphi_{ij}(r)$ is a non-decreasing function of $r \in \mathcal{R}_{ij}$ (Corollary 3.7). We now apply Corollary 3.4 and infer the following: There exists an integer $r^* \in \mathcal{R}_{ij}$ such that

$$x^*_{ij}(r) = \begin{cases} 0 & \text{if } r < r^*, r \in \mathcal{R}_{ij}, \\ 1 & \text{if } r > r^*, r \in \mathcal{R}_{ij}. \end{cases}$$

This implies that the optimal solution to (LP2) can be written in the fashion implied by the lemma, with $r^*_{ij} = r^*$ and $w_{ij} = x^*_{ij}(r^*)$. To prove equation (3.12), we first note that

$$\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\varphi_{ij}(r)x^*_{ij}(r) = w_{ij} \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\varphi_{ij}(r)\gamma^*_{ij}(r) + (1 - w_{ij}) \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\varphi_{ij}(r)\lambda^*_{ij}(r). \tag{3.13}$$

Let the range of values of the random variable $\mathcal{V}_{ij}$ be given by $\mathcal{R}_{ij} = \{1, \dots, k\}$. Recall that $\gamma^*_{ij}(r) = 0$ if $1 \le r < r^*_{ij}$ and $\gamma^*_{ij}(r) = 1$ if $r^*_{ij} \le r \le k$. Hence, we have

$$\begin{aligned} \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot \gamma^*_{ij}(r) &= \sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \left(r \cdot \gamma^*_{ij}(r) - \sum_{s\in\mathcal{R}_{ij}: s<r} \gamma^*_{ij}(s) \right) && \text{(Lemma 3.11)} \\ &= \sum_{r=r^*_{ij}}^{k} g_{ij}(r) \left(r - (r - r^*_{ij}) \right) \\ &= r^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \ge r^*_{ij}]. \end{aligned}$$

Similarly, we can show that

$$\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot \lambda^*_{ij}(r) = s^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \ge s^*_{ij}].$$

Now, the lemma follows from equation (3.13). □

3.2.2 Posted-price mechanism and analysis

The posted-price mechanism is described in Figure 2. In Step 1, we select an arbitrary ordering of the bidders. In Step 2, we find the optimal solution to (LP2). To explain Step 3 and Step 4, we first recall Lemma 3.13, which has the following simple interpretation. For all bidders $i \in \mathcal{I}$ and items $j \in \mathcal{J}$, the optimal solution to (LP2) treats the pair (i, j) as a separate entity who is interested only in item j, and who is willing to pay at most $\mathcal{V}_{ij}$ (Definition 3.2) for item j. Furthermore, the pair (i, j) is offered item j at a posted price that is chosen randomly from the set $\{r^*_{ij}, s^*_{ij}\}$. Note that this is not a feasible mechanism: Since each pair (i, j) is considered separately, an item may be allocated more than once, or a bidder may exhaust her budget or demand. The LP solution, however, is extremely useful, as it allows us to "uncouple" all the (i, j) pairs and find the prices $\{r^*_{ij}, s^*_{ij}\}$ for each of them.

Motivated by the above interpretation, in Steps 3 and 4, we set $\tilde{r}_{ij} = r^*_{ij}$ with probability w_{ij}, and with the remaining probability $(1 - w_{ij})$, we set $\tilde{r}_{ij} = s^*_{ij}$. Furthermore, we ensure that these random events corresponding to all possible (bidder, item) pairs are mutually independent. Under the assumption that every (i, j) pair behaves as a separate entity who is willing to pay at most $\mathcal{V}_{ij}$ for an item, let $\tilde{X}_{ij}$ be the indicator random variable that is set to 1 if she is allocated the item, and let $\tilde{P}_{ij}$ denote her payment. Note that $\tilde{P}_{ij} = \tilde{r}_{ij}$ if $\tilde{X}_{ij} = 1$, and $\tilde{P}_{ij} = 0$ if $\tilde{X}_{ij} = 0$. It follows that

$$\begin{aligned}
\mathbf{E}[\tilde{X}_{ij}] &= w_{ij} \cdot \Pr[\mathcal{V}_{ij} \geq r^*_{ij}] + (1 - w_{ij}) \cdot \Pr[\mathcal{V}_{ij} \geq s^*_{ij}] \\
&= w_{ij} \cdot \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \gamma^*_{ij}(r) + (1 - w_{ij}) \cdot \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \lambda^*_{ij}(r) \\
&= \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot x^*_{ij}(r) .
\end{aligned}$$

We can also write:

$$\mathbf{E}[\tilde{P}_{ij}] = w_{ij} \cdot r^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \geq r^*_{ij}] + (1 - w_{ij}) \cdot s^*_{ij} \cdot \Pr[\mathcal{V}_{ij} \geq s^*_{ij}] = \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x^*_{ij}(r) .$$

The last equality follows from equation (3.12). Hence, the optimal objective value of (LP2) is given by $\sum_{i,j} \mathbf{E}[\tilde{P}_{ij}]$. Furthermore, Constraints 1, 2 and 3 of (LP2) imply that

$$\sum_{j \in \mathcal{J}} \mathbf{E}[\tilde{X}_{ij}] \leq n_i , \qquad \sum_{j \in \mathcal{J}} \mathbf{E}[\tilde{P}_{ij}] \leq B_i , \qquad \text{and} \qquad \sum_{i \in \mathcal{I}} \mathbf{E}[\tilde{X}_{ij}] \leq 1 . \tag{3.14}$$

To summarize, if every pair (i, j) were behaving as a separate entity and we posted the (random) price $\tilde{r}_{ij}$ for the pair (i, j), then the total expected revenue will be equal to the optimal objective of (LP2), and the demand, budget, and supply constraints will hold in expectation.

There are two difficulties in implementing the above scheme: (1.) In reality, the pairs (i, j) do not behave as separate entities. (2.) Since the budget, demand and supply constraints hold *only* in expectation, there may be occasions where one or more of these constraints are violated.

Step 5 circumvents these difficulties. We visit the bidders according to the pre-determined ordering. While visiting bidder i, if an item j is available (i. e., has not been purchased by any bidder $i' < i$), then we offer item j to bidder i with probability $1/4$, at a (random) posted price $\tilde{r}_{ij}$. Intuitively, this step ensures that with constant probability, each pair (i, j) behaves as a separate agent (see the proof of Theorem 3.14),

Posted-Price Mechanism

1. Choose an arbitrary but fixed ordering of all bidders and denote it by $1, \ldots, n$.
2. Solve LP2, and let x^*_{ij} denote the value assigned to variable x_{ij}, for all $i \in \mathcal{I}, j \in \mathcal{J}$.
3. FOR each (i, j)
   Independently pick one of the two solutions in the convex combination (see Lemma 3.13) with probability equal to its weight in the combination.
4. Let $\tilde{r}_{ij} \in \{r^*_{ij}, s^*_{ij}\}$ denote the threshold in the chosen solution where the allocation function $\tilde{\gamma}_{ij}$ jumps from zero to one.
5. Let $\{Y_{ij}\}_{i\in\mathcal{I}, j\in\mathcal{J}}$ be a set of mn mutually independent $0/1$ random variables such that
   $\Pr[Y_{ij} = 1] = 1/4$ and $\Pr[Y_{ij} = 0]$, for all $i \in \mathcal{I}, j \in \mathcal{J}$.

   $Q \leftarrow \mathcal{J}$ (Initialize Q to be the set of all items).
   FOR $i = 1, 2, \ldots, n$
   Initialize $W_i \leftarrow \emptyset$;
   FOR each item $j \in Q$
   $W_i \leftarrow W_i \cup \{j\}$ iff $Y_{ij} = 1$;
   Only the set of items W_i is offered to bidder i;
   Each $j \in W_i$ is offered at a price $\tilde{r}_{ij}$;
   Bidder i buys a subset of items $S_i \subseteq W_i$;
   $Q \leftarrow Q \setminus S_i$.

Figure 2: DSIC Mechanism for independent valuations

so that the expected revenue remains within a constant factor of the optimal objective of (LP2). Applying Lemma 3.3 and Lemma 3.12, we infer that the revenue of the sequential posted-price mechanism is a $O(1)$-approximation to the revenue of the optimal Bayesian incentive-compatible mechanism.

Note that since an item is offered to a bidder *only if* it has not been purchased by anyone else, and since the bidder herself decides the subset of items she wants to buy at the posted prices, this mechanism is clearly feasible (i. e., satisfy the demand, budget and supply constraints), and universally truthful DSIC.

Theorem 3.14. *The posted-price mechanism in Figure 2 is a universally truthful DSIC mechanism and its revenue is a $O(1)$-approximation to the revenue of the optimal truthful-in-expectation BIC mechanism, when the valuations of a bidder follow product distributions that satisfy the monotone hazard-rate condition.*

Proof. For all bidders $i \in \mathcal{I}$ and items $j \in \mathcal{J}$, let X_{ij} be the $0/1$ random variable denoting whether item j is taken by bidder i, and let P_{ij} denote the price at which it is taken (which is 0 if the item is not taken by the bidder). Recall equation (3.14), where $\tilde{X}_{ij}$ (respectively, $\tilde{P}_{ij}$) is the random variable for the allocation (respectively, payment) corresponding to the pair (i, j), under the assumption that every (bidder, item) pair behaves as a separate entity. Since the probability that any pair (i, j) is considered at all is $1/4$ (see Step 5, Figure 2), we infer that $X_{ij} \le \tilde{X}_{ij}/4$ and $P_{ij} \le \tilde{P}_{ij}/4$. By linearity of expectation, the inequalities of equation (3.14) imply that

$$\mathbf{E}\left[\sum_{k\neq j} P_{ik}\right] \leq B_i/4\,, \quad \mathbf{E}\left[\sum_{k\neq j} X_{ik}\right] \leq d_i/4\,, \quad \text{and} \quad \mathbf{E}\left[\sum_{k\neq i} X_{kj}\right] \leq 1/4 \quad \text{for all } i \in \mathcal{I}, j \in \mathcal{J}\,.$$

Since the valuation v_{ij} is independent of other valuations, this implies the above statements hold regardless of v_{ij}. Now applying Markov's inequality, we have for all i, j:

$$\Pr\left[\sum_{k\neq j} P_{ik} \geq 3B_i/4\right] \leq 1/3\,, \quad \Pr\left[\sum_{k\neq j} X_{ik} \geq d_i\right] \leq 1/4\,, \quad \text{and} \quad \Pr\left[\sum_{k\neq i} X_{kj} \geq 1\right] \leq 1/4\,.$$

By union bounds, this implies that with probability at least $1/6$, we have

$$\sum_{k\neq j} P_{ik} < \frac{3B_i}{4}\,, \quad \sum_{k\neq j} X_{ik} < d_i\,, \quad \text{and} \quad \sum_{k\neq i} X_{kj} = 0\,.$$

In this event, item j is offered to bidder i with probability $1/4$ (using one of two random choices of the posted price from Lemma 3.13). Furthermore, in this event, the bidder will take the item if $v_{ij} \geq \tilde{r}_{ij}$, since $\tilde{r}_{ij} \leq B_i/4$ (so that the bidder has sufficient budget to purchase this item), and the bidder has not exhausted his demand d_i. Since the valuation v_{ij} itself is independent of the event that the item j is offered to bidder i, this implies that with probability at least $1/6 \times 1/4$, the posted-price mechanism obtains revenue $\sum_r g_{ij}(r)\varphi_{ij}(r)x^*_{ij}(r)$ along each (i, j) (using the definition of the latter quantity from equation (3.12)). By linearity of expectation over all (i, j), we have a $O(1)$-approximation to the objective of (LP2). The theorem follows from Lemma 3.3 and Lemma 3.12. □

Corollary 3.15. *Under the assumptions of Theorem 3.14, if there are no budget constraints, then the revenue of the sequential posted-price mechanism is a constant-factor approximation to the optimal social welfare.*

Proof. If the Budget Constraint (2) is removed from (LPREV), it is an upper bound on the optimal social welfare that can be obtained by any mechanism. The rest of the analysis remains the same. □

One issue with mechanisms where bidders have budget and demand constraints is that given posted prices for the items, the bidder needs to solve a two-dimensional KNAPSACK problem to determine her optimal bundle, and this is NP-hard in general. However, this does not change our results, for the following reasons. First, the analysis of the algorithm in Figure 2 simply shows that with constant probability, the bidder solves a trivial KNAPSACK instance where all items fit into the KNAPSACK. Another interesting aspect of our mechanism is that the LP relaxations (both (LPREV) and (LP2)) do not encode any incentive-compatibility constraint. Furthermore, in a sequential posted-price mechanism, a bidder $i \in \mathcal{I}$ does not have to report her valuations to the auctioneer. Depending on the publicly known valuation distributions, the choices made by the bidders $i' < i$, and the randomness of the mechanism, bidder i is offered a subset W_i of available items, and she has to pay a posted price $\tilde{r}_{ij}$ for buying any item $j \in W_i$. The subset of items W_i and the posted prices $\{\tilde{r}_{ij}\}$ do not depend on the valuations of bidder i. Hence, the scheme remains incentive compatible regardless of the KNAPSACK heuristic used by the bidder. Therefore, our results hold as long as the bidder uses any reasonable KNAPSACK heuristic that is maximal, in the sense that the bidder goes on buying more and more items until she exhausts her demand or budget constraints.

3.3 General distributions

The key point shown by Theorem 3.14 is that the revenue of the optimal truthful-in-expectation BIC mechanism is approximated to a constant factor by a simple sequential posted-price mechanism, where the prices are posted for each item and each bidder. This crucially used the monotone hazard-rate condition. The natural question to ask is whether such a simple mechanism is a good approximation for more general classes of distributions. Note that even for one bidder, the optimal universally truthful DSIC mechanism may not be a sequential posted-price scheme. We show below that the optimal universally truthful DSIC mechanism has a logarithmic advantage over sequential posted pricing even for regular distributions (see Definition 3.8), and that this gap is tight.

Theorem 3.16. *Suppose that $v_{ij} \in [1,L]$, are independent for different (i,j) but do not necessarily satisfy the monotone hazard-rate condition. Then, there is a $\Theta(\log L)$ gap between the revenues of the optimal sequential posted-price scheme and the optimal universally truthful DSIC mechanism. The lower bound holds for regular distributions and one bidder.*

Proof. To show the lower bound, consider the following scenario: There is only one bidder, m items, and no budget or demand constraints. The valuations are i. i. d. for each item j, and follow the common distribution with $\Pr[v \geq r] = 1/r$, for $r = 1,2,\ldots,m$, so that $m = L$. This distribution is regular, since $\varphi(r) = 0$ for $r < m$, and $\varphi(m) = m$. We have $\mathbf{E}[v_{1j}] = \sum_{r=1}^{m}(1/r) = H_m$ for every item j, and by Chernoff bounds,

$$\Pr\left[\sum_{j=1}^{m} v_{1j} \leq (1-\varepsilon)mH_m\right] \leq e^{\frac{-mH_m\varepsilon^2}{2m}} = o(1)\,.$$

Therefore, a scheme that sets a price of $mH_m(1-\varepsilon)$ for the bundle of m items sells the bundle with probability $1-o(1)$, so that the expected revenue is $\Omega(m\log m)$. Any sequential posted-price scheme can extract a revenue of $\max_r r\cdot\Pr[v \geq r] = 1$ from each item, so that the total expected revenue from m items is at most $m \times 1 = m$. This shows a gap of $\Omega(\log m) = \Omega(\log L)$.

The upper bound follows from a standard scaling argument. We start with (LPREV) which upper bounds the optimal achievable revenue. For each (i,j), we group the r values in powers of 2 so that there are $(\log L)$ groups. Let group $\mathcal{G}_k$ denote the interval $[2^k, 2^{k+1})$. We find a group $\mathcal{G}_{k^*_{ij}}$ that contributes at least $(1/\log L)$ fraction to the sum $\sum_r r\, g_{ij}(r) x_{ij}(r)$. Specifically, we require that

$$\sum_{r\in\mathcal{G}_{k^*_{ij}}} r\cdot g_{ij}(r)\cdot x_{ij}(r) \geq (1/\log L)\sum_{r\in\mathcal{R}_{ij}} r\cdot g_{ij}(r)\cdot x_{ij}(r)\,.$$

Note that if we use the price $2^{k^*_{ij}}$ for the pair (i,j), and every (bidder, item) pair behaves as a separate entity, then our expected revenue will be at least $(1/2\log L)$ times the optimal objective value of (LPREV). Suppose that we modify the Steps 2, 3, and 4 of the algorithm in Figure 2: Instead of solving (LP2) and randomly selecting $\tilde{r}_{ij} \in \{r^*_{ij}, s^*_{ij}\}$, we solve (LPREV) and deterministically set $\tilde{r}_{ij} = 2^{k^*_{ij}}$. Step 5 of the algorithm is left unchanged. Using similar arguments as in the proof of Theorem 3.14, it is now easy to see that the resulting sequential posted-price mechanism will give a $\Omega(1/\log L)$-approximation to the optimal revenue.

$\square$

3.3.1 Adaptive posted-price schemes

The above implies that we cannot generalize the result in Section 3.2 even to the class of regular distributions (see Definition 1.1). This is because an $\omega(1)$ gap is introduced in going from (LPREV) to (LP2). However, we can show that if the space of mechanisms is restricted, (LP2) itself is a good relaxation to the optimal revenue in this space.

In particular, suppose that we restrict the space of mechanisms to be those that are posted price, and sequential by bidder: Depending on the subset of items and bidders left, the mechanism adaptively chooses the next bidder and posts prices for a subset of the remaining items. The bidder, being a utility maximizer, solves a knapsack problem to choose the optimal subset of items. Once this bidder is dealt with, the mechanism again adaptively chooses the next bidder depending on the acceptance strategy of this bidder. Since the optimization problem the bidders need to solve is a two-dimensional KNAPSACK problem, which is NP-hard, we assume that the bidders are *monotone optimizers* in the following sense.

Assumption 3.17. *Suppose that we are given some adaptive posted-price scheme. Consider any pair (i, j), where $i \in \mathcal{I}, j \in \mathcal{J}$. Fix the valuations $v_{i'j'}$ corresponding to all other (bidder, item) pairs $(i', j') \neq (i, j)$. Also fix the* KNAPSACK *heuristics used by all the other bidders $i' \neq i$. It follows that we can uniquely determine the subset of items that will be offered to bidder i, and the prices at which those items will be offered to bidder i. Under these circumstances, we assume that the quantity of item j taken by bidder i is a monotonically non-decreasing function of v_{ij}.*

The above assumption holds if the bidder solves KNAPSACK optimally. Similar to the posted-price scheme in Section 3.2.2, our algorithm itself will allow arbitrary KNAPSACK heuristics, as long as it satisfies Assumption 3.17 and the bidder buys all the items if she is not constrained by either demand or budget. Next, we state another assumption that will be used throughout the rest of Section 3.3.1.

Assumption 3.18. *The valuations of different bidders for different items follow independent distributions, that is, the random variables $\{v_{ij}\}_{i\in\mathcal{I},j\in\mathcal{J}}$ are mutually independent. Furthermore, the distribution $f_{ij}(\cdot)$, from which the random variable v_{ij} is drawn, is regular for all $i \in \mathcal{I}, j \in \mathcal{J}$.*

Under Assumptions 3.17 and 3.18, we will show a $O(1)$-approximation to the revenue maximizing adaptive posted-price mechanism. First, we consider the following nonlinear program.

$$\text{Maximize} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot 4\cdot \tilde{p}_{ij}(r)\cdot x_{ij}(r) \qquad (\text{SEQ})$$

$$\begin{array}{rcllr}
\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) & \leq & d_i & \forall i\in\mathcal{I} & (1)\\
\sum_{j\in\mathcal{J}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot \tilde{p}_{ij}(r)\cdot x_{ij}(r) & \leq & B_i & \forall i\in\mathcal{I} & (2)\\
\sum_{i\in\mathcal{I}}\sum_{r\in\mathcal{R}_{ij}} g_{ij}(r)\cdot x_{ij}(r) & \leq & 1 & \forall j\in\mathcal{J} & (3)\\
\tilde{p}_{ij}(r)\cdot x_{ij}(r) & \leq & r\cdot x_{ij}(r) - \sum_{s=1}^{r-1} x_{ij}(s) & \forall i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij} & (4)\\
x_{ij}(r) & \in & [0,1] & \forall i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij} & (5)\\
\tilde{p}_{ij}(r) & \in & [0,r] & \forall i\in\mathcal{I}, j\in\mathcal{J}, r\in\mathcal{R}_{ij} & (6)
\end{array}$$

Lemma 3.19. *If the bidders are monotone optimizers, then the objective of the nonlinear program* (SEQ) *upper bounds the revenue of the optimal adaptive posted-price mechanism.*

Proof. Fix any adaptive posted-price mechanism, and let $x_{ij}(r)$ denote the probability that bidder $i \in \mathcal{I}$ obtains item $j \in \mathcal{J}$ when $\mathcal{V}_{ij} = r \in \mathcal{R}_{ij}$ (Definition 3.2). Similarly, let $p_{ij}(r)$ denote the expected price paid by bidder i for item j, conditioned on the events that $\mathcal{V}_{ij} = r$ and bidder i obtains item j. Define $\tilde{p}_{ij}(r) = p_{ij}(r)/4$. It follows that the auctioneer's revenue is equal to the objective of (SEQ). Constraint 1 of (SEQ) holds since no bidder i buys more than d_i items, constraint 2 holds since no bidder pays more than her budget, and constraint 3 holds since no item is allocated more than once.

To interpret constraint 4, fix a bidder $i \in \mathcal{J}$, an item $j \in \mathcal{J}$, and also fix all the valuations $v_{i'j'}$ corresponding to all other (bidder, item) pairs $(i', j') \neq (i, j)$, and the KNAPSACK heuristics used by every bidder $i' \neq i$. This uniquely determines the subset of items (and their prices) that will be offered to bidder i. If the adaptive posted-price mechanism offers item j to bidder i, then it posts a unique price p^* for the item. Given these conditions, let $X_{ij}(r) \in [0,1]$ denote the probability that bidder i takes item j when $\mathcal{V}_{ij} = r \in \mathcal{R}_{ij}$, and let $P_{ij}(r) = p^* \cdot X_{ij}(r)$ denote the expected price paid by bidder i for item j when $\mathcal{V}_{ij} = r \in \mathcal{R}_{ij}$. We will show:

$$\frac{P_{ij}(r)}{4} = \left(\frac{p^*}{4}\right) \cdot X_{ij}(r) \leq r \cdot X_{ij}(r) - \sum_{s=1}^{r-1} X_{ij}(s) \quad \text{for all } r \in \mathcal{R}_{ij}\,. \tag{3.15}$$

To verify equation (3.15), we consider two possible cases.

Case 1: $p^*/4 > B_i/4$.

In this case, bidder i never takes item j since the price exceeds her budget, implying that $X_{ij}(r) = P_{ij}(r) = 0$ for all $r \in \mathcal{R}_{ij}$. Hence, equation (3.15) holds.

Case 2: $p^*/4 \leq B_i/4$.

To analyze this case, first recall that we have already fixed both the valuations of bidder i for every item $j' \neq j$, and the subset of items (and their prices) offered to bidder i. Next, note that if $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4) < B_i/4$, then v_{ij} is also uniquely determined. Hence, we see that $X_{ij}(r)$ is either 0 or 1 for all $r < B_i/4$. However, at $r = B_i/4$, the valuation v_{ij} is not unique and $X_{ij}(r)$ can take a fractional value. Since the bidder is a monotone utility maximizer (Assumption 3.17), we infer that $X_{ij}(r)$ is a monotone step function of r, and this step function has at most one jump. If the step function has no jump at all, then $X_{ij}(r) = 0$ for all $r \in \mathcal{R}_{ij}$, and equation (3.15) holds. Consequently, we assume that the step function jumps at $r^* \leq B_i/4$, that is, $X_{ij}(r) = 0$ for all $r < r^*$, and $X_{ij}(r) = X_{ij}(r^*) > 0$ for all $r \geq r^*$. It is easy to see that equation (3.15) holds for all $r < r^*$.

Recall that $p^*/4 \leq B_i/4$. As a consequence, if $\mathcal{V}_{ij} = \min(v_{ij}, B_i/4) < p^*/4$, then $v_{ij} < p^*/4 < p^*$. Since bidder i never takes item j when $v_{ij} < p^*$, and since $X_{ij}(r^*) > 0$, it follows that $p^*/4 \leq r^*$. To summarize, it remains to verify equation (3.15) only if $p^*/4 \leq r^* \leq r \leq B_i/4$. In this situation, we get

$$\left(\frac{p^*}{4}\right) \cdot X_{ij}(r) = \left(\frac{p^*}{4}\right) \cdot X_{ij}(r^*) \leq r^* \cdot X_{ij}(r^*)\,. \tag{3.16}$$

Furthermore, we have

$$r \cdot X_{ij}(r) - \sum_{s=1}^{r-1} X_{ij}(s) = r \cdot X_{ij}(r^*) - (r - r^*) X_{ij}(r^*) = r^* \cdot X_{ij}(r^*) \,. \tag{3.17}$$

Equation (3.15) follows from equations (3.16) and (3.17). Having verified equation (3.15), we proceed with the proof of the lemma.

Taking expectation over all the remaining valuations $\{v_{i'j'}\}$, where $(i', j') \neq (i, j)$, and recalling the definitions of $x_{ij}(r), p_{ij}(r)$, we see that $x_{ij}(r) = \mathbf{E}[X_{ij}(r)]$ and $p_{ij}(r) x_{ij}(r) = \mathbf{E}[P_{ij}(r)]$. Hence, equation (3.15) implies:

$$\tilde{p}_{ij}(r) \cdot x_{ij}(r) = \frac{p_{ij}(r)}{4} \cdot x_{ij}(r) = \mathbf{E}\left[\frac{P_{ij}(r)}{4}\right] \leq r \cdot x_{ij}(r) - \sum_{s=1}^{r-1} x_{ij}(s) \,. \tag{3.18}$$

Thus, constraint 4 of (SEQ) holds, and the revenue maximizing adaptive posted-price mechanism is a feasible solution to the nonlinear program (SEQ). □

Lemma 3.20. *If the bidders are monotone optimizers, then* (LP2) *is a* 4*-approximation to the revenue of the optimal adaptive posted-price scheme.*

Proof. First, we take the optimal solution to the non-linear program (SEQ). Next, for all $i \in \mathcal{I}, j \in \mathcal{J}$, let the range of $\mathcal{V}_{ij}$ be given by $\mathcal{R}_{ij} = \{1, \ldots, |\mathcal{R}_{ij}|\}$. We do the following:

FOR all $i \in \mathcal{I}, j \in \mathcal{J}$
  FOR $r = 1, 2, \ldots, |\mathcal{R}_{ij}|$
    Increase the value of $\tilde{p}_{ij}(r)$ and decrease the value of $x_{ij}(r)$, ensuring that their product $\tilde{p}_{ij}(r) \cdot x_{ij}(r)$ remains the same, and continue doing this until the corresponding constraint 4 (defined by the tripe (i, j, r)) of the non-linear program (SEQ) becomes tight.

This preserves the objective of (SEQ) and all the other constraints. When this process terminates, we get an optimal solution to (SEQ) where the constraint 4 is tight for all $i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij}$. Hence, we can replace $\tilde{p}_{ij}(r) \cdot x_{ij}(r)$ by the right hand side of constraint 4, and rewrite the non-linear program (SEQ) as a linear program (LP3).

$$\text{Maximize} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot 4 \cdot \left(r \cdot x_{ij}(r) - \sum_{s=1}^{r-1} x_{ij}(s) \right) \qquad \text{(LP3)}$$

$$\begin{array}{rcll}
\sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & d_i & \forall i \in \mathcal{I} \\
\sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \left(r x_{ij}(r) - \sum_{s=1}^{r-1} x_{ij}(s) \right) & \leq & B_i & \forall i \in \mathcal{I} \\
\sum_{i \in \mathcal{I}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & 1 & \forall j \in \mathcal{J} \\
x_{ij}(r) & \in & [0,1] & \forall i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij}
\end{array}$$

Now, we invoke Lemma 3.11 and deduce that (LP3) is equivalent to the following linear program.

$$\text{Maximize} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} 4 \cdot g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x_{ij}(r) \qquad \text{(LP4)}$$

$$\begin{array}{rcll}
\sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & d_i & \forall i \in \mathcal{I} \\
\sum_{j \in \mathcal{J}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot \varphi_{ij}(r) \cdot x_{ij}(r) & \leq & B_i & \forall i \in \mathcal{I} \\
\sum_{i \in \mathcal{I}} \sum_{r \in \mathcal{R}_{ij}} g_{ij}(r) \cdot x_{ij}(r) & \leq & 1 & \forall j \in \mathcal{J} \\
x_{ij}(r) & \in & [0,1] & \forall i \in \mathcal{I}, j \in \mathcal{J}, r \in \mathcal{R}_{ij}
\end{array}$$

Hence, the objective of (LP4) upper bounds the revenue of the optimal adaptive posted-price scheme when the bidders are monotone optimizers. The lemma follows from comparing (LP4) with (LP2). □

The final mechanism and analysis are the same as in Section 3.2: Solve (LP2), and making use of Assumption 3.18, decompose the LP solution into a convex combination of posted prices per edge, and sequentially post these prices for every bidder. To complete the analysis, note that Lemma 3.13 only requires that the distribution $g_{ij}(\cdot)$ be regular. This shows the following theorem:

Theorem 3.21. *There is a polynomial time $O(1)$-approximation to the revenue of the optimal adaptive posted-price scheme, when the bidders are monotone optimizers, and the valuations of different bidders for different items are drawn from mutually independent regular distributions.*

Acknowledgments: We thank Shuchi Chawla, Vincent Conitzer, and Peng Shi for helpful discussions, and the anonymous reviewers for detailed reviews.

References

[1] ZOË ABRAMS: Revenue maximization when bidders have budgets. In *Proc. 17th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'06)*, pp. 1074–1082. ACM Press, 2006. [doi:10.1145/1109557.1109676] 434

[2] AARON ARCHER, CHRISTOS PAPADIMITRIOU, KUNAL TALWAR, AND ÉVA TARDOS: An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2004. Preliminary version in SODA'03. [doi:10.1080/15427951.2004.10129086] 437

[3] LAWRENCE M. AUSUBEL: An efficient ascending-bid auction for multiple objects. *American Economic Review*, 94(5):1452 – 1475, 2004. [doi:10.1257/0002828043052330] 434

[4] LAWRENCE M. AUSUBEL AND PAUL R. MILGROM: Ascending auctions with package bidding. *Frontiers of Theoretical Economics*, 1(1):1, 2002. de Gruyter. 435

[5] JEAN-PIERRE BENOÎT AND VIJAY KRISHNA: Multiple-object auctions with budget constrained bidders. *Review of Economic Studies*, 68(1):155 – 179, 2001. [doi:10.1111/1467-937X.00164] 434, 435

[6] SAYAN BHATTACHARYA, VINCENT CONITZER, KAMESH MUNAGALA, AND LIRONG XIA: Incentive compatible budget elicitation in multi-unit auctions. In *Proc. 21st Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'10)*, pp. 554–572. ACM Press, 2010. [ACM:1873648] 434, 435, 436

[7] KIM C. BORDER: Reduced form auctions revisited. *Economic Theory*, 31(1):167–181, 2007. [doi:10.1007/s00199-006-0080-z] 434

[8] CHRISTIAN BORGS, JENNIFER CHAYES, NICOLE IMMORLICA, MOHAMMAD MAHDIAN, AND AMIN SABERI: Multi-unit auctions with budget-constrained bidders. In *Proc. 6th ACM Conf. Electronic Commerce (ACM-EC'05)*, pp. 44–51. ACM Press, 2005. [doi:10.1145/1064009.1064014] 434

[9] PATRICK BRIEST: Uniform budgets and the envy-free pricing problem. In *Proc. 35th Internat. Colloq. on Automata, Languages and Programming (ICALP'08)*, pp. 808–819. Springer, 2008. [doi:10.1007/978-3-540-70575-8_66] 433, 440

[10] PATRICK BRIEST, SHUCHI CHAWLA, ROBERT KLEINBERG, AND S. MATTHEW WEINBERG: Pricing randomized allocations. In *Proc. 21st Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'10)*, pp. 585–597. ACM Press, 2010. [ACM:1873601.1873650] 433

[11] SANDRO BRUSCO AND GIUSEPPE LOPOMO: Budget constraints and demand reduction in simultaneous ascending-bid auctions. *Journal of Industrial Economics*, 56(1):113–142, 2008. [doi:10.1111/j.1467-6451.2008.00335.x] 430, 434, 435

[12] SHUCHI CHAWLA, JASON D. HARTLINE, AND ROBERT KLEINBERG: Algorithmic pricing via virtual valuations. In *Proc. 8th ACM Conf. Electronic Commerce (ACM-EC'07)*, pp. 243–251. ACM Press, 2007. [doi:10.1145/1250910.1250946] 433, 435

[13] SHUCHI CHAWLA, JASON D. HARTLINE, DAVID L. MALEC, AND BALASUBRAMANIAN SIVAN: Multi-parameter mechanism design and sequential posted pricing. In *Proc. 42nd STOC*, pp. 311–320. ACM Press, 2010. [doi:10.1145/1806689.1806733] 433

[14] YEON-KOO CHE AND IAN GALE: Expected revenue of all-pay auctions and first-price sealed-bid auctions with budget constraints. *Economics Letters*, 50:373–379, 1996. [doi:10.1016/0165-1765(95)00766-0] 434

[15] YEON-KOO CHE AND IAN GALE: The optimal mechanism for selling to a budget-constrained buyer. *J. Economic Theory*, 92(2):198–233, 2000. [doi:10.1006/jeth.1999.2639] 434

[16] NING CHEN, NICOLE IMMORLICA, ANNA R. KARLIN, MOHAMMAD MAHDIAN, AND ATRI RUDRA: Approximating matches made in heaven. In *Proc. 36th Internat. Colloq. on Automata, Languages and Programming (ICALP'09)*, pp. 266–278. Springer, 2009. [doi:10.1007/978-3-642-02927-1_23] 435

[17] EDWARD H. CLARKE: Multipart pricing of public goods. *Public Choice*, 11(1):17–33, 1971. [doi:10.1007/BF01726210] 430

[18] SHAHAR DOBZINSKI, RON LAVI, AND NOAM NISAN: Multi-unit auctions with budget limits. *Games and Economic Behavior*, 74(2):486–503, 2012. Preliminary version in FOCS'08. [doi:10.1016/j.geb.2011.08.003] 434

[19] BENJAMIN EDELMAN, MICHAEL OSTROVSKY, AND MICHAEL SCHWARZ: Internet advertising and the generalized second price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007. [doi:10.1257/aer.97.1.242] 435

[20] EDITH ELKIND AND SHAHEEN FATIMA: Maximizing revenue in sequential auctions. In *Proc. 3rd Internat. Workshop on Internet and Network Economics (WINE'07)*, pp. 491–502. Springer, 2007. [doi:10.1007/978-3-540-77105-0_53] 435

[21] ANDREW V. GOLDBERG, JASON D. HARTLINE, ANNA R. KARLIN, MICHAEL SAKS, AND ANDREW WRIGHT: Competitive auctions. *Games and Economic Behavior*, 55(2):242–269, 2006. [doi:10.1016/j.geb.2006.02.003] 434

[22] ANDREW V. GOLDBERG, JASON D. HARTLINE, AND ANDREW WRIGHT: Competitive auctions and digital goods. In *Proc. 12th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'01)*, pp. 735–744. ACM Press, 2001. [ACM:365411.365768] 434

[23] THEODORE GROVES: Incentives in teams. *Econometrica*, 41(4):617–631, 1973. JSTOR. 430

[24] VENKATESAN GURUSWAMI, JASON D. HARTLINE, ANNA R. KARLIN, DAVID KEMPE, CLAIRE KENYON, AND FRANK MCSHERRY: On profit-maximizing envy-free pricing. In *Proc. 16th Ann. ACM-SIAM Symp. on Discrete Algorithms (SODA'05)*, pp. 1164–1173. ACM Press, 2005. [ACM:1070432.1070598] 432, 433

[25] ISA EMIN HAFALIR, R. RAVI, AND AMIN SAYEDI: Sort-cut: A pareto optimal and semi-truthful mechanism for multi-unit auctions with budget-constrained bidders. Technical report, 2009. [arXiv:0903.1450v1] 434

[26] JASON D. HARTLINE AND TIM ROUGHGARDEN: Simple versus optimal mechanisms. In *Proc. 10th ACM Conf. on Electronic Commerce (ACM-EC'09)*, pp. 225–234. ACM Press, 2009. [doi:10.1145/1566374.1566407] 435

[27] JEAN-JACQUES LAFFONT AND JACQUES ROBERT: Optimal auction with financially constrained buyers. *Economics Letters*, 52(2):181–186, 1996. [doi:10.1016/S0165-1765(96)00849-X] 434

[28] RON LAVI AND CHAITANYA SWAMY: Truthful and near-optimal mechanism design via linear programming. *J. ACM*, 58(6/25):1–24, 2011. Preliminary version in FOCS'05. [doi:10.1145/2049697.2049699] 437

[29] ALEJANDRO M. MANELLIA AND DANIEL R. VINCENT: Multidimensional mechanism design: Revenue maximization and the multiple-good monopoly. *J. Economic Theory*, 137(1):153–185, 2007. [doi:10.1016/j.jet.2006.12.007] 434

[30] ROGER B. MYERSON: Optimal auction design. *Math. Oper. Res.*, 6(1):58–73, 1981. [doi:10.1287/moor.6.1.58] 431, 433, 434, 444, 446

[31] NOAM NISAN, JASON BAYER, DEEPAK CHANDRA, TAL FRANJI, ROBERT GARDNER, YOSSI MATIAS, NEIL RHODES, MISHA SELTZER, DANNY TOM, HAL VARIAN, AND DAN ZIGMOND: Google's auction for TV ads. In *Proc. 36th Internat. Colloq. on Automata, Languages and Programming (ICALP'09)*, pp. 309–327. Springer, 2009. See also ESA'09 and SODA'10. [doi:10.1007/978-3-642-02930-1_26] 430, 434

[32] MALLESH M. PAI AND RAKESH VOHRA: Optimal auctions with financially constrained bidders. *Working Paper*, 2008. CiteSeerX. 434

[33] JOHN THANASSOULIS: Haggling over substitutes. *J. Economic Theory*, 117(2):217–245, 2004. [doi:10.1016/j.jet.2003.09.002] 434

[34] WILLIAM VICKREY: Counterspeculation, auctions and competitive sealed tenders. *J. Finance*, 16(1):8–37, 1961. [doi:10.1111/j.1540-6261.1961.tb02789.x] 430

[35] ROBERT B. WILSON: *Nonlinear Pricing*. Oxford University Press, 1997. 434

AUTHORS

Sayan Bhattacharya
Graduate Student
Department of Computer Science
Duke University, Durham, NC 27708
bsayan@cs.duke.edu
http://www.cs.duke.edu/~bsayan

Gagan Goel
Research Scientist
Google Inc.
New York, NY
gagangoel@google.com
http://research.google.com/pubs/GaganGoel.html

Sreenivas Gollapudi
Senior Researcher
Microsoft Search Labs
Mountain View, CA 94043
sreenig@microsoft.com
http://research.microsoft.com/en-us/people/sreenig/

Kamesh Munagala
Associate Professor
Department of Computer Science
Duke University, Durham, NC 27708
kamesh@cs.duke.edu
http://www.cs.duke.edu/~kamesh

ABOUT THE AUTHORS

SAYAN BHATTACHARYA is a Ph. D. student in the Department of Computer Science at Duke University. His adviser is Kamesh Munagala. He obtained his B. E. in Computer Science and Engineering from Jadavpur University, Kolkata in July 2008. He is interested in Approximation Algorithms and Computational Microeconomics.

GAGAN GOEL is a research scientist at Google. His research interests lie in the design and analysis of algorithms and its applications to optimization, market design, and game theory. He received his Ph. D. in Algorithms, Combinatorics, and Optimization from Georgia Tech in August 2009 under the supervision of Vijay Vazirani. Before that he received his B. Tech. in Computer Science and Engineering from the Indian Institute of Technology, Delhi in August 2004.

SREENIVAS GOLLAPUDI is a Senior Researcher at Microsoft Search Labs which he joined in July 2006. Prior to moving to Search Labs, he held senior engineering roles at Ebrary in their Advanced Technologies Group and in the distributed database group at Oracle. He obtained his B. Tech. from IIT Bombay, and his Ph. D., in 2004, from the State University of New York at Buffalo where his advisor was Aidong Zhang. He is a holder of twenty-nine US patents. His current research interests include web algorithms, algorithms for large data sets, and game theoretic approaches to social network analysis.

KAMESH MUNAGALA is Associate Professor of Computer Science Science at Duke University, where he has been employed since 2004. He received his M. S. degree in 2002 and his Ph. D. in 2003, both from Stanford University. His advisor was Serge Plotkin. He received his B. Tech. degree in 1998 from the Indian Institute of Technology, Bombay where his advisor was Abhiram Ranade. He spent a year as a post-doctoral scholar in Pat Brown's Lab doing computational biology at the Department of Biochemistry, Stanford University School of Medicine. His research interests are in approximation algorithms, combinatorial optimization, computational economics, scheduling theory, social networks, and data mining. He received an NSF CAREER award in 2008 and an Alfred P. Sloan research fellowship in 2009.