For a given choice of n and q let now the set of vectors $\mathcal{V}_q^n = \{v_1, \ldots, v_{(\ell(q))}\}$ with $\ell(q) = \Omega(q^{(n-1)/2})$ and $v_i \cdot v_j \leq 1/n - 1/q$ for any $i \neq j$ be given. We will define a unit-demand pricing instance based on these vectors. Consumer distribution $\mathcal{C}$ will be defined as a finite support distribution. For each v_j we define a consumer type with valuation vector $\tilde{v}_j = 2^j \cdot v_j$ and probability $\mu_j = 2^{-j}$. Since $||v_j||_1 \leq \sqrt{n}||v_j||_2 = 1$, vectors from $\mathcal{V}_q^n$ can also be interpreted as lotteries. We define lotteries λ_j with probability vectors $\phi_j = v_j$ and assign them price $p_j = (1/q) \cdot 2^j$.

Consider the utility $u(\tilde{v}_j, \lambda_j)$ of consumer type $\tilde{v}_j$ when purchasing lottery λ_j. We may write that

$$\begin{aligned} u(\tilde{v}_j, \lambda_j) &= \tilde{v}_j \cdot \phi_j - (1/q) \cdot 2^j \\ &= 2^j (v_j)^2 - (1/q) \cdot 2^j = (1/n - 1/q)2^j, \end{aligned}$$

by the fact that $||v_j||_2 = 1/\sqrt{n}$. On the other hand, the consumer type's utility from any other lottery λ_i is bounded above by

$$u(\tilde{v}_j, \lambda_i) = \tilde{v}_j \cdot \phi_i = (v_j \cdot v_i)\, 2^j \leq (1/n - 1/q)2^j,$$

since $v_j \cdot v_i \leq 1/n - 1/q$ for all $v_j, v_i \in \mathcal{V}_q^n$. Thus, given the lottery system defined above each consumer $\tilde{v}_j$ will choose to purchase lottery λ_j and we obtain revenue

$$\sum_{j=1}^{\ell(q)} 2^{-j} \cdot (1/q) \cdot 2^j = \ell(q)/q = \Omega(q^{(n-3)/2}),$$

using that $\ell(q) = \Omega(q^{(n-1)/2})$. It remains to estimate the optimal item pricing revenue. Consider a single item priced at $p \in \mathbb{R}_+$ and all other items priced at $+\infty$. For $2^{k-1} < p \leq 2^k$ consumer types $\tilde{v}_1, \ldots, \tilde{v}_{k-1}$ surely have valuations of less than p for all items. It follows that the total probability mass of consumers who are able to afford the item is bounded above by $\sum_{j=k}^{\ell(q)} 2^{-j} \leq 2^{-k+1}$ and, thus, total revenue is at most $p \cdot 2^{-k+1} \leq 2$. It follows that the optimal item pricing results in revenue of at most $2n$ and for any constant n we obtain a lower bound of $\Omega(q^{(n-3)/2})$ on the revenue gap. In particular, we can make this gap arbitrarily large by choosing a sufficiently large q in any dimension $n \geq 4$. □

4 Results in the Buy-Many Model

As we have argued before, the assumption made in the buy-one model that a consumer purchases a single utility maximizing lottery from any given system is not justifiable in general. We now continue by investigating the more realistic buy-many model, in which we allow consumers to buy any combination of lotteries maximizing their expected utility. As we will see in Section 4.1, this reduces the advantage lottery systems have over pure item pricings drastically. In particular, this implies that known inapproximability results for the item pricing problem yield similar bounds in the lottery setting and algorithmic results similar to the buy-one model cannot be obtained. In Section 4.2 we prove that our bound on the revenue gap is asymptotically tight.

4.1 Upper Bound and Hardness of Approximation Let an arbitrary system of lotteries in the buy-many model over n distinct items be given. We assume throughout this section that the utility maximizing collection of lotteries for each consumer type given this system consists of a single lottery. This assumption is w.l.o.g., as we can add lotteries corresponding to the joint distribution of some collection of lotteries to the system until it holds. The following randomized algorithm turns the lottery system into a pure item pricing:

(1) For each item i, let p_i be the price of the cheapest lottery with probability at least $1/(130n^3)$ for item i ($p_i = +\infty$ if no such lottery exists).

(2) With probability 1/2, uniformly sample t from $\{-1, 0, \ldots, 3\lfloor \log n \rfloor + 9\}$ and assign price $2^t p_i$ to every item i.

(3) Else sample a single item i uniformly at random. Assign price $+\infty$ to all items other than i. Price item i at $130n^3 \mathrm{e}^j p_i$ with probability $(1 - 1/\mathrm{e})\mathrm{e}^{-j}$ for all $j \in \mathbb{N}_0$.

The core idea of the algorithm is the following: Every lottery with some minimum probability of allocating some specific item defines an upper bound on the payment of consumer types prefering this item, since by buying multiple copies of the lottery at hand they can make the probability of receiving the desired item approach 1 exponentially fast. Thus, for each item we let the cheapest lottery with some minimum probability for it define its *base price* and assign a random item price via a carefully tailored two stage stochastic process. We are going to argue that the above algorithm outputs an item pricing that is an expected $\mathcal{O}(\log n)$-approximation to the revenue of an optimal lottery system in the buy-many model. Note, that the algorithm is easily derandomized via exhaustive search over the entire range of (relevant) random coin flips. Throughout this section we assume w.l.o.g. that $\sum_{i=1}^n \phi_i = 1$ for every lottery. This is easily achieved by adding a dummy item valued at 0 by all consumers to the instance.

THEOREM 4.1. *Given a distribution $\mathcal{C}$ of unit-demand consumers and an optimal lottery system in the buy-many model, the above algorithm returns an item pricing with revenue $r \geq 1/\mathcal{O}(\log n) \cdot r_L^*(\mathcal{C})$.*

It is known that the unit-demand item pricing problem cannot be approximated within $\mathcal{O}(n^{\varepsilon})$ for some $\varepsilon > 0$ [5]. Thus, we immediately obtain the following inapproximability result for lottery pricing in the buy-many model.

COROLLARY 4.1. *The unit-demand lottery pricing problem in the buy-many model cannot be approximated within $\mathcal{O}(n^{\varepsilon})$ for some $\varepsilon > 0$, unless $NP \subseteq \bigcap_{\delta>0}$ $BPTIME(2^{\mathcal{O}(n^{\delta})})$.*

Proof of Theorem 4.1: We will show that in going from the optimal lottery system to a pure item pricing, the expected loss in revenue is bounded by $\mathcal{O}(\log n)$ for every single consumer type. So let a single consumer type from $\mathcal{C}$ with values $(v_1, \ldots, v_n)$ be given. Furthermore, assume that this consumer buys a lottery with probabilities $(\phi_1, \ldots, \phi_n)$ and price p when offered the optimal lottery system in the buy-many model. By $j = \text{argmax}\{v_i \mid i : \phi_i \geq 1/(130n^3)\}$ we denote the consumer's favorite item among those for which he has at least a $1/(130n^3)$-chance of receiving them.

Finally, recall that for every item i, p_i denotes the price of the cheapest lottery that has probability at least $1/(130n^3)$ for item i. We start with the observation that in any optimal lottery system a consumer's utility does not depend significantly on items she receives with negligible probability.

PROPOSITION 4.1. *We may assume without loss of generality that $\sum_{i=1}^{n} \phi_i(v_i - v_j) \leq p/4$.*

Proof. Let $\mathcal{C}'$ be the set of all consumer types for which the above does not hold. In particular, for each consumer type in $\mathcal{C}'$, if we let j again denote her favorite item with probability at least $1/(130n^3)$ in the lottery $(\phi_1, \ldots, \phi_n)$ she buys at price p, there exists an item k with $\phi_k(v_k - v_j) > p/(4n)$. Let $\mathcal{C}'_k$ be the set of all consumer types for which item k satisfies this inequality and consider a single class $\mathcal{C}'_k$.

Remove all lotteries except for the ones bought by consumers in $\mathcal{C}'_k$. For the remaining lotteries, set their probabilities for all items other than k to 0 and reduce their prices by a factor of $8n$. Now consider a single consumer type in $\mathcal{C}'_k$ with values $(v_1, \ldots, v_n)$ buying lottery $(\phi_1, \ldots, \phi_n)$ at price p in the original lottery system and favorite item j among those with minimum probability $1/(130n^3)$. We want to lower bound the revenue from this consumer given the modified system of lotteries.

If the consumer purchases the modified version of the lottery she bought originally, revenue has decreased by at most a factor of $8n$. If she does not, she now buys some other lottery (or a combination of lotteries) with some probability μ_k for item k and price $q/8n$. With lottery $(\phi_1, \ldots, \phi_n)$, the consumer had a chance of at least $1 - (n-1)/(130n^3) \geq 1 - 1/(130n^2)$ of receiving an item valued at v_j or less. Thus, the marginal utility of adding a copy of $(\mu_1, \ldots, \mu_n)$ at its original price q would have been $\mu_k(1 - 1/(130n^2))(v_k - v_j) - q$. Since the consumer chooses not to buy a copy, we have

$$\mu_k \left(1 - \frac{1}{130n^2}\right)(v_k - v_j) - q \leq 0.$$

We know that $\phi_k v_k \geq \phi_k(v_k - v_j) \geq p/(4n)$ and, thus, buying the modified version of lottery $(\phi_1, \ldots, \phi_n)$ at price $p/(8n)$ yields utility $\phi_k v_k - p/(8n) \geq \frac{1}{2}\phi_k v_k$. Thus, it must be the case that $\mu_k \geq \phi_k/2$. Finally, we obtain

$$\begin{aligned} q &\geq \mu_k \left(1 - \frac{1}{130n^2}\right)(v_k - v_j) \\ &\geq \frac{\phi_k}{2}\left(1 - \frac{1}{130n^2}\right)(v_k - v_j) \\ &\geq \frac{1}{8n}\left(1 - \frac{1}{130n^2}\right)p. \end{aligned}$$

Thus, if the consumer purchases at price $q/8n$ the reduction in revenue is bounded below by $1/(64n^2)(1 - 1/(130n^2)) \geq 1/(65n^2)$. Now observe that all lotteries in our modified lottery system have probability at most $1/(130n^3)$. If we multiply both probabilities and prices of all lotteries by $130n^3$, the consumer's utility from the lottery she currently purchases remains unchanged, while the utility obtainable by buying any other bundle of lotteries cannot increase and, thus, this does not affect the consumer's buying decision and effectively increases the revenue from every consumer in $\mathcal{C}'_k$ by a factor of $(130n^3)/(65n^2) = 2n$ compared to the original optimal lottery system.

Now assume that more than half the revenue of the original optimal lottery system was due to consumer types in $\mathcal{C}'$. Then there must exist a class $\mathcal{C}'_k$ that carries more than a $1/(2n)$-fraction of the overall revenue, which we have just shown how to increase by a factor of $2n$, a contradiction. Hence, at most half the revenue stems from consumer types in $\mathcal{C}'$ and we may ignore these consumer types. □

Consider the price assignments defined in Step (2) of the algorithm for different values of $t \in \{-1, 0, \ldots, 3\lfloor \log n \rfloor + 9\}$. We know that

$$\begin{aligned} v_j &= \sum_{i=1}^{n} \phi_i v_i - \sum_{i=1}^{n} \phi_i(v_i - v_j) \\ &\geq p - \frac{p}{4} = \frac{3}{4}p \end{aligned}$$

and, thus, for $t = -1$ the consumer has strictly positive utility from buying j and will consequently purchase some item. Denote by i_0 the item bought for $t = -1$. For increasing values of t, the consumer might switch to other items that yield higher utility. Refer to these items as $i_1, \ldots, i_\ell$ in the order they are bought.

FACT 4.1. *It holds that* $p_{i_0} > p_{i_1} > \cdots > p_{i_\ell}$ *and* $v_{i_0} \geq v_{i_1} \geq \cdots \geq v_{i_\ell}$.

Proof. Note that the consumer buys only if this results in non-negative utility. In going from t to $t+1$ the utility from buying any item i decreases by $(2^{t+1} - 2^t)p_i = 2^t p_i$ and, thus, decreases strictly less on cheaper items. It follows that $p_{i_0} > p_{i_1} > \cdots > p_{i_\ell}$.

Assume then that $p_{i_j} > p_{i_{j+1}}$, but $v_{i_j} < v_{i_{j+1}}$ for some i_j, i_{j+1}. Then buying item i_{j+1} yields strictly higher utility than i_j for any value of t and i_j is never bought. ☐

We proceed by deriving an upper bound on the consumer's utility given the original lottery system. For $t = -1$, her utility from the item pricing is at least $v_j - p_j/2 \geq v_j - p/2$, since $\phi_j \geq 1/(130n^3)$ and, thus, $p_j \leq p$ by definition. Since she decides to purchase item i_0, it must be the case that $v_{i_0} \geq v_j - p/2$. We may then write that

$$\begin{aligned}\sum_{i=1}^{n} \phi_i v_i - p &= \sum_{i=1}^{n} \phi_i (v_i - v_j) + v_j - p \\ &\leq v_j - \frac{3}{4}p \text{ , since } \sum_{i=1}^{n} \phi_i (v_i - v_j) \leq \frac{p}{4} \\ &\leq v_{i_0} - \frac{p}{4} \text{ , since } v_j \leq v_{i_0} + \frac{p}{2}. \qquad (4.2)\end{aligned}$$

Next, we are going to bound the utility the consumer can achieve by focusing on item i_ℓ from below. Let $(\mu_1, \ldots, \mu_n)$ be the lottery defining price p_{i_ℓ}, i.e., the cheapest lottery in the original system with $\mu_{i_\ell} \geq 1/(130n^3)$. Now consider a strictly worse lottery with probability exactly $1/(130n^3)$ for item i_ℓ, probability 0 for all other items and price p_{i_ℓ}. Assume that this was the only available lottery and let k denote the number of copies our consumer would choose to purchase. Then, by the fact that the $(k+1)$-th copy does not yield positive marginal utility for her, we may conclude that $1/(130n^3)(1 - 1/(130n^3))^k v_{i_\ell} - p_{i_\ell} \leq 0$ and, thus, $(1 - 1/(130n^3))^k v_{i_\ell} \leq 130n^3 p_{i_\ell}$. Consequently, the utility from buying k copies is at least

$$\begin{aligned}&\left(1 - \left(1 - \frac{1}{130n^3}\right)^k\right) v_{i_\ell} - kp_{i_\ell} \\ \geq\ &\max\Big\{v_{i_\ell} - (k + 130n^3)p_{i_\ell} \text{ , } 0\Big\}. \qquad (4.3)\end{aligned}$$

By the fact that a strictly better lottery was part of the original lottery system, combining (4.2) and (4.3) yields

$$\begin{aligned}v_{i_0} - \frac{p}{4} &\geq \sum_{i=1}^{n} \phi_i v_i - p \\ &\geq \max\Big\{v_{i_\ell} - (k + 130n^3)p_{i_\ell} \text{ , } 0\Big\}.\end{aligned}$$

Finally, rearranging for p we obtain

$$p \leq 4\left(v_{i_0} - v_{i_\ell} + \min\Big\{(k + 130n^3)p_{i_\ell} \text{ , } v_{i_\ell}\Big\}\right) \qquad (4.4)$$

as an upper bound on the price paid by the consumer given the original lottery system.

We proceed by proving a lower bound on the expected price paid given the item pricing returned by our algorithm. We distinguish the following three cases.

Case (1): $v_{i_\ell} \leq (k + 130n^3)p_{i_\ell}$. Let t_j be the highest value of t in Step (2) of the algorithm at which the consumer purchases item i_j. Observe that as long as item i_j is priced at $v_{i_j} - v_{i_{j+1}}$ or less it yields higher utility than item i_{j+1}. It follows that $2^{t_j+1} p_{i_j} \geq v_{i_j} - v_{i_{j+1}}$ or, equivalently,

$$2^{t_j} p_{i_j} \geq (1/2)(v_{i_j} - v_{i_{j+1}}).$$

Using that $v_{i_\ell} \leq (k + 130n^3)p_{i_\ell}$ and $v_{i_\ell} > 130n^3(1 - 1/(130n^3))^{-(k-1)} p_{i_\ell}$ (by the fact that the k-th copy of the lottery with probability $1/(130n^3)$ for item i_ℓ at price p_{i_ℓ} has positive marginal utility), it readily follows that $v_{i_\ell} \leq 260n^3 p_{i_\ell}$. In particular, this implies that the range of t in Step (2) of the algorithm includes $\lfloor \log v_{i_\ell} \rfloor$. Let now R denote the the price paid by the consumer given the item pricing and note that Step (2) of the algorithm is performed with probability 1/2. We have

$$\begin{aligned}\mathbb{E}[R] &\geq \frac{1}{2}\left(\sum_{j=0}^{\ell-1} \text{Prob}(t = t_j)\frac{1}{2}(v_{i_j} - v_{i_{j+1}}) \right. \\ &\qquad \left. + \text{Prob}(t = \lfloor \log v_{i_\ell} \rfloor)\frac{v_{i_\ell}}{2}\right) \\ &\geq \frac{1}{12\lfloor \log n \rfloor + 44}\left(\sum_{j=0}^{\ell-1}(v_{i_j} - v_{i_{j+1}}) + v_{i_\ell}\right) \\ &= \frac{1}{12\lfloor \log n \rfloor + 44} v_{i_0}.\end{aligned}$$

Case (2): $v_{i_\ell} > (k + 130n^3)p_{i_\ell}$, $k \leq 130n^3 + 2n$. In this case we know that $\lfloor \log k + 130n^3 \rfloor$ lies within the range of t in Step (2) of the algorithm. Similar to Case (1) above we obtain

$$\mathbb{E}[R] \geq \frac{1}{12\lfloor \log n \rfloor + 44}\Big(v_{i_0} - v_{i_\ell} + (k + 130n^3)p_{i_\ell}\Big).$$

Case (3): $v_{i_\ell} > (k + 130n^3)p_{i_\ell}$, $k > 130n^3 + 2n$. Once more, recall that our consumer has positive marginal utility from buying the k-th copy of a lottery that offers a $1/(130n^3)$-chance of receiving item i_ℓ at price p_{i_ℓ}. Thus,

$$\begin{aligned} v_{i_\ell} &\geq 130n^3\Big(1 - \frac{1}{130n^3}\Big)^{-(k-1)} p_{i_\ell} \\ &\approx 130n^3 e^{k-130n^3-1} p_{i_\ell}, \end{aligned}$$

where the above holds with arbitrary precision for large values of n. Let A_{i_ℓ} denote the event that the algorithm chooses to perform the random experiment in Step (3) and picks item i_ℓ to assign a price different from $+\infty$ to. For the expected payment of our consumer conditioned on A_{i_ℓ} we have

$$\begin{aligned} &\mathbb{E}[R \mid A_{i_\ell}] \\ &\geq \sum_{j=0}^{k-130n^3-1} \Big(1 - \frac{1}{e}\Big) e^{-j} 130n^3 e^j p_{i_\ell} \\ &= (k - 130n^3)\Big(1 - \frac{1}{e}\Big) 130n^3 p_{i_\ell} \\ &= (nk + (130n^3 - n)k - 130^2 n^6)\Big(1 - \frac{1}{e}\Big) p_{i_\ell} \\ &\geq \Big(1 - \frac{1}{e}\Big) nk p_{i_\ell}. \end{aligned}$$

Finally, since event A_{i_ℓ} has probability $1/(2n)$, we obtain

$$\begin{aligned} \mathbb{E}[R] &\geq \frac{1}{2}\left(\sum_{j=0}^{\ell-1} \text{Prob}(t = t_j) \frac{1}{2}(v_{i_j} - v_{i_{j+1}})\right) \\ &\quad + \frac{1}{2n}\mathbb{E}[R \mid A_{i_\ell}] \\ &\geq \frac{1}{12\lfloor \log n \rfloor + 44}(v_{i_0} - v_{i_\ell}) + \frac{e-1}{2e} k p_{i_\ell} \\ &\geq \frac{1}{12\lfloor \log n \rfloor + 44}\Big(v_{i_0} - v_{i_\ell} + (k + 130n^3)p_{i_\ell}\Big). \end{aligned}$$

So we have that $\mathbb{E}[R] = 1/\mathcal{O}(\log n) \cdot p$ in each case, which finishes the proof. □

4.2 A Lower Bound Finally, we will show that the bound derived in the previous section is tight. In fact, it turns out that this is even true for the restricted case of consumers with uniform valuations. This is somewhat surprising, as this distinction is quite significant in the buy-one model, as we have seen before.

THEOREM 4.2. *For all n there exist uniform-valuation consumer distributions $\mathcal{C}$ with $r_L^*(\mathcal{C})/r^*(\mathcal{C}) = \Omega(\log n)$.*

Proof of Theorem 4.2: Let $n \in \mathbb{N}$ be prime and $\mathcal{P}$ denote the set of all distinct polynomials of degree 2 over the field $\mathbb{Z}/n\mathbb{Z}$. We identify each polynomial $P \in \mathcal{P}$ with the set

$$S_P = \big\{(0, P(0)), \ldots, (n-1, P(n-1))\big\} \subset \big(\mathbb{Z}/n\mathbb{Z}\big)^2.$$

Observe that $|\mathcal{P}| = n^3$ and $|S_P| = n$, $|S_P \cap S_Q| \leq 2$ for all $P \neq Q \in \mathcal{P}$ by the fact that polynomials of maximum degree 2 over $\mathbb{Z}/n\mathbb{Z}$ for n prime have at most 2 zeroes.

We define a random pricing instance based on these polynomials as follows. The set of items corresponds to the elements of $(\mathbb{Z}/n\mathbb{Z})^2$. Let $k = \lfloor \log n \rfloor$. For each $P \in \mathcal{P}$ we define a corresponding class $\mathcal{C}_P$ of 2^{k-j} identical consumers with a non-zero value $v_P = 2^{j-k}$ for the items in S_P, where j is drawn uniformly at random from $\{1, \ldots, k\}$. Let $\mathcal{C} = \bigcup_P \mathcal{C}_P$ denote the complete instance.

The proof of the lower bound proceeds in two steps. We first argue that with non-zero probability our random experiment creates a pricing instance on which every pure item pricing yields revenue $\mathcal{O}(n^3/\log n)$. We then show that for every instance created by the experiment we can find a lottery system that yields revenue $\Omega(n^3)$.

LEMMA 4.1. *Let $\mathcal{C}$ be a random pricing instance as defined above. It holds that $r^*(\mathcal{C}) = \mathcal{O}(n^3/\log n)$ with positive probability.*

Proof. Let p be a price vector resulting in revenue r on some instance created by our random experiment. Then there must exist another price vector p' that assigns only prices from $\{2^{1-k}, 2^{2-k}, \ldots, 2^0\}$ and makes revenue at least $r/2$ by sales to consumers buying at a price equal to their value. To see this, note that it is w.l.o.g. to assume that prices are powers of 2 and as long as more than half the revenue generated by a price vector p comes from consumers buying at a price no more than half their values, we can increase revenue by doubling all prices. Since the identical consumers in each class $\mathcal{C}_P$ contribute a total revenue of 1 if they buy at their full values, we only need to prove that with positive probability no price vector extracts the full value of consumers in more than $\mathcal{O}(n^3/\log n)$ different classes $\mathcal{C}_P$.

Consider a fixed price vector p and let T_j denote the set of items priced at 2^{j-k} for $j = 1, \ldots, k$. By $\mathcal{B}_j$ we denote the set of consumer classes with a non-zero value for at least one item in T_j and value zero for all items in $T_1, \ldots, T_{j-1}$, formally,

$$\mathcal{B}_j = \{\mathcal{C}_P \mid S_P \cap T_j \neq \emptyset\} \setminus \Big(\bigcup_{i=1}^{j-1} \mathcal{B}_i\Big).$$

A set $\mathcal{C}_P \in \mathcal{B}_j$ of consumers yields revenue 1 if and only if their random value is $v_P = 2^{j-k}$. Let random variable Y_j denote the number of consumer classes in $\mathcal{B}_j$ that pay their full values given price vector p. We will bound Y_j in two steps. If $|\mathcal{B}_j| \leq n^3/(\log n)^2$, then it trivially holds that $Y_j \leq |\mathcal{B}_j| \leq n^3/(\log n)^2$. If $|\mathcal{B}_j| > n^3/(\log n)^2$, let X_P^j be a random indicator variable with $X_P^j = 1$ if $v_P = 2^{j-k}$ and $X_P^j = 0$ else. Thus, $\text{Prob}(X_P^j = 1) = 1/k$. It follows that $Y_j = \sum_{P:\mathcal{C}_P \in \mathcal{B}_j} X_P^j$ and $\mathbb{E}[Y_j] = \frac{1}{k}|\mathcal{B}_j|$. Applying the Chernoff bound [14] and using that $|\mathcal{B}_j| > n^3/(\log n)^2$ and $k \leq \log n$ we obtain

$$\begin{aligned} \text{Prob}\left(Y_j \geq \frac{6}{k}|\mathcal{B}_j|\right) &\leq 2^{-\frac{6}{k}|\mathcal{B}_j|} \\ &\leq 2^{-6\left(\frac{n}{\log n}\right)^3}. \end{aligned}$$

Thus, $Y_j \leq \max\{n^3/(\log n)^2, 12|\mathcal{B}_j|/(\log n)\}$ with probability at least $1-2^{-6\left(\frac{n}{\log n}\right)^3}$. Let $Y = \sum_j Y_j$. Applying the union bound yields

$$\begin{aligned} Y &= \sum_{j=1}^{k} Y_j \leq \sum_{j=1}^{k}\left(\frac{n^3}{(\log n)^2} + \frac{12|\mathcal{B}_j|}{\log n}\right) \\ &= k\frac{n^3}{(\log n)^2} + \frac{12\sum_{j=1}^{k}|\mathcal{B}_j|}{\log n} \leq \frac{13n^3}{\log n} \end{aligned}$$

with probability at least $1-2^{-5\left(\frac{n}{\log n}\right)^3}$. Finally, observe that there are at most k^{n^2} different price vectors with prices in $\{2^{1-k}, \ldots, 2^0\}$. Applying the union bound once again the probability that any of these extracts full values from more than $13n^3/\log n$ consumer classes is at most $(\log n)^{n^2} 2^{-5\left(\frac{n}{\log n}\right)^3} < 1$ for $n \geq 2$. □

LEMMA 4.2. *Let $\mathcal{C}$ be any pricing instance as defined above. It holds that $r_L^*(\mathcal{C}) = \Omega(n^3)$.*

Proof. We construct a system of lotteries as follows. For every consumer class $\mathcal{C}_P$ with value v_P for items in set S_P, we introduce a lottery ϕ_P with probability $1/n$ for each of the items in S_P at price $v_P/2$. A consumer from class $\mathcal{C}_P$ has utility $v_P - (1/2)v_P = (1/2)v_P$ from purchasing lottery ϕ_P. By the fact that $|S_P \cap S_Q| \leq 2$ for all $Q \neq P$, any other lottery ϕ_Q has probability of at most $2/n$ of allocating an item in set S_P. We also know that, since $k \leq \log n$, we have $v_Q \geq 2^{-k}v_P \geq (1/n)v_P$ and, thus, a consumer from $\mathcal{C}_P$ has utility at most $(2/n)v_P - (1/2)v_Q \leq (2/n)v_P - (1/(2n))v_P = (3/(2n))v_P$. Since the same bounds the marginal utility of any lottery other than ϕ_P being bought in addition to some other set of lotteries, it follows that consumers in $\mathcal{C}_P$ either purchase lottery ϕ_P at price $v_P/2$ or some set of at least $((1/2)v_P)/((3/(2n))v_P) = (1/3)n$ lotteries at price at least $(1/n)v_P$ each, resulting in overall revenue at least $(1/3)v_P$ from each of these consumers. Summing over all classes $\mathcal{C}_P$ and the consumers in each class yields overall revenue of $(1/3)n^3$. □

5 Acknowledgments

We are very grateful to Jason Hartline for numerous enlightening discussions in the early stages of this research. The result in Theorem 3.3 has independently been discovered by David Malec.

References

[1] Gagan Aggarwal, Tomás Feder, Rajeev Motwani, and An Zhu. Algorithms for Multi-Product Pricing. In *Proc. of the 31st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 253–266, 2004.

[2] Aaron Archer, Christos Papadimitriou, Kunal Talwar, and Éva Tardos. An Approximate Truthful Mechanism for Combinatorial Auctions with Single Parameter Agents. In *Proc. of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 205–214, 2003.

[3] Mark Armstrong. Price Discrimination by a Many-Product Firm. *Review of Economic Studies*, 66(1):151–68, 1999.

[4] Maria-Florina Balcan and Avrim Blum. Mechanism Design, Machine Learning, and Pricing Problems. *SIGecom Exch.*, 7(1):34–36, 2007.

[5] Patrick Briest. Uniform Budgets and the Envy-Free Pricing Problem. In *Proc. of the 35th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 808–819, 2008.

[6] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic Pricing via Virtual Valuations. In *Proc. of the 8th ACM Conference on Electronic Commerce (EC)*, pages 243–251, 2007.

[7] Peerapong Dhangwatnotai, Shahar Dobzinski, Shaddin Dughmi, and Tim Roughgarden. Truthful Approximation Schemes for Single-Parameter Agents. In *Proc. of the 49th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 15–24, 2008.

[8] Shahar Dobzinski, Noam Nisan, and Michael Schapira. Truthful Randomized Mechanisms for Combinatorial Auctions. In *Proc. of the 38th ACM Symposium on Theory of Computing (STOC)*, pages 644–652, 2006.

[9] Venkatesan Guruswami, Jason D. Hartline, Anna R. Karlin, David Kempe, Claire Kenyon, and Frank McSherry. On Profit-Maximizing Envy-Free Pricing. In *Proc. of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1164–1173, 2005.

[10] Jason Hartline and Anna Karlin. Profit Maximization in Mechanism Design. In Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay Vazirani, editors, *Algorithmic Game Theory*, chapter 13, pages 331–362. Cambridge University Press, 2007.

[11] Alejandro Manelli and Daniel Vincent. Bundling as an Optimal Selling Mechanism for a Multiple-Good

Monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.

[12] Alejandro Manelli and Daniel Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory*, 137(1):153–185, 2007.

[13] Preston McAfee and John McMillan. Multidimensional Incentive Compatibility and Mechanism Design. *Journal of Economic Theory*, 46(2):335–354, 1988.

[14] Michael Mitzenmacher and Eli Upfal. *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*. Cambridge University Press, 2005.

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

[16] John Riley and Richard Zeckhauser. Optimal Selling Strategies: When to Haggle, When to Hold Firm. *Quarterly J. Economics*, 98(2):267–289, 1983.

[17] Jean-Charles Rochet and Philippe Chone. Ironing, Sweeping, and Multidimensional Screening. *Econometrica*, 66(4):783–826, 1998.

[18] John Thanassoulis. Haggling Over Substitutes. *J. Economic Theory*, 117:217–245, 2004.

A Vector Packing

LEMMA 3.1 *Let $n \geq 1$ be given. For every $q \geq 2n$ there exists a set $\mathcal{V}_q^n$ of vectors in $S^{n+}_{1/\sqrt{n}}$, such that $v \cdot w \leq 1/n - 1/q$ for all $v, w \in \mathcal{V}_q^n$ with $v \neq w$ and $|\mathcal{V}_q^n| = \Omega(q^{(n-1)/2})$.*

Proof. For vectors $v, w \in S^{n+}_{1/\sqrt{n}}$ we may write that

$$v \cdot w = \frac{1}{2}\Big(v^2 + w^2 - (v-w)^2\Big) = \frac{1}{n} - \frac{1}{2}||v-w||_2^2.$$

Consequently, the condition that $v \cdot w \leq 1/n - 1/q$ for all $v, w \in \mathcal{V}_q^n$ is equivalent to asking that a ball of radius $\sqrt{2/q}$ around the tip of any vector from $\mathcal{V}_q^n$ does not contain the tip of any of the other vectors. For the remainder of this proof we will associate vectors in $\mathcal{V}_q^n$ with n-dimensional balls of radius $\sqrt{2/q}$ centered at their tips.

We construct the set $\mathcal{V}_q^n$ by the following simple greedy approach. While there exists a point in $S^{n+}_{1/\sqrt{n}}$ that is not covered by previously selected balls, we choose it as the center of a new ball of radius $\sqrt{2/q}$.

We now have to lower bound the number of vectors found in this fashion. When the procedure terminates, it must be the case that $S^{n+}_{1/\sqrt{n}}$ is completely covered by the selected balls. Choose any $\varepsilon > 0$. By B_r^n we denote the n-dimensional ball of radius r centered at the origin, B_r^{n+} its part in the all-positive orthant. If we increase the radius of all balls chosen by our packing procedure by $\varepsilon\sqrt{2/q}$, we know that they completely cover the set

$$B^{n+}_{1/\sqrt{n}+\varepsilon\sqrt{2/q}} - B^{n+}_{1/\sqrt{n}-\varepsilon\sqrt{2/q}},$$

as all points in this set are at distance at most $\varepsilon\sqrt{2/q}$ from $S^{n+}_{1/\sqrt{n}}$. The n-dimensional Lebesgue measure of this set is

$$\begin{aligned}&\lambda^n\Big(B^{n+}_{1/\sqrt{n}+\varepsilon\sqrt{2/q}} - B^{n+}_{1/\sqrt{n}-\varepsilon\sqrt{2/q}}\Big)\\ = \;& 2^{-n}\Big((\sqrt{1/n}+\varepsilon\sqrt{2/q})^n \\ &\qquad -(\sqrt{1/n}-\varepsilon\sqrt{2/q})^n\Big)\lambda^n(B_1^n),\end{aligned}$$

where $\lambda^n(B_1^n)$ denotes the measure of the n-dimensional unit ball. Similarly, each ball of radius $(1+\varepsilon)\sqrt{2/q}$ has measure

$$\Big((1+\varepsilon)\sqrt{2/q}\Big)^n \lambda^n(B_1^n),$$

and it follows that the number of balls selected by our procedure is at least

$$\begin{aligned}& \frac{2^{-n}\left(\left(\sqrt{1/n}+\varepsilon\sqrt{2/q}\right)^n - \left(\sqrt{1/n}-\varepsilon\sqrt{2/q}\right)^n\right)}{\left((1+\varepsilon)\sqrt{2q}\right)^n}\\ \geq\;& \frac{2^{-n+1}\varepsilon\sqrt{2/q}\cdot n\left(\sqrt{1/n}-\varepsilon\sqrt{2/q}\right)^{n-1}}{\left((1+\varepsilon)\sqrt{2q}\right)^n}\\ \geq\;& \frac{2^{-n+1}\varepsilon\sqrt{2/q}\cdot n\left((1-\varepsilon)\sqrt{1/n}\right)^{n-1}}{\left((1+\varepsilon)\sqrt{2q}\right)^n}\\ =\;& \frac{\varepsilon n\left((1-\varepsilon)\sqrt{1/n}\right)^{n-1}}{(1+\varepsilon)^n 2^{(3n-3)/2}} q^{(n-1)/2} = \Omega(q^{(n-1)/2}),\end{aligned}$$

where the first inequality follows by lower-bounding the difference between the two terms of the numerator by $2\varepsilon\sqrt{(2/q)}$ times the derivative of the convex function x^n in $\sqrt{1/n}-\varepsilon\sqrt{2/q}$ and the second inequality uses the fact that $\sqrt{2/q} \leq \sqrt{1/n}$. □