Approximate Revenue Maximization with Multiple Items

Nir Shabbat - 05305311

December 5, 2012

Introduction

The paper I read is called "Approximate Revenue Maximization with Multiple Items" by Sergiu Hart and Noam Nisan from 2012. It addresses the question of how a single seller can maximize its revenue when selling multiple items to a single buyer. The buyer values for the items are assumed to be independently distributed and additive. The two kinds of simple auctions that are considered in the paper are selling each of the items separately and selling all the items together as a bundle.

In this work I will try to summarize the paper in my own words, giving my intuitions on why certain claims are correct, and elaborate on some of the points that the paper goes quickly over. I will not be going into the definitions in the paper, or the mathematical aspects of the proofs, since these are fully presented in the original paper.

In the last section, I present some of the questions that came to my head when reading the paper.

Results

The paper is divided into five sections, where the first section is an introduction, and the second section presents the notation and preliminaries that will be used.

Warm up: Selling Separately vs. Bundling

The third section of the paper analyzes the gaps between the two kinds of auctions: selling items separately and selling them as a bundle. For deriving an upper bound on the bundling revenue with respect to separate sells, a distribution called the equal-revenue distribution is used.

The CDF of the Pareto distribution is given by:

$$F\left(x\right)=\begin{cases}1-\left(\frac{x_m}{x}\right)^{\alpha} & x\geq x_m\\ 0 & x<x_m\end{cases}$$

The equal-revenue distribution (denoted ER) is a Pareto distribution with $x_m = 1$ and $\alpha = 1$, and so it has a CDF $F\left(x\right) = 1-\frac{1}{x}$, and a density function $f\left(x\right) = \frac{1}{x^2}$ (with support $[1, +\infty)$). Given that a buyer's valuation for an item is distributed according to the ER distribution, the seller will obtain the same (maximum) revenue of 1 by offering the item at any price $p \geq 1$ (assuming a take-it-or-leave-it mechanism), since $\frac{d}{dp}\left(p\left(1-F\left(p\right)\right)\right) = \frac{d}{dp}p\cdot\frac{1}{p} = 1$ for every $p \geq 1$ (hence the name equal-revenue).

The ER distribution is chosen for computing an upper bound on the ratio between these two auctions, since this distribution has the heaviest possible tail. Roughly speaking, this means that the probability that the valuation of the buyer for an item will be higher than p, for an infinitely large p, is higher than any other distribution. Intuitively, the ER distribution maximizes the uncertainty of the seller with regards to the question of "What are the odds that the valuation of the buyer for the item higher than x?". Selling ER-distributed items as a bundle, somewhat decreases this uncertainty (this can be thought of as being similar to the fact that the variance of the average of two i.i.d random variables is less than each variable variance).

In section 3.1, the paper shows that for 2 items distributed according to the ER distribution, the bundling auction maximal revenue (denoted $\mathrm{BREV}\left(ER\times ER\right)$) is $2.5569\ldots$ (meaning, it is $2.5569\ldots/2$ times the maximal revenue of separate selling, which obviously has a revenue of 2 in this case). It then goes on to show (Lemma 8) that for k items i.i.d ER the ratio between selling as a bundle and selling separately is $\Theta\left(logk\right)$, meaning that selling separately may yield, as k increases, an arbitrarily small proportion of the optimal revenue (there seem to be a typo in the proof of Lemma 8, where it should have been written $p/\left(k\left(\log p+1\right)\right)\geq 2$ instead of ≤ 2).

In section 3.2, an upper bound on the bundling revenue (in respect to selling separately) is derived. First, the definition (first-order) stochastic dominance is presented, followed by two simple lemmas. Lemma 9, stating that if a one-dimensional X is stochastically dominated by a one-dimensional Y then $\mathrm{REV}(X) \leq \mathrm{REV}(Y)$, and Lemma 10, stating that for every one-dimensional X and every $r \geq 0$, $\mathrm{REV}(X) \leq r$ iff X is stochastically dominated by $r \cdot Y$ where $Y \sim ER$.

In Lemma 11 it is shown (using the expression for a linear combination of 2 independent ER-distributed variables derived previously in Lemma 6), that equalizing the scaling factors in a scaled sum of two independent ER-distributed variables stochastically dominants any other scaled sum of these two variables (giving that the sum of the scaling factors stays the same). Intuitively, this is true since the density of ER-distributed variables decrease exponentially, and so giving an equal "weight" to each of the two ER random variables allows us to extract more of each of them on the average. Corollary 12 uses Lemma 11 to show that every scaled sum of k i.i.d ER variables, is stochastically dominated by a scaled sum of the same variables where the original scale factors are replaced by their average.

Lemma 13 uses Corollary 12 to derive an upper-bound on the ratio of revenue between bundling and selling-separately for k independently distributed items. It turns out, that that the ratio calculated in 3.1 using items which have ER-distributed valuation, is in fact the worst-case scenario, and the ratio is at most$1.278\ldots$ for 2 items and $O(\log k)$ for k items.

Section 3.3 presents some lower bounds on the bundling revenue (again, with respect to selling separately). First, the case of k items which are independently distributed according to different distributions is considered. Lemma 14 shows that in this case, the bundling revenue can be more than k times worst than the selling separately revenue. This is simple corollary to the fact that the seller can offer the bundle at a price equal to the price of the item which has the maximal revenue when each of the items are sold separately. Example 15 shows that this bound is tight. In this example, the buyer valuates item i at M^i with probability M^{-i}, and at 0 with probability $1 - M^{-i}$. The revenue for each item i when sold separately is 1 (sold at price M^i), and so the total revenue is k, where the revenue of the bundle is at most $1 + 1/(M-1)$ (that is, the ratio is arbitrarily close to k as M gets larger). The intuition for choosing this distribution is in a way similar to choosing the ER distribution for deriving an upper bound. For very large M values, the odds of separately selling an item i are very slim, but since the item is priced very high the expected revenue becomes 1. On the other hand, when all of those items are sold as a bundle, the odds that the valuation of n items out of the k is positive decreases much faster in n than the sum of values for those n items.

The rest of section 3.3 deals with the case in which the k items are distributed according to identical distributions. Lemma 16 states that the maximal revenue obtained from selling a bundle of two items i.i.d F is at least the 2/3 of the revenue obtained by selling them separately. This bound is shown to be tight in Example 17 where F is a distribution with support $\{0, 1\}$ and $\mathrm{P}(1) = 2/3$, in which case the maximal revenue attained by separately selling each items optimally is 4/3 and the maximal revenue attained by optimally selling the a bundle of both is 8/9.

Lemma 18 then gives a lower bound in the case of k items, and claims that for k items, the maximal bundling revenue is at least 1/4 of the separate selling revenue. To achieve this bound, the same method that was used in the proof of Lemmas 16 is used, which is setting the price for the bundle based on the pricing of the single item. This lower bound of 1/4 on the ratio between bundling and separate-selling is not considered tight by the writers of the papers. They give an example (Example 19) of a distribution F where bundling k items distributed according to F extracts less than 0.57 of the maximal revenue obtained by separately selling those k items. The writers believe that this bound of 0.57 is in fact the tight bound. In this example, each of the items is distributed according to a distribution F on $\{0, 1\}$ with $\mathrm{P}(1) = \frac{c}{k}$. The maximal revenue of selling separately in this case is $c = k \cdot \frac{c}{k}$, and clearly the maximal revenue of the bundling auction is attained at integral prices for the bundle. They then calculate the revenue of bundling for prices $1, 2, 3$...to show the revenue decreases as the price increases. Then using differentiation of the ratios, they arrive at the value of c that gives the maximal ratio of between the two type of auctions ($c = 1.256\ldots$ is the solution of $1 - e^{-c} = 2\left(1 - (c + 1)\, e^{-c}\right)$).

Two Items

This section presents Theorem 20, which the main result of the paper. The theorem is a generalization of Theorem 1 stated in the introduction of the paper.

Theorem 20 *Let X and Y be mutli-dimensional random variables. If X and Y are independent then*

$$\mathrm{REV}\,(X,Y) \leq 2 \cdot (\mathrm{REV}\,(X) + \mathrm{REV}\,(Y))$$

To prove the above theorem, the paper first prove a series of lemmas. Lemma 21 (Marginal Mechanism Lemma) shows that the revenue from selling the two sets of items X and Y is bounded by the sum of expected valuations for one set of items, plus the expected maximal revenue of the induced mechanism on the other variable, i.e.:

$$\mathrm{REV}\,(X,Y) \leq \mathrm{VAL}\,(Y) + \mathrm{E_Y}\,[\mathrm{REV}\,(X \mid Y)]$$

To see this is true, take an optimal mechanism for selling X, Y (one which achieves $\mathrm{REV}\,(X,Y)$), and then look at the selling mechanism it induces on selling the items in X for some fixed values of Y: $y_1 \dots y_k$ (i.e. how to sell the items of X given some valuations for the items in Y). The problem is that this induced mechanism for X, also specifies the quantities of Y to give to the buyer, yet we only want a mechanism for X items. In order to solve this and also keep the induced mechanism IR and IC, instead of giving the buyer q_i units of each y_i (or each y_i with probability q_i) we pay it the buyer a sum of $\sum_{i=1}^{k} q_i y_i$. This mechanism is called the marginal mechanism and its revenue is the revenue of the original optimal mechanism (for X, Y) with $Y = y_i$, minus the payment to the buyer which is $\sum_{i=1}^{k} q_i y_i$. The payment to the buyer is bounded by $\sum_{i=1}^{k} y_i$ (since $q_i \leq 1$, and by taking the expectations over the values of Y we prove the lemma.

The bound presented in the Marginal Mechanism Lemma can't be used directly in order to bound $\mathrm{REV}\,(X,Y)$, since the sum of expected values of Y ($\mathrm{VAL}\,(Y)$) can be infinite, even when $\mathrm{REV}\,(X,Y)$ is finite. For example, if $X, Y \sim ER$, then $\mathrm{VAL}\,(X) = \mathrm{VAL}\,(\mathrm{Y}) = +\infty$, but $\mathrm{REV}\,(X,Y) = \mathrm{BREV}\,(X,Y) = 2.5569\dots$ (In section 4.3 of the article we see that bundling is optimal in this case).

In order to facilitate this bound, the domain (X, Y) is divided into two sub-domains, each of them bounded from above using this Lemma, and then the results are combined or "stitched" together.

If Z is a random variable, and S is set of values of Z, then $\mathbb{I}_{Z \in S} \cdot Z$ (where $\mathbb{I}$ is the indicator variable) is a sub-domain of Z. Lemma 22 claims that for every such S, we have $\mathrm{REV}\,(\mathbb{I}_{Z \in S} \cdot Z) \leq \mathrm{REV}\,(Z)$, which in words means that the maximal revenue achievable from the sub-domain distribution is no higher than the maximal revenue on the entire domain. This is clearly true, since any sub-domain can only increase the probability that the buyer's valuation will be zero, and since we know this mechanism is NPT, the optimal mechanism for $\mathbb{I}_{Z \in S} \cdot Z$ will extract at least as much from Z.

Next, Lemma 23 is presented. This lemma concerns the combination, or the "stitching", of two sub-domains, and it states that the sum of maximal revenues achieved from two sub-domain which cover Z is at least the maximal revenue:

$$\mathrm{REV}\left(\mathbb{I}_{Z \in S} \cdot Z\right) + \mathrm{REV}\left(\mathbb{I}_{Z \in T} \cdot Z\right) \geq \mathrm{REV}\left(Z\right)$$

Where S and T are sets of values of Z, and $S \cup T$ contains the support of Z. This is true, since the optimal mechanism on Z (with revenue $\mathrm{REV}\left(Z\right)$), extracts some of its revenue from $S \backslash T$, and some of its revenue from $T \backslash S$. The same optimal mechanism will extract at least $S \backslash T$ revenue when run on $\mathbb{I}_{Z \in S} \cdot Z$, and at least $T \backslash S$ revenue when run on $\mathbb{I}_{Z \in T} \cdot Z$.

Lemma 24 simply extends the Marginal Mechanism presented in Lemma 21 to sub-domains, in the case X and Y are independent, using the Sub-Domain Restriction Lemma (Lemma 22).

Lemma 25, the Smaller Value Lemma is presented next. This lemma states that for each X and Y one-dimensional independent random variables:

$$\mathbb{E}\left(\mathbb{I}_{Y \leq X} \cdot Y\right) \leq \mathrm{REV}\left(X\right)$$

To see this, we show a mechanism for selling X with a revenue of $\mathbb{E}\left(\mathbb{I}_{Y \leq X} \cdot Y\right)$. The mechanism offers the item at a price y, which is randomly selected according to distribution Y. The expected revenue of this mechanism is:

$$\mathbb{E}_{y \sim Y}\left(y \cdot \mathbb{P}\left(X \geq y\right)\right) = \mathbb{E}_{y \sim Y}\left(\mathbb{E}_{x \sim X}\left(y \cdot \mathbb{I}_{y \leq x}\right)\right) = \mathbb{E}_{x \sim X, y \sim Y}\left(\mathbb{I}_{y \leq x} \cdot y\right) = \mathbb{E}\left(\mathbb{I}_{Y \leq X} \cdot Y\right)$$

And so the optimal mechanism achieves at least this revenue and the lemma is proved.

We can now use Lemmas 23-25 to prove theorem 20 for the one-dimensional case. From Lemma 23 we have:

$$\mathrm{REV}\left(X, Y\right) \leq \mathrm{REV}\left(\mathbb{I}_{Y \leq X}\left(X, Y\right)\right) + \mathrm{REV}\left(\mathbb{I}_{X \leq Y}\left(X, Y\right)\right) \tag{1}$$

Using Lemma 24 we can bound the first term on the right side:

$$\mathrm{REV}\left(\mathbb{I}_{Y \leq X}\left(X, Y\right)\right) \leq \mathbb{E}\left(\mathbb{I}_{Y \leq X} \cdot Y\right) + \mathrm{REV}\left(X\right)$$

And using Lemma 25 we get:

$$\mathrm{REV}\left(\mathbb{I}_{Y\leq X}\left(X,Y\right)\right)\leq 2\mathrm{REV}\left(X\right)$$

Similarly, for the the second term on the right side of (1) it holds:

$$\mathrm{REV}\left(\mathbb{I}_{X\leq Y}\left(X,Y\right)\right)\leq 2\mathrm{REV}\left(Y\right)$$

Which leads to:

$$\mathrm{REV}\left(X,Y\right)\leq 2\left(\mathrm{REV}\left(X\right)+\mathrm{REV}\left(Y\right)\right)$$

The rest of this section extends Theorem 20 to the multi-dimensional case. Theorem 27 shows that for any independent multi-dimensional X,Y random variables:

$$\mathrm{REV}\left(X,Y\right)\leq \mathrm{REV}\left(X\right)+\mathrm{REV}\left(Y\right)+\mathrm{BREV}\left(X\right)+B\mathrm{REV}\left(Y\right)$$

and since for any X we have (by definition) $\mathrm{REV}\left(X\right)\geq \mathrm{BREV}\left(X\right)$, we get that Theorem 20 also holds for multi-dimensional variables (sets of items).

The Rest

In the rest of the paper a few more results are presented.

A tighter bound on the two items scenario, in the case where their distribution is not only independent, but also identical is provided. In this case Theorem 2 (presented in the introduction) states that selling those items separately yields at least $e/\left(e+1\right)=0.73\ldots$ of the optimal revenue, or formally:

$$2\mathrm{REV}\left(X\right)\geq \frac{e}{e+1}\mathrm{REV}\left(X,X\right)$$

A class of distributions where the bundling auction is optimal is characterized. Basically, these are distributions in which the density function f satisfies (starting at some x) that: $\left(x^{3/2}f\left(x\right)\right)'\leq 0$. Note that the ER distribution satisfies this requirement and so for two items i.i.d-ER bundling is optimal. The largest gap the write have found, between selling separately and the optimal one is $2/2.559\ldots$ which is the gap between selling separately and bundling two i.i.d-ER items.

Section 4.4 deals with multiple buyers, and generalizes Theorem 20 for the case of n buyers. This generalization works for both the main notions of IC: *dominant-strategy* IC and *Bayes-Nash* IC. Again, as in Theorem 20, selling the items separately to the n buyers, yields at least half the maximal revenue:

$$\mathrm{SREV}^{[n]}(X,Y) \geq \frac{1}{2} \cdot \mathrm{REV}^{[n]}(X,Y)$$

In section 5, the proof of Theorem 3 is presented. The theorem gives a lower bound of $O\left(\log^2 \mathrm{k}\right)$ on the revenue of the selling separately mechanism in relation to the optimal auction when selling k independent items. The proof is done by inductevly applying the equations of Theorem 27 and Lemma 13.

In contrast, the bundling auction may yield very low revenue in comparison to the optimal revenue. Example 15 gave a scenario in which the bundling action extracts only $\frac{1}{k}$ of the optimal revenue. Lemma 31 in this section proves that this lower bound is tight, and that the bundling revenue has a lower bound of $O\left(\frac{1}{k}\right)$ in comparison to the optimal revenue. The proof is similar to that of Theorem 3, only that this time we don't need to use Lemma 13, we just inductively apply Theorem 27.

Questions

1. The paper mentions (page 4, footnote 8) that there is an easy proof for showing that selling two independent items separately always yields at least half of the optimal revenue for the special case of deterministic auctions. What is that proof?

2. For two items with i.i.d valuation F, the paper shows that selling them separately yields at least $\frac{e}{e+1} = 0.73\ldots$ of the optimal revenue (page 5, Theorem 2). The worst case example that was found by the authors yields as low as 0.78 of the optimal revenue. They assume that 0.78 is indeed the tight bound. Can we prove that the tight bound is0.78, or can we find an example where selling separately yields less than 0.78 (somewhere between 0.73 and 0.78)?

3. The main result of the paper is that separately selling two items (independent but not necessarily having the same distribution) yields at

least 0.5 of the optimal revenue. Is this tight, that is, can we find an example where separately selling two (independent) items yields at most 0.5 of the optimal revenue?

4. Can we characterize (partially or completely) the cases in which selling separately is optimal?

5. The paper shows a family of distributions in which bundling is optimal (Page 20, Theorem 28). Can we find other cases? Can we completely characterize cases where bundling is optimal?

6. Can we characterize the cases in which deterministic auctions are optimal?

7. Can we characterize the cases in which bundling is at least as good as selling separately (beyond the case where bundling is optimal) for two i.i.d items? For two independent items? For two items? for k-items?

8. The paper proves that for k items i.i.d F, bundling yields at least $\frac{1}{4}$ of the revenue obtained by selling them separately. It also gives an example where bundling yields at most 0.57 of the revenue obtained by selling separately. The authors suspect 0.57 is the indeed the tight bound. Can we prove this or can we find an example that contradicts this assumption?