

Approximating the Revenue Maximization Problem with Sharp Demands

Vittorio Bilò¹ Michele Flammini² and Gianpiero Monaco²

¹Dipartimento di Matematica e Fisica “Ennio De Giorgi”, Università del Salento, Italy

²Dipartimento di Ingegneria e Scienze dell’Informazione e Matematica, Università di L’Aquila, Italy

A major decisional process in many business activities concerns whom to sell products (or services) to and at what price, with the goal of maximizing the total revenue. On the other hand, consumers would like to buy at the best possible prices and experience fair sale criteria.

In this work, we address such a problem from a computational point of view, considering a two-sided market in which the supply side consists of m indivisible items and the demand one is populated by n potential buyers (in the following also called consumers or customers), where each buyer i has a demand d_i (the number of items that i requests) and valuations v_{ij} representing the benefit i gets when owing item j . As several papers on this topic, we assume that, by means of market research or interaction with the consumers, the seller knows each customer’s valuation for each item.

The seller sets up a price p_j for each item j and assigns (i.e., sells) bundle of items to buyers with the aim of maximizing her revenue, that is the sum of the prices of all the sold items. When a consumer is assigned (i.e., buys) a set of items, her utility is the difference between the total valuation of the items she gets (valuations being additive) and the purchase price.

The sets of the sold items, the purchasing customers and their purchase prices are completely determined by the allocation of bundles of items to customers unilaterally decided by the seller. Nevertheless, we require such an allocation to meet two basic fairness constraints: (i) each customer i is allocated at most one bundle not exceeding her demand d_i and providing her a non-negative utility, otherwise she would not buy the bundle; (ii), the allocation must be envy-free [7], i.e., each customer i does not prefer any subset of d_i items different from the bundle she is assigned.

Notice that in our scenario a trivial envy-free solution always exists that lets $p_j = \infty$ for each item j and does not assign any item to any buyer.

Many papers considered the *unit demand case* in which $d_i = 1$ for each consumer i . Arguably, the *multi-demand case*, where $d_i \geq 1$ for each consumer i , is more general and finds much more applicability. To this aim, we can identify two main multi-demand schemes. The first one is the

relaxed multi-demand model, where each buyer i requests at most $d_i \geq 1$ items, and the second one is the *sharp multi-demand model*, where each buyer i requests exactly $d_i \geq 1$ items and, therefore, a bundle of size less than d_i has no value for buyer i .

For relaxed multi-demand models, a standard technique can reduce the problem to the unit demand case in the following way: each buyer i with demand d_i is replaced by d_i copies of buyer i , each requesting a single item. However, such a trick does not apply to the sharp demand model. Moreover, as also pointed out in [2], the sharp multi-demand model exhibits a property that unit demand and relaxed multi-demand ones do not possess. In fact, while in the latter model any envy-free pricing is such that the price p_j is always at most the value of v_{ij} , in the sharp demand model, a buyer i may pay an item j more than her own valuation for that item, i.e., $p_j > v_{ij}$ and compensate her loss with profits from the other items she gets (see section 3.1 of [2]). Such a property, also called *overpricing*, clearly adds an extra challenge to find an optimal revenue.

The sharp demand model is quite natural in several settings. Consider, for instance, a scenario in which a public organization has the need of buying a fixed quantity of items in order to reach a specific purpose (i.e. locations for offices, cars for services, bandwidth, storage, or whatever else), where each item might have a different valuation for the organization because of its size, reliability, position, etc. Yet, suppose a user wants to store on a remote server a file of a given size s and there is a memory storage vendor that sells slots of fixed size c , where each cell might have different features depending on the server location and speed and then yielding different valuations for the user. In this case, a number of items smaller than $\lceil \frac{s}{c} \rceil$ has no value for the user. Similar scenarios also apply to cloud computing. In [2], the authors used the following applications for the sharp multi-demand model. In TV (or radio) advertising [5], advertisers may request different lengths of advertising slots for their ads programs. In banner (or newspaper) advertising, advertisers may request different sizes or areas for their displayed ads, which may be decomposed

into a number of base units. Also, consider a scenario in which advertisers choose to display their advertisement using medias (video, audio, animation) [1,6] that would usually need a fixed number of positions, while text ads would need only one position each. An example of formulation sponsored search using sharp multi-demands can be found in [4]. Other results concerning the sharp multi-demand model in the Bayesian setting can be found in [3].

We consider the revenue maximization problem with sharp multi-demand and limited supply. We first prove that, for related valuations, the problem cannot be approximated to a factor $O(m^{1-\epsilon})$, for any $\epsilon > 0$, unless $P = NP$ and that such result is asymptotically tight. In fact we provide a simple m -approximation algorithm even for unrelated valuations.

Our inapproximability proof relies on the presence of some buyers not being able to receive any bundle of items in any envy-free outcome. Thus, it becomes natural to ask oneself what happens for instances of the problem, that we call *proper*, where no such pathological buyers exist. For proper instances, we design an interesting 2-approximation algorithm and show that the problem cannot be approximated to a factor $2 - \epsilon$ for any $0 < \epsilon \leq 1$ unless $P = NP$. Therefore, also in this subcase, our results are tight. We remark that it is possible to efficiently decide whether an instance is proper. Moreover, if discarding useless buyers is allowed, an instance can be made proper in polynomial time, without worsening the value of its optimal solution.

REFERENCES

1. A. Bezjian-Avery, B. Calder, D. Iacobucci. New Media Interactive Advertising vs. Traditional Advertising. *Journal of Advertising Research*, pp. 23-32, 1998.
2. N. Chen, X. Deng, P. W. Goldberg, J. Zhang. On Revenue Maximization with Sharp Multi-Unit Demands. In *CoRR*, arXiv:arXiv:1210.0203, 2012.
3. X. Deng, P.W. Goldberg, B. Tang, J. Zhang. Multi-unit Bayesian Auction with Demand or Budget Constraints. In *Proceedings of the ACM EC Workshop on Incentives and Trust in E-Commerce (WIT-EC 2012)*.
4. X. Deng, Y. Sun, M. Yin, Y. Zhou. Mechanism Design for Multi-slot Ads Auction in Sponsored Search Markets. In *Proceedings of The Fourth International Frontiers of Algorithmics Workshop (FAW 2010)*, pp. 11-22.
5. N. Nisan, J. Bayer, D. Chandra, T. Franji, R. Gardner, Y. Matias, N. Rhodes, M. Seltzer, D. Tom, H. Varian, D. Zigmund. Google's Auction for TV Ads. In *Proceedings of The 36th International Colloquium on Automata, Languages and Programming (ICALP 2009)*, pp. 309-327.
6. G. Rosenkrans. The Creativeness and Effectiveness of Online Interactive Rich Media Advertising. *Journal of interactive advertising*, 9(2), 2009.
7. L. Walras. Elements of Pure Economics. *Allen and Unwin*, 1954.