A Simple and Approximately Optimal Mechanism for an Additive Buyer

MOSHE BABAIOFF
Microsoft Research
and
NICOLE IMMORLICA
Microsoft Research
and
BRENDAN LUCIER
Microsoft Research
and
S. MATTHEW WEINBERG
Princeton University

In this letter we briefly survey our main result from [Babaioff el al. 2014]: a simple and approximately revenue-optimal mechanism for a monopolist who wants to sell a variety of items to a single buyer with an additive valuation.

Categories and Subject Descriptors: J.4 [**Social and Behavioral Science**]: Economics

General Terms: Economics, Theory

Additional Key Words and Phrases: Optimal Mechanisms, Simple Mechanisms, Revenue, Approximation

---

1. INTRODUCTION

Imagine that a monopolist seller has a collection of n indivisible items for sale. How should he sell the items to maximize revenue given that the buyers are strategic? With just a single item for sale to a single buyer with value drawn from a distribution F, Myerson [1981] shows that the optimal protocol is straightforward: the seller should post a fixed take-it-or-leave-it price p chosen to maximize $p \cdot (1 - F(p))$, the expected revenue (price times probability of sale). Despite the simplicity of the single-item case, extending this solution to handle multiple items remains the primary open challenge in mechanism design.

Consider even the simplest multi-item scenario [Hart and Nisan 2012]: there is a single buyer[1] with item values drawn independently from distributions $D_1, \ldots, D_n$, and whose value for a set of items is additive. Even when there are only two items for sale, it is known that the revenue-optimal mechanism may involve randomiza-

[1]Note that if the seller has unlimited copies of each item for sale, then a mechanism for a single buyer directly extends to the case of multiple buyers.

Authors' addresses: `moshe@microsoft.com`, `nicimm@microsoft.com`, `brlucier@microsoft.com`, `smweinberg@csail.mit.edu`

tion [Thanassoulis 2004], even to the extent of offering the buyer a choice among infinitely many lotteries [Daskalakis et al. 2014; Hart and Nisan 2013]. This is troubling not only from the perspective of analyzing optimal mechanisms, but also from the point of view of their usefulness. For a mechanism to be useful in practice, it should be simple to describe and transparent in its execution. Indeed, Myerson's single-item auction is exciting not only for its optimality, but also its practicality. The danger, then, is that revenue-optimal but complex mechanisms for multiple items may share the fate of other mathematically optimal designs, such as the Vickrey-Clarke-Groves mechanism, which are only rarely used in practice [Ausubel and Milgrom 2006]. It is therefore crucial to pair the study of revenue optimization with an exploration of the power of simple auctions. In other words, *what is the relative strength of simple versus complex mechanisms?*

2. THE FAILURE OF SELLING SEPARATELY OR BY BUNDLING

The above question was posed in general by Hartline and Roughgarden [2009], and by Hart and Nisan [2012] specifically for the setting of a single additive buyer. They proposed the following suggestion for a simple multi-item auction: sell each item separately, posting a fixed price on each one. The optimal price to set on item i would be $\arg\max_p p(1 - D_i(p))$, mirroring the single-item scenario. At first glance, one might expect this simple approach to be optimal. After all, the buyer's value for each item is sampled independently, and her value for item i doesn't depend at all on what other items she receives due to additivity. There is absolutely no interaction between the items at all from the buyer's perspective, so why not sell the items separately? Somewhat counter-intuitively, it turns out that this mechanism might not be optimal. For example, suppose that there are 2 items and the buyer's value for each item is distributed uniformly on $\{1, 2\}$. Then the optimal price to set on a single item is either 1 or 2, yielding a per-item revenue of 1 and a total revenue of 2. However, the seller could instead post a single take-it-or-leave-it price of 3 on the bundle of both items. Now the consumer is willing to purchase with probability $3/4$, leading to a total revenue of $9/4 > 2$. Hart and Nisan [2012] show that replacing two items whose values are drawn from the uniform distribution with n items whose values are drawn from the Equal-Revenue distribution[2] yields an example with a gap of $\Omega(\log(n))$.

What is going on in this example? The inherent problem is that the buyer's value for all items concentrates around its expectation. This is potentially helpful for revenue generation, but the strategy of selling items separately cannot exploit this property. On the other hand, the mechanism designed to target such concentration (selling only the grand bundle at a fixed price) does very poorly in settings where concentration doesn't occur; Hart and Nisan show that this grand-bundle mechanism achieves only an $\Omega(n)$ approximation to the optimal revenue in general. For example, consider an n-item instance where the buyer's value for item i is 2^i with probability $1/2^i$, and 0 otherwise. Then one can see that for all p, the buyer is only willing to purchase the grand bundle at price p with probability at most $2/p$ and therefore selling the grand bundle yields expected revenue at most 2. Yet, selling each item separately at price 2^i yields expected revenue 1 per item and n in

[2]The Equal-Revenue distribution has CDF $F(x) = 0$ for $x \leq 1$, and $F(x) = 1 - 1/x$ for $x \geq 1$.

total. Unfortunately, we must conclude that neither selling separately nor selling together can always approximate the optimal revenue to within a constant factor.

3. OUR RESULT

Our main result is that the *maximum* of the revenue generated by these two approaches — either selling all items separately or selling only the grand bundle — is a constant-factor approximation to the optimal revenue. In other words, for any product distribution of buyer values, either selling items separately approximates the optimal revenue to within a constant factor, or else bundling all items together does. Since a good approximation to the expected revenue of each approach can be computed in polynomial time given an appropriate access to the distribution, our results furthermore imply a polytime constant-factor approximation mechanism for the case of an additive buyer with independently (and non-identically) distributed values, even without the restriction of simplicity. Moreover, prior to our work, it was not even known if *any* deterministic mechanism could achieve a constant-factor approximation to the optimal mechanism, even without regard for simplicity or computational efficiency.

Main Result (Informal). *In any market with a single additive buyer and arbitrary independent item value distributions, either selling every item separately or selling all items together as a grand bundle generates at least a constant fraction of the optimal revenue.*[3]

This result complements an active research area aimed at characterizing distributions and valuations in which simple mechanisms are *precisely* optimal [Alaei et al. 2013; Hart and Nisan 2012; Pavlov 2011; Tang and Wang 2014]. In contrast to that literature, we show that a maximum over simple mechanisms is *approximately* optimal, for *arbitrary* product distributions and *additive* valuations. Our result also echoes a similar line of investigation for markets with unit-demand valuations in which a buyer's value for a set of items is his maximum value for an item in the set. In this setting, it is known [Chawla et al. 2007; Chawla et al. 2010; Chawla et al. 2010] that selling items separately achieves a constant approximation to the optimal revenue. Our result illustrates that a similar approximation can be achieved for additive buyers, provided that we also consider selling all items together as a grand bundle.

To obtain some intuition into our result, recall the example above with n items and uniformly-distributed values. This example illustrates that selling all items separately may be a poor choice when the value for the grand bundle concentrates around its expectation. What we show is that, in fact, this is the *only* scenario in which selling all items separately is a poor choice. We prove that if the total value for all items does *not* concentrate, then selling separately must generate a constant fraction of optimal revenue.

Our argument makes use of a core-tail decomposition technique introduced by Li and Yao [2013] to study the revenue of selling items separately. Roughly speaking, the idea is to split the support of each item's value distribution at some cutoff into

[3]The best constant currently known is 6. A lower bound of 2 is also known.

a "tail" (those values that are sufficiently large), and a "core" (the remainder). One can then try to analyze the revenue contributed from items that lie in the tail separately from those that lie in the core. As the cutoffs get larger and larger, the contribution from items in the tail becomes easier to approximate. When the cutoffs are sufficiently large, we can in fact show that selling separately obtains a constant fraction of the tail's contribution, as it becomes extremely unlikely that multiple items are simultaneously in the tail (and selling separately is optimal for a single item). Similarly, as the cutoffs get smaller and smaller, the contribution from items in the core becomes easier to approximate. When the cutoffs are sufficiently small, we can in fact show that the better of selling separately and together obtains a constant fraction of the core's contribution, as the contribution from items in the core must concentrate whenever the expected sum of values greatly exceeds the cutoffs. Our result follows by delicately balancing asymmetrically the cutoffs for all items so that both arguments hold simultaneously.

4. CONCLUSIONS AND FUTURE WORK

Our work provides a simple, approximately optimal mechanism for a monopolist seller facing an additive buyer [Babaioff et al. 2014]. Specifically, we show that the better of selling all items separately or all items together achieves a constant-factor approximation. Our analysis shows that this constant factor is at most 6, yet the worst known examples only exhibit a gap of 2. It is important to get a better understanding of where the gap actually lies.

A more pressing follow-up question is extending our results to multiple bidders. Exciting recent work by Yao [2015] provides a lookahead reduction for the case of many additive bidders. Applying our single-buyer mechanisms within his framework, Yao is able to also develop simple constant-factor approximations for auctions with many additive buyers, largely resolving this direction. However, the techniques at hand and in Yao's reduction appear highly specialized to additive bidders, and generalizing even slightly seems highly non-trivial. Therefore, an intriguing open question is to generalize our results and Yao's reduction beyond the case of additive bidders.

REFERENCES

Alaei, S., Fu, H., Haghpanah, N., and Hartline, J. 2013. The Simple Economics of Approximately Optimal Auctions. In *the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

Ausubel, L. M. and Milgrom, P. 2006. The lovely but lonely vickrey auction. In *Combinatorial Auctions, chapter 1*. MIT Press.

Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. 2014. A simple and aproximately optimal mechanism for an additive buyer. In *the 55th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.

Chawla, S., Hartline, J. D., and Kleinberg, R. 2007. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*. EC '07. 243–251.

Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. 2010. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *the 42nd ACM Symposium on Theory of Computing (STOC)*.

Chawla, S., Malec, D. L., and Sivan, B. 2010. The Power of Randomness in Bayesian Optimal Mechanism Design. In *the 11th ACM Conference on Electronic Commerce (EC)*.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2014. The Complexity of Optimal Mechanism Design. In *the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.

Hart, S. and Nisan, N. 2012. Approximate revenue maximization with multiple items. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. EC '12. 656–656.

Hart, S. and Nisan, N. 2013. The menu-size complexity of auctions. In *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*. EC '13. 565–566.

Hartline, J. and Roughgarden, T. 2009. Simple versus optimal mechanisms. In *EC'09*. 225–234.

Li, X. and Yao, A. C.-C. 2013. On revenue maximization for selling multiple independently distributed items. *Proceedings of the National Academy of Sciences 110,* 28, 11232–11237.

Myerson, R. 1981. Optimal auction design. *Mathematics of Operations Research 6,* 1, 58–73.

Pavlov, G. 2011. A Property of Solutions to Linear Monopoly Problems. *B. E. Journal of Theoretical Economics 11,* 1. Article 4.

Tang, P. and Wang, Z. 2014. Optimal Mechanisms with Simple Menus. In *the 15th ACM Conference on Economics and Computation (EC)*.

Thanassoulis, J. 2004. Haggling over substitutes. *Journal of Economic Theory 117,* 2, 217–245.

Yao, A. C.-C. 2015. An n-to-1 Bidder Reduction for Multi-Item Auctions and its Applications. In *the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*.