Buy-Many Mechanisms: What Are They and Why Should You Care?

SHUCHI CHAWLA
University of Wisconsin-Madison
and
YIFENG TENG
University of Wisconsin-Madison
and
CHRISTOS TZAMOS
University of Wisconsin-Madison

Multi-item mechanisms can be very complex, offering many different bundles to the buyer that could even be randomized. Such complexity is thought to be necessary as the revenue gaps between randomized and deterministic mechanisms, or deterministic and simple mechanisms are huge even for additive valuations. Furthermore, the optimal revenue displays strange properties such as non-continuity: changing valuations by tiny multiplicative amounts can change the optimal revenue by an arbitrarily large multiplicative factor. Our work shows that these strange properties do not apply to most natural situations as they require that the mechanism overcharges the buyer for a bundle while selling individual items at much lower prices. In such cases, the buyer would prefer to break his order into smaller pieces paying a much lower price overall.

We advocate studying a new revenue benchmark, namely the optimal revenue achievable by "buy-many" mechanisms, that is much better behaved. In a buy-many mechanism, the buyer is allowed to interact with the mechanism multiple times instead of just once. We show that the optimal buy-many revenue for *any* n item setting is at most $O(\log n)$ times the revenue achievable by item pricing. Furthermore, a mechanism of finite menu-size $(n/\epsilon)^{2^{O(n)}}$ suffices to achieve $(1+\epsilon)$-approximation to the optimal buy-many revenue. Both these results are tight in a very strong sense, as any family of mechanisms with description complexity sub-doubly-exponential in n cannot achieve better than logarithmic approximation in revenue. In contrast, for buy-one mechanisms, no simple mechanism of finite menu-size can get a finite-approximation in revenue. Moreover, the revenue of buy-one mechanisms can be extremely sensitive to multiplicative changes in values, while as we show optimal buy-many mechanisms satisfy revenue continuity.

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

General Terms: Economics, Theory

Additional Key Words and Phrases: Multi-item Mechanisms, Buy-many Mechanisms, Menu-size Complexity, Revenue Maximization, Approximation

1. INTRODUCTION

It is well-known that revenue-optimal mechanisms can be complicated when a seller has more than one item to sell to a buyer. Even in the simplest possible setting where the buyer has additive values, the optimal mechanism may offer the buyer infinitely many options, each of which corresponds to a randomized allocation over bundles of items at a certain price. Furthermore, when the buyer's values for the

Authors' addresses: shuchi@cs.wisc.edu, yifengt@cs.wisc.edu, tzamos@wisc.edu

items are correlated, no mechanism described by a finite number of menu entries can get any finite approximation in revenue [Briest et al. 2015; Hart and Nisan 2019]. In fact, these negative results hold already when the seller has only two items to sell, exhibiting an extreme "curse of dimensionality".

These negative results are regarded as unavoidable barriers in understanding the trade-off between optimality and simplicity for revenue maximization. Consequently, the focus has shifted to understanding this trade-off when the buyer's valuation distribution has a specific structure. For example, it is known that if the buyer's values for different items are independent and the buyer's value is unit-demand [Chawla et al. 2007; Chawla et al. 2015] or additive [Hart and Nisan 2012; Li and Yao 2013; Babaioff et al. 2014], either selling each item separately or selling all items as a grand bundle gets a constant fraction of optimal revenue. Such a result can be extended to a more general setting with subadditive valuations [Rubinstein and Weinberg 2018; Chawla and Miller 2016] as well as multiple buyers [Chawla et al. 2010; Yao 2015; Chawla and Miller 2016; Cai and Zhao 2017], but continues to require some degree of independence across individual item values. A different approach by Psomas et al. [2019] studies the smoothed complexity for revenue optimal mechanism design over correlated value distributions, but arrives at similar conclusions.

Our work [Chawla et al. 2019; 2020] presents an alternate view of the complexity of revenue-optimal mechanisms for buyers with arbitrary valuations. Our thesis is that the curse of dimensionality only happens when the buyer's actions are restricted. Let us illustrate through an example adapted from Hart and Nisan [2019]. Consider a seller with two items and a single buyer with additive values. With probability 4/7, the buyer values the first item at \$1 and the second item at \$0. With probability 2/7, the buyer values the first item at \$0 and the second item at \$2. With the remaining probability, 1/7, the buyer values each item at \$4 (and the pair at \$8). The optimal mechanism for such a setting is as follows. The seller offers a menu of 4 options to the buyer: the buyer can choose to pay \$5 and get both items, or pay \$1 and get only the first item, or pay \$2 and get only the second item, or pay \$0 and get nothing. Observe that in such an optimal mechanism, the buyer offers deterministic pricing that is superadditive over the items. In fact, all of the examples in the literature achieving large gaps between optimal mechanisms and simple mechanisms exhibit similar kinds of superadditivity in prices. However, to correctly implement such mechanisms with superadditive price, the seller must make sure that the buyer only interacts with the mechanism once. In the above example, a buyer that values each item \$4 would purchase both items and pay \$5 if he is only allowed to participate once. If we allow the buyer to interact with the mechanism for multiple times, he would first purchase the first item at a price of \$1, then purchase the second item at a price of \$2. Thus the revenue of the mechanism decreases when the buyer is allowed to participate in the mechanism more than once.

We advocate the study of revenue maximization under the **buy-many constraint**, which allows the buyer to interact with the mechanism multiple times and effectively disallows the seller from charging super-additive prices as in the above example. By the Taxation Principle, any single-buyer mechanism can be described by a menu of possible outcomes, each of which is a lottery that assigns a price to a randomized allocation. In a **buy-many** mechanism the buyer is allowed to adaptively purchase any multi-set of menu options. Such a definition is in contrast to the traditional **buy-one** mechanisms, where the buyer is only allowed to interact with the mechanism once.

The buy-many constraint is a natural property that most real-world mechanisms satisfy. Most classes of simple mechanisms studied in the literature, such as item pricing, grand bundle pricing, two-part tariffs, etc. also satisfy this constraint. Furthermore, as we show below, the optimal buy-many revenue exhibits many nice properties – bounded menu size, continuity, and approximability by simple mechanisms – that the optimal unconstrained (buy-one) revenue doesn't. As such we believe that the optimal buy-many revenue is worthy of further study as well as a good benchmark for revenue maximization.

In this note we describe the basic properties of the optimal buy-many revenue as well as simplicity versus optimality results under the buy-many constraint.

2. PROPERTIES OF BUY-MANY MECHANISMS

2.1 Revenue Approximation via Simple Mechanisms

It is known that an infinite revenue gap exists between optimal mechanisms and simple mechanisms when the buyer's values for different items are arbitrarily correlated:

THEOREM 2.1. [Hart and Nisan 2019] *There exists a distribution $\mathcal{D}$ over additive values on 2 items, such that for any finite menu size $m < \infty$,*

$$\frac{\textsc{Rev}(\mathcal{D})}{\textsc{Rev}_m(\mathcal{D})} = \infty.$$

Here $\textsc{Rev}_m(\mathcal{D})$ denotes the maximum revenue achievable by mechanisms with menu size at most m.

We notice that the large-revenue-gap examples established by Briest et al. and Hart and Nisan are no longer valid since their optimal revenue is achieved by superadditive pricings, which are not buy-many. On the other hand, Babaioff et al. [2018] observed that there exists a revenue gap between the optimal buy-one mechanism and the optimal buy-many mechanism even for additive buyers with independent item values. The first question we want to ask is whether the revenue gap between optimal mechanisms and simple mechanisms can still be large when the optimal mechanism satisfies the buy-many constraint. Our first result shows that it cannot.

THEOREM 2.2. *For any distribution $\mathcal{D}$ over arbitrary valuation functions,*

$$\frac{\textsc{BuyManyRev}(\mathcal{D})}{\textsc{Srev}(\mathcal{D})} = O(\log n).$$

Here $\textsc{BuyManyRev}(\mathcal{D})$ denotes the optimal revenue obtained by any buy-many mechanism; $\textsc{Srev}(\mathcal{D})$ denotes the optimal revenue obtained by item pricing. This theorem generalizes a previous result by Briest et al. [2015], who established the same revenue gap for a unit-demand buyer. We emphasize that the theorem requires no assumption over the distribution, except that the valuation function should be monotone non-decreasing in the set of items allocated: allocating extra items will never lower the buyer's value.

Can we do better than $O(\log n)$ in the approximation ratio? We show that the above result is tight in a very strong sense, even if we are allowed to use much more complex mechanisms.

Theorem 2.3. *There exists a distribution over XOS valuation functions for which no mechanism with description complexity at most* $2^{2^{o(n^{1/4})}}$ *can obtain a* $o(\log n)$ *fraction of the optimal revenue obtained by buy-many mechanisms.*

Observe that item pricing can be described using $O(n \log R)$ bits when the values are integers in the range $[1, R]$. Our construction sets $R = \text{poly}(n)$. This implies that to obtain $O(\log n)$-approximation in revenue, we only need a mechanism with description complexity almost linear in n. However, if we want to do even slightly better asymptotically, we need a mechanism that has to be described by doubly-exponential in n number of bits.

2.2 Menu-size Complexity for Near-optimal Mechanisms

The menu size of a mechanism, defined as the number of different outcomes the seller offers to the buyer, was first introduced by Hart and Nisan [2019]. It has been studied extensively in literature as a measure of complexity for single-buyer mechanisms (see, e.g., [Babaioff et al. 2017; Gonczarowski 2018; Kothari et al. 2019]). It has a direct correspondence with the communication complexity of the interaction protocol between the buyer and the seller.

For the unconstrained setting, it is known that even for selling two items to an additive buyer, the menu-size complexity of the optimal buy-one mechanism can be infinite [Hart and Nisan 2019]. The same is true for any mechanism that achieves a bounded approximation to the optimal revenue. Positive results for getting near-optimal revenue using mechanisms with finite menu-size complexity have only been established for buyers with subadditive valuation functions and independent values for each item [Babaioff et al. 2017; Kothari et al. 2019].

For buy-many mechanisms, we define their menu-size complexity to be similar to the "additive menu size" introduced by Hart and Nisan [2019], which corresponds to the number of "basic" options the buy-many mechanism offers. An example in [Daskalakis et al. 2017] shows that the optimal buy-many mechanism may still have an infinite menu size. Can we approximate the optimal revenue obtained by buy-many mechanisms with bounded menu-size? We give an affirmative answer to this question in the following theorem.

Theorem 2.4. *For any distribution* $\mathcal{D}$ *over arbitrary valuation functions, there exists a buy-many mechanism* $\mathcal{M}$ *generated by* $(n/\epsilon)^{2^{O(n)}}$ *menu entries, such that* $\text{Rev}_{\mathcal{D}}(\mathcal{M}) \geq (1-\epsilon)\text{BuyManyRev}(\mathcal{D})$.

The above theorem shows that a mechanism with menu-size complexity doubly-exponential in n can get near-optimal revenue obtained by any buy-many mechanism. Like Theorem 2.2, this theorem does not make any assumption on the value distribution and allows it to be arbitrarily correlated without structure. By Theorem 2.3, the doubly-exponential dependency on n is also tight, since no mechanism with sub-doubly-exponential description complexity can get $o(\log n)$-approximation in revenue.

2.3 Revenue Continuity of Optimal Mechanisms

We are interested in understanding the extent to which the optimal revenue changes if the value distribution is perturbed slightly. This extent depends, of course, on the manner in which the value distribution is perturbed. It is known [Daskalakis and Weinberg 2012; Rubinstein and Weinberg 2018; Gonczarowski and Weinberg 2018; Kothari et al. 2019; Brustle et al. 2019]), for example, that when each valuation function in the support of the distribution gets perturbed additively by some small $\epsilon > 0$, the optimal revenue does not change too much multiplicatively. However, when the buyer's values are unbounded, it is natural to ask what happens if the valuation function is perturbed multiplicatively, since such a perturbation is scale-invariant.

Formally, let $\mathcal{D}$ be a distribution over valuation functions, and let $\mathcal{D}'$ be another distribution obtained by taking each valuation function in the support of $\mathcal{D}'$ and changing each component of this function multiplicatively by some factor in $[1-\epsilon, 1+\epsilon]$ for some small $\epsilon > 0$. Can we then show that $\text{Rev}_{\mathcal{D}'} \geq (1-\epsilon')\text{Rev}_{\mathcal{D}}$ where ϵ' goes to 0 as ϵ goes to 0? We call such a property *revenue continuity.*

While revenue continuity is inherently interesting, it also has practical implications. Continuity implies that revenue estimates established on the basis of market analysis will be robust to errors in estimating demand. Furthermore, the accuracy of these estimates will improve directly with a reduction in measurement error. From an algorithmic standpoint, revenue continuity allows discretizing the values to their most significant digits through a sufficiently fine multiplicative grid. This is possible to do without a significant drop in revenue.

A surprising result based on an example by Psomas et al. [2019] shows that optimal revenue obtained by buy-one mechanisms does not exhibit revenue continuity, even for additive buyers. It is even possible that the optimal revenue is infinite before the perturbation, but finite afterward. However, such revenue discontinuity does not happen to buy-many mechanisms. We state the result in the following theorem.

Theorem 2.5. *Let $\mathcal{D}$ be a distribution over arbitrary valuation functions, and $\mathcal{D}'$ a $(1 \pm \epsilon)$-multiplicative-perturbation of $\mathcal{D}$. Then*

$$\text{BuyManyRev}(\mathcal{D}') \geq (1 - \mathit{poly}(n, \epsilon))\,\text{BuyManyRev}(\mathcal{D}).$$

The above theorem implies that in sharp contrast to the unconstrained setting, buy-many mechanisms always satisfy revenue continuity, while again we do not need to have an assumption over the value distribution. The polynomial dependency on n is necessary even when the buyer is unit-demand.

3. CONCLUSION AND OPEN QUESTIONS

In this note, we advocate studying the revenue achievable by buy-many mechanisms as an alternative benchmark for the study of revenue-maximizing multi-item mechanisms. The optimal buy-many revenue has many nice properties that the unconstrained optimal buy-one revenue lacks for general value distributions: simple mechanisms such as item pricing can get a good approximation to the optimal buy-many revenue; mechanisms of bounded menu-size can give a near-optimal approximation; the optimal buy-many revenue satisfies revenue-continuity.

There are several interesting directions for future work. One direction is to investigate other revenue maximization problems using buy-many revenue as the benchmark. For example, for special cases such as an additive buyer with independent item values, is there an efficient approximation scheme computing the optimal buy-many revenue? Can we find improved constant factors approximations using simple mechanisms for other special cases? Another direction is to investigate similar revenue benchmarks for multi-buyer limited supply settings. How should the definition of buy-many mechanisms be extended to those settings?

REFERENCES

Babaioff, M., Gonczarowski, Y. A., and Nisan, N. 2017. The menu-size complexity of revenue approximation. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 869–877.

Babaioff, M., Immorlica, N., Lucier, B., and Weinberg, S. M. 2014. A simple and approximately optimal mechanism for an additive buyer. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. IEEE, 21–30.

Babaioff, M., Nisan, N., and Rubinstein, A. 2018. Optimal deterministic mechanisms for an additive buyer. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. ACM, 429–429.

Briest, P., Chawla, S., Kleinberg, R., and Weinberg, S. M. 2015. Pricing lotteries. *Journal of Economic Theory 156*, 144–174.

Brustle, J., Cai, Y., and Daskalakis, C. 2019. Multi-item mechanisms without item-independence: Learnability via robustness. *arXiv preprint arXiv:1911.02146*.

Cai, Y. and Zhao, M. 2017. Simple mechanisms for subadditive buyers via duality. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2017. 170–183.

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

Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. ACM, 311–320.

Chawla, S., Malec, D., and Sivan, B. 2015. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior 91*, 297–317.

Chawla, S. and Miller, J. B. 2016. Mechanism design for subadditive agents via an ex ante relaxation. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 579–596.

Chawla, S., Teng, Y., and Tzamos, C. 2019. Buy-many mechanisms are not much better than item pricing. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 237–238.

Chawla, S., Teng, Y., and Tzamos, C. 2020. Menu-size complexity and revenue continuity of buy-many mechanisms. *arXiv preprint arXiv:2003.10636*.

Daskalakis, C., Deckelbaum, A., and Tzamos, C. 2017. Strong duality for a multiple-good monopolist. *Econometrica 85,* 3, 735–767.

Daskalakis, C. and Weinberg, S. M. 2012. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 370–387.

Gonczarowski, Y. A. 2018. Bounding the menu-size of approximately optimal auctions via optimal-transport duality. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2018. 123–131.

GONCZAROWSKI, Y. A. AND WEINBERG, S. M. 2018. The sample complexity of up-to-ε multi-dimensional revenue maximization. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 416–426.

HART, S. AND NISAN, N. 2012. Approximate revenue maximization with multiple items. In *ACM Conf. on Electronic Commerce*. 656.

HART, S. AND NISAN, N. 2019. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory 183*, 991–1029.

KOTHARI, P., SINGLA, S., MOHAN, D., SCHVARTZMAN, A., AND WEINBERG, S. M. 2019. Approximation schemes for a unit-demand buyer with independent items via symmetries. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 220–232.

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.

PSOMAS, A., SCHVARTZMAN, A., AND WEINBERG, S. M. 2019. Smoothed analysis of multi-item auctions with correlated values. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. ACM, 417–418.

RUBINSTEIN, A. AND WEINBERG, S. M. 2018. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Transactions on Economics and Computation (TEAC) 6,* 3-4, 19.

YAO, A. C.-C. 2015. An n-to-1 bidder reduction for multi-item auctions and its applications. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics, 92–109.