

An $O(\log \log m)$ Prophet Inequality for Subadditive Combinatorial Auctions

PAUL DÜTTING

Google Research and London School of Economics

and

THOMAS KESSELHEIM

University of Bonn

and

BRENDAN LUCIER

Microsoft Research

We survey the main results from [Dütting, Kesselheim, and Lucier 2020]:¹ a simple posted-price mechanism for subadditive combinatorial auctions with m items that achieves an $O(\log \log m)$ approximation to the optimal welfare, plus a variant with entry fees that approximates revenue. These are based on a novel subadditive prophet inequality.

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

General Terms: Economics, Theory

Additional Key Words and Phrases: Simple Mechanisms, Combinatorial Auctions, Subadditive, Revenue, Welfare, Approximation

1. INTRODUCTION

Consider the problem of welfare maximization in an online Bayesian combinatorial auction. There is a set M of m objects to be divided among n agents. Each agent i has a valuation function v_i that assigns a value to every subset of objects. These valuation functions are random, drawn independently from known (but not necessarily identical) distributions. Agents arrive one by one in an arbitrary order. When an agent arrives she reveals her valuation, and the decision-maker must choose which subset of items to allocate to the agent and how much she should pay. The goal is to design a dominant strategy truthful mechanism that approximately maximizes the total expected value of the assignment.

If agent valuations are submodular (exhibit decreasing marginal values), then there is a truthful mechanism that obtains at least half of the expected optimal welfare in hindsight [Feldman et al. 2015].² This mechanism has a simple form: each item j is assigned a precomputed fixed price p_j , then each agent in sequence is allocated the subset S of the remaining items that maximizes her utility

¹FOCS 2020.

²This result actually applies to the more general class of XOS valuations.

Authors' addresses: p.d.duetting@lse.ac.uk, thomas.kesselheim@uni-bonn.de, brlucier@microsoft.com

$v_i(S) - \sum_{j \in S} p_j$. In the special case with exactly one item, this result is the prophet inequality due to Krenkel and Sucheston [1977; 1978] and Samuel-Cahn [1984]. The extension to submodular valuations over multiple items is part of a line of literature extending prophet inequalities to more general allocation problems, yielding applications to truthful mechanism design for both welfare and revenue (e.g., [Chawla et al. 2010; Kleinberg and Weinberg 2012; Feldman et al. 2015; Dütting and Kleinberg 2015; Feldman et al. 2016; Rubinstein and Singla 2017; Chawla et al. 2019; Dütting et al. 2020; Cai and Zhao 2017]).

One of the more vexing open problems in this space is whether these results extend to subadditive valuations. Valuation v is subadditive if $v(S) + v(T) \geq v(S \cup T)$ for all sets S and T . Subadditive allocation has received considerable attention from both the algorithmic and economic perspectives. For the former, there is a known $O(1)$ -approximate polynomial-time algorithm for the offline problem [Feige 2009]. It is also known that running a sealed-bid auction for each object separately yields an $O(1)$ approximation to the optimal welfare at any Bayes-Nash equilibrium [Feldman et al. 2013]. It's tempting to guess that an $O(1)$ -approximate prophet inequality is possible as well, but prior to this work the best-known prophet inequality bound (and truthful approximation to either welfare or revenue for the Bayesian setting) was $O(\log m)$ [Feldman et al. 2015; Cai and Zhao 2017].

2. AN $O(\log \log m)$ APPROXIMATION FOR SUBADDITIVE VALUATIONS

We make progress on this problem by obtaining an $O(\log \log m)$ -approximate price-based prophet inequality for subadditive combinatorial auctions. This means we can find item prices so that the expected welfare of the corresponding dominant strategy incentive compatible posted-price mechanism will be an $O(\log \log m)$ approximation to the expected welfare that could be achieved by an optimal offline algorithm.

THEOREM 2.1 (WELFARE, EXISTENTIAL). *For subadditive valuations drawn independently from known distributions, there exist static anonymous item prices that yield an $O(\log \log m)$ approximation to the optimal expected welfare.*

Theorem 2.1 is existential: it only shows the existence of appropriate prices. We discuss how to turn this result into a computational (polytime) result in Section 5. We also show how to apply our prophet inequality to revenue maximization, following a framework due to Cai and Zhao [2017]; we discuss revenue in Section 6. At the heart of our approach is the following lemma, which for a given and fixed subadditive valuation v_i asserts the existence of item prices p_j for a given set U that satisfy a useful inequality.

LEMMA 2.2 (KEY LEMMA). *For every $i \in N$, subadditive function v_i , and set $U \subseteq M$ there exist prices p_j for $j \in U$ and a probability distribution λ over $S \subseteq U$ such that for all $T \subseteq U$*

$$\sum_{S \subseteq U} \lambda_S \left(v_i(S \setminus T) - \sum_{j \in S \setminus T} p_j \right) + \sum_{j \in T} p_j \geq \frac{v_i(U)}{\gamma}, \quad (1)$$

where $\gamma \in O(\log \log m)$.

Given this lemma it is relatively straightforward to show Theorem 2.1. The idea is to let (U_1, \dots, U_n) be the welfare-maximizing allocation, and for each U_i and

$j \in U_i$ use the prices p_j from Lemma 2.2 with $U = U_i$. The welfare argument then proceeds by rewriting the welfare as the sum of buyer utilities and revenue, with Lemma 2.2 providing a tool to lower bound the buyer utilities.

In this lower bound argument the set T from Lemma 2.2 can be interpreted as the set of items which are already gone when we consider agent i , and λ is a distribution over sets of items S that agent i considers to buy. Of course, agent i can only buy items that are still available, so she only derives value from $S \setminus T$. The lemma therefore establishes that the utility that can be obtained by agent i , plus the revenue obtained from selling the items in T , is at least a factor γ of her contribution to the optimal welfare.

Before describing our approach to proving Lemma 2.2, let's take a brief detour to compare with previous results and the barrier at $O(\log m)$.

3. ASIDE: BALANCED PRICES

For intuition, consider a simpler case of Lemma 2.2 where the valuation v is additive, and where λ is constrained to always choose $S = U$ with probability 1. In this case we could set price $p_j = \frac{1}{2}v(j)$ for each item j . That is, prices mirror a scaled-down version of the valuation itself. By additivity, the left-hand side of (1) becomes $v(U \setminus T) - \frac{1}{2}v(U \setminus T) + \frac{1}{2}v_t(T) = \frac{1}{2}v_i(U \setminus T) + \frac{1}{2}v_i(T)$. Since v is additive, inequality (1) holds with $\gamma = 1/2$.

If the valuation is not additive, this pricing strategy requires some modification. We'd still like to find prices that somehow approximate the valuation v . It suffices to find prices such that (a) the total price of U is approximately $v(U)$, and (b) the total price of any $T \subseteq U$ is (approximately) at least $v(U) - v(U \setminus T)$, the value lost if T is removed. Prices that satisfy these properties are said to be (α, β) -balanced, where α and β capture the approximation in (a) and (b) respectively [Kleinberg and Weinberg 2012; Dütting et al. 2020]. If prices are (α, β) -balanced, then (1) will be satisfied with $\gamma = O(\alpha\beta)$ even when λ chooses U with probability 1.

The $O(1)$ -approximation results for submodular valuations follow because submodular valuations admit $(1, 1)$ -balanced prices [Feldman et al. 2015]. For subadditive valuations, it is known that $(O(\log m), 1)$ -balanced prices always exist, but there are examples in which no better approximation is possible [Feldman et al. 2015]. So breaking the $O(\log m)$ barrier requires a different approach.

Luckily for us, balancedness is stronger than what's needed to satisfy Lemma 2.2, and for our welfare result more generally. For example, suppose v is a unit-demand valuation that assigns value 1 to each item in U . This valuation admits $(1, 1)$ -balanced prices: for example, $p_j = 1/m$ for each j . These prices are low enough that the entire set U is "affordable" – the total price of U is equal to $v(U) = 1$. But this feels unnecessary, since the unit-demand buyer would never actually purchase more than one item. Suppose instead we set $p_j = 1/2$ for each j . These prices are not very balanced, since the sum of all prices is much larger than $v(U)$. But they nevertheless enable high-welfare outcomes: if even a single item in U is available when agent i arrives, she can purchase it and obtain high utility. Such high prices suggest and facilitate an outcome where the buyer targets a small number of items that carry most of the value. In Lemma 2.2, this added flexibility is captured by the distribution λ : for this unit-demand example, one can verify that setting $p_j = 1/2$

and choosing λ to be a uniform distribution over all singletons would satisfy (1) with $\gamma = 1/2$. Our general construction uses this additional flexibility to find improved prices for subadditive valuations, where good balanced prices may not exist.

4. FINDING PRICES FOR SUBADDITIVE VALUATIONS

We now return to Lemma 2.2 and give a brief overview of the proof techniques. To prove Lemma 2.2 we write down an LP and use strong LP duality to show the following equivalent condition: there exists a probability distribution λ over set of items S so that for every probability distribution μ with $\sum_{T:j \in T} \mu_T \leq \sum_{S:j \in S} \lambda_S$, i.e., that puts at most the same probability mass on each item j as distribution λ , it holds that

$$\sum_{S,T} \lambda_S \cdot \mu_T \cdot v_i(S \setminus T) \geq \frac{1}{\gamma} \cdot v_i(U). \quad (2)$$

We interpret the left-hand side of (2) as a zero-sum game, in which the protagonist chooses λ and the antagonist chooses μ , and the protagonist's goal is to maximize $\sum_{S,T} \lambda_S \cdot \mu_T \cdot v_i(S \setminus T)$. This has a natural interpretation: the designer's goal is to find a purchasing strategy for the buyer that maximizes the value of the set they obtain, and the adversary's goal is to arrange the purchasing outcomes so that removing all previously-sold items (i.e., T) steals most of the value from the buyer, leaving their realized value $v_i(S \setminus T)$ as small as possible.

We prove a lower bound on the value of this game by restricting attention to distributions λ that put the same probability mass q on each item. The crux of our argument is that for each such “equal-marginals distribution” λ with corresponding probability q , the value of the zero-sum game is at least $f(q) - f(q^2)$, where $f(q)$ is the optimal expected social welfare that can be achieved by a distribution over sets of items $S \subseteq U$ that puts probability mass at most q on each item. Intuitively, if it's possible for the adversary to choose some distribution μ over T that is guaranteed to “steal the value” from the buyer's distribution λ over S , then it must be that the set $S \cap T$ has high expected value. But if λ and μ each place probability at most q on each item, then the distribution over $S \cap T$ places probability at most q^2 on each item. Thus, if the adversary can perform well in the zero-sum game for some q , this directly implies that we should consider the game with the significantly smaller marginal probability q^2 .

To turn this intuition into an $O(\log \log m)$ bound, we let the protagonist consider such “equal marginal distributions” for $q = 2^{-2^k}$ for $k = 0$ to $k = O(\log \log m)$, and obtain a lower bound on the value of the zero-sum game by taking the average of the sum of the corresponding lower bounds $f(q) - f(q^2)$. Now by the choice of the q this telescoping sum has $O(\log \log m)$ terms, and it evaluates to $f(1/2) - f(1/m^2)$. The proof is completed by observing that the latter is at least $(1/2 - o(1)) \cdot v_i(U)$.

5. COMPUTING PRICES IN POLYNOMIAL TIME

Theorem 2.1 shows that appropriate item prices exist. In the paper we also show how to turn our existential proof into a polytime construction, assuming appropriate

demand query access to the valuation functions.³

THEOREM 5.1 (WELFARE, COMPUTATIONAL). *For subadditive valuations drawn independently from known distributions and any $\epsilon > 0$, there is a polytime (in n , m , and $1/\epsilon$) algorithm to compute static and anonymous item prices that yield an $O(\log \log m)$ approximation to the optimal expected welfare up to an additive error of ϵ .*

The additive ϵ term in Theorem 5.1 comes from errors introduced by sampling from the distributions over valuations. But the construction is non-trivial even in the full information case where the valuations are fixed and no sampling is required. Remember that our existential analysis used a dual formulation where we argued about equilibrium distributions λ and μ of a zero-sum game, rather than prices. We prove Theorem 5.1 by reformulating our optimization problem in a way that avoids having to compute these equilibrium distributions directly. This reformulation involves taking yet another dual, bringing us back into the space of prices.

Specifically, for a given constraint q on the marginal allocation probabilities, we note that the value $f(q^2)$ can be encoded as the solution to a configuration LP—that is, a fractional allocation problem—under the constraint that no more than a q^2 fraction of any item can be allocated in expectation. We then take a dual of this configuration LP, resulting in a new program whose variables can be interpreted as item prices (but different from the original optimization over item prices that we started with). We can compute a solution to this dual LP (and hence item prices) in polynomial time with the Ellipsoid method since a separation oracle can be implemented with demand queries. The final prices we use for Theorem 5.1 are the prices from this dual LP solution scaled by q .

It might seem strange that we start with an LP for $f(q^2)$, rather than $f(q)$, then scale prices by q at the end. The intuition behind this construction is as follows. In Section 4 we bounded the value of the zero-sum game by $f(q) - f(q^2)$. Here $f(q)$ is the highest expected value that the protagonist could obtain from a choice of λ if the adversary abstained, and $f(q^2)$ is an upper bound on how much value the antagonist can take away by choosing μ optimally. By taking the dual prices for the configuration LP for $f(q^2)$ and scaling them by q , we are effectively setting prices that approximate the welfare loss due to the antagonist’s strategy, which is to say the worst-case loss from excluding items that have already been sold.

6. REVENUE MAXIMIZATION

We also show how to apply our new prophet inequality to the design of simple and dominant strategy incentive compatible mechanisms that approximate the Bayesian optimal revenue. Our revenue approximation makes use of a framework for constructing simple mechanisms due to Cai and Zhao [2017], which builds upon a recent literature applying a duality approach to revenue maximization [Cai et al. 2016]. Cai and Zhao established an $O(\log m)$ revenue approximation for subadditive valuations under a natural item independence assumption. A key step in their analysis invokes a posted-price-based prophet inequality for welfare maximization,

³A *demand query* returns, for a given valuation v and choice of item prices (p_1, \dots, p_m) , a subset of items S that maximizes $v(S) - \sum_{j \in S} p_j$.

augmented to allow ex ante allocation probability constraints. We show that our prophet inequality, extended to handle arbitrary (not necessarily equal) constraints on the marginal allocation probabilities, can be used to improve the revenue approximation from $O(\log m)$ to $O(\log \log m)$.

7. GOING BEYOND $O(\log \log m)$

In the paper we demonstrate that the $O(\log \log m)$ -factor that shows up in all our bounds is best possible using our approach. In particular, our constructions use only distributions λ that set the same marginal probability q of allocating each item. We show by way of example that such distributions (and their associated dual prices) can suffer loss as high as $\Omega(\log \log m)$.

Our restriction to equal-marginal distributions was crucial for our approach to optimizing over distributions. It is natural to wonder whether our bound could be improved by relaxing the equal-marginals assumption and using an arbitrary profile of marginal distributions. We conjecture that an $O(1)$ -approximate prophet inequality can be achieved using item prices that are dual to a distribution with unequal marginals. We leave resolving this conjecture as an open problem.

REFERENCES

- CAI, Y., DEVANUR, N. R., AND WEINBERG, S. M. 2016. A duality based unified approach to bayesian mechanism design. In *STOC*. 926–939.
- CAI, Y. AND ZHAO, M. 2017. Simple mechanisms for subadditive buyers via duality. In *STOC*. 170–183.
- CHAWLA, S., HARTLINE, J. D., MALEC, D. L., AND SIVAN, B. 2010. Multi-parameter mechanism design and sequential posted pricing. In *STOC*. 311–320.
- CHAWLA, S., MILLER, J. B., AND TENG, Y. 2019. Pricing for online resource allocation: Intervals and paths. In *SODA*. 1962–1981.
- DÜTTING, P., FELDMAN, M., KESSELHEIM, T., AND LUCIER, B. 2020. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.* 49, 3, 540–582.
- DÜTTING, P. AND KLEINBERG, R. 2015. Polymatroid prophet inequalities. In *ESA*. 437–449.
- FEIGE, U. 2009. On maximizing welfare when utility functions are subadditive. *SIAM J. Comput.* 39, 1, 122–142.
- FELDMAN, M., FU, H., GRAVIN, N., AND LUCIER, B. 2013. Simultaneous auctions are (almost) efficient. In *STOC*. 201–210.
- FELDMAN, M., GRAVIN, N., AND LUCIER, B. 2015. Combinatorial auctions via posted prices. In *SODA*. 123–135.
- FELDMAN, M., SVENSSON, O., AND ZENKLUSEN, R. 2016. Online contention resolution schemes. In *SODA*. 1014–1033.
- KLEINBERG, R. AND WEINBERG, S. M. 2012. Matroid prophet inequalities. In *STOC*. 123–136.
- KRENGEL, U. AND SUCHESTON, L. 1977. Semiamarts and finite values. *Bull. Am. Math. Soc.* 83, 745–747.
- KRENGEL, U. AND SUCHESTON, L. 1978. On semiamarts, amarts, and processes with finite value. *Adv. in Prob. and Relat. Top.* 4, 197–266.
- RUBINSTEIN, A. AND SINGLA, S. 2017. Combinatorial prophet inequalities. In *SODA*. 1671–1687.
- SAMUEL-CAHN, E. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Ann. Probab.* 12, 1213–1216.