ARTICLE IN PRESS
YGAME:2390

Games and Economic Behavior ••• (••••) •••–•••

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Extreme value theorems for optimal multidimensional pricing

Yang Cai [a,1], Constantinos Daskalakis [b,*,2]

[a] Computer Science, McGill University, Canada
[b] EECS, MIT, United States

ARTICLE INFO

Article history:
Received 6 March 2012
Available online xxxx

JEL classification:
C70
D44

Keywords:
Multidimensional pricing
Revenue maximization
Extreme value theory

ABSTRACT

We provide near-optimal, polynomial-time algorithms for pricing n items to optimize revenue against a unit-demand buyer whose values are independent from known distributions. For any chosen $\epsilon > 0$ and values in $[0, 1]$, our algorithm's revenue is optimal up to an additive ϵ. For values sampled from monotone hazard rate (MHR) or regular distributions, we achieve a $(1-\epsilon)$-fraction of the optimal revenue in polynomial time and quasi-polynomial time, respectively.

Our algorithm for bounded distributions applies probabilistic techniques to understand the statistical properties of revenue distributions, obtaining a reduction in the algorithm's search space via dynamic programming. Adapting this approach to MHR and regular distributions requires the proof of novel extreme-value theorems for such distributions. As a byproduct, we show that, for all n, a constant or a polylogarithmic (in n) number of distinct prices suffice for near-optimal revenue for MHR and regular distributions, respectively.

1. Introduction

We study the following pricing problem. A seller has n items to sell to a buyer who is looking to buy a single item. The seller wants to maximize profit from the sale, leveraging stochastic knowledge she has about the buyer to achieve this goal. In particular, we assume that the seller has access to a distribution $\mathcal{F}$ from which the values $(v_1, \dots, v_n)$ of the buyer for the items are drawn. Given this information, the seller wants to compute prices $p_1, \dots, p_n$ for the items to maximize her revenue, assuming that the buyer is *quasi-linear*—i.e. will buy the item i maximizing $v_i - p_i$, as long as this difference is positive. That is, the seller's expected revenue from a price vector $P = (p_1, \dots, p_n)$ is

$$\mathcal{R}_P = \sum_{i=1}^{n} p_i \cdot \Pr\left[(i = \arg\max\{v_j - p_j\}) \wedge (v_i - p_i \geq 0)\right], \tag{1}$$

where we assume that the arg max breaks ties in favor of a single item, when there are multiple maximizers. A more sophisticated seller could try to improve her revenue by pricing lotteries over items, that is also price randomized allocations of items (Briest et al., 2010), albeit this may be less natural than item pricing, and we will not study it extensively in this paper.

* Corresponding author.
E-mail addresses: cai@cs.mcgill.ca (Y. Cai), costis@csail.mit.edu (C. Daskalakis).

[1] Work done while the author was a Ph.D. student at MIT, supported by NSF Awards CCF-0953960 (CAREER) and CCF-1101491.
[2] Supported by a Sloan Foundation Fellowship, a Microsoft Research Faculty Fellowship, and NSF Awards CCF-0953960 (CAREER) and CCF-1101491.

http://dx.doi.org/10.1016/j.geb.2015.02.003
0899-8256/

Please cite this article in press as: Cai, Y., Daskalakis, C. Extreme value theorems for optimal multidimensional pricing. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.02.003

While our problem has a simple statement, it exhibits rich behavior depending on the nature of $\mathcal{F}$. For example, if $\mathcal{F}$ assigns the same value to all the items with probability 1, i.e. when the buyer always values all items equally, the problem becomes *single-dimensional*. In this setting, it is clear that lotteries do not improve the revenue and that the optimal price vector can assign the same price to all the items. This observation is a special case of the more general, celebrated result of Myerson (1981) on optimal mechanism design, i.e. the multi-buyer version of our problem, and generalizations thereof. Myerson's result provides a closed-form solution to the multi-buyer problem in a single sweep that covers many settings, but only works under the same limiting assumption that every buyer is single-dimensional, i.e. receives the same value from all the items. (More generally, every buyer receives the same value from all outcomes of the mechanism that provide her service.)

Following Myerson, a large body of research in both Economics and Engineering has been devoted to extending his result to the *multi-dimensional setting*, where the buyers' values come from general distributions. And, while there has been sporadic progress (see survey Manelli and Vincent, 2007 and its references), an optimal multi-dimensional mechanism, generalizing Myerson's result, does not seem to be in sight. Indeed, there is not even an optimal solution known for the single-buyer item pricing problem. Even the ostensibly easier version of that problem, where the values of the buyer for the items are independent and supported on a set of cardinality 2 is unresolved.[3] Our main contribution in this paper is to develop near-optimal polynomial-time algorithms for this problem, when the buyer's values for the items are independent.

1.1. Main results

We partition our results into algorithmic and structural. The former provide efficient algorithmic procedures for computing near-optimal price vectors. The latter shed light into the structure of optimal solutions.

Algorithmic results. Previous work on the item pricing problem has provided constant factor approximation algorithms. The best known polynomial-time algorithm obtains revenue that is at least 1/2 of the revenue of the optimal price vector (Chawla et al. 2007, 2010a). We discuss these approaches in Section 1.3, also noting that they are limited to constant factor approximations. We are aiming instead for item pricing mechanisms that come arbitrarily close to the optimal revenue, obtaining the following results. Their proofs are overviewed in Sections 4 through 9, while complete details are provided in Appendices A–J.

Theorem 1 *(**Main algorithmic result:** additive PTAS for bounded distributions). Suppose that the values of the buyer for n items are independent and normalized to lie in* $[0, 1]$. *Then, for all* $\epsilon > 0$, *there exists an algorithm that computes a price vector whose revenue is within an additive* ϵ *of optimal, and whose running time is polynomial in* $n^{\frac{\log^3 1/\epsilon}{\epsilon^4}}$.

Theorem 2 *(General algorithm). Suppose that the values of the buyer for n items are independent and supported on some interval* $[u_{min}, r \cdot u_{min}]$ *for some* $u_{min} > 0$ *and* $r \geq 1$. *Then, for all* $\epsilon > 0$, *there is an algorithm that computes a price vector whose revenue is at least a* $(1-\epsilon)$*-fraction of the optimal revenue, and whose running time is polynomial in* $\max\left\{n^{\log^{11} r \cdot \log\log r}, n^{\frac{\log^3 r \cdot \log\frac{1}{\epsilon}}{\epsilon^8}}\right\}$.[4]

Theorem 3 *(Multiplicative PTAS for MHR distributions). There is a Polynomial-Time Approximation Scheme*[5] *for computing an optimal price vector, when the values of the buyer are independently drawn from Monotone Hazard Rate distributions.*[6]

For any accuracy $\epsilon > 0$, *the algorithm runs in time polynomial in* $n^{\frac{1}{\epsilon^7}}$, *and outputs a price vector whose revenue is at least a* $(1-\epsilon)$*-fraction of the optimal revenue, where n is the number of items.*

Theorem 4 *(Multiplicative Quasi-PTAS for regular distributions). There is a Quasi-Polynomial-Time Approximation Scheme*[7] *for computing an optimal price vector, when the values of the buyer are independent and drawn from regular distributions.*[8]

---

[3] Incidentally, the problem is trickier than it originally seems, and various intuitive properties that one would expect from the optimal solution fail to hold. See Appendix J for an interesting example.

[4] We note that a natural approach for computing approximately optimal price vectors is to discretize the domain of price vectors and show that searching over the discretized domain suffices for approximating the optimal revenue. However, a straightforward application of the discretizations proposed by Nisan (Chawla et al., 2007) and Hartline and Koltun (2005) to our problem would result in running time of $\left(\frac{1}{\epsilon}\log r\right)^{O(n)}$. The purpose of our theorem is to remove the exponential dependence of the running time on the number of items n.

[5] A Polynomial-Time Approximation Scheme (PTAS) is a family of algorithms $\{\mathcal{A}_\epsilon\}_\epsilon$, indexed by the accuracy parameter $\epsilon > 0$, such that for every fixed $\epsilon > 0$, $\mathcal{A}_\epsilon$ runs in time polynomial in the size of its input. See Section 2 for a formal definition.

[6] Monotone Hazard Rate (MHR) distributions are a commonly studied class of distributions that contain such familiar distributions as the Uniform, Gaussian and Exponential distributions. See Section 2 for a formal definition.

[7] A Quasi-Polynomial-Time Approximation Scheme (Quasi-PTAS) is a family of algorithms $\{\mathcal{A}_\epsilon\}_\epsilon$, indexed by the accuracy parameter $\epsilon > 0$, such that for every fixed $\epsilon > 0$, $\mathcal{A}_\epsilon$ runs in time quasi-polynomial in the size of its input. See Section 2 for formal definition.

[8] Regular distributions are another widely studied class of distributions that contain MHR distributions. See Section 2 for a formal definition.

Please cite this article in press as: Cai, Y., Daskalakis, C. Extreme value theorems for optimal multidimensional pricing. Games Econ. Behav. (2015), http://dx.doi.org/10.1016/j.geb.2015.02.003