Theoretical Computer Science 581 (2015) 83–96

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

Bounding the optimal revenue of selling multiple goods ☆

Yiannis Giannakopoulos

University of Oxford, Department of Computer Science, Wolfson Building, Parks Road, Oxford, OX1 3QD, UK

ARTICLE INFO

Article history:
Received 15 October 2014
Received in revised form 28 February 2015
Accepted 3 March 2015
Available online 6 March 2015
Communicated by A. Fiat

Keywords:
Optimal auctions
Revenue maximization
Duality
Mechanism design

ABSTRACT

Using duality theory techniques we derive simple, closed-form formulas for bounding the optimal revenue of a monopolist selling many heterogeneous goods, in the case where the buyer's valuations for the items come i.i.d. from a uniform distribution and in the case where they follow independent (but not necessarily identical) exponential distributions. We apply this in order to get in both these settings specific performance guarantees, as functions of the number of items m, for the simple deterministic selling mechanisms studied by Hart and Nisan [1], namely the one that sells the items separately and the one that offers them all in a single bundle.

We also propose and study the performance of a natural randomized mechanism for exponential valuations, called Proportional. As an interesting corollary, for the special case where the exponential distributions are also identical, we can derive that offering the goods in a single full bundle is the optimal selling mechanism for *any number of items*. To our knowledge, this is the first result of its kind: finding a revenue-maximizing auction in an additive setting with arbitrarily many goods.

1. Introduction

What selling mechanism should a monopolist with many heterogeneous goods deploy when facing a buyer, in order to maximize his revenue? The buyer has private valuations for the items but the seller can only have an incomplete prior knowledge of these values, in the form of a probability distribution over them. Furthermore, the buyer is strategic and selfish, meaning that if asked to submit her valuations for the goods she will lie if this is to improve her own personal gain. This is one of the most fundamental problems in auction theory, and although it has received a lot of attention from both the Economics and Computer Science communities, it still remains elusive.

1.1. Related work

The problem is famous for demonstrating a striking dichotomy; the case of a single good is fully resolved by the seminal work of Myerson [2]: the optimal selling strategy is to just set a take-it-or-leave it price for the item, and this price is given by a very simple formula involving the probability distribution. However, for more goods, in fact even for just two and three items, our understanding of optimal mechanisms is totally insufficient. The problem seems to be qualitatively completely different and it is widely believed that no closed-form, elegant solutions are within reach (see e.g. [3,4]).

☆ Supported by the European Union FP7-ICT grant 284731 (UaESMC) and ERC Advanced Grant 321171 (ALGAME).
E-mail address: ygiannak@cs.ox.ac.uk.

http://dx.doi.org/10.1016/j.tcs.2015.03.010

First of all, it is known that in general determinism (i.e. setting selling prices for various bundles of goods) is not enough any more and that lotteries need to be deployed for optimality [5–9]. Manelli and Vincent [7] provide some sufficient conditions for deterministic mechanisms to be optimal, but these are arguably rather involved, so they were able to instantiate and verify them only for the case of two and three goods with valuations i.i.d. according to a uniform distribution. Hart and Nisan [1] have also presented a very simple sufficient condition, in the special case of two i.i.d. items, for the full-bundle mechanism (that just sets a single selling price for all the items together) to be optimal and deploy it to show that this holds for the equal-revenue distribution and, more generally, for all Pareto distributions with parameter $\alpha \geq 1/2$. Finally, Daskalakis et al. [9] have shown that full-bundling is also an optimal selling strategy for two exponentially i.i.d. goods. Nothing is clearly known in this front for more than three items or other distributions, apart from the recent work of Giannakopoulos and Koutsoupias [10] where a closed-form description of a deterministic mechanism that is shown to be optimal for up to six uniformly i.i.d. items is given, and conjectured that this holds for any number of goods.

In the current paper we contribute to this line of work, by generalizing the result of [9] from 2 *to* any number of goods, *showing that determinism is optimal for arbitrarily many identical valuations distributed according to an exponential distribution, and in fact optimality is achieved by the simple full-bundling strategy.* To our knowledge, this is the first result that provides specific description of an optimal auction for an arbitrary number of items. In fact, together with [10], they are the first such results to break the boundary of three items.

This path of discovering "simple" descriptions of optimal selling mechanisms is further being narrowed by a recent computational hardness result from Daskalakis et al. [4], where it is shown that even for independent (but not identical) valuations with finite support of size 2, it is #P-hard to compute exactly the optimal mechanism. However this does not exclude the possibility of a PTAS and for i.i.d. settings Cai and Huang [11] and Daskalakis and Weinberg [12] have indeed already provided such efficient *algorithmic* approximations.

So, given the previous discussion, it is essential to try to *approximate* the optimal revenue by selling mechanisms that are as *simple* as possible. We may lose something with respect to the total revenue objective, but on the other hand these mechanisms are much easier to understand, describe, analyze and implement, and such results may in fact enrich our understanding of the character of exact optimal auctions in general. Hart and Nisan [1] provide such solid and elegant approximation ratio guarantees (logarithmic with respect to the number of items) that hold "universally" for all product (independent) distributions, without even assuming standard regularity conditions (like e.g. in [2,7,13]), by studying the two most natural deterministic mechanisms: the one that treats every item independently and sells them separately to the buyer and, on the other end, the one that treats all items as a single full bundle. In particular, they prove an $O(\log^2 m)$-approximation ratio for the former and, with the extra requirement of identical distributions, an $O(\log m)$-approximation for the latter, where m is the number of goods. They also provide slightly improved guarantees for the special case of two i.i.d. goods. Li and Yao [14] further improved their results, by bringing these down to $\Theta(\log m)$ and $\Theta(1)$ respectively, which are also proved to be tight up to constants.[1]

In this paper we try to specialize these general probabilistic results for the case of specific "canonical" distributions, namely the uniform and exponential ones, and we show that by doing so one can get good constant-factor, almost optimal in many cases, performance guarantees. Furthermore, in keeping up with the spirit of the line of work that studies simple but still well-performing mechanisms *we also propose a very natural* randomized *mechanism for exponential distributions and provide good approximation guarantees for it.*

Daskalakis et al. [9] and Giannakopoulos and Koutsoupias [10] have developed duality-theory frameworks for the general problem of multidimensional optimal auctions, the former having a strong measure-theoretic flavor by using classic results from optimal transport theory combined with Strassen's theorem for stochastic dominance, and the latter resembling more in spirit traditional linear programming theory formulations.

1.2. Our results and techniques

Our model is the standard single-buyer multi-item additive valuations Bayesian model of McAfee and McMillan [17] which is used in many other works [1,7,9,14]. Critical to the exposition is the analytical characterization of truthfulness through subgradients of convex functions given by Rochet [18].

Inspired by the elegant approach of Hart and Nisan [1], we take the opposite direction to their universal approximation guarantees for general independent distributions, and try to give better, specialized bounds for the case of uniform and exponential distributions. Our main strategy is driven by the standard technique in approximation algorithms, to use weak-duality from traditional linear programming to upper-bound the optimal objective and then use this to calculate approximation upper bounds for particular algorithms. Since the optimal revenue problem cannot be fully captured by means of traditional combinatorial LPs, we use the duality-theory framework developed in [10]. In particular, we use the weak-duality theorem (see Theorem 1) in order to get specific closed-form bounds for our settings (Theorems 2 and 3), by constructing and plugging-in appropriate feasible dual solutions (Theorems 4 and 5). This is the most technical part of the

[1] In a very exciting recent result, after our paper was first made publicly available, Babaioff et al. [15] showed that combining these two simple selling mechanisms one can guarantee a constant approximation ratio of 6 by just assuming independence of the item valuations. Yao [16] generalized their results to multi-bidder settings.