Electronic Commerce Research and Applications 13 (2014) 128–138

Contents lists available at ScienceDirect

Electronic Commerce Research and Applications

journal homepage: www.elsevier.com/locate/ecra

Approximations to auctions of digital goods with share-averse bidders

Jinxiang Pei *, Diego Klabjan, Wei Xie

Industrial Engineering & Management Sciences, Northwestern University, United States

ARTICLE INFO

Article history:
Received 25 September 2012
Received in revised form 2 November 2013
Accepted 2 November 2013
Available online 6 December 2013

Keywords:
Auction
Share-averse bidders
Constant-approximation algorithms
Prior-free algorithms

ABSTRACT

We consider the case of a digital product for share-averse bidders, where the product can be sold to multiple buyers who experience some disutility from other firms or consumers owning the same product. We model the problem of selling a digital product to share-averse bidders as an auction and apply a Bayesian optimal mechanism design. We also design constant-approximation algorithms in the prior-free setting including both average- and worst-case analyses.

1. Introduction

Consider the problem of selling a valuable piece of information or dataset. In principle, the information or dataset can be sold to all bidders at no marginal cost to the seller. In general, many digital goods share the same property. In other words, digital goods are expensive to produce but cheap to reproduce, since the unit cost of reproduction is negligible and virtually zero. They can be consumed by more than one user at the same time. However, in reality, the value of the information or dataset to a bidder decreases as increasing numbers of bidders obtain the information because the competitive advantage of possessing the information becomes weaker with more receivers of the information. The seller needs to know:

1. What is the optimal number of copies, k that they should sell to?
2. What is the profit maximizing price to charge the k buyers?

Digital product is available in unlimited supply. The firm can sell as many copies as there are buyers. On the other side of this tug-of-war, each of the buyers that obtains a copy incurs some disutility from others obtaining the same product. This indicates that the buyers are willing to pay less if more copies are sold. This represents the basic trade-off to the seller. Since the firms can sell to multiple parties who obtain some disutility from sharing, we call this a digital product with share-averse bidders.

We model this example as an auction. The possible setup is selling either to an individual buyer (standard single item auction) or to multiple buyers depending upon the valuations of all buyers. We introduce a deterministic function to capture the decreasing valuations when sharing the product. To make the problem tractable, we assume that this deterministic function is known to all bidders and hence it is not part of bidder's private information. This assumption implies a single-parameter auction. A distinct attribute of our problem is the task of modeling the number of winners. This can be computed by applying the VCG auction on the virtual values in the Bayesian setting. We apply the well-known Myerson mechanism for maximizing the expected revenue in the Bayesian setting where the values are drawn from a prior distribution. Unlike an auction for standard physical goods, the number of items sold or the number of winners cannot be determined in advance in the auction for digital goods with share-averse bidders. We need to use the prior distribution to compute the number of winners in this auction. In prior-free settings, it is challenging to model the number of winners. We study prior-free auctions and establish the approximation ratio in both average- and worst-cases of the appropriately designed algorithms. We design an algorithm in the prior-free average case where the number of winners can be obtained by the VCG auction on bids. The proposed algorithm approximately maximizes the revenue against a certain benchmark. The algorithm provides good techniques to handle the prior-free approximations and to resolve the issue of coping with the number of winners. We design another prior-free auction/algorithm in the worst case where the bidders are divided into sample and market groups and the number of winners is determined by computing the winning price from the sample group.

* Corresponding author.
E-mail addresses: peipauj@u.northwestern.edu (J. Pei), d-klabjan@northwestern.edu (D. Klabjan), weixie@u.northwestern.edu (W. Xie).

1567-4223/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.elerap.2013.11.001

We first show a single-sample approximation algorithm in the prior-free average case where the revenue benchmark is the expected optimal revenue as shown in the Bayesian setting. By single-sample, we mean that the algorithm is developed over only one sample of bids from all bidders. The average case analysis can be used to compare the revenue performance in this prior-free setting with the Bayesian optimal revenue. The second prior-free approximation algorithm designed is a random sampling algorithm in the worst case. The benchmark in the second approximation algorithm is the revenue of the optimal single price auction. The optimal single price auction is optimal among all auctions where the price by which the winners need pay is unique. In other words, in the optimal single price auction the winners pay the same price. The worst case approximation is more challenging since the approximation ratio is applied to every possible realization of valuations.

An important contribution of our work is the design of prior-free algorithms for auctioning digital goods with share-averse bidders. The single-sample algorithm is a constant-approximation algorithm to the revenue benchmark. In the single-sample auction, the number of winners is a decision variable. Another major contribution of this work is the technique to deal with the number of winners in a complicated single-parameter auction setting. We need to sacrifice the revenue performance for determining the number of winners in the approximations. For example, in the prior-free worst-case analysis, determining the number of winners makes the approximation even less "optimal" in comparison with the standard single-parameter problem without shareability. This is not a problem in the standard single-parameter auction because of the indivisible nature of standard products.

2. Literature review

We mainly differentiate auctions by two features: externality and shareability, see Fig. 1. By externality, we mean the allocation externality where a party obtaining the product influences the remaining parties. Shareability is the property of one unit of the product consumed by more than one party at the same time. At the top level, we divide auctions by prior distributions: Bayesian auctions with known prior distributions and prior-free auctions without any prior distributions. The traditional approach to auction design is to study optimal auctions (i.e., revenue-maximizing auctions) in the Bayesian setting, Myerson (1981). There are cases in the Bayesian setting where shareability has been taken into account, e.g., the patent licensing. The problem of licensing an innovation to firms that are competitors in a downstream market has been well studied. Kamien (1992) provides an excellent survey of patent licensing. Katz and Shapiro (1986) show a licensing game in which the bidders are identical and their signals are publicly observable. In our work, bidder's signal of willingness to pay is private. Schmitz (2002) analyzes a revenue-maximizing auction for a sale of multiple licenses where each bidder's signal is private. All these papers assume no allocation externality, i.e., a firm who gets a license does not affect other firms obtaining no licenses. Our study adds allocation externality into the setting. Our work is also related to the literature on sales with externality. Both Jehiel et al. (1996) and Jehiel and Moldovanu (2000) discuss auctions with externality. However, neither considers the shareability of a product at the same time, i.e., it is impossible to share the product in the auction. We allow multiple bidders to share a product. Salek and Kempe (2008) is closest to our work. They study auctions in which items being auctioned can be shared among multiple winners, and the valuation of winners decreases in the number of winners. They exhibit an optimal truthful auction for a single item in the sense of Myerson. We advance this by studying the prior-free auctions. In addition, in the Myerson's setting our model is different because we do not allow fractional allocations. On the flip side of externality, i.e., positive externality or share-attraction, Haghpanah et al. (2011) and Bhalgat et al. (2012) show how to model bidder's preference regarding positive externality in a social network setting.

Literature on prior-free auction design, the focus of our work, is rare. In practice, the Bayesian approach is restrictive since the prior distribution is usually unknown. A prior-free mechanism design improves understanding of the auction without the assumption of prior distributions. Goldberg et al. (2006) investigate such a prior-free mechanism design problem where the monopolist has a constant marginal cost of supplying units. They completely eliminate the prior distribution assumption in their analysis. Dhangwantnotai et al. (2010) propose single-sample approximations for a prior-free mechanism design. We apply the single-sample techniques to the auctions for digital goods with share-averse bidders and analyze their performance.

The paper is structured as follows. In Section 3 we state the model for auctioning digital goods with share-averse bidders and show the Bayesian optimal mechanism design in Section 4. We then focus on the approximation algorithms and their analysis in Section 5. We conclude the introduction with a literature review.

3. Model

We model the sales of a digital good with share-averse bidders by an auction. In the models that follow we use the following notation. We use bold letters to denote vectors. Let $\mathcal{N}$ be the set of all bidders and n be the total number of bidders, i.e., $|\mathcal{N}| = n$. Let $\mathbf{v} = (v_1, \ldots, v_n)$ be the vector of strict valuations of all bidders for the single digital good with share-averse bidders. Namely, v_i is the valuation of bidder i as the individual winner and v_i is drawn from continuous distribution G_i, which we assume are i.i.d. Let $g(\cdot)$ be the corresponding probability density function of $G = G_i$. We denote by $\mathbf{v}_{-i} = (v_1, \ldots, v_{i-1}, v_{i+1} \ldots, v_n)$ the mask vector after removing bidder i's value and $\mathbf{v}_{-i-j} = (v_1, \ldots, v_{i-1}, v_{i+1}, \ldots, v_{j-1}, v_{j+1}, \ldots, v_n)$ the mask vector after removing both bidder i and j's values. Similarly, we denote the joint distribution function without bidder i by mask vector G_{-i}. We assume that the distribution has the monotone hazard rate property, i.e., $h(\cdot) = \frac{g(\cdot)}{1-G(\cdot)}$ is increasing. We denote $\phi(v_i) = v_i - \frac{1-G(v_i)}{g(v_i)}$ as the virtual value. Let $f(\cdot) : \mathbb{Z} \to \mathbb{R}$ be a decreasing function such that $0 \leqslant f(\cdot) \leqslant 1$. This function models the disutility a bidder obtains from the product being shared with other bidders. For example, if the auction awards 2 copies of the product, then the value to agent i is $v_i \cdot f(2)$. Since we have a finite number of bidders, there exists a (possibly non-unique) optimal number of bidders k such that the revenue obtained by selling to those k bidders is greater than any other number. In case of a tie, we always pick the smaller k. In this setting, we can model bidder's utility using utility function $u_i(\mathbf{v}, k) = x_i(\mathbf{v}) \cdot v_i \cdot f(k) - p_i(\mathbf{v})$, where $x_i \in \{0, 1\}$ is the allocation of an item to agent i or not and $p_i \in \mathbb{R}_+$ is the payment of bidder i to the seller. Finally, for simplicity we assume that all feasible allocations are single units, i.e., there is no additional utility in a bidder receiving multiple copies or a fraction of a copy. This is a reasonable assumption for digital goods with share-averse bidders, given that they are allocated as discrete units.

We denote by

$$Q_i(v_i) = \int_{\mathbf{v}_{-i}} x_i(\mathbf{v}) f\left(\sum_j x_j(\mathbf{v})\right) dG_{-i}(\mathbf{v}_{-i})$$

the conditional allocation to bidder i by which we model the incentive constraint in the Bayesian setting. The seller's problem is to determine a subset $S \subseteq \mathcal{N}$ of bidders to allocate to while maintaining the Bayesian incentive constraint and individual