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

Contents lists available at SciVerse ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Approximately optimal auctions for correlated bidders

Shahar Dobzinski[a,*], Hu Fu[b], Robert Kleinberg[b]

[a] Weizmann Institute of Science, Department of Computer Science and Applied Mathematics, Rehovot, Israel
[b] Cornell University, Department of Computer Science, Ithaca, NY, USA

ARTICLE INFO

Article history:
Received 10 October 2011
Available online xxxx

JEL classification:
D82

Keywords:
Bayesian mechanism design
Approximation Algorithms

ABSTRACT

We consider the design of dominant strategy incentive compatible, revenue-maximizing auctions for an indivisible good, when bidders' values are drawn from a *correlated* distribution. For *independent* distributions, Myerson showed that the optimal auction for risk-neutral bidders remains incentive compatible regardless of bidders' risk attitudes. We show that, for correlated distributions, the same is true when only two bidders are involved, whereas for more bidders, randomization can generate strictly more revenue. However, for risk-neutral bidders, we show this gain is never more than a $\frac{5}{3}$-factor.
This is a consequence of two results of independent interest: (1) a polynomial-time derandomization of auctions for two bidders; (2) a polynomial-time computable deterministic auction that $\frac{5}{3}$-approximates the optimal revenue extractable from risk-neutral bidders.
Moreover, we give a polynomial-time algorithm to compute optimal auctions for a constant number of bidders, and for any number of bidders we give polynomial-time algorithms with approximation factors arbitrarily close to $\frac{3}{2}$.

1. Introduction

We consider the problem of designing a revenue-maximizing auction in dominant strategies for one indivisible item and n players, when players have private values that are drawn from a *correlated* distribution $\mathcal{D}$. Specifically, we seek a dominant strategy incentive compatible auctions that are ex-post individually rational and analyze their revenue in equilibrium.[1]

We will be interested in particular in the risk attitude of the players. We will design auctions for environments where players are risk-neutral (that is, maximize their expected profit) and environments when nothing is assumed regarding the risk attitude of the players (we achieve this by considering deterministic allocation rules). As for terminology, we follow the computer science literature and call a mechanism *truthful in expectation* if bidding truthfully maximizes the expected profit of each player given any valuation profile of other players (and hence incentive compatible for risk-neutral players), and (*deterministic*) *truthful* if its allocation function has a $\{0, 1\}$ range (hence it does not make any assumption regarding attitudes toward risk of the players).

Myerson, in his seminal paper (Myerson, 1981), studies this problem with independent valuations and proves in particular the remarkable property that the optimal auction for risk-neutral players is in fact deterministic.[2] The bulk of this paper studies the following two questions:

* Corresponding author.
E-mail addresses: shahar.dobzinski@weizmann.ac.il (S. Dobzinski), hufu@cs.cornell.edu (H. Fu), rdk@cs.cornell.edu (R. Kleinberg).

[1] Please refer to our Related Work section for a survey on a rich literature discussing optimal auctions in other models, e.g., Cremer and McLean (1985) and Rahman (forthcoming).

[2] In fact, Myerson even shows that the optimal deterministic mechanism is in fact optimal even among Bayesian incentive compatible mechanisms. In this paper we want to guarantee incentive compatibility even in situations in which players are not necessarily aware of the underlying distribution or have an inaccurate knowledge of it, hence we will not be considering Bayesian mechanisms here.

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.03.010

1. Can we obtain a "weak Myerson theorem" for the correlated distribution case? That is, is the optimal auction for risk-neutral players deterministic? If not, can we quantify the revenue gap between truthful-in-expectation auctions and their deterministic counterparts?
2. Can the optimal auction (either deterministic or truthful in expectation) be computed in polynomial time? If not, can we efficiently find incentive compatible auctions that are *approximately* optimal?

We will draw connections between these questions and progress in the study of both. Our investigation start with the beautiful paper of Ronen (2001), in which he presents the following elegant truthful mechanism, the *lookahead auction*. First, find the $n-1$ players with the lowest values. Given their values, consider the conditional distribution of the value v_i of player i with the highest valuation, and calculate a price p that maximizes the expected revenue. If $v_i \geqslant p$, player i is assigned the item and is charged p, otherwise no one is assigned the item and no one pays anything. Ronen (2001) shows that *for every distribution* the lookahead provides in expectation at least half of the revenue of the optimal auction. We consider the k*-lookahead auction*: find the $n-k$ players with the lowest values. Given their values, consider the conditional distribution of the values of the k players with the highest values, and run the optimal auction for these k players. We show that *the k-lookahead auction is a $\frac{3k-1}{2k-1}$-approximation to the revenue of the optimal auction.*[3] Our revenue guarantee is robust in the sense that if the optimal auction that we run for k players is truthful in expectation then the revenue guarantee will be with respect to the optimal (unconstrained) truthful-in-expectation auction, and if the k-players auction is constrained to be deterministic, then the guarantee is with respect to the optimal deterministic auction.

The k-lookahead auction teaches us a valuable lesson: even in markets with no key players, limited competition between a few players suffices to generate almost the same revenue as the optimal auction (up to a small multiplicative factor). In addition, one advantage of the k-lookahead auction is its simplicity: in general, it is easier to design an optimal auction for small number of players than to design an optimal auction for many players. Our next steps are to formalize this intuition and then make use of it.

We observe the following simple fact: the optimal truthful-in-expectation auction can be found by solving a natural linear program that encodes allocation probabilities, expected payments, and the incentive constraints that link them.[4] Can we exploit this to obtain polynomial-time algorithms for computing optimal auctions? To answer this question we have to concretely define the computational model in hand. One can first consider a model where the running time has to be polynomial in the support size of the distribution. We call this the *explicit model*. Papadimitriou and Pierrakos (2011) prove that it is possible to exactly compute the optimal deterministic auction for 2 players, but NP-hard to do so for more players. In contrast, we show that using the LP approach computing the optimal truthful-in-expectation auction is easy in the explicit model, for any number of players.

Unfortunately, while the explicit model might be useful for settings with a small number of players, for large numbers of players the distribution usually has exponentially large support, e.g. when each player has two possible values and valuations are independent. Therefore, in this paper we mainly use the *oracle model* of Ronen and Saberi (2002): the distribution is given to us as a black box and we are allowed to ask *conditional-distribution* queries. That is, given the values of $n-k$ players, what is the conditional distribution of the values of the remaining k players?

In the oracle model, the optimal truthful-in-expectation auction can be found in polynomial time using the LP approach *for a constant number of players*. Combining this with our results regarding the revenue guarantee of the k-lookahead auction, we have that for any distribution we can compute in polynomial time an auction for risk-neutral players that generates in expectation $\frac{2k-1}{3k-1}$ of the revenue of the optimal auction for risk-neutral players.

Our next result is the following general implementation result: *for every truthful-in-expectation mechanism for two players and a single item, there exists a mechanism that is a probability distribution over deterministic truthful mechanisms (a* universally truthful *mechanism) with the same allocation function and the same (expected) payments*. The result is in the same spirit of a result by Manelli and Vincent (2010) that show that every Bayesian mechanism with *independent* valuations (for any number of players) can be implemented as a truthful-in-expectation mechanism. Our result shows that for two players one can *always* take any incentive compatible mechanism for risk-neutral players and make it incentive compatible without assuming any risk attitude of the players.

In particular we get that the optimal auction for risk-neutral players can be implemented as a probability distribution over deterministic mechanisms. In this case we can further strengthen this result: since all mechanisms in the distribution must have the same expected revenue, a standard purification argument gives us that by fixing the random coins we get a deterministic mechanism with the same revenue as the optimal mechanism for risk-neutral players. That is, we answer question 1 from the beginning of the introduction in the affirmative: there is a weak Myerson theorem for environments with two players. Unfortunately, we also show an example that proves where this weak Myerson theorem fails for three or more bidders.

[3] Ronen (2001) claims that the approximation ratio of the k-lookahead auction is no better than 2, for every k, but his proof is incorrect. He essentially claims (without a proof) that all truthful mechanisms achieve an expected revenue of at most 1 for a certain distribution for two players. However, the pivot auction we define in Section 5 provides an expected revenue of 1.5 for that distribution.

[4] Although some caution is needed as we have to define the mechanism for all possible values in the domain, not just for values in the support of the distribution.