On Approximating Optimal Auctions

(extended abstract)

Amir Ronen*
Department of Computer Science
Stanford University
(amirr@Robotics.Stanford.EDU)

Abstract

We study the following problem: A seller wishes to sell an item to a group of self-interested agents. Each agent i has a *privately* known valuation v^i for the object. Given a distribution on these valuations, our goal is to construct an auction that maximizes the seller's expected revenue (optimal auction). The auction must be incentive compatible and satisfy individual rationality. We present a simple generic auction that **guarantees** at least half of the optimal revenue. We generalize this result in several directions, in particular, for the case of multiple copies with unit demand. Our auction requires the ability to learn (or compute) in polynomial time the conditional distribution of the agent with the maximal valuation, given the valuations of the other agents. We show that this ability is in some sense essential. Finally we suggest a generalization of our auction and argue that it will generate a revenue which is close to optimal for reasonable distributions. In particular we show this under an independence assumption.

1 Introduction

Auctions of various kinds play a major role in economics and electronic commerce. They also give rise to many interesting theoretical questions. Of particular interest, both practical and theoretical, is the issue of revenue maximization, also known as optimal auction design.

In this paper we consider the following problem: A seller wishes to sell one item (e.g a house) to a group of self-interested agents. Each agent has

*This research was supported by Darpa grants number F30602-98-C-0214 and F30602-00-2-0598.

a *privately known* valuation $v^i \in [1, h]$ for the object. If agent i wins the auction, her total profit is $u^i = v^i - p$, where p denotes the agent's payment. If she looses, her valuation is zero. The goal of each agent is to maximize her own profit. An auction is a protocol that decides who wins the item and for what price. Since rational agents may *manipulate* the given protocol if it is beneficial for them, or may simply refuse to participate in it, we focus on auctions that satisfy two standard requirements:

Incentive compatibility (IC) Each agent has a dominant strategy, i.e. a strategy which is always better than any other strategy.

Individual rationality (IR) The profit of an agent who behaves according to her dominant strategy is always non negative.

Individual rationality is desirable and even necessary to many applications. Incentive compatibility guarantees that rational agents will indeed behave according to their dominant strategies. We say that an auction is *valid* if it satisfies both requirements.

The optimal auction problem is described by the following: Given a distribution of the agents' valuations, find a valid auction that maximizes the seller's expected revenue provided that the agents behave according to their dominant strategies.

This problem is a subject of long and intensive research in micro economics. The major thread of research focuses in *characterizing* the optimal auction. The problem is solved ([10] and others) for the case where the agents' valuations are *independent* (and the distributions obey some regularity conditions). Unfortunately, this is barely the case and little progress has been made for general distributions. Under much weaker requirements of IR and IC (and under additional assumptions), beautiful results of [10, 1] show that the seller can extract the expected first order statistics. The assumptions of these theorems are very strong and seem unrealistic for most applications (see e.g. a discussion in [6]). A comprehensive recent survey of auction theory can be found in [7].

In this paper we study the optimal auction problem using a different and more computer science oriented approach. Instead of characterizing the optimal auction, we look for an algorithm that *computes* it efficiently. Since the corresponding problem cannot be solved in polynomial time, we look for an *approximation* – a solution that guarantees a revenue of at least $1/c$ of the optimum (for a fixed $c \geq 1$). We note that when such approximations are available, they often perform much better in practice than their worst case performance. Moreover, in cases where the seller knows an auction which is not an approximation but seems to function well on many distributions, she can compare it to the approximation and take the better one. By doing this, the seller can both guarantee to have at least $1/c$ of the optimal revenue, and at the same time, exploit her heuristic knowledge. Such comparison can be done e.g. by sampling the distribution and comparing the average revenues.

We introduce a simple generic auction called *1-lookahead*. In this auction the bidders are simply requested to report their true valuations. The auction is designed in such a way that reporting the true valuations is a dominant strategy. Let w^i denote the declaration of agent i. W.l.o.g. assume that $w^1 \geq w^2, \ldots, \geq w^n$. Our auction rejects all bids except the highest one w^1. According to the rejected bids, it computes a price $p = p(w^2, \ldots, w^n)$. If $w^1 \geq p$, the corresponding agent wins and pays p. Otherwise, nobody wins. We show that this auction is valid and extracts a revenue of **at least half** of the optimum. The auction is simple and can be implemented as a standard English (or Japanese) auction with one final price increment by the seller. Similar results are obtained for the case of multiple copies with unit demand, risk averse or risk seeking seller and buyers and for auctions that instead of incentive compatibility demand only the existence of Bayesian equilibrium.

Our auction is "fair" in the sense that the highest offer is always preferred over lower ones. It is also simple and traditional. Our approximation ratio shows that the loss because of these properties is not too large.

We require that the computation of the auction and the auction itself are executed in polynomial time. Let the conditional first order statistics be the distribution of the maximal valuation v^1 given the other valuations $(v^2, \ldots, v^n)$. In order to compute the 1-lookahead auction, the auctioneer must be able to compute this distribution in polynomial time. We show that this assumption is in some sense critical. We consider a setting where the designer can only *sample* the underlying distribution. We show that if pseudo random functions exist, it is impossible to do better than a $1/2 \cdot \log(h)$-approximation. A simple auction which does not use the conditional first order statistics was proposed in [3]. Using that paper's technique, we show that this auction is a $2 \cdot \log(h)$-approximation. These results are given in the context of learning.

The approximation ratio of our 1-lookahead auction is tight. Moreover, in many interesting cases, other intuitive auctions will perform better. Thus, we suggest a natural generalization of our auction called *k-lookahead* . Although the approximation ratio of this auction is also 2, we argue that it will extract a revenue which is close to optimal for most reasonable distributions. For the case where $(v^1, \ldots, v^k)$ are independent, we show that this is a $(k+1)/k$-approximation.

Optimal auction design is a fundamental problem in micro economics. A guide to the extensive literature on this topic as well as a collection of important papers can be found in [7]. The first auction that satisfies IR and IC was formally introduced in [12]. Two recent papers [3, 2] pioneer the study of optimal auctions from the computer science point of view. These papers consider the case of unlimited supply which is essentially different from our case. Nevertheless, some of the tools established in these papers can be applied to our problem as well. More general mechanism design problems were studied by many researchers in recent years, but as far as we know, not in the context of revenue maximizing.

Organization of the paper: Section 2 formally defines our problem. Section 4 defines the 1-lookahead auction and proves that it is a 2-approximation. Section 5 shows the necessity of the conditional first order statistics and section 6 describes the k-lookahead auction.

2 The Problem

In this section we formally present our problem and notations. A seller wishes to sell[1] one item (e.g a house) to a group of self-interested agents. Each agent has a *privately known* valuation $v^i \in [1, h]$ for the object. If agent i wins the auction, her total profit is $u^i = v^i - p$, where p denotes the agent's payment. If she looses, her valuation is zero. The goal of each agent is to maximize her own profit.

Notations: We let $[n] = \{0, 1, 2, \ldots, n\}$. We denote the possible types (valuations) of each agent i, by $W^i = \{1, 1+\epsilon, 1+2\epsilon, \ldots, 2, 2+\epsilon, \ldots, h\}$ and let $W = (W^1, \ldots, W^n)$. We use the following vectorial notation: given a vector $a = (a^1, \ldots, a^n)$ we let $a^{-i} = (a^1, \ldots, a^{i-1}, a^{i+1}, \ldots, a^n)$ and (b^i, a^{-i}) denote the vector $(a^1, \ldots, a^{i-1}, b^i, a^{i+1}, \ldots, a^n)$.

An auction is a protocol that decides who wins the item and for what price. The simplest type of auctions are protocols in which the agents are simply required to declare their types (revelation auctions). According to these declarations the auction determines who has won and for what price. We say that an agent i is *truthful* if she declares her actual type. We call an auction truthful if truth-telling is dominant for all agents. By the revelation principle [9, pp. 871] for every valid auction there exists a truthful revelation

[1] All our results apply for reverse auctions as well.

auction that yields the same revenue. Thus, we can limit ourselves to such auctions. Formally:

Definition 1 (auction) *An* auction *is a pair of functions* (k, p) *such that:*

- $k : W \to [n]$ *is an allocation algorithm determining who wins the object. (a zero value means that nobody won.)*
- $p : W \to R_+$ *is a payment function determining how much the winner must pay.*

As agents may manipulate the auction or refuse to participate in it we impose the following standard requirements on the set of auctions that we allow:

Definition 2 (valid auction) *We call an auction* (k, p) valid *if it satisfies the following conditions:*

Individual rationality (IR) *The profit of a truthful agent is* always *non negative. I.e.* $p(w) \leq w_{k(w)}$.

Incentive compatibility (IC) *Truth-telling is a dominant*[2] *strategy for each agent.*

Given the distribution ϕ on the type space, we can compare between valid auctions according to their expected revenue. In particular we can define optimal and approximating auctions.

Definition 3 (c-approximation) *Let* ϕ *be a distribution on the type space of the agents,* $m = (k, p)$ *a valid auction and* $c \geq 1$*. The* revenue $\bar{R}_m$ *of* m *is the expected payment* $E_{v \in \phi}[p(v)]$*.* m *is called* c-approximation *if for every valid auction* m'*,* $\bar{R}_m \geq 1/c \cdot \bar{R}'_m$*. It is called* optimal *if it is a* 1-*approximation.*

Consider an algorithm that accepts as input a probability distribution on the type space and returns an auction. We say that the algorithm solves the **c-approximation optimal auctions** problem if it returns a c-approximation for every distribution ϕ. Both, the algorithm and the auction, must have a polynomial computational time. Until section 5 we ignore the actual representation of the distribution ϕ.

[2]Most results in this paper apply for Bayesian incentive compatibility as well.

Note that there are many valid auctions. Let k be an allocation algorithm with the following property: If $i = k(w)$ and w^i increases then i keeps winning. One can define the payment $p(w)$ as the minimal w^i such that $i = k(w^i, w^{-i})$. It is not difficult to show that this characterizes the family of valid auctions [11, 8]. Since this family of auctions is huge and complex, it is unlikely that the optimal auction can be found in polynomial time. Currently, we do not know what is the complexity of the optimal auction problem.

It is possible to formulate the optimal auction problem as an integer program. For every possible $v = (v^1, \dots v^n)$, let $p_k^i(v) \in \{0,1\}$ be 1 iff agent i wins the auction and pays k. One possible formulation is:

$$\max \sum_{v,i,k} k \cdot p_k^i \cdot \phi(v) \text{ such that:}$$

$$\textbf{IR} \quad \text{For all } i, v, k\text{:} \quad p_k^i(v) \cdot (v^i - k) \geq 0$$

$$\textbf{IC} \quad \text{For all } i, v^i, w^i \geq v^i, v^{-i}, k\text{:} \quad p_k^i(v^i, v^{-i}) \leq p_k^i(w^i, v^{-i})$$

$$\textbf{Uniqueness} \quad \text{For all } v\text{:} \quad \sum_{i,k} p_k^i(v) = 1$$

Unfortunately, the number of variables of this program $(n \cdot ((h-1)/\epsilon)^{(n+1)})$ is *exponential* in the number of agents. Thus, this approach is feasible only when the number of agents is small. It is not difficult to see that the above matrix is not totally unimodular.

3 Some Basic examples

This section gives a few toy examples of optimal auctions. All the examples consist of two agents Agent1 and Agent2, a simple distribution and an optimal valid auction. Real life problems are of course much more complicated.

Independent valuations: Consider the following distribution: Each agents i's valuation v^i is uniformly distributed in $[0, 100]$. The valuations are independent. Consider the following solution known as the second price auction [12]. Each agent declares her type. Let w^i denote their declarations. (Note that they can lie!) The object is given to the agent with the maximal declaration for the price of the second one. The reader can verify that this is a valid auction. It is possible however to generate higher revenue. Consider the same auction where the seller has a reserved price of \$50. A convenient way to describe this auction is as a second price auction with the addition of one dummy agent Agent0. This agent's valuation is fixed to \$50. The reader may verify that the revenue of this auction is indeed higher. The optimality of this auction is implied by classic results in auction theory (see e.g. [10]).

Correlation: Consider the following distribution: Agent1's valuation v^1 is uniformly distributed in $[0, 100]$; $v^2 = 2 \cdot v^1$. Clearly a second price auction, even with reserved prices, is not optimal here. Consider the following solution: Denote the higher declaration by $\overline{w}$ and the lower declaration by $\underline{w}$. Let $p = 2 \cdot \underline{w}$. The lower agent is rejected. If $\overline{w} \geq p$, the high agent wins and pays p. Otherwise nobody wins. The reader can verify that this is indeed a valid auction. When both agents are truthful the auction extracts all possible revenue (i.e. $p = \overline{w}$). Thus, this auction is clearly optimal among the valid ones.

Anti correlation: Agent1's valuation v^1 is uniformly distributed in $[0, 100]$; $v^2 = 100 - v^1$. The optimal auction in this case is similar except that $p = \max(100 - \underline{w}, \underline{w}$. Note that while in a "classic" auction it is unlikely that such an anti-correlation will occur, it is not hard to imagine it in the context of reverse auctions and resource allocation problems.

4 A 2-approximation auction

This section presents a simple generic 2-approximation auction. From now on we order the agents in a decreasing order of their bids, i.e. $w^1 > w^2 > \ldots > w^n$. (Ties are broken arbitrarily.) Given the distribution ϕ and $(v^2, \ldots, v^n)$ let ϕ_1 denote the conditional distribution on the agent with the maximal valuation.

Definition 4 (1-lookahead auction) *The auction computes the price $p^1 = p^1(w^2, \ldots, w^n)$ that* maximizes *its revenue from agent1 (according to ϕ_1!). If $w^1 \geq p^1$, agent1 wins and pays p^1, otherwise nobody wins.*

In other words, all the agents except the one with the maximal offer are rejected. Based on the rejected agents *only*, the auction proposes a price p^1 to agent1. Agent1 wins iff $w^1 \geq p^1$. Since by setting $p^1 = w^2$ the auction can guarantee itself a revenue of w^2, it must be that $p^1 \geq w^2$. Such an auction can be implemented as a standard English (or Japanese) auction[3]

[3] This is highly desirable for many applications.

with one final price increment by the seller. We now show that this is a two approximation.

Theorem 4.1 *The* 1*-lookahead auction is a* 2*-approximation.*

Proof: Clearly our auction satisfies IR. We now show that it satisfies IC as well. Fix the declarations of the other agents and consider agent i. We need to show that declaring v^i is a dominant strategy. If $v^i > w^2$ and the agent is truthful, she wins and pays p^1 iff $v^1 \geq p^1$. In the case of $v^1 \geq p^1$, the agent would like to win and therefore truth-telling is optimal for her. If $w^2 < v^1 < p^1$ then the only way the agent can affect her profit is by declaring $w^1 \geq p^1$. In this case her profit will be $v^1 - p^1 < 0$, therefore she is better of losing the auction. Consider the case of $v^i < w^2$. Since $p^1 \geq w^2$, in order to win the agent must pay *at least* w^2 and therefore lose money. Thus in this case, truthfulness is optimal for the agent as well. It remains to show that this auction is a 2-approximation. Let $\bar{R}$ denote the revenue of the 1-lookahead auction, let m' be another mechanism and let $\bar{R}'$ denote its revenue. We need to show that $\bar{R} \geq 2 \cdot \bar{R}'$. We present the expected revenue of m' as the sum of two disjointed cases: when it picks the highest bid and when it is not. We denote the contribution of each of these cases by $\bar{R}'_1$ and $\bar{R}'_2$ respectively. Clearly $\bar{R}' = \bar{R}'_1 + \bar{R}'_2$.

Claim 4.2 $\bar{R} \geq \bar{R}'_1$

Proof: We first note that in the case of a single bidder, it is known that the auction must offer the agent a threshold p such that the agent wins and pays p iff $w^1 \geq p$. Fix $(v^2, \ldots, v^n)$ and consider agent1. Since we choose p^1 as the price that maximizes the expected payment of agent1, we get that it is at least the expected payment of agent1 in m'. By integrating over all possible tuples $(v^2, \ldots, v^n)$, we prove our claim. □

Claim 4.3 $\bar{R} \geq \bar{R}'_2$.

Proof: Fix $(v^2, \ldots, v^n)$. Because of the individual rationality, m' cannot get more than v^2 from agents $(2 \ldots n)$. On the other hand, since by setting $p^1 = v^2$, our mechanism can guarantee a revenue of v^2, we get that it extracts from agent1 at least v^2 in expectation (over v^1). We prove our claim by integrating over all possible $(v^2, \ldots, v^n)$. □

As clearly, $\bar{R}' = \bar{R}'_1 + \bar{R}'_2 \leq 2 \cdot \bar{R}$, our theorem is proven. □

This simple principle can be generalized in several ways. For the case of k copies of the same item with unit demand, the seller can reject all but the

k-highest offers and then construct the optimal auction on the remaining agents. Similar analysis shows that it is a 2-approximation as well. The analysis remains true in the case of a risk-averse or risk-seeking seller and even when truth-telling is only required to be a Bayesian equilibrium (but IR remains a hard constraint) and the buyers have any risk profile. A first price version of this auction is possible as well. We do not know if similar principles can be applied to more complex problems and in particular to combinatorial auctions.

Our auction is "fair" in the sense that the highest offer is always preferred over lower ones. It is also simple and traditional. Theorem 4.1 shows that the loss because of these properties is not too large.

The approximation ratio of 2 is tight for the 1-lookahead auction. Consider a case of two agents. Agent2's type is fixed to 1. v^1 is determined according to the following distribution:

$$Pr[v^1 = k] = \begin{cases} 1/h & k = h \\ 1 - 1/h & k = 1 + \epsilon \end{cases}$$

The optimal revenue in this case is about 2 while our auction extracts a revenue of around 1. The approximation ratio of our auction improves when $E[v^1/v^2]$ is small.

5 Computational issues

Our main goal in this paper is to ensure that both the computation of our auction and the auction itself are performed in polynomial time. This section shows that the 1-lookahead auction satisfies this requirement provided that the distribution $v^1|(v^2, \ldots, v^n)$ can be computed in polynomial time. We then show that this assumption is in some sense essential: if the above distribution is hard to learn or compute, then no auction can do better than a $(1/2 \cdot \log h)$-approximation. A simple auction which does not require the above distribution was suggested in [3]. Using that paper's technique, we show that this auction is a $2 \cdot \log(h)$-approximation.

Definition 5 (first order statistics) *Let ϕ be a distribution over the type space. A* first order statistics *is a polynomial time algorithm that gets a price k and computes $Pr_\phi[max(v^1, \ldots, v^n) \geq k]$.*

Definition 6 (conditional first order statistics) *Let ϕ be a distribution on the type space. A* conditional first order statistics *is a polynomial time algorithm that gets a price k and $(v^2, \ldots, v^n)$ and computes $Pr[v^1 \geq k|(v^2, \ldots, v^n)]$.*

When the actual distribution ϕ can be computed polynomially by the designer, she can also compute the conditional first order statistics. A natural interpretation of the results of this section is in a learning setup. In such a setup, the designer can only *sample* ϕ. Based on these samples, she computes the desired auction. This section is greatly inspired by [3].

5.1 Computing the 1-lookahead auction

When the conditional first order statistics is available, the computation of the 1-lookahead auction is trivial. The designer can simply try all possible prices k and take the one that maximizes $k \cdot Pr[v^1 \geq k|(v^2, \ldots, v^n)]$. If h is large, it is also possible to try only the prices $(v^2, \alpha \cdot v^2, \alpha^2 \cdot v^2, \ldots, h)$ and get an α-approximation of the optimal price. Similarly, the designer can compute the optimal price (with high probability) when she can only sample the conditional statistics.

5.2 Without the conditional first order statistics

In this subsection we study a setup where instead of having an oracle for the distribution ϕ, the designer can only sample it. Both the auction construction and the auction itself must be polynomial time computable. For such a setting we show that if pseudo random functions exist [4], it is impossible to do better than a $1/2 \cdot \log h$-approximation. A simple $2 \cdot \log h$-approximation is shown in theorem 5.6.

Consider the following distribution[4] ϕ:

$$Pr[v^1 = k] = \begin{cases} 1/2 + 1/2h & k = 1 \\ 1/2k & k = (2, 4 \ldots, h) \end{cases}$$

For simplicity we assume that h is a power of 2. Consider the case of one bidder with this distribution. By asking a price of 1, the seller gets a revenue of 1. It is not difficult to see that this is optimal. We now show that this remains true even if we weaken the IC requirement by allowing probabilistic auctions and require that truth telling only maximizes the agent's expected profit[5].

Lemma 5.1 *No auction that satisfies IC in expectation and IR can extract revenue greater than* 1 *on* ϕ.

[4] ϕ was introduced at [3] as an example of a distribution which is bad for auctions.

[5] Alternative proofs rely on mechanism design theory.

Proof: Let θ_i and m_i denote the winning probability and the expected payment when the agent declares 2^i, $(i = 0, \ldots, \log(h))$. Let v denote the actual valuation of the agent. The profit of the agent when declaring 2^i is $\theta_i \cdot v^1 - m_i$. Note that because of the IR requirement, it must be that $m^i \leq 2^i$. Let $R_i = m_i/2^{i+1}$. Let $R' = \sum_i R_i$. Note that for $i \geq 1$, R_i is equal to the contribution of the type 2^i to the total revenue. Thus $R' = 1/2h + \sum_i R_i$.

When the type of the agent is 2^{i+1}, it is not beneficial for her to declare 2^i. Therefore:

$$2^{i+1} \cdot \theta_{i+1} - m_{i+1} \geq 2^{i+1} \cdot \theta_i - m_i$$

Thus:

$$\begin{aligned} \theta_1 - \theta_0 &\geq 2 \cdot R_1 - R_0 \\ \theta_2 - \theta_1 &\geq 2 \cdot R_2 - R_1 \\ &\vdots \\ \theta_h - \theta_{h-1} &\geq 2 \cdot R_h - R_{h-1} \end{aligned}$$

By summation we get that $\theta_h - \theta_0 \geq R' + (R_h - 2 \cdot R_0)$ and thus $R' \leq (\theta_h - R_h) + (2 \cdot R_0 - \theta_0) \leq 1 - 1/2h$. Therefore $R \leq 1$. $\square$

We now have a distribution in which any valid auction can obtain a revenue of no more than 1. On the other hand, if the designer knew the agent's type, she would be able to extract a revenue of more than $1/2 \cdot \log h$. Consider the following setting: There are n agents. Agent1's type is always maximal and is distributed according to ϕ. The types of agents $(2, \ldots, n)$ are always smaller than Agent1's smallest type. For convenience we assume that $v^{i \geq 2} \in \{\epsilon, 2\epsilon\}$ (although these types are less than 1). Consider the case where $(v^2, \ldots, v^n)$ uniquely determine v^1 but this dependency is **hard** to learn (or compute). If this dependency is hard enough, a polynomial auction will not be able to *distinguish* between this case and the case where v^1 is *independent* of $(v^2 \ldots, v^n)$. As v^1 is distributed according to ϕ, any polytime auction must have a revenue around 1. In the rest of this section we formally define hard functions, show that if hard functions exist, polytime auctions cannot do better than $(1/2 \cdot \log h)$-approximation, and present an auction that achieves a $(2 \cdot \log h)$ approximation ratio. In order to avoid the discussion of pseudo randomness as much as possible, we use non-standard definitions and notions. A good reference for pseudo randomness is [5].

Some conventions: When considering algorithms that accept a distribution as input we assume that they can *sample* this distribution in one unit of time (and therefore learn it). We call an algorithm T-algorithm if its running time is bounded by T. We let $\circ$ denote the string concatenation operator and if not stated otherwise let the operator $\in$ denote the uniform choice.

Definition 7 (T-pseudo randomization) *Let θ be a distribution and $\mathcal{D}$ a family of polytime distributions (over the same space as θ). We say that $\mathcal{D}$ is a T-*pseudo randomization *of θ if for every T-algorithm A, $|Pr_{d \in D}[A(d) = 1] - Pr[A(\theta) = 1]| \leq 1/3$.*

Let P denote the distribution where $(v^2, \ldots, v^n)$ are chosen uniformly from $(\epsilon, 2\epsilon)$ and v^1 is independently drawn from ϕ (of lemma 5.1). The existence of a pseudo-randomization for P follows naturally from standard assumptions in cryptography. One possible way to see this is to construct a pseudo randomization of P from a pseudo random function [4]. Roughly speaking, $\psi(k, x)$ is pseudo random if it can be computed in polynomial time but looks random for any polytime algorithm which does not have k. We say that ψ is injective of for every k, $\psi(k, .)$ is injective. This requirement is satisfied by the standard constructions of pseudo random functions.

Lemma 5.2 *Let $T << 2^n$. If injective pseudo random functions exist then there exists a T-pseudo randomization of P.*

Proof:(sketch) Let $\psi(k, x)$ be a pseudo random function. For every key k, define $l_k(x) = \psi(k, x)$ where x is a string of length $n - 1$ and $l_k(x)$ is of length $h+1$. Let $g : \{0, 1\}^{h+1} \to \{0, 1, 2, 4, \ldots, h\}$ such that when y is chosen uniformly, $g(y)$ is distributed according to ϕ. Define a family $\mathcal{D} = \{d_k(.)\}$ by $d_k(x) = g(x \circ l_k(x))$. It is not difficult to see that as long as $T << 2^{(n-1)}$, $\mathcal{D}$ is a T-pseudo randomization of P. □

For small values of n, it is possible to obtain a similar construction by enlarging the set of possible types for agents $(2 \ldots n)$.).

Given a distribution d on the type space and auction A, let $R_A(d)$ be a random variable that denotes A's revenue and let $\bar{R}_A(d)$ denote its expectation. As before we let $R_{opt}(d)$ to denote the revenue of the optimal (possibly exponential) valid auction on d. Note that a designer with an unlimited computational power can fully learn d and construct the optimal auction.

Theorem 5.3 *Let $T > 0$, $0 < \epsilon < 1$ and let $m = 3/\epsilon^2$. If an $((m + 1) \cdot T)$-pseudo randomization for P exists, then for every T-auction A, there exists a distribution d such that $\bar{R}_A(d) \leq 1 + 5 \cdot \epsilon$, but $R_{opt}(d) > 1/2 \cdot \log(h) - \epsilon$.*

Proof: Let $\mathcal{D}$ be an $((m+1)\cdot T)$-pseudo randomization for P and assume by contradiction an auction A that violates the theorem. We will define an $((m+1)\cdot T)$ algorithm A' that distinguishes between P and $\mathcal{D}$ as follows: Given an input distribution d, run the auction m times. If the average revenue is better than $1+3\cdot\epsilon$, output 1, otherwise 0.

Claim 5.4 $Pr_{d\in\mathcal{D}}[A'(d)=1]\geq 0.75$

Proof: Let $d\in\mathcal{D}$. Denote by $\bar{R}^1_A(d)$ the expected payment of Agent1. Since A violates the theorem's assumption, $\bar{R}^1_A(d)\geq 1+3\cdot\epsilon$. From the Chernoff bound $Pr[R^1_A(d)<1+\epsilon]\leq Pr[R^1_A(d)<(1-\epsilon)(1+3\cdot\epsilon)]\leq \exp(-\epsilon^2\cdot m/2)<0.25$ $\square$

Claim 5.5 $Pr[A'(P)=1]\leq 0.4$

Proof: By lemma 5.1, $\bar{R}^1_A(P)\leq 1$. Thus, by the other direction of the Chernoff bound, $Pr[R^1_A(P)\geq 1+\epsilon]\leq \exp(-\epsilon^2\cdot m/3)<0.4$ $\square$

Thus, A' separates between P and $\mathcal{D}$ – a contradiction. Therefore for the majority of the distributions $d\in\mathcal{D}$, $\bar{R}_A(d)\leq 1+5\cdot\epsilon$. Finally, most $d's$ have a distribution of v^1 which is close to ϕ. Otherwise it is possible to distinguish between P and $\mathcal{D}$ according to v^1. Therefore there exists a d such that $E[v^1]\geq 1/2\cdot\log(h)-\epsilon$ and $\bar{R}_A(d)\leq 1+5\cdot\epsilon$ as requested. $\square$

Although, this example may look somewhat artificial, it represents a "real" phenomenon. If the dependency between the agents is hard to learn or compute, we cannot expect to approximate the optimal auction. Finally, we show that a trivial auction achieves an approximation factor of $2\cdot\log(h)$. This result was proven in [3] under a different setting.

Definition 8 (Vickrey auction with reserved price) *Let* $r\geq 0$. *The* Vickrey auction with reserved price r *is the following auction: If* $v^1<r$, *all agents are rejected. Otherwise agent1 wins and pays* $\max(v^2,r)$.

Proposition 5.6 *There exists a price r such that the Vickrey auction with reserved price r is a $2\cdot\log(h)$ approximation.*

Proof:[3] Let d be a distribution and let $\bar{v}^1(d)$ denote the expectation of v^1. Consider the intervals $I_i=[2^i,2^{i+1})$ and the (mutually exclusive) events that $v^1\in I_i$. There are $\log(h)$ such intervals. Let I_i be the interval that contributes most to $\bar{v}^1(d)$. Define $r=2^i$ and let $\bar{R}(d)$ denote the revenue of the Vickrey auction with reserved price r. Clearly $\bar{R}(d)\geq 1/(2\log(h))\cdot\bar{v}^1(d)\geq 1/(2\log(h))\cdot R_{opt}(d)$. $\square$

In order to compute such an r, the designer needs to have only the first order statistics available. For learning this statistics, it is sufficient to learn $\log(h)$ Boolean variables, each converges exponentially to its expectation.

It is possible to show a matching bound of $\ln(h)$ by considering the distribution that satisfies $k \cdot Pr[v^1 \geq k] = 1$ for all k. The number of Boolean variables that need to be learnt increases from $\log(h)$ to h though.

6 A k-lookahead auction

The 1-lookahead auction rejected all but the highest offer and then computed the optimal auction for agent1. For many distributions of interest, this does not look like a reasonable solution. A natural generalization of this auction is to reject all but the k-highest bidders. Then, according to the conditional distribution of the remainders, compute the optimal auction for them. We call this auction k-lookahead. Formally:

Definition 9 (k-lookahead auction) *Let ϕ' be the conditional distribution $\phi((v^1, \ldots, v^k)|(v^{k+1}, \ldots, v^n))$. The* ($k$-lookahead auction) *is the optimal auction on agents $(1, \ldots, k)$ according to ϕ'.*

Clearly, this auction is at least as good as the 1-lookahead auction and therefore a 2-approximation. Surprisingly, this ratio is tight. We first demonstrate this for the case of $k = 2$. Consider the case of three agents with the following distribution:

- Agent3's type is always 1.
- Agent2's type is uniformly drawn from $\{1 + j \cdot \epsilon\}$ where $j = (1, \ldots \log(h))$.
- Agent2 determines the probability which the type of Agent1 is drawn from. If $v^2 = 1 + j \cdot \epsilon$ then $v^1 = 2^j$ with probability $1/2^{j+1}$ and $1 + (j + 1) \cdot \epsilon$ with probability $1 - 1/2^{j+1}$.

It is not difficult to show that our auction extracts a revenue around $1 + 1/\log(h)$. The optimal auction on the other hand uses Agent2 as an "indicator" of Agent1. It asks from Agent1 a price of 2^j. If $v^1 < 2^j$, it gives the item to Agent3 for the price of 1. The revenue of this auction is of around 2. To generalize this example for larger values of k, consider a similar case where agents $(2, \ldots, k)$ determine v^1's distribution in a fully sensitive way – each change in v^i flips v^1's distribution. Similar analysis shows that the revenue of the k-lookahead is around 1 while the optimal revenue is about 2. This example, however, is very artificial. Moreover, when the agents are highly dependent, the seller takes a lot of risk when relying on the agents' information. We argue that the k-lookahead auction will extract revenues which are close to optimal on reasonable distributions. In particular, we show this for the case where $(v^1, \ldots, v^k)$ are independent (but not necessarily independent of $(v^{k+1}, \ldots, v^n)$).

Theorem 6.1 *If $(v^1, \ldots, v^k)$ are independent, the k-lookahead auction is a $(k+1)/k$-approximation.*

Proof: Fix the values of the $n-k$ lowest valuations and the agents who have them. For convenience assume that these are agents $(k+1, \ldots, n)$. Let A_{opt} be the optimal auction and let $\bar{R}_{opt}(v^{k+1}, \ldots, v^n)$ denote its revenue in this case (i.e. the expected revenue over agents $(1, \ldots, k)$). Similarly let $\bar{R}(v^{k+1}, \ldots, v^n)$ denote the revenue of the k-lookahead auction. Note that the order among the types of agents $(1, \ldots, k)$ is not fixed. We will show that $\bar{R}(v^{k+1}, \ldots, v^n) \geq (k/(k+1)) \cdot \bar{R}_{opt}(v^{k+1}, \ldots, v^n)$. Since, this holds for every $n-k$ lowest agents and valuations, the theorem follows. Consider the optimal auction A_{opt}. Let m_{k+1} denote the contribution of agents $(k+1, \ldots, n)$ to the optimal revenue $\bar{R}_{opt}(v^{k+1}, \ldots, v^n)$. Because of the IR requirement it must be that $m_{k+1} \leq v^{k+1}$. For $j \leq k$, let m_j denote the contribution of agent j to $\bar{R}_{opt}(v^{k+1}, \ldots, v^n)$. Clearly $\bar{R}_{opt}(v^{k+1}, \ldots, v^n) = \sum_j m_j$. If for all $j \leq k$, $m_{k+1} \leq m_j$, then we are done as $\bar{R}(v^{k+1}, \ldots, v^n) \geq m_1 + \ldots + m_k$. Otherwise, let $\hat{j}$ denote the agent with the minimal m_j. Consider the following auction: pretend that agent $\hat{j}$ declares v_{k+1} and run the optimal auction A_{opt}. In case where one of the $n-k$ lowest agents would have won A_{opt}, give the item to agent $\hat{j}$ for the price of v^{k+1}. Clearly, the contribution of agent $\hat{j}$ now is at least m_{k+1}. Because of the *independence* assumption, the distribution of the other agents remains the same. (I.e. the auction does not lose information because it ignores the declaration of agent $\hat{j}$.) Since we only reduced the type of agent $\hat{j}$, the contribution of the other agents can only increase. Thus this auction extracts a revenue of at least $\bar{R}_{opt}(v^{k+1}, \ldots, v^n) - m_{\hat{j}} = \sum_{j \neq \hat{j}} m_j \geq (k/(k+1)) \cdot \bar{R}_{opt}(v^{k+1}, \ldots, v^n)$.

□

This technique may also be used to show bounds for weakly dependent agents. Currently we do not know if it is possible to achieve approximation ratios better than 2 for general distributions. We leave this as an intriguing open problem.

Acknowledgments: We thank Moni Naor, Yoav Shoham and Bob Wilson for useful discussions. We thank Andrew Goldberg, Eiichiro Kazumori, Noam Nisan and Inbal Ronen for commenting on an earlier draft of this paper.

References

[1] Jacques Crémer and McLean Richard. Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interpendent. *Econometrica*, 53:345–61, 1985.

[2] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions for multiple digital goods. Technical report, Technical Report STAR-TR-00.05-02, InterTrust Technologies Corp, 2000.

[3] A. V. Goldberg, J. D. Hartline, and A. Wright. Competitive auctions and digital goods. In *12th Annual ACM-SIAM Symposium on Discrete Algorithms*, January 2001.

[4] O. Goldreich, S. Goldwasser, and S. Michali. How to construct pseudo random functions. In *The 25th Symp. on foundations of computer science*, pages 464–479, 1984.

[5] Oded Goldreich. The foundations of cryptography. http://www.wisdom.weizmann.ac.il/ oded/foc-book.html.

[6] Paul Klemperer. Auction theory: a guide to the literature. *Journal of economic surveys*, pages 227–286, 1999.

[7] Paul Klemperer. *The economic theory of auctions.* Edward Elgar Publishing, 2000.

[8] Daniel Lehmann, Liadan O'Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *ACM Conference on Electronic Commerce (EC-99)*, pages 96–102, November 1999.

[9] A. Mas-Collel, Whinston W., and J. Green. *Microeconomic Theory.* Oxford university press, 1995.

[10] R. B. Myrson. Optimal auction design. *Mathematics of operational research*, 6:58–73, 1981.

[11] Amir Ronen. Algorithms for rational agents. In *Proc. of the 27th Annual Conference on Current Trends in Theory and Practice of Informatics*, 2000.

[12] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.