Optimal Auctions are Hard

(extended abstract, draft)

Amir Ronen[*] Amin Saberi[†]

April 29, 2002

Abstract

We study a fundamental problem in micro economics called optimal auction design: 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, the goal is to construct an optimal auction, i.e. a truth revealing protocol that maximizes the seller's expected revenue.

We study this problem from a computational perspective and show several lower bounds. In particular we prove that no deterministic polynomial time auction can achieve an approximation ratio better than $\frac{3}{4}$. This result applies to many other problems, such as the design of combinatorial auctions. The probability distribution constructed in our example has sensitive dependencies among the agents. On the flip side, we show that if the agents are not strongly dependent, the problem can be approximated within a factor close to 1.

1 Introduction

The design of optimal selling or buying mechanisms (a.k.a. optimal auction design) plays a major role in both theoretical and practical economics. It is also a source of many interesting mathematical problems. A comprehensive survey of the literature in this field can be found in [9]. In this paper, we study a problem which is perhaps the most basic in optimal auction design:

A seller wishes to sell an item to a group of self-interested agents. Each agent i has a *privately known* valuation $1 \leq v_i \leq h$ for the object. If the

[*]Stanford CS and ICSI Berkeley. Email: amirr@icsi.berkeley.edu.

[†]College of Computing, Georgia Institute of Technology. This work was done while visiting U.C. Berkeley and ICSI Berkeley. Email: saberi@cc.gatech.edu.

agent 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. Each agent *selfishly* tries to maximize her own profit. An auction is a protocol that decides who wins the item (if at all) and for what price. Given a probability distribution D on the valuations of the agents, the goal is to design an auction that maximizes the *expected revenue* of the seller.

Since rational agents may manipulate the given protocol if it is beneficial for them, or may simply refuse to participate in it, the auction must satisfy two standard game theoretic requirements: *Incentive Compatibility (IC)* meaning that the dominant strategy of each agent is to declare her value truthfully, and *Individual Rationality (IR)* meaning that the profit of a truthful agent is always non-negative. We call such an auction *valid.*

Given the probability distribution D, the objective is to construct an auction that, among all valid auctions, yields the highest expected revenue. The distribution is provided through an oracle.

This problem can be reduced to a pure optimization problem. From a computational point of view, it is particularly interesting because its type of constraints are very different from constraints of "ordinary" optimization problems.

The above problem has been a subject of long and intensive research in micro economics. The major thread of research focuses on *characterizing* the optimal auction (e.g. [15, 19, 4]). The problem is solved [15] for the case where the agents' valuations are independent (and the distribution obeys some regularity conditions). Unfortunately, this is barely the case and only little progress [4, 19] has been made for general probability distributions.

In this paper we study this problem from a computational perspective. We focus on three types of auctions: pre-committed, ascending, and general. For each type we show a lower bound on the approximation ratio that can be obtained by any polynomial time auction. In particular we prove that no deterministic polynomial time auction can achieve an approximation ratio better than $\frac{3}{4}$. A $\frac{1}{2}$-approximation was given at [18] which leaves a small gap between the upper and the lower bounds. We note that this bound is explicit and is not based on any complexity assumptions. (This is possible due to the oracle setup.) The lower bound immediately implies that it is impossible to characterize the optimal auction for the general case.

It is often the case that the seller needs to commit to the protocol in advance, i.e. the auction can only query the oracle before the bids are submitted. We call such auctions *pre-committed.* Using an argument similar to [18], we prove that if pseudo-random functions exist, no polynomial time pre-committed auction can have a revenue more than $1/ln(h)$ of the optimal revenue, (h is the highest possible value). On the other hand, we show that a simple fixed price auction can achieve this bound.

These results together with the $\frac{1}{2}$-approximation of [18] have interesting consequences: Firstly, in contrast to the traditional approach of theoretical economics, in which the designer (seller) commits to the protocol, we show that in many situations, she has a lot of incentive *not* to declare the protocol in advance. Secondly, theoretical economists mostly look for solutions which can be described as closed formula. We show that not only can the optimal auction not be described this way, but any auction which can be described by a closed formula, cannot achieve an expected revenue better than a simple fixed price auction (which is far from optimum).

Another natural special case of auctions are ascending auctions. Such auctions maintain a minimum threshold for every agent. At each stage, the auction can either raise the threshold of one of the agents or declare one of them as the winner. In the first case, if the agent's declared type is still less than or equal to the threshold, she will remain in the auction. Otherwise, the agent drops out and she will never be specified as a winner. In the second case, the winning agent pays her threshold. We allow the auction to query the oracle at each stage. Using a probabilistic argument, we show that if the number of queries is polynomial, an adversary can fool the auction so that it recognizes the agents' valuations as independent although crucial information is hidden in the probability distribution. Thus, we can prove that no polynomial-time ascending auction can have an expected revenue better than $\frac{3}{4}$ of the optimal ascending auction.

Finally, we generalize the proof of the ascending case, showing that any deterministic polynomial auction can be at most a $\frac{3}{4}$-approximation for general auctions.

Our lower bounds immediately generalize to many economic problems and in particular to combinatorial auctions (see [21] for a survey).

The probability distributions that we use for our lower bounds are very sensitive in a sense that the agents are $(n-2)$-wise independent but $(n-1)$-wise dependent (n is the number of agents). In contrast to the results mentioned above, we show that if the dependency between the agents is bounded (for a precise definition see section 7), it is possible to approximate the problem within a factor close to 1.

Related Work: Optimal auction design is a long studied 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 [10].

The optimal auction problem is solved ([15] and others) for the case where the agents' valuations are *independent* (and the probability distributions obey some regularity conditions). Unfortunately, this is barely the case and little progress has been made for general probability distributions. Under much weaker requirements of IR and IC (and under additional assumptions), beautiful results of [15, 5] 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 [9]).

In computer science the problem studied here has been first investigated in [18]. In particular, that paper shows a simple $\frac{1}{2}$-approximation to this problem. Various optimal auction design problems were studied in recent years, mostly in the context of digital goods [7, 6, 2] and online auctions [11, 3]. These problems are essentially different from ours in two aspects. Firstly, they adopt a worst case approach which is not natural for our problem. Secondly, since the agents do not compete with each other on the object, the constraints imposed on the algorithm are significantly weaker.

2 A Game Theoretic Model

In this section we formally present our problem and notations. We try as much as possible to combine the standard notions of the optimal auction literature with basic computational concepts.

We consider the following scenario. 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 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 $V_i = \{1, 1+\Delta, 1+2\Delta, \ldots, 2, 2+\Delta, \ldots, h\}$ and let $V = \prod_i V_i$. 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 any 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. Of course, the agents may *lie* if it is beneficial for them to do so. We

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

denote the declaration of the agents by $w = (w_1, \ldots, w_n)$. An important observation known as the revelation principle [14, pp. 871] says that w.l.o.g. we can limit ourselves to revelation auctions. Formally:

Definition 1 (auction) *An* auction *is an algorithm that accepts a type declaration $w = (w_1, \ldots, w_n)$ and returns two numbers k and p where k denotes who wins the auction (zero value means that nobody won), and p denotes the winner's payment.*

We say that the auction is *polynomial* if its computational time is polynomial in the number of possible types $(h-1)/\Delta$ and in the number of agents n.

We say that an agent i is *truthful* if she declares her actual type, i.e. $w_i = v_i$. As the agents may manipulate the auction or refuse to participate in it, we impose the following standard game theoretic requirements on the set of auctions that we allow:

Definition 2 (valid auction) *We call an auction* 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 strategy for each agent. I.e. it is never beneficial for the agent to mis-report her type.*

Unless stated otherwise, from now on the word auction will refer to a valid auction.

Given the probability distribution ϕ on the type space, we define the revenue of an auction as the expected payment when the agents are truthful. Formally:

Definition 3 (revenue) *Let ϕ be a probability distribution on the type space of the agents, and let $m = (k, p)$ be a valid auction. The* revenue *of m is the expected payment $E_{v \in \phi}[p(v)]$.*

We say that an auction is a *c-approximation* if its revenue is at least c times the revenue of any valid auction.

A naïve description of the distribution requires exponential space. In this paper we assume that ϕ is given as an oracle that accepts queries of the form $Pr[a_1 \leq v_1 \leq b_1, \ldots, a_n \leq v_n \leq b_n]$. In particular, conditional distributions can be queried. We note that our lower bounds apply even to an oracle which can be queried about *any* event, i.e. $Pr[s \subseteq V]$. We allow the auction to *query* the oracle in time $O(1)$.

Note that there are many valid auctions. Let k be an allocation function 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 see that this characterizes the family of valid auctions [17, 13].

It is possible to formulate the optimal auction problem as an integer program [18]. Unfortunately, the number of variables of this program ($n \cdot ((h-1)/\Delta)^{(n+1)}$) is *exponential* in the number of agents. Thus, we can use this approach only when the number of agents is small.

3 Two Examples

This section presents two examples of optimal auctions. Both examples consist of two agents Agent 1 and Agent 2, a simple probability distribution and an optimal valid auction. Real life problems are of course much more complicated.

Independent valuations: Consider the following probability distribution: Each agents i's valuation v_i is uniformly distributed in $(1, 2, \ldots, 100)$. The valuations are independent.

Consider the following solution known as the second price auction [20]. Each agent declares her type. (Note that the agents can lie!). The object is given to the agent with the higher declaration for the price of the lower one. It is not difficult to 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 \$51: If both agents are below \$51, nobody wins; If one of them is above the reserved price, she pays \$51; Otherwise the higher agent wins paying the declaration of the lower agent.

The optimality of this auction is implied by a classic result in auction theory (see e.g. [15]).

Correlation: Consider the following probability distribution: v_1 is uniformly distributed in $(1, 2, \ldots, 100)$ and $v_2 = 2 \cdot v_1$.

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. Since $E[\overline{w}]$ is an upper bound on the revenue of any valid auction, this auction is clearly optimal.

4 Pre-committed Auctions

It is often the case that the designer must commit to the protocol before the bidding phase. Therefore, the auction can use the oracle only before the declarations of the agents are submitted. In this case we can exploit techniques of [18] to show a matching bound of $\ln(h)$ on the approximation ratio of any polynomial auction. The proofs use ideas similar to [18, section 5.2]. Thus, we only sketch them.

Definition 4 (pre-committed auction) *We say that an auction is* pre-commited *if its queries are independent of the declared types.*

Theorem 4.1 *There exists a* $\ln(h)$ *pre-committed auction.*

Proof: (sketch) Let D be the probability distribution and let $\bar{B}$ denote the expected maximum $E_{v \in D}[\max_i v_i]$. Clearly $\bar{B}$ is an upper bound on the revenue of any valid auction. Let r be a price that maximizes $k \cdot Pr[\max_i v_i \geq k]$. Consider an auction that fixes the price to r and chooses the winner arbitrarily between the agents with $v_i \geq r$ (if any). It is not difficult to see that this auction is valid and has a revenue of at least $\bar{B}/\ln(h)$. □

We now show that under a natural hardness assumption, this upper bound is tight. Recall that a function $F(X \to Y)$ is called pseudo random [8] (p.r.f.) if any (randomized, sampling) algorithm that sees a polynomial sequence of the form $(x, F(x))$ (x is sampled uniformly) will not be able to distinguish between it and a totally random sequence of pairs (x, y).

Theorem 4.2 *If pseudo random functions exist then any polynomial pre-committed auction is at most a* $\ln(h)$ *approximation.*

Proof: (sketch) Consider the following probability distribution d_1:

$$Pr[v_1 = k] = \begin{cases} \frac{1}{h} & k = h \\ \frac{1}{k} - \frac{1}{k+1} & k = 1, 2, \ldots, h-1. \end{cases}$$

Clearly, no valid auction can extract more than 1 from this distribution. On the other hand, the expected value of v_1 is around $\ln(h)$. (d_1 was introduced at [7] as a distribution which is bad for auctions.)

We let Agent 1 be distributed according to v_1. Now consider a set of $(n-1)$ "small" agents with types uniformly in $\{\epsilon, 2\cdot\epsilon\}$. (We can scale everything by $1/\epsilon$ to get types greater than 1.) Consider a p.r.f. F that map the type vector of the small agents to the type of Agent 1, such that any polynomial algorithm sees the distribution d_1.

Let D be a probability distribution in which Agent 1 is independent of agents $(2,\ldots,n)$, and let D' be a distribution in which Agent 1's type is determined according to $F(v_{-1})$. Clearly, any auction will yield a revenue of at most 1 for D (neglecting the epsilons).

Assume by contradiction that an auction A that yields a revenue better than $R_{D'} > 1 + O(1)$ on D' exists. Note that A cannot use the oracle. By the Chernoff bound, if we sample D' polynomially many times, we get w.h.p. an average revenue close to $R_{D'}$. But then we can distinguish between D and D' – a contradiction. Thus the revenue of any polynomial algorithm on D' is ~ 1.

On the other hand, an optimal exponential algorithm can offer a price of $F(v_{-1})$ to Agent 1, extracting a revenue of $\sim \ln(h)$. $\square$

These results together with the 0.5-approximation of [18] have interesting consequences: Firstly, in a contrast to the traditional approach of theoretical economics, in which the designer commits to the protocol, we show that in many situations, the designer may have a lot of incentive *not* to declare the protocol in advance. Secondly, theoretical economists almost always look for solutions which can be described as closed formula. We show that not only the optimal auction can not be described this way, but a closed formula can not achieve an expected revenue better that a simple fixed price auction which is very far from optimum.

5 Ascending Auctions

Perhaps the most natural and traditional way of constructing valid auctions is by ascending auctions. Formally, such an auction maintains a set of thresholds $(t_1,\ldots,t_n)$. At the beginning all thresholds are set to 1. At each stage, the auction can either raise the threshold of one of the agents by Δ or declare one of them as the winner. In the first case, if the agent's type v_i is still less than or equal to t_i, the agent remains in the auction. Otherwise, the agent drops out, so she cannot win anymore. In the second case, the winning agent i pays her threshold t_i. We allow the auction to *query* the oracle at each stage. Note that when an agent drops out, her type v_i is revealed to the auction (it is exactly the point in which the agent dropped out). Beforehand, the auction only knows that $v_i \geq t_i$.

The most common example of such an auction is the well known English auction. This auction raises all the thresholds at the same pace, until all but the highest agent have dropped out. Currently, around 95% of the auctions on the Internet are English.

Interestingly, not every valid auction is equivalent to an ascending one. (We omit a counter example due to lack of space.). Thus, we are both strengthening and weakening the auction relatively to pre-committed auctions.

In [18], a $\frac{1}{2}$-approximation ascending auction was given. This auction is similar to the English auction, except that after all but the highest agent have dropped out, the auction computes the optimal price to offer to this agent.

The following theorem shows that no polynomial ascending auction can do better than $\frac{3}{4}$. Even more so, there exists an ascending auction which is $\frac{4}{3}$ times better than this auction!

Theorem 5.1 *Any deterministic polynomial ascending auction is an at most $\frac{3}{4}$-approximation. Moreover, it is an at most $\frac{3}{4}$-approximation against ascending auctions.*

The proof is described in subsections 5.1 through 5.4.

5.1 High Level Description of the Proof

The idea of the proof is to use sensitive dependencies between the agents. We construct these dependencies such that the algorithm cannot distinguish between them and the independent distribution. We consider a probability distribution in which one of the agents (Agent 1) has high values, and many small agents which always have values around 1 ($1, 1+\epsilon, \ldots, 1+n^4 \cdot \epsilon$). We take ϵ to be small enough and ignore it in the calculation of the revenue. The main function of the small agents is to provide *information* about Agent 1.

We fool the algorithm as follows. We start with a probability distribution D in which the agents are independent. After some polynomial time the algorithm must first advance one of the agents. Note that it is always better for the algorithm to advance one of the small agents. W.l.o.g. let n be the agent that is moved. Consider the case where n is rejected (i.e. she has a value of 1).

Let $(B_1, \ldots, B_k)$ be the queries of the algorithm at this stage. We now change the distribution to D' such that: n is independent of all the others; Agent 1 is highly dependent on agents $(2, \ldots, n-1)$; The values of $(B_1, \ldots, B_k)$ remain the same.

We will show that w.h.p, the algorithm must either use the information of all the other agents (henceforth losing the chance of allocating to them), or lose important information about Agent 1. In both cases, the ratio between its revenue and the optimal revenue is at most $\frac{3}{4}$.

5.2 Proof Preliminaries

Given a probability distribution D, we denote the revenue of the algorithm by R_D and the optimal revenue by $\bar{R}_D$. Given probability distributions $(d_1, \ldots, d_n)$ of single agents, we denote by $[d_1, \ldots, d_n]$ the joint distribution where the agents are independent. We denote by 1, the constant probability distribution of 1.

Let h be arbitrary high. We define the following probability distribution d for Agent 1.

$$Pr[v_1 = k] = \begin{cases} 2/h & k = h \\ 2/h & k = h/2 \\ 1 - 4/h & k = 1 \end{cases}$$

We also define d_+ and d_- to be the probability distributions when we shift the mass up or down between h and $h/2$. I.e. d_+ is defined by $(1 - 4/h, 0, 4/h)$ and d_- is defined by $(1-4/h, 4/h, 0)$. Note that $d = \frac{1}{2} \cdot d_+ + \frac{1}{2} \cdot d_-$.

We now define the probability distribution d_i for the small agents. There are $n^4 + 1$ types of the form $(1, 1+\epsilon, \ldots, 1+n^4 \cdot \epsilon)$. Let p be a constant arbitrary close to 1. We put a mass of p on 1 and divide the rest equally. We define d_i by

$$Pr[v_1 = k] = \begin{cases} p & k = 1 \\ (1-p)/n^4 & k > 1 \end{cases}$$

Finally, we let $D = [d, d_2, \ldots, d_n]$. We say that a type z_i of a small agent i is *short* if $z_i \leq 1 + (n^4 - n^2) \cdot \epsilon$, and that $z = (z_2, \ldots, z_n)$ is short if all the z_is are short. Note that when we choose a type according to D, the type vector of the small agents is short w.h.p.

We now calculate the optimal revenue for the two agent distributions $[d, 1]$, $[d_+, 1]$, and $[d_-, 1]$. In these cases, the valid auctions are as follows:

A Offer Agent 1 a price of 4. If she rejects, give the object to Agent 2 for a price of 1.
B Offer Agent 1 a price of 2. If she rejects, give the object to Agent 2 for a price of 1.
C Offer Agent 1 a price of 4, and if she rejects do not allocate.
D Offer Agent 1 a price of 2, and if she rejects do not allocate.
E Give the item to one of the agents for a price of 1.

Let $R_0 = \bar{R}_{[d,1]}$. It is not difficult to see that the optimal auction here is **A**. Thus $R_0 = 2 + (1 - 2/h) \sim 3$.

Let $R_+ = \bar{R}_{[d_+,1]}$. The optimal auction here is **A** as well. Thus $R_+ \sim 5$. Any other auction cannot extract a revenue of more than 4.

Let $R_- = \bar{R}_{[d_-,1]}$. Similarly to the previous case we get that the optimal auction is **B**, and that $R_- \sim 3$. Any other auction cannot extract a revenue of more than 2.

As we ignore the epsilons, the same calculations apply for D, $[d_+, d_2, \ldots, d_n]$, and $[d_-, d_2, \ldots, d_n]$.

5.3 Adding Dependencies

Recall that in order to fool the algorithm, we start with the probability distribution D. After a polynomial time, the algorithm moves agent n. Let $(B_1, \ldots, B_k)$ denote the queries of the algorithm. We create a distribution D' such that: n is independent of all the others, Agent 1 is very dependent on agents $(2, \ldots, n-1)$ and the values of $(B_1, \ldots, B_k)$ remain the same. In this subsection we construct D'. For convenience we denote $n' = n - 1$ and $n'' = n - 2$.

The following combinatorial lemma says that we can make Agent 1 almost fully dependent on agents $(2, \ldots, n')$ and the algorithm will not notice. Moreover, the dependency is sensitive. That is, we have to know the type of *all* the agents in order to know the distribution of Agent 1.

Lemma 5.2 (labeling lemma) *Let H be the grid of the types of agents $(2, \ldots, n')$ with the uniform distribution. Let $(B_1, \ldots, B_k)$ be subsets of H (k is polynomial). There exists a labeling $L : H \to \{-1, 0, +1\}$ such that*

Indistinguishablity *For all j, $E_{x \in B_j}[L(x)] = 0$.*

Fooling *for almost every short type vector* z*, for all agents* i*,* $Pr_{x_i \geq z_i}[L((x_i, z_{-i})) = 1] \sim \frac{1}{2}$ *and* $Pr_{x_i \geq z_i}[L((x_i, z_{-i})) = -1] \sim \frac{1}{2}$*.*

Proof: Let $m = |H| = n''^{k \cdot n''}$. Note that $m >> n$. Recall that for a labeling L, the discrepancy of L is defined as $Disc(L) = \max_j |\sum_{x \in B_j} L(x)|$. Let $r = m^c$ where $\frac{1}{2} < c < 1$.

Let L be a random $\{+1, -1\}$ labeling of H. Standard counting arguments show that w.h.p. $Disc(L) < r$ (see e.g. [1, Chapter 12] for a similar argument). Thus, for every subset B_j, we can change up to r coordinates and get a discrepancy of zero.

Let z be a type vector for agents $(2, \ldots, n')$. We say that z is *not corrupted* if for all i and $x_i \geq z_i$ we did not change the value of $L((x_i, z_{-i}))$. The overall number of corrupted type vectors is clearly bounded by $r \cdot k \cdot n^4 \cdot n << m$. Thus, almost all the type vectors are not corrupted (and short). Consider a non corrupted short vector z. For each agent i, we have at least n^2 random $\{+1, -1\}$ entries. Thus, by applying the Chernoff bound and then the union bound we get that w.h.p. for every agent i, $Pr_{x_i \geq z_i}[L((x_i, z_{-i})) = 1] \sim \frac{1}{2}$ and $Pr_{x_i \geq z_i}[L((x_i, z_{-i})) = -1] \sim \frac{1}{2}$ as requested. □

W.l.o.g. we can assume that $B_1 = H$. Thus, with probability $\sim \frac{1}{2}$, $L(x) = 1$ and with probability $\sim \frac{1}{2}$, $L(x) = -1$. We are now in a position to construct D'.

Lemma 5.3 (fooling lemma) *Let* $(B_1, \ldots, B_k)$ *be the a polynomial set of queries. For sufficiently high* n*, there exists a probability distribution* D' *such that:*

Indistinguishablity $D'(B_j) = D(B_j)$ *for all* j*.*

Sensitivity *For almost every short type* y *of agents* $(2, \ldots, n')$*, for all* i*, and for every value* k *of Agent 1,* $Pr[x_1 = k]|[(z_i \geq y_i), y_{-i}] \sim d(k)$*. In other words, in order to know if Agent 1 is shifted, the auction must know all the types of agents* $(2, \ldots, n)$*.*

Fooling *For almost every short type* y *of agents* $(2, \ldots, n')$*, Agent 1 is either distributed by* d_+ *or by* d_-*. (Both probabilities are* $\sim \frac{1}{2}$*.)*

Independence *Agent* n *is independent of the others.*

Proof: We first comment that, since the the probability that a small agent has a type of 1 is a constant $p < 1$, all the high probability events in our labeling lemma 5.2 are also high probability with respect to D, and that w.h.p. o(n) small agents have types greater than 1.

Consider the set of queries $(B_1, \ldots, B_k)$. By projecting it on agents $(2, \ldots, n')$ we get at most k subsets for the labeling lemma.

Let $L()$ be a labeling constructed in lemma 5.2. For each type vector y of $(2, \ldots, n')$ we determine how Agent 1 is shifted according to $L(y)$. We get D' by adding agent n as an independent agent.

¿From the indistinguishablity condition of the labeling lemma, and the fact that d is the average of d_+ and d_- we have the indistinguishablity (we need a slight correction for types of 1 which we omit). The sensitivity is obvious from the sensitivity condition of the labeling lemma and the fact that d is the average of d_+ and d_-. The fooling condition follows immediately from the fooling condition of the labeling lemma.

□

5.4 Bounding the approximation ratio

Let D' be the probability distribution on the agents. We first bound the revenue $R_{D'}$ of our auction. With probability arbitrarily close to 1 the auction fails in its first step (i.e., agent n is rejected). This is because the agent's type is 1 with a probability p – a constant arbitrarily close to one. Consider the other small agents $(2, \ldots, n')$. W.h.p., their type vector y is short and the sensitivity and fooling properties of lemma 5.3 hold. In this case the algorithm can either reject all the small agents, or ignore the information of at least one small agent. In the first case the revenue is clearly bounded by $2, 3$, or, 4 depending on the direction in which the mass is shifted. Thus, on the average we have ~ 3. In the second case, the algorithm sees a distribution which is very close to d. Thus, its revenue is $\sim R_0$. Therefore we obtain that $R_{D'} \leq 3 + o(1)$.

The optimal auction will use **all** the information of agents $(2, \ldots, n')$. It will then conduct the optimal auction on Agent 1 and Agent n (according to the conditional distribution). Therefore, $\bar{R} \sim \frac{1}{2} \cdot R_+ + \frac{1}{2} \cdot R_- = 4$. Note that this auction can be implemented as an ascending auction.

Thus, we showed that the approximation ratio is bounded by $\frac{3}{4} - o(1)$.

This completes the proof of theorem 5.1. □

6 General Auctions

We now consider the most general type of valid auctions. Such an auction can be described as follows [16]: For each agent i the auction computes a price which is independent of her declaration $p_i(w_{-i})$. If there exists an agent i such that $p_i < w_i$, she must win. (So only one agent can be above her threshold.) To simplify the exposition we assume that no ties are possible. In the rest of the section we show that the previous lower bound of $\frac{3}{4}$ holds for this type of auctions as well. The proof uses the same ideas of the ascending case and we only sketch it.

Theorem 6.1 *No deterministic polynomial time auction can achieve an approximation ratio better than* $\frac{3}{4}$.

Proof: (sketch) Consider the probability distribution D as defined in subsection 5.2. Among the small agents, there exists at least one whose winning probability is at most $1/n$. W.l.o.g. let it be agent n. We now apply fooling lemma 5.3, letting n be the independent agent. Let D' denote the resulting distribution. We comment that D' can be constructed such that agent n still wins with a probability of at most $1/n$.

We now focus only on the cases where agent n is not winning. Let x be a type for the small agents. W.h.p. x is short, and O(n) agents have a type $x_i > 1$. Recall that the type of Agent 1 can be either 1, $h/2$, or h.

Suppose that the algorithm offers a threshold $p_1(x) > h/2$ to Agent 1. Consider the case where the type of agent 1 is $h/2$ and that a small agent $\hat{i}$ is the winner. Since $\hat{i}$ must keep winning in all the cases where her type is above $x_{\hat{i}}$, we have that the auction must offer a threshold of $p_1 > h/2$ in all these cases! As in subsection 5.4, this means that the expected revenue in all these cases is $\sim R_0 = 3$.

If whenever $p_1(x) > h/2$, no small agent is winning, or, if the auction never raises the threshold $p_1(x)$ above $h/2$, it cannot get a revenue of above 3 either.

Therefore, the revenue of the auction (in all the cases where n is not winning) is bounded by ~ 3.

In addition, no valid auction can have a revenue of more than 5 and the optimal revenue remains ~ 4. Thus we have: $R_{D'}/\bar{R}_{D'} \leq (((n-1)/n \cdot 3) + 1/n \cdot 5)/(4 - o(1)) \sim \frac{3}{4}$ □

7 Reasonable Distributions

In the previous section we showed a bound of $\frac{3}{4}$ on the approximation ratio of any polynomial auction. The probability distribution D' that we used was very sensitive. It was both $(n-2)$-wise independent and fully $(n-1)$-wise dependent. In reality, we do not expect such cases.

In this section, we show that "reasonable" distributions can be approximated by a factor close to 1.

Definition 5 *[18]* **(k-lookahead auction)** *The auction rejects all but the k highest agents. It then conducts the optimal auction on the high agents (according to the conditional probability distribution given the types of the others) .*

In [18] it was shown that when the agents are independent this auction extracts at least $k/(k+1)$ of the optimal revenue. We show that if the probability distribution is not too sensitive, a variant of this auction is close to optimal.

Let D and D' be two probability distributions on the agents' type space V. We say that D and D' are *β-distant* if their variation distance [12] is at least β. I.e. $\max_{s \subset V} |D(s) - D'(s)| \geq \beta$. Otherwise we say that the distributions are *β-close.*

Definition 6 (β-dependent) *Let D be a probability distribution on the agents. We say that agent i is β-*dependent *on agent j if there exists a type v_{-i} such that the conditional distributions $v_i|v_{-i}$ and $v_i|v_{-i,-j}$ are β-*distant.

In other words, there exists a type for all agents except i and j, in which agent i is highly dependent on agent j.

Definition 7 (reasonable distribution) *We say that a probability distribution D is (β, l)-*reasonable *if: Each agent i is β-dependent on at most l agents, and for each agent j, at most l agents are β-dependent on her.*

Theorem 7.1 *Let D be a (β, l)-reasonable distribution. For all $k > l$, the $(k \cdot l + 1)$-lookahead auction has a revenue of at least $\bar{R}_D \cdot (1 - (l+1)/k) - h \cdot \beta$.*

Proof: We start with a simple lemma.

Lemma 7.2 *Let R be an auction, and D and D' two β-close probability distributions on the agent types. Then, $|R_D - R_{D'}| \leq h \cdot \beta$.*

Proof: Let $p(x)$ denote the payment of the auction when the type is x. W.l.o.g. $R_D \geq R_{D'}$. Let E be the set of points x s.t. $D(x) \geq D'(x)$. Clearly, $R_D - R_{D'} \leq \sum_x p(x) \cdot (D(E) - D'(E)) \leq h \cdot \beta$ □

Let $r' = k \cdot l + 1$. Fix the types of the $n - r'$ lowest agents. W.l.o.g. suppose that agents $(1, \ldots, r')$ are the highest. Let $\tilde{R}$ denote the revenue of the optimal auction in this case. It is sufficient to show that there exists an auction on the high agents, which yields a revenue R of at least $\tilde{R} \cdot 1 - (l+1)/k - h \cdot \beta$.

Consider the optimal auction. Denote by $m_1, \ldots, m_{r'}$ the contributions (expected payments) of agents $(1, \ldots, r')$. Also denote by $m_{r'+1}$, the total expected payment of the rest.

If $m_j \geq m_{r+1}$ for each high agent, then $R \geq r/(r+1) \cdot \tilde{R}$.

Otherwise, consider an auction on the high agents that "replaces" the low agents by a high agent j, pretending that her type equals the highest among low agent types, and then simulates the optimal auction (according to the conditional distribution). As in [18], we lose the contribution m_j and the information of this agent; However we gain the contribution m_{r+1}. Let A denote the k agents with the highest contributions among agents $(1, \ldots, r')$. Since the distribution is (β, l)-reasonable, there exists at least one agent j among the r' agents such that all the agents of A are β-independent of her (each agent is dependent on at most l agents). The contribution of the agents which are dependent on j is at most $l/k \cdot \tilde{R}$ (because they are lower than the contributions of the agents in A). We also lose the contribution of agent j which is at most $1/k \cdot \tilde{R}$ and gain m_{r+1}. The distribution on the remaining agents is β-close to the original.

Thus, by lemma 7.2 we have that $\tilde{R} - R \leq (l+1)/k - h \cdot \beta$ as requested.

□

For most auction applications, it is possible to give a reasonable bound on the maximum valuation h, regardless of the number of agents. We comment that for $k = O(\log n)$, it is possible to compute the k-lookahead auction in polynomial time. It is possible to strengthen theorem 7.1 in several ways. In particular we can allow a small $(O(\log(n)))$ number $l_2(n)$ of influential agents (i.e. agents on which more than l agents are dependent). In this case we always include these agents in the optimal auction stage, obtaining a revenue of at least $\bar{R}_D \cdot (1 - l/k) - h \cdot \beta$ in polynomial time. We argue, that reasonable probability distributions will have β-values close to zero, for reasonable l and l_2.

8 Open problems

In this section we mention open problems which directly stem from this work.

Besides closing the gap between the upper and the lower bounds, perhaps the most intriguing open questions are to analyze distributions in which the agents are only positively correlated, and the approximation ratio of randomized algorithms. In both cases we can show, using similar techniques, that the optimal solution is not obtainable, but we currently do not even know whether a PTAS is possible.

Acknowledgments: We thank Christos Papadimitriou, Luca Trevisan, Vijay Vazirani, and Felix Wu for useful discussions. We thank Christos Papadimitriou, Inbal Ronen, and Vijay Vazirani for commenting on an earlier draft of this paper.

References

[1] Noga Alon and Spencer Joel. *The probabelistic Method.* John Wiley and Sons Inc., 2000.

[2] Ziv Bar-Yossef, Kris Hildrum, and Felix Wu. Incentive-compatible online auctions for digital goods. In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 964–970, 2002.

[3] Avrim Blum, Tuomas Sandholm, and Martin Zinkevich. Online algorithms for market clearing. In *Proc. of the 13th SIAM Symposium on Discrete Algorithms (SODA 2002)*, pages 971–980, 2002.

[4] Kim-Sau Chung and Jeffrey C. Ely. Ex-post incentive compatible mechanism design. Mimeo, Northwestern University, 2001.

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

[6] 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.

[7] 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.

[8] O. Goldreich, S. Goldwasser, and S. Michali. How to construct pseudo random functions. In *The 25th Symp. on the Foundations of Computer Science*, pages 464–479, 1984.

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

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

[11] Ron Lavi and Noam Nisan. Competitive analysis of online auctions. Proc. of the second ACM Conference on Electronic Commerce (EC00), 2000.

[12] Lucien Le Cam and Grace Lo Yang. *Asymptotics in statistics: Some Basic Concepts*. New York: Springer, 1990.

[13] 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.

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

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

[16] Noam Nisan and Amir Ronen. Algorithmic mechanism design (extended abstract). In *The Thirty First Annual ACM symposium on the Theory of Computing (STOC99)*, pages 129–140, May 1999.

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

[18] Amir Ronen. On approximating optimal auctions (extended abstract). In *The Third ACM Conference on Electronic Commerce (EC01)*, 2001.

[19] Ilya R. Segal. Optimal pricing mechanisms with unknown demand. Working paper, 2002.

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

[21] Rakesh Vohra and Sven de Vries. Combinatorial auctions: A survey. *INFORMS J. of Computing*, forthcoming.

A Not All Valid Auctions are Ascending

This is a simple example of an auction which doesn't have an ascending implementation:

Suppose that we have three agents and their types can be either $(X, 3, 4)$, $(5, X, 6)$, and $(7, 8, X)$ with probability 1/3 each, where X is a uniform random variable from [1..10]. The auction which allocates the good to the agent with the valuation X with the price of 0 is clearly valid. But it's not possible to implement this auction as an ascending auction because in the first step advancing the threshold of any agent might result in loosing the agent whom the good should be allocated to.