

Derandomization of Auctions*

Gagan Aggarwal[†]
Stanford University
Stanford, CA
gagan@cs.stanford.edu

Amos Fiat
Tel Aviv University
Tel Aviv, Israel
fiat@tau.ac.il

Andrew V. Goldberg
Microsoft Research
Mountain View, CA
goldberg@microsoft.com.

Jason D. Hartline
Microsoft Research
Mountain View, CA 94043.
hartline@microsoft.com

Nicole Immorlica
M.I.T.
Cambridge, MA.
nickle@mit.edu

Madhu Sudan
M.I.T.
Cambridge, MA.
madhu@mit.edu

ABSTRACT

We study the problem of designing seller-optimal auctions, i.e. auctions where the objective is to maximize revenue. Prior to this work, the only auctions known to be approximately optimal in the worst case employed randomization. Our main result is the existence of deterministic auctions that approximately match the performance guarantees of these randomized auctions. We give a fairly general derandomization technique for turning any randomized mechanism into an *asymmetric* deterministic one with approximately the same revenue. In doing so, we bypass the impossibility result for symmetric deterministic auctions and show that asymmetry is nearly as powerful as randomization for solving optimal mechanism design problems. Our general construction involves solving an exponential-sized flow problem and thus is not polynomial-time computable. To complete the picture, we give an explicit polynomial-time construction for derandomizing a specific auction with good worst-case revenue. Our results are based on toy problems that have a flavor similar to the hat problem from [3].

Categories and Subject Descriptors

G.2 [Mathematics of Computing]: Discrete Mathematics

General Terms

Algorithms, Economics, Theory

*This work was done when the authors were visiting Microsoft Research, Mountain View, CA.

[†]Supported in part by a Stanford Graduate Fellowship, NSF Grants EIA-0137761 and ITR-0331640, and a grant from SNRC.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

STOC'05, May 22-24, 2005, Baltimore, Maryland, USA.
Copyright 2005 ACM 1-58113-960-8/05/0005 ...\$5.00.

Keywords

Mechanism design, Auctions, Derandomization

1. INTRODUCTION

In [8], Goldberg et al. proposed the study of profit maximization in auctions using a worst-case competitive analysis. They focus on the *unlimited supply, unit demand, single-item* auction problem where an effectively infinite number of identical units of an item are for sale to consumers that each desire at most one unit. In their worst-case competitive framework, they gave a randomized auction that achieves a constant fraction of the optimal single-price revenue. Further, they prove that such randomization is necessary for *symmetric* auctions, ones whose outcome is not a function of the order of the input bids. Our main result is to show that this result does not hold for asymmetric auctions; we give an asymmetric deterministic auction that approximates the revenue of the optimal single-price sale in the worst case.

In general, design and analysis of auctions and other mechanisms requires a game-theoretic treatment; in order to understand the performance of an auction, the behavior of the bidders in the auction must be understood. To handle this problem, we adopt the solution concept of *truthful mechanism design*; we only consider mechanisms where each bidder has a dominant strategy of bidding their true value for the good regardless of the actions of any of the other bidders. It is well-known that truthful auctions are precisely those auctions that compute an offer price for a bidder that is not a function of their bid value, but may be a function of other bidders' bids. This bidder is then allocated the good if their bid is above the offer price, otherwise the bidder is rejected. The simplest truthful auction that is approximately optimal in worst-case is the *randomized sampling auction* in [8] which randomly partitions the bidders into two sets and uses the optimal sale price for each set as the offer price for all bidders in the opposite set.

Randomization in mechanism design, in a spirit similar to randomization in online algorithms, allows a mechanism to choose the right course of action with some positive probability. Standard algorithmic derandomization techniques do not directly apply to mechanism design because it is not possible to simply run the mechanism with all possible outcomes of a randomized decision making procedure; the decisions

made in mechanisms are generally irrevocable. Instead, we explore derandomization through asymmetry. Here we look for mechanisms that will make up for making wrong choices for some bidders by making right choices for others. A key challenge in this endeavor is the coordination of the choices.

Our main result is to show that on any input bid vector, \mathbf{b} , with bids $b_i \in [1, h]$, any randomized truthful auction \mathcal{A} that obtains an expected profit of $\mathbf{E}[\mathcal{A}]$ can be converted into a truthful deterministic asymmetric auction with profit $\mathbf{E}[\mathcal{A}] / 4 - 2h$. Given any auction that always obtains an expected profit that is within a constant fraction of optimal, such as the randomized sampling auction, this gives a deterministic auction that gets within a constant fraction of the optimal profit with a small additive loss.

This general derandomization technique involves solving an exponentially large flow problem. We introduce a different technique to obtain a polynomial-time constructible derandomization of an auction from [6] with good worst-case revenue guarantees. This resolves the question about the existence of good deterministic worst-case auctions that can be run efficiently.

The problem of asymmetric coordination of decisions in mechanism design is exemplified by a related toy problem that is similar in nature to hat-guessing problem of [3]. We cast this problem as follows. Players interact in a three-stage game. First, the players collectively decide on their strategies. Second, colored hats are placed on each player's head such that each player can see the color of the hats of all other players except for their own. Third, each player must independently try to guess the color of their own hat. The question is what strategy should the players adopt in order to ensure that many of the players correctly guess their hats' color. Assume that the available hat colors are blue or red. There is a natural randomized strategy where each player flips a coin and guesses red or blue with equal probability. With this strategy, in expectation, half of the people with red hats and half of those with blue hats guess their hat color correctly. We would like to devise a deterministic strategy that achieves approximately the same bound.

We solve this problem via a flow-based derandomization technique. When applied to the simple randomized color-guessing algorithm above, it gives an asymmetric deterministic (exponential-time) algorithm that guesses red for exactly half of the reds and guesses blue for exactly half of the blues (rounded down if the numbers are odd). This flow-based derandomization extends naturally to the case in which there are k different hat colors and we would like to guess correctly on about $1/k$ fraction of each color. For the $k = 2$ case, we show that there is actually a polynomial-time computable derandomization that matches the performance of the flow-based algorithm; however, we do not know if the same is possible for $k \geq 3$.

This coloring problem is related to the auction problem as follows. Consider the case where there are only two types of bidders, those with a high valuation for the item, h ; and those with a low valuation for the item, 1. Mapping h to the color red and 1 to the color blue, a solution to the color-guessing problem would offer half the h bids a price of h and half the 1 bids a price of 1 and thus, the profit of such an auction would be at least half of optimal revenue (because either a price of h or a price of 1 is the optimal single price).

Problems related to the above hat problem were first studied in the context of coding theory [3, 13]. That these types

of problems are relevant to mechanism design makes explicit a connection between the two fields. A similar hat coloring problem to the one we consider was recently proposed by Feige [4] who independently derived a similar flow-based construction. We are unaware of any previous applications of the hat problem to truthful mechanism design.

This paper is organized as follows. In Section 3 we formally define the hat-guessing problem and give the flow-based solution. In Section 4 we show how the flow-based technique can be generalized to convert any randomized auction into an asymmetric deterministic auction with approximately the same performance bound. Finally, in Section 5 we give a polynomial-time deterministic auction with good worst-case revenue bounds.

2. PRELIMINARIES

We consider single-round, sealed-bid auctions for selling an item in unlimited supply to any of n unit-demand bidders. As mentioned in the introduction, we adopt the game-theoretic solution concept of truthful mechanism design. A useful simplification of the problem of designing truthful auctions is obtained through the following algorithmic characterization. Related formulations to the one given here have appeared in numerous places in recent literature (e.g., [1, 12, 5, 9]). To the best of our knowledge, the earliest dates back to the 1970s [10].

DEFINITION 1. *Given a bid vector of n bids, $\mathbf{b} = (b_1, \dots, b_n)$, let \mathbf{b}_{-i} denote the vector of bids with b_i replaced with a '?', i.e.,*

$$\mathbf{b}_{-i} = (b_1, \dots, b_{i-1}, ?, b_{i+1}, \dots, b_n).$$

DEFINITION 2 (BID-INDEPENDENT AUCTION, BI_f). *Let f be a function from masked bid vectors (with a '?') to prices (non-negative real numbers). The deterministic bid-independent auction defined by f , BI_f , works as follows. For each bidder i :*

1. Set $t_i = f(\mathbf{b}_{-i})$.
2. If $t_i < b_i$, bidder i wins at price t_i .
3. If $t_i > b_i$, bidder i loses.
4. Otherwise, ($t_i = b_i$) the auction can either accept the bid at price t_i (in which case bidder i is a winner) or reject it.

A randomized bid-independent auction is a distribution over deterministic bid-independent auctions.

The proof of the following theorem can be found, for example, in [5].

THEOREM 1. *An auction is truthful if and only if it is equivalent to a bid-independent auction.*

Given this equivalence, we will use the terms *bid-independent* and *truthful* interchangeably. We denote the profit of a truthful auction \mathcal{A} on input \mathbf{b} as $\mathcal{A}(\mathbf{b})$. This profit is given by the sum of the prices charged to the winning bidders. For a randomized bid-independent auction, $f(\mathbf{b}_{-i})$ and $\mathcal{A}(\mathbf{b})$ are random variables.

It is natural to consider a worst-case competitive analysis of truthful auctions. In the competitive framework of [5] and

subsequent papers, the performance of a truthful auction is gauged in comparison to the profit of the *optimal single price sale of at least two units*, denoted by $\mathcal{F}^{(2)}$ in previous literature. There are a number of reasons to choose this metric for comparison; interested readers should see [5] or [7] for a more detailed discussion.

Unfortunately, as we show in Section 4.3, it is not possible for a deterministic auction to always perform well in comparison to such an optimal sale. Instead, we assume that all bids are between 1 and h and define OPT as the profit of the *optimal single price sale*. Following [8, 2] we look for auctions that obtain a profit of at least $\text{OPT}/\beta - \gamma h$ for small constants β and γ . We refer to β as the *approximation ratio* and γh as the *additive loss*. Such an approximation framework is tantamount to considering a *promise* problem. If we are promised that $\text{OPT} \gg \gamma h$ then our auction is constant fraction of optimal. This motivates the following formal definition.

DEFINITION 3. *We say an auction is approximately optimal if its expected profit on any input, $\mathbf{b} \in [1, h]^n$, is at least $\text{OPT}(\mathbf{b})/\beta - \gamma h$ for fixed constants β and γ .*

3. A HAT PROBLEM

In this section we consider the problem of asymmetric coordination through a hat problem. Here n players must devise strategies such that when each of them has a hat of one of k colors placed on her head and can only observe the hats on others' heads but not her own; about a $1/k$ fraction of the players wearing each color correctly guess their hat color. We refer to this game as a "bid-independent" hat-guessing problem as player i 's viewpoint, seeing an n -dimensional vector of hat colors with their own missing, is similar to the view point a truthful mechanism has when considering a price to offer the i th bidder.

The following randomized hat-guessing scheme achieves this desired bound in expectation: *each player i guesses each of the colors with equal probability $1/k$* . We now give a technique that uses an arbitrary ordering of the players in place of randomness to achieve the same bound deterministically. It is instructive to view this technique as a derandomization of the simple randomized algorithm proposed above.

Let $\mathbf{c} = (c_1, \dots, c_n)$ represent the array of colors. Let \mathbf{c}_{-i} represent the array of colors with the i 'th color hidden, i.e., $\mathbf{c}_{-i} = (c_1, \dots, c_{i-1}, ?, c_{i+1}, \dots, c_n)$. Note that \mathbf{c}_{-i} is precisely the view of player i .

Consider a flow problem as shown in Figure 3. The graph has a source vertex s and a sink vertex t . For each of the nk^{n-1} possible values of \mathbf{c}_{-i} , we have a vertex, $v_{\mathbf{c}_{-i}}$. We place an arc from s to each of these vertices. For each of the k^{n+1} possible values of (χ, \mathbf{c}) (where χ is one of the k colors), we have a vertex $v_{\chi, \mathbf{c}}$. We place an arc between each of these vertices and t . We also add an arc between $v_{\mathbf{c}_{-i}}$ and $v_{\chi, \mathbf{c}}$ signifying that we get \mathbf{c} when we reveal that at position i in \mathbf{c}_{-i} is a hat with color c_i . Notice that the in-degree due to such arcs of a vertex $v_{\chi, \mathbf{c}}$ is precisely the number of hats of color χ in \mathbf{c} . The out-degree of a vertex $v_{\mathbf{c}_{-i}}$ is exactly k , one for each possible color of the hat at position i .

Now, imagine the following flow on this graph that represents the randomized hat-guessing algorithm above. Between s and each $v_{\mathbf{c}_{-i}}$ place a flow of 1. This corresponds to the randomized algorithm, upon seeing \mathbf{c}_{-i} , having a to-

tal probability of 1 to spend on guessing colors for the i th player's hat. On each of the outgoing arcs from $v_{\mathbf{c}_{-i}}$ we place a flow of $1/k$ corresponding to the probability with which the randomized algorithm picks each color. Now notice that the incoming flow to $v_{\chi, \mathbf{c}}$ is precisely $1/k$ times the number of hats in \mathbf{c} from color class χ . Send all of this flow on the arc from $v_{\chi, \mathbf{c}}$ to t . This is a feasible flow.

Next, we set capacities on the arcs as follows. For each χ and \mathbf{c} , set the capacity of the arc $(v_{\chi, \mathbf{c}}, t)$ to $\lfloor n_\chi(\mathbf{c})/k \rfloor$, where $n_\chi(\mathbf{c})$ represents the number of hats in \mathbf{c} that are colored χ . On all other arcs, place a capacity of 1. The fractional flow set up in the previous paragraph respects these capacities except on the arcs into t . Thus, the minimum cut in this graph separates t from all other vertices, and a maximum flow saturates the capacities of all arcs into t . Furthermore, since all capacities are integral, there is a maximum flow that is integral. Revisiting the analogy between flow and probability, since each of the $v_{\mathbf{c}_{-i}}$ has at most 1 unit of incoming flow, an integral flow places the entirety of this flow on a single outgoing arc corresponding to deterministically guessing a color for the i 'th hat in \mathbf{c}_{-i} . Thus, such a flow specifies a deterministic "bid-independent" hat-guessing algorithm.

We now analyze the performance of this deterministic hat-guessing algorithm on \mathbf{c} . Given $n_\chi(\mathbf{c})$ hats with color χ in \mathbf{c} , the capacity of the outgoing arc from $v_{\chi, \mathbf{c}}$ to t is $\lfloor n_\chi(\mathbf{c})/k \rfloor$. Since this arc is saturated in an integral maximum flow, it must be that $\lfloor n_\chi(\mathbf{c})/k \rfloor$ of the $n_\chi(\mathbf{c})$ incoming arcs have one unit of flow on them. This corresponds to the deterministic algorithm correctly guessing χ when considering \mathbf{c}_{-i} for $\lfloor n_\chi(\mathbf{c})/k \rfloor$ positions i in \mathbf{c} colored χ out of a total of $n_\chi(\mathbf{c})$ such positions. This holds true for all colors χ ; thus, this deterministic hat-guessing algorithm correctly guesses about a $1/k$ fraction from each color class.

For the case that $k = 2$, there is a polynomial-time construction of a coordination strategy for n players that meets the same guarantees as the above flow-based technique. We omit the details here as it follows as a special case from the *deterministic coin flipping* algorithm that we give in Section 5.1. An obvious next step is to obtain a polynomial-time construction for $k \geq 3$ colors. This problem seems much more difficult even for $k = 3$, and we leave it as an open question.

4. AUCTION DERANDOMIZATION

The main goal of this paper is to design deterministic auctions that are approximately optimal. We show that in fact *any* randomized auction has a deterministic counterpart that achieves approximately the same profit. As a corollary of this result, known approximately-optimal randomized auctions imply the existence of approximately-optimal deterministic auctions. Our proof first reduces any auction to a special type of auction that we define, called a *guessing auction*, and then uses a flow-based construction similar to that in Section 3 to derandomize the guessing auction.

4.1 Guessing Auctions

The flow-based construction for the hat-guessing problem in Section 3 works for the case where our goal is to consider \mathbf{c}_{-i} and guess what \mathbf{c} is. An auction gets revenue not only when it guesses bid values correctly, but also when it guesses a value below a bid value. In order to resolve this discrepancy, we define the notion of a *guessing auction* that




Figure 1: Flow for randomized hat-guessing. The label on an edge represents the amount of flow on the edge.

uses only powers of two as prices and gets credit for revenue from a bidder only when it offers her a price equal to her bid rounded down to the nearest power of two. Not only are guessing auctions approximately as powerful as standard auctions (in terms of approximating the optimal profit), but it is possible to convert any auction into a guessing auction while only losing a factor of four from the profit.

DEFINITION 4 (\mathcal{G}_A). *The guessing auction, \mathcal{G}_A , for an auction \mathcal{A} simulates \mathcal{A} on \mathbf{b} . Suppose \mathcal{A} offers bidder i price p_i and let 2^k be the largest power of two less than p_i . Then \mathcal{G}_A offers bidder i price 2^{k+j} for integer $j \geq 0$ with probability 2^{-j-1} .*

LEMMA 1. *For any auction \mathcal{A} with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector \mathbf{b} , there is a corresponding guessing auction \mathcal{G}_A whose expected profit on any input bid vector \mathbf{b} is at least $\mathbf{E}[\mathcal{A}(\mathbf{b})]/4$.*

PROOF. To see that the guessing auction achieves a profit of $\mathbf{E}[\mathcal{A}]/4$, we show that the expected profit of \mathcal{G}_A from bidder i , given that bidder i bids above p_i , is at least $2^{k-1} \geq p_i/4$. Suppose $b_i \in [2^{k+j}, 2^{k+j+1})$, then the probability that \mathcal{G}_A guesses bid i is 2^{-j-1} . The payoff on correctly guessing is 2^{k+j} . Thus, the expected payment of bidder i is 2^{k-1} . Since each bidder's expected payment in \mathcal{G}_A is a fourth of their payment in \mathcal{A} , we have the desired bound. \square

We note that if we are constructing a guessing auction from an auction that only uses prices that are powers of two, then we only lose a factor of two of the profit instead a factor of four.

4.2 The Flow Construction

We now show how to derandomize any guessing auction \mathcal{G}_A .

LEMMA 2. *Corresponding to any guessing auction \mathcal{G}_A with expected profit $\mathbf{E}[\mathcal{G}_A(\mathbf{b})]$ on bid vector \mathbf{b} , there is a deterministic auction whose profit on any input bid vector \mathbf{b} is at least $\mathbf{E}[\mathcal{G}_A(\mathbf{b})] - 2h$, where h is the highest bid value in \mathbf{b} .*

PROOF. First, round all bid values down to the nearest power of two. Making an analogy between the k colors in the hat-guessing problem and the $\log h$ powers of two that are the possible (rounded) values of bids, we proceed by setting up a flow construction identical to that for the k -color guessing problem, except that the fractional flow on an arc from $v_{\mathbf{b}_{-i}}$ to $v_{2^j, \mathbf{b}}$ is the probability that \mathcal{G}_A on seeing \mathbf{b}_{-i} guesses 2^j . Furthermore, the flow from $v_{2^j, \mathbf{b}}$ to t is the expected number of times \mathcal{G}_A guesses one of the bids with (rounded) value 2^j correctly. We represent this quantity by $E_j(\mathbf{b})$. We then set the capacities as before such that the capacity on the arc between $v_{2^j, \mathbf{b}}$ and t is $\lfloor E_j(\mathbf{b}) \rfloor$; all other capacities are set to one.

Once again, the above fractional flow implies the existence of an integer-valued flow, and this integer-valued flow corresponds to an auction making a deterministic bid-independent offer upon seeing \mathbf{b}_{-i} . The flow out of $v_{2^j, \mathbf{b}}$ is precisely the number of indices i such that the auction, upon seeing \mathbf{b}_{-i} , correctly guesses 2^j ; since this arc is in a minimum cut, it is saturated and the flow out of it is precisely $\lfloor E_j(\mathbf{b}) \rfloor$. Thus, considering a bid vector \mathbf{b} where the expected profit of \mathcal{G}_A is $\mathbf{E}[\mathcal{G}_A] = \sum_j 2^j E_j(\mathbf{b})$, the deterministic auction obtains $\sum_j 2^j \lfloor E_j(\mathbf{b}) \rfloor \geq \sum_j [2^j E_j(\mathbf{b}) - 2^j] \geq \mathbf{E}[\mathcal{G}_A] - 2h$. \square

The following theorem follows directly from Lemmas 1 and 2.

THEOREM 2. *Corresponding to any single-round sealed-bid auction \mathcal{A} with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{b})]$ on input bid vector \mathbf{b} , there is a deterministic auction \mathcal{A}' whose expected profit on any input bid vector \mathbf{b} is at least $\mathbf{E}[\mathcal{A}(\mathbf{b})]/4 - 2h$.*

As a corollary, using this derandomization result with the approximately optimal auctions in [8, 5, 6], we obtain deterministic auctions that are approximately optimal. In this construction, we assumed that the range of bid values $[1, h]$ is known. This assumption is not necessary. When considering \mathbf{b}_{-i} , we can compute h , which is the maximum bid value scaled such that the minimum bid value is 1 on the new scale, correctly for all but the minimum and maximum

bid value. Assuming the worst, i.e., the auction fails to get any profit from the highest and lowest bid, we only lose an additional additive $h + 1$.

4.3 Additive Loss Term

One discrepancy between the bounds given in this paper and the bounds given in [5, 6] is that our bounds are interesting only if the profit from the optimal single price sale is larger than $2h$. This is not true of all bid vectors, i.e., those with h larger than the profit from the *optimal single-price sale that sells at least two items* (this quantity is denoted $\mathcal{F}^{(2)}$ in [5]). For this case, the bounds obtained in [5] are better because they prove the auction's performance to be a constant fraction of $\mathcal{F}^{(2)}$ without any additive loss term. We can view these two types of analysis as the difference between solving a worst-case problem and a promise problem. Given the promise that the optimal single price sale achieves a large profit in comparison to h , our auction gets a constant fraction of optimal; otherwise, it may not.

In this section, we show that such a promise is necessary for obtaining a deterministic auction that performs well in the worst case. In particular, we show that there is no deterministic auction that obtains a profit that is a constant fraction of $\mathcal{F}^{(2)}$ on all inputs.

LEMMA 3. *No deterministic truthful auction obtains a constant fraction of $\mathcal{F}^{(2)}$ on all bid vectors.*

PROOF. We will show this by contradiction. Assume that we are given a deterministic auction BI_f , specified by bid-independent function f , that obtains a profit of $\mathcal{F}^{(2)}/\beta$ on all inputs.

Let $\mathbf{b} = (1, \dots, 1)$ be the all-ones bid vector. Assume, without loss of generality, that $f(\mathbf{b}_{-1})$, the price offered the first bidder, is 1. Now, for any $\alpha > \beta$ and $i \in I = \{1, 2, \dots, \lfloor \frac{n+1}{n/\alpha-1} \rfloor\}$, consider $\mathbf{b}^{(i)}$ as the all-ones bid vector except for $b_1 = n\alpha^i$. Let S_i be the set of other bidders (not including bidder 1) that are offered price 1 when the input to the auction is $\mathbf{b}^{(i)}$, i.e., $S_i = \{j > 1 : f(\mathbf{b}_{-j}^{(i)}) = 1\}$.

Fact 1: $|S_i| \geq n/\alpha - 1$.

This follows directly from the fact that otherwise BI_f 's profit would be at most $\mathcal{F}^{(2)}/\alpha$ which would contradict our assumption.

Fact 2: $\bigcap_{i \in I} S_i \neq \emptyset$.

Assuming the contrary, let the intersection of the S_i s be empty. Then the union of the S_i s is of size

$$\sum_i |S_i| \geq |I|(n/\alpha - 1) \geq n + 1 > n.$$

However, we know that all the S_i s are subsets of $\{1, 2, \dots, n\}$. This implies that $|\bigcup_i S_i| \leq n$, leading to a contradiction.

From Fact 2, there exists i, j , and k with $i < j$ and $k \in S_i \cap S_j$. Pick some $h \gg n^n$ (a number bigger than any of the $n\alpha^i$'s) and let $p_1 = f(\mathbf{b}'_{-1})$, where \mathbf{b}' is the all-ones input except for $b'_k = h$.

Case 1: $p_1 \leq n\alpha^i$. Then on the input that is all ones except for $b_1 = n\alpha^j$ and $b_k = h$, $\mathcal{F}^{(2)} = 2n\alpha^j$, but auction profit is at most $n\alpha^i + n < 2n\alpha^j/\beta$.

Case 2: $p_1 > n\alpha^i$. Then on the input that is all ones except for $b_1 = n\alpha^i$ and $b_k = h$, $\mathcal{F}^{(2)} = 2n\alpha^i$, but the auction profit is at most n and is therefore not $\mathcal{F}^{(2)}/\beta$. \square

4.4 Limited Supply

While we presented these results in terms of the unlimited supply auction problem, they also apply to the limited supply auction. Note that the number of items sold by the derandomized auction is no more than the expected number of items sold by the randomized auction. Thus, if the randomized auction never oversells, neither does its derandomized equivalent.

5. A POLYNOMIAL-TIME DETERMINISTIC AUCTION

In this section, we describe a competitive deterministic asymmetric auction, the outcome of which can be easily computed. There are three key ingredients in this auction, (a) a truthful profit extractor, (b) a pair of consensus estimate functions, and (c) a deterministic coin-flipping algorithm. The first two of these were used previously by Goldberg and Hartline along with a random coin flip to get an approximately optimal auction [6]. The main result of this section is to show how to derandomize this coin flip to obtain a deterministic auction with roughly the same performance guarantees. First we describe and solve another hat problem which is related to the problem of derandomizing a coin flip, then we review profit extraction and consensus estimates, and finally we combine these techniques to give the first polynomial-time computable deterministic auction that is approximately optimal.

5.1 A Deterministic Coin Flip

Consider the following continuous hat problem where the each of the hats is colored a shade of red. We would like each of the players to simulate a coin flip with the collective property that, for any particular shade of red, at least half the players with darker hats choose heads and at least half choose tails (rounding down). We assume, without loss of generality, that the hats are all distinct shades; if not, then we can use an arbitrary ordering to break ties (for example by choosing unique identifiers for each player).

Note that we can reduce the 2-color hat-guessing problem to the problem of deterministic coin flipping as follows. Run the algorithm with the two colors – light red and dark red. Interpret a *heads* coin as “light red” and a *tails* coin as “dark red”. The resulting algorithm, modulo rounding, guesses half of the light reds and half of the dark red hats correctly.

The algorithm we are about to propose solves this deterministic coin flip problem. In fact, our solution satisfies the following stronger property: the coins the players determine are actually perfectly alternating with the shade of the hat color. First, some definitions.

DEFINITION 5. *Given a vector of n hat shades, \mathbf{c} , the sign of \mathbf{c} (shorthand for “the sign of the permutation of the ordering of hats”) is the parity of the number of transpositions of adjacent hats it takes to sort \mathbf{c} , notated $\text{sgn}(\mathbf{c})$.¹*

DEFINITION 6. *Given a vector of n hat shades, \mathbf{c} , the rank of i , denoted $\text{rank}(\mathbf{c}, i)$, is the number hats in \mathbf{c} that are darker than c_i .*

We now propose the following *deterministic coin flip algorithm*, ϕ : Given \mathbf{c}_{-i} as the shades of the hats that player i

¹While the number of transpositions performed in sorting \mathbf{c} is not unique, the parity of the number of transpositions is.

sees, player i computes her coin, $\phi(\mathbf{c}_{-i})$, by imagining that her own hat is the darkest shade, ∞ , and computing the sign of this imagined vector of hat colors, $(\mathbf{c}_{-i}, \infty)$, as her coin flip.

LEMMA 4. *The deterministic coin flip algorithm, ϕ , is perfectly alternating with the shades of the hats' colors.*

PROOF. This result is implied by the fact that

$$\phi(\mathbf{c}_{-i}) \equiv \text{sgn}(\mathbf{c}) + \text{rank}(\mathbf{c}, i) \pmod{2},$$

which is evident because one way to sort $(\mathbf{c}_{-i}, \infty)$ would be to first sort \mathbf{c} and then replace hat i with ∞ which would require $\text{rank}(\mathbf{c}, i)$ additional transpositions to move the ∞ to the front of the array. \square

In this solution to the deterministic coin-flipping problem, each player can compute their own coin by simply executing ϕ ; however, no player can compute the coin of any other player. Clearly, each player can compute her coin in $O(n \log n)$ time; furthermore, as is evident from the above proof, the coins of all the players can be computed in total of $O(n \log n)$ time.

5.2 Profit Extraction

We briefly review the truthful profit extraction mechanism. This mechanism is a special case of a general cost-sharing schema due to Moulin and Shenker [11].

The goal of profit extraction is, given bids \mathbf{b} , to extract a target value R of profit from some subset of the bidders.

ProfitExtract $_R$: Given bids \mathbf{b} , find the largest k such that the highest k bidders can equally share the cost R . Charge each of these bidders R/k . If no subset of bidders can cover the cost, the mechanism has no winners.

Important properties of this auction are as follows:

- ProfitExtract $_R$ is truthful.
- If $R \leq \text{OPT}(\mathbf{b})$, ProfitExtract $_R(\mathbf{b}) = R$; otherwise it has no winners and no revenue.

Since ProfitExtract $_R$ is truthful, we let pe_R be its bid-independent function.

5.3 Consensus Estimators

A pair of consensus estimators is a pair of functions, r_{even} and r_{odd} , having the following properties:

1. (consensus) For any V , either r_{even} or r_{odd} is a consensus. A function r is a consensus for V if for all $v \in [V/2, V]$, $r(v) = r(V)$.
2. (estimate) For any V and $r \in \{r_{\text{even}}, r_{\text{odd}}\}$ that is a consensus on V , $r(V) \in [V/2, V]$.

It is easy to see that the following functions form such a pair of consensus estimators [6].

$r_{\text{even}}(v) = 2v$ rounded down to the nearest even power of two.

$r_{\text{odd}}(v) = 2v$ rounded down to the nearest odd power of two.

We will apply these consensus estimators to the values taken by $\text{OPT}(\mathbf{b}_{-i})$, which denotes the optimal single price profit from all the bidders except for bidder i , in order to obtain a consensus on an approximate value for OPT .

5.4 An Efficient Auction

In this section, we describe an auction, DCORE, that is a derandomization of a variant of the consensus revenue estimate (CORE) auction of [6]. This auction is built from the three components discussed above – the deterministic coin flip, the profit extractor, and our pair of consensus estimate functions.

DEFINITION 7. *DCORE is the bid-independent auction implemented by the following function, f :*

$$f(\mathbf{b}_{-i}) = \text{pe}_{R_i}(\mathbf{b}_{-i}),$$

with $R_i = r_{\phi(\mathbf{b}_{-i})}(\text{OPT}(\mathbf{b}_{-i}))$.

DCORE is bid-independent and therefore truthful. We now show that DCORE is approximately optimal.

THEOREM 3. *The profit of DCORE is at least $\text{OPT}/4 - h$.*

PROOF. If the optimal single price sale has exactly one winner, then the optimal revenue is h and approximating it within an additive h is trivial. Otherwise, let $\text{OPT} = \text{OPT}(\mathbf{b})$ be the revenue from the optimal single price sale. Then, for every i , we have $\text{OPT}/2 \leq \text{OPT}(\mathbf{b}_{-i}) \leq \text{OPT}$. Since r_{odd} and r_{even} are a pair of consensus estimates, one of them is a consensus on OPT . Suppose, without loss of generality, that it is r_{even} . Now consider the following thought experiment. Suppose we had set $R_i = r_{\text{even}}(\text{OPT}(\mathbf{b}_{-i}))$ for all i . Then the profit of our auction would be $r_{\text{even}}(\text{OPT})$ because r_{even} is a consensus at $r_{\text{even}}(\text{OPT}) \leq \text{OPT}$ and therefore the profit extraction technique will extract it. Let the price charged to the k winning bidders in the profit extraction technique be p , generating a total profit of $pk \in [\text{OPT}/2, \text{OPT}]$. In reality, by the deterministic coin-flipping procedure, at least $k/2 - 1$ of these k bidders had $\phi(\mathbf{b}_{-i}) = \text{even}$ and thus these bidders all pay p , exactly as they would have in the thought experiment. The total profit thus accounted for is $pk/2 - p \geq \text{OPT}/4 - h$, which proves the theorem. \square

6. CONCLUSIONS

We have shown the existence of deterministic auctions that are approximately optimal in the worst case. By necessity, these auctions are asymmetric. This gives an affirmative answer to the question left open in [8]. The construction we developed uses a guessing auction as an intermediate and suffers a performance loss because of it. It would be nice to show a more direct derandomization; this is made difficult due to the fact that the natural analogy to the flow problem that takes into account the fact that the auction obtains profit from all bidders above the offered price, is not a totally unimodular linear program and thus is not guaranteed to have an integral optimal solution.

The flow-based technique allows derandomization of any auction. However, it takes exponential time. We use a different polynomial-time technique to derandomize the CORE auction. The existence of a general derandomization technique with a polynomial runtime remains open.

7. REFERENCES

- [1] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Proc. of the 42nd IEEE Symposium on Foundations of Computer Science*, 2001.

- [2] A. Blum and J. Hartline. Near-Optimal Online Auctions. In *Proc. 16th Symp. on Discrete Alg.* ACM/SIAM, 2005.
- [3] T. Ebert. *Applications of recursive operators to randomness and complexity*. PhD thesis, University of California at Santa Barbara, 1998.
- [4] U. Feige. You Can Leave Your Hat On (if you guess its color). Technical Report MCS04-03, Computer Science and Applied Mathematics, The Weizmann Institute of Science, 2004.
- [5] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proc. 34th ACM Symposium on the Theory of Computing*. ACM Press, New York, 2002.
- [6] A. Goldberg and J. Hartline. Competitiveness via Consensus. In *Proc. 14th Symp. on Discrete Algorithms*. ACM/SIAM, 2003.
- [7] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 2002. Submitted for publication. An earlier version available as InterTrust Technical Report STAR-TR-99.09.01.
- [8] A. Goldberg, J. Hartline, and A. Wright. Competitive Auctions and Digital Goods. In *Proc. 12th Symp. on Discrete Algorithms*, pages 735–744. ACM/SIAM, 2001.
- [9] D. Lehmann, L. O’Callaghan, and Y. Shoham. Truth Revelation in Approximately Efficient Combinatorial Auctions. In *Proc. of 1st ACM Conf. on E-Commerce*, pages 96–102. ACM Press, New York, 1999.
- [10] J. Mirrlees. An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies*, 38:175–208, April 1971.
- [11] H. Moulin and S. Shenker. Strategyproof Sharing of Submodular Costs: Budget Balance vs. Efficiency. *Economic Theory*, 18:511–533, 2001.
- [12] I. Segal. Optimal Pricing Mechanism with Unknown Demand. *American Economic Review*, 93:509–29, 2003.
- [13] P. Winkler. Games People don’t Play. In D. Wolfe and T. Rogers, editors, *Puzzler’s tribute. A feast for the mind*. A K Peters, 2001.