Hats, auctions and derandomization*

Oren Ben-Zwi† Ilan Newman‡ Guy Wolfovitz§

Abstract

We investigate derandomizations of digital good randomized auctions. We propose a general derandomization method which can be used to show that for every random auction there exists a deterministic auction having asymptotically the same revenue. In addition, we construct an explicit optimal deterministic auction for bi-valued auctions.

1 introduction

Marketing a digital good may suffer from a low revenue due to incomplete knowledge of the marketer. Consider, for example, a major sport event with some 10^8 potential TV viewers. Assume farther that all potential viewers are willing to pay 10$ or more each, and that no more than 10^6 are willing to pay 100$ each. If the concessionaire will charge 1$ or 100$ as a pay per view price for the event, the overall collected revenue will be $10^8$$ at the most. This is worse than the $10^9$$ that can be collected, having known the valuations beforehand.

This lack of knowledge motivates the study of an *unlimited supply, unit demand, single item* auction. Goldberg et al. [17] studied these auctions and suggested, in order to obtain a prior free, worst case analysis framework, to use the optimal fixed price auction as a benchmark to compare with. They adopted the *online algorithms* terminology [25] and named the revenue of the fixed price auction the *offline revenue* and the revenue of a multi-price *truthful* auction, i.e., an auction for which every bidder has an incentive to bid its own value, *online revenue*. The *competitive ratio* of an auction for a bid vector b is defined to be the ratio between the best offline revenue for b to the revenue of that auction on b. The competitive ratio of an auction is just the worst competitive ratio of that auction on all possible bid vectors. For random auctions, a similar notion is defined by taking the expected revenue. If an auction has a constant competitive ratio it is said to be *competitive*. If an auction has a constant competitive ratio, possibly with some small additive loss, it is said to be *general competitive* (see Section 2 for definitions). Later, Koutsoupias and Pierrakos [20] used *online auctions* in the usual context of online algorithms, but here we will stick to Goldberg et al.'s [17] notation.

Although this optimal fixed price benchmark may seem unintuitive at first, it was a-posteriori explained by a work of Hartline and Roughgarden as an important general template. If we charac-

*Part of the results presented in this work are based on works which appeared in [5–7].

†Faculty of Computer Science, University of Vienna, Vienna, Austria. Email: oren.benzwi@gmail.com

‡Department of Computer Science, University of Haifa, Haifa, Israel. Email: ilan@cs.haifa.ac.il.

§Department of Computer Science, University of Haifa, Haifa, Israel. Email: gwolfovi@cs.haifa.ac.il.

terize the set of all mechanisms that are Bayesian optimal for some i.i.d. distribution over possible valuations, and then define a worst case performance benchmark that corresponds to competing simultaneously with all on a fixed, worst case, valuation profile. It turns out that for this setting this defines the optimal first price auction [18].

It can be shown [15] that the online revenue is no more than the offline revenue, even though the former uses more than one optional price. In fact, there even exists a lower bound of 2.42 on the competitive ratio of any auction [16].

Note, however, that the optimal offline revenue is unknown in truthful auctions.

It is well known (see for example [22]) that in order to achieve truthfulness one can use only the set of *bid independent auctions*, i.e., auctions in which the computation of the price offered to a bidder is done while ignoring its own bid value. Hence, an intuitive auction that often comes to one's mind is the *Deterministic Optimal Price (DOP) auction* [2, 15, 24]. In this auction the mechanism computes and offers each bidder the price of an optimal offline auction for all other bids. This auction preforms well on most bid vectors. In fact, it was even proved by Segal [24] that if the input is chosen uniformly at random, then this auction is asymptotically optimal. For a worst case analysis, however, it preforms very poor. Consider, for example, an auction in which there are n bidders and only two possible bid values: 1 and h, where $h \gg 1$. We denote this setting as *bi-valued* auctions. Let n_h be the number of bidders who bid h. Applying *DOP* on a bid vector for which $n_h = n/h$ will result in a revenue of n_h instead of the n revenue of an offline auction. This is because every "h-bidder" is offered 1 (since $n - 1 > h \cdot (n_h - 1)$) and every "1-bidder" is offered h (since $n - 1 < h \cdot n_h$). Here an "h-bidder" refers to a bidder that bids h and a "1-bidder" refers to a bidder that bids 1. Therefore, the competitive ratio of *DOP* is unbounded. Similar examples regarding the performances of *DOP* in the bi-valued auction setting appeared already in Goldberg et al. [15], and in Aggarwal et al. [2].

Goldberg et al. [17] showed that there exist random competitive auctions. Other works with different random competitive auctions, competitive lower bounds, and better analysis of existing auctions were presented, see for example [13, 14, 16]. For a survey see the work of Hartline and Karlin that appeared in [23, Chapter 13] (Profit Maximization In Mechanism Design). In all these works, no deterministic auction was presented. In fact, Goldberg et al. [15, 17] even proved that randomization is essential if the auction is symmetric (aka anonymous), i.e., if the outcome of the auction does not depend on the order of the input bids.

Aggarwal et al. [1,2] later showed how to construct from any randomized auction a deterministic, asymmetric auction with approximately the same revenue. In order to establish the result the authors used *guessing auctions*, in which the bidder gets the good only if the price equals *exactly* the bid (rather than lower or equal, as in a the standard setting). The guessing auctions are then "solved" using a *hat guessing game* which they introduce. Therefore for every random auction a deterministic "dual" auction can be constructed, though not in polynomial time, where the deterministic one *guarantees* a revenue which is close to the expected revenue of the random auction. Following is a more formal claim of their result: given a randomized auction A which accepts bid-vectors in $[1, h]^n$, there exists a deterministic, asymmetric auction A_D satisfying $P_{A_D}(b) \geq P_A(b)/4 - O(h)$ for every $b \in [1, h]^n$; here $P_{A_D}(b)$ is the revenue of A_D given a bid-vector b and $P_A(b)$ is the expected revenue of A given a bid-vector b. The same result also holds in the more restrictive case where A accepts only discrete bid-vectors in $[h]^n$. In addition, Aggarwal et al. showed that if the bid-vectors are restricted to be vectors of powers of 2 then the multiplicative factor of 4 above can be improved to 2.

1.1 Our Results

We show how to eliminate the multiplicative factor of 4. We use *Lovász's Local Lemma* [11] to show that for every random auction there exists a deterministic *auction* that guarantees the expected revenue of the random one, on any bid vector. More formally, for a randomized auction A and bid vector b let $P_A(b)$ be the expected revenue of A on b. We show that given a random auction A, there exists a deterministic auction which, given a bid-vector $b \in [h]^n$, guarantees a revenue of $P_A(b) - O(h\sqrt{n \ln hn})$. As is the case with the construction of Aggarwal et al. [1], our construction is also not polynomial time computable.

For bi-valued auctions, with bid values $\{1, h\}$ and n bidders, we show a polynomial time deterministic auction, for which we guarantee a revenue of $\max\{n, h \cdot n_h\} - O(\sqrt{n \cdot h})$, where n_h is just the number of bidders that bid h. We then show that this bound is unconditionally optimal by showing that every auction (including a random superpolynomial one), cannot guarantee more than $max\{n, h \cdot n_h\} - \Omega(\sqrt{n \cdot h})$. That is, there exists an auction with no multiplicative loss and with only $O(\sqrt{n \cdot h})$ additive loss, and every auction has at-least these losses. Let us note here, that if we restrict ourselves to anonymous auctions (symmetric) then we have a multiplicative loss of $\Omega(h)$ and an additive loss of $\Omega(n/h)$ over the $\max\{n, h \cdot n_h\}$ revenue of the best offline [2, 15].

In order to find a polynomial time deterministic auction for bi-valued auctions, we solve a certain hat guessing puzzle. *Hat guessing games* is an emerging research field in combinatorics. It was broadly brought to the attention of researchers by a work of Peter Winkler [27], and since then was studied in many works such as [3, 8–10, 12]. The beautiful work of Aggarwal et al. [2] established a connection between hat guessing games and unlimited supply, unit demand, single item auctions. In fact, the authors introduced three different hat games and used them to establish their derandomizations. We show how a different hat game can help in forming an answer to the bi-valued auctions setting. This hat game was previously studied by Doerr and by Feige [9, 12]. Our deterministic hat strategy improves Doerr's result and answers an open question of Feige.

Consider the following game. There are n players, each wearing a hat colored red or blue. Each player does not see the color of its own hat but does see the colors of all other hats. Simultaneously, each player has to guess the color of its own hat, without communicating with the other players. The players are allowed to meet beforehand, hats-off, in order to coordinate a strategy. We give a polynomial time deterministic strategy which guarantees that the number of correct guesses is at least $\max\{n_\mathrm{r}, n_\mathrm{b}\} - O(n^{1/2})$, where n_r is the number of players with a red hat and $n_\mathrm{b} = n - n_\mathrm{r}$ is the number of players with a blue hat.

1.2 Additive Loss

Already in the work that suggested using competitive analysis, namely Goldberg et al. [17], a major obstacle arises. It was indicated that no auction can be competitive against bids with one high value (see Goldberg et al. [15] for details). The first solution that was suggested to this problem was taking a different benchmark as the offline auction. This different benchmark was again the maximum single price auction, only now the number of winning bidders is bounded to be at-least two. The term competitive was then used to indicate an auction that has a constant ratio on every bid vector against any single price auction that sells at-least two items. A few random competitive auctions were indeed suggested using this definition over the years, but, as noted before, no deterministic (asymmetric) auction was ever found. In fact, this was proved not to be a coincidence when Aggarwal et al. [1] showed that no deterministic auction can be competitive even on this weaker benchmark.

Given this lower bound a new solution should be considered, and indeed Aggarwal et al. [1] suggested such. The new definition suggested generalizing the competitive notion to include also additive losses on top of the multiplicative ones considered before.

We argue that our results, and in particular the auction and the tight lower bound for the bi-valued setting, indicate that this second approach of considering also the additive loss is more accurate, as it shows how analyzing with a finer granularity turns an uncompetitive auction to an optimal one. We elaborate on this agenda in the discussion section 5.

1.3 Organization

After a short preliminaries section we introduce our general derandomization on section 3. Section 4 presents the deterministic polynomial time algorithm for bi-valued auctions together with a tight lower bound that applies even for random auctions. It starts, however, in the presentation of the new deterministic hat guessing bound that serves as a building block for the auction. We conclude with closing remarks and open problems in section 5.

2 Preliminaries

For a natural number k, let $[k]$ denote the set $\{1, 2, ..., k\}$. A bid-vector $b \in [h]^n$ is a vector of n bids, each taking a value in $[h]$. For $b \in [h]^n$ and $i \in [n]$ we denote by b_{-i} the vector which is the result of replacing the ith bid in b with a question mark; that is, b_{-i} is the vector $(b_1, b_2, \dots, b_{i-1}, ?, b_{i+1}, \dots b_n)$. For every $i \in [n]$, we let $[h]^n_{-i} = \{b_{-i} | \ b \in [h]^n\}$.

Definition 1 (Unlimited supply, unit demand, single item auction)**.** *An* unlimited supply, unit demand, single item auction *is a mechanism in which there is one* item *of unlimited supply to sell by an* auctioneer *to* n bidders*. The bidders place* bids *for the item according to their* valuation *of the item. The auctioneer then sets prices for every bidder. If the price for a bidder is lower than or equal to its bid, then the bidder is considered as a* winner *and gets to buy the item for its price. A bidder with price higher than its bid does not pay nor gets the item. The auctioneer's* revenue *is the sum of the winners prices.*

A *truthful auction* is an auction in which every bidder bids its true valuation for the item. Truthfulness can be established through *bid-independent* auctions (see for example [22]). A *bid-independent* auction is an auction for which the auctioneer computes the price for bidder i using only the vector b_{-i} (that is, without the ith bid). Two models have been proposed for describing random truthful auctions. The first, being the *truthful in expectation*, refers to auctions for which a bidder maximizes its expected utility by bidding truthfully. The second model, the *universally truthful* is merely a probability distribution over deterministic auctions. Our results uses this second definition, however, it is known that the two models collide in this setting [21].

Definition 2 (fixed price, offline auction)**.** *The* fixed price, offline auction *is the auction that on each bid vector $b \in [h]^n$ fixes a* single *price, $\alpha = \alpha(b)$ for all bidders, so to maximize the revenue given that price. Namely, α is chosen such that $\sum_{b_i \geq \alpha} \alpha$ is maximized.*

Definition 3 (General competitive auction)**.** *Let $OPT(b)$ be the best fixed-price (offline) revenue for an n-bid vector b with maximum bid h. An auction A is a* general competitive auction *if its revenue (expected revenue) from every bid vector b, $P_A(b)$ is $\geq \alpha \cdot OPT(b) - o(nh)$ where α is a constant not depending on n or h.*

2.1 A structural lemma

Let A be a randomized truthful auction that accepts bid-vectors from $[h]^n$. We think of A as a distribution over deterministic auctions. Hence, we may view A's execution in the following manner. The auction maintains a set of nm functions $\{g_{i,j}\colon\ i \in [n], j \in [m]\}$, where $g_{i,j}$ is a function from $[h]^n_{-i}$ to $[h]$. This corresponds to a collection of m deterministic auctions, where the jth is defined by the set of functions $\{g_{i,j} |\ i \in [n]\}$. On a bid-vector $b \in [h]^n$, the auction tosses some coins, and chooses, accordingly, an integer $j \in [m]$. The auction then offers bidder i the price $g_{i,j}(b_{-i})$. Let $\text{accept}_{i,j}(b)$ be 1 if $g_{i,j}(b_{-i}) \leq b_i$ and 0 otherwise. Let p_j be the probability that $j \in [m]$ was chosen. The expected revenue of the auction on input b is then:

$$P_A(b) = \sum_j p_j \sum_i \text{accept}_{i,j}(b)\ \cdot\ g_{i,j}(b_{-i}).$$

Note that for every $j \in [m]$, the set of functions $\{g_{i,j} |\ i \in [n]\}$ is just a deterministic strategy, denoted A_j, and A is, as explained before, a distribution on deterministic strategies.

Note also that given A, namely the set of m deterministic auctions, and a distribution, $\mathcal{D} = (p_1, \ldots, p_n)$ on $[m]$, A induces another randomized auction A' as follows: for a given b, it chooses *for each $i \in [n]$, independently*, a $j_i \in [m]$ according to $\mathcal{D}$ (namely $\forall i, Pr(j_i = j) = p_j$), and acts according to the set of functions thus chosen, namely $\{g_{i,j_i}, i = 1, \ldots, n\}$.

By definition, the expected revenue of A' on input b is given by:

$$P_{A'}(b) = \sum_i \sum_j p_j\ \cdot\ \text{accept}_{i,j}(b)\ \cdot\ g_{i,j}(b_{-i}).$$

We call A' the *bidder-self-randomness-dual* of A (as the function for different bidders are "not coordinated"). Comparing the revenue of A and that of A' immediately implies

Lemma 2.1. *Let A be a randomized auction and A' be its bidder-self-randomness-dual auction. Then A and A' have the same expected revenue on every bid-vector.*

This corresponds with the minimax Theorem [26] and with Yao's Lemma [28]. We note that A' is concentrated on possibly many more deterministic algorithms than A. Not only may g_1 be chosen from the jth copy while g_2 from the ℓ's copy, (namely $g_1 = g_{1,j}$ while $g_2 = g_{2,\ell}$ with $j \neq \ell$), it could also be that for different bid vectors b, b', g_1 for b is $g_{1,j}$ while g_1 for b' is $g_{1,\ell}$. This works since no consistency requirement between different b's, and/or different i's, is required in the expression for the expectation above.

2.2 Probabilistic Tools

The following two well known lemmata are used in the proofs. We explicitly state both for completeness. The first is just the famous Lovász Local Lemma [11]. We will need the following version of it [4].

Lemma 2.2 (The local lemma; symmetric case). *Let Bad_i, $1 \leq i \leq N$, be events in an arbitrary probability space. Suppose that each event Bad_i is mutually independent of a set of all the other events Bad_j but at most d, and that $\Pr[Bad_i] \leq p$ for all $1 \leq i \leq N$. If $ep(d+1) \leq 1$, where e is the base of the natural logarithm, then $\Pr[\bigwedge_{i=1}^{N} \neg Bad_i] > 0$.*

The second lemma is just a tail bound inequality proved by Hoeffding [19].

Lemma 2.3 (Hoeffding). *Let X be the average of n independent random variables X_i, where $X_i \in [a_i, b_i]$ for all i. Then: $\Pr[X < \mathbb{E}[X] - t] \leq 2 \exp\left(\frac{-2n^2t^2}{\sum_{i=1}^{n}(b_i - a_i)^2}\right)$*

3 A General Derandomization

This section is devoted for the proof of the following Theorem.

Theorem 3.1. *Let A be a randomized auction which accepts bid-vectors in $[h]^n$. Assume that A has expected revenue $P_A(b)$ for every bid-vector $b \in [h]^n$. Then there exists a deterministic auction A_D that guarantees a revenue of $P_{A_D}(b) \geq P_A(b) - O(h\sqrt{n \ln hn})$ for every bid-vector $b \in [h]^n$.*

The proof of Theorem 3.1 can be outlined as follows. Given a randomized auction A we first move to the bidder-self-randomness-dual auction A' that has the same revenue as A. Let A_D be a deterministic auction that is chosen according to the distribution that A' induces on deterministic auctions. We show that the event Bad_b, defined by $P_{A_D}(b) < P_{A'}(b) - t$, depends on a relatively few number of other events $Bad_{b'}$. Moreover, for every b, we have that the probability of Bad_b is sufficiently small. We then apply the Lovász Local Lemma to show that there exists a single deterministic auction A_D, namely a choice of a collection of functions $\{g_{i,j_i}, i \in [n]\}$, for which none of the events Bad_b occur. This will conclude the proof of the theorem.

We stress the fact that the result of Aggarwal et al. [1] is more general in the sense that it deals with bid-vectors in $[1, h]^n$, while Theorem 3.1 only deals with discrete bid-vectors. However, discrete bid-vectors make sense in real life auctions where bids are monetary bids, being made with discrete valued currency. We stipulate that the construction used in the proof of Theorem 3.1 is not known to be polynomial time computable and that this is also the case in the construction of Aggarwal et al. [1].

We now formally present the proof.

Proof of Theorem 3.1. Let A be a randomized auction which accepts bid-vectors in $[h]^n$, using a distribution over m deterministic auctions. Let $\{g_{i,j}\colon\ i \in [n], j \in [m]\}$ be the set of functions that A maintains. Let $(p_1, \ldots, p_n)$ be the distribution over $[m]$ that is used by A, and let A' be the bidder-self-randomness-dual of A. Namely, in which for each b, g_i is chosen independently for each i, among all $g_{i,j},\ j \in [m]$, with the corresponding probabilities $\{p_j, j \in [m]\}$.

For every vector b let $P_A(b)$ be the revenue expected by A on b. By Lemma 2.1, the revenue of A' on a bid vector b is

$$P_A(b) = \sum_i \sum_j p_j \ \cdot\ \text{accept}_{i,j}(b)\ \cdot\ g_{i,j}(b_{-i})$$

In the following, all events are with respect to the distribution defined by the runs of the random auction A'. Namely, the probability space contains, for each bid vector b, an n-tuple of independently chosen values $(g_{i,j})_{i\in[n]}$ as defined above by A'.

Let $t := h\sqrt{n \ln 2hn}$. For a random run of A', namely, for a deterministic auction A'_D that is chosen at random according to the distribution induced by A' on deterministic auctions, let Bad_b be the event that $P_{A'_D}(b) < P_A(b) - t$.

We need the following two claims.

Claim 3.2. *For all $b \in [h]^n$, $Pr[Bad_b] < 1/(2h^2n^2)$.*

Proof. Fix $b \in [h]^n$ and let X_i be the revenue extracted from bidder i in a run of A', namely, for g_i chosen for b. Then, $X_i = \text{accept}_{i,j}(b) \cdot g_i(b_{-i})$. Note that $X_i \in [1, h]$ for all i and that the X_i's are independent random variables. Let X be the sum of the X_i's. Namely, X is the revenue on b for that specific run. We have already argued that $\mathbb{E}[X] = P_A(b)$, thus,

$$\Pr[Bad_b] = \Pr[X < \mathbb{E}[X] - t],$$

which by Lemma 2.3 is at most $2\exp\left(\frac{-2t^2}{h^2n}\right)$. The claim now follows since $t = h\sqrt{n \ln 2hn}$. ■

Claim 3.3. *For all $b \in [h]^n$, Bad_b depends on at most hn other events $Bad_{b'}$.*

Proof. Let $b \in [h]^n$ be fixed. The tuple $(g_i(b_{-i})), i = 1, \ldots, n$, defines the revenue on b for a strategy that chooses this tuple. Thus for a vector $b' \neq b$, the events $Bad_b, Bad_{b'}$ may be dependent only if for some $i \in [n]$, $b_{-i} = b'_{-i}$. This is so, since otherwise, once the tuple for b is chosen, that leaves complete freedom in choosing the tuple for b'. For a fixed b, and fixed i there are h possible vectors b', for which $b_{-i} = b'_{-i}$. Hence the claim follows. ■

Combining the two claims above with the Lovász Local Lemma, we get that with positive probability Bad_b does not occur for all $b \in [h]^n$. Hence, there is a set of tuples $(g_i(b_{-i}))_{i\in[n]}, b \in [h]^n$, for which Bad_b does not occur for every $b \in [h]^n$. This set of tuples is the deterministic strategy A_D for which for every bid-vector $b \in [h]^n$, $P_{A_D}(b) \geq P_A(b) - t$. This completes the proof of the theorem. ■

4 Bi-valued Auctions and a Hat Game

We establish a connection between bi-valued auctions and a specific hat guessing game known as the *majority hat game*. This game was studied by Doerr [9] and later by Feige [12]. We derive new results regarding this game, which enable us to solve the bi-valued auction problem optimally.

4.1 A Hat Game

A group of n players is gathered, n_{r} of which wear a red hat and $n_{\mathrm{b}} = n - n_{\mathrm{r}}$ wear a blue hat. Every player in the group can see the colors of the hats of the other players, but cannot see and does not know the color of its own hat, a color which has been picked by an adversary. No form of communication is allowed between the players. At the mark of an unseen force, each player simultaneously guesses the color of its hat. The objective of the players as a group is to make the total number of correct guesses as large as possible. In order to achieve this goal, the players are allowed to meet beforehand, hats-off, and agree upon some strategy.

Theorem 4.1. *There exists a polynomial time deterministic strategy which guarantees at least* $\max\{n_r, n_b\} - O(n^{1/2})$ *correct guesses.*

Let us give a few remarks. First, this result is optimal, in the sense that any deterministic strategy can guarantee only $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - \Omega(n^{1/2})$ correct guesses in the worst case; this was proved by Feige [12] and Doerr [9]. Second, this result improves a result of Doerr [9] who gave a polynomial time deterministic strategy which guarantees at least $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{2/3})$ correct guesses, and a result of Feige [12] who gave a non-polynomial time deterministic strategy which guarantees at least $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$ correct guesses. Feige further asked whether there exists a polynomial time deterministic strategy which guarantees this last bound, and our result answers this question affirmatively. Lastly it should be noted that Winkler [27], who brought the problem to light, gave a simple polynomial time deterministic strategy which guarantees $\lfloor n/2 \rfloor$ correct guesses.

The proof of Theorem 4.1 has two parts. First, we design a polynomial time randomized strategy for the players, a strategy which guarantees that under any hat assignment, the expected number of correct guesses is $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$. We then derandomize this strategy by giving a polynomial time deterministic strategy that always achieves, up to another $O(n^{1/2})$ additive loss, the expected number of correct guesses of the randomized strategy.

4.1.1 Randomized strategy

Let the players agree in advance on some ordering so that the ith player is well defined and known to all. Under a given hat assignment, let $\chi_{\mathrm{r}}(i)$ be the number of red hats that the ith player sees. Analogously, let $\chi_{\mathrm{b}}(i)$ be the number of blue hats that the ith player sees. Say that a player is red (respectively blue) if she wears a red (respectively blue) hat.

Our strategy is a collection of randomized strategies, one for each player. We describe the strategy of the ith player, Paula. First Paula computes two positive integers $a(i)$ and $b(i)$, and sets $p(i) = a(i)/b(i)$. If $|\chi_{\mathrm{r}}(i) - \chi_{\mathrm{b}}(i)| \leq 1$, then Paula takes $a(i) = 1$ and $b(i) = 2$, so that $p(i) = 1/2$. Otherwise, $|\chi_{\mathrm{r}}(i) - \chi_{\mathrm{b}}(i)| \geq 2$ and so we have either $\chi_{\mathrm{r}}(i) = n/2 + c$ for some

$c > 0$ or $\chi_{\mathrm{b}}(i) = n/2 + c$ for some $c > 0$ (but not both). In the former case Paula takes $a(i) = \min\{\lfloor n^{1/2}\rfloor, \lceil c\rceil\}$ and $b(i) = \lfloor n^{1/2}\rfloor$, so that $p(i) = \min\{1, \lceil c\rceil/\lfloor n^{1/2}\rfloor\}$ and in the latter case she takes $a(i) = \lfloor n^{1/2}\rfloor - \min\{\lfloor n^{1/2}\rfloor, \lceil c\rceil\}$ and $b(i) = \lfloor n^{1/2}\rfloor$, so that $p(i) = 1 - \min\{1, \lceil c\rceil/\lfloor n^{1/2}\rfloor\}$. Note that $a(i), b(i)$ and $p(i)$ can be computed in polynomial time. Having $p(i)$ at hand, Paula draws a uniformly random real p in the unit interval, guesses red if $p \leq p(i)$ and blue otherwise.

Lemma 4.2. *If each player follows the above strategy then the expected number of correct guesses is at least* $\max\{n_r, n_b\} - O(n^{1/2})$.

Proof. We shall assume throughout the proof that $n_{\mathrm{r}} \geq n_{\mathrm{b}}$; the argument for the other case is symmetric. We consider the following cases.

- $n_{\mathrm{r}} = n_{\mathrm{b}}$. In that case, every player guesses correctly with probability $1/2$. Thus the expected number of correct guesses is $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\}$.
- $n_{\mathrm{r}} \in \{n_{\mathrm{b}} + 1, n_{\mathrm{b}} + 2\}$. In that case, every red player guesses red with probability $1/2$ and every blue player guesses blue with probability $1 - O(n^{-1/2})$. Thus, the expected number of correct guesses is $n_{\mathrm{r}}/2 + n_{\mathrm{b}}(1 - O(n^{-1/2}))$, which is clearly at least $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$.
- $n_{\mathrm{r}} \geq n_{\mathrm{b}} + 3$. In the last case we examine the gap between n_{r} and $n/2$. If this gap is small ($\ll \sqrt{n}$) then enough blue players will guess blue. Otherwise enough red players will guess red. Formally, let $x > 1$ satisfy $n_{\mathrm{r}} = n/2+x$, so that $n_{\mathrm{b}} = n/2-x$. First assume that $\lceil x\rceil \leq \lfloor n^{1/2}\rfloor$. In that case, every red player guesses red with probability $\lceil x-1\rceil/\lfloor n^{1/2}\rfloor = (\lceil x\rceil - 1)/\lfloor n^{1/2}\rfloor$, and every blue player guesses blue with probability $1 - \lceil x\rceil/\lfloor n^{1/2}\rfloor$. Therefore, the expected number of correct guesses is

$$\begin{aligned}(n/2+x)(\lceil x\rceil - 1)/\lfloor n^{1/2}\rfloor + (n/2-x)(1 - \lceil x\rceil/\lfloor n^{1/2}\rfloor) &= \\ (n/2+x)\lceil x\rceil/\lfloor n^{1/2}\rfloor - (n/2+x)/\lfloor n^{1/2}\rfloor + (n/2-x)(1 - \lceil x\rceil/\lfloor n^{1/2}\rfloor) &\geq \\ (n/2-x)\lceil x\rceil/\lfloor n^{1/2}\rfloor - (n/2+x)/\lfloor n^{1/2}\rfloor + (n/2-x)(1 - \lceil x\rceil/\lfloor n^{1/2}\rfloor) &\geq \\ (n/2-x) - (n/2+x)/\lfloor n^{1/2}\rfloor &\geq n/2 - 4n^{1/2},\end{aligned}$$

  which is at least $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$, since $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} \leq n/2 + O(n^{1/2})$. Next assume that $\lceil x\rceil > \lfloor n^{1/2}\rfloor$. In that case, every red player guesses its hat correctly with probability 1 and so the expected number of correct guesses is at least $n_{\mathrm{r}} \geq \max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$. ■

4.1.2 Derandomization

The randomized strategy we gave above has two phases. In the first phase the ith player computes in *deterministic* polynomial time some number $p(i)$ in the unit interval. Moreover, the strategy is symmetric namely, for some p_{r} and p_{b} that depend only on the number of red hats and the number of blue hats, we have $p(i) = p_{\mathrm{r}}$ if the ith player is red and $p(i) = p_{\mathrm{b}}$ if the ith player is blue. Given the first phase, the second phase guarantees that the expected number of correct guesses is $p_{\mathrm{r}}n_{\mathrm{r}} + (1 - p_{\mathrm{b}})n_{\mathrm{b}}$, which was shown by Lemma 4.2 to be at least $\max\{n_{\mathrm{r}}, n_{\mathrm{b}}\} - O(n^{1/2})$. We show in the following that, if for all $1 \leq i \leq n$ the ith player has determined $p(i)$, we can replace the second phase of the randomized strategy by a non-symmetric, polynomial time, *deterministic* strategy which guarantees that at least $p_{\mathrm{r}}n_{\mathrm{r}} - O(n^{1/2})$ red players make a correct guess and at least $(1-p_{\mathrm{b}})n_{\mathrm{b}} - O(n^{1/2})$ blue players make a correct guess. By Lemma 4.2 this will imply Theorem 4.1.

Suppose that for all $1 \le i \le n$, the ith player has determined $a(i), b(i)$ and $p(i)$. The following is the strategy that the ith player follows in order to determine its guess.

1. Let $X(i) = \sum_j j$, where the sum ranges over all $j \neq i$ such that the jth player is red.
2. Let $Y(i) = \sum_j 1$, where the sum ranges over all $j < i$ such that the jth player is red.
3. Let $Z(i) = i + X(i) + (b(i) - 1)Y(i) \pmod{b(i)}$.
4. Guess red if $Z(i) < a(i)$, blue otherwise.

Note that the above deterministic strategy can be implemented so that its running time is polynomial in n. This fact together with the next lemma proves Theorem 4.1.

Lemma 4.3. *Suppose that for all $1 \le i \le n$, the ith player has computed $a(i), b(i)$ and $p(i)$. If each player follows the above strategy, then the number of red players that make a correct guess is at least $p_r n_r - O(n^{1/2})$ and the number of blue players that make a correct guess is at least $(1 - p_b)n_b - O(n^{1/2})$.*

Proof. In what follows we make use of the following facts, which follow from the definition of $a(i)$ and $b(i)$ in the previous section. If the ith player and the jth player both have a hat of the same color, then $a(i) = a(j)$ and $b(i) = b(j)$. Furthermore, for all $1 \le i \le n$, $1 \le b(i) \le 2n^{1/2}$.

Let us first consider the red players. Let $1 \le i < j \le n$ be two indices of players so that the ith player's hat and the jth player's hat are both red and furthermore, for all $i < k < j$ we have that the kth player's hat is blue. Let $a(i) = a(j) = a$ and $b(i) = b(j) = b$ so that $p_{\mathrm{r}} = a/b$. We have $i + X(i) = j + X(j)$ and $Y(j) - Y(i) = 1$. Thus $Z(j) - Z(i) = b - 1 \pmod{b}$. This implies that out of each b consecutive red players, a guess red. Thus, since $b \le 2n^{1/2}$, at least $p_{\mathrm{r}}n_{\mathrm{r}} - O(n^{1/2})$ red players guess red.

Next consider the blue players. Let $1 \le i < j \le n$ be two indices of players so that the ith player's hat and the jth player's hat are both blue and furthermore, for all $i < k < j$ we have that the kth player's hat is red. Let $a(i) = a(j) = a$ and $b(i) = b(j) = b$ so that $p_{\mathrm{b}} = a/b$. We have $X(i) = X(j)$ and $Y(j) - Y(i) = j - i - 1$. Thus $Z(j) - Z(i) = j - i + (b-1)(j - i - 1) \pmod{b} = b(j - i) - b + 1 \pmod{b} \equiv 1 \pmod{b}$. This implies that out of each b consecutive blue players, $b - a$ guess blue. Thus, since $b \le 2n^{1/2}$, at least $(1 - p_{\mathrm{b}})n_{\mathrm{b}} - O(n^{1/2})$ blue players guess blue. ■

4.2 A Bi-valued Auction

Consider *bi-valued auctions*, in which there are n bidders, each can select a value from $\{1, h\}$. The auction's revenue equals the number of bidders it offers 1 plus h times the number of bidders it offers h if indeed their value is h. Let $n_h(b)$ denote the number of bidders who bids h in a bid vector b. Recall that the best offline revenue on vector b is $\max\{n, h \cdot n_h(b)\}$. In this section we will prove the following.

Theorem 4.4. *For bi-valued auctions with n bidders and values from $\{1, h\}$*

1. *There exists a polynomial time deterministic auction that for all bid vector b has revenue*

$$\max\{n, h \cdot n_h(b)\} - O(\sqrt{n \cdot h})$$

2. *There is no auction that for all bid vector b has revenue*

$$\max\{n, h \cdot n_h(b)\} - o(\sqrt{n \cdot h})$$

Note that the lower bound result is unconditional and applies also for randomized superpolynomial auctions. We proceed with a proof for the upper bound in the next section and a proof for the lower bound in section 4.2.2.

4.2.1 An Auction

We present next a solution to the bi-valued auction problem, namely we show an optimal polynomial time deterministic auction. We start again by describing a random auction. A derandomization will be built later using the same methods we presented in the former section for the hat guessing problem.

A Random Bi-valued Auction

For a fixed input b, let n_h be the number of h-bids in b and $n_h(i)$ be the number of h-bids in b_{-i}. Let $p'(i) = \frac{h \cdot n_h(i) - n}{h \cdot \sqrt{n_h(i)}}$. If $p'(i) \leq 0$ set $p(i) = 0$ and if $p'(i) \geq 1$ set $p(i) = 1$. Otherwise, $(0 < p'(i) < 1)$, set $p(i) = p'(i)$. The auction offers value h for bidder i with probability $p(i)$ and 1 otherwise.

Lemma 4.5. *The expected revenue of the auction described above is* $\max\{n, h \cdot n_h\} - O(\sqrt{n \cdot h})$

Proof. If $\exists i, p(i) \neq p'(i)$ then either $h \cdot n_h(i) \leq n$ so the auction will offer 1 to any 1-bidder and the revenue will be at-least $n = \max\{n, h \cdot n_h\}$, or $h \cdot n_h(i) - n \geq h \cdot \sqrt{n_h(i)}$ so every h-bidder will be offered h with probability $1 - O(1/\sqrt{n_h})$ and the expected revenue thus is $hn_h \cdot (1 - 1/\sqrt{n_h}) = \max\{n, h \cdot n_h\} - O(\sqrt{n \cdot h})$. In either case our auction's revenue is $\max\{n, h \cdot n_h\} - O(\sqrt{n \cdot h})$. Assume now that $\forall i, p(i) = p'(i)$, note that in this case $|n - h \cdot n_h| = O(h\sqrt{n_h})$. The expected revenue for any bid vector with n_h bids of value h is then:

$$h \cdot n_h \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}} + n_h \cdot (1 - \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}}) + (n - n_h) \cdot (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h}})$$

$$\geq h \cdot n_h \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}} + n_h \cdot (1 - \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}}) + (n - n_h) \cdot (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h - 1}})$$

$$\geq h \cdot n_h \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}} + n \cdot (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h - 1}}) + n_h \cdot (1 - \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}}) - n_h \cdot (1 - \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}})$$

$$= h \cdot n_h \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}} + n \cdot (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h - 1}})$$

$$= h \cdot n_h \cdot \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h - 1}} + n \cdot (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h - 1}}) - \frac{h \cdot n_h}{\sqrt{n_h - 1}}$$

Observe that the sum of the first two terms in the last expression above is $\max\{n, h \cdot n_h\} - O(\sqrt{n \cdot h})$. This is because $|n - h \cdot n_h| = O(\sqrt{n \cdot h})$. The third term however, can be absorbed also into the $O(\sqrt{n \cdot h})$, which completes the proof of the lemma. ■

Hence our auction's expected revenue is within an additive loss of $O(\sqrt{n \cdot h})$ from the revenue of the best offline as promised. As noted before, a derandomization for this auction can be built using the same ideas appeared in the hat guessing game. This derandomization produces an auction which has for the worst case only another additive loss of $O(\sqrt{n \cdot h})$ over the expected revenue of the random auction. Hence, in total, an additive loss of $O(\sqrt{n \cdot h})$ over the best offline revenue is achieved. We sketch this derandomization here for completeness.

Derandomization

Let $a(i) = h \cdot n_h(i) - n$ and $b(i) = h \cdot \sqrt{n_h(i)}$. The auction computes the value offered to bidder i according to the following.

1. Let $X(i) = \sum_j j$, where the sum ranges over all $j \neq i$ such that the jth bidder bids h.
2. Let $Y(i) = \sum_j 1$, where the sum ranges over all $j < i$ such that the jth bidder bids h.
3. Let $Z(i) = i + X(i) + (b(i) - 1)Y(i) \pmod{b(i)}$.
4. Offer h to bidder i if $Z(i) < a(i)$. Otherwise offer 1 to the i's bidder.

Note that for the random auction whenever $p(i) = p'(i)$ (or as stated here $a(i)/b(i) \in [0,1]$) we can define the probability that a 1-bidder will be offered 1, $p_{1,1} = (1 - \frac{h \cdot n_h - n}{h \cdot \sqrt{n_h}})$, the probability that an h-bidder will be offered 1, $p_{h,1} = (1 - \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}})$ and the probability that an h-bidder will be offered h, $p_{h,h} = \frac{h \cdot (n_h - 1) - n}{h \cdot \sqrt{n_h - 1}}$. The proof of the following lemma resembles the proof of lemma 4.3, noting that we should also consider the case of "wrong" offers for h-bidders.

Lemma 4.6. *An auction that follows the above formulation gains revenue of* $n_h \cdot (h \cdot p_{h,h} + p_{h,1}) - O(\sqrt{n \cdot h})$ *from all h-bidders. From the* 1*-bidders, the auction collects* $(n - n_h)p_{1,1} - O(\sqrt{n \cdot h})$*.*

Proof. Let $a(1)$ be the (identical) value $a(i)$ computed by all 1-bidders. In the same manner let $b(1)$, $a(h)$, $b(h)$ be the (identical) values computed by all bidders. The lemma follows Lemma 4.5 and the following claim:

Claim 4.7.

- *For every* $b(1)$ *consecutive* 1*-bidders the auction will offer* h *to* $a(1)$ *of them and* 1 *to* $b(1) - a(1)$
- *For every* $b(h)$ *consecutive* h*-bidders the auction will offer* h *to* $a(h)$ *of them and* 1 *to* $b(h) - a(h)$

Proof. Consider the h-bidders first. Let $1 \leq i < j \leq n$ be the indices of two consecutive h-bidders. We have $i + X(i) = j + X(j)$ and $Y(j) - Y(i) = 1$. Thus $Z(j) - Z(i) = b(h) - 1 \pmod{b(h)}$. This implies that out of each $b(h)$ consecutive h-bidders, $a(h)$ will be offered h and $b(h) - a(h)$ will be offered 1.

Next consider the 1-bidders. Let $1 \leq i < j \leq n$ be the indices of two consecutive 1-bidders. We have $X(i) = X(j)$ and $Y(j) - Y(i) = j - i - 1$. Thus $Z(j) - Z(i) = j - i + (b(1) - 1)(j - i - 1) = b(1)(j - i) - b(1) + 1 \equiv 1 \pmod{b(1)}$. This implies that out of each $b(1)$ consecutive 1-bidders, $b(1) - a(1)$ are offered 1. ■

■

It is clear that this auction can be implemented in polynomial time as claimed, hence the upper bound of Theorem 4.4 follows.

Informal Remark: A natural critic that should arise at first glance of our "complicated" suggested auction is its being "unintuitive". How can one explain/excuse suboptimal actions whenever $n_h \neq n/h$? Why not deploy *DOP* in these settings? Note, however, that the proposed auction does exactly the same. On most inputs it acts as the *DOP* and only on inputs where $n_h \approx n/h$ it deploys the "sophisticated" auction. In particular, the auction "sacrifices the accuracy" of results whenever for the bid vector b we have that $n \leq hn_h(b) \leq n + h\sqrt{n_h(b)}$. This "sophisticated sacrifice", however, results in turning an unbounded competitive auction into an optimal one.

4.2.2 A Lower Bound

We prove optimality of the suggested auction in the previous section. For this we prove a lower bound on the additive loss of any bi-valued auction. The lower bound is unconditional and holds also for the expected revenue of random auctions. Furthermore, the bound does not depend on the computation time needed for the auction, therefore, randomness has no significant effect on the revenue gained in this setting.

Lemma 4.8. *Let A be an auction for the bi-valued $\{1, h\}$ setting and let $P_A(b)$ be A's revenue on bid vector b. Then $P_A(b)$ equals* $\max\{n, h \cdot n_h\} - \Omega(\sqrt{h \cdot n})$, *where b is of size n and n_h is the number of bids of value h in b.*

Proof. To prove a lower bound on the difference between the offline revenue $\max\{n, h \cdot n_h\}$, and any auction we define a distribution $\mathcal{D}$ on the possible two-values bid vectors $\{1, h\}^n$. We then show that for any deterministic auction, the expected revenue for a random bid vector b (expectation now is with respect to $\mathcal{D}$), is at most P. On the other hand, we show that the expected revenue of the offline single price (over the distribution $\mathcal{D}$) is at least $P + \Delta$, for some Δ. This implies (by standard averaging argument, see for example [28]), that for any auction, (including randomized ones), there must be some vector b for which the auction's revenue is Δ less than the fixed-price offline optimal auction.

The distribution $\mathcal{D}$ in our case is quite simple: for every bidder $i \in [n]$ independently, set $b_i = h$ with probability $1/h$ and $b_i = 1$ with probability $1 - 1/h$. Now, for every deterministic truthful auction, knowing $\mathcal{D}$, the price for every element should better be in $[h]$, otherwise there is another auction that assign prices in $[h]$ and achieves at least the same revenue for every bid vector (the one that assigns 1 for every value less than 1 and h for every value higher than h). Further, for such auction, the revenue is the sum of revenues obtained from the n bidders. Thus the expectation is the sum of expectations of the revenue obtained from the single bidders. Since for bidder i the expectation is exactly 1 (since for any fixed b_i the auction must set a constant price $\alpha \in [h]$ independent of b_i. Hence for $\alpha > 1$, the expected revenue from bidder i is $\frac{1}{h} \cdot h = 1$, and for $\alpha = 1$ the expected value is clearly 1). We conclude that for every deterministic truthful auction as above, the expected revenue (with respect to $\mathcal{D}$), is exactly n.

We now want to prove that the expected revenue of the fixed-price offline auction, that knows b, is $n + \Omega(\sqrt{hn})$. We know, however, the exact revenue of such auction for every bid vector b. It is just $M(b) = \max\{n, h \cdot n_h(b)\}$, where $n_h(b)$ is the number of h-bids in b.

Thus the expected revenue is

$$\mathbb{E}_{\mathcal{D}}[M(b)] := \sum_{i<n/h} n \cdot \binom{n}{i} \cdot (1/h)^i \cdot (1-1/h)^{n-i} + \sum_{i>n/h} h \cdot i \cdot \binom{n}{i} \cdot (1/h)^i \cdot (1-1/h)^{n-i} \qquad (1)$$

$$+ n \cdot \binom{n}{n/h} (1/h)^{n/h} (1-1/h)^{n-n/h}$$

To estimate this sum, it is instructive to examine the following deterministic auction which we note before as DOP. On each vector b, DOP assigns value h for every bidder i for which the number of h-bids in b_{-i}, is at least n/h (we assume n/h is an integer), and 1 otherwise.

On one side, as argued before, the expected revenue of DOP with respect to $\mathcal{D}$ is

$$\mathbb{E}[P_{DOP}] = n \qquad (2)$$

On the other hand, the same expression, is by definition,

$$\mathbb{E}[P_{DOP}] = \sum_{i<n/h} n \cdot \binom{n}{i} \cdot (1/h)^i \cdot (1-1/h)^{n-i} + \sum_{i>n/h} h \cdot i \cdot \binom{n}{i} \cdot (1/h)^i \cdot (1-1/h)^{n-i} \qquad (3)$$

$$+(n/h) \cdot \binom{n}{n/h} (1/h)^{n/h} (1-1/h)^{n-n/h}$$

Comparing the expression in Equation (1) and Equation (3), and using Equation (2), we get:

$$\mathbb{E}_D[M(b)] = n + (n - n/h) \cdot \binom{n}{n/h} (1/h)^{n/h} (1-1/h)^{n-n/h}$$

Hence we conclude that the difference in expectation between offline revenue $\mathbb{E}_D[M(b)]$ and the expected revenue on any deterministic auction, which is n, is,

$$\mathbb{E}_D[M(b)] - n = n(1-1/h) \cdot \binom{n}{n/h} (1/h)^{n/h} (1-1/h)^{n-n/h}$$

By Stirling's approximation we know that

$$\binom{n}{n/h} = \Theta\left(\frac{\sqrt{h/n}}{\sqrt{(1-1/h)}(1/h)^{n/h}(1-1/h)^{n-n/h}} \right)$$

Therefore, the additive lost is at least $\Omega(\sqrt{h \cdot n})$ as claimed. ∎

5 Discussion

We have presented an existential general *derandomization* for *unlimited supply, unit demand, single item* auctions. This derandomization produces an auction with the same asymptotic revenue guarantee as the expected revenue of the randomized. Furthermore, this derandomization is *direct* (in the sense that no intermediate like "guessing auction" is involved). This answers an open question posed by Aggarwal et al. [1] about the existence of direct derandomizations. Another interesting open question posed in the same work on the existence of a general polynomial time derandomization, remains open and challenging.

Bi-valued auctions appeared as examples in several works, such as [2, 15]. We present here a connection between these auctions and a certain *hat guessing game* [9, 12]. Solving this puzzle optimally results in an optimal deterministic auction for bi-valued auctions. Surprisingly, the establishment of the tight lower bound for these auctions involves analyzing the *DOP*, the deterministic optimal auctions for i.i.d. inputs.

Our general derandomization suffers from an additive loss of $\widetilde{O}(h\sqrt{n})$ over the expected revenue of a random auction. Aggarwal et al. [2] proved that every deterministic auction will suffer from an additive loss over the *best offline* auction, hence did not rule out exact derandomizations. We showed, by the lower bound on bi-valued auctions, that *every* auction (including a random one) suffers from an additive loss of $\Omega(\sqrt{nh})$ over the best offline auction. Clearly, our understanding of the additive loss is not complete yet and needs some further investigation.

Farther research should ask whether there exists more cases of exact derandomization? Is there a general exact derandomization? And of-course, try to deploy these derandomization techniques to other mechanism design settings. Another interesting future direction, noticing that the connection between truthful auctions and hat guessing games was not a coincidence, is to reinforce these connection, maybe with different kind of auctions.

References

[1] G. Aggarwal, A. Fiat, A. V. Goldberg, J. D. Hartline, N. Immorlica, and M. Sudan, *Derandomization of auctions*, Proceedings of the 37th annual acm symposium on theory of computing (stoc), 2005, pp. 619–625.

[2] G. Aggarwal, A. Fiat, A. V. Goldberg, J. D. Hartline, N. Immorlica, and M. Sudan, *Derandomization of auctions*, Games and Economic Behavior (to appear) (2010).

[3] N. Alon, *Problems and results in extremal combinatorics - ii.*, Discrete Mathematics **308** (2008), no. 19, 4460–4472.

[4] N. Alon and J. Spencer, *The probabilistic method (3rd edition)*, Wiley, New-York, 2008.

[5] O. Ben-Zwi, I. Newman, and G. Wolfovitz, *A new derandomization of auctions*, Proceedings of algorithmic game theory, second international symposium (sagt), 2009, pp. 233–237.

[6] O. Ben-Zwi and G. Wolfovitz, *A hat trick*, Fun, 2010, pp. 37–40.

[7] O. Ben-Zwi and I. Newman, *Optimal bi-valued auctions*, CoRR **abs/1106.4677** (2011).

[8] S. Butler, M. T. Hajiaghayi, R. D. Kleinberg, and T. Leighton, *Hat guessing games*, SIAM J. Discrete Math. **22** (2008), no. 2, 592–605.

[9] B. Doer, *Integral approximation*, Habilitationsschrift. Christian-Albrechts-Universita"t zu Kiel (2005).

[10] T. Ebert, *Applications of recursive operators to randomness and complexity*, Ph.D. Thesis, University of California at Santa Barbara (1988).

[11] P. Erdős and L. Lovász, *Problems and results on 3-chromatic hypergraphs and some related questions*, Infinite and Finite Sets (A. Hajnal et al., eds.) (1975), 609–628.

[12] 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).

[13] U. Feige, A. Flaxman, J. D. Hartline, and R. D. Kleinberg, *On the competitive ratio of the random sampling auction*, Internet and network economics, first international workshop, wine, 2005, pp. 878–886.

[14] A. Fiat, A. V. Goldberg, J. D. Hartline, and A. R. Karlin, *Competitive generalized auctions*, The 34th acm symposium on theory of computing, stoc, 2002, pp. 72–81.

[15] A. V. Goldberg, J. D. Hartline, A. Karlin, M. Saks, and A. Wright, *Competitive auctions*, Games and Economic Behavior **55** (2006), no. 2, 242–269.

[16] A. V. Goldberg, J. D. Hartline, A. R. Karlin, and M. E. Saks, *A lower bound on the competitive ratio of truthful auctions*, Stacs, 21st annual symposium on theoretical aspects of computer science, 2004, pp. 644–655.

[17] A. V. Goldberg, J. D. Hartline, and A. Wright, *Competitive auctions and digital goods*, Proceedings of the 12th annual acm-siam symposium on discrete algorithms (soda), 2001, pp. 735–744.

[18] J. D. Hartline and T. Roughgarden, *Optimal mechanism design and money burning*, Proceedings of the 40th annual acm symposium on theory of computing, 2008, pp. 75–84.

[19] W. Hoeffding, *Probability inequalities for sums of bounded random variables*, Journal of the American Statistical Association **58** (1963), no. 301, 13–30.

[20] E. Koutsoupias and G. Pierrakos, *On the competitive ratio of online sampling auctions*, Internet and network economics - 6th international workshop, wine, 2010, pp. 327–338.

[21] A. Mehta and V. V. Vazirani, *Randomized truthful auctions of digital goods are randomizations over truthful auctions*, Acm conference on electronic commerce, 2004, pp. 120–124.

[22] J. A. Mirrlees, *An exploration in the theory of optimum income taxation*, The Review of Economic Studies **38** (1971), no. 2, 175–208.

[23] N. Nisan, T. Roughgarden, É. Tardos, and V. V. V. (editors), *Algorithmic game theory*, Cambridge University Press, 2007.

[24] I. Segal, *Optimal pricing mechanisms with unknown demand*, American Economic Review **93** (2003).

[25] D. D. Sleator and R. E. Tarjan, *Amortized efficiency of list update and paging rules*, Commun. ACM **28** (1985), no. 2, 202–208.

[26] J. von Neumann, *Zur theorie der gesellschaftsspiele*, Mathematische Annalen **100** (1928), 295–320.

[27] P. Winkler, *Games people don't play*, In D. Wolfe and T. Rogers editors. Puzzlers tribute. A feast for the mind A. K. Peters (2001).

[28] A. Yao, *Probabilistic computations: Toward a unified measure of complexity (extended abstract)*, 18th annual symposium on foundations of computer science, (focs), 1977, pp. 222–227.