

Derandomization of Auctions*

Gagan Aggarwal[†] Amos Fiat[‡] Andrew V. Goldberg[§]
Jason D. Hartline[¶] Nicole Immorlica[¶] Madhu Sudan^{||}

June 29, 2010

Abstract

We study the role of randomization in seller optimal (i.e., profit maximization) auctions. Bayesian optimal auctions (e.g., Myerson [19]) assume that the valuations of the agents are random draws from a distribution and prior-free optimal auctions either are randomized (e.g., Goldberg et al. [10]) or assume the valuations are randomized (e.g., Segal [22]). Is randomization fundamental to profit maximization in auctions? Our main result is a general approach to derandomize single-item multi-unit unit-demand auctions while approximately preserving their performance (i.e., revenue). Our general technique is constructive but not computationally tractable. We complement the general result with the explicit and computationally-simple derandomization of a particular auction. Our results are obtained through analogy to *hat puzzles* that are interesting in their own right.

1 Introduction

We consider the paradigmatic problem of a monopolist seller wishing to maximize her profit. What selling institution should she employ? The Bayesian optimal auction literature shows that when the buyers' valuations are drawn at random from a distribution the optimal auction rule is deterministic; it maps a bid profile to an outcome and payments. The prior-free optimal auction literature gives auction rules that achieve approximately optimal revenues for any deterministic, even worst-case, buyer valuations and these auction rules are randomized; they map a bid profile to a distribution over outcomes and payments. In contrast, for welfare maximization, the multi-unit Vickrey auction is symmetric, deterministic, and optimal in both Bayesian and prior-free settings (up to tie breaking).

*Based on a conference paper of the same name [1] and chapter 2.2 of a co-author's thesis [15].

[†]Google Inc., Mountain View, CA.

[‡]Tel Aviv University, Tel Aviv, Israel.

[§]Microsoft Research Silicon Valley Campus, Mountain View, CA

[¶]Northwestern University, Evanston, IL.

^{||}Massachusetts Institute of Technology, Cambridge, MA.

This leaves the question: *Is there something inherent about profit maximization that requires randomization?* This question is interesting as a mathematical curiosity. Additionally, deterministic auctions are often more pleasing as mechanisms: they do not require bidders to trust the randomization of the mechanism designer (i.e., deterministic mechanisms are easier to verify), and the revenue guarantees are entirely robust (i.e., the revenue guarantee is with probability one as opposed to in expectation). Motivated by such considerations, we show how to turn any randomized mechanism into a deterministic one while retaining, with probability one, a quarter of the expected revenue guarantee. We do so by replacing randomization with explicit price discrimination. We note that price discrimination is often undesirable in context of mechanism design as it encourages agents to try and misrepresent their identity. By randomly permuting bidder identities, our results can alternatively be interpreted as a way to take a randomized mechanism with an expected revenue guarantee and turn it into a robust and anonymous randomized mechanism whose revenue guarantee is a quarter of that of the original mechanism *with probability one*.

We focus on the fundamental setting where the seller has multiple identical units of a single item and wishes to sell them to a number of unit-demand buyers. Both the Bayesian (under standard assumptions) and prior-free auction literature reduce the case where the seller’s supply is limited and with a fixed marginal cost (or value for keeping units) to the zero-marginal-cost unlimited-supply setting.¹ This problem, recently referred to as *the digital good auction problem* (cf. [11]), is our focus.

Our approach to investigating the question of randomization in profit maximization is to start from the prior-free auctions for digital goods.² In this setting, there are a number of randomized mechanisms that are approximately optimal with respect to a natural revenue benchmark. The most well known auction in this class is the *random sampling optimal price* auction, which randomly partitions the buyers into two groups, computes the optimal single price on each set, and then offers that price to the opposite set. Goldberg et al. [11] showed that this auction achieves a constant approximation. *Does there exist a deterministic counterpart to this (or another) randomized auction procedure?*

In part to address this question, Goldberg et al. [11] define the *deterministic optimal price* auction, which considers each buyer in turn, computes the optimal single price for the set of all other buyers, and offers this buyer the computed price. They show that this auction can not be approximately optimal on all valuation profiles.³ To see this, consider an example with ten buyers with a high value of ten and ninety buyers with a low value of one. The deterministic optimal price auction offers all low-value buyers a high price of ten

¹This reduction is done as follows. Subtract the marginal cost from all buyer valuations. If there are only k -units available, run a k -unit Vickrey auction to get a sale price equal to the $k + 1$ st highest valuation. Run the optimal unlimited supply auction on the winners of the Vickrey auction. Modify the prices of all winners so that they are at least Vickrey price. Finally, add the marginal cost to the prices of all winners.

²An investigation on the role of randomization in Bayesian optimal auctions does not seem to be sensible as randomization is inherent in the Bayesian design and analysis framework.

³Segal [22] re-analyzed the deterministic optimal price auction to show that it is approximately optimal as the number of buyers grows when the valuations are drawn i.i.d. from any distribution with bounded support. Note, however, that such a result is not relevant to our study of derandomization.

and all high-value buyers a low price of one thereby achieving a total revenue of only ten. In contrast, with just a single fixed price of either one or ten, a seller could have easily obtained a revenue of 100. The deterministic optimal price auction indeed finds two suitable prices for the buyers, but offers these prices to precisely the wrong buyers. This is essentially due to a lack of ability to coordinate offer prices between buyers – i.e., to achieve high revenue, the price offered to a buyer depends in a delicate way on the price offered to other buyers; the deterministic optimal price auction is unable to achieve this balance. Goldberg et al. [11] show that this short coming is inherent in all anonymous (a.k.a., symmetric) deterministic auctions.

Our question thus focuses on the role of asymmetry in profit maximization and whether it is as powerful as randomization. It is easy to use randomization to coordinate offer prices in order to circumvent the problem illustrated above. The random sampling optimal price auction, for example, achieves this by computing an approximately optimal price through market research on a random sample of buyers and offering that same price to all remaining buyers. Randomization guarantees (in expectation) that the remaining buyers include about half the high-value buyers, and hence coordination is achieved. In a deterministic auction, one may hope to use asymmetry instead of randomization to coordinate offered prices.

As a warm up exercise we pose the following *hat puzzle* that addresses the question of using asymmetry to coordinate decentralized decisions. There are n players numbered one through n . Each player wears a hat of one of k colors. A player can see the color of all other players’ hats but not their own. In absence of any communication between players they must all simultaneously attempt to guess the color of their own hat. The players win if approximately (i.e., rounding down) a k th fraction of players of each hat color guess correctly. Of course there is a trivial randomized hat guessing solution to this puzzle: each player should guess their hat color by picking uniformly at random from the set of k colors. We will show how the randomized hat guessing solution can be derandomized using a general construction. This construction is central to our technique for turning any randomized digital good auction into a deterministic one while approximately preserving its revenue.

Goldberg and Hartline [9] consider the related question of how much randomness, measured in coins (i.e., random bits), is needed for an anonymous auction to be approximately optimal. While the deterministic impossibility discussed above suggests that at least one coin is necessary, Goldberg and Hartline show that one coin is sufficient. Essentially they propose two auctions and prove that for any fixed valuation profile one of the auctions must be *good*, i.e., a constant approximation. Therefore, it is enough to flip a single coin to choose between the two; the expected performance of the combined auction is also a constant approximation (albeit a factor of two worse). In order to turn this into a deterministic auction, one must use asymmetry to determine which buyers receive prices from which auction. This must be done in a coordinated way in order to satisfy the necessary property that enough of the high-valued buyers receive prices from the good auction.

To capture this kind of asymmetric coin flipping we propose a second hat puzzle. There are n players numbered one through n . Players wear a hats that are various shades of red (e.g., representing positive real numbers). A player can see the color of all other players’

hats but not their own. In absence of any communication between players they must each individually and simultaneously choose one of *heads* or *tails*. The players win if for all hat shades, approximately (i.e., rounding down) half the players with darker hats chose heads (and thus the other half chose tails). Again, of course, there is a simple randomized strategy that achieves the desired result in expectation: each player flips a fair coin and reports the outcome as their choice. We solve this puzzle by designing a deterministic asymmetric coin flipping strategy. The solution of this puzzle leads to an explicit derandomization of the single-coin auction of Goldberg and Hartline [9].

One final desiderata for auction design is computational tractability. How useful is an auction that could never be implemented because not even a computer could determine its outcome in a reasonable amount of time? Unfortunately, although we give a constructive proof that any auction can be derandomized with minimal loss in performance, it is not tractably constructive. In other words, its natural implementation would take computational time exponential in the number of buyers. We leave open the question of whether there exist faster implementations than the natural one. Fortunately, the explicit derandomization of the one-coin auction is both constructive and tractable (the required computation is simple). Thus, there are deterministic, tractable digital goods auctions that are approximately optimal.

Organization. In Section 2 we review the basics of prior-free optimal auction design. As a warm up exercise in Section 3 we adapt to our setting the result of Goldberg et al. [10] that no anonymous deterministic auction can be approximately optimal. In Section 4 we describe our first hat puzzle. In Section 5 we give our general construction for derandomizing any digital good auction. Section 6 gives our explicit derandomization of the one-coin auction from Goldberg and Hartline [9]. We conclude in Section 7 with some open questions.

2 Prior-Free Optimal Auctions and Digital Goods

In the digital good auction problem a seller has n identical units of an item for sale, and there are n unit-demand agents. Each agent i has a private value v_i for receiving a single unit of the item. Each agent's objective is to maximize their quasi-linear utility, i.e., the difference between their value and their payment. We attempt to design a single-round, sealed-bid auction that satisfies *ex post incentive compatibility* (IC), i.e, it is a dominant strategy for agents to bid their true values, *ex post individually rationality* (IR), i.e., by reporting their true value no agent receives a negative utility, and *no positive transfers* (NPT), i.e., the agent payments to the auctioneer are non-negative. In such auctions we assume agents bid their valuation (and equate values with bids).

A useful simplification of the IC mechanism design problem is attained through the following algorithmic characterization. Related formulations to the one given here have appeared in numerous places in recent literature (e.g., [2, 22, 8, 16]). To the best of our knowledge, the earliest dates back to the 1970s [17].

Definition 2.1 Given a valuation profile, $\mathbf{v} = (v_1, \dots, v_n)$, let $\mathbf{v}_{-i} \in \mathbb{R}_+^{n-1} \times n$ denote the masked valuation profile with v_i replaced with a ‘?’ , i.e.,

$$\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, ?, v_{i+1}, \dots, v_n).$$

Definition 2.2 (bid-independent pricing function) A bid-independent pricing function f is a function that maps masked valuation profiles to take-it-or-leave-it offer prices, i.e., $f : \mathbb{R}_+^{n-1} \times n \rightarrow \mathbb{R}_+$.

The following theorem follows directly from [17].

Theorem 2.1 A deterministic auction is IC, IR, and NPT if and only if there exists a bid-independent pricing function f such that on bids \mathbf{v} for each agent i :

1. let $t_i = f(\mathbf{v}_{-i})$;
2. if $t_i < b_i$, agent i wins at price t_i ;
3. if $t_i > b_i$, agent i loses; and
4. otherwise, ($t_i = b_i$) the auction can either accept the bid at price t_i or reject it.

A randomized auction is IC, IR, and NPT if and only if it is a randomization over deterministic IC, IR, and NPT auctions.

Our goal is to design an incentive-compatible deterministic auction that is optimal or approximately optimal in a prior-free setting. There is no prior-free mechanism that is optimal in an absolute sense (i.e., achieves a higher revenue than any other mechanism on all valuation profiles) so instead we consider relative optimality. In looking at relative optimality we compare the revenue for a given valuation profile to a revenue benchmark for that profile. A mechanism is good if it compares well to the benchmark across all valuation profiles. We adopt the following benchmark as proposed in [11] and further motivated in [10].

Definition 2.3 Given the valuation profile $\mathbf{v} = (v_1, \dots, v_n)$ of the agents, let $v_{(i)}$ denote the i th largest valuation and k be the index at which $v_{(i)}$ is maximized. Then v_k is the optimal price for \mathbf{v} and $kv_{(k)}$ is the optimal revenue, denoted $\text{OPT}(\mathbf{v})$.

This benchmark has a very compelling economic justification: suppose the agent valuations, \mathbf{v} , are drawn i.i.d. from some distribution, then the Bayesian optimal auction would post the monopoly price (i.e., the one that optimizes revenue for the distribution). $\text{OPT}(\mathbf{v})$ is also posts a price, but it posts the the optimal price in hind-sight for the specific valuation profile. Therefore it is simultaneously an upper bound on the revenue of every Bayesian optimal auction. A mechanism that approximates $\text{OPT}(\mathbf{v})$ point-wise for all \mathbf{v} also approximates the Bayesian optimal auction for any distribution. This observation was recently made concrete by Hartline and Roughgarden [13].

While most of the literature on prior-free optimal auctions looks at a slightly weaker benchmark and gives mechanisms that are constant multiplicative factor approximations; such an approximation is not possible for deterministic auctions (See Appendix A). Instead we will look to approximate the strong benchmark OPT up to a constant multiplicative factor and less some additive multiple of the highest bid, h . This analysis framework is identical to that of [3, 5].

Definition 2.4 *An auction \mathcal{A} is (β, γ) -approximately optimal if its expected profit on any valuation profile, $\mathbf{v} \in [1, h]^n$, is at least $\text{OPT}(\mathbf{v})/\beta - \gamma h$ for fixed constants β and γ . \mathcal{A} 's approximation ratio is β and its additive loss is γh .*

While we approach the problem of derandomizing auctions with the objective of obtaining an (β, γ) -approximately optimal auction, our main result is a generic derandomization that takes any randomized auction and gives a deterministic one that guarantees a fraction of its revenue always. Therefore, if the original randomized auction was good by some expected revenue criteria, then the resulting auction is approximately as good with respect to the same criteria.

3 Anonymous Deterministic Auctions

We start by considering the design of anonymous, a.k.a., symmetric, deterministic auctions. In these auctions the outcome for a permutation of \mathbf{v} is the permutation of the outcome for \mathbf{v} . Goldberg et al. [10] showed that no anonymous deterministic auction can approximate the optimal auction. For completeness we reiterate a version the proof here, adapted to our approximation framework.

Theorem 3.1 [10] *For any anonymous deterministic auction, \mathcal{A} , and any constants β and γ , there exists a valuation profile $\mathbf{v} \in [1, h]^n$ such that \mathcal{A} 's revenue on \mathbf{v} is less than $\text{OPT}(\mathbf{v})/\beta - \gamma h$*

Proof: We will restrict our attention to $\mathbf{v} \in \{1, h\}^n$. Let $n_h(\mathbf{v})$ denote the number of agents with value h in \mathbf{v} . $\mathcal{A}(\mathbf{v})$ denote the auction's revenue on \mathbf{v} .

1. $\text{OPT}(\mathbf{v}) = \max(n, h \cdot n_h(\mathbf{v}))$.

A revenue of n can be achieved with the single price of 1. A revenue of $h \cdot n_h(\mathbf{v})$ can be achieved with a single price of h .

2. There exists \mathbf{v}^* for which $\mathcal{A}(\mathbf{v}^*) \leq n_h(\mathbf{v}^*)$.

Let f be the bid-independent function for \mathcal{A} . First, by anonymity, for a fixed n , f is a function only of the number of h 's in \mathbf{v}_{-i} . Thus, we adopt short-hand $f(k)$ to denote the value of $f(\mathbf{v})$ when $n_h(\mathbf{v}) = k$.

If $f(0) > 1$ then the auction revenue on the all-ones profile is zero. If $f(n - 1) \leq 1$ then the auction revenue on the all-high profile is at most $n = n_h(\mathbf{v})$. Otherwise,

$f(0) \leq 1$ and $f(n-1) > 1$. In this case there is some value k^* such that $f(k^*) > 1$ and $f(k^* - 1) \leq 1$; for instance, it can be identified by increasing k from zero (where $f(0) \leq 1$) until the first moment at which $f(k) > 1$ (which must happen by $k = n-1$ as $f(n-1) > 1$). Consider the n -agent valuation profile \mathbf{v}^* with $n_h(\mathbf{v}^*) = k^*$. The auction offers price $f(k^*) > 1$ to all low-valued agents; they decline. The auction then offers price $f(k^* - 1) \leq 1$ to all high-valued agents; they accept. Thus, $\mathcal{A}(\mathbf{v}^*) \leq k^* = n_h(\mathbf{v}^*)$.

3. There exists an n and h for which $\mathcal{A}(\mathbf{v}^*) < \text{OPT}(\mathbf{v}^*)/\beta - \gamma h$.

First note that for any n and h , one can always find valuations \mathbf{v}^* and the corresponding number of high values k^* satisfying the property in (2). For any such setting of parameters, (1) implies $\text{OPT}(\mathbf{v}^*) \geq (hk^* + n)/2$ and (2) implies $\mathcal{A}(\mathbf{v}^*) \leq k^*$. Thus, it is enough to argue that for some n and h , the corresponding \mathbf{v}^* and k^* satisfy $k^* < (hk^* + n)/(2\beta) - \gamma h$. This can be ensured by a careful choice of n and h (e.g., observing that since $k^* \leq n$ it is sufficient to pick $n/(2\gamma\beta) > h > 2\beta$).

This completes the proof. □

4 A Hat Puzzle

A *hat puzzle* is a (usually, simultaneous-move, single-round) game in which players attempt to coordinate their actions to achieve a desired objective. The players in the game have asymmetric and incomplete information. Each player is wearing a hat with a particular color or shade; each player knows the colors of all players' hats except for their own. Often, the objective would be trivial to attain if the players knew their own hat color. The challenge is for agents to coordinate their strategies (functions from the $(n-1)$ -tuple of other players' hats to an action) to meet the desired objective. Such an objective might be for some of the players to guess their hat color correctly.

Hat puzzles of various forms have been contemplated in the mathematics community, partially due to their connections to coding theory, discrepancy problems, and auto-reducibility of random sequences; and often simply because they make interesting brain-teasers [6, 7, 20, 23]. Notice that the strategies in a hat puzzle have the same structure as the pricing function in the bid-independent representation of an incentive-compatible auction. This connection motivates us to consider several hat puzzles related to auction design.

Hat Puzzle 1 (balanced k -color) *In the balanced k -color hat puzzle the hats are one of k distinct colors. The action space is the set of colors. We interpret the action of a player as a guess of their own hat color. The objective is to have at least a k th fraction (rounded down) of the players with each hat color guess correctly.*

Example 1 *Suppose there are $n = 5$ people, Jackie, Tito, Jermaine, Marlon and Michael, and $k = 2$ hat colors, red and blue. If Jackie and Tito are wearing red hats, and Jermaine, Marlon, and Michael are wearing blue hats, then they would win if at least one red-hatted player guesses red and one blue-hatted player guesses blue. In other words, the players win*




Figure 1: Hat puzzle flow graph.

if Jackie or Tito guesses red and at least one of the remaining three; Jermaine, Marlon, or Michael; guesses blue.

The balanced k -color hat puzzle is easy to solve in expectation if the players have access to randomization. Each player simply guesses each of the colors with equal probability $1/k$. We call this the *uniform random strategy*. As our goal is a general approach to derandomizing auctions, we first consider derandomizing this random hat guessing strategy. Notice that standard algorithmic derandomization techniques, i.e., trying all possible coin flips and selecting the best, can not be implemented within the rules of the puzzle. Instead, we give a flow-based technique that uses the asymmetry of the players information in place of randomness to solve the puzzle deterministically.⁴

Solution. First we define a bipartite graph representing the game (see Figure 1). The nodes are defined as follows. Each node on the left-hand side represents a possible viewpoint of player i . Let $\mathbf{c} = (c_1, \dots, c_n)$ represent the profile of players' colors and for any index i , \mathbf{c}_{-i} the masked color profile (Definition 2.1). Note that \mathbf{c}_{-i} is precisely the view of player i . For each of the nk^{n-1} possible values of \mathbf{c}_{-i} , we have a vertex $v_{\mathbf{c}_{-i}}$ on the left-hand side. Each node on the right-hand side represents a possible scenario (a hat color for each player) and a corresponding guess of some player. Let χ be one of the k colors. Then a scenario and corresponding guess is a pair (χ, \mathbf{c}) (note this pair does not specify which player guesses χ). For each of the k^{n+1} possible values of the pair (χ, \mathbf{c}) , we have a vertex $v_{\chi, \mathbf{c}}$ on the right-hand side. We also include a source vertex s and sink vertex t . The arc set is defined as follows. We place an arc from the source s to each vertex on the left-hand side, and another arc from each vertex on the right-hand side to t . We also add an arc between $v_{\mathbf{c}_{-i}}$ and $v_{\chi, \mathbf{c}}$ signifying

⁴Concurrently and independently from our work Feige [7] developed similar flow-based algorithms for related hat puzzles.

that the hidden hat at position i in c_{-i} has color c_i . Notice that the in-degree of a vertex $v_{\chi, \mathbf{c}}$ is precisely the number $n_\chi(\mathbf{c})$ 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 . The structure of the graph is sketched in Figure 1 (all arcs are directed from left to right in this figure).

Next, we set upper and lower capacities on the arcs as follows. For each χ and \mathbf{c} , we lower bound the flow of the arc $(v_{\chi, \mathbf{c}}, t)$ to $\lfloor n_\chi(\mathbf{c})/k \rfloor$ (recall $n_\chi(\mathbf{c})$ represents the number of hats in \mathbf{c} that are colored χ). This represents the objective that at least $n_\chi(\mathbf{c})$ players should guess χ in scenario \mathbf{c} . For every other arc, we upper bound its flow by 1. This represents the requirement that each player can submit at most one guess.

We can interpret the randomized hat-guessing algorithm as a feasible flow on this graph. 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 total probability of 1 to spend on guessing a color 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 (this is possible since each $v_{\mathbf{c}_{-i}}$ has k outgoing arcs). Now notice that the incoming flow to $v_{\chi, \mathbf{c}}$ is precisely $1/k$ times the number $n_\chi(\mathbf{c})$ of hats in \mathbf{c} from color class χ . Thus by sending all of this flow on the outgoing arc to t , we satisfy all capacities. The flow is sketched in Figure 1 (the labels on the arcs represent the amount of flow on that arc).

Similarly, we can interpret an integral flow in the this graph as a deterministic hat-guessing algorithm: each player observes a vector \mathbf{c}_{-i} of hat colors and identifies the corresponding node $v_{\mathbf{c}_{-i}}$ of the graph. If the integral flow sends flow from this node to $v_{\chi, \mathbf{c}}$, then the player guesses χ as his hat color (if the integral flow does not send flow through this node, then we let the player guess arbitrarily). We now analyze the performance of such an algorithm on \mathbf{c} . Given $n_\chi(\mathbf{c})$ hats with color χ in \mathbf{c} , the lower bound on the capacity of the outgoing arc from $v_{\chi, \mathbf{c}}$ to t is $\lfloor n_\chi(\mathbf{c})/k \rfloor$. Therefore, 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. For each such arc $(v_{\mathbf{c}_{-i}}, v_{\chi, \mathbf{c}})$, player i correctly guesses χ in the game setting \mathbf{c} . Thus, players guess the correct color χ for $\lfloor n_\chi(\mathbf{c})/k \rfloor$ positions out of a total of $n_\chi(\mathbf{c})$ such positions. This holds true for all colors χ , and so the players have solved the puzzle.

A classic result on integrality of network flows (see, for example, Hoffman's circulation theorem in the book by Schrijver [21, Theorem 11.2]) states that in a graph with integral capacities, if there is a feasible fractional flow, then there is a feasible integral flow. Therefore, the existence of the randomized hat-guessing algorithm implies the existence of a deterministic one. Unfortunately, our constructed flow graph is exponentially big so it will generally take exponential computational time find an integral flow in it; therefore, our reduction is not tractably constructive.

5 General Auction Derandomization

By analogy to the hat-guessing technique of Section 4, we can show that *any* randomized auction has a deterministic counterpart that achieves approximately the same profit.

Theorem 5.1 *Corresponding to any single-round sealed-bid truthful auction \mathcal{A} with expected profit $\mathbf{E}[\mathcal{A}(\mathbf{v})]$ on input bid vector \mathbf{v} , there is a deterministic truthful auction \mathcal{A}' whose profit on any input bid vector \mathbf{v} is at least $\mathbf{E}[\mathcal{A}(\mathbf{v})] / 4 - 2h$, where $h = \max_i b_i$ is the highest bid.*

The constructive proof of Theorem 5.1 converts any (randomized) auction to a special type of (randomized) auction, called a *guessing auction*, and then uses a flow-based construction similar to that in Section 4 to derandomize the guessing auction. The conversion to a guessing auction loses a factor of four relative to the randomized auction’s performance and the derandomization of the guessing auction has an additional additive loss of $2h$. The proof of the theorem follows directly from Lemmas 5.2 and 5.3, below.

5.1 Guessing Auctions

The flow-based construction for balanced k -color in Section 4 tries to reconstruct c_i from the vector \mathbf{c}_{-i} . In order to use this construction in the auction setting, we would like to draw an analogy between a player’s hat color and an agent’s bid, and between a player’s guess for his hat color and an agent’s price. However, an auction gets revenue not only when it offers a price equal to the bid value, but also when it offers a price below a bid value. In order to resolve this discrepancy, we define the notion of a *guessing auction* that uses only powers of two as prices. Payments from agents with values at least twice their offered price are collected but ignored (i.e., they do not contribute to the guessing auction’s revenue).

Definition 5.1 *The revenue of a guessing auction that offers price t_i to agent i is the sum of the t_i for which $\log t_i = \lfloor \log b_i \rfloor$.*

It is possible to convert any incentive-compatible auction into a truthful guessing auction while only losing a factor of four from the revenue. Denote the revenue of an incentive-compatible auction \mathcal{A} on valuation profile \mathbf{v} as $\mathcal{A}(\mathbf{v})$. This revenue is given by the sum of the prices charged to the winning agents. For a randomized auction $\mathcal{A}(\mathbf{v})$ is a random variable. The revenue of a guessing auction is as defined above.

Definition 5.2 (Guessing Auction $\mathcal{G}_{\mathcal{A}}$) *The guessing auction, $\mathcal{G}_{\mathcal{A}}$, for auction \mathcal{A} works as follows. Simulate \mathcal{A} on input bid vector \mathbf{v} . Suppose \mathcal{A} offers agent i price t_i and let 2^k be the largest power of two less than t_i . Then for integer $j \geq 0$, $\mathcal{G}_{\mathcal{A}}$ offers agent i price 2^{k+j} with probability 2^{-j-1} .*

Lemma 5.2 *For any IC, IR, NPT auction \mathcal{A} with expected revenue $\mathbf{E}[\mathcal{A}(\mathbf{v})]$ on valuation profile \mathbf{v} , the corresponding guessing auction $\mathcal{G}_{\mathcal{A}}$ is IC, IR, and NPT and attains an expected revenue of at least $\mathbf{E}[\mathcal{A}(\mathbf{v})] / 4$ on \mathbf{v} .*

Proof: It is clear that the existence of a bid-independent pricing function for \mathcal{A} implies one exists for $\mathcal{G}_{\mathcal{A}}$. Therefore, by Theorem 2.1, $\mathcal{G}_{\mathcal{A}}$ is IC, IR, and NPT. We now bound the expected revenue (in the sense of Definition 5.1) of $\mathcal{G}_{\mathcal{A}}$. Consider an agent i who is offered price $q_i \leq v_i$ by \mathcal{A} . The revenue of \mathcal{A} from i is q_i . We show that the revenue (in the sense

of Definition 5.1) of $\mathcal{G}_{\mathcal{A}}$ from i is at least $q_i/4$. Let k be such that $2^k \leq q_i < 2^{k+1}$. Suppose agent i 's value v_i is in the interval $[2^{k+j}, 2^{k+j+1})$. Then the probability that $\mathcal{G}_{\mathcal{A}}$ offers i price 2^{k+j} is 2^{-j-1} , and the revenue from this offer is 2^{k+j} . Thus, the expected revenue attained from agent i is $2^{j-1} \cdot 2^{k+j} = 2^{k-1}$. As $q_i < 2^{k+1}$ by assumption, the revenue attained from i is at least a quarter of his price q_i in \mathcal{A} . This holds for all i completing the proof. \square

We note that if auction \mathcal{A} only uses prices that are powers of two, then the revenue of the corresponding guessing auction $\mathcal{G}_{\mathcal{A}}$ is actually within a factor of two of the revenue of \mathcal{A} instead of a factor of four.

5.2 The Flow Construction

We now show how to derandomize any guessing auction $\mathcal{G}_{\mathcal{A}}$. Our derandomization draws an analogy between bids and hat colors to deterministically compute prices for the guessing auction using the flow-based technique of Section 4.

Lemma 5.3 *For any IC, IR, NPT guessing auction \mathcal{G} , there is an IC, IR, and NPT deterministic auction \mathcal{A} with (deterministic) revenue $\mathcal{A}(\mathbf{v})$ at least $\mathbf{E}[\mathcal{G}(\mathbf{v})] - 2h$ on $\mathbf{v} \in [1, h]^n$.*

Proof: First, round all valuations down to the nearest power of two. We draw an analogy between the k colors in the balanced k -color puzzle and the $\log h$ powers of two that are the possible (rounded) valuations. Set up a flow construction identical to that for the balanced k -color puzzle, except that the fractional flow on an arc from $v_{\mathbf{v}_{-i}}$ to $v_{2^j, \mathbf{v}}$ is the probability that \mathcal{G} on seeing \mathbf{v}_{-i} offers price 2^j to i . Furthermore, the flow from $u_{2^j, \mathbf{v}}$ to t is the expected number of times \mathcal{G} offers one of the agents with (rounded) bid 2^j a price equal to 2^j . We represent this quantity by $E_j(\mathbf{v})$. We then set the capacities as before: we require flow on an arc between $u_{2^j, \mathbf{v}}$ and t to be at least $\lfloor E_j(\mathbf{v}) \rfloor$ and all other arc flows to be at most 1. Once again, the above fractional flow implies the existence of an integer-valued flow [21, Theorem 11.2], and this integer-valued flow corresponds to an auction making a deterministic bid-independent offer upon seeing \mathbf{v}_{-i} . The flow out of $u_{2^j, \mathbf{v}}$ is precisely the number of indices i such that the auction, upon seeing \mathbf{v}_{-i} , correctly guesses 2^j ; since the flow is feasible, the flow out of $u_{2^j, \mathbf{v}}$ is at least $\lfloor E_j(\mathbf{v}) \rfloor$. Thus, considering a valuation profile \mathbf{v} where the expected revenue of \mathcal{G} is $\mathbf{E}[\mathcal{G}(\mathbf{v})] = \sum_{j=1}^{\log h} 2^j E_j(\mathbf{v})$, the deterministic auction obtains $\sum_{j=1}^{\log h} 2^j \lfloor E_j(\mathbf{v}) \rfloor \geq \sum_{j=1}^{\log h} [2^j E_j(\mathbf{v}) - 2^j] \geq \mathbf{E}[\mathcal{G}(\mathbf{v})] - 2h$. \square

As a corollary of Theorem 5.1, known approximately-optimal randomized auctions [8, 9, 11, 12] imply the existence of approximately-optimal deterministic auctions. Using a the best known randomized auction from Hartline and McGrew [12], we obtain the following result (which we improve in the next section).

Theorem 5.4 *There is a deterministic auction with revenue at least $\text{OPT}(\mathbf{v})/13 - 2h$ for any valuation profile \mathbf{v} .*

Proof: The theorem follows from Theorem 5.1 and the Hartline-McGrew auction that achieves a 3.25-approximation to $\text{OPT}(\mathbf{v})$ on valuation profiles $\mathbf{v} \in [1, h]^n$ with $\text{OPT}(\mathbf{v}) > h$.

□

In the above construction we assumed that the range of valuations $[1, h]$ is known. This assumption is not necessary. When considering \mathbf{v}_{-i} , we can compute h , which is the maximum valuation in \mathbf{v}_{-i} scaled such that the minimum valuation is one. This calculation is based on \mathbf{v}_{-i} and therefore bid-independent. We then perform the above construction separately for each agent based on the calculated h . As our calculation of h is correct for all but the minimum and maximum valued agents, these agents will be offered prices according to the above construction and so the profit guarantee applies for these agents. Assuming the worst, that is the auction fails to get any profit from the minimum and maximum valued agents, we only lose an additional additive $h + 1$.

6 A Polynomial-time Deterministic Auction

The flow construction used to derandomize auctions in Section 5 has exponential size. Therefore the direct derandomization from Section 5 is not computationally tractable. In this section, we describe an approximately optimal deterministic asymmetric auction, the outcome of which can be computed in polynomial time. In particular, our deterministic auction guarantees a revenue which is at least $\text{OPT}(\mathbf{v})/4 - h$ on all valuation profiles $\mathbf{v} \in [1, h]^n$.

There are three key ingredients in this auction: (a) an incentive-compatible auction, called a *profit extractor*, for extracting a given feasible target revenue from a set of agents; (b) a pair of bid-independent functions, called *consensus estimates*, with the property that at least one of them is both a consensus for all agents and an estimate of the optimal revenue OPT ; and (c) a deterministic coin-flipping algorithm. To see how these pieces fit together, first suppose we knew an estimate of the optimal revenue $R \leq \text{OPT}$. Could we then design an incentive-compatible auction to recover revenue R ? The goal of the profit extractor is to do just that: given bids \mathbf{v} , a profit extractor extracts a target revenue R from some subset of the agents. Although this mechanism is incentive compatible, deterministic, and extracts the target revenue, it requires an estimate of the the optimal revenue as input. We can not hope to estimate the optimal revenue bid-independently; rather we compute n bid-independent estimates of the optimal revenue, one for each agent. If these estimates are coordinated appropriately (namely, if enough high-valued agents compute the same estimate), then we can use these estimates as inputs to our profit extractor and generate sufficient revenue. As we will show, combining the consensus estimates with the deterministic coin-flipping algorithm does exactly this.

The profit extractor and consensus estimates were used previously by Goldberg and Hartline [9] along with a single random coin flip to get an approximately optimal randomized auction. Even though their auction uses just one bit of randomness, it is difficult to derandomize using standard algorithmic techniques since such techniques tend not to be bid-independent. For example, one might consider trying to derandomize the auction of Goldberg and Hartline [9] by running it twice – once with “heads” and once with “tails” – and outputting the higher-revenue solution. However, this clearly can not be represented bid-independently and therefore is not IC. Instead, as the final ingredient of our construc-

tion, we design a deterministic coin flip which can be calculated bid-independently and use it to derandomize the auction of Goldberg and Hartline [9].

6.1 Profit Extractor

The profit extractor we present here is a special case of a general cost-sharing scheme due to Moulin [18].

Definition 6.1 (ProfitExtract_R) *The profit extractor, ProfitExtract_R, works as follows. Given bids \mathbf{v} , find the largest k such that the highest k agents can equally share the cost R (that is, each of their valuations is at least R/k). These agents are the winners and the rest are losers. Charge each of the winners R/k and the losers nothing. If no such k exists, then all agents are losers and pay nothing.*

As we will base our deterministic auction on this mechanism, it is important to note that it is incentive compatible and actually extracts revenue R whenever $R \leq \text{OPT}(\mathbf{v})$.

Lemma 6.1 [18] *ProfitExtract_R is incentive compatible.*

Proof: We define a bid-independent pricing function for ProfitExtract_R. Let the bid-independent function $\text{pe}_{i,R}(\mathbf{v}_{-i})$ equal $\frac{R}{l+1}$ where l is the largest number such that the highest l agents in \mathbf{v}_{-i} all have a value at least $\frac{R}{l+1}$. If no such l exists, let $\text{pe}_{i,R}(\mathbf{v}_{-i})$ be ∞ , a number larger than any valuation. Let k be the number of winners in ProfitExtract_R. Then, by definition, each of these k winners has value at least $\frac{R}{k}$ while each of the losers in ProfitExtract_R has value strictly less than $\frac{R}{k+1}$. Thus, for a winner i , $\text{pe}_{i,R}(\mathbf{v}_{-i})$ equals $\frac{R}{k}$, i 's price in ProfitExtract_R. For a loser j , $\text{pe}_{i,R}(\mathbf{v}_{-j})$ equals $\frac{R}{k+1}$, implying that j is a loser in the bid-independent implementation as well. \square

Lemma 6.2 *If $R \leq \text{OPT}(\mathbf{v})$, ProfitExtract_R(\mathbf{v}) = R ; otherwise it has no winners and no revenue.*

Proof: Let the number of items sold by OPT be k , i.e., $\text{OPT} = kv_{(k)}$. Since $\text{OPT} \geq R$, k satisfies $kv_{(k)} \geq R$. Thus, ProfitExtract_R can find a set of winners who can equally share R , i.e., these k highest-valued agents can; therefore, it does. Otherwise, if $R > \text{OPT}$, then by the definition of OPT this revenue can not be shared equally. Hence ProfitExtract_R(\mathbf{v}) has no winners and thus no revenue. \square

6.2 Consensus Estimator

The goal of the consensus estimator is to compute an estimate of the optimal revenue for each agent bid-independently. A *pair of consensus estimators* is a pair of functions, r_0 and r_1 , having the following properties:

1. For any real number V , there exists an $r \in \{r_0, r_1\}$ such that for all $v \in [V/2, V]$, $r(v) = r(V)$. This r is called a *consensus* on V .

2. For any real number V and $r \in \{r_0, r_1\}$ that is a consensus on V , $r(V) \in [V/2, V]$. In this case, r is said to *estimate* V .

It is easy to see that the following functions form such a pair of consensus estimators [9].

$$\begin{aligned} r_0(v) &= 2v \text{ rounded down to the nearest } \textit{even} \text{ power of two.} \\ r_1(v) &= 2v \text{ rounded down to the nearest } \textit{odd} \text{ power of two.} \end{aligned}$$

We will apply these consensus estimators to the value $\text{OPT}(\mathbf{v}_{-i})$ in order to obtain a consensus on an approximate value for $\text{OPT}(\mathbf{v})$.

6.3 Deterministic Coin Flipping

Our deterministic coin-flipping algorithm is best described via an analogy to another hat puzzle.

Hat Puzzle 2 (deterministic coin flip) *In the deterministic coin flip hat puzzle the hats are a distinct shade of red. The action space consists of two actions, heads and tails (corresponding to sides of the coin). The objective is that among players with the m darkest hats at least $\lfloor m/2 \rfloor$ should chose each action, for all m .*

Again there is a natural randomized strategy that achieves our desired objective in expectation – each player simply chooses heads or tails uniformly at random. The algorithm we are about to present achieves the property deterministically. In fact, our solution satisfies the following stronger property: lining the players up from darkest hat to lightest hat, the sequence of coin choices alternates. We call such a set of choices *perfectly alternating*. Our algorithm is based on the notion of the *sign of a permutation*.

Definition 6.2 *Given a vector of n numbers, \mathbf{c} , the sign of \mathbf{c} is the parity of the number of transpositions of adjacent numbers it takes to sort \mathbf{c} , notated $\text{sign}(\mathbf{c})$.*

Theorem 6.3 *(E.g., [14]). The parity of the number of transpositions needed to sort a vector is unique.*

The *deterministic coin-flipping algorithm*, ϕ , works as follows. Each player has an identity i . Given \mathbf{c}_{-i} as the shades of the hats that player i sees, player i computes his coin choice, $\phi(\mathbf{c}_{-i})$, by imagining that his own hat is the darkest shade, denoted ∞ . As his coin flip, he chooses the sign of his imagined vector of hat colors, denoted $(\mathbf{c}_{-i}, \infty)$, whose i 'th entry is ∞ and j 'th entry for $j \neq i$ is c_j . To prove that this algorithm solves the puzzle, we will use the notion of the *rank* of a player.

Definition 6.3 *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 .*

Lemma 6.4 *The deterministic coin-flipping algorithm, ϕ , is perfectly alternating with the shades of the hats' colors.*

Proof: This result is implied by the fact that

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

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 ∞ to the front of the vector. As the parity of the ranks alternate, this implies the lemma. \square

In this solution to the deterministic coin-flipping problem, each player can compute his own coin by simply executing ϕ ; however, no player can compute the coin of anyone else as he does not know his own hat shade. Clearly, each player can compute his 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 $O(n \log n)$ time.

6.4 Deterministic Auction

Finally, we have developed the tools necessary to describe our auction, the *Deterministic Consensus Estimate Profit Extraction* (DCEPE) auction. This auction is built from the three components discussed above – the profit extractor, the pair of consensus estimators, and the deterministic coin-flipping algorithm. The auction uses the deterministic coin-flipping algorithm to pick a consensus estimator and corresponding estimate for each agent – namely, an agent i with a deterministic coin flip of heads (tails) calculates the optimum revenue R'_i on \mathbf{v}_{-i} and then estimates OPT to be $2R'_i$ rounded down to the nearest even (odd) power of two. Agent i is then offered a price equal to the price computed for i by the profit extractor on \mathbf{v} with target estimate, R_i .

Definition 6.4 (DCEPE) *The Deterministic Consensus Estimate Profit Extraction (DCEPE) auction is implemented by the following bid-independent function:*

$$f(\mathbf{v}_{-i}) = pe_{i,R_i}(\mathbf{v}_{-i}),$$

where $R_i = r_{\phi(\mathbf{v}_{-i})}(\text{OPT}(\mathbf{v}_{-i}))$ and $pe_{i,R}$ is the bid-independent function for the mechanism ProfitExtract_R defined in the proof of Lemma 6.1.

DCEPE is bid-independent and therefore truthful. We now show that DCEPE is approximately optimal. Our proof is similar to the proof of the revenue of the “CORE” auction in Goldberg and Hartline [9].

Theorem 6.5 *The profit of DCEPE is at least $\text{OPT}/4 - h$.*

Proof: Let $\text{OPT} = \text{OPT}(\mathbf{v})$ be the revenue from the optimal single price sale. 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, for every i , we have $\text{OPT}/2 \leq \text{OPT}(\mathbf{v}_{-i}) \leq \text{OPT}$.

Since r_1 and r_0 are a pair of consensus estimates, one of them is a consensus on OPT. Suppose, without loss of generality, that it is r_0 . Then $r_0(\text{OPT}(\mathbf{v}_{-i})) = r_0(\text{OPT})$ for all i . Now consider the following thought experiment. Suppose we had chosen consensus function r_0 for all i and so $R_i = r_0(\text{OPT}) \in [\text{OPT}/2, \text{OPT}]$ for all i . In this case, our auction is equivalent to the profit extraction mechanism on input $r_0(\text{OPT})$. Let p be the price charged to the k winning agents in this thought experiment. Note that by Lemma 6.2, as $r_0(\text{OPT}) \leq \text{OPT}$, the total profit is $pk = r_0(\text{OPT}) \in [\text{OPT}/2, \text{OPT}]$. In reality, in the deterministic coin-flipping procedure, at least $k/2 - 1$ of these k agents computed $\phi(\mathbf{v}_{-i}) = 0$ and thus these agents all pay p , exactly as they would have in the thought experiment. The total profit accounted for is $pk/2 - p \geq \text{OPT}/4 - h$, which proves the theorem. \square

7 Discussion

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 [11]. Our general derandomization technique has the benefit of being universal – it can be applied to *any* multi-unit auction for unit-demand agents to get a deterministic auction with approximately the same revenue guarantees. This is quite interesting from a theoretical standpoint, as it is the first known incentive-preserving derandomization technique.

Unfortunately, our general technique will not yield asymptotically optimal results. For example, when the number of winners is large, there are randomized auctions which generate revenue arbitrarily close to that of the benchmark. However, our general derandomization technique fails to preserve this asymptotic behavior. In our analysis, we lose a factor of four. We lose a factor of two when we round the bids to the nearest power of two, and an additional factor of two through the reduction to guessing auctions (which only get “credit” for “correct” guesses). A more direct derandomization technique might allow us to find asymptotically optimal deterministic auctions.

Alternatively, one could try to directly develop new asymptotically optimal deterministic auctions. Recall that our explicit derandomization of the one-coin auction is also not asymptotically optimal. A factor of two is lost from profit extracting an estimate that can be as far as a factor of two from OPT; a second factor of two is lost from the fact that the deterministic coin flip results in only half of the agents using the estimate for which there is a consensus, i.e., the good one. Both of these losses could be reduced if we could deterministically generate more random bits. This approach suggests the following generalization of the deterministic coin-flipping hat puzzle; we do not know if a solution exists for any $k \geq 3$.

Hat Puzzle 3 (deterministic dice) *In the deterministic dice hat puzzle the hats are a distinct shade of red. The action space consists of k actions (corresponding to sides of a k -sided die). The objective is that among players with the m darkest hats at least $\lfloor m/k \rfloor$ should chose each action, for all m .*

Recent followup work to this paper followed a different approach to the one discussed above to show that there are in fact deterministic, asymmetric auctions that are asymp-

totically optimal [4]. This work gives an existence proof, via that Lovász Local Lemma, that that shows that for any randomized auction that achieves an good expected revenue on some profile of valuations, there is a deterministic, asymmetric auction that achieves this revenue deterministically less an additive loss of $O(h\sqrt{n \log hn})$ (for integer valued valuation profiles with maximum value h). This approach can be applied to the (asymptotically optimal) random sampling optimal price auction, discussed in the introduction, to give an asymptotically optimal deterministic auction.

Finally, we note that our derandomization technique is computationally intractable as stated. Computationally tractable solutions to the balanced k -color hats puzzle would imply a computationally tractable derandomization technique. For the special case of two hat colors, say red and blue, the puzzle can be solved in polynomial time by reduction to the deterministic coin-flipping puzzle. Simply define a total ordering on players where player i with hat color χ_i is less than player j with hat color χ_j if and only if χ_i is red and χ_j is blue or $\chi_i = \chi_j$ and $i < j$. Then use the deterministic coin-flipping algorithm to generate a guess of *heads* or *tails* for each player and equate *heads* with *red* and *tails* with *blue*. The resulting algorithm, modulo rounding, guesses half of the red hats and blue hats correctly. We do not know how to solve this problem in polynomial time for general k .

References

- [1] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of auctions. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, 2005. 619–625.
- [2] A. Archer and E. Tardos. Truthful mechanisms for one-parameter agents. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2001.
- [3] M.-F. Balcan, A. Blum, J. Hartline, and Y. Mansour. Mechanism Design via Machine Learning. In *46th Annual IEEE Symposium on Foundations of Computer Science*, pages 605 – 614, 2005.
- [4] O. Ben-Zwi, I. Newman, and G. Wolfvitz. A new derandomization of auctions. In *SAGT '09: Proceedings of the 2nd International Symposium on Algorithmic Game Theory*, pages 233–237, 2009.
- [5] A. Blum and J. Hartline. Near-Optimal Online Auctions. In *Proc. 16th Symp. on Discrete Alg.* ACM/SIAM, 2005.
- [6] T. Ebert, W. Merkle, and H. Vollmer. On the autoreducibility of random sequences. *SIAM Journal on Computing*, 32:1542–1569, 2003.
- [7] 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.

- [8] A. Fiat, A. Goldberg, J. Hartline, and A. Karlin. Competitive Generalized Auctions. In *Proceedings of 34th ACM Symposium on the Theory of Computing*, pages 72–81, 2002.
- [9] A. Goldberg and J. Hartline. Competitiveness via Concensus. In *Proceedings of 14th ACM Symposium on Discrete Algorithms*, pages 215–222, 2003.
- [10] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions and digital goods. *Games and Economic Behavior*, 55:242–269, May 2006.
- [11] A.V. Goldberg, J.D. Hartline, and A. Wright. Competitive auctions and digital goods. *Proceedings of 12th ACM Symposium on Discrete Algorithms*, pages 735–744, 2001.
- [12] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proceedings of 6th ACM Conference on Electronic Commerce*, pages 175–182, 2005.
- [13] J. Hartline and T. Roughgarden. Optimal Mechanism Design and Money Burning. In *Proceedings of the 40th annual ACM Symposium on Theory of Computing*, 2008.
- [14] T.W. Hungerford. *Algebra*. Springer, 1997.
- [15] N. Immorlica. *Computing with Strategic Agents*. PhD thesis, Massachusetts Institute of Technology, 2005.
- [16] 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.
- [17] J. Mirrlees. An Exploration into the Theory of Optimal Income Taxation. *Review of Economics Studies*, 38:175–208, April 1971.
- [18] H. Moulin. Incremental cost sharing: Characterization by coalition strategy-proofness. *Social Choice and Welfare*, 16:279–320, 1999.
- [19] R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [20] S. Robinson. Why mathematicians now care about their hat color. *New York Times*, page F5, April 10, 2001.
- [21] A. Schrijver. *Combinatorial Optimization*. Springer, 2003.
- [22] I. Segal. Optimal Pricing Mechanism with Unknown Demand. *American Economic Review*, 93:509–29, 2003.
- [23] 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.

A Necessity of additive loss

Much of the work on prior-free optimal auctions for digital goods considers a slightly different performance benchmark than ours (See, e.g., [8, 10, 12]). They define $\mathcal{F}^{(2)}(\mathbf{v})$ to be the revenue of the optimal single-price sale to at least two agents. That is, if $v_{(i)}$ denotes the i th highest valuation, then $\mathcal{F}^{(2)}(\mathbf{v}) = \max_{k \geq 2} kv_{(k)}$. This literature uses this weaker benchmark because no auction can achieve a constant approximation to $\text{OPT}(\mathbf{v})$, e.g., in the case where there is one extremely high valued agent. As the literature shows, it is possible for a randomized auction to get a constant approximation to this weaker benchmark, $\mathcal{F}^{(2)}$. This approach allows for good revenue guarantees even when there is one extremely high valued agent. Our benchmark, on the other hand, offers no guarantee in such a setting as $\text{OPT}(\mathbf{v})/\beta - \gamma h$ is negative (supposing $\beta, \gamma \geq 1$).

We show below that when looking for deterministic auctions the additive loss of $\Omega(h)$ is unavoidable even for the weaker benchmark of $\mathcal{F}^{(2)}$. Thus, our choice to stick with the more natural benchmark OPT is justified.

Lemma A.1 *No deterministic truthful auction obtains a constant fraction of $\mathcal{F}^{(2)}$ on all bid vectors.*

Proof: Assume for contradiction that we are given a deterministic auction with bid-independent function f , that obtains a profit of $\mathcal{F}^{(2)}/\beta$ on all inputs. Let $\mathbf{v} = (1, \dots, 1)$ be the all-ones bid vector. Assume without loss of generality that $f(\mathbf{v}_{-1})$, the price offered the first agent, is 1. Now, for any $\alpha > \beta$ and $i \in I = \{1, 2, \dots, \lceil (n+1)/(n/\alpha - 1) \rceil\}$, consider $\mathbf{v}^{(i)}$ as the all-ones bid vector except for $b_1 = n\alpha^i$. Let S_i be the set of other agents (not including agent 1) that are offered price 1 when the input to the auction is $\mathbf{v}^{(i)}$, i.e., $S_i = \{j > 1 : f(\mathbf{v}_{-j}^{(i)}) = 1\}$.

Fact 1: $|S_i| \geq n/\alpha - 1$. This follows directly from the fact that otherwise the auction's profit would be at most $\mathcal{F}^{(2)}/\alpha$ which would contradict our assumption.

Fact 2: $\exists i, j : S_i \cap S_j \neq \emptyset$. Clearly $|\bigcup_i S_i| \leq n$. But for a contradiction, if the intersection of every pair of S_i s is empty then the union of the S_i s is of size $\sum_i |S_i| \geq |I|(n/\alpha - 1) \geq n + 1 > n$.

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 $p_1 = f(\mathbf{v}'_{-1})$ where \mathbf{v}' is the all-ones input except for $b'_k = h$. If $p_1 \leq n\alpha^i$, 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 $n\alpha^i + n < 2n\alpha^j/\beta$. On the other hand, if $p_1 > n\alpha^i$, 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