



Bayesian and Randomized Clock Auctions

MICHAL FELDMAN, Tel Aviv University, Israel and Microsoft Research, Israel

VASILIS GKATZELIS, Drexel University, USA

NICK GRAVIN, Shanghai University of Finance and Economics, China

DANIEL SCHOEPLIN, Drexel University, USA

In a single-parameter mechanism design problem, a provider is looking to sell some service to a group of potential buyers. Each buyer i has a private value v_i for receiving this service, and some feasibility constraint restricts which subsets of buyers can be served simultaneously. Recent work in economics introduced (deferred-acceptance) *clock auctions* as a superior class of auctions for this problem, due to their transparency, simplicity, and very strong incentive guarantees. Subsequent work in computer science focused on evaluating these auctions with respect to their social welfare approximation guarantees, leading to strong impossibility results: in the absence of prior information regarding the buyers' values, no deterministic clock auction can achieve a bounded approximation, even for simple feasibility constraints with only two maximal feasible sets.

We show that these negative results can be circumvented either by using access to *prior information* or by leveraging *randomization*. In particular, we provide clock auctions that give a $O(\log \log k)$ approximation for general downward-closed feasibility constraints with k maximal feasible sets, for three different information models, ranging from full access to the value distributions to complete absence of information. The more information the seller has, the simpler and more practical these auctions are. Under full access, we use a particularly simple deterministic clock auction, called a *single-price clock auction*, which is only slightly more complex than posted price mechanisms. In this auction, each buyer is offered a single price, then a feasible set is selected among those who accept their offers. In the other extreme, where no prior information is available, this approximation guarantee is obtained using a complex randomized clock auction. In addition to our main results, we propose a parameterization that interpolates between single-price clock auctions and general clock auctions, paving the way for an exciting line of future research.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; **Computational pricing and auctions**.

Additional Key Words and Phrases: clock auctions, obvious strategyproofness

ACM Reference Format:

Michal Feldman, Vasilis Gkatzelis, Nick Gravin, and Daniel Schoepflin. 2022. Bayesian and Randomized Clock Auctions. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC '22), July 11–15, 2022, Boulder, CO, USA*. ACM, New York, NY, USA, 26 pages. <https://doi.org/10.1145/3490486.3538247>

1 INTRODUCTION

Our goal in this paper is to design auctions for the following well-studied class of mechanism design problems: given some set N of n buyers that request a service, and a feasibility constraint $\mathcal{F} \subseteq 2^N$ that restricts the subsets of buyers that can be served simultaneously, we need to decide which feasible subset of buyers $F \in \mathcal{F}$ should be served, and how much each served buyer should pay for the service. A crucial obstacle is that we do not know the *value* v_i of each buyer i for the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EC '22, July 11–15, 2022, Boulder, CO, USA

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9150-4/22/07...\$15.00

<https://doi.org/10.1145/3490486.3538247>

service, i.e., the amount that they are willing to pay for it. Therefore, unless our auction is carefully designed, the buyers can choose to misrepresent this value, aiming to minimize their payment.

For a simple concrete example, consider a setting where some ticket broker is looking to sell 100 tickets for a concert and each potential buyer is interested in purchasing a specific number of tickets so that they can go with their family, or group of friends. The buyers would go to the concert only if they are able to secure their desired number of tickets (depending on their group size), but the broker has just 100 tickets, so it would only be feasible to satisfy subsets of buyers whose total demand for tickets is at most 100. The broker needs to use some mechanism that decides who should get tickets and how much they should each pay, and there are several different types of mechanisms to choose from. For example, she could ask each potential buyer to privately submit a bid with the amount that they are willing to pay, and then use this information to decide the outcome. A much simpler, but possibly inefficient, option would be to just post a fixed price and sell the tickets at that price on a first-come-first-served basis.

Milgrom and Segal [52, 53] recently provided a compelling argument that an ideal solution for this class of problems is a (deferred-acceptance) *clock auction*, and they designed one for reallocating radio frequency licenses, generating almost 20 billion dollars in revenue [34]. An ascending clock auction takes place over a sequence of rounds: each buyer is offered a *personalized* price that weakly increases over time and if the price exceeds the amount they are willing to pay, they can permanently drop out of the auction and pay nothing (note that this is unlike other ascending auctions that assign prices to items instead). When the set of buyers that remain active (the ones that have not dropped out) becomes feasible, the auction can terminate and serve these buyers at the cost of the last price they accepted. Therefore, designing a clock auction reduces to the algorithmic problem of carefully choosing the vector of prices offered to buyers in each round.

Any mechanism that follows the clock auction format automatically satisfies a list of very appealing properties that make clock auctions a highly practical solution that is well-suited for real-world applications. For example, not only are they *strategyproof* (meaning buyers have no reason to unilaterally misrepresent their value); they are actually *obviously strategyproof* [49], which implies significantly stronger incentive guarantees, including *weak group-strategyproofness*, i.e., that even if a coalition of buyers misrepresented their values in a coordinated way, they would not all benefit from this deviation. Furthermore, clock auctions guarantee (i) *transparency* (there is no way in which the auctioneer can mishandle, behind the scenes, the information provided by the buyers), (ii) *unconditional winner privacy* (the winners of the auction never need to reveal their true value), and (iii) *simplicity* (the buyers do not need to understand the inner workings of the auction; all they need to know is the price offered to them in each round). These are properties that most strategyproof auctions do not satisfy (see Section 1.2 for additional discussion on these properties).

Motivated by the strong evidence in favor of clock auctions, subsequent work focused on analyzing their performance in a variety of different settings (e.g., [8, 10, 14, 26, 29, 38, 39, 45, 50]). For the problems studied in this paper, these efforts were overshadowed by a strong impossibility result: even for seemingly very simple instances, all prior-free deterministic clock auctions are bound to perform poorly [26]. However, the restriction to prior-free auctions, i.e., ones where the seller has no prior information regarding the amount that the buyers may be willing to pay, is rather unrealistic, given the vast amounts of historical data being gathered and stored nowadays. In fact, the standard model in auction design is in the *Bayesian setting*, where the value v_i of each buyer i is drawn independently from some distribution D_i and, although the seller does not know the realization of the values, she knows the distributions. Also, although deterministic auctions are more appealing in practice, using randomization to overcome adversarially constructed instances can be a very effective tool. Despite the long literature on Bayesian auctions and randomized auctions, the power of these tools in the context of clock auctions remains largely unexplored. In

this paper, we address this gap in the literature by analyzing both deterministic Bayesian clock auctions and randomized prior-free clock auctions.

1.1 Our Results

We study the problem of maximizing the social welfare, i.e., the sum of the values of the served buyers, and we compare the expected welfare of our auctions to the expected optimal welfare. We consider problem instances with general downward-closed feasibility constraints $\mathcal{F}$ and provide three alternative clock auctions (depending on the amount of information available to the seller) that all achieve a $O(\log \log k)$ approximation guarantee, where k is the number of maximal sets in $\mathcal{F}$. The more information the seller has, the simpler the proposed auction, providing sellers with three options that exhibit an interesting trade-off between access to information and auction simplicity. These results are in stark contrast to the fact that prior-free deterministic clock auctions cannot achieve any bounded approximation even for very simple instances where $k = 2$ [26].

Result 1: Deterministic single-price clock auction under full access to priors (Section 3). In the standard Bayesian setting, where buyer values are drawn from known distributions, we show that the $O(\log \log k)$ approximation guarantee can be provided by a deterministic clock auction that has a particularly simple structure. It is a member of a class of auctions, which we call *single-price clock auctions*, which use only a small portion of the power that clock auctions possess and, as we discuss in Section 3, are closely related to the well-studied class of posted-price mechanisms. Like posted price mechanisms, single-price clock auctions offer just a single price to each buyer. However, single-price clock auctions can defer the decision regarding which subset of buyers to serve until every buyer has responded to the price they were offered, whereas posted-price mechanisms need to serve every buyer who accepts the offered price. Our results show that this advantage allows us to achieve an exponential improvement over posted-price mechanisms, which cannot achieve better than a $O(\log k / \log \log k)$ approximation, even for a simple feasibility constraint where all the k maximal feasible sets are disjoint [6, 54].

Result 2: Deterministic clock auction under restricted access to priors (Section 4). Our second result is a deterministic auction which achieves the same approximation guarantee using only limited information regarding the distributions: it has access to the expected value of each bidder and to the expected value of the optimal solution, but has no additional information regarding the moments of D_i . This is in line with recent work focusing on mechanism design with limited information (e.g., [3–5, 15, 16, 20, 33, 42, 57]). In contrast to our first auction, which leverages its unrestricted access to the distributions to offer only a single price to each bidder, our second auction may need to interact multiple times with each bidder, over a sequence of rounds. Nevertheless, its structure remains rather simple: every bidder faces the *same* price p , which is gradually increased until either the set A of active bidders is feasible, i.e., $A \in \mathcal{F}$, or some feasible subset of the active bidders $F \subset A$ reaches an appropriate revenue target.

Result 3: Randomized prior-free clock auction (Section 5). Our third result applies to settings without any access to prior information. Using the tools and intuition from the analysis of the first two auctions, we show that the “hedging auction” introduced in [14] gives an $O(\log \log k)$ approximation as well. This provides an exponential improvement over the $O(\sqrt{\log k})$ bound shown in [14]. The hedging auction is randomized and significantly more complicated than the deterministic auctions above. It randomly chooses between two alternatives. The first alternative uses non-uniform pricing (i.e., the price offered to each buyer in each round can be different from what is offered to others) and the price trajectory is determined by a complex price update process. The second alternative randomly “samples” a subset of the buyers and then uses their values to decide which of the non-sampled buyers to serve. Although neither one of these alternatives

achieves a good performance on its own, we show that the best of the two is guaranteed to perform well. The hedging auction depends on randomization in two ways: i) it hedges between the two alternatives, and ii) it uses randomized sampling.

In Section 6, we complement our positive results with an inapproximability result that applies to all three settings we study. Our lower bound holds for a very simple instance with two maximal feasible sets and applies even to Bayesian settings with independent values. This is in contrast to a previous lower bound which holds only under correlated priors [14].

A hierarchy of clock auctions (Section 7). We conclude with a proposed “hierarchy” of clock auctions, depending on the amount of power that they use. This hierarchy sets the stage for an exciting new line of research with a broad range of open problems that would help us develop a more refined understanding of the guarantees that are achievable by a spectrum of gradually more powerful clock auctions. As an initial step toward that direction, we provide some results and observations, and discuss interesting connections between the construction of lower bounds for single-price clock auctions and the literature on large deviations in probability theory.

1.2 The Compelling Advantages of (Deferred-Acceptance) Clock Auctions

The recent surge of interest in clock auctions, starting with the work of Milgrom and Segal [52, 53], has largely been motivated by the long list of appealing properties that these auctions satisfy, which make them a better fit for many real-world applications than the widely-studied sealed-bid auctions. In a sealed-bid auction, the buyers are asked to directly report their values to the auctioneer through a bid, and the auction then takes all this information into consideration in deciding who should be served and what the payments should be. The mechanism design literature has produced several sealed-bid auctions that are strategyproof, but experimental evidence suggests that people often misreport their values even when faced with simple strategyproof sealed-bid auctions, whereas they do follow their dominant strategy in clock auctions [43].

Obvious strategyproofness. To provide theoretical justification for this phenomenon, Li [49] introduced the much more demanding property of *obvious strategyproofness* (OSP). In simple terms, an auction is OSP if it provides every bidder with an *obviously dominant* strategy: at any point during the auction, each bidder’s utility in the best-case scenario if she deviates from this strategy is no more than her utility in the worst-case scenario if she follows this strategy. Using both experimental evidence and theoretical arguments, Li argued that even non-experts can recognize that an obviously dominant strategy is their optimal choice. He also showed that in clock auctions it is an obviously dominant strategy for every bidder to remain in the auction until the price she is offered exceeds her value. In fact, for single-parameter domains like the ones we study here, Li showed that clock auctions are essentially the *only* class of auctions that can achieve the stronger incentive property of OSP. Furthermore, an auction that satisfies OSP, like clock auctions do, is also guaranteed to be *weakly group-strategyproof*, i.e., even if the buyers got together and manipulated the auction in a coordinated way, they would not all benefit from that deviation. This is an incentive property that very few sealed-bid auctions satisfy. Apart from their improved incentives, clock auctions also satisfy properties related to transparency, privacy, and simplicity.

Transparency. Another shortcoming of sealed-bid auctions is that the bidders need to trust that the allocation and prices that the auction outputs are, in fact, the result of the precise computations dictated by the auction’s rules. For example, they need to trust that the auctioneer will not charge them more, even though they have revealed exactly how much they are willing to pay, through their bid. Also, even if the bidders believe that the auctioneer is not malicious, they need to trust that there were no unintended errors in computing the outcome or the prices, as even very small errors could violate the incentive properties. This is in stark contrast to clock auctions where every bidder sees the prices in each round and need not even worry about how they were computed, as

this does not affect their incentives in any way. In fact, clock auctions are *credible*, an important property introduced by Akbargpour and Li [2], meaning that it is optimal for the auctioneer to follow the stated rules of the auction. When selling a single item Akbargpour and Li [2] showed that an ascending clock auction is the *unique* optimal auction that is both credible and strategyproof.

Unconditional Winner Privacy. Sealed-bid auctions also require every bidder to reveal to the auctioneer exactly how much they are willing to pay for the service. In clock auctions, bidders reveal their true value only if they choose to drop out, so the winners do not need to ever reveal how much they were actually willing to pay for the service (which is often very sensitive information). Milgrom and Segal [53] termed this property *unconditional winner privacy*, based on a notion defined earlier by Brandt and Sandholm [11], and demonstrate that the set of auctions which preserve unconditional winner privacy corresponds to the set of clock auctions.

Simplicity without Sacrificing Sophistication. A particularly exciting property of clock auctions from an algorithm designer’s standpoint is that they achieve simplicity without sacrificing sophistication. In the line of work motivated by simplicity in mechanism design (e.g., [7, 41, 55]), simplicity often implies that the algorithmic aspect of the auction is rather straightforward so that the participants can understand it. However, clock auctions provide a very simple interface to the buyers (the ascending prices), and it is always an obviously dominant strategy for them to remain in the auction while their price does not exceed their value, *no matter how the prices are chosen*. Thus, the designer can still make elaborate algorithmic decisions in the “back-end”, in choosing the sequence of prices that it offers, aiming to achieve good performance guarantees, while maintaining the incentives of these buyers, and a simple “front-end” that the buyers interact with.

1.3 Other Related Work

In one of the first papers on the performance of clock auctions, Dütting et al. [26] proved a rather pessimistic result: there exists a seemingly simple family of instances with n bidders, where no deterministic prior-free clock auction can achieve an approximation of $O(\log^{1-\epsilon} n)$ for any constant $\epsilon > 0$. In these instances, the auction needs to decide between serving a single bidder or any subset of the remaining $n - 1$ bidders, so it has just $k = 2$ disjoint maximal feasible sets. This simple feasibility constraint clearly exhibits the impact of the clock auctions’ information limitations, leading to an inapproximability result that grows arbitrarily with the number of agents, even though k is constant. In recent work, Christodoulou et al. [14] designed a deterministic prior-free clock auction that guarantees a $O(\log n)$ approximation for general downward-closed feasibility constraints, showing that this is essentially the class of worst-case instances. To overcome the obstacle posed by these instances, they proposed a randomized clock auction that achieves a $O(\sqrt{\log k})$ approximation for downward-closed feasibility constraints, implying a $O(\sqrt{\log n})$ approximation for the interesting class of feasibility constraints with disjoint maximal feasible sets (generalizing the lower bound instances). Our results in this paper achieve an exponential improvement over both of these results.

Apart from the work focusing on the design of clock auctions, some recent work has focused more broadly on the design of auctions that satisfy just one of the clock auction properties, such as obvious strategyproofness (OSP) [22, 36] or credibility [21, 32, 37]. Most relevant to our results, Ferraioli et al. [36] showed that the aforementioned setting of two disjoint maximal sets remains an obstacle for any prior-free deterministic OSP mechanism by demonstrating that no such mechanism can achieve better than a $\sqrt{\log n}$ -approximation. We thus show an exponential gap between clock auctions with priors or randomization and prior-free deterministic OSP mechanisms for this setting.

We defer additional discussion on related work to the relevant sections; we discuss related work on posted-price mechanisms in Section 3 and on auctions with limited information in Section 4.

2 PRELIMINARIES

A set N of n buyers request some service and each buyer i has a private value v_i indicating how much they are willing to pay for it; the vector of all values is denoted by $\mathbf{v} = (v_i)_{i \in N}$. A feasibility constraint $\mathcal{F} \subseteq 2^N$ contains the subsets of buyers that can be served simultaneously. Throughout the paper we focus on feasibility constraints that are downward-closed, i.e., if $F \in \mathcal{F}$ then $F' \in \mathcal{F}$ for every $F' \subseteq F$. That is, if the set of buyers F can be served, then any subset of F can also be served; a very mild assumption satisfied by virtually all interesting constraint structures.

Two examples of downward-closed feasibility sets are Disjoint-Maximal-Sets and Knapsack. In Disjoint-Maximal-Sets, the buyers are partitioned into k disjoint groups, $S_1, \dots, S_k$, and a set F of buyers is feasible if and only if it is a subset of one of these groups, i.e., $\mathcal{F} = \{F \subseteq N \mid \exists S_j \text{ s.t. } F \subseteq S_j\}$. In Knapsack, each buyer $i \in N$ has a demand of size c_i and the auctioneer has a supply constraint restricting her to serve only subsets whose total demand does not exceed the supply, which is normalized to 1, i.e., $\mathcal{F} = \{F \subseteq N \mid \sum_{i \in F} c_i \leq 1\}$.

Our goal is to design auctions that serve a feasible subset of buyers $F \in \mathcal{F}$, aiming to maximize the social welfare, $SW(F) = \sum_{i \in F} v_i$. In the Bayesian setting, the value v_i of each bidder i is independently drawn from a distribution D_i and the product distribution over all bidders is denoted by $\mathbf{D} = \times_{i \in N} D_i$. The social welfare of a Bayesian auction is, therefore, a random variable and, for simplicity, we use $\mathbb{E}_{\mathbf{v} \sim \mathbf{D}} [\text{AUC}]$ to denote the expected social welfare of an auction AUC. We evaluate the performance of our Bayesian clock auctions using the demanding benchmark of the expected optimal social welfare, i.e., $\text{OPT} = \mathbb{E}_{\mathbf{v} \sim \mathbf{D}} [\max_{F \in \mathcal{F}} \{\sum_{i \in F} v_i\}]$. We say an auction AUC approximates the optimal welfare within a factor α if $\text{OPT} \leq \alpha \cdot \text{AUC}$ for every distribution $\mathbf{D}$. Recall that a Bayesian auction can use information regarding the distribution $\mathbf{D}$. When clear in the context, we omit the subscript from the expectation. In the prior-free setting, the values of the bidders are chosen adversarially and the expected social welfare of a randomized auction is over its own randomness. Our goal in both settings is to design auctions with small approximation factors.

A clock auction is a multi-round mechanism in which each bidder faces a personalized “clock” price, which weakly increases over time. At the outset of the auction, all bidders are “active”. An initial vector of prices $\mathbf{p}^1 = \{p_i^1\}_{i=1}^n$ is posted to the bidders, and every bidder may choose to permanently exit the auction (this bidder is no longer active) or stay active at the current price. In each round t , the auctioneer posts a new price p_i^t to each active bidder i , where $p_i^t \geq p_i^{t-1}$. Each active bidder can then choose to remain active or exit permanently. If in round τ the remaining active bidders form a feasible set $F \in \mathcal{F}$, the auction can terminate and serve each active bidder i at price p_i^τ (all bidders who have exited the auction are not served and pay 0). The prices in each round are chosen using any public information, such as the history of prices offered, the bidders’ responses to these prices (whether they accepted them or not), and the feasibility structure. In a setting with prior information, the prior information can also be used to inform this choice.

3 DETERMINISTIC SINGLE-PRICE CLOCK AUCTION WITH PRIORS

We begin by proposing a particularly simple class of clock auctions which we call *single-price clock auctions*. Rather than asking each bidder to respond to multiple (increasing) price offers, these auctions offer a single price to each. As a result, they can be implemented in a decentralized and asynchronous fashion: each bidder could arrive at a different time, respond to the price that they are offered, and depart. The mechanism first collects the responses of all bidders and then, using this information, decides which subset to serve. Single-price clock auctions take the following form:

MECHANISM 1: Single-price clock auction template

- 1 Calculate some vector of prices $\mathbf{p}$ using public information
 - 2 Set the clock of each bidder i to p_i (bidders then choose to remain active or exit)
 - 3 Among the bidders that chose to remain in the auction, select a “preferred” feasible subset F
 - 4 Reject all bidders $i \notin F$, accept all bidders in $i \in F$ and charge them p_i
-

In this section, we first discuss interesting connections between these auctions and the very well-studied class of posted-price mechanisms. We then show our main technical result of the section. Given a downward-closed feasibility constraint $\mathcal{F}$, let $\mathcal{S} = (S_1, \dots, S_k)$ denote the collection of maximal feasible sets in $\mathcal{F}$. We demonstrate that, using full access to the priors, we can design a single-price clock auction that achieves a $O(\log \log k)$ -approximation to the benchmark welfare.

3.1 Connections to posted-price mechanisms

Single-price clock auctions are closely related to the class of (*sequential*) *posted-price mechanisms*, which has attracted a lot of attention in the algorithmic mechanism design literature (e.g., [13, 24, 27, 28, 35, 40, 46, 54, 56]). These mechanisms approach the buyers sequentially and offer each of them a take-it-or-leave-it price for being served (possibly depending on the history of responses). They are then served if and only if they accept the price offered to them. These mechanisms are typically given access to the full value distributions that the buyer values are drawn from.

This line of work draws upon the “prophet inequality” literature in optimal stopping theory [47, 48]. The literature on sequential posted pricing (and prophet inequalities) includes different models, depending on the assumption regarding the arrival order of the buyers. The most common assumption is that buyers arrive in an adversarial order, but other models have been studied too. For example, in the so-called *ordered prophets* model, the mechanism designer can choose the order in which she approaches the buyers (see, e.g., [1, 9, 13, 59]). Similarly, in the *prophet secretary* model, buyers are assumed to arrive in a random order (see, e.g., [1, 17, 19, 30, 31]).¹

From an implementation standpoint, posted-price mechanisms are appealing because they are decentralized, like single-price clock auctions. However, single-price clock auctions do not commit to serving *all* the buyers who accept the prices offered to them (they are not *online* mechanisms). They instead defer their decision regarding which subset of buyers to serve after all the buyers have responded to their offers, which provides these auctions with some additional information.

These mechanisms also differ in the level of adaptivity they may employ in setting prices (see, e.g., [18, 51]). A non-adaptive mechanism commits to the prices before receiving any responses, whereas an adaptive mechanism can offer prices sequentially, conditioning the offers on the responses of previous buyers. A single-price clock auction uses non-adaptive pricing (every buyer’s price is determined in advance). Thus, comparing the performance of single-price clock auctions to that of posted-price mechanisms has interesting implications regarding the power of *deferred-acceptance*, that the former have, relative to the power of *price-adaptivity*, possessed by the latter.

3.2 Achieving the $O(\log \log k)$ bound

In what follows, we demonstrate that if the auctioneer has access to the entire product distribution $\mathbf{D} = \times_{i \in N} D_i$ of bidder values, she can design a single-price clock auction which obtains a $O(\log \log k)$ approximation guarantee. This is in contrast to posted-price mechanisms which we know cannot achieve better than a $O\left(\frac{\log k}{\log \log k}\right)$ approximation even for Disjoint-Maximal-Sets [6, 54]. Thus, our single-price clock auction outperforms all posted-price mechanisms. In this

¹It is easy to verify that even ordered prophets can be simulated as multi-round clock auctions.

light, our result suggests that, in our setting, the ability to defer the acceptance decisions can be significantly more useful than the ability to adapt the prices.

THEOREM 3.1. *Every downward-closed setting with k maximal sets admits a deterministic single-price clock auction that obtains an $O(\log \log k)$ approximation to the expected optimal welfare if the full distribution $\mathbf{D}$ is known.*

Before deriving our single-price clock auction, we first decompose the benchmark welfare into two parts. For each realization of bidder values there is some maximal feasible set $S \in \mathcal{S}$ that achieves the maximum social welfare. This set could have a lot of “low” value bidders, a few “high” value bidders, or a mix of the two. To prove that our auction approximates OPT, we can then divide OPT into two parts: the portion of this expected welfare that comes from “low” value bidders and the portion that comes from “high” value bidders. We show that if the low-value contribution is greater, the approximation factor is achieved by offering all bidders a price of 0 and by accepting only the bidders from the set with highest expectation. On the other hand, if the low-value contribution is not within $\log \log k$ factor of the benchmark, then the high-value part must make up nearly all of the expected optimal welfare. We show that the vast majority of the high-value contribution comes from sets with a logarithmic (or fewer) number of high value bidders. To define the partition of bidders into high and low values, we set a threshold t_S for each maximal feasible set $S \in \mathcal{S}$ so that the expected number of bidders above the threshold is $\log k$. That is, let

$$S(t, \mathbf{v}) \stackrel{\text{def}}{=} \{i \in S \mid v_i > t\}$$

Then, the threshold t_S is the value satisfying²

$$\mathbf{E}_{\mathbf{v}} [|S(t_S, \mathbf{v})|] = \log k.$$

The choice of $\log k$ is a critical point at which it is still likely that for all k maximal sets, the number of bidders exceeding their corresponding set’s threshold t_S is within a constant factor of the expectation. This is cast in the following lemma.

LEMMA 3.2. *Let t be a threshold for a set S such that $\mathbf{E}_{\mathbf{v}} [|S(t, \mathbf{v})|] = \log k$. Then*

$$\Pr_{\mathbf{v}} [\exists x \in [0, t) : |S(x, \mathbf{v})| > 10 \cdot \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]] = o(1/k^2).$$

PROOF. First consider the threshold t such that $\mathbf{E}_{\mathbf{v}} [|S(t, \mathbf{v})|] = \log k$ and an interval of threshold values $[t', t)$ such that $\mathbf{E}_{\mathbf{v}} [|S(t', \mathbf{v})|] = 2 \cdot \mathbf{E}_{\mathbf{v}} [|S(t, \mathbf{v})|]$. We would like to have $|S(x, \mathbf{v})| < 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]$ for all $x \in [t', t)$. Consider a set of independent $\{0, 1\}$ -value random variables $\{\xi_i\}_{i \in S}$ where $\xi_i \stackrel{\text{def}}{=} \mathbb{I}[v_i \geq t']$. Note that $\sum_{i \in S} \xi_i = |S(t', \mathbf{v})|$. By Chernoff bound, it follows that

$$\Pr_{\mathbf{v}} \left[|S(t', \mathbf{v})| > 5 \mathbf{E}_{\mathbf{v}} [|S(t', \mathbf{v})|] \right] < \left(\frac{e^4}{5^5} \right)^{2 \cdot \log k} = o(1/k^2).$$

For all $x \in [t', t)$ we have $|S(x, \mathbf{v})| \leq |S(t', \mathbf{v})|$ and $\mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] \geq \mathbf{E}_{\mathbf{v}} [|S(t, \mathbf{v})|] = \frac{1}{2} \cdot \mathbf{E}_{\mathbf{v}} [|S(t', \mathbf{v})|]$. Thus, if $|S(t', \mathbf{v})| < 5 \mathbf{E}_{\mathbf{v}} [|S(t', \mathbf{v})|]$, then $|S(x, \mathbf{v})| < 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]$ for all $x \in [t', t)$. Hence,

$$\Pr_{\mathbf{v}} \left[\exists x \in [t', t) : |S(x, \mathbf{v})| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] \right] = o(1/k^2).$$

Consider the list of intervals $[r_0, r_1), [r_1, r_2), \dots, [r_{\ell-1}, r_{\ell})$, where $\mathbf{E}_{\mathbf{v}} [|S(r_j, \mathbf{v})|] = 2 \mathbf{E}_{\mathbf{v}} [|S(r_{j-1}, \mathbf{v})|]$ for all $j = 1, \dots, \ell$ and such that $r_0 = 0$ (or alternatively the lowest value in the support of any bidder) and $r_{\ell} = t$. We may apply the same Chernoff bound argument that we did for the interval

²For simplicity of presentation we assume that such a value exists (e.g., in cases where the expected value is continuous). Otherwise, we set the threshold t_S to be the value satisfying $\inf_{t_S} \mathbf{E}_{\mathbf{v}} [|S(t_S, \mathbf{v})|] \leq \log k$. The same is true for the t' value described in the proof of Lemma 3.2.

$[t', t)$ to any of these intervals and get a bound of $o(1/k^{2^{\ell-j+1}})$ for the interval $[r_{j-1}, r_j)$. Then, by the union bound over all the intervals we get an upper bound of

$$\Pr_{\mathbf{v}} \left[\exists x \in \bigcup_{j=1}^{j=\ell} [r_{j-1}, r_j) : |S(x, \mathbf{v})| > 10 \cdot \mathbb{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] \right] \leq \sum_{j=1}^{j=\ell} o \left(\frac{1}{k^{2^{\ell-j+1}}} \right) = o \left(\frac{1}{k^2} \right).$$

Consequently, $\Pr_{\mathbf{v}} [\exists x \in [0, t) : |S(x, \mathbf{v})| > 10 \mathbb{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]] = o(1/k^2)$. $\square$

We now write the following low-high decomposition of the benchmark.

$$\text{OPT} = \mathbb{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} v_i \right] \leq \underbrace{\mathbb{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} \min(v_i, t_S) \right]}_{\text{LOW}} + \underbrace{\mathbb{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} v_i \cdot \mathbb{I} [v_i > t_S] \right]}_{\text{HIGH}}, \quad (1)$$

where for every bidder i we define two regimes with respect to every set S such that $i \in S$:

$$\text{low-value part: } \check{v}_{i,S} \stackrel{\text{def}}{=} \min\{t_S, v_i\}, \quad \text{high-value part: } \hat{v}_{i,S} \stackrel{\text{def}}{=} v_i \cdot \mathbb{I} [v_i > t_S],$$

so $\check{v}_{i,S} \in [0, t_S]$, and if $\hat{v}_{i,S} > 0$ then it must be that $\hat{v}_{i,S} > t_S$. Also, one can easily verify that $v_i \leq \check{v}_{i,S} + \hat{v}_{i,S}$ for any bidder i and any set S such that $i \in S$.

We first establish the following core-tail decomposition of $\text{HIGH} = \mathbb{E}_{\mathbf{v}} [\max_{S \in \mathcal{S}} \sum_{i \in S} \hat{v}_i]$ that allows us to reduce the problem to the case where all sets $S \in \mathcal{S}$ have size of at most $O(\log k)$. We define the core part of HIGH as

$$\text{HIGH-CORE} \stackrel{\text{def}}{=} \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \hat{v}_i \cdot \mathbb{I} [|S(t_S, \mathbf{v})| \leq 10 \log k + 1] \right) \right].$$

It follows that

$$\begin{aligned} \text{HIGH} &\leq \text{HIGH-CORE} + \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \hat{v}_i \cdot \mathbb{I} [|S(t_S, \mathbf{v})| > 10 \log k + 1] \right) \right] \\ &\leq \text{HIGH-CORE} + \sum_{S \in \mathcal{S}} \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in S} \hat{v}_i \cdot \mathbb{I} [|S(t_S, \mathbf{v})| > 10 \log k + 1] \right] \\ &\leq \text{HIGH-CORE} + \sum_{S \in \mathcal{S}} \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in S} v_i \cdot \mathbb{I} [|S(t_S, \mathbf{v}) - \{i\}| > 10 \log k] \right] \\ &= \underbrace{\text{HIGH-CORE} + \sum_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{E} [v_i] \cdot \Pr_{\mathbf{v}} [|S(t_S, \mathbf{v}) - \{i\}| > 10 \log k]}_{\text{HIGH-TAIL}}. \end{aligned} \quad (2)$$

In the second inequality, we replace max by sum. In the third, we modify the event in the right hand side so that it is independent of the value of bidder i ; we also use the fact that $\hat{v}_i \leq v_i$. In the last inequality we use the independence in the right hand side to write the expectation as a product.

We claim that HIGH-TAIL can be covered by the simple single-price clock auction which sets zero prices to all bidders and selects the set with maximum expected value. Note that it is extremely unlikely that there are more than $10 \log k$ bidders above the threshold t_S in any set S , i.e., the event that $|S(t_S, \mathbf{v})| > 10 \mathbb{E}[|S(t_S, \mathbf{v})|] + 1 = 10 \log k + 1$ for each $S \in \mathcal{S}$ has probability of $o(1/k^2)$. Indeed, Lemma 3.2 gives an even stronger tail probability bound. Note that a simple single-price clock auction with prices at zero would select the feasible set with maximum expected value

and get expected social welfare of $\max_{S \in \mathcal{S}} \mathbf{E}[\sum_{i \in S} v_i] \geq \frac{1}{k} \sum_{S \in \mathcal{S}} \sum_{i \in S} \mathbf{E}[v_i]$, which is more than $\sum_{S \in \mathcal{S}} \sum_{i \in S} \mathbf{E}[v_i] \cdot o(1/k^2)$. Recall that

$$\text{LOW} = \mathbf{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} \check{v}_{i,S} \right] = \mathbf{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} \min(v_i, t_S) \right].$$

Next we show that the single-price clock auction with zero prices is a $O(1)$ -approximation to the LOW term in (1). To this end, it is useful to rewrite $\sum_{i \in S} \check{v}_{i,S} = \int_0^{t_S} |S(x, \mathbf{v})| dx$. By linearity of expectations we also have

$$\mathbf{E}_{\mathbf{v}} \left[\sum_{i \in S} \check{v}_{i,S} \right] = \mathbf{E}_{\mathbf{v}} \left[\int_0^{t_S} |S(x, \mathbf{v})| dx \right] = \int_0^{t_S} \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] dx$$

We use a slightly more complex core-tail decomposition than (2) for the LOW term of (1).

$$\text{LOW-CORE} \stackrel{\text{def}}{=} \mathbf{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \check{v}_{i,S} \cdot \mathbb{I} \left[\forall x \in [0, t_S] : |S(x, \mathbf{v})| \leq 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1 \right] \right) \right]$$

It follows that

$$\begin{aligned} \text{LOW} &\leq \text{LOW-CORE} + \mathbf{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \check{v}_{i,S} \cdot \mathbb{I} \left[\exists x \in [0, t_S] : |S(x, \mathbf{v})| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1 \right] \right) \right] \\ &\leq \text{LOW-CORE} + \sum_{S \in \mathcal{S}} \mathbf{E}_{\mathbf{v}} \left[\sum_{i \in S} \check{v}_{i,S} \cdot \mathbb{I} \left[\exists x \in [0, t_S] : |S(x, \mathbf{v})| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1 \right] \right] \\ &\leq \text{LOW-CORE} + \sum_{S \in \mathcal{S}} \mathbf{E}_{\mathbf{v}} \left[\sum_{i \in S} v_i \cdot \mathbb{I} \left[\exists x \in [0, t_S] : |S(x, \mathbf{v}) - \{i\}| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] \right] \right] \\ &= \text{LOW-CORE} + \underbrace{\sum_{S \in \mathcal{S}} \sum_{i \in S} \mathbf{E}[v_i] \cdot \Pr_{\mathbf{v}} [\exists x \in [0, t_S] : |S(x, \mathbf{v}) - \{i\}| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]]}_{\text{LOW-TAIL}} \quad (3) \end{aligned}$$

We again may cover the tail (second) term on the right hand side of (3) since by Lemma 3.2 $\Pr_{\mathbf{v}} [\exists x \in [0, t_S] : |S(x, \mathbf{v}) - \{i\}| > 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]] = o(1/k^2)$. Hence, $\max_{S \in \mathcal{S}} \mathbf{E}[\sum_{i \in S} v_i] \geq \frac{1}{k} \sum_{S \in \mathcal{S}} \sum_{i \in S} v_i$ is at least $\Theta(k)$ times the value of the LOW-TAIL term in the right hand side of (3). Let us now estimate the LOW-CORE term in (3). Let

$$\mathbb{I}_{\mathcal{E}(S)}(\mathbf{v}) \stackrel{\text{def}}{=} \mathbb{I} \left[\forall x \in [0, t_S] : |S(x, \mathbf{v})| \leq 10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1 \right]$$

be the indicator of the event that for all thresholds x in the range from 0 to t_S the number of bidders in S with values above x does not exceed the expected number by a factor of 10. Then

$$\begin{aligned} \text{LOW-CORE} &= \mathbf{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \check{v}_{i,S} \cdot \mathbb{I}_{\mathcal{E}(S)}(\mathbf{v}) \right) \right] = \mathbf{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \left(\mathbb{I}_{\mathcal{E}(S)}(\mathbf{v}) \cdot \int_0^{t_S} |S(x, \mathbf{v})| dx \right) \right] \\ &\leq \mathbf{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{S}} \int_0^{t_S} (10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1) dx \right] = \max_{S \in \mathcal{S}} \int_0^{t_S} (10 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|] + 1) dx \\ &\leq \max_{S \in \mathcal{S}} \int_0^{t_S} (11 \mathbf{E}_{\mathbf{v}} [|S(x, \mathbf{v})|]) dx = 11 \max_{S \in \mathcal{S}} \mathbf{E}_{\mathbf{v}} \left[\int_0^{t_S} |S(x, \mathbf{v})| dx \right] = 11 \max_{S \in \mathcal{S}} \mathbf{E}_{\mathbf{v}} \left[\sum_{i \in S} v_i \right]. \end{aligned}$$

In other words, we get that the single-price clock auction with zero prices gets a constant approximation to the LOW term of (1).

Inequalities (1) and (2) give us an upper bound $\text{OPT} \leq \text{LOW} + \text{HIGH-TAIL} + \text{HIGH-CORE}$, where the welfare of a single-price clock auction with zero prices covers the terms LOW and HIGH-TAIL.

It remains to show we can obtain a $O(\log \log k)$ approximation to the HIGH-CORE term of the benchmark using a single-price clock auction. Our next lemma gives an upper bound on the HIGH-CORE term. In HIGH-CORE, we may effectively assume that each set $S \in \mathcal{S}$ is of size at most $10 \log k + 1$.

LEMMA 3.3. *Let m be any positive integer, and let*

$$\Delta \stackrel{\text{def}}{=} \mathbf{E} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \hat{v}_i \cdot \mathbb{I} \left[|S(t_S, \mathbf{v})| \leq m \right] \right) \right].$$

There exists a uniform price p such that

$$\Delta \leq O(\log m) \cdot \mathbf{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} \mathbf{E} [v_i \mid v_i \geq p] \cdot \mathbb{I} [v_i \geq p] \right].$$

Note that $\mathbf{E} [\max_{S \in \mathcal{S}} \sum_{i \in S} \mathbf{E} [v_i \mid v_i \geq p] \cdot \mathbb{I} [v_i \geq p]]$ is precisely the expected welfare we get in a single-price clock auction with a uniform price p . Lemma 3.3 shows that the best choice of a single uniform price will give an $O(\log m)$ -approximation to Δ , which translates to $O(\log \log k)$ -approximation to HIGH-CORE³ with $m = 10 \log k + 1$. We give a proof sketch here; the complete proof appears in Appendix A.

PROOF SKETCH FOR LEMMA 3.3. Consider the contribution of bidders with value less than $2 \cdot \Delta$ to the benchmark (equals to Δ). The $O(\log m)$ approximation guarantee is achieved by the revenue (not the welfare) of a single-price clock auction. Observe that the contribution to the benchmark of bidders with $\hat{v}_i \leq \frac{\Delta}{2m}$ is at most $\frac{\Delta}{2}$ so by ignoring all such bidders we lose at most a constant factor of Δ . Consider partitioning all remaining bidders in N into $\log m + 2$ buckets $[\frac{\Delta}{2m}, \frac{\Delta}{m}]$, $[\frac{\Delta}{m}, \frac{2\Delta}{m}]$, $\dots$, $[\Delta, 2\Delta]$ based on their value. Let B be the bucket with the largest expected contribution among all these $O(\log m)$ buckets. The revenue obtained by posting the lower bound of the interval for B is a 2-approximation to the welfare contained in B for each set and, in particular, the one with maximum realized value.

On the other hand, the uniform price auction with $p = 2\Delta$ covers the constant fraction of the contribution from the bidders with $\hat{v}_i \geq 2\Delta$. Indeed, if $v_i \geq 2\Delta$ for some i , then with a constant probability i is the only such bidder. I.e., a single-price clock auction with $p = 2\Delta$ takes all such bidders with constant probability and achieves $O(1)$ approximation to the contribution of bidders with $\hat{v}_i \geq 2\Delta$. $\square$

Thus, a single-price clock auction achieves a $O(\log \log k)$ approximation to the HIGH term in the right side of (1). Combined with the constant approximation given by a single-price clock auction with zero prices to the LOW term, we are ready to present the proof of Theorem 3.1.

PROOF OF THEOREM 3.1. By the analysis above, we can achieve the desired approximation using one of the following uniform price clock auctions (the one with the highest expected social welfare):

- (1) Either let $p_i = 0$ for all $i \in [n]$,

³This is not very surprising given similar statements in the literature, see, e.g., Corollary 1 in [54].

(2) or choose a $j \in [0, \log(10 \log k + 1)]$ and let $p_i = \Delta \cdot 2^{1-j}$ for all $i \in [n]$, where

$$\Delta = \mathbf{E} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} v_i \cdot \mathbb{I} \left[v_i > t_S \right] \cdot \mathbb{I} \left[|S(t_S, \mathbf{v})| \leq 10 \log k + 1 \right] \right) \right].$$

Note that these prices can be found using only the thresholds t_S and the value of Δ . Both t_S and Δ can be inferred from the bidder distributions and constraint system, which are public information. The uniform price of zero covers the LOW term of our benchmark decomposition (1) and the HIGH-TAIL portion of (2), whereas the uniform prices of $\Delta \cdot 2^{1-j}$ cover the HIGH-CORE portion of (2) (which, together, cover the HIGH term of (1)). $\square$

3.3 Alternative Parameterizations

Similar to [14], our bounds are parameterized by k , the number of maximal feasible sets of the instance. Alternative parameterizations that have been considered include, e.g., the number of agents, n , or the size of the largest feasible set, r . For the class of instances with Disjoint-Maximal-Sets, which are known to pose an obstacle for both posted-price mechanisms and deterministic prior-free clock auctions, our bounds improve prior results with respect to both parameters.

When using the number of agents as a parameter, it is known that even for Disjoint-Maximal-Sets, no posted-price mechanism can achieve better than a $O(\log n / \log \log n)$ approximation [6, 54], and no deterministic prior-free clock auction can achieve a $O(\log^{1-\epsilon} n)$ approximation for any constant $\epsilon > 0$ [26]⁴. Since, for Disjoint-Maximal-Sets the number of maximal feasible sets is $k \leq n$, our auctions achieve a $O(\log \log n)$ approximation, which is an exponential improvement compared to both posted-price mechanisms and deterministic prior-free clock auctions.

When using the size of the largest feasible set, r , as the parameter, Rubinstein [54] demonstrated that posted-price mechanisms can achieve a $O(\log n \log r)$ approximation for general downward-closed feasibility constraints. Applying the mechanism described in Theorem 3.1, but with $\Delta = \text{OPT}$ and $m = r$ (where m is from the statement of Lemma 3.3), we get an improved bound with respect to this parameter as well. Indeed, Lemma 3.3 holds for any threshold t_S , in particular for $t_S = 0$. Since the size of every feasible set is at most r , we have $\mathbb{I} \left[|S(t_S, \mathbf{v})| \leq r \right] = 1$ for all $\mathbf{v}$. Lemma 3.3 then shows the existence of a single price p such that the corresponding single-price clock auction achieves a $O(\log r)$ approximation to $\Delta = \text{OPT}$.

4 DETERMINISTIC CLOCK AUCTION WITH LIMITED INFORMATION

Our main result in this section is a deterministic uniform price clock auction that achieves an approximation of $O(\log \log k)$ for any downward-closed $\mathcal{F}$, using only limited prior information regarding the distributions: the expected value of each bidder and the expected value of the optimal solution $\text{OPT} = \mathbf{E}_{\mathbf{v}}[\max_{S \in \mathcal{S}} \sum_{i \in S} v_i]$. This is in stark contrast to deterministic clock auctions without this information, which cannot achieve any constant approximation even for $k = 2$ [26]. In other words, by employing a slightly more complicated auction procedure, i.e., a uniform ascending price procedure, one can circumvent the need for full distributional knowledge by “discovering” the appropriate price from Section 3 which gave the approximation guarantee.

This clock auction contributes to the growing literature aiming to design more robust auctions that do not require full access to the priors. For example, in a “parametric auction” [3, 5] the seller only knows some parameters of the distributions (e.g., the mean). Another related literature is the one focusing on the design of auctions that instead have access to a limited number of samples from each distribution (e.g., [4, 15, 16, 20, 25, 33, 42, 57]).

⁴This result and all the results in Sections 3 and 4 focus on deterministic mechanisms.

Our proposed auction starts by checking whether there exists some maximal feasible set $S \in \mathcal{S}$ whose expected value is at least $\text{OPT}/\log \log k$. If such a set exists, then the desired approximation can easily be achieved by accepting all the bidders of that set, for a price of 0, and rejecting all other bidders (we refer to this part of the auction as “AUC_o”). In the more interesting case where no such set exists, our auction uniformly raises the price of every active bidder using some arbitrarily small step $\delta > 0$, giving each bidder the opportunity to drop out after each price increase. While we present the auction as offering this price simultaneously to all bidders for simplicity, we instead offer each bidder the increased price one at a time in a consistent, but arbitrary, order. The auction terminates when the remaining active bidders become feasible, or the revenue of some feasible subset of active bidders reaches a desired goal of $g = \text{OPT}/(4\alpha)$ where α is the approximation guarantee given by Corollary 4.2 in Appendix A (we refer to this part of the auction as “U-PRICE_g”). Given its similarity to the proof of Theorem 3.1, we defer the full proof of Theorem 4.1 to Appendix A and instead provide a sketch below, highlighting the key differences.

MECHANISM 2: A uniform-price clock auction with limited information

- 1 Let all bidders be active $A = N$ and set clock price $p = 0$
 - 2 **if** $\max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{E}[v_i] \geq \text{OPT}/\log \log k$ **then**
 - 3 Accept the bidders in $\arg \max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{E}[v_i]$, charge each of them $p = 0$, and reject everyone else
 - 4 **else**
 - 5 **Repeat:** increase p by some $\delta > 0$ and offer p to all $i \in A$; remove from A bidders who reject p
 - 6 **Until:** A is feasible (i.e., $A \in \mathcal{F}$) **or** there exists some $F \subset A$ such that $F \in \mathcal{F}$ and $|F| \cdot p \geq g$
 - 7 Accept the largest feasible set $F \subseteq A$, charge each $i \in F$ a price of p , and reject everyone else
-

THEOREM 4.1. *Mechanism 2 achieves a $O(\log \log k)$ approximation to the expected optimal welfare for any downward-closed $\mathcal{F}$ using only $\mathbb{E}[v_i]$ for each $i \in N$ and $\text{OPT} = \mathbb{E}_{\mathbf{v}}[\max_{S \in \mathcal{S}} \sum_{i \in S} v_i]$.*

PROOF SKETCH. Our proof proceeds similarly as in the last section, again employing our benchmark decomposition, inequality (1). Note that AUC_o is exactly the single-price clock auction with all prices equal to zero, which, from our previous section, we know covers the LOW-CORE, LOW-TAIL, and HIGH-TAIL portions of our benchmark. We can then assume that HIGH-CORE is a 2 approximation to OPT, as otherwise AUC_o would be a $O(\log \log k)$ approximation to OPT. Thus it suffices to show that U-PRICE_g gives a $O(\log \log k)$ approximation to HIGH-CORE.

From the proof of Lemma 3.3 in Appendix A we may conclude the following related corollary.

COROLLARY 4.2. *Let m be any positive integer, and let*

$$\Delta \stackrel{\text{def}}{=} \mathbb{E} \left[\max_{S \in \mathcal{S}} \left(\sum_{i \in S} \hat{v}_i \cdot \mathbb{I} \left[|S(t_S, \mathbf{v})| \leq m \right] \right) \right].$$

There exist an $\alpha = O(\log m)$ and a uniform price p such that

$$\Delta \leq \alpha \cdot p \cdot \mathbb{E} \left[\max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{I} \left[v_i \geq p \right] \right] \quad \text{or} \quad \Delta \leq \alpha \cdot \mathbb{E} \left[\max_{i \in N} \mathbb{E} \left[v_i \mid v_i \geq p \right] \cdot \mathbb{I} \left[v_i \geq p \right] \right].$$

Corollary 4.2, thus, proves the existence of a uniform price p such that either the expected revenue obtained by serving the set with the largest number of bidders with value above p is a $O(\log \log k)$ approximation to HIGH-CORE or the expected welfare from serving the single highest value bidder is a $O(\log \log k)$ approximation to HIGH-CORE. Without access to the full prior distributions, however, we cannot calculate this price in advance. Instead, we use U-PRICE_g to gradually sweep through a full range of prices, aiming to find one that yields revenue at least $g = \frac{\text{OPT}}{4\alpha}$. However, this could fail if U-PRICE_g reaches a lower price that also satisfies the revenue

target and stops too soon, selecting a set with low social welfare. We address this by a more careful analysis of the expected welfare of U-PRICE_g , conditioning on the expected size $\mathbb{E}[|T(v)|]$ of the largest feasible set $T(v)$ of bidders who would accept price p . When $\mathbb{E}[|T(v)|] \geq 8$ we show that with constant probability we serve at least half the expected number of bidders and obtain the $O(\log \log k)$ approximation from the *revenue* of U-PRICE_g . On the other hand, when $\mathbb{E}[|T(v)|] < 8$ we show that the expected *social welfare* of U-PRICE_g is within a constant factor of the single highest value bidder, which then gives us the desired approximation. $\square$

5 RANDOMIZED PRIOR-FREE CLOCK AUCTION

In this section, we consider prior-free settings; i.e., settings where no information is given about the bidder values. Using the analytical tools and intuition developed in the previous sections, we can now provide a tighter analysis of the randomized clock auction (the “hedging auction”) introduced in [14], showing that it achieves a $O(\log \log k)$ approximation (an exponential improvement over the $O(\sqrt{\log k})$ bound shown in [14]). This auction provides guarantees even in prior-free settings, but this is at the expense of simplicity and practicality; namely, it is randomized and uses non-uniform pricing. Note that our randomized clock auctions are distributions over deterministic clock auctions, and thus they satisfy all the properties of deterministic clock auctions *ex post* (rather than in expectation).

For completeness, we first explain how this auction works and provide some intuition regarding its approximation guarantee. The hedging function runs either the “water-filling clock auction” (WFCA) or the “sampling auction”, each with probability $1/2$. We begin by defining the “water-filling clock auction” (WFCA). Let $P = \{p = x \cdot \epsilon : x \in \mathbb{N}\}$ denote the set of possible prices that the WFCA considers. As was demonstrated in [14], the WFCA obtains welfare at least $r^*/2$ where $r^* = \max_{p \in P, F \in \mathcal{F}} \{p \cdot |i \in F : v_i \geq p|\}$, i.e., the maximum possible revenue an auction could obtain by offering one of the prices in P to all bidders. This then translates to an $O(\log |S|)$ -approximation for any given set S .

MECHANISM 3: The deterministic *water-filling clock auction* (WFCA)

```

1 let  $t \leftarrow 0$ ,  $A \leftarrow N$ , and  $p_i \leftarrow 0$  for all  $i \in N$ 
2 while  $A \notin \mathcal{F}$  do
3    $t \leftarrow t + 1$ 
4   let  $W \leftarrow \arg \max_{S \in \mathcal{F}: S \subseteq A} \{\sum_{i \in S} p_i\}$  be the latest set of conditional winners
5   let  $\ell \leftarrow \min_{i \in A \setminus W} \{p_i\}$  be the lowest price among conditional losers
6   foreach bidder  $i \in A \setminus W$  with  $p_i = \ell$  do
7     update  $p_i \leftarrow p_i + \epsilon$ 
8     if  $i$  rejects updated price then
9       let  $A \leftarrow A \setminus \{i\}$ 
10 return  $A$ 

```

We next describe the sampling auction. The sampling auction “samples” bidders independently with probability $1/2$ by raising their clock until they reject. We let $T \subseteq N$ denote the sampled bidders and $U = N \setminus T$ denote the unsampled bidders. This sampling process ideally allows us to estimate the value of each S since $v(S) = 2 \cdot \mathbb{E}[v(S \cap U)] = 2 \cdot \mathbb{E}[v(S \cap T)]$.

MECHANISM 4: The Sampling Auction: A randomized clock auction**Input:** The set $\mathcal{S}$ of all maximal sets in $\mathcal{F}$

- 1 let $T \leftarrow \emptyset$
- 2 **for** each bidder $i \in N$ **do**
- 3 With probability $1/2$: increase the clock of i until she rejects and let $T \leftarrow T \cup \{i\}$
- 4 Let $R \leftarrow \operatorname{argmax}_{S \in \mathcal{S}} v(S \cap T)$
- 5 **return** $R \setminus T$

Intuitively, the water-filling auction and sampling auction work when sets of high value are small and large, respectively. This is because when sets are small the optimal revenue in hindsight r^* is close in value to the optimal welfare and when sets are large we expect with reasonably high probability that a random sample of the bidders will give a good estimate of the total value of the set. The “Hedging Auction” leverages this intuition. It combines the benefits of the water-filling clock auction (WFCA) and the sampling auction by running one or the other with equal probability.

MECHANISM 5: The Hedging Auction**Input:** The feasibility constraint $\mathcal{F}$, and set $\mathcal{S}$ of all maximal sets in $\mathcal{F}$

- 1 With probability $1/2$: Run the WFCA (Mechanism 3) on input $\mathcal{F}$
- 2 With the remaining probability $1/2$: Run the Sampling Auction (Mechanism 4) on input $\mathcal{F}$

THEOREM 5.1. *The Hedging Auction obtains a $O(\log \log k)$ -approximation to the social welfare for any downward-closed feasibility constraint $\mathcal{F}$, where k is the number of maximal feasible sets in $\mathcal{F}$.*

PROOF. We fix the vector of bidder values $\mathbf{v}$. For a set $S \in \mathcal{S}$ let $\tau(S)$ be the threshold p such that $|\{i \in S : v_i \geq p\}| = 60 \log k$ and S_{top} denote the subset of bidders with value at least $\tau(S)$ in S (i.e., the $60 \log k$ highest value bidders in S)⁵. Let O denote the optimal set. We then have that the WFCA obtains a $O(\log \log k)$ -approximation to the welfare in O_{top} . Since we run the WFCA with probability $1/2$ in the hedging auction, it just remains to show that the sampling auction obtains a $O(\log \log k)$ -approximation to $v(O) - v(O_{\text{top}})$. We can upper bound $v(S) - v(S_{\text{top}})$ for any S by using our $S(t, \mathbf{v})$ notation from Section 3 as $v(S) - v(S_{\text{top}}) = v(S \setminus S_{\text{top}}) \leq \int_0^{\tau(S)} |S(x, \mathbf{v})| dx$ since every bidder in $S \setminus S_{\text{top}}$ has value no more than $\tau(S)$ by definition. To show that the sampling auction obtains a $O(\log \log k)$ -approximation to $\int_0^{\tau(O)} |O(x, \mathbf{v})| dx$, we use the following lemma. We defer the proof to Appendix A as it is similar to that of Lemma 3.2.

LEMMA 5.2. *When running the sampling auction on any instance we have that for any $S \in \mathcal{S}$,*

$$\Pr [\exists x \in [0, \tau(S)] : |T \cap S(x, \mathbf{v})| \notin [1/9, 8/9] \cdot |S(x, \mathbf{v})|] = o(1/k^2).$$

By taking a union bound over the k maximal sets in $\mathcal{S}$, we obtain the following corollary.

COROLLARY 5.3. *When running the sampling auction on an instance with k maximal sets in $\mathcal{S}$,*

$$\Pr [\forall S \in \mathcal{S} \forall x \in [0, \tau(S)] : |T \cap S(x, \mathbf{v})| \in [1/9, 8/9] \cdot |S(x, \mathbf{v})|] = 1 - o(1/k).$$

We can now complete the proof. We first assume that $v(O)/100 \geq \max_{S \in \mathcal{S}} v(S_{\text{top}})$ as otherwise the WFCA would give a $O(\log \log k)$ -approximation. Let $S^* = \operatorname{argmax}_{S \in \mathcal{S}} v(S \cap T)$. We then have that the random sampling auction obtains welfare $v(S^* \cap T)$. Since S^* was the set of highest sampled welfare, with probability $1 - o(1/k)$ the event in Corollary 5.3 holds and we have that

$$v(S^* \cap T) \geq v(O \cap T) \geq \int_0^{\tau(O)} |T \cap O(x, \mathbf{v})| dx \geq \int_0^{\tau(O)} \frac{1}{9} |O(x, \mathbf{v})| dx \geq \frac{1}{9} \cdot (v(O) - v(O_{\text{top}})),$$

⁵We assume consistent (lexicographical) tie-breaking in the event that there isn't a price for which exactly $60 \log k$ bidders have value above the threshold.

where the third inequality is due to the lower bound on $|T \cap O(x, \mathbf{v})|$ from Corollary 5.3. But then, since $v(O_{\text{top}}) \leq v(O)/100$ we have that $v(S^* \cap T) \geq \frac{11v(O)}{100}$. We can bound $v(S^* \cap U)$ similarly by

$$\begin{aligned} v(S^* \cap U) &\geq \int_0^{\tau(S^*)} |U \cap S^*(x, \mathbf{v})| dx \geq \int_0^{\tau(S^*)} \frac{1}{9} |S^*(x, \mathbf{v})| dx \geq \frac{1}{9} \cdot (v(S^*) - v(S_{\text{top}}^*)) \\ &\geq \frac{1}{9} \cdot \left(v(S^*) - \frac{v(O)}{100} \right) \geq \frac{1}{9} \cdot \left(v(S^* \cap T) - \frac{v(O)}{100} \right) \geq \frac{1}{9} \cdot \left(\frac{11v(O)}{100} - \frac{v(O)}{100} \right) \geq \frac{v(O)}{900}, \end{aligned}$$

where the second inequality is due to the fact that $|U \cap S^*(x, \mathbf{v})| + |T \cap S^*(x, \mathbf{v})| = |S^*(x, \mathbf{v})|$ and the upper bound on $|T \cap S^*(x, \mathbf{v})|$ from Corollary 5.3. Thus, when the WFCA does not give a $O(\log \log k)$ -approximation, the sampling auction gives a $O(1)$ -approximation, as desired. $\square$

6 LOWER BOUND

To complement our positive results, in this section we show that no clock auction (single-price or not) can achieve better than a $29/27$ -approximation to the expected optimal welfare even under full information and independent values. This is in contrast to the lower bound in [14], which applies only with respect to correlated priors. Moreover, our bound applies in a very simple feasibility instance that is a special case of both Disjoint-Maximal-Sets and Knapsack (which is the simplest setting beyond matroid feasibility, for which an optimal solution can be obtained using a deterministic prior-free clock auction [53].) This result demonstrates that in the process of discovering the high-valued bidders, it may be inevitable that the auction loses some low-valued bidders along the way.

THEOREM 6.1. *For any constant $\epsilon > 0$, no clock auction can obtain $29/27 - \epsilon$ approximation to the expected optimal welfare, even in simple Disjoint-Maximal-Sets or Knapsack feasibility constraints and even under full access to the independent value distributions.*

PROOF. Consider an instance with two disjoint maximal feasible sets S and T . S consists of a single bidder with deterministic value 1. T consists of two bidders, each has value $2/5$ with probability $p = 2/3$ and value 1 with probability $(1 - p) = 1/3$, independently. The expected optimal welfare is $1 \cdot p^2 + 7/5 \cdot 2p(1 - p) + 2 \cdot (1 - p)^2 = 4/9 + 28/45 + 2/9 = 58/45$. We show that no clock auction can achieve a higher expected welfare than $54/45$, concluding the assertion of the theorem.


A clock auction that always serves the bidder in S can only obtain value 1. A clock auction that always serves the bidders in T without increasing either clock above $2/5$ obtains value $4/5 \cdot p^2 + 7/5 \cdot 2p(1 - p) + 2 \cdot (1 - p)^2 = 16/45 + 28/45 + 2/9 = 54/45$. It remains to consider the case where the clock auction increases the clock of at least one bidder in T above $2/5$. Fix one of the bidders in T and suppose its clock increases above $2/5$. If its value is $2/5$ (this happens with probability p), then the auction loses that bidder and can obtain value of at most 1. If its value is 1 (this happens with probability $1 - p$), the auction obtains value 1 from this bidder, and an expected value $p \cdot 2/5 + (1 - p) \cdot 1$ from the other bidder. Together, the auction obtains an expected welfare of at most $p \cdot 1 + (1 - p) \cdot (1 + p \cdot 2/5 + (1 - p) \cdot 1) = 54/45$, as desired. $\square$

7 A HIERARCHY OF CLOCK AUCTIONS AND ADDITIONAL RESULTS

Our positive results from the previous sections using simple clock auctions provide hope that clock auctions with prior information can achieve even stronger performance guarantees. To better understand their power and limitations, we propose the study of these auctions through a hierarchy of gradually increasing power.

In the figure below, we provide a sketch of this hierarchy among classes of clock auctions, with the arrows pointing from less general to more general classes. The class of *adaptive single-price*

clock auctions includes all mechanisms that can combine the strengths of posted-prices mechanisms and single-price clock auctions: an adaptive single-price clock auction can offer a price to each buyer in any desired order, and the price offered to each buyer can adapt to the responses provided by other buyers, earlier in the ordering. Furthermore, the decision regarding which subset of buyers should be served can be deferred until all buyers have responded to their offers.



In the remainder of this section, we first present some initial results and observations for single-price clock auctions (§7.1). We then propose the refinement of multi-round clock auctions based on the number of distinct prices offered to each bidder, or the extent to which the choice of the prices can adapt to previous bidder responses (§7.2).

7.1 Single-Price Clock Auctions

Binary valuations. A class of instances that emphasizes the benefits of single-price clock auctions over posted-price mechanisms is that of binary valuations, i.e., where the value of each bidder can be either 0 or 1. Although this class is seemingly simple, it actually poses some of the most significant obstacles for posted-price mechanisms. Many of the lower bounds in the prophet inequality literature come from this class, and even for the special case of Disjoint-Maximal-Sets with n binary value bidders, one cannot obtain a better than $O(\log n / \log \log n)$ approximation [46, 54]. This is in stark contrast to single-price clock auctions, which can trivially achieve the optimal social welfare:

OBSERVATION 7.1. *For the special case of binary valuations (i.e., $v_i \in \{0, 1\}$ for all $i \in N$), every downward-closed set system admits a single-price clock auction that achieves OPT .*

Indeed, the auction can just set a price of, say, $1/2$ for all bidders⁶, at which point the bidders whose value is 0 will drop out, revealing the bidders with value 1. Then, the auction can choose to serve the largest feasible subset of the remaining bidders, leading to the optimal social welfare.

Distributions with bounded support size. A more general class of instances where a single-price clock auction performs well is instances with bounded support size. Specifically, if the support size $\text{supp}(D_i)$ is at most ℓ for every $i \in N$, then we can achieve an approximation of ℓ (proof deferred to Appendix A.4).

CLAIM 7.2. *Suppose $|\text{supp}(D_i)| \leq \ell$ for all $i \in N$, then there is a single-price clock auction using the full distributional prior $\mathbf{D}$ that obtains ℓ approximation to the expected optimal social welfare.*

Connections to large deviation theory. In cases beyond constant support size and binary valuations, the question of whether single-price clock auctions can achieve a constant approximation is quite non-trivial. In fact, the problem remains non-trivial even if we restrict attention to the very special case of Disjoint-Maximal-Sets with equal size maximal sets (i.e., k disjoint maximal sets of $m = n/k$ bidders each) and i.i.d. bidders (i.e., $D_i = D$ for all $i \in N$). As we discuss in greater detail in Appendix B, this question appears to be related to the *large deviation theory* [23] and, particularly, the distributions of the values of sums of random variables when the values of the sums far exceed their expectations.

⁶Clearly, any price $p \in (0, 1)$ would achieve the same result.

7.2 Multi-Round Clock Auctions

As we demonstrate in Sections 4 and 5, moving beyond single-price clock auctions opens up an even greater set of design options. Gradually increasing prices over time allows a clock auction to more carefully discover the preferences of the bidders. In particular, in Section 4, we leverage multiple rounds of uniform prices to reduce the prior information required, and, in Section 5, we leverage randomization and multiple rounds of non-uniform prices to eschew the need for prior information altogether. Examining the trade-off between auction complexity versus prior information access is a fundamental question that follows naturally from our work.

An interesting example where multiple rounds of prices are clearly more powerful than a single price with distributional information is the case of selling a single item: the well-known second-price auction is implementable as an ascending price clock auction (with uniform prices for all bidders as in our auction from Section 4). In this case, even without prior information, a multi-round clock auction obtains the optimal welfare. On the other hand, a single-price clock auction with access to full prior information cannot obtain the optimal welfare in the single item setting even if bidders' values are drawn from i.i.d. distributions. The relative power of offering multiple rounds of prices versus distributional information also then seems to depend on the feasibility constraint in question, suggesting a rich and fundamental line of further research.

The multi-round clock auctions we describe in this work use prior information in a very minimal way (or use no information at all). General multi-round clock auctions with full distributional information, however, can follow highly complicated price trajectories which leverage this information without sacrificing any of their appealing properties. One could then use this price discovery process to update the posterior beliefs regarding bidder values over time to possibly obtain even better performance as a function of the underlying “complexity” of the clock auction itself.

We believe that an interesting way to gradually evaluate the power of more complicated clock auctions is through an even more refined hierarchy depending on (i) the number of prices offered to each bidder, and (ii) whether or not they can adapt these queries to the responses of the bidders. The class of clock auctions maintains its appealing properties no matter how many prices they offer to the bidders, yet, as we reflect in this work, increasing the number of prices comes at a cost. Minimizing the number of prices offered to each bidder reduces the burden on them and speeds up the auction which may be desirable from an implementation perspective, further motivating a deeper analysis of the hierarchy we propose.

8 CONCLUSIONS AND FUTURE WORK

In this paper, we achieve significantly improved social welfare approximation guarantees using Bayesian and randomized clock auctions; our results suggest many directions for future work. Most notably, understanding the best approximation guarantees that can be obtained by deterministic Bayesian clock auctions or randomized prior-free clock auctions under various feasibility constraints is an open line of research. Even for the conceptually simple Disjoint-Maximal-Sets setting, there is an asymptotic gap between the best known upper and lower bounds. A particularly exciting question is whether or not clock auctions can achieve constant approximation guarantees for all downward-closed settings, via access to priors or randomization.

Another direction for future research is toward a refined understanding of the impact that limited prior information has on the performance of clock auctions. For example, how would the results change if the auctioneer has only a limited number of samples from the value distribution of each bidder? This direction has been studied in the prophet inequality literature (see, e.g., [4, 12, 16, 44, 57]).

Finally, it would be interesting to gain a better understanding of the hierarchy of clock auctions discussed in Section 7. For example, is there an asymptotic separation between the best approximation that can be guaranteed by general clock auctions with full distributional priors and single-price clock auctions (with the same distributional information) in various settings? Moreover, one can define a hierarchy of clock auctions based on the number of different prices they are allowed to offer to each bidder. Results in this paper are obtained at the two extremes: single-price clock auctions or clock auctions that use an arbitrary number of rounds. What is the power of clock auctions that are allowed to offer up to $q > 1$ prices to each bidder?

ACKNOWLEDGMENTS

The first author is partially supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), and by the Israel Science Foundation (grant number 317/17). The second and fourth authors were supported in part by NSF grants CCF-2008280 and CCF-1755955. The third author is supported in part by Science and Technology Innovation 2030 "New Generation of Artificial Intelligence" Major Project No.(2018AAA0100903), Innovation Program of Shanghai Municipal Education Commission, Program for Innovative Research Team of Shanghai University of Finance and Economics (IRTSHUFE) and the Fundamental Research Funds for the Central Universities, and by NSFC grant No. 62150610500. We would like to thank Aviad Rubinstein for his feedback that helped us generalize our results.

REFERENCES

- [1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert Kleinberg, and Brendan Lucier. 2017. Beating $1-1/e$ for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 61–71.
- [2] Mohammad Akbarpour and Shengwu Li. 2020. Credible auctions: A trilemma. *Econometrica* 88, 2 (2020), 425–467.
- [3] Pablo Azar, Silvio Micali, Constantinos Daskalakis, and S Matthew Weinberg. 2013. Optimal and efficient parametric auctions. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 596–604.
- [4] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. 2019. Prior independent mechanisms via prophet inequalities with limited information. *Games Econ. Behav.* 118 (2019), 511–532.
- [5] Pablo Daniel Azar and Silvio Micali. 2013. Parametric digital auctions. In *Innovations in Theoretical Computer Science, ITCS '13, Berkeley, CA, USA, January 9–12, 2013*, Robert D. Kleinberg (Ed.). ACM, 231–232.
- [6] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. 2007. Matroids, Secretary Problems, and Online Mechanisms. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. Society for Industrial and Applied Mathematics, USA, 434–443.
- [7] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. 2020. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)* 67, 4 (2020), 1–40.
- [8] Eric Balkanski, Pranav Garimidi, Vasilis Gkatzelis, Daniel Schoepflin, and Xizhi Tan. 2022. Deterministic Budget-Feasible Clock Auctions. In *Proceedings of the Thirty-Third Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2022)*. Society for Industrial and Applied Mathematics, USA.
- [9] Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pál, and Balasubramanian Sivan. 2021. Improved Revenue Bounds for Posted-Price and Second-Price Mechanisms. *Operations Research* 69, 6 (2021), 1805–1822.
- [10] Martin Bichler, Zhen Hao, Richard Littmann, and Stefan Waldherr. 2020. Strategyproof auction mechanisms for network procurement. *OR Spectrum* 42, 4 (2020), 965–994.
- [11] Felix Brandt and Tuomas Sandholm. 2005. Unconditional privacy in social choice. In *Proceedings of the 10th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-2005), Singapore, June 10–12, 2005*, Ron van der Meyden (Ed.). National University of Singapore, 207–218.
- [12] Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. 2022. Single-Sample Prophet Inequalities via Greedy-Ordered Selection. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1298–1325.
- [13] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*.

311–320.

- [14] Giorgos Christodoulou, Vasilis Gkatzelis, and Daniel Schoepflin. 2022. Optimal Deterministic Clock Auctions and Beyond. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [15] Richard Cole and Tim Roughgarden. 2014. The sample complexity of revenue maximization. In *Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014*, David B. Shmoys (Ed.). ACM, 243–252. <https://doi.org/10.1145/2591796.2591867>
- [16] José Correa, Paul Dütting, Felix Fischer, and Kevin Schewior. 2019. Prophet inequalities for iid random variables from an unknown distribution. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC*. ACM, 3–17.
- [17] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 169–186.
- [18] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2019. Recent developments in prophet inequalities. *ACM SIGecom Exchanges* 17, 1 (2019), 61–70.
- [19] Jose Correa, Raimundo Saona, and Bruno Ziliotto. 2019. Prophet secretary through blind strategies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1946–1961.
- [20] José R Correa, Andrés Cristi, Boris Epstein, and José A Soto. 2020. The two-sided game of googol and sample-based prophet inequalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2066–2081.
- [21] Constantinos Daskalakis, Maxwell Fishelson, Brendan Lucier, Vasilis Syrgkanis, and Santhoshini Velusamy. 2020. Simple, credible, and approximately-optimal auctions. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 713–713.
- [22] Bart De Keijzer, Maria Kyropoulou, and Carmine Ventre. 2020. Obviously Strategyproof Single-Minded Combinatorial Auctions. In *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020*. 71.
- [23] F. den Hollander. 2008. *Large Deviations*. American Mathematical Society.
- [24] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. 2020. Prophet inequalities made easy: stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.* 49, 3 (2020), 540–582.
- [25] Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. 2021. Efficient two-sided markets with limited information. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. 1452–1465.
- [26] Paul Dütting, Vasilis Gkatzelis, and Tim Roughgarden. 2017. The Performance of Deferred-Acceptance Auctions. *Math. Oper. Res.* 42, 4 (2017), 897–914.
- [27] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. 2020. An $O(\log \log m)$ prophet inequality for subadditive combinatorial auctions. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 306–317.
- [28] Paul Dütting and Robert Kleinberg. 2015. Polymatroid prophet inequalities. In *Algorithms-ESA 2015*. Springer, 437–449.
- [29] Paul Dütting, Inbal Talgam-Cohen, and Tim Roughgarden. 2017. Modularity and greed in double auctions. *Games and Economic Behavior* 105 (2017), 59–83.
- [30] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. 2018. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 700–714.
- [31] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. 2017. Prophet secretary. *SIAM Journal on Discrete Mathematics* 31, 3 (2017), 1685–1701.
- [32] Meryem Essaidi, Matheus V. X. Ferreira, and S. Matthew Weinberg. 2022. Credible, Strategyproof, Optimal, and Bounded Expected-Round Single-Item Auctions for All Distributions. In *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA (LIPIcs)*, Mark Braverman (Ed.), Vol. 215. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 66:1–66:19.
- [33] Tomer Ezra, Michal Feldman, and Ilan Nehama. 2018. Prophets and Secretaries with Overbooking. In *Proceedings of the 2018 ACM Conference on Economics and Computation, EC*. ACM, 319–320.
- [34] FCC. 2017. Broadcast Incentive Auction. <https://www.fcc.gov/about-fcc/fcc-initiatives/incentive-auctions>.
- [35] Michal Feldman, Nick Gravin, and Brendan Lucier. 2014. Combinatorial auctions via posted prices. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 123–135.
- [36] Diodato Ferraioli, Paolo Penna, and Carmine Ventre. 2021. Two-way greedy: Algorithms for imperfect rationality. In *International Conference on Web and Internet Economics*. Springer, 3–21.
- [37] Matheus VX Ferreira and S Matthew Weinberg. 2020. Credible, truthful, and two-round (optimal) auctions via cryptographic commitments. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 683–712.
- [38] Vasilis Gkatzelis, Evangelos Markakis, and Tim Roughgarden. 2017. Deferred-acceptance auctions for multiple levels of service. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 21–38.

- [39] Vasilis Gkatzelis, Rishi Patel, Emmanouil Pountourakis, and Daniel Schoepflin. 2021. Prior-Free Clock Auctions for Bidders with Interdependent Values. In *Algorithmic Game Theory - 14th International Symposium, SAGT 2021, Aarhus, Denmark, September 21-24, 2021, Proceedings (Lecture Notes in Computer Science)*, Ioannis Caragiannis and Kristoffer Arnsfelt Hansen (Eds.), Vol. 12885. Springer, 64–78.
- [40] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. 2007. Automated Online Mechanism Design and Prophet Inequalities. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1 (AAAI 2007)*. AAAI Press, 58–65.
- [41] Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM Conference on Electronic Commerce*. ACM, 225–234.
- [42] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. 2018. Making the Most of Your Samples. *SIAM J. Comput.* 47, 3 (2018), 651–674. <https://doi.org/10.1137/16M1065719>
- [43] John H Kagel, Ronald M Harstad, and Dan Levin. 1987. Information impact and allocation rules in auctions with affiliated private values: A laboratory study. *Econometrica: Journal of the Econometric Society* (1987), 1275–1304.
- [44] Haim Kaplan, David Naori, and Danny Raz. 2022. Online Weighted Matching with a Sample. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 1247–1272.
- [45] Anthony Kim. 2015. Welfare maximization with deferred acceptance auctions in reallocation problems. In *Algorithms-ESA 2015*. Springer, 804–815.
- [46] Robert Kleinberg and Seth Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing*. ACM, 123–136.
- [47] Ulrich Krengel and Louis Sucheston. 1977. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83, 4 (1977), 745–747.
- [48] Ulrich Krengel and Louis Sucheston. 1978. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics* 4 (1978), 197–266.
- [49] Shengwu Li. 2017. Obviously Strategy-Proof Mechanisms. *American Economic Review* (2017).
- [50] Simon Loertscher and Leslie M Marx. 2020. Asymptotically optimal prior-free clock auctions. *Journal of Economic Theory* (2020), 105030.
- [51] Brendan Lucier. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1 (2017), 24–47.
- [52] Paul Milgrom and Ilya Segal. 2014. Deferred-acceptance auctions and radio spectrum reallocation. In *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, Moshe Babaioff, Vincent Conitzer, and David A. Easley (Eds.). ACM, 185–186.
- [53] Paul Milgrom and Ilya Segal. 2019. Clock Auctions and Radio Spectrum Reallocation. *Journal of Political Economy* (2019).
- [54] Aviad Rubinstein. 2016. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. 324–332.
- [55] Aviad Rubinstein. 2016. On the computational complexity of optimal simple mechanisms. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*. ACM, 21–28.
- [56] Aviad Rubinstein and Sahil Singla. 2017. Combinatorial Prophet Inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2017)*. Society for Industrial and Applied Mathematics, USA, 1671–1687.
- [57] Aviad Rubinstein, Jack Z Wang, and S Matthew Weinberg. 2020. Optimal Single-Choice Prophet Inequalities from Samples. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [58] Jan Vondrák. 2010. A note on concentration of submodular functions. *CoRR* abs/1005.2791 (2010). [arXiv:1005.2791](http://arxiv.org/abs/1005.2791)
- [59] Qiqi Yan. 2011. Mechanism design via correlation gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 710–719.

A OMITTED PROOFS

A.1 Proof of Lemma 3.3 and Corollary 4.2

We note that the proofs of Lemma 3.3 and Corollary 4.2 are the same. We thus provide one unified proof for both below.

PROOF. Let $m = 10 \log k + 1$ and $\Delta = \mathbb{E}[\max_{S \in \mathcal{S}} (\sum_{i \in S} \hat{v}_i \cdot \mathbb{I}[|S(t_S, \mathbf{v})| \leq m])]$. Consider an event that there exists a bidder with value greater than or equal to 2Δ . By posting the price of 2Δ for every bidder, we identify all such bidders and would serve at least one if they exist. Following the

same reasoning as in [54], let q_i be the probability that bidder i has value $\hat{v}_i \geq 2\Delta$. Then

$$\Delta \geq 2 \cdot \Delta \cdot \Pr_{\mathbf{v}} [\exists i : \hat{v}_i \geq 2\Delta] = 2 \cdot \Delta \cdot \left(1 - \prod_i (1 - q_i)\right).$$

Rearranging gives $\frac{1}{2} \leq \prod_i (1 - q_i) \leq e^{-\sum_i q_i}$ but then $\sum_i q_i \leq \ln 2$. The probability that we would want to accept a bidder and cannot (because there is some other bidder in another feasible set we are taking) is less than $\ln 2$ so we must obtain at least a $(1 - \ln 2)$ -fraction of the contribution from bidders with value above $2 \cdot \Delta$ by serving at most one such bidder.

Consider how much the small bidders with value less than $2 \cdot \Delta$ contribute to the benchmark. The approximation guarantee is achieved by the revenue (not the welfare) of a single-price clock auction. Observe that the contribution to the benchmark of bidders with $\hat{v}_i \leq \frac{\Delta}{2m}$ is at most $\frac{\Delta}{2}$ so by ignoring all such bidders we lose at most a constant factor against the benchmark. Consider partitioning all remaining bidders into $\log m + 2$ buckets $[\frac{\Delta}{2m}, \frac{\Delta}{m}]$, $[\frac{\Delta}{m}, \frac{2\Delta}{m}]$, $\dots$, $[\Delta, 2\Delta]$ based on their value. We can decompose the remaining half of the benchmark (among bidders with value less than $2 \cdot \Delta$) by upper bounding it by the sum of the best set for each of these buckets. By the pigeonhole principle, the largest expected contribution of these buckets is a $O(\log m)$ fraction of the contribution to the benchmark among bidders with value less than $2 \cdot \Delta$. Consider the best of these buckets B . The revenue obtained by posting the lower bound of the interval for B is a 2-approximation to the welfare contained in B for each set (by construction) and, in particular, the one with maximum realized value. Thus, posting the price equal to the lower bound of B for each bidder and then selecting the set with the greatest number of accepting bidders obtains revenue which is a $O(\log m)$ fraction of the expected contribution from values less than or equal to 2Δ . Finally, choosing the better of the core (prices below 2Δ) and the tail ($p = 2\Delta$) then gives a $O(\log m)$ -approximation to Δ . $\square$

A.2 Proof of Lemma 5.2

PROOF. Consider an arbitrary set S and the threshold $\tau(S)$ for which $|S(\tau(S), \mathbf{v})| = 60 \log k$. Since in the sampling auction each bidder is sampled independently with probability $1/2$ we have that $\mathbb{E}[T \cap S(\tau(S), \mathbf{v})] = 30 \log k = \mathbb{E}[U \cap S(\tau(S), \mathbf{v})]$. Now consider the threshold t'_S such that $|S(t'_S, \mathbf{v})| = 80 \log k = 4/3 \cdot |S(\tau(S), \mathbf{v})|$. Similarly as above, we have that $\mathbb{E}[T \cap S(t'_S, \mathbf{v})] = \mathbb{E}[U \cap S(t'_S, \mathbf{v})] = 40 \log k$. By a Chernoff bound, it follows that

$$\Pr \left[|T \cap S(t'_S, \mathbf{v})| > \frac{4}{3} \mathbb{E}[|T \cap S(t'_S, \mathbf{v})|] \right] < \left(\frac{e^{1/3}}{(4/3)^{4/3}} \right)^{40 \log k} < 0.951^{40 \log k} = o(1/k^2). \quad (4)$$

Now, for all $p \in [t'_S, \tau(S)]$ we know that $|T \cap S(p, \mathbf{v})| \leq |T \cap S(t'_S, \mathbf{v})|$ and $\mathbb{E}[|T \cap S(p, \mathbf{v})|] \geq \mathbb{E}[|T \cap S(\tau(S), \mathbf{v})|] = \frac{3}{4} \cdot \mathbb{E}[|T \cap S(t'_S, \mathbf{v})|]$. Therefore, if $|T \cap S(t'_S, \mathbf{v})| < \frac{4}{3} \mathbb{E}[|T \cap S(t'_S, \mathbf{v})|]$, then $|T \cap S(p, \mathbf{v})| < \frac{16}{9} \mathbb{E}[|T \cap S(p, \mathbf{v})|]$ for all $p \in [t'_S, \tau(S)]$. Hence, we have

$$\Pr \left[\exists p \in [t'_S, \tau(S)) : |T \cap S(p, \mathbf{v})| > \frac{16}{9} \mathbb{E}[|T \cap S(p, \mathbf{v})|] \right] = o(1/k^2). \quad (5)$$

Consider intervals $[r_0, r_1), [r_1, r_2), \dots, [r_{\ell-1}, r_\ell)$ where $\mathbb{E}[|T \cap S(r_j, \mathbf{v})|] = \frac{3}{4} \cdot \mathbb{E}[|T \cap S(r_{j-1}, \mathbf{v})|]$ for all $j = 1, \dots, \ell$ and such that $r_0 = 0$ and $r_\ell = \tau(S)$. We can apply the same Chernoff bound argument that we did for the interval $[t'_S, \tau(S))$ to any such interval and get a bound of $o\left(1/k^{2 \cdot (4/3)^{\ell-j+1}}\right)$ for

the interval $[r_{j-1}, r_j]$. But then, we may take a union bound over all of these intervals to obtain that

$$\Pr \left[\exists p \in \cup_{j=1}^{j=\ell} [r_{j-1}, r_j] : |T \cap S(p, \mathbf{v})| > \frac{16}{9} \cdot \mathbb{E}[|T \cap S(p, \mathbf{v})|] \right] \leq \sum_{j=1}^{j=\ell} o \left(\frac{1}{k^{2 \cdot (4/3)^{\ell-j+1}}} \right) = o \left(\frac{1}{k^2} \right).$$

Consequently, $\Pr[\exists p \in [0, \tau(S)) : |T \cap S(p, \mathbf{v})| > \frac{16}{9} \cdot \mathbb{E}[|T \cap S(p, \mathbf{v})|]] = o(1/k^2)$. Since we have $\mathbb{E}[|T \cap S(p, \mathbf{v})|] = \frac{1}{2} \cdot |S(p, \mathbf{v})|$ we obtain $\Pr[\exists p \in [0, \tau(S)) : |T \cap S(p, \mathbf{v})| > \frac{8}{9} \cdot |S(p, \mathbf{v})|] = o(1/k^2)$. We can repeat the argument replacing the sampled set of bidders T with the unsampled set $U = N \setminus T$ and obtain that $\Pr[\exists p \in [0, \tau(S)) : |U \cap S(p, \mathbf{v})| > \frac{16}{9} \cdot \mathbb{E}[|U \cap S(p, \mathbf{v})|]] = o(1/k^2)$ and thus $\Pr[\exists p \in [0, \tau(S)) : |U \cap S(p, \mathbf{v})| > \frac{8}{9} \cdot |S(p, \mathbf{v})|] = o(1/k^2)$. Finally, since $U \cap S(p, \mathbf{v}) = S(p, \mathbf{v}) \setminus (T \cap S(p, \mathbf{v}))$, we obtain $\Pr[\exists p \in [0, \tau(S)) : |T \cap S(p, \mathbf{v})| < \frac{1}{9} \cdot |S(p, \mathbf{v})|] = o(1/k^2)$, completing the proof. $\square$

A.3 Proof of Theorem 4.1

PROOF. Our proof proceeds similarly as the one in Section 3, again employing our benchmark decomposition, inequality (1). Note that AUC_o is exactly the single-price clock auction with all prices equal to zero, which, from Section 3, we know covers the LOW-CORE, LOW-TAIL, and HIGH-TAIL portions of our benchmark. In other words, from our analysis in the previous section, we have the following two claims:

CLAIM A.1. $\text{HIGH-TAIL} \leq \mathbb{E}_{\mathbf{v}}[\text{AUC}_o(\mathbf{v})] = \max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{E}[v_i]$

and

CLAIM A.2. $\text{LOW} \leq 12 \cdot \mathbb{E}_{\mathbf{v}}[\text{AUC}_o(\mathbf{v})] = 12 \cdot \max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{E}[v_i]$.

Inequalities (1) and (2) give us an upper bound $\text{OPT} \leq \text{LOW} + \text{HIGH-TAIL} + \text{HIGH-CORE}$, where the welfare of AUC_o covers the terms LOW and HIGH-TAIL by claims A.1 and A.2: $13 \cdot \mathbb{E}_{\mathbf{v}}[\text{AUC}_o(\mathbf{v})] \geq \text{LOW} + \text{HIGH-TAIL}$. Thus

$$\text{OPT} - 13 \cdot \mathbb{E}_{\mathbf{v}}[\text{AUC}_o(\mathbf{v})] \leq \text{HIGH-CORE} \leq \text{OPT}.$$

We can now assume that HIGH-CORE is a 2 approximation to OPT, as otherwise AUC_o would be a $26 < \log \log k$ approximation to OPT. Thus it suffices to show that U-PRICE_g gives a $O(\log \log k)$ approximation to HIGH-CORE. From the proof of Lemma 3.3 we obtain Corollary 4.2 which shows that there is a single price p such that either our *revenue* will be a $O(\log \log k)$ approximation to HIGH-CORE when we uniformly post p for each bidder $i \in N$ and serve a maximal feasible set of accepting bidders, or the price p is so high that by just serving a single bidder we can get the expected *welfare* to be a $O(\log \log k)$ approximation to HIGH-CORE.

Note that $p \cdot \mathbb{E}[\max_{S \in \mathcal{S}} \sum_{i \in S} \mathbb{I}[v_i \geq p]]$ is precisely the expected revenue we get in a clock auction with a uniform price p . Corollary 4.2 shows that the best choice of a single uniform price will give an $O(\log m)$ -approximation to Δ , which translates to $O(\log \log k)$ -approximation to HIGH-CORE with $m = 10 \log k + 1$. The proof of this corollary directly follows the proof of Lemma 3.3.

Note that Corollary 4.2 proves the existence of a uniform price p that would give us the desired approximation of $O(\log \log k)$. Since we cannot calculate this price in advance, U-PRICE_g gradually sweeps through a full range of prices, aiming to find one that yields the desired revenue. However, this could fail if U-PRICE_g reaches a lower price that also satisfies the revenue target and stops too soon, selecting a set with low social welfare. We get around this problem by more careful analysis of the expected welfare of U-PRICE_g . We choose revenue goal $g = \frac{\text{OPT}}{4\alpha}$ in U-PRICE_g based on the approximation guarantee α in Corollary 4.2. Recall that

$$g = \frac{\text{OPT}}{4\alpha} \leq \frac{\text{HIGH-CORE}}{2\alpha}. \quad (6)$$

We consider three cases for U-PRICE_g . Let $T(\mathbf{v})$ be a feasible set of maximum size among the bidders who accept the price p from Corollary 4.2.

Case 1: $\mathbb{E}_{\mathbf{v}}[|T(\mathbf{v})|] \geq 8$, i.e., the expected size of the set T is large. Note that

$$|T(\mathbf{v})| = \max_{S \in \mathcal{S}} \sum_{i \in S} X_i, \text{ where } X_i \stackrel{\text{def}}{=} \mathbb{I}[v_i \geq p] \text{ are independent Bernoulli r.v. for } i \in N.$$

Thus, one can think of $|T(\mathbf{v})|$ as an XOS (maximum of additive) function $f(X_1, \dots, X_n)$ with marginal contribution of each coordinate i in $[0, 1]$. According to [58], such functions enjoy strong concentration with Chernoff-type bounds. In particular, Corollary 3.2 in [58] gives us

$$\Pr[|T| \leq \mathbb{E}[|T|]/2] \leq e^{-\mathbb{E}[|T|]/8} \leq e^{-1}.$$

Let us assume that $|T(\mathbf{v})| \geq \mathbb{E}[|T|]/2$. According to Corollary 4.2 and (6) we have $2g \leq \frac{\text{HIGH-CORE}}{\alpha} \leq p \cdot \mathbb{E}[|T|]$, i.e., $g \leq p \cdot |T(\mathbf{v})|$. In this case, U-PRICE_g obtains welfare of at least g . Indeed, the auction must stop at price p or earlier, as the revenue goal of g would have been attained at price p . Note that the auction can stop only when the revenue goal is reached, or all active bidders can be served. In the former case, the revenue of U-PRICE_g is at least g . In the latter case, all the bidders in $T(\mathbf{v})$ will be served before reaching the price p , i.e., everyone in $T(\mathbf{v})$ with the welfare of at least g . Therefore, the expected social welfare of U-PRICE_g is at least

$$\Pr[|T| \geq \mathbb{E}[|T|]/2] \cdot g \geq (1 - e^{-1})g = \frac{\text{OPT}}{O(\log \log k)}.$$

Case 2: $\Pr_{\mathbf{v}}[\text{U-PRICE}_g \text{ serves set } F : \sum_{i \in F} v_i \geq g] \geq \frac{1}{2}$. Then the expected welfare of U-PRICE_g is at least $g/2 = \frac{\text{OPT}}{O(\log \log k)}$.

Case 3: $\mathbb{E}_{\mathbf{v}}[|T(\mathbf{v})|] < 8$ and $\Pr_{\mathbf{v}}[\text{U-PRICE}_g \text{ serves set } F : \sum_{i \in F} v_i \geq g] < \frac{1}{2}$. Thus, we stop at set F with $\sum_{i \in F} v_i < g$ with probability more than $1/2$. When $\sum_{i \in F} v_i < g$, U-PRICE_g must be serving the whole set A of active bidders. In particular, it means that a bidder $j \in \arg\max_{i \in N} v_i$ must be served in A , as we can assume without loss of generality that each bidder in N can be served as a singleton. We show in the following Claim A.3 that the expected welfare of U-PRICE_g must be at least half of the expected maximum.⁷

CLAIM A.3. Assume $\Pr_{\mathbf{v}}[\text{U-PRICE}_g \text{ serves set } F : \sum_{i \in F} v_i \geq g] < \frac{1}{2}$. Then the expected social welfare of U-PRICE_g is at least $\frac{1}{2} \mathbb{E}_{\mathbf{v}}[\max_i v_i]$.

PROOF. Sample $\mathbf{v} \sim \mathbf{D}$ and fix the bidder $i = \arg\max_j v_j$ with the maximum value $v_i = t$. Assuming that i is the bidder with maximum value t , it suffices to show that U-PRICE_g serves i with probability at least $1/2$ for any $i \in N$ and $t \geq 0$. To this end, consider independently sampling valuation profile $\mathbf{v}'_{-i} \sim \mathbf{D}'_{-i} : \forall j \neq i, v'_j \sim D_j \mid v'_j < t$. Note that (i) $\mathbf{v}'_{-i}$ has exactly the same distribution as $\mathbf{v}_{-i}$ conditioned on $v_i = t$ being the maximum in $\mathbf{v}$ and (ii) $\mathbf{D}'_{-i}$ is stochastically dominated by $\tilde{\mathbf{v}}_{-i} \sim \mathbf{D}_{-i} = \times_{j \neq i} D_j$, i.e., we can do probability coupling between $\mathbf{v}'_{-i} \sim \mathbf{D}'_{-i}$ and $\tilde{\mathbf{v}}_{-i} \sim \mathbf{D}_{-i}$, so that each $\mathbf{v}'_{-i}$ is coordinate-wise smaller than $\tilde{\mathbf{v}}_{-i}$ ($\forall j \neq i, v'_j \leq \tilde{v}_j$).

By the Claim's assumption the social welfare of U-PRICE_g on a random valuation profile $(\tilde{\mathbf{v}}_{-i}, \tilde{v}_i) \sim \mathbf{D}$ is less than g with probability more than $1/2$. Now, if this welfare is less than g , then the revenue of offering a uniform price q must be smaller than g for any price $q \geq 0$, i.e.,

$$\forall q \geq 0, q \cdot |\{j \in N : \tilde{v}_j \geq q\}| < g \implies \forall q \geq 0, q \cdot |\{j \neq i : v'_j \geq q\}| < g. \quad (7)$$

⁷Despite the fact that we serve the maximum bidder with probability more than $1/2$, we still need to show that U-PRICE_g gets sufficiently high expected welfare.

Therefore, U-PRICE_g must serve bidder i on the valuation profile $(\mathbf{v}'_{-i}, t)$, as (a) i has the maximum bid and would always stay active; (b) we cannot reach the revenue goal of g with a uniform price q without using bidder i according to (7). This concludes the proof, as the probability of having a $(\tilde{\mathbf{v}}_{-i}, \tilde{v}_i) \sim \mathbf{D}$ such that the welfare of $\text{U-PRICE}_g(\tilde{\mathbf{v}}_{-i}, \tilde{v}_i) < g$ is greater than half and we will serve i on each corresponding profile $(\mathbf{v}'_{-i}, t)$. $\square$

Now we can conclude the analysis of case 3. Claim A.3 gives a lower bound of $\frac{1}{2} \mathbb{E}_{\mathbf{v}}[\max_i v_i]$ on the expected welfare of U-PRICE_g . When $\mathbb{E}[|T|] < 8$, the expected maximum $\mathbb{E}_{\mathbf{v}}[\max_i v_i]$ is a constant approximation to the revenue $p \cdot \mathbb{E}[|T|]$ of U-PRICE_g . On the other hand, $\frac{1}{2} \mathbb{E}_{\mathbf{v}}[\max_i v_i]$ is also a constant approximation to the other term $\mathbb{E}[\max_{i \in N} \mathbb{E}[v_i \mid v_i \geq p] \cdot \mathbb{I}[v_i \geq p]]$ in Lemma 4.2. $\square$

A.4 Single-price clock auctions for the distributions with constant support sizes

Let $\mathcal{F}$ be any downward-closed family of feasible sets over the set of N bidders. We assume that all possible values v_i of each bidder $i \in N$ can be described by a list of size at most ℓ , i.e., the support $\text{supp}(D_i)$ of each distribution D_i has size at most ℓ . We now give the proof of Claim 7.2.

PROOF. Let us write all possible types in each distribution D_i in the increasing order $(\theta_{i,1}, \dots, \theta_{i,\ell})$. Consider the following ℓ single-price clock auctions: for a given $j \in \{1, 2, \dots, \ell\}$

- (1) set price $p_i = \theta_{i,j}$ for each $i \in N$.
- (2) among the bidders who accepted their price select $S \subset \{i : v_i \geq \theta_{i,j}\}$ with maximal revenue

$$\max_{S \in \mathcal{F}} \sum_{i \in S} \theta_{i,j}.$$

Our single-price clock auction is the one among these ℓ auctions with the highest expected revenue⁸, which can be written as

$$\max_{j \in [\ell]} \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{F}} \sum_{i \in S} \theta_{i,j} \cdot \mathbb{I}[v_i \geq \theta_{i,j}] \right].$$

We have the following upper bound on the benchmark.

$$\begin{aligned} \text{OPT} &= \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i \right] = \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i \sum_{j=1}^{\ell} \mathbb{I}[v_i = \theta_{i,j}] \right] \leq \mathbb{E}_{\mathbf{v}} \left[\sum_{j=1}^{\ell} \max_{S \in \mathcal{F}} \sum_{i \in S} v_i \cdot \mathbb{I}[v_i = \theta_{i,j}] \right] \\ &= \sum_{j=1}^{\ell} \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{F}} \sum_{i \in S} \theta_{i,j} \cdot \mathbb{I}[v_i = \theta_{i,j}] \right] \leq \ell \cdot \max_{j \in [\ell]} \mathbb{E}_{\mathbf{v}} \left[\max_{S \in \mathcal{F}} \sum_{i \in S} \theta_{i,j} \cdot \mathbb{I}[v_i = \theta_{i,j}] \right], \quad (8) \end{aligned}$$

where to get the first inequality we used that for each valuation profile $\mathbf{v}$ the maximal sum $\sum_{i \in S} v_i$ for $S \in \mathcal{F}$ can be covered by ℓ sums $\sum_{i \in S} v_i \cdot \mathbb{I}[v_i = \theta_{i,j}]$. To conclude the proof we notice that for each valuation profile $\mathbf{v}$ and feasible $S \in \mathcal{F}$

$$\sum_{i \in S} \theta_{i,j} \cdot \mathbb{I}[v_i = \theta_{i,j}] \leq \sum_{i \in S} \theta_{i,j} \cdot \mathbb{I}[v_i \geq \theta_{i,j}],$$

i.e., the last term in (8) is not more than ℓ times the revenue of our single-price clock auction, which is not more than the welfare. $\square$

⁸Note that we can write the revenue of the auction in this way because $\mathcal{F}$ is downward-closed set system.

B CONNECTIONS TO LARGE DEVIATION THEORY

As we mentioned in Section 7, in cases beyond constant support size and binary valuations, the question of whether single-price clock auctions can achieve a constant approximation is quite non-trivial. In fact, the problem remains non-trivial even if we restrict attention to the very special case of Disjoint-Maximal-Sets with equal size maximal sets (i.e., k disjoint maximal sets of $m = n/k$ bidders each) and i.i.d. bidders (i.e., $D_i = D$ for all $i \in N$). This question appears to be related to the *large deviation theory* [23] and, particularly, the distributions of the values of sums of random variables when the values of the sums far exceed their expectations.

For instances where all k maximal disjoint sets have the same size m , and all $n = km$ bidders have their v_i values drawn i.i.d., the expected optimal welfare is equal to the maximum taken over k independent samples of $\xi(m) = \sum_{i=1}^m v_i$. We are interested in the case where this maximum is significantly larger than the expectation of $\mathbb{E}[\xi(m)]$ which only happens for large $k = \Theta(2^m)$. In this case, the expectation of this maximum is roughly equal to the value of the top k -th quantile of $\xi(m)$. Note that in the limit when $k = \Theta(2^m)$ and $m \rightarrow \infty$, the tail probability for the distribution of $\xi(m)$ is captured by Cramer's theorem (a fundamental result in large deviation theory). According to this theorem, there is a certain rate function I such that $\xi(m)$ satisfies a *large deviation principle with rate function I* , meaning that $\Pr\left[\frac{\xi(m)}{m} \geq x\right] = e^{-I(x) \cdot m \cdot (1+o(1))}$.

On the other hand, in a single-price clock auction with price vector $\mathbf{p}$, the agents with $v_i < p_i$ drop out, so the event that $v_i \sim D_i$ lies in the “tail” of D_i ($v_i \geq p_i$) is instead captured by the Bernoulli random variable $\hat{v}_i = \mathbb{I}[v_i \geq p_i] \cdot \mathbb{E}[v_i \mid v_i \geq p_i]$. Given the set of bidders that accepted the prices offered to them, the obvious choice is then to serve the feasible subset with the largest sum of prices, leading to an expected social welfare of $\max_{S \in \mathcal{S}} \sum_{i \in S} \hat{v}_i$.

We can think of the variable $\hat{v}_i$ as a *binary approximation* of the random variable v_i up to an additive error⁹ of $O(\mathbb{E}[v_i])$. Formally, a *binary approximation* of a random variable $\xi_i : \Omega_i \rightarrow \mathbb{R}$ on a probability space Ω_i is the following (generalized) Bernoulli random variable $\bar{\xi}_i$ defined for an event $\mathcal{E}_i \subset \Omega_i$.

$$\bar{\xi}_i : \mathbb{E}[\xi_i \mid \mathcal{E}_i] \text{ w.p. } \Pr[\mathcal{E}_i] \quad , \quad \mathbb{E}[\xi_i \mid \bar{\mathcal{E}}_i] \text{ w.p. } \Pr[\bar{\mathcal{E}}_i]$$

Then, to achieve $O(1)$ approximation single-price clock auctions, we need to cover the tail probabilities of $\frac{1}{m} \sum_{i=1}^m v_i$ in the large deviation regimes by the constantly related tail probabilities of $\frac{1}{m} \sum_{i=1}^m \hat{v}_i$. In the language of large deviation theory, we need to cover the tail probability of the sum of i.i.d. random variables by the tail probability of the sum of binary approximations. This is captured in the following claim.

CLAIM B.1. *Assume that there is a single-price clock auction for the Disjoint-Maximal-Sets with k sets of size m and $n = k \cdot m$ i.i.d. bidders that is $O(1)$ approximation $\forall k, m \in \mathbb{N}$. Then for any collection $(\xi_i)_{i=1}^m$ of m non-negative i.i.d. random variables and for every $x > O(\mathbb{E}[\xi_i])$ there must exist binary approximations $(\bar{\xi}_i)_{i=1}^m$ of $(\xi_i)_{i=1}^m$ such that*

$$\Pr\left[\frac{\sum_{i=1}^m \bar{\xi}_i}{m} \geq x \cdot \Omega(1)\right] \geq \Pr\left[\frac{\sum_{i=1}^m \xi_i}{m} \geq x\right].$$

We note that the binary approximations $\bar{\xi}_i$ in Claim B.1 may not be identical. If, on the other hand, we could achieve the constant approximation using a uniform single-price clock auction (which is reasonable to expect in the i.i.d. setting), then this would imply a stronger Claim B.1 with identical binary approximations $(\bar{\xi}_i)_{i=1}^m$.

⁹As a single-price clock auction with 0 prices achieves in expectation the welfare of $m \cdot \mathbb{E}[v_i]$, we can safely ignore the additive terms of order $O(\mathbb{E}[v_i])$ for each bidder i .