



The Competition Complexity of Prophet Inequalities

JOHANNES BRÜSTLE, Sapienza University of Rome, Italy

JOSÉ CORREA, Universidad de Chile, Chile

PAUL DÜTTING, Google Research, Switzerland

TOMER EZRA, Harvard University, USA

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

VICTOR VERDUGO, Pontificia Universidad Católica de Chile, Chile

We study the classic single-choice prophet inequality problem through a resource augmentation lens. Our goal is to bound the $(1 - \epsilon)$ -competition complexity of different types of online algorithms. This metric asks for the smallest k such that the expected value of the online algorithm on k copies of the original instance, is at least a $(1 - \epsilon)$ -approximation to the expected offline optimum on a single copy.

We show that block threshold algorithms, which set one threshold per copy, are optimal and give a tight bound of $k = \Theta(\log \log 1/\epsilon)$. This shows that block threshold algorithms approach the offline optimum doubly-exponentially fast. For single threshold algorithms, we give a tight bound of $k = \Theta(\log 1/\epsilon)$, establishing an exponential gap between block threshold algorithms and single threshold algorithms.

Our model and results pave the way for exploring resource-augmented prophet inequalities in combinatorial settings. In line with this, we present preliminary findings for bipartite matching with one-sided vertex arrivals, as well as in XOS combinatorial auctions. Our results have a natural competition complexity interpretation in mechanism design and pricing applications.

Additional Key Words and Phrases: Prophet Inequalities, Competition Complexity

ACM Reference Format:

Johannes Brüstle, José Correa, Paul Dütting, Tomer Ezra, Michal Feldman, and Victor Verdugo. 2024. The Competition Complexity of Prophet Inequalities. In *The 25th ACM Conference on Economics and Computation (EC '24)*, July 8–11, 2024, New Haven, CT, USA. ACM, New York, NY, USA, 24 pages. <https://doi.org/10.1145/3670865.3673467>

1 INTRODUCTION

Studying algorithms beyond the worst case is an increasingly popular approach in theoretical computer science, and is obviously of great practical interest. In this work, we bring together two such directions, which are particularly relevant for online algorithms: The prophet inequality paradigm and resource augmentation. The former restricts the input to be stochastic —rather than worst-case— thus enhancing the online algorithm by giving it better information. The latter restricts the optimal algorithm to use fewer resources, thus degrading its performance.

The standard prophet inequality [Krengel and Sucheston, 1977, 1978] can be interpreted as an allocation problem. There is a single item, which we need to allocate to one of n players that arrive one-by-one in an online fashion. Each player has a value $v_i \geq 0$ for the item. The values v_1, \dots, v_n are drawn independently from known distributions F_1, \dots, F_n . We compare the expected value



This work is licensed under a Creative Commons Attribution International 4.0 License.

EC '24, July 8–11, 2024, New Haven, CT, USA

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0704-9/24/07

<https://doi.org/10.1145/3670865.3673467>

achievable by an online algorithm that has to allocate the item in an online fashion, to the expected offline optimum, which can simply choose the maximum value in the sequence.

The classic result in this area is a tight factor 2 approximation, which can be achieved with a simple (single) threshold algorithm [Samuel-Cahn, 1984]. By interpreting the threshold as a price and the players as buyers, this result has immediate implications for several economic and mechanism design applications. It implies that a simple (sequential) posted-price mechanism yields a constant factor approximation to the optimal mechanism [Chawla et al., 2010, Hajiaghayi et al., 2007]. Driven in part by this connection to market pricing, the last decade has seen an impressive amount of work on combinatorial extensions of the prophet inequality problem (e.g., for matroids [Alaei, 2011, Chawla et al., 2010, Feldman et al., 2016, Kleinberg and Weinberg, 2012], general downward-closed set systems [Rubinstein and Singla, 2017], or combinatorial auctions [Correa and Cristi, 2023, Dütting et al., 2017, Dütting et al., 2020, Ezra et al., 2020, Feldman et al., 2015]).

Resource augmentation, in turn, has a long history of success stories in the design and analysis of algorithms. It is particularly popular in scheduling [Kalyanasundaram and Pruhs, 2000, Phillips et al., 1997], but it has also been studied in many other areas such as paging [Sleator and Tarjan, 1984], bin packing [Csirik and Woeginger, 2000], or selfish routing [Roughgarden and Tardos, 2002]. There are typically several plausible notions of resource augmentation, which illuminate different aspects of the problem at hand. In the scheduling literature, for example, we grant the online algorithm access to additional machines. One can also assume it has access to faster machines. In paging, one can analyze the online algorithm with a larger cache relative to that used by the offline algorithm; while in bin packing, one may consider giving the online algorithms larger bins.

In the context of online allocation described above, we grant the online algorithm access to more samples from the same distributions, say the online algorithm can choose to allocate the item to one of nk players, where for each distribution F_i , there are k players whose value is drawn from this distribution. The question we seek to answer is that of understanding how many additional resources (samples) the online algorithm needs in order to effectively approximate the performance of the offline optimum in this stochastic setting.

1.1 The Model

In this work, we take a resource augmentation approach to prophet inequalities. We study the prophet inequality problem in a setting, where the offline algorithm is handicapped by having less allocation opportunities. This is particularly well motivated in the mechanism design and pricing applications, where it is very likely that the comparatively simpler (sequential) posted-price mechanism attracts additional buyers compared to the number of buyers that would show up if one was to sell the item through an auction.

In our model, the online algorithm sees k copies of a prophet inequality instance, where every instance has n values v_1, \dots, v_n drawn independently from distributions F_1, \dots, F_n . We compare the expected value achieved by the online algorithm on k copies, to the expected maximum value on a single copy. Our default assumption is a fixed order model, in which the variables arrive in the same order in each block. Our results continue to hold when variables arrive in arbitrary order within each block, and each block may have its own arrival order.

We also consider a further relaxation, the γ -displacement model, which is parameterized by an integer $\gamma \geq 1$. In this model, the algorithm faces γk copies of the original instance, and an adversary can determine the arrival order, but is constrained by the fact that in each meta-block of γn variables, each type of variable should appear at least once. This model interpolates between the block model with arbitrary intra-block arrivals and the fully adversarial model.

Similarly to Eden et al. [2017] and Beyhaghi and Weinberg [2019], for $\varepsilon \geq 0$, we define the $(1 - \varepsilon)$ -competition complexity as the smallest k such that the expected value of the online algorithm

is guaranteed to be at least a $(1 - \varepsilon)$ fraction of the expected maximum value for every instance. This complexity measure has previously been studied by Brustle et al. [2022] in the context of prophet inequalities and posted pricing, but only for the case of i.i.d. distributions. They show that while for $\varepsilon = 0$ the competition complexity is unbounded, for $\varepsilon > 0$ it scales as $\Theta(\log \log 1/\varepsilon)$. This shows that the optimal online algorithm (dynamic pricing policy), approaches the optimal offline algorithm (optimal auction) doubly-exponentially fast.

Of course, the competition complexity metric is also interesting when restricted to a certain class of online algorithms. We consider algorithms from three different classes: (1) Single threshold algorithms, which set a single threshold and accept the first value that exceeds the threshold. (2) Block threshold algorithms, which set a threshold for each copy, and accept the first value that exceeds the threshold for its copy. (3) General threshold algorithms, which set a threshold for each of the nk steps and accept the first value that exceeds its threshold.

Naturally, a simple backwards induction argument shows that threshold algorithms are optimal (in a per instance sense) among all online algorithms. More interestingly, as a first structural insight, we show that with respect to the competition complexity metric, block threshold algorithms are optimal (Proposition 1 in Section 2). We thus generalize the classic result of Samuel-Cahn [1984], which shows that for $k = 1$ a single threshold algorithm is optimal. Apart from this, threshold algorithms are of particular importance because of their simplicity and natural interpretation as posted-price mechanisms.

1.2 Our Contribution

As our main contribution, we resolve the $(1 - \varepsilon)$ -competition complexity for the classic (single-choice) prophet inequality problem with non-identical distributions. We show that the competition complexity for general distributions exhibits the same asymptotics as in the i.i.d. case. Moreover, the optimal asymptotics are attained by block threshold algorithms.

Main Result 1 (Theorem 1): For every $\varepsilon > 0$, the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms is $\Theta(\log \log(1/\varepsilon))$.

Our second main result shows a tight bound for single threshold algorithms.

Main Result 2 (Theorem 2): For every $\varepsilon > 0$, the $(1 - \varepsilon)$ -competition complexity of the class of single threshold algorithms is $\Theta(\log(1/\varepsilon))$.

We thus show that also in the case of general, non-identical distributions the best online algorithm (dynamic pricing policy) approaches the best offline algorithm (optimal auction) doubly exponentially fast. Moreover, this is true even when we don't use the full power of dynamic pricing, and prices remain constant within each block.

In addition, our results show that there is an exponential gap between single threshold algorithms (static pricing policies) and block threshold algorithms (policies that update prices only periodically).

1.3 Our Techniques

Our main technical contribution is the upper bound of $O(\log \log(1/\varepsilon))$ on the $(1 - \varepsilon)$ -competition complexity for block threshold algorithms. The matching lower bound of $\Omega(\log \log(1/\varepsilon))$ already applies in the i.i.d. setting and follows from Brustle et al. [2022].

The main technical ingredient in our proof of the upper bound is an approximate stochastic dominance inequality that allows us to evaluate the performance of any block threshold algorithm, with decreasing thresholds $\tau_1 > \tau_2 > \dots, \tau_k$ (Lemma 1). The approximation factor of this stochastic dominance inequality is parameterized by τ_1, \dots, τ_k , and it is obtained through a careful analysis

of the minimum stochastic-dominance approximation factor achievable on each of the $k + 1$ sub-intervals $[0, \tau_k), [\tau_k, \tau_{k-1}), \dots, [\tau_1, \infty)$ defined by the block thresholds. Then, to find the best possible approximation guarantee for block threshold algorithms that can be obtained in this way, we need to solve a (high-degree polynomial) max-min optimization problem. Rather than solving this problem exactly, we provide a tight double-exponentially fast increasing lower bound on the value of this max-min problem by constructing an explicit set of thresholds (Lemma 2). In combination with the lower bound, this shows that our stochastic dominance approach provides an optimal competition complexity bound up to a constant factor.

Our upper bound of $O(\log(1/\varepsilon))$ for single threshold algorithms follows in a rather direct way from the “median rule” proof of Samuel-Cahn [1984]. Our key insight for this case is that the upper bound is asymptotically tight, which we show by providing an explicit lower bound construction.

1.4 Other Arrival Orders

We also explore the robustness of the $(1 - \varepsilon)$ -competition complexity metric to different assumptions about the arrival order. To study this effect, we introduce the γ -displacement model, where $\gamma \geq 1$ is a parameter. In this model, the algorithm faces γk copies of the original instance. The arrival order is determined by an adversary, but the adversary is constrained by the requirement that, within each meta-block of γn variables, each type of variable should appear at least once. While for $\gamma = 1$ the adversary is restricted to move variables within their block, for $\gamma > 1$ the adversary can also move variables across blocks.

For this model, we show that for every $\gamma \geq 1$, the $(1 - \varepsilon)$ -competition complexity of block threshold algorithms is $O(\gamma \log \log(1/\varepsilon))$ (Proposition 6). Comparing this with the $(1 - \varepsilon)$ -competition complexity of block threshold algorithms in the default model, this shows that the competition complexity is increased by a multiplicative factor of γ , but the scaling behavior in $1/\varepsilon$ remains the same. This shows that the competition complexity of block threshold algorithms degrades gracefully as we move away from the default model.

We also show that a comparable result cannot be achieved for single threshold algorithms. Namely, for this class of algorithms, we show that there exists some $\gamma > 1$, such that the $(1 - \varepsilon)$ -competition complexity of single threshold algorithms is least $\Omega(1/\varepsilon^{1/3})$ (Proposition 7). This shows that the $(1 - \varepsilon)$ -competition complexity of this type of algorithms transitions from growing logarithmically in $1/\varepsilon$ when $\gamma = 1$ to growing (at least) polynomially in $1/\varepsilon$ when $\gamma > 1$.

We complement these results, with a lower bound on the $(1 - \varepsilon)$ -competition complexity of general threshold algorithms in the fully adversarial model, showing that the $(1 - \varepsilon)$ -competition complexity is at least $\Omega(1/\varepsilon)$ (Proposition 8).

1.5 Extensions

Our work opens up the question of studying resource-augmented prophet inequalities for richer combinatorial settings. We present some preliminary results for submodular and XOS combinatorial auctions. The proper generalization of threshold algorithms for this setting are prices, and similar to the distinction between block threshold algorithms and single threshold algorithms, we can distinguish between block-consistent prices which stay fixed within a block and static prices that remain fixed throughout.

Additional Result 1 (in full version): The $(1 - \varepsilon)$ -competition complexity of block-consistent prices for submodular and XOS combinatorial auctions is $O(\log(1/\varepsilon))$.

Additional Result 2 (in full version): The $(1 - \varepsilon)$ -competition complexity of static prices for submodular and XOS combinatorial auctions is $O(1/\varepsilon)$.

Our results present a first glimpse at a potentially rich theory and already show that the constant-factor that can be shown in the single-shot setting [Dütting et al., 2017, Feldman et al., 2015] vanishes exponentially fast with block-consistent prices with additional resources. Whether this can also be achieved with static prices remains open, just as the question of whether a double-exponentially fast approach is possible with dynamic prices. More generally, it would be interesting to establish a formal separation between static and dynamic prices. It also remains open whether comparable results can be obtained for subadditive combinatorial auctions [Correa and Cristi, 2023, Dütting et al., 2020].

Of course, the same question can be studied for other combinatorial settings such as matroids [Chawla et al., 2010, Kleinberg and Weinberg, 2012] or matching constraints [Ezra et al., 2022, Gravin and Wang, 2019]. Finally, we hope that our paper will spark work on other notions of resource augmentation.

1.6 Related Work

A seminal paper in Economic Theory that is very relevant to our work is a paper by Bulow and Klemperer [1996]. This paper argues that in an i.i.d. auction setting rather than going with the complex revenue-maximizing auction, one can attract one additional buyer and simply run a second price auction. The revenue that results from the simpler auction with one additional buyer will be higher than that from the optimal auction.

This work obviously found its audience in computer science, where it inspired work on the competition complexity of simple mechanisms, see for example the excellent work of Eden et al. [2017], Beyhaghi and Weinberg [2019], and Babaioff et al. [2020].

Our work is also related to that of Abolhassani et al. [2017] and Liu et al. [2021] through the notion of “frequent instances”. Both of these papers study the *frequent prophets* problem, in which each distribution must be repeated at least some number of times: An instance is m -frequent if each distribution appears at least m times. The inspiration behind frequent prophets comes from the difficulty of directly analyzing the general independent case in random and free-order models, in which the values either arrive in uniform random order or the algorithm is free to choose the order of observation, and the idea is to bring the instance closer to the i.i.d. case. Abolhassani et al. [2017] show how to design a 0.738-competitive policy for $O(1)$ -frequent instances in the free-order model and $\Theta(\log n)$ -frequent instances in the random-order model, while Liu et al. [2021] improve on the random-order results by showing how to obtain a $(\beta - \epsilon)$ -competitive policy for $O_\epsilon(1)$ -frequent instances, where $\beta \approx 0.745$ is the optimal guarantee for the i.i.d. setting [Correa et al., 2021]. The difference to our work is that they study both ALG and OPT on “frequent instances”, while we compare the algorithm in a k -frequent instance with the prophet on a standard (1-frequent) instance. Another difference is that we work in the fixed order setting, whereas their results apply to the random and free-order models.

1.7 Organization

The paper is organized as follows. We formally state our model and the competition complexity metric in Section 2. We also introduce single threshold algorithms and block threshold algorithms, and establish the optimality of the latter. In Section 3 we present our results for block threshold algorithms. Afterwards, in Section 4, we turn our attention to single threshold algorithms. We discuss combinatorial extensions in Section 5 and the full version. We defer the discussion of additional arrival orders to Appendix B.

2 PRELIMINARIES

The Prophet Inequality Problem. Consider the following game between an online algorithm (“gambler”) and an offline algorithm (“the prophet”). The online algorithm gets to observe a sequence of n (non-negative) numbers v_1, \dots, v_n one-by-one. Each v_i is drawn independently from a known distribution $F_i \in \Delta$, where Δ is the set of all distributions over $\mathbb{R}_{\geq 0}$. We call a sequence of distributions F_1, \dots, F_n an instance.

We refer to the online algorithm as ALG. Upon seeing a value v_i the online algorithm has to immediately and irrevocably decide whether to accept the current value v_i and stop the game, or to proceed to the next value v_{i+1} . Every online algorithm induces a stopping time $\rho \in [n] \cup \{\text{null}\}$, where ρ is the index of the value chosen by the online algorithm (to handle the case where the online algorithm does not accept any value, we set $\rho = \text{null}$ and interpret v_{null} as zero). The algorithm’s expected reward is $\mathbb{E}[\text{ALG}(v)] = \mathbb{E}[v_\rho]$. The offline algorithm, in contrast, can see the entire sequence of values v_1, \dots, v_n at once and can simply choose the maximum value $\max_{i \in [n]} v_i$. The offline algorithm’s expected reward is $\mathbb{E}[\max_{i \in [n]} v_i]$.

In the Prophet Inequality problem the online algorithm is evaluated by its *competitive ratio*, defined as the worst-possible ratio (over all instances) between the online algorithm’s expected reward $\mathbb{E}[\text{ALG}(v)] = \mathbb{E}[v_\rho]$ and the offline algorithm’s expected reward $\mathbb{E}[\max_{i \in [n]} v_i]$. Let $\alpha \in [0, 1]$. An online algorithm is α -competitive if

$$\inf_{F_1, \dots, F_n \in \Delta} \frac{\mathbb{E}_{v \sim (F_1, \dots, F_n)} [\text{ALG}(v)]}{\mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i]} \geq \alpha.$$

The Competition Complexity Benchmark. Our goal is to compare the expected performance of an online algorithm on k independent copies of the original instance, to the expected value of the offline algorithm on a single instance. This can be done under different assumptions on how the nk values of the k copies are presented to the online algorithm.

Our default model is what we call the *block model*. In this model, the online algorithm sees the k copies of the original instance one after the other. We refer to each copy as a block. Within each block, the n values arrive in the same order as in the original instance.

More formally, when (F_1, \dots, F_n) is the original instance, we denote with $(F_1, \dots, F_n)^k$ the instance with k copies. The input to the online algorithm consists of kn numbers

$$v_1^{(1)}, \dots, v_n^{(1)}, v_1^{(2)}, \dots, v_n^{(2)}, \dots, v_1^{(k)}, \dots, v_n^{(k)},$$

where each $v_i^{(j)}$ for $i \in [n]$ and $j \in [k]$ is an independent draw from F_i . The offline algorithm receives only n numbers v_1, \dots, v_n where each v_i for $i \in [n]$ is an independent draw from F_i .

Let \mathcal{A} be a family of online algorithms. Let $\mathcal{A}_{n,k}$ be defined as all the online algorithms in \mathcal{A} that are defined on $\Delta_{n,k} = \{(F_1, \dots, F_n)^k \mid F_1, \dots, F_n \in \Delta\}$.

Definition 1 (Competition complexity). *Given $\varepsilon \geq 0$, the $(1 - \varepsilon)$ -competition complexity with respect to a class of algorithms \mathcal{A} is the smallest positive integer number $k(\varepsilon)$ such that for every n , every $F_1, \dots, F_n \in \Delta$, and every $k \geq k(\varepsilon)$, it holds that*

$$\max_{\text{ALG} \in \mathcal{A}_{n,k}} \mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}(v)] \geq (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i].$$

The case $\varepsilon = 0$ is also referred to as *exact competition complexity*; it was shown by Brüstle et al. [Brüstle et al., 2022, Theorem 2.1] that the exact competition complexity is unbounded even for the i.i.d. case (namely, where $F_i = F_j$ for all i, j). So, naturally, our focus will be on the case $\varepsilon > 0$.

We remark that, as is common in the literature, we focus on a worst-case notion of competition complexity, which asks for the minimum number of copies that suffices for a worst-case approximation guarantee. In Appendix C, we explore an alternative, which asks for the expected number of copies that are required to achieve this.

Classes of Online Algorithms. We are particularly interested in three types of online algorithms. In order of increasing generality these are:

- A *single threshold algorithm* is defined by a single threshold τ and it accepts the first value v_ℓ (indexed according to arrival order) that is at least τ .
- A *block threshold algorithm* sets k thresholds $\tau = (\tau_1, \dots, \tau_k)$, i.e., one threshold per block. It accepts the first value $v_i^{(j)}$ that is larger than the threshold τ_j for its block.
- A *general threshold algorithm* sets nk thresholds $\tau = (\tau_1, \dots, \tau_{nk})$, and accepts the first value v_ℓ (again, indexed according to arrival order) that is at least τ_ℓ .

In all of the above cases, we allow the algorithm to accept a value only with a certain probability in case it exactly meets the threshold. This ability to randomize is relevant only for distributions with point masses. A standard backward-induction argument shows that the optimal online algorithm is a general threshold algorithm.

In the classic prophet inequality problem, a single threshold algorithm attains the best possible competitive ratio [Samuel-Cahn, 1984]. Proposition 1 generalizes this result, and shows that for the more general competition complexity benchmark, it is without loss to focus on block threshold algorithms. The main idea behind this reduction is that finding the competition complexity of general threshold algorithms is, essentially, equivalent to understanding the competition complexity for instances where every distribution is a weighted Bernoulli. For these instances, and within each block, only non-zero values of a suffix of the block are chosen, and therefore it is sufficient to implement a single threshold per block. This implies that block threshold algorithms are as powerful as general threshold algorithms for the block model.

Proposition 1. *For every $\varepsilon \in (0, 1)$, the $(1 - \varepsilon)$ -competition complexity with respect to the class of block threshold algorithms is the same as the $(1 - \varepsilon)$ -competition complexity with respect to the class of general threshold algorithms.*

PROOF. Let k' be the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms, and let k be the $(1 - \varepsilon)$ -competition complexity of the class of general threshold algorithms. Clearly $k' \geq k$, since every block threshold algorithm is a general threshold algorithm. We next prove the other inequality. By the definition of k' , there exists an instance (F_1, \dots, F_n) , such that for every block threshold algorithm ALG, it holds that

$$\mathbb{E}_{v \sim (F_1, \dots, F_n)^{k'-1}} [\text{ALG}(v)] < (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i]. \tag{1}$$

Given a block threshold algorithm ALG, let

$$\alpha = \frac{1}{4} \left((1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i] - \mathbb{E}_{v \sim (F_1, \dots, F_n)^{k'-1}} [\text{ALG}(v)] \right),$$

and let

$$\beta^* = \inf \left\{ \beta \geq 0 : \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \cdot \mathbf{1}[\max_{i \in [n]} v_i \geq \beta]] \leq \alpha \right\}.$$

Let $m = \lceil \beta^* / \alpha \rceil$, and for every $i \in [m]$, let $p_i = \Pr[(i - 1) \cdot \alpha \leq \max_{i \in [n]} v_i < i \cdot \alpha]$.

For every $i \in \{1, \dots, m - 1\}$, let D_i be the weighted Bernoulli distribution that takes the value $i \cdot \alpha$ with probability $p_{i+1} / (1 - \sum_{j=i+2}^m p_j)$, and zero otherwise. Let $M_1 = \max_{i \in [m-1]} w_i$ with

$w \sim (D_1, \dots, D_{m-1})$, and $M_2 = \max_{i \in [n]} v_i$ with $v \sim (F_1, \dots, F_n)$. It holds that M_1 has the same distribution as $\lfloor M_2 \cdot \mathbf{1}[M_2 < m \cdot \alpha] / \alpha \rfloor \cdot \alpha$, since for every $r \in [m-1]$ we have

$$\begin{aligned} \Pr[M_1 = r \cdot \alpha] &= \frac{p_{r+1}}{1 - \sum_{j=r+2}^m p_j} \cdot \prod_{r'=r+1}^{m-1} \left(1 - \frac{p_{r'+1}}{1 - \sum_{j=r'+2}^m p_j} \right) \\ &= \frac{p_{r+1}}{1 - \sum_{j=r+2}^m p_j} \cdot \prod_{r'=r+1}^{m-1} \left(\frac{1 - \sum_{j=r'+1}^m p_j}{1 - \sum_{j=r'+2}^m p_j} \right) \\ &= p_{r+1} = \Pr \left[\left\lfloor \frac{M_2 \cdot \mathbf{1}[M_2 < m \cdot \alpha]}{\alpha} \right\rfloor \cdot \alpha = r \cdot \alpha \right]. \end{aligned} \quad (2)$$

Thus, for the instance (D_1, \dots, D_{m-1}) , it holds that

$$\mathbb{E}_{w \sim (D_1, \dots, D_{m-1})} [\max_{i \in [m-1]} w_i] \geq \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i] - 2\alpha, \quad (3)$$

since an α term is lost due to values above $m \cdot \alpha \geq \beta^*$, and another α term is lost due to the flooring. On the other hand, since the thresholds calculated by the best general threshold algorithm for $(D_1, \dots, D_{m-1})^{k'-1}$ are monotonically decreasing, and since the distributions of each block are weighted Bernoulli variables with increasing weights, it holds that within each block, only non-zero values of a suffix of the block are chosen. Thus, the optimal algorithm ALG^* for instance $(D_1, \dots, D_{m-1})^{k'-1}$ is a block threshold algorithm.

Let $\tau_1^* \geq \dots \geq \tau_{k'-1}^*$ be the monotone decreasing thresholds used by ALG^* . Consider the algorithm ALG' that selects a value in block ℓ if it is in the interval $[\alpha \cdot (1 + \lfloor \tau_\ell^* / \alpha \rfloor), m \cdot \alpha)$. By Equation (2), the probability that ALG' selects an element in every iteration is the same as ALG^* , and given that both algorithms stop at block ℓ , the expectation of ALG' is at least the expectation of ALG^* . Thus,

$$\mathbb{E}_{w \sim (D_1, \dots, D_{m-1})^{k'-1}} [\text{ALG}^*(w)] \leq \mathbb{E}_{v \sim (F_1, \dots, F_n)^{k'-1}} [\text{ALG}'(v)] \leq \mathbb{E}_{v \sim (F_1, \dots, F_n)^{k'-1}} [\text{ALG}(v)],$$

where the second inequality is by noticing that if we select values above $m \cdot \alpha$, it can only improve the performance, and the algorithm becomes a block threshold algorithm. Therefore,

$$\begin{aligned} &(1 - \varepsilon) \cdot \mathbb{E}_{w \sim (D_1, \dots, D_{m-1})} [\max_{i \in [m-1]} w_i] - \mathbb{E}_{w \sim (D_1, \dots, D_{m-1})^{k'-1}} [\text{ALG}^*(w)] \\ &\geq (1 - \varepsilon) \cdot \mathbb{E}_{w \sim (D_1, \dots, D_{m-1})} [\max_{i \in [m-1]} w_i] - \mathbb{E}_{v \sim (F_1, \dots, F_n)^{k'-1}} [\text{ALG}(v)] \\ &= (1 - \varepsilon) \cdot \mathbb{E}_{w \sim (D_1, \dots, D_{m-1})} [\max_{i \in [m-1]} w_i] - (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i] + 4\alpha \\ &\geq (1 - \varepsilon)(-2\alpha) + 4\alpha > 0, \end{aligned}$$

where the equality is by definition of α , the first inequality is by Equation (3), and the last inequality is since $\varepsilon < 1$, and since by Equation (1), $\alpha > 0$. Thus, $k > k' - 1$, which concludes the proof. \square

Beyond the Block Model. Our default model is the block model, according to which we repeat an instance k times, keeping the arrival order within each block the same. Our results continue to hold, even if the arrival order within each block is arbitrary. We discuss this and additional models along with implications for our results in Appendix B.

As is standard in the literature, to simplify the presentation, we generally assume that the distributions do not admit point masses (i.e., F_i is continuous for every $i \in [n]$). In Appendix A, we discuss how to adjust our algorithms for cases with point masses.

The no point masses assumption simplifies the exposition because in the absence of point masses every quantile is associated with a threshold. With point masses, this is not necessarily the case, but the problem can be resolved using randomization.

3 BLOCK THRESHOLD ALGORITHMS

In this section, we study the competition complexity of the class of block threshold algorithms. The following is the main result of this section.

THEOREM 1. *For every $\varepsilon > 0$ the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms is $\Theta(\log \log(1/\varepsilon))$.*

In what follows, given $\tau = (\tau_1, \dots, \tau_k)$ we denote by ALG_τ the block threshold algorithm such that for every copy $j \in \{1, \dots, k\}$ it selects the first element that exceeds τ_j , if such element exists. We denote $\tau_0 = \infty$ and $\tau_{k+1} = 0$.

The remainder of this section is organized as follows. As a warm-up, we start by resolving the case of $k = 2$, which motivates and presents the main ideas of our approach. In Section 3.1 we establish a lemma that allows us to evaluate the performance of any block threshold algorithm ALG_τ by establishing an approximate stochastic dominance inequality that extends the case of $k = 2$ to every k . This inequality is parameterized by the thresholds $\tau = (\tau_1, \dots, \tau_k)$. In order to get the best possible approximation factor, we need to solve the induced max-min problem. In Section 3.2 we solve the corresponding max-min problem by presenting an explicit feasible solution. Finally, in Section 3.3, we present the full proof of Theorem 1, based on the ingredients established in previous sections.

Warm-up: The case of $k = 2$. To design a block threshold algorithm for $k = 2$, we compute the thresholds τ_1, τ_2 by finding the appropriate quantiles of the distribution of $\max_{i \in [n]} v_i$, with $v_i \sim F_i$ for every $i \in \{1, \dots, n\}$. More specifically, for $j \in \{1, 2\}$, let $p_j = \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq \tau_j]$, where τ_1 and τ_2 will be determined by specifying the values of p_1 and p_2 . By our assumption of F_1, \dots, F_n having no point masses, any $p_1, p_2 \in [0, 1]$ corresponds to a pair of thresholds τ_1 and τ_2 . To establish the competition complexity result involving the expectations of ALG_τ and $\max_{i \in [n]} v_i$, our goal is to state an approximate stochastic dominance result between these two random variables. By setting $\tau_1 \geq \tau_2$ and such that $p_1 > 0$, we can show that

$$\begin{aligned} \Pr_{v \sim (F_1, \dots, F_n)^2} [\text{ALG}_\tau(v) \geq x] &\geq \phi_1(p_1, p_2) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] && \text{when } x \geq \tau_1, \\ \Pr_{v \sim (F_1, \dots, F_n)^2} [\text{ALG}_\tau(v) \geq x] &\geq \phi_2(p_1, p_2) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] && \text{when } \tau_1 > x \geq \tau_2, \\ \Pr_{v \sim (F_1, \dots, F_n)^2} [\text{ALG}_\tau(v) \geq x] &\geq \phi_3(p_1, p_2) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] && \text{when } \tau_2 > x, \end{aligned} \quad (4)$$

where $\phi_1(p_1, p_2) = 1 - p_1 + (1 - p_1)(1 - p_2)$, $\phi_2(p_1, p_2) = p_1/p_2 + (1 - p_1)(1 - p_2)$, and $\phi_3(p_1, p_2) = p_1 + p_2(1 - p_1)$. In particular, we get that

$$\Pr_{v \sim (F_1, \dots, F_n)^2} [\text{ALG}_\tau(v) \geq x] \geq \min \left\{ \phi_1(p_1, p_2), \phi_2(p_1, p_2), \phi_3(p_1, p_2) \right\} \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]$$

for every $x \geq 0$. The above inequality is stated in its general form for every k in Lemma 1. Then, in order to get the best possible approximation factor in this way, we solve the following max-min optimization problem:

$$\max \left\{ \min \left\{ \phi_1(p_1, p_2), \phi_2(p_1, p_2), \phi_3(p_1, p_2) \right\} : 0 < p_1 \leq p_2, p_1, p_2 \in [0, 1] \right\}.$$

It can be shown that the optimal solution for this problem is attained at $p_1 = 2/5$ and $p_2 = 2/3$, which yields a factor of $4/5$. For the case of general k , we do not solve exactly this max-min problem as it is a high-dimensional polynomial optimization problem, but we construct explicitly a feasible solution that yields the optimal competition complexity guarantee up to a constant factor. This is formalized in Lemma 2.

3.1 Reduction to Max-Min Problem via Approximate Stochastic Dominance

We start by showing a lemma that allows us to evaluate the performance of any block threshold algorithm ALG_τ by establishing an approximate stochastic dominance inequality that extends (4) to every k . When $k = 2$, the inequalities in (4) are obtained by splitting the domain of x according to the three different sub-intervals defined by the two thresholds. We exploit this idea to get an approximate stochastic dominance inequality by studying each of the $k + 1$ ranges defined by the k block thresholds.

More specifically, given $F_1, \dots, F_n \in \Delta$, and $\tau = (\tau_1, \dots, \tau_k)$ such that $p_0 = 0$ and $p_\ell = \Pr_{v \sim (F_1, \dots, F_n)}[\max_{j \in [n]} v_j \geq \tau_\ell] > 0$ for every $\ell \in \{1, 2, \dots, k + 1\}$, and let

$$\begin{aligned}\Phi_1(F_1, \dots, F_n, \tau) &= \sum_{\ell=1}^k \prod_{j=1}^{\ell} (1 - p_j), \\ \Phi_i(F_1, \dots, F_n, \tau) &= \frac{1}{p_i} \sum_{\ell=1}^{i-1} p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) + \sum_{\ell=i}^k \prod_{j=0}^{\ell} (1 - p_j) \quad \text{for every } i \in \{2, \dots, k\}, \text{ and} \\ \Phi_{k+1}(F_1, \dots, F_n, \tau) &= \sum_{\ell=1}^k p_\ell \prod_{j=0}^{\ell-1} (1 - p_j).\end{aligned}$$

Let $\Phi(F_1, \dots, F_n, \tau) = \min_{i \in \{1, \dots, k+1\}} \Phi_i(F_1, \dots, F_n, \tau)$. We say that $\tau = (\tau_1, \dots, \tau_k)$ is *decreasing* if $\tau_j > \tau_{j+1}$ for every $j \in \{1, \dots, k\}$.

Lemma 1. *For every $F_1, \dots, F_n \in \Delta$, every $x \geq 0$, and every decreasing $\tau = (\tau_1, \dots, \tau_k)$ such that $\Pr_{v \sim (F_1, \dots, F_n)}[\max_{j \in [n]} v_j \geq \tau_1] > 0$, we have*

$$\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] \geq \Phi(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x].$$

PROOF. Given $x \geq 0$, let $\ell(\tau, x) = \min\{\ell \in \{1, \dots, k+1\} : x \geq \tau_\ell\}$, that is, $\ell(\tau, x)$ is the first block ℓ for which x is at least the threshold τ_ℓ . Let \mathcal{B} be defined as follows: If there exists $\ell \in \{1, \dots, k\}$ such that $\max_{i \in [n]} v_i^{(\ell)} > \tau_\ell$, then

$$\mathcal{B} = \min \left\{ \ell \in \{1, \dots, k\} : \max_{i \in [n]} v_i^{(\ell)} > \tau_\ell \right\},$$

and $\mathcal{B} = k + 1$ otherwise. That is, \mathcal{B} is equal to the first block for which there exists a value in the block that surpasses the block threshold, and it is equal $k + 1$ in case such value does not exist. Recall that we denote $p_0 = 0$ and $p_\ell = \Pr_{v \sim (F_1, \dots, F_n)}[\max_{j \in [n]} v_j \geq \tau_\ell]$ for every $\ell \in \{1, 2, \dots, k + 1\}$; in particular $p_{k+1} = 1$. The following holds:

$$\begin{aligned}& \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] \\ &= \sum_{\ell=1}^k \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x \text{ and } \mathcal{B} = \ell] \\ &= \sum_{\ell=1}^k \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell \mid \mathcal{B} > \ell - 1] \Pr_{v \sim (F_1, \dots, F_n)^k} [\mathcal{B} > \ell - 1] \\ &\geq \sum_{\ell=1}^k \Pr_{v \sim (F_1, \dots, F_n)^k} [v_j^{(\ell)} \notin [\tau_\ell, x] \text{ for all } j \in [n] \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell \mid \mathcal{B} > \ell - 1] \prod_{j=0}^{\ell-1} (1 - p_j)\end{aligned}$$

$$= \sum_{\ell=1}^k \Pr_{v^{(\ell)} \sim (F_1, \dots, F_n)} [v_j^{(\ell)} \notin [\tau_\ell, x] \text{ for all } j \in [n] \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell] \prod_{j=0}^{\ell-1} (1 - p_j), \quad (5)$$

where in the second equality, we use that $\mathcal{B} = \ell$ is equivalent to the algorithm not stopping before reaching the block ℓ , and in block ℓ having $\max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell$. For the inequality, we observe that conditioned on $\mathcal{B} > \ell - 1$, the intersection of the event $\text{ALG}_\tau(v) \geq x$ with $\max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell$ contains $v_j^{(\ell)} \notin [\tau_\ell, x]$ for all $j \in [n]$. The final equality follows from independence across copies.

Suppose that $\ell(\tau, x) \geq 2$ and let $\ell \in \{1, \dots, \ell(\tau, x) - 1\}$. In particular, we have $\tau_\ell > x$, and therefore,

$$\begin{aligned} & \Pr_{v^{(\ell)} \sim (F_1, \dots, F_n)} [v_j^{(\ell)} \notin [\tau_\ell, x] \text{ for all } j \in [n] \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell] \\ &= \Pr_{v^{(\ell)} \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell] \\ &= \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \cdot \frac{\Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq \tau_\ell]}{\Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]} \\ &\geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \cdot \frac{\Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq \tau_\ell]}{\Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq \tau_\ell(\tau, x)]} \\ &= \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \cdot \frac{p_\ell}{p_{\ell(\tau, x)}}, \end{aligned} \quad (6)$$

where the inequality holds since $x \geq \tau_{\ell(\tau, x)}$.

Now suppose that $\ell(\tau, x) \leq k$ and let $\ell \in \{\ell(\tau, x), \dots, k\}$. In particular, we have $x \geq \tau_\ell$, and therefore,

$$\begin{aligned} & \Pr_{v^{(\ell)} \sim (F_1, \dots, F_n)} [v_j^{(\ell)} \notin [\tau_\ell, x] \text{ for all } j \in [n] \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq \tau_\ell] \\ &\geq \Pr_{v^{(\ell)} \sim (F_1, \dots, F_n)} [v_j^{(\ell)} \notin [\tau_\ell, x] \text{ for all } j \in [n] \text{ and } \max_{i \in [n]} v_i^{(\ell)} \geq x] \\ &\geq \sum_{i=1}^n \Pr_{v \sim (F_1, \dots, F_n)} [v_i \geq x \text{ and } v_j < \tau_\ell \text{ for all } j \in [n] \setminus \{i\}] \\ &= \sum_{i=1}^n \Pr_{v \sim (F_1, \dots, F_n)} [v_i \geq x] \Pr_{v \sim (F_1, \dots, F_n)} [v_j < \tau_\ell \text{ for all } j \in [n] \setminus \{i\}] \\ &\geq \sum_{i=1}^n \Pr_{v \sim (F_1, \dots, F_n)} [v_i \geq x] \Pr_{v \sim (F_1, \dots, F_n)} [\max_{j \in [n]} v_j < \tau_\ell] \\ &= (1 - p_\ell) \sum_{i=1}^n \Pr_{v \sim (F_1, \dots, F_n)} [v_i \geq x] \\ &\geq (1 - p_\ell) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x], \end{aligned} \quad (7)$$

where the first inequality holds since $x \geq \tau_\ell$; the second inequality holds since the summation in the third line is made of disjoint events whose union has a probability that lower bounds the probability of the second line; the first equality holds by independence across the values; the third inequality holds since $\max_{j \in [n]} v_j < \tau_\ell$ is equivalent to $v_j < \tau_\ell$ for all j , and the last inequality holds by the union bound.

Then, when $\ell(\tau, x) = 1$, from (5) and (7) we get

$$\begin{aligned} \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] &\geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \sum_{\ell=1}^k (1 - p_\ell) \prod_{j=0}^{\ell-1} (1 - p_j) \\ &= \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \sum_{\ell=1}^k \prod_{j=0}^{\ell-1} (1 - p_j) \\ &= \Phi_1(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]. \end{aligned}$$

When $\ell(\tau, x) \in \{2, \dots, k\}$, from (5), (6) and (7) we get

$$\begin{aligned} &\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] \\ &\geq \left(\sum_{\ell=1}^{\ell(\tau, x)-1} \frac{p_\ell}{p_{\ell(\tau, x)}} \prod_{j=0}^{\ell-1} (1 - p_j) + \sum_{\ell=\ell(\tau, x)}^k (1 - p_\ell) \prod_{j=0}^{\ell-1} (1 - p_j) \right) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \\ &= \left(\frac{1}{p_{\ell(\tau, x)}} \sum_{\ell=1}^{\ell(\tau, x)-1} p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) + \sum_{\ell=\ell(\tau, x)}^k \prod_{j=0}^{\ell-1} (1 - p_j) \right) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \\ &= \Phi_{\ell(\tau, x)}(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]. \end{aligned}$$

Finally, when $\ell(\tau, x) = k + 1$, from (5) and (6) we get

$$\begin{aligned} \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] &\geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \sum_{\ell=1}^k \frac{p_\ell}{p_{k+1}} \prod_{j=0}^{\ell-1} (1 - p_j) \\ &= \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \sum_{\ell=1}^k p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) \\ &= \Phi_{k+1}(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x], \end{aligned}$$

where the first equality holds since $p_{k+1} = 1$.

Overall, we conclude that for every $x \geq 0$ we have

$$\begin{aligned} \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_\tau(v) \geq x] &\geq \Phi_{\ell(\tau, x)}(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \\ &\geq \min_{i \in \{1, \dots, k+1\}} \Phi_i(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] \\ &= \Phi(F_1, \dots, F_n, \tau) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x], \end{aligned}$$

which finishes the proof of the lemma. \square

Since Lemma 1 provides an approximate stochastic dominance inequality parameterized in the thresholds $\tau = (\tau_1, \dots, \tau_k)$, in order to get the best possible approximation factor by using this approach we have to solve the following max-min problem:

$$\begin{aligned} &\max \left\{ \min_{i \in \{1, \dots, k+1\}} \Phi_i(F_1, \dots, F_n, \tau) : \tau_1 > \tau_2 > \dots > \tau_k \geq 0 \right\} \\ &= \max \left\{ \Phi(F_1, \dots, F_n, \tau) : \tau_1 > \tau_2 > \dots > \tau_k \geq 0 \right\}. \end{aligned} \tag{8}$$

In the following subsection we study this max-min problem.

3.2 Lower Bound on the Max-Min Problem via an Explicit Feasible Solution

When $k = 2$, we can compute exactly the optimal solution of the max-min problem (8). The problem becomes much harder for general k , but our following lemma establishes a double-exponentially fast increasing lower bound on the optimal value of (8). This lower bound holds by providing a specific set of thresholds, obtained after finding a well-chosen set of quantiles p_1, \dots, p_k .

Lemma 2. *For every $F_1, F_2, \dots, F_n \in \Delta$ and every $k \geq 2$, let $\bar{\tau} = (\bar{\tau}_1, \dots, \bar{\tau}_k)$ such that*

$$\Pr_{v \sim (F_1, \dots, F_n)} [\max_{j \in [n]} v_j \geq \bar{\tau}_\ell] = 1 - (5/4)^{-(5/4)^\ell}$$

for every $\ell \in \{1, \dots, k\}$. Then, we have $\Phi(F_1, \dots, F_n, \bar{\tau}) \geq 1 - (5/4)^{-(5/4)^k}$.

PROOF. To prove the lemma, we show that the following holds:

$$\begin{aligned} \Phi_i(F_1, \dots, F_n, \bar{\tau}) &\geq 1 \text{ for every } i \in \{1, \dots, k\}, \text{ and} \\ \Phi_{k+1}(F_1, \dots, F_n, \bar{\tau}) &\geq 1 - (5/4)^{-(5/4)^k}. \end{aligned}$$

Then, the lemma follows since, overall, we get

$$\Phi(F_1, \dots, F_n, \bar{\tau}) \geq 1 - (5/4)^{-(5/4)^k}.$$

Let $p_0 = 0$, and $p_\ell = \Pr_{v \sim (F_1, \dots, F_n)} [\max_{j \in [n]} v_j \geq \bar{\tau}_\ell]$ for every $\ell \in \{1, \dots, k+1\}$. When $i = 1$, we have

$$\begin{aligned} \Phi_1(F_1, \dots, F_n, \bar{\tau}) &= \sum_{\ell=1}^k \prod_{j=1}^{\ell} (1 - p_j) = \sum_{\ell=1}^k \prod_{j=1}^{\ell} (5/4)^{-(5/4)^j} \\ &= \sum_{\ell=1}^k (5/4)^{5(1-(5/4)^\ell)} \geq (5/4)^{5(1-(5/4))} + (5/4)^{5(1-(5/4)^2)} > 1, \end{aligned}$$

where the first inequality holds since $k \geq 2$. For $i = 2$,

$$\begin{aligned} \Phi_2(F_1, \dots, F_n, \bar{\tau}) &= \frac{p_1}{p_2} + \sum_{\ell=2}^k \prod_{j=0}^{\ell} (1 - p_j) \geq \frac{p_1}{p_2} + (1 - p_1)(1 - p_2) \\ &= \frac{1 - (5/4)^{-5/4}}{1 - (5/4)^{-(5/4)^2}} + (5/4)^{-5/4} \cdot (5/4)^{-(5/4)^2} > 1, \end{aligned}$$

where the first inequality is since $k \geq 2$.

For every $i \in \{3, \dots, k+1\}$, we have

$$\begin{aligned} \sum_{\ell=1}^{i-1} p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) &= \sum_{\ell=1}^{i-1} \left(\prod_{j=0}^{\ell-1} (1 - p_j) - (1 - p_\ell) \prod_{j=0}^{\ell-1} (1 - p_j) \right) \\ &= \sum_{\ell=1}^{i-1} \left(\prod_{j=0}^{\ell-1} (1 - p_j) - \prod_{j=0}^{\ell} (1 - p_j) \right) \\ &= 1 - \prod_{j=0}^{i-1} (1 - p_j) = 1 - \prod_{j=1}^{i-1} (5/4)^{-(5/4)^j} \\ &= 1 - (5/4)^{-\sum_{j=1}^{i-1} (5/4)^j} = 1 - (5/4)^{5(1-(5/4)^{i-1})}. \end{aligned}$$

Thus, for every $i \in \{3, \dots, k\}$ we have

$$\Phi_i(F_1, \dots, F_n, \bar{\tau}) \geq \frac{1}{p_i} \sum_{\ell=1}^{i-1} p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) = \frac{1 - (5/4)^{5(1-(5/4)^{i-1})}}{1 - (5/4)^{-(5/4)^i}} \geq 1,$$

where the last inequality holds since $i \geq 3$, and when $i = k + 1$, we have

$$\Phi_{k+1}(F_1, \dots, F_n, \bar{\tau}) = \sum_{\ell=1}^k p_\ell \prod_{j=0}^{\ell-1} (1 - p_j) = 1 - (5/4)^{5(1-(5/4)^k)} \geq 1 - (5/4)^{-(5/4)^k},$$

where the last inequality holds for every $k \geq 2$. This concludes the proof. \square

3.3 Putting it All Together

With Lemma 1 and Lemma 2 at hand, we are now ready to prove our main theorem.

PROOF OF THEOREM 1. Given $\varepsilon > 0$, by Lemma 1 and Lemma 2, we have that for every $F_1, \dots, F_n \in \Delta$ and every $k \geq \max(2, \log_{5/4} \log_{5/4}(1/\varepsilon))$, by taking $\bar{\tau} = (\bar{\tau}_1, \dots, \bar{\tau}_k)$ as defined in Lemma 2, the following holds:

$$\begin{aligned} \mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\bar{\tau}}(v)] &= \int_0^\infty \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\bar{\tau}}(v) \geq x] dx \\ &\geq \Phi(F_1, \dots, F_n, \bar{\tau}) \int_0^\infty \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] dx \\ &\geq \left(1 - (5/4)^{-(5/4)^k}\right) \int_0^\infty \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] dx \\ &\geq (1 - \varepsilon) \int_0^\infty \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x] dx \\ &= (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i], \end{aligned}$$

which implies that the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms is $O(\log \log(1/\varepsilon))$.

Then, the theorem follows by the fact that [Brüstle et al., 2022, Theorem 2.2] show that the $(1 - \varepsilon)$ -competition complexity of the class of general threshold algorithms when $F_1 = F_2 = \dots = F_n$ is $\Omega(\log \log(1/\varepsilon))$. \square

We note that our proof of Theorem 1 actually shows a $\Theta(\log \log(1/\varepsilon))$ competition complexity guarantee in an approximate stochastic dominance sense, which is stronger than the regular competition complexity result, based on comparing expectations. In particular, our argument strengthens the result of Brüstle et al. [Brüstle et al., 2022] even in the i.i.d. case.

We know from the lower bound on the competition complexity of [Brüstle et al., 2022] that one cannot strengthen Lemma 2 to show that for every c one can devise a series of quantiles for which $\Phi \geq 1 - c^{-c^k}$. However, the proof in [Brüstle et al., 2022] is not explicit, and we now present a simpler proof that shows that one cannot obtain better than $1 - \Omega(c^{-c^k})$ approximation to the value of the prophet using k blocks in the stochastic dominance sense for $c = 3$.

Proposition 2. *For every $k \geq 1$, there exists a positive integer value n_k and $F_1, \dots, F_{n_k} \in \Delta$, such that for every $\tau = (\tau_1, \dots, \tau_k)$, there exists $x \in \mathbb{R}_{\geq 0}$ such that*

$$\Pr_{v \sim (F_1, \dots, F_{n_k})^k} [\text{ALG}_\tau(v) \geq x] < (1 - \varepsilon_k) \Pr_{v \sim (F_1, \dots, F_{n_k})} [\max_{i \in [n_k]} v_i \geq x],$$

where $\varepsilon_k = 3^{-3^k}$ for every k .

PROOF. For every $k \geq 1$, let $n_k = 3^{k+2}$, and for every $i \in [n_k]$, let F_i be the distribution of the random variable $v_i = i \cdot \text{Bernoulli}(1/i) - \text{Uniform}[0, 1]$. For ease of notation, let $F = (F_1, \dots, F_{n_k})$. We first observe that since $\max_{i \in [n_k]} i \cdot \text{Bernoulli}(1/i)$ is distributed uniformly on the set $\{1, \dots, n_k\}$, the random variable $\max_{i \in [n_k]} v_i$ is uniformly distributed over the interval $[0, n_k]$, and for every $x \in [0, n]$, it holds that $\Pr[\max_{i \in [n_k]} v_i \geq x] = (n - x)/n$. For this instance, the optimal block threshold algorithm uses monotone decreasing thresholds, since otherwise sorting them leads to thresholds that stochastically dominates the original ones. Given $\tau = (\tau_1, \dots, \tau_k)$, let $A = \{i \in [k] : \tau_i \leq n_k \cdot 3^{-3^i}\}$. If $A = \emptyset$, then for $x = 0$, it holds that

$$\begin{aligned} \Pr_{v \sim F^k} [\text{ALG}_\tau(v) \geq x] &= 1 - \prod_{\ell=1}^k \Pr_{v^{(\ell)} \sim F} [\max_{i \in [n_k]} v_i^{(\ell)} < \tau_\ell] \\ &\leq 1 - \prod_{\ell=1}^k \frac{\tau_\ell}{n_k} \leq 1 - \prod_{\ell=1}^k 3^{-3^\ell} < (1 - \varepsilon_k) \Pr_{v \sim F} [\max_{i \in [n_k]} v_i \geq x] \end{aligned}$$

where the first inequality holds since $\max_{i \in [n_k]} v_i^{(\ell)}$ is uniformly distributed on the interval $[0, n]$, the second inequality is since $A = \emptyset$, and the last inequality holds by the definition of ε_k , together with $\Pr_{v \sim F} [\max_{i \in [n_k]} v_i \geq x] = 1$.

Else, if $A \neq \emptyset$, let $i^* = \min A$, and let $x = n_k \cdot 3^{-2 \cdot 3^{i^*-1}}$. Then, we denote by B_ℓ (and by \bar{B}_ℓ its complement) the event in which there exists a value in block ℓ exceeding the threshold τ_ℓ , and by B_ℓ^x the event that in block ℓ , for all $i \in [n_k]$ it holds that $v_i^{(\ell)} \notin [\tau_\ell, x]$. For $\ell \geq i^*$, it holds that

$$\Pr[B_\ell^x \text{ and } B_\ell] = \frac{\tau_\ell - \lfloor \tau_\ell \rfloor}{\tau_\ell} \cdot \left(\prod_{j=\lceil \tau_\ell \rceil}^{x-1} \frac{i-1}{i} \right) \cdot \left(1 - \prod_{j=x}^{n_k} \frac{i-1}{i} \right) = \frac{n_k - x}{n_k} \cdot \frac{\tau_\ell}{x}. \quad (9)$$

Thus, we have

$$\begin{aligned} \Pr_{v \sim F^k} [\text{ALG}_\tau(v) \geq x] &= \sum_{\ell=1}^k \left(\prod_{j=1}^{\ell-1} \Pr[\bar{B}_j] \right) \cdot \Pr[B_\ell] \cdot \Pr[\text{ALG}_\tau(v) \geq x \mid B_\ell, \bar{B}_1, \dots, \bar{B}_{\ell-1}] \\ &\leq \sum_{\ell=1}^{i^*-1} \left(\prod_{j=1}^{\ell-1} \Pr[\bar{B}_j] \right) \cdot \Pr[B_\ell] + \sum_{\ell=i^*}^k \left(\prod_{j=1}^{\ell-1} \Pr[\bar{B}_j] \right) \cdot \Pr[B_\ell^x \wedge B_\ell] \\ &\leq \sum_{\ell=1}^{i^*-1} \left(\prod_{j=1}^{\ell-1} 3^{-3^j} \right) \cdot (1 - 3^{-3^\ell}) + \sum_{\ell=i^*}^k \left((3^{-3^{i^*}})^{\ell-i^*} \prod_{j=1}^{i^*-1} 3^{-3^j} \right) \cdot \frac{n_k - x}{n_k} \cdot \frac{\tau_\ell}{x} \\ &< (1 - \varepsilon_k) \cdot \frac{n_k - x}{n_k} = (1 - \varepsilon_k) \Pr_{v \sim F} [\max_{i \in [n_k]} v_i \geq x], \end{aligned}$$

where the first inequality holds since $\Pr[\text{ALG}_\tau(v) \geq x \mid B_\ell, \bar{B}_1, \dots, \bar{B}_{\ell-1}]$ is 1 for $\ell < i^*$, and is $\Pr[B_\ell^x \mid B_\ell]$ otherwise; the second inequality is by the definition of A , since the expression is monotone decreasing in τ_ℓ for $\ell < i^*$, monotone increasing in τ_ℓ for $\ell \geq i^*$, and by Equation (9); the last inequality holds by the definitions of ε_k for every k , and the last equality is since $\max_{i \in [n_k]} v_i$ is uniformly distributed over the interval $[0, n_k]$. \square

4 SINGLE THRESHOLD ALGORITHMS

In this section, we study the competition complexity of the class of single threshold algorithms. Let ALG_τ denote the single threshold algorithm with threshold $\tau \in \mathbb{R}_{\geq 0}$. We show the following tight bound on the competition complexity of single threshold algorithms.

THEOREM 2. *For every $\varepsilon > 0$ the $(1 - \varepsilon)$ -competition complexity of the class of static single threshold algorithms is $\Theta(\log(1/\varepsilon))$.*

We prove Theorem 2 through Lemma 3 and Proposition 3 below. Lemma 3 establishes the upper bound claimed in the theorem, and Proposition 3 shows a matching lower bound.

We begin with Lemma 3, which shows that for every instance, there exists a single threshold τ^* such that for every $\varepsilon > 0$, the $(1 - \varepsilon)$ -competition complexity of the single threshold algorithm ALG_{τ^*} with respect to the instance, is $O(\log(1/\varepsilon))$. The upper bound follows rather directly from the celebrated “median rule” proof of Samuel-Cahn [Samuel-Cahn, 1984]; we work a bit harder to show that it also holds in a stochastic dominance sense.

Lemma 3. *For every F_1, \dots, F_n , there exists a threshold τ^* such that for every $\varepsilon > 0$, and for every $k \geq \log_2(1/\varepsilon)$*

$$\mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v)] \geq (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i].$$

Furthermore, the stronger approximate stochastic-dominance inequality

$$\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \geq x] \geq (1 - \varepsilon) \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]$$

holds for all $x \in \mathbb{R}_{\geq 0}$.

PROOF. Consider the unique threshold τ^* satisfying the following equation

$$\prod_{i=1}^n \Pr_{v_i \sim F_i} [v_i \leq \tau^*] = \frac{1}{2}. \tag{10}$$

In what follows, for ease of notation we consider the input sequence of nk values as v_1, \dots, v_{nk} , where $v_{n(\ell-1)+i} \sim F_i$ for every $\ell \in [k]$ and every $i \in [n]$. We next show that for every number of copies $k \geq 1$, given that the algorithm accepts a value, its expectation is larger than the value of the prophet for a single block, i.e., for every $k \geq 1$,

$$\mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \mid \text{there exists } i \in [nk] : v_i \geq \tau^*] \geq \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i]. \tag{11}$$

In fact we show the following stronger claim, that for every $x \geq 0$ it holds that

$$\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \geq x \mid \text{there exists } i \in [nk] : v_i \geq \tau^*] \geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x]. \tag{12}$$

Let \hat{i} be the first index of a value v_i that exceeds τ^* (and $\hat{i} = 0$ if no such v_i exists), and let $\hat{r} = \lceil \hat{i}/n \rceil$, i.e., the block from which a value is chosen. The LHS of Equation (11) (respectively, Equation (12)) can be rewritten as $\mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \mid \hat{r} \neq 0]$ (respectively, $\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \geq x \mid \hat{r} \neq 0]$). Thus, it is sufficient to prove that Equation (11) holds for every non-zero realization of \hat{r} , i.e., for every $k \geq 1$ and every $r \in [k]$,

$$\mathbb{E}_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \mid \hat{r} = r] \geq \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i]. \tag{13}$$

Since the values of v that don't belong to the r -th n -tuple of values can be ignored, one can observe that Equation (13) is equivalent to the proof of the original prophet inequality, provided by Samuel-Cahn [Samuel-Cahn, 1984]. We next prove the stronger claim of Equation (12). To this end, we show that for every $x \in \mathbb{R}_{\geq 0}$, and for $r \neq 0$ it holds that

$$\Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \geq x \mid \hat{r} = r] \geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x].$$

Note, that for $x \leq \tau^*$, the LHS is 1, thus the inequality holds, and it is sufficient to prove it for $x > \tau^*$. Thus,

$$\begin{aligned}
 & \Pr_{v \sim (F_1, \dots, F_n)^k} [\text{ALG}_{\tau^*}(v) \geq x \mid \hat{r} = r] \\
 &= \sum_{i=(r-1)n+1}^{rn} \Pr[\text{for all } j \in [i-1] : v_j < \tau^*, \text{ and } v_i \geq x \mid \hat{r} = r] \\
 &= \sum_{i=(r-1)n+1}^{rn} \Pr[\text{for all } j \in [i-1] : v_j < \tau^*, v_i \geq x \text{ and } \hat{r} = r] \cdot \frac{1}{\Pr[\hat{r} = r]} \\
 &= 2^r \sum_{i=(r-1)n+1}^{rn} \Pr[\text{for all } j \in [i-1] : v_j < \tau^*, \text{ and } v_i \geq x] \\
 &\geq 2^r \sum_{i=(r-1)n+1}^{rn} \frac{1}{2^r} \cdot \Pr[v_i \geq x] \geq \Pr_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i \geq x],
 \end{aligned}$$

where the first equality is since $x > \tau^*$, the third equality is since by definition of τ^* , $\hat{r} = r$ with probability $1/2^r$ and since $\hat{r} = r$ follows from the other events in the probability, the first inequality is since the $\Pr[\text{for all } j \in \{1, \dots, (r-1)n\} : v_j < \tau^*] = 1/2^{r-1}$, and since $\Pr[\text{for all } j \in \{(r-1)n+1, \dots, i-1\} : v_j < \tau^*] \geq 1/2$, and all the events are independent. The last inequality follows since the sum of probabilities that v_i exceeds x , is at least the probability that the maximum exceeds x .

The proof of the lemma then follows by combining Equation (11) (respectively for the “furthermore” part, Equation (12)) with the observation that, by definition of τ^* , it holds that

$$\begin{aligned}
 \Pr[\text{there exists } i \in [nk] : v_i \geq \tau^*] &= 1 - \Pr[\text{for all } i \in [nk] : v_i \leq \tau^*] \\
 &= 1 - \Pr[\text{for all } i \in [n] : v_i \leq \tau^*]^k = 1 - \frac{1}{2^k} \geq 1 - \varepsilon,
 \end{aligned}$$

where the inequality holds for $k \geq \log_2(1/\varepsilon)$. □

Our next result, Proposition 3, shows a matching lower bound of $\Omega(\log(1/\varepsilon))$ on the competition complexity of single threshold algorithms. This lower bound holds even with respect to the case of i.i.d. distributions.

Proposition 3. *For every $n \geq 2$, and for every $\varepsilon \in (0, 1)$ there exists an instance with n i.i.d. values that are distributed according to some distribution F , such that for every τ , and every $k < \frac{\log_2(1/\varepsilon)}{804}$, it holds that*

$$\mathbb{E}_{v \sim F^{k \cdot n}} [\text{ALG}_\tau(v)] < (1 - \varepsilon) \cdot \mathbb{E}_{v \sim F^n} [\max_{i \in [n]} v_i]. \tag{14}$$

PROOF. It is sufficient to consider $\varepsilon < 1/20$, since otherwise $k = 0$, and the claim holds trivially. Consider the distribution F in which $v_i = 1 + 200\varepsilon \cdot \text{Bernoulli}(1/(20n)) + \text{Uniform}[0, \varepsilon]$ for every $i \in [n]$ and every $\ell \in [k]$. The RHS of Equation (14) satisfies that

$$(1 - \varepsilon) \cdot \mathbb{E}_{v \sim F^n} [\max_{i \in [n]} v_i] \geq (1 - \varepsilon) \cdot \left(1 + 200\varepsilon(1 - e^{-1/20})\right) \geq 1 + 8\varepsilon,$$

where the first inequality is since the maximum is at least $1 + 200\varepsilon$ with probability greater than $1 - e^{-1/20}$, and otherwise it is at least 1. The second inequality holds for every $\varepsilon \in (0, 1/20)$. Now consider a static threshold algorithm with a threshold τ .

Case 1: If $\tau > 1 + \frac{(20n-29)\varepsilon}{20n}$, then the LHS of Equation (14) satisfies that

$$\begin{aligned} \mathbb{E}_{v \sim F^{k \cdot n}} [\text{ALG}_\tau(v)] &\leq \Pr[\text{there exists } i \in [nk] : v_i \geq \tau] \cdot (1 + 201\varepsilon) \\ &\leq (1 - (1 - 3/2n)^{nk})(1 + 201\varepsilon) \leq (1 - 16^{-k})(1 + 201\varepsilon) \leq 1, \end{aligned}$$

where the first inequality is since the support of F is bounded by $1 + 201\varepsilon$, the second inequality is since the probability that $v_i > \tau$ is at most $3/(2n)$, the third inequality is since $n \geq 2$, and the last inequality is since $k < \frac{\log_2(1/\varepsilon)}{804}$. Thus, Equation (14) holds.

Case 2: If $\tau \leq 1 + \frac{(20n-29)\varepsilon}{20n}$, then let $q = \frac{1+\varepsilon-\tau}{\varepsilon} \geq \frac{29}{20n}$. The LHS of Equation (14) satisfies that

$$\begin{aligned} \mathbb{E}_{v \sim F^{k \cdot n}} [\text{ALG}_\tau(v)] &\leq \mathbb{E}_{v_i \sim F} [v_i \mid v_i \geq \tau] \\ &= \mathbb{E}_{v_i \sim F} [v_i \cdot \mathbf{1}[v_i \geq \tau]] / \Pr[v_i \geq \tau] \\ &\leq \left(\frac{1}{20n} (1 + 201\varepsilon) + \frac{20n-1}{20n} q(1 + \varepsilon) \right) / \left(\frac{1}{20n} + \frac{20n-1}{20n} q \right) \\ &= ((1 + 201\varepsilon) + (20n-1)q(1 + \varepsilon)) / (1 + (20n-1)q) \\ &< 1 + 8\varepsilon, \end{aligned}$$

where the last inequality holds for every $n \geq 2$, and $q \geq \frac{29}{20n}$. Thus, Equation (14) holds. \square

5 COMBINATORIAL EXTENSIONS

We conclude with a short discussion of additional results for combinatorial settings, which we present in the full version. In addition to formulating a general model, we obtain some preliminary results for combinatorial auctions and matching feasibility constraints. In what follows, we briefly elaborate on these results.

A combinatorial auction is defined by a set of m items and n buyers. Each buyer i has a (set-valued) valuation function $v_i : 2^{[m]} \rightarrow \mathbb{R}_{\geq 0}$. The valuation functions of the buyers, v_1, \dots, v_n , are sampled independently from known distributions over valuation functions, F_1, \dots, F_n . We consider the block arrival model, where the online algorithm sees k copies of the instance while the prophet sees just one copy. Generalizing the notion of block-threshold algorithms and single-threshold algorithms, we consider *block-consistent pricing schemes* and *static pricing schemes*. The former set a price $p_j \geq 0$ on each item $j \in [m]$, which may be updated after each block. The latter uses the same prices throughout the k copies.

We show that for submodular and XOS valuations (see full version for definitions), there exists a block-consistent pricing scheme that achieves a $(1 - \varepsilon)$ -competition complexity of $O(\log(1/\varepsilon))$. We show this result through a variation of the balanced pricing approach [Feldman et al., 2015]. In our algorithm, an item's price is set proportional to that item's contribution to the offline optimum allocation, with prices in block $t \in [k]$ adjusted by a factor of $1 - 1/(2^{k+1-t})$. We also show a $O(1/\varepsilon)$ -competition complexity result for static pricing schemes. For this result, we also set item prices proportional to the items' contributions to the expected offline optimum, but apply a single scalar, namely $1 - 1/(k + 1)$ throughout all blocks.

Our result for matching feasibility constraints concerns bipartite matching problems and shows a $O(\log(1/\varepsilon))$ competition complexity. This result is subsumed by the result for XOS combinatorial auctions, but it is based on a different algorithmic approach; namely the OCRS-based algorithm of Ezra et al. [2022].

ACKNOWLEDGMENTS

J. Correa was partially funded by Anillo ICMD (ANID grant ACT210005) and the Center for Mathematical Modeling (ANID grant FB210005). T. Ezra is supported by the Harvard University Center of Mathematical Sciences and Applications. M. Feldman was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement No. 866132), by an Amazon Research Award, by the NSF-BSF (grant number 2020788) and by a grant from TAU Center for AI and Data Science (TAD). V. Verdugo was partially funded by Anillo ICMD (ANID grant ACT210005).

REFERENCES

- M. Abolhassani, S. Ehsani, H. Esfandriari, M. Hajiaghayi, R. Kleinberg, and B. Lucier. 2017. Beating $1-1/e$ for Ordered Prophets. In *STOC 2017*. 61–71.
- S. Alaei. 2011. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *FOCS 2011*. 512–521.
- N. Arnosti and W. Ma. 2022. Tight guarantees for static threshold policies in the prophet secretary problem. *Operations Research* (2022).
- M. Babaioff, K. Goldner, and Y. A. Gonczarowski. 2020. Bulow-Klemperer-Style Results for Welfare Maximization in Two-Sided Markets. In *SODA 2020*. 2452–2471.
- H. Beyhaghi and S. M. Weinberg. 2019. Optimal (and Benchmark-Optimal) Competition Complexity for Additive Buyers over Independent Items. In *STOC 2019*. 686–696.
- J. Brustle, J. Correa, P. Dütting, and V. Verdugo. 2022. The Competition Complexity of Dynamic Pricing. In *EC 2022*. 303–320.
- J. Bulow and P. Klemperer. 1996. Auctions versus Negotiations. *American Economic Review* 86, 1 (1996), 180–94.
- S. Chawla, J. Hartline, D. Malec, and B. Sivan. 2010. Multiparameter Mechanism Design and Sequential Posted Pricing. In *STOC 2010*. 311–320.
- J. Correa and A. Cristi. 2023. A constant-factor prophet inequality for online combinatorial auctions. In *STOC 2023*. 686–697.
- J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. 2021. Posted Price Mechanisms and Optimal Threshold Strategies for Random Arrivals. *Math. Oper. Res.* 46, 4 (2021), 1452–1478.
- J. Csirik and G.J. Woeginger. 2000. Resource Augmentation for Online Bounded Space Bin Packing. In *ICALP 2000*. 296–304.
- P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. 2017. Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs. In *FOCS 2017*. 540–551.
- P. Dütting, T. Kesselheim, and B. Lucier. 2020. An $O(\log \log m)$ Prophet Inequality for Subadditive Combinatorial Auctions. In *FOCS 2020*. 306–317.
- A. Eden, M. Feldman, O. Friedler, I. Talgam-Cohen, and S. M. Weinberg. 2017. The Competition Complexity of Auctions: A Bulow-Klemperer Result for Multi-Dimensional Bidders. In *EC 2017*. 343.
- T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang. 2022. Prophet matching with general arrivals. *Mathematics of Operations Research* 47, 2 (2022), 878–898.
- Tomer Ezra, Michal Feldman, Tim Roughgarden, and Warut Suksompong. 2020. Pricing multi-unit markets. *ACM Transactions on Economics and Computation* 7, 4 (2020), 1–29.
- M. Feldman, N. Gravin, and B. Lucier. 2015. Combinatorial Auctions via Posted Prices. In *SODA 2015*. 123–135.
- M. Feldman, O. Svensson, and R. Zenklusen. 2016. Online Contention Resolution Schemes. In *SODA 2016*, Robert Krauthgamer (Ed.). 1014–1033.
- N. Gravin and H. Wang. 2019. Prophet Inequality for Bipartite Matching: Merits of Being Simple and Non Adaptive. In *EC 2019*. 93–109.
- M. Hajiaghayi, R. Kleinberg, and T. Sandholm. 2007. Automated Mechanism Design and Prophet Inequalities. In *AAAI 2007*. 58–65.
- B. Kalyanasundaram and K. Pruhs. 2000. Speed is as powerful as clairvoyance. *Journal of the ACM* 47, 4 (2000), 617–643.
- R. Kleinberg and S. M. Weinberg. 2012. Matroid Prophet Inequalities. In *STOC 2012*. 123–136.
- U. Krengel and L. Sucheston. 1977. Semiamarts and finite values. *Bull. Amer. Math. Soc.* 83 (1977), 745–747.
- U. Krengel and L. Sucheston. 1978. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics* 4 (1978), 197–266.
- A. Liu, R. Paes Leme, M. Pál, J. Schneider, and B. Sivan. 2021. Variable Decomposition for Prophet Inequalities and Optimal Ordering. In *EC 2021*. 692.
- C. A. Phillips, C. Stein, E. Torgn, and J. Wein. 1997. Optimal Time-Critical Scheduling via Resource Augmentation (Extended Abstract). In *STOC 1997*. 140–149.
- T. Roughgarden and E. Tardos. 2002. How bad is selfish routing? *J. ACM* 2 (2002), 236–259. Issue 49.

A. Rubinstein and S. Singla. 2017. Combinatorial Prophet Inequalities. In *SODA 2017*. 1671–1687.

E. Samuel-Cahn. 1984. Comparison of Threshold Stop Rules and Maximum for Independent Nonnegative Random Variables. *The Annals of Probability* 12, 4 (1984), 1213 – 1216.

D. D. Sleator and R. E. Tarjan. 1984. Amortized Efficiency of List Update Rules. In *STOC 1984*. 488–492.

A DEALING WITH POINT MASSES

In this appendix, we discuss the adaptations of our results to the case where the distributions have point masses. To deal with point masses, we can use standard techniques (e.g., [Arnosti and Ma, 2022]) by allowing our algorithms to select, whenever observing a value that is equal to the threshold τ , with a fixed probability p , and when the value exceeds τ , with probability one. In particular, every threshold strategy can be described by two parameters, τ and p .

To generalize our proofs to the case where there are point masses, we do the following: Given distributions F_1, \dots, F_n , and a quantile $g \in (0, 1)$, let $\tau^* = \sup_{\tau \geq 0} \prod_{i=1}^n \Pr_{v_i \sim F_i} [v_i \leq \tau] \leq g$ and $q = \prod_{i=1}^n \Pr_{v_i \sim F_i} [v_i \leq \tau^*]$. For every $i \in [n]$, let $q_i = \Pr_{v_i \sim F_i} [v_i \leq \tau^*]$, and let $p_i = \Pr_{v_i \sim F_i} [v_i = \tau^*]$.

Then, it holds that $\prod_{i=1}^n q_i = q \geq g \geq \prod_{i=1}^n (q_i - p_i)$, and both are strict inequalities, unless $p_i = 0$ for all $i \in [n]$. If $p_i = 0$ for all $i \in [n]$, then each threshold should be interpreted as to select the first element that exceeds τ^* . Otherwise, let p^* be the unique value satisfying $\prod_{i=1}^n (q_i - p^* \cdot p_i) = g$. Observe that p^* must be in $[0, 1]$. Then, the algorithm should select the first element that exceeds τ^* , and should accept the value τ^* with probability p^* . Note that p^* is independent of the order of the distributions, and can be computed by only using distributional information.

We next show that if there are point masses, and if the algorithm must use a *deterministic* single threshold (i.e., that selects the first value that exceeds it), then the $(1 - \varepsilon)$ -competition complexity is in $\Theta(1/\varepsilon)$. To prove this we first show that the $(1 - \varepsilon)$ -competition complexity is bounded from below by $\Omega(1/\varepsilon)$.

Proposition 4. *For every $n \geq 2$, and for every $\varepsilon \in (0, 1)$ there exists an instance with n i.i.d. values that are distributed according to some distribution F with point masses, such that for every τ , and every $k < \frac{1}{60\varepsilon}$, it holds that*

$$\mathbb{E}_{v \sim F^{k \cdot n}} [\text{ALG}_\tau(v)] < (1 - \varepsilon) \cdot \mathbb{E}_{v \sim F^n} [\max_{i \in [n]} v_i]. \quad (15)$$

PROOF. It is sufficient to consider $\varepsilon < 1/20$, since otherwise $k = 0$, and the claim holds trivially. Consider the distribution F such that $v_i = 1 + 10 \cdot \text{Bernoulli}(\varepsilon/(3n))$ for every $i \in [n]$. The RHS of Equation (15) satisfies that

$$\begin{aligned} (1 - \varepsilon) \cdot \mathbb{E}_{v \sim F^n} [\max_{i \in [n]} v_i] &= (1 - \varepsilon) \left(1 + 10 \left(1 - \left(1 - \frac{\varepsilon}{3n} \right)^n \right) \right) \\ &\geq (1 - \varepsilon) \left(1 + 10 \left(1 - e^{-\varepsilon/3} \right) \right) \geq 1 + 2\varepsilon, \end{aligned}$$

where the last inequality holds for every $\varepsilon \leq 1/20$. Now consider a static threshold algorithm with a threshold τ . We have two cases.

Case 1: If $\tau < 1$, then since the algorithm will select the first value v_1 , and the LHS of Equation (15) satisfies that

$$\mathbb{E}_{v \sim F^{k \cdot n}} [\text{ALG}_\tau(v)] = \mathbb{E}_{v \sim F} [v] = 1 + \frac{10\varepsilon}{3n} < 1 + 2\varepsilon,$$

where the inequality is since $n \geq 2$. Thus, Equation (15) holds.

Case 2: If $\tau \geq 1$, then the algorithm will only accept the value 11, and the the LHS of Equation (15) satisfies that

$$\begin{aligned}\mathbb{E}_{v \sim F^{k \cdot n}}[\text{ALG}_\tau(v)] &= \Pr_{v \sim F^{k \cdot n}}[\text{there exists } i \in [nk] : v_i > 1] \cdot 11 \\ &= 11 \cdot \left(1 - \left(1 - \frac{\varepsilon}{3n}\right)^{nk}\right) \leq 1,\end{aligned}$$

where the last inequality holds for every $n \geq 2$, every $\varepsilon \leq 1/20$, and every $k \leq 1/(60\varepsilon)$. Thus, (15) holds, and this concludes the proof. \square

We next show that there is deterministic single threshold algorithm that guarantees a $(1 - \varepsilon)$ -competition complexity of $O(1/\varepsilon)$.

Proposition 5. *For every $\varepsilon > 0$ the $(1 - \varepsilon)$ -competition complexity of the class of deterministic single threshold algorithms for the case where there are point masses is at most $\lceil 1/\varepsilon \rceil$.*

PROOF. Given an instance (F_1, \dots, F_n) , consider the threshold $\tau = (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)}[\max_{i \in [n]} v_i]$, and let $p = \Pr_{v \sim (F_1, \dots, F_n)}[\max_{i \in [n]} v_i \leq \tau]$. Then, for $k = \lceil 1/\varepsilon \rceil$, it holds that

$$\begin{aligned}\mathbb{E}_{v \sim (F_1, \dots, F_n)^k}[\text{ALG}_\tau(v)] &= (1 - p^k)\tau + \sum_{\ell=1}^k \sum_{i=1}^n \mathbb{E}_{v \sim (F_1, \dots, F_n)^k}[\max(v_i^{(\ell)} - \tau, 0)] \cdot p^{\ell-1} \prod_{j < i} \Pr[v_j^{(\ell)} \leq \tau] \\ &\geq (1 - p^k)\tau + p^k \sum_{\ell=1}^k \sum_{i=1}^n \mathbb{E}_{v_i \sim F_i}[\max(v_i - \tau, 0)] \\ &\geq (1 - p^k)\tau + p^k \sum_{\ell=1}^k \mathbb{E}_{v \sim (F_1, \dots, F_n)}[\max(\max_{i \in [n]} v_i - \tau, 0)] \\ &\geq (1 - p^k)\tau + p^k \cdot \frac{1}{\varepsilon} \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)}[\max_{i \in [n]} v_i - \tau] \\ &= (1 - p^k)\tau + p^k \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)}[\max_{i \in [n]} v_i] \\ &\geq (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)}[\max_{i \in [n]} v_i],\end{aligned}$$

which concludes the proof. \square

B ALTERNATIVE ARRIVAL ORDERS

In this appendix, we study the competition complexity of the single-choice problem beyond the block model. We first observe that our results continue to hold if the variables in each block arrive in arbitrary order. We then consider a version where variables can be moved across blocks, but in a limited way. Finally, we establish a lower bound for the fully adversarial case where the order can be arbitrary.

Displacement Model. In our basic model — the block model — we assumed that variables arrive in the same order within each block. In the *displaced* block model, within each of the k blocks the n values sampled from the distributions F_1, \dots, F_n are presented in arbitrary order.

A first observation is that our proofs of the competition complexity results for block threshold algorithms and single threshold algorithms continue to hold in the displaced block model.

Observation 1. *In the displaced block model, for every $\varepsilon > 0$ the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms is $\Theta(\log \log(1/\varepsilon))$. Furthermore, the $(1 - \varepsilon)$ -competition complexity of the class of single threshold algorithms is $\Theta(\log(1/\varepsilon))$.*

Naturally, this motivates the question of studying the competition complexity in a setting where the values can be displaced across the different blocks. In what follows, for a pair of natural numbers (n, k) we denote by S_n the set of permutations of $[n]$ and by $\sigma(F_1, \dots, F_n)^k$ the instance that permutes $(F_1, \dots, F_n)^k$ according to the permutation $\sigma \in S_{nk}$.

We introduce the γ -displacement model in which the adversary may change the order of γk copies of an instance $F = (F_1, \dots, F_n)$, in a way such that if we split up the γnk variables into γ -blocks of size γn , then within each γ -block every distribution appears at least once. In what follows, we define a few objects to formally introduce this model.

Given a positive integer γ and an integer i , let $B_{\gamma n}(i) = \{(i-1)\gamma n + j : j \in [\gamma n]\}$ be the i -th γ -block. Given an order $\sigma \in S_{\gamma nk}$, and $j \in [n]$, let $I(\sigma, j)$ be the indices of σ where there is a variable of type j . Finally, let $D_\gamma(n, k)$ be the subset of permutations in $S_{\gamma nk}$ satisfying that, after permutation according to σ , each γ -block, contains at least one of each index in $[n]$, that is

$$D_\gamma(n, k) = \left\{ \sigma \in S_{\gamma nk} : \text{for every } j \in [n] \text{ and every } i \in [k], |I(\sigma, j) \cap B_{\gamma n}(i)| \geq 1 \right\}.$$

In the γ -displacement model, we compare $\min_{\sigma \in D_\gamma(n, k)} \mathbb{E}_{v \sim \sigma F \gamma k} [\text{ALG}_\tau(v)]$ to $\mathbb{E}_{v \sim F} [\max_{i \in [n]} v_i]$, where $F = (F_1, \dots, F_n)$. We remark that the 1-displacement model corresponds to the displaced model discussed in Observation 1. In the following proposition, we show that the competition complexity in the γ -displacement model increases by at most a factor γ in comparison to the 1-displacement model.

Proposition 6. *In the γ -displacement model, for every $\varepsilon > 0$ the $(1 - \varepsilon)$ -competition complexity of the class of block threshold algorithms is $O(\gamma \log \log(1/\varepsilon))$.*

PROOF. To prove the result, we provide a reduction of the γ -displacement model to the block model. In the block model with k copies, let $\tau = (\tau_1, \dots, \tau_k)$ be the thresholds corresponding to the quantiles $p_\ell = \Pr_{v \sim (F_1, \dots, F_n)} [\max_{j \in [n]} v_j \geq \tau_\ell]$ for $\ell \in [k]$. Given some permutation $\sigma \in D_\gamma(n, k)$, define the thresholds $\tau' = (\tau'_1, \dots, \tau'_k)$ to be such that $\Pr_{v \sim \sigma F \gamma k} [\max_{j \in B_{\gamma n}(\ell)} v_j \geq \tau'_\ell] = p_\ell$ for $\ell \in [k]$. Then, in the instance $\sigma(F_1, \dots, F_n)^{\gamma k}$ with γk copies, consider the block thresholds $\tau' = (\tau'_1, \dots, \tau'_1, \tau'_2, \dots, \tau'_2, \dots, \tau'_k, \dots, \tau'_k)$ where each threshold τ'_ℓ for every $\ell \in [k]$ is repeated consecutively γ times. Since $\sigma \in D_\gamma(n, k)$, for every $x \geq 0$ and every $\ell \in [k]$,

$$\Pr_{v \sim \sigma F \gamma k} [\max_{j \in B_{\gamma n}(\ell)} v_j \geq x] \geq \Pr_{v \sim F} [\max_{j \in [n]} v_j \geq x],$$

that is, we have a stochastic dominance inequality. We can then adjust the block model proof in Section 3 on γ -blocks instead of single blocks to get that $O(\gamma \log \log(1/\varepsilon))$ copies provides a $(1 - \varepsilon)$ -competition complexity in the γ -displacement model. \square

We next show that an equivalent guarantee for the case of single-threshold algorithms cannot be achieved: namely, the $(1 - \varepsilon)$ -competition complexity of this type of algorithms in the γ -displacement model has to scale at least polynomially in $1/\varepsilon$.

Proposition 7. *For the γ -displacement model, there exists a $\gamma > 0$ such that the $(1 - \varepsilon)$ -competition complexity of any single-threshold algorithm is at least $1/(6\varepsilon^{1/3})$.*

PROOF. Consider $\varepsilon > 0$ small enough, such that $0 < \varepsilon < 0.076$. For ease of presentation, suppose that $1/\varepsilon$ is an integer. Consider the following problem instance with $n = \frac{1}{\varepsilon} + 2$ random variables of three different types: There is one random variable of type 1, with distribution $1 + U[0, \varepsilon^{10}]$. There is one random variable of type 2, whose distribution is $1 + \varepsilon^{1/3} + U[0, \varepsilon^{10}]$ with probability $\varepsilon^{1/3}$ and $U[0, \varepsilon^{10}]$ otherwise. There are $1/\varepsilon$ many random variables of type 3, with distribution $U[0, \varepsilon^{10}]$. So we basically have a deterministic one (type 1), a high value with low probability (type 2), and many

zeros (type 3) – plus some random noise. We denote the distributions of these random variables by F_1, F_2 , and F_3 . Let $F = F_1 \times F_2 \times F_3 \times \dots \times F_3$.

The prophet can take the high value (the $\approx 1 + \varepsilon^{1/3}$ value) from the second random variable if it realizes, or the guaranteed value of ≈ 1 from the first random variable if it doesn't. It can thus achieve an expected value of at least

$$\text{OPT} \geq \varepsilon^{1/3} \cdot (1 + \varepsilon^{1/3}) + (1 - \varepsilon^{1/3}) \cdot 1 = 1 + \varepsilon^{2/3}.$$

Consider the γ -displacement model for $\gamma = 3$ and k copies. We want to show that any single-threshold algorithm, in order to achieve a $(1 - \varepsilon)$ -approximation to the prophet, must have $k \geq 1/(6\varepsilon^{1/3})$. Assume towards contradiction that $k < 1/(6\varepsilon^{1/3})$.

Recall that in the γ -displacement model we seek a guarantee that applies for any possible grouping of the γnk random variables, into k many γ -blocks of size γn each. We construct a hard instance as follows. The first γ -block is:

$$\underbrace{(1, \dots, 1, 2, 2, 2, 3, \dots, 3)}_{2k+1} \quad \underbrace{\hspace{10em}}_{\frac{3}{\varepsilon} + 2 - 2k}$$

All the remaining γ -blocks are:

$$(1, 2, 2, 2, 3, \dots, 3) \quad \underbrace{\hspace{10em}}_{\frac{3}{\varepsilon} + 2}$$

Here 1, 2, and 3 indicate that the respective random variable is of that type. So in the first γ -block we have $2k + 1$ random variables of type 1, followed by 3 random variables of type 2, followed by $3/\varepsilon + 2 - 2k$ random variables of type 3. Similarly, in the remaining γ -blocks we have one random variable of type 1, followed by 3 random variables of type 2, followed by $3/\varepsilon + 2$ random variables of type 3. Note that both types of blocks consist of $\gamma n = 3/\varepsilon + 6$ random variables, and that variables of type 1 and type 2 occur for a total of $3k$ times while variables of type 3 occur for a total of $3k/\varepsilon$ times. So this forms a valid instance.

Let's now analyze the performance of an arbitrary single-threshold algorithm ALG. Let $T \geq 0$ be the algorithm's threshold, and let $q = \Pr_{X \sim F_1}[X \leq T]$.

$$\Pr[\text{ALG} \geq 1] \leq 1 - q^{3k} (1 - \varepsilon^{1/3})^{3k} \leq 1 - q^{3k} (1 - 3k\varepsilon^{1/3}),$$

where the last inequality uses that $(1 - x)^{3k} \geq (1 - 3kx)$ for $x = \varepsilon^{1/3}$.

Next we show an upper bound on the probability that the algorithm stops on a value that is at least $1 + \varepsilon^{1/3}$. For this the algorithm must skip over all the random variables of type 1 in the first γ -block. Therefore,

$$\Pr[\text{ALG} \geq 1 + \varepsilon^{1/3}] \leq q^{2k+1} \leq q^{2k},$$

where we used that $q \leq 1$. We thus obtain,

$$\begin{aligned} \mathbb{E}[\text{ALG}] &\leq \Pr[\text{ALG} \geq 1] + \varepsilon^{1/3} \cdot \Pr[\text{ALG} \geq 1 + \varepsilon^{1/3}] + \varepsilon^{10} \\ &= 1 - q^{3k} (1 - 3k\varepsilon^{1/3}) + \varepsilon^{1/3} q^{2k} + \varepsilon^{10}. \end{aligned}$$

We need that $\mathbb{E}[\text{ALG}] \geq (1 - \varepsilon)\text{OPT}$. This gives us the following inequality

$$1 - q^{3k} (1 - 3k\varepsilon^{1/3}) + \varepsilon^{1/3} q^{2k} + \varepsilon^{10} \geq (1 - \varepsilon)(1 + \varepsilon^{2/3}),$$

which by rearrangement gives us that

$$\underbrace{\varepsilon + \varepsilon^{5/3} + \varepsilon^{10}}_{< \frac{\varepsilon^{2/3}}{2}} + \underbrace{\varepsilon^{1/3} q^{2k}}_{\leq q^{1.5k} \varepsilon^{1/3}} + \underbrace{3k\varepsilon^{1/3} q^{3k}}_{\leq \frac{q^{3k}}{2}} \geq q^{3k} + \varepsilon^{2/3},$$

where the underlined inequalities hold since $\varepsilon < 0.076$, $q \leq 1$, and $k < 1/(6\varepsilon^{1/3})$.

However, by the short multiplication formula¹ we get that

$$q^{3k} + \varepsilon^{2/3} = \frac{\varepsilon^{2/3}}{2} + \frac{q^{3k}}{2} + \left(\sqrt{\frac{\varepsilon^{2/3}}{2}} + \sqrt{\frac{q^{3k}}{2}} \right) \geq \frac{\varepsilon^{2/3}}{2} + \frac{q^{3k}}{2} + q^{1.5k} \varepsilon^{1/3}.$$

This yields a contradiction, which implies that $k \geq 1/(6\varepsilon^{1/3})$. □

Fully Adversarial Model. Beyond bounded displacement, a natural question is to ask what can be said in the case where the adversary is not restricted to choosing an order among a limited family of permutations, but can instead arrange the nk distributions in any order. We refer to this as the *fully adversarial model*. It corresponds to the case where the instance of the online algorithm is given by $\sigma(F_1, \dots, F_n)^k$, for an arbitrary $\sigma \in S_{nk}$. That is, each distribution is used k times to sample values, but the order in which they are presented to the online algorithm is arbitrary. We show that in this case the $(1 - \varepsilon)$ -competition complexity is lower bounded by $\Omega(1/\varepsilon)$.

Proposition 8. *For all $n \geq 2$, $\varepsilon \in (0, 1)$ and every $k < 1/(4\varepsilon)$, there exist distributions F_1, \dots, F_n and a permutation $\sigma \in S_{nk}$ such that for every sequence of thresholds τ we have*

$$\mathbb{E}_{v \sim \sigma(F_1, \dots, F_n)^k} [\text{ALG}_\tau(v)] < (1 - \varepsilon) \cdot \mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i].$$

PROOF. First note that for $\varepsilon \in [1/2, 1)$, the condition $k < 1/(4\varepsilon)$ implies $k = 0$. Hence we may assume $\varepsilon \in (0, 1/2)$. Let $F_1 = \dots = F_{n-1} = 1$ with probability 1, and $F_n = 2$ with probability 2ε and $F_n = 0$ otherwise. Let σ be the permutation that sends the k copies of F_n to be the last k distributions seen by the algorithm. Observe that a single threshold algorithm is optimal for this instance. Then, consider any single threshold $\tau \in (1, 2]$. Then for every $\varepsilon \in (0, 1/2)$ and every $k < 1/(4\varepsilon)$, we have

$$\mathbb{E}_{v \sim \sigma(F_1, \dots, F_n)^k} [\text{ALG}_\tau(v)] = (1 - (1 - 2\varepsilon)^k) \cdot 2 \leq 7/8.$$

Hence, the optimal threshold is at most 1, and yields $\mathbb{E}_{v \sim \sigma(F_1, \dots, F_n)^k} [\text{ALG}_\tau(v)] = 1$. On the other hand,

$$\mathbb{E}_{v \sim (F_1, \dots, F_n)} [\max_{i \in [n]} v_i] = 2\varepsilon \cdot 2 + (1 - 2\varepsilon) \cdot 1 = 1 + 2\varepsilon.$$

This concludes the proof, as for every $\varepsilon \in (0, 1/2)$, we have $(1 - \varepsilon)(1 + 2\varepsilon) > 1$. □

C EXPECTED NUMBER OF BLOCKS

In this appendix, we make a remark on the expected performance of the single threshold algorithm. Recall that its competition complexity in the block model is $\Theta(\log(1/\varepsilon))$ (Theorem 2).

Proposition 9. *There is a single threshold algorithm that, in expectation, terminates after two blocks and achieves an expected value of at least $\mathbb{E}_{v \sim F} [\max_{i \in [n]} v_i]$.*

PROOF. We consider the same algorithm as in Lemma 3, namely a single-threshold algorithm, with threshold τ^* satisfying $\prod_{i=1}^n \Pr_{v_i \sim F_i} [v_i \leq \tau^*] = \frac{1}{2}$. Equation (11) shows that the expected value achieved by this algorithm, conditional on stopping, is at least $\mathbb{E}_{v \sim F} [\max_{i \in [n]} v_i]$. We additionally observe that the expected number of blocks after which it stops is given by $\sum_{i=1}^\infty i (1/2)^i = 2$. □

We remark that the preceding result is essentially tight: With k copies, no algorithm can get a better approximation to $\mathbb{E}_{v \sim F} [\max_{i \in [n]} v_i]$ than $k/2$. Thus, the expected number of blocks to achieve a $(1 - \varepsilon)$ -approximation, is at least $2 - 2\varepsilon$.

¹The short multiplication formula that we are using here is $(a - b)^2 = a^2 - 2ab + b^2$. Since $(a - b)^2 \geq 0$, this shows that $a^2 + b^2 \geq 2ab$. We apply this to $a = \sqrt{\varepsilon^{2/3}/2}$ and $b = \sqrt{q^{3k}/2}$.