Optimisation for Prophet Inequalities

Johannes Brüstle

Department of Mathematics
London School of Economics and Political Science

This dissertation is submitted for the degree of
Doctor of Philosophy

January 2025

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work, with the following exceptions.

- Chapter 2 is based on [13] *"The competition complexity of dynamic pricing"*, co-authored with Jose Correa, Paul Dütting and Victor Verdugo. It has appeared in *Proceedings of the 23rd ACM Conference on Economics and Computation* and was accepted to *Mathematics of Operations Research*.
- Chapter 3 is based on [12] *"The Competition Complexity of Prophet Inequalities"*, co-authored with Jose Correa, Paul Dütting, Tomer Ezra, Michal Feldman and Victor Verdugo. It is to appear in *"Proceedings of the 25th ACM Conference on Economics and Computation"*.
- Chapter 4 is based on *"Splitting Guarantees for Prophet Inequalities via Nonlinear Systems"*, co-authored with Sebastian Perez-Salazar and Victor Verdugo. It has been accepted at The 20th Conference on Web and Internet Economics (WINE 2024).
- Chapters 1 and 5 contain small parts of all works mentioned above: [13], [12] and *"Splitting Guarantees for Prophet Inequalities via Nonlinear Systems"*, co-authored with Sebastian Perez-Salazar and Victor Verdugo.

Johannes Brüstle
January 2025

Acknowledgements

First and foremost, I would like to thank my advisors Paul Dütting and László Végh for their support throughout my PhD. Working together with you, I feel that I learned much on different ways to conduct research in mathematics.

Thank you as well to all my co-authors in this thesis. I greatly enjoyed our collaborations and meetings. After a somewhat difficult period caused by Covid-19, it was great that I had a chance to meet most of you in person. In particular, I had a very nice time during my visit to Santiago, Chile, where I got to meet a number of great people.

My stay at LSE wouldn't have been the same without everyone I met at the department. You made my time in London memorable and I hope to see you again sometime.

I am grateful to LSE who by awarding me the PhD Studentship, supported me throughout my degree.

Thank you to my Canadian friends for sticking around.

Finally, I would like to thank my parents and my brother for supporting me over the years.

Abstract

Prophet inequalities have been extensively studied within optimal stopping theory, in part for their applicability to many online decision-making processes. In this thesis, we explore the benefits of viewing prophet inequalities as optimisation problems. We present three main results.

First, we study the classic single-choice prophet inequality problem through a resource augmentation lens. In this setting, the algorithm faces an additional number of online independent and identically distributed (i.i.d.) bidders compared to the prophet. The optimal algorithm has a simple description as it sets carefully chosen thresholds $T(j)$ for each incoming bidder j and accept a given value from the distribution if and only if it surpasses $T(j)$. Our goal is to analyse the competition complexity, which relates to the number of extra resources required in order to approximate the benchmark by a given factor.

Next, we generalise to arbitrary, independent distributions. Now, the 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 $(1-\varepsilon)$-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 $k = \Theta(\log(\log 1/\varepsilon))$. This shows that block threshold algorithms approach the offline optimum doubly-exponentially fast. For single threshold algorithms, which set the same threshold throughout, we give a tight bound of $k = \Theta(\log(1/\varepsilon))$ establishing an exponential gap between block and single threshold algorithms.

Finally, we move on to the i.i.d. k-selection prophet inequality problem, which is a different extension of the single choice setting in the case of i.i.d. distributions. At each time step, a decision is made to accept or reject the value, under the constraint of accepting at most k in total. Our work proposes an infinite-dimensional linear programming formulation that fully characterises the worst-case tight approximation ratio of the k-selection prophet inequality problem, complementing the recent semi-infinite linear programming general approach by Jiang et al. [EC 2023]. Notably, we

introduce a nonlinear system of differential equations that generalises Hill and Kertz's equation. For small k, we observe that this approach yields the best approximation ratios to date.

Table of contents

List of figures

List of tables

Chapter 1

Introduction

1.1 Main concepts

1.1.1 Prophet inequalities

The standard *prophet inequality* as introduced in [57, 58] is about a gambler who is given the choice of opening one of n treasure boxes, each of which contains a single value. The values of all boxes are distributed according to known non-negative independent random variables. As the gambler opens each box in sequence, a value is drawn from the distribution of that box. Upon seeing this value, the gambler has to make the choice of accepting it, or forgoing it forever and moving onto the next box. [57, 58] showed that the optimal strategy for the gambler recovers a tight fraction of $1/2$ of the largest value among all boxes, in expectation. Formally, n agents with values $v_1, \ldots, v_n \geq 0$ distributed independently according to $F_1, \ldots F_n$ arrive in online fashion. For each agent $i \in [n]$, upon arrival v_i is revealed to the algorithm, which has to make an immediate and irrevocable decision to either accept it or move on to the following agent. The first time a value is accepted for some agent $i \in [n]$, the algorithm terminates and hence v_i is the welfare it has generated. The expected value of the algorithm is then compared to the performance of a so-called prophet who sees all values in advance and thus scores $\mathbb{E}[\max_{i\in[n]} v_i]$. We wish to find an algorithm ALG and the smallest possible constant c for which

$$c \cdot \mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\text{ALG}(v)] \geq \mathbb{E}[\max_{i\in[n]} v_i].$$

Whereas [57, 58] accurately analyzed the optimal stopping rule, [73] proposed a simple single threshold algorithm that also recovers the tight solution to the problem. More explicitly, assuming continuous distributions, they choose threshold τ such that $\Pr[\max_{i\in[n]} v_i > \tau] = \frac{1}{2}$. The algorithm accepts a value v_i if and only if $v_i \geq \tau$. Although this algorithm is clearly not optimal for every given instance, it is sufficient to recover the optimal worst case ratio of $c = 2$ when compared to the benchmark. More precisely, for some sequences of distributions F, the optimal algorithm achieves a better approximation ratio than the simple $1/2$ quantile of maximum algorithm. However, there exists an instance where the optimal algorithm also does no better than $c = 2$. This instance is given by two agents where the first has value 1 with probability one and the second has value $\frac{1}{\varepsilon}$ with probability ε and 0 otherwise. In this case, any algorithm will get expected welfare equal to 1, whereas the offline optimum achieves $2 - \varepsilon$.

By viewing the agents as bidders and the threshold τ as a price set by a seller, it becomes clear that this result has applications in economics and mechanism design. Partly due to this connection, there has been much recent work done on extensions to combinatorial settings of the prophet inequality problem, such as matroids [2, 18, 37, 55], general downward-closed set systems [72] and combinatorial auctions [20, 27, 29, 34, 36]). Let us consider more closely the setting of combinatorial extensions given by feasibility constraints and additive valuations. Let $\mathcal{F}$ be a downwards closed set system defined on $\{1, \ldots, n\}$ along with corresponding known non-negative independent distributions $F_1, \ldots, F_n$. The perspective of the gambler is the same as for the standard prophet inequality where values drawn from $F_1, \ldots, F_n$ arrive sequentially and a decision must be made on each one in online fashion. However, instead of stopping after accepting a single value, the gambler now collects a set $S \in \mathcal{F}$. The value achieved by the gambler is $\mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\sum_{i\in S} v_i]$. It is compared to the offline benchmark, $\mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\max_{T\in\mathcal{F}}\{\sum_{i\in T} v_i\}]$. [55] give a remarkable generalization of the standard prophet inequality by showing that when $\mathcal{F}$ is a matroid, the same ratio of 2 is achieved, that is

$$2 \cdot \mathbb{E}_{v\sim(F_1,\ldots,F_n)}\left[\sum_{i\in S} v_i\right] \geq \mathbb{E}_{v\sim(F_1,\ldots,F_n)}\left[\max_{T\in\mathcal{F}}\{\sum_{i\in T} v_i\}\right].$$

Another major reason for the renewed interest in prophet inequalities is their relevance to auctions, specifically posted priced mechanisms (PPM) in online sales [3, 18, 20, 27, 42, 55]. It was implicitly shown by [18] and [42] that every prophet-type inequality implies a corresponding approximation guarantee in a PPM, and the converse holds as well [24].

1.1.2 Competition complexity

In any setting where the goal is to compare the optimal offline algorithm against an online algorithm, resource augmentation is the general concept of realizing this analysis while granting only the online algorithm some additional amount of resources. For example, in the case of an online scheduling problem, the online algorithm could have access to faster processors compared to the clairvoyant optimum. This type of resource augmentation was considered in [51], who applied their results to obtain new insights on the classic uniprocessor CPU scheduling problem. They showed that at the expense of granting the online algorithm a CPU that is slightly faster compared to the one used by the optimal offline algorithm, the performance of these two are within a constant factor of each other, with respect to the number of jobs. This is remarkable since it is known not to be true when both CPUs have the same performance. Resource augmentation has a long history of success stories in the design and analysis of algorithms. It is particularly popular in scheduling [51, 69], but it has also been studied in many other areas such as paging [75], bin packing [25], or selfish routing [71]. 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 may also grant the online algorithm access to additional 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 selfish routing, comparing the equilibrium flow to the optimal flow can lead to an arbitrarily large price of anarchy. However, one can consider the equilibrium flow on an augmented version of the original instance, with an increased traffic rate. Comparing this to the optimal flow on the original instance yields interesting results.

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. This concept is called *competition complexity* and has been the subject of study in some excellent recent work by [30], [11] and [8]. To the best of our knowledge, our work, published in [13], is the first to consider competition complexity for posted price auctions. As we will further discuss in the following section, for example in the setting of allocating an item to several agents with the intent of maximizing welfare, online allocation mechanisms have practical advantages over say a more traditional second price auction. Being faced with a take it or leave it price is more time efficient and easier to understand since the price paid is fixed and doesn't depend on other agent's bids. Thus it is fair to assume that an online posted price mechanism may recruit more bidders than the optimal, second price mechanism.

1.2 Overview of our results

In this thesis, we contribute to different extensions of the classic prophet inequality.

1.2.1 Competition complexity of sequential posted pricing, i.i.d distributions

In the standard prophet inequality setting, the benchmark $\mathbb{E}[\max_{i\in[n]} v_i]$ can alternatively be viewed as the expected welfare of the optimal, not necessarily sequential, incentive compatible mechanism. In fact a welfare of $\mathbb{E}[\max_{i\in[n]} v_i]$ can be achieved by a second price auction. That is, an auction where upon receiving all bids b_i, the winner is declared to be the agent with the highest bid and their payment is equal to the second highest bid. It is well known that this auction is strongly incentive compatible, meaning rational bidders will submit their true value as their bid, $b_i = v_i$ [19, 41, 76]. In many practical scenarios where such auctions would be deployed, it is often the case that a simple posted price sequential auction is more appealing than the second price auction to a number of buyers. One reason is, having the price depend on another agent's bid adds an element of uncertainty. Another is that the result of a second price auction can only be announced once all agents have submitted their bid, whereas in the sequential auction, bidders are immediately notified of the outcome. Soliciting bids can potentially be a time consuming process when it is not an automated procedure. Thus due to the increased simplicity of the online sequential posted price auction, it is fair to assume that in a significant portion of instances, it will attract more participants than the second price auction. In this context, it becomes important to be able to compare the performance of the second price auction on n agents with the performance of the optimal sequential posted price auction on $m \geq n$ agents. We refer to the dependence of m on n and ε as the competition complexity of dynamic pricing.

In Chapter 2, this motivation leads us to study the following problem. Let $A_m(F)$ be the expected social welfare achievable by the optimal sequential posted price policy on m i.i.d. draws from a random variable X with cumulative distribution function F. We compare it to the expected maximum $M_n(F)$ of $n \leq m$ i.i.d. draws. We note that the optimal online policy is indeed sequential posted price. This is easy to see when F is continuous. Any online policy is equivalent to a set of prices P_t at time each time $t \leq m$ for which the item is sold. Let $q_t = \Pr[X \in P_t]$. Then we see that the posted price strategy that posts $z(t) := \Pr[X \geq z(t)] = q_t$ has larger or equal expected welfare. An instance where F may by discontinuous reduces to one with a continuous distribution in that, if we allow an algorithm to allocate at a given price with any chosen probability (rather than either allocating with probabilities 0 or 1), then it can execute any quantile strategy $z(t)$ as above. We find that for any $m \geq n$, there exists F such that $A_m(F) < M_n(F)$. In light of this result, we investigate how large the ratio between these two quantities can become, given the size of m relative to n. For fixed $\varepsilon \geq 0$ and fixed n, we want to find the smallest $m \geq n$ such that for every F we have

$$(1+\varepsilon) \cdot A_m(F) \geq M_n(F).$$

The solution is expressed in terms of the function $\phi : \mathbb{R}_+ \to \mathbb{R}_+$ given by

$$\phi(\varepsilon) = \int_0^1 \frac{1}{y(1-\log(y))+\varepsilon} \mathrm{d}y$$

where log refers to the natural logarithm. We find that for any $\varepsilon > 0$ and any n, we have $(1+\varepsilon)A_m(F) \geq M_n(F)$ for every F if $m \geq \phi(\varepsilon)n$, and for large n this is tight. To get some intuition on this function, it can be shown that $\phi(\epsilon)$ is $\Theta(\log(\log(1/\varepsilon)))$ as $\varepsilon \to 0$, with very small constants hidden in the big-O notation.

We now give a brief summary of how the results were obtained, which in part illustrates the role optimization played in our work. The general approach in this chapter is to formulate our question as an optimization problem, which we then analyze precisely, without making much use of approximations. The values of the sequence $(A_i(F))_{i \in [m]}$ have a simple description in that they satisfy the following recurrence:

$$A_0(F) = 0, \quad A_1(F) = \mathbb{E}(X), \; \ldots, \; A_{i+1}(F) = \mathbb{E}(\max\{X, A_i(F)\}) \quad \forall i \in [m-1].$$

where X is a random variable distributed according to F as noted and analyzed in [45]. Further, this formula makes a lot of sense by observing the following. From the perspective of the algorithm, let us start at the end and assume we have not allocated the item and there is only a single agent, agent m, remaining. In this case, we definitely want to allocate the item, thus pricing it at 0. Given that our strategy is clear once we're faced with the last agent, what to do at the previous step? For agent $m-1$, we know that if we don't allocate to this agent, we can always recover expected welfare $\mathbb{E}[X]$ from the following agent. Hence we are interested in allocating the item if and only if agent $m-1$ values it more than $\mathbb{E}[X]$, and so we sell it at this price, $\mathbb{E}[X]$. Backtracking in this way, we always price the item for a given agent at price exactly equal to the expected welfare we could obtain from remaining agents. This yields the formula.

We then observe that the values $(A_i(F))_{i\in\mathbb{N}}$, of the optimal posted prices given some F only depend on the value of the integral of F on certain intervals, as we will show in section 2.3. In other words, we find that while every feasible sequence $A_i(\cdot)$ may be realized by many different probability distributions, there is one distribution that gives the worst case with respect to our question. As a consequence, for fixed $\varepsilon > 0$ and $m \geq n > 0$, we are able to reduce the question of finding if

$$\min_{F\in\Delta}(1+\varepsilon)A_m(F) - M_n(F) > 0, \tag{1.1}$$

where Δ is the space of all cumulative distribution functions. Instead, we obtain the following infinite-dimensional, non-linear optimization problem and ask if it has a non-negative objective.

$$\begin{aligned} \text{minimize} \quad & (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i - \sum_{i=0}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i \\ \text{subject to} \quad & \delta_{j+1} \leq \delta_j \qquad \text{for every integer } j \geq 0, \\ & \delta_j^2 \leq \delta_{j-1}\delta_{j+1} \qquad \text{for every integer } j \geq 1, \\ & \delta_0 = 1 \text{ and } \delta_j > 0 \quad \text{for every integer } j \geq 1. \end{aligned} \tag{1.2}$$

In this optimization problem, we should interpret δ_i as $A_{i+1}(F) - A_i(F)$. It follows that clearly the sum $\sum_{i=0}^{m-1}\delta_i = A_m(F)$. A significant observation that we describe carefully in Section 2.3 is that in expression (1.1), we only need to consider a simple subset of Δ for which indeed $M_n(F) = \sum_{i=0}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i$.
This new formulation of the problem turns out to be useful because although it is non-linear, we do have convexity, allowing us to use standard optimization tools. Indeed, first-order methods yield a solution by means of a recurrence relation. Given $\varepsilon > 0$ and a positive integer $n \geq 2$, let $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$ be the sequence defined by the following recurrence:

$$\rho_{\varepsilon,1} = 1, \text{ and } (n-1)\rho_{\varepsilon,j-1}^n - \varepsilon = n\rho_{\varepsilon,j}^{n-1} \text{ for every } j \geq 2 \text{ for which } (n-1)\rho_{\varepsilon,j-1}^n - \varepsilon > 0.$$

These terms are decreasing in j and we will show that our original question is equivalent to finding the last term of the sequence, or the largest m such that $\rho_{\varepsilon,m}$ is well-defined. Finally, we return to continuous methods as we analyze the recurrence $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$ by accurately approximating it by an ordinary differential equation (ODE).

It may seem as if the appearance of an ODE is a bit surprising and not an essential part of the proof. However, we know that already [45] used the recursion from the optimal dynamic program in [45] to provide an ODE where the approximation ratio β is embedded as a unique constant that guarantees crucial analytical properties of the solution of the ODE:

$$y' = y(\ln y - 1) - 1/\beta + 1, \quad y(0) = 1, \quad y(1) = 0.$$

We also observe that in Chapter 4, which addresses a very closely related problem, ODEs play a central role.

In order to prove existence of a unique solution to the ODE, we will make use of a simple global extension to the well known Picard-Lindelöf theorem [70] in the case of Lipschitz functions. We consider solutions to the initial value problem

$$\begin{aligned} x' &= f(t,x) \\ x(t_0) &= x_0 \end{aligned}$$

Theorem 1. *Let $f(t,x)$ be a continuous function $f : D \to \mathbb{R}$, where $D \subset \mathbb{R} \times \mathbb{R}$ open, and f is locally Lipschitz with respect to the x variable. Then for any $(t_0, x_0) \in D$, there is a unique local solution to $\varphi(t)$ to the intial value problem. More precisely, there exists some small interval (a,b) including t_0 where $\varphi(t)$ exists and is unique.*

1.2.2 Competition complexity for sequential posted pricing, independent distributions.

In Chapter 3, we explore a more general setting, where we drop the assumption that distributions must be identical. We start by defining the standard prophet inequality setting, with bidders having independent values drawn respectively from distributions $F_1 \dots, F_n$. The benchmark becomes $\mathbb{E}_{v \sim (F_1,\dots,F_n)}[\max_{i \in [n]} v_i]$. We ask how many copies of $(F_1, \dots, F_n)$ must a sequential algorithm see in order to recover a $(1-\varepsilon)$ fraction of the benchmark. Let us a give a more detailed description of the setting.

The Prophet Inequality Setting. 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 $\rho \in [n] \cup \{\mathsf{null}\}$, 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 $\rho = \mathsf{null}$ 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 $v_1, \ldots, v_n$ 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

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

We seek to give an upper bound on k such that

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

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

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

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 $v_1, \ldots, v_n$ 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*

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

The case $\varepsilon = 0$ is also referred to as *exact* competition complexity; it was shown in [12, 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 B.3, 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 $\boldsymbol{\tau} = (\tau_1, \dots, \tau_k)$, 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 $\boldsymbol{\tau} = (\tau_1, \dots, \tau_{nk})$, 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 [73]. Proposition 8 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.

Approximate Stochastic Dominance. We give an introduction to this notion here because it plays a central role in proving the main result of the chapter. A simplifying explanation for this is, we find that for the competition complexity setting, it is very convenient to work with quantile based algorithms. In order to analyze the latter, the approximate stochastic dominance notion is a natural choice. First-order stochastic dominance between two random variables can be defined as follows. Consider two random variables X, Y. Then we say X stochastically dominates Y if for any outcome x, $\Pr[X \geq x] \geq \Pr[Y \geq x]$. We refer to approximate stochastic dominance as this property being satisfied up to a fixed multiplicative constant. That is, X approximately stochastically dominates Y up to factor $c \leq 1$ if for any outcome x, $\Pr[X \geq x] \geq c \Pr[Y \geq x]$. We note that an immediate consequence is that the relation holds in expectation as well. That is $\mathbb{E}[X] \geq c \cdot \mathbb{E}[Y]$.

A notable difference to our approach in Chapter 2 is that we no longer explicitly analyze the optimal sequential mechanism. We note that although such optimal sequential mechanism has a simple, very similar description to that seen in Chapter 2, a direct analysis seems much more complicated. In particular, finding the worst case F for every sequence $A_i(F)$, leading to an optimization problem such as 1.2 becomes problematic. Instead, we choose a fixed threshold for each of the k copies of $(F_1, \ldots, F_n)$. We call these block threshold algorithms. First, we show that the competition complexity of such block threshold prices is asymptotically the same as for the larger class of dynamic pricing policies. Second, to pin down the asymptotic behaviour of the competition complexity in this setting, we choose a set of prices $\boldsymbol{\tau}$ that is not optimal, but more straightforward to analyze. Here it is worth noting that the results of chapter 3 do not imply those of chapter 2. In 2 we obtain a tight description of the competition complexity given by $\phi(\varepsilon)$. In particular, for $\varepsilon = \phi^{-1}(1)$, we recover the optimal approximation factor of ≈ 0.745 of the standard i.i.d prophet inequality (without resource augmentation). From this perspective, 2 is more relevant in terms of practical applications. The results of 3 are only asymptotically optimal in terms of $\frac{1}{\varepsilon}$.

More specifically, given $F_1, \ldots, F_n \in \Delta$, and $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_k)$ such that $p_0 = 0$ and $p_\ell = \Pr_{v \sim (F_1, \ldots, F_n)}[\max_{j \in [n]} v_j \geq \tau_\ell] > 0$ for every $\ell \in \{1, 2, \ldots, k\}$. That is, p_ℓ is the quantile of threshold τ_ℓ with respect to the distribution of the largest value. Define

$$\Phi_1(F_1,\ldots,F_n,\boldsymbol{\tau}) = \sum_{\ell=1}^{k}\prod_{j=1}^{\ell}(1-p_j),$$

$$\Phi_i(F_1,\ldots,F_n,\boldsymbol{\tau}) = \frac{1}{p_i}\sum_{\ell=1}^{i-1}p_\ell\prod_{j=0}^{\ell-1}(1-p_j) + \sum_{\ell=i}^{k}\prod_{j=0}^{\ell}(1-p_j) \quad \text{for every } i \in \{2,\ldots,k\}, \text{ and}$$

$$\Phi_{k+1}(F_1,\ldots,F_n,\boldsymbol{\tau}) = \sum_{\ell=1}^{k}p_\ell\prod_{j=0}^{\ell-1}(1-p_j).$$

The interpretation of the Φ_i for $i \in [n]$ is

$$\Pr_{v\sim(F_1,\ldots,F_n)^k}[\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x] \geq \Phi_i(F_1,\ldots,F_n,\boldsymbol{\tau}) \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x].$$

for $x \in [\tau_{i-1},\tau_i]$, where we define $\tau_0 = 0$. That is, it provides a lower bound on the algorithm's performance for the case when the maximum value lies in the interval $[\tau_{i-1},\tau_i]$.

Let $\Phi(F_1,\ldots,F_n,\boldsymbol{\tau}) = \min_{i\in\{1,\ldots,k+1\}}\Phi_i(F_1,\ldots,F_n,\boldsymbol{\tau})$. We say that $\boldsymbol{\tau} = (\tau_1,\ldots,\tau_k)$ is decreasing if $\tau_j > \tau_{j+1}$ for every $j \in \{1,\ldots,k\}$. Our question becomes how to choose decreasing threshold prices τ_i such that $\Phi(F_1,\ldots,F_n,\boldsymbol{\tau})$ is maximized. In this sense, we have again reduced the initial formulation to an optimization problem over a tangible search space. However, it remains unclear how to find the optimal solution to this reduced problem.

Nonetheless, by making a careful choice of quantile thresholds, we conclude again that it suffices to let $k = \Theta(\log\log(1/\varepsilon))$. Indeed, we obtain the same asymptotic result as in Chapter 2 for a strictly more general setting, and one could argue that the methods used are significantly more simple.

We also analyze the competition complexity of the even simpler class of algorithms that set a single fixed threshold for all of $(F_1,\ldots,F_n)^k$. In this case, we find the asymptotically tight result of $k = \Theta(\log(1/\varepsilon))$.

1.2.3 A linear programming approach to the i.i.d k-selection prophet inequality problem

In Chapter 4, we revisit the well studied i.i.d k-selection prophet inequality problem. The setting is the same as in the original prophet inequality problem, with the only difference being that the auctioneer sells up to k items instead of a single one. In recent years, there has been substantial progress in understanding the approximation limits for prophet inequality problems, mainly driven by their applicability in mechanism design [62]. One of the most prominent settings is the i.i.d. k-selection prophet inequality problem, where the decision-maker selects at most k values from the n observed and aims to maximize the expected sum of values selected. The offline benchmark in this case is $\sum_{t=n-k+1}^{n} \mathbb{E}[X_{(i)}]$ where $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$ are the ordered statistics of the random values $X_1, \ldots, X_n$. Observe that when $k = 1$, this setting corresponds to the classic i.i.d. prophet inequality problem. We refer to this problem as the *k-selection prophet inequality*. When the length of the sequence is n and k selections can be made, we refer for short to this problem as (k, n)-PIP.

In a line of work, the value of $\gamma_{n,k}$ has been proven to be at least $1 - k^k e^{-k}/k! \approx 1 - 1/\sqrt{2\pi k}$ (see, e.g., [5, 10, 16, 27, 77]). Using a different approach, [47] introduced very recently a general optimization framework to characterize worst-case approximation ratios for prophet inequality problems, including the i.i.d. k-selection setting, but it is unclear how to use their framework to obtain provable analytical lower bounds for $k \geq 2$. We propose a method to obtain improved approximation ratios, with results for small k displayed in the following table.

k	1	2	3	4	5
Our approach	0.7454	0.8290	0.8648	0.8875	0.9035
[10]	0.6543	0.7427	0.7857	0.8125	0.8311
[47]	0.7489	0.8417	0.8795	0.9006	0.9143

Table 1.1 Ratio lower bounds for $k \in \{1, \ldots, 5\}$ and comparison with other approaches. [10] offer provable lower bounds, while [47] reports $\gamma_{n,k}$ using their LP approach for $n = 8000$. We acknowledge the gap between our bounds and $\gamma_{n,k}$ for $n = 8000$. We are uncertain whether the gap arises from suboptimality in our approach or from a possible slow convergence rate of $\gamma_{n,k}$ as a sequence in n.

We note that in the table above, while [47] compute numerical approximations of $\gamma_{n,k}$ for up to $n = 8000$ that seem to be decreasing in n, our results are lower bounds on $\liminf_n \gamma_{n,k}$.

Our proof has some high level elements in common with those seen in Chapters 2 and 3. By interpreting the fixed prices chosen by an algorithm in quantile space, we give a description of the problem that is the infinite dimensional linear program $[\mathrm{P}]_{n,k}$. That is, $[\mathrm{P}]_{n,k}$ has value $\gamma_{n,k}$. However, this time we do not analyze the program directly. Instead, we establish a weak duality relation with respect to a dual program, $[\mathrm{D}]_{n,k}$. Our task becomes to find a good feasible solution to the dual $[\mathrm{D}]_{n,k}$. It turns out to be a natural choice to describe such a candidate solution by using the solution to a closed form system of nonlinear ordinary differential equations, given by equations (4.1) - (4.3). Then, two major parts remain to be proven. First, we show that this system of nonlinear ODEs has a solution. Second, it is still non-trivial to show that our candidate LP solution is feasible. Once these two facts are established, we immediately obtain the lower bound on $\liminf_n \gamma_{n,k}$. By finding solutions to the system of non-linear ODE's (4.1) - (4.3) computationally, we find that for small k, our method provides the best provable lower bounds on $\liminf_n \gamma_{n,k}$ to date.

1.3 Related work

The competition complexity of auctions has proven to be an important line of work at the intersection of Economics and Computation [11, 15, 30, 35]. The resource augmentation approach it proposes originates in a seminal paper by Bulow and Klemperer [15], who asked this question for the revenue achievable by the simple but suboptimal second-price auction and Myerson's optimal auction. They showed that for i.i.d. bidders whose valuations are drawn from a regular distribution F, the second-price auction with $n + 1$ bidders is guaranteed to achieve at least the expected revenue of the optimal auction with n bidders. They concluded that rather than going for the more complicated auction mechanism, one could simply attract one more buyer to the simpler auction mechanism.

Subsequent work has extended this basic result to a variety of more complex auction settings [11, 30, 61], and also introduced the idea of approximate competition complexity where instead of shooting for optimality, one aims at 99% or 99.9% of optimum [35].

Chapters 2 and 3 examine the relative power of a simple mechanism (dynamic pricing) to that of an optimal mechanism (the optimal auction) and thus fits under the broader umbrella of *simple vs. optimal mechanisms* (e.g., [43, 44]). At the technical core of our work, we rely on a connection between posted-price mechanisms and prophet inequalities that was pioneered and explored in the last fifteen years [18, 24, 42]. This line of work motivated work on prophet inequalities more generally. Most relevant for us is the work on the i.i.d. single-item prophet inequality [1, 21, 23, 54, 60, 74], but there is also exciting work on combinatorial extensions such as [27, 29, 36, 55]. A closely related line of work has examined the gap between various simple mechanisms including posted-price mechanisms and the optimal mechanism on the same number of bidders [4, 28, 48–50].

Chapter 3 is more specifically related to the recent work of [1, 60] 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. [1] 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. The difference to our work is that they study 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 the results in [1, 60] apply to the random and free-order model.

In Chapter 4 we return to the i.i.d. version of the prophet inequality, introduced by [45], and more specifically to their approach of the problem. Although the Hill and Kertz equation has been used in various recent works [23, 60, 67], to the best of the authors' knowledge, our result for multiple selections, where the approximation ratio is embedded in a nonlinear system, has not been previously explored.

Lastly, linear and convex programming have been a powerful tool for the design of online algorithms. For instance, in online and Bayesian matching problems [39, 63], online knapsack [9, 53], secretary problem [14, 17], factor-revealing linear programs [37, 59], and competition complexity. [67], similar to us, uses a quantile-based linear programming formulation to provide optimal policies in the context of decision-makers with a limited number of actions.

Chapter 2

The Competition Complexity of Dynamic Pricing

2.1 Introduction

In this chapter, we study competition complexity in the context of posted pricing. We focus on the fundamental single-item i.i.d case and compare optimal dynamic pricing versus the optimal auction. While we study the social welfare case, all our results translate to revenue maximization under the standard regularity assumption (see Section 2.2 for a detailed discussion).

Since we are focusing on social welfare, the simplest way to state our question is in prophet inequality terminology. Our goal is to compare the expected reward $A_m(F)$ achievable by the optimal policy found by backward induction on $m \geq n$ i.i.d. draws from a distribution F, to the expected maximum $M_n(F)$ of n i.i.d. draws from F. For fixed $\varepsilon \geq 0$ and fixed n, we want to find the smallest $m \geq n$ such that for every F we have

$$(1+\varepsilon) \cdot A_m(F) \geq M_n(F).$$

We refer to the functional dependence of m on n and ε as the *competition complexity of dynamic pricing*. We sometimes refer to the case $\varepsilon = 0$ as *exact* competition complexity and to the case $\varepsilon > 0$ as the *approximate* version.

2.1.1 Warm-up: the uniform case

As a warm-up and to illustrate some of the key ideas in our general competition complexity analysis, consider the case where $F = U[0,1]$ is a uniform distribution over $[0,1]$, and convince ourselves that in this case $A_{2n} \geq M_n$ for all n, so the exact competition complexity is linear. We have that M_n is just the maximum of n i.i.d. draws from a uniform distribution over [0,1], and therefore $M_n = n/(n+1)$. On the other hand, we can compute A_n through the usual backward induction: The recursion is $A_{n+1} = \mathbb{E}(\max\{X, A_n\})$ for $n \geq 1$ and $A_1 = \mathbb{E}(X)$ where $X \sim U[0,1]$. That is, $A_1 = 1/2$, and for $n \geq 1$,

$$\begin{aligned} A_{n+1} &= \mathbb{E}(\max\{X, A_n\}) \\ &= A_n \Pr(X < A_n) + \mathbb{E}(X \mid X \geq A_n) \Pr(X \geq A_n) \\ &= A_n^2 + \frac{(1+A_n)}{2}(1-A_n) = \frac{1}{2}(1+A_n^2). \end{aligned}$$

Observe that apart from getting an exact formula for the recurrence, we get a simple expression for $A_{n+1} - A_n$, that is, the marginal gain of the optimal algorithm when we add one more buyer: $A_{n+1} - A_n = (1-A_n)^2/2$ for $n \geq 1$. In particular, this idea will be further exploited to understand the competition complexity of general distributions.

To analyze the competition complexity for the uniform case, we proceed by induction. It is easy to verify that the claim holds for $n = 1$ since $A_2 = 5/8 > 1/2 = M_1$. So we assume $A_{2n} \geq M_n$, and we want to show $A_{2n+2} \geq M_{n+1}$. Note that if $A_{2n+1} \geq M_{n+1}$ then also $A_{2n+2} \geq A_{2n+1} \geq M_{n+1}$, and there we are done, so we consider the case $A_{2n+1} < M_{n+1}$. We have

$$\begin{aligned} A_{2n+2} &= A_{2n} + (A_{2n+2} - A_{2n+1}) + (A_{2n+1} - A_{2n}) \\ &= A_{2n} + \frac{1}{2}(1 - A_{2n+1})^2 + \frac{1}{2}(1 - A_{2n})^2. \end{aligned}$$

Since the function $f(x) = x + \frac{1}{2}(1-x)^2$ is increasing in $\mathbb{R}_+$, and given that $A_{2n} \geq M_n$, we obtain a lower bound that together with $A_{2n+1} < (n+1)/(n+2)$ yields

$$A_{2n+2} \geq M_n + \frac{1}{2}\left(\left(\frac{1}{n+1}\right)^2 + \left(\frac{1}{n+2}\right)^2\right).$$

The argument is completed by observing that what we add to M_n on the right-hand side is at least $M_{n+1} - M_n = 1/((n+1)(n+2))$. We conclude that for the uniform distribution, it suffices to choose $m \geq 2n$. A closer examination of the asymptotic behavior of A_m and M_n shows that this analysis is in fact tight. Indeed for large m and n, $A_m \approx 1 - 2/(m + \log(m) + 1.76799)$ [38, 64] while $M_n \approx 1 - 1/n$ which roughly shows that we need $m = 2n + o(n)$.

2.1.2 Our contribution

The above analysis of the uniform case already rules out a "plus constant" result as in Bulow and Klemperer [15]. It leaves some hope that the exact competition complexity of dynamic pricing may be linear or, if not, then at least polynomial with a small polynomial. Our first main result shows that this hope is unfounded. Indeed, the exact competition complexity is not only "large," it is in fact unbounded.

Main Result 1 (exact competition complexity): For any $m \geq n$, there exists a distribution F such that $A_m(F) < M_n(F)$.

In light of this strong impossibility, a natural question is whether this impossibility persists if we relax our goals and aim for 99% or 99.99% of optimal. It turns out that things change, and quite drastically so. This is formalized by our second main result, which nails down the approximate competition complexity in terms of function $\phi : \mathbb{R}_+ \to \mathbb{R}_+$ given by

$$\phi(\varepsilon) = \int_0^1 \frac{1}{y(1-\log(y))+\varepsilon} \mathrm{d}y.$$

Main Result 2 (approximate competition complexity): Consider $\varepsilon > 0$ and any n. Then, we have $(1+\varepsilon)A_m(F) \geq M_n(F)$ for every F if $m \geq \phi(\varepsilon)n$, and for large n this is tight.

While our first main result shows that the exact competition complexity of dynamic pricing is unbounded, our second main result shows that if we aim for approximate optimality, then the competition complexity not only drops from being unbounded to being linear, it is actually *linear with a very small constant*.

We illustrate this in Figure 2.1. In the technical part of the chapter, we show that the function $\phi(\epsilon)$ grows as $\Theta(\log\log 1/\varepsilon)$ as $\varepsilon \to 0$, with very small constants hidden in the big-O notation. For example, to obtain 99% of optimal it is sufficient to have $m \geq 2.30 \cdot n$, and to obtain 99.99% of optimal it is sufficient to have $m \geq 2.53 \cdot n$.

An interesting implication of our analysis is that it yields the factor 0.745 i.i.d. prophet inequality [23, 54, 60, 74] and its tightness [45] as a special case. Here is how: Rather than fixing ε and finding $m(n,\varepsilon)$, we may fix $m(n,\varepsilon) = n$ and find ε. The equality $m(n,\varepsilon) = \phi(\varepsilon)n$ corresponds to solving $\phi(\varepsilon) = 1$. This yields $\varepsilon = \phi^{-1}(1)$ and corresponds to an approximation guarantee of $1/(1+\phi^{-1}(1)) \approx 0.745$.

2.1.3 Our techniques

Our argument for the uniform distribution $F = U[0,1]$ that we presented above relied on a formula for the differences between two consecutive terms A_{n+1} and A_n, and at its

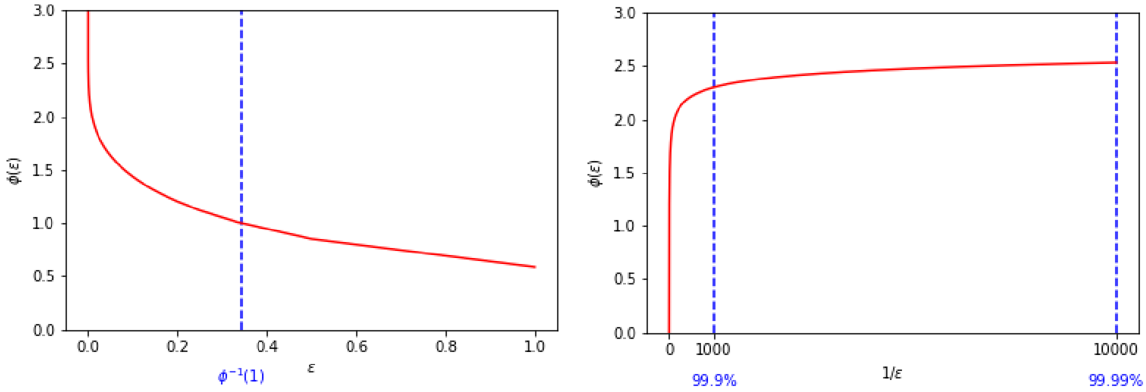

Fig. 2.1 Plot of $\phi(\varepsilon)$ as a function of ε on the left, and as a function of $1/\varepsilon$ on the right. Plotting $\phi(\varepsilon)$ as a function of $1/\varepsilon$ serves to illustrate the very slow growth of $\phi(\varepsilon)$ as $\Theta(\log\log 1/\varepsilon)$ when $\varepsilon \to 0$. The dashed blue line in the left plot is at $\varepsilon = \phi^{-1}(1) \approx 0.342$ which implies the optimal factor $1/(1+\phi^{-1}(1)) \approx 0.745$ for the i.i.d. prophet inequality. In the other plot the two blue dashed lines are at $1/\varepsilon = 100$ and $1/\varepsilon = 1000$ which correspond to approximation ratios of 99.9% and 99.99%. The value of $\phi(\varepsilon)$ at these points is the constant required to obtain these approximation ratios.

core compared $A_{2(n+1)} - A_{2n}$ to $M_{n+1} - M_n$. Intuitively, we explored properties of the rate of growth and curvature of the two sequences $A_1, A_2, \ldots, A_m$ and $M_1, M_2, \ldots, M_n$.

Our general argument builds on this intuition. Our first key observation characterizes the sequences $A_1, A_2, \ldots, A_m$ that can arise. Namely, we show that for any distribution F, the corresponding infinite sequence $(A_i(F))_{i\in\mathbb{N}}$ satisfies the following three properties. Moreover, for any infinite sequence $(A_i)_{i\in\mathbb{N}}$ satisfying these properties there is a distribution F that leads to this sequence. The three properties are:

(1) The sequence $(A_i)_{i\in\mathbb{N}}$ is non-decreasing,

(2) The sequence $(A_{i+1} - A_i)_{i\in\mathbb{N}}$ is non-increasing, and

(3) The sequence $((A_{i+2} - A_{i+1})/(A_{i+1} - A_i))_{i\in\mathbb{N}}$ is non-decreasing.

Our second key observation is that given a *fixed* infinite sequence $(A_i)_{i\in\mathbb{N}}$ with these properties, we can identify the compatible distribution F that maximizes M_n. This worst-case distribution is a simple piece-wise constant distribution, and allows us to express the largest possible M_n as a function of the $(A_i)_{i\in\mathbb{N}}$. We thus reduce the problem of checking whether for a fixed n and m, $(1+\varepsilon)A_m(F) - M_n(F) \geq 0$ for all F, to an infinite dimensional optimization problem that seeks to minimize $(1+\varepsilon)A_m(F) - M_n(F)$ over all infinite sequences satisfying properties (1)–(3): The inequality is satisfied by all F if and only if the objective value of this infinite-dimensional optimization problem is non-negative. To show our two main results, we then solve this infinite-dimensional optimization problem optimally. This reduces the problem to the analysis of a recursion, which can be pointwise bounded by a differential equation, which, by a careful analysis, leads to the function $\phi(\varepsilon)$.

2.2 Formal statement of our results

For our analysis, it will be convenient to consider $\mathbb{N} = \{0,1,2,\dots\}$, the natural numbers including zero. We consider distributions F over the non-negative reals with finite expectation. For a distribution F, we let $M_0(F) = 0$ and for $n \geq 1$ we let $M_n(F) = \mathbb{E}(\max\{X_1, X_2, \dots, X_n\})$, where $X_1, \dots, X_n$ is an i.i.d. sample distributed according to F. We denote by $A_n(F)$ the value of the optimal policy and the sequence $(A_n(F))_{n\in\mathbb{N}}$ satisfies the following recurrence: $A_0(F) = 0$, $A_1(F) = \mathbb{E}(X)$ and $A_{n+1}(F) = \mathbb{E}(\max\{X, A_n(F)\})$, where X is a random variable distributed according to F. We now formally state our main results.

Theorem 2. *For every positive integer $n > 1$, and every positive integer $m \geq n$, there exists a distribution F such that $A_m(F) < M_n(F)$.*

Theorem 3. *Let $\varepsilon > 0$ and let n be a positive integer. Then, for every $m \geq \phi(\varepsilon)n = \Theta(\log\log 1/\varepsilon)n$, and every distribution F we have $(1+\varepsilon)A_m(F) \geq M_n(F)$. Conversely, for any $\delta > 0$, there exists a distribution G such that for n sufficiently large and $m < (\phi(\varepsilon) - \delta)n$, we have $(1+\varepsilon)A_m(G) < M_n(G)$.*

While Theorem 2 shows that the exact competition complexity of dynamic pricing is unbounded, Theorem 3 shows that the approximate competition complexity not only drops from being unbounded to being linear, it is actually linear with a very small constant (see Figure 2.1).

As mentioned in the introduction, Theorems 2 and 3 translate to the case of revenue by using standard reductions between social welfare and revenue optimization for the i.i.d. case [18, 24, 42]. Given a distribution F, the *virtual valuation* of F is the function $\phi_F(x) = x - (1 - F(x))/f(x)$, where f is the probability density function of F. To construct an algorithm for the revenue setting in the i.i.d. case with distribution F and n buyers, we reduce to the social welfare case as follows: We run the optimal dynamic welfare policy for an instance with n buyers identically and independently distributed according to F^ϕ, where F^ϕ is the distribution of the random variable $\tilde{\phi}_F(X) = \max(0, \phi_F(X))$ when X is distributed according to F. By doing so, the optimal dynamic welfare policy is defined by thresholds $\tau_1, \dots, \tau_n$, which can be converted into optimal dynamic revenue prices (a posted price mechanism) with $p_i = \tilde{\phi}_F^{-1}(\tau_i)$, for every $i \in \{1, \dots, n\}\}$, when F is *regular*, i.e., ϕ_F is monotone non-decreasing [42]. We remark that this reduction is based in the classic result of Myerson for revenue maximizing single-item auctions [65].

2.3 An equivalent optimization problem

In this section, we develop the main building block of our analysis. The key result of this section, Theorem 4, shows that the question of whether for a given $\varepsilon \geq 0$, $n \geq 1$, and $m \geq 1$ it holds that $(1+\varepsilon)A_m(F) \geq M_n(F)$ for all F reduces to showing whether the following infinite-dimensional, non-linear optimization problem has a non-negative objective.

$$\text{minimize} \quad (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i - \sum_{i=0}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i \tag{2.1}$$

$$\text{subject to} \quad \delta_{j+1} \leq \delta_j \quad \text{for every integer } j \geq 0, \tag{2.2}$$

$$\delta_j^2 \leq \delta_{j-1}\delta_{j+1} \quad \text{for every integer } j \geq 1, \tag{2.3}$$

$$\delta_0 = 1 \text{ and } \delta_j > 0 \quad \text{for every integer } j \geq 1. \tag{2.4}$$

Theorem 4. *Let $\varepsilon \geq 0$, and let n and m be two positive integers. Then, we have $(1+\varepsilon)A_m(F) \geq M_n(F)$ for every distribution F if and only if the optimal value of the optimization problem* (2.1)-(2.4) *is non-negative.*

We prove this theorem by characterizing the sequences $(A_j(F))_{j\in\mathbb{N}}$ that can result from distributions F and by relating the value of $M_n(F)$ to the values of the sequence $(A_j(F))_{j\in\mathbb{N}}$. The characterization uncovers the properties of the sequences that can arise. Given a sequence of non-negative real values $(S_n)_{n\in\mathbb{N}}$, we denote by $(\partial S_n)_{n\in\mathbb{N}}$ the sequence such that $\partial S_n = S_{n+1} - S_n$ for every non-negative integer n. Consider the following properties:

(a) The sequence $(S_n)_{n\in\mathbb{N}}$ is strictly increasing.

(b) The sequence $(\partial S_n)_{n\in\mathbb{N}}$ is non-increasing.

(c) The sequence $(\partial S_{n+1}/\partial S_n)_{n\in\mathbb{N}}$ is non-decreasing.

Observe that the properties (b)-(c) imply that the sequence $(\partial S_{n+1}/\partial S_n)_{n\in\mathbb{N}}$ is not only non-decreasing, but also bounded with $\partial S_{n+1}/\partial S_n \leq 1$ for every $n \in \mathbb{N}$, and therefore it is convergent to a limit value of at most one. In what follows, given a distribution F, let $\omega_0(F) = \inf\{y \in \mathbb{R} : F(y) > 0\}$ and $\omega_1(F) = \sup\{y \in \mathbb{R} : F(y) < 1\}$ be the left and right endpoints of the support of F.

We need a few lemmas to prove Theorem 4. We also use the following proposition about the optimal policy.

Proposition 1. *For every distribution F the following holds:*

(i) $A_{n+1}(F) = A_n(F) + \int_{A_n(F)}^{\infty}(1 - F(y))\mathrm{d}y$ *for every* $n \in \mathbb{N}$.

(ii) $A_{n+2}(F) = A_{n+1}(F) + \int_{A_n(F)}^{A_{n+1}(F)} F(y)\mathrm{d}y$ *for every* $n \in \mathbb{N}$.

(iii) $\lim_{n\to\infty} A_n(F) = \omega_1(F)$.

(iv) *If* $\omega_0(F) < \omega_1(F)$ *and* F *has finite expectation, then* $A_n(F) < A_{n+1}(F)$ *for every* $n \in \mathbb{N}$.

Proof. Since $A_{n+1}(F) = \mathbb{E}(\max\{A_n(F), X\})$, where X is distributed according to F, we get

$$A_{n+1}(F) = A_n(F)F(A_n(F)) + \int_{A_n(F)}^{\infty} sf(s)\mathrm{d}s.$$

By integrating by parts, we have

$$\int_{A_n(F)}^{\infty} sf(s)\mathrm{d}s = (1 - F(A_n(F)))A_n(F) + \int_{A_n(F)}^{\infty}(1 - F(s))\mathrm{d}s,$$

and therefore (i) holds since we have

$$\begin{aligned} A_{n+1}(F) &= A_n(F)F(A_n(F)) + (1 - F(A_n(F))A_n(F) + \int_{A_n(F)}^{\infty}(1 - F(s))\mathrm{d}s \\ &= A_n(F) + \int_{A_n(F)}^{\infty}(1 - F(s))\mathrm{d}s. \end{aligned}$$

To prove (ii), observe that

$$\begin{aligned} \int_{A_n(F)}^{\infty}(1 - F(s))\mathrm{d}s &= \int_{A_n(F)}^{A_{n+1}(F)}(1 - F(s))\mathrm{d}s + \int_{A_{n+1}}^{\infty}(1 - F(s))\mathrm{d}s \\ &= A_{n+1}(F) - A_n(F) - \int_{A_n(F)}^{A_{n+1}(F)} F(s)\mathrm{d}s + \int_{A_{n+1}(F)}^{\infty}(1 - F(s))\mathrm{d}s, \end{aligned}$$

and therefore, (i) implies that

$$\int_{A_n(F)}^{A_{n+1}(F)} F(s)\mathrm{d}s = \int_{A_{n+1}(F)}^{\infty}(1 - F(s))\mathrm{d}s = A_{n+2}(F) - A_{n+1}(F),$$

where the last equality holds also by (i).

We now show (iii). Let $L = \lim_{n\to\infty} A_n(F)$ and assume, for the sake of contradiction, that $L < \omega_1(F)$. Since $(A_n(F))_{n\in\mathbb{N}}$ is non-decreasing, we have $A_n(F) \leq L$ for every n. Let $U = \min\{L + 1, (L + \omega_1(F))/2\}$. From (i), we have

$$\begin{aligned}
A_{n+1}(F) - A_n(F) &= \int_{A_n(F)}^{\infty} (1 - F(y))dy \\
&= \int_{A_n(F)}^{\omega_1(F)} (1 - F(y))dy \\
&\geq \int_{L}^{U} (1 - F(y))dy \geq (U - L)(1 - F(U)) > 0,
\end{aligned}$$

where the first inequality holds since $A_n(F) \leq L < U \leq \omega_1(F)$, and the second inequality holds since F is non-decreasing. The last inequality follows by the definition of $\omega_1(F)$ and using that $L < U < \omega_1(F)$. Since this inequality holds for all $n \in \mathbb{N}$, it implies that

$$A_{n+1}(F) = \sum_{j=0}^{n} (A_{j+1}(F) - A_j(F)) \geq \frac{n+1}{2}(U - L)(1 - F(U)) \to \infty$$

as $n \to \infty$, which contradicts that $L < \omega_1(F) \leq \infty$. Finally, we show (iv). Since F has a finite expectation, $\omega_0(F) < \omega_1(F)$ and the support is contained in the non-negative reals, we have that $A_1(F) = \mathbb{E}(X) > 0 = A_0(F)$. Then, the property holds by induction on n and property (ii). $\square$

An important implication of Proposition 1(iv) is that the sequence $(A_j(F))_{j \in \mathbb{N}}$ is strictly increasing unless F is a distribution that puts probability one on a single value. For these distributions F, however, $A_m(F) = M_n(F)$ for all $m, n \geq 1$, so they trivially satisfy $(1 + \varepsilon)A_m(F) \geq M_n(F)$.

In the remainder, we will consider distributions F with $\omega_0(F) < \omega_1(F)$. We begin by showing that for such distributions F the sequence $(A_j(F))_{j \in \mathbb{N}}$ satisfies properties (a)-(c).

This lemma is a first step towards showing that the optimization problem or equivalently the inequality $(1 + \varepsilon)A_m(F) \geq M_n(F)$ we are studying may be simplified. That is, we do not need to consider the space of all probability distributions, because for one, the set of sequences $A_m(F)$ actually has a particular structure for any distribution F. In lemma 2, we strengthen this argument by showing that the other term that appears, $M_n(F)$, can also be bounded by a rather simple expression.

Lemma 1. *For every distribution F with $\omega_0(F) < \omega_1(F)$, the sequence $(A_n(F))_{n \in \mathbb{N}}$ satisfies properties (a)-(c).*

Proof. Consider a distribution F with $\omega_0(F) < \omega_1(F)$ and a non-negative integer n. Observe that property (a) holds directly for the sequence $(A_n(F))_{n \in \mathbb{N}}$ from Proposition

1(iv). By Proposition 1(ii), it holds that

$$A_{n+2}(F) - A_{n+1}(F) = \int_{A_n(F)}^{A_{n+1}(F)} F(y)\mathrm{d}y \leq A_{n+1}(F) - A_n(F),$$

where the inequality holds since $F(y) \leq 1$ for every $y \in \mathbb{R}$. Therefore, property (b) holds. Observe that thanks to Proposition 1(ii) again, we have

$$\frac{A_{n+2}(F) - A_{n+1}(F)}{A_{n+1}(F) - A_n(F)} = \frac{1}{A_{n+1}(F) - A_n(F)} \int_{A_n(F)}^{A_{n+1}(F)} F(y)\mathrm{d}y,$$

and since F is monotone non-decreasing, we therefore have

$$F(A_n(F)) \leq \frac{A_{n+2}(F) - A_{n+1}(F)}{A_{n+1}(F) - A_n(F)} \leq F(A_{n+1}(F)),$$

from where we conclude that that $(A_n(F))_{n\in\mathbb{N}}$ satisfies property (c). □

Next we show that for the type for distributions we are interested in, it is possible to prove an upper bound on the value of $M_n(F)$ in terms of the values of the sequence $(A_j(F))_{j\in\mathbb{N}}$.

Lemma 2. *For every distribution F with $\omega_0(F) < \omega_1(F)$, we have*

$$M_n(F) \leq \sum_{j=0}^{\infty} \left(1 - \left(\frac{\partial A_{j+1}(F)}{\partial A_j(F)}\right)^n\right) \partial A_j(F).$$

Proof. Consider the concave function $\varphi : \mathbb{R} \to \mathbb{R}$ given by $\varphi(x) = 1 - x^n$, and for every non-negative integer j let $\mu_j(y) = 1/\partial A_j(F)$ for every $y \in [A_j(F), A_{j+1}(F))$ and zero elsewhere. In particular, μ_j is a probability density function over $[A_j(F), A_{j+1}(F))$. Then, by Jensen's inequality, we have

$$\begin{aligned}
\frac{1}{\partial A_j(F)} \int_{A_j(F)}^{A_{j+1}(F)} (1 - F(y)^n)\mathrm{d}y &= \int_{\mathbb{R}} \varphi(F(y))\mu_j(y)\mathrm{d}y \\
&\leq \varphi\left(\int_{\mathbb{R}} F(y)\mu_j(y)\mathrm{d}y\right) \\
&= 1 - \left(\frac{1}{\partial A_j(F)} \int_{A_j(F)}^{A_{j+1}(F)} F(y)\mathrm{d}y\right)^n = 1 - \left(\frac{\partial A_{j+1}(F)}{\partial A_j(F)}\right)^n,
\end{aligned}$$

where the last equality holds by Proposition 1(ii). In particular, for every non-negative integer j we have

$$\int_{A_j(F)}^{A_{j+1}(F)} (1 - F(y)^n)\mathrm{d}y \le \left(1 - \left(\frac{\partial A_{j+1}(F)}{\partial A_j(F)}\right)^n\right) \partial A_j(F). \tag{2.5}$$

Therefore, we have

$$\begin{aligned} M_n(F) = \int_0^\infty (1 - F(y)^n)\mathrm{d}y = \sum_{j=0}^{\infty} \int_{A_j(F)}^{A_{j+1}(F)} (1 - F(y)^n)\mathrm{d}y \\ \le \sum_{j=0}^{\infty} \left(1 - \left(\frac{\partial A_{j+1}(F)}{\partial A_j(F)}\right)^n\right) \partial A_j(F), \end{aligned}$$

where the second equality holds by Proposition 1(iii) and the inequality comes from (2.5). □

Our final ingredient is a reverse to the previous two lemmas. It shows that for any sequence satisfying properties (a)-(c) we can construct a distribution G for which $(A_j(G))_{j\in\mathbb{N}}$ matches the values of the sequence and $M_n(G)$ matches the upper bound on $M_n(G)$ in terms of the values of the sequence.

Lemma 3. *For every $(B_n)_{n\in\mathbb{N}}$ with $B_0 = 0$, and satisfying (a)-(c), there exists a distribution G such that $A_n(G) = B_n$ for every non-negative integer n. Furthermore, we have*

$$M_n(G) = \sum_{j=0}^{\infty} \left(1 - \left(\frac{\partial B_{j+1}}{\partial B_j}\right)^n\right) \partial B_j.$$

Proof. We construct explicitly the distribution G satisfying the statement of the lemma. Recall that since $(B_n)_{n\in\mathbb{N}}$ satisfies properties (b)-(c) the sequence $(\partial B_{n+1}/\partial B_n)_{n\in\mathbb{N}}$ converges to a value $\rho \in (0,1]$. We prove the following claim.

Claim 1. *Suppose that $\rho < 1$. Then, there exists a value $\mathcal{B} > 0$ such that $\lim_{n\to\infty} B_n = \mathcal{B}$.*

Since the sequence $(B_n)_{n\in\mathbb{N}}$ satisfies property (c), we have that $\partial B_n \le \rho \partial B_{n-1}$, and therefore $\partial B_n \le \rho^n \partial B_0 = \rho^n B_1$ for every $n \in \mathbb{N}$. On the other hand, we have

$$B_n = \sum_{j=0}^{n-1} (B_{j+1} - B_j) = \sum_{j=0}^{n-1} \partial B_j \le B_1 \sum_{j=0}^{n-1} \rho^n \le \frac{B_1}{1-\rho},$$

which implies that the sequence $(B_n)_{n\in\mathbb{N}}$ is upper bounded. Since by property (a) the sequence $(B_n)_{n\in\mathbb{N}}$ is strictly increasing, we conclude that $(B_n)_{n\in\mathbb{N}}$ is a convergent sequence and we call $\mathcal{B}$ the value of this limit. This establishes the claim.

We now construct the distribution G satisfying the conditions of the statement. Consider $G : \mathbb{R} \to \mathbb{R}$ defined as follows: $G(x) = 0$ for every $x \in (-\infty, 0)$, for every non-negative integer j and every $x \in [B_j, B_{j+1})$ we have $G(x) = \partial B_{j+1}/\partial B_j$, and let $G(x) = 1$ for every $x \geq \lim_{n\to\infty} B_n$. Since the sequence $(B_n)_{n\in\mathbb{N}}$ satisfies property (a), the function G is well defined for every non-negative integer n. Furthermore, since the sequence $(B_n)_{n\in\mathbb{N}}$ satisfies (c), we have that G is non-decreasing, and property (b) implies that $G(x) \leq 1$ for every $x \in \mathbb{R}_+$. If $\rho = 1$ then $\lim_{x\to\infty} G(x) = 1$. Otherwise, if $\rho < 1$, by Claim 1 there exists a value $\mathcal{B} > 0$ such that $\lim_{n\to\infty} B_n = \mathcal{B}$, and therefore $G(x) = 1$ for every $x \geq \mathcal{B}$. Therefore, we conclude that G is a distribution.

In what follows, we show that $A_n(G) = B_n$ for every non-negative integer n. We proceed by induction. By construction, we have $A_0(G) = 0 = B_0$. Suppose that $B_i = A_i(G)$ for every $i \in \{0, 1, \ldots, n\}$. By Proposition 1, for every positive integer n it holds that

$$\int_{A_{n-1}(G)}^{A_n(G)} G(y)\mathrm{d}y = \int_{A_n(G)}^{\infty} (1 - G(y))\mathrm{d}y = A_{n+1}(G) - A_n(G),$$

and therefore the inductive step implies that

$$\int_{B_{n-1}}^{B_n} G(y)\mathrm{d}y = A_{n+1}(G) - B_n. \tag{2.6}$$

On the other hand, by construction of G it holds that

$$\int_{B_{n-1}}^{B_n} G(y)\mathrm{d}y = \frac{B_{n+1} - B_n}{B_n - B_{n-1}} \cdot (B_n - B_{n-1}) = B_{n+1} - B_n = \partial B_n,$$

and therefore together with (2.6) we conclude that $A_{n+1}(G) = B_{n+1}$. Finally, we have

$$M_n(G) = \int_0^{\infty} (1 - G(y)^n)\mathrm{d}y = \sum_{j=0}^{\infty} \int_{A_j(G)}^{A_{j+1}(G)} (1 - G(y)^n)\mathrm{d}y = \sum_{j=0}^{\infty} \left(1 - \left(\frac{\partial B_{j+1}}{\partial B_j}\right)^n\right) \partial B_j,$$

where the second equality holds since $\lim_{j\to\infty} A_j(G) = \omega_1(G)$, by Proposition 1(iii). □

We are now ready to prove Theorem 4.

Proof. Proof of Theorem 4.

We start by showing that if for some $\varepsilon \geq 0$, $n \geq 1$, and $m \geq 1$, there exists a distribution F such that $(1 + \varepsilon)A_m(F) < M_n(F)$ then the objective value of the optimization problem (2.1)-(2.4) must be negative. Note that for this distribution F it must hold that $\omega_0(F) < \omega_1(F)$ because otherwise $A_m(F) = M_n(F)$, and so we must have $A_{j+1}(F) > A_j(F)$ for all $j \in \mathbb{N}$ by Proposition 1(iv).

We construct a solution $(\delta_j)_{j\in\mathbb{N}}$ for the optimization problem as follows. For every non-negative integer j, let $\delta_j(F) = \partial A_j(F)/\partial A_0(F)$. We begin by showing that the sequence $(\delta_j)_{j\in\mathbb{N}}$ satisfies (2.2)-(2.4). By construction we have $\delta_0(F) = \partial A_0(F)/\partial A_0(F) = 1$, that is, (2.4) holds. By Lemma 1, the sequence $(A_j(F))_{j\in\mathbb{N}}$ satisfies properties (a)-(c). In particular, the sequence $(\partial A_j(F))_{j\in\mathbb{N}}$ is non-increasing and therefore $\delta_{j+1}(F) \leq \delta_j(F)$ for every integer $j \geq 0$, that is, (2.2) is satisfied. The sequence $(\partial A_{j+1}(F)/\partial A_j(F))_{j\in\mathbb{N}}$ is non-decreasing, and therefore $\delta_{j+1}(F)/\delta_j(F) \geq \delta_j(F)/\delta_{j-1}(F)$ for every integer $j \geq 1$, that is, $\delta_j(F)^2 \leq \delta_{j-1}(F)\delta_{j+1}(F)$, and therefore (2.3) is satisfied. Finally, observe that

$$\begin{aligned}
0 &> \frac{1}{\partial A_0(F)}\Big((1+\varepsilon)A_m(F) - M_n(F)\Big) \\
&= (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i(F) - \frac{M_n(F)}{\partial A_0(F)} \\
&\geq (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i(F) - \sum_{j=0}^{\infty}\left(1-\left(\frac{\partial A_{j+1}(F)}{\partial A_j(F)}\right)^n\right)\frac{\partial A_j(F)}{\partial A_0(F)} \\
&= (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i(F) - \sum_{j=0}^{\infty}\left(1-\left(\frac{\delta_{j+1}(F)}{\delta_j(F)}\right)^n\right)\delta_j(F),
\end{aligned}$$

where the first inequality holds by assumption and the second inequality comes from Lemma 2. So, in particular, the last expression of the above chain, which coincides with the objective in (2.1) must be negative.

Conversely, suppose that the value of the optimization problem (2.1)-(2.4) is negative. That is, there exists a sequence $(\delta_j^\star)_{j\in\mathbb{N}}$ satisfying (2.2)-(2.4) such that

$$(1+\varepsilon)\sum_{i=0}^{m-1}\delta_i^\star - \sum_{i=0}^{\infty}\left(1-\left(\frac{\delta_{i+1}^\star}{\delta_i^\star}\right)^n\right)\delta_i^\star < 0. \tag{2.7}$$

Consider the sequence $(B_j)_{j\in\mathbb{N}}$ defined as follows: $B_0 = 0$ and $B_j = \sum_{i=0}^{j-1}\delta_i^\star$ for every $j \geq 1$. In particular, we have

$$B_{j+1} = \sum_{i=0}^{j}\delta_i^\star > \sum_{i=0}^{j-1}\delta_i^\star = B_j$$

for every integer $j \geq 1$, and therefore the sequence $(B_j)_{j\in\mathbb{N}}$ satisfies (a). Since the sequence $(\delta_j^\star)_{j\in\mathbb{N}}$ satisfies (2.2)-(2.3), by construction it holds directly that $(B_j)_{j\in\mathbb{N}}$ satisfies (b)-(c), and therefore by Lemma 3 there exists a distribution G such that

$A_j(G) = B_j$ for every non-negative integer j, and we have

$$\begin{aligned}(1+\varepsilon)A_m(G) &= (1+\varepsilon)B_m \\ &= (1+\varepsilon)\sum_{i=0}^{m-1}\delta_i^\star \\ &< \sum_{i=0}^{\infty}\left(1-\left(\frac{\delta_{i+1}^\star}{\delta_i^\star}\right)^n\right)\delta_i^\star = \sum_{i=0}^{\infty}\left(1-\left(\frac{\partial B_{i+1}}{\partial B_i}\right)^n\right)\partial B_i = M_n(G),\end{aligned}$$

where the last equality also holds by Lemma 1. This finishes the proof of the theorem. □

2.4 Exact competition complexity: proof of theorem 2

We show next how to use Theorem 4 to prove the impossibility result in Theorem 2 about the exact competition complexity.

Proof. Proof of Theorem 2. Letting $\varepsilon = 0$ in Theorem 4, it suffices to show that the value of the optimization problem (2.1)-(2.4) is strictly negative. Consider the sequence $(b_i)_{i\in\mathbb{N}}$ defined as follows: $b_0 = 1$, and $b_1 \in (0,1)$ to be specified later. For every $i \in \{1,\ldots,m-1\}$ let

$$b_{i+1} = b_i\left(\frac{n}{n-1}\right)^{\frac{1}{n}}\left(\frac{b_i}{b_{i-1}}\right)^{\frac{n-1}{n}}, \tag{2.8}$$

and for every $i \geq m$ let $b_{i+1} = b_i^2/b_{i-1}$. We first show that $(b_i)_{i\in\mathbb{N}}$ is feasible for the optimization problem (2.1)-(2.4). By construction the sequence satisfies (2.4). We start with the monotonicity property (2.2). Consider the function $h(x) = (n/(n-1))^{1/n}x^{(n-1)/n}$ and let $h^{(i)}$ be the function obtained from the composition of h with itself i times. From (2.8), we get $b_{i+1}/b_i = h^{(i)}(b_1/b_0) = h^{(i)}(b_1)$ for every $i \in \{0,1,\ldots,m-1\}$. Observe that $h(x)$ is monotone increasing and continuous on $x \in [0,1]$, with $h(0) = 0$, and therefore $h^{(i)}$ is also monotone increasing, continuous and $h^{(i)}(0) = 0$, for every $i \in \{0,1,\ldots,m-1\}$. Since we also know $b_{j+1}/b_j = b_m/b_{m-1}$ for every $j \geq m$, it suffices to prove $b_i/b_{i-1} \leq 1$ for every $i \in \{1,\ldots,m\}$ in order to show that the sequence $(b_i)_{i\in\mathbb{N}}$ satisfies (2.2). To this end, we make any choice of $b_1 \in (0,1)$ in a way that $\max_{i\in\{0,1,\ldots,m-1\}} h^{(i)}(b_1) < 1$. This implies that property (2.2) is satisfied.

Claim 2. *For every $x \in (0,1]$ we have $\left(\frac{n}{n-1}\right)^{\frac{1}{n}} x^{\frac{n-1}{n}} > x$.*

To see this, consider the function $g : \mathbb{R} \to \mathbb{R}$ given by $g(x) = \left(\frac{n}{n-1}\right)^{\frac{1}{n}} x^{\frac{n-1}{n}} - x$. This function is concave in the interval $[0,1]$ and therefore the minimum is attained in either zero or one. Since $g(0) = 0$ and $g(1) = (n/(n-1))^{1/n} - 1 > 0$, we conclude that

$g(x) > 0$ for every $x \in (0,1]$, proving the claim.

In particular, for every $i \in \{1,\dots,m-1\}$ we have

$$\frac{b_{i+1}}{b_i} = \left(\frac{n}{n-1}\right)^{\frac{1}{n}} \left(\frac{b_i}{b_{i-1}}\right)^{\frac{n-1}{n}} = g\left(\frac{b_i}{b_{i-1}}\right) + \frac{b_i}{b_{i-1}} > \frac{b_i}{b_{i-1}},$$

where we used the fact that $0 < b_i/b_{i-1} < 1$. Since $b_{i+1}/b_i = b_m/b_{m-1} < 1$ for every $i \geq m$, we conclude that (2.3) is also satisfied, and therefore the sequence $(b_i)_{i\in\mathbb{N}}$ is feasible for the optimization problem (2.1)-(2.4). We now show that the objective value of the sequence $(b_i)_{i\in\mathbb{N}}$ is strictly negative. We first observe that the objective value is equal to

$$\begin{aligned}
&\sum_{i=0}^{m-1} b_i - \sum_{i=0}^{m-1}\left(1-\left(\frac{b_{i+1}}{b_i}\right)^n\right) b_i - \sum_{i=m}^{\infty}\left(1-\left(\frac{b_{i+1}}{b_i}\right)^n\right) b_i \\
&= \sum_{i=0}^{m-1}\left(\frac{b_{i+1}}{b_i}\right)^n b_i - \sum_{i=m}^{\infty}\left(1-\left(\frac{b_{i+1}}{b_i}\right)^n\right) b_i
\end{aligned}$$

By construction of the sequence we have

$$\begin{aligned}
\sum_{i=0}^{m-1} b_i\left(\frac{b_{i+1}}{b_i}\right)^n &= b_1^n + \frac{n}{n-1}\sum_{i=1}^{m-1} b_i\left(\frac{b_i}{b_{i-1}}\right)^{n-1} \\
&= b_1^n + \frac{n}{n-1}\sum_{i=1}^{m-1} b_{i-1}\left(\frac{b_i}{b_{i-1}}\right)^n = b_1^n + \frac{n}{n-1}\sum_{i=0}^{m-2} b_i\left(\frac{b_{i+1}}{b_i}\right)^n,
\end{aligned}$$

and therefore

$$\begin{aligned}
b_1^n &= \sum_{i=0}^{m-1} b_i\left(\frac{b_{i+1}}{b_i}\right)^n - \frac{n}{n-1}\sum_{i=0}^{m-2} b_i\left(\frac{b_{i+1}}{b_i}\right)^n \\
&= b_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n + \sum_{i=0}^{m-2} b_i\left(\frac{b_{i+1}}{b_i}\right)^n - \frac{n}{n-1}\sum_{i=0}^{m-2} b_i\left(\frac{b_{i+1}}{b_i}\right)^n \\
&= b_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n - \frac{1}{n-1}\sum_{i=0}^{m-2} b_i\left(\frac{b_{i+1}}{b_i}\right)^n.
\end{aligned}$$

By rearranging terms we conclude that

$$\begin{aligned}\sum_{i=0}^{m-1} b_i \left(\frac{b_{i+1}}{b_i}\right)^n &= b_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n + \sum_{i=0}^{m-2} b_i \left(\frac{b_{i+1}}{b_i}\right)^n \\ &= b_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n + (n-1)\left(b_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n - b_1^n\right) \\ &= nb_{m-1}\left(\frac{b_m}{b_{m-1}}\right)^n - (n-1)b_1^n.\end{aligned}$$

Let $\gamma = b_m/b_{m-1}$. We have $\gamma < 1$, $b_m = \gamma b_{m-1}$ and inductively $b_{m+i} = \gamma^{i+1} b_{m-1}$ for every non-negative i. Therefore, overall, the objective value of the sequence is equal to

$$\begin{aligned}&\sum_{i=0}^{m-1}\left(\frac{b_{i+1}}{b_i}\right)^n b_i - \sum_{i=m}^{\infty}\left(1-\left(\frac{b_{i+1}}{b_i}\right)^n\right) b_i \\ &= nb_{m-1}\gamma^n - (n-1)b_1^n - (1-\gamma^n)\sum_{i=0}^{\infty}\gamma^{i+1}b_{m-1} \\ &= nb_{m-1}\gamma^n - (n-1)b_1^n - \frac{(1-\gamma^n)\gamma}{1-\gamma}b_{m-1} \\ &= nb_{m-1}\gamma^n - (n-1)b_1^n - b_{m-1}\sum_{i=1}^{n}\gamma^i \\ &= -(n-1)b_1^n - b_{m-1}\left(\sum_{i=1}^{n}\gamma^i - n\gamma^n\right) < 0,\end{aligned}$$

which concludes the proof. □

We note that the sequence $(b_n)_{n\in\mathbb{N}}$ defined in the proof of Theorem 2 gives one possible construction of a distribution such that $(1+\varepsilon)A_m(F) \geq M_n(F)$. More precisely, $(b_n)_{n\in\mathbb{N}}$ is a sequence such that the value of the optimization problem (2.1)-(2.4) is negative. In other words, it satisfies the properties of $(\delta_j^\star)_{j\in\mathbb{N}}$ (2.7) as defined in (the converse direction of) the proof of Theorem 4.

2.5 Approximate Competition Complexity: Proof of Theorem 3

In this section we show how to use Theorem 4 to derive Theorem 3 about the approximate competition complexity. In particular, we show how to optimally solve the optimization problem (2.1)-(2.4). For every m, n and $\varepsilon > 0$, we show how to reduce the task to finding the minimum of a real convex function in finite dimension. Then, using this reduction, we show that the optimal value of (2.1)-(2.4) is obtained by a recursive formula. As a final step, we analyze this recurrence by considering a continuous counterpart defined by a differential equation.

Consider the function $\Gamma^{\varepsilon}_{n,m} : \mathbb{R}^{m-1}_{+} \to \mathbb{R}$ defined by

$$\Gamma^{\varepsilon}_{n,m}(x) = \varepsilon + x_1^n + \sum_{i=1}^{m-2} x_i \left(\varepsilon + \left(\frac{x_{i+1}}{x_i}\right)^n\right) - x_{m-1}(n-1-\varepsilon).$$

Given $\varepsilon > 0$ and positive integer $n \geq 2$, let $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$ be the sequence defined by the following recurrence:

$$\rho_{\varepsilon,1} = 1, \text{ and } (n-1)\rho^n_{\varepsilon,j-1} - \varepsilon = n\rho^{n-1}_{\varepsilon,j} \text{ for every } j \geq 2 \text{ such that } (n-1)\rho^n_{\varepsilon,j-1} - \varepsilon > 0. \tag{2.9}$$

For fixed ε and n we say that $\rho_{\varepsilon,j}$ is well defined if $(n-1)\rho^n_{\varepsilon,j-1} - \varepsilon > 0$. Observe that by letting $x = \rho_{\varepsilon,j}$ in Claim 2, we get that $\rho_{\varepsilon,j}$ is decreasing in j. It follows that if $\rho_{\varepsilon,m}$ is well defined, then so is $\rho_{\varepsilon,j}$ for $j \leq m$. As a first step, we will show that the optimal value of (2.1)-(2.4) can be obtained in terms of the sequence $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$. To prove this result we require a few propositions.

Proposition 2. *Let $\varepsilon > 0$, and let $n \geq 2$ and $m \geq 3$ be two positive integers such that $\rho_{\varepsilon,m}$ is well defined. Then, $\Gamma^{\varepsilon}_{n,m}$ is convex over $\mathbb{R}^{m-1}_{+}$ and it has a unique minimizer Y in this region, given by*

$$Y_1 = \rho_{\varepsilon,m} \text{ and } Y_j = \prod_{k=0}^{j-1} \rho_{\varepsilon,m-k} \text{ for every } j \in \{2, \ldots, m-1\}. \tag{2.10}$$

Furthermore, $\Gamma^{\varepsilon}_{n,m}(Y) = \varepsilon - (n-1)\rho^n_{\varepsilon,m}$.

Proof. We begin by proving (strict) convexity of $\Gamma^{\varepsilon}_{n,m}$. We proceed by induction on m. Observe first that when $m = 3$, we have that $\Gamma^{\varepsilon}_{n,3}(x_1, x_2) = \varepsilon + x_1^n + p(x_1, x_2) - x_2(n - 1 - \varepsilon)$, where $p(y,z) = y(\varepsilon + (z/y)^n)$. The Hessian of p is

$$\nabla^2 p(y,z) = n(n-1)z^{n-2}y^{1-n} \begin{pmatrix} z^2/y^2 & -z/y \\ -z/y & 1 \end{pmatrix},$$

and this is a positive semidefinite matrix for every $(y,z) \in \mathbb{R}^2_{+}$, since one eigenvalue is equal to zero, and the other is $n(n-1)z^{n-2}y^{1-n}((z/y)^2 + 1) > 0$. In particular, p is convex over $\mathbb{R}^2_{+}$. Since the function $\varepsilon + x_1^n - x_2(n-1-\varepsilon)$ is also convex over $\mathbb{R}^2_{+}$, we conclude that $\Gamma^{\varepsilon}_{n,3}$ is convex over $\mathbb{R}^2_{+}$. Now consider an integer value $m > 3$, and observe that

$$\Gamma^{\varepsilon}_{n,m+1}(x_1, \ldots, x_m) = p(x_{m-1}, x_m) - (x_m - x_{m-1})(n-1-\varepsilon) + \Gamma^{\varepsilon}_{n,m}(x_1, \ldots, x_{m-1}),$$

and therefore the convexity follows by the inductive step, that is, $\Gamma^{\varepsilon}_{n,m}$ convex over $\mathbb{R}^{m-1}_{+}$, together with p convex over $\mathbb{R}^{2}_{+}$. Every minimizer y of $\Gamma^{\varepsilon}_{n,m}$ over $\mathbb{R}^{m-1}_{+}$ is a solution to the system given by $\nabla\Gamma^{\varepsilon}_{n,m} = 0$, that is,

$$(n-1)\left(\frac{y_2}{y_1}\right)^n - \varepsilon = n y_1^{n-1}, \tag{2.11}$$

$$(n-1)\left(\frac{y_{i+1}}{y_i}\right)^n - \varepsilon = n\left(\frac{y_i}{y_{i-1}}\right)^{n-1} \quad \text{for every } i \in \{2,\ldots,m-2\}, \tag{2.12}$$

$$n-1-\varepsilon = n\left(\frac{y_{m-1}}{y_{m-2}}\right)^{n-1}, \text{ and } y \in \mathbb{R}^{m-1}_{+}. \tag{2.13}$$

The above system has a unique solution and therefore this proves the first part.

To finish the proof, we show that Y defined in (2.10) is strictly positive, satisfies the system (2.11)-(2.13), and $\Gamma^{\varepsilon}_{n,m}(Y) = \varepsilon - (n-1)\rho^{n}_{\varepsilon,m}$. Since $\rho_{\varepsilon,j}$ is well-defined for all $j \leq m$, we have $\rho_{\varepsilon,j} = ((n-1)\rho^{n}_{\varepsilon,j-1} - \varepsilon)^{1/(n-1)} > 0$. This implies that $Y \in \mathbb{R}^{m-1}_{+}$. Next observe that $Y_2 = \rho_{\varepsilon,m}\rho_{\varepsilon,m-1}$ and therefore $Y_2/Y_1 = \rho_{\varepsilon,m-1}$. Then, we have

$$(n-1)(Y_2/Y_1)^n - \varepsilon = (n-1)\rho^{n}_{\varepsilon,m-1} - \varepsilon = n\rho^{n-1}_{\varepsilon,m} = nY_1^{n-1},$$

and therefore (2.11) is satisfied. Similarly, for every $j \in \{2,\ldots,m-2\}$, we have $Y_j/Y_{j-1} = \rho_{m-j+1}$ and $Y_{j+1}/Y_j = \rho_{m-j}$. Then, we have

$$(n-1)(Y_{j+1}/Y_j)^n - \varepsilon = (n-1)\rho^{n}_{\varepsilon,m-j} - \varepsilon = n\rho^{n-1}_{\varepsilon,m-j+1} = n(Y_j/Y_{j-1})^{n-1},$$

and therefore (2.12) is satisfied. Finally, since $Y_{m-1}/Y_{m-2} = \rho_{\varepsilon,2}$, we have

$$n-1-\varepsilon = (n-1)\rho^{n}_{\varepsilon,1} - \varepsilon = n\rho^{n-1}_{\varepsilon,2} = n(Y_{m-1}/Y_{m-2})^{n-1},$$

and therefore (2.13) is satisfied. We now evaluate $\Gamma^{\varepsilon}_{n,m}(Y)$. The vector Y satisfies (2.11)-(2.13) and therefore

$$\begin{aligned}(n-1)\sum_{i=1}^{m-2} Y_i\left(\frac{Y_{i+1}}{Y_i}\right)^n + (n-1)Y_{m-1} - \varepsilon\sum_{i=1}^{m-1} Y_i &= nY_1^n + n\sum_{i=2}^{m-1} Y_i\left(\frac{Y_i}{Y_{i-1}}\right)^{n-1} \\ &= nY_1^n + n\sum_{i=2}^{m-1} Y_{i-1}\left(\frac{Y_i}{Y_{i-1}}\right)^{n} \\ &= nY_1^n + n\sum_{i=1}^{m-2} Y_i\left(\frac{Y_{i+1}}{Y_i}\right)^{n}.\end{aligned}$$

By subtracting the first term of the left hand side we get

$$\sum_{i=1}^{m-2} Y_i \left(\frac{Y_{i+1}}{Y_i}\right)^n = (n-1)Y_{m-1} - \varepsilon \sum_{i=1}^{m-1} Y_i - nY_1^n,$$

and by rearranging terms we obtain that

$$\sum_{i=1}^{m-2} Y_i \left(\varepsilon + \left(\frac{Y_{i+1}}{Y_i}\right)^n\right) = (n-1-\varepsilon)Y_{m-1} - nY_1^n.$$

Therefore, the minimum of $\Gamma_{n,m}^{\varepsilon}$ over $\mathbb{R}_+^{m-1}$ is equal to

$$\varepsilon + Y_1^n + (n-1-\varepsilon)Y_{m-1} - nY_1^n - (n-1-\varepsilon)Y_{m-1} = \varepsilon - (n-1)Y_1^n.$$

The proof follows since we have $Y_1 = \rho_{\varepsilon,m}$. □

Now that we have shown $\Gamma_{n,m}^{\varepsilon}(Y) = \varepsilon - (n-1)\rho_{\varepsilon,m}^n$, we still need to establish the connection between $\Gamma_{n,m}^{\varepsilon}(Y)$ and our objective (2.1) in order to complete the picture and show that the value of our optimization problem (2.1) − (2.4) is effectively given by $\varepsilon - (n-1)\rho_{\varepsilon,m}^n$. The following proposition will complete the feasibility of Y of the constraints in (2.2) − (2.4). Once this is established, in Lemma 4 we find that plugging Y into the objective (2.1) gives an expression that is upper bounded by $\Gamma_{n,m}^{\varepsilon}(Y)$ plus some remainder.

Proposition 3. *Let $\varepsilon > 0$, let $n \geq 2$ and $m \geq 3$ be two positive integers such that $\rho_{\varepsilon,m}$ is well defined, and let Y be as defined in* (2.10)*. Then, the following holds:*

(a) For every $j \in \{1, \ldots, m-1\}$ we have that $Y_{j+1} \leq Y_j$.

(b) For every $j \in \{2, \ldots, m-1\}$ we have that $Y_j^2 \leq Y_{j-1}Y_{j+1}$.

Proof. Observe that for every $k \in \{1, \ldots, m-1\}$, we have $Y_{m-k+1}/Y_{m-k} = \rho_{\varepsilon,k}$. For $k=1$ we have $Y_m/Y_{m-1} = \rho_\varepsilon = 1$. From the definition of the recurrence, we have

$$(n-1)\rho_{\varepsilon,k-1}^n \geq (n-1)\rho_{\varepsilon,k-1}^n - \varepsilon = n\rho_{\varepsilon,k}^{n-1}$$

for every $k \in \{2, \ldots, m-1\}$. By induction, if $\rho_{\varepsilon,k-1} \leq 1$, we have $\rho_{\varepsilon,k}^{n-1} \leq (n-1)/n$ and therefore $\rho_{\varepsilon,k} \leq 1$. This concludes part (a). Since for every $j \in \{1, \ldots, m-1\}$ we have $Y_{j+1}/Y_j = \rho_{\varepsilon,m-j}$, to prove part (b) it suffices to show $\rho_{\varepsilon,k+1} \leq \rho_{\varepsilon,k}$ for every $k \in \{1, \ldots, m-2\}$. From the construction of the recurrence, for every $k \in \{1, \ldots, m-2\}$ it holds that

$$\rho_{\varepsilon,k} \geq \left(\frac{n}{n-1}\right)^{\frac{1}{n}} \rho_{\varepsilon,k+1}^{(n-1)/n}.$$

By (a) we have $\rho_{\varepsilon,k+1} \in [0,1]$, which together with Claim 2 implies that

$$\left(\frac{n}{n-1}\right)^{\frac{1}{n}} \rho_{\varepsilon,k+1}^{(n-1)/n} \geq \rho_{\varepsilon,k+1}.$$

Therefore we conclude that $\rho_{\varepsilon,k+1} \leq \rho_{\varepsilon,k}$. This proves part (b). □

The other part of Lemma 4 is to give a lower bound on (2.1) − (2.4). That is, we need to bound the objective (2.1) for any candidate δ that satisfies the constraints. We find that the objective function can be written as $\Gamma^{\varepsilon}_{n,m}(\delta_1,\ldots,\delta_{m-1})$ plus some remainder in function of δ. To arrive at the desired conclusion, we need to bound this remainder for any feasible δ. The following proposition gives a convenient way to go about this task.

Proposition 4. *For every sequence $(\delta_j)_{j\in\mathbb{N}}$ satisfying* (2.2)-(2.4) *for which $\delta_m/\delta_{m-1} < 1$, there exists a sequence $(\beta_j)_{j\in\mathbb{N}}$ satisfying* (2.2)-(2.4)*, and such that the following holds:*

(a) For every $j \in \{0,1,\ldots,m-1\}$ we have $\delta_j = \beta_j$, and $\beta_m/\beta_{m-1} < 1$.

(b) $\displaystyle\sum_{i=m-1}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i \leq \beta_{m-1}\sum_{i=0}^{n-1}\left(\frac{\beta_m}{\beta_{m-1}}\right)^i.$

Proof. Suppose we are given $(\delta_j)_{j\in\mathbb{N}}$ satisfying (2.2)-(2.4) for which $\delta_m/\delta_{m-1} < 1$. We claim that then there exists a sequence $(\beta_j)_{j\in\mathbb{N}}$ satisfying (2.2)-(2.4) such that (a) holds and furthermore (i) $\beta_j \geq \delta_j$ for all $j \geq m$ and (ii) $\beta_j/\beta_{j-1} = \beta_m/\beta_{m-1}$ for all $j \geq m$.

If $(\delta_j)_{j\in\mathbb{N}}$ does not already satisfy these properties, then it must be because of (ii). In particular, there must be a smallest index $j \geq m$ such that $\delta_{j+1}/\delta_j > \delta_m/\delta_{m-1}$. We next describe a procedure that maintains all properties, but extends (ii) so that it holds for one more index. Applying this procedure iteratively, we obtain $(\beta_j)_{j\in\mathbb{N}}$.

Given $(\delta_j)_{j\in\mathbb{N}}$ satisfying (2.2)-(2.4), let $k(\delta) \geq m$ be the first value j such that $\delta_{j+1}/\delta_j > \delta_m/\delta_{m-1}$. In particular, we have $\delta_j/\delta_{j-1} = \delta_m/\delta_{m-1}$ for every $j \in \{m,\ldots,k(\delta)\}$. Consider the sequence $(D_j)_{j\in\mathbb{N}}$ defined as follows: $D_j = \delta_j$ for every $j \in \{0,1,\ldots,m-1\}$,

$$D_m = \delta_{m-1}\left(\frac{\delta_{k(\delta)+1}}{\delta_{m-1}}\right)^{\frac{1}{k(\delta)-m+2}},$$

$D_j = D_m(D_m/\delta_{m-1})^{j-m}$ for every $j \in \{m+1,\ldots,k(\delta)\}$, and $D_j = \delta_j$ for every $j \geq k(\delta)+1$. Observe that from the construction it holds directly that $D_{j+1}/D_j = D_m/\delta_{m-1}$ for every $j \in \{m,\ldots,k(\delta)-1\}$. Furthermore, we have

$$\frac{\delta_{k(\delta)+1}}{D_{k(\delta)}} = D_m\left(\frac{D_m}{\delta_{m-1}}\right)^{k(\delta)-m+1}\cdot\frac{1}{D_m}\left(\frac{\delta_{m-1}}{D_m}\right)^{k(\delta)-m} = \frac{D_m}{\delta_{m-1}},$$

and therefore, we have $D_{j+1}/D_j = D_m/D_{m-1}$ for every $j \in \{m-1, \ldots, k(\delta)\}$. By construction, the sequence $(D_j)_{j\in\mathbb{N}}$ satisfies (2.2)-(2.4) and $D_m/D_{m-1} < 1$. We show next that $D_j \geq \delta_j$ for every $j \in \{m, \ldots, k(\delta)\}$. Since $\delta_{j+1}/\delta_j \geq \delta_m/\delta_{m-1}$ for every $j \in \{m, \ldots, k(\delta)\}$, we have

$$\left(\frac{\delta_m}{\delta_{m-1}}\right)^{k(\delta)-m+1} \leq \prod_{j=m}^{k(\delta)} \frac{\delta_{j+1}}{\delta_j} = \frac{\delta_{k(\delta)+1}}{\delta_m},$$

which implies that $\delta_m \leq \delta_{m-1}(\delta_{k(\delta)+1}/\delta_{m-1})^{\frac{1}{k(\delta)-m+2}} = D_m$. For $j \in \{m+1, \ldots, k(\delta)\}$ we proceed by induction:

$$D_j = D_{j-1}\frac{D_m}{\delta_{m-1}} \geq \delta_{j-1}\frac{D_m}{\delta_{m-1}} \geq \delta_{j-1}\frac{\delta_m}{\delta_{m-1}} = \delta_j \cdot \frac{\delta_{j-1}}{\delta_j} \cdot \frac{\delta_m}{\delta_{m-1}} = \delta_j,$$

where the first equality holds by construction of the sequence, the first inequality holds by the inductive hypothesis, the second inequality holds since $D_m \geq \delta_m$, and the last equality follows since $\delta_j/\delta_{j-1} = \delta_m/\delta_{m-1}$ for every $j \in \{m, \ldots, k(\delta)\}$. This finishes the proof of part (a).

In the remainder we will prove part (b) using the existence of a sequence $(\beta_j)_{j\in\mathbb{N}}$ for which (a) holds as well as (i) and (ii). To this end we need the following definition and claim. For every sequence $(\delta_j)_{j\in\mathbb{N}}$ let

$$\mathcal{R}(\delta) = \sum_{i=m-1}^{\infty} \left(1 - \left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i.$$

Claim 3. *$\mathcal{R}$ is non-decreasing in δ_i for every $i \geq m$.*

Before proving Claim 3, we show how together with the properties of the sequence $(\beta_j)_{j\in\mathbb{N}}$ it implies property (b). Namely,

$$\begin{aligned}
\sum_{i=m-1}^{\infty} \left(1 - \left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i &\leq \sum_{i=m-1}^{\infty} \left(1 - \left(\frac{\beta_{i+1}}{\beta_i}\right)^n\right)\beta_i \\
&= \left(1 - \left(\frac{\beta_m}{\beta_{m-1}}\right)^n\right) \sum_{i=m-1}^{\infty} \beta_i \\
&= \beta_{m-1}\left(1 - \left(\frac{\beta_m}{\beta_{m-1}}\right)^n\right) \sum_{i=0}^{\infty} \left(\frac{\beta_m}{\beta_{m-1}}\right)^i = \beta_{m-1} \sum_{i=0}^{n-1} \left(\frac{\beta_m}{\beta_{m-1}}\right)^i,
\end{aligned}$$

where the inequality holds by Claim 3 and (i), the first equality holds by (ii), and the second equality holds because (ii) implies $\beta_i = \beta_{m-1}(\beta_m/\beta_{m-1})^{i-m+1}$ for $i \geq m-1$.

It remains to prove Claim 3. Consider the function $\varphi : \mathbb{R}_+^2 \to \mathbb{R}_+$ such that $\varphi(x,y) = (1-(y/x)^n)x$. In particular, the derivative of $\mathcal{R}$ with respect to δ_i, with $i \geq m$, is equal to

$$\begin{aligned}
&\frac{\partial \varphi}{\partial y}(\delta_{i-1},\delta_i) + \frac{\partial \varphi}{\partial x}(\delta_i,\delta_{i+1}) \\
&= -n\left(\frac{\delta_i}{\delta_{i-1}}\right)^{n-1} + 1 + (n-1)\left(\frac{\delta_{i+1}}{\delta_i}\right)^n \\
&= n\left(\frac{\delta_{i+1}}{\delta_i}\right)^{n-1}\left(\frac{1}{n}\left(\frac{\delta_i}{\delta_{i+1}}\right)^{n-1} + \left(1-\frac{1}{n}\right)\frac{\delta_{i+1}}{\delta_i} - \left(\frac{\delta_i^2}{\delta_{i+1}\delta_{i-1}}\right)^{n-1}\right) \\
&\geq n\left(\frac{\delta_{i+1}}{\delta_i}\right)^{n-1}\left(1 - \left(\frac{\delta_i^2}{\delta_{i+1}\delta_{i-1}}\right)^{n-1}\right) \geq 0,
\end{aligned}$$

The first inequality holds because for any $p \in [0,1]$ we have that $(1-1/n)p + 1/(np^{n-1}) \geq 1$, and $\delta_{i-1} \leq \delta_i \leq \delta_{i+1}$ for every $i \geq m$, and the second holds since $(\delta_j)_{j\in\mathbb{N}}$ satisfies constraint (2.3). This concludes the proof of the claim. □

The following lemma relates the optimal value of the optimization problem (2.1)-(2.4) with the sequence $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$. Using Lemma 4 and Theorem 4 we can numerically find the competition complexity by computing the recurrence (2.9) (see Figure 2.2). More specifically, given n and ε, we just have to find the last value m for which the value of the optimization problem is non-negative, and this can be found by numerically computing the values of the recurrence (2.9).

Lemma 4. *Let $\varepsilon > 0$, and let $n \geq 2$ and $m \geq 3$ be two positive integers such that $\rho_{\varepsilon,m}$ is well defined. Then, the value of the optimization problem* (2.1)-(2.4) *is equal to $\varepsilon - (n-1)\rho_{\varepsilon,m}^n$.*

Proof. Consider $Y \in \mathbb{R}_+^{m-1}$ as defined in (2.10). For every $\alpha \in (0,1)$, consider the sequence $(\mathcal{Y}_j(\alpha))_{j\in\mathbb{N}}$ defined as follows: $\mathcal{Y}_0(\alpha) = 1$, $\mathcal{Y}_j(\alpha) = Y_j$ for every $j \in \{1,\ldots,m-1\}$ and $\mathcal{Y}_j(\alpha) = \alpha^{j-m+1}Y_{m-1}$ for every $j \geq m$. Thanks to Proposition 2 and Proposition 3, for every $\alpha \in (0,1)$ the sequence $(\mathcal{Y}_j(\alpha))_{j\in\mathbb{N}}$ satisfies (2.2)-(2.4). The objective value

(2.1) of the sequence is equal to

$$
\begin{aligned}
&(1+\varepsilon)\sum_{i=0}^{m-1}\mathcal{Y}_i(\alpha)-\sum_{i=0}^{\infty}\left(1-\left(\frac{\mathcal{Y}_{i+1}(\alpha)}{\mathcal{Y}_i(\alpha)}\right)^n\right)\mathcal{Y}_i(\alpha)\\
&=\varepsilon+Y_1^n+(1+\varepsilon)Y_{m-1}+\sum_{i=1}^{m-2}\left(\varepsilon+\left(\frac{Y_{i+1}}{Y_i}\right)^n\right)Y_i-\sum_{i=m-1}^{\infty}\left(1-\left(\frac{\mathcal{Y}_{i+1}(\alpha)}{\mathcal{Y}_i(\alpha)}\right)^n\right)\mathcal{Y}_i(\alpha)\\
&=\varepsilon+Y_1^n+(1+\varepsilon)Y_{m-1}+\sum_{i=1}^{m-2}\left(\varepsilon+\left(\frac{Y_{i+1}}{Y_i}\right)^n\right)Y_i-(1-\alpha^n)\sum_{i=0}^{\infty}\mathcal{Y}_{m+i-1}(\alpha)\\
&=\varepsilon+Y_1^n+\sum_{i=1}^{m-2}\left(\varepsilon+\left(\frac{Y_{i+1}}{Y_i}\right)^n\right)Y_i-Y_{m-1}\left(\left((1-\alpha^n)\sum_{i=0}^{\infty}\alpha^i\right)-1-\varepsilon\right)\\
&=\Gamma^{\varepsilon}_{n,m}(Y)+Y_{m-1}\left(n-(1-\alpha^n)\sum_{i=0}^{\infty}\alpha^i\right)\\
&=\varepsilon-(n-1)\rho^n_{\varepsilon,m}+Y_{m-1}\left(n-(1-\alpha^n)\sum_{i=0}^{\infty}\alpha^i\right),
\end{aligned}
$$

where the last equality holds by Proposition 2. In particular, the feasibility of $(\mathcal{Y}_j(\alpha))_{j\in\mathbb{N}}$ for every $\alpha\in(0,1)$ implies that the value of the optimization problem (2.1)-(2.4) is upper bounded by

$$
\varepsilon-(n-1)\rho^n_{\varepsilon,m}+Y_{m-1}\inf_{\alpha\in(0,1)}\left\{n-(1-\alpha^n)\sum_{i=0}^{\infty}\alpha^i\right\}=\varepsilon-(n-1)\rho^n_{\varepsilon,m}. \tag{2.14}
$$

Let $(\delta_j)_{j\in\mathbb{N}}$ be any sequence satisfying (2.2)-(2.4). We denote by $\mathcal{V}(\delta)$ the objective value (2.1), which by rearranging terms, is equal to

$$
\mathcal{V}(\delta)=\Gamma^{\varepsilon}_{n,m}(\delta_1,\ldots,\delta_{m-1})+n\delta_{m-1}-\sum_{i=m-1}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i.
$$

Now either $\delta_{i+1}/\delta_i=1$ for all $i\geq m-1$ in which case $\mathcal{V}(\delta)=\Gamma^{\varepsilon}_{n,m}(\delta_1,\ldots,\delta_{m-1})+n\delta_{m-1}\geq\min_{x\in\mathbb{R}^{n-1}_+}\Gamma^{\varepsilon}_{n,m}(x)=\varepsilon-(n-1)\rho^n_{\varepsilon,m}$, where the last inequality holds by Proposition 2. Otherwise, by Proposition 4, there exists a sequence $(\beta_j)_{j\in\mathbb{N}}$ satisfying

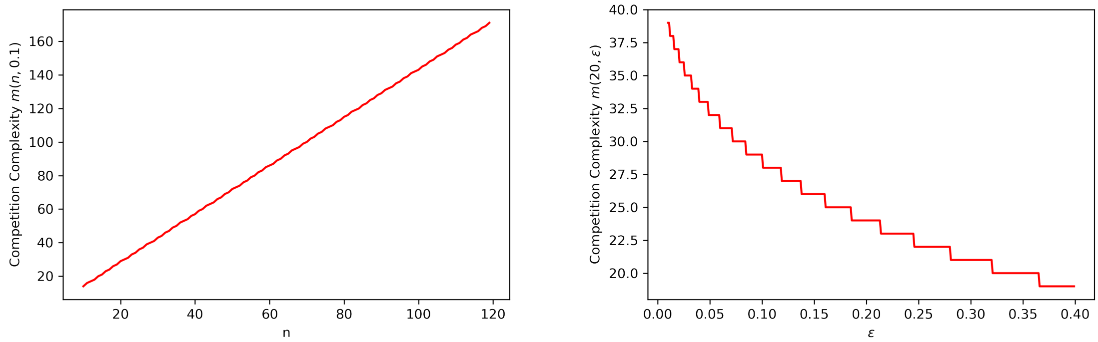

Fig. 2.2 On the left, we have a plot of the competition complexity as a function of n when $\varepsilon = 0.1$. On the right, we have a plot of the competition complexity as a function of ε when $n = 20$.

(2.2)-(2.4) for which the following holds:

$$
\begin{aligned}
\mathcal{V}(\delta) &= \Gamma^{\varepsilon}_{n,m}(\delta_1,\ldots,\delta_{m-1}) + n\delta_{m-1} - \sum_{i=m-1}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i \\
&= \Gamma^{\varepsilon}_{n,m}(\beta_1,\ldots,\beta_{m-1}) + n\beta_{m-1} - \sum_{i=m-1}^{\infty}\left(1-\left(\frac{\delta_{i+1}}{\delta_i}\right)^n\right)\delta_i \\
&\geq \Gamma^{\varepsilon}_{n,m}(\beta_1,\ldots,\beta_{m-1}) + n\beta_{m-1} - \beta_{m-1}\sum_{i=0}^{n-1}\left(\frac{\beta_m}{\beta_{m-1}}\right)^i \\
&\geq \min_{x\in\mathbb{R}^{m-1}_+}\Gamma^{\varepsilon}_{n,m}(x) + \beta_{m-1}\left(n - \sum_{i=0}^{n-1}\left(\frac{\beta_m}{\beta_{m-1}}\right)^i\right) \\
&\geq \varepsilon - (n-1)\rho^n_{\varepsilon,m} + \beta_{m-1}\left(n - \sum_{i=0}^{n-1}\left(\frac{\beta_m}{\beta_{m-1}}\right)^i\right),
\end{aligned}
$$

where the second equality holds by property (a) in Proposition 4, the first inequality holds by property (b) in Proposition 4, and the last inequality again holds by Proposition 2. Observe that for every $(\beta_j)_{j\in\mathbb{N}}$, the last term of the above inequality can be lower bounded by zero, and therefore, we get that $\mathcal{V}(\delta) \geq \varepsilon - (n-1)\rho^n_{\varepsilon,m}$ also in this case. This, together with the upper bound in (2.14), concludes the proof of the lemma. □

As a second step, we study the recurrence $(\rho_{\varepsilon,j})_{j\in\mathbb{N}}$ to find the point in which it becomes non-positive. More specifically, by Theorem 4 and Lemma 4, our aim is to find the greatest index m for which $\rho_{\varepsilon,m}$ is well defined, or equivalently the unique m for which $(n-1)\rho^n_{\varepsilon,m} - \varepsilon \leq 0$. To understand this problem we consider a differential equation that will serve as an upper bound to our recurrence relation. Recall the definition of $\phi(\varepsilon) = \int_0^1 1/(y(1-\log(y))+\varepsilon)dy$. Given a value $\varepsilon > 0$, consider the following ordinary differential equation:

$$y'(t) = y(t)(\log(y(t)) - 1) - \varepsilon \text{ for every } t \in (0, \phi(\varepsilon)), \tag{2.15}$$

$$y(0) = 1. \tag{2.16}$$

We define $y(\phi(\varepsilon)) = \lim_{t \uparrow \phi(\varepsilon)} y(t)$ as the continuous extension of y in $\phi(\varepsilon)$. The following lemma summarizes our results for the differential equation and $\phi(\varepsilon)$.

Lemma 5. *For every $\varepsilon > 0$, the differential equation* (2.15)-(2.16) *has a unique solution y_ε. Furthermore, the following holds:*

(a) For every $t \in [0, \phi(\varepsilon))$ we have $y'_\varepsilon(t) < 0$. In particular, y_ε is decreasing and invertible on $[0, \phi(\varepsilon))$ and $y_\varepsilon(\phi(\varepsilon)) = 0$.

(b) For every integer $n \geq 2$, and every $j \in \mathbb{N}$ for which $\rho_{\varepsilon,j}$ is well-defined, we have

$$\frac{n-1}{n} \rho^n_{\varepsilon,j} - \frac{\varepsilon}{n} \leq y_\varepsilon\left(\frac{j}{n}\right).$$

(c) For every $\delta \in (0, \phi(\varepsilon))$, there exists n_0 such that for every $n \geq n_0$ we have $(n-1)\rho^n_{\varepsilon,k} - \varepsilon > 0$, where $k = \lfloor (\phi(\varepsilon) - \delta) n \rfloor$.

(d) We have $\phi(\varepsilon) = \Theta(\log \log 1/\varepsilon)$ for $\varepsilon < 1$.

Before proving Lemma 5, we show how to use it to establish Theorem 3.

Proof. Proof of Theorem 3. Fix $\varepsilon > 0$ and consider the non-trivial case where $n \geq 2$. We begin with the first part of the theorem. By Lemma 4 it suffices to find the largest index j for which $\rho_{\varepsilon,j}$ is well defined. Suppose for a contradiction that for some $m \geq \phi(\varepsilon)n$, $\rho_{\varepsilon,m}$ is well defined but $(n-1)\rho^n_{\varepsilon,m} - \varepsilon > 0$. Define $\varepsilon' > 0$ such that $m/n = \phi(\varepsilon')$. Note that such an ε' exists and $\varepsilon' \leq \varepsilon$ because ϕ is monotone and continuous.

Claim 4. *For every positive integer j, $\rho_{\varepsilon',j}$ is well defined when $\rho_{\varepsilon,j}$ is well defined, and $\rho_{\varepsilon',j} \geq \rho_{\varepsilon,j}$.*

Using Claim 4, we have

$$\frac{n-1}{n} \rho^n_{\varepsilon,m} - \frac{\varepsilon}{n} \leq \frac{n-1}{n} \rho^n_{\varepsilon',m} - \frac{\varepsilon'}{n} \leq y_{\varepsilon'}\left(\phi(\varepsilon')\right) = 0,$$

where the second inequality holds by Lemma 5(b) and the final equality holds by Lemma 5(a). This yields a contradiction. To prove the claim, we consider an inductive argument. The claim clearly holds for $j = 1$, and assume that it holds for every

$k \le j-1$. If $\rho_{\varepsilon,j}$ is well defined, that is $(n-1)\rho^n_{\varepsilon,j-1} - \varepsilon > 0$, by the inductive step we have $\rho_{\varepsilon',j-1} \ge \rho_{\varepsilon,j-1}$ and therefore,

$$(n-1)\rho^n_{\varepsilon',j-1} - \varepsilon' \ge (n-1)\rho^n_{\varepsilon,j-1} - \varepsilon > 0,$$

meaning $\rho_{\varepsilon',j}$ is also well defined. Furthermore, in this case we have

$$n\rho^{n-1}_{\varepsilon,j} = (n-1)\rho^n_{\varepsilon,j-1} - \varepsilon \le (n-1)\rho^n_{\varepsilon',j-1} - \varepsilon' = n\rho^{n-1}_{\varepsilon',j},$$

which implies that $\rho_{\varepsilon',j} \ge \rho_{\varepsilon,j}$.

By Lemma 5(d) we have $\phi(\varepsilon) = \Theta(\log\log 1/\varepsilon)$. The second part of the theorem holds by Lemma 4 and Lemma 5(c). This finishes the proof of the theorem. □

We prove Lemma 5. For (a), (b) and (c) we first need a few propositions.

The following proposition simply proves 5(a).

Proposition 5. *For every $\varepsilon > 0$, there exists a unique solution of the differential equation* (2.15)-(2.16)*, that we denote y_ε. Furthermore, for every $t \in [0, \phi(\varepsilon))$ we have $y'_\varepsilon(t) < 0$. In particular, y_ε is decreasing and invertible in $[0, \phi(\varepsilon)]$, and $y_\varepsilon(\phi(\varepsilon)) = 0$.*

Proof. Since $y(t)(\log(y(t)) - 1) - \varepsilon$ as a function of y and t is continuous when $y > 0$, in this case a solution exists by Peano's existence theorem [66]. Observe that for any solution y of the differential equation (2.15)-(2.16), we have $y'(0) = -\varepsilon < 0$. Furthermore, for every $y \in (0,1]$, since $\log(y) \le 0$ and $\varepsilon > 0$, $y' < 0$. Now we show uniqueness. Assume the solution is not unique, that is there exist $y_1(t) \ne y_2(t)$ satisfying the ODE conditions. Denote by $[0,T)$ the largest interval on which y_1 and y_2 exist. Then let $t^* = \inf_{t\in[0,T)}\{t : y_1(t) \ne y_2(t)\}$. By continuity, we know that $y_1(t^*) = y_2(t^*)$. Now by Picard-Lindelöf theorem [70] that we give for convenience as theorem 1, there exists some interval $[t^* - \delta, t^* + \delta]$ where the solution to the ODE is unique. In particular, $y_1(t) = y_2(t)$ on $[t^* - \delta, t^* + \delta]$, contradicting the choice of t^*.

Denote this unique solution as y_ε and let $y_\varepsilon(T)$ be the continuous extension of y_ε. Now let us find the value of T. In particular, the function y_ε is strictly decreasing and thus invertible and we know it has a differentiable inverse in $[0,1]$. That is, we can write t as a function of y_ε. Also, denote $T = y_\varepsilon^{-1}(0)$. By standard integration rule, we may write

$$t(1) = t(0) + \int_0^1 \frac{dt}{dy_\varepsilon} dy_\varepsilon = 1 + \int_0^1 \frac{1}{\frac{dy_\varepsilon}{dt}} dy_\varepsilon = y_\varepsilon^{-1}(0) - \int_0^1 \frac{1}{y_\varepsilon(1-\log(y_\varepsilon)) + \varepsilon} dy_\varepsilon$$

In other words, we get

$$y_\varepsilon^{-1}(1) = T - \int_0^1 \frac{1}{s(1-\log(s))+\varepsilon} \mathrm{d}s = T - \phi(\varepsilon)$$

and since $y_\varepsilon^{-1}(1) = 0$ we conclude that $y_\varepsilon^{-1}(0) = T = \phi(\varepsilon)$. □

The following two propositions are technical and their sole purpose is to help with the proof of Lemma 5(b). In this lemma, we analyse $y_\varepsilon(\frac{j}{n})$ inductively by expressing it as a Taylor expansion centered at previous point $(\frac{j-1}{n})$. Although this is conceptually straightforward, in order to obtain the desired bounds, we need to go all the way up to the third derivative in the Taylor expansion. There does not seem to be a single clean proof as to why bounds are achieved for all regimes of j and ε. However, we are able to make a complete argument by going through an involved case analysis and achieving the desired bounds in each case, by different methods.

Given $\varepsilon > 0$, consider the function $M_\varepsilon : \mathbb{R} \to \mathbb{R}$ given by

$$M_\varepsilon(x) = \left(\log(x) - 1 - \frac{\varepsilon}{x}\right)\left(x\log^2(x) + x\log(x) - x - \varepsilon\right).$$

Proposition 6. *Let $\varepsilon > 0$ and let $\alpha_\varepsilon = y_\varepsilon^{-1}(\exp(-\frac{1}{2}(1+\sqrt{5})))$. Then, the following holds:*

(a) For every $t \in [0, \phi(\varepsilon)]$ we have $y_\varepsilon'''(t) = M_\varepsilon(y_\varepsilon(t))$.

(b) For every $t \in [0, \alpha_\varepsilon]$ we have $y_\varepsilon'''(t) \geq 0$.

(c) When $\varepsilon \leq 0.25$, we have $y_\varepsilon'''(t) \geq -1.173$ for every $t \in [\alpha_\varepsilon, \phi(\varepsilon)]$.

(d) When $\varepsilon \leq 0.25$, there exists $x_\varepsilon \in (0.01, 0.067)$ such that y_ε''' is increasing in $[y_\varepsilon^{-1}(x_\varepsilon), \phi(\varepsilon)]$.

(e) When $\varepsilon \geq 0.25$, we have $y_\varepsilon'''(t) \geq 0$ for every $t \in [0, \phi(\varepsilon)]$.

Proof. By a direct computation, we have that

$$y_\varepsilon''(t) = y_\varepsilon'(t)(\log(y_\varepsilon(t)) - 1) + y_\varepsilon(t) \cdot y_\varepsilon'(t)/y_\varepsilon(t) = y_\varepsilon'(t)\log(y_\varepsilon(t)),$$

and therefore,

$$\begin{aligned}
y_\varepsilon'''(t) &= y_\varepsilon''(t)\log(y_\varepsilon(t)) + y_\varepsilon'(t)\cdot\frac{y_\varepsilon'(t)}{y_\varepsilon(t)}\\
&= y_\varepsilon'(t)\log^2(y_\varepsilon(t)) + y_\varepsilon'(t)\cdot\frac{y_\varepsilon'(t)}{y_\varepsilon(t)}\\
&= y_\varepsilon'(t)\left(\log^2(y_\varepsilon(t)) + \log(y_\varepsilon(t)) - 1 - \frac{\varepsilon}{y_\varepsilon(t)}\right)\\
&= \Big(y_\varepsilon(t)(\log(y_\varepsilon(t))-1)-\varepsilon\Big)\left(\log^2(y_\varepsilon(t)) + \log(y_\varepsilon(t)) - 1 - \frac{\varepsilon}{y_\varepsilon(t)}\right)\\
&= M_\varepsilon(y_\varepsilon(t)),
\end{aligned}$$

which proves (a). Consider the function $g(x) = x\log^2(x) + x\log(x) - x$. We have that $g(x) \le 0$ for every $\exp(-\frac{1}{2}(1+\sqrt{5})) \le x \le 1$, and together with Proposition 5, implies that $g(y_\varepsilon(t)) - \varepsilon \le 0$ for every $t \in [0, \alpha_\varepsilon]$. Furthermore, by Proposition 5 we have that $y_\varepsilon'(t)/y_\varepsilon(t) \le 0$ for every $t \in [0, \alpha_\varepsilon]$ and therefore

$$y_\varepsilon'''(t) = M_\varepsilon(y_\varepsilon(t)) = \frac{y_\varepsilon'(t)}{y_\varepsilon(t)}(g(y_\varepsilon(t)) - \varepsilon) \ge 0,$$

which proves (b). To prove (c), observe that by Proposition 5 we have that $y_\varepsilon(\alpha_\varepsilon) \ge y_\varepsilon(t) \ge 0$ for every $t \in [\alpha_\varepsilon, \phi(\varepsilon)]$, and since $0.199 > y_\varepsilon(\alpha_\varepsilon) = \exp(-(1+\sqrt{5})/2) > 0.198$, we have that

$$\min_{\varepsilon\in(0,0.25)}\min_{t\in[\alpha_\varepsilon,\phi(\varepsilon)]} y_\varepsilon'''(t) = \min_{\varepsilon\in(0,0.25)}\min_{t\in[\alpha_\varepsilon,\phi(\varepsilon)]} M_\varepsilon(y_\varepsilon(t)) \ge \min_{\substack{\varepsilon\in[0,0.25],\\ x\in[0,0.199]}} M_\varepsilon(x) \approx -1.1722,$$

where the first equality comes from part (a). We now prove (d). By a direct computation, we have

$$\begin{aligned}
M_\varepsilon'(x) &= -\frac{\varepsilon^2}{x^2} - \frac{2\varepsilon}{x} - \frac{2\varepsilon\log(x)}{x} + \log^3(x) + 3\log^2(x) - 2\log(x) - 1,\\
M_\varepsilon''(x) &= \frac{2\varepsilon^2}{x^3} + \frac{2\varepsilon\log(x)}{x^2} - \frac{2}{x} + \frac{3\log^2(x)}{x} + \frac{6\log(x)}{x}.
\end{aligned}$$

Furthermore, we have

$$\min_{\substack{\varepsilon\in[0,0.25],\\ x\in[0,067]}} M_\varepsilon''(x) \approx 0.716,$$

and therefore the function M'_ε is increasing in $(0,0.067]$ for every $\varepsilon \in (0,0.25]$. On the other hand, we have

$$M'_\varepsilon(0.067) > -\frac{\varepsilon^2}{(0.067)^2} - \frac{2\varepsilon}{0.067} - \frac{2\varepsilon\log(0.067)}{x} + 6.57,$$

and this is a quadratic concave function over $[0,0.25]$ that attains the minimum at $\varepsilon = 0.25$ with a value of ≈ 5.36. Furthermore, we have

$$M'_\varepsilon(0.01) < -\frac{\varepsilon^2}{(0.01)^2} - \frac{2\varepsilon}{0.01} - \frac{2\varepsilon\log(0.01)}{x} - 25.83,$$

and this is a quadratic concave function over $[0,0.25]$ that attains the maximum at ≈ 0.036 with a value of ≈ -12.83. Therefore, for every $\varepsilon \in (0,0.25]$, the continuity of M'_ε implies the existence of a value $x_\varepsilon \in (0.01, 0.0067)$ such that $M'_\varepsilon(x_\varepsilon) = 0$. Since the function M'_ε is increasing in $[0,0.067]$, we have $M'_\varepsilon(x) \leq M_\varepsilon(x_\varepsilon) = 0$ for every $x \in [0,x_\varepsilon]$, and therefore the function M_ε is decreasing in the interval $[0,x_\varepsilon]$. By Proposition 5 we have that y_ε is decreasing in $[0,\phi(\varepsilon)]$, and therefore we conclude that $y'''_\varepsilon = M_\varepsilon \circ y_\varepsilon$ is increasing in the interval $[y_\varepsilon^{-1}(x_\varepsilon), \phi(\varepsilon)]$.

Finally, we prove (e). Recall that $g(x) = x\log^2(x) + x\log(x) - x$. It is sufficient to verify that $g(x) \leq \varepsilon$ for every $x \in (0,1]$ when $\varepsilon \geq 0.25$, since we have $y'''_\varepsilon(t) = y'_\varepsilon(t)(g(y_\varepsilon(t)) - \varepsilon)/y_\varepsilon(t)$, and $y'_\varepsilon \leq 0$ in $[0,\phi(\varepsilon)]$. We have

$$g'(x) = \log^2(x) + 2x\log(x) \cdot \frac{1}{x} + \log(x) + x \cdot \frac{1}{x} - 1 = \log(x)(\log(x) + 3).$$

We have $g'(x) \geq 0$ when $x \in (0, e^{-3}]$ and $g'(x) \leq 0$ when $x \in [e^{-3}, 1]$. Therefore, the maximum of g in $(0,1]$ is attained at e^{-3} and we conclude that $g(x) \leq 5e^{-3} - \varepsilon \leq 5e^{-3} - 0.25 < 0$ for every $x \in (0,1]$. This concludes the proof of the proposition. $\square$

Given $\varepsilon > 0$ and a positive integer $n \geq 2$, consider the function $F_{n,\varepsilon} : \mathbb{R} \to \mathbb{R}$ given by

$$F_{n,\varepsilon}(x) = x + \frac{x(\log(x) - 1)}{n} + \frac{\log(x)(x(\log(x) - 1) - \varepsilon)}{2n^2}.$$

Proposition 7. *Let $n \geq 2$ be an integer value and let $\varepsilon \in (0,0.25]$. Then, the following holds:*

(a) For every $x \in (0,1]$ we have $F_{n,\varepsilon}(x) \geq \left(\frac{n-1}{n}\right) x^{\frac{n}{n-1}}$.

(b) For every $x \in [0.01, 0.199]$ we have $F_{n,\varepsilon}(x) \geq \left(\frac{n-1}{n}\right) x^{\frac{n}{n-1}} + \frac{1.173}{6n^6}$.

(c) For every $x \in [0, 0.07]$ we have $F_{n,\varepsilon}(x) + \frac{M_\varepsilon(x)}{6n^6} \geq \left(\frac{n-1}{n}\right) x^{\frac{n}{n-1}}$.

Proof. The inequality in (a) holds by [23, Proposition D.1.]. Consider the function $G_n : \mathbb{R} \to \mathbb{R}$ given by

$$G_n(x) = 1 + \frac{\log(x) - 1}{n} + \frac{\log(x)(\log(x) - 1)}{2n^2} - \left(\frac{n-1}{n}\right) x^{\frac{1}{n-1}} - \frac{1.173}{6xn^6}.$$

To prove (b) it suffices to show that $G_n(x) \geq 0$ for every $x \in [0.01, 0.199]$, since $-\varepsilon \log(x) \geq 0$ for every $x \in [0.01, 0.199]$, and therefore

$$F_{n,\varepsilon}(x) - \left(\frac{n-1}{n}\right) x^{\frac{n}{n-1}} - \frac{1.173}{6n^6} \geq x \cdot G_n(x) \geq 0.$$

We have that $\{G_n(0.199)\}_{n\in\mathbb{N}}$ is a strictly positive and decreasing sequence, and therefore it is sufficient to show that G_n is decreasing in the interval $[0.01, 0.199]$. We have

$$\begin{aligned} G_n'(x) &= \frac{1}{nx} + \frac{\log(x)}{n^2 x} - \frac{1}{2n^2 x} - \frac{1}{n} x^{\frac{1}{n-1} - 1} + \frac{1.173}{6n^6 x^2} \\ &= \frac{1}{nx^2}\left(x + \frac{x\log(x) - x/2}{n} - x^{\frac{n}{n-1}} + \frac{1.173}{6n^5}\right), \end{aligned}$$

and let

$$h_n(x) = x + \frac{x\log(x) - x/2}{n} - x^{\frac{n}{n-1}} + \frac{1.173}{6n^5}.$$

It is sufficient to show that h_n is non-positive in $[0.01, 0.199]$. We have

$$\begin{aligned} h_n'(x) &= 1 + \frac{\log(x) + 1/2}{n} - \left(1 - \frac{1}{n}\right) x^{\frac{1}{n-1}}, \\ h_n''(x) &= \frac{1}{nx} - \frac{1}{n} x^{\frac{1}{n-1} - 1} = \frac{1}{nx}\left(1 - x^{\frac{1}{n-1}}\right), \end{aligned}$$

and therefore $h_n''(x) > 0$ for every $x \in [0.01, 0.199]$. This implies that h_n is convex in the interval $[0.01, 0.199]$, and therefore it is sufficient to verify that $h_n(0.01) < 0$ and $h_n(0.199) < 0$. In fact, we have

$$\begin{aligned} h_n(0.01) &= 0.01 + \frac{0.01\log(0.01) - 0.005}{n} - 0.01^{\frac{n}{n-1}} + \frac{1.173}{6n^5} \\ &\leq \frac{0.01\log(0.01) - 0.005}{2} + \frac{1.173}{6 \cdot 2^5} < -0.019, \\ h_n(0.199) &= 0.199 + \frac{0.199\log(0.199) - 0.0995}{n} - 0.199^{\frac{n}{n-1}} + \frac{1.173}{6n^5} \\ &\leq \frac{0.199\log(0.199) - 0.0995}{2} + \frac{1.173}{6 \cdot 2^5} < -0.2, \end{aligned}$$

and therefore we conclude that h_n is non-positive in $[0.01, 0.199]$, which implies that G_n is positive in $[0.01, 0.199]$. This proves (b). Finally, to prove (c), consider the function $\Psi : \mathbb{R}_+^3 \to \mathbb{R}$ given by

$$\Psi(x,y,\varepsilon) = x + \frac{x(\log(x)-1)}{y} + \frac{\log(x)(x(\log(x)-1)-\varepsilon)}{2y^2} + \frac{M_\varepsilon(x)}{6y^6} - \left(1-\frac{1}{y}\right)x^{\frac{y}{y-1}}.$$

Then, we have

$$\inf_{\substack{n\geq 2,\\ \varepsilon\in(0,0.25],\\ x\in[0,0.07]}} \left(F_{n,\varepsilon}(x) + \frac{M_\varepsilon(x)}{6n^6} - \left(\frac{n-1}{n}\right)x^{\frac{n}{n-1}}\right) \geq \min_{\substack{y\geq 2,\\ \varepsilon\in(0,0.25],\\ x\in[0,0.07]}} \Psi(x,y,\varepsilon) \geq 0,$$

which concludes the proof. □

Proof. Proof of Lemma 5. Part (a) holds by Proposition 5. To prove (b) we proceed by induction. When $j=1$, we have $\rho_{\varepsilon,1}^{n-1} = 1 = y_\varepsilon(0)$. For every $j \geq 1$, we take the Taylor expansion around point $\frac{j-1}{n}$ to express the function value at $\frac{j}{n}$. We obtain:

$$\begin{aligned}
& y_\varepsilon\Big(\frac{j}{n}\Big) \\
&= y_\varepsilon\Big(\frac{j-1}{n}\Big) + \frac{1}{n}y_\varepsilon'\Big(\frac{j-1}{n}\Big) + \frac{1}{2n^2}y_\varepsilon''\Big(\frac{j-1}{n}\Big) + \frac{1}{6n^6}y_\varepsilon'''(\xi) \\
&= y_\varepsilon\Big(\frac{j-1}{n}\Big) + \frac{1}{n}y_\varepsilon'\Big(\frac{j-1}{n}\Big)\Big(1+\frac{1}{2n}\log\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big)\Big) + \frac{1}{6n^6}y_\varepsilon'''(\xi) \\
&= y_\varepsilon\Big(\frac{j-1}{n}\Big) + \frac{y_\varepsilon(\frac{j-1}{n})(\log(y_\varepsilon(\frac{j-1}{n}))-1)-\varepsilon}{n}\Big(1+\frac{1}{2n}\log\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big)\Big) + \frac{1}{6n^6}y_\varepsilon'''(\xi) \\
&= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6}y_\varepsilon'''(\xi)
\end{aligned}$$

where $\xi \in ((j-1)/n, j/n)$, and the second and third equalities come from the ODE definition. The induction is simply on j, starting at $j=1$. However, we need to control the error term of the Taylor expansion, $\frac{1}{6n^6}y_\varepsilon'''(\xi)$. Depending on the value of ε and j (since $\xi \in (\frac{j-1}{n}, \frac{j}{n})$), we have different guarantees that will help us with the induction step, established in Proposition 6. We consider four different cases.

Case 1: Suppose that $\varepsilon \geq 0.25$. By Proposition 6(e) we have $y_\varepsilon'''(\xi) \geq 0$, and therefore

$$\begin{aligned} y_\varepsilon\Big(\frac{j}{n}\Big) &= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6} y_\varepsilon'''(\xi) \\ &\geq F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} \\ &\geq \left(\frac{n-1}{n}\right) y_\varepsilon\Big(\frac{j-1}{n}\Big)^{\frac{n}{n-1}} - \frac{\varepsilon}{n} \geq \left(\frac{n-1}{n}\right) \rho_{\varepsilon,j}^{n} - \frac{\varepsilon}{n} = \rho_{\varepsilon,j+1}^{n-1}, \end{aligned}$$

where the second inequality holds from Proposition 7(a), and in the third inequality we used the inductive step.

Case 2: Suppose that $\varepsilon \leq 0.25$ and $2 \leq j \leq \alpha_\varepsilon n$. In particular, we have $j/n \in [0, \alpha_\varepsilon]$. By Proposition 6(b) we have $y_\varepsilon'''(\xi) \geq 0$, and therefore

$$\begin{aligned} y_\varepsilon\Big(\frac{j}{n}\Big) &= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6} y_\varepsilon'''(\xi) \\ &\geq F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} \\ &\geq \left(\frac{n-1}{n}\right) y_\varepsilon\Big(\frac{j-1}{n}\Big)^{\frac{n}{n-1}} - \frac{\varepsilon}{n} \geq \left(\frac{n-1}{n}\right) \rho_{\varepsilon,j}^{n} - \frac{\varepsilon}{n} = \rho_{\varepsilon,j+1}^{n-1}, \end{aligned}$$

where the second inequality holds from Proposition 7(a), and in the third inequality we used the inductive step.

Case 3: Suppose that $\varepsilon \leq 0.25$ and $\alpha_\varepsilon n + 1 \leq j \leq y_\varepsilon^{-1}(x_\varepsilon)n + 1$, where x_ε is the value guaranteed by Proposition 6(d). In particular, we have $(j-1)/n \in [\alpha_\varepsilon, y_\varepsilon^{-1}(x_\varepsilon)]$, and by Proposition 6(d) we have $0.01 < x_\varepsilon$, which implies that $0.01 < y_\varepsilon((j-1)/n) \leq 0.199$. By Proposition 6(c) we have $y_\varepsilon'''(\xi) \geq -1.173$, and therefore

$$\begin{aligned} y_\varepsilon\Big(\frac{j}{n}\Big) &= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6} y_\varepsilon'''(\xi) \\ &\geq F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} - \frac{1.173}{6n^6} \\ &\geq \left(\frac{n-1}{n}\right) y_\varepsilon\Big(\frac{j-1}{n}\Big)^{\frac{n}{n-1}} - \frac{\varepsilon}{n} \geq \left(\frac{n-1}{n}\right) \rho_{\varepsilon,j}^{n} - \frac{\varepsilon}{n} = \rho_{\varepsilon,j+1}^{n-1}, \end{aligned}$$

where the second inequality holds from Proposition 7(b), and in the third inequality we used the inductive step.

Case 4: Suppose that $\varepsilon \leq 0.25$ and $j \geq y_\varepsilon^{-1}(x_\varepsilon)n + 1$. In particular, we have $(j-1)/n \geq y_\varepsilon^{-1}(x_\varepsilon)$ and $y_\varepsilon((j-1)/n) \leq x_\varepsilon < 0.067$. By Proposition 6(d), y_ε''' is increasing in

$[y_\varepsilon^{-1}(x_\varepsilon), \phi(\varepsilon)]$, and therefore $y_\varepsilon'''(\xi) \geq y_\varepsilon'''((j-1)/n)$. Then, we have

$$\begin{aligned}
y_\varepsilon\Big(\frac{j}{n}\Big) &= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6} y_\varepsilon'''(\xi) \\
&\geq F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} + \frac{1}{6n^6} y_\varepsilon'''\Big(\frac{j-1}{n}\Big) \\
&= F_{n,\varepsilon}\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) + \frac{1}{6n^6} M_\varepsilon\Big(y_\varepsilon\Big(\frac{j-1}{n}\Big)\Big) - \frac{\varepsilon}{n} \\
&\geq \left(\frac{n-1}{n}\right) y_\varepsilon\Big(\frac{j-1}{n}\Big)^{\frac{n}{n-1}} - \frac{\varepsilon}{n} \geq \left(\frac{n-1}{n}\right) \rho_{\varepsilon,j}^n - \frac{\varepsilon}{n} = \rho_{\varepsilon,j+1}^{n-1},
\end{aligned}$$

where the second inequality holds from Proposition 7(c), and in the third inequality we used the inductive step.

Part (c) is a direct extension of [52, Corollary 6.9]. Finally we prove (d). By definition, recall that

$$\phi(\varepsilon) = \int_0^1 \frac{1}{y(1-\log(y)) + \varepsilon} \mathrm{d}y.$$

We apply the change of variables $x = -\log(y)$ to get that

$$\phi(\varepsilon) = \int_0^\infty \frac{1}{1 + x + \varepsilon e^x} \mathrm{d}x.$$

Note that the function $f(x) = 1 + x - \varepsilon e^x$ has a unique root in $x \in [0, \infty)$ for $\varepsilon < 1$, that we denote r_ε (i.e., $f(r_\varepsilon) = 0$). In particular, we have $1 + x \geq \varepsilon e^x$ for every $x \leq r_\varepsilon$, and $1 + x \leq \varepsilon e^x$ for every $x \geq r_\varepsilon$. Then, we have

$$\int_0^{r_\varepsilon} \frac{1}{2(1+x)} \mathrm{d}x \leq \int_0^{r_\varepsilon} \frac{1}{1 + x + \varepsilon e^x} \mathrm{d}x$$

and

$$\int_{r_\varepsilon}^\infty \frac{1}{2\varepsilon e^x} \mathrm{d}x \leq \int_{r_\varepsilon}^\infty \frac{1}{1 + x + \varepsilon e^x} \mathrm{d}x.$$

By adding both inequalities we get

$$\frac{1}{2}\left(\int_0^{r_\varepsilon} \frac{1}{1+x} \mathrm{d}x + \int_{r_\varepsilon}^\infty \frac{1}{\varepsilon e^x} \mathrm{d}x\right) \leq \phi(\varepsilon).$$

On the other hand, we have

$$\phi(\varepsilon) \leq \int_0^{r_\varepsilon} \frac{1}{1+x} \mathrm{d}x + \int_{r_\varepsilon}^\infty \frac{1}{\varepsilon e^x} \mathrm{d}x,$$

and therefore, by evaluating the integrals, we have

$$\frac{1}{2}\left(\log(1+r_\varepsilon)+\frac{\exp(-r_\varepsilon)}{\varepsilon}\right)\leq\phi(\varepsilon)\leq\log(1+r_\varepsilon)+\frac{\exp(-r_\varepsilon)}{\varepsilon}.$$

Observe that $r_\varepsilon=\log(1+r_\varepsilon)+\log(1/\varepsilon)$, and therefore $r_\varepsilon\geq\log(1/\varepsilon)$. Furthermore, when ε is sufficiently small, we have $f(2\log(1/\varepsilon))=1+2\log(1/\varepsilon)-1/\varepsilon<0$, and therefore $r_\varepsilon\leq 2\log(1/\varepsilon)$. Then, for ε sufficiently small we have $\log(1/\varepsilon)\leq r_\varepsilon\leq 2\log(1/\varepsilon)$, which implies that

$$\begin{aligned}\log\left(1+\log\left(\frac{1}{\varepsilon}\right)\right)+\frac{\exp(-2\log(\frac{1}{\varepsilon}))}{\varepsilon}&\leq\log(1+r_\varepsilon)+\frac{\exp(-r_\varepsilon)}{\varepsilon}\\&\leq\log\left(1+2\log\left(\frac{1}{\varepsilon}\right)\right)+\frac{\exp(-\log(\frac{1}{\varepsilon}))}{\varepsilon}.\end{aligned}$$

The result now follows from the fact that the leftmost expression is lower bounded as

$$\log\log\left(\frac{1}{\varepsilon}\right)\leq\log\left(1+\log\left(\frac{1}{\varepsilon}\right)\right)+\frac{\exp(-2\log(\frac{1}{\varepsilon}))}{\varepsilon},$$

and the rightmost is upper bounded as $\log\left(1+2\log\left(\frac{1}{\varepsilon}\right)\right)+\frac{\exp(-\log(\frac{1}{\varepsilon}))}{\varepsilon}\leq 2\log\log\left(\frac{1}{\varepsilon}\right)$. □

Chapter 3

The Competition Complexity of Prophet Inequalities

3.1 Introduction

The standard prophet inequality [57, 58] 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 $v_1, \ldots, v_n$ are drawn independently from known distributions $F_1, \ldots, F_n$. We compare the expected value 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.

3.1.1 The model

In this chapter, 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. For convenience, we now give a summary of essential concepts and definitions required for understanding the chapter.

A more detailed explanation is given in chapter 1, section 1.2.2.

Prophet Inequality Setting. An online algorithm (ALG) observes a sequence of n non-negative values $v_1, \ldots, v_n$, drawn independently from known distributions $F_1, \ldots, F_n$. ALG must decide immediately upon seeing v_i whether to accept it and stop, or proceed to v_{i+1}. The stopping time $\rho \in [n] \cup \{\mathsf{null}\}$ determines the algorithm's reward $\mathbb{E}[\text{ALG}(v)] = \mathbb{E}[v_\rho]$. The offline algorithm (prophet) selects $\max_{i\in[n]} v_i$, achieving $\mathbb{E}[\max_{i\in[n]} v_i]$. ALG's *competitive ratio* is the worst-case ratio between its reward and the prophet's:

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

where $\alpha \in [0,1]$. The goal is to find an upper bound on k such that

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

Competition Complexity. In the *block model*, ALG observes k independent copies of the instance $(F_1,\ldots,F_n)$. Each copy contains n values arriving in sequence. The input is

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

where $v_i^{(j)}$ is drawn independently from F_i. The offline algorithm sees only the n values from one instance.

The *competition complexity* of a class of algorithms $\mathcal{A}$ is the smallest $k(\varepsilon)$ such that for any n, $F_1,\ldots,F_n$, and $k \geq k(\varepsilon)$:

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

For $\varepsilon = 0$, the competition complexity is unbounded even in the i.i.d. case [12].

Online Algorithms. We consider three types of algorithms:

- *Single threshold algorithms*: Use a fixed threshold τ and accept the first $v_i \geq \tau$.
- *Block threshold algorithms*: Assign thresholds $\boldsymbol{\tau} = (\tau_1,\ldots,\tau_k)$, one per block.
- *General threshold algorithms*: Use thresholds $\boldsymbol{\tau} = (\tau_1,\ldots,\tau_{nk})$, one per value.

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 [30] and [11], 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 [12] 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 8 in Section 3.2). We thus generalize the classic result of [73], 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.

3.1.2 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 5): 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 6): 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).

3.1.3 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 [12].

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 $\tau_1 > \tau_2 > \ldots, \tau_k$ (Lemma 6). The approximation factor of this stochastic dominance inequality is parameterized by $\tau_1, \ldots, \tau_k$, and it is obtained through a careful analysis of the minimum stochastic-dominance approximation factor achievable on each of the $k+1$ subintervals $[0, \tau_k), [\tau_k, \tau_{k-1}), \ldots, [\tau_1, \infty)$ 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 7). 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 [73]. Our key insight for this case is that the upper bound is asymptotically tight, which we show by providing an explicit lower bound construction.

3.1.4 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 23). 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 24). 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 25).

3.1.5 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 (Theorem 7): The $(1-\varepsilon)$-competition complexity of block-consistent prices for submodular and XOS combinatorial auctions is $O(\log(1/\varepsilon))$.

Additional Result 2 (Theorem 8): 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 [27, 36] 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 [20, 29].

Of course, the same question can be studied for other combinatorial settings such as matroids [18, 55] or matching constraints [33, 40]. Finally, we hope that our work will spark ideas on other notions of resource augmentation.

3.1.6 Organization

The chapter is organized as follows. In Section 3.3 we present our results for block threshold algorithms. Afterwards, in Section 3.4, we turn our attention to single threshold algorithms. We discuss combinatorial extensions in Section 3.5. We defer the discussion of additional arrival orders to Appendix B.2.

3.2 Block threshold reduction

We first show that block threshold algorithms are in fact worst case optimal.

Proposition 8. *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

$$\mathbb{E}_{v \sim (F_1,\dots,F_n)^{k'-1}}[\mathrm{ALG}(v)] < (1-\varepsilon) \cdot \mathbb{E}_{v \sim (F_1,\dots,F_n)}[\max_{i \in [n]} v_i]. \tag{3.1}$$

Given a block threshold algorithm ALG, let

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

and let

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

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. This is a valid probability since

$$p_{i+1}/(1-\sum_{j=i+2}^{m} p_j) \le 1 \Leftrightarrow \sum_{j=i+1}^{m} p_j \le 1.$$

and

$$\sum_{j=i+1}^{m} p_j = \Pr[i \cdot \alpha \le \max_{\ell \in [n]} v_\ell < m \cdot \alpha]$$

by definition of p_j's. Let $M_1 = \max_{i \in [m-1]} w_i$ with $w \sim (D_1, \ldots, D_{m-1})$, and $M_2 = \max_{i \in [n]} v_i$ with $v \sim (F_1, \ldots, F_n)$. 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

$$\begin{aligned}
\Pr[M_1 = r \cdot \alpha] &= \frac{p_{r+1}}{1-\sum_{j=r+2}^{m} p_j} \cdot \prod_{r'=r+1}^{m-1} \left(1 - \frac{p_{r'+1}}{1-\sum_{j=r'+2}^{m} p_j}\right) \\
&= \frac{p_{r+1}}{1-\sum_{j=r+2}^{m} p_j} \cdot \prod_{r'=r+1}^{m-1} \left(\frac{1-\sum_{j=r'+1}^{m} p_j}{1-\sum_{j=r'+2}^{m} p_j}\right) \\
&= p_{r+1} = \Pr\left[\left\lfloor \frac{M_2 \cdot \mathbf{1}[M_2 < m \cdot \alpha]}{\alpha} \right\rfloor \cdot \alpha = r \cdot \alpha\right]. \qquad (3.2)
\end{aligned}$$

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

$$\mathbb{E}_{w \sim (D_1, \ldots, D_{m-1})}[\max_{i \in [m-1]} w_i] \ge \mathbb{E}_{v \sim (F_1, \ldots, F_n)}[\max_{i \in [n]} v_i] - 2\alpha, \qquad (3.3)$$

since an α term is lost due to values above $m \cdot \alpha \ge \beta^*$, 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 $(D_1, \ldots, D_{m-1})^{k'-1}$ 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 $\mathrm{ALG}^\star$ for instance $(D_1, \ldots, D_{m-1})^{k'-1}$ is a block threshold algorithm.

Let $\tau_1^\star \ge \ldots \ge \tau_{k'-1}^\star$ be the monotone decreasing thresholds used by $\mathrm{ALG}^\star$. Consider the algorithm ALG' that selects a value in block ℓ if it is in the interval $[\alpha \cdot (1 + \lfloor \tau_\ell^\star/\alpha \rfloor), m \cdot \alpha)$. By Equation (3.2), the probability that ALG' selects an element in every iteration is the same as $\mathrm{ALG}^\star$, and given that both algorithms stop at block ℓ, the expectation of ALG' is at least the expectation of $\mathrm{ALG}^\star$. Thus,

$$\mathbb{E}_{w \sim (D_1, \ldots, D_{m-1})^{k'-1}}[\mathrm{ALG}^\star(w)] \le \mathbb{E}_{v \sim (F_1, \ldots, F_n)^{k'-1}}[\mathrm{ALG}'(v)] \le \mathbb{E}_{v \sim (F_1, \ldots, F_n)^{k'-1}}[\mathrm{ALG}(v)],$$

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,

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

where the equality is by definition of α, the first inequality is by Equation (3.3), and the last inequality is since $\varepsilon < 1$, and since by Equation (3.1), $\alpha > 0$. Thus, $k > k' - 1$, which concludes the proof. □

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.2.

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 B.1, 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.

3.3 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 5. *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 $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_k)$ we denote by $\mathrm{ALG}_{\boldsymbol{\tau}}$ the block threshold algorithm such that for every copy $j \in \{1, \ldots, k\}$ 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 3.3.1 we establish a lemma that allows us to evaluate the performance of any block threshold algorithm $\mathrm{ALG}_{\boldsymbol{\tau}}$ by establishing an approximate stochastic dominance inequality that extends the case of $k = 2$ to every k. This inequality is parameterized by the thresholds $\tau = (\tau_1, \ldots, \tau_k)$. In order to get the best possible approximation factor, we need to solve the induced max-min problem. In Section 3.3.2 we solve the corresponding max-min problem by presenting an explicit feasible solution. Finally, in Section 3.3.3, we present the full proof of Theorem 5, 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 $i \in \{1, \ldots, n\}$. More specifically, for $j \in \{1,2\}$, let $p_j = \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq \tau_j]$, where τ_1 and τ_2 will be determined by specifying the values of p_1 and p_2. By our assumption of $F_1, \ldots, F_n$ 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 $\text{ALG}_{\boldsymbol{\tau}}$ 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

$$\begin{aligned}
\Pr_{v\sim(F_1,\ldots,F_n)^2}[\text{ALG}_{\boldsymbol{\tau}}(v) \geq x] &\geq \phi_1(p_1,p_2) \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \quad \text{when } x \geq \tau_1, \\
\Pr_{v\sim(F_1,\ldots,F_n)^2}[\text{ALG}_{\boldsymbol{\tau}}(v) \geq x] &\geq \phi_2(p_1,p_2) \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \quad \text{when } \tau_1 > x \geq \tau_2, \\
\Pr_{v\sim(F_1,\ldots,F_n)^2}[\text{ALG}_{\boldsymbol{\tau}}(v) \geq x] &\geq \phi_3(p_1,p_2) \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \quad \text{when } \tau_2 > x, \qquad (3.4)
\end{aligned}$$

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

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

for every $x \geq 0$. The above inequality is stated in its general form for every k in Lemma 6. Then, in order to get the best possible approximation factor in this way, we solve the following max-min optimization problem:

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

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 7.

3.3.1 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 $\mathrm{ALG}_{\boldsymbol{\tau}}$ by establishing an approximate stochastic dominance inequality that extends (3.4) to every k. When $k = 2$, the inequalities in (3.4) 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 $F_1, \ldots, F_n \in \Delta$, and $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_k)$ such that $p_0 = 0$ and $p_\ell = \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{j\in[n]} v_j \geq \tau_\ell] > 0$ for every $\ell \in \{1,2,\ldots,k\}$, let

$$\Phi_1(F_1,\ldots,F_n,\boldsymbol{\tau}) = \sum_{\ell=1}^{k}\prod_{j=1}^{\ell}(1-p_j),$$

$$\Phi_i(F_1,\ldots,F_n,\boldsymbol{\tau}) = \frac{1}{p_i}\sum_{\ell=1}^{i-1} p_\ell \prod_{j=0}^{\ell-1}(1-p_j) + \sum_{\ell=i}^{k}\prod_{j=0}^{\ell}(1-p_j) \quad \text{for every } i \in \{2,\ldots,k\}, \text{ and}$$

$$\Phi_{k+1}(F_1,\ldots,F_n,\boldsymbol{\tau}) = \sum_{\ell=1}^{k} p_\ell \prod_{j=0}^{\ell-1}(1-p_j).$$

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

Lemma 6. *For every $F_1,\ldots,F_n \in \Delta$, every $x \geq 0$, and every decreasing $\boldsymbol{\tau} = (\tau_1,\ldots,\tau_k)$ such that $\Pr_{v\sim(F_1,\ldots,F_n)}[\max_{j\in[n]} v_j \geq \tau_1] > 0$, we have*

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

Proof. Given $x \geq 0$, let $\ell(\boldsymbol{\tau},x) = \min\{\ell \in \{1,\ldots,k+1\} : x \geq \tau_\ell\}$, that is, $\ell(\boldsymbol{\tau},x)$ 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 $\ell \in \{1,\ldots,k\}$ such that $\max_{i\in[n]} v_i^{(\ell)} > \tau_\ell$, then

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

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 $p_\ell = \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{j\in[n]} v_j \geq \tau_\ell]$ for every $\ell \in \{1,2,\ldots,k+1\}$; in particular $p_{k+1} = 1$. The following holds:

$$
\begin{aligned}
&\Pr_{v\sim(F_1,\ldots,F_n)^k}[\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x] \\
&= \sum_{\ell=1}^{k} \Pr_{v\sim(F_1,\ldots,F_n)^k}[\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x \text{ and } \mathcal{B} = \ell] \\
&= \sum_{\ell=1}^{k} \Pr_{v\sim(F_1,\ldots,F_n)^k}[\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell \mid \mathcal{B} > \ell-1] \Pr_{v\sim(F_1,\ldots,F_n)^k}[\mathcal{B} > \ell-1] \\
&\geq \sum_{\ell=1}^{k} \Pr_{v\sim(F_1,\ldots,F_n)^k}[v_j^{(\ell)} \notin [\tau_\ell, x) \text{ for all } j\in[n] \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell \mid \mathcal{B} > \ell-1] \prod_{j=0}^{\ell-1}(1-p_j) \\
&= \sum_{\ell=1}^{k} \Pr_{v^{(\ell)}\sim(F_1,\ldots,F_n)}[v_j^{(\ell)} \notin [\tau_\ell, x) \text{ for all } j\in[n] \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell] \prod_{j=0}^{\ell-1}(1-p_j),
\end{aligned} \tag{3.5}
$$

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 $\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x$ 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 $\ell(\boldsymbol{\tau}, x) \geq 2$ and let $\ell \in \{1, \ldots, \ell(\boldsymbol{\tau}, x) - 1\}$. In particular, we have $\tau_\ell > x$, and therefore,

$$
\begin{aligned}
&\Pr_{v\sim(F_1,\ldots,F_n)^k}[v_j^{(\ell)} \notin [\tau_\ell, x) \text{ for all } j\in[n] \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell] \\
&= \Pr_{v^{(\ell)}\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell] \\
&= \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \cdot \frac{\Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq \tau_\ell]}{\Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x]} \\
&\geq \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \cdot \frac{\Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq \tau_\ell]}{\Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq \tau_{\ell(\boldsymbol{\tau},x)}]} \\
&= \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x] \cdot \frac{p_\ell}{p_{\ell(\boldsymbol{\tau},x)}},
\end{aligned} \tag{3.6}
$$

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

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

$$\begin{aligned}
&\Pr_{v^{(\ell)} \sim (F_1,\ldots,F_n)}[v_j^{(\ell)} \notin [\tau_\ell, x) \text{ for all } j \in [n] \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq \tau_\ell] \\
&\geq \Pr_{v^{(\ell)} \sim (F_1,\ldots,F_n)}[v_j^{(\ell)} \notin [\tau_\ell, x) \text{ for all } j \in [n] \text{ and } \max_{i\in[n]} v_i^{(\ell)} \geq x] \\
&\geq \sum_{i=1}^{n} \Pr_{v \sim (F_1,\ldots,F_n)}[v_i \geq x \text{ and } v_j < \tau_\ell \text{ for all } j \in [n] \setminus \{i\}] \\
&= \sum_{i=1}^{n} \Pr_{v \sim (F_1,\ldots,F_n)}[v_i \geq x] \Pr_{v \sim (F_1,\ldots,F_n)}[v_j < \tau_\ell \text{ for all } j \in [n] \setminus \{i\}] \\
&\geq \sum_{i=1}^{n} \Pr_{v \sim (F_1,\ldots,F_n)}[v_i \geq x] \Pr_{v \sim (F_1,\ldots,F_n)}[\max_{j\in[n]} v_j < \tau_\ell] \\
&= (1 - p_\ell) \sum_{i=1}^{n} \Pr_{v \sim (F_1,\ldots,F_n)}[v_i \geq x] \\
&\geq (1 - p_\ell) \Pr_{v \sim (F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x],
\end{aligned} \tag{3.7}$$

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 the upper bound on $\max_{j\in[n]} v_j$ implies the corresponding event in the third line, and the last inequality holds by the union bound.

Then, when $\ell(\boldsymbol{\tau}, x) = 1$, from (3.5) and (3.7) we get

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

When $\ell(\boldsymbol{\tau},x) \in \{2,\dots,k\}$, from (3.5), (3.6) and (3.7) we get

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

Finally, when $\ell(\boldsymbol{\tau},x) = k+1$, from (3.5) and (3.6) we get

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

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

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

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

which finishes the proof of the lemma. □

Since Lemma 6 provides an approximate stochastic dominance inequality parameterized in the thresholds $\boldsymbol{\tau} = (\tau_1,\dots,\tau_k)$, in order to get the best possible approximation factor by using this approach we have to solve the following max-min problem:

$$\begin{aligned}
&\max\left\{ \min_{i\in\{1,\dots,k+1\}} \Phi_i(F_1,\dots,F_n,\boldsymbol{\tau}) : \tau_1 > \tau_2 > \dots > \tau_k \geq 0 \right\} \\
&= \max\left\{ \Phi(F_1,\dots,F_n,\boldsymbol{\tau}) : \tau_1 > \tau_2 > \dots > \tau_k \geq 0 \right\}. \qquad (3.8)
\end{aligned}$$

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

3.3.2 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 (3.8). 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 (3.8). This lower bounds holds by providing a specific set of thresholds, obtained after finding a well-chosen set of quantiles $p_1, \dots, p_k$.

Lemma 7. *For every $F_1, F_2, \dots, F_n \in \Delta$ and every $k \geq 2$, let $\overline{\boldsymbol{\tau}} = (\overline{\tau}_1, \dots, \overline{\tau}_k)$ such that*

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

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

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

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

Then, the lemma follows since, overall, we get

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

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

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

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

$$\Phi_2(F_1,\ldots,F_n,\overline{\boldsymbol{\tau}}) = \frac{p_1}{p_2} + \sum_{\ell=2}^{k}\prod_{j=0}^{\ell}(1-p_j) \geq \frac{p_1}{p_2} + (1-p_1)(1-p_2)$$
$$= \frac{1-(5/4)^{-5/4}}{1-(5/4)^{-(5/4)^2}} + (5/4)^{-5/4}\cdot(5/4)^{-(5/4)^2} > 1,$$

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

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

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

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

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

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

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

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

3.3.3 Putting it all together

With Lemma 6 and Lemma 7 at hand, we are now ready to prove our main theorem.

Proof of Theorem 5. Given $\varepsilon > 0$, by Lemma 6 and Lemma 7, we have that for every $F_1,\ldots,F_n \in \Delta$ and every $k \geq \max(2, \log_{5/4}\log_{5/4}(1/\varepsilon))$, by taking $\overline{\boldsymbol{\tau}} = (\overline{\tau}_1,\ldots,\overline{\tau}_k)$ as defined in Lemma 7, the following holds:

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

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 [12, Theorem 2.2] show that the $(1-\varepsilon)$-competition complexity of the class of general threshold algorithms when $F_1 = F_2 = \cdots = F_n$ is $\Omega(\log\log(1/\varepsilon))$. □

We note that our proof of Theorem 5 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 [12] even in the i.i.d. case.

We know from the lower bound on the competition complexity of [12] that one cannot strengthen Lemma 7 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 [12] 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 9. *For every $k \geq 1$, there exists a positive integer value n_k and $F_1,\ldots,F_{n_k} \in \Delta$, such that for every $\boldsymbol{\tau} = (\tau_1,\ldots,\tau_k)$, there exists $x \in \mathbb{R}_{\geq 0}$ such that*

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

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

Proof. For every $k \geq 1$, let $n_k = 3^{3^{k+2}}$, and for every $i \in [n_k]$, let F_i be the distribution of the random variable $v_i = (i - \mathrm{Uniform}[0,1]) \cdot \mathrm{Bernoulli}(1/i)$. For ease of notation, let $F = (F_1,\ldots,F_{n_k})$. We first observe that since $\max_{i\in[n_k]} i \cdot \mathrm{Bernoulli}(1/i)$ is distributed uniformly on the set $\{1,\ldots,n_k\}$, 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 $\boldsymbol{\tau} = (\tau_1, \ldots, \tau_k)$, let $A = \{i \in [k] : \tau_i \leq n_k \cdot 3^{-3^i}\}$. If $A = \emptyset$, then for $x = 0$, it holds that

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

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 $\overline{B}_\ell$ 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

$$\Pr[\overline{B}_\ell^x] = 1 - \frac{\tau_\ell}{\lceil \tau_\ell \rceil} \cdot \left(\prod_{j=\lceil \tau_\ell \rceil + 1}^{x} \frac{i-1}{i} \right) = 1 - \frac{\tau_\ell}{x}. \tag{3.9}$$

Thus, we can bound the probability of the algorithm selecting a value of at least x as follows:

$$\begin{aligned}\Pr_{v\sim F^k}[\mathrm{ALG}_{\boldsymbol{\tau}}(v) \geq x] &\leq 1 - \Pr[\overline{B}_{i^*}^x \wedge \overline{B}_1 \wedge \ldots \overline{B}_{i^*-1}] = 1 - \left(1 - \frac{\tau_\ell}{x}\right) \prod_{j=1}^{i^*-1} \Pr[\overline{B}_j] \\ &\leq 1 - \left(1 - \frac{3^{-3^{i^*}}}{3^{-2\cdot 3^{i^*-1}}}\right) \prod_{j=1}^{i^*-1} 3^{-3^j} < \left(1 - 3^{-3^{i^*}}\right)\left(1 - 3^{-2\cdot 3^{i^*-1}}\right) \\ &\leq \left(1 - 3^{-3^k}\right)\left(1 - 3^{-2\cdot 3^{i^*-1}}\right) = (1-\varepsilon_k) \Pr_{v\sim F}[\max_{i\in[n_k]} v_i \geq x],\end{aligned}$$

where the first inequality holds since if the algorithm does not select a value in the first $i^* - 1$ iterations, and in iteration i^* there exists a value in $[\tau_{i^*,x})$ then the algorithm selects it; the equality is by Equation (3.9) and by independence of the events; the second inequality is since $\tau_j > n_k \cdot 3^{-3^j}$ for $j < i^*$, and $\tau_{i^*} \leq n_k \cdot 3^{-3^{i^*}}$; the third inequality holds for every $i^* \geq 1$; the fourth inequality holds since $i^* \leq k$; the last equality is by definition of ε_k and since $\max_{i\in[n_k]} v_i$ is uniformly distributed over the interval $[0, n_k]$. □

3.4 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 6. *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 6 through Lemma 8 and Proposition 10 below. Lemma 8 establishes the upper bound claimed in the theorem, and Proposition 10 shows a matching lower bound.

We begin with Lemma 8, 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 [73]; we work a bit harder to show that it also holds in a stochastic dominance sense.

Lemma 8. *For every $F_1, \ldots, F_n$, there exists a threshold τ^* such that for every $\varepsilon > 0$, and for every $k \geq \log_2(1/\varepsilon)$*

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

Furthermore, the stronger approximate stochastic-dominance inequality

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

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

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

$$\prod_{i=1}^{n} \Pr_{v_i \sim F_i}[v_i \leq \tau^*] = \frac{1}{2}. \tag{3.10}$$

In what follows, for ease of notation we consider the input sequence of nk values as $v_1, \ldots, v_{nk}$, 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

$k \geq 1$,

$$\mathbb{E}_{v\sim(F_1,\ldots,F_n)^k}[\text{ALG}_{\tau^*}(v) \mid \text{there exists } i \in [nk] :\ v_i \geq \tau^*] \geq \mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i]. \quad (3.11)$$

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

$$\Pr_{v\sim(F_1,\ldots,F_n)^k}[\text{ALG}_{\tau^*}(v) \geq x \mid \text{there exists } i \in [nk] :\ v_i \geq \tau^*] \geq \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \geq x]. \quad (3.12)$$

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 (3.11) (respectively, Equation (3.12)) can be rewritten as $\mathbb{E}_{v\sim(F_1,\ldots,F_n)^k}[\text{ALG}_{\tau^*}(v) \mid \hat{r} \neq 0]$ (respectively, $\Pr_{v\sim(F_1,\ldots,F_n)^k}[\text{ALG}_{\tau^*}(v) \geq x \mid \hat{r} \neq 0]$). Thus, it is sufficient to prove that Equation (3.11) holds for every non-zero realization of $\hat{r}$, i.e., for every $k \geq 1$ and every $r \in [k]$,

$$\mathbb{E}_{v\sim(F_1,\ldots,F_n)^k}[\text{ALG}_{\tau^*}(v) \mid \hat{r} = r] \geq \mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i]. \quad (3.13)$$

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 (3.13) is equivalent to the proof of the original prophet inequality, provided by Samuel-Cahn [73]. We next prove the stronger claim of Equation (3.12). To this end, we show that for every $x \in \mathbb{R}_{\geq 0}$, and for $r \neq 0$ it holds that

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

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,

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

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 $\Pr[\text{for all } j \in \{1, \ldots (r-1)n\} : v_j < \tau^*] = 1/2^{r-1}$, and since $\Pr[\text{for all } j \in \{(r-1)n+1, \ldots, i-1\} : v_j < \tau^*] \geq 1/2$, 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 (3.11) (respectively for the "furthermore" part, Equation (3.12)) 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 10, 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 10. *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*

$$\mathbb{E}_{v \sim F^{k \cdot n}}[\text{ALG}_\tau(v)] < (1-\varepsilon) \cdot \mathbb{E}_{v \sim F^n}[\max_{i \in [n]} v_i]. \tag{3.14}$$

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 (3.14) 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 (3.14) 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 (3.14) 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 (3.14) satisfies that

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

where the last inequality holds for every $n \geq 2$, and $q \geq \frac{29}{20n}$. Thus, Equation (3.14) holds. □

3.5 Combinatorial extensions

In this section, we present a generalization of our model to combinatorial settings, and discuss some initial results. We present the general model in Section 3.5.1, and some preliminary results for combinatorial auctions in Section 3.5.2. In Section B.4 we give additional results for bipartite matching with one-sided vertex arrivals. This latter result is actually implied by the result in Section 3.5.2. The purpose of presenting an alternative proof is to show how a different technique, in this case online contention resolution schemes, can also be used to study the competition complexity in combinatorial settings.

3.5.1 A general model

In online combinatorial Bayesian selection problems, there is a series of n decisions that need to be made. Each decision $i \in [n]$ is associated with a set of alternatives A_i from which the decision-maker needs to choose, and with an information v_i drawn independently from some F_i on support S_i that is revealed to the decision-maker at the time of decision i. We denote by $S = \bigtimes_{i \in [n]} S_i$, $F = \bigtimes_{i \in [n]} F_i$, and $A = \bigtimes_{i \in [n]} A_i$. Additionally, there is a non-empty feasibility constraint $\mathcal{F} \subseteq A$, such that the decision-maker, must select a tuple of alternatives $(a_1, \ldots, a_n)$ that is in $\mathcal{F}$, and there is a reward function $f : A \times S \to \mathbb{R}_{\geq 0}$.

At each step $i \in [n]$, the algorithm observes v_i, and needs to select an alternative $a_i \in A_i$ in an immediate and irrevocable way. The algorithm's performance is measured against the offline optimum. By fixing a class of feasibility constraints $\mathbb{C}_{\mathcal{F}}$, a class of valuation functions $\mathbb{C}_f$, and a class of distributions $\mathbb{C}_F$, an online algorithm ALG is α-competitive if

$$\inf_{F \in \mathbb{C}_F} \inf_{\mathcal{F} \in \mathbb{C}_{\mathcal{F}}} \inf_{f \in \mathbb{C}_f} \frac{\mathbb{E}_{v \sim F}[f(\mathrm{ALG}(v), v) \cdot \mathbf{1}[\mathrm{ALG}(v) \in \mathcal{F}]]}{\mathbb{E}_{v \sim F}[\max_{a \in \mathcal{F}} f(a, v)]} \geq \alpha,$$

where $\mathrm{ALG}(v)$ is the (possibly random) tuple of alternatives chosen by the algorithm ALG when observing a sequence of information v.

A simple example of this setting is when for all $i \in [n]$ we have $A_i = \{0,1\}$, $\mathbb{C}_{\mathcal{F}} = \{\{(a_1, \ldots, a_n) \mid \sum_{i \in [n]} a_i \leq 1\}\}$, $\mathbb{C}_f = \{\{f(a,v) = \sum_{i \in [n]} a_i \cdot v_i\}\}$, and $\mathbb{C}_F = \Delta^n$, which corresponds to the standard (single-choice) prophet inequality setting.

We next describe a generalization of the block model to the combinatorial Bayesian selection framework. The input to the algorithm is given by k copies of an online combinatorial Bayesian selection problem, so there are kn decisions in total. For every $j \in [kn]$, the j-th decision is of type i if $j \equiv i \mod n$, and for each decision j of type i, the information v_j is sampled independently according to F_i. For every decision j of type i we must select an alternative in A_i.

The output of the algorithm is a kn-dimensional vector of alternatives. To define feasibility, and to evaluate the reward achieved by an output, we require that there is an infinite series of classes of feasibility constraints $(\mathbb{C}^i_{\mathcal{F}})_{i \in \mathbb{N}}$, and an infinite series of classes of reward functions $(\mathbb{C}^i_f)_{i \in \mathbb{N}}$ so that we can evaluate feasibility via $\mathcal{F}_k \in \mathbb{C}^k_{\mathcal{F}}$ and the reward via $f_k \in \mathbb{C}^k_f$. For a concrete problem it is typically clear how to define $(\mathbb{C}^i_{\mathcal{F}})_{i \in \mathbb{N}}$ and $(\mathbb{C}^i_f)_{i \in \mathbb{N}}$ based on $\mathbb{C}_{\mathcal{F}}$ and $\mathbb{C}_f$.

We are interested in comparing the expected reward of the algorithm on k copies to the expected optimal reward on a single copy.

Definition 2 (Combinatorial competition complexity)**.** *Given a series of decisions associated with alternatives $A = A_1 \times \ldots \times A_n$, a class of distributions $\mathbb{C}_F$, an infinite series of classes of feasibility constraints $(\mathbb{C}^i_{\mathcal{F}})_{i \in \mathbb{N}}$, and an infinite series of classes of reward functions $(\mathbb{C}^i_f)_{i \in \mathbb{N}}$, for every $\varepsilon \geq 0$, 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 $k \geq k(\varepsilon)$, every $F \in \mathbb{C}_F$, $\mathcal{F}_k \in \mathbb{C}^k_{\mathcal{F}}$, $f_k \in \mathbb{C}^k_f$, it holds that*

$$\max_{\mathrm{ALG} \in \mathcal{A}_{n,k}} \mathbb{E}_{v \sim F^k}[f_k(\mathrm{ALG}(v), v) \cdot \mathbf{1}[\mathrm{ALG}(v) \in \mathcal{F}_k]] \geq (1-\varepsilon) \cdot \mathbb{E}_{v \sim F}[\max_{a \in \mathcal{F}_1} f_1(a, v)], \quad (3.15)$$

where $\mathcal{A}_{n,k}$ are all algorithms in $\mathcal{A}$ that are defined on $\mathbb{C}_F$ and $\mathbb{C}^k_{\mathcal{F}}$, and ALG$(v)$ *is the (possibly random) tuple of alternatives chosen by* ALG *when observing a sequence of values v.*

For example, to obtain the competition complexity for the standard prophet inequality setting, we can let $A_i = \{0,1\}$ for all $i \in [n]$, $\mathbb{C}^k_{\mathcal{F}} = \{\{(a_1, \ldots, a_{n \cdot k}) \mid \sum_{i \in [n \cdot k]} a_i \leq 1\}\}$, $\mathbb{C}^k_f = \{\{f(a,v) = \sum_{i \in [n \cdot k]} a_i \cdot v_i\}\}$, and $\mathbb{C}_F = \Delta^n$.

3.5.2 Combinatorial auctions

In this section, we generalize the static pricing scheme by Feldman et al. [36] for XOS markets, to a dynamic pricing scheme that has a $(1-\varepsilon)$-competition complexity of $O(\log(1/\varepsilon))$, and a static pricing scheme that has a $(1-\varepsilon)$-competition complexity of $O(1/\varepsilon)$.

In the combinatorial auction setting, there is a set M of m items, and n agents. Each agent $i \in [n]$ is associated with a valuation function $v_i : 2^M \to \mathbb{R}_{\geq 0}$ that is drawn independently from a distribution F_i. We consider price-based algorithms, which set a price p_j for each item $j \in M$. The agents then arrive one-by-one, and purchase a set of available items in their demand. That is, agent i buys a set T_i that maximizes the utility $v_i(T) - \sum_{j \in T} p_j$ over all $T \subseteq M'$, where $M' \subseteq M$ are the items that are available when agent i arrives. We evaluate the performance of a pricing scheme by the expected social welfare it achieves, that is $\mathbb{E}_{v \sim F}[\sum_{i \in [n]} v_i(T_i)]$.

We consider a repeated version of the combinatorial auction problem where we see k independent copies of the buyers. As before we refer to each copy as a block. The valuation of buyer $i \in [nk]$ of type $r \in [n]$ is drawn independently according to the distribution F_r. Each buyer $i \in [nk]$ purchases a set of available items T_i that maximizes their utility. We compare the expected social welfare of a price-based algorithm on k copies to the expected maximum social welfare on a single copy.

A pricing scheme is called *static* if the prices of the items are fixed in advance before the arrival of the agents, and is called *dynamic*, if the prices may adapt after each agent has made a purchase (before the arrival of the next agent). We define *block-consistent prices* as prices that are static throughout each block but can change between blocks.

A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is called XOS (a.k.a. fractionally subadditive) if there exists a non-empty set of additive functions $G = \{g_1, \ldots g_s\}$ for some positive integer s, such that $v(T) = \max_{t \in [s]} \sum_{j \in T} g_t(j)$ for every $T \subseteq M$. A valuation function $v : 2^M \to \mathbb{R}_{\geq 0}$ is submodular if for every T, T' with $T \subseteq T'$ and every $j \in M \setminus T' \subseteq M$ it holds that $v_i(T \cup \{j\}) - v_i(T) \geq v_i(T' \cup \{j\}) - v_i(T')$. Every submodular valuation function is XOS, but not vice versa [7]. We refer to combinatorial auctions in which all valuation functions are submodular (resp. XOS), as submodular (resp. XOS) combinatorial auctions.

Theorem 7. *For every $k \geq 1$, and every XOS combinatorial auction defined by M, n, and a product distribution $F = F_1 \times \ldots \times F_n$, there exists a block-consistent pricing scheme on k copies such that*

$$\mathbb{E}_{v \sim F^k}\Big[\sum_{i \in [nk]} v_i(T_i)\Big] \geq \left(1 - \frac{1}{2^k}\right) \cdot \mathbb{E}_{v \sim F}\Big[\max_{(T_1, \ldots, T_n) \in P(M)} \sum_{i \in [n]} v_i(T_i)\Big],$$

where T_i is the demanded set purchased by each agent $i \in [n]$, and $P(M)$ is the set of all partitions of M into n sets. In particular, the $(1-\varepsilon)$-competition complexity of block-consistent prices for submodular and XOS combinatorial auctions is $O(\log(1/\varepsilon))$.

Proof. For every item $j \in M$, and agent i let $\mathrm{OPT}_{j,i} = E_{v \sim F}[g_{i(j)}(j) \cdot \mathbf{1}[i = i(j)]]$, where $i(j)$ is the agent that receives item j in the welfare-maximizing allocation according to valuation profile $v = (v_1, \ldots, v_n)$, and $g_{i(j)}$ is the additive function that corresponds to the definition of XOS valuation that maximizes the value of the set agent $i(j)$ receives in the optimal allocation according to valuation profile $v = (v_1, \ldots, v_n)$.

Let p be the the block-consistent pricing scheme where the price during block $t \in [k]$ for item $j \in M$ is $p_{t,j} = (1 - \frac{1}{2^{k+1-t}}) \cdot \sum_{i \in [n]} \mathrm{OPT}_{j,i}$. Let X_j be the block in which item j is purchased (if it is not purchased, then let $X_j = k+1$), and let $q_{t,j} = \Pr[X_j = t]$.

The welfare of the mechanism can be decomposed into revenue and surplus. For the revenue, we have

$$\mathbb{E}_{v \sim F^k}[\text{Revenue}] = \sum_{j \in M} \sum_{t \in [k]} q_{t,j} \cdot p_{t,j} = \sum_{j \in M} \sum_{t \in [k]} q_{t,j} \cdot \left(1 - \frac{1}{2^{k+1-t}}\right) \cdot \sum_{i \in [n]} \mathrm{OPT}_{j,i}.$$

For every agent $i \in [nk]$, let's denote by AV_i the set of items that is available when i arrives. For each type $r \in [n]$ let's denote by T_r^* the set that agent i of type r receives in the welfare-maximizing allocation on profile $v \sim F$, and by g_r^* the additive function that maximizes agent r's value with respect to v_i and T_r^*. Then, if agent i is of type r,

$$\begin{aligned}
\mathbb{E}_{v\sim F^k}[\text{Surplus}_i] &= \sum_{M'\subseteq M} \Pr[\text{AV}_i = M'] \cdot \mathbb{E}_{v_i\sim F_r}\Big[\max_{T\subseteq M'}(v_i(T) - p(T))\Big] \\
&\geq \sum_{M'\subseteq M} \Pr[\text{AV}_i = M'] \cdot \mathbb{E}_{v\sim F}[v_r(T_r^* \cap M') - p(T_r^* \cap M')] \\
&\geq \sum_{M'\subseteq M} \Pr[\text{AV}_i = M'] \cdot \mathbb{E}_{v\sim F}\Big[\sum_{j\in T_r^*\cap M'} (g_r^*(j) - p(j))\Big] \\
&\geq \sum_{j\in M} \Big(1 - \sum_{t=1}^{\lceil i/n \rceil} q_{j,t}\Big) \cdot \frac{\text{OPT}_{j,r}}{2^{k+1-\lceil \frac{i}{n} \rceil}},
\end{aligned}$$

where the first inequality is since v_i and v_r are drawn from the same distribution, and since T_i is the demand set; the second inequality is by the definition of an XOS function; and the last inequality is by the definition of $\text{OPT}_{j,i}$, and because the probability that $j \in AV_i$ is at least 1 minus the sum of probabilities that j was sold up to block $\lceil i/n \rceil$ included.

Since $\mathbb{E}_{v\sim F^k}[\text{Welfare}] = \mathbb{E}_{v\sim F^k}[\text{Revenue}] + \sum_{i\in[nk]} \mathbb{E}_{v\sim F^k}[\text{Surplus}_i]$, it is sufficient to lower bound for each $j \in M$, and $i \in [n]$, the sum of the coefficient of $\text{OPT}_{j,i}$ in the revenue with the sum of coefficients of $\text{OPT}_{j,i}$ in the surplus, which is

$$\begin{aligned}
&\sum_{t=1}^{k} q_{t,j}\left(1 - \frac{1}{2^{k+1-t}}\right) + \sum_{t=1}^{k} \frac{1 - \sum_{t'=1}^{t} q_{j,t'}}{2^{k+1-t}} \\
&= 1 - \frac{1}{2^k} + \sum_{t=1}^{k} q_{t,j}\left(1 - \frac{1}{2^{k+1-t}}\right) - \sum_{t=1}^{k}\sum_{t'=t}^{k} \frac{q_{j,t}}{2^{k+1-t'}} \\
&= 1 - \frac{1}{2^k}.
\end{aligned}$$

This concludes the proof. □

We next show a similar result with worse guarantees that uses static pricing.

Theorem 8. *For every $k \geq 1$, and every XOS combinatorial auction defined by M, n, and a product distribution $F = F_1 \times \ldots \times F_n$, there exists a static pricing scheme on k copies such that*

$$\mathbb{E}_{v\sim F^k}\Big[\sum_{i\in[nk]} v_i(T_i)\Big] \geq \left(1 - \frac{1}{k+1}\right) \cdot \mathbb{E}_{v\sim F}\Big[\max_{(T_1,\ldots,T_n)\in P(M)} \sum_{i\in[n]} v_i(T_i)\Big],$$

where T_i is the demanded set purchased by each agent $i \in [n]$, and $P(M)$ is the set of all partitions of M into n sets. In particular, the $(1-\varepsilon)$-competition complexity of static prices for submodular and XOS combinatorial auctions is $O(1/\varepsilon)$.

Proof. We define price $p_j = (1 - \frac{1}{k+1}) \sum_{i \in [n]} \mathrm{OPT}_{j,i}$, where $\mathrm{OPT}_{j,i}$ is defined as in the proof of Theorem 7. Let X_j be the indicator whether item j was sold to one of the nk agents, and let $q_j = \Pr[X_j = 1]$. It holds that

$$\mathbb{E}_{v \sim F^k}[\text{Revenue}] = \sum_{j \in M} q_j \cdot p_j = \sum_{j \in M} q_j \cdot \left(1 - \frac{1}{k+1}\right) \cdot \sum_{i \in [n]} \mathrm{OPT}_{j,i}.$$

For every agent $i \in [nk]$, let's denote by AV_i the set of items that is available when i arrives. For each type $r \in [n]$ let's denote by T_r^* the set that agent i of type r receives in the welfare-maximizing allocation on profile $v \sim F$, and by g_r^* the additive function that maximizes agent r's value with respect to v_i and T_r^*. Then, if agent i is of type r,

$$\begin{aligned}
\mathbb{E}_{v \sim F^k}[\text{Surplus}_i] &= \sum_{M' \subseteq M} \Pr[\mathrm{AV}_i = M'] \cdot \mathbb{E}_{v_i \sim F_r}\left[\max_{T \subseteq M'} (v_i(T) - p(T))\right] \\
&\geq \sum_{M' \subseteq M} \Pr[\mathrm{AV}_i = M'] \cdot \mathbb{E}_{v \sim F}[v_r(T_r^* \cap M') - p(T_r^* \cap M')] \\
&\geq \sum_{M' \subseteq M} \Pr[\mathrm{AV}_i = M'] \cdot \mathbb{E}_{v \sim F}\left[\sum_{j \in T_r^* \cap M'} (g_r^*(j) - p(j))\right] \\
&\geq \sum_{j \in M} (1 - q_j) \cdot \frac{\mathrm{OPT}_{j,r}}{k+1}.
\end{aligned}$$

Since $\mathbb{E}_{v \sim F^k}[\text{Welfare}] = \mathbb{E}_{v \sim F^k}[\text{Revenue}] + \sum_{i \in [nk]} \mathbb{E}_{v \sim F^k}[\text{Surplus}_i]$, it is sufficient to lower bound for each $j \in M$, and $i \in [n]$, the sum of the coefficient of $\mathrm{OPT}_{j,i}$ in the revenue with the sum of coefficients of $\mathrm{OPT}_{j,i}$ in the surplus, which is

$$q_j \left(1 - \frac{1}{k+1}\right) + \sum_{t=1}^{k} \frac{(1 - q_j)}{k+1} = 1 - \frac{1}{k+1}.$$

This concludes the proof. □

Chapter 4

Splitting Guarantees for Prophet Inequalities via Nonlinear Systems

4.1 Introduction

Hill and Kertz [45] provided an algorithm that guarantees an approximation ratio of $1 - 1/e$ and an upper bound of $\gamma \approx 0.745$ on the approximation ratio by studying the optimal dynamic program for the worst-case distributions. Later, [45] used the recursion from the optimal dynamic program in [45] to provide an ordinary differential equation (ODE)—that we termed Hill and Kertz equation for simplicity and in honor to both authors—where the γ bound is embedded as a *unique constant* that guarantees crucial analytical properties of the solution of the ODE: $y' = y(\ln y - 1) - 1/\gamma + 1$, $y(0) = 1$, $y(1) = 0$. However, the lower bound on the approximation remained $1 - 1/e$ for many years until [23] used the Hill and Kertz equation to provide an algorithm that attains an approximation ratio of at least γ for any n.

Over the years, it has remained elusive for $k \geq 2$ how to get a result on the line of the Hill and Kertz equation, that is, to obtain provable approximation ratios via studying a closed-form differential equation related to the optimal dynamic programming solution. This motivates the central question of this work: *Can we find a closed-form nonlinear system of differential equations to lower bound the optimal asymptotic approximation ratio for the i.i.d. k-selection prophet inequality?*

4.1.1 Our contributions and techniques

Our first main result characterizes the optimal approximation ratio via a new infinite-dimensional linear program. In our second main result, we provide a closed-form nonlinear system of differential equations that gives provable lower bounds on the approximation ratio for (k, n)-PIP as defined in section 1.2.3, when n is large enough.

In our third result, applying our new provable lower bounds for (k,n)-PIP, we find a tight approximation ratio for the stochastic assignment problem. Below, we present more details about our results.

Exact Formulation for (k,n)-PIP. Our first step towards a characterization of the asymptotic approximation ratio is a new infinite-dimensional linear program that *characterizes* the optimal approximation ratio for (k,n)-PIP. This formulation is inspired by writing the optimal dynamic programming formulation in the quantile space. We take a minimax approach where we search the worst-case distribution while optimizing for the dynamic program's value. We show that in the continuous space of quantiles $[0,1]$, such a problem is linear; see formulation $[\mathrm{P}]_{n,k}$ in Section 4.3 for the details of the formulation. While there exist other linear programming formulations for (k,n)-PIP (see, e.g., [47]), using our formulation, we can provide an analysis as n grows that organically provides the nonlinear system of differential equations that we later use to get provable lower bounds. In Section 4.3, we provide the exact linear programming formulation and the proof that characterizes the optimal approximation ratio for (k,n)-PIP.

Approximation via a Nonlinear System. The analysis of our infinite-dimensional program as n approaches infinity leads us to introduce a system of k coupled nonlinear differential equations, extending the Hill and Kertz equation ($k=1$). This new nonlinear system is parameterized by k nonnegative values $\theta_1,\ldots,\theta_k$, and we look for functions $y_1,\ldots,y_k$ satisfying the following in the interval $[0,1)$:

$$(\Gamma_k(-\ln y_k))' = k!\,(1-1/(k\theta_k)) - \Gamma_{k+1}(-\ln y_k), \tag{4.1}$$

$$(\Gamma_k(-\ln y_j))' = k! - \Gamma_{k+1}(-\ln y_j) - \frac{\theta_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln y_{j+1})) \text{ for every } j \in [k-1], \tag{4.2}$$

$$y_j(0) = 1 \text{ and } \lim_{t\uparrow 1} y_j(t) = 0 \quad \text{for every } j \in [k], \tag{4.3}$$

where $\Gamma_\ell(x) = \int_x^\infty t^{\ell-1}e^{-t}\,\mathrm{d}t$ is the upper incomplete gamma function. Note that for $k=1$ in (4.1)-(4.3) we recover the Hill and Kertz differential equation. Now, if $\gamma_{n,k}$ denotes the optimal approximation ratio for (k,n)-PIP, we can prove that there is $n_0 = n_0(k)$ such that for $n \geq n_0$ we have

$$\gamma_{n,k} \geq \left(1 - 24k\frac{\ln(n)^2}{n}\right)\sum_{j=1}^{k}\theta_j^\star, \tag{4.4}$$

where $\theta_1^\star,\ldots,\theta_k^\star$ are the values for which there exists a solution to the nonlinear system of differential equations (4.1)-(4.3). We remark that the lower bound on the approximation ratio given by the nonlinear system in inequality (4.4) is simply obtained by the summation of the k constants that define it. For instance, when $k = 2$, the two constants for which the nonlinear system has a solution are $\theta_1^\star \approx 0.346$ and $\theta_2^\star \approx 0.483$, and therefore we get a provable lower bound of ≈ 0.829 on $\gamma_{n,2}$, the optimal approximation ratio for 2 selections, for any n large enough.

To prove inequality (4.4) we employ a dual-fitting approach within our infinite-dimensional linear program. Namely, we introduce a dual infinite-dimensional program of our exact formulation, and using the solution of the nonlinear system (4.1)-(4.3), we explicitly construct feasible solutions for this dual. Then, the theorem is obtained by a weak-duality argument. The details are presented in Section 4.4. We remark that our analysis requires a careful study of the nonlinear system (4.1)-(4.3); which we also provide in Section 4.4. The small multiplicative loss $(1 - 24k\ln(n)^2/n)$ appears when we construct the dual feasible solution, and it's needed in our analysis to ensure feasibility in the dual problem; note that this loss vanishes as n grows. We finally note that any feasible solution to our dual program can be implemented using a quantile-based algorithm; therefore, we can implement an algorithm that has an approximation ratio at least $(1 - 24k\ln(n)^2/n)\sum_{i=1}^k \theta_i^\star$.

k	1	2	3	4	5
Our approach ($\sum_{i=1}^k \theta_i^\star$)	0.7454	0.8293	0.8648	0.8875	0.9035
[10]	0.6543	0.7427	0.7857	0.8125	0.8311

Table 4.1 Comparison of known provable lower bounds for $\gamma_{n,k}$ when n is large and $k \in \{1,\ldots,5\}$. The bounds of [10] hold for every n.

Application to the Stochastic Assignment Problem. Finally, in Section 4.5, as an application of our new provable lower bounds for (k,n)-PIP we provide the tight optimal approximation ratio for the classic stochastic sequential assignment problem (SSAP for short) by [26]. In the SSAP, there are n non-negative values $r_1,\ldots,r_n$ and n i.i.d. non-negative values $X_1,\ldots,X_n$ that are observed one at the time by a decision-maker. At each time t, the decision-maker must assign irrevocably X_t to one of the remaining available r_i values that have not been assigned yet. Assigning X_t to r_i provides a reward of $r_i \cdot X_t$, and the goal is to maximize the expected sum of rewards. This problem extends (k,n)-PIP and relates to several online matching problems [39, 63].

We revisit the SSAP through the lens of prophet inequalities and provide an exact value of its asymptotic approximation ratio. Specifically, we first characterize the optimal approximation ratio for SSAP to be equal to $\alpha_n = \min_{k\in[n]} \gamma_{n,k}$. This immediately implies that $\limsup_n \alpha_n \leq \gamma \approx 0.745$. To the best of the authors' knowledge, the best current provable lower bounds over $\gamma_{n,k}$ are the following: (1) for $k = 1$,

$\gamma_{n,1} \geq 0.745$ [23]; (2) for any $k \geq 1$, $\gamma_{n,k} \geq 1 - k^k e^{-k}/k!$ (see, e.g., [10, 27]); (3) the values reported by [10] in Table 4.1. These three results together imply in principle that $\alpha_n \geq 0.7427$; hence, there is a constant gap between the lower and the upper bound on α_n. The $1 - k^k e^{-k}/k!$ lower bound for the k-selection prophet inequality problem is at least 0.78 for $k \geq 3$ which is in particular larger than $\liminf_n \gamma_{n,1} \approx 0.745$. Nevertheless, to the best of the authors' knowledge, no monotonicity in k is known for the values $\gamma_{n,k}$. Therefore, our new provable 0.829 lower bound for $k = 2$ allows us to conclude that the approximation ratio for the SSAP is exactly $\gamma \approx 0.745$ for n sufficiently large, fully characterizing the approximation ratio of the problem.

4.1.2 Related work

The basic prophet inequality problem, as introduced by Krengel and Sucheston [57], was resolved by using a dynamic program that gave a tight approximation ratio of 1/2. Samuel-Cahn [73] later showed that a simple threshold algorithm yields the same guarantee. Since then, there have been several generalizations spanning combinatorial constraints, different valuation functions and arrival orders, resource augmentation, and limited knowledge of the distributions [21, 22, 31, 55].

A major reason for the renewed interest in prophet inequalities is their relevance to auctions, specifically posted priced mechanisms (PPM) in online sales [3, 18, 20, 27, 42, 55]. It was implicitly shown by [18] and [42] that every prophet-type inequality implies a corresponding approximation guarantee in a PPM, and the converse holds as well [24]. Using these well-known reductions, our lower bounds for the i.i.d. k-selection prophet inequality problem also yield PPM's for the problem of selling k homogeneous goods to n unit-demand buyers who arrive sequentially with independent and identically distributed valuations.

Linear and convex programming have been a powerful tool for the design of online algorithms. For instance, in online and Bayesian matching problems [39, 63], online knapsack [9, 53], secretary problem [14, 17, 68], factor-revealing linear programs [37, 59], and competition complexity [12]. Similar to us, Perez-Salazar et al. [67] use a quantile-based linear programming formulation to provide optimal policies in the context of decision-makers with a limited number of actions.

Our analysis provides a new nonlinear system of differential equations, which extends the ordinary differential equation by Hill and Kertz for $k = 1$, and provides provable lower bounds on the asymptotic approximation ratio. Although the Hill and Kertz equation has been used in various recent works [12, 23, 60, 67] and ODE methods, have been used in other online selection problems [6, 32, 68] to provide asymptotic guarantees, to the best of the authors' knowledge, our result for *multiple* *selections*, where the approximation ratio is embedded in a nonlinear system, has not been previously explored.

4.2 Preliminaries

An instance of (k,n)-PIP is given by a tuple (n,k,F), where n is the number of values $X_1,\dots,X_n$ that are drawn i.i.d according to the continuous distribution F supported on $\mathbb{R}_+$. This assumption is commonly made in the literature (e.g., Liu et al. [60]), as we can perturb a discrete distribution by introducing random noise at the cost of a negligible loss in the objective. Given an instance of (k,n)-PIP, observe that we can always scale the values $X_1,\dots,X_n$ by a positive factor so the optimal value is equal to 1, and the reward of the optimal policy is scaled by the same amount. In particular, the approximation ratio of the optimal policy remains the same.

Given an instance (n,k,F), we use dynamic programming to compute the optimal reward of the optimal sequential policy. Let $A_{t,\ell}(F)$ be the reward of the optimal policy when $\ell \le k$ choices are still to be made in periods $\{t,\dots,n\}$. Then, for every $t \in [n]$ and $\ell \in [k]$, the following holds:

$$A_{t,\ell}(F) = \sup_{x\ge 0}\Big\{(\mathbb{E}[X \mid X \ge x] + A_{t+1,\ell-1}(F))\Pr[X \ge x] + A_{t+1,\ell}(F)\Pr[X < x]\Big\}, \tag{4.5}$$

$$A_{n+1,\ell}(F) = 0, \text{ and } A_{t,0}(F) = 0. \tag{4.6}$$

Equation (4.5) corresponds to the continuation value condition in optimality; the term in the braces is the expected value obtained when a threshold x is chosen when at period t and ℓ choices can still be made. In (4.6) we have the border conditions. In particular, it holds

$$\gamma_{n,k} = \inf\Big\{A_{1,k}(F) : \text{instances } (n,k,F) \text{ with } \mathrm{OPT}_{n,k}(F) = 1\Big\}. \tag{4.7}$$

4.3 An infinite-dimensional formulation

In this section, we provide the characterization of the optimal approximation ratio for (k,n)-PIP, $\gamma_{n,k}$, via an infinite-dimensional linear program. For every positive integers

k and n, with $n \geq k$, consider the following infinite dimensional linear program:

$$\begin{aligned}
&\inf \quad d_{1,k} && [\mathrm{P}]_{n,k} \\
&\text{s.t.} \quad d_{t,\ell} \geq \int_0^q h(u)\,\mathrm{d}u + q d_{t+1,\ell-1} + (1-q) d_{t+1,\ell}, \text{ for every } t \in [n], \ell \in [k], \text{ and } q \in [0,1], && (4.8) \\
&\int_0^1 g_{n,k}(u) h(u)\,\mathrm{d}u \geq 1, && (4.9) \\
&h(u) \geq h(w) \quad \text{for every } u \leq w, \text{with } u,w \in [0,1], && (4.10) \\
&d_{t,\ell} \geq 0 \quad \text{for every } t \in [n+1] \text{ and every } \ell \in \{0\} \cup [k], && (4.11) \\
&h(u) \geq 0 \quad \text{for every } u \in [0,1] \text{ and } h \text{ is continuous in } (0,1), && (4.12)
\end{aligned}$$

where $g_{n,k}(u) = \sum_{j=n-k+1}^{n} j\binom{n}{j}(1-u)^{j-1}u^{n-j}$ for every $u \in [0,1]$. $d_{t,\ell}$ represents that the welfare of the optimal dynamic policy after having seen $t \leq n$ values and $\ell \leq k$ of them remain to be selected. The family of constraints (4.8) asserts that for any quantile q, $d_{t,\ell}$ is at least the welfare obtained by using q as a threshold for selection in the current round, t. That is, the first two terms of the right hand side cover the case when the value is indeed selected, whereas the third term covers the case when it is not. Constraint (4.9) represents a normalization to instances where the optimal welfare $\sum_{t=n-k+1}^{n} \mathbb{E}[X_{(i)}] \geq 1$. In the program $[\mathrm{P}]_{n,k}$, the variables $h(u)$ represent the values of a non-negative and non-increasing function h in $[0,1]$, and therefore we have infinitely many of them, whereas the variables d are finitely many. More specifically, $h(u)$ represents $F^{-1}(1-u)$, where F is the c.d.f of a probability distribution and minimizing over all possible choices of h is equivalent to search the worst-case distribution for (k,n)-PIP. We remark the continuity for h in the program is mainly for the sake of simplicity in our analysis but does not represent a strict requirement.

The following structural result formalizes the interpretation of the variables provided in the previous paragraph. The remainder of the section is dedicated to its proof.

Theorem 9. *The optimal approximation ratio for* (k,n)-PIP *is equal to the optimal value of* $[\mathrm{P}]_{n,k}$.

We denote by $v_{n,k}$ the optimal value of $[\mathrm{P}]_{n,k}$. We show that $v_{n,k} = \gamma_{n,k}$ in Theorem 9 by proving both inequalities, $v_{n,k} \leq \gamma_{n,k}$ and $v_{n,k} \geq \gamma_{n,k}$, separately. For the first inequality, we argue that any instance of (k,n)-PIP with $\mathrm{OPT}_{n,k}(F) = 1$, produces a feasible solution to $[\mathrm{P}]_{n,k}$ with an objective value equal to the reward of the optimal sequential policy. For the second, we show that any feasible solution (d,f) to $[\mathrm{P}]_{n,k}$ produces an instance of (k,n)-PIP such that the reward of the optimal policy is no larger than the objective value of the instance (d,f).

Before we prove the inequalities, we leave a proposition with some preliminary properties that we use in our analysis. The proof can be found in Appendix C.1.

Proposition 11. *Let F be a continuous and strictly increasing distribution over the non-negative reals. Then, the following properties hold:*

(i) *For every n and k with $n \geq k$, we have* $\mathrm{OPT}_{n,k}(F) = \int_0^1 g_{n,k}(u) F^{-1}(1-u)\,\mathrm{d}u$.

(ii) *Suppose that X is a random variable distributed according to F. Then, for every $x \geq 0$, it holds* $\mathbb{E}[X \mid X \geq x] \Pr[X \geq x] = \int_0^q F^{-1}(1-u)\,\mathrm{d}u$, *where* $q = \Pr[X \geq x]$.

We use the following two lemmas to prove Theorem 9.

Lemma 9. *Let F be a continuous and strictly increasing distribution over the non-negative reals such that* $\mathrm{OPT}_{n,k}(F) = 1$, *and let $h(u) = F^{-1}(1-u)$ for every $u \in [0,1]$. Then, $(A(F), h)$ is feasible for* $[\mathrm{P}]_{n,k}$, *where $A(F) = (A_{t,\ell}(F))_{t,\ell}$ is defined according to* (4.5)-(4.6).

Proof. By construction, we have $h \geq 0$ and h is non-increasing, therefore constraints (4.10)-(4.12) are satisfied by $(A(F), h)$. Furthermore, we have

$$\int_0^1 g_{n,k}(u) h(u)\,\mathrm{d}u = \int_0^1 g_{n,k}(u) F^{-1}(1-u)\,\mathrm{d}u = \mathrm{OPT}_{n,k}(F) = 1,$$

where the second equality holds by Proposition 11(i). Then, constraint (4.9) is satisfied by $(A(F), h)$. Let $q \in [0,1]$, and $x \geq 0$ such that $q = \Pr[X \geq x]$. Then, for every $t \in [n]$ and every $\ell \in [k]$, we have

$$\begin{aligned} A_{t,\ell}(F) &\geq (\mathbb{E}[X \mid X \geq x] + A_{t+1,\ell-1}(F)) \Pr[X \geq x] + A_{t+1,\ell}(F) \Pr[X < x] \\ &= \int_0^q h(u)\,\mathrm{d}u + q A_{t+1,\ell-1}(F) + (1-q) A_{t+1,\ell}(F), \end{aligned}$$

where in the first inequality we used condition (4.5), while in the second inequality, we used that $q = \Pr[X \geq x]$ and Proposition 11(ii). Then, $(A(F), h)$ satisfies constraint (4.8), and we conclude that $(A(F), h)$ is feasible for $[\mathrm{P}]_{n,k}$. ☐

Lemma 10. *Let (d, h) be any feasible solution for $[\mathrm{P}]_{n,k}$. Then, there exists a probability distribution G, such that $d_{t,\ell} \geq A_{t,\ell}(G)$ for every $t \in [n]$ and $\ell \in [k]$.*

Proof. Given a feasible solution (d, h) for $[\mathrm{P}]_{n,k}$, consider the random variable $h(1-Q)$, where Q is a uniform random variable over the interval $[0,1]$, and let G be the probability distribution of $h(1-Q)$. Then, G is continuous and non-decreasing. Since h is a non-increasing function, we have $\Pr[h(1-Q) \leq h(u)] \geq \Pr[1-Q \geq u] = \Pr[1-u \geq Q] = 1-u$, and therefore $G^{-1}(1-u) \leq h(u)$. We prove the inequalities in the lemma statement for d and the probability distribution G via backward induction in $t \in \{1, \ldots, n+1\}$. For $t = n+1$ we have $d_{n+1,\ell} = 0 = A_{n+1,\ell}(G)$ for any $\ell \in [k]$.

Assume the result holds true for $\{t+1,\dots,n+1\}$. If $\ell = 0$, we have $d_{t,\ell} = 0 = A_{t,\ell}(G)$, and for $\ell \in [k]$, we have

$$
\begin{aligned}
d_{t,\ell} &\geq \sup_{q\in[0,1]} \left\{ \int_0^q h(u)\,\mathrm{d}u + q d_{t+1,\ell-1} + (1-q) d_{t+1,\ell} \right\} \\
&\geq \sup_{q\in[0,1]} \left\{ \int_0^q h(u)\,\mathrm{d}u + q A_{t+1,\ell-1}(G) + (1-q) A_{t+1,\ell}(G) \right\} \\
&\geq \sup_{q\in[0,1]} \left\{ \int_0^q G^{-1}(1-u)\,\mathrm{d}u + q A_{t+1,\ell-1}(G) + (1-q) A_{t+1,\ell}(G) \right\} \\
&= \max_{x\geq 0} \{ (\mathbb{E}[X \mid X \geq x] + A_{t+1,\ell-1}(G)) \Pr[X \geq x] + A_{t+1,\ell}(G) \Pr[X < x] \} \\
&= A_{t,\ell}(G),
\end{aligned}
$$

where the first inequality holds since (d,h) satisfies constraint (4.8); the second inequality holds by induction; the third holds since $G^{-1}(1-u) \leq h(u)$, and the first equality by Proposition 11(ii). This concludes the proof of the lemma. □

Proof of Theorem 9. By Lemma 9, for every probability distribution F we have that $(A(F),h)$ is feasible for $[\mathrm{P}]_{n,k}$, and its objective value is equal to $A_{1,k}(F)$. This implies that $A_{1,k}(F) \geq v_{n,k}$ for every F, and therefore $\gamma_{n,k} \geq v_{n,k}$. Since in $[\mathrm{P}]_{n,k}$, h only appears in integrals over some interval, we observe that in fact that the value of $[\mathrm{P}]_{n,k}$ does not change if we restrict the condition on h to be strictly increasing instead of non-decreasing. In this case, by Lemma 10, for every feasible solution (d,h) in $[\mathrm{P}]_{n,k}$ there exists a probability distribution G such that $d_{1,k} \geq A_{1,k}(G)$ and $G^{-1}(1-u) = h(u)$. Thus $\mathrm{OPT}_{n,k}(G) \geq 1$. Therefore, the optimal value of $[\mathrm{P}]_{n,k}$ is lower bounded by the infimum in (4.7), which is equal to $\gamma_{n,k}$. We conclude that $v_{n,k} \geq \gamma_{n,k}$, and therefore both values are equal. □

4.4 Lower bound on the approximation ratio

In this section, we prove the following result.

Theorem 10. *For every $k \geq 1$, there exists $n_0 \in \mathbb{N}$, such that for every $n \geq n_0$ we have*

$$
\gamma_{n,k} \geq \left(1 - 24k\frac{\ln(n)^2}{n}\right) \sum_{j=1}^{k} \theta_j^\star,
$$

where $\theta_1^\star,\dots,\theta_k^\star$ are the values for which there exists a solution to the nonlinear system of differential equations (4.1)-(4.3).

Before proving this theorem, we provide a warm-up for the case of $k = 2$, i.e., when 2 selections are possible. The goal of Section 4.4.1 is to provide the main insights into the derivation of the nonlinear system (4.1)-(4.3) and to sktech the main steps in the proof of Theorem 10. In Section 4.4.2 we provide the full detailed proof of Theorem 10 which holds for every k.

4.4.1 Warm-up: the case of $k = 2$

In this subsection, we sketch the deduction of the nonlinear system (4.1)-(4.3). We focus on the case $k = 2$ and provide a weak dual formulation of $[\mathrm{P}]_{n,k}$. From this weak dual, we can find a recursion that in the limit converges to a solution of (4.1)-(4.3) with $k = 2$. While this approach only gives us a lower bound on $\liminf_n \gamma_{n,2}$, we strengthen this asymptotic result by showing how to transform a solution of the system (4.1)-(4.3) into a feasible solution to our dual LP by incurring in a slight loss when n is large.

A Dual for $k = 2$. Consider the following infinite-dimensional linear program:

$$
\begin{aligned}
&\sup \quad v && [\mathrm{D}]_{n,2}\\
\text{s.t.}\quad &\int_0^1 \beta_{1,\ell}(q)\,\mathrm{d}q \le \mathbf{1}_2(\ell), \quad \text{for all } \ell \in [2], && (4.13)\\
&\int_0^1 \beta_{t+1,2}(q)\,\mathrm{d}q \le \int_0^1 (1-q)\beta_{t,2}(q)\,\mathrm{d}q, \quad \text{for all } t \in [n-1], && (4.14)\\
&\int_0^1 \beta_{t+1,1}(q)\,\mathrm{d}q \le \int_0^1 (1-q)\beta_{t,1}(q)\,\mathrm{d}q + \int_0^1 q\beta_{t,2}(q)\,\mathrm{d}q, \ \text{for all } t \in [n-1], && (4.15)\\
&v g_{n,2}(u) \le \sum_{t=1}^{n} \int_u^1 \beta_{t,1}(q) + \beta_{t,2}(q)\,\mathrm{d}q, \ \text{for } u \in [0,1], && (4.16)\\
&\beta_{t,\ell}(q) \ge 0 \quad \text{for all } q \in [0,1], t \in [n] \text{ and } \ell \in [2], && (4.17)
\end{aligned}
$$

where $\mathbf{1}_2(\ell) \in \{0,1\}$ and $\mathbf{1}_2(\ell) = 1$ if and only if $\ell = 2$. This linear program can be interpreted as follows. Fix an algorithm that makes decisions based on quantiles. Now, the variables $\beta_{t,\ell}(q)$ can be interpreted as the probability densities of the events that *the algorithm chooses quantile q for time t when ℓ items remain to be chosen, observes the t-th value x_t, and selects it if $x_t \ge F^{-1}(1-q)$*. Variable $v \ge 0$ captures the approximation ratio of such algorithm. Constraints (4.13)-(4.15) capture the valid transitions from time t to $t+1$. Specifically, for the algorithm to observe a value at time t when ℓ items can still be chosen, it must have observed a value at time $t-1$ under one of the following conditions: Either (1) there are ℓ items that can be chosen at $t-1$ but the algorithm did not select the $t-1$ observed value; or (2) there are $\ell+1$ items that can be chosen at $t-1$ and the algorithm selected the $t-1$ observed value. Constraint (4.16) relates the offline density of least two out of n values being in the top u quantiles with the density that the algorithm selects values in the same quantile. Among all possible quantiles $u \in [0,1]$, the largest v such that the ratio of these densities is at least v corresponds to a lower bound on the approximation ratio of the algorithm.

In Lemma 11, we show that $[\mathrm{D}]_{n,2}$ is a weak dual to the exact formulation $[\mathrm{P}]_{n,k}$, for $k = 2$. That is, the optimal value $w_{n,2}$ of $[\mathrm{D}]_{n,2}$ is at most $\gamma_{n,2}$, which is the optimal approximation ratio for 2 selections. Thus, finding solutions to the weak dual provide a mechanism to give provable lower bounds on $\gamma_{n,2}$. We remark that $[\mathrm{D}]_{n,2}$ is not necessarily a strong dual to $[\mathrm{P}]_{n,k}$. Variables $\beta_{t,\ell}$ correspond to constraints (4.8) and variable v corresponds to constraint (4.9); however, in $[\mathrm{D}]_{n,2}$ there are no dual variables for constraints (4.10).

A Feasible Solution and Its Limit. We now construct a particular feasible solution to the weak dual $[\mathrm{D}]_{n,2}$ and show that it produces a set of points that converge to a solution of (4.1)-(4.3) for $k = 2$. Inspired by the quantile-based solution of [23] and the LP characterization by [67], we propose the following solution to $[\mathrm{D}]_{n,2}$:

$$\beta_{t,2}(q) = \theta_2 \cdot (-g'_{n,2}(q))\,\mathbf{1}_{(\varepsilon_{t-1},\varepsilon_t)}(q), \qquad \beta_{t,1}(q) = \theta_1 \cdot (-g'_{n,2}(q))\,\mathbf{1}_{(\mu_{t-1},\mu_t)}(q), \quad (4.18)$$

for $t = 1,\ldots,n$, where $\theta_1,\theta_2 \geq 0$ and $0 = \varepsilon_0 < \varepsilon_1 < \cdots < \varepsilon_n$ and $0 = \mu_0 = \mu_1 < \mu_2 < \cdots < \mu_n$. Note that if $\varepsilon_n = \mu_n = 1$, then $(\beta,\theta_1 + \theta_2)$ is a feasible solution to $[\mathrm{P}]_{n,k}$; hence, $\theta_1 + \theta_2 \leq \gamma_{n,2}$. Now, assuming that constraints (4.13)-(4.15) are tightened by β, we can deduce the following implicit recursions for ε_t and μ_t:

$$g_{n,2}(\varepsilon_{t+1}) - g_{n,2}(\varepsilon_t) = -\frac{1}{\theta_2} - \frac{2}{n+1}\left(g_{n+1,3}(\varepsilon_t) - (n+1)\right) \qquad t \in \{0\} \cup [n], \quad (4.19)$$

$$\begin{aligned} g_{n,2}(\mu_{t+1}) - g_{n,2}(\mu_t) = {} & \frac{\theta_2}{\theta_1}\frac{2}{n+1}\left(g_{n+1,3}(\epsilon_t) - (n+1)\right) \qquad t \in \{0\} \cup [n], \quad (4.20) \\ & - \frac{2}{n+1}\left(g_{n+1,3}(\mu_t) - (n+1)\right) \end{aligned}$$

In other words, by fixing $\theta_2 > 0$, we can find a sequence of increasing ε_t's. Likewise, by fixing $\theta_1,\theta_2 > 0$, we can find a sequence of increasing μ_t's. Our goal is to find a pair $(\theta_1^\star,\theta_2^\star)$ such that $\varepsilon_n(\theta_2) = 1$ and $\mu_n(\theta_1,\theta_2) = 1$. A simple inductive argument shows that ε_t is decreasing in θ_2. Indeed, by differentiating (4.19) in θ_2, we obtain $g'_{n,2}(\varepsilon_{t+1}(\theta_2))\varepsilon'_{t+1}(\theta_2) = 1/\theta_2^2 + \varepsilon'_t(\theta_2)(1-\varepsilon_t(\theta_2))g'_{n,2}(\varepsilon_t(\theta_2))$. From here, and also using (4.19), we can show that: (1) ε_t blows up when θ_1 goes to 0 (2) ε_t tends to 0 when θ_2 goes to $+\infty$ and (2) there this a unique $\theta_2^\star > 0$ such that $\varepsilon_n(\theta_2^\star) = 1$. Analogously, having already found and fixed $\theta_2^\star > 0$, we can show that there is a unique $\theta_1^\star > 0$ such that $\mu_n(\theta_1^\star,\theta_2^\star) = 1$. This shows that our solution $(\beta,\theta_1^\star + \theta_2^\star)$ is feasible to $[\mathrm{D}]_{n,2}$.

The pair $(\theta_1^\star, \theta_2^\star)$ depends non-trivially on n. Thus, to tackle this dependency, we perform a limit analysis on (4.19) and (4.20). For this, we first perform a change of variable $y_{1,t} = e^{-n\mu_t}$ and $y_{2,t} = e^{-n\varepsilon_t}$ for all t. We see that for all $\ell \in [2]$, $y_{\ell,t} \in [0,1]$, $y_{\ell,1} = 1$, $y_{\ell,n} = e^{-n}$ for all $\ell \in [2]$. Now, consider the function $\hat{y}_\ell : [0,1] \to [0,1]$ such that, $\hat{y}_\ell(0) = 1$, $\hat{y}_\ell(t/n) = y_{\ell,t}$, for all $t \in \{0\} \cup [n]$, and $\hat{y}_\ell$ is linear between in $[(t-1)/n, t/n]$ for $t \in [n]$. Now using (4.19) and (4.20), the approximation $g_{n,k}(u) \approx n\Gamma_k(nu)/(k-1)!$,[1] and assuming that $(\theta_1^\star, \theta_2^\star)$ converges when $n \to \infty$, we obtain that $(\hat{y}_1, \hat{y}_2)$ converges to a solution of the nonlinear system (4.1)-(4.3) with $k = 2$. A formal proof of the existence of a solution to the nonlinear system (4.1)-(4.3) for general k appears in Lemma 12.

From the Limit to a Feasible Solution. The limit analysis of the solution β only bounds $\liminf_n \gamma_{n,2}$. To provide an analysis for finite n as in Theorem 10, we construct an explicit solution to $[\mathrm{D}]_{n,2}$ from the solution (y_1, y_2) of (4.1)-(4.3). While we are tempted to take $\varepsilon_t = -\ln(y_2(t/n))/n$ and $\mu_t = -\ln(y_1(t/n))/n$ and defining β as in (4.18), this poses two nontrivial challenges. Firstly, ε_n and μ_n are larger than 1 since $y_1(1) = y_2(1) = 0$. Secondly, μ_1 is positive; yet for β to be feasible to $[\mathrm{D}]_{n,2}$, μ_1 must be 0, due to constraint (4.13).

To address the first challenge, we take $\bar{n} \le n$, and define $\varepsilon_t = -\ln(y_2(t/\bar{n}))/\bar{n}$ for $t \in \{0\} \cup [\bar{n}-1]$ and $\varepsilon_{\bar{n}} = 1$ for some $\bar{n} \le n$. Similarly, $\mu_0 = \mu_1 = 0$, $\mu_t = -\ln(y_1(t/\bar{n}))/\bar{n}$, for $t \in [\bar{n}-1] \setminus \{1\}$, and $\mu_{\bar{n}} = 1$. The value $\bar{n}$ appears when approximating $g'_{n,k}(u)$ with $\bar{n}\Gamma_k(\bar{n}u)'/(k-1)!$ (see Proposition 17); for $k = 2$ we have $\bar{n} = n - 3$. For n large enough, we show that $-\ln(y_\ell(1 - 1/\bar{n}))/\bar{n} < 1$ for $\ell \in [2]$; hence, our new sequence of ε's and μ's is well-defined.

For the second challenge, on a first read, it might seem that defining $\mu_1 = 0$ solves the problem. However, this is not the case because, in constraint (4.15) for $t = 2$, the left-hand side could be larger than the right-hand side as our choice of ε_t and μ_t mimics an Euler approximation of the nonlinear system (4.1)-(4.3); hence, μ_2 depends on μ_1 in the Euler approximation; however, we defined it to be 0. To address this, we add enough mass to $\beta_{1,2}$ so constraint (4.15) approximately holds. A complete description of the solution appears in (4.36).

In general, these two fixes introduce two sources of multiplicative loss in the objective value $\theta_1^\star + \theta_2^\star$. The first is via $\bar{n}$, and the second is via the mass addition to $\beta_{1,2}$. We show that these multiplicative losses are in the order of $1 - c\ln(n)^2/n$ for some constant c, and they vanish as n grows.

[1] in Proposition 17 we prove formally this approximation for the derivative of $g_{n,k}$.

4.4.2 Main Analysis

We now provide the proof of Theorem 10. Our proof is organized into three main steps. In the first step, we introduce a *maximization* infinite-dimensional linear program that we call $[\mathrm{D}]_{n,k}$, and we prove that weak duality holds for the pair $[\mathrm{P}]_{n,k}$ and $[\mathrm{D}]_{n,k}$. Namely, the optimal value of $[\mathrm{D}]_{n,k}$ provides a lower bound on the optimal value of $[\mathrm{P}]_{n,k}$. The program $[\mathrm{D}]_{n,k}$ can be formally deduced from $[\mathrm{P}]_{n,k}$ by dropping Constraint (4.10); however, the weak duality still requires a proof as we are dealing with infinite-dimensional programs.

In the second step, we introduce a second maximization infinite-dimensional linear program parametrized by a value $\bar{n} \le n$, namely $[\mathrm{D}]_{n,k}(\bar{n})$. This program is akin to $[\mathrm{D}]_{n,k}$, but is described by a set of constraints that become more handy when analyzing the nonlinear system. Furthermore, $[\mathrm{D}]_{n,k}(\bar{n})$ restricts the time horizon until $\bar{n}$. We show that the optimal value of $[\mathrm{D}]_{n,k}(\bar{n})$ provides a lower bound on the program $[\mathrm{D}]_{n,k}$, and therefore, it gives a lower bound on the optimal value of $[\mathrm{P}]_{n,k}$ as well.

In the third step, we build an explicit feasible solution to the problem $[\mathrm{D}]_{n,k}(\bar{n})$ starting from a solution of the nonlinear system of differential equations (4.1)-(4.3). Using the valid bounds found in the previous two steps, we can provide a lower bound on the value $\gamma_{n,k}$ for n large enough. In particular, we show, as n grows, that the sequence of lower bounds provides a lower bound on the optimal asymptotic approximation factor.

First step: Weak duality. For every $\ell \in [k]$, let $\mathbf{1}_k(\ell) = 1$ if $\ell = k$ and $\mathbf{1}_k(\ell) = 0$ for $\ell \neq k$. Consider the following infinite-dimensional linear program:

$$\sup \quad v \qquad [\mathrm{D}]_{n,k}$$

$$\text{s.t.} \quad \int_0^1 \beta_{1,\ell}(q)\,\mathrm{d}q \le \mathbf{1}_k(\ell), \quad \text{for all } \ell \in [k], \tag{4.21}$$

$$\int_0^1 \beta_{t+1,k}(q)\,\mathrm{d}q \le \int_0^1 (1-q)\beta_{t,k}(q)\,\mathrm{d}q, \quad \text{for all } t \in [n-1], \tag{4.22}$$

$$\int_0^1 \beta_{t+1,\ell}(q)\,\mathrm{d}q \le \int_0^1 (1-q)\beta_{t,\ell}(q)\,\mathrm{d}q + \int_0^1 q\beta_{t,\ell+1}(q)\,\mathrm{d}q, \text{ for all } t \in [n-1], \ell \in [k-1], \tag{4.23}$$

$$v g_{n,k}(u) \le \sum_{t=1}^{n}\sum_{\ell=1}^{k}\int_u^1 \beta_{t,\ell}(q)\,\mathrm{d}q, \text{ for } u \in [0,1], \tag{4.24}$$

$$\beta_{t,\ell}(q) \ge 0 \quad \text{for all } q \in [0,1], t \in [n] \text{ and } \ell \in [k]. \tag{4.25}$$

The variables $\beta_{t,\ell}(q)$ represent the probability density of an optimal algorithm choosing quantile q at time t when ℓ items can still be chosen, and variable v captures the approximation factor of the policy. (See also the interpretation of the linear program in subsection 4.4.1.) We denote by $w_{n,k}$ the optimal value of $[\mathrm{D}]_{n,k}$. The following lemma shows that weak duality holds for the pair of infinite-dimensional programs $[\mathrm{P}]_{n,k}$ and $[\mathrm{D}]_{n,k}$.

Lemma 11. *For every $n \geq 1$ and every $k \in \{1,\ldots,n\}$, we have $v_{n,k} \geq w_{n,k}$.*

Proof. Consider a feasible solution (d,h) for $[\mathrm{P}]_{n,k}$ and a feasible solution (β,v) for $[\mathrm{D}]_{n,k}$. Since $[\mathrm{P}]_{n,k}$ is a minimization problem, we can assume that $d_{n+1,\ell} = 0$ for every $\ell \in [k]$ and $d_{0,\ell} = 0$ for every $t \in [n]$; if they are non-zero, we can easily make them zero without changing the objective value of (d,h). In what follows, we show that $v \leq d_{1,k}$. Since (β,v) satisfies constraint (4.24) and $h(u) \geq 0$ by constraint (4.12), we get

$$\begin{aligned} v\int_0^1 g_{n,k}(u)h(u)\,\mathrm{d}u &\leq \int_0^1 \sum_{t=1}^{n}\sum_{\ell=1}^{k}\left(\int_u^1 \beta_{t,\ell}(q)\,\mathrm{d}q\right)h(u)\,\mathrm{d}u \\ &= \sum_{t=1}^{n}\sum_{\ell=1}^{k}\int_0^1\int_u^1 \beta_{t,\ell}(q)h(u)\,\mathrm{d}q\,\mathrm{d}u \\ &= \sum_{t=1}^{n}\sum_{\ell=1}^{k}\int_0^1 \beta_{t,\ell}(q)\int_0^q h(u)\,\mathrm{d}u\,\mathrm{d}q, \end{aligned} \tag{4.26}$$

where the first equality holds by exchanging the summation and the integrals, and the second equality holds by exchanging the integration order for u and q. Then, inequality (4.26) together with constraint (4.9) imply that

$$v \leq \sum_{t=1}^{n}\sum_{\ell=1}^{k}\int_0^1 \beta_{t,\ell}(q)\int_0^q h(u)\,\mathrm{d}u\,\mathrm{d}q. \tag{4.27}$$

On the other hand, from constraint (4.8), for every $t \in [n]$ and every $\ell \in [k]$ we have

$$\begin{aligned} &\int_0^1 \beta_{t,\ell}(q)\int_0^q h(u)\,\mathrm{d}u\,\mathrm{d}q \\ &\leq d_{t,\ell}\int_0^1 \beta_{t,\ell}(q)\,\mathrm{d}q - d_{t+1,\ell-1}\int_0^1 q\beta_{t,\ell}(q)\,\mathrm{d}q - d_{t+1,\ell}\int_0^1 (1-q)\beta_{t,\ell}(q)\,\mathrm{d}q, \end{aligned} \tag{4.28}$$

where we used that $\beta_{t,\ell} \geq 0$, and then we integrated over $q \in [0,1]$. When $\ell = k$, note that

$$
\begin{aligned}
&\sum_{t=1}^{n} d_{t,k} \int_0^1 \beta_{t,k}(q)\, \mathrm{d}q - \sum_{t=1}^{n} d_{t+1,k} \int_0^1 (1-q)\beta_{t,k}(q)\, \mathrm{d}q \\
&= d_{1,k} \int_0^1 \beta_{1,k}(q)\, \mathrm{d}q + \sum_{t=2}^{n} d_{t,k} \int_0^1 \beta_{t,k}(q)\, \mathrm{d}q - \sum_{t=1}^{n} d_{t+1,k} \int_0^1 (1-q)\beta_{t,k}(q)\, \mathrm{d}q \\
&\leq d_{1,k} + \sum_{t=1}^{n-1} d_{t+1,k} \int_0^1 \beta_{t+1,k}(q)\, \mathrm{d}q - \sum_{t=1}^{n} d_{t+1,k} \int_0^1 (1-q)\beta_{t,k}(q)\, \mathrm{d}q \\
&\leq d_{1,k} + \sum_{t=1}^{n-1} d_{t+1,k} \int_0^1 (1-q)\beta_{t,k}(q)\, \mathrm{d}q - \sum_{t=1}^{n-1} d_{t+1,k} \int_0^1 (1-q)\beta_{t,k}(q)\, \mathrm{d}q \leq d_{1,k}, \qquad (4.29)
\end{aligned}
$$

where the first inequality holds by constraint (4.21) and by changing the index range in the first summation, and the second inequality holds by inequality (4.22) and the fact that $d_{n+1,k} = 0$. On the other hand, note that

$$
\begin{aligned}
&\sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t,\ell} \int_0^1 \beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k} d_{t+1,\ell-1} \int_0^1 q\beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t+1,\ell} \int_0^1 (1-q)\beta_{t,\ell}(q)\, \mathrm{d}q \\
&= \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t,\ell} \int_0^1 \beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t+1,\ell} \int_0^1 q\beta_{t,\ell+1}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t+1,\ell} \int_0^1 (1-q)\beta_{t,\ell}(q)\, \mathrm{d}q \\
&= \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t,\ell} \int_0^1 \beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n-1} \sum_{\ell=1}^{k-1} d_{t+1,\ell} \left(\int_0^1 q\beta_{t,\ell+1}(q)\, \mathrm{d}q + \int_0^1 (1-q)\beta_{t,\ell}(q)\, \mathrm{d}q \right) \\
&\leq \sum_{t=1}^{n} \sum_{\ell=1}^{k-1} d_{t,\ell} \int_0^1 \beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n-1} \sum_{\ell=1}^{k-1} d_{t+1,\ell} \int_0^1 \beta_{t+1,\ell}(q)\, \mathrm{d}q \\
&= \sum_{\ell=1}^{k-1} d_{1,\ell} \int_0^1 \beta_{1,\ell}(q)\, \mathrm{d}q \leq \sum_{\ell=1}^{k-1} d_{1,\ell} \mathbf{1}_k(\ell) = 0, \qquad (4.30)
\end{aligned}
$$

where the first equality holds by changing the index range of the second summation and $d_{t+1,0} = 0$ for every $t \in [n]$, the second equality holds by factoring the summations and $d_{n+1,\ell} = 0$ for every $\ell \in [k-1]$, the first inequality holds by inequality (4.23), and the last inequality by constraint (4.21).

Then, by summing over $t \in [n]$ and $\ell \in [k]$ in inequality (4.28), we get

$$
\begin{aligned}
&\sum_{t=1}^{n} \sum_{\ell=1}^{k} \int_0^1 \beta_{t,\ell}(q) \int_0^q h(u)\, \mathrm{d}u\, \mathrm{d}q \\
&\leq \sum_{t=1}^{n} \sum_{\ell=1}^{k} d_{t,\ell} \int_0^1 \beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k} d_{t+1,\ell-1} \int_0^1 q\beta_{t,\ell}(q)\, \mathrm{d}q - \sum_{t=1}^{n} \sum_{\ell=1}^{k} d_{t+1,\ell} \int_0^1 (1-q)\beta_{t,\ell}(q)\, \mathrm{d}q, \\
&\leq d_{1,k}, \qquad (4.31)
\end{aligned}
$$

where the second inequality comes from (4.29) and (4.30) together. Finally, (4.27) and (4.31) imply that $v \le d_{1,k}$, which concludes the proof of the lemma. □

Second step: A truncated LP with a useful structure. Consider the following infinite-dimensional linear program: For $\overline{n} \le n$, we consider the following LP

$$\sup \quad v \qquad [\mathrm{D}]_{n,k}(\bar{n})$$

$$\text{s.t.} \quad \int_0^1 \alpha_{t,k}(q)\,\mathrm{d}q + \int_0^1 \sum_{\tau<t} q\alpha_{\tau,k}(q)\,\mathrm{d}q \le 1, \quad \text{for all } t \in [\bar{n}], \tag{4.32}$$

$$\int_0^1 \alpha_{t,\ell}(q)\,\mathrm{d}q + \int_0^1 \sum_{\tau<t} q\alpha_{\tau,\ell}(q)\,\mathrm{d}q \le \int_0^1 \sum_{\tau<t} q\alpha_{\tau,\ell+1}(q)\,\mathrm{d}q, \quad \text{for all } t \in [\bar{n}], \ell \in [k-1], \tag{4.33}$$

$$v g_{n,k}(u) \le \sum_{t=1}^{\bar{n}_k} \sum_{\ell=1}^{k} \int_u^1 \alpha_{t,\ell}(q)\,\mathrm{d}q, \text{ for } u \in [0,1], \tag{4.34}$$

$$\alpha_{t,\ell}(q) \ge 0 \quad \text{for all } q \in [0,1], t \in [\bar{n}] \text{ and } \ell \in [k], \tag{4.35}$$

We will prove that the optimal value of $[\mathrm{D}]_{n,k}(\bar{n})$ is a lower bound to the optimal value of $[\mathrm{D}]_{n,k}$. The following technical proposition allows us to ensure that any feasible solution to $[\mathrm{D}]_{n,k}(\bar{n})$ induces a feasible solution to $[\mathrm{D}]_{n,k}$. We present the proof of the proposition in Appendix C.2.

Proposition 12. *For every feasible solution* (α, v) *to* $[\mathrm{D}]_{n,k}(\bar{n})$*, there is* (α', v) *feasible to* $[\mathrm{D}]_{n,k}(\bar{n})$ *for which all constraints* (4.32) *and* (4.33) *are tightened.*

The following proposition states the lower bound we need in the rest of our analysis.

Proposition 13. *For every* $k < n$*, and every* $\bar{n} \le n$*, the optimal value of* $[\mathrm{D}]_{n,k}(\bar{n})$ *is at most the optimal value of* $[\mathrm{D}]_{n,k}$*.*

Proof. Let (α, v) be a feasible solution to $[\mathrm{D}]_{n,k}(\bar{n})$. Moreover, define $\alpha_{0,\ell} = 0\ \forall \ell \in [k]$. By the previous proposition, we can assume that α tightens all Constraints (4.32) and (4.33). From here, we can deduce that for $0 \le t < \bar{n}_k$

$$\int_0^1 \alpha_{t+1,k}(q)\,\mathrm{d}q = 1 - \sum_{\tau\le t} \int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q = \int_0^1 \alpha_{t,k}\,\mathrm{d}q - \int_0^1 q\alpha_{t,k}\,\mathrm{d}q.$$

The second equality is equivalent to

$$1 - \sum_{\tau<t} \int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q = \int_0^1 \alpha_{t,k}\,\mathrm{d}q.$$

where for $t = 0$ this is the empty sum. This equality holds by Proposition 12 and (4.32). For $\ell < k$, we have

$$\begin{aligned}\int_0^1 \alpha_{t+1,\ell}(q)\,\mathrm{d}q &= \sum_{\tau \le t} \int_0^1 q\alpha_{\tau,\ell+1}(q)\,\mathrm{d}q - \sum_{\tau \le t} \int_0^1 q\alpha_{\tau,\ell}(q)\,\mathrm{d}q \\ &= \int_0^1 \alpha_{t,\ell}(q)\,\mathrm{d}q + \int_0^1 q\alpha_{t,\ell+1}(q)\,\mathrm{d}q - \int_0^1 q\alpha_{t,\ell}(q)\,\mathrm{d}q \\ &= \int_0^1 (1-q)\alpha_{t,\ell}(q)\,\mathrm{d}q + \int_0^1 q\alpha_{t,\ell+1}(q)\,\mathrm{d}q.\end{aligned}$$

The second equality is equivalent to

$$\sum_{\tau < t} \int_0^1 q\alpha_{\tau,\ell+1}(q)\,\mathrm{d}q - \sum_{\tau < t} \int_0^1 q\alpha_{\tau,\ell}(q)\,\mathrm{d}q = \int_0^1 \alpha_{t,\ell}(q)\,\mathrm{d}q$$

where for $t = 0$ this is the empty sum. This equality holds by Proposition 12 and (4.33). Hence, α satisfies constraint (4.21) by the definition of $\alpha_{0,\ell}$ and the equalities above. (4.22) (4.23) follow from the equalities as well, for $t < \bar{n}_k$. If we define $\bar{\alpha}$ as follows

$$\bar{\alpha}_{t,\ell}(q) = \begin{cases} \alpha_{t,\ell}(q), & t \le \bar{n}_k, \\ 0, & t > \bar{n}_k, \end{cases}$$

then, $(\bar{\alpha}, v)$ is a feasible solution to $[\mathrm{D}]_{n,k}$. From here, the result follows immediately. □

Third step: From the nonlinear system to LP. Since $\gamma_{n,k}$ is equal to the value of $[\mathrm{P}]_{n,k}$, which in turn is at least the value of $[\mathrm{D}]_{n,k}$, the previous result implies that we only need to provide a feasible solution to $[\mathrm{D}]_{n,k}(\bar{n})$ to provide a lower bound on $\gamma_{n,k}$. The latter will be defined by a solution to the non-linear system of equations (4.1)-(4.3). For a given θ, we denote it by $\mathrm{NLS}_k(\theta)$. The following lemma summarizes some properties of $\mathrm{NLS}_k(\theta)$ that we use in our analysis.

Lemma 12. *For every positive integer k, the following holds:*

(i) *There exists $\theta^\star$ for which $\mathrm{NLS}_k(\theta^\star)$ has a solution. We denote such a solution by $(Y_1, \dots, Y_k)$.*

(ii) *The vector $\theta^\star$ satisfies that $0 < \theta_1^\star < \theta_2^\star < \cdots < \theta_k^\star < 1/k$.*

(iii) *For every $j \in [k]$, the function Y_j is non-increasing.*

We defer the proof of Lemma 12 to Subsection 4.4.3. Let $\bar{n}_k = n - k - 1$ and let $y_{j,t} = Y_j(t/\bar{n}_k)$. Let us define $\varepsilon_{j,t} = -\ln(y_{j,t})/\bar{n}_k$, for $t \in \{0, 1, \dots, \bar{n}_k - 1\}$, $j \in \{1, \dots, k\}$, and $\varepsilon_{j,\bar{n}_k} = 1$. We can show that for n large enough, $-\ln(y_{j,t})/\bar{n}_k \le 1$ for $t \in \{0, \dots, \bar{n}_k - 1\}$

(see Proposition 18); hence, $0 \le \varepsilon_{1,j} \le \cdots \le \varepsilon_{j,\bar{n}_k} \le 1$. Let $B_\ell = (\ell - 1) \cdot (4c_k^k + c_k/k!)$ for $\ell \in \{1,\ldots,k\}$, where $c_k = 24k! \max\{\theta^\star_{\ell+1}/\theta^\star_\ell : \ell \in \{1,\ldots,k-1\}\}$. Now, consider the following family of functions:

$$\alpha^\star_{t,\ell}(q) = \begin{cases} 0, & t \le k-\ell, \\ \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)^{-(k-\ell+1)} \left(B_\ell \ln(\bar{n}_k) \mathbf{1}_{[0,1/\bar{n}_k]}(q) - \theta^\star_\ell g'_{n,k}(q) \mathbf{1}_{(0,\varepsilon_{\ell,t})}(q)\right), & t = k-\ell+1, \\ \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)^{-(k-\ell+1)} \left(-\theta^\star_\ell g'_{n,k}(q)\right) \mathbf{1}_{(\varepsilon_{\ell,t-1},\varepsilon_{\ell,t})}(q), & t \ge k-\ell+2. \end{cases} \tag{4.36}$$

Note that for all $u \in [0,1]$, we have

$$\begin{aligned} \sum_{\ell=1}^{k} \int_u^1 \alpha^\star_{t,\ell}(q)\,\mathrm{d}q &\ge \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)^{-k} \left(\sum_{\ell=1}^{k} \theta^\star_\ell\right) g_{n,k}(u) \\ &\ge \left(1 - 12k\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \left(\sum_{\ell=1}^{k} \theta^\star_\ell\right) g_{n,k}(u), \end{aligned} \tag{4.37}$$

where in the first inequality we used that $(1 + 12\ln(\bar{n}_k)^2/\bar{n}_k)^{-(k-\ell+1)}$ is increasing in ℓ and in the second inequality we used the standard Bernoulli inequality. Inequality (4.37) guarantees that $(\alpha^\star, v^\star)$ satisfies constraint (4.34) with $v^\star = \left(1 - 12k \cdot \ln(\bar{n}_k)^2/\bar{n}_k\right) \sum_{\ell=1}^k \theta^\star_\ell$. Before proving Theorem 10 we need the following lemma; we defer the proof to section 4.4.4.

Lemma 13. *For $\bar{n}_k = n - k - 1$, and n large enough, $\alpha^\star$ satisfies constraints* (4.32) *and* (4.33).

Proof of Theorem 10. As a consequence of Lemma 13, we have that $(\alpha^\star, v^\star)$ is a feasible solution to $[\mathrm{D}]_{n,k}(\bar{n})$. In particular, we obtain the approximation

$$\gamma_{n,k} \ge v^\star = \left(1 - 12k\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \sum_{\ell=1}^{k} \theta^\star_\ell \ge \left(1 - 24k\frac{\ln(n)^2}{n}\right) \sum_{\ell=1}^{k} \theta^\star_\ell,$$

when n is sufficiently large. □

4.4.3 Analysis of $\mathrm{NLS}_k(\theta)$ and proof of lemma 12

In this section, we analyze the nonlinear system $\mathrm{NLS}_k(\theta)$ in terms of the existence of solutions. Given functions $y_1, \ldots, y_k : \mathbb{R} \to \mathbb{R}_+$, let $y = (y_1, \ldots, y_k)$, and for each pair $r, \ell \in [k]$, define $\phi_{r,\ell,y}(t) = \Gamma_r(-\ln y_\ell(t))$ for every $t \in [0,1)$. Observe that by simple differentiation, we have

$$\phi'_{r,\ell,y}(t) = -\Gamma'_r(-\ln y_\ell(t))\frac{y'_\ell(t)}{y_\ell(t)} = (-\ln y_\ell(t))^{r-1}y'_\ell(t), \tag{4.38}$$

since $\Gamma'_r(x) = -x^{r-1}e^{-x}$. Furthermore, when $r \geq 2$, observe that $\phi'_{r,\ell,y}(t) = -\phi'_{r-1,\ell,y}(t)\ln y_\ell(t)$, which is a consequence of the derivative formula in (4.38). For a vector $\theta_{\ell:k} = (\theta_\ell,\ldots,\theta_k)$, we define the system $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$ to be the subsystem of $\mathrm{NLS}_k(\theta)$ that only consider the differential equations from $\ell,\ldots,k$ and the terminal conditions, that is,

$$\begin{aligned}
&(\Gamma_k(-\ln y_k))' = k!\,(1-1/(k\theta_k)) - \Gamma_{k+1}(-\ln y_k),\\
&(\Gamma_k(-\ln y_j))' = k! - \Gamma_{k+1}(-\ln y_j) - \frac{\theta_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln y_{j+1})) \text{ for every } j \in \{\ell,\ldots,k-1\},\\
&\qquad y_j(0) = 1 \text{ and } \lim_{t\uparrow 1} y_j(t) = 0 \quad \text{for every } j \in \{\ell,\ldots,k\}.
\end{aligned}$$

When $\ell = 1$, the system $\mathrm{NLS}_{1,k}(\theta)$ is exactly the system $\mathrm{NLS}_k(\theta)$. We also remark that, by replacing, any solution y of $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$ satisfies the following conditions:

$$\begin{aligned}
\phi'_{k,k,y} &= k!\Big(1-\frac{1}{k\theta_k}\Big) - \phi_{k+1,k,y},\\
\phi'_{k,j,y} &= k! - \phi_{k+1,j,y} - \frac{\theta_{j+1}}{\theta_j}(k! - \phi_{k+1,j+1,y}) \text{ for every } j \in \{\ell,\ldots,k-1\}.
\end{aligned} \tag{4.39}$$

We will use $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$ to inductively show that $\mathrm{NLS}_k(\theta)$ satisfies all properties of Lemma 12. One key step to showing the existence of a solution to $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$ is to first establish some properties that any solution of $\mathrm{NLS}_{\ell+1,k}(\theta_{\ell+1:k})$, provided by the induction hypothesis, must satisfy. This helps since, a priori, we don't have a handle on such solutions and definitely no explicit form. Then we will use these properties to show the induction step. In Proposition 16 we show 12(ii) and 12(iii). Proposition 15 is useful for showing 12(iii) as it gives a simple sufficient criterion for monotonicity to hold. The following proposition gives an equivalent formulation of some expressions used to understand the behaviour of $\phi'_{k,j,y}$.

Proposition 14. *Consider $\ell \in [k-1]$, and let $\theta_{\ell,k}$ be such that there is a solution $y = (y_\ell,\ldots,y_k)$ for $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$, and such that $y'_j(s) \neq 0$ for every j and every $s \in (0,1)$. Then, the following holds:*

(i) $\phi'_{k,k,y}(t)\exp\Big(\int_t^1 \ln y_k(s)\,\mathrm{d}s\Big) = k!\left(1-\frac{1}{k\theta_k}\right).$

(ii) For every $j \in \{\ell, \ldots, k-1\}$, we have that $\phi'_{k,j,y}(t) \exp\Big(\int_t^1 \ln y_j(s)\, \mathrm{d}s\Big)$ is equal to

$$k!\left(1 - \frac{\theta_{j+1}}{\theta_j}\right) + \frac{\theta_{j+1}}{\theta_j} \int_t^1 \phi'_{k,j+1,y}(\tau) \ln y_{j+1}(\tau) \exp\Big(\int_\tau^1 \ln y_j(s)\, \mathrm{d}s\Big)\, \mathrm{d}\tau.$$

Proof. We start by observing the following: $\phi''_{k,k,y} = -\phi'_{k+1,k,y} = -(-\ln y_k)^k y'_k = \phi'_{k,k,y} \ln y_k$, where the first equality holds from the first identity in (4.39), and the other two equalities come from (4.38). From here, by integrating, we have that for every $t, r \in (0,1)$ with $r \geq t$, it holds

$$\begin{aligned}
\phi'_{k,k,y}(t) \exp\Big(\int_t^r \ln y_k(s)\, \mathrm{d}s\Big) &= \phi'_{k,k,y}(t) \exp\Big(\int_t^r \frac{\phi''_{k,k,y}(s)}{\phi'_{k,k,y}(s)}\, \mathrm{d}s\Big) \\
&= \phi'_{k,k,y}(t) \exp\Big(\ln \phi'_{k,k,y}(r) - \ln \phi'_{k,k,y}(t)\Big) = \phi'_{k,k,y}(r).
\end{aligned}$$

We conclude part (i) by doing $r \to 1$: We use that $y_k(r) \to 0$ in $\mathrm{NLS}_k(\theta)$, therefore $\phi_{k+1,k,y}(r) \to 0$, and then $\phi'_{k,k,y}(r) \to k!(1 - 1/(k\theta_k))$, using the first equality in (4.39).

For $j \in \{\ell, \ldots, k-1\}$, we proceed in a similar way. From the second equality in (4.39) we get

$$\phi''_{k,j,y} = -\phi'_{k+1,j,y} + \frac{\theta_{j+1}}{\theta_j} \phi'_{k+1,j+1,y} = \phi'_{k,j,y} \ln y_k - \frac{\theta_{j+1}}{\theta_j} \phi'_{k,j+1,y} \ln y_{j+1}, \tag{4.40}$$

where the last equality comes from the observation after the derivative formula in (4.38). On the other hand, for every $r, \tau \in (0,1)$ with $r \geq \tau$, we have

$$\begin{aligned}
&\frac{\partial}{\partial \tau}\left(\phi'_{k,j,y}(\tau) \exp\Big(\int_\tau^r \ln y_j(s)\, \mathrm{d}s\Big)\right) \\
&= \phi''_{k,j,y}(\tau) \exp\Big(\int_\tau^r \ln y_j(s)\, \mathrm{d}s\Big) - \phi'_{k,j,y}(\tau) \exp\Big(\int_\tau^r \ln y_j(s)\, \mathrm{d}s\Big) \ln y_j(\tau) \\
&= \Big(\phi''_{k,j,y}(\tau) - \phi'_{k,j,y}(\tau) \ln y_j(\tau)\Big) \exp\Big(\int_\tau^r \ln y_j(s)\, \mathrm{d}s\Big) \\
&= -\frac{\theta_{j+1}}{\theta_j} \phi'_{k,j+1,y}(\tau) \ln y_{j+1}(\tau) \exp\Big(\int_\tau^r \ln y_j(s)\, \mathrm{d}s\Big),
\end{aligned}$$

where the last equality comes from the equality in (4.40). We conclude part (ii) by doing $r \to 1$ and then integrating τ between t and one: We use that $y_j(r) \to 0$ in $\mathrm{NLS}_k(\theta)$, therefore $\phi_{k+1,j,y}(r) \to 0$, $\phi_{k+1,j+1,y}(r) \to 0$, and then $\phi'_{k,j,y}(r) \to k!(1 - \theta_{j+1}/\theta_j)$, using the second equality in (4.39). □

Proposition 15. *Consider $\ell \in [k-1]$, and let $\theta_{\ell:k}$ be such that there is a solution $y = (y_\ell, \ldots, y_k)$ for $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$, and let $j \in \{\ell, \ldots, k-1\}$. If y_{j+1} is non-increasing, and if there is $t_1 \in [0,1)$ such that $y_j'(t_1) < 0$ and $y_j(t_1) < 1$, then $y_j'(t) < 0$ for all $t \in [t_1, 1)$.*

Proof. We prove the result by contradiction. Suppose there exists $t_2 \in (t_1, 1)$ such that $y_j'(t_2) \geq 0$. By the continuity of y_j, the value $\min\{y_j(t) : t \in [t_1, t_2]\}$ is well-defined, the minimum in $[t_1, t_2]$ is attained at $t' \in [t_1, t_2]$, and $y_j(t') < 1$ since $y_j(t_1) < 1$ and we are assuming $y_j'(t_1) < 0$. Hence there definitely exists some points t'' in $[t_1, t_2]$ where $y_j(t'') < y_j(t_1)$. Then, in a neighborhood of t', there is $t_1' < t_2'$ such that $y_j(t_1') = y_j(t_2') < 1$ and $y_j'(t_1') < 0$ and $y_j'(t_2') \geq 0$. Then,

$$
\begin{aligned}
0 > (-\ln y_j(t_1'))^{k-1} y_j'(t_1') = \phi_{k,j,y}'(t_1') &= k!\left(1 - \frac{\theta_{j+1}}{\theta_j}\right) - \phi_{k+1,j,y}(t_1') + \frac{\theta_{j+1}}{\theta_j}\phi_{k+1,j+1,y}(t_1') \\
&\geq k!\left(1 - \frac{\theta_{j+1}}{\theta_j}\right) - \phi_{k+1,j,y}(t_2') + \frac{\theta_{j+1}}{\theta_j}\phi_{k+1,j+1,y}(t_2') \\
&= \phi_{k,j,y}'(t_2'),
\end{aligned}
$$

where in the second inequality we used that y_{j+1} is non-increasing. Also, recall that $\phi_{k+1,j,y}(t_1') = \Gamma_{k+1}(-\ln y_j(t_1')) = \Gamma_{k+1}(-\ln y_j(t_2')) = \phi_{k+1,j,y}(t_2')$. From here, the contradiction follows since $\phi_{k,j,y}'(t_2') = (-\ln y_j(t_2))^{k-1} y_j'(t_2') \geq 0$. □

Proposition 16. *Consider $\ell \in [k-1]$, and let $\theta_{\ell:k}$ be such that there is a solution $y = (y_\ell, \ldots, y_k)$ for $\mathrm{NLS}_{\ell,k}(\theta_{\ell:k})$. Then, the following conditions are necessary: For every $j \in \{\ell, \ldots, k\}$, y_j is strictly decreasing in $[0,1)$, $\theta_j < \theta_{j+1}$ for all $j < k$, and $\theta_k < 1/k$.*

Proof. We proceed by induction. Since $y_k(0) = 1$, $y_k(t) \to 0$ for $t \to 1$ from the left, and y_k is differentiable in $(0,1)$, there is a value $t_k \in (0,1)$ such that $y_k'(t_k) < 0$ and $y_k(t_k) < 1$. Since $\phi_{k,k,y}'(t) = y_k'(t)(-\ln y_k(t))^{k-1}$, we have $\phi_{k,k,y}'(t_k) < 0$. Then, from Proposition 14(i), it must be that $1 - 1/k\theta_k < 0$, that is, $\theta_k < 1/k$. Together with Proposition 14(i), this implies that for every $t \in (0,1)$ it holds $\phi_{k,k,y}'(t) < 0$, that is, $\phi_{k,k,y} = \Gamma_k(-\ln y_k)$ is strictly decreasing in $(0,1)$. We conclude that y_k strictly decreases in $[0,1)$.

Assume inductively that $y_{j+1}, \ldots, y_k$ are strictly decreasing for some $j < k$. We will show that $\theta_j < \theta_{j+1}$ and y_j is strictly decreasing. Note that for every $t \in (0,1)$, we have

$$
\phi_{k,j+1,y}'(t) \ln y_{j+1}(t) = -y_{j+1}'(t)(-\ln y_{j+1}(t))^k > 0, \tag{4.41}
$$

where the inequality follows by our inductive assumption and the equality by the derivative formula in (4.38). Now, if $\theta_{j+1} \leq \theta_j$, for every $t \in (0,1)$ we have

$$
\begin{aligned}
&\phi'_{k,j,y}(t) \exp\Big(\int_t^1 \ln y_j(s)\,\mathrm{d}s\Big) \\
&= k!\left(1 - \frac{\theta_{j+1}}{\theta_j}\right) + \frac{\theta_{j+1}}{\theta_j}\int_t^1 \phi'_{k,j+1,y}(\tau)\ln y_{j+1}(\tau)\exp\Big(\int_\tau^1 \ln y_j(s)\,\mathrm{d}s\Big)\,\mathrm{d}\tau \\
&\geq \frac{\theta_{j+1}}{\theta_j}\int_t^1 \phi'_{k,j+1,y}(\tau)\ln y_{j+1}(\tau)\exp\Big(\int_\tau^1 \ln y_j(s)\,\mathrm{d}s\Big)\,\mathrm{d}\tau \geq 0,
\end{aligned}
$$

where the first equality holds by Proposition 14(ii), the first inequality holds by $\theta_{j+1} \leq \theta_j$, and in the last inequality we used inequality (4.41) and the inductive assumption. Therefore, for every $t \in (0,1)$, we have $\phi'_{k,j,y}(t) \geq 0$, which cannot happen since the differentiability of y_j and the border conditions imply that we can always find $t_j \in (0,1)$ such that $y'_j(t_j) < 0$ and $y_j(t_j) < 1$, i.e.,

$$\phi'_{k,j,y}(t_j) = y'_j(t_j)(-\ln y_j(t_j))^{k-1} < 0.$$

We conclude that $\theta_j < \theta_{j+1}$.

We prove next the monotonicity of y_j. Consider $t' = \inf\{t \in [0,1] : y'_j(t) < 0, y_j(t) < 1\}$, which is well-defined since the set is non-empty. If $t' = 0$, then there is a sequence $(t_n)_{n\in\mathbb{N}}$ in $(0,1)$ such that $t'_n \to 0$, $y'_j(t_n) < 0$, and $y_j(t_n) < 1$ for all $n \in \mathbb{N}$. Then, since y_{j+1} is strictly decreasing, by Proposition 15 we get that for all $n \in \mathbb{N}$ and every $t \in [t'_n, 1)$ we have $y'_j(t) < 0$. Since $t'_n \to 0$, we conclude that y'_j is strictly decreasing in $[0,1)$. Otherwise, suppose that $t' > 0$. Then, $y'_j(t) \geq 0$ or $y_j(t) \geq 1$ for every $t \in (0,t')$. Assume that $y_j(s) > 1$ for some $s \in (0,t')$. Then, since $\lim_{q\to 1} y_j(q) = 0$, the continuity of y_j and the fact that $y_j(t') \leq 1$, imply the existence of a value $t'' \in (0,t']$ such that $y_j(t'') = 1$. Note that $0 = y'_j(t'')(-\ln y_j(t''))^{k-1} = \phi'_{k,j,y}(t'')$, and

$$
\begin{aligned}
\phi'_{k,j,y}(t'') &= k! - \phi_{k+1,j,y}(t'') - \frac{\theta_{j+1}}{\theta_j}\left(k! - \phi_{k+1,j+1}(t'')\right) \\
&= k! - \Gamma_{k+1}(0) - \frac{\theta_{j+1}}{\theta_j}\left(k! - \phi_{k+1,j+1}(t'')\right) \\
&= -\frac{\theta_{j+1}}{\theta_j}\left(k! - \Gamma_{k+1}(-\ln y_{j+1}(t''))\right) < 0,
\end{aligned}
$$

which is a contradiction; the first equality holds from (4.39), the second holds since $y_j(t'') = 1$, the third since $\Gamma_{k+1}(0) = k!$, and the inequality follows from y_{j+1} being strictly decreasing. Therefore, we have $y_j(t) \leq 1$ for all $t \in (0,t')$, which further implies that $y_j(t) \leq 1$ for all $t \in (0,1)$.

If $y_j(s) < 1$ for some $s \in [0,t']$, then there exists $t''' \in (0,t')$ such that $y_j'(t''') < 0$ and $y_j(t''') < 1$, which contradicts the minimality of t'. Then, $y_j(t) = 1$ for every $t \in (0,t']$. But this implies that $\phi'_{k,j,y}(t) = 0$ and $\phi_{k+1,j,y}(t) = k!$ for every $t \in (0,t']$, and therefore from (4.39) we get that $k! = \phi_{k+1,j+1,y}(t) = \Gamma_{k+1}(-\ln y_{j+1}(t))$ for every $t \in (0,t']$. This implies that $y_{j+1}(t) = 1$ for every $t \in (0,t']$, which contradicts the fact that y_{j+1} is strictly decreasing. We conclude that $t' = 0$ and, therefore, y_j is strictly decreasing. This finishes the proof of the proposition. □

Proof of Lemma 12. In what follows, we show that there is a choice of θ such that the system $\text{NLS}_k(\theta)$ has a solution. We proceed inductively. We show that there is a solution to this system for an appropriate choice of $\theta_{j:k}$. Using this solution, we can extend it to a solution for $\text{NLS}_{j-1,k}(\theta_{j-1:k})$, where $\theta_{j-1:k} = (\theta_{j-1}, \theta_{j:k})$ for an appropriate choice of θ_{j-1}. For every ℓ we denote by $y_\ell(1)$ the value $\lim_{t\uparrow 1} y_\ell(t)$.

We start with $j = k$. In this case, the $\text{NLS}_{k,k}(\theta_k)$ is the following system:

$$\Gamma_k(-\ln y_k)' = k! - k!/(\theta_k k) - \Gamma_{k+1}(-\ln y_k),$$
$$y_k(0) = 1 \text{ and } y_k(1) = 0.$$

We could analyze this system in the same way as the Hill and Kertz differential equation when $k = 1$. However, we will not show uniqueness for all other θ_j, hence for the sake of simplicity, we only claim existence. As mentioned earlier, we can also write the first equality as

$$\phi'_{k,k,y} = k!\Big(1 - \frac{1}{k\theta_k}\Big) - \Gamma_{k+1}(-\ln y_k)$$

where

$$\phi'_{k,k,y}(t) = -\Gamma_k'(-\ln y_k(t))\frac{y_k'(t)}{y_k(t)} = (-\ln y_k(t))^{k-1} y_k'(t)$$

Hence we have a continuous, explicit first-order differential equation for the domain of $y > 0$. Moreover, from Proposition 16, we already know y_k must be strictly decreasing. Moreover, $y_k(t) \to 0$ as $t \to 1$. Note that we used $\theta_k < \frac{1}{k}$. Thus we may apply Peano's existence theorem [66] and conclude there exists a solution to $\text{NLS}_{k,k}(\theta_k)$ on $[0,1]$ for any $\theta_k < \frac{1}{k}$.

Assume inductively that we have found a $\theta^\star_{j+1:k}$ where we have a solution $(Y_{j+1}, \dots, Y_k)$ to $\text{NLS}_{j+1,k}(\theta^\star_{j+1:k})$ for some $j < k$. We now show that the system $\text{NLS}_{j,k}(\theta_j, \theta^\star_{j+1}, \dots, \theta^\star_k)$

is feasible for a choice of θ_j. This boils down to finding a solution for the following:

$$\Gamma_k(-\ln Y_j)' = k! - \Gamma_{k+1}(-\ln Y_j) - \frac{\theta^\star_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln Y_{j+1})) \tag{4.42}$$

$$Y_j(0) = 1, \text{ and } Y_j(1) = 0. \tag{4.43}$$

where $\theta^\star_{j+1}$ and Y_{j+1} are given and satisfy $Y_{j+1}(0) = 1, Y_{j+1}(1) = 0$. By Proposition 16 we have that $Y_{j+1}, \ldots, Y_k \in [0,1]$ are strictly decreasing and $\theta^\star_{j+1} < \cdots < \theta^\star_k < 1/k$.

The difference is that these ODE's depend on the solutions of previous systems, Y_{j+1}, of which we know little. In order to show existence, our plan is to try to find a solution to the Euler approximation first. Then taking the limit of this approximation will give existence of the original system. Let $\theta_j > 0$ and consider the following Euler approximation to a candidate solution to (4.42)-(4.43). Let $m \in \mathbb{N}$ be non-negative and consider the following recursion: $y_{m,j,0} = 1$, and $\Gamma_k(-\ln y_{m,j,t+1})$ is equal to

$$\Gamma_k(-\ln y_{m,j,t}) + \frac{1}{m}\left(k! - \Gamma_{k+1}(-\ln y_{m,j,t}) - \frac{\theta^\star_{j+1}}{\theta_j}\Big(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m))\Big)\right). \tag{4.44}$$

Note that the sequence is well-defined for $y_{m,j,t} > 0$. Let $t' = \max\{t \in [m] \cup \{0\} : y_{m,j,t} > 0\}$. For $t = 0$, we have $\Gamma_k(-\ln y_{m,j,1}) = (k-1)!$ and therefore $y_{m,j,1} = 1$. For $t = 1$, we have

$$\Gamma_k(-\ln y_{m,j,2}) = (k-1)! - \frac{\theta^\star_{j+1}}{\theta_j m}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(1/m))) < (k-1)!, \tag{4.45}$$

which implies that $y_{m,j,2} < 1$ for any $\theta_j > 0$. We note that if $\theta_j \to \infty$, then $\Gamma_k(-\ln y_{m,j,2}) \to (k-1)!$. Inductively, we can show that for $\theta_j \to \infty$, $y_{m,j,t} = 1$ for all t; in particular, $t' = m$.

We proceed to show that we have sufficient understanding of the Euler approximation. That is, $y_{m,j,t}$ is decreasing in t. Moreover, we want to show that we can adjust parameter θ_j such that $y_{m,j,t}$ decreases at a rate such that it tends to 0 for large m and t close to m. In short, if these claims hold, then we can find suitable θ for the Euler approximations. Then, standard arguments show convergence of the Euler approximation, to a limit that satisfies the constraints of our system of ODE's. The last part holds just by definition of the approximation, (4.44). We start by showing that $y_{m,j,t}$ is decreasing in t as long as $y_{m,j,t} \geq 1/m$ and m is such that $m/\ln(m) \geq 1$ which holds for $m \geq 2$. We know this is true for $t \in \{1,2\}$. We assume the result holds from 1 up to t, and we show next the result holds for $t+1$, with $t \geq 2$. Observe that

$$\Gamma_k(-\ln y_{m,j,t+1}) - \Gamma_k(-\ln y_{m,j,t}) \tag{4.46}$$
$$= \frac{1}{m}\left(k! - \Gamma_{k+1}(-\ln y_{m,j,t}) - \frac{\theta^\star_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))\right) \tag{4.47}$$
$$= \frac{1}{m}\sum_{\tau=0}^{t-1}(\Gamma_{k+1}(-\ln y_{m,j,\tau}) - \Gamma_{k+1}(-\ln y_{m,j,\tau+1}))$$
$$- \frac{\theta^\star_{j+1}}{m\theta_j}\sum_{\tau=0}^{t-1}(\Gamma_{k+1}(-\ln Y_{j+1}(\tau/m)) - \Gamma_{k+1}(-\ln Y_{j+1}((\tau+1)/m)))$$
$$= \Gamma_k(-\ln y_{m,j,t}) - \Gamma_k(-\ln y_{m,j,t-1}) + \frac{1}{m}(\Gamma_{k+1}(-\ln y_{m,j,t-1}) - \Gamma_{k+1}(-\ln y_{m,j,t})) \tag{4.48}$$
$$- \frac{\theta^\star_{j+1}}{m\theta_j}(\Gamma_{k+1}(-\ln Y_{j+1}((t-1)/m)) - \Gamma_{k+1}(-\ln Y_{j+1}(t/m))),$$

where the first equality holds by writing the previous expression using two telescopic sums, and the third equality holds by rearranging terms and using (4.44). Note that $y_{m,j,t+1} < y_{m,j,t}$ if and only if $\Gamma_k(-\ln y_{m,j,t+1}) < \Gamma_k(-\ln y_{m,j,t})$. Since Y_{j+1} is strictly decreasing, the result follows after the following claim. The proof of Claim 5 is in Appendix C.2.

Claim 5. *$\Gamma_k(-\ln y_{m,j,t}) - \Gamma_k(-\ln y_{m,j,t-1}) + \frac{1}{m}(\Gamma_{k+1}(-\ln y_{m,j,t-1}) - \Gamma_{k+1}(-\ln y_{m,j,t})) \leq 0$ for $m \geq 2$ and $y_{m,j,t} \geq \frac{1}{m}$.*

As noted above, we only need the claim to hold in this regime.

We now show that $\partial y_{m,j,t}/\partial\theta_j \geq 0$ for all $t \leq t'$ and such that $y_{m,j,t} \geq 1/m$. Furthermore, we show that for $t \geq 1$ as before, we have $\partial y_{m,j,t}/\partial\theta_j > 0$. We proceed by induction in t. The result is clearly true for $t = 0$. Suppose that $\partial y_{m,j,t}/\partial\theta_j \geq 0$ and let's show the result for $t+1$. By taking the derivative with respect to θ_j in the Euler recursion (4.44), we have

$$(-\ln y_{m,j,t+1})^{k-1}\frac{\partial y_{m,j,t+1}}{\partial\theta_j}$$
$$= \frac{\partial}{\partial\theta_j}\Gamma_k(-\ln y_{m,j,t}) - \frac{1}{m}\frac{\partial}{\partial\theta_j}\Gamma_{k+1}(-\ln y_{m,j,t}) + \frac{\theta^\star_{j+1}}{\theta_j^2 m}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))$$
$$= \frac{\partial}{\partial\theta_j}\Gamma_k(-\ln y_{m,j,t})\left(1 + \frac{1}{m}\ln y_{m,j,t}\right) + \frac{\theta^\star_{j+1}}{\theta_j^2 m}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))$$
$$\geq \frac{\theta^\star_{j+1}}{\theta_j^2 m}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(1/m)))$$

where we used that $y_{m,j,t} \geq 1/m$ for the inequality. Since $y_{m,j,t+1} \in (0,1)$, it follows that $\partial y_{m,j,t+1}/\partial\theta_j > 0$. Notice that the right-hand side of the inequality is independent of t and grows as $1/\theta_j^2$. Hence, as $\theta_j \to 0$, we have that $t' \to 1$ and so $y_{m,j,t} \to 1$ for $t \leq t'$. As a byproduct of this analysis, we also see that t' is strictly increasing in θ_j. Now, let $\theta_j(m)$ be such that $2/m \geq y_{m,j,m-\sqrt{m}} \geq 1/m$. By the work done so far, we can conclude that such $\theta_j(m)$ exists. In particular, $\theta_j(m) > 0$. The next claim shows that $\theta_j(m) \leq \theta^\star_{j+1}$. We defer its proof to Appendix C.2.

Claim 6. *We have* $\theta_j(m) \leq \theta^\star_{j+1}$.

Thus, we have that $\{\theta_j(m)\}_m$ is bounded. Hence we may invoke Bolzano-Weierstrass: if we let m tend to infinity, we can find a convergent subsequence $\{\theta_j(m_\ell)\}_\ell$ with a limit denoted as $\theta^\star_j$.

Finally, using the Euler approximations, we may construct a family of functions that contain a converging subsequence to one that satisfies the conditions of our system. That is, let $y_{\ell,j} : [0,1] \to [0,1]$ be the piece-wise linear interpolation of the points $\{y_{m_\ell,j,t}\}_t$, where $y_{m_\ell,j,t}$ is assigned as the image to the point $t/m_\ell \leq 1$. That is, $y_{\ell,j}(t) = y_{m_\ell,j,t}$ for $t \in [m_\ell]$. This defines $y_{\ell,j}$ on an equally spaced set of points. All other values of the function are given by the straight line segments connecting these points. By a standard analysis argument on the convergence of sequences of functions, we can show that the sequence $\{y_{\ell,j}\}_\ell$ has a uniformly convergent subsequence to a function $Y_j : [0,1] \to [0,1]$ (see, e.g., [56, Chapter 3]). Furthermore, this function Y_j is differentiable and satisfies (4.42)-(4.43). Hence, we have found $\theta^\star_{j:k}$ such that the system $\mathrm{NLS}_{k,j}(\theta^\star_{j:k})$ is feasible. □

4.4.4 Feasibility Analysis and Proof of Lemma 13

In this subsection, we prove Lemma 13. The crux of the proof follows by analyzing the functions $\alpha_{t,j}(q) = (1 + 12\ln(\bar{n}_k)^2/\bar{n}_k)^{k-j+1}\alpha^\star_{t,j}(q)$. These functions hold the following two claims:

Claim 7. *There is $n_0 \geq 1$ such that for any $n \geq n_0$ and for any $t \in \{0,\ldots,\bar{n}_k - 1\}$, we have*

$$\int_0^1 \alpha_{t+1,k}(q)\,\mathrm{d}q + \sum_{\tau \leq t} \int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q \leq 1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}.$$

Claim 8. *There is $n_0 \geq 1$ such that for any $n \geq n_0$, for any $j < k$, and for any $t \in \{0,\ldots,\bar{n}_k - 1\}$, we have*

$$\int_0^1 \alpha_{t+1,j}(q)\,\mathrm{d}q + \sum_{\tau \leq t} \int_0^1 q\alpha_{\tau,j}(q)\,\mathrm{d}q \leq \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \sum_{\tau \leq t} \int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q.$$

Using these two claims, we show how to conclude Lemma 13 and then prove them.

Proof of Lemma 13. First, we have,

$$\begin{aligned}&\int_0^1 \alpha^\star_{t+1,k}(q)\,\mathrm{d}q + \sum_{\tau\le t}\int_0^1 q\alpha^\star_{\tau,k}(q)\,\mathrm{d}q\\&= \frac{1}{1+12\ln(\bar{n}_k)^2/\bar{n}_k}\left(\int_0^1 \alpha_{t+1,k}(q)\,\mathrm{d}q + \sum_{\tau\le t}\int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q\right) \le 1,\end{aligned}$$

where we used Claim 7, which shows that $\alpha^\star$ satisfies constraints (4.32). Additionally,

$$\begin{aligned}&\int_0^1 \alpha^\star_{t,j}(q)\,\mathrm{d}q + \sum_{\tau\le t}\int_0^1 q\alpha^\star_{\tau,j+1}(q)\,\mathrm{d}q\\&= \frac{1}{(1+12\ln(\bar{n}_k)^2/\bar{n}_k)^{k-j+1}}\left(\int_0^1 \alpha_{t,j}(q)\,\mathrm{d}q + \sum_{\tau\le t}\int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q\right)\\&\le \frac{1}{(1+12\ln(\bar{n}_k)^2/\bar{n}_k)^{k-j}}\sum_{\tau\le t}\int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q\\&= \sum_{\tau\le t}\int_0^1 q\alpha^\star_{\tau,j+1}(q)\,\mathrm{d}q,\end{aligned}$$

where in the inequality we used Claim 8; which shows that $\alpha^\star$ satisfies constraints (4.33). This concludes the lemma. □

We devote the rest of this section to prove Claims 7 and 8. The two claims follow by a careful analysis of the solution to $\mathrm{NLS}_{n,k}(\theta^\star)$ as well as the function $g_{n,k}$. In the following proposition, we leave some useful properties satisfied by the function $g_{n,k}$. Recall that we set $\bar{n}_k = n-k-1$. The proof can be found in Appendix C.2.

Proposition 17. *For every $u \in (0,1)$, the following holds:*

(i) $g'_{n,k}(u) = -(n-k+1)(n-k)\binom{n}{k-1}(1-u)^{n-k-1}u^{k-1}$.

(ii) $g'_{n+1,k+1}(u) = \frac{n+1}{k}u g'_{n,k}(u)$.

(iii) *If* $n > (k+1)+2(k+1)^2$, $-g'_{n,k}(u) \le -n\left(1+4\frac{k^2}{n}\right)\frac{\Gamma_k(\bar{n}_k u)'}{(k-1)!}$.

(iv) *If* $n > 4k$ *and* $u \in (0,s)$, *with* $s \le \frac{1}{2\sqrt{\bar{n}_k}}$, *then* $-n\left(1-4\frac{k^2}{n}\right)\left(1-\frac{\bar{n}_k s^2}{1-s}\right)\frac{\Gamma_k(\bar{n}_k u)'}{(k-1)!} \le -g'_{n,k}(u)$.

Recall that for every $r,\ell \in [k]$ we defined $\Phi_{r,\ell} = \Gamma_r(-\ln(Y_\ell))$. Observe that conditions (4.1)-(4.2) for the nonlinear system $\mathrm{NLS}_k(\theta^\star)$ can be rewritten to get the following identities in $[0,1)$:

$$\Phi'_{k,k} = k!\,(1 - 1/(k\theta^\star_k)) - \Phi_{k+1,k}, \tag{4.49}$$

$$\Phi'_{k,\ell} = k! - \Phi_{k+1,\ell} - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}(k! - \Phi_{k+1,\ell+1}) \text{ for every } \ell \in [k-1],. \tag{4.50}$$

Furthermore, since the functions Y_j are non-increasing, $\Phi_{k,j}$ are also non-increasing. Recall that $\Gamma'_r(x) = -x^{r-1}e^{-x}$ and therefore, when $r \geq 2$, we have $\Phi'_{r,\ell}(t) = -\Phi'_{r-1,\ell}(t)\ln Y_\ell(t)$. We use the following technical proposition in the rest of our analysis. For the sake of presentation, we defer its proof to Appendix C.2.

Proposition 18. *For every positive integer k, the following holds:*

(i) *Let $b_k = 4k!\max\{\theta^\star_{\ell+1}/\theta^\star_\ell : \ell \in \{1,\dots,k-1\}\}$. Then, for every $t \in (0,1)$, every $\ell \in [k]$, and every $r \in \{0,\dots,k-1\}$, we have $Y_\ell(t)(-\ln Y_\ell(t))^r \leq b_k(1-t)$.*

(ii) *Let $d_k = \min\{\theta^*_{\ell+1}/\theta^*_\ell : \ell \in \{1,\dots,k-1\}\} - 1 > 0$. There exists $\Delta_k > 0$ such that for every $t \in (\Delta_k, 1]$ and $\ell \in [k]$, it holds $Y_j(t) \geq d_k(1-t)^2$.*

(iii) *Let $c_k = 6b_k$. We have $\Phi''_{k,k}(t) \geq 0$ for every $t \in (0,1)$. Furthermore, there is an integer N_k, such that for every $n \geq N_k$, every $\ell \in [k-1]$, and every $t \in (0, 1-1/n]$, we have $|\Phi''_{k,\ell}(t)| \leq c_k \ln(n)$.*

(iv) *Let $\bar{c}_k = (kc_k)^{1/k}$ and N_k as in* (iii). *There is $\delta_k > 0$ such that for any $n \geq N_k$, $j < k$, and every $t \leq \min\{\delta_k, 1-1/n\}$, we have $Y_j(t) \geq 1 - \bar{c}_k \ln(n)^{1/k} t^{2/k}$.*

Proof of Claim 7. For $t \in \{0,1,\dots,\bar{n}_k - 1\}$, we have

$$\begin{aligned} &\int_0^1 \alpha_{t+1,k}(q)\,\mathrm{d}q + \sum_{\tau \leq t}\int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q \\ &= \theta^\star_k\left(\int_{\varepsilon_{k,t}}^{\varepsilon_{k,t+1}}(-g'_{n,k}(q))\,\mathrm{d}q + \int_0^{\varepsilon_{k,t}} q(-g'_{n,k}(q))\,\mathrm{d}q\right) + B_k\frac{\ln(\bar{n}_k)}{\bar{n}_k}. \end{aligned} \tag{4.51}$$

Now, we bound the term in parenthesis:

$$
\begin{aligned}
&\int_{\varepsilon_{k,t}}^{\varepsilon_{k,t+1}} (-g'_{n,k}(q))\,\mathrm{d}q + \int_0^{\varepsilon_{k,t}} q(-g'_{n,k}(q))\,\mathrm{d}q \\
&= \int_{\varepsilon_{k,t}}^{\varepsilon_{k,t+1}} (-g'_{n,k}(u))\,\mathrm{d}u + \frac{k}{n+1}\int_0^{\varepsilon_{k,t}} (-g'_{n+1,k+1}(u))\,\mathrm{d}u \\
&\le \left(1+16\frac{k^2}{n}\right)\left(\frac{n}{(k-1)!}\int_{\varepsilon_{k,t}}^{\varepsilon_{k,t+1}} -(\Gamma_k(\bar{n}_k u))'\,\mathrm{d}u + \frac{k}{n+1}\frac{n+1}{k!}\int_0^{\varepsilon_{k,t}} -(\Gamma_{k+1}(\bar{n}_k u))'\,\mathrm{d}u\right) \\
&= \frac{n}{(k-1)!}\left(1+16\frac{k^2}{n}\right)\left(\Gamma_k(-\ln y_{k,t}) - \Gamma_k(-\ln y_{k,t+1}) + \frac{1}{n}(k! - \Gamma_{k+1}(-\ln y_{k,t})\right) \\
&\le \frac{n}{(k-1)!}\left(1+16\frac{k^2}{n}\right)\left(\frac{1}{\bar{n}_k}(k! - \frac{(k-1)!}{\theta_k^\star} - \Gamma_{k+1}(-\ln y_{k,t})) - (\Gamma_k(-\ln y_{k,t+1}) - \Gamma_k(-\ln y_{k,t}))\right) \\
&\quad + \frac{n}{\bar{n}_k}\left(1+16\frac{k^2}{n}\right)\frac{1}{\theta_k^\star} \\
&= \frac{n}{(k-1)!}\left(1+16\frac{k^2}{n}\right)\left(\Phi'_{k,k}(t/\bar{n}_k) - \bar{n}_k\Big(\Phi_{k,k}((t+1)/\bar{n}_k)) - \Phi_{k,k}(t/\bar{n}_k)\Big)\right) \\
&\quad + \frac{n}{\bar{n}_k}\left(1+16\frac{k^2}{n}\right)\frac{1}{\theta_k^\star} \\
&\le \frac{n}{\bar{n}_k(k-1)!}\left(1+16\frac{k^2}{n}\right)\left(\Phi'_{k,k}\left(\frac{t}{\bar{n}_k}\right) - \frac{\Phi_{k,k}((t+1)/\bar{n}_k)) - \Phi_{k,k}(t/\bar{n}_k)}{1/\bar{n}_k}\right) \\
&\quad + \frac{1}{\theta_k^\star}\left(1+20\frac{k^2}{n}\right). \qquad (4.52)
\end{aligned}
$$

The first equality and inequality follow by Proposition 17, where we used implicitly that $1 + 4(k+1)^2/(n+1) \le 1 + 16k^2/n$ for any $k \ge 1$. The next equality follows by computing the integrals. The next inequality follows by bounding $1/n \le 1/\bar{n}_k$ and adding and subtracting $1/\theta_k^\star$. The last equality follows by rearranging terms and the last inequality follows by bounding $n/\bar{n}_k(1 + 16k^2/n) \le 1 + 20k^2/n$ for $n \ge 20k^2(k+1)/(4k^2 - k - 1)$ and any $k \ge 1$.

The following claim allows us to bound the first term in (4.52). We defer the proof of the claim to Appendix C.2.

Claim 9. *It holds that* $\Phi'_{k,k}\left(\frac{t}{\bar{n}_k}\right) - \frac{\Phi_{k,k}((t+1)/\bar{n}_k)) - \Phi_{k,k}(t/\bar{n}_k))}{1/\bar{n}_k} \le 0.$

Then, in (4.51), we have

$$
\int_0^1 \alpha_{t+1,k}(q)\,\mathrm{d}q + \sum_{\tau\le t}\int_0^1 q\alpha_{\tau,k}(q)\,\mathrm{d}q \le \left(1+20\frac{k^2}{n}\right) + B_k\frac{\ln(\bar{n})}{\bar{n}} \le 1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k},
$$

where the last inequality holds for n large enough. This concludes the proof of Claim 7. □

Proof of Claim 8. For $j < k$, and $t = k - j$, we have

$$\int_0^1 \alpha_{t+1,j}(q)\,\mathrm{d}q + \sum_{\tau \le t} \int_0^1 q\alpha_{\tau,j}(q)\,\mathrm{d}q \le B_j \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \theta_j^\star \int_0^{\varepsilon_{j,k-j+1}} (-g_{n,k})'(q)\,\mathrm{d}q$$
$$= B_j \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \theta_j^\star \int_{\varepsilon_{j,k-j}}^{\varepsilon_{j,k-j+1}} (-g_{n,k})'(q)\,\mathrm{d}q + \theta_j^\star \int_0^{\varepsilon_{j,k-j}} (-g_{n,k})'(q)\,\mathrm{d}q,$$

and for $t > k - j$, we have

$$\int_0^1 \alpha_{t+1,j}(q)\,\mathrm{d}q + \sum_{\tau \le t} \int_0^1 q\alpha_{\tau,j}(q)\,\mathrm{d}q$$
$$\le B_j \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \theta_j^\star \left(\int_{\varepsilon_{j,t}}^{\varepsilon_{j,t+1}} (-g_{n,k})'(q)\,\mathrm{d}q + \int_0^{\varepsilon_{j,t}} q(-g_{n,k})'(q)\,\mathrm{d}q \right).$$

Then, for any $t \ge k - j$, we obtain the inequality

$$\int_0^1 \alpha_{t+1,j}(q)\,\mathrm{d}q + \sum_{\tau \le t} \int_0^1 q\alpha_{\tau,j}(q)\,\mathrm{d}q$$
$$\le B_j \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \theta_j^\star \left(\int_{\varepsilon_{j,t}}^{\varepsilon_{j,t+1}} (-g_{n,k})'(q)\,\mathrm{d}q + \int_0^{\varepsilon_{j,t}} q(-g_{n,k})'(q)\,\mathrm{d}q \right) + \theta_j^\star \int_0^{\varepsilon_{j,k-j}} (-g_{n,k})'(q)\,\mathrm{d}q. \tag{4.53}$$

We upper bound separately the last two terms in (4.53). For the first term, we have

$$\left[\frac{n}{(k-1)!} \left(1 + 16 \frac{k^2}{n} \right) \right]^{-1} \left(\int_{\varepsilon_{j,t}}^{\varepsilon_{j,t+1}} (-g_{n,k})'(u)\,\mathrm{d}u + \frac{k}{n+1} \int_0^{\varepsilon_{j,t}} (-g_{n+1,k+1})'(u)\,\mathrm{d}u \right)$$
$$\le \Gamma_k(-\ln y_{j,t}) - \Gamma_k(-\ln y_{j,t+1}) + \frac{1}{\bar{n}_k} (k! - \Gamma_{k+1}(-\ln y_{j,t}))$$
$$= \frac{1}{\bar{n}_k} \left(\Gamma_k(-\ln y_j)' \left(\frac{t}{\bar{n}_k} \right) - \frac{\Gamma_k(-\ln y_{j,t+1}) - \Gamma_k(-\ln y_{j,t})}{1/\bar{n}_k} + \frac{\theta_{j+1}^\star}{\theta_j^\star} (k! - \Gamma_{k+1}(-\ln y_{j+1,t})) \right)$$
$$= \frac{1}{\bar{n}_k} \left(\Phi_{k,\ell}'(t/\bar{n}_k) - \frac{\Phi_{k,\ell}((t+1)/\bar{n}_k) - \Phi_{k,\ell}(t/\bar{n}_k)}{1/\bar{n}_k} + \frac{\theta_{j+1}^\star}{\theta_j^\star} (k! - \Gamma_{k+1}(-\ln y_{j+1,t})) \right). \tag{4.54}$$

The first inequality follows by Proposition 17. The following claim allows us to guarantee that $\varepsilon_{\ell,k}$ is close to zero, for all ℓ, which allows us to use 17(iv). The proof simply uses Proposition 18(iv) for $\bar{n}_k \ge N_k$ and we skip it for brevity.

Claim 10. *For any ℓ, we have $\varepsilon_{\ell,k} \le 2\bar{c}_k \ln(\bar{n}_k)^{1/k} / \bar{n}_k^{1+2/k}$, where $\bar{c}_k$ is defined in Proposition 18.*

Note that the claim implies that for n large, $\varepsilon_{\ell,k} \leq (k-1)/\bar{n}_k \leq 2/\sqrt{\bar{n}_k}$. In addition to this claim, the following claims allow us to bound the terms in the parenthesis in (4.54). We defer their proof to Appendix C.2.

Claim 11. *It holds that* $\bar{n}_k(\Phi_{k,\ell}(t/\bar{n}_k) - \Phi_{k,\ell}((t+1)/\bar{n}_k)) + \Phi'_{k,\ell}(t/\bar{n}_k) \leq c_k \ln(\bar{n}_k)/\bar{n}_k$, *where* $c_k > 0$ *is defined in Proposition 18.*

Claim 12. *For n sufficiently large, we have*

$$\left(1 - 4\frac{(k+1)^2}{n+1}\right)^{-1} \left(1 - \frac{\bar{n}_k \varepsilon_{j+1,t}^2}{1-\varepsilon_{j+1,t}}\right)^{-1} \leq 1 + 10\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}.$$

Hence, we can further bound (4.54) as follows:

$$\leq \frac{1}{\bar{n}_k}\left(c_k \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \frac{\theta^\star_{j+1}}{\theta^\star_j}(k! - \Gamma_{k+1}(-\ln y_{j+1,t}))\right) \qquad \text{(Using Claim 11)}$$

$$\leq \frac{1}{\bar{n}_k}\left(c_k \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \left(1 - \frac{4(k+1)^2}{n+1}\right)^{-1}\left(1 - \frac{\bar{n}_k \varepsilon_{j+1,t}^2}{1-\varepsilon_{j+1,t}}\right)^{-1} \frac{\theta^\star_{j+1}}{\theta^\star_j}(k-1)! \int_0^{\varepsilon_{j+1,t}} (-g'_{n,k}(u))u \, \mathrm{d}u\right)$$

(Using 17(iv) and Claim 10)

$$\leq \frac{1}{\bar{n}_k}\left(c_k \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \left(1 + 10\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \frac{\theta^\star_{j+1}}{\theta^\star_j}(k-1)! \int_0^{\varepsilon_{j+1,t}} (-g'_{n,k}(u))u \, \mathrm{d}u\right)$$

(Using Claim 12)

$$\leq \frac{1}{\theta^\star_j \bar{n}_k}\left(1 + 10\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\left(\theta^\star_j c_k \frac{\ln \bar{n}_k}{\bar{n}_k} + \theta^\star_{j+1}(k-1)! \int_0^{\varepsilon_{j+1,t}} (-g'_{n,k}(u))u \, \mathrm{d}u\right).$$

From here, we obtain

$$\theta^\star_j \left(\int_{\varepsilon_{j,t}}^{\varepsilon_{j+1,t}} (-g_{n,k})'(q) \, \mathrm{d}q + \int_0^{\varepsilon_{j,t}} q(-g_{n,k})'(q) \, \mathrm{d}q\right)$$
$$\leq \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \frac{c_k}{k!} \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \int_0^{\varepsilon_{j+1,t}} \theta^\star_{j+1} q(-g_{n,k})'(q) \, \mathrm{d}q$$
$$\leq \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right) \frac{c_k}{k!} \frac{\ln(\bar{n}_k)}{\bar{n}_k} + \left(1 + 12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\left(\sum_{\tau \leq t} \int_0^1 q\alpha_{\tau,j+1}(q) \, \mathrm{d}q - B_{j+1}\frac{\ln(\bar{n}_k)}{\bar{n}_k}\right).$$

We now bound the last term in (4.53). Note that the function $-\Gamma_k(\bar{n}_k u)' = \bar{n}_k(\bar{n}_k u)^{k-1}e^{-\bar{n}_k u}$ is increasing in $[0,(k-1)/\bar{n}_k]$ and decreasing in $[(k-1)/\bar{n}_k, +\infty)$. Then,

$$
\begin{aligned}
&\theta_j^\star \int_0^{\varepsilon_{j,k-j}} (-g_{n,k})'(q)\,\mathrm{d}q \\
&\leq \frac{1}{k}\int_0^{\varepsilon_{j,k}} (-g_{n,k})'(q)\,\mathrm{d}q && \text{(Since } \varepsilon_{j,k-j} \leq \varepsilon_{j,k} \text{ and } \theta_j^\star \leq 1/k\text{)} \\
&\leq \frac{n}{k!}\left(1+4\frac{k^2}{n}\right)\int_0^{\varepsilon_{j,k}} (-\Gamma_k(\bar{n}_k u))'\,\mathrm{d}u && \text{(Using Proposition 17)} \\
&\leq \frac{n}{k!}\left(1+4\frac{k^2}{n}\right)\bar{n}_k^k\left(2\bar{c}_k\frac{\ln(\bar{n}_k)^{1/k}}{\bar{n}_k^{1+2/k}}\right)^k && \text{(Using Claim 10)} \\
&\leq \frac{2^k k^{1/k}}{k!}c_k^k\left(1+8\frac{k^2}{n}\right)\frac{\ln(\bar{n}_k)}{\bar{n}_k} \\
&\leq 4c_k^k\left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\frac{\ln(\bar{n}_k)}{\bar{n}_k},
\end{aligned}
$$

where we used that $2^k k^{1/k} \leq 4k!$ for all $k \geq 1$ and the bound $8k^2/n \leq 12\ln(\bar{n}_k)^2/\bar{n}_k$ for n large. Then,

$$
\begin{aligned}
&\int_0^1 \alpha_{t+1,j}(q)\,\mathrm{d}q + \sum_{\tau\leq t}\int_0^1 q\alpha_{\tau,j}(q)\,\mathrm{d}q \\
&\leq \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\sum_{\tau\leq t}\int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q \\
&\quad + 4c_k^k\left(1+8\frac{k^2}{n}\right)\frac{\ln(\bar{n}_k)}{\bar{n}_k} + B_j\frac{\ln(\bar{n}_k)}{\bar{n}_k} + \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\frac{c_k}{k!}\frac{\ln(\bar{n}_k)}{\bar{n}_k} \\
&\quad - \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\frac{\ln(\bar{n}_k)}{\bar{n}_k}B_{j+1} \\
&\leq \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\sum_{\tau\leq t}\int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q + \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\frac{\ln(\bar{n}_k)}{\bar{n}_k}\left(4c_k^k + \frac{c_k}{k!} + B_j - B_{j+1}\right) \\
&= \left(1+12\frac{\ln(\bar{n}_k)^2}{\bar{n}_k}\right)\sum_{\tau\leq t}\int_0^1 q\alpha_{\tau,j+1}(q)\,\mathrm{d}q,
\end{aligned}
$$

where we used that $B_j = (4c_k^k + c_k/k!)\cdot(j-1)$ for $j \geq 1$. This finishes the proof of Claim 8. □

4.5 A tight prophet inequality for sequential assignment

In this section, we show that our new provable lower bounds for the k-selection prophet inequality imply a tight approximation ratio for the i.i.d. sequential stochastic assignment problem by Derman et al. [26], that we call SSAP in what follows. We provide the proof in two steps. Firstly, we show that SSAP is more general than (k,n)-PIP in the sense that any policy for the sequential stochastic assignment problem with n time periods implies a policy for (k,n)-PIP for any $k \in [n]$ (Proposition 19). This shows that the approximation ratio cannot be larger than $\min_{k\in[n]} \gamma_{n,k}$. Secondly, we match the upper bound by using the structure of the optimal policy for SSAP (Proposition 20).

In the SSAP, the input is given by n non-negative values (rewards) $r_1 \le r_2 \le \cdots \le r_n$ and we observe exactly n non-negative values, presented one after the other in n time periods, and drawn independently from a distribution F. For notational convenience, we assume that time starts at $t = n$ and decreases all the way down to $t = 1$, i.e., the value t represents the number of time periods that remain before the next value is presented. For every period t, we observe the value $X_t \sim F$, and we have to irrevocably assign the value X_t to one of the unassigned rewards r_τ's. The goal is to find a sequential policy π that maximizes $v_{n,F,r}(\pi) = \mathbb{E}\left[\sum_{t=1}^{n} X_t r_{\pi(t)}\right]$ where π is a permutation of $[n]$. Note that the optimal offline value corresponds to $\sum_{t=1}^{n} r_t \mathbb{E}\left[X_{(t)}\right]$. We denote by α_n the largest approximation ratio that any policy can attain in SSAP for instances with n time periods.

Proposition 19. *For every n, it holds that $\alpha_n \le \min_{k\in[n]} \gamma_{n,k}$.*

Proof. Let π be a policy for SSAP with approximation ratio α. Given $k \in [n]$, we use π to construct a policy π' for (k,n)-PIP. Without loss of generality, we can assume that $\mathrm{OPT}_{n,k} = 1$. Fix $\varepsilon \in (0, 1/n^2)$ and consider the following instance for SSAP: $r_i = \varepsilon i$ for each $i \in \{1,\dots,n-k\}$ and $r_i = 1$ for each $i \in \{n-k+1,\dots,n\}$. The policy π' simulates π by creating $r_1 \le r_2 \le \cdots \le r_n$ as defined before. When π assigns X_t to some $r_i = 1$, then π' selects the value X_t, while if π assigns X_t to some $r_i = \varepsilon i$, then π' discards X_t. Then, we have that

$$\begin{aligned}
\mathbb{E}\left[\sum_{t=1}^{n} X_t \mathbf{1}_{\{t \text{ selected by } \pi\}}\right] &\ge \alpha \mathbb{E}\left[\sum_{t=1}^{n} r_t X_{(t)}\right] - \varepsilon n \\
&\ge \alpha \mathbb{E}\left[\sum_{t=1}^{k} X_{(n-t+1)}\right] - (1+\alpha)\varepsilon n = \alpha - (1+\alpha)\varepsilon n.
\end{aligned}$$

This shows that the approximation ratio of π' is at least $\alpha - (1+\alpha)\varepsilon n$. Since this holds for any $\varepsilon \in (0, 1/n^2)$, we conclude that π' has an approximation ratio of at least α. Since this holds for any SSAP policy for n time periods and any $k \in [n]$, we conclude the proof. □

Proposition 20. *For every n, it holds that $\alpha_n \ge \min_{k\in[n]} \gamma_{n,k}$.*

To prove this proposition, we need to use the structure of the optimal dynamic programming policy for SSAP shown by Derman et al. [26], which we describe in what follows. For each time distribution F and each time t, there exist values $0 = \mu_{0,t}(F) \le \mu_{1,t}(F) \le \cdots \le \mu_{t,t}(F)$, where the value $\mu_{i,t}$ is the optimal expected value in problem with $t-1$ time periods in which the reward r_i is assigned under the optimal policy. If $X_t \in [\mu_{\tau-1,t}(F), \mu_{\tau,t}(F)]$ then, the optimal policy assigns X_t with the τ-th smallest available reward for $\tau \in \{1,\ldots,t\}$. Furthermore, Derman et al. [26] show that $v_{n,F,r}(\pi^\star) = \sum_{t=1}^n r_t \mu_{t,n+1}(F)$. Note that the values μ are completely independent of the rewards, and they just depend on the distribution F and n.

Proof of Proposition 20. For every $\ell \in [n]$, let $d_\ell = r_\ell - r_{\ell-1}$ where $r_0 = 0$. Since the rewards r_t are non-decreasing in t, we have $d_\ell \ge 0$ for every $\ell \in [n]$, and $r_j = \sum_{\ell=1}^j d_\ell$. Then, for every distribution F, we have

$$\frac{\sum_{t=1}^n r_t \mu_{t,n+1}(F)}{\sum_{t=1}^n r_t \mathbb{E}[X_{(t)}]} = \frac{\sum_{\tau=1}^n d_\tau \sum_{t=\tau}^n \mu_{t,n+1}(F)}{\sum_{\tau=1}^n d_\tau \sum_{t=\tau}^n \mathbb{E}[X_{(t)}]} \ge \min_{\tau\in[n]} \frac{\sum_{t=\tau}^n \mu_{t,n+1}(F)}{\sum_{t=\tau}^n \mathbb{E}[X_{(t)}]}.$$

Note that $\sum_{t=\tau}^n \mu_{t,n+1}(F)$ is the reward collected by the optimal policy $\pi^\star$ in the instance $r_1 = \cdots = r_{\tau-1} = 0 < 1 = r_\tau = \cdots = r_n$. Furthermore, $\sum_{t=\tau}^n \mathbb{E}[X_{(t)}]$ is the sum of the $n-\tau+1$ largest values in a sequence of n i.i.d. samples from F, i.e., $\sum_{t=\tau}^n \mathbb{E}[X_{(t)}] = \mathrm{OPT}_{n,n-\tau+1}(F)$. Therefore, the ratio inside the minimization operator can be interpreted as the ratio in a (k,n)-PIP with $k = n-\tau+1$. Since $\pi^\star$ is optimal for the instance r described above, then it must be the case that $v_{n,F,r}(\pi^\star) = \sum_{t=\tau}^n \mu_{t,n+1}(F) \ge \gamma_{n,n-\tau+1}\mathrm{OPT}_{n,n-\tau+1}(F)$. The proof follows since this holds for every $\tau \in [n]$. □

Proposition 19 and Proposition 20 imply that $\alpha_n = \min_{k\in[n]} \gamma_{n,k}$ for every n. The $1 - k^k e^{-k}/k!$ lower bound on $\gamma_{n,k}$ imply that $\gamma_{n,k}$ is at least 0.78 for $k \ge 3$ (see, e.g., [10, 27]) which is in particular larger than $\liminf_n \gamma_{n,1} \approx 0.745$. Since our results imply that $\liminf_n \gamma_{n,2} \ge 0.829$, we conclude that $\liminf_n \alpha_n = \liminf_n \gamma_{n,1} \approx 0.745$.

Chapter 5

Discussion and Future Work

5.1 Competition complexity for i.i.d values

An additional set of questions that fits the wider theme of chapter 2 concerns the competition complexity of *static pricing*. Here—unlike in the case of dynamic pricing—there are two questions we could ask. The first comparison is between static pricing A'_m and the optimal auction M_n; the other is between static pricing A'_m and dynamic pricing A_n.

For the first comparison between A'_m and M_n, we observe the following. First, since $A'_m \leq A_m$ for all m, our impossibility (Main Result 1) implies that the exact competition complexity of static pricing is unbounded. Moreover, while the approximate competition complexity of static pricing may be linear (similar to our Main Result 2 for dynamic pricing), the dependence on ε certainly has to be worse. This follows from considering the uniform case: For m sufficiently large, we have that $1 - 2\log(m)/m \leq A'_m \leq 1 - \log(m)/(3m)$ (see Appendix A for a derivation of these inequalities). Since $M_n \approx 1 - 1/n$, for large m and n, this means that to ensure that $(1+\varepsilon)A'_m \geq M_n$, we approximately need that $(1+\varepsilon)(1 - \log(m)/(3m)) \geq 1 - 1/n$. Then, for ε small with respect to n, say $\varepsilon = 1/n^2$, we can approximate by subtracting ε from the left-hand side. We get $1 - (1+\varepsilon)\log(m)/(3m) \geq 1 - 1/n$, which happens if and only if $3m/\log(m) \geq n(1+\varepsilon)$, which for ε of this order implies that we need at least $m = cn$ with $c = \Omega(\log(1/\varepsilon))$.

For the other comparison, between A'_m and A_n, observe that for the exact version, we need $m = \Omega(n\log(n))$, even for the uniform distribution. This again follows from the asymptotic formulas for $A'_m \approx 1 - 2\log(m)/m$ and $A_n \approx 1 - 2/(n + \log(n) + 1.76799)$, which show that roughly what we need is that $m/\log(m) \geq n$ and therefore $m = \Omega(n\log(n))$. We leave the full resolution of these gaps, which will shed additional light on the relative power of static and dynamic pricing, to future work.

5.2 Competition complexity for independent values

In chapter 3, section 3.5.1 obviously leaves more open than answered questions. In fact, we introduce a quite general model of what we believe is a good way to extend the study of competition complexity to a combinatorial setting. In theorem 7, we prove that the $(1-\varepsilon)$-competition complexity of block-consistent prices for submodular and XOS combinatorial auctions is $O(\log(1/\varepsilon))$. It would be very interesting to find, maybe even for simpler subclasses, if it is possible to get asymptotic results that lie somewhere between $O(\log(1/\varepsilon))$ and the main result of this chapter for a single item, $O(\log\log(1/\varepsilon))$. For example, we could study the competition complexity for the vertex arrival model in bipartite graphs with one-sided arrival. This problem falls in the more general XOS combinatorial auctions setting of Section 3.5.2. More specifically, we want to consider an underlying bipartite graph $G=(U,V,E)$, and the feasibility constraint is given by the set of matchings in G, that is, $\mathcal{F}=\{S\subseteq E : \text{for all } i\in U\cup V, |\{e\in S : i\in e\}|\leq 1\}$, and the valuation function is additive, i.e., $f(v,S)=\sum_{e\in S} v_e$. The vertices of V arrive online, one by one, and upon their arrival, their edges to all vertices in U are revealed. For every edge $e\in E$, the value v_e is sampled according to a distribution F_e, and we denote $F=\bigtimes_{e\in E} F_e$. This setting is more structured, and it feels quite natural to consider resource augmentation in an online vertex arrival model. We conclude that it could be of interest and moreover possible, to further explore this direction.

5.3 K item i.i.d prophets

In chapter 4, we provide a new exact formulation for (k,n)-PIP. From our formulation, we can derive the nonlinear system of differential equations (4.1)-(4.3) as n tends to infinity. Using this system, we can obtain provable guarantees for the approximation ratio of (k,n)-PIP when n is large. We use the nonlinear system (4.1)-(4.3) to provide new improved bounds for small values of $k\in\{2,\ldots,5\}$. As a direct application of our new bounds, we provide a tight approximation ratio for the SSAP.

We also remark that our linear programming formulation offers several characteristics that make it suitable for a nonlinear analysis. [46] use a different formulation to provide bounds on the k-selection problem, and for the particular case of $k=1$, they also connect to the Hill and Kertz equation. Nevertheless, they go through an extra intermediate formulation in their limit analysis. With our approach, we can directly provide a feasible solution in quantile space that converges to a solution of the nonlinear system (see Subsection 4.4.1), and furthermore, it works for general values of k.

Finding an analytical formula for the approximation ratio of (k,n)-PIP remains an open problem, but our findings offer a potential avenue toward this goal. Using the system (4.1)-(4.3), we believe it is possible to characterize the value $\theta_k^\star$ using an integral equation and show that $\theta_k^\star \geq (1-e^{-k})/k$, though we don't yet have all details for this. Providing a similar lower bound for $j < k$ becomes nontrivial due to the dependency with $j+1$.

We also remark that, in principle, our solution is suboptimal as we only construct a feasible solution to the weak dual $[\mathrm{D}]_{n,k}$. Another place where suboptimality could appear is in the the weak duality between $[\mathrm{D}]_{n,k}$ and $[\mathrm{P}]_{n,k}$. However, we believe these two programs satisfy strong duality. In fact, the following LP:

$$
\begin{aligned}
& \sup \quad v & & [\mathrm{D}_{\text{strong}}]_{n,k} \\
\text{s.t.} \quad & \int_0^1 \beta_{1,\ell}(q)\,\mathrm{d}q \leq \mathbf{1}_k(\ell), \quad \text{for all } \ell \in [k], \\
& \int_0^1 \beta_{t+1,k}(q)\,\mathrm{d}q \leq \int_0^1 (1-q)\beta_{t,k}(q)\,\mathrm{d}q, \quad \text{for all } t \in [n-1], \\
& \int_0^1 \beta_{t+1,\ell}(q)\,\mathrm{d}q \leq \int_0^1 (1-q)\beta_{t,\ell}(q)\,\mathrm{d}q + \int_0^1 q\beta_{t,\ell+1}(q)\,\mathrm{d}q, \ \text{ all } t \in [n-1], \ell \in [k-1], \\
& v g_{n,k}(u) + \frac{\mathrm{d}\eta}{\mathrm{d}u}(u) \leq \sum_{t=1}^{n}\sum_{\ell=1}^{k} \int_u^1 \beta_{t,\ell}(q)\,\mathrm{d}q, \text{ for } u \in [0,1], \\
& \beta_{t,\ell}(q) \geq 0 \quad \text{for all } q \in [0,1], t \in [n] \text{ and } \ell \in [k], \\
& \eta(q) \geq 0 \quad \text{for all } q \in [0,1], \\
& \eta(0) = \eta(1) = 0,
\end{aligned}
$$

can be shown to be a strong dual to $[\mathrm{P}]_{n,k}$. The weak duality proof is analogous to the proof of Lemma 11, and the strong duality holds using a discretization argument as in [67]. We remark that $[\mathrm{D}]_{n,k}$ is a restriction of $[\mathrm{D}_{\text{strong}}]_{n,k}$ with $\eta = 0$.

Our results are valid for n large enough ($n \geq n_0$, with n_0 depending only on k); hence, providing a lower bound for all n remains an open problem and will require additional structural results over the solution of the nonlinear system (4.1)-(4.3). For instance, for $k > 1$ and $j < k$, we observe that the higher-order derivatives of y_j change signs over $[0,1]$, hence, ruling out techniques that work in the case $k = 1$ (see, e.g., [12, 23]). We leave the tightness of our approximation result in Theorem 10 through our nonlinear system (4.1)-(4.3) as an open question, i.e., whether $\inf_{n\geq 1} \gamma_{n,k} = \sum_{\ell=1}^{k} \theta_\ell^\star$.

References

[1] Abolhassani, M., Ehsani, S., Esfandiari, H., Hajiaghayi, M., Kleinberg, R. D., and Lucier, B. (2017). Beating $1 - 1/e$ for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 61–71, New York, NY, USA. ACM.

[2] Alaei, S. (2011). Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In Papadimitriou, C. H. and Arora, S., editors, *Proceedings of the 52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 512–521, Palm Springs, CA, USA. IEEE.

[3] Alaei, S. (2014). Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972.

[4] Alaei, S., Hartline, J. D., Niazadeh, R., Pountourakis, E., and Yuan, Y. (2015). Optimal auctions vs. anonymous pricing. In Charikar, M., Cohen, E., and King, V., editors, *Proceedings of the IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1463, Washington, DC, USA. IEEE.

[5] Arnosti, N. and Ma, W. (2023). Tight guarantees for static threshold policies in the prophet secretary problem. *Operations research*, 71(5):1777–1788.

[6] Assaf, D., Goldstein, L., and Samuel-Cahn, E. (2002). Ratio prophet inequalities when the mortal has several choices. *The Annals of Applied Probability*, 12(3):972–984.

[7] B. Lehmann, D. L. and Nisan, N. (2006). Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior*, 55:270–296.

[8] Babaioff, M., Goldner, K., and Gonczarowski, Y. A. (2020). Bulow-klemperer-style results for welfare maximization in two-sided markets. In *SODA 2020*, pages 2452–2471.

[9] Babaioff, M., Immorlica, N., Kempe, D., and Kleinberg, R. (2007). A knapsack secretary problem with applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2007)*, volume 4627 of *Lecture Notes in Computer Science*, pages 16–28. Springer.

[10] Beyhaghi, H., Golrezaei, N., Leme, R. P., Pál, M., and Sivan, B. (2021). Improved revenue bounds for posted-price and second-price mechanisms. *Operations Research*, 69(6):1805–1822.

[11] Beyhaghi, H. and Weinberg, S. M. (2019). Optimal (and benchmark-optimal) competition complexity for additive buyers over independent items. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC '19, pages 686–696, New York, NY, USA. Association for Computing Machinery.

[12] Brustle, J., Correa, J., Dütting, P., Ezra, T., Feldman, M., and Verdugo, V. (2024). The competition complexity of prophet inequalities. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC 2024)*, EC '24, pages 1–24, New York, NY, USA. Association for Computing Machinery.

[13] Brustle, J., Correa, J., Dütting, P., and Verdugo, V. (2022). The competition complexity of dynamic pricing. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC 2022)*, EC '22, pages 303–320, New York, NY, USA. Association for Computing Machinery.

[14] Buchbinder, N., Jain, K., and Singh, M. (2014). Secretary problems via linear programming. *Mathematics of Operations Research*, 39(1):190–206.

[15] Bulow, J. and Klemperer, P. (1996). Auctions versus negotiations. *American Economic Review*, 86(1):180–94.

[16] Chakraborty, T., Even-Dar, E., Guha, S., Mansour, Y., and Muthukrishnan, S. (2010). Approximation schemes for sequential posted pricing in multi-unit auctions. In *WINE 2010*, pages 158–169.

[17] Chan, T. H., Chen, F., and Jiang, S. H.-C. (2014). Revealing optimal thresholds for generalized secretary problem via continuous lp: impacts on online k-item auction and bipartite k-matching with random arrival order. In *SODA 2014*, pages 1169–1188.

[18] Chawla, S., Hartline, J. D., Malec, D. L., and Sivan, B. (2010). Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 311–320, New York, NY, USA. Association for Computing Machinery.

[19] Clarke, E. H. (1971). Multipart pricing of public goods. *Public Choice*, 11:17–33.

[20] Correa, J. and Cristi, A. (2023). A constant-factor prophet inequality for online combinatorial auctions. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing (STOC 2023)*, pages 686–697, New York, NY, USA. Association for Computing Machinery.

[21] Correa, J., Dütting, P., Fischer, F., and Schewior, K. (2019a). Prophet inequalities for i.i.d. random variables from an unknown distribution. In *Proceedings of the 20th ACM Conference on Economics and Computation (EC)*, pages 3–17, New York, NY, USA. Association for Computing Machinery.

[22] Correa, J., Foncea, P., Hoeksma, R., Oosterwijk, T., and Vredeveld, T. (2019b). Recent developments in prophet inequalities. *SIGecom Exchanges*, 17(1):61–70.

[23] Correa, J. R., Foncea, P., Hoeksma, R., Oosterwijk, T., and Vredeveld, T. (2021). Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of Operations Research*, 46(4):1452–1478.

[24] Correa, J. R., Foncea, P., Pizarro, D., and Verdugo, V. (2019c). From pricing to prophets, and back! *Operations Research Letters*, 47(1):25–29.

[25] Csirik, J. and Woeginger, G. J. (2000). Resource augmentation for online bounded space bin packing. In auf der Heide, F. M. and Schenkel, R., editors, *Proceedings of the 27th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 296–304, Berlin, Heidelberg. Springer.

[26] Derman, C., Lieberman, G. J., and Ross, S. M. (1972). A sequential stochastic assignment problem. *Management Science*, 18(7):349–355.

[27] Dütting, P., Feldman, M., Kesselheim, T., and Lucier, B. (2020). Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582.

[28] Dütting, P., Fischer, F. A., and Klimm, M. (2016). Revenue gaps for discriminatory and anonymous sequential posted pricing. *CoRR*, abs/1607.07105.

[29] Dütting, P., Kesselheim, T., and Lucier, B. (2020). An $o(\log\log m)$ prophet inequality for subadditive combinatorial auctions. In *Proceedings of the IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 306–317, Piscataway, NJ, USA. IEEE.

[30] Eden, A., Feldman, M., Friedler, O., Talgam-Cohen, I., and Weinberg, S. M. (2017). The competition complexity of auctions: A Bulow-Klemperer result for multi-dimensional bidders. In *Proceedings of the 2017 ACM Conference on Economics and Computation (EC)*, page 343.

[31] Ehsani, S., Hajiaghayi, M., Kesselheim, T., and Singla, S. (2018). Prophet secretary for combinatorial auctions and matroids. In *SODA 2018*, pages 700–714.

[32] Ekbatani, F., Niazadeh, R., Nuti, P., and Vondrák, J. (2024). Prophet inequalities with cancellation costs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1247–1258.

[33] Ezra, T., Feldman, M., Gravin, N., and Tang, Z. G. (2022). Prophet matching with general arrivals. *Mathematics of Operations Research*, 47(2):878–898.

[34] Ezra, T., Feldman, M., Roughgarden, T., and Suksompong, W. (2020). Pricing multi-unit markets. *ACM Transactions on Economics and Computation*, 7(4):1–29.

[35] Feldman, M., Friedler, O., and Rubinstein, A. (2018). 99% revenue via enhanced competition. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 443–460.

[36] Feldman, M., Gravin, N., and Lucier, B. (2015). Combinatorial auctions via posted prices. In Charikar, M., Cohen, E., and King, V., editors, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 123–135, Philadelphia, PA. Society for Industrial and Applied Mathematics (SIAM).

[37] Feldman, M., Svensson, O., and Zenklusen, R. (2016). Online contention resolution schemes. In Krauthgamer, R., editor, *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1014–1033. SIAM.

[38] Gilbert, J. and Mosteller, F. (1966). Recognizing the maximum of a sequence. *Journal of the American Statistical Association*, 61:35–73.

[39] Goyal, V. and Udwani, R. (2023). Online matching with stochastic rewards: Optimal competitive ratio via path-based formulation. *Operations Research*, 71(2):563–580.

[40] Gravin, N. and Wang, H. (2019). Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In *EC 2019*, pages 93–109.

[41] Groves, T. (1973). Incentives in teams. *Econometrica*, 41(4):617–631.

[42] Hajiaghayi, M. T., Kleinberg, R., and Sandholm, T. (2007). Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 58–65, Palo Alto, CA, USA. AAAI Press.

[43] Hartline, J. (2020). *Mechanism Design and Approximation*. http://jasonhartline.com/MDnA/.

[44] Hartline, J. D. and Roughgarden, T. (2009). Simple versus optimal mechanisms. In *Proceedings 10th ACM Conference on Electronic Commerce (EC)*, pages 225–234.

[45] Hill, T. P. and Kertz, R. P. (1982). Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables. *The Annals of Probability*, 10(2):336–345.

[46] Jiang, J., Ma, W., and Zhang, J. (2022). Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack. In *SODA 2022*, pages 1221–1246.

[47] Jiang, J., Ma, W., and Zhang, J. (2023). Tightness without counterexamples: A new approach and new results for prophet inequalities. In *EC 2023*, page 909.

[48] Jin, Y., Jiang, S., Lu, P., and Zhang, H. (2021). Tight revenue gaps among multi-unit mechanisms. In Immorlica, N., Lucier, B., and Johari, R., editors, *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, pages 654–673. ACM.

[49] Jin, Y., Lu, P., Qi, Q., Tang, Z. G., and Xiao, T. (2019a). Tight approximation ratio of anonymous pricing. In Charikar, M., Cohen, E., and King, V., editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 674–685, New York, NY, USA. ACM.

[50] Jin, Y., Lu, P., Tang, Z. G., and Xiao, T. (2019b). Tight revenue gaps among simple mechanisms. In Charikar, M., Cohen, E., and King, V., editors, *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 209–228, Philadelphia, PA. Society for Industrial and Applied Mathematics (SIAM).

[51] Kalyanasundaram, B. and Pruhs, K. (2000). Speed is as powerful as clairvoyance. *Journal of the ACM*, 47(4):617–643.

[52] Kertz, R. P. (1986). Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality. *Journal of Multivariate Analysis*, 19(1):88–112.

[53] Kesselheim, T., Tönnis, A., Radke, K., and Vöcking, B. (2014). Primal beats dual on online packing lps in the random-order model. In *STOC 2014*, pages 303–312.

[54] Kleinberg, J. M. and Kleinberg, R. (2018). Delegated search approximates efficient search. In Éva Tardos, Elkind, E., and Vohra, R., editors, *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 287–302, New York, NY, USA. Association for Computing Machinery.

[55] Kleinberg, R. and Weinberg, S. M. (2012). Matroid prophet inequalities. In Charikar, M., Cohen, E., and King, V., editors, *Proceedings of the 44th Symposium on Theory of Computing Conference (STOC)*, pages 123–136, New York, NY, USA. ACM.

[56] Kolmogorov, A. N. and Fomin, S. V. (1975). *Introductory real analysis*.

[57] Krengel, U. and Sucheston, L. (1977). Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83:745–747.

[58] Krengel, U. and Sucheston, L. (1978). On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics*, 4:197–266.

[59] Lee, E. and Singla, S. (2018). Optimal online contention resolution schemes via ex-ante prophet inequalities. In *ESA 2018*.

[60] Liu, A., Leme, R. P., Pál, M., Schneider, J., and Sivan, B. (2021). Variable decomposition for prophet inequalities and optimal ordering. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC)*, page 692, New York, NY, USA. Association for Computing Machinery.

[61] Liu, S. and Psomas, C. (2018). On the competition complexity of dynamic mechanism design. In *Proceedings of SODA'18*, pages 2008–2025.

[62] Lucier, B. (2017). An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47.

[63] Mehta, A., Saberi, A., Vazirani, U., and Vazirani, V. (2007). Adwords and generalized online matching. *Journal of the ACM*, 54(5):22–es.

[64] Moser, L. (1956). On a problem of Cayley. *Scripta Mathematica*, 22:289–292.

[65] Myerson, R. B. (1981). Optimal auction design. *Math. Oper. Res.*, 6(1):58–73.

[66] Peano, G. (1886). Sull'integrabilità delle equazioni differenziali del primo ordine. *Atti Accad. Sci. Torino*, 21:437–445.

[67] Perez-Salazar, S., Singh, M., and Toriello, A. (2022). The iid prophet inequality with limited flexibility. *arXiv preprint arXiv:2210.05634*.

[68] Perez-Salazar, S., Singh, M., and Toriello, A. (2024). Robust online selection with uncertain offer acceptance. *Mathematics of Operations Research*. Forthcoming.

[69] Phillips, C. A., Stein, C., Torng, E., and Wein, J. (1997). Optimal time-critical scheduling via resource augmentation (extended abstract). In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC)*, pages 140–149. ACM.

[70] Picard, E. (1893). Sur l'application des méthodes d'approximations successives à l'étude de certaines équations différentielles ordinaires. *Journal de Mathématiques Pures et Appliquées*, 9:217–272.

[71] Roughgarden, T. and Tardos, E. (2002). How bad is selfish routing? *Journal of the ACM*, 49(2):236–259.

[72] Rubinstein, A. and Singla, S. (2017). Combinatorial prophet inequalities. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1671–1687. SIAM.

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

[74] Singla, S. (2018). *Combinatorial Optimization Under Uncertainty: Probing and Stopping-Time Algorithms*. PhD thesis, Carnegie Mellon University, Pittsburgh, PA, USA.

[75] Sleator, D. D. and Tarjan, R. E. (1984). Amortized efficiency of list update rules. In *Proceedings of the 16th Annual ACM Symposium on Theory of Computing (STOC)*, pages 488–492. ACM.

[76] Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.

[77] Yan, Q. (2011). Mechanism design via correlation gap. In *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2011)*, pages 710–719, San Francisco, CA, USA. SIAM.

Appendix A

Missing proofs of Chapter 2

In this appendix we show that for sufficiently large values of m, we have $1 - 2\log(m)/m \le A'_m \le 1 - \log(m)/(3m)$. First observe that the expected welfare with m buyers, obtained by a static price of T, can be lower bounded by the expected revenue, which is equal to $T \cdot \mathbb{P}(\max_{j\in\{1\dots,m\}} X_j > T) = T(1 - T^m)$. Then, the optimal welfare with static prices can be lower bounded by the revenue of the static price $T^\star_m$ that maximizes the expected revenue $R(T) = T - T^{m+1}$, which is equal to $T^\star_m = (\frac{1}{m+1})^{1/m}$. In particular, the expected revenue with price $T^\star_m$ satisfies that

$$R(T^\star_m) = T^\star_m - (T^\star_m)^{m+1} = \left(\frac{1}{m+1}\right)^{1/m} - \left(\frac{1}{m+1}\right)^{(m+1)/m}.$$

For every $m \ge 1$ we have $R(T^\star_m) \ge R((1/m)^{1/m}) = (1/m)^{1/m} - (1/m)^{1+1/m}$, and we have $(1/m)^{1+1/m} \le 2/m$. Therefore, $R(T^\star_m) \ge (1/m)^{1/m} - 2/m$. Furthermore, observe that $(1/m)^{1/m} = \exp(-\log(m)/m) \ge 1 - \log(m)/m$, where the last inequality holds since $\exp(-x) \ge 1 - x$ for every $x \ge 0$. Hence, we conclude that the optimal revenue is at least $1 - \log(m)/m - 2/m \ge 1 - 2\log(m)/m$. For a given static price $T > 0$, the expected welfare with m buyers is equal to

$$W(T) = (1 - T^m)T + (1 - T^m)(1 - T)/2 = \frac{1}{2}(1 + T)(1 - T^m) = \frac{1}{2}\left(\frac{1}{T} + 1\right)R(T),$$

If $T > T^\star_m$, we have $R(T) < R(T^\star_m)$, and therefore

$$W(T) = \frac{1}{2}\left(\frac{1}{T} + 1\right)R(T) < \frac{1}{2}\left(\frac{1}{T^\star_m} + 1\right)R(T^\star_m) = W(T^\star_m).$$

Let $\overline{T}_m$ be the maximizer of the welfare W. The previous inequality implies that $\overline{T}_m \leq T^\star_m$, and therefore

$$
\begin{aligned}
w(\overline{T}) &= \frac{1}{2}(1+T^\star_m)(1-T^\star_m) \\
&\leq \frac{1}{2}(1+T^\star_m) = \frac{1}{2}\left(1+\left(\frac{1}{m+1}\right)^{1/m}\right) \leq 1 - \frac{\log(m)}{3m},
\end{aligned}
$$

where the last inequality holds since the function $f(x) = 1 - \frac{\log(x)}{3x} - \frac{1}{2}(1+(\frac{1}{x+1})^{1/x})$ is strictly decreasing in $[1,\infty)]$ and $\lim_{x\to\infty} f(x) = 0$.

Appendix B

Missing proofs of Chapter 3

B.1 Dealing with Point Masses

In this appendix, we discuss the adaptions of our results to the case where the distributions have point masses. To deal with point masses, we can use standard techniques (e.g., [5]) 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 $F_1, \ldots, F_n$, 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 21. *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*

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

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

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

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 (B.1) satisfies that

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

where the inequality is since $n \geq 2$. Thus, Equation (B.1) holds.

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

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

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

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

Proposition 22. *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 $(F_1,\ldots,F_n)$, consider the threshold $\tau = (1-\varepsilon)\cdot\mathbb{E}_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i]$, and let $p = \Pr_{v\sim(F_1,\ldots,F_n)}[\max_{i\in[n]} v_i \leq \tau]$. Then, for $k = \lceil 1/\varepsilon\rceil$, it holds that

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

which concludes the proof. □

B.2 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 $F_1,\ldots,F_n$ 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

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

In the γ-displacement model, we compare $\min_{\sigma \in D_\gamma(n,k)} \mathbb{E}_{v\sim\sigma F^{\gamma k}}[\mathrm{ALG}_{\boldsymbol{\tau}}(v)]$ 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 23. *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 $\boldsymbol{\tau} = (\tau_1,\dots,\tau_k)$ 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 $\boldsymbol{\tau}' = (\tau_1',\dots,\tau_k')$ 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 $\boldsymbol{\tau}' = (\tau_1',\dots,\tau_1',\tau_2',\dots,\tau_2',\dots,\tau_k',\dots,\tau_k')$ 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 3.3 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 24. *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 $F = F_1 \times F_2 \times F_3 \times \ldots \times F_3$.

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

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

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 $\gamma n k$ random variables, into k many γ-blocks of size γn each. We construct a hard instance as follows. The first γ-block is:

$$(\underbrace{1, \ldots, 1}_{2k+1}, 2, 2, 2, \underbrace{3, \ldots, 3}_{\frac{3}{\varepsilon}+2-2k})$$

All the remaining γ-blocks are:

$$(1, 2, 2, 2, \underbrace{3, \ldots, 3}_{\frac{3}{\varepsilon}+2})$$

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 formula[1] we get that

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

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

[1]The 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 $b = \sqrt{q^{3k}/2}$.

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 $\sigma(F_1, \ldots, F_n)^k$, 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 25. *For all $n \geq 2$, $\varepsilon \in (0,1)$ and every $k < 1/(4\varepsilon)$, there exist distributions $F_1, \ldots, F_n$ and a permutation $\sigma \in S_{nk}$ such that for every sequence of thresholds $\boldsymbol{\tau}$ we have*

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

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 $F_1 = \ldots = F_{n-1} = 1$ 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

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

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

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

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

B.3 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 6).

Proposition 26. *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 8, 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 (3.11) 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 $2 - 2\varepsilon$.

B.4 Matching Feasibility Constraints

In this appendix, we study the competition complexity for the vertex arrival model in bipartite graphs with one-sided arrival. This problem falls in the more general XOS combinatorial auctions setting of Section 3.5.2.

More specifically, in this setting there is an underlying bipartite graph $G = (U, V, E)$, and the feasibility constraint is given by the set of matchings in G, that is, $\mathcal{F} = \{S \subseteq E : \text{for all } i \in U \cup V, |\{e \in S : i \in e\}| \leq 1\}$, and the valuation function is additive, i.e., $f(v, S) = \sum_{e \in S} v_e$. The vertices of V arrive online, one by one, and upon their arrival, their edges to all vertices in U are revealed. For every edge $e \in E$, the value v_e is sampled according to a distribution F_e, and we denote $F = \bigtimes_{e \in E} F_e$.

Theorem 7 implies that the $(1-\varepsilon)$-competition complexity of block-consistent prices for this matching setting is $O(\log(1/\varepsilon))$. In what follows, we show how to obtain this result by a different approach based on online contention resolution schemes. Our Algorithm 1 extends the one proposed by [33] for vertex arrival model in bipartite graphs with one-sided arrival[2].

Theorem 11. *For every $k \geq 1$, for every bipartite graph $G = (U, V, E)$, and every $F = \bigtimes_{e \in E} F_e$, Algorithm 1 always returns a matching in G, and it holds that*

$$\mathbb{E}_{v \sim F^k}[\textstyle\sum_{e \in \mathrm{ALG}(v)} v_e] \geq \left(1 - \frac{1}{2^k}\right) \cdot \mathbb{E}_{v \sim F}[\max_{S \in \mathcal{F}} \textstyle\sum_{e \in S} v_e],$$

where $\mathcal{F}$ is the set of matchings in G. In particular, the $(1-\varepsilon)$-competition complexity of Algorithm 1 for the online matching problem with one-sided vertex arrival is $O(\log(1/\varepsilon))$.

Proof. First, observe that the algorithm is well defined since $2 - \sum_{j' < j} x_{(j', u^*)} \geq 2 - \sum_{j'} x_{(j', u^*)} \geq 2 - 1 = 1$. We assume without loss of generality that for every $u \in U$, it holds that $\sum_{j \in V} x_{(j,u)} = 1$. This assumption can be made by adding $|U|$ auxiliary vertices that have edges to all edges in U, and all their edges always have a value of

[2]Ezra et al. [33] showed that the result holds also with respect to non-bipartite graphs.

Algorithm 1: Online contention resolution scheme

```
Initialize ALG ← ∅;
for i ∈ {1,...,k} do
    for j ∈ {1,...,|V|} do
        1. Observe v_{j,u}^{(i)} for every u ∈ U;
        2. Sample ṽ_e ~ F_e for every e ∈ E not incident to j;
        3. Compute a maximum matching μ* in G with edge-weights w as
           follows:
        w_{j,u} = v_{j,u}^{(i)} for every u ∈ U, and w_e = ṽ_e for every other edge;
        4. if u* ∈ U, the partner of j in μ*, exists and is available then
            Match j with u* with probability 1/(2 − ∑_{j'<j} x_{(j',u*)}), where
              x_e = Pr[e ∈ argmax_{S∈F} ∑_{e∈S} v_e] for every e ∈ E;
            Update ALG ← ALG ∪ {(j,u*)};
        end
    end
end
return ALG;
```

zero. It holds that

$$\mathbb{E}_{v\sim F}\left[\max_{S\in\mathcal{F}}\sum_{e\in S}v_e\right]=\sum_{e\in E}\Pr_{v\sim F}\left[e\in\arg\max_{S\in\mathcal{F}}\sum_{e\in S}v_e\right]\cdot\mathbb{E}_{v\sim F}\left[v_e\mid e\in\arg\max_{S\in\mathcal{F}}\sum_{e\in S}v_e\right]. \quad \text{(B.2)}$$

On the other hand, for every $i\in[k]$, and every $e\in E$, it holds that

$$\mathbb{E}_{v\sim F^k}[v_e^{(i)}\mid v_e^{(i)}\in \mathrm{ALG}(v)]=\mathbb{E}_{v\sim F}\left[v_e\mid e\in\arg\max_{S\in\mathcal{F}}\sum_{e\in S}v_e\right], \quad \text{(B.3)}$$

since $(v^{(i)}_{(j,u)})_{u\in U}$, $(\tilde{v}_{(j',u)})_{j'\in V\setminus\{j\},u\in U}$ are distributed the same as $v\sim F$. We refer to time (j,i) to the arrival of the j-th node in block i. In what follows, we prove by induction that for every $u\in U$

$$\Pr[u \text{ is available at time } (j,i)]=\frac{2-\sum_{j'<j}x_{(j',u)}}{2^i}. \quad \text{(B.4)}$$

The base case is when $i=1$, and j is the first vertex to arrive, and then it holds that $(2-\sum_{j'<j}x_{(j',u)})/2^i=(2-0)/2=1$. Assume by induction that it is true if j is not the first vertex to arrive at block i, then for i and $j-1$, and else j is the first to arrive at block $i>1$, then for $i-1$ and $j_{\text{last}}\in V$, where j_{last} is the last vertex to arrive in block $i-1$. If j is not the first vertex of block i, then if we denote by j_{pre} the vertex that arrives before j in block i, then

$$\begin{aligned}
&\Pr[u \text{ is available at time } (j,i)] \\
&= \Pr[u \text{ is available at time } (j_{\text{pre}},i)] \left(1 - \Pr[u \text{ is the partner of } j_{\text{pre}}] \cdot \frac{1}{2 - \sum_{j'<j_{\text{pre}}} x_{(j',u)}}\right) \\
&= \frac{2 - \sum_{j'<j_{\text{pre}}} x_{(j',u)}}{2^i} \left(1 - x_{(j_{\text{pre}},u)} \cdot \frac{1}{2 - \sum_{j'<j_{\text{pre}}} x_{(j',u)}}\right) \\
&= \frac{2 - \sum_{j'<j_{\text{pre}}} x_{(j',u)}}{2^i} \left(\frac{2 - \sum_{j'<j} x_{(j',u)}}{2 - \sum_{j'<j_{\text{pre}}} x_{(j',u)}}\right) = \frac{2 - \sum_{j'<j} x_{(j',u)}}{2^i}.
\end{aligned}$$

Else, if j is the first vertex of block i, then

$$\begin{aligned}
&\Pr[u \text{ is available at time } (j,i)] \\
&= \Pr[u \text{ is available at time } (j_{\text{last}},i-1)] \left(1 - \Pr[u \text{ is the partner of } j_{\text{last}}] \cdot \frac{1}{2 - \sum_{j'<j_{\text{last}}} x_{(j',u)}}\right) \\
&= \frac{2 - \sum_{j'<j_{\text{last}}} x_{(j',u)}}{2^{i-1}} \left(1 - x_{(j_{\text{last}},u)} \cdot \frac{1}{2 - \sum_{j'<j_{\text{last}}} x_{(j',u)}}\right) \\
&= \frac{2 - \sum_{j'<j_{\text{last}}} x_{(j',u)}}{2^{i-1}} \left(\frac{2 - \sum_{j'\leq j_{\text{last}}} x_{(j',u)}}{2 - \sum_{j'<j_{\text{last}}} x_{(j',u)}}\right) = \frac{2 - \sum_{j'<j} x_{(j',u)}}{2^i},
\end{aligned}$$

where the last equality is since $\sum_{j'<j} x_{(j',u)} = 0$, and $\sum_{j'\leq j_{\text{last}}} x_{(j',u)} = 1$, which concludes the proof of the induction. Finally, for every $e \in E$, we have

$$\begin{aligned}
\sum_{i\in[k]} \Pr[v_e^{(i)} \in \text{ALG}(v)] &= x_e \cdot \frac{1}{2 - \sum_{j'<j} x_{(j',u^*)}} \cdot \sum_{i\in[k]} \Pr[u^* \text{ is available at time } (j,i)] \\
&= x_e \cdot \frac{1}{2 - \sum_{j'<j} x_{(j',u^*)}} \cdot \sum_{i\in[k]} \frac{2 - \sum_{j'<j} x_{(j',u^*)}}{2^i} = x_e \cdot \left(1 - \frac{1}{2^k}\right),
\end{aligned} \tag{B.5}$$

where the second equality holds by Equation (B.4). The theorem then holds by combining (B.2), (B.3), and (B.5), together with $\Pr_{v\sim F}[e \in \arg\max_{S\in\mathcal{F}} \sum_{e\in S} v_e] = x_e$. □

Appendix C

Missing proofs of Chapter 4

C.1 Missing Proof from Section 4.3

Proof of Proposition 11. We show that for any $j \in [n]$, $\int_0^1 j\binom{n}{j}(1-u)^{j-1}u^{n-j}F^{-1}(1-u)\,\mathrm{d}u = \mathbb{E}[X_{(j)}]$. This is sufficient since summing over all $j \in \{n-k+1,\dots,n\}$ will then complete the proof. By performing a change of variables $x = F^{-1}(1-u)$, we get

$$\begin{aligned}
&\int_0^1 j\binom{n}{j}(1-u)^{j-1}u^{n-j}F^{-1}(1-u)\,\mathrm{d}u \\
&= \int_\infty^0 j\binom{n}{j}(F(x))^{j-1}(1-F(x))^{n-j}x(-f(x))\,\mathrm{d}x \\
&= \int_0^\infty \frac{n!}{(j-1)!(n-j)!}f(x)(F(x))^{j-1}(1-F(x))^{n-j}x\,\mathrm{d}x = \mathbb{E}[X_{(j)}],
\end{aligned}$$

where $f(x) = F'(x)$. The final equality simply follows from the known fact that the probability density function $f_{X_{(j)}}(x) = n!f(x)(F(x))^{j-1}(1-F(x))^{n-j}/((j-1)!(n-j)!)$. This finishes part (i).

For part (ii), recall that $\mathbb{E}[X|X \geq x]\Pr[X \geq x] = \mathbb{E}[X\mathbb{1}_{\{X\geq x\}}]$. On the other hand, we have that

$$\int_0^q F^{-1}(1-u)\,\mathrm{d}u = \int_\infty^{F^{-1}(1-q)} z(-f(z))\,\mathrm{d}z = \int_x^\infty zf(z)\mathrm{d}z = \mathbb{E}[X\mathbb{1}_{\{X\geq x\}}],$$

where we used the change of variable $z = F^{-1}(1-u)$, and in the second to last equality, we use that $q = \Pr[X \geq x] = 1 - F(x)$. This finished the proof. □

C.2 Missing Proofs from Section 4.4

Proof of Claim 5. By induction $y_{m,j,t-1} > y_{m,j,t}$. Then, we can compare the following ratio

$$\begin{aligned} m \cdot \frac{\Gamma_k(-\ln y_{m,j,t-1}) - \Gamma_k(-\ln y_{m,j,t})}{\Gamma_{k+1}(-\ln y_{m,j,t-1}) - \Gamma_{k+1}(-\ln y_{m,j,t})} &= m \cdot \frac{\int_{-\ln y_{m,j,t-1}}^{-\ln y_{m,j,t}} x^{k-1} e^{-x} \,\mathrm{d}x}{\int_{-\ln y_{m,j,t-1}}^{-\ln y_{m,j,t}} x^{k} e^{-x} \,\mathrm{d}x} \\ &\geq m \inf_{x \in [-\ln y_{m,j,t-1}, -\ln y_{m,j,t}]} \frac{1}{x} \\ &= m \cdot \frac{1}{-\ln y_{m,j,t}} = \frac{m}{\ln m} \geq 1 \end{aligned}$$

From here, the claim follows. □

Proof of Claim 6. By contradiction, assume that $\theta_j > \theta^\star_{j+1}$. Let $t \leq t'$. Now, note that

$$\begin{aligned} &\frac{1}{m}\left(k! - \Gamma_{k+1}(-\ln y_{m,j,t}) - \frac{\theta^\star_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))\right) \\ \geq &\frac{\theta^\star_{j+1}}{\theta_j m}\left(k! - \Gamma_{k+1}(-\ln y_{m,j,t}) - (k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))\right) \\ = &-\frac{\theta^\star_{j+1}}{m\theta_j}\left(\Gamma_{k+1}(-\ln y_{m,j,t}) - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))\right) \end{aligned}$$

On the other hand, using (4.47) and (4.48) and Claim 5, we obtain

$$\begin{aligned} &\frac{1}{m}\left(k! - \Gamma_{k+1}(-\ln y_{m,j,t}) - \frac{\theta^\star_{j+1}}{\theta_j}(k! - \Gamma_{k+1}(-\ln Y_{j+1}(t/m)))\right) \\ \leq &-\frac{\theta^\star_{j+1}}{m\theta_j}(\Gamma_{k+1}(-\ln Y_{j+1}((t-1)/m)) - \Gamma_{k+1}(-\ln Y_{j+1}(t/m))) \end{aligned}$$

From here, we deduce $\Gamma_{k+1}(-\ln Y_{j+1}((t-1)/m)) \leq \Gamma_{k+1}(-\ln y_{m,j,t})$ or equivalently $Y_{j+1}((t-1)/m) \leq y_{m,j,t}$. For $t = 1$, this last inequality implies $y_{m,j,1} \geq 1$ and this is impossible as we always have $y_{m,j,1} < 1$ for any $\theta_j > 0$. We conclude that $\theta_j \leq \theta^\star_{j+1}$. □

Proof of Proposition 12. In what follows, we use the convention that an empty sum is equal to zero. We also avoid writing the limits in the integrals and the differentials "dq" as they are clear from the context. More specifically, we write $\int_0^1 h(q)\,\mathrm{d}q = \int h$

for any integrable function h in $[0,1]$. We simply say that (α, v) is *feasible* if (α, v) is feasible to $[\mathrm{D}]_{n,k}(\bar{n})$. We also use the notation $\bar{n} = \bar{n}_k$ to avoid notational clutter.

We fix $v \geq 0$ such that (α, v) is a feasible solution to $[\mathrm{D}]_{n,k}(\bar{n})$. Let's define the sets

$$J_k^{(\alpha,v)} = \{t \in [\bar{n}_k] : (\alpha, v) \text{ does not tighten constraint (4.32) for } t\}$$
$$J_\ell^{(\alpha,v)} = \{t \in [\bar{n}_k] : (\alpha, v) \text{ does not tighten constraint (4.33) for } t, \ell\}, \quad \ell < k$$

Let $t'_{(\alpha,v)} = \max\{t \in J_1^{(\alpha,v)} \cup \cdots \cup J_k^{(\alpha,v)}\}$. If

$$J_1^{(\alpha,v)} \cup \cdots \cup J_k^{(\alpha,v)} = \emptyset,$$

then we define $t'_{(\alpha,v)} = \bar{n} + 1$; otherwise, $t' \in [\bar{n}]$. Let $p'_{(\alpha,v)} = |\{\ell \in [k] : t'_{(\alpha,v)} \in J_\ell^{(\alpha,v)}\}|$ be the number of constraints of type (4.32)-(4.33) for which the $t'_{(\alpha,v)}$-th constraint is not tight.

Now, among all possible feasible solution (α, v), for a fixed v, choose the one that maximizes $t'_{(\alpha,v)}$. If $t'_{(\alpha,v)} = \bar{n} + 1$, then we are done. Otherwise, let $t' \in [\bar{n}]$ be the maximum value for such a solution. Among all feasible solutions (α, v) such that $t'_{(\alpha,v)} = t'$ choose the one that minimizes $p' = p'_{(\alpha,v)}$. Note that $p' \geq 1$. Now, we will modify (α, v) by finitely many mass transfers and additions yielding a new feasible solution (α', v) such that either $t'_{(\alpha',v)} > t'$ or either $t'_{(\alpha',v)} = t'$ and $p'_{(\alpha',v)} < p'$. In any case, we will obtain a contradiction.

Let's assume first that $t' \in J_k^{(\alpha,v)}$—the general case is handled similarly; we explain at the end the minor changes. We analyze two different cases:

Case 1. If $t' = \bar{n}$, then, we consider the solution $\bar{\alpha}_{t,k} = \alpha_{t,k}$ for $t < \bar{n}$ and $\bar{\alpha}_{\bar{n},k} = \alpha_{\bar{n},k} + \varepsilon \mathbf{1}_{(0,1)}$ with $\varepsilon > 0$ such that $\int \alpha_{\bar{n},k} + \varepsilon + \sum_{\tau<\bar{n}} \int q\alpha_{\tau,k} = 1$. Also, $\bar{\alpha}_{t,\ell} = \alpha_{t,\ell}$ for $\ell < k$. Note that $(\bar{\alpha}, v)$ remains feasible and tightens one more constraint in (4.32); this contradicts our choice of p'.

Case 2. If $t' < n$, we define

$$\bar{\alpha}_{t,k} = \begin{cases} \alpha_{t,k}, & t < t', \\ \alpha_{t',k} + \sum_{\tau > t'} \omega_t \alpha_{t,k}, & t = t', \\ (1 - \omega_t)\alpha_{t,k}, & t > t', \end{cases}$$

where $\omega_{t'+1}, \ldots, \omega_n \in [0,1]$. Let $\bar{\alpha}_{t,\ell} = \alpha_{t,\ell}$ for $\ell < k$. Note that $(\bar{\alpha}, v)$ satisfies (4.34), it satisfies (4.32) for $t < t'$, and for $t > t'$, we have $\int \bar{\alpha}_{t,k} + \sum_{\tau<t} \int q\bar{\alpha}_{\tau,k} = \int (1 - \omega_t)\alpha_{t,k} + \sum_{\tau<t} \int q\alpha_{\tau,k} + \sum_{\tau \geq t} \omega_\tau \int q\alpha_{\tau,k}$, which is increasing in ω_τ for $\tau > t$ and decreasing in ω_t.

We start with the values $\omega_{t'+1},\ldots,\omega_n = 0$ and at this point $(\bar{\alpha},v)$ is feasible. By the choice of t', we have $\int \alpha_{t',k} + \sum_{\tau>t'} \omega_\tau \int \alpha_{\tau,k} + \sum_{\tau<t} \int q\alpha_{\tau,k} \leq 1$ for $\omega_{t'+1},\ldots,\omega_n > 0$ small enough. Now, we increment $\omega_{t'+1}$ as much as possible while keeping feasibility of $(\bar{\alpha},v)$. We repeat the same process in the order $\omega_{t'+2},\ldots,\omega_n$. We note that $\omega_{t'+1},\ldots,\omega_n$ are not all 0's.

Suppose that we have $\int \bar{\alpha}_{t',k} + \sum_{\tau<t'} \int q\bar{\alpha}_{\tau,k} = \int \alpha_{t',k} + \sum_{\tau>t'} \omega_\tau \int \alpha_{\tau,k} + \sum_{\tau<t} \int q\alpha_{\tau,k} < 1$. Then, we claim that $\omega_{t'+1},\ldots,\omega_{\bar{n}} = 1$. Indeed, let $\tau' > t'$ be the smallest τ such that $\omega_\tau < 1$. Then, $\bar{\alpha}_{t,k} = 0$, for $t \in \{t'+1,\ldots,\tau-1\}$. Furthermore, constraints (4.32) for $t \in \{t', t'+1, \ldots \tau - 1\}$ are not tight, because they are dominated by constraint (4.32) for $t = t'$ which is not tight. Since increasing ω_τ does not affect constraints (4.32) for $t \geq \tau$, we can increase slightly ω_τ and contradict the choice of $\omega_{t'+1},\ldots,\omega_{\bar{n}}$. From this analysis, we also deduce that every constraint (4.32) for $t = t',\ldots,\bar{n}$ is not tight. Furthermore, $\bar{\alpha}_{t,k} = 0$ for $t > t'$. Define

$$\hat{\alpha}_{t,k}(q) = \begin{cases} \bar{\alpha}_{t,k}(q)\,(=\alpha_{t,k}(q)), & t < t', \\ \bar{\alpha}_{t,k}(q) + c_t\delta_{\{1\}}(q), & t \geq t', \end{cases}$$

where $\delta_{\{1\}}(\cdot)$ is the Dirac delta at one. We define $\hat{\alpha}_{t,\ell} = \bar{\alpha}_{t,\ell} = \alpha_{t,\ell}$ for $\ell < k$. Note that $(\hat{\alpha},v)$ satisfies constraints (4.33) and constraints (4.34), and constraints (4.32) for $t < t'$. If we define $c_{t'} = 1 - \int \bar{\alpha}_{t',k} - \sum_{\tau<t} \int q\bar{\alpha}_{\tau,k} > 0$ we have that $\hat{\alpha}$ satisfies constraint (4.32) at $t = t'$ with equality. For $t > t'$ we define $c_t = 1 - \sum_{\tau<t} \int q\hat{\alpha}_{\tau,k} \geq 0$. A small computation shows that $(\hat{\alpha},v)$ is again feasible and tightens (4.32) for t'. Furthermore, all the other constraints (4.33) remain unchanged for $t \leq t'$ as they are only affected by terms $\alpha_{j,\tau}$ with $\tau < t'$. This implies that either $p'_{(\hat{\alpha},v)} < p'$ if $p' > 1$ or $t'_{(\hat{\alpha},v)} > t'$ if $p' = 1$. In any case, this leads again to a contradiction to our choice of (α,v).

When $t' \notin J_k^{(\alpha,v)}$, we have $t' \in J_\ell^{(\alpha,v)}$ for some $\ell < k$. In this case, the analysis is the same with the only difference that the constraints will have the value $\sum_{\tau<t} \int q\alpha_{\tau,\ell+1}$ on the right-hand side instead of 1. It is crucial to notice that this value is a non-negative constant when modifying $\alpha_{t,\ell}$; hence, our mass transfers and additions still work. □

Proof of Proposition 17. Part (i) follows directly by computing the derivative:

$$\begin{aligned}
g'_{n,k}(u) &= \sum_{j=n-k+1}^{n} j\binom{n}{j}\left(-(j-1)(1-u)^{j-2}u^{n-j} + (n-j)(1-u)^{j-1}u^{n-j-1}\right) \\
&= \sum_{j=n-k+1}^{n-1} j\binom{n}{j}(n-j)(1-u)^{j-1}u^{n-j-1} - \sum_{j=n-k+1}^{n} j(j-1)\binom{n}{j}(1-u)^{j-2}u^{n-j} \\
&= \sum_{j=n-k+2}^{n} (j-1)\binom{n}{j-1}(n-j+1)(1-u)^{j-2}u^{n-j} - \sum_{j=n-k+1}^{n} j(j-1)\binom{n}{j}(1-u)^{j-2}u^{n-j} \\
&= \sum_{j=n-k+2}^{n} \frac{n!}{(j-2)!(n-j)!}(1-u)^{j-2}u^{n-j} - \sum_{j=n-k+1}^{n} \frac{n!}{(j-2)!(n-j)!}(1-u)^{j-2}u^{n-j} \\
&= -(n-k+1)(n-k)\binom{n}{k-1}(1-u)^{n-k-1}u^{k-1}.
\end{aligned}$$

In the same line, part (ii) follows by evaluating directly $g'_{n+1,k+1}$ using the previous formula:

$$g'_{n+1,k+1}(u) = -(n-k+1)(n-k)\binom{n+1}{k}(1-u)^{n-k-1}u^{k} = \frac{n+1}{k}ug'_{n,k}(u).$$

For part (iii), we have

$$\begin{aligned}
-g'_{n,k}(u) &= (n-k+1)(n-k)\binom{n}{k-1}(1-u)^{n-k-1}u^{k-1} \\
&= \frac{(n-k+1)}{\bar{n}_k^{k-1}}\binom{n}{k-1}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1}(\bar{n}_k+1) \\
&= \frac{(\bar{n}_k+2)}{\bar{n}_k^{k-1}}\frac{n\cdot(n-1)\cdots(n-(k-1)+1)}{(k-1)!}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1}(\bar{n}_k+1) \\
&\le \frac{n}{(k-1)!}\left(\frac{n-k/2}{\bar{n}_k}\right)^{k-1}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1}(\bar{n}_k+1) \\
&\le \frac{n}{(k-1)!}\left(1+\frac{4k^2}{n}\right)(\bar{n}_k u)^{k-1}e^{-\bar{n}_k u}(\bar{n}_k+1).
\end{aligned}$$

The final inequality follows by the bound $(1-u)^x \le e^{-ux}$ for $u \in [0,1]$ and observing that

$$\begin{aligned}
\left(\frac{n-k/2}{\bar{n}_k}\right)^{k-1} = \left(1+\frac{k/2+1}{n-k-1}\right)^{k-1} &\le \left(1+\frac{k+1}{n-(k+1)}\right)^{k+1} \\
&\le \exp((k+1)^2/(n-(k+1)))
\end{aligned}$$

Rewrite $n = (k+1) + c(k+1)^2$ for some $c > 1$. We get $\exp((k+1)^2/(n-(k+1))) = \exp(1/c)$ and we can compute $\left(1 + 4k^2/n\right) \geq 1 + 2/c$. Thus the inequality holds when $\frac{1}{c} \leq \ln(1+2/c)$, which is true for any $c > 2$. That is, letting $n \geq (k+1) + 2(k+1)^2$ suffices. This proves the claim as the above inequality can be slightly strengthened by a factor of $\bar{n}_k/(\bar{n}_k+1)$.

For (iv), we have

$$\begin{aligned}
-g'_{n,k}(u) &= (n-k+1)(n-k)\binom{n}{k-1}(1-u)^{n-k-1}u^{k-1} \\
&\geq \frac{\bar{n}_k^2}{\bar{n}_k^{k-1}}\binom{n}{k-1}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1} \\
&= \frac{\bar{n}_k^2}{\bar{n}_k^{k-1}}\frac{n\cdot(n-1)\cdots(n-(k-1)+1)}{(k-1)!}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1} \\
&\geq \frac{n\cdot\bar{n}_k}{(k-1)!}\left(1-\frac{k}{\bar{n}_k}\right)^{k-1}(\bar{n}_k u)^{k-1}(1-u)^{n-k-1} \\
&\geq \frac{n\cdot\bar{n}_k}{(k-1)!}\left(1-4\frac{k}{n}\right)^{k}(\bar{n}_k u)^{k-1}e^{-\bar{n}_k u/(1-u)} \\
&\geq \frac{n}{(k-1)!}\left(1-4\frac{k^2}{n}\right)\left(1-\frac{\bar{n}_k u^2}{1-u}\right)(\bar{n}_k u)^{k-1}e^{-\bar{n}_k u}\bar{n}_k
\end{aligned}$$

where in the third equality we use that $(1-u)^{-1} = 1 + u/(1-u) \leq \exp(u/(1-u))$, and in last inequality we use $\exp(-\bar{n}_k u/(1-u)) = \exp(-\bar{n}_k u - \bar{n}_k u^2/(1-u)) \geq \exp(-\bar{n}_k u)(1-\bar{n}_k u^2/(1-u))$. Observe that $\Gamma_k(\bar{n}_k u)' = \Gamma'_k(\bar{n}_k u)\bar{n}_k = -(\bar{n}_k u)^{k-1}e^{-\bar{n}_k u}\bar{n}_k$. We conclude by noting that the function $1 - \bar{n}_k x^2/(1-x)$, in $[0,1]$, is decreasing, positive at zero, and it has a unique root in the value $(\sqrt{4\bar{n}_k+1}-1)/(2\bar{n}_k)$. Since this value is larger than $1/(2\sqrt{\bar{n}_k})$, the conclusion follows. □

Proof of Proposition 18. Note that by (4.50), for every $t \in (0,1)$ we have

$$|\Phi'_{k,\ell}(t)| \leq k! + |\Phi_{k+1,\ell}(t)| + \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}k! + \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}|\Phi_{k+1,\ell+1}(t)| \leq 4k!\frac{\theta^\star_{\ell+1}}{\theta^\star_\ell},$$

where the last inequality holds since $\Phi_{k+1,r} \leq k!$ in $(0,1)$ for every $r \in [k]$, and $\theta^\star_\ell < \theta^\star_{\ell+1}$ by Proposition 12(ii). Let $b_k = 4k!\max_\ell \theta^\star_{\ell+1}/\theta^\star_\ell$. Then, since $\Phi_{k,\ell}(1) = 0$, using the Taylor first-order approximation for $\Phi_{k,\ell}$ in one, for every $t \in (0,1)$ we have $\Phi_{k,\ell}(t) \leq b_k(1-t)$. For each $\ell \in [k]$, by the formula $\Gamma_k(x) = (k-1)!\cdot e^{-x}\sum_{r=0}^{k-1} x^r/r!$ applied with $x = -\ln(Y_\ell(t))$ we conclude that $b_k(1-t) \geq \Phi_{k,\ell}(t) = (k-1)!\sum_{r=0}^{k-1} Y_\ell(t)(-\ln Y_\ell(t))^r/r!$, and then $Y_\ell(t)(-\ln Y_\ell(t))^r \leq b_k\cdot r!/(k-1)! \leq b_k$, where the first inequality holds since $Y_\ell(t) \in [0,1]$ for every $t \in (0,1)$. This concludes part (i).

For the second part, by (4.50), for each $\ell \neq k$ we have

$$-\Phi_{k,\ell}(t) = \int_t^1 \Phi'_{k,\ell}(\tau)\,\mathrm{d}\tau = \int_t^1 \left(k! - \Phi_{k+1,\ell}(\tau) - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}(k! - \Phi_{k+1,\ell+1}(\tau))\right)\mathrm{d}\tau$$
$$\leq (1-t)k!\left(1 - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}\right) + \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}\int_t^1 \Phi_{k+1,\ell+1}(\tau)\,\mathrm{d}\tau.$$

For each $\ell \neq k$, choose $\delta_\ell > 0$ such that $\Phi_{k+1,\ell+1}(t) \leq k!(1-\theta^\star_\ell/\theta^\star_{\ell+1})/2$ and $Y_\ell(t)^{1/2}(-\ln Y_k(t))^{k-1} \leq 1$ for $t \in (\delta_\ell, 1)$. Using the first inequality, get

$$(1-t)k!\left(1 - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}\right) + \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}\int_t^1 \Phi_{k+1,\ell+1}(\tau)\,\mathrm{d}\tau$$
$$\leq (1-t)k!\left(1 - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}\right) + (1-t)k!\frac{1}{2}\left(\frac{\theta^\star_{\ell+1}}{\theta^\star_\ell} - 1\right) = -k!(1-t)\frac{1}{2}\left(\frac{\theta^\star_{\ell+1}}{\theta^\star_\ell} - 1\right).$$

Hence for this interval, using the bound $\Phi_{k,\ell}(t) = \Gamma_k(-\ln Y_\ell(t)) \leq k! Y_\ell(t)(-\ln Y_\ell(t))^{k-1}$, we have $Y_\ell(t) \geq (1-t)^2(\theta^\star_{\ell+1}/\theta^\star_\ell - 1)^2/4$. For $\ell = k$, we have

$$(1-t)k!\left(\frac{1}{k\theta_k} - 1\right) \leq \Gamma_k(-\ln Y_k(t)) \leq k! Y_k(t)(-\ln Y_k(t))^{k-1}.$$

By part (i), we know that $Y_k(t) \leq b_k(1-t)$. Hence, for some $\delta_k > 0$, $Y_k(t)^{1/2}(-\ln Y_k(t))^{k-1} \leq 1$ for all $t \in (\delta_k, 1)$. Hence, $Y_k(t) \geq (1-t)^2(1/k\theta_k - 1)^2$. Part (ii) follows by taking $\Delta_k = \max\{\delta_1, \ldots, \delta_k\}$ and $d_k = \min\{(\theta^\star_{j+1}/\theta^\star_j) - 1 : j \in \{1, \ldots, k-1\}\}/4$. Using Lemma 12, we can conclude that $d_k > 0$.

From (4.49) we have $\Phi''_{k,k}(t) = -\Phi'_{k+1,k}(t) \geq 0$, since $\Phi_{k+1,k}$ is non-increasing. For $\ell \neq k$,

$$|\Phi''_{k,\ell}(t)| = |-\Phi'_{k+1,\ell}(t) + \Phi'_{k+1,\ell+1}(t)|$$
$$\leq |\Phi'_{k,\ell}(t)(-\ln Y_\ell(t))| + |\Phi'_{k,\ell+1}(t)(-\ln Y_{\ell+1}(t))|$$
$$\leq b_k\left(-\ln Y_\ell(t) - \ln Y_{\ell+1}(t)\right).$$

Let $N_k = \max\{1/(1-\Delta_k) + 1, 1/d_k\}$. Then, for every $n \geq N_k$ we have $1 - 1/n > \Delta_k$. By the previous part we have that $Y_\ell(1-1/n) \geq d_k n^{-2}$ for all ℓ. Let $c_k = 6b_k$. Since $-\ln Y_\ell(t) - \ln Y_{\ell+1}(t)$ is increasing as a function of t, for every $t \in (0, 1-1/n)$ we have $|\Phi''_{k,\ell}(t)| \leq b_k \cdot 2\ln(n^2/d_k) \leq b_k \cdot 2\ln(n^3) = c_k \ln(n)$. This concludes the proof of part (iii).

For (iv), for $\ell \geq 2$, using a Taylor expansion around zero, for some $\xi \in (0,t)$ and $t < 1 - 1/n$, we have

$$\begin{aligned}\Phi_{k,\ell}(t) &= \Gamma_k(-\ln Y_\ell(0)) + \Gamma_k(-\ln Y_\ell)'(0)t + \frac{1}{2}\Gamma_k(-\ln Y_\ell)''(\xi)t^2 \\ &\geq (k-1)! - \frac{\theta^\star_{\ell+1}}{\theta^\star_\ell}k!t - \frac{c_k \ln(n)}{2}t^2 \\ &\qquad \text{(Using the previous part and } \Gamma_k(-\ln Y_\ell)(0) = 0 \text{ using NLS)} \\ &\geq (k-1)! - \frac{c_k \ln(n)}{2}t^2,\end{aligned}$$

where we used the properties of $\mathrm{NLS}_k(\theta^\star)$ and the definition of b_k. Since $Y_\ell(0) = 1$, for some $\delta_k > 0$ we have that $Y_\ell(t) = 1 - \varepsilon_\ell(t)$ for $t \in [0,\delta]$ with $\varepsilon_\ell(t) \leq 1/2$ for $t \in [0,\delta]$ and $\varepsilon_\ell(t) \to 0$ when $t \to 1$. We simply write $\varepsilon = \varepsilon_\ell(t)$ for convenience. Then, using the characterization of the gamma function Γ_k as a Poisson distribution, we can deduce that

$$\begin{aligned}\frac{c_k \ln(n)}{2}t^2 &\geq \int_0^{-\ln Y_\ell(t)} s^{k-1}e^{-s}\,\mathrm{d}s \\ &\geq \int_0^{\varepsilon} s^{k-1}e^{-s}\,\mathrm{d}s \\ &= (k-1)!\sum_{j\geq k} e^{-\varepsilon}\frac{\varepsilon^j}{j!} \\ &\geq \frac{(k-1)!}{k!}\varepsilon^k e^{-\varepsilon} \geq \frac{1}{2}\frac{(k-1)!}{k!}\varepsilon^k,\end{aligned}$$

where in the second inequality we used that $\ln(1-\varepsilon) \leq -\varepsilon$ and the other inequalities follow by straightforward computations. From here, we obtain that $\varepsilon \leq \bar{c}_k \ln(n)^{1/k}t^{2/k}$, where $\bar{c}_k = (kc_k)^{1/k}$. This concludes (iv). □

Proof of Claim 9. Using a Taylor expansion, we have

$$\Phi'_{k,k}\left(\frac{t}{\bar{n}_k}\right) - \frac{\Phi_{k,k}((t+1)/\bar{n}_k)) - \Phi_{k,k}(t/\bar{n}_k)}{1/\bar{n}_k} = -\frac{1}{2\bar{n}_k}\Phi''_{k,k}(\xi)$$
$$\text{(For some } \xi \in (t/\bar{n}_k, (t+1)/\bar{n}_k))$$

We have $\Phi''_{k,k} = -\Gamma_{k+1}(-\ln y_k)' = -(-\ln y_k)^k y'_k \geq 0$. This concludes the proof of the claim. □

Proof of Claim 11. Using a Taylor expansion, we have

$$\begin{aligned}
&\bar{n}_k(\Phi_{k,\ell}((t-1)/\bar{n}_k) - \Phi_{k,\ell}(t/\bar{n}_k)) + \Phi'_{k,\ell}((t-1)/\bar{n}_k) \\
&= \bar{n}_k \left(\frac{1}{\bar{n}_k}\Phi'_{k,\ell}\left(\frac{t-1}{\bar{n}_k}\right) + \Phi_{k,\ell}\left(\frac{t-1}{\bar{n}_k}\right) - \Phi_{k,\ell}\left(\frac{t}{\bar{n}_k}\right)\right) = -\bar{n}_k \cdot \frac{1}{2\bar{n}_k^2}\Phi''_{k,\ell}(\xi) = -\frac{1}{2\bar{n}_k}\Phi''_{k,\ell}(\xi),
\end{aligned}$$

for some value $\xi \in ((t-1)/\bar{n}_k, t/\bar{n}_k)$. Since $t/\bar{n}_k \leq (\bar{n}_k - 1)/\bar{n}_k = 1 - 1/\bar{n}_k$, by Proposition 18(iii) we have $-\Phi''_{k,\ell}(\xi) \leq c_k \ln(\bar{n}_k)$, which concludes the proof of the claim. □

Proof of Claim 12. We first verify that for every $\ell \in [k]$, for $\bar{n}_k \geq 1/d_k$, and $t \leq \bar{n}_k - 1$ we have $\varepsilon_{\ell,t} \leq 3\ln(\bar{n}_k)/\bar{n}_k$, where d_k is defined in Proposition 18. Indeed, using Proposition 18(ii) we obtain

$$-\ln Y_j(1 - 1/\bar{n}_k) \leq -\ln(d_k)/\bar{n}_k + 2\ln(\bar{n}_k)/\bar{n}_k.$$

For $\bar{n}_k \geq 1/d_k = 4/\min\{\theta^\star_{\ell+1}/\theta^\star_\ell - 1 : \ell \in \{1, \ldots, k-1\}\}$, we obtain the desired result. Then

$$\begin{aligned}
\left(1 - 4\frac{(k+1)^2}{n+1}\right)^{-1}\left(1 - \frac{\bar{n}_k \varepsilon^2_{j+1,t}}{1 - \varepsilon_{j+1,t}}\right)^{-1} &\leq \left(1 + 40\frac{k^2}{n}\right)\left(1 - \frac{\ln(\bar{n}_k)^2}{\bar{n}_k - \ln(\bar{n}_k)}\right)^{-1} \\
&\leq 1 + 10\frac{\ln(\bar{n}_k)^2}{\bar{n}_k},
\end{aligned}$$

which holds for n large. □