Minimization is Harder in the Prophet World*

Vasilis Livanos † Ruta Mehta ‡

Abstract

We study I.I.D. prophet inequalities for *cost minimization*, where the problem is to pick a *cost* from a sequence $X_1, \ldots, X_n$ drawn independently from a known distribution in an online manner, and compete against the prophet who can see all the realizations upfront and select the minimum. In contrast to the well-studied *rewards maximization* setting where a simple threshold strategy achieves a competitive ratio of ≈ 0.745 for all distributions, the cost minimization setting turns out to be much more complex.

In our main result, we obtain a complete and nuanced characterization of the I.I.D. cost prophet inequality: if the expected value of the given distribution is infinite, then the competitive ratio is also infinite. On the other hand, when the expected value is finite, we show that the competitive ratio of the *optimal stopping strategy* is a (distribution-dependent) constant, which we characterize precisely as the solution to a simple inequality. Furthermore, we obtain a closed form for this constant for a broad class of distributions we call *entire distributions*, we show that the constant is 2 for MHR distributions and obtain matching lower bounds for all our results. To the best of our knowledge, these are the first optimal distribution-sensitive guarantees for the prophet inequality setting.

We then focus on single-threshold strategies and design a single threshold that achieves a tight $O(\text{polylog}\, n)$-factor approximation.

In terms of techniques, we characterize the expected value of order statistics using the *hazard rate* of the distribution, which facilitates our analysis. Our results may also be used to design approximately optimal posted price-style mechanisms. We believe both of these may be of independent interest.

Keywords: prophet inequalities, cost minimization, online algorithms, hazard rate, MHR distributions, order statistics

1 Introduction

The classical prophet inequality due to Krengel, Sucheston, and Garling [KS77] concerns the setting where one is presented with *take-it-or-leave-it* rewards $X_1, \ldots, X_n$ in an online manner, drawn independently from known distributions, and can "stop" at any point and collect the last reward seen. Given that the distributions are known, the inequality ensures the existence of a *stopping strategy* S (online algorithm) for any arrival order of the random variables, with expected reward at least half that of a *prophet* who can see the realizations of all the X_i's upfront (offline optimum), *i.e.*, $\mathbb{E}[S] \geq \frac{1}{2}\mathbb{E}[\max_i X_i]$. This result, and its variations and generalizations, have found extensive applications to online optimization and mechanism design, particularly, in the design of simple yet approximately optimal sequential posted price mechanisms, both online and offline, for *revenue (rewards) maximization* while *selling items* [HKS07, CHMS10, CHK07, KW19] (see Section 1.2 for a detailed discussion). A fundamental special case is the setting in which $X_1, \ldots, X_n$ are *independent* and *identically distributed* (I.I.D.). The study of this setting, in the rewards maximization case, was initiated by Hill and Kertz [HK82]. Kertz [Ker86] showed that the competitive ratio in the I.I.D. case approaches ≈ 0.745 as n goes to infinity, via a recursive approach, and conjectured that this is the best bound possible. Later, the bound of ≈ 0.745 was shown to be tight by Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld [CFH+21]. We refer the reader to two excellent surveys [HK92] and [CFH+18] for more results about I.I.D. prophet inequalities.

However, what if the X_i's are *costs* and the goal is *cost minimization*, like in the case of *procuring items* while *minimizing the payment*? For example, consider a house buyer trying to decide when to buy a house in a sellers' market, where houses are selling fast. When a house arrives with its price (cost) listed, she may have to decide the

*The full version of the paper can be accessed at `https://arxiv.org/abs/2209.07988`. Work on this paper partially supported by NSF grant CCF-1750436.

†Dept. of Computer Science, University of Illinois at Urbana-Champaign, IL 61801. `livanos3@illinois.edu`.

‡Dept. of Computer Science, University of Illinois at Urbana-Champaign, IL 61801. `rutameht@illinois.edu`.

same day whether to buy it or not. Given that the buyer may have only distributional knowledge of future house prices, the goal is to devise a buying strategy so that the price paid is minimized.

Towards this, we study the *cost* counterpart of the prophet inequality, where the X_i's represent costs arriving in an online manner, and one *must* "stop" at some point and select the last cost seen. Note that here the constraint is *upwards-closed*, *i.e.,* one of the X_i's *has to be* selected. In particular, if one makes it to X_n, they are forced to pick its realization regardless of how high it is. The goal is to design a stopping strategy (online algorithm) ALG that *minimizes the expected cost*, and is comparable to the cost of an all-knowing prophet who can always select the minimum realization and thus has expected cost $\mathbb{E}[min_i X_i]$. For an $\alpha \geq 1$, we say that algorithm ALG achieves an α-factor *cost prophet inequality*, or is α-*competitive/approximate*, if

$$\mathbb{E}[ALG] \leq \alpha \cdot \mathbb{E}\left[\min_i X_i\right]. \tag{1.1}$$

For the rewards maximization setting, the competitive ratio of $1/2$ in the classical prophet inequality is achievable through simple single-threshold algorithms [SC84, KW19] of the form "accept the first $X_i \geq \tau$ for some threshold τ", and is known to be tight. Furthermore, there exist simple online algorithms that achieve constant-factor approximations even for general multi-dimensional settings with complicated constraints (e.g. matroids, matchings, etc) [KW19, Ala14, JMZ22a, GW23, EFGT22]. Motivated by these works, we ask:

For the cost minimization setting, can we obtain similar results to the rewards setting?
What is the factor achieved by the optimal online algorithm, i.e., the smallest possible α of any online algorithm in (1.1)?
Is the factor achieved a constant? Is it achievable by simple single-threshold algorithms?

In this paper, we study the above questions for the case of independent and identically distributed (I.I.D.) random variables. At first glance, one may wonder why the cost minimization is not equivalent to reward maximization with *negative* X_i's. The reason is that a strategy for rewards maximization is allowed to not pick any of the X_i's, namely downward closed constraints, and hence will pick nothing if X_i's are negative. In contrast, in the cost minimization setting, one of the X_i's has to be selected. This difference turns out to be crucial, as we demonstrate that upwards-closed constraints are harder, and lead to qualitatively different guarantees.

Our main contribution is a complete and nuanced characterization of the best-possible competitive ratio, namely the one achieved by the optimal stopping strategy. Contrary to the rewards maximization setting, in which the optimal online algorithm achieves a tight ≈ 0.745-competitive ratio for all distributions [HK82, Ker86, CFH+21], bounding the competitive ratio for cost minimization depends on whether the distribution has a finite expected value or not.

First, if the expected value of the given distribution is infinite, e.g. the equal-revenue distribution[1], then there exists a family of instances, one for every $n \geq 2$, for which the competitive ratio is infinite (Proposition A.1 due to Lucier [Luc]). This is because, for any algorithm, regardless of which realization it selects, their expected cost will be infinite, whereas the expected minimum cost is finite. This prevents any bounded factor approximation for *all* distributions.

Second, for distributions with finite expected value, we provide a characterization of the tight competitive ratio, and show it is a (distribution-dependent) constant. We demonstrate how to easily calculate it for any distribution and also show it has a closed form for a broad class of distributions which we call *entire distributions*. This class contains several well-known distributions, including the uniform, exponential, Gaussian, arcsine, beta, gamma, Rayleigh and Weibull distribution, among others. Moreover, we show that the competitive ratio for all entire MHR distributions is 2, and is tight, providing a uniform bound to complement our distribution-dependent main result. Our optimal distribution-sensitive guarantees may be of interest in initiating a similar analysis for the rewards maximization and other settings.

We then focus on single-threshold strategies and design a single threshold that achieves a tight $O(\text{polylog}\, n)$-factor approximation for entire distributions, where, the exponent of the logarithm depends on the distribution. We also briefly discuss how our results may also be used to design approximately optimal posted price-style mechanisms.

[1]The CDF of the equal-revenue distribution is $F(x) = 1 - 1/x$ and is supported on $[1, +\infty)$.

In what follows, we give an overview of our approach and techniques, which in turn provides several characterizations of the first moment of order statistics via the hazard rate of their distribution. These may be of independent interest.

1.1 Overview of our Approach and Techniques: Analysis through Hazard Rate In the I.I.D. cost prophet inequality setting, all the X_i's are drawn from a common non-negative known distribution $\mathcal{D}$. Since single-threshold algorithms have proven very successful in providing constant-factor approximations in the rewards setting, we start by asking whether we can achieve similar results for cost prophet inequalities. The intuition behind this is that if n is very large, one could set a single threshold close to $\mathbb{E}[\min_i X_i]$ and with good probability there will be at least one realization below the threshold.

Unfortunately, this intuition turns out to be wrong, even for simple distributions, like the exponential. We present an example in Appendix A explaining why this intuition fails. In fact, in Section 4 we show that no single threshold can achieve a constant competitive ratio; we discuss what one can achieve by a single threshold later.

Thus, in our search for a constant-factor competitive algorithm, we study *optimal online algorithms*, i.e., algorithms that achieve the smallest possible α in (1.1). It follows that, much like the classical I.I.D. prophet inequality [HK82, CFH+21, LLP+21], it suffices to only consider threshold-based algorithms with oblivious thresholds: set thresholds $\tau_1, \ldots, \tau_n$ upfront, and accept the first $X_i \leq \tau_i$. Intuitively, this is because the process is *memory-less*, i.e. the decision in the i-th round is independent of the past realizations and depends only on the realization of X_i and the distribution of the future costs.

Since there cannot be a bounded competitive ratio for all distributions, as discussed earlier (unbounded ratio for the *equal revenue distribution*), the natural next step is to search for the largest class of *tractable distributions.*

As it turns out, the analysis of the optimal online algorithm is quite tricky, a fact that mirrors the involved analysis in the rewards maximization setting [CFH+21]. We tackle this issue via a new technique that utilizes the hazard rate of the given distribution.

Hazard Rate. For a given distribution $\mathcal{D}$ supported on $[a, b)$ (where $a \geq 0$ and b can be infinite), with probability density and cumulative distribution functions f and F respectively, the *hazard rate* of $\mathcal{D}$ is defined for all x in the support of $\mathcal{D}$ as $h(x) \triangleq \frac{f(x)}{1-F(x)}$. Also referred to as the *failure rate*, it is a fundamental quantity within several fields of economics and mathematics and has found a lot of applications in survival analysis [KP02], reliability theory [RH04], pricing [HR09, GPZ21, BBDS17, BGGM12, CD15, DW12, DRY15, GKL17] and even forensic analysis [KAnA11]. For our results, we utilize the antiderivative of the hazard rate, $H(x) = \int_0^x h(z)\, dz$, often called the *cumulative hazard rate* of $\mathcal{D}$.

Analysing the Optimal Online Algorithm. Via the cumulative hazard rate of $\mathcal{D}$, we are able to bridge the gap between the expressions of the prophet's cost and the cost of the optimal online algorithm. Let $\mathcal{P}_n$ and $\mathcal{A}_n$ denote the cost of the prophet and the optimal online algorithm on n random variables, respectively. In particular, we show that

$$\mathbb{E}[\mathcal{P}_n\text{'s cost}] = a + \int_a^b e^{-nH(x)}\, dx$$

whereas

$$\mathbb{E}[\mathcal{A}_n\text{'s cost}] = a + \int_a^{\mathbb{E}[\mathcal{A}_{n-1}\text{'s cost}]} e^{-H(x)}\, dx.$$

The main difficulty then is to show that

$$\exists \text{ a constant } \lambda, \text{ independent of } n, \text{ for which } \quad \frac{a + \int_a^{\mathbb{E}[\mathcal{A}_{n-1}\text{'s cost}]} e^{-H(x)}\, dx}{a + \int_a^b e^{-nH(x)}\, dx} \leq \lambda.$$

Since to achieve our objective, we need to analyze $\mathbb{E}[\mathcal{P}_n\text{'s cost}]$, it is useful to introduce some terminology from extreme value theory (for more information see [HF06]). Let $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$ denote the non-decreasing rearrangement of $X_1, X_2, \ldots, X_n$. The $X_{(i)}$ is called the *ith order statistics of* $\mathcal{D}$. Let μ_n denote the expected value of $X_{(1)}$, *i.e.,* the expected minimum of n I.I.D. random variables drawn from $\mathcal{D}$.

Via an upper bound on $\mathbb{E}[\mathcal{A}_n\text{'s cost}]$ that depends on $H^{\leftarrow}(x)$, the (generalized) inverse of $H(x)$, we reduce the problem of calculating λ to the analysis of the tail behaviour of $\frac{\mu_n}{\mu_{n-1}}$, as $n \to \infty$ and this ratio converges to 1. We capture this tail behaviour via a function $\phi(\lambda, n)$.

Utilizing the hazard rate again, we derive in Lemma 3.4 a simple expression for $\phi(\lambda, n)$ and show in Lemma 3.3 that $\phi(\lambda, n)$ is increasing in λ. We then use the fact that the moments of order statistics, and in particular the expected minimum, of a distribution satisfy nice recurrence properties (Lemma 3.6). Using this, we precisely characterize λ via a simple inequality that needs to be satisfied (Lemma 3.5), namely

$$n\,\phi(\lambda, n) \geq \frac{\mu_{2:n}}{\mu_n} - 1, \tag{1.2}$$

where $\mu_{2:n}$ denotes the expected value of the second smallest random variable, out of n random variables following $\mathcal{D}$, i.e. the first moment of the second order statistic of $\mathcal{D}$. Finally, notice that $\frac{\mu_{2:n}}{\mu_n} - 1$ is upper bounded by $\frac{\mu_{2,2}}{\mu_2} - 1$, a constant. By analyzing the hazard rate of the first order statistic, we utilize the fact that it is concentrated around its expected value to show that the competitive ratio is a constant independent of n, it is equal to the infimum among all $\lambda > 1$ that satisfy (1.2) and satisfies this, simplified, version of (1.2):

$$\lambda - \frac{F^{\leftarrow}\left(\frac{c\left(1+\frac{1}{\lambda-1}\right)}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} \geq 0, \tag{1.3}$$

where $F^{\leftarrow}(x)$ denotes the generalized inverse of the CDF F of $\mathcal{D}$, otherwise known as the *quantile function* (see Section 2 for its definition).

This leads us to our main theorem.

THEOREM 1.1. *Let $\mathcal{D}$ be a distribution supported on $[a, b]$ or $[a, +\infty)$, where $a \geq 0$, with cdf F, quantile function $F^{\leftarrow}$ and for which $\mu_1 = \int_a^b (1 - F(x))\,dx < +\infty$. There exist constants $n_0 \in \mathbb{N}$ and λ such that, for every $n \geq n_0$, there exists a λ-factor cost prophet inequality for the I.I.D. setting under $\mathcal{D}$. λ can be characterized as*

$$\lambda := \inf_{\lambda' > 1} \left\{ \lambda' \;\middle|\; \lambda' - \frac{F^{\leftarrow}\left(\frac{c\left(1+\frac{1}{\lambda'-1}\right)}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} \geq 0, \quad \forall n \geq n_0 \right\}, \qquad \text{for } c = \frac{\mu_{2:2}}{\mu_2} - 1.$$

And this bound is tight.

The hazard rate function is known to be useful in quantifying how "heavy-tailed" a distribution is [BPH96]. Thus, the observant reader might wonder whether λ is related to the *extreme value index* γ from the celebrated Fisher-Tippett-Gnedenko theorem in the field of extreme value theory (for more information see [HF06]), which is also known to capture information about the tail of a distribution. We demonstrate that this is not the case, and the extreme value index cannot characterize the competitive ratio on its own, in Proposition A.2, which can be found in Appendix A.

We present, in Table 1 the explicit competitive ratios computed via Theorem 1.1 for a selection of common distributions.

Distribution Class	$F(x)$	$H(x)$	Competitive Ratio
Exponential	$1 - e^{-x}$	x	2
Uniform	x	$\log\left(\frac{1}{1-x}\right)$	2
Weibull	$1 - e^{-x^2}$	x^2	$\sqrt{\frac{6}{\pi}}$
Pareto ($\alpha > 1$)	$1 - \left(\frac{c}{x}\right)^{\alpha}$	$\alpha \log \frac{x}{c}$	1[a]
Pareto ($\alpha \leq 1$)	$1 - \left(\frac{c}{x}\right)^{\alpha}$	$\alpha \log \frac{x}{c}$	$+\infty$[b]

Table 1: A summary of the competitive ratio for several common distributions.

[a]By $\lambda = 1$, we denote the fact that asymptotically, as $n \to \infty$, $\frac{\mathbb{E}[\mathcal{A}_n\text{'s cost}]}{\mathbb{E}[\mathcal{P}_n\text{'s cost}]}$ goes to 1, i.e. for any $c > 0$, there exists a $n_0 \in \mathbb{N}$ such that one can obtain a $(1 + c)$-approximation to the prophet's cost for any $n \geq n_0$.

[b]The Pareto distribution has infinite expected value for $\alpha \leq 1$.

Entire Distributions and Closed Form. Interestingly, we are able to strengthen Theorem 1.1, providing a closed form of the optimal competitive ratio, for a broad class of distributions, which we call *entire*[2]. A distribution $\mathcal{D}$ is called *entire* if the cumulative hazard rate function H of $\mathcal{D}$ has a series expansion[3] $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i} \neq 0$, where $0 < d_1 < d_2 < \dots$, and this series is convergent everywhere in the support of $\mathcal{D}$. This class includes almost all commonly used distributions, including the uniform, exponential, Gaussian, Weibull, Rayleigh, arcsine, beta and gamma distributions, among many others.

In particular, we show that the competitive ratio is $\lambda(d_1)$, a function that depends only on d_1. Furthermore, $\lambda(d_1)$ is a decreasing function, a fact which is intuitive since, as d_1 increases, H grows rapidly, and thus $\mathcal{D}$ has a less heavy tail which leads to a better approximation.

THEOREM 1.2. *Let $\mathcal{D}$ be an entire distribution supported on $[a, b]$ or $[a, +\infty)$, where $a \geq 0$, with cdf F. There exists $n_0 \in \mathbb{N}$ such that, for every $n \geq n_0$, there exists a $\lambda(d)$-factor cost prophet inequality for the I.I.D. setting under $\mathcal{D}$, where*

$$\lambda(d) = \frac{(1 + 1/d)^{1/d}}{\Gamma(1 + 1/d)},$$

d is the smallest degree of the series expansion of H, and $\Gamma(\cdot)$ is the Gamma function.

Moreover, this constant is tight for the distribution with cumulative hazard rate $H(x) = x^d$.

For completeness and to assist the reader, we have included some relevant background on the Gamma function in Appendix C.

To understand how $\lambda(d)$ grows with d, consider Stirling's approximation for the Gamma function, $\Gamma(z) \approx \frac{\sqrt{2\pi}}{z}\left(\frac{z}{e}\right)^z$. Replacing this in the expression of $\lambda(d)$, we have

$$\lambda(d) = \frac{(1 + 1/d)^{1/d}}{\Gamma(1 + 1/d)} \approx \frac{(1 + 1/d)^{1/d}}{\frac{\sqrt{2\pi}}{1+1/d}\left(\frac{1+1/d}{e}\right)^{1+1/d}} = \frac{e}{\sqrt{2\pi}}\ e^{1/d}.$$

Thus the dependence of $\lambda(d)$ on d is (approximately) inversely exponential.

An interesting question is whether this phenomenon of the competitive ratio being arbitrarily large is due to some technicality that exists for distributions with unbounded support. We answer this negatively in Appendix A.4, where, for any $\alpha > 0$, we provide a family of distributions, each supported on $[0, 1]$, for which the competitive ratio of the optimal algorithm is exactly $\lambda(\alpha)$.

MHR Distributions. Distributions with monotonically increasing hazard rate have been extensively studied in the mechanism design literature due to their sought after properties and applications (e.g., see [GPZ21, BBDS17, BGGM12, CD15, DW12, DRY15, GKL17, HR09]). These are known as *monotone hazard rate (MHR)* distributions. For entire MHR distributions, we are able to show that $d \geq 1$, and since $\lambda(1) = 2$, we show that the optimal algorithm is 2-factor competitive. In addition, we show that the factor of 2 is tight for the exponential distribution, which has constant hazard rate.

THEOREM 1.3. *For every entire MHR distribution, there exists a 2-competitive cost prophet inequality, for large enough n.*

This factor is tight, since there is no $(2 - \varepsilon)$-cost prophet inequality for any $\varepsilon > 0$ for the exponential distribution, which has constant hazard rate.

Single-Threshold Algorithms. Given the success of single-threshold algorithms in the rewards maximization setting, where they are able to achieve the best-possible competitive ratio for non-I.I.D. random variables and they are only a small constant away from the best-possible competitive ratio in the I.I.D. case, we ask whether there exists a single-threshold algorithm that achieves a constant-factor competitive ratio for the cost prophet inequality setting as well. The answer turns out to be negative.

We show that, for entire distributions, no single-threshold algorithm can achieve a better than poly-logarithmic competitive ratio. We also obtain a matching upper bound. In particular, given an entire distribution $\mathcal{D}$, by

[2]The term is analogous to the notion of entire functions, which are functions that are analytic everywhere, i.e. their series expansion converges everywhere.

[3]Notice that the exponents d_i need not be integers.

analyzing how the competitive ratio of a single-threshold algorithm grows as n increases, we design a threshold T such that the algorithm that selects the first $X_i \le T$ for $i < n$ and X_n otherwise, yields a $\mathrm{O}(\operatorname{polylog} n)$-factor cost prophet inequality. Here, the power in the poly-logarithmic factor inversely depends on the smallest degree of the series expansion of H.

THEOREM 1.4. *Given $X_1, \ldots, X_n$ drawn independently from a non-negative entire distribution $\mathcal{D}$, there exists a single-threshold algorithm that is $\mathrm{O}(\operatorname{polylog} n)$-competitive, for large enough n. Moreover, this factor is tight, i.e. there exist distributions for which no single-threshold algorithm is $\mathrm{o}(\operatorname{polylog} n)$-competitive.*

Application to Mechanism Design. Finally we note that, similar to the extensive application of classical prophet inequalities in designing simple yet approximately optimal posted-price mechanisms for selling items (see Section 1.2), our algorithms and results for the *cost prophet inequality* can be used for the design and analysis of posted-price-style mechanisms for *procuring items*.

Consider a procurement auction (also known as a reverse auction), in which the auctioneer (buyer) wants to procure a single item sold by n different sellers, with an I.I.D. distribution governing the sellers' valuation for selling the item to the auctioneer or, in other words, the sellers' costs/price. If the sellers arrive in an online manner with take-it-or-leave-it offers, for example as is the case in a seller's housing market, then the standard reduction of a posted-price mechanism to a prophet inequality due to Hajiaghayi, Kleinberg and Sandholm [HKS07] applies directly to the cost setting, if one wants to minimize the social cost.

To minimize the *procurement price* paid by the buyer (auctioneer), one simply needs to use the *virtual costs* $\phi(c) = c + \frac{F(c)}{f(c)}$, since Myerson's optimal auction [Mye81] applies to any single-parameter environment. This holds only if $\mathcal{D}$ is a regular distribution (a class of distributions which includes MHR distributions among others). For non-regular $\mathcal{D}$, one simply needs to "iron" the social cost function, just as in the classical rewards setting, and proceed similarly afterwards. For more details on this see [MS83] (Theorem 1).

1.2 Related Work Given the intractability of the optimal (revenue-maximizing) mechanisms for selling items [BCKW15, HN19, DDT15, DDT14], the focus turned to designing *approximately optimal* yet simple mechanisms where prophet inequalities for *reward maximization* have been extensively studied. The works of Hajiaghayi, Kleinberg and Sandholm [HKS07] and Chawla, Hartline, Malec and Sivan [CHMS10] pioneered the use of prophet inequalities to analyze (sequential) posted price mechanisms for *selling items*. Specifically, [HKS07] observed that the problem of designing posted price mechanisms that maximize welfare can be reduced to an appropriate optimal stopping theory problem, and this was extended to revenue-maximizing posted price mechanisms in [CHMS10]. This result led to a significant effort to understand how the expected revenue of an optimal posted price mechanism compares to that of the optimal auction [CHK07, Yan11, BH08, ABF+17, Ala14, FGL15, DFKL20, BBK21, Dob21, DV16, AKS21, AS20, DKL0]. In a surprising result, Correa, Foncea, Pizarro and Verdugo [CFPV19] showed that the reverse direction also holds, establishing an equivalence between finding stopping rules in an optimal stopping problem and designing optimal posted price mechanisms – for more information on these applications see a survey by Lucier [Luc17]. Recently, [CCD+23] initiated the study of buy-and-sell prophet inequalities, named *trading prophets*, and obtains constant factor guarantees even with single-threshold algorithms.

Our work is most closely related to the long line of work that considers the case of I.I.D. random variables drawn from a known distribution, which dates back to Hill and Kertz [HK82]. As stated previously, Kertz [Ker86] showed that the competitive ratio in the I.I.D. case approaches ≈ 0.745 as n goes to infinity and conjectured its tightness. A simpler proof of this can be found in [SM02]. The bound of ≈ 0.745 was shown to be tight by Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld [CFH+21]. The proofs of both the upper and lower bounds were recently simplified, by [JMZ22b] and [LLP+21] respectively. Esfandiari, Hajiaghayi, Liaghat and Monemizadeh [EHLM17] considered prophet inequalities for cost minimization in the non-I.I.D. setting, and showed that no bound exists on the competitive ratio. Furthermore, the *large market* domain of our paper, where the distribution is independent of the number of random variables, is well-established and has also been studied for the rewards maximization setting [AEE+17, CFPV19, ABG+20].

The $1/2$-competitive factor guaranteed by the classical prophet inequality for adversarial arrival order has been shown to hold for more general classes of downwards-closed constraints, all the way up to matroids [KW19]. For the special case of k-uniform matroids, where one can select up to k values, Alaei [Ala14] showed a $\left(1 - \frac{1}{\sqrt{k+3}}\right)$-competitive ratio. This was recently improved for small k via the use of a static threshold [CDL21] and later made tight for all k [JMZ22a]. Ezra, Feldman, Gravin and Tang [EFGT22] showed a 0.337-prophet inequality for matching constraints. Rubinstein [Rub16] considered general downwards-closed feasibility constraints and obtained logarithmic approximations. The standard setting has been extended to combinatorial valuation functions [RS17, CL21], where one can obtain a constant competitive ratio when maximizing a submodular function but a logarithmic hardness is known for subadditive functions [RS17]. Recently, [CGKM20] studied non-adaptive threshold algorithms for matroid constraints and gave the first constant-factor competitive algorithm for graphic matroids. Qin, Rajagopal, Vardi and Wierman [QRVW19] study the related problem of *convex prophet inequalities*, in which, instead of costs, one sequentially observes random *cost functions* and needs to assign a total mass of 1 across all functions in an online manner. Our work is a special case of their work, for constant cost functions and identical distributions, which is why we are able to obtain constant competitive ratios compared to the ratio obtained in [QRVW19], which is polynomial in the number of cost functions n.

Esfandiari, Hajiaghayi, Liaghat and Monemizadeh [EHLM17] introduced *prophet secretary*, in which the arrival order of the random variables is chosen uniformly at random, instead of by an adversary. They gave an adaptive-threshold algorithm that achieves a $(1-1/e)$-competitive ratio and showed no algorithm can achieve a factor better than 0.75. Ehsani, Hajiaghayi, Kesselheim and Singla [EHKS18] extend this result to matroid constraints. The factor of $1-1/e$ was recently beaten, first for the case where the algorithm is allowed to choose the arrival order, called the *free order setting* [AEE+17] and later for random arrival order [ACK18]. The best currently known ratio is obtained by Correa, Saona and Ziliotto [CSZ20], where they also improve the upper bound to 0.732. When one can select up to k values, Arnosti and Ma [AM22] recently gave a surprising and quite beautiful single-threshold algorithm that achieves the best competitive ratio of $1-e^{-k}\frac{k^k}{k!}$. More general feasibility constraints have also been studied in the random arrival order case, i.e. for matroids [AW20] and matchings [BGMS21, PRSW22].

Several of these results are described in the context of an *online contention resolution scheme (OCRS)*, which is an algorithm used to round a fractional solution of a linear program in an online manner. Originally introduced by Chekuri, Vondrák and Zenklusen [CVZ14] for the offline case, Feldman, Svensson and Zenklusen [FSZ21] showed the existence of constant-factor approximate OCRSs for several classes of interesting constraints and demonstrated that an α-approximate OCRS for a constraint implies an α-competitive prophet inequality for the same constraint. This connection was proved to be deeper, as Lee and Singla [LS18] used *ex-ante prophet inequalities* to design optimal OCRSs for matroids. Recently, in a beautiful series of works, Dughmi [Dug20, Dug22] showed that the design of particular (offline) contention resolution schemes is equivalent to another problem in optimal stopping theory, the matroid secretary problem. Whether such connections exists between cost prophet inequalities and OCRSs for upwards-closed constraints is an interesting open question.

Organization. Section 2 introduces the cost prophet inequality setting and contains relevant definitions as well as important observations. Section 3 characterizes the optimal algorithm and contains our main characterization results. In Section 4, we focus on single-threshold algorithms and design a fixed threshold that yields a tight $\mathrm{O}(\operatorname{polylog} n)$-competitive cost prophet inequality for entire distributions. Finally, we conclude with some interesting open problems in Section 5.

Due to space constraints and to improve the readability, all missing proofs can be found in Appendix B, several examples which motivate our analysis can be found in Appendix A and some technical background about the Gamma function that we use, can be found in Appendix C.

2 Preliminaries

In this section we formalize the *I.I.D. cost prophet inequality* setting, and define several important quantities. We are given as input a distribution $\mathcal{D}$ supported on $[a,b) \subseteq [0,+\infty)$, where b can be infinite, and we sequentially observe the independent realizations of n random *costs* $X_1,\ldots,X_n \sim \mathcal{D}$. We *must* "stop" at some point and take the last cost seen. In particular, at any point after observing an X_i, we can choose to select or discard it. If we select X_i, then the process ends and we receive a cost equal to X_i. Otherwise X_i gets *discarded* forever and the process continues. An all-knowing prophet, who can see the realizations of all X_i's upfront can always select the minimum realized cost and hence their expected cost is

$$\text{Offline-OPT} = \mathbb{E}\left[\min_i X_i\right].$$

Let $F:[0,+\infty)\to[0,1]$, where $F(x)=\Pr_{X\sim\mathcal{D}}[X\le x]$, and $f:[0,+\infty)\to[0,1]$ denote the *Cumulative*

Distribution Function (CDF) and *Probability Density Function (PDF)* of $\mathcal{D}$, respectively. Given $\mathcal{D}$, the goal is to design a *stopping strategy* that minimizes the expected cost. That is, design an (online) algorithm ALG to decide when to "stop" and select the last cost seen, such that the expected cost incurred is minimized, and ideally is comparable to the prophet's cost. Formally, for $\alpha \geq 1$, we say that ALG is α-factor approximate/competitive, or achieves an *α-cost prophet inequality* if

$$\mathbb{E}\left[ALG\right] \leq \alpha \cdot \mathbb{E}\left[\min_i X_i\right] = \alpha \cdot \text{Offline-OPT}. \tag{2.4}$$

Given an algorithm $\mathcal{A}$, let $G_{\mathcal{A}}(i)$ denote its expected cost, when it observes i I.I.D. random variables drawn from $\mathcal{D}$. Thus, the expected cost of $\mathcal{A}$ is denoted by $\mathbb{E}\left[\mathcal{A}\right] = G_{\mathcal{A}}(n)$. Let $X_{(1)} \leq X_{(2)} \leq \cdots \leq X_{(n)}$ denote the order statistics of $X_1, X_2, \ldots, X_n$ (i.e. the non-increasing rearrangement of $X_1, X_2, \ldots, X_n$). We denote the expected value of $X_{(i)}$ by $\mu_{i:n}$. Since our main focus is the expected minimum of $X_1, \ldots, X_n$, i.e. $X_{(1)}$, we write $\mu_n \coloneqq \mu_{1:n}$ for brevity. Notice that μ_n is exactly the expected cost of the prophet who can always select the minimum of the n realizations. Let $R_{\mathcal{A}}(n)$ denote the competitive ratio of $\mathcal{A}$ against the prophet's cost μ_n, i.e. $R_{\mathcal{A}}(n) = \frac{G_{\mathcal{A}}(n)}{\mu_n}$. Whenever the algorithm is clear from context, we drop the subscript and just use $G(n)$ and $R(n)$.

The following well-known observation characterizes μ_n.

Observation 2.1. *For $n \geq 1$,*

$$\mu_n = \mathbb{E}\left[\min_{i=1}^{n} X_i\right] = \int_0^\infty \left(1 - F(s)\right)^n ds.$$

Hazard Rate. All of our results make heavy use of the *hazard (failure) rate* of a distribution. We refer the reader to [BPH96] for an extensive overview. Intuitively, for discrete distributions, the hazard rate at a point t represents the probability that an event occurs at time t, given that the event has not occurred up to time t. For continuous distributions, the hazard rate instead quantifies the instantaneous rate of the event's occurrence at time t.

Definition 2.1. (Hazard Rate) *For a distribution $\mathcal{D}$ with cumulative distribution function F and probability density function f, the* hazard rate *of $\mathcal{D}$ is defined as*

$$h(x) \triangleq \frac{f(x)}{1 - F(x)},$$

for all x in the support of $\mathcal{D}$. Furthermore, let H denote the antiderivative of h, which we call the cumulative hazard rate *of $\mathcal{D}$,*

$$H(x) \triangleq \int_0^x h(u)\, du.$$

Next, we express μ_n in terms of the hazard rate of $\mathcal{D}$.

Observation 2.2.

$$\mu_n = \int_0^\infty e^{-nH(u)}\, du.$$

Proof. Notice that,

$$H(x) = \int_0^x h(u)\, du = \int_0^x \frac{f(u)}{1 - F(u)}\, du = -\int_0^x \left(\ln\left(1 - F(u)\right)\right)' du = -\ln\left(1 - F(x)\right),$$

which implies that $1 - F(x) = e^{-H(x)}$, and thus, from Observation 2.1, we have $\mu_n = \int_0^\infty e^{-nH(u)}\, du$. ∎

Distributions with monotonically increasing hazard rate have found a special place within mechanism design literature, originally introduced for the study of revenue maximization. They are known as MHR (or IFR) distributions.

Definition 2.2. (Monotone Hazard Rate Distribution) *A distribution $\mathcal{D}$ is called a* Monotone Hazard Rate (MHR) *distribution if and only if the hazard rate function h (Definition 2.1) of $\mathcal{D}$ is monotonically increasing.*

Quantile Functions. For a monotone function f from (a, b) (where b may be infinite) to (c, d), its generalized inverse $f^{\leftarrow} : (c, d) \to (a, b)$ is defined as

$$f^{\leftarrow}(y) = \inf \{x : a < x < b, f(x) \geq y\}.$$

For more properties of this transformation, see [HF06]. In particular, if F denotes the cdf of a distribution $\mathcal{D}$, then $F^{\leftarrow}(y)$ denotes the quantile function of $\mathcal{D}$, i.e. $F^{\leftarrow}(y)$ is the smallest value τ for which $\Pr_{X\sim\mathcal{D}}[X \leq \tau] \geq y$.

Entire Distributions. While our main result holds for all distributions with finite expected value, we are able to provide an explicit closed form of the competitive ratio for a broad class of distributions, which we call *entire distributions* and define below. This class contains several well-known distributions, including the uniform, exponential, Gaussian, arcsine, beta, gamma, Rayleigh and Weibull distribution, among others.

Before we proceed, we need a generalization of the concept of a Taylor series for a function.

Definition 2.3. (Puiseux Series) *We say that a function $f : \mathbb{R} \to \mathbb{R}$ has a Puiseux series expansion if there exist integers $n > 0$ and $i_0 \in \mathbb{Z}$ as well as coefficients $a_1, a_2, \ldots$ where $a_1 \neq 0$, such that*

$$f(x) = \sum_{i=i_0}^{\infty} a_i x^{i/n}.$$

In other words, Puiseux series are a generalization of Taylor series in that they allow for fractional exponents in the indeterminate, as long as they have a bounded denominator. For the remainder of the paper, we will use the simpler form $f(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$ to denote the Puiseux series of a function f, with the understanding that $d_1, d_2, \ldots$ have a bounded denominator. Furthermore, we only focus on Puiseux series for functions $f : [0, +\infty) \to [0, +\infty)$.

The smallest exponent of the indeterminate in the Puiseux series, $d_1 = {}^{i_0}/_{n}$ is called the *valuation* of f and plays a significant role in our results for entire distributions. The *radius of convergence* of a Puiseux series around 0 is the largest number $r \geq 0$ such that the series converges if x is substituted for a non-zero real number $t \leq r$. A Puiseux series is *convergent* at a point x if $x \leq r$.

Now we are ready to define the class of entire distributions.

Definition 2.4. (Entire Distribution) *A continuous distribution $\mathcal{D}$ with support in $[0, +\infty)$ and cumulative hazard rate H is called* entire *if $\mathbb{E}_{X\sim\mathcal{D}}[X] < +\infty$, H has a Puiseux series around 0, i.e. $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$, the Puiseux series is not identically zero and is convergent for every point in the support of $\mathcal{D}$.*

Since h is a non-negative function, H is a non-negative and monotonically non-decreasing function. Using this we obtain the following observation.

Observation 2.3. *Consider an entire distribution $\mathcal{D}$ supported on $[0, +\infty)$ with cumulative hazard rate $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$, where $d_1 < d_2 < \ldots$. Then, $a_1 > 0$ and $d_1 > 0$.*

Gamma function. The Gamma (Γ) function – which is an extension of the factorial function over the reals – and its relatives arise in our analysis of entire distributions. For $x > 0$, it is defined as $\Gamma(x) = \int_0^{\infty} t^{x-1} e^{-t}\, dt$. Of particular interest to us is the *lower incomplete Gamma function* γ, which is defined for $s > 0, x \geq 0$ as $\gamma(s, x) = \int_0^{x} t^{s-1} e^{-t}\, dt$.

To assist the reader, we include a primer on the Gamma function and its relatives in Appendix C, along with a few technical lemmas used in our analysis.

3 Optimal Algorithm: Constant Approximation via Multiple Thresholds

In this section, we focus on *optimal algorithms* for the cost prophet inequality (CPI) setting, i.e., algorithms that achieve the smallest possible α in (1.1). We show that, if the distribution has finite expected value, these algorithms achieve a (distribution-dependent) constant-factor CPI. Moreover, this factor is 2 for entire MHR distributions.

We first observe that, just as in the rewards maximization setting, it suffices to focus on threshold-based algorithms to achieve the optimal competitive ratio. A threshold-based algorithm decides thresholds $\tau_1, \ldots, \tau_n$ upfront using only the knowledge of the underlying distribution $\mathcal{D}$, and selects the first $X_i \leq \tau_i$. Since the thresholds do not depend on the realizations of the X_i's, the optimal threshold-based algorithm is an *oblivious* algorithm.

Proposition 3.1. *For any instance of the cost prophet inequality setting, one can achieve the optimal competitive ratio with a threshold-based oblivious algorithm.*

Intuitively, this is because the algorithm's decision in round i is independent of past realizations and only depends on the realization of X_i and the number of remaining random variables, i.e. it is a *memoryless* process.

Given Proposition 3.1, we focus on threshold-based algorithms. Let $\tau_1, \ldots, \tau_n$ denote the thresholds of the *optimal algorithm.* As it turns out, the optimal thresholds have a very natural interpretation: the algorithm should select the next random variable X_i if and only if its value is smaller than the value it expects to receive by ignoring X_i and continuing the process. That is, the optimal threshold for the next random variable when we have k realizations left to see is exactly the expected cost incurred by an optimal algorithm when its input is $k-1$ I.I.D. random variables.

We then analyze the performance of the optimal-threshold algorithm and show it attains a constant-factor competitive ratio if $\mathcal{D}$ has finite expected value. For this, we identify each $\mathcal{D}$ by its cumulative hazard rate $H(x)$ and provide a characterization of the constant factor achieved by the optimal algorithm as the minimum constant that satisfies a simple inequality. Moreover, we obtain a closed form for the constant for entire distributions, and show that it only depends on the growth rate of H[4].

Finally, we focus on the special case of MHR distributions and show that if $\mathcal{D}$ is an *entire MHR distribution*, the competitive ratio is at most 2, a factor which is made tight by the exponential distribution for which the competitive ratio is exactly 2.

3.1 Characterizing the Optimal Thresholds In this section we obtain an exact formulation for the optimal thresholds and, using these, design an optimal threshold-based algorithm. In what follows, we use $G(i)$ to denote $G_{OPTALG}(i)$ for brevity, where $OPTALG$ is an optimal algorithm.

Lemma 3.1. *For the cost prophet inequality problem with random variables $X_1, X_2, \ldots, X_n$, $\tau_n = +\infty$ for every algorithm. For $1 \le i \le n-1$, the optimal threshold for the random variable X_i is*

$$\tau_i = G(n-i).$$

Lemma 3.1 implies that the following threshold-based algorithm is an optimal algorithm; it achieves the best possible competitive ratio for the cost prophet inequality (CPI) problem.

Algorithm 3.1. Optimal Threshold Algorithm($\mathcal{D}$)
Set $\tau_n \leftarrow +\infty$ and $\tau_{n-1} \leftarrow \mathbb{E}_{X\sim\mathcal{D}}[X]$.
for $i \leftarrow n-2$ **to** 1 **do**
    $\tau_i \leftarrow F(\tau_{i+1})\,\mathbb{E}[X \mid X \le \tau_{i+1})] + (1-F(\tau_{i+1}))\,\tau_{i+1}$.
for $i \leftarrow 1$ **to** n **do**
    Let z_i be the realization of X_i.
    if $z_1, \ldots, z_{i-1}$ were not selected and $z_i \le \tau_i$ **then**
        Select z_i.

Before we continue, we provide an expression for the expected cost of the optimal algorithm, based on Lemma 3.1, in the following lemma.

Lemma 3.2. *The expected cost incurred by Algorithm 3.1 is*

$$G(n) = a + \int_a^{G(n-1)} e^{-H(u)}\,du.$$

Proof. Recall that $G(n)$ satisfies the recurrence relation in (B.1)

$$G(n) = F(\tau_1)\,\mathbb{E}[X \mid X \le \tau_1] + (1-F(\tau_1))\,G(n-1).$$

Substituting the optimal thresholds from Lemma 3.1 into the recurrence above, we obtain

$$G(n) = F(G(n-1))\,\mathbb{E}[X \mid X \le G(n-1)] + (1-F(G(n-1)))\,G(n-1)$$

[4]Recall that H is non-decreasing, since its derivative, the hazard rate function h, is non-negative.

$$
\begin{aligned}
&= F(G(n-1))\frac{\int_a^{G(n-1)} uf(u)\,du}{F(G(n-1))} + (1 - F(G(n-1)))\,G(n-1)\\
&= \int_a^{G(n-1)} uf(u)\,du + (1 - F(G(n-1)))\,G(n-1)\\
&= [uF(u)\,du]_a^{G(n-1)} - \int_a^{G(n-1)} F(u)\,du + (1 - F(G(n-1)))\,G(n-1)\\
&= G(n-1)F(G(n-1)) - \int_a^{G(n-1)} F(u)\,du + G(n-1) - G(n-1)F(G(n-1))\\
&= a + G(n-1) - a - \int_a^{G(n-1)} F(u)\,du - G(n-1)F(G(n-1))\\
&= a + \int_a^{G(n-1)} (1 - F(u))\,du.
\end{aligned}
$$

Next, recall that $H(x) = -\log(1 - F(x))$, and thus we obtain

$$G(n) = a + \int_a^{G(n-1)} e^{-H(u)}\,du.$$

□

3.2 Constant Factor Competitive Ratio In this section, we show Theorem 1.1, restated below. Recall that $\mu_{i:n}$ denotes the expected value of the i-th order statistic out of n draws from $\mathcal{D}$ and, for brevity, we write $\mu_n = \mu_{1:n}$.

Theorem 3.1. *Let $\mathcal{D}$ be a distribution supported on $[a,b]$ or $[a,+\infty)$, where $a \geq 0$, with cdf F, quantile function $F^{\leftarrow}$ and for which $\mu_1 = \int_a^b (1 - F(x))\,dx < +\infty$. There exist constants $n_0 \in \mathbb{N}$ and λ such that, for every $n \geq n_0$, there exists a λ-factor cost prophet inequality for the I.I.D. setting under $\mathcal{D}$. λ can be characterized as*

$$\lambda := \inf_{\lambda'>1} \left\{ \lambda' \;\middle|\; \lambda' - \frac{F^{\leftarrow}\left(\frac{c\left(1+\frac{1}{\lambda'-1}\right)}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} \geq 0, \quad \forall n \geq n_0 \right\}, \qquad \text{for } c = \frac{\mu_{2:2}}{\mu_2} - 1.$$

And this bound is tight.

The goal of the proof of Theorem 1.1 is to show that λ is a finite constant independent of n. We begin with certain lemmas about the tail behaviour of order statistics that are crucial in our analysis. Let

$$\phi(\lambda, n) = 1 - e^{-H(\lambda\mu_{n-1})} - \frac{\int_0^{H(\lambda\mu_{n-1})} H^{\leftarrow}(u)e^{-u}\,du}{\lambda\mu_{n-1}}.$$

Intuitively, $\phi(\lambda, n)$ captures the tail behavior of $\frac{\mu_n}{\mu_{n-1}}$ which converges to 1 as $n \to \infty$, for specific values of λ. Irrespective of this, we proceed to show that, for fixed n, as λ grows, $\phi(\lambda, n)$ approaches 1.

Lemma 3.3. *$\phi(\lambda, n)$ is strictly increasing in λ for every $n \geq 1$, and*

$$\lim_{\lambda\to\infty} \phi(\lambda, n) = 1.$$

Proof. To show that $\phi(\lambda, n)$ is strictly increasing in λ for every $n \geq 1$, fix n and observe that

$$\frac{d\phi(\lambda,n)}{d\lambda} = \mu_{n-1}h(\lambda\mu_{n-1})e^{-H(\lambda\mu_{n-1})} - \frac{\lambda\mu_{n-1}^2 e^{-H(\lambda\mu_{n-1})}h(\lambda\mu_{n-1}) \cdot \lambda\mu_{n-1} - \mu_{n-1}\int_0^{H(\lambda\mu_{n-1})} H^{\leftarrow}(u)e^{-u}\,du}{\lambda^2\mu_{n-1}^2}$$

$$= \frac{1}{\lambda^2 \mu_{n-1}} \int_0^{H(\lambda\mu_{n-1})} H^{\leftarrow}(u) e^{-u}\, du,$$

where the last expression is strictly positive since both H and $H^{\leftarrow}$ are increasing functions, and $\lambda\mu_{n-1} > a$.

Next, we analyze $\lim_{\lambda\to\infty} \phi(\lambda, n)$. Notice that

$$\lim_{\lambda\to\infty} 1 - e^{-H(\lambda\mu_{n-1})} = 1,$$

since $\lim_{x\to\infty} H(x) = +\infty$ by definition. Also,

$$\lim_{\lambda\to\infty} -\frac{\int_0^{H(\lambda\mu_{n-1})} H^{\leftarrow}(u) e^{-u}\, du}{\lambda\mu_{n-1}} = 0,$$

since $\int_0^{+\infty} H^{\leftarrow}(u) e^{-u}\, du$ converges, due to $\mathcal{D}$ having finite expected value.

Thus,

$$\lim_{\lambda\to\infty} \phi(\lambda, n) = 1.$$

□

Next, we provide a surprisingly simple expression for $\phi(\lambda, n)$, which will be used to characterize the optimal λ.

Lemma 3.4.

$$\phi(\lambda, n) = \frac{\int_a^{\lambda\mu_{n-1}} F(u)\, du}{\lambda\mu_{n-1}}.$$

Proof. Let $\psi(x) = 1 - e^{-H(x)} - \frac{\int_0^{H(x)} H^{\leftarrow}(u) e^{-u}\, du}{x}$, and observe that $\phi(\lambda, n) = \psi(\lambda\mu_{n-1})$.

Next, notice that $1 - e^{-H(x)} = [-e^{-u}]_0^{H(x)}$, and thus

$$\begin{aligned}
\psi(x) &= \left[-e^{-u}\right]_0^{H(x)} - \frac{\int_0^{H(x)} H^{\leftarrow}(u) e^{-u}\, du}{x} \\
&= \frac{\int_0^{H(x)} \left(x \frac{d(-e^{-u})}{du} - H^{\leftarrow}(u) e^{-u}\right) du}{x} \\
&= \frac{\int_0^{H(x)} \left(x \frac{d(-e^{-u})}{du} - H^{\leftarrow}(u) \frac{d(-e^{-u})}{du}\right) du}{x} \\
&= \frac{\int_0^{H(x)} \frac{d(-e^{-u})}{du} \left(x - H^{\leftarrow}(u)\right) du}{x} \\
&= \int_0^{H(x)} \frac{d(-e^{-u})}{du} \left(1 - \frac{H^{\leftarrow}(u)}{x}\right) du \qquad (3.5)
\end{aligned}$$

Next, let $y = H^{\leftarrow}(u) \iff u = H(y)$, which implies that $du = h(y)\, dy$ and that $d\left(-e^{-u}\right) = h(y) e^{-H(y)}\, dy$, and thus, (3.5) becomes

$$\psi(x) = \int_a^x h(y) e^{-H(y)} \left(1 - \frac{y}{x}\right) dy.$$

Finally, notice that

$$\begin{aligned}
\psi(x) &= \int_a^x \frac{d - e^{-H(y)}}{dy} \left(1 - \frac{y}{x}\right) dy \\
&= \left[-e^{-H(y)} \left(1 - \frac{y}{x}\right)\right]_a^x - \int_a^x \frac{e^{-H(y)}}{x}\, dy \\
&= 1 - \frac{a}{x} - \int_a^x \frac{e^{-H(y)}}{x}\, dy
\end{aligned}$$

$$= \frac{\int_a^x \left(1 - e^{-H(y)}\right) dy}{x}$$

$$= \frac{\int_a^x F(y)\, dy}{x}. \tag{3.6}$$

Substituting $x = \lambda\mu_{n-1}$ into (3.6) yields

$$\phi(\lambda, n) = \frac{\int_a^{\lambda\mu_{n-1}} F(y)\, dy}{\lambda\mu_{n-1}}.$$

□

We continue by showing that $\phi(\lambda, n)$, for specific values of λ, can match the tail behaviour of a specific function of moments of order statistics of the distribution.

Lemma 3.5. *There exist constants $n_0 \in \mathbb{N}$ and*

$$\lambda = \inf \left\{ \lambda' \mid \lambda - \frac{F^{\leftarrow}\left(\frac{c\left(1+\frac{1}{\lambda-1}\right)}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} \geq 0, \quad \forall n \geq n_0 \right\}$$

such that for every $n \geq n_0$, we have

$$\phi(\lambda, n) \geq \frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n},$$

where $\mu_{2:n}$ is the expected value of the second order statistic of F.

Proof. Let $c = \frac{\mu_{2:2}}{\mu_2} - 1$, and notice that $\frac{\mu_{2:n}}{\mu_n} \leq \frac{\mu_{2:2}}{\mu_2} = c + 1$. Thus,

$$\frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n} \leq \frac{c}{n}.$$

Next, by Lemma (3.3), we know that $\phi(\lambda, n)$ is strictly increasing in λ for every $n \geq 1$, and

$$\lim_{\lambda \to \infty} \phi(\lambda, n) = 1,$$

which implies that for every $n > c$, there exists a λ such that

$$\phi(\lambda, n) \geq \frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n}. \tag{3.7}$$

Next, we show that $\lambda = \inf \left\{ \lambda' \mid \lambda' - \frac{F^{\leftarrow}\left(\frac{c\left(1+\frac{1}{\lambda-1}\right)}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} \geq 0, \quad \forall n \geq n_0 \right\}$ satisfies (3.7).

By Lemma 3.4, we know that

$$\phi(\lambda, n) = \frac{\int_a^{\lambda\mu_{n-1}} F(u)\, du}{\lambda\mu_{n-1}}.$$

Taking the first-order Taylor approximation of $\int_a^x F(u)\, du$ around $x = a$, we obtain

$$\frac{\int_a^x F(u)\, du}{x} \approx \frac{x - a}{x} F(x),$$

and thus

$$\phi(\lambda, n) \approx \frac{\lambda\mu_{n-1} - a}{\lambda\mu_{n-1}} F(\lambda\mu_{n-1}).$$

Therefore, for large enough n, it suffices to show that

$$F(\lambda\mu_{n-1}) \geq \frac{c}{n} \cdot \frac{\lambda\mu_{n-1}}{\lambda\mu_{n-1} - a}.$$

If $a = 0$, we know that $\frac{\lambda\mu_{n-1}}{\lambda\mu_{n-1}-a} = 1$, and if $a \neq 0$, for λ bounded away from 1, we know that $\frac{\lambda\mu_{n-1}}{\lambda\mu_{n-1}-a} \leq \frac{\lambda a}{\lambda a - a} = \frac{\lambda}{\lambda-1} = 1 + \frac{1}{\lambda-1}$.

Thus, we need that

$$F(\lambda\mu_{n-1}) \geq \frac{c\left(1 + \frac{1}{\lambda-1}\right)}{n}. \tag{3.8}$$

For λ bounded away from 1, let $c' = c\left(1 + \frac{1}{\lambda-1}\right)$. Thus, (3.8) becomes

$$F(\lambda\mu_{n-1}) \geq \frac{c'}{n}. \tag{3.9}$$

By [ABN08] (Eq. (5.5.3)), for large enough n, we obtain

$$\mu_{n-1} \approx F^{\leftarrow}\left(\frac{1}{n}\right),$$

which turns (3.9) into

$$F\left(\lambda F^{\leftarrow}\left(\frac{1}{n}\right)\right) \geq \frac{c'}{n} \iff \lambda F^{\leftarrow}\left(\frac{1}{n}\right) \geq F^{\leftarrow}\left(\frac{c'}{n}\right) \iff \lambda \geq \frac{F^{\leftarrow}\left(\frac{c'}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)}, \tag{3.10}$$

which is true, as λ is the infimum across all numbers satisfying (3.10). Finally, λ is a constant, since $\frac{F^{\leftarrow}\left(\frac{c'}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)}$ is the ratio of two quantiles of $\mathcal{D}$ of the same order that converges to a constant as $n \to \infty$ [ABN08] (Eq. (5.5.2), (5.5.3)). ∎

Finally, we will need the following useful recurrence for moments of order statistics.

Lemma 3.6. ([ABN08], Theorem 5.3.13) *Let $\mu_{i:n}$ denote the expected value of the i-th order statistic out of n random variables. Then,*

$$\mu_{2:n} = n\mu_{1:n-1} - (n-1)\mu_{1:n}.$$

We are now ready to prove Theorem 1.1.

Proof. [Proof of Theorem 1.1] Fix a $n_0 \in \mathbb{N}$ to be defined later, and let $\xi = \inf\left\{\xi' \mid \frac{G(n)}{\mu_n} \leq \xi' \quad \forall n \geq n_0\right\}$. Note that, at this point, no bound on ξ has been established, and hence it could, in principle, be that $\xi = +\infty$. We show this is not the case. Recall that $G(n)$ denotes the expected value of the optimal threshold algorithm on n random variables. We start with the following upper bound on $G(n)$.

Claim 3.1. *For all $n > n_0$,*

$$G(n) \leq \xi\mu_{n-1}e^{-H(\xi\mu_{n-1})} + \int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(u)\, e^{-u}\, du.$$

Proof. By Lemma 3.2, we know that

$$G(n) = a + \int_a^{G(n-1)} e^{-H(u)}\, du \leq a + \int_a^{\xi\mu_{n-1}} e^{-H(u)}\, du, \tag{3.11}$$

where the last inequality follows from $\frac{G(n-1)}{\mu_{n-1}} \leq \xi$ by the definition of ξ.

Let $x = H(u) \iff u = H^{\leftarrow}(x)$, which implies $du = \frac{dH^{\leftarrow}(x)}{dx}\,dx$. Thus, (3.11) becomes

$$\begin{aligned} G(n) \le a + \int_0^{H(\xi\mu_{n-1})} \frac{dH^{\leftarrow}(x)}{dx}\, e^{-x}\,dx \\ &= a + \left[H^{\leftarrow}(x)\, e^{-x}\right]_0^{H(\xi\mu_{n-1})} + \int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(x)\, e^{-x}\,dx \\ &= a + \xi\mu_{n-1}e^{-H(\xi\mu_{n-1})} - H^{\leftarrow}(0) + \int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(x)\, e^{-x}\,dx \\ &= \xi\mu_{n-1}e^{-H(\xi\mu_{n-1})} + \int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(x)\, e^{-x}\,dx, \end{aligned}$$

since $H(a) = 0 \iff a = H^{\leftarrow}(0)$. ∎

Next, using Claim 3.1, we have

$$\begin{aligned} \frac{G(n)}{\mu_n} &\le \frac{\xi\mu_{n-1}e^{-H(\xi\mu_{n-1})} + \int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(u)\, e^{-u}\,du}{\mu_n} \\ &= \xi\frac{\mu_{n-1}}{\mu_n}\left(e^{-H(\xi\mu_{n-1})} + \frac{\int_0^{H(\xi\mu_{n-1})} H^{\leftarrow}(u)e^{-u}\,du}{\xi\mu_{n-1}}\right) \\ &= \xi\frac{\mu_{n-1}}{\mu_n}\left(1-\phi(\xi,n)\right). \end{aligned} \tag{3.12}$$

We need (3.12) to be at most ξ. This happens if and only if

$$\frac{\mu_{n-1}}{\mu_n}\left(1-\phi(\xi,n)\right) \le 1 \iff \left(1 + \frac{\mu_{n-1}}{\mu_n} - 1\right)\left(1-\phi(\xi,n)\right) \le 1.$$

This is true, for large enough n, if

$$\phi(\xi,n) \ge \frac{\mu_{n-1}}{\mu_n} - 1. \tag{3.13}$$

By Lemma 3.6 we know that

$$\mu_{2:n} = n\mu_{n-1} - (n-1)\mu_n,$$

which implies that

$$\frac{\mu_{n-1}}{\mu_n} - 1 = \frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n},$$

and (3.13) becomes

$$\phi(\xi,n) \ge \frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n}. \tag{3.14}$$

Finally, Lemma 3.5 shows the existence of constants ξ, namely $\xi = \lambda$, and n_0, and our result follows.

The tightness of our result is immediate from the fact that the λ guaranteed by Lemma 3.5 is taken to be the infimum among all such λ satisfying (3.14). ∎

3.3 Closed Form for Entire Distributions

If $\mathcal{D}$ is an entire distribution (see Definition 2.4), we provide a closed form for λ, via Theorem 1.2, restated here.

Theorem 3.2. *Let $\mathcal{D}$ be an entire distribution supported on $[a,b]$ or $[a,+\infty)$, where $a \ge 0$, with cdf F. There exists $n_0 \in \mathbb{N}$ such that, for every $n \ge n_0$, there exists a $\lambda(d)$-factor cost prophet inequality for the I.I.D. setting under $\mathcal{D}$, where*

$$\lambda(d) = \frac{(1+1/d)^{1/d}}{\Gamma(1+1/d)},$$

d is the smallest degree of the Puiseux series of H, and $\Gamma(\cdot)$ is the Gamma function.

Moreover, this constant is tight for the distribution with cumulative hazard rate $H(x) = x^d$.

3.3.1 Upper Bound

THEOREM 3.3. *Let $\mathcal{D}$ be an entire distribution on $[0,+\infty)$ with cumulative hazard rate H, which has a Puiseux series $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$ where $d_1 < d_1 < \dots$, and let*

$$\lambda(d_1) = \frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}.$$

Then, Algorithm 3.1 achieves a $\lambda(d_1)$-competitive ratio with respect to μ_n, for large enough n.

Proof. By Observation 2.3 we have that $d_1 > 0$ and $a_1 > 0$. For the competitive ratio of Algorithm 3.1, we start by analyzing its expected cost with respect to the cumulative hazard rate $H(x)$. Recall that, by Lemma 3.2, the expected cost incurred by Algorithm 3.1 is

$$G(n) = \int_0^{G(n-1)} e^{-H(u)}\,du.$$

Recall that $R(n)$ denotes the competitive ratio of Algorithm 3.1 for n random variables, and that our algorithm compares against the prophet who always selects the minimum value out of all realizations, i.e. μ_n on expectation. We want to show that $R(n)$ is upper bounded by a constant for all $n \geq 1$. By Lemma 3.2, we have

$$R(n) = \frac{G(n)}{\mu_n} = \frac{1}{\mu_n}\int_0^{G(n-1)} e^{-H(u)}\,du = \frac{1}{\mu_n}\int_0^{G(n-1)} e^{-\sum_{i=1}^{\infty} a_i u^{d_i}}\,du.$$

Before we proceed, we analyze μ_n.

LEMMA 3.7. *For every $n \geq 1$.*

$$\mu_n = \frac{\Gamma(1+1/d_1)}{(a_1\,n)^{1/d_1}} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right).$$

Proof.

$$\mu_n = \int_0^{\infty} e^{-nH(u)}\,du = \int_0^{\infty} e^{-n\,\sum_{i=1}^{\infty} a_i u^{d_i}}\,du = \int_0^{\infty} e^{-n\,a_1 u^{d_1}} \cdot e^{-n\,\sum_{i=2}^{\infty} a_i u^{d_i}}\,du$$

(3.15)
$$= \int_0^{\infty} e^{-n\,a_1 u^{d_1}} \cdot \prod_{i=2}^{\infty} e^{-n\,a_i u^{d_i}}\,du = \int_0^{\infty} e^{-n\,a_1 u^{d_1}} \cdot \prod_{i=2}^{\infty}\sum_{\ell_i \geq 0} \frac{\left(-n\,a_i u^{d_i}\right)^{\ell_i}}{\ell_i!}\,du.$$

Let $x = n\,a_1 u^{d_1} \iff u = \left(\frac{x}{n\,a_1}\right)^{1/d_1}$. Then,

$$dx = n\,a_1 d_1 u^{d_1-1}\,du \iff du = \frac{u^{1-d_1}}{n\,a_1\,d_1}\,dx = \frac{x^{1/d_1-1}}{n\,{a_1}^{1/d_1}\,d_1}\,dx,$$

and (3.15) becomes

$$\mu_n = \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\int_0^{\infty} e^{-x}x^{1/d_1-1} \cdot \prod_{i=2}^{\infty}\sum_{\ell_i\geq 0} \frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\,dx.$$

As in the proof of Lemma 3.8, each term $\sum_{\ell_i\geq 0}\frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}$ converges to $e^{-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}}$, and thus we have

$$\prod_{i=2}^{\infty}\sum_{\ell_i\geq 0}\frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!} = \sum_{\ell_2,\ell_3,\ldots\geq 0}\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}$$

$$
\begin{aligned}
\mu_n &= \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\int_0^\infty e^{-x}x^{1/d_1-1}\cdot\sum_{\ell_2,\ell_3,\ldots\geq 0}\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\,dx\\
&= \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\sum_{\ell_2,\ell_3,\ldots\geq 0}\int_0^\infty e^{-x}x^{1/d_1-1}\cdot\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\left(\frac{x}{n\,a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\,dx\\
&= \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\sum_{\ell_2,\ell_3,\ldots\geq 0}\int_0^\infty e^{-x}x^{1/d_1+1/d_1\sum_{j=2}^{\infty}d_j\ell_j-1}\cdot\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\,(n\,a_1)^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\,dx\\
&= \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\sum_{\ell_2,\ell_3,\ldots\geq 0}\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\,(n\,a_1)^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\cdot\int_0^\infty e^{-x}x^{1/d_1+1/d_1\sum_{j=2}^{\infty}d_j\ell_j-1}\,dx\\
&= \frac{1}{n\,{a_1}^{1/d_1}\,d_1}\sum_{\ell_2,\ell_3,\ldots\geq 0}\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\,(n\,a_1)^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\cdot\Gamma\left(1/d_1+1/d_1\sum_{j=2}^{\infty}d_j\ell_j\right) \qquad (3.16)\\
&= \frac{\Gamma(1/d_1)}{d_1\,(n\,a_1)^{1/d_1}}+\\
&\quad+\frac{1}{d_1\,(n\,a_1)^{1/d_1}}\sum_{\substack{\ell_2,\ell_3,\ldots\geq 0\\ \ell_2,\ell_3\ldots\neq(0,0,\ldots)}}\prod_{i=2}^{\infty}\frac{\left(-n\,a_i\,(n\,a_1)^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\Gamma\left(1/d_1+1/d_1\sum_{j=2}^{\infty}d_j\ell_j\right)\\
&= \frac{\Gamma(1+1/d_1)}{(n\,a_1)^{1/d_1}}+\frac{1}{d_1\,(n\,a_1)^{1/d_1}}\sum_{\substack{\ell_2,\ell_3,\ldots\geq 0\\ \ell_2,\ell_3,\ldots\neq(0,0,\ldots)}}n^{\sum_{j=2}^{\infty}\ell_j(1-d_j/d_1)}\\
&\qquad\cdot\prod_{i=2}^{\infty}\frac{\left(-a_i a_1^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\Gamma\left(1/d_1+1/d_1\sum_{j=2}^{\infty}d_j\ell_j\right). \qquad (3.17)
\end{aligned}
$$

where (3.16) follows by the definition of the Gamma function.

Notice that, since $d_i > d_1$ for all $i \geq 2$, we have $n^{1-d_i/d_1} = \mathrm{o}(1)$ for all $i \geq 2$. Also, in the exponent of n in the second summand, there is always at least one ℓ_j that is not 0, and thus

$$\mu_n = \frac{\Gamma(1+1/d_1)}{(a_1\,n)^{1/d_1}} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right).$$

∎

We are now ready to upper bound $R(n)$.

Lemma 3.8. *For every $n \geq 1$, we have*

$$R(n) \leq \frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}.$$

Proof. We show that $R(n) \leq \frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}$ via induction on n. For $n=1$, $R(1)=1$ and $\frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)} \geq 1$ for all $d_1 > 0$. For the induction hypothesis, assume $R(k) \leq \frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}$ for all $k \leq n$, and let $\lambda(d_1) = \frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}$ for brevity. For $n+1$ we have

$$R(n+1) = \frac{1}{\mu_{n+1}}\int_0^{G(n)} e^{-\sum_{i=1}^{\infty}a_i u^{d_i}}\,du$$

$$
\begin{aligned}
&\le \frac{1}{\mu_{n+1}} \int_0^{\lambda(d_1)\mu_n} e^{-\sum_{i=1}^{\infty} a_i u^{d_i}}\, du \\
&= \frac{1}{\mu_{n+1}} \int_0^{\lambda(d_1)\mu_n} e^{-a_1 u^{d_1}} \cdot \prod_{i=2}^{\infty} e^{-a_i u^{d_i}}\, du \\
&= \frac{1}{\mu_{n+1}} \int_0^{\lambda(d_1)\mu_n} e^{-a_1 u^{d_1}} \cdot \prod_{i=2}^{\infty} \sum_{\ell_i \ge 0} \frac{\left(-a_i u^{d_i}\right)^{\ell_i}}{\ell_i!}\, du. \qquad (3.18)
\end{aligned}
$$

where the second inequality follows by our induction hypothesis, since $G(n) \le \lambda(d_1)\mu_n$. Let $x = a_1 u^{d_1} \iff u = \left(\frac{x}{a_1}\right)^{1/d_1}$. Also,

$$
dx = a_1 d_1 u^{d_1 - 1}\, du \iff du = \frac{u^{1-d_1}}{a_1 d_1}\, dx = \frac{x^{1/d_1 - 1}}{{a_1}^{1/d_1} d_1}\, dx\,.
$$

Thus, (3.18) becomes

$$
R(n+1) \le \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \int_0^{a_1(\lambda(d_1)\mu_n)^{d_1}} e^{-x} x^{1/d_1 - 1} \cdot \prod_{i=2}^{\infty} \sum_{\ell_i \ge 0} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\, dx
$$

Recall that $\mathcal{D}$ is an entire distribution and thus the Puiseux series for H converges everywhere in the support of $\mathcal{D}$. Therefore, each term $\sum_{\ell_i \ge 0} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}$ converges to $e^{-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}}$, and thus we can use the distributive law for infinite products [DP02] and obtain

$$
\prod_{i=2}^{\infty} \sum_{\ell_i \ge 0} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!} = \sum_{\ell_2, \ell_3, \ldots \ge 0} \prod_{i=2}^{\infty} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}.
$$

Therefore,

$$
\begin{aligned}
R(n+1) &\le \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \int_0^{a_1(\lambda(d_1)\mu_n)^{d_1}} e^{-x} x^{1/d_1 - 1} \cdot \sum_{\ell_2, \ell_3, \ldots \ge 0} \prod_{i=2}^{\infty} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\, dx \\
&= \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2, \ell_3, \ldots \ge 0} \int_0^{a_1(\lambda(d_1)\mu_n)^{d_1}} e^{-x} x^{1/d_1 - 1} \cdot \prod_{i=2}^{\infty} \frac{\left(-a_i \left(\frac{x}{a_1}\right)^{d_i/d_1}\right)^{\ell_i}}{\ell_i!}\, dx \\
&= \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2, \ell_3, \ldots \ge 0} \int_0^{a_1(\lambda(d_1)\mu_n)^{d_1}} e^{-x} x^{1/d_1 + 1/d_1 \sum_{j=2}^{\infty} d_j \ell_j - 1} \cdot \prod_{i=2}^{\infty} \frac{\left(-a_i {a_1}^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!}\, dx \\
&= \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2, \ell_3, \ldots \ge 0} \prod_{i=2}^{\infty} \frac{\left(-a_i {a_1}^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!} \int_0^{a_1(\lambda(d_1)\mu_n)^{d_1}} e^{-x} x^{1/d_1 + 1/d_1 \sum_{j=2}^{\infty} d_j \ell_j - 1}\, dx \\
&= \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2, \ell_3, \ldots \ge 0} \prod_{i=2}^{\infty} \frac{\left(-a_i {a_1}^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!} \gamma\left(1/d_1 + 1/d_1 \sum_{j=2}^{\infty} d_j \ell_j,\, a_1 \left(\lambda(d_1)\mu_n\right)^{d_1}\right), \qquad (3.19)
\end{aligned}
$$

where the third inequality follows by multiplying together the terms of each sum, the fourth inequality follows by exchanging the order of summation and integration, the fifth inequality follows because the product does not depend on x, and the last inequality follows by the definition of $\gamma(s, x)$.

Now, using Fact C.2, (3.19) becomes

$$R(n+1) \leq \frac{1}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2,\ell_3,\ldots\geq 0} \prod_{i=2}^{\infty} \frac{\left(-a_i {a_1}^{-d_i/d_1}\right)^{\ell_i}}{\ell_i!} \left(a_1 \left(\lambda(d_1)\mu_n\right)^{d_1}\right)^{1/d_1 + 1/d_1 \sum_{j=2}^{\infty} d_j \ell_j}$$

$$\cdot \sum_{\ell_1 \geq 0} \frac{\left(-a_1 \left(\lambda(d_1)\mu_n\right)^{d_1}\right)^{\ell_1}}{\ell_1! \left(1/d_1 + 1/d_1 \sum_{j=2}^{\infty} d_j \ell_j + \ell_1\right)}$$

$$= \frac{{a_1}^{1/d_1} \lambda(d_1)\mu_n}{d_1 {a_1}^{1/d_1} \mu_{n+1}} \sum_{\ell_2,\ell_3,\ldots\geq 0} \prod_{i=2}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i!} \sum_{\ell_1 \geq 0} \frac{\left(-a_1 \left(\lambda(d_1)\mu_n\right)^{d_1}\right)^{\ell_1}}{\ell_1! \left(1/d_1 + 1/d_1 \sum_{j=1}^{\infty} d_j \ell_j\right)}$$

$$= \frac{\lambda(d_1)\mu_n}{\mu_{n+1}} \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=2}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i!} \frac{\left(-a_1 \left(\lambda(d_1)\mu_n\right)^{d_1}\right)^{\ell_1}}{\ell_1! \left(1 + \sum_{j=1}^{\infty} d_j \ell_j\right)}$$

$$= \lambda(d_1) \frac{\mu_n}{\mu_{n+1}} \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=1}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i! \left(1 + \sum_{j=1}^{\infty} d_j \ell_j\right)}. \tag{3.20}$$

Claim 3.2. *For large enough n,*

$$\frac{\mu_n}{\mu_{n+1}} \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=1}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i! \left(1 + \sum_{j=1}^{\infty} d_j \ell_j\right)} \leq 1.$$

Proof. By Lemma 3.7, we know that, for large enough n, there exist a constant $c \geq 0$ such that

$$\mu_n = \frac{\Gamma\left(1 + 1/d_1\right)}{\left(a_1 \ n\right)^{1/d_1}} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right) \leq \frac{\Gamma\left(1 + 1/d_1\right)}{\left(a_1 \ n\right)^{1/d_1}} \left(1 + \mathrm{o}\left(1\right)\right).$$

Therefore, we have

$$\frac{\mu_n}{\mu_{n+1}} = \left(\frac{n+1}{n}\right)^{1/d_1} \left(1 + \mathrm{o}\left(1\right)\right).$$

Thus

$$\frac{\mu_n}{\mu_{n+1}} \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=1}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i! \left(1 + \sum_{j=1}^{\infty} d_j \ell_j\right)} = \left(1 + \frac{1}{n}\right)^{1/d_1} \left(1 + \mathrm{o}\left(1\right)\right) \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=1}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i! \left(1 + \sum_{j=1}^{\infty} d_j \ell_j\right)}$$

$$\leq \left(1 + \frac{1}{n}\right)^{1/d_1} \left(1 + \mathrm{o}\left(1\right)\right) \left(1 - \sum_{i=1}^{\infty} \frac{a_i}{1 + d_i} \left(\lambda(d_1)\mu_n\right)^{d_i} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right)\right). \tag{3.21}$$

Notice that

$$\sum_{i=1}^{\infty} \frac{a_i}{1 + d_i} \left(\lambda(d_1)\mu_n\right)^{d_i} = \frac{a_1}{1 + d_1} \left(\lambda(d_1)\right)^{d_1} \mu_n^{d_1} + \sum_{i=2}^{\infty} \frac{a_i}{1 + d_i} \left(\lambda(d_1)\mu_n\right)^{d_i}.$$

Also, $\mu_n^{d_i} = \mathrm{O}\left(\frac{1}{n^{d_i/d_1}}\right)$, and for $i \geq 2$, we have $d_i > d_1$, which implies that $\mu_n^{d_i} = \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right)$. Thus, (3.21) becomes

$$\frac{\mu_n}{\mu_{n+1}} \sum_{\ell_1,\ell_2,\ldots\geq 0} \prod_{i=1}^{\infty} \frac{\left(-a_i \ \left(\lambda(d_1)\mu_n\right)^{d_i}\right)^{\ell_i}}{\ell_i! \left(1 + \sum_{j=1}^{k} d_j \ell_j\right)} \leq \left(1 + \frac{1}{n}\right)^{\frac{1}{d_1}} \left(1 + \mathrm{o}\left(1\right)\right) \left(1 - \frac{a_1}{1 + d_1} \left(\lambda(d_1)\mu_n\right)^{d_1} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right)\right)$$

$$\leq \left(\left(1+\frac{1}{n}\right)^{1/d_1} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right)\right)\left(1-\frac{a_1}{1+d_1}\left(\lambda(d_1)\mu_n\right)^{d_1} + \mathrm{o}\left(\frac{1}{n^{1/d_1}}\right)\right)$$

$$= 1+\frac{1}{d_1\, n}-\frac{a_1}{1+d_1}\left(\lambda(d_1)\mu_n\right)^{d_1}+\mathrm{o}\left(\frac{1}{n}\right)$$

$$= 1+\frac{1}{d_1\, n}-\frac{a_1}{1+d_1}\left(\lambda(d_1)\right)^{d_1}\frac{\left(\Gamma\left(1+1/d_1\right)\right)^{d_1}}{a_1 n}+\mathrm{o}\left(\frac{1}{n}\right) \tag{3.22}$$

For (3.22) to be bounded above by 1, we need

$$\frac{1}{d_1\, n} \leq \frac{a_1}{1+d_1}\left(\lambda(d_1)\right)^{d_1}\frac{\left(\Gamma\left(1+1/d_1\right)\right)^{d_1}}{a_1 n} \iff \left(\lambda(d_1)\right)^{d_1} \geq \frac{1+1/d_1}{\left(\Gamma\left(1+1/d_1\right)\right)^{d_1}} \iff$$

$$\lambda(d_1) \geq \frac{\left(1+1/d_1\right)^{1/d_1}}{\Gamma\left(1+1/d_1\right)},$$

which holds, since $\lambda(d_1) = \frac{(1+1/d_1)^{d_1}}{\Gamma(1+1/d_1)}$. ∎

Thus, it follows that Algorithm 3.1 achieves a $\frac{(1+1/d_1)^{1/d_1}}{\Gamma(1+1/d_1)}$-competitive ratio with respect to μ_n. ∎

∎

REMARK 3.1. *The astute reader might observe that throughout the paper we've assumed that the support of $\mathcal{D}$ begins at 0, which implies that $H(x) = \int_0^x h(u)\,du$, and thus $H(0) = 0$, which in turn implies that $d_1 > 0$. This is without loss of generality. Specifically, if the support of $\mathcal{D}$ begins at $a > 0$, one can "shift" it to the origin to find the approximation factor. Formally, we have $H(x) = \int_a^x h(u)du = \int_0^{x-a} h(u+a)du$, and thus $H(a) = 0$. Define $H'(x) = H(x+a)$. We have $H'(0) = 0$ and the approximation factor of the original distribution depends on $d_1' > 0$. Thus, this dependence is a technicality that does not affect the approximation factor.*

3.3.2 Lower Bound In this section, we show that there exist entire distributions for which the upper bounds given by λ of the previous section is tight. Notice that a cumulative hazard rate $H(x) = x^d$ defines, for $d > 0$, a distribution on $[0, +\infty)$, with CDF $F(x) = 1 - e^{-x^d}$, since $F(0) = 0$ and $\lim_{x\to\infty} F(x) = 1$. For $d = 1$, the resulting distribution is the exponential with rate 1.

THEOREM 3.4. *Consider the distribution $\mathcal{D}$ for which $H(x) = x^d$ for $d \geq 0$. For any $\varepsilon > 0$, there is no $\left(\frac{(1+1/d)^{1/d}}{\Gamma(1+1/d)} - \varepsilon\right)$-competitive cost prophet inequality for the single-item setting and I.I.D. random variables drawn from $\mathcal{D}$.*

Proof. Let $\lambda(d) = \frac{(1+1/d)^{1/d}}{\Gamma(1+1/d)}$.

LEMMA 3.9. *For every $n \geq 1$,*

$$\mu_n = \frac{\Gamma\left(1+1/d\right)}{n^{1/d}}.$$

Proof. The proof follows immediately from the proof of Lemma 3.7. In particular, we have

$$\mu_n = \int_0^\infty e^{-nH(u)}\,du = \int_0^\infty e^{-n\,u^d}\,du, \tag{3.23}$$

and, by (3.17) of Lemma 3.7, since $a_1 = 1$ and $a_2 = \cdots = a_k = 0$, we get that

$$\mu_n = \frac{\Gamma\left(1+1/d\right)}{n^{1/d}}.$$

∎

Using Lemma 3.9, we have that

$$R(n) = \frac{G(n)}{\mu_n} = \frac{n^{1/d}}{\Gamma(1+1/d)} \int_0^{G(n-1)} e^{-H(u)}\, du = \frac{n^{1/d}}{\Gamma(1+1/d)} \int_0^{G(n-1)} e^{-u^d}\, du. \tag{3.24}$$

Let $x = u^d \iff u = x^{1/d}$. Also,

$$dx = d\, u^{d-1}\, du \iff du = \frac{u^{1-d}}{d}\, dx = \frac{x^{1/d-1}}{d}\, dx,$$

and thus (3.24) becomes

$$R(n) = \frac{n^{1/d}}{d\,\Gamma(1+1/d)} \int_0^{(G(n-1))^d} e^{-x} x^{1/d-1}\, dx = \frac{n^{1/d}}{\Gamma(1+1/d)}\, \frac{1}{d}\, \gamma\left(1/d, (G(n-1))^d\right). \tag{3.25}$$

where the second equality follows from the definition of the lower incomplete Gamma function.

Lemma 3.10. *$R(n)$ is increasing in n.*

Proof. Recall that, by (3.25), we have

$$\begin{aligned} R(n) &= \frac{n^{1/d}}{\Gamma(1/d)} \gamma\left(1/d, (G(n-1))^d\right) \\ &= \frac{1}{d} \frac{n^{1/d}}{\Gamma\,(1+1/d)} \gamma\left(1/d, (G(n-1))^d\right) \\ &= \frac{1}{d\,\mu_n} \gamma\left(1/d, (G(n-1))^d\right). \end{aligned}$$

However, by Fact C.2, we have

$$\gamma\left(1/d, (G(n-1))^d\right) = G(n-1) \sum_{k=0}^{\infty} \frac{\left(-\,(G(n-1))^d\right)^k}{k!\ (1/d+k)}.$$

Thus

$$\begin{aligned} R(n) &= \frac{G(n-1)}{d\,\mu_n} \sum_{k=0}^{\infty} \frac{\left(-\,(G(n-1))^d\right)^k}{k!\ (1/d+k)} \\ &= \frac{G(n-1)}{\mu_{n-1}}\ \frac{\mu_{n-1}}{\mu_n} \sum_{k=0}^{\infty} \frac{\left(-\,(G(n-1))^d\right)^k}{k!\ (1+d\,k)} \\ &= R(n-1)\ \frac{\mu_{n-1}}{\mu_n} \sum_{k=0}^{\infty} \frac{\left(-\,(G(n-1))^d\right)^k}{k!\ (1+d\,k)}. \end{aligned}$$

It suffices to show that

$$\frac{\mu_{n-1}}{\mu_n} \sum_{k=0}^{\infty} \frac{\left(-\,(G(n-1))^d\right)^k}{k!\ (1+d\,k)} \geq 1.$$

Notice that

$$\frac{\mu_{n-1}}{\mu_n} = \left(\frac{n}{n-1}\right)^{1/d} = \left(1 + \frac{1}{n-1}\right)^{1/d} = \sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell}.$$

Thus, it suffices to show that

$$\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \sum_{k=0}^{\infty} \frac{\left(-\left(G(n-1)\right)^{d}\right)^{k}}{k!\ (1+d\,k)} \geq 1.$$

We use the fact that $G(n-1) \leq \lambda(d)\mu_{n-1} = \left(\frac{1+1/d}{n-1}\right)^{1/d}$ and get

$$\begin{aligned}
\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \sum_{k=0}^{\infty} \frac{\left(-\left(G(n-1)\right)^{d}\right)^{k}}{k!\ (1+d\,k)} &= \sum_{k=0}^{\infty}\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \frac{\left(-\left(G(n-1)\right)^{d}\right)^{k}}{k!\ (1+d\,k)} \\
&\geq \sum_{k=0}^{\infty}\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \frac{(-(1+1/d))^{k}}{k!\,(n-1)^{k}\,(1+d\,k)} \\
&= \sum_{k=0}^{\infty}\sum_{\ell=0}^{1/d} \binom{1/d}{\ell} \frac{(-(1+1/d))^{k}}{k!\ (1+d\,k)} \cdot \frac{1}{(n-1)^{\ell+k}} \\
&= 1 + \frac{1}{d(n-1)} - \frac{1+1/d}{(d+1)(n-1)} + \mathrm{O}\left(\frac{1}{n^2}\right)
\end{aligned}$$

Thus, for this quantity to be greater than 1, it suffices to have

$$\frac{1}{d(n-1)} \geq \frac{1+1/d}{(d+1)(n-1)} \iff \frac{d+1}{d} \geq 1+1/d,$$

which is true. ∎

Assume, towards contradiction, that $\lim_{n\to\infty} R(n) = \lambda^* < \lambda(d_1) = \frac{(1+1/d)^{1/d}}{\Gamma(1+1/d)}$.

We know that $G(n-1) = R(n-1)\mu_{n-1} = \frac{\Gamma(1+1/d)}{(n-1)^{1/d}} R(n-1)$. Thus we get

$$R(n) = \frac{n^{1/d}}{\Gamma(1+1/d)}\,\frac{1}{d}\,\gamma\left(1/d, \frac{\left(\Gamma\left(1+1/d\right)\right)^{d}}{(n-1)} \left(R(n-1)\right)^{d}\right). \tag{3.26}$$

Recall that, by Fact C.2, $\gamma(s,x) = x^s \sum_{k=0}^{\infty} \frac{(-x)^k}{k!\,(s+k)}$, and thus (3.26) becomes

$$\begin{aligned}
R(n) &= \left(\frac{n}{n-1}\right)^{1/d} R(n-1)\frac{1}{d}\sum_{k=0}^{\infty} \frac{\left(-\frac{(\Gamma(1+1/d))^{d}}{(n-1)} \left(R(n-1)\right)^{d}\right)^{k}}{k!\,(1/d+k)} \\
R(n) &= \left(\frac{n}{n-1}\right)^{1/d} R(n-1)\sum_{k=0}^{\infty} \frac{\left(-\frac{(\Gamma(1+1/d))^{d}}{(n-1)} \left(R(n-1)\right)^{d}\right)^{k}}{k!\,(1+d\,k)} \\
&= R(n-1)\left(1+\frac{1}{n-1}\right)^{1/d}\sum_{k=0}^{\infty} \frac{\left(-\frac{(\Gamma(1+1/d))^{d}}{(n-1)} \left(R(n-1)\right)^{d}\right)^{k}}{k!\,(1+d\,k)}
\end{aligned}$$

Notice that

$$\left(1+\frac{1}{n-1}\right)^{1/d} = \sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell}.$$

Thus,

$$\left(1+\frac{1}{n-1}\right)^{1/d}\sum_{k=0}^{\infty} \frac{\left(-\frac{(\Gamma(1+1/d))^{d}}{(n-1)} \left(R(n-1)\right)^{d}\right)^{k}}{k!\,(1+d\,k)}$$

$$
\begin{aligned}
&= \sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \sum_{k=0}^{\infty} \frac{\left(-\frac{(\Gamma(1+1/d))^d}{(n-1)} \left(R(n-1)\right)^d\right)^k}{k!\,(1+d\,k)} \\
&= \sum_{k=0}^{\infty}\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell}} \binom{1/d}{\ell} \cdot \frac{\left(-\frac{(\Gamma(1+1/d))^d}{(n-1)} \left(R(n-1)\right)^d\right)^k}{k!\,(1+d\,k)} \\
&= \sum_{k=0}^{\infty}\sum_{\ell=0}^{1/d} \frac{1}{(n-1)^{\ell+k}} \binom{1/d}{\ell} \cdot \frac{\left(-\left(\Gamma\left(1+1/d\right) \cdot R(n-1)\right)^d\right)^k}{k!\,(1+d\,k)} \\
&\approx 1 + \frac{1}{d\,(n-1)} - \frac{\left(\Gamma\left(1+1/d\right) \cdot R(n-1)\right)^d}{(d+1)\,(n-1)},
\end{aligned}
$$

where, for large enough n, we can ignore higher order terms and we also have $R(n-1) \approx \lambda^*$. Thus, for $R(n) \leq \lambda^*$, it must be that

$$
\frac{1}{d} - \frac{\left(\Gamma\left(1+1/d\right) \cdot \lambda^*\right)^d}{(d+1)} \leq 0 \iff \left(\Gamma\left(1+1/d\right) \cdot \lambda^*\right)^d \geq 1 + 1/d \iff \lambda^* \geq \frac{(1+1/d)^{1/d}}{\Gamma\left(1+1/d\right)},
$$

and we arrive at a contradiction.

Therefore, for any $\varepsilon > 0$, there is no $\left(\frac{(1+1/d)^{1/d}}{\Gamma(1+1/d)} - \varepsilon\right)$-competitive cost prophet inequality for the single-item setting and I.I.D. random variables drawn from $\mathcal{D}$. ∎

Now, Theorem 1.2 follows by Theorems 3.3 and 3.4.

3.4 Special Case: MHR Distributions Even though the constant-factor competitive ratio obtained by Algorithm 3.1 is distribution-dependent, it turns out that we can show a uniform factor of 2 when the distributions are MHR and entire. This factor is also tight, and it provides a nice parallel to the standard 1/2-competitive prophet inequality in the rewards setting [KS77, KS78, SC84, KW19].

THEOREM 3.5. *For every entire MHR distribution, there exists an algorithm that achieves a 2-competitive ratio in the cost prophet inequality setting.*

Proof. Let $\mathcal{D}$ be an entire MHR distribution with cumulative hazard rate H where H has a Puiseux series $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$ and $d_1 < d_2 < \ldots$. Notice that since $\mathcal{D}$ has a monotonically increasing hazard rate, we have $h'(x) = H''(x) \geq 0$ everywhere in $[0, +\infty)$. Thus,

$$
\left(\sum_{i=1}^{\infty} a_i x^{d_i}\right)'' \geq 0 \iff \left(\sum_{i=1}^{\infty} a_i\, d_i x^{d_i-1}\right)' \geq 0 \iff \sum_{i=1}^{\infty} a_i\, d_i\, \left(d_i-1\right) x^{d_i-2} \geq 0,
$$

for all $x \geq 0$. Recall that, by Observation 2.3, for H to be the cumulative hazard rate of a distribution $\mathcal{D}$, it must be an increasing function in x, and thus $a_1 > 0$.

Assume towards contradiction that $d_1 < 1$, which implies that the first term of H is negative. We use this to contradict the fact that $H''(x) \geq 0$ everywhere. In particular, consider a point y where

$$
\begin{aligned}
a_1\, d_1(1-d_1) y^{d_1} > \sum_{i=2}^{\infty} a_i\, d_i\, (d_i-1) y^{d_i} &\iff \\
y^2 \left(a_1\, d_1(1-d_1) y^{d_1-2} - \sum_{i=2}^{\infty} a_i\, d_i\, (d_i-1) y^{d_i-2}\right) > 0 &\iff \\
-y^2 H''(y) > 0 \implies H''(y) < 0.&
\end{aligned}
$$

Such a point can always be found because, for any choice of $a_1, a_2, \ldots$ and $d_1 < d_2 < \ldots$, one can pick a small enough y that ensures $a_1\, d_1(1-d_1) y^{d_1}$ dominates the term $\sum_{i=2}^{\infty} a_i\, d_i\, (d_i-1) y^{d_i}$.

Therefore, for all entire MHR distributions, it must be the case that $d_1 \geq 1$. This implies that for every entire MHR distribution $\mathcal{D}$, $\lambda(d_1) \leq \lambda(1) = 2$, and thus Algorithm 3.1 obtains a 2-factor approximation to the prophet's cost. ∎

Furthermore, notice that if we consider the distribution with $H(x) = x$, i.e. the exponential distribution, then, as a corollary of Theorem 1.2 for $d = 1$, we get that the factor of 2 is tight. The exponential distribution is MHR as it has a constant hazard rate, and hence we obtain the following result.

Corollary 3.1. *For any $\varepsilon > 0$, there exists no $(2-\varepsilon)$-factor cost prophet inequality for the exponential distribution.*

Now, Theorem 1.3 follows by Theorem 3.5 and Corollary 3.1.

4 Single Threshold Algorithm

This section is dedicated to proving Theorem 1.4. We design an algorithm which sets a fixed threshold T and selects the first realization that is below T. If our algorithm ever reaches X_n and has not selected any value, it is forced to pick the realization of X_n regardless of its cost. Our choice of T is

$$T = \Theta\left(\left(\frac{\log n}{n}\right)^k\right),$$

for an appropriate value of k that depends on the given distribution.

We analyze our algorithm's performance for an entire distribution with Puiseux series for the cumulative hazard rate $H(x) = \sum_{i=1}^{\infty} a_i\, x^{d_i}$, where $d_1 < d_2 < \dots$, and obtain a $\mathrm{O}\left((\log n)^{1/d_1}\right)$-competitive ratio. We then proceed to show that this ratio is asymptotically tight, as we show that no single threshold algorithm can achieve a competitive ratio better than $\Omega\left((\log n)^{1/d}\right)$ for the distribution with $H(x) = x^d$. Our results imply a $\mathrm{O}\left(\operatorname{polylog} n\right)$-factor single-threshold cost prophet inequality for the single-item setting and entire distributions.

Theorem 1.4 follows by Theorems 4.1 and 4.2.

4.1 Upper Bound

Theorem 4.1. *Let $\mathcal{D}$ be an entire distribution on $[0, +\infty)$ for which the cumulative hazard rate function H has Puiseux series $H(x) = \sum_{i=1}^{\infty} a_i x^{d_i}$, where $d_1 < d_2 < \dots$. Then, there exists a single threshold $T = T(n, d_1, a_1)$ such that the algorithm that selects the first value $X_i \leq T$ for $i < n$ and X_n otherwise, achieves a $\mathrm{O}\left((\log n)^{1/d_1}\right)$-competitive ratio compared to μ_n, for large enough n.*

Proof. We start by analyzing the algorithm's performance for an arbitrary choice of T. We have

$$\begin{aligned}
\mathbb{E}[ALG] &= \left(1 - (1-F(T))^{n-1}\right)\mathbb{E}\left[X \mid X \leq T\right] + (1-F(T))^{n-1}\,\mathbb{E}[X] \qquad (4.27)\\
&= \left(1 - e^{-(n-1)H(T)}\right)\int_0^T \left(1 - \frac{F(x)}{F(T)}\right) dx + e^{-(n-1)H(T)}\int_0^\infty e^{-H(x)}\,dx\\
&= \left(1 - e^{-(n-1)H(T)}\right)\int_0^T \left(1 - \frac{1-e^{-H(x)}}{1-e^{-H(T)}}\right) dx + e^{-(n-1)H(T)}\int_0^\infty e^{-H(x)}\,dx\\
&= \frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-H(x)}\,dx - Te^{-H(T)}\right) + e^{-(n-1)H(T)}\int_0^\infty e^{-H(x)}\,dx \iff
\end{aligned}$$

$$R(n) = \frac{1}{\mu_n}\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-H(x)}\,dx - Te^{-H(T)}\right) + e^{-(n-1)H(T)}\mu_1\right). \qquad (4.28)$$

Notice that,

$$\int_0^T e^{-H(x)}\,dx - Te^{-H(T)} \leq T\left(1 - e^{-H(T)}\right).$$

Using the above, (4.28) becomes

$$R(n) \leq \frac{1}{\mu_n}\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}T\left(1-e^{-H(T)}\right)+e^{-(n-1)H(T)}\mu_1\right)$$
$$(4.29) \qquad = \frac{1}{\mu_n}\left(\left(1-e^{-(n-1)H(T)}\right)T+e^{-(n-1)H(T)}\mu_1\right).$$

By Lemma 3.7, we know that there exist constants $c_1, c_2 > 0$ such that for large enough n, we have

$$c_1\frac{\Gamma\left(1+1/d_1\right)}{n^{1/d_1}} \leq \mu_n \leq c_2\frac{\Gamma\left(1+1/d_1\right)}{n^{1/d_1}}.$$

Thus, (4.29) becomes

$$(4.30) \qquad R(n) \leq \frac{n^{1/d_1}}{c_1\,\Gamma\left(1+1/d_1\right)}\left(\left(1-e^{-(n-1)H(T)}\right)T+e^{-(n-1)H(T)}\,c_2\,\Gamma\left(1+1/d_1\right)\right).$$

$$(4.31) \qquad = \frac{n^{1/d_1}}{c_1\,\Gamma\left(1+1/d_1\right)}\left(1-e^{-(n-1)H(T)}\right)T+e^{-(n-1)H(T)}\frac{c_2}{c_1}n^{1/d_1}.$$

Let

$$T = \left(\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}\right)^{1/d_1}.$$

Since $H(T)=\sum_{i=1}^{\infty}a_iT^{d_i}$, we have

$$H(T) = a_1\cdot\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}+\sum_{i=2}^{\infty}a_i\left(\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}\right)^{d_i/d_1} = \frac{\log\left(\frac{n}{\log n}\right)}{d_1\,(n-1)}+\sum_{i=2}^{\infty}a_i\left(\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}\right)^{d_i/d_1}.$$

Since $d_i > d_1$ for all $i \geq 2$, we have that, for large enough n,

$$H(T) \approx a_1\,T^{d_1},$$

as $\sum_{i=2}^{\infty}a_iT^{d_i} = \mathrm{o}\left(T^{d_1}\right)$. Thus, (4.31) becomes

$$\begin{aligned}
R(n) &\leq \frac{n^{1/d_1}}{c_1\,\Gamma\left(1+1/d_1\right)}\left(1-e^{-(n-1)\,a_1\,T^{d_1}}\right)T+e^{-(n-1)\,a_1\,T^{d_1}}\frac{c_2}{c_1}n^{1/d_1}\\
&= \frac{n^{1/d_1}}{c_1\,\Gamma\left(1+1/d_1\right)}\left(1-e^{-(n-1)\,a_1\,\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}}\right)\left(\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}\right)^{\frac{1}{d_1}}+e^{-(n-1)\,a_1\,\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}}\frac{c_2}{c_1}n^{\frac{1}{d_1}}\\
&= \frac{n^{1/d_1}}{c_1\,\Gamma\left(1+1/d_1\right)}\left(1-\left(\frac{\log n}{n}\right)^{1/d_1}\right)\left(\frac{\log\left(\frac{n}{\log n}\right)}{d_1\,a_1\,(n-1)}\right)^{1/d_1}+\frac{c_2}{c_1}\left(\frac{\log n}{n}\right)^{1/d_1}n^{1/d_1}\\
&= \frac{1}{c_1\,\Gamma\left(1+1/d_1\right)\,\left(d_1\,a_1\right)^{1/d_1}}\left(\frac{n}{n-1}\right)^{\frac{1}{d_1}}\left(1-\left(\frac{\log n}{n}\right)^{\frac{1}{d_1}}\right)\left(\log\left(\frac{n}{\log n}\right)\right)^{\frac{1}{d_1}}+\frac{c_2}{c_1}\left(\log n\right)^{\frac{1}{d_1}}
\end{aligned}$$

However, there exists a constant $c_3 > 0$ such that for large enough n,

$$\left(1+\frac{1}{n-1}\right)^{1+1/d_1}\left(1-\left(\frac{\log n}{n}\right)^{1/d_1}\right) \leq c_3,$$

and also $\left(\log\left(\frac{n}{\log n}\right)\right)^{1/d_1} \leq (\log n)^{1/d_1}$. Thus,

$$R(n) \leq \frac{c_3}{c_1\,\Gamma\left(1+1/d_1\right)\,\left(d_1\,a_1\right)^{1/d_1}} \cdot (\log n)^{1/d_1} + \frac{c_2}{c_1}\,(\log n)^{1/d_1} = \mathrm{O}\left((\log n)^{1/d_1}\right).$$

∎

4.2 Lower Bound

THEOREM 4.2. *Consider the distribution $\mathcal{D}$ for which $H(x) = x^d$ for $d \geq 0$. There is no $\mathrm{o}\left((\log n)^{1/d}\right)$-competitive single-threshold cost prophet inequality for the single-item setting and I.I.D. random variables drawn from $\mathcal{D}$.*

Proof. Recall by (4.28) that

$$R(n) = \frac{1}{\mu_n}\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-H(x)}\,dx - Te^{-H(T)}\right) + e^{-(n-1)H(T)}\mu_1\right).$$

Assume, towards contradiction, that $R(n) = \mathrm{o}\left((\log n)^{1/d}\right)$. For this to be the case, it must be that

$$e^{-(n-1)H(T)}\,\frac{\mu_1}{\mu_n} = \mathrm{o}\left((\log n)^{1/d}\right), \tag{4.32}$$

and also that

$$\frac{1}{\mu_n}\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-H(x)}\,dx - Te^{-H(T)}\right)\right) = \mathrm{o}\left((\log n)^{1/d}\right). \tag{4.33}$$

By (4.32) and the definition of $\mathrm{o}(\cdot)$, we have that for every $\varepsilon > 0$, there must exist a $n_0 \geq 1$ such that for all $n \geq n_0$, we have

$$e^{-(n-1)H(T)}\,\frac{\mu_1}{\mu_n} \leq \varepsilon\,(\log n)^{1/d} \iff e^{-(n-1)H(T)} \leq \varepsilon\,\frac{\mu_n}{\mu_1}\,(\log n)^{1/d} \iff$$

$$-(n-1)H(T) \leq \log\left(\varepsilon\,\frac{\mu_n}{\mu_1}\,(\log n)^{1/d}\right) \iff H(T) \geq \frac{\log\left(\frac{\mu_1}{\varepsilon\,\mu_n\,(\log n)^{1/d}}\right)}{n-1} \iff$$

$$T^d \geq \frac{\log\left(\frac{\mu_1}{\varepsilon\,\mu_n\,(\log n)^{1/d}}\right)}{n-1} \iff T \geq \left(\frac{\log\left(\frac{\mu_1}{\varepsilon\,\mu_n\,(\log n)^{1/d}}\right)}{n-1}\right)^{1/d}. \tag{4.34}$$

However, by (4.33), we have that for every $\varepsilon' > 0$, there must exist a $n_1 \geq 1$ such that for all $n \geq n_1$, we have

$$\frac{1}{\mu_n}\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-H(x)}\,dx - Te^{-H(T)}\right)\right) \leq \varepsilon'\,(\log n)^{1/d}$$

$$\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\int_0^T e^{-x^d}\,dx - Te^{-H(T)}\right)\right) \leq \varepsilon'\,\mu_n\,(\log n)^{1/d}$$

$$\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(\frac{1}{d}\gamma\left(1/d, T^d\right) - Te^{-H(T)}\right)\right) \leq \varepsilon'\,\mu_n\,(\log n)^{1/d}. \tag{4.35}$$

where the last equality follows by substituting $t = x^d$ in the integral, as seen several other times in the paper.

Notice that T has to be decreasing in n, since, if not, one can easily see from (4.27) that the algorithm is too eager to select a value and its performance degrades rapidly as n increases. Therefore, we know that $\lim_{n\to\infty} T = 0$. Furthermore, by Fact C.3, we know that for small T, i.e. large enough n, we have

$$\gamma\left(1/d, T^d\right) \approx d\,T,$$

and thus (4.35) becomes

$$\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}\left(T-Te^{-H(T)}\right)\right)\le \varepsilon'\,\mu_n\,\left(\log n\right)^{1/d} \iff$$
$$\left(\frac{1-e^{-(n-1)H(T)}}{1-e^{-H(T)}}T\,\left(1-e^{-H(T)}\right)\right)\le \varepsilon'\,\mu_n\,\left(\log n\right)^{1/d} \iff$$
$$T\,\left(1-e^{-(n-1)H(T)}\right)\le \varepsilon'\,\mu_n\,\left(\log n\right)^{1/d}.$$

However, by (4.34) we know that we must have

$$T\ge\left(\frac{\log\left(\frac{\mu_1}{\varepsilon\,\mu_n\,(\log n)^{1/d}}\right)}{n-1}\right)^{1/d},$$

and if

$$T\,\left(1-e^{-(n-1)H(T)}\right)\le \varepsilon'\,\mu_n\,\left(\log n\right)^{1/d},$$

then it also must be the case that

$$T\,\left(1-\frac{\varepsilon\,\mu_n\,\left(\log n\right)^{1/d}}{\mu_1}\right)\le \varepsilon'\,\mu_n\,\left(\log n\right)^{1/d}.$$

Notice that by Lemma 3.7

$$\mu_n=\frac{\Gamma\left(1+1/d\right)}{n^{1/d}}\qquad\text{and}\qquad \mu_1=\Gamma\left(1+1/d\right),$$

and thus

$$T\,\left(1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}\right)\le \varepsilon'\,\mu_1\,\left(\frac{\log n}{n}\right)^{1/d}.$$

For every ε, for n large enough, we have $1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}>0$, and thus

$$T\le \varepsilon'\,\frac{\mu_1\,\left(\frac{\log n}{n}\right)^{1/d}}{1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}}. \tag{4.36}$$

To arrive at a contradiction, we use (4.34) and (4.36) to show that it suffices to find, for every $\varepsilon>0$, a constant $\varepsilon'>0$ such that

$$\varepsilon'\,\frac{\mu_1\,\left(\frac{\log n}{n}\right)^{1/d}}{1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}}<\left(\frac{\log\left(\frac{\mu_1}{\varepsilon\,\mu_n\,(\log n)^{1/d}}\right)}{n-1}\right)^{1/d}=\left(\frac{\log\left(\frac{1}{\varepsilon}\,\left(\frac{n}{\log n}\right)^{1/d}\right)}{n-1}\right)^{1/d}.$$

Indeed, rearranging the terms above, we get

$$\varepsilon'<\frac{1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}}{\mu_1\,\left(\frac{\log n}{n}\right)^{1/d}}\cdot\left(\frac{\log\left(\frac{1}{\varepsilon}\,\left(\frac{n}{\log n}\right)^{1/d}\right)}{n-1}\right)^{1/d}$$
$$=\frac{1}{\mu_1}\cdot\frac{1-\varepsilon\,\left(\frac{\log n}{n}\right)^{1/d}}{\left(\frac{\log n}{n}\right)^{1/d}}\cdot\left(\frac{\frac{1}{d}\cdot\log\left(\frac{1}{\varepsilon^d}\,\frac{n}{\log n}\right)}{n-1}\right)^{1/d}$$

$$= \frac{1}{\mu_1} \cdot \left(1 - \varepsilon \left(\frac{\log n}{n}\right)^{1/d}\right) \cdot \left(\frac{\frac{1}{d} \cdot \frac{n}{\log n} \log\left(\frac{1}{\varepsilon^d} \frac{n}{\log n}\right)}{n-1}\right)^{1/d}$$

$$= \frac{1}{\mu_1} \cdot \left(1 - \varepsilon \left(\frac{\log n}{n}\right)^{1/d}\right) \cdot \left(\frac{1}{d} \cdot \frac{n}{n-1} \cdot \frac{\log\left(\frac{1}{\varepsilon^d} \frac{n}{\log n}\right)}{\log n}\right)^{1/d}$$

$$= \frac{1}{\mu_1 \, d^{1/d}} \cdot \left(1 - \varepsilon \left(\frac{\log n}{n}\right)^{1/d}\right) \cdot \left(\frac{n}{n-1} \cdot \frac{\log n - \log\left(\varepsilon^d \, \log n\right)}{\log n}\right)^{1/d}.$$

Notice, however, that for any fixed $\varepsilon > 0$, we have

$$\lim_{n\to\infty} \left(1 - \varepsilon \left(\frac{\log n}{n}\right)^{1/d}\right) \cdot \left(\frac{n}{n-1} \cdot \frac{\log n - \log\left(\varepsilon^d \, \log n\right)}{\log n}\right)^{1/d} = 1,$$

and thus, for every $\varepsilon > 0$ there exists a large enough n and a constant $0 < \varepsilon' \frac{1}{\mu_1 d^{1/d}}$ such that (4.34) and (4.36) cannot simultaneously hold, and we arrive at a contradiction. □

5 Conclusion

Our work initiates the study of prophet inequalities for cost minimization, and provides a nuanced distribution-sensitive analysis. It opens up a number of interesting questions.

- The optimal online algorithm has n distinct thresholds, one for each X_i, which is at the other extreme compared to the single-threshold algorithms. What if we are allowed to use at most k-thresholds for $k > 1$? How does the competitive ratio improve with k, starting with the poly-logarithmic factor we show for $k = 1$?
- If one has only sample access to $\mathcal{D}$, how does the competitive ratio of the optimal algorithm change with the number of samples? This question, with importance in practical applications when the distributions are not fully known, has been studied extensively in the rewards maximization setting [AKW19, CDFS22, RWW20].
- Can we derive optimal distribution-sensitive guarantees for the rewards maximization setting to complement the uniform bound of ≈ 0.745 due to [CFH+21]?

Acknowledgements: The authors would like to thank Matt Weinberg for sharing helpful references on prophet inequalities and Thomas Kesselheim and Brendan Lucier for constructive and helpful feedback on an earlier version of the paper. The first author would also like to thank Chandra Chekuri, Efstathios Konstantinos Chrontsios Garitsis, Kesav Krishnan, Shubhang Kulkarni, Rhea Jain, Robert Andrews and Maria Paula Parga Nina for their insightful comments and helpful discussions.

References

[ABF+17] Marek Adamczyk, Allan Borodin, Diodato Ferraioli, Bart De Keijzer, and Stefano Leonardi. Sequential posted-price mechanisms with correlated valuations. *ACM Trans. Econ. Comput.*, 5(4), 12 2017.

[ABG+20] Reza Alijani, Siddhartha Banerjee, Sreenivas Gollapudi, Kamesh Munagala, and Kangning Wang. Predict and match: Prophet inequalities with uncertain supply. *Proc. ACM Meas. Anal. Comput. Syst.*, 4(1), may 2020.

[ABN08] Barry C. Arnold, N. Balakrishnan, and H. N. Nagaraja. *A First Course in Order Statistics.* Society for Industrial and Applied Mathematics, 2008.

[ACK18] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the 1-1/e barrier. In Éva Tardos, Edith Elkind, and Rakesh Vohra, editors, *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 303–318. ACM, 2018.

[AEE+17] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert D. Kleinberg, and Brendan Lucier. Beating 1-1/e for ordered prophets. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 61–71. ACM, 2017.

[AKS21] Sepehr Assadi, Thomas Kesselheim, and Sahil Singla. *Improved Truthful Mechanisms for Subadditive Combinatorial Auctions: Breaking the Logarithmic Barrier*, pages 653–661. SIAM, 2021.

[AKW19] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prior independent mechanisms via prophet inequalities with limited information. *Games Econ. Behav.*, 118:511–532, 2019.

[Ala14] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.

[AM22] Nick Arnosti and Will Ma. Tight guarantees for static threshold policies in the prophet secretary problem. In David M. Pennock, Ilya Segal, and Sven Seuken, editors, *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, page 242. ACM, 2022.

[AS20] Sepehr Assadi and Sahil Singla. Improved truthful mechanisms for combinatorial auctions with submodular bidders. *SIGecom Exch.*, 18(1):19–27, 2020.

[AW20] Marek Adamczyk and Michal Wlodarczyk. Multi-dimensional mechanism design via random order contention resolution schemes. *SIGecom Exch.*, 17(2):46–53, 1 2020.

[BBDS17] Moshe Babaioff, Liad Blumrosen, Shaddin Dughmi, and Yaron Singer. Posting prices with unknown distributions. *ACM Trans. Econ. Comput.*, 5(2), 3 2017.

[BBK21] Alexander Braun, Matthias Buttkus, and Thomas Kesselheim. Asymptotically optimal welfare of posted pricing for multiple items with MHR distributions. In Petra Mutzel, Rasmus Pagh, and Grzegorz Herman, editors, *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, volume 204 of *LIPIcs*, pages 22:1–22:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[BCKW15] Patrick Briest, Shuchi Chawla, Robert Kleinberg, and S. Matthew Weinberg. Pricing lotteries. *Journal of Economic Theory*, 156:144–174, 2015. Computer Science and Economic Theory.

[BGGM12] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget-constrained auctions with heterogeneous items. *Theory Comput.*, 8(1):429–460, 2012.

[BGMS21] Brian Brubach, Nathaniel Grammel, Will Ma, and Aravind Srinivasan. Improved guarantees for offline stochastic matching via new ordered contention resolution schemes. In *NeurIPS*, 2021.

[BH08] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: An asymptotic analysis. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, EC '08, page 49, New York, NY, USA, 2008. Association for Computing Machinery.

[BPH96] R.E. Barlow, F. Proschan, and L.C. Hunter. *Mathematical Theory of Reliability.* Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 1996.

[CCD+23] José Correa, Andrés Cristi, Paul Dütting, Mohammad Hajiaghayi, Jan Olkowski, and Kevin Schewior. Trading prophets, 2023.

[CD15] Yang Cai and Constantinos Daskalakis. Extreme value theorems for optimal multidimensional pricing. *Games Econ. Behav.*, 92:266–305, 2015.

[CDFS22] José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for independent and identically distributed random variables from an unknown distribution. *Math. Oper. Res.*, 47(2):1287–1309, 2022.

[CDL21] Shuchi Chawla, Nikhil R. Devanur, and Thodoris Lykouris. Static pricing for multi-unit prophet inequalities (extended abstract). In Michal Feldman, Hu Fu, and Inbal Talgam-Cohen, editors, *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings*, volume 13112 of *Lecture Notes in Computer Science*, pages 545–546. Springer, 2021.

[CFH+18] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *SIGecom Exch.*, 17(1):61–70, 2018.

[CFH+21] José R. Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms and optimal threshold strategies for random arrivals. *Math. Oper. Res.*, 46(4):1452–1478, 2021.

[CFPV19] José Correa, Patricio Foncea, Dana Pizarro, and Victor Verdugo. From pricing to prophets, and back! *Operations Research Letters*, 47(1):25–29, 2019.

[CGKM20] Shuchi Chawla, Kira Goldner, Anna R. Karlin, and J. Benjamin Miller. Non-adaptive matroid prophet inequalities. *CoRR*, abs/2011.09406, 2020.

[CHK07] Shuchi Chawla, Jason D. Hartline, and Robert Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, EC '07, page 243–251, New York, NY, USA, 2007. Association for Computing Machinery.

[CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the Forty-Second ACM Symposium on Theory of Computing*, STOC '10, page 311–320, New York, NY, USA, 2010. Association for Computing Machinery.

[CL21] Chandra Chekuri and Vasilis Livanos. On submodular prophet inequalities and correlation gap. In Ioannis Caragiannis and Kristoffer Arnsfelt Hansen, editors, *Algorithmic Game Theory - 14th International Symposium, SAGT 2021, Aarhus, Denmark, September 21-24, 2021, Proceedings*, volume 12885 of *Lecture Notes in Computer Science*, page 410. Springer, 2021.

[CSZ20] Jose Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. *Mathematical Programming*, 08 2020.

[CVZ14] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.

[DDT14] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. The complexity of optimal mechanism design. In Chandra Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1302–1318. SIAM, 2014.

[DDT15] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Strong duality for a multiple-good monopolist. In Tim Roughgarden, Michal Feldman, and Michael Schwarz, editors, *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, pages 449–450. ACM, 2015.

[DFKL20] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582, 2020.

[DKL0] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. *SIAM Journal on Computing*, 0(0):FOCS20–239–FOCS20–275, 0.

[DLMF22] *NIST Digital Library of Mathematical Functions*. http://dlmf.nist.gov/, Release 1.1.7 of 2022-10-15, 2022. F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, eds.

[Dob21] Shahar Dobzinski. Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders. *SIAM J. Comput.*, 50(3), 2021.

[DP02] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, 2 edition, 2002.

[DRY15] Peerapong Dhangwatnotai, Tim Roughgarden, and Qiqi Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.

[Dug20] Shaddin Dughmi. The outer limits of contention resolution on matroids and connections to the secretary problem. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, volume 168 of *LIPIcs*, pages 42:1–42:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[Dug22] Shaddin Dughmi. Matroid secretary is equivalent to contention resolution. In Mark Braverman, editor, *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, volume 215 of *LIPIcs*, pages 58:1–58:23. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[DV16] Shahar Dobzinski and Jan Vondrák. Impossibility results for truthful combinatorial auctions with submodular valuations. *J. ACM*, 63(1):5:1–5:19, 2016.

[DW12] Constantinos Daskalakis and Seth Matthew Weinberg. Symmetries and optimal multi-dimensional mechanism design. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, page 370–387, New York, NY, USA, 2012. Association for Computing Machinery.

[EFGT22] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet matching with general arrivals. *Mathematics of Operations Research*, 47(2):878–898, 2022.

[EHKS18] Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In Artur Czumaj, editor, *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 700–714. SIAM, 2018.

[EHLM17] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.

[FGL15] Michal Feldman, Nick Gravin, and Brendan Lucier. *Combinatorial Auctions via Posted Prices*, pages 123–135. SIAM, 2015.

[FSZ21] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM J. Comput.*, 50(2):255–300, 2021.

[Gau98] Walter Gautschi. The incomplete gamma functions since tricomi. In *In Tricomi's Ideas and Contemporary Applied Mathematics, Atti dei Convegni Lincei, n. 147, Accademia Nazionale dei Lincei*, pages 203–237, 1998.

[GKL17] Yiannis Giannakopoulos, Elias Koutsoupias, and Philip Lazos. Online Market Intermediation. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 47:1–47:14, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[GPZ21] Yiannis Giannakopoulos, Diogo Poças, and Keyu Zhu. Optimal pricing for mhr and λ-regular distributions. *ACM Trans. Econ. Comput.*, 9(1), 1 2021.

[GW23] Nick Gravin and Hongao Wang. Prophet inequality for bipartite matching: Merits of being simple and nonadaptive. *Mathematics of Operations Research*, 48(1):38–52, 2023.

[HF06] Laurens Haan and Ana Ferreira. *Extreme Value Theory: An Introduction*. Springer, 01 2006.

[HK82] T. P. Hill and Robert P. Kertz. Comparisons of stop rule and supremum expectations of i.i.d. random variables.

Ann. Probab., 10(2):336–345, 05 1982.
[HK92] Theodore P Hill and Robert P Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.
[HKS07] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22Nd National Conference on Artificial Intelligence - Volume 1*, AAAI'07, pages 58–65. AAAI Press, 2007.
[HN19] Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *J. Econ. Theory*, 183:991–1029, 2019.
[HR09] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. *SIGecom Exch.*, 8(1), 2009.
[JMZ22a] Jiashuo Jiang, Will Ma, and Jiawei Zhang. *Tight Guarantees for Multi-unit Prophet Inequalities and Online Stochastic Knapsack*, pages 1221–1246. SIAM, 2022.
[JMZ22b] Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tightness without counterexamples: A new approach and new results for prophet inequalities, 2022.
[KAnA11] M. Kayid, H. Al-nahawati, and I.A. Ahmad. Testing behavior of the reversed hazard rate. *Applied Mathematical Modelling*, 35(5):2508–2515, 2011.
[Ker86] Robert P Kertz. 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, 1986.
[KP02] J. D. Kalbfleisch and R. L. Prentice. *The Statistical Analysis of Failure Time Data*. John Wiley & Sons, 2nd edition, 2002.
[KS77] Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.*, 83(4):745–747, 07 1977.
[KS78] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.
[KW19] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities and applications to multi-dimensional mechanism design. *Games Econ. Behav.*, 113:97–115, 2019.
[LLP+21] Allen Liu, Renato Paes Leme, Martin Pál, Jon Schneider, and Balasubramanian Sivan. Variable decomposition for prophet inequalities and optimal ordering. In Péter Biró, Shuchi Chawla, and Federico Echenique, editors, *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, page 692. ACM, 2021.
[LS18] Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In Hannah Bast, Grzegorz Herman, and Yossi Azar, editors, *26th European Symposium on Algorithms, ESA 2018*, Leibniz International Proceedings in Informatics, LIPIcs, Germany, August 2018. Schloss Dagstuhl- Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing.
[Luc] Brendan Lucier. personal communication.
[Luc17] Brendan Lucier. An economic view of prophet inequalities. *SIGecom Exch.*, 16(1):24–47, September 2017.
[MS83] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.
[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
[PRSW22] Tristan Pollner, Mohammad Roghani, Amin Saberi, and David Wajc. Improved online contention resolution for matchings and applications to the gig economy. In David M. Pennock, Ilya Segal, and Sven Seuken, editors, *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, pages 321–322. ACM, 2022.
[QM99] Feng Qi and Jia-Qiang Mei. Some inequalities of the incomplete gamma and related functions. *Zeitschrift für Analysis und ihre Anwendungen*, 18(3):793–799, 1999.
[QRVW19] Junjie Qin, Ram Rajagopal, Shai Vardi, and Adam Wierman. Convex prophet inequalities. *SIGMETRICS Perform. Eval. Rev.*, 46(2):85–86, 1 2019.
[RH04] Marvin Rausand and Arnljot Høyland. *System Reliability Theory: Models, Statistical Methods and Applications*. Wiley-Interscience, Hoboken, NJ, 2004.
[RS17] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687. SIAM, 2017. Longer ArXiv version is at http://arxiv.org/abs/1611.00665.
[Rub16] Aviad Rubinstein. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 324–332, New York, NY, USA, 2016. Association for Computing Machinery.
[RWW20] Aviad Rubinstein, Jack Z. Wang, and S. Matthew Weinberg. Optimal single-choice prophet inequalities from samples. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPIcs*, pages 60:1–60:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
[SC84] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random

variables. *The Annals of Probability*, 12(4):1213–1216, 1984.
[SM02] Uwe Saint-Mont. A simple derivation of a complicated prophet region. *J. Multivar. Anal.*, 80(1):67–72, 1 2002.
[Yan11] Qiqi Yan. Mechanism design via correlation gap. In Dana Randall, editor, *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, pages 710–719. SIAM, 2011.

Appendix

A Counterexamples

A.1 Single-Threshold Counterexample For the cost-prophet inequality setting, a natural approach that is seemingly intuitive is to set a single threshold T close to $\mu_n = \mathbb{E}[\min_i X_i]$ since, if n is large enough, with good probability there will be a realization below the threshold and, this way, one would achieve a very good competitive ratio.

We present an example that shows why this natural intuition fails.

Example. Consider the exponential distribution, for which $F(x) = 1 - e^{-x}$, $f(x) = e^{-x}$, $H(x) = x$, $E[X] = 1$ and

$$\mu_n = \int_0^\infty e^{-nx}\, dx = \frac{1}{n}.$$

In our attempt to achieve a constant competitive ratio, we set a threshold $T = \frac{c}{n}$ for some constant $c > 0$. If there exists a realization of $X_1, \ldots, X_{n-1}$ that is below T, then we would select it; otherwise we are forced to select X_n and obtain a cost equal to $\mathbb{E}[X]$.

The probability that there exists a realization of $X_1, \ldots, X_{n-1}$ that is below T is $1 - (1 - F(T))^{n-1}$. Thus, the expected cost of our algorithm is

$$\begin{aligned}
\mathbb{E}[ALG_n] &= \left(1 - (1 - F(T))^{n-1}\right) \mathbb{E}\left[X \mid X \le T\right] + (1 - F(T))^{n-1}\, E[X] \\
&= \left(1 - e^{-(n-1)T}\right) \mathbb{E}\left[X \mid X \le T\right] + e^{-(n-1)T} \cdot 1 \\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \mathbb{E}\left[X \mid X \le c/n\right] + e^{-c\frac{n-1}{n}} \\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{\int_0^{\frac{c}{n}} x f(x)\, dx}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} \\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{\int_0^{\frac{c}{n}} x e^{-x}\, dx}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} \\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{1 - e^{-c/n} - \frac{c}{n}\, e^{-c/n}}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} \\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) + e^{-c\frac{n-1}{n}}.
\end{aligned}$$

Thus, the competitive ratio is

$$R(n) = \frac{\mathbb{E}[ALG_n]}{\mu_n} = n\left(\left(1 - e^{-c\frac{n-1}{n}}\right)\left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) + e^{-c\frac{n-1}{n}}\right).$$

Notice that, as $n \to +\infty$, we have

$$\lim_{n \to +\infty} n\left(1 - e^{-c\frac{n-1}{n}}\right)\left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) = \frac{c\, e^{-c}\, (e^c - 1)}{2},$$

but

$$\lim_{n \to +\infty} n\, e^{-c\frac{n-1}{n}} = +\infty,$$

and thus the competitive ratio of this algorithm is infinite. ∎

A.2 I.I.D. Counterexample In the I.I.D. cost prophet inequality setting, one cannot hope to obtain a bounded competitive factor for *all* distributions. The following counterexample is due to Lucier [Luc].

PROPOSITION A.1. ([Luc]) *For any $n \geq 2$, there exists an instance of the I.I.D. cost prophet inequality problem for which no algorithm is α-factor competitive for any $\alpha > 0$.*

Proof. Let $n = 2$ and consider the equal-revenue distribution, with support $[1, +\infty)$ and $F(x) = 1 - 1/x$. For this distribution, we have

$$\mathbb{E}[X] = \int_0^\infty (1 - F(x))\, dx = 1 + \int_1^\infty (1 - F(x))\, dx = 1 + \int_1^\infty \frac{1}{x}\, dx = +\infty.$$

In this case, the expected cost of any algorithm is $\mathbb{E}[ALG] = +\infty$, regardless of whether it stops at X_1 or at X_2. However, the prophet is always able to select the minimum of X_1 and X_2, which is

$$\text{OPT} = \mu_2 = \int_0^\infty (1 - F(x))^2\, dx = 1 + \int_1^\infty \frac{1}{x^2}\, dx = 2.$$

Therefore, no algorithm can achieve a finite competitive ratio.

Notice that the above counterexample can be easily extended to any n. Due to the recursive nature of the optimal online algorithm, we have $\mathbb{E}[ALG] = +\infty$ regardless of which X_i the algorithm chooses to stop at. However, μ_n is finite for any $n \geq 2$. ◻

A.3 Extreme Value Index

PROPOSITION A.2. *The tight competitive ratio for the exponential distribution ($F_1(x) = 1 - e^{-x}$, $x \in [0, +\infty)$) and the uniform distribution ($F_2(x) = x$, $x \in [0, 1]$) is 2, whereas for the Weibull distribution ($F_3(x) = 1 - e^{-\frac{x^2}{2}}$, $x \in [0, +\infty)$) it is $\sqrt{\frac{6}{\pi}}$.*

Notice that the exponential and Weibull distributions belong to the Gumbel class of asymptotic distributions, for which the extreme value index $\gamma = 0$. However, the uniform distribution belongs to the reversed Weibull class of asymptotic distributions and has extreme value index $\gamma = -1$.

Proof. [Proof of Proposition A.2] We analyze Lemma 3.5 via the use of Lemma 3.4 for each of the aforementioned distributions:

1. Exponential: We have

$$\phi(\lambda, n) = (n-1)\frac{\int_0^{\frac{\lambda}{n-1}} 1 - e^{-x}\, dx}{\lambda} = 1 - \frac{n-1}{\lambda}\left(1 - e^{-\frac{\lambda}{n-1}}\right).$$

   Also,

$$\frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n} = \frac{n\left(\frac{1}{n-1} + \frac{1}{n}\right) - 1}{n} = \frac{1}{n-1}.$$

   The infimum over all $\lambda > 1$ for which

$$1 - \frac{n-1}{\lambda}\left(1 - e^{-\frac{\lambda}{n-1}}\right) \geq \frac{1}{n-1}$$

   is 2.

2. Uniform: We have

$$\phi(\lambda, n) = (n)\frac{\int_0^{\frac{\lambda}{n}} x\, dx}{\lambda} = \frac{\lambda}{2n}.$$

   Also,

$$\frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n} = \frac{1}{n}.$$

The infimum over all $\lambda > 1$ for which

$$\frac{\lambda}{2n} \geq \frac{1}{n}$$

is 2.

3. Weibull: We have

$$\phi(\lambda, n) = \sqrt{\frac{2(n-1)}{\pi}} \frac{\int_0^{\lambda\sqrt{\frac{\pi}{2(n-1)}}} 1 - e^{-\frac{x^2}{2}}\, dx}{\lambda} = \frac{\sqrt{2(n-1)}}{\lambda\sqrt{\pi}} \left(\lambda\sqrt{\frac{\pi}{2(n-1)}} - \sqrt{\frac{\pi}{2}}\mathrm{Erf}\left(\lambda\frac{\sqrt{\pi}}{2\sqrt{n-1}}\right)\right),$$

where $\mathrm{Erf}(x)$ is the error function, defined as $\mathrm{Erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2}\, dt$. Also,

$$\frac{\frac{\mu_{2:n}}{\mu_n} - 1}{n} = \sqrt{\frac{n}{n-1}} - 1.$$

The infimum over all $\lambda > 1$ for which

$$\frac{\sqrt{2(n-1)}}{\lambda\sqrt{\pi}} \left(\lambda\sqrt{\frac{\pi}{2(n-1)}} - \sqrt{\frac{\pi}{2}}\mathrm{Erf}\left(\lambda\frac{\sqrt{\pi}}{2\sqrt{n-1}}\right)\right) \geq \sqrt{\frac{n}{n-1}} - 1$$

is $\sqrt{\frac{6}{\pi}}$.

□

A.4 Bounded Support Distributions

Observation A.1. *For any $\alpha > 0$, there exists a distribution $\mathcal{D}_\alpha$, supported on $[0, 1]$ such that for the I.I.D. cost prophet inequality setting with random variables drawn from $\mathcal{D}_\alpha$*

1. *there exists an α-competitive cost prophet inequality, and*
2. *there does not exist an $(\alpha - \varepsilon)$-competitive cost prophet inequality for any constant $\varepsilon > 0$.*

Proof. Consider the Beta distribution, which is supported on $[0, 1]$ and is parameterized by $\alpha > 0$, and for which $F_\alpha(x) = x^\alpha$. For this distribution, we have

$$H_\alpha(x) = -\log\left(1 - F_\alpha(x)\right) = \log\left(\frac{1}{1 - x^\alpha}\right).$$

The Puiseux series of H_α around $x = 0$ is

$$H_\alpha(x) = \sum_{k \geq 1} \frac{x^{k\alpha}}{k},$$

which converges for $x \in [0, 1)$[5]. Thus, we observe that, for this distribution, $d_1 = \alpha$, and from Theorem 1.2, we know that there exists a tight $\frac{(1+1/\alpha)^{1/\alpha}}{\Gamma(1+1/\alpha)}$-cost prophet inequality. □

B Missing Proofs

B.1 Proof of Observation 2.1

Observation B.1. *For $n \geq 1$,*

$$\mu_n = \mathbb{E}\left[\min_{i=1}^{n} X_i\right] = \int_0^\infty \left(1 - F(s)\right)^n ds.$$

[5] The Puiseux series and H are equal also for $x \to 1$, since both diverge to $+\infty$.

Proof. Let $Y_n = \min_{i=1}^n X_i$, then the CDF F_{Y_n} of Y_n is

$$F_{Y_n}(x) = \Pr[Y_n \le x] = 1 - \Pr[Y_n > x] = 1 - \prod_{i=1}^{n} \Pr[X_i > x] = 1 - (1 - F(x))^n, \forall x \in [0, +\infty).$$

Recall that for a random variable X, we have

$$\mathbb{E}[X] = \int_0^\infty x f_X(x)\, dx = \int_0^\infty f_X(x) \int_0^x dt\, dx.$$

By changing the order of integration, we obtain

$$\mathbb{E}[X] = \int_0^\infty \int_0^x f_X(x)\, dx\, dt = \int_0^\infty \Pr[X \ge t]\, dt = \int_0^\infty (1 - F_X(t))\, dt.$$

Using this, we get that the expected cost of the prophet (offline optimum), denoted by μ_n is,

$$\mu_n = \mathbb{E}[Y_n] = \int_0^\infty (1 - F_{Y_n}(s))\, ds = \int_0^\infty (1 - (1 - (1 - F(s))^n))\, ds = \int_0^\infty (1 - F(s))^n\, ds.$$

∎

B.2 Proof of Observation 2.3

Observation B.2. *Consider an entire distribution $\mathcal{D}$ supported on $[0, +\infty)$ with cumulative hazard rate $H(x) = \sum_{i=1}^\infty a_i x^{d_i}$, where $d_1 < d_2 < \dots$. Then, $a_1 > 0$ and $d_1 > 0$.*

Proof. Once can easily see that $a_1 > 0$ since H is non-negative. Note that, for any choice of $a_1, a_2, \dots$ and $d_1 < d_2 < \dots$, since the Puiseux series of H is convergent for every x in the support of $\mathcal{D}$, there exists a small enough $x_* \in [0, 1)$ such that,

$$\left|a_1 x_*^{d_1}\right| > \sum_{i=2}^\infty \left|a_i x_*^{d_i}\right|.$$

Thus, if $a_1 < 0$, we have $H(x_*) < 0$, a contradiction.

Next we show that $d_1 \ge 0$. Consider the derivative of H, namely $h(x) = \sum_{i=1}^\infty a_i d_i x^{d_i - 1}$. Again, given that fact that $d_1 < d_i$ for all $i \ge 2$, there exists y_* such that

$$\left|a_1 d_1 y_*^{d_1 - 1}\right| > \sum_{i=2}^\infty \left|a_i d_i y_*^{d_i - 1}\right|.$$

Thus since $a_1 > 0$, we have $a_1 \cdot d_1 < 0$ which implies $h(y_*) < 0$, a contradiction to h being non-negative. ∎

B.3 Proof of Proposition 3.1

Proposition B.1. *For any instance of the cost prophet inequality setting, one can achieve the optimal competitive ratio with a threshold-based oblivious algorithm.*

Proof. Since every algorithm has to select a value, if an algorithm observes the realization of X_n, it is forced to select it. When an algorithm sees X_{n-1}, it has to decide whether to select it or not. Whatever the decision process of the algorithm, let $p^{\mathcal{A}}(r \mid X_{n-1} = z)$ be the probability that algorithm $\mathcal{A}$ selects the realization of X_i, given X_i. Then, the expected cost of $\mathcal{A}$ is

$$\sum_{z \ge 0} z\, p^{\mathcal{A}}(r \mid X_{n-1} = z) + \left(1 - \sum_{z \ge 0} p^{\mathcal{A}}(r \mid X_{n-1} = z)\right) \mathbb{E}[X_n].$$

For a fixed choice of $L = \sum_{z \ge 0} p^{\mathcal{A}}(r \mid X_{n-1} = z)$, to maximize this quantity, $\mathcal{A}$ will greedily assign all the probability mass of L to the lowest values z. Thus, the only choice $\mathcal{A}$ has to make is L itself, which is equal to $\Pr\left[X_{n-1} \le F^{-1}(L)\right]$. Therefore, every choice of L implies a threshold, namely $F^{-1}(L)$.

Finally, for the remaining random variables, the observation holds via induction, since the random variables are I.I.D. ∎

B.4 Proof of Lemma 3.1

Lemma B.1. *For the cost prophet inequality problem with random variables $X_1, X_2, \ldots, X_n$, $\tau_n = +\infty$ for every algorithm. For $1 \le i \le n-1$, the optimal threshold for the random variable X_i is*

$$\tau_i = G(n-i).$$

Proof. The lemma follows by backwards induction on n.
Base case. Since we are forced to select a single value, if the algorithm ever observes X_n, it must select its realization. This is equivalent to $\tau_n = +\infty$. It then follows that $G(1) = \mathbb{E}_{X \sim \mathcal{D}}[X]$.
Induction. Consider the i-th step, where $i < n$. For our induction hypothesis, assume that $\tau_j = G(n-j)$ for all $i < j < n$. Conditioned to the fact that the algorithm has reached the i-th step, the expected cost of the optimal algorithm is $G(n-i+1)$; i.e. the cost that the optimal algorithms expects to receive from the remaining $n-i+1$ variables. Since τ_i is the optimal threshold for X_i, we obtain the following recurrence for $G(n-i+1)$.

$$G(n-i+1) = F(\tau_i)\,\mathbb{E}\left[X \mid X \le \tau_i\right] + (1 - F(\tau_i))\,G(n-i). \tag{B.1}$$

This is because with probability $F(\tau_i)$ we select X_i and therefore receive cost $\mathbb{E}\left[X \mid X \le \tau_i\right]$, and with probability $1 - F(\tau_i)$, we ignore X_i and we receive cost equal to the expected value of the optimal algorithm on $X_{i+1}, \ldots, X_n$, i.e. $G(n-i)$. Thus, it suffices to show that setting $\tau_i = G(n-i)$ minimizes $G(n-i+1)$.

We rearrange (B.1) and obtain

$$\begin{aligned}
G(n-i+1) &= F(\tau_i)\,\mathbb{E}\left[X \mid X \le \tau_i\right] + (1 - F(\tau_i))\,G(n-i) \\
&= F(\tau_i) \cdot \frac{\int_0^{\tau_i} u f(u)\,du}{F(\tau_i)} + (1 - F(\tau_i))\,G(n-i) \\
&= \int_0^{\tau_i} u f(u)\,du + (1 - F(\tau_i))\,G(n-i) \\
&= \int_0^{\tau_i} u\,(F(u))'\,du + (1 - F(\tau_i))\,G(n-i) \\
&= \left[uF(u)\right]_0^{\tau_i} - \int_0^{\tau_i} F(u)\,du + (1 - F(\tau_i))\,G(n-i) \\
&= \tau_i F(\tau_i) - \int_0^{\tau_i} F(u)\,du + (1 - F(\tau_i))\,G(n-i).
\end{aligned}$$

where the second equality follows by the definition of $\mathbb{E}\left[X \mid X \le \tau_i\right]$ and the second-to-last equality follows via integration by parts.

We will show that the optimal threshold at the i-th step is

$$\tau_i = G(n-i).$$

In other words, we will show that

$$\begin{aligned}
G(n-i)F(G(n-i)) - \int_0^{G(n-i)} F(u)\,du + (1 - F(G(n-i)))\,G(n-i) \\
\le \tau_i F(\tau_i) - \int_0^{\tau_i} F(u)\,du + (1 - F(\tau_i))\,G(n-i),
\end{aligned} \tag{B.2}$$

for any $\tau_i \ne G(n-i)$. Rearranging (B.2), we get

$$\begin{aligned}
G(n-i)F(G(n-i)) - \int_0^{G(n-i)} F(u)\,du + (1 - F(G(n-i)))\,G(n-i) \\
\le \tau_i F(\tau_i) - \int_0^{\tau_i} F(u)\,du + (1 - F(\tau_i))\,G(n-i) \iff
\end{aligned}$$

$$G(n-i) - \int_0^{G(n-i)} F(u)\,du \le \tau_i F(\tau_i) - \int_0^{\tau_i} F(u)\,du + G(n-i) - F(\tau_i)G(n-i) \iff$$
$$F(\tau_i)\left(G(n-i) - \tau_i\right) \le \int_0^{G(n-i)} F(u)\,du - \int_0^{\tau_i} F(u)\,du \iff$$
$$F(\tau_i)\left(G(n-i) - \tau_i\right) \le \int_{\tau_i}^{G(n-i)} F(u)\,du. \tag{B.3}$$

We distinguish between two cases: $\tau_i < G(n-i)$ and $\tau_i > G(n-i)$. In the case where $\tau_i < G(n-i)$, (B.3) becomes

$$F(\tau_i) \le \frac{\int_{\tau_i}^{G(n-i)} F(u)\,du}{G(n-i) - \tau_i},$$

which is true by the mean value theorem, since F is increasing and $\tau_i < G(n-i)$. Similarly, in the case where $\tau_i > G(n-i)$, (B.3) becomes

$$F(\tau_i) \ge \frac{\int_{\tau_i}^{G(n-i)} F(u)\,du}{G(n-i) - \tau_i} = \frac{\int_{G(n-i)}^{\tau_i} F(u)\,du}{\tau_i - G(n-i)},$$

which is again true by the mean value theorem, since F is increasing and $\tau_i > G(n-i)$.

We conclude that the optimal threshold for X_i is

$$\tau_i = G(n-i).$$

∎

C Background on the Gamma Function

The Gamma function $\Gamma(x)$ extends the factorial function to complex numbers. In particular,

$$\Gamma(n+1) = n!$$

for every $n \in \mathbb{N}$.

Here we give a brief and incomplete primer on the Gamma function, to assist the reader. However, for a more extensive treatment along with many folklore results about the function, see [Gau98].

Definition C.1. (Gamma (Γ) Function) *For every $x > 0$, the* Gamma function *is defined as*

$$\Gamma(x) = \int_0^\infty t^{x-1} e^{-t}\,dt.$$

Like the factorial function, the Gamma function also satisfies the following recurrence

$$\Gamma(x+1) = x\Gamma(x).$$

The following fact is closely related to Stirling's approximation for the Gamma function and is due to [DLMF22, Eq. 5.11.E7].

Fact C.1. *For $a > 0$ and $b \in \mathbb{R}$, we have*

$$\Gamma(a+b) \le \sqrt{2\pi}\left(\frac{a}{e}\right)^a \cdot a^b.$$

Of particular use to us are the following special functions that are related to the Gamma function.

Definition C.2. (Upper ($\Gamma(\cdot,\cdot)$ and Lower $\gamma(\cdot,\cdot)$ Incomplete Gamma Functions) *For every $s > 0, x \ge 0$, the* Upper Incomplete Gamma function *is defined as*

$$\Gamma(s,x) = \int_x^\infty t^{s-1} e^{-t}\,dt,$$

whereas the Lower Incomplete Gamma function *is defined as*

$$\gamma(s,x) = \int_0^x t^{s-1} e^{-t}\,dt.$$

For every $s > 0, x \geq 0$, we have

$$\Gamma(s,x) + \gamma(s,x) = \Gamma(s).$$

Next, we describe a few known results about the lower incomplete Gamma function that we use throughout the paper.

FACT C.2. *For the lower incomplete Gamma function $\gamma(s,x)$ with $s, x > 0$, we have*

$$\gamma(s,x) = x^s \sum_{k=0}^{\infty} \frac{(-x)^k}{k!\ (s+k)}.$$

Proof. By the definition of the lower incomplete Gamma function, we have

$$\gamma(s,x) = \int_0^x t^{s-1} e^{-t}\, dt = \int_0^x \sum_{k=0}^{\infty} (-1)^k \frac{t^{s+k-1}}{k!} = \sum_{k=0}^{\infty} (-1)^k \frac{x^{s+k}}{k!\ (s+k)} = x^s \sum_{k=0}^{\infty} \frac{(-x)^k}{k!\ (s+k)}.$$

∎

The following fact follows easily via Fact C.2.

FACT C.3. *We have that, as $x \to 0$,*

$$\frac{\gamma(s,x)}{x^s} \to s^{-1}.$$

The following claim is due to Qi and Mei [QM99].

CLAIM C.1. *[See 3.1 in [QM99]] For small enough x, we have*

$$\gamma(s,x) \leq s^{-1}\, x^{s-1}\, e^{-x}.$$