

IID Prophet Inequality with Random Horizon: Going Beyond Increasing Hazard Rates

Giordano Giambartolomei ✉

Department of Informatics, King's College London, UK

Frederik Mallmann-Trenn ✉ 🏠

Department of Informatics, King's College London, UK

Raimundo Saona ✉ 🏠 

Institute of Science and Technology Austria, Klosterneuburg, Austria

Abstract

Prophet inequalities are a central object of study in optimal stopping theory. In the *iid* model, a gambler sees values in an online fashion, sampled independently from a given distribution. Upon observing each value, the gambler either accepts it as a reward, or irrevocably rejects it and proceeds to observe the next value. The goal of the gambler, who cannot see the future, is to maximise the expected value of the reward while competing against the expectation of a prophet (the offline maximum). In other words, one seeks to maximise the gambler-to-prophet ratio of the expectations.

This model has been studied with infinite, finite and unknown number of values. When the gambler faces a random number of values, the model is said to have a random horizon. We consider the model in which the gambler is given *a priori* knowledge of the horizon's distribution. Alijani et al. (2020) designed a single-threshold algorithm achieving a ratio of $1/2$ when the random horizon has an increasing hazard rate and is independent of the values. We prove that with a single threshold, a ratio of $1/2$ is actually achievable for several larger classes of horizon distributions, with the largest being known as the \mathcal{G} class in reliability theory. Moreover, we show that this does not extend to its dual, the $\overline{\mathcal{G}}$ class (which includes the decreasing hazard rate class), while it can be extended to low-variance horizons. Finally, we construct the first example of a family of horizons, for which multiple thresholds are necessary to achieve a nonzero ratio. We establish that the Secretary Problem optimal stopping rule provides one such algorithm, paving the way towards the study of the model beyond single-threshold algorithms.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Online algorithms, Prophet Inequality, Random Horizon, Secretary Problem

Digital Object Identifier 10.4230/LIPIcs.ICALP.2025.87

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: <https://arxiv.org/abs/2407.11752>

Funding *Giordano Giambartolomei*: EPSRC grants EP/W005573/1 and EP/X021696/1.

Frederik Mallmann-Trenn: EPSRC grant EP/W005573/1.

Raimundo Saona: ERC grant CoG 863818 (ForM-SMArt), ANID Chile grant ACT210005, French Agence Nationale de la Recherche (ANR) grant ANR-21-CE40-0020 (CONVERGENCE), and Austrian Science Fund (FWF) grant 10.55776/COE12.

Acknowledgements We would like to thank José Correa for his precious advice, Bruno Ziliotto and Vasilis Livanos for early conversations.

1 Introduction

Prophet inequalities are a central object of study in optimal stopping theory. A gambler sees nonnegative values in an online fashion, sampled from an instance of independent random variables $\{X_i\}$ with known distributions $\{\mathcal{D}_i\}$, in adversarial, random or selected order,



© Giordano Giambartolomei, Frederik Mallmann-Trenn, and Raimundo Saona;
licensed under Creative Commons License CC-BY 4.0

52nd International Colloquium on Automata, Languages, and Programming (ICALP 2025).

Editors: Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis

Article No. 87; pp. 87:1–87:21



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



depending on the particular model. When observing each value, the gambler either accepts it as a reward, or irrevocably rejects it and proceeds with observing the next value. The goal of the gambler, who cannot see the future, is to maximise the expected value of the reward while competing against the expectation of a prophet (out of metaphor, the offline maximum or supremum, depending on whether the instance is finite or not). In other words, one seeks to maximise the gambler-to-prophet ratio of the expectations.

The gambler represents any online decision maker, such as an algorithm or stopping rule. Probabilistically, we will refer to it as a *stopping time* τ . Informally, the term *online* implies that the gambler, unable to see the future, will always stop at a time τ such that the event $\{\tau = i\}$ depends only on the first i values observed.

1.1 Prophet inequality models

Several models and extensions of prophet inequalities are present in the literature. We introduce the variants upon which our model is built, briefly reviewing the state-of-the-art.

The very first prophet inequality model, typically referred to as classical Prophet Inequality (PI), is due to Krengel and Sucheston [25, 26]. The given instance is composed of countably many integrable independent nonnegative random variables $\{X_i\}$ with known distributions $\{\mathcal{D}_i\}$ in a fixed given order, usually referred to as *adversarial*. The gambler-to-prophet ratio to be maximised will therefore be $\mathbb{E}X_\tau$ over $\mathbb{E}\sup_i X_i$. When working with *infinite* instances, one often considers only finite stopping rules ($\mathbb{P}(\tau < \infty) = 1$) and variables such that $\mathbb{E}\sup_i X_i < \infty$. The $1/2$ -hardness of PI (shown by Garlin [25]) is tight [26]. A single-threshold 2-approximation is also possible [35].

IID PI is specialised to $\{X_i\}$ independent and identically distributed (*iid*) according to a given distribution \mathcal{D} . The hardness of this problem is $1/\beta \approx 0.7451$, where $\beta \in \mathbb{R}$ is theoretically derived as solution to an integral equation [21]. The β -approximation *quantile strategy* devised in [10] makes the aforementioned hardness tight.

IID Prophet Inequality with Random Horizon

The IID PI with Random Horizon (RH) was first introduced in [17]. Consider a random variable H , which we will call the (random) *horizon*, with given discrete distribution \mathcal{H} , that is, supported on an arbitrary subset of \mathbb{N} . This assumption comes with no loss of generality (formally, it rules out that the horizon has mass at zero). H will be assumed finite ($\mathbb{P}(H < \infty) = 1$) and integrable ($\mathbb{E}H < \infty$), which we denote $H \in \mathcal{L}^1(\Omega)$. RH is a relaxation of IID PI, considering integrable *iid* random variables X_1, \dots, X_H , with given distribution \mathcal{D} independent of \mathcal{H} . This setup, supported on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, models the game where the gambler, facing an unknown number of values X_1, \dots, X_H , can still use *a priori* knowledge of \mathcal{H} when maximising returns. More precisely, denote $[n] = \{1, \dots, n\}$. The goal is to maximise the gambler-to-prophet ratio of $\mathbb{E}X_\tau$ over $\mathbb{E}M_H$, where $M_H := \max_{i \in [H]} X_i$ and both expectations run over the randomness of the *iid* copies of $X \sim \mathcal{D}$ and $H \sim \mathcal{H}$. The gambler must select the value in ignorance of whether it is the last or not. If the gambler fails to stop by the time the last value is inspected, the return is zero. This ignorance, added to the usual nonanticipative constraint, yields a no easier model than IID PI. On top of the usual disadvantage of playing against a prophet, which knows all future realisations of the values, and chooses the largest; the gambler now competes against a prophet, which can also foresee the random number of values.

1.2 Previous bounds on IID Prophet Inequality with Random Horizon

No constant-approximations are possible for RH [2, Theorem 1.6] (we refer to this by saying that RH is hard). Distributional knowledge of the horizon is not enough to guarantee, in expectation, that the gambler's return can attain a worst-case constant multiple of the prophet's return. Therefore, we impose additional restrictions on the distribution of the horizon. The largest distributional class for which a constant-approximation has been found so far is that of horizons with increasing hazard rates $\lambda(h)$. For simplicity, we will use *increasing* and *decreasing* instead of *nondecreasing* and *nonincreasing* respectively. Recall that the hazard rate $\lambda(h)$ of a horizon H is defined as follows. Denote the survival function $S(h) = \mathbb{P}(H \geq h)$. If $|\text{supp}(H)| = \infty$, for every $h \in \mathbb{N}$, $\lambda(h) := \mathbb{P}(H = h)/S(h)$. If $|\text{supp}(H)| < \infty$, for every $h \in \mathbb{N}$, $\lambda(h)$ is defined analogously for all $h \leq \sup \text{supp}(H)$, while it is set to 1 for all $h > \sup \text{supp}(H)$.

► **Definition 1** (IHR and DHR class). H is Increasing Hazard Rate (IHR) if for every $h \in \mathbb{N}$, $\lambda(h) \leq \lambda(h+1)$. H is Decreasing Hazard Rate (DHR) if for every $h \geq \inf \text{supp}(H)$, $\lambda(h) \geq \lambda(h+1)$.

One says that IHR and DHR are *dual* classes, and the geometric distribution is the only discrete distribution that belongs to both of them.

The state-of-the-art for RH consists in the findings of [2]. In [2, Theorem 3.2], through the use of *second order stochastic dominance*, it is shown that if $H \in \text{IHR}$, $\mathbb{E}X_\tau \geq (2^{-1/\mu})^{-1} \mathbb{E}M_H$, where $\mu := \mathbb{E}H \geq 1$. This ensures a (uniform) 2-approximation on the IHR class. Studying monotone hazard rates stems from interpreting RH from a reliability theory point of view: an optimal stopping game under evolving risk (or ageing) that it ends by the next step, and, despite more than 15 years have passed, no progress has been made on classes larger than the IHR class.

1.3 Our contributions

In this paper we contribute to the study of RH by:

- Extending existing characterisations of the optimal algorithm to unbounded horizons.
- Complementing and extending key ideas from [2] to yield the existence of single-threshold 2-approximations on important superclasses of the IHR class, culminating in the \mathcal{G} class.
- Showing that the $\overline{\mathcal{G}}$ class (dual of the \mathcal{G} class) admits hard horizons for single-threshold algorithms.
- Deriving single-threshold constant-approximations for sufficiently concentrated horizons with finite second moments.
- Showing that RH admits families of horizons for which single-threshold algorithms are not competitive, despite the Secretary Problem (SP) optimal stopping rule providing a constant-approximation (see Section 1.5 for details on SP and its variants).

In this section we give formal statements of the above results, with the exception of Theorem 14 in Section 2, due to the overall degree of technicalities involved. Informally speaking, this theorem characterises the optimal algorithm in terms of a discounted infinite optimal stopping problem, with a special focus on unbounded horizons. As a technical tool, the equivalence of stopping rules for RH with stopping rules for the discounted problem is leveraged not only to obtain hardness results for single-threshold algorithms, but also to formalise heuristic arguments involving the optimal algorithm under geometric horizon. More specifically: we exploit the equivalence between stopping rules for RH and corresponding rules for the discounted problem to extend hardness results to neighbouring instances; we show that with a geometric horizon, if the value distribution is bounded, the optimal algorithm of RH is single-threshold, pinpointing the structural equation for the threshold.

The next three results involve stochastic orders, which we now introduce. Note that several standard facts regarding stochastic orders are often relied upon throughout this work. They will be recalled as remarks, whose proof can always be found in [39]. In the following definitions, H and G always refer to horizons.

► **Definition 2** (Probability Generating Function Order). *Given two horizons H, G , we say that G is dominated by H in the probability generating function (pgf) order, and denote it as $G \prec_{pgf} H$, if for every $t \in (0, 1)$, $\mathbb{E}(t^G) \geq \mathbb{E}(t^H)$.*

► **Remark 3.** $G \prec_{pgf} H$ if and only if for all $t \in (0, 1)$, $\sum_{i=1}^{\infty} S_G(i)t^i \leq \sum_{i=1}^{\infty} S_H(i)t^i$.

The \mathcal{G} class, consisting of every positive discrete distribution that dominates, in the pgf order, the (\mathbb{N} -valued) geometric distribution with the same mean, has been introduced in [24].

► **Definition 4** (\mathcal{G} class). *The \mathcal{G} class consists of every horizon that dominates, in the pgf order, the geometric distribution with the same mean: $\mathcal{G} := \{H : G \prec_{pgf} H, G \sim \text{Geom}(1/\mathbb{E}H)\}$.*

The dual class is obtained by reversing the pgf ordering, and is denoted as $\bar{\mathcal{G}}$. It is well known that $\text{IHR} \subset \mathcal{G}$, with inclusion being strict. Similarly, $\text{DHR} \subset \bar{\mathcal{G}}$. In applications of reliability theory, the classes within \mathcal{G} and its dual are a staple of modelling many aspects of ageing and more generally of an evolving risk of failure.

► **Theorem 5.** *For all $H \in \mathcal{G}$ with $\mathbb{E}H =: \mu$, a single-threshold $2 - 1/\mu$ -approximation exists.*

From [2, Theorem 3.5] it follows that any 2-approximation on the \mathcal{G} class is tight. More specifically, the 2-approximation yielding Theorem 5 is essentially the same as the one used on the IHR class in [2, Theorem 3.1]. The key feature of the stopping rule is to determine the value p such that $\mathbb{P}(X > p) = 1/\mu$, and then to accept the first value exceeding p . The \mathcal{G} class may be somewhat abstract, compared to its subclasses, which have immediate intuitive interpretations in terms of ageing concepts, such as the IHR class. Its possible interpretations are quite general and go beyond mere ageing, as can be read in [24]. To gain some intuition of how much more general a class it is, the reader is referred to Section 3, where we will describe some of its largest and most well-known subclasses, as evidence that Theorem 5 generalises significantly [2, Theorem 3.1, Theorem 3.2]. We conclude by mentioning that statistical tests for the \mathcal{L} class (the continuous equivalent of the \mathcal{G} class) have also been developed [23]. The possibility of adapting them to tests for the \mathcal{G} class adds to the practical significance of the result.

Our next result requires introducing the stopping rule for SP with deterministic horizon m . This consists of rejecting the first $r - 1$ values, and subsequently accepting the first value ranked better than all the previous ones. The waiting time $r_m \sim m/e$ as $m \rightarrow \infty$ yields the optimal stopping rule (see Section 1.5 for details). In the following, we provide the first example of a family of horizons, for which no single-threshold constant-approximation is possible, and yet an adaptation of the SP stopping rule to random horizons is proven to provide a constant-approximation. The family of horizons we will consider is of the form $\{H_m\}_{m \geq M}$, with pmf's parametrised by $\varepsilon > 0$, $M \in \mathbb{N}$ large enough and $m := \sup \text{supp}(H_m) < \infty$, and will therefore be denoted as $\mathcal{H}_M(\varepsilon)$. Let ζ_{r_m} be the (optimal) stopping rule for SP with deterministic horizon m . Taking the minimum $\tau_m := \zeta_{r_m} \wedge (H_m + 1)$ produces a natural adaptation of the SP stopping rule to the random horizon H_m .

► **Theorem 6.** *There exists a family of horizons $\mathcal{H}_M(\varepsilon)$, such that for all fixed ε small enough and arbitrary M , no single-threshold constant-approximation is possible. Nonetheless, for every fixed ε , no matter how small, the random horizon SP rule τ_m is an approximate constant-approximation, as long as $M = M(\varepsilon)$ is fixed large enough.*

For brevity, we say that the family of horizons $\mathcal{H}_M(\varepsilon)$ is *hard* for single-threshold algorithms, but the random horizon SP rule is *competitive* on it. The SP stopping rule only needs information regarding the relative ranks of the values, not the values themselves. Our result opens up a fresh line of investigation into new classes of distributions, to be handled via multiple threshold algorithms, which cannot be accessed via single-threshold algorithms.

Since we showed that single-threshold 2-approximations exist for the \mathcal{G} class, finding whether single-threshold constant-approximations exist for the $\overline{\mathcal{G}}$ class turns into an interesting problem. The equivalent characterisation of stopping rules for RH in terms of a discounted stopping problem, combined with a perturbation argument on $\mathcal{H}_M(\varepsilon)$, yields a negative answer.

► **Theorem 7.** *No single-threshold constant-approximation is possible on the $\overline{\mathcal{G}}$ class.*

In Section 4.2 we will describe more in detail some of the largest and most well-known subclasses of the $\overline{\mathcal{G}}$ class. They are no less important than the corresponding subclasses of the \mathcal{G} class. For example, discrete DHR distributions are crucial in several areas: they have been shown to describe well the number of seasons a show is run before it gets cancelled and the number of periods until failure of a device (a system, or a component, that functions until first failure) governed by a continuous DHR life distribution in the grouped data case [27]. Theorem 7 motivates our conjecture that single-threshold 2-approximations are not possible on the DHR class either. Since our characterisation result of stopping rules for RH extends to unbounded horizons, it will be a crucial tool in proving this conjecture within our framework, as DHR horizons have support of the form $[a, \infty) \cap \mathbb{N}$. Furthermore, suitably combining the techniques yielding Theorems 5 and 7 with the fact that the tight geometric instance for 2-approximations belongs to the \mathcal{G} class, motivates our second conjecture that the \mathcal{G} class is essentially optimal for 2-approximations on \mathcal{L}^1 horizons.

In our last result, we go back to single-threshold algorithms under horizons H satisfying only $\mathbb{E}(H^2) < \infty$, denoted $H \in \mathcal{L}^2(\Omega)$. We show that it is possible to exploit concentration bounds, in order to ensure that the same single-threshold algorithm used for the \mathcal{G} and $\overline{\mathcal{G}}$ class is still a 2-approximation. In the following result, W_0 denotes the principal branch of the Lambert function.

► **Theorem 8.** *For every H having both $\mu := \mathbb{E}H$ and $\sigma^2 := \text{Var } H$ finite, such that*

$$\sigma^2 \leq \mu^2 \left[1 - \frac{1}{\mu} + W_0 \left(- \left(2 - \frac{1}{\mu} \right) e^{-(2 - \frac{1}{\mu})} \right) \right], \quad (1)$$

a single-threshold $2 - 1/\mu$ -approximation exists.

The Lambert function is readily simulated and one can reliably compute numerically the magnitude of concentration required by (1), so that a 2-approximation is ensured. To exemplify, we consider the large-market limit, that is, $\mu \rightarrow \infty$. Under this computationally simpler hypothesis, taking the square root on both sides of (1) yields approximately $\text{CV} := \sigma/\mu \leq 0.770$, where CV is the *coefficient of variation* of the horizon. Variations on this procedure show that both weaker and stronger than 2 constant-approximations are possible for all suitably concentrated horizons. Consider again the upper bound on CV provided by square-rooting (1) in the large-market limit: $\text{CV} \leq \sqrt{1 + W_0(-2e^{-2})}$. Replacing 2 with $C > 2$ yields $\text{CV} \leq \sqrt{C - 1 + W_0(-Ce^{-C})}$. This will ensure a weaker C -approximation in the large-market limit. As $x \geq 1$ grows, $W_0(-xe^{-x})$ is negative strictly increasing (valued -1 at $x = 1$) and vanishing. Thus the upper bound on the CV ensuring a C -approximation for C growing large scales roughly as $\sqrt{C - 1}$, meaning that there will be, although progressively

weaker, constant-approximations, as long as $CV < \sqrt{C-1}$. In the other direction, stronger C -approximations are ensured when taking *admissible* $1 < C < 2$. When reducing C , we clearly cannot go past $e/(e-1)$, as the upper bound on the CV vanishes (approaching a deterministic horizon). At this point, our estimates hit the optimum competitive ratio for single-threshold algorithms on IID PI, $1 - 1/e$ [13]. This generalises as follows.

► **Corollary 9.** *For every H having both $\mu := \mathbb{E}H$ and $\sigma^2 := \text{Var} H$ finite, and every constant $C \geq [1 - (1 - 1/\mu)^\mu]^{-1}$, if $CV \leq \sqrt{C-1 + W_0(-Ce^{-C})}$, then a single-threshold C -approximation exists.*

The impact of our results concerns online auctions and, more generally, stochastic matching problems, where random horizons are considered and the connection to our model is well-established under perishable inventory assumptions (see Section 1.5).

1.4 Our techniques

The toolbox we assembled can be described in three key points.

Discounted infinite problem. The optimal algorithm is characterised via a correspondence between stopping rules for RH and stopping rules for a discounted optimal stopping problem, with the discount factors being the survival function of the horizon. More general optimal stopping problems (unbounded random horizon) do not always admit an optimal algorithm [37]. We show that for RH, this is always possible. This is crucial for geometric and DHR horizons, which are unbounded. The generalisation is the product of an original measure-theoretic construction, leveraging trace σ -algebras.

Stochastic orders. The extension of the 2-approximation from the IHR class to the \mathcal{G} class and \mathcal{L}^2 horizons is obtained by exploiting stochastic orders, which have never before been applied to prophet inequality problems. Several distributional classes, starting from the IHR class, can be equivalently defined through some stochastic-ordering with respect to the geometric distribution (as previously seen, the largest of such classes exploits the pgf order). Thus our Theorem 5 not only extends to the \mathcal{G} class, but simplifies the arguments of [2]. The proof of Theorem 8 leverages extremal properties of specific Bernoulli distributions with respect to another stochastic order, the Laplace transform order (see Section 5 for details), for distributions having finite variance.

Hard instance. Our analysis of the hard instance for single-threshold algorithms is also novel. The literature on prophet inequalities has focused only on bounded discrete instances that are hard for every algorithm. On the other hand, the instance we provide exploits a continuous Pareto distribution, which, combined with a suitable family of horizons, is both hard for single-threshold algorithms, and competitive for the SP stopping rule, against the offline maximum. This requires a more sophisticated and analytic approach. The phenomenon in which certain algorithms fail on an instance but others do not is subtle. To the best of our knowledge, a competitive analysis of the SP stopping rule against the offline maximum with a random horizon is also new (see Section 1.5). It is worth noting that this hard instance is designed to live in a close neighbourhood of the $\overline{\mathcal{G}}$ -class, so that perturbing it suitably yields, at the same time, a hard family in the $\overline{\mathcal{G}}$ class.

1.5 Additional related work and motivation

The literature on prophet inequalities models with deterministic horizon is vast (see [11, 12, 29, 19] for surveys). We only mention a few models closely related to PI and IID PI for context. In the Random Order PI values arrive in a uniform random order, it is 0.688-competitive [7] and 0.7235-hard [15]). In the Order Selection PI the order can be chosen by the gambler, it is 0.7258-competitive [6] and inherits $1/\beta$ -hardness from IID PI. Very recently, Oracle-augmented PI have also been introduced [18]. The existing literature strictly related to RH is limited to [2, 17, 30], which focuses on online automated mechanism design (see [38] for a survey), with emphasis on the multiple-items setting. In [2] a standard machinery has been put in place, which extends results for single-item RH to multiple items, provided the horizons for each item are uniformly bounded. This applies to our results, too. [30] considers the problem of unknown market size (that is, random horizon with unknown distribution, motivated by search keyword auctions such as Google). As previously mentioned, this assumption is hard [17]. In practice, distributional knowledge is available via the history of past transactions, from which RH arises as a natural model. On a technical level, RH is also well recognised as connected to other online stochastic matching problems with a random number of queries, which have recently been shown to admit 2-approximations [3]. A potential ground for applications of our techniques would be improving such guarantees.

The Secretary Problem (SP) has been extensively studied (see [16, 14] for variations). The classical SP is an online problem where there are n secretaries for hire, who show up in a uniform random order at interviews. An employer, who wants to maximise the probability of hiring the best solely based on the relative ranks up to the present, makes the decision of hiring or rejecting and keep interviewing, at each interview. The optimal stopping rule rejects all the first $r - 1$ of the n secretaries, and accepts the first secretary with relative rank 1 thereafter. The value $r = r_n$ is such that as $n \rightarrow \infty$, both r/n and the probability of selecting the best tend to $1/e \approx 0.3679$ [28].

When the number of secretaries is random, denoted by N , the secretaries are interviewed in a uniform random order, conditionally on N . Depending on the horizon's distribution, we can have a more complicated set Γ composed of multiple islands, in which it is optimal to stop at a secretary with relative rank 1, and outside which it is not [33]. The model has been shown to be hard [1].

SP has been studied also in its *full-information* variant, with and without a random horizon. The common continuous distribution, from which the quality measurement of each of the n secretaries is sampled independently, is known to the employer, who directly inspects, at each interview, the secretary's true quality [16, 4, 32]. Both no- and full-information SP have also been studied under the harder hypothesis of *random freeze* [36, 37]. More recent developments extend SP to combinatorial structures and unknown horizons [20, 31].

1.6 Organisation of the paper

In Section 2 we give preliminaries on the formal approach to random horizon optimal stopping problems, and prove the equivalence with discounted optimal stopping problem. In Section 3 we prove Theorem 5. In Section 4 we prove Theorem 6. In Section 4.2 we prove Theorem 7. In Section 5 we prove Theorem 8.

2 Preliminaries

In this section we set up our notation, give the details of the probabilistic model, and characterise the optimal algorithm.

2.1 Notation

We denote $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$, $\overline{\mathbb{N}} := \mathbb{N} \cup \{\infty\}$, $\overline{\mathbb{R}} := \mathbb{R} \cup \{\infty, -\infty\}$, $[n] := \{1, 2, \dots, n\} \subset \mathbb{N}$, $[0] = \{0\}$, $[n]_0 := \{0\} \cup [n]$, $a \wedge b := \min(a, b)$ and $a \vee b := \max(a, b)$ for all $a, b \in \mathbb{R}$. We follow standard asymptotic notation for nonnegative functions (and sequences, with trivial extensions to signed functions when the use of absolute value is consistent), here we point out the only slightly less common $f(x) \asymp g(x)$ as $x \rightarrow a \in \overline{\mathbb{R}}$ if $f(x) = \mathcal{O}(g(x))$ and $f(x) = \Omega(g(x))$. We will adopt the convention of omitting the dependence on all parameters not involved in the main limiting process (which will be clear from the context), regardless of whether uniformity (when not relevant to the context) holds within the parametric range. *Absolutely continuous* distributions will be referred to as *continuous*. *Discontinuous* distributions are assumed to satisfy the usual regularity assumption (isolated jumps). The convergence in distribution of random variables is denoted as weak convergence: \xrightarrow{w} . *Almost surely* is abbreviated a. s. and $H \sim G$ denotes that H, G are identically distributed random variables. The essential supremum of a collection of random variables is denoted by ess sup . $S(h) := \mathbb{P}(H \geq h)$ is called *survival function* of H ($S_H(h)$ if necessary). We follow the standard probabilistic functional notation for powers, such as $\mathbb{P}^2(\cdot) := [\mathbb{P}(\cdot)]^2$ and $\mathbb{E}^2(\cdot) := [\mathbb{E}(\cdot)]^2$.

2.2 Probabilistic model

As aforementioned in this section (and this section alone), RH is studied from the point of view of optimal stopping, hence horizons admitting mass at zero will also be considered. In Section 1.1 we ruled out any mass at zero, since in all subsequent sections, concerned with competitive analysis, this will be the standard assumption. A formal convention, in fact, restores the equivalence with a scenario admitting mass at zero, from a competitive analysis point of view. Add a value $X_0 := 0$, and set $X_\tau := X_0 =: M_{[0]}$ for all $\omega \in \{H = 0\}$. Informally, when the game does not start, both the gambler and the prophet receive no return. Adopt the convention $0/0 = 1$. Having set $X_0 = 0$ ensures that the game always starts, unless $\omega \in \{H = 0\}$. Assuming absence of penalty for entering the game even though it does not start yields the equivalence: the ratio being 1 does not affect the worst case competitive ratio.

In RH the gambler must select the value in ignorance of whether it is the last or not: at each step $i \in \mathbb{N}_0$ the gambler only learns whether the inspected value x_i is the last or not (that is, whether $\omega \in \{H = i\}$ or $\omega \in \{H > i\}$) *after* the decision of accepting or rejecting it has been made (that is, in the $i + 1$ st step). Probabilistically, it is standard to model this as follows. The random list X_1, \dots, X_H is constructed from the infinite underlying process $\mathbf{X} := \{X_i\}$ of *iid* copies of X , whose natural filtration (intuitively, the information associated with the history of the process up to the present) is the sequence of $\sigma(X_1, \dots, X_i) := \mathcal{B}_i$, for every $i \in \mathbb{N}_0$, where $\mathcal{B}_0 := \{\emptyset, \Omega\}$ (no information at the start: this is the implicit starting point of all filtration we will introduce). Let $\nu := \mathbb{E}X < \infty$. The nonanticipative condition and the ignorance aforementioned are modelled by requiring that $\{\tau = i\} \in \sigma(\mathbb{1}_{\{H=0\}}, X_1, \dots, \mathbb{1}_{\{H=i-1\}}, X_i) =: \mathcal{F}_i$ for every $i \in \mathbb{N}_0$. Intuitively, this second filtration represents the information associated with the history of the game, which combines both the history of the process of the values and the state of the game in all preceding

steps (that is, whether it ended or is still running). We recover the original X_1, \dots, X_H by imposing that the reward is zero if the gambler fails to stop by time H on the infinite sequence, that is, defining the reward sequence $\{Y_i\}$, where $Y_i := X_i \mathbb{1}_{\{H \geq i\}}$, to which we add $Y_0 := X_0$.

For technical reasons, we consider the slightly less common class of all possibly infinite stopping times for the problem $\mathbf{X}(\mathbf{Y})$, denoted as $\mathcal{T}(\mathcal{T}^*)$. Formally, $\tau \in \mathcal{T}(\mathcal{T}^*)$, means that $\{\tau = i\} \in \mathcal{B}_i(\mathcal{F}_i)$ for all $i \in \mathbb{N}$, also admitting $\mathbb{P}(\tau = \infty) > 0$. This requires that we *compactify* the problem, by prescribing a conventional, yet natural, reward value, in the event that the gambler never stops. We set $Y_\infty := \limsup_{i \rightarrow \infty} Y_i = 0$, since $\mathbb{P}(H < \infty) = 1$, so if the gambler never stops, a. s. the horizon will be reached, so there can be no reward. In this way infinite stopping rules do not increase the largest possible expected reward, and the equivalence with the original formulation is maintained. In conclusion, RH has been reformulated as an infinite optimal stopping problem with reward sequence $\mathbf{Y} := \{Y_i\}_{i \in \bar{\mathbb{N}}_0}$, underlying process \mathbf{X} and filtration $\mathcal{F} := \{\mathcal{F}_i\}_{i \in \bar{\mathbb{N}}_0}$, where $\mathcal{F}_\infty := \sigma(\bigcup_{i \in \mathbb{N}_0} \mathcal{F}_i)$ (the complete information about the history of the process, which is the ending point of all infinite filtration we will introduce), with respect to the class \mathcal{T}^* . We seek to find the value of the problem $\mathcal{V}(\mathbf{Y}) := \sup_{\tau \in \mathcal{T}^*} \mathbb{E}Y_\tau$ and, if it exists, an optimal stopping rule $\bar{\tau} \in \mathcal{T}^*$. *Optimal* means that $\mathbb{E}Y_{\bar{\tau}} = \mathcal{V}(\mathbf{Y})$. The reference to the horizon H can be made explicit, whenever the need arises, via the notation $\mathbf{Y}^{(H)} := \{Y_i^{(H)}\}_{i \in \bar{\mathbb{N}}_0}$.

2.3 Optimal algorithm

Under suitable hypotheses, an optimal stopping problem with independent horizon H , having survival function $S(i)$, is equivalent to a finite (if H is bounded) or infinite (if H is unbounded) optimal stopping problem *discounted* by factors $S(i)$. The discounted problem is $\mathbf{Z} := \{Z_i\}_{i \in \bar{\mathbb{N}}_0}$, where $Z_0 := Y_0$, $Z_i := S(i)X_i$ for every $i \in \mathbb{N}$, $Z_\infty := \limsup_{i \rightarrow \infty} Z_i = 0 = Y_\infty$. Thus, for \mathbf{Z} we use the filtration $\mathcal{B} := \{\mathcal{B}_i\}$, having set \mathcal{B}_0 and \mathcal{B}_∞ as usual. The latter is used when the problem is infinite. Thus $\mathcal{V}(\mathbf{Z}) := \sup_{\zeta \in \mathcal{T}} \mathbb{E}Z_\zeta$, and the equivalence aforementioned means that $\mathcal{V}(\mathbf{Y}) = \mathcal{V}(\mathbf{Z})$.

The bounded horizon case is solved in [37] by *backward induction*, which does not extend to the unbounded case where an optimal stopping rule is needed explicitly. The main contribution of our Theorem 14 is showing that RH with unbounded horizon H admits an explicit optimal stopping rule of the form $\bar{\tau} := \bar{\zeta} \wedge (H + 1)$, where we will take $\bar{\zeta}$ to be the *Snell stopping rule* for \mathbf{Z} . Let us recall a few facts, adapted from [8, 9] to our compactified framework. Let \mathcal{T}_i be the class of all stopping rules in \mathcal{T} that do not stop before step i .

► **Definition 10** (Snell envelope). *The Snell envelope of \mathbf{Z} is $\{V_i\}$, where $V_i := \text{ess sup}_{\zeta \in \mathcal{T}_i} \mathbb{E}_{\mathcal{B}_i} Z_\zeta$.*

Note that $\mathcal{V}(\mathbf{Z}) = V_0$. The intuitive interpretation of the stochastic process $\mathbf{V} = \{V_n\}$ is the best expected return available at step i among the rules that have reached step i .

► **Remark 11.** Since $\mathbb{E} \sup_{i \in \mathbb{N}} Z_i < \infty$ for every $i \in \mathbb{N}_0$ the Snell envelope satisfies

$$V_i = Z_i \vee \mathbb{E}_{\mathcal{B}_i} V_{i+1}. \tag{2}$$

Informally, (2) represents an extension of the dynamic programming principle for the infinite process $\{V_i\}$: at time i it is optimal to stop whenever the currently inspected value Z_i is greater than the highest expected future return as per the previous informal characterisation of V_i . We use (2) to compute V_0 . Note that for finite problems it coincides with backward induction.

► **Definition 12** (Snell stopping rule). *The Snell rule is $\bar{\zeta} := \inf\{i \in \mathbb{N}_0 : Z_i \geq \mathbb{E}_{\mathcal{B}_i} V_{i+1}\}$.*

87:10 IID Prophet Inequality with Random Horizon

► **Remark 13.** If $\mathbb{E} \sup_{i \in \mathbb{N}} Z_i < \infty$ and $Z_\infty := \limsup_{i \rightarrow \infty} Z_i \in \mathbb{R}$, then the Snell rule $\bar{\zeta}$ is optimal.

► **Theorem 14.** Let H be a finite horizon with $\mu := \mathbb{E}H$ and possibly mass at zero, and let $m := \sup \text{supp}(H)$, \mathbf{Y} and \mathbf{Z} as previously defined.

It holds that: $\mathcal{V}(\mathbf{Y}) = \mathcal{V}(\mathbf{Z})$; there exists $\bar{\zeta} \in \mathcal{T}$ such that $\mathbb{E}Z_{\bar{\zeta}} = \mathcal{V}(\mathbf{Z})$; $\bar{\tau} := \bar{\zeta} \wedge (H + 1)$ is such that $\mathbb{E}Y_{\bar{\tau}} = \mathcal{V}(\mathbf{Y})$.

Furthermore, $\bar{\zeta}$ is characterised as: the backward induction stopping rule for $\mathbf{Z}^{(m)}$ if $m < \infty$; the Snell rule for \mathbf{Z} , if $m = \infty$.

Idea of the proof. The case $m < \infty$ leads to the equivalent characterisation as backward induction on $\{Z_i\}_{i \in [m]_0}$ [37, Theorem 2.1]. For the case $m = \infty$, a crucial step is to show that for every stopping rule $\sigma \in \mathcal{T}^*$, we can construct a stopping rule $\zeta \in \mathcal{T}$, such that $\mathbb{E}Y_\sigma = \mathbb{E}Y_\tau$ (that is σ and τ are equivalent), where $\tau := \zeta \wedge (H + 1) \in \mathcal{T}^*$. This requires an inductive measure-theoretic construction. For every $i \in \mathbb{N}$, let $E_i := \{\sigma = i\} \cap \{H \geq i\}$. We have that $E_i \in \mathcal{F}_i \cap \{H \geq i\} = \mathcal{B}_i \cap \{H \geq i\}$, where the intersections denote, as standard, the corresponding trace σ -algebras. Next, we construct inductively mutually disjoint events $A_i \in \mathcal{B}_i$, such that $E_i = A_i \cap \{H \geq i\}$. These events $\{A_i\}$ are then relied upon to define ζ as, informally speaking, a stopping rule that stops when σ successfully stops by the time the horizon has realised, and does not stop otherwise. Formally, denote $C := (\bigcup_{i \in \mathbb{N}} A_i)^c \in \mathcal{B}_\infty$, and define for every $i \in \mathbb{N}$,

$$\zeta(\omega) := i \quad \text{if } \omega \in A_i, \quad (3)$$

while setting ζ to ∞ otherwise. Following the standard convention $\mathbb{E}Y_\sigma := \mathbb{E} \sum_{i \in \mathbb{N}} Y_i \mathbb{1}_{\{\sigma=i\}} + \mathbb{E}Y_\infty$, we verify by direct computation that $\mathbb{E}Y_\sigma = \mathbb{E}Y_\tau$ is ensured by $\mathbb{E}Y_\infty = 0$.

We show that $\mathcal{V}(\mathbf{Y}) = \mathcal{V}(\mathbf{Z})$ as follows. First, note that $Y_\infty = Z_\infty = 0$. Second, for every $\sigma \in \mathcal{T}^*$, $\mathbb{E}Y_\sigma = \mathbb{E}Z_\zeta$, where ζ is defined as in (3). Third, for every $\zeta \in \mathcal{T}$,

$$\mathbb{E}Y_\zeta = \mathbb{E}Y_\tau, \quad (4)$$

with $\tau := \zeta \wedge (H + 1)$. These facts imply that $\sup_{\sigma \in \mathcal{T}^*} \mathbb{E}Y_\sigma = \sup_{\zeta \in \mathcal{T}} \mathbb{E}Z_\zeta$.

Given an optimal stopping rule $\bar{\zeta} \in \mathcal{T}$, we have that $\bar{\tau} := \bar{\zeta} \wedge (H + 1) \in \mathcal{T}^*$ provides an optimal rule based on (4). Furthermore, our assumptions ensure that $\mathbb{E} \sup_{i \in \mathbb{N}} Z_i < \infty$ and $\limsup_{i \rightarrow \infty} Z_i = 0$. Hence, the Snell envelope \mathbf{V} of \mathbf{Z} satisfies (2) by Remark 11, and yields the optimal stopping. ◀

We also derive a result, which gives the intuitive interpretation of $\mathbb{E}_{\mathcal{B}_i} V_{i+1}$ as the best expected return available at step i among the rules that will not stop by or at step i .

► **Lemma 15.** Let \mathbf{V} be the Snell envelope of \mathbf{Z} . Then $\mathbb{E}_{\mathcal{B}_i} V_{i+1} = \text{ess sup}_{\zeta \in \mathcal{T}_{i+1}} \mathbb{E}_{\mathcal{B}_i} Z_\zeta$ a. s.

3 Warm-up: a 2-approximation on the \mathcal{G} class

We build familiarity with the model by proving Theorem 5, which derives a 2-approximation on the \mathcal{G} class, extending the results of [2, Section 3]. Recall that $X \sim \mathcal{D}$ is the distribution of the values and $H \sim \mathcal{H}$ of the horizon (with no mass at zero and $\mu := \mathbb{E}H$). The expected return of a single-threshold algorithm is readily computed.

► **Lemma 16.** Let $\pi > 0$ and consider the stopping rule $\tau_\pi := \inf\{i \in \mathbb{N} : Y_i \geq \pi\} \in \mathcal{T}^*$, where $Y_i := X_i \mathbb{1}_{\{H \geq i\}}$. Then $\mathbb{E}X_{\tau_\pi} = c_\pi \mathbb{E}(X|X \geq \pi)$, where $c_\pi = c_\pi(\mathcal{H}, \mathcal{D}) := 1 - \mathbb{E}[\mathbb{P}^H(X < \pi)]$.

The quantity c_π yields the competitive ratio. With this in mind, we upper-bound $\mathbb{E}M_H$ by *ex-ante* relaxation of the prophet, combined with a simple variational argument.

► **Lemma 17.** *Let X be continuous and $p > 0$ such that $\mathbb{P}(X \geq p) = 1/\mu$. Then we have that $\mathbb{E}M_H \leq \mathbb{E}(X|X \geq p)$.*

Idea of the proof. Let $f(x)$ be the probability density function (pdf) of M_H and $v(x)$ the pdf of X . In order to upper-bound $\mathbb{E}M_H = \int_0^\infty xf(x)dx =: I(f)$, we solve the variational problem of maximising the functional $I(f)$ with constraints: $\int_0^\infty f(x)dx = 1$; $0 \leq f(x) \leq \mu v(x)$ for all $x \in [0, \infty)$; f is nonnegative Lebesgue integrable on $[0, \infty)$, the constraint follows from a union bound. The maximal solution is $\tilde{f}(x) = \mu v(x)$ for all $x \geq p$, and zero otherwise, with p as in the claim. Then $I(\tilde{f}) = \mathbb{E}(X|X \geq p)$. ◀

To prove Theorem 5 we extend Lemma 17 to discontinuous X via *stochastic tie-breaking*. The technique is standard with deterministic horizons, but with random horizons, it is not yet clear that the assumptions imposed on RH are robust enough. We will show that an underlying property of the prophet, Uniform Integrability (UI), is the key. Before that, we derive the equivalent of Lemma 16 for a randomised algorithm (randomisation does not increase the optimal value of optimal stopping problems). Attach to every X_i a biased coin flip, that is *iid* Bernoulli random variables $B_i \sim B \sim \text{Ber}(q)$ independent of X and H , q denoting the probability of 1 (heads). Denote the class of adapted, thus randomised, stopping rules as $\overline{\mathcal{T}}$.

► **Lemma 18.** *Let $\pi > 0$ and consider the stopping rule $\tau_{\pi,q} := \inf\{i \in \mathbb{N} : Y_i \geq \pi\} \in \overline{\mathcal{T}}$. Then $\mathbb{E}X_{\tau_{\pi,q}} = c_{\pi,q}\mathbb{E}(X|X \geq \pi)$, where $c_{\pi,q} = c_{\pi,q}(\mathcal{H}, \mathcal{D}) := 1 - \mathbb{E}\{[1 - q\mathbb{P}(X \geq \pi)]^H\}$.*

Rewriting $c_\pi = 1 - \mathbb{E}\{[1 - \mathbb{P}(X \geq \pi)]^H\}$ in Lemma 16 helps to see the statement of Lemma 18 as its natural generalisation, although combinatorially tedious to prove.

Idea of the proof of Theorem 5. Let X be continuous. By Lemmas 16 and 17,

$$\mathbb{E}X_{\tau_p} \geq c_p \mathbb{E}M_H,$$

where $c_p := 1 - \mathbb{E}[\mathbb{P}^H(X < p)]$. For any $H \in \mathcal{G}$, $G \prec_{pgf} H$ with $G \sim \text{Geom}(1/\mu)$, $\mu = \mathbb{E}H$. Let $t := \mathbb{P}(X < p) = 1 - 1/\mu$, so that $c_p = 1 - \mathbb{E}t^H$, then

$$\mathbb{E}t^H \leq \mathbb{E}t^G = \frac{1}{\mu} \sum_{h \in \mathbb{N}} t^h (1 - 1/\mu)^{h-1} = \frac{1 - 1/\mu}{2 - 1/\mu}.$$

This implies that $c_p \geq (2 - 1/\mu)^{-1}$ and the result follows.

Let X be discontinuous with cumulative distribution function (cdf) $V(x)$ having $1 - 1/\mu$ in correspondence of a jump at p , say $\lim_{\varepsilon \rightarrow 0^+} V(p - \varepsilon) < 1 - 1/\mu < V(p)$. Let $D := \{j_k\}$ be the set of discontinuities (increasingly enumerated), one of which is p . Let $\{\varepsilon(l)\}$ be a small enough positive monotonically vanishing sequence. Linearly interpolate $V(x)$ on disjoint intervals of length $\varepsilon(l)$ and keep it unaltered otherwise. This yields an approximating sequence of continuous and piecewise differentiable cdf's $V_l(x)$ of continuous random variables $X^{(l)}$. Upon showing that the family $\{X^{(l)}\}$ is UI, define $M_H^{(l)} := \max\{X_1^{(l)}, \dots, X_H^{(l)}\}$. RH assumptions ensure that $\{M_H^{(l)}\}$ is UI too. Since $X^{(l)} \xrightarrow{w} X$, $M_H^{(l)} \xrightarrow{w} M_H$. By Lemma 17, for every l there is p_l such that $\mathbb{P}(X^{(l)} \geq p_l) = 1/\mu$ and $\mathbb{E}M_H^{(l)} \leq \mathbb{E}(X^{(l)}|X^{(l)} \geq p_l)$. By construction, $p_l \rightarrow p$ as $l \rightarrow \infty$, so $\mathbb{E}(X^{(l)}|X^{(l)} \geq p_l) \rightarrow \mathbb{E}(X|X \geq p)$. All these facts ensure that

$$\mathbb{E}M_H \leq \mathbb{E}(X|X \geq p).$$

87:12 IID Prophet Inequality with Random Horizon

To conclude, for any $H \in \mathcal{G}$, since $G \prec_{pgf} H$, by Lemma 18 the single-threshold algorithm $\tau_{p,\bar{q}} \in \overline{\mathcal{T}}$, with $\bar{q} := [\mu\mathbb{P}(X \geq p)]^{-1}$, is such that

$$\mathbb{E}X_{\tau_{p,\bar{q}}} \geq c_{p,\bar{q}}\mathbb{E}M_H,$$

where $c_{p,\bar{q}} := 1 - \mathbb{E} \left\{ [1 - \bar{q}\mathbb{P}(X \geq p)]^H \right\} = 1 - \mathbb{E} \left[(1 - 1/\mu)^H \right] \geq (2 - 1/\mu)^{-1}$. \blacktriangleleft

► Remark 19. If X is discontinuous with $1/\mu$ in correspondence of a discontinuity jump of the survival function of X , the $2 - 1/\mu$ -approximation is a single-threshold randomised algorithm, where the threshold is the value p at which the discontinuity jump occurs, and the randomisation parameter is $[\mu\mathbb{P}(X \geq p)]^{-1}$, but the argument above shows that if X is discontinuous and $1/\mu$ does not occur in correspondence of a discontinuity jump, the randomisation parameter is 1, meaning that the nonrandomised algorithm for the continuous case is enough to yield a $2 - 1/\mu$ -approximation.

In the rest of the paper we will omit any explicit adaptation of similar tie-breaking arguments, as we showed that RH is robust enough. Some of the most well-known subclasses of the \mathcal{G} class, which earned considerable interest in applications, are: IHRE (Increasing Hazard Rate in Expectation); HIHRE (Harmonically Increasing Hazard Rate in Expectation); NBU (New Better than Used); NBUE (New Better than Used in Expectation). The following tower of inclusions (and its dual, which will follow in the next section) is well-established [5, 34, 22] and known to be strict: $\text{IHR} \subset \text{IHRE} \subset \text{HIHRE} \subset \text{NBU} \subset \text{NBUE} \subset \text{HNBUE} \subset \mathcal{G}$.

The tight instance for the $2 - 1/\mu$ -approximation involves a geometric horizon and was studied in [2, Theorem 3.5]. Their basic assumptions, concerning a two-point distribution, are intuitive facts: that an optimal algorithm with geometric horizon should exist and that it has a single threshold. We show mathematically that they both hold, with the second extending, more in general, to any bounded value distribution. As a bonus, we obtain an equation for the optimal threshold. Recall that $\text{Geom}(1 - q)$ denotes the geometric distribution with failure probability $0 < q < 1$.

► Lemma 20. *Let $H \sim \text{Geom}(1 - q)$, and X such that $\text{supp}(X) = [a, b]$, $0 \leq a \leq b$. Then there exists a threshold V_0 such that the corresponding single-threshold algorithm is optimal.*

Idea of the proof. By Theorem 14, $\mathcal{V}(\mathbf{X}) = \mathcal{V}(\mathbf{Y}) = \mathcal{V}(\mathbf{Z}) =: V_0$. By (2), it holds that $V_0 = \mathbb{E}V_1 = \mathbb{E}(Z_1 \vee \mathbb{E}_{\mathcal{B}_1} V_2)$. Since $Z_i := q^{i-1}X_i$, by time invariance and Lemma 15 it follows that $\mathbb{E}_{\mathcal{B}_1} V_2 = q\mathbb{E}V_1 = qV_0$. Therefore, $V_0 = \mathbb{E}(X \vee qV_0)$; equivalently

$$f(V_0) := V_0 + \int_{qV_0}^b F(x)dx = b.$$

The function $f(V)$ is strictly increasing and differentiable on $[0, b]$, with $f(0) \leq b \leq f(b)$. Hence the value V_0 exists. By Remark 13 and Definition 12 and time invariance again, it is optimal to stop at the first value greater than or equal to V_0 (the expected return). \blacktriangleleft

4 A first step beyond single-threshold algorithms

In this section we construct a parametric family of horizons, which is hard for any single-threshold algorithm, and yet it allows for an adaptive competitive algorithm. We also show that it can be perturbed so that it enters the $\overline{\mathcal{G}}$ class, while remaining hard for single-threshold algorithms.

4.1 The hard instance

Let us start with the hard instance X for the values, which is a Pareto (Type I) distribution with scale parameter 1 and shape parameter $1 + \varepsilon$, where $\varepsilon > 0$ is fixed small enough. The corresponding cdf, denoted as $V(x)$, vanishes on $(-\infty, 1)$, whereas on $[1, \infty)$

$$V(x) = 1 - x^{-(1+\varepsilon)}. \quad (5)$$

Since M_n has cdf $F_n(x) = V^n(x)$, which is differentiable for all $x \neq 1$, we obtain that its density (defined almost everywhere) $f_n(x)$ vanishes on $(-\infty, 1)$, whereas $f_n(x) = nV^{n-1}(x)V'(x)$ on $(1, \infty)$. By direct computation, relying on the properties of the Gamma and Beta function, it follows that $\mathbb{E}M_n \sim \Gamma(1 - (1 + \varepsilon)^{-1}) n^{1/(1+\varepsilon)}$, as $n \rightarrow \infty$, uniformly in $\varepsilon > 0$. Consequently, in conjunction with the instance X for the values, we consider a family of horizons $\{H_m\}$, with $m \in \mathbb{N}$ large enough, such that for every fixed such m , $\text{supp}(H_m) = [\ell, m] \cap \mathbb{N}$, with the lower limit $1 < \ell \in \mathbb{N}$. For every m fixed, we define the pmf of each horizon by letting, for every $\ell \leq h \leq m$,

$$p_h^{(m)} := \mathbb{P}(H_m = h) = Z_m^{-1} h^{-\frac{1}{1+2\varepsilon}}, \quad (6)$$

having defined the normalising constant $Z_m := \sum_{h=\ell}^m h^{-\frac{1}{1+2\varepsilon}}$. Through standard integral estimates of summations, it is shown that

$$Z_m \sim [1 + (2\varepsilon)^{-1}] m^{\frac{2\varepsilon}{1+2\varepsilon}}, \quad (7)$$

and that the following holds.

► **Lemma 21.** *As $m \rightarrow \infty$, uniformly in $\varepsilon > 0$ small enough, $\mathbb{E}M_{H_m} \sim Nm^{\frac{1}{1+\varepsilon}}$, where $N = N(\varepsilon) := \Gamma\left(1 - \frac{1}{1+\varepsilon}\right) \left[1 + \frac{\varepsilon}{(1+\varepsilon)(1+2\varepsilon)}\right]^{-1} \left[1 + \frac{1}{2\varepsilon}\right]^{-1}$.*

By (7) and integral estimates akin to those leading to Lemma 21 it is straightforward to compute that for every $n \in \mathbb{N}$, as $m \rightarrow \infty$,

$$\mathbb{E}H_m^n \sim \frac{\sum_{h=\ell}^m h^{n+1-\frac{1}{1+2\varepsilon}}}{[1 + (2\varepsilon)^{-1}] m^{\frac{2\varepsilon}{1+2\varepsilon}}} \sim \frac{m^{n-\frac{1}{1+2\varepsilon}}}{(n+1 - (1+2\varepsilon)^{-1}) [1 + (2\varepsilon)^{-1}] m^{\frac{2\varepsilon}{1+2\varepsilon}}} = \frac{2\varepsilon m^n}{n + 2(n+1)\varepsilon}. \quad (8)$$

The properties of the Gamma function ensure that as $x \rightarrow 0^+$, $\Gamma(x) \sim x^{-1}$. As a result, as ε vanishes, $N \sim 2\varepsilon(1 - (1 + \varepsilon)^{-1})^{-1} \rightarrow 2$ and we have the following.

► **Remark 22.** *As $m \rightarrow \infty$, uniformly in $\varepsilon > 0$ small enough, $\mathbb{E}M_{H_m} \asymp m^{\frac{1}{1+\varepsilon}}$.*

We are now ready to prove the hardness of the family of horizons $\{H_m\}$ for single-threshold algorithms. A stopping rule with a single threshold π is denoted as τ_π . A crucial point is that the algorithm facing H_m will set a threshold π_m exploiting all distributional knowledge regarding X and H_m . Yet, this is not enough to obtain a constant-approximation on the family $(X, \{H_m\})$ for all suitably small ε .

► **Proposition 23.** *For fixed $\varepsilon > 0$ small enough, given the instance of copies of X and the family of horizons $\{H_m\}$ of (5) and (6) respectively, we have that for every sequence of thresholds $\{\pi_m\}$ $\liminf_{m \rightarrow \infty} \frac{\mathbb{E}X_{\tau_{\pi_m}}}{\mathbb{E}M_{H_m}} = 0$.*

Idea of the proof. Our strategy will be to characterise the sequence of optimal single thresholds $\{\bar{\pi}_m\}$ (corresponding to horizons $\{H_m\}$ for m large enough and ε small enough), where $\bar{\pi}_m$ maximizes the value obtained $\mathbb{E}X_{\tau_{\bar{\pi}_m}} := \mathbb{E}Y_{\tau_{\bar{\pi}_m}}^{(H_m)}$ over all single threshold strategies

87:14 IID Prophet Inequality with Random Horizon

π_m . Denote with r_m the gambler-to-prophet ratio of $\mathbb{E}Y_{\tau_{\pi_m}^{(H_m)}}$ and $\mathbb{E}M_{H_m}$. As a consequence of the characterisation obtained, it will follow that there is always a subsequence $\{m_j\}$, such that as $j \rightarrow \infty$, r_{m_j} vanishes. This yields the claim, since r_m is an upper bound on any gambler-to-prophet ratio for arbitrary single-threshold algorithms.

Since X is continuous, by Lemma 16 we compute $\mathbb{E}Y_{\tau_\pi^{(H_m)}}$, which is nontrivial only if $\pi > 1$. This is expressed as the product of

$$\mathbb{E}(X|X \geq \pi) = (1 + 1/\varepsilon)\pi \quad \text{and} \quad c_\pi = 1 - \mathbb{E}[V^{H_m}(\pi)] = 1 - \mathbb{E}\left[(1 - 1/\pi^{1+\varepsilon})^{H_m}\right].$$

Next, we show analytically that the optimal sequence of thresholds $\{\bar{\pi}_m\}$ satisfies the *optimality equation*

$$g_m(\bar{\pi}_m) := \mathbb{E}\left[(1 - 1/\bar{\pi}_m^{1+\varepsilon})^{H_m}\right] + \frac{1 + \varepsilon}{\bar{\pi}_m^{1+\varepsilon}} \mathbb{E}\left[H_m (1 - 1/\bar{\pi}_m^{1+\varepsilon})^{H_m-1}\right] = 1.$$

To conclude, by these facts and Lemma 21 and Remark 22 we have that, as $m \rightarrow \infty$, uniformly in $\varepsilon > 0$ small enough,

$$r_m = c_{\bar{\pi}_m} \frac{\mathbb{E}(X|X \geq \bar{\pi}_m)}{\mathbb{E}M_{H_m}} \sim \left(1 + \frac{1}{\varepsilon}\right) \frac{c_{\bar{\pi}_m} \bar{\pi}_m}{Nm^{\frac{1}{1+\varepsilon}}} \asymp \frac{c_{\bar{\pi}_m} \bar{\pi}_m}{\varepsilon m^{\frac{1}{1+\varepsilon}}}.$$

Consider that $\{\bar{\pi}_m\}$ is either bounded or unbounded.

Bounded case. Since $c_{\bar{\pi}_m} < 1$, r_m vanishes as $m \rightarrow \infty$, for any fixed ε small enough. Note that this case covers also the trivial case for which there exists a subsequence $\{\bar{\pi}_{m_j}\}$ such that $\bar{\pi}_{m_j} \leq 1$, for which the numerator of r_{m_j} becomes upper-bounded by $\nu := \mathbb{E}X$.

Intermediate case. Before moving on to the general unbounded case, we consider the subcase of divergence to infinity with $m^{\frac{1}{1+\varepsilon}} = o(\bar{\pi}_m)$. Note that by Bernoulli's inequality we have that

$$c_{\bar{\pi}_m} := 1 - \mathbb{E}\left[(1 - 1/\bar{\pi}_m^{1+\varepsilon})^{H_m}\right] \leq \frac{\mathbb{E}H_m}{\bar{\pi}_m^{1+\varepsilon}}.$$

By (8) with $n = 1$ it holds that as $m \rightarrow \infty$, for every fixed ε small enough r_m vanishes, since

$$\frac{c_{\bar{\pi}_m} \bar{\pi}_m}{m^{\frac{1}{1+\varepsilon}}} \leq \frac{\mathbb{E}H_m}{m^{\frac{1}{1+\varepsilon}} \bar{\pi}_m^\varepsilon} \asymp \varepsilon \left(\frac{m^{\frac{1}{1+\varepsilon}}}{\bar{\pi}_m}\right)^\varepsilon \rightarrow 0.$$

Unbounded case. More in general, we can rely on the previous point to show that the unbounded case overall only allows for r_m to vanish. Consider that since $0 \leq r_m \leq 1$, if all of its convergent subsequences $\{r_{m_j}\}$ vanish for all ε small enough, then $\{r_m\}$ vanishes for all ε small enough. We show the sufficient condition by contradiction: assume that as $j \rightarrow \infty$, $r_{m_j} \rightarrow \rho = \rho(\varepsilon) \in (0, 1]$, with $\rho(\varepsilon)$ being bounded away from 0 no matter how small ε is taken. Then we have that as $j \rightarrow \infty$, for $\varepsilon > 0$ fixed small enough, $\mu_j := m_j/\bar{\pi}_{m_j}^{1+\varepsilon}$ is asymptotically equivalent to a bounded sequence, so it must be bounded too. Thus, we can consider a convergent subsequence of $\{\mu_j\}$, denoted $\{\mu_{j_k}\}$, such that as $k \rightarrow \infty$, $\mu_{j_k} \rightarrow \alpha = \alpha(\varepsilon) \geq 0$. A sharp asymptotic analysis of $c_{\bar{\pi}_{m_{j_k}}}$ reveals that the *optimality equation* satisfied by $\bar{\pi}_{m_{j_k}}$ allows only for α being bounded away from zero, unless it is identically zero. In what follows, it is useful to rewrite the *optimality equation* as

$$\mathbb{E}\left[\phi(\bar{\pi}_{m_{j_k}}, H_{m_{j_k}})\right] = 1, \quad \text{where} \quad \phi(\pi, H) := (1 - 1/\pi^{1+\varepsilon})^{H-1} [1 - 1/\pi^{1+\varepsilon} + (1 + \varepsilon)H/\pi^{1+\varepsilon}],$$

and denote $g_{m_{j_k}}(\bar{\pi}_{m_{j_k}}) := \mathbb{E}[\phi(\bar{\pi}_{m_{j_k}}, H_{m_{j_k}})]$. By estimating asymptotically the series expansion of $\phi(\bar{\pi}_{m_{j_k}}, H_{m_{j_k}})$ with respect to the first argument, uniformly in $\alpha > 0$, the following bound is shown to hold as $k \rightarrow \infty$:

$$g_{m_{j_k}}(\bar{\pi}_{m_{j_k}}) \leq 1 + \varepsilon [(2 + \alpha)e^{-\alpha} - 2] + \mathcal{O}(\varepsilon^2) + \mathcal{O}(1/m_{j_k}).$$

It follows that for $\varepsilon > 0$ small enough, as $k \rightarrow \infty$, $g_{m_{j_k}}(\bar{\pi}_{m_{j_k}}) < 1$ if the strictly decreasing function $h(\alpha) := (2 + \alpha)e^{-\alpha} - 2$ is bounded away from zero. Note that $h(\alpha)$ is negative and vanishes as $\alpha \geq 0$ vanishes. Since we have previously shown that if $\alpha > 0$ exists, it must be bounded away from zero as ε vanishes (and thus $h(\alpha)$ is bounded away from zero too), in this case the optimality equation cannot be satisfied, as long as ε is taken sufficiently small. Thus only $\alpha = 0$ would not yield a contradiction at this point, meaning that for all ε small enough, $\mu_{j_k} \rightarrow 0$ as $k \rightarrow \infty$. By boundedness, also $\mu_j \rightarrow 0$ as $j \rightarrow \infty$. This corresponds to the scenario in the previous point (replace, in that argument, m with m_j) and therefore r_{m_j} vanishes. This is in contradiction with the assumption, which ensures $r_{m_j} \rightarrow \rho > 0$, and therefore we must have $\rho = 0$.

The sequence $\{r_m\}$ is thus shown to always admit a vanishing subsequence. \blacktriangleleft

4.2 Hardness of the $\bar{\mathcal{G}}$ class

The most notable subclasses of the $\bar{\mathcal{G}}$ class are the dual of those introduced at the end of Section 3: $\text{DHR} \subset \text{DHRE} \subset \text{HDHRE} \subset \text{NWU} \subset \text{NWUE} \subset \text{HNWUE} \subset \bar{\mathcal{G}}$. A hint towards the hardness of the $\bar{\mathcal{G}}$ class could be that the duality mirrors the estimates of Section 3, meaning that the lower-bound, which provides the competitive ratio ensured, turns into an upper-bound, informally speaking. This mirroring has significance, because the ex-ante relaxation of the prophet is tight on geometric horizons, and is essentially maxed-out by the very definition of the \mathcal{G} class.

Idea of the proof of Theorem 7. We show that a perturbation of the hard instance from the previous section prevents any single-threshold constant-approximation on the $\bar{\mathcal{G}}$ class.

Step 1. We start by constructing a family $\tilde{\mathcal{H}}_M(\varepsilon) \subset \bar{\mathcal{G}}$, which is a sequence of horizons \tilde{H}_m arising as a perturbation of horizons H_m in the hard family $\mathcal{H}_M(\varepsilon)$ of (6), with fixed $\ell = 2$, as m grows, for all fixed $0 < \varepsilon < 1/4$ small enough. This is done by adding to the pmf of said H_m , a (to be suitably rescaled) mass at one, $\delta_m := Cm^{\frac{2\varepsilon}{1+2\varepsilon}}$, where $C = C(\varepsilon) := (3 + 10\varepsilon)/(2\varepsilon)$. We denote the perturbed horizon \tilde{H}_m for every m, ε considered, and observe that the corresponding normalizing constants for the pmf's \tilde{Z}_m and Z_m ; and expectations $\tilde{\mu}_m := \mathbb{E}\tilde{H}_m$ and $\mu_m := \mathbb{E}H_m$; satisfy, as m grows, the following:

$$\tilde{Z}_m \sim \tilde{C}Z_m, \quad \tilde{C} := C + 1 + \frac{1}{2\varepsilon} \tag{9}$$

$$\tilde{\mu}_m \sim \frac{1}{\tilde{C}} \left(C + \frac{1 + 2\varepsilon}{1 + 4\varepsilon} m \right). \tag{10}$$

Trivially, these follow from (7) and (8) with $n = 1$ and the definitions, respectively $\tilde{Z}_m := Z_m + \delta_m$ and $\tilde{\mu}_m = \delta_m/\tilde{Z}_m + \mu_m Z_m/\tilde{Z}_m$. To show that for all m large enough, $\tilde{H}_m \in \bar{\mathcal{G}}$, we have to establish that eventually (short for all m large enough from now on), for all $t \in (0, 1)$,

$$\frac{1}{\tilde{Z}_m} \left[\delta_m t + \sum_{h=2}^m t^h h^{-\frac{1}{1+2\varepsilon}} \right] \geq \frac{1}{\tilde{\mu}_m} \sum_{h=1}^{\infty} (1 - 1/\tilde{\mu}_m)^{h-1} t^h,$$

87:16 IID Prophet Inequality with Random Horizon

which we recast as

$$\sum_{h=2}^m t^h h^{-\frac{1}{1+2\varepsilon}} \geq -\delta_m t + \frac{\varepsilon_m t}{1-t(1-1/\tilde{\mu}_m)}, \quad (11)$$

having set $\varepsilon_m := \tilde{Z}_m/\tilde{\mu}_m$. Let $\eta_m := (1 - \varepsilon_m/\delta_m)/(1 - 1/\tilde{\mu}_m)$. We establish that eventually for all $t \in (0, \eta_m]$

$$-\delta_m t + \frac{\varepsilon_m t}{1-t(1-1/\tilde{\mu}_m)} \leq 0, \quad (12)$$

and that eventually for all $t \in [\eta_m, 1)$

$$\sum_{h=2}^m h(h-1)t^{h-2}h^{-\frac{1}{1+2\varepsilon}} < \frac{2\varepsilon_m(1-1/\tilde{\mu}_m)}{[1-t(1-1/\tilde{\mu}_m)]^3}. \quad (13)$$

This implies (11) for all $t \in (0, 1)$, since eventually on $(0, \eta_m]$ (11) holds with strict inequality directly by (12), while for all $t \in [\eta_m, 1)$ we have that the difference of the left-hand and right-hand side of (11), denoted $f_m(t)$, is strictly concave ((13) states precisely that $f_m''(t) < 0$). Since $f_m(t) > 0$ for all $t \in (0, \eta_m]$ and since it vanishes at both ends of the unit interval, it cannot vanish on $(\eta_m, 1)$, meaning that there cannot be any crossings of the two sides of (11) on $(\eta_m, 1)$, which therefore holds everywhere.

Step 2. Next, we show that with X being set to be the hard instance of values X of (5), with $\varepsilon > 0$ fixed small enough as in *Step 1*, $\mathbb{E}M_{\tilde{H}_m} = \Omega(\mathbb{E}M_{H_m})$ and $\mathbb{E}X_{\tau_{\pi_m}} := \mathbb{E}Y_{\tau_{\pi_m}}^{(\tilde{H}_m)} \leq \mathbb{E}Y_{\tau_{\pi_m}}^{(H_m)}$ as m grows, for every sequence of thresholds $\{\pi_m\}$. The first fact follows from Remark 22 and (9) and observing that the law of total expectation yields

$$\mathbb{E}M_{\tilde{H}_m} = \frac{\delta_m}{\tilde{Z}_m} \mathbb{E}X + \frac{Z_m}{\tilde{Z}_m} \mathbb{E}M_{H_m} > \frac{Z_m}{\tilde{Z}_m} \mathbb{E}M_{H_m} \sim \frac{1}{\tilde{C}} \mathbb{E}M_{H_m}.$$

The second fact follows from considering the survival functions of \tilde{H}_m and H_m , respectively $\tilde{S}_m(i)$ and $S_m(i)$. For all $i \geq 3$, $\tilde{S}_m(i) = S_m(i)Z_m/\tilde{Z}_m$, whereas $\tilde{S}_m(2) = S_m(2) - \delta_m/\tilde{Z}_m$. Consider next a single-threshold algorithm τ_π , where $\pi = \pi_m$ on horizon \tilde{H}_m . By Theorem 14 (note that the notations τ and σ used there are now swapped, as our focus here is on the original stopping rule) for any horizon \tilde{H}_m , there is a stopping rule equivalent to τ_π for the problem $\mathbf{Y}^{(\tilde{H}_m)}$, denoted as $\sigma_\pi := \zeta_\pi \wedge (\tilde{H}_m + 1)$, where ζ_π is, as per (3), an algorithm for the discounted problem $\mathbf{Z} = \{\tilde{S}_m(i)X_i\}$, which stops only when τ_π successfully stops by the time the horizon has realised. Intuitively, we can construct ζ_π as the stopping rule with single threshold π on the events $\{\tau_\pi = i\} \cap \{\tilde{H}_m \geq i\}$. Otherwise it stops at m on $\{\tau_\pi > \tilde{H}_m\}$. The proof of Theorem 14 shows that this can be done, so that ζ_π is adapted to \mathbf{X} , and as a result, as m grows,

$$\mathbb{E}Y_{\tau_\pi}^{(\tilde{H}_m)} = \mathbb{E} \sum_{i=1}^m \tilde{S}_m(i)X_i \mathbb{1}_{\{\zeta_\pi=i\}} \leq \mathbb{E} \sum_{i=1}^m S_m(i)X_i \mathbb{1}_{\{\zeta_\pi=i\}} = \mathbb{E}Y_{\tau_\pi}^{(H_m)}.$$

Step 3. Finally, suppose by contradiction that a constant-approximation exists on the $\bar{\mathcal{G}}$ class. Then there exist a sequence of thresholds $\{\pi_m\}$ and $c > 0$ such that τ_{π_m} is a $1/c$ -approximation on $\tilde{\mathcal{H}}_M(\varepsilon) \subset \bar{\mathcal{G}}$ by *Step 1*. By Proposition 23 for all $\varepsilon > 0$ fixed small enough, there exists a subsequence $\{\pi_{m_j}\}$ such that r_{m_j} , the gambler-to-prophet ratio of $\mathbb{E}Y_{\tau_{\pi_{m_j}}}^{(H_{m_j})}$ over $\mathbb{E}M_{H_{m_j}}$, vanishes. Let \tilde{r}_{m_j} be the gambler-to-prophet ratio of $\mathbb{E}Y_{\tau_{\pi_{m_j}}}^{(\tilde{H}_{m_j})}$ over $\mathbb{E}M_{\tilde{H}_{m_j}}$. Combining this with *Step 2*'s asymptotics yields the following contradiction as j grows: $c \leq \tilde{r}_{m_j} = \mathcal{O}(r_{m_j}) \rightarrow 0$. \blacktriangleleft

4.3 The competitive Secretary Problem rule

In this section we show that a straightforward adaptation of the SP stopping rule with deterministic horizon m is competitive on the horizons H_m defined in (6). Consider any instance of the values X which is a nonnegative integrable continuous random variable. Consider the process $\mathbf{X} := (X_1, \dots, X_m)$, where X_i are *iid* copies of X . We equivalently characterise it in the context of SP via a uniform random permutation $\pi = (\pi_1, \dots, \pi_m)$. Consider $\mathbf{X}^\pi := (X_{\pi_1}, \dots, X_{\pi_m})$. Since \mathbf{X} is exchangeable, $\mathbf{X}^\pi \sim \mathbf{X}$. Consider now the waiting time $r_m - 1$ for the SP stopping rule, denoted as ζ_{r_m} , and recall that $r_m \sim m/e$. Consider \mathbf{X}^π conditionally on \mathbf{X} , denoted as $(\mathbf{X}^\pi | \mathbf{X})$. This is the process of the realised values arriving in a uniform random order and reformulates the usual setup for SP. In this reformulation, the classical result of [28] can be restated as

$$\mathbb{P}\left(X_{\pi_{\zeta_{r_m}}} = M_m | \mathbf{X}\right) \geq \frac{1}{e}. \quad (14)$$

Given the stopping rule ζ_{r_m} , denote the event of winning (that is, choosing the best) as $W := \{X_{\pi_{\zeta_{r_m}}} = M_m\}$ and that of winning at time i as $W_i := \{i = \inf\{j \geq r_m : X_{\pi_j} = M_m\}\}$. Note that by partitioning W conditionally on \mathbf{X} , we have that

$$\mathbb{P}(W | \mathbf{X}) = \sum_{i=r_m}^m \mathbb{P}(W_i | \mathbf{X}). \quad (15)$$

The proof of (14) relies on the following fact, which we will also exploit:

$$\mathbb{P}(i = \inf\{j \geq r_m : X_{\pi_j} = M_m\} | \mathbf{X}) = \frac{r_m - 1}{m} \frac{1}{i - 1}. \quad (16)$$

When considering RH, instead of working with \mathbf{X} as usual, we start from \mathbf{X}^π , by defining accordingly $Y_i^{(H_m)} := X_{\pi_i} \mathbb{1}_{\{H_m \geq i\}}$. This is equivalent, since for every stopping rule $\tau \in \mathcal{T}^*$, $\mathbb{E}X_\tau = \mathbb{E}X_\tau^\pi = \mathbb{E}Y_\tau^{(H_m)}$. Yet it is formally advantageous in the following argument.

Idea of the proof of Theorem 6. Consider the stopping rule $\tau_m := \zeta_{r_m} \wedge (H_m + 1) \in \mathcal{T}^*$. By (4) in Theorem 14 firstly, and the law of total expectation with respect to \mathbf{X} secondly, we have that

$$\mathbb{E}Y_{\tau_m}^{(H_m)} = \mathbb{E}\left[S_m(\zeta_{r_m})X_{\pi_{\zeta_{r_m}}}\right] = \mathbb{E}\left\{\mathbb{E}\left[S_m(\zeta_{r_m})X_{\pi_{\zeta_{r_m}}} | \mathbf{X}\right]\right\}.$$

Since at any given step (greater than $r_m - 1$) stopping at the overall maximum implies stopping at the relative maximum,

$$\mathbb{E}\left[S_m(i)X_{\pi_i} \mathbb{1}_{\{\zeta_{r_m}=i\}} | \mathbf{X}\right] \geq \mathbb{E}\left[S_m(i)X_{\pi_i} \mathbb{1}_{W_i} | \mathbf{X}\right] = S_m(i)M_m \mathbb{P}(W_i | \mathbf{X}).$$

Therefore,

$$\mathbb{E}\left[S_m(\zeta_{r_m})X_{\pi_{\zeta_{r_m}}} | \mathbf{X}\right] \geq M_m \sum_{i=r_m}^m S_m(i) \mathbb{P}(W_i | \mathbf{X}).$$

By (7) and comparison of the summation with the corresponding integral as in Lemma 21, we estimate the survival function of H_m as $m \rightarrow \infty$: since $r_m \sim m/e$, for all $i > r_m \sim m/e$ we have that

$$S_m(i) > 1 - \left[\frac{i-1}{m+1}\right]^{\frac{2\epsilon}{1+2\epsilon}}.$$

87:18 IID Prophet Inequality with Random Horizon

Define the sequence of functions

$$f_m(\varepsilon) := (m+1)^{-\frac{2\varepsilon}{1+2\varepsilon}} \frac{r_m}{m} \int_{r_m-2}^{m-1} x^{-\frac{1}{1+2\varepsilon}} dx, \quad g(\varepsilon) := \frac{1}{e} \left\{ 1 - \left[1 + \frac{1}{2\varepsilon} \right] \left(1 - e^{-\frac{2\varepsilon}{1+2\varepsilon}} \right) \right\}.$$

For every fixed $\varepsilon > 0$,

$$\sum_{i=r_m}^m S_m(i) \mathbb{P}(W_i | \mathbf{X}) \geq \sum_{i=r_m}^m \mathbb{P}(W_i | \mathbf{X}) - \sum_{i=r_m}^m \left(\frac{i-1}{m+1} \right)^{\frac{2\varepsilon}{1+2\varepsilon}} \mathbb{P}(W_i | \mathbf{X}) \geq \frac{1}{e} - f_m(\varepsilon) \longrightarrow g(\varepsilon)$$

as $m \rightarrow \infty$, where we used (14)–(16) and the usual integral estimate in the second inequality, whereas the asymptotics come from $r_m \sim m/e$. It is crucial that $g(\varepsilon)$ is positive and strictly increasing on $(0, \infty)$: even though as $\varepsilon \rightarrow 0$, $g(\varepsilon) = \mathcal{O}(\varepsilon)$, no matter how small, it still provides a constant approximation for every ε fixed.

Putting all the above together yields

$$\mathbb{E}[S_m(\zeta_{r_m}) X_{\pi_{\zeta_{r_m}}} | \mathbf{X}] \geq (g(\varepsilon) - \delta) M_m,$$

where $\delta \rightarrow 0^+$, as $m \rightarrow \infty$. This implies the final estimate

$$\mathbb{E}X_{\tau_m} := \mathbb{E}Y_{\tau_m}^{(H_m)} \geq (g(\varepsilon) - \delta) \mathbb{E}M_m.$$

Fix $\varepsilon > 0$ small and consider horizons $\{H_m\}_{m \geq M}$ as defined in (6), where $M = M(\varepsilon) \in \mathbb{N}$ is large enough. By Proposition 23, the RH problem under this family does not admit single-threshold constant-approximations for the fixed ε , assuming it is small enough (the hard instance for the values being the Pareto distribution defined in (5)). On the other hand, by our final estimate, τ_m provides an approximate $g(\varepsilon)$ -approximation: for every (small) $\varepsilon > 0$ fixed, setting M large enough, using $H_m \leq m$ yields that for every instance X ,

$$\mathbb{E}X_{\tau_m} \geq (g(\varepsilon) - \delta) \mathbb{E}M_{H_m},$$

where $\delta = \delta(M) \rightarrow 0$ as $M \rightarrow \infty$. ◀

5 A 2-approximation for concentrated \mathcal{L}^2 -horizons

Finally, we prove Theorem 8, extending the 2-approximation for the \mathcal{G} class to sufficiently concentrated horizons. We will make use of the following stochastic order.

► **Definition 24** (Laplace transform order). *Given nonnegative random variables H and G , G is dominated by H in the Laplace transform (Lt) order, denoted as $G \prec_{Lt} H$, if $\forall s > 0$, $\mathbb{E}e^{-sG} \geq \mathbb{E}e^{-sH}$.*

► **Remark 25.** Consider a horizon H , having finite expectation μ and variance σ^2 , and random variables $B \sim \text{Ber}\left(\frac{\mu^2}{\mu^2 + \sigma^2}\right)$ and $G := \frac{\mu^2 + \sigma^2}{\mu} B$. Then $G \prec_{Lt} H$.

Idea of the proof of Theorem 8. Recall that the single-threshold 2-approximation of Theorem 5 obtains at least competitive ratio

$$c_{p,\bar{q}} = 1 - \mathbb{E}[(1 - 1/\mu)^H] = 1 - \mathbb{E}(e^{-\bar{s}H})$$

with $\bar{s} := -\log(1 - 1/\mu)$. Since $G \prec_{Lt} H$, where G is the two-point distribution of Remark 25, it follows that

$$c_{p,\bar{q}} \geq 1 - \mathbb{E}(e^{-\bar{s}G}) = \frac{\mu^2}{\mu^2 + \sigma^2} \left[1 - (1 - 1/\mu)^{\frac{\mu^2 + \sigma^2}{\mu}} \right]. \quad (17)$$

As a result, an inequality that yields a single-threshold $2 - 1/\mu$ -approximation is

$$\frac{\mu^2}{\mu^2 + \sigma^2} \left[1 - (1 - 1/\mu)^{\frac{\mu^2 + \sigma^2}{\mu}} \right] \geq \frac{1}{2 - 1/\mu}, \quad (18)$$

and straightforward estimates show that this is ensured by $e^{-x} \leq 1 - x(2 - 1/\mu)^{-1}$, having defined $x := 1 + \sigma^2/\mu^2$. Furthermore, defining the constraints $y := x - (2 - 1/\mu) > -1$ and $-2 < \bar{y} := -(2 - 1/\mu) < -1$, we can rewrite the condition above as $ye^y \leq \bar{y}e^{\bar{y}}$. Under the constraints given, this is more explicitly rewritten as $y \leq W_0(\bar{y}e^{\bar{y}})$, where W_0 denotes the principal branch of the Lambert function. The definition of y and x yields (1). ◀

It is possible to substitute suitable values of $C > 1$ for $2 - 1/\mu$ in (1), the only issue would arise near the branching point of the Lambert function. With this caveat in mind, we derive Corollary 9 by substituting *admissible* values of $C > 1$ for $2 - 1/\mu$ in (18). Since we obtain

$$CV \leq \sqrt{1 - 1/\mu + W_0(-Ce^{-C})} = \sqrt{C - 1 + W_0(-Ce^{-C})},$$

this ensures a C -approximation. The degenerate cases establish insightful connections with the literature. The deterministic case ($CV = 0$) can be recovered directly from (17), which can be reformulated as a guarantee that the competitive ratio ($1/C$) is always greater than the left-hand side of

$$\frac{1 - (1 - 1/\mu)^{\mu(1+CV^2)}}{1 + CV^2} \geq \frac{1 - e^{-(1+CV^2)}}{1 + CV^2}.$$

Therefore, if $CV = 0$, the optimal single-threshold algorithm exploited ensures a competitive ratio of $1 - 1/e$. This value is the global maximum of the function on the right-hand side, and coincides with the optimal competitive ratio for single-threshold algorithms for (deterministic) IID PI [13][Theorem 21]. The left-hand side achieves global maximum $1 - (1 - 1/\mu)^\mu$ at $CV = 0$, reflecting that a deterministic horizon is no harder than a random horizon. Hence any admissible C is constrained to be greater than $[1 - (1 - 1/\mu)^\mu]^{-1}$. This proves Corollary 9 and justifies the admissibility condition.

References

- 1 A. R. Abdel-Hamid, J. A. Bather, and G. B. Trustrum. The secretary problem with an unknown number of candidates. *J. Appl. Probab.*, 19(3):619–630, 1982.
- 2 R. Alijani, S. Banerjee, S. Gollapudi, K. Munagala, and K. Wang. Predict and match: prophet inequalities with uncertain supply. *Proc. ACM Meas. Anal. Comput. Syst.*, 4(1):Article 4, 2020.
- 3 A. Aouad and W. Ma. A nonparametric framework for online stochastic matching with correlated arrivals. In *EC*, page 114. ACM, 2023. doi:10.1145/3580507.3597773.
- 4 T. Bojdecki. On optimal stopping of a sequence of independent random variables - probability maximizing approach. *Stoch. Process. Their Appl.*, 6:153–163, 1978.
- 5 C. Bracquemond, D. Roy, and M. Xie. On some discrete notions of aging. In *System and bayesian reliability. Essays in honor of Prof. E. Barlow on his 70th birthday*, pages 185–197. World Scientific, 2001.
- 6 A. Bubna and A. Chiplunkar. Prophet inequality: order selection beats random order. In *EC*, pages 302–336. ACM, 2023. doi:10.1145/3580507.3597687.
- 7 Z. Chen, Z. Huang, D. Li, and Z. G. Tang. Prophet secretary and matching: the significance of the largest item. In *SODA*, pages 1371–1401, 2025.
- 8 Y. S. Chow and H. Robbins. On optimal stopping rules. *Z. Wahrscheinlichkeitstheorie*, 2:33–49, 1963.

- 9 Y. S. Chow, H. Robbins, and D. Siegmund. *Great expectations: the theory of optimal stopping*. Houghton Mifflin, 1971.
- 10 J. Correa, P. Foncea, R. Hoeksma, R. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *EC*, pages 169–189. ACM, 2017.
- 11 J. Correa, P. Foncea, R. Hoeksma, R. Oosterwijk, and T. Vredeveld. Recent developments in prophet inequalities. *SIGecom Exch.*, 17(1):61–70, 2018. doi:10.1145/3331033.3331039.
- 12 J. Correa, P. Foncea, D. Pizarro, and V. Verdugo. From pricing to prophets, and back! *Oper. Res. Lett.*, 47:25–29, 2019. doi:10.1016/J.ORL.2018.11.010.
- 13 S. Eshani, M. Hajiaghayi, M. Kesselheim, and S. Singla. Prophet secretary for combinatorial auctions and matroids. In *SODA*, pages 700–714. SIAM, 2018. doi:10.1137/1.9781611975031.46.
- 14 P. R. Freeman. The secretary problem and its extensions: a review. *Int. Stat. Rev.*, 51(2):189–206, 1983.
- 15 G. Giambartolomei, F. Mallmann-Trenn, and S. Saona. Prophet inequalities: Separating Random Order from Order Selection. arXiv:2304.04024 [cs.DS].
- 16 J. P. Gilbert and F. Mosteller. Recognizing the maximum of a sequence. *J. Am. Stat. Assoc.*, 61(313):35–73, 1966.
- 17 M. T. Hajiaghayi, R. D. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, pages 58–65, 2007.
- 18 S. Har-Peled, E. Harb, and V. Livanos. Oracle-augmented prophet inequalities. In *ICALP*, pages 81:1–81:19, 2024.
- 19 T. P. Hill and R. P. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math.*, 125:191–207, 1992.
- 20 M. Hoefer and B. Kodric. Combinatorial secretary problems with ordinal information. In *ICALP*, pages 133:1–133:14, 2017.
- 21 R. P. Kertz. Stop rule and supremum expectations of i.i.d. random variables: a complete comparison by conjugate duality. *J. Multivar. Anal.*, 19(1):88–112, 1986.
- 22 B. Klefsjö. The hnbue and hnwue classes of life distributions. *Naval Res. Logist. Quart.*, 29:331–344, 1982.
- 23 B. Klefsjö. Testing exponentiality against hnbue. *Scand. J. Stat.*, 10(2):65–75, 1983.
- 24 B. Klefsjö. A useful ageing property based on the laplace transform. *J. Appl. Prob.*, 20:615–626, 1983.
- 25 U. Krengel and L. Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.*, 83(4):745–747, 1977.
- 26 U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Adv. Probab. Related Topics*, 4:197–266, 1978.
- 27 N. A. Langberg, R. V. Léon, J. Lynch, and F. Proschan. Extreme points of the class of discrete decreasing failure life distributions. *Math. Oper. Res.*, 5(1):35–42, 1980. doi:10.1287/MOOR.5.1.35.
- 28 D. V. Lindley. Dynamic programming and decision theory. *J. R. Stat. Soc. Ser. C Appl. Stat.*, 10(1):39–51, 1961.
- 29 B. Lucier. An economic view of prophet inequalities. *SIGecom Exch.*, 16(1):24–47, 2017. doi:10.1145/3144722.3144725.
- 30 M. Mahdian and A. Saberi. Multi-unit auctions with unknown supply. In *CEC*, pages 243–249. ACM, 2006. doi:10.1145/1134707.1134734.
- 31 S. Oveis Gharan and J. Vondrák. On variants of the matroid secretary problem. *Algorithmica*, 67:472,497, 2013.
- 32 Z. Poroziński. The full-information best choice problem with a random number of observations. *Stoch. Process. Their Appl.*, 24:293–307, 1987.
- 33 E. L. Presman and I. M. Sonin. The best choice problem for a random number of objects. *Theory Probab. its Appl.*, 17(4):657–668, 1972.
- 34 T. Rolski. Mean residual life. *Bull. Int. Statist. Inst.*, 46:220–266, 1975.

- 35 E. Samuel-Cahn. Comparisons of threshold stop rule and maximum for independent nonnegative random variables. *Ann. Probab.*, 12(4):1213–1216, 1984.
- 36 E. Samuel-Cahn. The best-choice secretary problem with random freeze on jobs. *Stoch. Process. Their Appl.*, 55(2):315–327, 1995.
- 37 E. Samuel-Cahn. Optimal stopping with random horizon with application to the full-information best-choice problem with random freeze. *J. Am. Stat. Assoc.*, 91(433):357–364, 1996.
- 38 T. Sandholm. Automated mechanism design: A new application area for search algorithms. In *CP*, pages 19–36, 2003.
- 39 M. Shaked and J. G. Shanthikumar. *Stochastic orders*. Springer, 2007.