
Lookback Prophet Inequalities

Ziyad Benomar
ENSAE, Ecole Polytechnique,
FairPlay joint team
ziyad.benomar@ensae.fr

Dorian Baudry
Department of Statistics,
University of Oxford
dorian.baudry@ox.ac.uk

Vianney Perchet
CREST, ENSAE, Criteo AI LAB
Fairplay joint team
vianney.perchet@normalesup.org

Abstract

Prophet inequalities are fundamental optimal stopping problems, where a decision-maker observes sequentially items with values sampled independently from known distributions, and must decide at each new observation to either stop and gain the current value or reject it irrevocably and move to the next step. This model is often too pessimistic and does not adequately represent real-world online selection processes. Potentially, rejected items can be revisited and a fraction of their value can be recovered. To analyze this problem, we consider general decay functions D_1, D_2, \dots , quantifying the value to be recovered from a rejected item, depending on how far it has been observed in the past. We analyze how lookback improves, or not, the competitive ratio in prophet inequalities in different order models. We show that, under mild monotonicity assumptions on the decay functions, the problem can be reduced to the case where all the decay functions are equal to the same function $x \mapsto \gamma x$, where $\gamma = \inf_{x>0} \inf_{j \geq 1} D_j(x)/x$. Consequently, we focus on this setting and refine the analyses of the competitive ratios, with upper and lower bounds expressed as increasing functions of γ .

1 Introduction

Optimal stopping problems constitute a classical paradigm of decision-making under uncertainty [Dynkin, 1963]. Typically, in online algorithms, these problems are formalized as variations of the secretary problem [Lindley, 1961] or the prophet inequality [Krengel and Sucheston, 1977]. In the context of the prophet inequality, the decision-maker observes a finite sequence of items, each having a value drawn independently from a known probability distribution. Upon encountering a new item, the decision-maker faces the choice of either accepting it and concluding the selection process or irreversibly rejecting it, with the objective of maximizing the value of the selected item. However, while the prophet inequality problem is already used in scenarios such as posted-price mechanism design [Hajiaghayi et al., 2007] or online auctions [Syrgkanis, 2017], it might present a pessimistic model of real-world online selection problems. Indeed, it is in general possible in practice to revisit previously rejected items and potentially recover them or at least recover a fraction of their value.

Consider for instance an individual navigating a city in search of a restaurant. When encountering one, they have the choice to stop and dine at this place, continue their search, or revisit a previously passed option, incurring a utility cost that is proportional to the distance of backtracking. In another example drawn from the real estate market, homeowners receive offers from potential buyers. The decision to accept or reject an offer can be revisited later, although buyer interest may have changed, resulting in a potentially lower offer or even a lack of interest. Lastly, in the financial domain, an

agent may choose to sell an asset at its current price or opt for a lookback put option, allowing them to sell at the asset's highest price over a specified future period. To make a meaningful comparison between the two, one must account for factors such as discounting (time value of money) and the cost of the option.

1.1 Formal problem and notation

To encompass diverse scenarios, we propose a general way to quantify the cost incurred by the decision-maker for retrieving a previously rejected value.

Definition 1.1 (Decay functions). *Let $\mathcal{D} = (D_1, D_2, \dots)$ be a sequence of non-negative functions defined on $[0, \infty)$. It is a sequence of decay functions if*

- (i) $D_1(x) \leq x$ for all $x \geq 0$,
- (ii) the sequence $(D_j(x))_{j \geq 1}$ is non-increasing for all $x \geq 0$,
- (iii) the function D_j is non-decreasing for all $j \geq 1$.

In the context of decay functions \mathcal{D} , if a value x is rejected, the algorithm can recover $D_j(x)$ after j subsequent steps. The three conditions defining decay functions serve as fundamental prerequisites for the problem. The first and second conditions ensure that the recoverable value of a rejected item can only diminish over time, while the final condition implies that an increase in the observed value x corresponds to an increase in the potential recovered value. Although the non-negativity of the decay functions is non-essential, we retain it for convenience, as we can easily revert to this assumption by considering that the algorithm has a reward of zero by not selecting any item.

The motivating examples that we introduced can be modeled respectively with decay functions of the form $D_j(x) = x - c_j$ where $(c_j)_{j \geq 1}$ is a non-decreasing positive sequence, $D_j(x) = \xi_j x$ with $\xi_j \sim \mathcal{B}(p_j)$ and $(p_j)_{j \geq 1}$ a non-increasing sequence of probabilities, and $D_j(x) = \lambda^j x$ with $\lambda \in [0, 1]$. In one of these examples (housing market), the natural model is to use *random decay functions*: the buyer makes the same offer if they are still interested, and offers 0 otherwise. Definition 1.1 can be easily extended to consider this case. However, to enhance the clarity of the presentation, we only discuss the deterministic case in the rest of the paper. In Appendix D, we explain how all the proofs and theorems can be generalized to that case.

The \mathcal{D} -prophet inequality. For any decay functions \mathcal{D} , we define the \mathcal{D} -prophet inequality problem, where the decision maker, knowing \mathcal{D} , observes sequentially the values X_1, \dots, X_n , with X_i drawn from a known distribution F_i for all $i \in [n]$. If they decide to stop at some step τ , then instead of gaining X_τ as in the classical prophet inequality, they can choose to select the current item X_τ and have its full value, or select any item X_i with $i < \tau$ among the rejected ones but only recover a fraction $D_{\tau-i}(X_i) \leq X_i$ of its value. Therefore, if an algorithm **ALG** stops at step τ its reward is

$$\begin{aligned} \text{ALG}^{\mathcal{D}}(X_1, \dots, X_n) &= \max\{X_\tau, D_1(X_{\tau-1}), D_2(X_{\tau-2}), \dots, D_{\tau-1}(X_1)\} \\ &= \max_{0 \leq i \leq \tau-1} \{D_i(X_{\tau-i})\}, \end{aligned}$$

with the convention $D_0(x) = x$. If the algorithm does not stop at any step before n , then its reward is $\text{ALG}^{\mathcal{D}}(X_1, \dots, X_n) = \max_{1 \leq i \leq n} \{D_i(X_{\pi(n-i+1)})\}$, which corresponds to $\tau = n + 1$.

Remark 1.1. *As in the standard prophet inequality, an algorithm is defined by its stopping time, i.e., the rule set to decide whether to stop or not. Hence, if \mathcal{D} and \mathcal{D}' are two different sequences of decay functions, any algorithm for the \mathcal{D} -prophet inequality, although its stopping time might depend on the particular sequence of functions \mathcal{D} , is also an algorithm for the \mathcal{D}' -prophet inequality. Consider for example an algorithm **ALG** with stopping time $\tau(\mathcal{D})$ that depends on \mathcal{D} . Its reward in the \mathcal{D}' -prophet inequality is $\text{ALG}^{\mathcal{D}'}(X_1, \dots, X_n) = \max_{0 \leq i \leq \tau-1} \{D'_i(X_{\tau(\mathcal{D})-i})\}$.*

Observation order. Several variants of the prophet inequality problem have been studied, depending on the order of observations. The standard model is the adversarial (or fixed) order: The instance of the distributions F_1, \dots, F_n is chosen by an adversary, and the algorithm observes the samples $X_1 \sim F_1, \dots, X_n \sim F_n$ in this order [Krengel and Sucheston, 1977, 1978]. In the *random order* model, the adversary can again choose the distributions, but the algorithm observes the samples in a

uniformly random order. Another setting in which the observation order is no longer important is the IID model [Hill and Kertz, 1982, Correa et al., 2021b], where all the values are sampled independently from the same distribution F . The \mathcal{D} -prophet inequality is well-defined in each of these different order models: if the items are observed in the order $X_{\pi(1)}, \dots, X_{\pi(n)}$ with π a permutation of $[n]$, then the reward of the algorithm is $\text{ALG}^{\mathcal{D}}(X_1, \dots, X_n) = \max_{0 \leq i \leq \tau-1} \{D_i(X_{\pi(\tau-i)})\}$. In this paper, we study the \mathcal{D} -prophet inequality in the three models we presented, providing lower and upper bounds in each of them.

Competitive ratio. In the \mathcal{D} -prophet inequality, an input instance I is a finite sequence of probability distributions (F_1, \dots, F_n) . Thus, for any instance I , we denote by $\mathbb{E}[\text{ALG}^{\mathcal{D}}(I)]$ the expected reward of ALG given I as input, and we denote by $\mathbb{E}[\text{OPT}(I)]$ the expected maximum of independent random variables $(X_i)_{i \in [n]}$, where $X_i \sim F_i$. With these notations, we define the competitive ratio, which will be used to measure the quality of the algorithms.

Definition 1.2 (Competitive ratio). *Let \mathcal{D} be a sequence of decay functions and ALG an algorithm. We define the competitive ratio of ALG by*

$$CR^{\mathcal{D}}(\text{ALG}) = \inf_I \frac{\mathbb{E}[\text{ALG}^{\mathcal{D}}(I)]}{\mathbb{E}[\text{OPT}(I)]},$$

with the infimum taken over the tuples of all sizes of non-negative distributions with finite expectation.

An algorithm is said to be α -competitive if its competitive ratio is at least α , which means that for any possible instance I , the algorithm guarantees a reward of at least $\alpha \mathbb{E}[\text{OPT}(I)]$. The notion of competitive ratio is used more broadly in competitive analysis as a metric to evaluate online algorithms [Borodin and El-Yaniv, 2005].

1.2 Contributions

It is trivial that non-zero decay functions \mathcal{D} guarantee a better reward compared to the classical prophet inequality. However, in general, this is not sufficient to conclude that the standard upper bounds or the competitive ratio of a given algorithm can be improved. Hence, a first key question is: what condition on \mathcal{D} is necessary to surpass the conventional upper bounds of the classical prophet inequality? Surprisingly, the answer hinges solely on the constant $\gamma_{\mathcal{D}}$, defined as follows,

$$\gamma_{\mathcal{D}} = \inf_{x>0} \inf_{j \geq 1} \left\{ \frac{D_j(x)}{x} \right\}. \quad (1)$$

In the adversarial order model, we demonstrate that the optimal competitive ratio achievable in the \mathcal{D} -prophet inequality is determined by the parameter $\gamma_{\mathcal{D}}$ alone. Additionally, in both the random order and IID models, we demonstrate the essential requirement of $\gamma_{\mathcal{D}} > 0$ for breaking the upper bounds of the classical prophet inequality. In particular, this implies that no improvement can be made with decay functions of the form $D_j(x) = x - c_j$ with $c_j > 0$, or $D_j(x) = \lambda^j x$ with $\lambda \in [0, 1)$. Subsequently, we develop algorithms and provide upper bounds in the \mathcal{D} -prophet inequality, uniquely dependent on the parameter $\gamma_{\mathcal{D}}$. We illustrate them in Figure 1, comparing them with the identity function $\gamma \mapsto \gamma$, which is a trivial lower bound.

1.3 Related work

Prophet inequalities. The first prophet inequality was proven by Krengel and Sucheston [Krengel and Sucheston, 1977, 1978] in the setting where the items are observed in a fixed order, demonstrating that the dynamic programming algorithm has a competitive ratio of $1/2$, which is the best possible. It was shown later that the same guarantee can be obtained with simpler algorithms [Samuel-Cahn, 1984, Kleinberg and Weinberg, 2012], accepting the first value above a carefully chosen threshold. For a more comprehensive and historical overview, we refer the interested reader to surveys on the problem such as [Lucier, 2017, Correa et al., 2019]. Prophet inequalities have immediate applications in mechanism design [Hajiaghayi et al., 2007, Deng et al., 2022, Psomas et al., 2022, Makur et al., 2024], auctions [Syrgkanis, 2017, Dütting et al., 2020], resource management [Sinclair et al., 2023], and online matching [Cohen et al., 2019, Ezra et al., 2020, Jiang et al., 2021, Papadimitriou et al., 2021, Brubach et al., 2021]. Many variants and related problems have been studied, including, for example, the matroid prophet inequality [Kleinberg and Weinberg, 2012, Feldman et al., 2016],

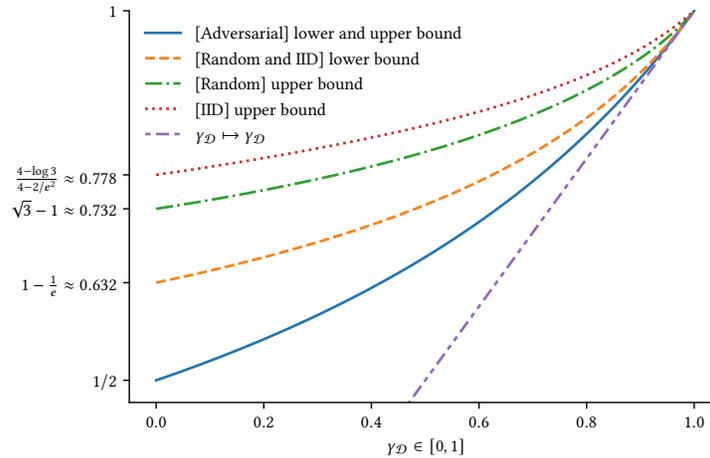


Figure 1: Lower and upper bounds on the competitive ratio in the \mathcal{D} -prophet inequality depending on $\gamma_{\mathcal{D}}$, in the adversarial order (Thm 4.3), random order (Thm 4.4) and IID (Thm 4.6) models

prophet inequality with advice [Diakonikolas et al., 2021], and variants with fairness considerations [Correa et al., 2021a, Arsenis and Kleinberg, 2022].

Random order and IID models. Esfandiari et al. [2017] introduced the *prophet secretary* problem, where items are observed in a uniformly random order, and they proved a $(1 - \frac{1}{e})$ -competitive algorithm. Correa et al. [2021c] showed later a competitive ratio of 0.669, and Harb [2024] enhanced it to 0.6724, which currently stands as the best-known solution for the problem. They also proved an upper bound of $\sqrt{3} - 1 \approx 0.732$, which was improved to 0.7254 in [Bubna and Chiplunkar, 2023] then 0.723 in [Giambartolomei et al., 2023]. Addressing the gap between the lower and upper bound remains an engaging and actively pursued open question. On the other hand, the study of prophet inequalities with IID random variables dates back to papers such as [Hill and Kertz, 1982, Kertz, 1986], demonstrating guarantees on the dynamic programming algorithm. The problem was completely solved in [Correa et al., 2021b], where the authors show that the competitive ratio of the dynamic programming algorithm is 0.745, thus it constitutes an upper bound on the competitive ratio of any algorithm, and they give a simpler adaptive threshold algorithm matching it. Another setting that we do not study in this paper, is the *order selection* model, where the decision-maker can choose the order in which the items are observed, knowing their distributions [Chawla et al., 2010, Beyhaghi et al., 2021, Peng and Tang, 2022].

Beyond the worst-case. In recent years, there has been increasing interest in exploring ways to exceed the worst-case upper bounds of online algorithms by providing the decision-maker with additional capabilities. A notable research avenue is learning-augmented algorithms [Lykouris and Vassilvtiskii, 2018], which equip the decision-maker with predictions or hints about unknown variables of the problem. Multiple problems have been studied in this framework, such as scheduling [Purohit et al., 2018, Lassota et al., 2023, Benomar and Perchet, 2024b], matching [Antoniadis et al., 2020, Dinitz et al., 2021, Chen et al., 2022], caching [Antoniadis et al., 2023, Chlkedowski et al., 2021, Christianson et al., 2023], the design of data structures [Kraska et al., 2018, Lin et al., 2022, Benomar and Coester, 2024], and in particular, online selection problems [Dütting et al., 2021, Sun et al., 2021, Benomar et al., 2023, Benomar and Perchet, 2024a, Diakonikolas et al., 2021]. More related to our setting, the ability to revisit items in online selection has been studied in problems such as the multiple-choice prophet inequality, where the algorithm can select up to k items and its reward is the maximum selected value [Assaf and Samuel-Cahn, 2000]. This allows for revisiting up to k items, chosen during the execution, for final acceptance or rejection decisions. Similarly, in Pandora’s box problem [Weitzman, 1978, Kleinberg et al., 2016] and its variants [Esfandiari et al., 2019, Gergatsouli and Tzamos, 2022, Atsidakou et al., 2024, Gergatsouli and Tzamos, 2024, Berger et al., 2024], the decision maker decides the observation order of the items, but a cost c_i is paid for observing each value X_i , with the gain being the maximum observed value minus the total opening costs. A very recent work investigates a scenario closely related to the lookback prophet inequality

[Ekbatani et al., 2024] where, upon selecting a candidate X_i , the decision-maker has the option to discard it and choose a new value X_j at any later step j , at a buyback cost of fX_i , where $f > 0$. The authors present an optimal algorithm for the case when $f \geq 1$, although the problem remains open for $f \in (0, 1)$. Other problems were studied in similar settings, such as online matching [Ekbatani et al., 2022] and online resource allocation [Ekbatani et al., 2023].

2 From \mathcal{D} -prophet to the D_∞ -prophet inequality

Let us consider a sequence \mathcal{D} of decay functions. By Definition 1.1, for any $x \in [0, \infty]$ the sequence $(D_j(x))_{j \geq 1}$ converges, since it is non-increasing and non-negative. Hence, there exists a mapping D_∞ such that for any $x \geq 0$, $\lim_{j \rightarrow \infty} D_j(x) = D_\infty(x)$. Furthermore, we can easily verify that D_∞ is non-decreasing and satisfies $D_\infty(x) \in [0, x]$ for all $x \geq 0$.

Thanks to these properties, we obtain that $(D_\infty)_{j \geq 1}$ also satisfies Definition 1.1, and is hence a valid sequence of decay functions. We thus refer to the corresponding problem as the D_∞ -prophet inequality. Since $D_j \geq D_\infty$ for any $j \geq 1$, it is straightforward that the stopping problem with the decay functions D_∞ would be less favorable to the decision-maker. More precisely, for any random variables X_1, \dots, X_n , observation order π , and algorithm ALG with stopping time τ , it holds that

$$\text{ALG}^{\mathcal{D}}(X_1, \dots, X_n) := \max\{X_{\pi(\tau)}, \max_{i < \tau} D_{\tau-i}(X_{\pi(i)})\} \geq \max\{X_{\pi(\tau)}, \max_{i < \tau} D_\infty(X_{\pi(i)})\},$$

which corresponds to the output of ALG (with the same decision rule) when all the decay functions are equal to D_∞ . Therefore, any guarantees established for algorithms in the D_∞ -prophet inequality naturally extend to the \mathcal{D} -prophet inequality. However, it remains uncertain whether the \mathcal{D} -prophet inequality can yield improved competitive ratios compared to the D_∞ -prophet inequality. In the following, we prove that this is not the case, for all the order models presented in Section 1.

Theorem 2.1. *Let D_∞ be the pointwise limit of the sequence of decay functions $\mathcal{D} = (D_j)_{j \geq 1}$. Then for any instance $I = (F_1, \dots, F_n)$ of non-negative distributions, it holds in the adversarial and the random order models that*

$$\forall \text{ALG} : \text{CR}^{\mathcal{D}}(\text{ALG}) \leq \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(I)]}{\mathbb{E}[\text{OPT}(I)]}, \quad (2)$$

where the supremum is taken over all the online algorithms \mathbf{A} . In the IID model, the same inequality holds with an additional $O(n^{-1/3})$ term, which depends only on the size n of the instance.

The main implication of Theorem 2.1 is the following corollary.

Corollary 2.1.1. *In the adversarial order and the random order models, if $\bar{\mathbf{A}}_\infty$ is an optimal algorithm for the D_∞ -prophet inequality, i.e. with maximal competitive ratio, then $\bar{\mathbf{A}}_\infty$ is also optimal for the \mathcal{D} -prophet inequality. Moreover, it holds that*

$$\text{CR}^{\mathcal{D}}(\bar{\mathbf{A}}_\infty) = \text{CR}^{D_\infty}(\bar{\mathbf{A}}_\infty).$$

A direct consequence of this result is that, in the adversarial and the random order models, the asymptotic decay D_∞ entirely determines the competitive ratio that is achievable and the upper bounds for the \mathcal{D} -prophet inequality. Therefore, we can restrict our analysis to algorithms designed for the problem with identical decay function. In the IID model, the same conclusion holds if the worst-case instances are arbitrarily large, making the additional $O(n^{-1/3})$ term vanish. This is the case in particular in the classical IID prophet inequality [Hill and Kertz, 1982].

2.1 Sketch of the proof of Theorem 2.1

While we use different techniques for each order model considered, all the proofs share the same underlying idea. Given any instance I of non-negative distributions, we build an alternative instance J such that the reward of any algorithm on I with decay functions $\mathcal{D} = (D_j)_j$ is at most its reward on J with decay functions all equal to D_∞ . To do this, we essentially introduce an arbitrarily large number of zero values between two successive observations drawn from distributions belonging to I . Hence, under J , the algorithm cannot recover much more than a fraction $D_\infty(X)$ for any past observation X collected from a distribution $F \in I$.

In the adversarial case, implementing this idea is straightforward, since nature can build J by directly inserting m zeros between each pair of consecutive values, and the result is obtained by making m arbitrarily large. For the random order model, we use the same instance J , but extra steps are needed to prove that the number of steps between two non-zero values is very large with high probability.

Moving to the IID model, an instance I is defined by a pair (F, n) , where F is a non-negative distribution, and n is the size of the instance. In this scenario, we consider an instance consisting of $m > n$ IID random variables $(Y_i)_{i \in [m]}$, each sampled from F with probability n/m , and equal to zero with the remaining probability. We again achieve the desired result by letting m be arbitrarily large compared to n . However, the number of variables sampled from F is not fixed; it follows a Binomial distribution with parameters $(m, n/m)$. We control this variability by using concentration inequalities, which causes the additional term $O(n^{-1/3})$.

3 From D_∞ -prophet to the γ_D -prophet inequality

As discussed in Section 2, Theorem 2.1 implies that, for either establishing upper bounds or guarantees on the competitive ratios of algorithms, it is sufficient to study the D_∞ -prophet inequality, where all the decay functions are equal to D_∞ . The remaining question is then to determine which functions D_∞ allow to improve upon the upper bounds of the classical prophet inequality. Before tackling this question, let us make some observations regarding algorithms in the D_∞ -prophet inequality.

In the D_∞ -prophet inequality, it is always possible to have a reward of $D_\infty(\max_{i \in [n]} X_i)$ by rejecting all the items and then selecting the maximum by the end. Thus, it is suboptimal to stop at a step i where $X_i \leq D_\infty(\max_{j < i} X_j)$. An algorithm respecting this decision rule is called *rational*.

Lemma 3.1. *For any rational algorithm ALG in the D_∞ -prophet inequality, if we denote by τ its stopping time, then for any instance $I = (F_1, \dots, F_n)$ and $X_i \sim F_i$ for all $i \in [n]$ we have*

$$ALG^{D_\infty}(X_1, \dots, X_n) = ALG^0(X_1, \dots, X_n) + D_\infty\left(\max_{i \in [n]} X_i\right) \mathbb{1}_{\tau=n+1},$$

where ALG^0 denotes the reward of the algorithm in the standard prophet inequality. Moreover, the optimal dynamic programming algorithm in the D_∞ -prophet inequality is rational.

The best competitive ratio in the D_∞ -prophet inequality is achieved, possibly among others, by the optimal dynamic programming algorithm, which is a rational algorithm by the previous Lemma. Hence, it suffices to prove upper bounds on rational algorithms. We use this observation to prove the next propositions.

Proposition 3.2. *In the D_∞ -prophet inequality, if $\inf_{x>0} \frac{D_\infty(x)}{x} = 0$, then it holds, in any order model, that*

$$\forall ALG : CR^{D_\infty}(ALG) \leq \sup_A CR^0(A), \quad (3)$$

where the supremum is taken over all the online algorithms A , and CR^0 denotes the competitive ratio in the standard prophet inequality.

Proposition 3.2 implies that if $\inf_{x>0} \frac{D_\infty(x)}{x} = 0$, then, in any order model, any upper bound on the competitive ratios of all algorithms in the classical prophet inequality is also an upper bound on the competitive ratios of all algorithm in the D_∞ -prophet inequality. Consequently, for surpassing the upper bounds of the classical prophet inequality, it is necessary to have, for some $\gamma > 0$, that $D_\infty(x) \geq \gamma x$ for all $x \geq 0$. Furthermore, the next Proposition allows giving upper bounds in the D_∞ -prophet inequality that depend only on $\inf_{x>0} \frac{D_\infty(x)}{x}$.

Proposition 3.3. *Let $\gamma = \inf_{x>0} D_\infty(x)/x$, and $0 < a < b$. Consider an instance I of distributions with support in $\{0, a, b\}$, then in any order model and for any algorithm ALG we have that*

$$CR^{D_\infty}(ALG) \leq \sup_A \frac{\mathbb{E}[A^\gamma(I)]}{\mathbb{E}[OPT(I)]},$$

where $\mathbb{E}[A^\gamma(I)]$ is the reward of A if all the decay functions were equal to $x \mapsto \gamma x$.

The core idea for proving this proposition is that rescaling an instance, i.e. considering $(rX_i)_{i \in [n]}$ instead of $(X_i)_{i \in [n]}$, has no impact in the classical prophet inequality. However, in the D_∞ -prophet

inequality, rescaling can be exploited to adjust the ratio $\frac{D_\infty(rx)}{rx}$. By considering instances with random variables taking values in $\{0, a, b\}$ almost surely, where $a < b$, a reasonable algorithm facing such an instance would never reject the value b . Consequently, the value it recovers from rejected items is either $D_\infty(0) = 0$ or $D_\infty(a)$. Rescaling this instance by a factor $r = s/a$ and taking the ratio to the expected maximum, the term $\frac{D_\infty(s)}{s}$ appears, with s a free parameter that can be chosen to satisfy $\frac{D_\infty(s)}{s} \rightarrow \inf_{x>0} \frac{D_\infty(x)}{x} = \gamma_{\mathcal{D}}$.

As a consequence, if $\inf_{x>0} \frac{D_\infty(x)}{x} = \gamma$, then any upper bound obtained in the γ -prophet inequality (when the decay functions are all equal to $x \mapsto \gamma x$) using instances of random variables (X_1, \dots, X_n) satisfying $X_i \in \{0, a, b\}$ a.s. for all i , is also an upper bound in the D_∞ -prophet inequality.

Implication Consider any sequence \mathcal{D} of decay functions, and define

$$\gamma_{\mathcal{D}} := \inf_{x>0} \left\{ \frac{D_\infty(x)}{x} \right\} = \inf_{x>0} \inf_{j \geq 1} \left\{ \frac{D_j(x)}{x} \right\}.$$

For any $x > 0$ and $j \geq 1$ it holds that $D_j(x) \geq \gamma_{\mathcal{D}}x$, therefore, any guarantees on the competitive ratio of an algorithm in the $\gamma_{\mathcal{D}}$ -prophet inequality are valid in the \mathcal{D} -prophet inequality, under any order model. Furthermore, combining Theorem 2.1 and Proposition 3.3, we obtain that for any instance I of random variables taking values in a set $\{0, a, b\}$ it holds that

$$\forall \text{ALG} : \text{CR}^{\mathcal{D}}(\text{ALG}) \leq \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^{\gamma_{\mathcal{D}}}(I)]}{\mathbb{E}[\text{OPT}(I)]},$$

with an additional term of order $O(n^{-1/3})$ in the IID model. In the particular case where $\gamma_{\mathcal{D}} = 0$, Proposition 3.3 with Theorem 2.1 give a stronger result, showing that no algorithm can surpass the upper bounds of the classical prophet inequality. This is true also for the IID model since the instances used to prove the tight upper bound of ≈ 0.745 are of arbitrarily large size [Hill and Kertz, 1982].

Therefore, by studying the γ -prophet inequality for $\gamma \in [0, 1]$, we can prove upper bounds and lower bounds on the \mathcal{D} -prophet inequality for any sequence \mathcal{D} of decay functions.

4 The γ -prophet inequality

We study in this section the γ -prophet inequality, where all the decay functions are equal to $x \mapsto \gamma x$, for some $\gamma \in [0, 1]$. For any algorithm ALG with stopping time τ and random variables X_1, \dots, X_n , if the observation order is π , we use the notation

$$\text{ALG}^\gamma(X_1, \dots, X_n) = \max\{X_{\pi(\tau)}, \gamma X_{\pi(\tau-1)}, \dots, \gamma X_{\pi(1)}\}.$$

and we denote by $\text{CR}^\gamma(\text{ALG})$ the competitive ratio of ALG in this setting. In the following, we provide theoretical guarantees for the γ -prophet inequality.

For each observation order, we first derive upper bounds on the competitive ratio of any algorithm, depending on γ , using only hard instances satisfying the condition of Proposition 3.3. This would guarantee that the upper bounds extend to the \mathcal{D} -prophet inequality if $\gamma_{\mathcal{D}} = \gamma$. Then, we design single-threshold algorithms with well-chosen thresholds depending on γ and the distributions, with competitive ratios improving with γ . A crucial property of single-threshold algorithms, which we use to estimate their competitive ratios, is that their reward satisfies

$$\text{ALG}^\gamma(X_1, \dots, X_n) = \text{ALG}^0(X_1, \dots, X_n) + \gamma(\max_i X_i) \mathbb{1}_{(\max_i X_i < \theta)}. \quad (4)$$

The additional term appearing due to γ depends only on $\max_{i \in [n]} X_i$, which is the reward of the prophet against whom we compete. This property is not satisfied by more general class of algorithms such as multiple-threshold algorithms, where each observation $X_{\pi(i)}$ is compared with a threshold θ_i .

Remark 4.1. *We only consider instances with continuous distributions in the proofs of lower bounds. The thresholds θ considered are such that $\Pr(\max_{i \in [n]} X_i \geq \theta) = g(\gamma, n, \pi)$, with g depending on γ , the order model π and the size of the instance n . Such a threshold is always guaranteed to exist when the distributions are continuous. However, as in the prophet inequality, the algorithms can be easily adapted to non-continuous distributions by allowing stochastic tie-breaking. A detailed strategy for doing this can be found for example in [Correa et al., 2021c].*

Before delving into the study of the different models, we provide generic lower and upper bounds, which depend solely on the bounds of the classical prophet inequality and γ .

Proposition 4.2. *In any order model, if α is a lower bound in the classical prophet inequality, and β an upper bound, then, in the γ -prophet inequality*

- (i) *there exists a trivial algorithm with a competitive ratio of at least $\max\{\gamma, \alpha\}$,*
- (ii) *the competitive ratio of any algorithm is at most $(1 - \gamma)\beta + \gamma$.*

4.1 Adversarial order

We first consider the adversarial order model, and prove the upper bound of $\frac{1}{2-\gamma}$. Then, we provide a single-threshold algorithm with a competitive ratio matching this upper bound, hence fully solving the γ -prophet inequality in this adversarial order model.

Theorem 4.3. *In the adversarial order model, the competitive ratio of any algorithm is at most $\frac{1}{2-\gamma}$. Furthermore, there exists a single threshold algorithm with a competitive ratio $\frac{1}{2-\gamma}$: given any instance (F_1, \dots, F_n) , this is achieved with the threshold θ satisfying*

$$\Pr_{X_1 \sim F_1, \dots, X_n \sim F_n}(\max_{i \in [n]} X_i \leq \theta) = \frac{1}{2-\gamma}.$$

The upper bound in the previous theorem is proved using instances satisfying the condition of Proposition 3.3. Hence it extends to the D_∞ - then to the \mathcal{D} -prophet inequality, with $\gamma = \gamma_{\mathcal{D}}$, by Proposition 3.3 and Theorem 2.1.

4.2 Random order

Consider now that the items are observed in a uniformly random order $X_{\pi(1)}, \dots, X_{\pi(n)}$, and $X^* = \max_{i \in [n]} X_i$. As for the adversarial model, we first prove an upper bound on the competitive ratio as a function of γ , and then prove a lower bound for a single-threshold algorithm. However, for this model, there is a gap between the two bounds, as illustrated in Figure 1.

We first prove an upper bound that depends on γ , matching the upper bound $\sqrt{3} - 1$ of Correa et al. [2021c] when $\gamma = 0$ and equal to 1 when $\gamma = 1$. Our single-threshold algorithm has a competitive ratio of at least $(1 - \frac{1}{e})$ when $\gamma = 0$, which is the best competitive ratio of a single threshold algorithm in the prophet inequality [Esfandiari et al., 2017, Correa et al., 2021c], and equal to 1 for $\gamma = 1$.

Theorem 4.4. *The competitive ratio of any algorithm ALG in the γ -prophet inequality with random order satisfies*

$$CR^\gamma(ALG) \leq (1 - \gamma)^{3/2}(\sqrt{3 - \gamma} - \sqrt{1 - \gamma}) + \gamma.$$

Furthermore, denoting by p_γ is the unique solution to the equation $1 - (1 - \gamma)p = \frac{1-p}{-\log p}$, the single-threshold algorithm ALG_θ with $\Pr_{X_1 \sim F_1, \dots, X_n \sim F_n}(\max_{i \in [n]} X_i \leq \theta) = p_\gamma$ satisfies

$$CR^\gamma(ALG) \geq 1 - (1 - \gamma)p_\gamma.$$

Similarly to the adversarial order model, we used instances satisfying the condition of Proposition 3.3 to prove the upper bound, thus it extends to the \mathcal{D} -prophet inequality with $\gamma = \gamma_{\mathcal{D}}$.

While the equation defining p_γ cannot be solved analytically, the solution can easily be computed numerically for any $\gamma \in [0, 1]$. Before moving to the IID case, we propose in the following a more explicit lower bound derived from Theorem 4.4.

Corollary 4.4.1. *In the random order model, the single threshold algorithm with a threshold θ satisfying $\Pr(\max_{i \in [n]} X_i \geq \theta) = \frac{1/e}{1 - (1-1/e)\gamma}$ has a competitive ratio of at least $1 - \frac{(1-\gamma)/e}{1 - (1-1/e)\gamma}$.*

4.3 IID Random Variables

In the classical IID prophet inequality, [Hill and Kertz, 1982] showed that the competitive ratio of any algorithm is at most ≈ 0.745 . The proof of this upper bound is hard to generalize for the IID γ -prophet inequality. As an alternative, we prove a weaker upper bound, which equals ≈ 0.778

for $\gamma = 0$ and 1 for $\gamma = 1$, and the proof relies on instances of arbitrarily large size satisfying the condition of Proposition 3.3, hence the upper bound can be extended to the \mathcal{D} -prophet inequality.

Subsequently, we present a single-threshold algorithm with the same competitive ratio as the random order algorithm. However, the proof is different, leveraging the fact that the variables are identically distributed. More precisely, we introduce a single-threshold algorithm with guarantees that depend on the size n of the instance, then we show that its competitive ratio is at least that of the algorithm presented in Theorem 4.4, with equality when n approaches infinity.

Although it might look surprising that the obtained competitive ratio in the IID model is not better than that of the random-order model, the same behavior occurs in the classical prophet inequality. Indeed, Li et al. [2022] established that no single-threshold algorithm can achieve a competitive ratio better than $1 - 1/e$ in the standard prophet inequality with IID random variables, which is also the best possible with a single-threshold algorithm in the random order. However, considering larger classes of algorithms, the competitive ratios achieved in the IID model are better than those of the random order model.

We describe the algorithm and give a first lower bound on its reward depending on the size of the instance in the following lemma.

Lemma 4.5. *Let $a_{n,\gamma}$ be the unique solution of the equation $\left(\frac{1}{(1-a/n)^n} - 1\right) \left(\frac{1}{a} - 1\right) = \gamma$, then for any IID instance X_1, \dots, X_n , the algorithm with threshold θ satisfying $\Pr(X_1 > \theta) = \frac{a_{n,\gamma}}{n}$ has a reward of at least*

$$\frac{1}{a_{n,\gamma}} \left(1 - \left(1 - \frac{a_{n,\gamma}}{n}\right)^n\right) \mathbb{E}[\max_{i \in [n]} X_i].$$

We can prove that the reward presented in the Lemma above is strictly better than that of the random order model. However, both are asymptotically equal as we show in the following theorem.

Theorem 4.6. *The competitive ratio of any algorithm in the IID γ -prophet inequality is at most*

$$U(\gamma) = 1 - (1 - \gamma) \frac{e^2 \log(3 - \gamma) - (2 - \gamma)}{2(2e^2 - 1) - (3e^2 - 1)\gamma}.$$

In particular, U is increasing, $U(0) = \frac{4 - \log 3}{2(2 - \frac{1}{e^2})} \approx 0.778$ and $U(1) = 1$. Furthermore, there exists a single-threshold algorithm ALG_θ satisfying

$$CR^\gamma(ALG_\theta) \geq 1 - (1 - \gamma)p_\gamma,$$

where p_γ is defined in Theorem 4.4.

To prove the upper bound, we used instances satisfying the condition of Proposition 3.3, guaranteeing that it remains true in the D_∞ -prophet inequality with $\gamma = \gamma_{\mathcal{D}}$. On the other hand, Theorem 2.1 ensures that the upper bound extends to the \mathcal{D} -prophet inequality, but with an additional $O(1/n^{1/3})$ term. The latter does change the result, as we considered instances of arbitrarily large size $n \rightarrow \infty$.

5 Conclusion

In this paper, we addressed the \mathcal{D} -prophet inequality problem, which models a very broad spectrum of online selection scenarios, accommodating various observation order models and allowing to revisit rejected candidates at a cost. The problem extends the classic prophet inequality, corresponding to the special case where all decay functions are zero. The main result of the paper is a reduction from the general \mathcal{D} -prophet inequality to the γ -prophet inequality, where all the decay functions equal to $x \mapsto \gamma x$ for some constant $\gamma \in [0, 1]$. Subsequently, we provide algorithms and upper bounds for the γ -prophet inequality, which remain valid, by the previous reduction, in the \mathcal{D} -prophet inequality. Notably, the proved upper and lower bounds match each other for the adversarial order model, hence completely solving the problem. Our analysis paves the way for more practical applications of prophet inequalities, and advances efforts towards closing the gap between theory and practice in online selection problems.

Limitations and future work

Better upper bounds in the D_∞ -prophet inequality. Proposition 3.3 establishes that upper bounds proved in the γ -prophet inequality using instances of random variables with support in some set $\{0, a, b\}$ remain true in the D_∞ -prophet inequality, hence in the \mathcal{D} -prophet inequality by Theorem 2.1. This is enough to establish a tight upper bound in the adversarial order model, but not in the random order and IID models. An interesting question to explore is if more general upper bounds can be extended, or not, from the γ - to the \mathcal{D} -prophet inequality.

Algorithms for the γ -prophet inequality. As explained in Section 4, our analysis of the competitive ratio of single-threshold algorithms relies on the identity (4), which is not satisfied for instance by multiple-threshold algorithms. In the adversarial order model, we proved that the optimal competitive ratio $1/(2 - \gamma)$ can be achieved with a single-threshold algorithm. However, this is not the case in the random order or IID models. An interesting research avenue is to study other classes of algorithms in the γ -prophet inequality.

Acknowledgements

This research was supported in part by the French National Research Agency (ANR) in the framework of the PEPR IA FOUNDRY project (ANR-23-PEIA-0003) and through the grant DOOM ANR-23-CE23-0002. It was also funded by the European Union (ERC, Ocean, 101071601). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

Dorian Baudry thanks the support of the French National Research Agency: ANR-19-CHIA-02 SCAI, ANR-22-SRSE-0009 Ocean, and ANR-23-CE23-0002 Doom. Dorian Baudry was partially funded by UK Research and Innovation (UKRI) under the UK government's Horizon Europe funding guarantee [grant number EP/Y028333/1].

References

- Antonios Antoniadis, Themis Gouleakis, Pieter Kleer, and Pavel Kolev. Secretary and online matching problems with machine learned advice. *Advances in Neural Information Processing Systems*, 33: 7933–7944, 2020.
- Antonios Antoniadis, Joan Boyar, Marek Eliás, Lene Monrad Favrholdt, Ruben Hoeksma, Kim S Larsen, Adam Polak, and Bertrand Simon. Paging with succinct predictions. In *International Conference on Machine Learning*, pages 952–968. PMLR, 2023.
- Makis Arsenis and Robert Kleinberg. Individual fairness in prophet inequalities. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 245–245, 2022.
- David Assaf and Ester Samuel-Cahn. Simple ratio prophet inequalities for a mortal with multiple choices. *Journal of applied probability*, 37(4):1084–1091, 2000.
- Alexia Atsidakou, Constantine Caramanis, Evangelia Gergatsouli, Orestis Papadigenopoulos, and Christos Tzamos. Contextual pandora's box. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 10944–10952, 2024.
- Ziyad Benomar and Christian Coester. Learning-augmented priority queues. *arXiv preprint arXiv:2406.04793*, 2024.
- Ziyad Benomar and Vianney Perchet. Advice querying under budget constraint for online algorithms. *Advances in Neural Information Processing Systems*, 36, 2024a.
- Ziyad Benomar and Vianney Perchet. Non-clairvoyant scheduling with partial predictions. In *Forty-first International Conference on Machine Learning*, 2024b.
- Ziyad Benomar, Evgenii Chzhen, Nicolas Schreuder, and Vianney Perchet. Addressing bias in online selection with limited budget of comparisons. *arXiv preprint arXiv:2303.09205*, 2023.

- Ben Berger, Tomer Ezra, Michal Feldman, and Federico Fusco. Pandora’s problem with deadlines. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 20337–20343, 2024.
- Hedyeh Beyhaghi, Negin Golrezaei, Renato Paes Leme, Martin Pál, and Balasubramanian Sivan. Improved revenue bounds for posted-price and second-price mechanisms. *Operations Research*, 69(6):1805–1822, 2021.
- Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge University Press, 2005.
- Brian Brubach, Nathaniel Grammel, Will Ma, and Aravind Srinivasan. Improved guarantees for offline stochastic matching via new ordered contention resolution schemes. *Advances in Neural Information Processing Systems*, 34:27184–27195, 2021.
- Archit Bubna and Ashish Chiplunkar. Prophet inequality: Order selection beats random order. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 302–336, 2023.
- 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*, pages 311–320, 2010.
- Justin Chen, Sandeep Silwal, Ali Vakilian, and Fred Zhang. Faster fundamental graph algorithms via learned predictions. In *International Conference on Machine Learning*, pages 3583–3602. PMLR, 2022.
- Jakub Chłkedowski, Adam Polak, Bartosz Szabucki, and Konrad Tomasz Żolna. Robust learning-augmented caching: An experimental study. In *International Conference on Machine Learning*, pages 1920–1930. PMLR, 2021.
- Nicolas Christianson, Junxuan Shen, and Adam Wierman. Optimal robustness-consistency tradeoffs for learning-augmented metrical task systems. In *International Conference on Artificial Intelligence and Statistics*, pages 9377–9399. PMLR, 2023.
- Alon Cohen, Avinatan Hassidim, Haim Kaplan, Yishay Mansour, and Shay Moran. Learning to screen. *Advances in Neural Information Processing Systems*, 32, 2019.
- Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Recent developments in prophet inequalities. *ACM SIGecom Exchanges*, 17(1):61–70, 2019.
- Jose Correa, Andres Cristi, Paul Duetting, and Ashkan Norouzi-Fard. Fairness and bias in online selection. In *International conference on machine learning*, pages 2112–2121. PMLR, 2021a.
- José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of operations research*, 46(4):1452–1478, 2021b.
- Jose Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. *Mathematical Programming*, 190(1-2):483–521, 2021c.
- Yuan Deng, Vahab Mirrokni, and Hanrui Zhang. Posted pricing and dynamic prior-independent mechanisms with value maximizers. *Advances in Neural Information Processing Systems*, 35: 24158–24169, 2022.
- Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, Ali Vakilian, and Nikos Zarifis. Learning online algorithms with distributional advice. In *International Conference on Machine Learning*, pages 2687–2696. PMLR, 2021.
- Michael Dinitz, Sungjin Im, Thomas Lavastida, Benjamin Moseley, and Sergei Vassilvitskii. Faster matchings via learned duals. *Advances in neural information processing systems*, 34:10393–10406, 2021.
- Paul Dütting, Thomas Kesselheim, and Brendan Lucier. An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. *ACM SIGecom Exchanges*, 18(2):32–37, 2020.

- Paul Dütting, Silvio Lattanzi, Renato Paes Leme, and Sergei Vassilvitskii. Secretaries with advice. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 409–429, 2021.
- Evgenii Borisovich Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics*, 4:627–629, 1963.
- Farbod Ekbatani, Yiding Feng, and Rad Niazadeh. Online matching with cancellation costs. *arXiv preprint arXiv:2210.11570*, 2022.
- Farbod Ekbatani, Yiding Feng, and Rad Niazadeh. Online resource allocation with buyback: Optimal algorithms via primal-dual. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 583–583, 2023.
- Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. Prophet inequalities with cancellation costs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1247–1258, 2024.
- Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- Hossein Esfandiari, MohammadTaghi HajiAghayi, Brendan Lucier, and Michael Mitzenmacher. Online pandora’s boxes and bandits. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 1885–1892, 2019.
- Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Online stochastic max-weight matching: prophet inequality for vertex and edge arrival models. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 769–787, 2020.
- Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.
- Evangelia Gergatsouli and Christos Tzamos. Online learning for min sum set cover and pandora’s box. In *International Conference on Machine Learning*, pages 7382–7403. PMLR, 2022.
- Evangelia Gergatsouli and Christos Tzamos. Weitzman’s rule for pandora’s box with correlations. *Advances in Neural Information Processing Systems*, 36, 2024.
- Giordano Giambartolomei, Frederik Mallmann-Trenn, and Raimundo Saona. Prophet inequalities: Separating random order from order selection. *arXiv preprint arXiv:2304.04024*, 2023.
- Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, pages 58–65, 2007.
- Elfarouk Harb. New prophet inequalities via poissonization and sharding, 2024.
- Theodore P Hill and Robert P Kertz. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability*, pages 336–345, 1982.
- Zhihao Jiang, Pinyan Lu, Zhihao Gavin Tang, and Yuhao Zhang. Online selection problems against constrained adversary. In *International Conference on Machine Learning*, pages 5002–5012. PMLR, 2021.
- Robert P Kertz. Stop rule and supremum expectations of iid random variables: a complete comparison by conjugate duality. *Journal of multivariate analysis*, 19(1):88–112, 1986.
- Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.
- Robert Kleinberg, Bo Waggoner, and E Glen Weyl. Descending price optimally coordinates search. *arXiv preprint arXiv:1603.07682*, 2016.
- Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. The case for learned index structures. In *Proceedings of the 2018 international conference on management of data*, pages 489–504, 2018.

- Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. 1977.
- Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.
- Alexandra Anna Lassota, Alexander Lindermayr, Nicole Megow, and Jens Schlöter. Minimalistic predictions to schedule jobs with online precedence constraints. In *International Conference on Machine Learning*, pages 18563–18583. PMLR, 2023.
- Bo Li, Xiaowei Wu, and Yutong Wu. Query efficient prophet inequality with unknown iid distributions. *arXiv preprint arXiv:2205.05519*, 2022.
- Honghao Lin, Tian Luo, and David Woodruff. Learning augmented binary search trees. In *International Conference on Machine Learning*, pages 13431–13440. PMLR, 2022.
- Denis V Lindley. Dynamic programming and decision theory. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 10(1):39–51, 1961.
- Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.
- Thodoris Lykouris and Sergei Vassilvtiskii. Competitive caching with machine learned advice. In *International Conference on Machine Learning*, pages 3296–3305. PMLR, 2018.
- Anuran Makur, Marios Mertzaniadis, Alexandros Psomas, and Athina Terzoglou. On the robustness of mechanism design under total variation distance. *Advances in Neural Information Processing Systems*, 36, 2024.
- Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 763–764, 2021.
- Bo Peng and Zhihao Gavin Tang. Order selection prophet inequality: From threshold optimization to arrival time design. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 171–178. IEEE, 2022.
- Alexandros Psomas, Ariel Schwartzman Cohenca, and S Weinberg. On infinite separations between simple and optimal mechanisms. *Advances in Neural Information Processing Systems*, 35:4818–4829, 2022.
- Manish Purohit, Zoya Svitkina, and Ravi Kumar. Improving online algorithms via ml predictions. *Advances in Neural Information Processing Systems*, 31, 2018.
- Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.
- Sean R Sinclair, Felipe Vieira Frujeri, Ching-An Cheng, Luke Marshall, Hugo De Oliveira Barbalho, Jingling Li, Jennifer Neville, Ishai Menache, and Adith Swaminathan. Hindsight learning for mdps with exogenous inputs. In *International Conference on Machine Learning*, pages 31877–31914. PMLR, 2023.
- Bo Sun, Russell Lee, Mohammad Hajiesmaili, Adam Wierman, and Danny Tsang. Pareto-optimal learning-augmented algorithms for online conversion problems. *Advances in Neural Information Processing Systems*, 34:10339–10350, 2021.
- Vasilis Syrgkanis. A sample complexity measure with applications to learning optimal auctions. *Advances in Neural Information Processing Systems*, 30, 2017.
- Martin Weitzman. *Optimal search for the best alternative*, volume 78. Department of Energy, 1978.

A From \mathcal{D} -prophet to the D_∞ -prophet inequality

In this section, we prove the reduction from the \mathcal{D} -prophet to the D_∞ -prophet inequality problem in the adversarial and random order models, and the reduction up to an additional $O(n^{-1/3})$ term in the IID model. First, we prove Corollary 2.1.1, which is the principal implication of Theorem 2.1.

A.1 Proof of Corollary 2.1.1

Proof. Let us denote $A_{*,\infty}$ the algorithm taking optimal decisions for any instance in the D_∞ -prophet inequality (obtained via dynamic programming). Then, by Theorem 2.1 we obtain for the adversarial and random order models that

$$\sup_{\text{ALG}} \text{CR}^{\mathcal{D}}(\text{ALG}) \leq \inf_I \sup_A \frac{\mathbb{E}[A^{D_\infty}(I)]}{\mathbb{E}[\text{OPT}(I)]} = \inf_I \frac{\mathbb{E}[A_{*,\infty}^{D_\infty}(I)]}{\mathbb{E}[\text{OPT}(I)]} = \text{CR}^{D_\infty}(A_{*,\infty}) = \sup_A \text{CR}^{D_\infty}(A). \quad (5)$$

Since $\text{CR}^{\mathcal{D}}(\text{ALG}) \geq \text{CR}^{D_\infty}(\text{ALG})$ for any algorithm, we deduce that (5) is an equality. If we consider now any algorithm \bar{A}_∞ that is optimal for the D_∞ -prophet inequality, not necessarily $A_{*,\infty}$, then Equation (5) provides

$$\text{CR}^{\mathcal{D}}(\bar{A}_\infty) \geq \text{CR}^{D_\infty}(\bar{A}_\infty) = \sup_{\text{ALG}} \text{CR}^{\mathcal{D}}(\text{ALG}).$$

The previous inequality is again an equality, and it implies that \bar{A}_∞ is also optimal, in the sense of the competitive ratio, for the \mathcal{D} -prophet inequality, and

$$\text{CR}^{\mathcal{D}}(\bar{A}_\infty) = \text{CR}^{D_\infty}(\bar{A}_\infty).$$

□

A.2 Auxiliary Lemma

The efficiency of the proof scheme introduced in Section 2.1 relies on the following key argument: if $(D_j)_{j \geq 1}$ converges pointwise to D_∞ , then for any algorithm A and any instance I , the output of A when all the decay functions are equal to D_m converges to its output when all the decay functions are equal to D_∞ . If X_1, \dots, X_n are the realizations of I observed by A and if σ is the order in which they are observed, then denoting τ the stopping time of A we can write that

$$\begin{aligned} & \mathbb{E}[A^{D_m}(I)] - \mathbb{E}[A^{D_\infty}(I)] \\ &= \mathbb{E}[\max\{X_{\sigma(\tau)}, D_m(\max_{i < \tau} X_{\sigma(i)})\}] - \mathbb{E}[\max\{X_{\sigma(\tau)}, D_\infty(\max_{i < \tau} X_{\sigma(i)})\}] \\ &\leq \max_{\substack{\pi \in \mathcal{S}_n \\ q \in [n]}} \left\{ \mathbb{E}[\max\{X_{\pi(q)}, D_m(\max_{i < q} X_{\pi(i)})\}] - \mathbb{E}[\max\{X_{\pi(q)}, D_\infty(\max_{i < q} X_{\pi(i)})\}] \right\}, \end{aligned}$$

where \mathcal{S}_n is the set of all permutations of $[n]$. The latter upper bound is independent of σ and A . We show in the following lemma that it converges to 0 when $m \rightarrow \infty$.

Lemma A.1. *Let \mathcal{S}_n be the set of all permutations of $[n]$. For any fixed instance $I = (F_1, \dots, F_n)$, considering $X_i \sim F_i$ for all $i \in [n]$, define for all $m \geq 1$*

$$\epsilon_m(I) = \max_{\substack{\pi \in \mathcal{S}_n \\ q \in [n]}} \left\{ \mathbb{E}[\max\{X_{\pi(q)}, D_m(\max_{i < q} X_{\pi(i)})\}] - \mathbb{E}[\max\{X_{\pi(q)}, D_\infty(\max_{i < q} X_{\pi(i)})\}] \right\}.$$

If $\mathbb{E}[\max_{i \in [n]} X_i] < \infty$, then $\lim_{m \rightarrow \infty} \epsilon_m(I) = 0$.

Proof. Let us denote f_1, \dots, f_n the respective probability density functions of X_1, \dots, X_n . For any $q \in [n]$ and $\pi \in \mathcal{S}_n$, let us define for all $m \geq 0$ the function $\varphi_m^{\pi,q} : [0, \infty)^q \rightarrow [0, \infty)$ by $\varphi_m^{\pi,q}(x_1, \dots, x_q) = \max\{x_{\pi(q)}, D_m(\max_{i < q} x_{\pi(i)})\} - \max\{x_{\pi(q)}, D_\infty(\max_{i < q} x_{\pi(i)})\}$. $\varphi_m^{\pi,q}$ is positive because $D_m \geq D_\infty$. The sequence $(\varphi_m^{\pi,q})_m$ is non-increasing, converges to 0 pointwise, and is dominated by $(x_1, \dots, x_q) \mapsto \max_{i \in [q]} x_i$, which is integrable with respect to the probability

measure $(x_1, \dots, x_q) \mapsto \prod_{i=1}^q f(x_i)$. Therefore, using the dominated convergence theorem, we deduce that $\lim_{m \rightarrow \infty} \mathbb{E}[\varphi_m^{\pi, q}(X_1, \dots, X_q)] = 0$. It follows that

$$\lim_{m \rightarrow \infty} \epsilon_m(I) = \lim_{m \rightarrow \infty} \left(\max_{\substack{\pi \in \mathcal{S}_n \\ q \in [n]}} \mathbb{E}[\varphi_m^{\pi, q}(X_1, \dots, X_q)] \right) = 0.$$

□

A.3 Proof of Theorem 2.1

Proof of Theorem 2.1. We provide a separate proof for each of the adversarial order, random order and IID models.

Adversarial order Let $I = (F_1, \dots, F_n)$ be any instance and $X_i \sim F_i$ for all $i \in [n]$. Consider the instance $I_m = (Y_1, \dots, Y_{mn})$, where $Y_{km} \sim F_k$ for any $k \in [n]$ and $Y_i = 0$ a.s. for all $i \notin \{m, 2m, \dots, mn\}$. It is clear that no reasonable algorithm would stop at a zero value: if the current observation is 0 it is preferable to wait for a non-null value, or it would have been preferable to stop at the previous non-null value. Hence, τ is a multiple of m : $\tau = \rho m$ for some $\rho \in \mathbb{N}^*$. Given that $D_j(0) = 0$ for all j and the sequence $(D_j)_j$ is non-increasing, we have that

$$\begin{aligned} \mathbb{E}[\text{ALG}^D(I_m)] &= \mathbb{E}[\max_{i \leq \tau} D_{\tau-i}(Y_i)] \\ &= \mathbb{E}[\max_{k \leq \rho} D_{\tau-km}(Y_{km})] \\ &= \mathbb{E}[\max_{k \leq \rho} D_{\rho m - km}(X_k)] \\ &= \mathbb{E}[\max\{X_\rho, \max_{k < \rho} D_{(\rho-k)m}(X_k)\}] \\ &\leq \mathbb{E}[\max\{X_\rho, \max_{k < \rho} D_m(X_k)\}] \\ &\leq \mathbb{E}[\max\{X_\rho, \max_{k < \rho} D_\infty(X_k)\}] + \epsilon_m(I), \end{aligned}$$

where $\epsilon_m(I)$ is defined in Lemma A.1. We can then use that the first right-hand term is the output of some other algorithm that would choose a stopping time ρ when facing I in the context of the D_∞ -prophet inequality. More precisely, consider the algorithm A_m which, given any instance $I = (F_1, \dots, F_n)$, simulates the behavior of ALG facing the sequence I_m , where at each step $i \in [mn]$

- if $i \notin \{m, \dots, nm\}$: ALG observes $Y_i = 0$,
- otherwise, if $i = km$ for some $k \in [n]$: A_m observes X_k and ALG observes $Y_{km} = X_k$
- if ALG stops on $Y_{\rho m}$, then A_m also stops, and its reward is X_ρ .

$A_m(X_1, \dots, X_n)$ stops at the same value as $\text{ALG}(Y_1, \dots, Y_m)$, their reward in the D_∞ -prophet inequality is the same, and since $\max_{i \in [n]} X_i = \max_{i \in [mn]} Y_i$ this yields to

$$\text{CR}^D(\text{ALG}) \leq \frac{\mathbb{E}[\text{ALG}^D(I_m)]}{\mathbb{E}[\text{OPT}(I_m)]} \leq \frac{\mathbb{E}[A_m^{D_\infty}(I)] + \epsilon_m(I)}{\mathbb{E}[\text{OPT}(I)]} \leq \sup_{A: \text{ algo}} \frac{\mathbb{E}[A^{D_\infty}(I)]}{\mathbb{E}[\text{OPT}(I)]} + \frac{\epsilon_m(I)}{\mathbb{E}[\text{OPT}(I)]},$$

and taking the limit when $m \rightarrow \infty$ gives the result, by making the second term vanish.

Random order Let $I = (F_1, \dots, F_n)$ be an instance of distributions and $X_i \sim F_i$ for $i \in [n]$. Using the notation δ_0 for the Dirac distribution in 0, we consider $I_m = (F_1, \dots, F_n, \delta_0, \dots, \delta_0)$ containing m copies of δ_0 so that the observations from this instance always contain at least m null values. Let Y_1, \dots, Y_m be a realization of this instance. For simplicity, say that $Y_i = X_i$ for $i \in [n]$, and $Y_i = 0$ for $i > n$.

We first show that when $m \rightarrow \infty$, since the observation order is drawn uniformly at random, the algorithm observes a large number of zeros between every two random variables drawn from (F_1, \dots, F_n) . Let us denote by π the uniformly random order in which the observations are received,

i.e. the algorithms observes $Y_{\pi(1)}, Y_{\pi(2)}, \dots$, and let $\ell \geq 1$ be some positive integer, and t_1, \dots, t_n be the increasing indices in which the variables Y_1, \dots, Y_n are observed, i.e. $t_1 < \dots < t_n$ and $\{t_1, \dots, t_n\} = \{\pi^{-1}(1), \dots, \pi^{-1}(n)\}$. Therefore, any observation outside $\{Y_{\pi(t_1)}, \dots, Y_{\pi(t_n)}\}$ is zero. Using the notation $L = \min_{i \in [n-1]} |t_{i+1} - t_i|$, we obtain that

$$\begin{aligned} \Pr(L \leq \ell) &= \Pr(\cup_{i=1}^{n-1} \{t_{i+1} - t_i \leq \ell\}) \\ &= \Pr(\cup_{k=1}^n \cup_{j=1}^{k-1} \{|\pi^{-1}(k) - \pi^{-1}(j)| \leq \ell\}) \\ &\leq \frac{n(n-1)}{2} \Pr(|\pi^{-1}(1) - \pi^{-1}(2)| \leq \ell) \\ &= \frac{n(n-1)}{2} \Pr((\cup_{k=1}^{n+m} (\pi^{-1}(1) = k, \pi^{-1}(2) \in \{k-\ell, \dots, k+\ell\} \setminus \{k\})) \\ &\leq \frac{n(n-1)}{2} \times (n+m) \times \frac{1}{n+m} \times \frac{2\ell}{n+m-1} \\ &\leq \frac{n^2\ell}{m}. \end{aligned}$$

Taking $\ell = \sqrt{m}$, we find that $\Pr(L \leq \ell) \leq n^2/\sqrt{m}$. Therefore, for any algorithm **ALG**, observing that the reward of **ALG**^D is at most $\max_{i \in [n]} X_i$ a.s., and by independence of $\max_{i \in [n]} X_i$ and L , we deduce that

$$\begin{aligned} \mathbb{E}[\mathbf{ALG}^D(I_m)] &= \mathbb{E}[\mathbf{ALG}^D(I_m) \mathbb{1}_{L > \ell}] + \mathbb{E}[\mathbf{ALG}^D(I_m) \mathbb{1}_{L \leq \ell}] \\ &\leq \mathbb{E}[\mathbf{ALG}^D(I_m) \mid L > \ell] + \mathbb{E}[(\max_{i \in [n]} X_i) \mathbb{1}_{L \leq \ell}] \\ &\leq \mathbb{E}[\mathbf{ALG}^D(I_m) \mid L > \ell] + \mathbb{E}[\max_{i \in [n]} X_i] \frac{n^2}{\sqrt{m}}. \end{aligned} \tag{6}$$

Let us denote τ the stopping time of **ALG** and $t_\rho = \max_{j \in [n]} \{t_j : t_j \leq \tau\}$ the last time when a variable $(X_j)_{j \in [n]}$ was observed by **ALG**. The sequence of functions $(D_j)_{j \geq 1}$ is non-increasing, hence

$$\begin{aligned} \mathbb{E}[\mathbf{ALG}^D(I_m) \mid L > \ell] &= \mathbb{E}[\max_{i \in [\tau]} D_{\tau-i}(Y_{\pi(i)}) \mid L > \ell] \\ &= \mathbb{E}[\max_{j \leq i} D_{\tau-t_j}(Y_{\pi(t_j)}) \mid L > \ell] \end{aligned} \tag{7}$$

$$\leq \mathbb{E}[\max_{j \leq i} D_{t_\rho-t_j}(Y_{\pi(t_j)}) \mid L > \ell] \tag{8}$$

$$\begin{aligned} &= \mathbb{E}[\max \{Y_{\pi(t_\rho)}, \max_{j < \rho} D_{t_\rho-t_j}(Y_{\pi(t_j)})\} \mid L > \ell] \\ &\leq \mathbb{E}[\max \{Y_{\pi(t_\rho)}, \max_{j < \rho} D_\ell(Y_{\pi(t_j)})\}], \end{aligned} \tag{9}$$

Equation (7) holds because the only non-zero values up to step τ are $(Y_{\pi(t_j)})_{j \in [\rho]}$. Inequality (8) uses that the sequence $(D_j)_{j \geq 1}$ is non-increasing, and (9) uses, in addition to that, the independence of L and $(Y_{\pi(t_j)})_{j \in [n]}$. We now argue that the term $\mathbb{E}[\max \{Y_{\pi(t_\rho)}, \max_{j < \rho} D_\ell(Y_{\pi(t_j)})\}]$ is the expected reward of an algorithm in the D_ℓ -inequality. Given that π is a uniform random permutation of $[n+m]$ and by definition of t_1, \dots, t_n , the application $\sigma : k \in [n] \mapsto \pi(t_k)$ is a random permutation of $[n]$. Therefore we consider the algorithm **A**_m that receives as input the instance $I = (F_1, \dots, F_n)$, then considers the array $u = (1, \dots, 1, 0, \dots, 0)$ composed of n values equal to 1 and m zero values, and a uniformly random permutation π of $[n+m]$, then simulates **ALG**^D(I_m) as follows: at each step $j \in [n+m]$

- if $u_{\pi(j)} = 0$, then **ALG** observes the value $Y_{\pi(j)} = 0$,
- if $u_{\pi(j)} = 1$, then **A**_m observes the next value $X_{\sigma(k)}$, and **ALG** observes $Y_{\pi(j)} = X_{\sigma(k)}$,
- when **ALG** decides to stop, **A**_m also stops, and its reward is the current value $X_{\sigma(k)}$.

With this construction, $(Y_j)_{j \in [n+m]}$ is indeed a realization of the instance I_m , and **A**_m stops on the last value sampled from F_1, \dots, F_n observed by **ALG**. Therefore, denoting ρ the stopping time of

\mathbf{A}_m , and $\epsilon_\ell(I)$ as defined in Lemma A.1, we have

$$\begin{aligned} \mathbb{E}[\mathbf{ALG}^{\mathcal{D}}(I_m) \mid L > \ell] &\leq \mathbb{E}[\max \{Y_{\pi(t_\rho)}, \max_{j < \rho} D_\ell(Y_{\pi(t_j)})\}] \\ &= \mathbb{E}[\max \{X_{\sigma(\rho)}, \max_{j < \rho} D_\ell(X_{\sigma(t_j)})\}] \\ &= \mathbb{E}[\mathbf{A}_m^{D_\ell}(I)] \\ &\leq \mathbb{E}[\mathbf{A}_m^{D_\infty}(I)] + \epsilon_\ell(I) \\ &\leq \sup_{\mathbf{A}: \text{algo}} \mathbb{E}[\mathbf{A}^{D_\infty}(I)] + \epsilon_\ell(I). \end{aligned}$$

Taking $\ell = \sqrt{m}$ and substituting into Equation (6), then observing that $\mathbb{E}[\mathbf{OPT}(I)] = \mathbb{E}[\mathbf{OPT}(I_m)]$, gives that

$$\text{CR}^{\mathcal{D}}(\mathbf{ALG}) \leq \frac{\mathbb{E}[\mathbf{ALG}^{\mathcal{D}}(I_m)]}{\mathbb{E}[\mathbf{OPT}(I_m)]} \leq \sup_{\mathbf{A}: \text{algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(I)]}{\mathbb{E}[\mathbf{OPT}(I)]} + \frac{\epsilon_{\sqrt{m}}(I)}{\mathbb{E}[\mathbf{OPT}(I)]} + \frac{n^2}{\sqrt{m}}.$$

Finally, taking $m \rightarrow \infty$ and using Lemma A.1, we deduce that

$$\text{CR}^{\mathcal{D}}(\mathbf{ALG}) \leq \sup_{\mathbf{A}: \text{algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(I)]}{\mathbb{E}[\mathbf{OPT}(I)]},$$

which completes the proof for the random order.

IID random variables For any probability distribution F on $[0, \infty)$ and for any $n \geq 1$ we denote $\mathbb{E}[\mathbf{OPT}(F, n)]$ the expected maximum of n independent random variables drawn from F , and for any algorithm \mathbf{ALG} we denote $\mathbb{E}[\mathbf{ALG}^{\mathcal{D}}(F, n)]$ its expected output when given n IID variable sampled from F as input. The proof of Theorem 2.1 for this last model is much more technical than for previous models, so we first prove several auxiliary results that we will later use to provide a concise proof of the last part of the theorem.

Lemma A.2. For any probability distribution F and $n \geq 1, \Delta \geq 0$, we have

$$\mathbb{E}[\mathbf{OPT}(F, n + \Delta)] \leq \left(1 + \frac{\Delta}{n}\right) \mathbb{E}[\mathbf{OPT}(F, n)].$$

Proof. We first write

$$\begin{aligned} \Pr(\mathbf{OPT}(F, n + \Delta) > x) &= 1 - F(x)^{n+\Delta} \\ &= \left(1 + F(x)^n \frac{1 - F(x)^\Delta}{1 - F(x)^n}\right) (1 - F(x)^n) \\ &= \left(1 + F(x)^n \frac{1 - F(x)^\Delta}{1 - F(x)^n}\right) \Pr(\mathbf{OPT}(F, n) > x), \end{aligned}$$

and then use that

$$F(x)^\Delta = e^{\Delta \log(F(x))} \geq 1 + \Delta \log(F(x)) = 1 - \frac{\Delta}{n} \log(1/F(x)^n) \geq 1 - \frac{\Delta}{n} \left(\frac{1 - F(x)^n}{F(x)^n}\right),$$

so we directly obtain

$$F(x)^n \frac{1 - F(x)^\Delta}{1 - F(x)^n} \leq \frac{\Delta}{n},$$

which gives that

$$\Pr(\mathbf{OPT}(F, n + \Delta) > x) \leq \left(1 + \frac{\Delta}{n}\right) \Pr(\mathbf{OPT}(F, n) > x).$$

As we consider non-negative random variables, it follows directly that

$$\mathbb{E}[\mathbf{OPT}(F, n + \Delta)] \leq \left(1 + \frac{\Delta}{n}\right) \mathbb{E}[\mathbf{OPT}(F, n)].$$

□

Lemma A.3. Let $N \sim \mathcal{B}(m, \varepsilon)$ and let $n := \mathbb{E}[N] = \varepsilon m$, then we have

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq n+n^{2/3}}] &\leq \frac{6}{n^{1/3}} \mathbb{E}[\text{OPT}(F, n+n^{2/3})], \\ \mathbb{E}[\text{OPT}(F, N)] &\leq \left(1 + \frac{3}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, n+n^{2/3})]. \end{aligned}$$

Proof. Let $\Delta, s > 0$ such that $\Delta \leq s$. For any $k \geq 1$ let $W_k = [s + (k-1)\Delta, s + k\Delta)$. $(W_k)_{k \geq 1}$ is a partition of $[s, \infty)$, thus we have

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq s}] &= \sum_{k=1}^{\infty} \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \in W_k}] \\ &\leq \sum_{k=1}^{\infty} \mathbb{E}[\text{OPT}(F, s+k\Delta) \mathbb{1}_{N \in W_k}] \\ &= \sum_{k=1}^{\infty} \mathbb{E}[\text{OPT}(F, s+k\Delta)] \Pr(N \in W_k) \\ &\leq \sum_{k=1}^{\infty} \left(1 + \frac{k\Delta}{s}\right) \mathbb{E}[\text{OPT}(F, s)] \Pr(N \in W_k) \\ &= \left(\Pr(N \geq s) + \frac{\Delta}{s} \sum_{k=1}^{\infty} k \Pr(N \in W_k)\right) \mathbb{E}[\text{OPT}(F, s)], \end{aligned}$$

where we used Lemma A.2 in the penultimate inequality. Furthermore, observing that

$$\sum_{k=1}^{\infty} k \Pr(N \in W_k) = \sum_{k=1}^{\infty} \sum_{\ell=0}^{k-1} \Pr(N \in W_k) = \sum_{\ell=0}^{\infty} \sum_{k=\ell+1}^{\infty} \Pr(N \in W_k) = \sum_{\ell=0}^{\infty} \Pr(N \geq s + \ell\Delta),$$

we obtain, given $\Delta \leq s$, that

$$\mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq s}] \leq \left(\Pr(N \geq s) + \frac{\Delta}{s} \sum_{k=0}^{\infty} \Pr(N \geq s + k\Delta)\right) \mathbb{E}[\text{OPT}(F, s)] \quad (10)$$

$$\leq \left(2 \sum_{k=0}^{\infty} \Pr(N \geq s + k\Delta)\right). \quad (11)$$

N is a Binomial random variable with expectation n . Therefore, Chernoff's inequality gives for any $\delta \geq 0$ that

$$\Pr(N \geq (1 + \delta)n) \leq \exp\left(-\frac{\delta^2 n}{2 + \delta}\right) \leq \exp\left(-\frac{\min(\delta, \delta^2)n}{3}\right),$$

where the second inequality can be derived by treating separately $\delta < 1$ and $\delta \geq 1$. In particular, for any $k \geq 1$, taking $\delta = \frac{k\Delta}{n}$ such that $\Delta \leq n$ yields

$$\Pr(N \geq n + k\Delta) \leq \exp\left(-\frac{\min(k\Delta, k^2\Delta^2/n)}{3}\right) \leq \exp\left(-\frac{k \min(\Delta, \Delta^2/n)}{3}\right) = \exp\left(-\frac{k\Delta^2}{3n}\right).$$

Substituting this Inequality into (11) with $s = n + \Delta$, we obtain

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq n+\Delta}] &\leq \left(2 \sum_{k=1}^{\infty} \Pr(N \geq n + k\Delta)\right) \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &\leq \left(2 \sum_{k=1}^{\infty} \exp\left(-\frac{k\Delta^2}{3n}\right)\right) \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &= \frac{2}{\exp\left(\frac{\Delta^2}{3n}\right) - 1} \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &\leq \frac{6n}{\Delta^2} \mathbb{E}[\text{OPT}(F, n + \Delta)], \end{aligned}$$

and taking $\Delta = n^{2/3}$ proves the first inequality of the lemma.

Let us move now to the second inequality. We have

$$\mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N < s}] \leq \mathbb{E}[\text{OPT}(F, s) \mathbb{1}_{N < s}] = \mathbb{E}[\text{OPT}(F, s)] \Pr(N < s),$$

and thus, using Inequality (10), again with $s = n + \Delta$ and $\Delta = n^{2/3}$, it follows that

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, N)] &= \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N < s}] + \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq s}] \\ &\leq \left(1 + \frac{\Delta}{s} \sum_{k=0}^{\infty} \Pr(N \geq s + k\Delta)\right) \mathbb{E}[\text{OPT}(F, s)] \\ &\leq \left(1 + \sum_{k=1}^{\infty} \Pr(N \geq n + k\Delta)\right) \mathbb{E}[\text{OPT}(F, s)] \\ &\leq \left(1 + \frac{3}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, n + \Delta)]. \end{aligned}$$

□

Lemma A.4. Let $\delta_1, \dots, \delta_m \stackrel{iid}{\sim} \mathcal{B}(\varepsilon)$, and $N = \sum_{i=1}^m \delta_i$. Denoting by $n = \mathbb{E}[N] = \varepsilon m$, if $n \geq 4$ then

$$\mathbb{E}[N^2 \text{OPT}(F, N)] \leq \left(1 + \frac{8}{n^3}\right) n^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})].$$

Proof. For all $k \in [m]$, denote by $N_k = \sum_{i=k}^m \delta_i$. We have that

$$\begin{aligned} N^2 \text{OPT}(F, N) &= \left(\sum_{i=1}^m \delta_i\right)^2 \text{OPT}(F, N) \\ &= \left(\sum_{i=1}^m \delta_i^2 + 2 \sum_{i < j} \delta_i \delta_j\right) \text{OPT}(F, N), \end{aligned}$$

and observing that $\delta_i^2 = \delta_i$ for all i , we obtain in expectation

$$\begin{aligned} \mathbb{E}[N^2 \text{OPT}(F, N)] &= m \mathbb{E}[\delta_1 \text{OPT}(F, \delta_1 + N_2)] + m(m-1) \mathbb{E}[\delta_1 \delta_2 \text{OPT}(\delta_1 + \delta_2 + N_3)] \\ &\leq m \mathbb{E}[\delta_1 \text{OPT}(F, 1 + N_2)] + m(m-1) \mathbb{E}[\delta_1 \delta_2 \text{OPT}(2 + N_3)] \\ &= m \varepsilon \mathbb{E}[\text{OPT}(F, 1 + N_2)] + m(m-1) \varepsilon^2 \mathbb{E}[\text{OPT}(2 + N_3)] \\ &= m \varepsilon \mathbb{E}[\text{OPT}(F, 1 + N_2)] + m^2 \varepsilon^2 \mathbb{E}[\text{OPT}(2 + N_3)]. \end{aligned} \tag{12}$$

For $j \in \{1, 2\}$, the proof of Lemma A.3 can be easily adjusted to prove an upper bound on $\mathbb{E}[\text{OPT}(F, j + N_{j+1})]$, by first bounding $\mathbb{E}[\text{OPT}(F, j + N_{j+1}) \mathbb{1}_{N_{j+1} \geq s}]$ then $\mathbb{E}[\text{OPT}(F, j + N_{j+1}) \mathbb{1}_{N_{j+1} < s}]$. The concentration arguments remain the same, replacing m by $m - j$. The expectation of N_{j+1} is $\varepsilon(m - j) = n - \varepsilon j$, hence we obtain

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, j + N_{j+1})] &\leq \left(1 + \frac{3}{(n - \varepsilon j)^{1/3}}\right) \mathbb{E}[\text{OPT}(F, j + (n - \varepsilon j) + (n - \varepsilon j)^{2/3})] \\ &= \left(1 + \frac{3}{(n - \varepsilon j)^{1/3}}\right) \mathbb{E}[\text{OPT}(F, j + n + n^{2/3})] \\ &\leq \left(1 + \frac{3}{(n - 2)^{1/3}}\right) \mathbb{E}[\text{OPT}(F, 2 + n + n^{2/3})] \\ &\leq \left(1 + \frac{4}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, 2 + n + n^{2/3})], \end{aligned}$$

where we used respectively in the last inequalities that $j \leq 2$ and $n \geq 4$. Furthermore, Lemma A.2 gives that

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, j + N_{j+1})] &\leq \left(1 + \frac{4}{n^{1/3}}\right) \left(1 + \frac{2}{n + n^{2/3}}\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &\leq \left(1 + \frac{6}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})], \end{aligned}$$

where the last inequality is true for $n \geq 4$. Finally, substituting into 12 yields

$$\begin{aligned} \mathbb{E}[N^2 \text{OPT}(F, N)] &\leq (m^2 \varepsilon^2 + m\varepsilon) \left(1 + \frac{6}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &= (n^2 + n) \left(1 + \frac{6}{n^{1/3}}\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &= \left(1 + \frac{1}{n}\right) \left(1 + \frac{6}{n^{1/3}}\right) n^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &\leq \left(1 + \frac{8}{n^{1/3}}\right) n^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})]. \end{aligned}$$

□

Lemma A.5. Let $\delta_1, \dots, \delta_m \stackrel{iid}{\sim} \mathcal{B}(\varepsilon)$, $N = \sum_{i=1}^m \delta_i$, $n = \mathbb{E}[N] = \varepsilon m$ and $L = \min_{i \neq j} \{|i - j| : \delta_i = 1, \delta_j = 1\}$, then for any $\ell \geq 0$ we have

$$\mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{L \leq \ell}] \leq 7m\ell\varepsilon^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})].$$

Proof. The random variables N and L are not independent, thus we need to adequately compute the distribution of L conditional to N . For any $\ell \geq 0$ and $k \geq 2$ we have

$$\begin{aligned} \Pr(L \leq \ell, N = s) &= \Pr\left(L \leq \ell, \sum_{i=1}^m \delta_k = s\right) \\ &= \Pr\left(\cup_{i=1}^m \cup_{j=\max(1, i-\ell)}^{i-1} (\delta_i = \delta_j = 1, \sum_{i=1}^m \delta_k = s)\right) \\ &\leq m\ell \Pr\left(\delta_1 = \delta_2 = 1, \sum_{i=3}^m \delta_k = s - 2\right) \\ &= m\ell \binom{m-2}{s-2} \varepsilon^s (1 - \varepsilon)^{m-s+2}, \end{aligned}$$

therefore

$$\begin{aligned} \Pr(L \leq \ell \mid N = s) &= \frac{\Pr(L \leq \ell, N = s)}{\Pr(N = s)} \\ &\leq \frac{m\ell \binom{m-2}{s-2} \varepsilon^s (1 - \varepsilon)^{m-s+2}}{\binom{m}{s} \varepsilon^s (1 - \varepsilon)^{m-s}} \\ &\leq m\ell \frac{\binom{m-2}{s-2}}{\binom{m}{s}} \\ &= m\ell \frac{s(s-1)}{m(m-1)} \\ &\leq \frac{\ell s^2}{m}. \end{aligned}$$

Using this inequality and Lemma A.4, we deduce that

$$\begin{aligned} \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{L \leq \ell}] &= \mathbb{E}[\text{OPT}(F, N) \Pr(L \leq \ell \mid N, \text{OPT}(F, N))] \\ &= \mathbb{E}[\text{OPT}(F, N) \Pr(L \leq \ell \mid N)] \\ &\leq \frac{\ell}{m} \mathbb{E}[N^2 \text{OPT}(F, N)] \\ &\leq \left(1 + \frac{8}{n^{1/3}}\right) \frac{\ell n^2}{m} \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &= \left(1 + \frac{8}{n^{1/3}}\right) m\ell\varepsilon^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})] \\ &= 7m\ell\varepsilon^2 \mathbb{E}[\text{OPT}(F, n + n^{2/3})]. \end{aligned}$$

where we used that $1 + \frac{8}{n^{1/3}} \leq 7$ for $n \geq 4$.

□

Using the previous lemmas, we can now prove the theorem. Let $m > n \geq 1$, $\Delta = n^{2/3}$, $\varepsilon = n/m$ and let Q be the probability distribution of a random variable that is equal to 0 with probability $1 - \varepsilon$, and drawn from F with probability ε .

Let us consider m i.i.d. variables $Y_1, \dots, Y_m \sim Q$, and for each $i \in [m]$ we denote by δ_i the indicator that Y_i is drawn from F . Define $N = \sum_{i=1}^m \delta_i \sim \mathcal{B}(m, \varepsilon)$ the number of random variables Y_i drawn from the distribution F . In the following, we upper bound the competitive ratio of any algorithm by analyzing its ratio on this particular instance. For this, we first provide a lower bound on $\mathbb{E}[\text{OPT}(Q, m)]$ using Lemma A.2, and obtain

$$\mathbb{E}[\text{OPT}(F, n - \Delta)] \geq \frac{1}{1 + \frac{2\Delta}{n}} \mathbb{E}[\text{OPT}(F, n + \Delta)] \geq \left(1 - \frac{2\Delta}{n}\right) \mathbb{E}[\text{OPT}(F, n + \Delta)],$$

thus we have

$$\begin{aligned} \mathbb{E}[\text{OPT}(Q, m)] &= \mathbb{E}[\text{OPT}(F, N)] \\ &\geq \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq n - \Delta}] \\ &\geq \mathbb{E}[\text{OPT}(F, n - \Delta) \mathbb{1}_{N \geq n - \Delta}] \\ &= \mathbb{E}[\text{OPT}(F, n - \Delta)] \Pr(N \geq n - \Delta) \\ &\geq \left(1 - \frac{2\Delta}{n}\right) \Pr(N \geq n - \Delta) \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &\geq \left(1 - \frac{2\Delta}{n} - \Pr(N < n - \Delta)\right) \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &\geq \left(1 - 2n^{-1/3} - \exp(-n^{1/3}/2)\right) \mathbb{E}[\text{OPT}(F, n + \Delta)] \\ &\geq \left(1 - 4n^{-1/3}\right) \mathbb{E}[\text{OPT}(F, n + \Delta)], \end{aligned} \tag{13}$$

where, for the last three inequalities, we used respectively Bernoulli's inequality, Chernoff bound, then $e^{-y} \leq 1/y$.

Then, we upper bound the reward of any algorithm given the instance (Q, m) as input. Let $L = \min_{i \neq j} \{|i - j| : \delta_i = 1, \delta_j = 1\}$ the smallest gap between two successive variables Y_i drawn from F , and let $t_1 < \dots < t_N$ the indices for which $\delta_i = 1$. We have for any algorithm ALG and positive integer ℓ that

$$\mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m)] = \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mathbb{1}_{N \geq n + \Delta \text{ or } L \leq \ell}] + \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mathbb{1}_{N < n + \Delta, L > \ell}]. \tag{14}$$

Using Lemma A.3 and Lemma A.5, the first term can be bounded as follows

$$\begin{aligned} \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mathbb{1}_{N \geq n + \Delta \text{ or } L \leq \ell}] &\leq \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq n + \Delta \text{ or } L \leq \ell}] \\ &\leq \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{N \geq n + \Delta}] + \mathbb{E}[\text{OPT}(F, N) \mathbb{1}_{L \leq \ell}] \\ &\leq \left(\frac{6}{n^{1/3}} + 7m\ell\varepsilon^2\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})]. \end{aligned}$$

Recalling that $\varepsilon = m/n$ and taking $\ell = \sqrt{m}$, we obtain

$$\mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mathbb{1}_{N \geq n + \Delta \text{ or } L \leq \ell}] \leq \left(\frac{6}{n^{1/3}} + \frac{7n^2}{\sqrt{m}}\right) \mathbb{E}[\text{OPT}(F, n + n^{2/3})]. \tag{15}$$

Regarding the second term in Equation (14), let τ be the stopping time of ALG and $t_\rho = \max\{j \leq \tau : \delta_j = 1\}$ the last value sampled from F and observed by ALG before it stops. We have

$$\begin{aligned} \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mathbb{1}_{N < n + \Delta, L > \ell}] &\leq \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mid N < n + \Delta, L > \ell] \\ &= \mathbb{E}[\max_{i \in [\tau]} D_{\tau-i}(Y_i) \mid N < n + \Delta, L > \ell] \\ &= \mathbb{E}[\max_{j \in [\rho]} D_{\tau-t_j}(Y_i) \mid N < n + \Delta, L > \ell] \\ &\leq \mathbb{E}[\max_{j \in [\rho]} D_{t_\rho-t_j}(Y_i) \mid N < n + \Delta, L > \ell] \\ &= \mathbb{E}[\max\{Y_{t_\rho}, \max_{j < \rho} D_{t_\rho-t_j}(Y_{t_j})\} \mid N < n + \Delta, L > \ell] \\ &\leq \mathbb{E}[\max\{Y_{t_\rho}, \max_{j < \rho} D_\ell(Y_{t_j})\} \mid N < n + \Delta]. \end{aligned}$$

We then prove that the last term is the reward of an algorithm A_m in the D_ℓ -prophet inequality. Let us A_m be the algorithm that takes as input an instance $X_1, \dots, X_{n+\Delta-1}$ of $n + \Delta$ IID random variables, then simulates $\text{ALG}^{\mathcal{D}}(Q, m) \mid N < n + \Delta$ as follows: let $\delta_1, \dots, \delta_m \stackrel{\text{iid}}{\sim} \mathcal{B}(n/m)$ set $N_A = 0$ and for each $i \in [m]$

- if $\delta_i = 0$: ALG observes the value $Y_i = 0$,
- if $\delta_i = 1$: increment N , then A_m observes the next value X_k , and ALG observes $Y_i = X_k$,
- if $N_A = n + \Delta - 1$ or ALG stops, then A_m also stops.

When ALG decides to stop, the current value observed by A_m is X_ρ : the last value Y_{t_ρ} observed by ALG such that $\delta_{t_\rho} = 0$. Observe that stopping when $N_A = n + \Delta + 1$, is equivalent to letting ALG observe zero values until the end, and stopping when ALG stops. Hence, the variables Y_1, \dots, Y_m have the same distribution as m IID samples from Q conditional to $N < n + \Delta$. Denoting ρ the stopping time of A_m and $\epsilon_\ell(F, n + \Delta)$ as defined in Lemma A.1, we deduce that

$$\begin{aligned} \mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m) \mid N < n + \Delta, L > \ell] &\leq \mathbb{E}[\max\{Y_{t_\rho}, \max_{j < \rho} D_\ell(Y_{t_j})\} \mid N < n + \Delta] \\ &= \mathbb{E}[\max\{X_\rho, \max_{j < \rho} D_\ell(X_j)\}] \\ &= \mathbb{E}[A_m^{D_\ell}(F, n + \Delta)] \\ &\leq \mathbb{E}[A_m^{D_\infty}(F, n + \Delta)] + \epsilon_\ell(F, n + \Delta). \end{aligned} \quad (16)$$

Substituting (15) and (16) in (14), with $\ell = \sqrt{m}$, yields

$$\mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m)] \leq \left(\frac{6}{n^{1/3}} + \frac{7n^2}{\sqrt{m}} \right) \mathbb{E}[\text{OPT}(F, n + \Delta)] + \mathbb{E}[A_m^{D_\infty}(F, n + \Delta)] + \epsilon_{\sqrt{m}}(F, n + \Delta),$$

and using Inequality 13, it follows that

$$\begin{aligned} \text{CR}^{\mathcal{D}}(\text{ALG}) &\leq \frac{\mathbb{E}[\text{ALG}^{\mathcal{D}}(Q, m)]}{\mathbb{E}[\text{OPT}(Q, m)]} \\ &\leq \frac{\frac{6}{n^{1/3}} + \frac{7n^2}{\sqrt{m}}}{1 - \frac{4}{n^{1/3}}} + \frac{\mathbb{E}[A_m^{D_\infty}(F, n + \Delta)] + \epsilon_{\sqrt{m}}(F, n + \Delta)}{(1 - \frac{4}{n^{1/3}})\mathbb{E}[\text{OPT}(F, n + \Delta)]} \\ &\leq \frac{\frac{6}{n^{1/3}} + \frac{7n^2}{\sqrt{m}}}{1 - \frac{4}{n^{1/3}}} + \frac{1}{1 - \frac{4}{n^{1/3}}} \left(\frac{\epsilon_{\sqrt{m}}(F, n + \Delta)}{\mathbb{E}[\text{OPT}(F, n + \Delta)]} + \sup_{A:\text{algo}} \frac{\mathbb{E}[A^{D_\infty}(F, n + \Delta)]}{\mathbb{E}[\text{OPT}(F, n + \Delta)]} \right), \end{aligned}$$

taking the limit $m \rightarrow \infty$ and using Lemma A.1 gives

$$\begin{aligned} \text{CR}^{\mathcal{D}}(\text{ALG}) &\leq \frac{6}{1 - \frac{4}{n^{1/3}}} + \frac{1}{1 - \frac{4}{n^{1/3}}} \left(\sup_{\text{A:algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(F, n + \Delta)]}{\mathbb{E}[\text{OPT}(F, n + \Delta)]} \right) \\ &= \frac{6}{n^{1/3} - 4} + \left(1 + \frac{4}{n^{1/3} - 4} \right) \left(\sup_{\text{A:algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(F, n + \Delta)]}{\mathbb{E}[\text{OPT}(F, n + \Delta)]} \right) \\ &\leq \frac{10}{n^{1/3} - 4} + \sup_{\text{A:algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(F, n + n^{2/3})]}{\mathbb{E}[\text{OPT}(F, n + n^{2/3})]}. \end{aligned}$$

where the last inequality holds because $\mathbb{E}[\mathbf{A}^{D_\infty}(F, n + \Delta)] \leq \mathbb{E}[\text{OPT}(F, n + \Delta)]$ for any algorithm \mathbf{A} . From here, the statement of the theorem can be deduced by observing that, for $k = n + n^{2/3}$, we have $n \geq (n + n^{2/3})/2 = k/2$, thus $n^{1/3} \geq k^{1/3}/2$, and we obtain

$$\begin{aligned} \text{CR}^{\mathcal{D}}(\text{ALG}) &\leq \frac{20}{k^{1/3} - 8} + \sup_{\text{A:algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(F, k)]}{\mathbb{E}[\text{OPT}(F, k)]} \\ &= \sup_{\text{A:algo}} \frac{\mathbb{E}[\mathbf{A}^{D_\infty}(F, k)]}{\mathbb{E}[\text{OPT}(F, k)]} + O\left(\frac{1}{k^{1/3}}\right). \end{aligned}$$

□

B From D_∞ -prophet to the $\gamma_{\mathcal{D}}$ -prophet inequality

B.1 Proof of Lemma 3.1

Proof. Let ALG be any rational algorithm in the D_∞ -prophet inequality. If ALG stops at some step $\tau \in [n]$, then by definition we have that $X_\tau > D_\infty(\max_{j < \tau} X_j)$, and thus $\text{ALG}^{D_\infty}(X_1, \dots, X_n) = \text{ALG}^0(X_1, \dots, X_n)$. Otherwise, if it stops at $\tau = n + 1$, then its reward is $\max_{i \in [n]} D_\infty(X_i) = D_\infty(\max_{i \in [n]} X_i)$, because D_∞ -is non increasing.

On the other hand, let \mathbf{A}_* be the optimal dynamic programming algorithm for the D_∞ -prophet inequality. At any step i , if $X_i < D_\infty(\max_{j < i} X_j)$, then stopping at i gives a reward of $D_\infty(\max_{j < i} X_j)$, while by rejecting X_i , the final reward is guaranteed to be at least $D_\infty(\max_{j < i} X_j)$. Thus rejecting X_i can only increase the reward, it is therefore the optimal decision. □

B.2 Proof of Proposition 3.2

Proof. Let us place ourselves in any order model, or in the IID model. Assume that $\inf_{x > 0} \frac{D_\infty(x)}{x} = 0$, then there exist a sequence $(s_k)_{k \geq 1}$ such that $\lim_{k \rightarrow \infty} \frac{D_\infty(s_k)}{s_k} = 0$.

Let $I = (F_1, \dots, F_n)$ an instance of non-negative random variables with finite expectation, and $X_i \sim F_i$ for all $i \in [n]$. Let ALG be a rational algorithm for the D_∞ -prophet inequality and let us denote τ its stopping time. Denoting $X^* := \max_{i \in [n]} X_i$ and using Lemma 3.1, we have for any constant $C > 0$ that that

$$\begin{aligned} \mathbb{E}[\text{ALG}^{D_\infty}(I)] &= \mathbb{E}[\text{ALG}^0(I) \mathbb{1}_{\tau \leq n}] + \mathbb{E}[D_\infty(X^*) \mathbb{1}_{\tau = n+1}] \\ &\leq \mathbb{E}[\text{ALG}^0(I)] + \mathbb{E}[D_\infty(C)] + \mathbb{E}[D_\infty(X^*) \mathbb{1}_{X^* > C}] \\ &\leq \sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)] + \mathbb{E}[D_\infty(C)] + \mathbb{E}[X^* \mathbb{1}_{X^* > C}]. \end{aligned} \tag{17}$$

Let $k \geq 1$ a positive integer, M a positive constant, and consider the instance I^k of random variables (X_1^k, \dots, X_n^k) with $X_i^k \sim \frac{s_k}{M} X_i$ for all $i \in [n]$. We have that $\text{OPT}(I^k) = \frac{s_k}{M} \text{OPT}(I)$ and $\sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I^k)] = \frac{s_k}{M} \sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)]$. Therefore, using Inequality (17) with I^k instead of I and

$C = \frac{s_k}{M}$, then dividing by $\text{OPT}(I^k)$ gives that

$$\begin{aligned} \text{CR}^{D_\infty}(\text{ALG}) &\leq \frac{\mathbb{E}[\text{ALG}^{D_\infty}(I^k)]}{\mathbb{E}[\text{OPT}(I^k)]} \leq \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]} + \frac{D_\infty(s_k)}{\frac{s_k}{M} \mathbb{E}[\text{OPT}(I)]} + \frac{\mathbb{E}[\frac{s_k}{M} X^* \mathbb{1}_{X^* > M}]}{\frac{s_k}{M} \mathbb{E}[\text{OPT}(I)]} \\ &= \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]} + \left(\frac{M}{\mathbb{E}[\text{OPT}(I)]} \right) \frac{D_\infty(s_k)}{s_k} + \frac{\mathbb{E}[X^* \mathbb{1}_{X^* > M}]}{\mathbb{E}[\text{OPT}(I)]}, \end{aligned}$$

and taking the limit $k \rightarrow \infty$, we obtain

$$\text{CR}^{D_\infty}(\text{ALG}) \leq \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]} + \frac{\mathbb{E}[X^* \mathbb{1}_{X^* > M}]}{\mathbb{E}[\text{OPT}(I)]},$$

and since X^* has finite expectation, taking the limit $M \rightarrow \infty$ gives

$$\text{CR}^{D_\infty}(\text{ALG}) \leq \sup_{\mathbf{A}} \frac{\mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]}.$$

Finally, taking the infimum over all instances, we obtain that

$$\text{CR}^{D_\infty}(\text{ALG}) \leq \sup_{\mathbf{A}} \text{CR}^0(\mathbf{A}).$$

As in the proof of Corollary 2.1.1, permuting \inf_I and $\sup_{\mathbf{A}}$ is possible because there is an algorithm (dynamic programming) achieving the supremum for any instance. By Lemma 3.1, the inequality above holds for in particular for the optimal dynamic programming algorithm, which has a maximal competitive ratio. Therefore, the inequality remains true for any other algorithm \mathbf{A} , not necessarily rational. \square

B.3 Proof of Proposition 3.3

Proof. Let us place ourselves in any order model or in the IID model. Since $\gamma = \inf_{x>0} \frac{D_\infty(x)}{x}$, there exists a sequence $(s_k)_{k \geq 1}$ of positive numbers such that $\lim_{k \rightarrow \infty} \frac{D_\infty(s_k)}{s_k} = \gamma$.

For the random variables X_1, \dots, X_n taking values in $\{0, a, b\}$ a.s., in any order model, it is clear that the optimal decision when observing a zero value is to reject it, and when observing the value b is to accept it. Let ALG be an algorithm satisfying this property and let τ be its stopping time. If $\tau = n + 1$ then necessarily $\max_{i \in [n]} X_i \neq b$, and the reward of ALG in that case is $D_\infty(a)$ if $\max_{i \in [n]} X_i$ and 0 otherwise. In particular, ALG is rational in the D_∞ -prophet inequality and we have by Lemma 3.1 that

$$\begin{aligned} \mathbb{E}[\text{ALG}^{D_\infty}(I)] &= \mathbb{E}[\text{ALG}^0(I)] + \mathbb{E}[D_\infty(\max_{i \in [n]} X_i) \mathbb{1}_{\tau = n+1}] \\ &= \mathbb{E}[\text{ALG}^0(I)] + D_\infty(a) \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a). \end{aligned} \quad (18)$$

Consider now the instance I^k of random variables (X_1^k, \dots, X_n^k) where $X_i^k = \frac{s_k}{a} X_i$ for all $i \in [n]$. I^k satisfies the same assumptions as I with $a^k = s_k$ and $b^k = \frac{s_k b}{a}$, and we have $\mathbb{E}[\text{OPT}(I^k)] = \frac{s_k}{a} \mathbb{E}[\text{OPT}(I)]$, $\mathbb{E}[\text{ALG}^0(I^k)] \leq \frac{s_k}{a} \sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)]$ and $(\max_{i \in [n]} X_i^k = a^k) \iff (\max_{i \in [n]} X_i = a)$. It follows that

$$\begin{aligned} \frac{\mathbb{E}[\text{ALG}^{D_\infty}(I^k)]}{\mathbb{E}[\text{OPT}(I^k)]} &\leq \frac{\sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]} + \frac{D_\infty(s_k)}{\frac{s_k}{a} \mathbb{E}[\text{OPT}(I)]} \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a) \\ &= \frac{\sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]} + \left(\frac{D_\infty(s_k)}{s_k} \right) \frac{a}{\mathbb{E}[\text{OPT}(I)]} \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a). \end{aligned}$$

Taking the limit $k \rightarrow \infty$ gives

$$\begin{aligned} \text{CR}^{D_\infty}(\text{ALG}) &\leq \frac{\sup_{\mathbf{A}} \mathbb{E}[\mathbf{A}^0(I)] + \gamma a \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a)}{\mathbb{E}[\text{OPT}(I)]} \\ &= \frac{\sup_{\mathbf{A}} (\mathbb{E}[\mathbf{A}^0(I)] + \mathbb{E}[\gamma (\max_{i \in [n]} X_i) \mathbb{1}_{\tau = n+1}])}{\mathbb{E}[\text{OPT}(I)]}. \end{aligned}$$

ALG is also rational in the γ -prophet inequality. Therefore, using Lemma 3.1, we deduce that

$$\text{CR}^{D_\infty}(\text{ALG}) \leq \frac{\mathbb{E}[\text{ALG}^\gamma(I)]}{\mathbb{E}[\text{OPT}(I)]} \leq \sup_A \frac{\mathbb{E}[\mathbf{A}^\gamma(I)]}{\mathbb{E}[\text{OPT}(I)]}.$$

This upper bound is true for the optimal dynamic programming algorithm, since it rejects all zeros and accepts b , therefore the upper bound also holds for any other algorithm. \square

C The $\gamma_{\mathcal{D}}$ -prophet inequality

C.1 Proof of Proposition 4.2

Proof. For the lower bound, it suffices to consider the following trivial algorithm: if $\alpha > \gamma$ then run \mathbf{A}_α , and if $\gamma > \alpha$ then observe all the items then select the one with maximum value.

For the upper bound, let $I = (F_1, \dots, F_n)$ be an instance of the problem and $X_i \sim F_i$ for all i , and let $\beta_I := \sup_A \frac{\mathbb{E}[\mathbf{A}^0(I)]}{\mathbb{E}[\text{OPT}(I)]}$. Let \mathbf{A} be any algorithm, τ its stopping time, and $Y_\tau = \max_{i < \tau} X_{\pi(i)}$, where π is the observation order. With the previous notations, we can write that $\mathbb{E}[\mathbf{A}^\gamma(I)] = \mathbb{E}[\max(X_{\pi(\tau)}, \gamma Y_\tau)]$. For any $x, y \geq 0$, the application $\gamma \mapsto \max(x, \gamma y)$ is convex on $[0, 1]$, hence it can be upper bounded by $(1 - \gamma)x + \gamma \max(x, y)$. Therefore

$$\begin{aligned} \mathbb{E}[\mathbf{A}^\gamma(I)] &\leq (1 - \gamma)\mathbb{E}[X_{\pi(\tau)}] + \gamma\mathbb{E}[\max(X_{\pi(\tau)}, Y_\tau)] \\ &\leq (1 - \gamma)\mathbb{E}[\mathbf{A}^0(I)] + \gamma\mathbb{E}[\text{OPT}(I)] \\ &\leq ((1 - \gamma)\beta_I + \gamma)\mathbb{E}[\text{OPT}(I)]. \end{aligned}$$

Therefore, $\text{CR}^\gamma(\text{ALG}) \leq ((1 - \gamma)\beta_I + \gamma)$. Taking the infimum over all the instances I gives the result. Indeed, if we denote \mathbf{A}_* the optimal dynamic programming algorithm for the standard prophet inequality, then

$$\inf_I \beta_I = \inf_I \frac{\mathbb{E}[\mathbf{A}_*^0(I)]}{\mathbb{E}[\text{OPT}(I)]} = \text{CR}^0(\mathbf{A}_*) \leq \beta.$$

\square

C.2 Proofs for the adversarial order model

Proof. We first prove the upper bound, and then analyze the single threshold algorithm proposed in the theorem.

Upper bound Let $\varepsilon \in (0, 1 - \gamma)$, and let $a = \frac{1}{1 - (1 - \varepsilon)\gamma}$ (such that $1 + (1 - \varepsilon)\gamma a = a$). Let X_1, X_2 the two random variables defined by $X_1 = a$ almost surely and

$$X_2 = \begin{cases} \frac{1}{\varepsilon} & \text{w.p. } \varepsilon \\ 0 & \text{w.p. } 1 - \varepsilon \end{cases}.$$

Stopping at the first step gives a reward of a , while stopping at the second step gives

$$\mathbb{E}[\max(\gamma a, X_2)] = \varepsilon \times \frac{1}{\varepsilon} + (1 - \varepsilon) \times \gamma a = 1 + (1 - \varepsilon)\gamma a = a,$$

hence the expected output of any algorithm is exactly equal to a . On the other hand

$$\mathbb{E}[\max(X_1, X_2)] = 1 + (1 - \varepsilon)a,$$

and we deduce that, for any algorithm ALG for the γ -prophet inequality, we have

$$\text{CR}(\text{ALG}) \leq \frac{\mathbb{E}[\text{ALG}(X_1, X_2)]}{\mathbb{E}[\max(X_1, X_2)]} = \frac{a}{1 + (1 - \varepsilon)a},$$

and this is true for any $\varepsilon \in (0, 1 - \gamma)$, thus taking $\varepsilon \rightarrow 0$ gives

$$\text{CR}(\text{ALG}) \leq \frac{\frac{1}{1 - \gamma}}{1 + \frac{1}{1 - \gamma}} = \frac{1}{2 - \gamma}.$$

Lower bound Consider an algorithm with an acceptance threshold θ , i.e. that accepts the first value larger than θ . Let $I = (F_1, \dots, F_n)$ be any instance, such that $X_i \sim F_i$ for all i , and let us define $X^* = \max_{i \in [n]} X_i$ and $p = \Pr(X^* \leq \theta)$. In the classical prophet inequality, if no value is larger than θ then the reward of the algorithm is zero, and we have the classical lower bound

$$\mathbb{E}[\text{ALG}^0(I)] \geq (1-p)\theta + p\mathbb{E}[(X^* - \theta)_+],$$

For the γ -prophet, if no value is larger than θ (i.e. $X^* \leq \theta$), then the algorithm gains γX^* instead of 0. Therefore, it holds that

$$\begin{aligned} \mathbb{E}[\text{ALG}^\gamma(I)] &= \mathbb{E}[\text{ALG}^0(I)] + \mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \\ &\geq (1-p)\theta + p\mathbb{E}[(X^* - \theta)_+] + \gamma\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \\ &= (1-p)\theta + p\mathbb{E}[(X^* - \theta) \mathbb{1}_{X^* > \theta}] + \gamma\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \\ &= (1-p)\theta + p(\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] - (1-p)\theta) + \gamma\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \\ &= (1-p)^2\theta + p\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] + \gamma\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}], \end{aligned}$$

and observing that

$$\theta = \frac{\mathbb{E}[\theta \mathbb{1}_{X^* \leq \theta}]}{p} \geq \frac{\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}]}{p},$$

we deduce the lower bound

$$\begin{aligned} \mathbb{E}[\text{ALG}^\gamma(I)] &\geq p\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] + \left(\gamma + \frac{(1-p)^2}{p}\right) \mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \\ &\geq \min\left\{p, \gamma + \frac{(1-p)^2}{p}\right\} \mathbb{E}[X^*]. \end{aligned}$$

The right term is maximized for p satisfying $p = \gamma + \frac{(1-p)^2}{p}$, that leads to

$$\begin{aligned} p = \gamma + \frac{(1-p)^2}{p} &\iff p^2 = \gamma p + 1 - 2p + p^2 \\ &\iff p = \frac{1}{2-\gamma}. \end{aligned}$$

Hence, by choosing a threshold θ satisfying $\Pr(X^* \leq \theta) = \frac{1}{2-\gamma}$ we obtain a competitive ratio of at least $\frac{1}{2-\gamma}$. \square

C.3 Proofs for the random order model

We prove here the upper and lower bounds stated in Theorem 4.4.

C.3.1 Proof of Theorem 4.4

Proof. We first prove the upper bound, and then derive the analysis for single threshold algorithms.

Upper bound Let $a > 0$, and let X_1, \dots, X_{n+1} be independent random variables such that $X_{n+1} = a$ a.s. and for $1 \leq i \leq n$

$$X_i \sim \begin{cases} n & \text{w.p. } \frac{1}{n^2} \\ 0 & \text{w.p. } 1 - \frac{1}{n^2} \end{cases}.$$

Any reasonable algorithm skips zero values and stops when observing the value n . The only strategic decision to make is thus to stop or not when observing $X_{n+1} = a$. By analyzing the dynamic programming solution ALG_* we obtain that the optimal decision rule is to skip a if it is observed before a certain step j , and select it otherwise. The step j corresponds to the time when the expectation of the future reward is less than a . Let π be the random order in which the variables are observed. Then, if $\pi^{-1}(n+1) < j$, i.e. if the value a is observed before time j , X_{n+1} is not selected. The output of this algorithm is hence n if at least one random variable equals n , and γa otherwise. This leads to

$$\begin{aligned} \mathbb{E}[\text{ALG}_*^\gamma(X) \mid \pi^{-1}(n+1) < j] &= n \left(1 - \left(1 - \frac{1}{n^2}\right)^n\right) + \gamma a \left(1 - \frac{1}{n^2}\right)^n \\ &\leq 1 + \gamma a, \end{aligned}$$

where we used the inequality $(1 - \frac{1}{n^2})^n \geq 1 - \frac{1}{n}$. On the other hand, if $\pi^{-1}(n) \geq j$, then a is selected if the value n has not been observed before it, hence for any $i \geq j$ we have

$$\begin{aligned} \mathbb{E}[\text{ALG}_*^\gamma(X) \mid \pi^{-1}(n+1) = i] &= n \left(1 - \left(1 - \frac{1}{n^2} \right)^{i-1} \right) + a \left(1 - \frac{1}{n^2} \right)^{i-1} \\ &\leq \frac{i-1}{n} + a, \end{aligned}$$

we deduce that

$$\begin{aligned} \mathbb{E}[\text{ALG}_*^\gamma(X)] &\leq (1 + \gamma a) \Pr(\pi^{-1}(n) \leq j-1) + \sum_{i=j}^{n+1} \left(\frac{i-1}{n} + a \right) \Pr(\pi^{-1}(n) = i) \\ &= \frac{j-1}{n+1} (1 + \gamma a) + \frac{1}{n+1} \sum_{i=j}^{n+1} \left(\frac{i-1}{n} + a \right) \\ &= (1 - (1 - \gamma)a) \frac{j}{n} - \frac{1}{2} \left(\frac{j}{n} \right)^2 + \frac{1}{2} + a + o(1) \\ &\leq 1 + 2\gamma a + (1 - \gamma)^2 a^2 + o(1), \end{aligned}$$

where the last inequality is obtained by maximizing over j/n . Finally, we directly obtain that

$$\begin{aligned} \mathbb{E}[\max_i X_i] &= n \left(1 - \left(1 - \frac{1}{n^2} \right)^n \right) + a \left(1 - \frac{1}{n^2} \right)^n \\ &= 1 + a + o(1), \end{aligned}$$

so for any algorithm **ALG** we obtain that

$$\text{CR}^\gamma(\text{ALG}) \leq \text{CR}^\gamma(\text{ALG}_*) \leq \frac{1 + 2\gamma a + (1 - \gamma)^2 a^2}{1 + a}.$$

The function above is minimized for $a = \sqrt{\frac{3-\gamma}{1-\gamma}} - 1$, which translates to

$$\text{CR}^\gamma(\text{ALG}) \leq (1 - \gamma)^{3/2} (\sqrt{3 - \gamma} - \sqrt{1 - \gamma}) + \gamma.$$

Lower bound We still denote by $I = (F_1, \dots, F_n)$ the input instance and $X_i \sim F_i$ for all $i \in [n]$. Let **ALG** be the algorithm with single threshold θ , then it is direct that

$$\text{ALG}^\gamma(I) = \text{ALG}^0(I) + \gamma X^* \mathbb{1}_{X^* < \theta}. \quad (19)$$

We start by giving a lower bound on $\mathbb{E}[\text{ALG}^0]$. We use from Correa et al. [2021c] (Theorem 2.1) that for any $x < \theta$ it holds that

$$\Pr(\text{ALG}^0(I) \geq x) = \Pr(\text{ALG}^0(I) \geq \theta) = \Pr(X^* \geq \theta) = 1 - p,$$

and for $x \geq \theta$ it holds that

$$\Pr(\text{ALG}^0(I) \geq x) \geq \frac{1-p}{-\log p} \Pr(X^* \geq x),$$

from which we deduce that

$$\begin{aligned} \mathbb{E}[\text{ALG}^0(I)] &= \int_0^\infty \Pr(\text{ALG}^0(I) \geq x) dx \\ &\geq (1-p)\theta + \frac{1-p}{-\log p} \int_\theta^\infty \Pr(X^* \geq x) dx. \end{aligned} \quad (20)$$

On the other hand, we obtain that

$$\begin{aligned} \mathbb{E}[X^* \mathbb{1}_{X^* < \theta}] &= \int_0^\infty \Pr(X^* \mathbb{1}_{X^* < \theta} \geq x) dx = \int_0^\infty \Pr(x \leq X^* < \theta) dx \\ &= \int_0^\theta (\Pr(X^* > x) - \Pr(X^* \geq \theta)) dx \\ &= \int_0^\theta \Pr(X^* > x) dx - (1-p)\theta. \end{aligned} \quad (21)$$

Using (19), (20) and (21) we deduce that

$$\begin{aligned}
 \mathbb{E}[\text{ALG}^\gamma(I)] &\geq (1-p)\theta + \frac{1-p}{-\log p} \int_\theta^\infty \Pr(X^* \geq x) dx + \gamma \int_0^\theta \Pr(X^* \geq x) dx - \gamma(1-p)\theta \\
 &= (1-\gamma)(1-p)\theta + \gamma \int_0^\theta \Pr(X^* \geq x) dx + \frac{1-p}{-\log p} \int_\theta^\infty \Pr(X^* \geq x) dx \\
 &\geq ((1-\gamma)(1-p) + \gamma) \int_0^\theta \Pr(X^* \geq x) dx + \frac{1-p}{-\log p} \int_\theta^\infty \Pr(X^* \geq x) dx \\
 &\geq \min \left\{ (1-\gamma)(1-p) + \gamma, \frac{1-p}{-\log p} \right\} \left(\int_0^\theta \Pr(X^* \geq x) dx + \int_\theta^\infty \Pr(X^* \geq x) dx \right) \\
 &= \min \left\{ 1 - (1-\gamma)p, \frac{1-p}{-\log p} \right\} \mathbb{E}[X^*].
 \end{aligned}$$

Finally, choosing $p = p_\gamma$ gives the result. \square

C.3.2 Proof of Corollary 4.4.1

Proof. For $p = \frac{1/e}{1-(1-1/e)\gamma}$, we have immediately that

$$1 - (1-\gamma)p = 1 - \frac{(1-\gamma)/e}{1 - (1-1/e)\gamma},$$

and $p \in [1/e, 1]$ for any $\gamma \in [0, 1]$. Since the function $x \mapsto (1-x)/\log(1/x)$ is concave, we can lower bound it on $[1/e, 1]$ by $x \mapsto 1 - 1/e + \frac{x-1/e}{e-1}$, which is the line intersecting it in $1/e$ and 1 . Therefore we have

$$\frac{1-p}{-\log p} \geq 1 - 1/e + \frac{p-1/e}{e-1} = 1 - \frac{(1-\gamma)/e}{1 - (1-1/e)\gamma}.$$

Finally, using the previous theorem, this choice of p guarantees a competitive ratio of at least

$$\min \left\{ 1 - (1-\gamma)p, \frac{1-p}{-\log p} \right\} = 1 - \frac{(1-\gamma)/e}{1 - (1-1/e)\gamma}.$$

\square

C.4 Proofs for the IID model

Proof of Lemma 4.5

Proof of Lemma 4.5. Let F be the cumulative distribution function of X_1 , $a > 0$ and ALG the algorithm with single threshold θ such that $1 - F(\theta) = \frac{a}{n}$. We denote $X^* = \max_{i \in [n]} X_i$. As in the previous proofs, we will begin by lower bounding $\text{ALG}^0(F, n)$. For any $i \in [n]$, ALG stops at step i if and only if $X_i > \theta$ and all the previous items were rejected, i.e. $X_j \leq \theta$ for all $j < i$. Thus we can write

$$\begin{aligned}
 \mathbb{E}[\text{ALG}^0(F, n)] &= \mathbb{E} \left[\sum_{i=1}^n (X_i \mathbb{1}_{X_i > \theta}) \mathbb{1}_{\forall j < i: X_j \leq \theta} \right] \\
 &= \sum_{i=1}^n F(\theta)^{i-1} \mathbb{E}[X_i \mathbb{1}_{X_i > \theta}] \\
 &= \frac{1 - F(\theta)^n}{1 - F(\theta)} \mathbb{E}[X_1 \mathbb{1}_{X_1 > \theta}] \\
 &= \frac{1 - F(\theta)^n}{a} \times n \mathbb{E}[X_1 \mathbb{1}_{X_1 > \theta}], \tag{22}
 \end{aligned}$$

where the second equality is true by the independence of the random variables $(X_i)_i$. On the other hand, we can upper bound $\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}]$ as follows

$$\begin{aligned} \mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] &\leq \Pr(X^* > \theta)\theta + \mathbb{E}[(X^* - \theta)_+] \\ &\leq \Pr(X^* > \theta)\theta + \mathbb{E}\left[\sum_{i=1}^n (X_i - \theta)_+\right] \\ &= \Pr(X^* > \theta)\theta + n(\mathbb{E}[X_1 \mathbb{1}_{X_1 > \theta}] - \Pr(X_1 > \theta)\theta) \\ &= (1 - F(\theta)^n - a)\theta + n\mathbb{E}[X_1 \mathbb{1}_{X_1 > \theta}]. \end{aligned}$$

Using the definition of θ , the independence of $(X_i)_i$ then Bernoulli's inequality we have that

$$\Pr(X^* > \theta) = 1 - F(\theta)^n = 1 - \left(1 - \frac{a}{n}\right)^n \leq 1 - (1 - n \times \frac{a}{n}) = a,$$

and observing that $\theta = \frac{\mathbb{E}[\theta \mathbb{1}_{X^* \leq \theta}]}{F(\theta)^n} \geq \frac{\mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}]}{F(\theta)^n}$, we deduce that

$$\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] \leq -\left(1 - \frac{1-a}{F(\theta)^n}\right) \mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] + n\mathbb{E}[X_1 \mathbb{1}_{X_1 > \theta}].$$

by substituting into (22), we obtain

$$\mathbb{E}[\text{ALG}^0(F, n)] \geq \frac{1 - F(\theta)^n}{a} \left(\mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] + \left(1 - \frac{1-a}{F(\theta)^n}\right) \mathbb{E}[X^* \mathbb{1}_{X^* \leq \theta}] \right).$$

Finally, the reward in the γ -prophet inequality is

$$\begin{aligned} \mathbb{E}[\text{ALG}^\gamma(F, n)] &= \mathbb{E}[\text{ALG}^0(F, n)] + \gamma \mathbb{E}[X^* \mathbb{1}_{X^* < \theta}] \\ &\geq \frac{1 - F(\theta)^n}{a} \mathbb{E}[X^* \mathbb{1}_{X^* > \theta}] + \left(\frac{1 - F(\theta)^n}{a} \left(1 - \frac{1-a}{F(\theta)^n}\right) + \gamma \right) \mathbb{E}[X^* \mathbb{1}_{X^* < \theta}] \\ &\geq \min \left\{ \frac{1 - F(\theta)^n}{a}, \frac{1 - F(\theta)^n}{a} \left(1 - \frac{1-a}{F(\theta)^n}\right) + \gamma \right\} \mathbb{E}[X^*]. \end{aligned}$$

The equation $\frac{1 - F(\theta)^n}{a} = \frac{1 - F(\theta)^n}{a} \left(1 - \frac{1-a}{F(\theta)^n}\right) + \gamma$, is equivalent to

$$\left(\frac{1}{(1 - a/n)^n} - 1 \right) \left(\frac{1}{a} - 1 \right) = \gamma, \quad (23)$$

and for any $n \geq 2$ the function $a \mapsto \left(\frac{1}{(1 - a/n)^n} - 1\right) \left(\frac{1}{a} - 1\right)$ is decreasing on $(0, 1]$ and goes from 1 to 0, thus Equation (23) admits a unique solution $a_{n,\gamma}$, and taking $a = a_{n,\gamma}$ guarantees a reward of $\frac{1 - F(\theta)^n}{a_{n,\gamma}} \mathbb{E}[X^*] = \frac{1 - (1 - \frac{a_{n,\gamma}}{n})^n}{a_{n,\gamma}} \mathbb{E}[X^*]$. \square

Proof of Theorem 4.6

Proof. We first prove the upper bound, and then we give the single-threshold algorithm satisfying the lower bound.

Upper bound We consider an instance similar to the one used in the proof of Theorem 4.4. Let $a, x > 0$, and let X_1, \dots, X_n be IID random variables with the following the distribution F defined by

$$X_1 \sim \begin{cases} n & \text{w.p. } \frac{1}{n^2} \\ a & \text{w.p. } \frac{x}{n} \\ 0 & \text{w.p. } 1 - \frac{x}{n} - \frac{1}{n^2} \end{cases}.$$

A reasonable algorithm would always reject the value 0 and accept the value n . However, if the algorithm faces an item with value a , it must decide to either accept it, or reject it with a guarantee of recovering γa at the end. By analyzing the dynamic programming algorithm ALG_* , we find that the optimal decision is to reject a if observed before a certain step j , and accept it otherwise. Let us denote τ the stopping time of ALG_* . By convention, we write $\tau = n + 1$ to say that no value was selected by the algorithm, in which case the reward is $\gamma \max_{i \in [n]} X_i$.

If $\tau \leq j - 1$ then necessarily $X_\tau = n$, because ALG_\star rejects the value a if it is met before step j , and if $\tau = n + 1$ then $\max_{i \in [n]} X_i \in \{0, a\}$, because otherwise the algorithm would have selected the value n and stopped earlier. It follows that the expected output of ALG_\star on this instance is

$$\begin{aligned} \mathbb{E}[\text{ALG}_\star^\gamma(F, n)] &= n \Pr(\tau < j) + \sum_{i=j}^n \mathbb{E}[X_i \mid \tau = i] \Pr(\tau = i) \\ &\quad + \gamma a \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a). \end{aligned} \tag{24}$$

Let us now compute the terms above one by one.

$$\Pr(\tau < j) = \Pr(\exists i \in [j - 1] : X_i = n) = 1 - \left(1 - \frac{1}{n^2}\right)^{j-1} \leq \frac{j}{n^2},$$

where we used Bernoulli's inequality $(1 - 1/n^2)^{j-1} \geq 1 - \frac{j-1}{n^2} > 1 - \frac{j}{n^2}$. For $i \in \{j, \dots, n\}$, ALG_\star stops at i if $X_i \in \{a, n\}$ and if it has not stopped before, i.e. $X_k \in \{0, a\}$ for all $k < j$ and $X_k = 0$ for all $k \in \{j, \dots, i - 1\}$, hence

$$\begin{aligned} \Pr(\tau = i) &= \Pr(\forall k < j : X_k \neq n \text{ and } \forall j \leq k \leq i - 1 : X_k = 0 \text{ and } X_i \neq 0) \\ &= \left(1 - \frac{1}{n^2}\right)^{j-1} \left(1 - \frac{x}{n} - \frac{1}{n^2}\right)^{i-j} \Pr(X_i \neq 0) \\ &\leq \left(1 - \frac{x}{n}\right)^{i-j} \Pr(X_i \neq 0), \end{aligned}$$

the second equality is true by independence, and the last inequality holds because $1 - \frac{1}{n^2} \leq 1$ and $1 - \frac{x}{n} - \frac{1}{n^2} \leq 1 - \frac{x}{n}$. By independence of the variables $(X_k)_k$, we also have that

$$\mathbb{E}[X_i \mid \tau = i] = \mathbb{E}[X_i \mid X_i \neq 0] = \frac{\mathbb{E}[X_i]}{\Pr(X_i \neq 0)} = \frac{1 + ax}{n \Pr(X_i \neq 0)}.$$

Finally, the event $(\tau = n + 1, \max_{i \in [n]} X_i = a)$ is equivalent $(\max_{i \in [j-1]} X_i = a, \forall k \geq j : X_k = 0)$. In fact, the algorithm does not stop before $n + 1$ if and only if $X_k \neq n$ for all $k < j$ and $X_k = 0$ for all $j \leq k \leq n$, and under these conditions, it holds that $\max_{i \in [n]} X_i = \max_{i \in [j-1]} X_i$. Therefore

$$\begin{aligned} \Pr(\tau = n + 1, \max_{i \in [n]} X_i = a) &= \Pr(\max_{i \in [j-1]} X_i = a, \forall k \geq j : X_k = 0) \\ &\leq \Pr(\max_{i \in [j-1]} X_i \neq 0) \Pr(\forall k \geq j : X_k = 0) \\ &= \left(1 - \left(1 - \frac{x}{n} - \frac{1}{n^2}\right)^{j-1}\right) \left(1 - \frac{x}{n} - \frac{1}{n^2}\right)^{n-j} \\ &= (1 - e^{-\frac{xj}{n}} + o(1))(e^{-x + \frac{xj}{n}} + o(1)) \\ &= (e^{\frac{xj}{n}} - 1)e^{-x} + o(1). \end{aligned}$$

All in all, we obtain by substituting into 24 that

$$\begin{aligned} \mathbb{E}[\text{ALG}_\star^\gamma(F, n)] &\leq \frac{j}{n} + \left(\frac{1 + ax}{n}\right) \sum_{i=j}^n \left(1 - \frac{x}{n}\right)^{i-j} + \gamma a e^{-x} (e^{\frac{xj}{n}} - 1) + o(1) \\ &= \frac{j}{n} + \left(\frac{1 + ax}{n}\right) \frac{1 - (1 - x/n)^{n-j+1}}{x/n} + \gamma a e^{-x} (e^{\frac{xj}{n}} - 1) + o(1) \\ &= \frac{j}{n} + \left(\frac{1}{x} + a\right) (1 - e^{-x + \frac{xj}{n}} + o(1)) + \gamma a e^{-x} (e^{\frac{xj}{n}} - 1) + o(1) \\ &= \frac{j}{n} - \left[\left(\frac{1}{x} + (1 - \gamma)a\right) e^{-x + \frac{xj}{n}} + \frac{1}{x} + a - \gamma a e^{-x} + o(1)\right] \\ &\leq \max_{s>0} \left\{s - \left[\left(\frac{1}{x} + (1 - \gamma)a\right) e^{-x + xs}\right]\right\} + \frac{1}{x} + (1 - \gamma e^{-x})a + o(1) \\ &= -\frac{1}{x} (\log(1 + (1 - \gamma)ax) + 1 - x) + \frac{1}{x} + (1 - \gamma e^{-x})a + o(1) \\ &= -\frac{1}{x} \log(1 + (1 - \gamma)ax) + 1 + (1 - \gamma e^{-x})a + o(1). \end{aligned}$$

On the other hand, we have that

$$\Pr(\max_{i \in [n]} X_i = n) = 1 - \left(1 - \frac{1}{n^2}\right)^n = \frac{1}{n} + o(1/n),$$

$$\Pr(\max_{i \in [n]} X_i = 0) = \left(1 - \frac{x}{n} - \frac{1}{n^2}\right)^n = e^{-x} + o(1),$$

$$\Pr(\max_{i \in [n]} X_i = a) = 1 - \Pr(\max_{i \in [n]} X_i = 0) - \Pr(\max_{i \in [n]} X_i = n) = 1 - e^{-x} + o(1),$$

therefore, the expected maximum value is

$$\begin{aligned} \mathbb{E}[\max_{i \in [n]} X_i] &= n \Pr(\max_{i \in [n]} X_i = n) + a \Pr(\max_{i \in [n]} X_i = a) \\ &= 1 + (1 - e^{-x})a + o(1). \end{aligned}$$

We deduce that

$$\begin{aligned} \frac{\mathbb{E}[\text{ALG}_*^\gamma(F, n)]}{\mathbb{E}[\max_{i \in [n]} X_i]} &\leq \frac{-\frac{1}{x} \log(1 + (1 - \gamma)ax) + 1 + (1 - \gamma)e^{-x}a}{1 + (1 - e^{-x})a} + o(1) \\ &= 1 - \frac{\frac{1}{x} \log(1 + (1 - \gamma)ax) - (1 - \gamma)ae^{-x}}{1 + (1 - e^{-x})a} + o(1). \end{aligned}$$

Consequently, for any $a, x > 0$ and for any algorithm **ALG** we have

$$\begin{aligned} \text{CR}(\text{ALG}) &\leq \text{CR}(\text{ALG}_*) \\ &\leq \lim_{n \rightarrow \infty} \frac{\mathbb{E}[\text{ALG}_*^\gamma(F, n)]}{\mathbb{E}[\max_{i \in [n]} X_i]} \\ &\leq 1 - \frac{\log(1 + (1 - \gamma)ax) - (1 - \gamma)axe^{-x}}{x + (1 - e^{-x})ax}. \end{aligned}$$

In particular, for $x = 2$ and $a = \frac{1-\gamma/2}{1-\gamma}$ we find that

$$\begin{aligned} \text{CR}(\text{ALG}) &\leq 1 - \frac{\log(3 - \gamma) - (2 - \gamma)e^{-2}}{2 + \frac{2-\gamma}{1-\gamma}(1 - e^{-2})} \\ &= 1 - (1 - \gamma) \frac{e^2 \log(3 - \gamma) - (2 - \gamma)}{2(2e^2 - 1) - \gamma(3e^2 - 1)} \\ &= U(\gamma). \end{aligned}$$

This proves the upper bound stated in the theorem, and we can verify that it is increasing, and satisfies $U(0) = \frac{4 - \log 3}{4 - 2/e^2}$ $U(1) = 1$

Lower bound on the competitive ratio We will prove that the algorithm presented in Lemma 4.5 has a competitive ratio of at least $(1 - (1 - \gamma)p_\gamma)$, where p_γ , first introduced in Theorem 4.4, is the unique solution of the equation $(1 - (1 - \gamma)p) = \frac{1-p}{-\log p}$, which is equivalent to $(\frac{1}{p} - 1)(\frac{1}{\log(1/p)} - 1) = \gamma$.

Let $a_\gamma = -\log(p_\gamma)$. It follows from the definition of p_γ that a_γ is the unique solution of the equation $(e^a - 1)(\frac{1}{a} - 1) = \gamma$. For any $n \geq 2$ and $x \geq 0$ we have that $(1 - x/n)^n \leq e^{-x}$, hence, by definition of $a_{n,\gamma}$ and a_γ

$$\begin{aligned} \left(\frac{1}{e^{-a_{n,\gamma}}} - 1\right) \left(\frac{1}{a_{n,\gamma}} - 1\right) &\leq \left(\frac{1}{(1 - \frac{a_{n,\gamma}}{n})^n} - 1\right) \left(\frac{1}{a_{n,\gamma}} - 1\right) \\ &= \gamma \\ &= \left(\frac{1}{e^{-a_\gamma}} - 1\right) \left(\frac{1}{a_\gamma} - 1\right). \end{aligned} \tag{25}$$

Moreover, the function $x \mapsto (e^x - 1)(1/x - 1)$ is decreasing on $(0, 1)$. In fact its derivative at any point $x \in (0, 1)$ is

$$\begin{aligned} \frac{d}{dx} \left[\left(\frac{1}{e^{-x}} - 1 \right) \left(\frac{1}{x} - 1 \right) \right] &= \left(\frac{1}{x} - 1 \right) e^x - \frac{e^x - 1}{x^2} \\ &= \frac{1}{x^2} (1 - x^2 - (1 - x)e^x) \\ &= \frac{1 - x}{x^2} (1 + x - e^x) < 0. \end{aligned}$$

It follows from (25) that $a_\gamma \leq a_{n,\gamma}$. Finally, given that $x \mapsto \frac{1 - e^{-x}}{x}$ is non-increasing on $(0, 1]$, we deduce that

$$\frac{1 - (1 - \frac{a_{n,\gamma}}{n})^n}{a_{n,\gamma}} \geq \frac{1 - e^{-a_{n,\gamma}}}{a_{n,\gamma}} \geq \frac{1 - e^{-a_\gamma}}{a_\gamma}.$$

We deduce that the competitive ratio of the algorithm described in Theorem 4.6 is at least $\frac{1 - e^{-a_\gamma}}{a_\gamma} = \frac{1 - p_\gamma}{\log(1/p_\gamma)} = 1 - (1 - \gamma)p_\gamma$. □

D Random decay functions

While we only studied deterministic decay functions in the paper, it is also possible to have scenarios with random decay functions. Consider for example that rejected items remain available after j steps with a probability p_j , this is modeled by $D_j(x) = \xi_j x$ with ξ_j a Bernoulli random variable with parameter p_j . We explain in this section how the definitions and our results extend to this case.

Definition D.1 (Random process). *Let \mathcal{X} is a non-empty set. A random process on \mathcal{X} is a collection of random variables $\{Z_x\}_{x \in \mathcal{X}}$. Two random processes $\mathcal{Z} = \{Z_x\}_{x \in \mathcal{X}}$ and $\mathcal{Z}' = \{Z'_x\}_{x \in \mathcal{X}'}$ are independent if any finite sub-process of \mathcal{Z} is independent of any sub-process of \mathcal{Z}' . For simplicity, let us say that the random processes $\{Z_x^1\}_{x \in \mathcal{X}^1}, \dots, \{Z_x^m\}_{x \in \mathcal{X}^m}$ are mutually independent if, for any $x_1 \in \mathcal{X}_1, \dots, x_m \in \mathcal{X}_m$, the random variables $Z_{x_1}^1, \dots, Z_{x_m}^m$ are mutually independent.*

Definition D.2 (Random decay functions). *Let $\mathcal{D} = (D_1, D_2, \dots)$ be a sequence of mutually independent random processes. We say that \mathcal{D} is a sequence of random decay functions if*

1. $\Pr(D_j(x) \notin [0, x]) = 0$ for any $x \geq 0$ and $j \geq 1$,
2. $j \in \mathbb{N}_{\geq 1} \mapsto \Pr(D_j(x) \geq a)$ is non-increasing for any $x, a \geq 0$,
3. $x \geq 0 \mapsto \Pr(D_j(x) \geq a)$ is non-decreasing for any $j \in \mathbb{N}_{\geq 1}$ and $a \geq 0$.

The second condition asserts that the random variable $D_{j-1}(x)$ has first-order stochastic dominance over $D_j(x)$. Along with the first condition, reflect that the distributions of the rejected values become progressively smaller. The last condition indicates that for any integer $j \geq 1$ and non-negative real numbers $x < y$, $D_j(y)$ has a first-order stochastic dominance over $D_j(x)$, which means that, as the value of x increases, so does the potential recovered value after j steps.

The decision-maker In the \mathcal{D} -prophet inequality with deterministic decay functions, we assumed that the decision-maker has full knowledge of the functions D_1, D_2, \dots . In the randomized setting, we assume instead that the decision-maker knows the distributions of the decay functions, i.e. knows the distribution of the random variables $D_j(x)$ for all $x \geq 0$ and $j \geq 1$. However, they do not observe their values until they decide to stop. The online selection process is therefore as follows: the algorithm knows beforehand the distributions of the decay functions, then at each step, it observes a new item with value X_i , and decides to stop or continue. Once they decide to stop at some time τ , they observe the values $D_1(X_{\tau-1}), \dots, D_\tau(X_1)$ and then they choose the maximal one. As a consequence, the stopping time τ is independent of the randomness induced by the decay functions. As in the deterministic case, the expected reward of any algorithm **ALG** can be written as

$$\mathbb{E}[\mathbf{ALG}^{\mathcal{D}}(X_1, \dots, X_n)] = \mathbb{E} \left[\max_{0 \leq i \leq \tau-1} \{D_i(X_{\tau-i})\} \right].$$

The limit decay A key result in our paper is the reduction of the problem to the case where all the decay functions are identical, and we prove this reduction by considering the pointwise limit of the decay functions. In the case of random decay functions, instead of the pointwise convergence, it holds for all $x \geq 0$ that the random variables $(D_j(x))_j$ converge in distribution to some random variable $D_\infty(x)$. In fact, for any $x \geq 0$ and $a \geq 0$, the sequence $(\Pr(D_j(x) \geq a))_{j \geq 1}$ is non-increasing and non-negative, thus it converges to some constant $G(x, a)$. Given that $x \mapsto \Pr(D_j(x) \geq a)$ is non-decreasing for any j , we obtain by taking the limit $j \rightarrow \infty$ that $x \mapsto G(x, a)$ is non-increasing, and with similar argument we obtain, for any $x \geq 0$, that $G(x, a) = 1$ for all $a \leq 0$ and $G(x, a) = 0$ for all $a > x$. Therefore, $a \mapsto 1 - G(x, a)$ defined the cumulative distribution of a random variable D_∞ such that

- $x \geq 0 \mapsto \Pr(D_\infty(x) \geq a)$ is non-decreasing for all $a \geq 0$,
- $\Pr(D_\infty(x) \notin [0, x]) = 0$ for all $x \geq 0$.

Therefore, for all $x \geq 0$, $D_\infty(x)$ is the limit in distribution of $(D_j(x))_j$, hence a sequence $\mathcal{D}' = (D'_1, D'_2, \dots)$ of mutually independent random processes such that $D'_j(x) \sim D_\infty(x)$ for any $j \geq 1$ and $x \geq 0$ defines a sequence of decay functions. We say in this case that all the decay functions are identically distributed as D_∞ . Moreover, it holds for all $x \geq 0$ that $\mathbb{E}[D_\infty(x)] = \lim_{j \rightarrow \infty} \mathbb{E}[D_j(x)] = \inf_{j \geq 1} \mathbb{E}[D_j(x)]$

From there, all the proofs of Section 2 can be easily generalized to the case of random decay functions, and it follows that we can restrict ourselves to studying identically distributed decay functions. Moreover, Proposition 3.2 can be generalized to the case of random decay functions, and the necessary condition for surpassing $1/2$ becomes $\inf_{x > 0} \frac{\mathbb{E}[D_\infty(x)]}{x} > 0$. Similarly, using that the stopping time τ of the algorithm is independent of randomness induced by D_∞ , Proposition 3.3 remains true with $\gamma = \inf_{x > 0} \frac{\mathbb{E}[D_\infty(x)]}{x}$.

Lower bounds For establishing lower bounds, observe that, for any random decay functions \mathcal{D} , if we denote $H_j(x) = \mathbb{E}[D_j(x)]$ for all x , then $\mathcal{H} = (H_1, H_2, \dots)$ defines a sequence of deterministic decay functions. Furthermore, for any instance X_1, \dots, X_n and any algorithm **ALG**, it holds that

$$\begin{aligned} \mathbb{E}[\mathbf{ALG}^{\mathcal{D}}(X_1, \dots, X_n)] &= \mathbb{E}\left[\max_{0 \leq i \leq \tau-1} \{D_i(X_{\tau-i})\}\right] \\ &= \mathbb{E}\left[\mathbb{E}\left[\max_{0 \leq i \leq \tau-1} \{D_i(X_{\tau-i})\} \mid \tau, X_1, \dots, X_n\right]\right] \\ &\geq \mathbb{E}\left[\max_{0 \leq i \leq \tau-1} \{\mathbb{E}[D_i(X_{\tau-i}) \mid \tau, X_1, \dots, X_n]\}\right] \\ &= \mathbb{E}\left[\max_{0 \leq i \leq \tau-1} \{H_i(X_{\tau-i})\}\right] \\ &= \mathbb{E}[\mathbf{ALG}^{\mathcal{H}}(X_1, \dots, X_n)]. \end{aligned}$$

It follows that lower bounds established for deterministic decay functions can be extended to random decay functions by considering their expectations.

Implications With the previous observations, both the lower and upper bounds, depending on $\gamma_{\mathcal{D}}$ that we proved in the deterministic \mathcal{D} -prophet inequality can be generalized to the random \mathcal{D} -prophet inequality, by taking

$$\gamma_{\mathcal{D}} = \inf_{x > 0} \inf_{j \geq 1} \frac{\mathbb{E}[D_j(x)]}{x}.$$

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: In the introduction, present the problem, and state the main contributions of the paper.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations are discussed in the conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: All the assumptions are clearly stated, and full proofs are in the appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification: No experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification: No experimental results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer:[NA]

Justification: No experiments.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We have reviewed the Code of Ethics, and the paper conforms to it.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The contribution of the paper is theoretical. We do not feel there is major societal impact that needs to be discussed.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The contribution of the paper is mainly theoretical.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: We do not use any assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not provide any new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: No such experiments are included in the paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: No such studies are included in the paper.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.