Online Selection with Uncertain Disruption

Yihua Xu[*†] Süleyman Kerimov[‡†] Sebastian Perez-Salazar[*†]

Abstract

In numerous online selection problems, decision-makers (DMs) must allocate on the fly limited resources to customers with uncertain values. The DM faces the tension between allocating resources to currently observed values and saving them for potentially better, unobserved values in the future. Addressing this tension becomes more demanding if an uncertain disruption occurs while serving customers. Without any disruption, the DM gets access to the capacity information to serve customers throughout the time horizon. However, with uncertain disruption, the DM must act more cautiously due to risk of running out of capacity abruptly or misusing the resources. Motivated by this tension, we introduce the Online Selection with Uncertain Disruption (OS-UD) problem. In OS-UD, a DM sequentially observes n non-negative values drawn from a common distribution and must commit to select or reject each value in real time, without revisiting past values. The disruption is modeled as a Bernoulli random variable with probability p each time DM selects a value. We aim to design an online algorithm that maximizes the expected sum of selected values before a disruption occurs, if any.

We evaluate online algorithms using the *competitive ratio*—the ratio between the expected value achieved by the algorithm and that of an optimal *clairvoyant* algorithm that knows all value realizations in advance but still faces uncertain disruption. Using a quantile-based approach, we devise a *non-adaptive* single-threshold algorithm that attains a competitive ratio of at least $1 - 1/e$, and an *adaptive* threshold algorithm characterized by a sequence of non-increasing thresholds that attains an asymptotic competitive ratio of at least 0.745. Both of these results are worst-case optimal within their corresponding class of algorithms. Our results reveal an interesting connection between the OS-UD problem and the i.i.d. prophet inequality problems as the number of customers grows large.

1 Introduction

Online selection models have gained increasing attention, as they capture key features of extensive applications such as online advertising (Mehta and Panigrahi (2012)), online resource allocation (Delong et al. (2024)), and applicant evaluation (Epstein and Ma (2024)). Generally speaking, in these problems, a decision-maker (DM) must irrevocably allocate limited resources to incoming customers with uncertain values. The DM faces a fundamental tension between allocating resources to currently observed values and saving them for potentially better, unobserved values in the future.

Alleviating this tension becomes more challenging if committing to a request may cause a *disruption*, potentially halting the remaining process. For example, in cloud computing, providers

[*]Department of Computational Applied Mathematics and Operations Research, Rice University, USA.
[†]Ken Kennedy Institute, Rice University, USA.
[‡]Jones Graduate School of Business, Rice University, USA.

such as Amazon Web Services and Azure often offer spot instances at lower prices than on-demand instances by utilizing under-utilized resources. While these providers aim to maximize resource utilization and generate additional revenue through the low-price spot instances, they must also manage the risk of negatively impacting the computational performance of on-demand users entitled to those resources. If the provider's decision affects the primary user entitled to those computational resources, then, as a result, the provider must withhold from using these resources until they are available again (e.g., see Cohen et al. (2019); Perez-Salazar et al. (2022)). In a similar vein, an owner of a reusable resource (e.g., a host renting her property on AirBnB, or a driver accepting a ride on Uber) might experience a disruption in operations due to property damage caused by the customer. In general, this additional uncertainty forces the DM to act more cautiously due to risk of misusing the resources. Motivated by this, we introduce the *Online Selection with Uncertain Disruption* (OS-UD) problem, which captures such disruptions when serving incoming requests.

In the OS-UD problem, a finite sequence of n independent and identically distributed (i.i.d.) non-negative random variables arrive sequentially. Upon observing a value, the DM must decide whether to (i) select it, or, (ii) reject it and observe the next value, if any, without the possibility of reconsidering past observations. However, with a known probability $p \in [0, 1]$, selecting a value disrupts the process, forfeiting this last selection, and halting the remaining process permanently.

Our objective is to design an online algorithm that maximizes the expected sum of selected values before a disruption occurs, if any. We evaluate the performance of algorithms using the *competitive ratio*, which is the fraction between the value of an algorithm and the value of an optimal *clairvoyant algorithm* that *knows the values upfront*, but still faces *unknown disruptions*. The competitive ratio is a number in the interval $[0, 1]$ that measures the "price" paid by the DM for not knowing all the sequence of values upfront. Our definition of competitive ratio differs slightly from the usual definition found in other online selection problems (e.g., see (Correa et al., 2021; Hill and Kertz, 1982; Krengel and Sucheston, 1977; Mehta and Panigrahi, 2012))—a formal description of our problem and discussion is presented in §1.1. Our goal is to shed light on the following questions: (i) What is the optimal (worst-case) competitive ratio for the OS-UD problem? (ii) How should a DM accept incoming requests to achieve the optimal competitive ratio?

In general, without imposing any structure on the disruption process, any online algorithm can perform arbitrarily bad and the optimal competitive ratio can be 0 (see Example 7.1). In this paper, we answer both questions by devising simple threshold policies, and we provide an exhaustive analysis for the case when the probability of disruption p is a fixed constant. We first show that the optimal online algorithm induced by a stochastic dynamic program is characterized by a sequence of non-increasing thresholds, where the algorithm accepts the incoming request if and only if its value is at least the current threshold. Motivated by this, in order to quantify the "price" that the DM must pay due uncertainties, we analyze two classes of algorithms that are easy-to-describe: *non-adaptive (fixed) threshold algorithms* and *adaptive threshold algorithms*. Such algorithms are also desirable in practice, as they are simple-to-implement, have economic interpretations (Arnosti and Ma, 2023; Naor, 1969; Van Mieghem, 1995) and are widely used in posted price mechanisms (Chawla et al., 2007, 2010, 2024; Correa et al., 2019).

For non-adaptive threshold algorithms, we offer a complete characterization of the optimal competitive ratio, which is $1 - 1/e$, through the design of an algorithm and a hard instance. For adaptive threshold algorithms, we present an algorithm that employs a non-increasing sequence of thresholds with a tight asymptotic competitive ratio of $\theta^* \approx 0.745$, where θ^* is a parameter appearing in the Hill and Kertz equation (Hill and Kertz, 1982; Kertz, 1986). Interestingly, θ^* coincides with the optimal competitive ratio of the i.i.d. prophet inequality problem (Correa et al., 2021; Hill and Kertz, 1982; Kertz, 1986).

The techniques derived for the fixed disruption probability case can also be adapted to the case when disruption is rare (where $p = \alpha/n$, for $\alpha < 1$). Notably, one can obtain asymptotic competitive ratios that are strictly larger than θ^* (see §6). After presenting the formal introduction of our problem in the next subsection, we provide the details of our results and techniques in §1.2.

1.1 Problem Formulation

Given a fixed disruption probability $p \in [0,1]$, an instance I of the OS-UD problem is given by a pair (n, F), where $n \geq 1$ is the number of values, and F is a continuous and increasing cumulative distribution function supported in the non-negative real numbers.[1] Let $\mathcal{F}$ be the set of all such cumulative distribution functions. An instance $I = (n, F)$ encodes two sequences of random variables $(X_i)_{i=1}^n$ and $(Y_i)_{i=1}^n$. The X_i's are non-negative and i.i.d. following the distribution F. The Y_i's are are i.i.d. 0/1-random variables with $\Pr(Y_i = 1) = p \in [0,1]$. The DM observes X_i's in the fixed order $1, \ldots, n$, sequentially. Upon observing X_i, the DM can either skip to the next value (if one remains) or select the current one. If the DM selects X_i, DM then observes Y_i indicating whether making this selection will disrupt the whole process. If $Y_i = 1$, we say that a *disruption occurred* and the process terminates. In this case, the DM receives the sum of selected values excluding X_i. If $Y_i = 0$, the process continues, and the DM receives X_i as part of the current round's value. We aim to find an online algorithm $\texttt{ALG}$ that maximizes the total expected value collected. We denote by $v(\texttt{ALG}(I))$ the expected value obtained by algorithm $\texttt{ALG}$ for an instance I.

We pursue a competitive analysis, where we compare the value of an algorithm to an *optimal clairvoyant algorithm* that has access to the value realizations $X_1, \ldots, X_n$, in advance, *but not to the sequence* $Y_1, \ldots, Y_n$. The clairvoyant can choose in which order to observe the X_i's, say $X_{\sigma(1)}, \ldots, X_{\sigma(n)}$, where σ is a permutation of $[n]$. Thus, the optimal clairvoyant algorithm obtains a value of $v(\texttt{OPT}(I)) = \mathbb{E}\left[\max_\sigma \sum_{i=1}^{\min\{D-1,n\}} X_{\sigma(i)}\right] = \mathbb{E}\left[\sum_{i=1}^{\min\{D-1,n\}} X_{(i)}\right]$, where D is a geometric random variable with parameter p,[2] and $X_{(j)}$ denotes the j^{th} largest order statistic. We abuse notation and simply write $\texttt{ALG}$ and $\texttt{OPT}$ for $\texttt{ALG}(I)$ and $\texttt{OPT}(I)$, respectively, when the instance I is clear from the context. For convenience, we will omit the word "online" when referring to online algorithms when the context is clear. Then, given an algorithm $\texttt{ALG}$ with a disruption probability

[1] We can assume that F is continuous and smooth by perturbing the observed values with a small continuous noise. The monotonicity assumption is less common, but follows a similar principle of perturbing the original F (e.g., see (Liu et al., 2020; Perez-Salazar and Verdugo, 2024)).

[2] That is, $\Pr(D = j) = p(1-p)^{j-1}$ for $j \geq 1$.

p, its *competitive ratio* is $\inf_{n\geq 1, F\in\mathcal{F}} v(\texttt{ALG})/v(\texttt{OPT})$, and we seek an algorithm with the largest competitive ratio possible. In the remainder of the paper, we focus on the case when $p \in (0,1)$, since when $p \in \{0,1\}$, any optimal algorithm obtains the same value as the optimal clairvoyant algorithm.

We remark that the clairvoyant algorithm is different from the *offline algorithm* that knows the realizations of X_i's and Y_i's upfront for all. The following example motivates us to use a more refined benchmark for our online algorithms, since it is impossible to obtain a constant competitive ratio when we attempt to replace OPT by the offline algorithm.

Example 1.1. *Consider the following $n \geq 1$ i.i.d. values $X_1, \dots, X_n$, where X_i's are uniformly distributed in $[0,n]$. Fix $p \in (0,1)$. Note that the offline algorithm will select all the values for which $Y_i = 0$. Thus, the expected value obtained by the offline algorithm is $v(\texttt{OPT}^{\textit{Offline}}) := \mathbb{E}[\textit{number of } i\textit{'s with } Y_i = 0] \cdot \mathbb{E}[X_1] = (1-p)n^2/2$. On the other hand, we have $v(\texttt{OPT}) \leq n/p$. Then, $v(\texttt{ALG})/v(\texttt{OPT}^{\textit{Offline}}) \leq v(\texttt{OPT})/v(\texttt{OPT}^{\textit{Offline}}) \leq 2/(p(1-p)n)$, for any algorithm ALG. This shows that, for any fixed $p \in (0,1)$, no constant competitive ratio is possible when we compare $v(\texttt{ALG})$ against the value of the offline algorithm that knows all the information in advance.*

Note that one can formulate the DM's problem as a stochastic dynamic program (see §3). However, analyzing competitive ratios is not straightforward as we are comparing two algorithms that operate with asymmetric information. Nevertheless, the dynamic program reveals a nice structural property that we will leverage in our analysis. Indeed, given an instance $I = (n, F)$ of the OS-UD problem, there exists an optimal algorithm that employs thresholds $\tau_1 \geq \cdots \geq \tau_n$ such that if the i^{th} observed value X_i is at least τ_i, then the algorithm selects it (see Proposition 3.1). This motivates us to focus on the two extremes for this class of threshold-based algorithms: (1) NA, which is the class of *non-adaptive algorithms*, where $\tau_1 = \cdots = \tau_n$; and (2) AD, which is the class of *adaptive algorithms*, where we do not constrain the thresholds to be the same, but the thresholds must be non-increasing, i.e., $\tau_1 \geq \cdots \geq \tau_n$. For a fixed $p \in (0,1)$ and a class of algorithms $\mathcal{C} \in \{\mathrm{NA}, \mathrm{AD}\}$, we define the competitive ratio for instances of length n in class $\mathcal{C}$ via:

$$\gamma_{n,p}^{\mathcal{C}} := \sup_{\texttt{ALG} \text{ in } \mathcal{C}} \inf_{F\in\mathcal{F}} \frac{v(\texttt{ALG}(n,F))}{v(\texttt{OPT}(n,F))}.$$

Given that $\mathcal{C}$ is a class of threshold-based algorithms, we can exchange the supremum with the infimum in the definition of $\gamma_{n,p}^{\mathcal{C}}$ without altering its value. Thus, the largest competitive ratio of any algorithm in class $\mathcal{C}$ for the OS-UD problem can be written as follows:

$$\gamma_p^{\mathcal{C}} := \sup_{\texttt{ALG} \text{ in } \mathcal{C}} \inf_{n\geq 1, F\in\mathcal{F}} \frac{v(\texttt{ALG})}{v(\texttt{OPT})} = \inf_{n\geq 1} \gamma_{n,p}^{\mathcal{C}}.$$

1.2 Our Technical Contributions

Our first result provides a tight performance bound for the non-adaptive threshold algorithms.

Theorem 1.1. *For any $p \in (0,1)$, we have $\gamma_p^{\mathrm{NA}} = 1 - 1/e$.*

To prove Theorem 1.1, we explicitly construct a fixed threshold algorithm with a competitive ratio of at least $1-1/e$, through focusing on a subclass of NA, the *quantile-based* threshold algorithms. Algorithms in this subclass are parametrized by $q \in [0,1]$, which we refer as the quantile of the algorithm. An algorithm with a quantile q computes a threshold τ via $q = \Pr(X \geq \tau) = 1 - F(\tau)$, and selects any value in the input sequence that is at least τ.

To prove that the competitive ratio of an algorithm ALG is at least θ, one can use standard stochastic dominance arguments, and it is sufficient to show that $\Pr[\texttt{ALG}$ obtains value of at least $\tau]$ is at least $\theta \cdot \Pr[\texttt{OPT}$ obtains value of at least $\tau]$ for any $\tau \geq 0$. However, due to possibility of making multiple selections, analyzing both of these quantities is not evident. Instead, we characterize the optimal values for both the clairvoyant and the quantile-based threshold algorithms as integrals of the inverse of F. From here, we can deduce a lower bound for $\gamma_{n,p}^{\mathrm{NA}}$ that only depends on n and p, and not dependent on F. This lower bound turns out to be monotonically decreasing in n and converging to $1-1/e$ providing the desired lower bound on γ_p^{NA}. To prove an upper bound on γ_p^{NA}, we explicitly construct an instance (n, F) such that $\gamma_{n,p}^{\mathrm{NA}} \leq 1 - 1/e + o(1)$, and since $\gamma_p^{\mathrm{NA}} = \inf_{n\geq 1} \gamma_{n,p}^{\mathrm{NA}}$, the result follows. We provide all the details for the proof in §4.

We then turn our attention to the performance of the general class of adaptive threshold algorithms, and we show that *adaptivity* is key. Our second result shows that a competitive ratio better than $1-1/e$ is possible for such algorithms.

Theorem 1.2. *For any $p \in (0,1)$, $\liminf_{n\to\infty} \gamma_{n,p}^{AD} \geq \theta^* \approx 0.745$, where $\beta = 1/\theta^*$ is the unique solution to the integral equation $\int_0^1 (y - y\ln y + (\beta - 1))^{-1}\,\mathrm{d}y = 1$.*

Theorem 1.2 shows that there is a strict separation between acting adaptive and non-adaptive. In order to prove this result, we provide a quantile-based algorithm similar to Correa et al. (2021) and Perez-Salazar et al. (2025). For each observed value i, we sample a quantile q_i from an appropriate distribution, and we compute the threshold τ_i such that $\Pr(X \geq \tau_i) = 1 - F(\tau_i) = q_i$. The algorithm then selects i if the observed value X_i is at least τ_i. By an appropriate choice of distribution for the quantiles, we show that the value of the algorithm is asymptotically a fraction $\theta^* \approx 0.745$ of $v(\texttt{OPT})$. We present the algorithm and its analysis in §5.

Our final result establishes a limit on the competitive ratio attainable by any algorithm for the OS-UD problem, contributing another piece to the puzzle of answering our first research question.

Theorem 1.3. *For any $p \in (0,1)$, we have $\gamma_p^{AD} \leq \theta^*$, where θ^* is defined in Theorem 1.2.*

We derive the upper bound by adapting the worst-case instance for the i.i.d. prophet inequality problem from (Hill and Kertz, 1982; Liu et al., 2020) to OS-UD. We present the details in §5.

2 Related Work

Online selection, as well as the tension between collecting short-term values and saving resources for long-term values, have been extensively studied in computer science and operations research literature through the lens of optimal stopping theory (e.g., see Hill and Kertz (1992); Krengel and

Sucheston (1977); Shiryaev (2007)), and for their broad applicability on various practical problems from crowd-sourcing (Mehta and Panigrahi, 2012) to capital investment problem (Goyal and Ravi, 2010). Here, we present several streams of literature that are closely related to our work.

Prophet Inequality. One related stream of literature to our work is the prophet inequality problem (e.g., see Krengel and Sucheston (1977); Samuel-Cahn (1984)), in particular, the case with i.i.d. values. In the i.i.d. prophet inequality problem (Hill and Kertz, 1982), a DM must select at most one value from a sequence of i.i.d. randomly generated values, and her goal is to design an algorithm with a large competitive ratio, where the offline benchmark is the expected maximum of the sequence of values. It is known that the optimal competitive ratio for the i.i.d. prophet inequality problem is $\theta^* \approx 0.745$, which is the unique parameter appearing in the Hill & Kertz equation (Correa et al., 2021; Hill and Kertz, 1982; Kertz, 1986). This optimal algorithm can be attained by a quantile-based algorithm that depends solely on n and is independent of the specific instance distribution. To facilitate the analysis of the algorithm, the competitive ratio can be retrieved through a unique solution of the Hill and Kertz equation. Similar techniques to using ODEs have also been explored in various recent work (Correa et al., 2021; Liu et al., 2020). By generalizing the Hill and Kertz equation, Brustle et al. (2025) introduces a novel non-linear system of differential equations and provide tight analysis for the k-prophet inequality problem. Our paper extends the single-selection prophet setting by incorporating an additional disruption indicator, and we use a similar quantile-based approach to develop an adaptive threshold algorithm, which achieves an asymptotic competitive ratio of θ^* as well.

Random horizon. There is large body of work in optimal stopping problems with random horizon (Hajiaghayi et al., 2007; Zhang and Jaillet, 2023). Here, the disruption is caused by not knowing the length n of values upfront. For example, the uncertain horizon setting has been extensively studied within the framework of the secretary problem. When no distributional information is available regarding the disruption time—beyond which no further applicants can be picked—it is known that no algorithm can achieve a constant competitive ratio (Hill and Krengel, 1991). However, if a random termination time with a known value-independent distribution exists, a conditionally optimal selection rule can be formulated (Samuel-Cahn, 1996). In our case, the disruption is caused potentially by selecting a value, and in principle we could observe the whole sequence of n values. Closer to our work is Alijani et al. (2020), which studies the prophet inequality problem with supply uncertainty. Even though the OS-UD problem has applications in settings with supply uncertainty, our main focus is in applications where serving a request can disrupt the remaining selection process. In a similar vein, our model is loosely related to the *stochastic knapsack* problems (Dean et al., 2008; Ma, 2018), where items have unknown stochastic sizes, items are packed sequentially, and if the knapsack is overflowed, then the whole process stops. We could regard the OS-UD problem as a knapsack problem, where we have a knapsack of capacity 1, and each item has two possible sizes: size $1/n$ with probability $1 - p$ and size $1 + 1/n$ with probability p. Nevertheless, the knapsack literature mostly deals with approximation algorithms as opposed to competitive analysis that we pursue in this work.

Dynamic matching and online resource allocation. The tension between committing to a decision now and delaying decisions in anticipation of better opportunities arises as an inherent trade-off in many stochastic models. In the context of dynamic matching and online resource allocation, recent work addressed this trade-off through the lens of an all-time regret notion (Gupta, 2024; He et al., 2025; Kerimov et al., 2024; Wei et al., 2023) and approximation algorithms (Aouad and Saritaç, 2020). In particular, the all-time regret notion implicitly deals with uncertain disruption by guaranteeing near-optimal performance throughout the time horizon. Our work studies the same trade-off in the context of online selection by explicitly introducing a disruption indicator, and while there are differences across these stochastic models, we hope that the modeling we propose in this work can be leveraged to study this fundamental tension under disruptions such as match rejections and item returns.

3 Preliminaries

In this section, we present some preliminaries needed in the remaining of the paper. All missing proofs in the main body are deferred to the end of the paper. The value maximization problem faced by the DM can be solved by means of stochastic dynamic programming. For any $n \geq 1$, disruption parameter $p \in (0,1)$, and distribution $F \in \mathcal{F}$, let $D_i(p,F)$ be the optimal value obtainable from the sequence $(X_j)_{j=i}^n$, when $i-1 \leq n$ values have been observed already (i.e., the i^{th} value is ready to be observed next) and no disruption has occurred yet. Then clearly $D_{n+1}(p,F) = 0$, and for $i \leq n$, we have

$$D_i(p,F) = \max_{\tau_i \geq 0} \left\{ (1-p) \Pr[X \geq \tau_i] \left(\mathbb{E}[X | X \geq \tau_i] + D_{i+1}(p,F) \right) + \Pr[X < \tau_i] D_{i+1}(p,F) \right\}. \quad \text{(DP)}$$

By solving the recursion (DP), we can obtain an optimal algorithm that at time $i \in [n]$, selects the observed value X_i if it is at least τ_i, where τ_i is the maximizer of (DP). The next proposition states that it is optimal to use a non-increasing sequence of thresholds, i.e., $\tau_1 \geq \cdots \geq \tau_n$.

Proposition 3.1. *There exists an optimal algorithm for the* OS-UD *problem that employs a non-increasing sequence of thresholds.*

Next, we provide two characterizations of $v(\texttt{OPT})$ in terms of the inverse of F and its derivative. These characterizations will be useful in the analysis of NA and AD classes.

Proposition 3.2. *The value achieved by the optimal algorithm can be characterized as follows:*

$$v(\texttt{OPT}) = \int_0^1 B_n(p,v) r(v) dv = \int_0^1 F^{-1}(1-q) g_n(p,q) dq, \tag{1}$$

where $B_n(p,v) := \mathbb{E}[\min\{X, D-1, n\}]$ *with* $X \sim Bin(n,v)$, $r(v) > 0$ *is a function such that* $\int_u^1 r(v)dv = F^{-1}(1-u)$,[3] *and* $g_n(p,q) := (1-p) \cdot n(1-pq)^{n-1}$.

[3] The existence of $r(v)$ is guaranteed by our assumption of F^{-1} being differentiable and strictly decreasing.

Finally, we reproduce the ordinary differential equation (ODE) system introduced by Hill and Kertz (1982), colloquially known as the Hill and Kertz equation. This will be used in the analysis of the AD class. We want to find a solution $y : [0, 1] \to [0, 1]$ to the following ODE:

$$
\begin{aligned}
&y'(t) = y(t)(\ln(y(t)) - 1) - \left(\frac{1}{\theta^*} - 1\right), \\
&y(0) = 1, \quad \lim_{t \uparrow 1} y(t) = 0.
\end{aligned}
$$

This system has a solution if and only if $\theta^* \approx 0.745$, where θ^* the unique solution to the integral equation in Theorem 1.2.

4 The Class of Non-Adaptive Algorithms

In this section, we analyze the performance of non-adaptive threshold algorithms. We fully characterize the competitive ratio achievable by this class, presenting an optimal algorithm that guarantees a competitive ratio of at least $1 - 1/e$ for any input to the OS-UD problem. Moreover, we show that no non-adaptive algorithm can achieve a competitive ratio exceeding $1 - 1/e + o(1)$ in the worst case. Together, these results imply Theorem 1.1.

4.1 Quantile-Based Non-Adaptive Algorithm

Our algorithm is quantile-based. It receives a quantile $q \in [0, 1]$, sets the threshold $\tau = F^{-1}(1-q)$, and selects any value of at least τ. For any $n \geq 1$, we denote the algorithm with quantile q_n by $\mathtt{ALG}^{q_n}$ and its expected value by $v(\mathtt{ALG}^{q_n})$. The main result of this section is the following:

Theorem 4.1. *For any $n \geq 1$, if $q_n = \min\{1, 1/pn\}$, we have $\frac{v(\mathtt{ALG}^{q_n})}{v(\mathtt{OPT})} \geq 1 - \frac{1}{e}$.*

Theorem 4.1 immediately implies $\gamma_p^{NA} \geq 1 - 1/e$, proving the first part of Theorem 1.1. The proof relies on two key lemmas. The first lemma provides a lower bound on $v(\mathtt{ALG}^{q_n})/v(\mathtt{OPT})$ that is independent of F.

Lemma 4.2. *For any $n \geq 1$ and $p \in (0, 1)$, we have*

$$
\frac{v(\mathtt{ALG}^{q_n})}{v(\mathtt{OPT})} \geq \frac{(1 - (1 - q_n p)^n)}{p} \cdot \min\left\{\frac{p}{1 - (1-p)^n}, \frac{1}{q_n n}\right\}, \tag{2}
$$

for any instance of the OS-UD *problem.*

Let $\eta_{n,p}^{q_n}$ denote the right-hand side of (2). The next lemma provides a useful monotonicity property.

Lemma 4.3. *For any $n \geq 1$, let $q_n = \min\{1, 1/pn\}$. Then, we have $\eta_{n,p}^{q_n} \geq \eta_{n+1,p}^{q_{n+1}}$. Furthermore, $\lim_{n\to\infty} \eta_{n,p}^{q_n} = 1 - 1/e$.*

With these two lemmas, we are ready to provide the proof of Theorem 4.1. Indeed, for any $k \geq 1$, we obtain $v(\mathtt{ALG}^{q_n})/v(\mathtt{OPT}) \geq \eta_{n,p}^{q_n} \geq \eta_{n+k,p}^{q_{n+k}}$, where the first inequality simply follows from Lemma

4.2, and the second inequality follows from applying the monotonicity in Lemma 4.3 iteratively k times. Thus, we obtain $v(\texttt{ALG}^{q_n})/v(\texttt{OPT}) \geq \lim_{k\to\infty} \eta_{n+k,p}^{q_{n+k}} = 1 - 1/e$ for any $n \geq 1$. We use the rest of this subsection providing the proofs of Lemma 4.2 and 4.3.

Proof of Lemma 4.2. Fix $n \geq 1$, and we drop the index n from quantile for simplicity. We start characterizing $v(\texttt{ALG}^q)$ as follows:

$$v(\texttt{ALG}^q) = \sum_{a=1}^{n} \Pr[\text{Bin}(n,q) = a] \left(\sum_{i=1}^{a} p \cdot (1-p)^{i-1} \cdot (i-1) + a \cdot (1-p)^a \right) \cdot \mathbb{E}[X | X \geq \tau]$$
$$= A_n(q,p) \frac{\int_\tau^\infty x \cdot f(x)dx}{\Pr[X \geq \tau]} = A_n(q,p) \frac{\int_0^q F^{-1}(1-u)du}{q} = A_n(q,p) \frac{\int_0^1 r(v) \min\{v,q\}dv}{q}, \tag{3}$$

where the second equality uses $A_n(q,p) = (1-p) \cdot (1-(1-qp)^n)/p$, which is a simplification from preceding expression; the third equality uses the assumption that F is strictly increasing and substitutes $F(x)$ by $1-u$, and the last equality is achieved by setting $F^{-1}(1-u) = \int_u^1 r(v)dv$, which is valid by our assumption on F. Then, per Proposition 3.2, we have

$$\frac{v(\texttt{ALG}^q)}{v(\texttt{OPT})} = \frac{(A_n(q,p)/q) \int_0^1 r(v) \min\{v,q\}dv}{\int_0^1 B_n(p,v) r(v) dv} \geq \inf_{v \in [0,1]} \frac{A_n(q,p) \cdot \min\{v,q\}}{B_n(p,v) \cdot q}. \tag{4}$$

For $v \geq q$, the ratio inside the infimum becomes $A_n(p,q)/B_n(p,v)$ which is a decreasing function in v. For $v < q$, the ratio now becomes

$$\frac{A_n(q,p) \cdot v}{B_n(p,v) \cdot q} = \frac{A_n(q,p)}{q} \frac{v}{\sum_{i=2}^n p \cdot (1-p)^{i-1} \sum_{j=1}^{i-1} \Pr[\text{Bin}(n,v) \geq j] + (1-p)^n \cdot nv}. \tag{5}$$

In Lemma 8.1, we show that (5) is non-decreasing in v. This implies that the infimum is attained either when $v \to 0$ or $v \to 1$. Thus,

$$\frac{v(\texttt{ALG}^q)}{v(\texttt{OPT})} \geq \min \left\{ \lim_{v\to 0} \frac{A_n(q,p)v}{B_n(p,q)q}, \lim_{v \to 1} \frac{A_n(p,q)}{B_n(p,q)} \right\}$$
$$= \frac{(1-p)(1-(1-qp)^n)}{p} \cdot \min \left\{ \frac{1}{\mathbb{E}[\min\{D-1,n\}]}, \frac{1}{(1-p)qn} \right\},$$

where the second line follows by a straightforward calculation of the corresponding limits. Here $D \sim \text{Geom}(p)$, and a direct calculation shows that $\mathbb{E}[\min\{D-1,n\}] = (1-p)(1-(1-p)^n)/p$. This concludes the proof. □

Proof of Lemma 4.3. We analyze three cases: (i) $pn < p(n+1) < 1$; (ii) $pn \leq 1 < p(n+1)$; and (iii) $1 < pn < p(n+1)$.

Case (i). Here, we have $q_n = q_{n+1} = 1$. Then, $\eta_{n,p}^1 = \min\left\{1, \frac{1-(1-p)^n}{pn}\right\} = \frac{1-(1-p)^n}{pn}$, where the first equality follows simply by definition, while the second equality follows from the Bernoulli's inequality $(1-p)^n \geq 1 - pn$. From here, it is immediate that $\eta_{n+1,p}^1 \leq \eta_{n,p}^1$.

Case (ii). Here, we have $q_n = 1$ and $q_{n+1} = 1/p(n+1)$. Then, $\eta_{n,p}^{q_n} = \frac{1-(1-p)^n}{pn}$, and

$$\eta_{n+1,p}^{q_{n+1}} = \left(1 - \left(1 - \frac{1}{n+1}\right)^{n+1}\right) \min\left\{1, \frac{1}{1-(1-p)^{n+1}}\right\} = 1 - \left(1 - \frac{1}{n+1}\right)^{n+1}$$

with calculations analogous to the previous case. Note that the function $(1-(1-p)^n)/p = \sum_{\ell=0}^{n-1}(1-p)^\ell$ is decreasing in p. Thus, we have $p \in [1/(n+1), 1/n]$, and

$$\eta_{n,p}^{q_n} \geq 1 - \left(1 - \frac{1}{n}\right)^n \geq 1 - \left(1 - \frac{1}{n+1}\right)^{n+1} = \eta_{n+1,p}^{q_{n+1}}.$$

Case (iii). In this case, we have $q_n = 1/pn$ and $q_{n+1} = 1/p(n+1)$. Thus, $\eta_{n,p}^{q_n} = 1 - \left(1 - \frac{1}{n}\right)^n$ and $\eta_{n+1,p}^{q_{n+1}} = 1 - \left(1 - \frac{1}{n+1}\right)^{n+1}$. As a byproduct of the analysis from the previous case, we have already proved $\eta_{n,p}^{q_n} \geq \eta_{n+1,p}^{q_{n+1}}$. To conclude, for $n > 1/p$ we have $q_n = 1/pn$, and so $\lim_n \eta_{n,p}^{q_n} = \lim_n 1 - \left(1 - \frac{1}{n}\right)^n = 1 - \frac{1}{e}$. ☐

4.2 Upper Bound on Competitive Ratio

In this subsection, we prove $\gamma_p^{NA} \leq 1 - 1/e$ to complete the proof of Theorem 1.1. To do so, we provide a family of instances with competitive ratios approaching $1 - 1/e$. For this hard instance, we define the distribution through its inverse as follows:

$$F^{-1}(1-u) = \frac{a_1}{n}\delta_{\{0\}}(u) + a_2\mathbb{I}_{(0,\beta/n]}(u), \tag{6}$$

where $a_1, a_2, \beta \in \mathbb{R}$, $\beta \leq n$,$\delta_{\{0\}}(u)$ denotes the Dirac delta function centered at $u = 0$. and $\mathbb{I}$ is a characteristic function. Denote by Poisson(β) a Poisson distribution with parameter β. We use the following two lemmas to characterize the value of the optimal algorithm and the value of the single threshold algorithm.

Lemma 4.4. *When $n \to \infty$, $v(\texttt{OPT}) \to a_1(1-p) + a_2\sum_{j=1}^\infty \Pr[Poisson(\beta) \geq j](1-p)^j$.*

Lemma 4.5. *Let $\lambda = \lambda^* > 0$ be a solution to the equation*

$$e^{-\lambda p}\left(a_1 + a_1p\lambda + a_2p\lambda^2\right) = a_1. \tag{7}$$

As n grows large, the expected value of $\texttt{ALG}^{\lambda/n}$ converges to

$$\lim_{n\to\infty} v(\texttt{ALG}^{\lambda/n}) = \max\left\{(1-p)a_1,\ \frac{(1-p)(1-e^{-\beta p})}{\beta p}(a_1 + a_2\beta),\ C_1\right\},$$

where $C_1 = \frac{(1-p)(1-e^{-\lambda^* p})}{\lambda^* p}(a_1 + a_2\lambda^*)$ if $\lambda^* \leq \beta$ and $C_1 = 0$ otherwise.

Let $f(a_1, a_2, \beta, p, \lambda) := v(\texttt{ALG}^{\frac{\lambda}{n}})/v(\texttt{OPT})$. We aim to bound $\min_{a_1,a_2,\beta,p}\max_\lambda f(a_1, a_2, \beta, p, \lambda)$, which will provide an upper bound on the competitive ratio.

Theorem 4.6. *Given any $\epsilon > 0$, for given inputs $a_1, a_2, \beta, p, \lambda$ which satisfy $a_2 = p(e-2)a_1$, and β sufficiently large, we have $v(\mathtt{ALG}^{\frac{\lambda}{n}})/v(\mathtt{OPT}) \leq 1 - 1/e + \epsilon$.*

Proof. Given $\beta > 1/p$ sufficiently large, plugging in $a_2 = p(e-2)a_1$ into (7) we obtain the following $e^{-\lambda^* p}\left(a_1 + a_1 p\lambda^* + p^2(e-2)a_1(\lambda^*)^2\right) = a_1$, which holds if and only if $\lambda^* = 1/p$ for $\lambda^* > 0$. Note that when $\lambda^* = 1/p$, through direct comparison and monotonicity analysis, we get

$$C_1 \geq \max\left\{(1-p)a_1, \frac{1-p}{\beta p}(1-e^{-\beta p})(a_1 + a_2\beta)\right\}.$$

Consequently, for the input parameters $a_1, a_2, \beta, p, \lambda$ satisfying $a_2 = p(e-2)a_1$, and β sufficiently large, the value $\lambda = 1/p$ serves as the maximizer of the function $f(a_1, a_2, \beta, p, \lambda)$. Then for any given $\epsilon > 0$, we have

$$\frac{v(\mathtt{ALG}^{\frac{1}{np}})}{v(\mathtt{OPT})} \leq \frac{(1-p)\cdot(1-1/e)\cdot(a_1 + a_2/p)}{a_1(1-p) + a_2\sum_{j=1}^{\infty}\Pr[\mathrm{Poisson}(\beta) \geq j](1-p)^j}.$$

Note that the desired upper bound of $1 - 1/e + \epsilon$ is only achieved when the necessary condition of $\sum_{j=1}^{\infty}\Pr[\mathrm{Poisson}(\beta) \geq j](1-p)^{j-1} \geq 1/p - \epsilon$ is met. We can find a sufficiently large β that ensures the validity of the above inequality. The process of finding such a β consists of the following two steps:

(i) Find an integer C_2 such that $(1-p)^{C_2} \leq \epsilon p/2$. C_2 could then be set as $\lceil \log_{1-p}(\epsilon p/2) \rceil$.

(ii) Using the integer C_2 found in the previous step, we find $\beta^* \geq \lambda$ such that $\Pr[\mathrm{Poisson}(\beta^*) \geq C_2] \geq 1 - \epsilon p/2$. The existence of such a β^* is guaranteed by the intermediate value theorem.

Using conditions (i) and (ii) above, we recover the necessary condition:

$$\begin{aligned}\sum_{j=1}^{\infty}\Pr[\mathrm{Poisson}(\beta) \geq j](1-p)^{j-1} &\geq \sum_{j=1}^{C_2}\left(1 - \frac{\epsilon p}{2}\right)(1-p)^{j-1} = \left(1 - \frac{\epsilon p}{2}\right)\frac{1-(1-p)^{C_2}}{p} \\ &\geq \left(1 - \frac{\epsilon p}{2}\right)\left(\frac{1}{p} - \frac{\epsilon}{2}\right) = \frac{1}{p} - \epsilon + \frac{\epsilon^2 p}{4} \geq \frac{1}{p} - \epsilon.\end{aligned}$$

□

5 The Class of Adaptive Algorithms

In this section, we focus on adaptive algorithms. In §5.1, we formally define our algorithm and present the proof of Theorem 1.2 to show that *adaptivity* is key, and one can improve the competitive ratio from the non-adaptive case in the limit. In §5.2, we present the proof of Theorem 1.3 to provide an upper bound on the competitive ratio, and we show that our derived asymptotic competitive ratio is tight.

5.1 Quantile-Based Threshold Algorithm

For each $i = 1, \dots, n$, our adaptive algorithm samples a quantile $q_i \in [0,1]$ from a density function with a support in $[\varepsilon_{i-1}, \varepsilon_i]$, where $0 = \varepsilon_0 < \varepsilon_1 < \cdots < \varepsilon_n = 1$. Then, if the observed value X_i is at least $F^{-1}(1-q_i)$, the algorithm selects the value; otherwise the value is rejected. Denote this algorithm by $\texttt{ALG}^{\text{AD}}$.

The construction of the densities over each interval $[\varepsilon_{i-1}, \varepsilon_i]$ follows a similar approach to the one outlined in Correa et al. (2021); Perez-Salazar et al. (2025). We take some time to explain the importance of these densities and the challenges involved in applying the method from Correa et al. (2021). From the second characterization of $v(\texttt{OPT})$ in Proposition 3.2, the function $g_n(p,q) = (1-p)n(1-pq)^{n-1}$ is the derivative of $A_n(p,q) = \mathbb{E}[\min\{\text{Bin}(n,q), D-1\}]$, i.e., the derivative of the expected number of values $\texttt{OPT}$ gets at quantile q . The goal of our analysis is to show that, for every $q \in [0,1]$, the derivative of the expected number of values at quantile q that the algorithm accepts is at least $\theta_n \cdot g_n(p,q)$. If this condition holds, then we can guarantee that $v(\texttt{ALG}^{\text{AD}}) \geq \theta_n \cdot v(\texttt{OPT})$. This analysis is quite stringent, as it requires specifying a *valid density for every* $q \in [0,1]$. However, the method from Correa et al. (2021) applied to the OS-UD probem only provides a density over the interval $[0,p]$, leaving $(1-p,1]$ unassigned. To address this limitation, we use a direct approach and construct densities that cover completely the interval $[0,1]$ in such a way that the competitive ratio of $\texttt{ALG}^{\text{AD}}$ converges to θ^*.

We now explain the density functions for $\texttt{ALG}^{\text{AD}}$. Let $0 = \varepsilon_0 < \varepsilon_1$. For $\theta_n > 0$, consider the following function $\beta_{1,n}(p,q) = -\theta_n \mathbb{I}_{[\varepsilon_0,\varepsilon_1]}(q) g_n'(p,q)/(1-p)$. Note that $\beta_{1,n} \geq 0$. If we want to sample q_1 from $\beta_{1,n}(p,\cdot)$, then we must have $1 = \int_0^1 \beta_{1,n}(p,q)\,dq$ which happens if and only if $\frac{1}{n\theta_n} = 1 - (1-p\varepsilon_1)^{n-1}$. From here, ε_1 is decreasing in θ_n; thus, there is θ_n such that $\varepsilon_1 \leq 1$. In general, let $\beta_{i,n}(p,q) = -\theta_n \mathbb{I}_{[\varepsilon_{i-1},\varepsilon_i]}(q) g_n'(p,q)/(1-p)$ such that the following system is satisfied for $0 = \varepsilon_1 < \varepsilon_2 < \cdots < \varepsilon_k \leq 1$, and for the largest k possible:

$$\int_0^1 \beta_{1,n}(p,q)\,dq = 1, \tag{8}$$

$$\int_0^1 \beta_{i+1,n}(p,q)\,dq = \int_0^1 \beta_{i,n}(p,q)(1-pq)\,dq \qquad \forall i < k. \tag{9}$$

We already know that we can satisfy this system with $k = 1$. We seek to satisfy this system for $k = n$ and $\varepsilon_n = 1$. Then, our densities for $\texttt{ALG}^{\text{AD}}$ become $\beta_{i,n}(p,q)/\int_0^1 \beta_{i,n}(p,q)\,dq$ for all $i \in [n]$.

Lemma 5.1. *There is a unique $\theta_n > 0$ such that the system (8) – (9) has a solution for $k = n$ and $\varepsilon_n = 1$.*

We present the proof of Lemma 5.1 after establishing the following guarantee on the competitive ratio of $\texttt{ALG}^{\text{AD}}$.

Theorem 5.2. *Let $\theta_n > 0$ as in Lemma 5.1. Using the densities $\beta_{i,n}/\int_0^1 \beta_{i,n}(p,q)\,\mathrm{d}q$ for $\texttt{ALG}^{AD}$ guarantees*

$$\frac{v(\texttt{ALG}^{AD})}{v(\texttt{OPT})} \geq \left(1 - \frac{(1-p)^{n-1}pn}{1-(1-p)^n}\right) \cdot \theta_n. \tag{10}$$

Proof. Let r_i be the probability that the algorithm observes X_i, $i \in [n]$. Then, we have

$$v(\texttt{ALG}^{\mathrm{AD}}) = \sum_{i=1}^{n} r_i \int_0^1 \frac{\beta_{i,n}(p,u)}{\int_0^1 \beta_{i,n}(p,q)\,\mathrm{d}q}(1-p)\int_{F^{-1}(1-u)}^{\infty} x\,dF(x)\,du,$$

and via induction, we have $r_1 = 1$ and for $i > 1$, $r_i = \int_0^1 \beta_{i-1,n}(p,q)(1-pq)\,dq$. Hence, using the system (8) – (9), we obtain

$$\begin{aligned}
v(\texttt{ALG}^{\mathrm{AD}}) &= \sum_{i=1}^{n} \int_0^1 \beta_{i,n}(p,u)\cdot(1-p)\int_0^q F^{-1}(1-q)dqdu && (11)\\
&= \theta_n \int_0^1 F^{-1}(1-q)g_n(p,q)dq - \theta_n \int_0^1 F^{-1}(1-q)\cdot n(1-p)^n dq, && (12)
\end{aligned}$$

where (11) follows from a change of variable and using $r_i = \int_0^1 \beta_{i,n}(p,q)\,dq$ for $i > 1$. Then, using a ratio comparison, we have

$$\frac{\int_0^1 F^{-1}(1-q)g_n(p,1)\,\mathrm{d}q}{\int_0^1 F^{-1}(1-q)g_n(p,q)\mathrm{d}q} \leq \frac{(1-p)^{n-1}pn}{1-(1-p)^n} \leq 1,$$

and applying this bound in (12), we obtain

$$v(\texttt{ALG}^{\mathrm{AD}}) \geq \theta_n \cdot \left(1 - \frac{(1-p)^{n-1}pn}{1-(1-p)^n}\right)\int_0^1 F^{-1}(1-u)g_n(p,u)\,\mathrm{d}u = \theta_n \cdot \left(1 - \frac{(1-p)^{n-1}pn}{1-(1-p)^n}\right) v(\texttt{OPT}),$$

which concludes the proof. □

We now present the proof of Lemma 5.1. The idea is to generalize the monotonicity of ε_1 as a function of θ_n to ε_i for all i. To this end, we first present an alternative characterization of the system (8) – (9).

Proof of Lemma 5.1. We start with an intermediate result. Given a fixed $k \leq n$, we claim that for all $i \leq k$, the following recursion holds:

$$g_n(p,\epsilon_i) - g_n(p,\epsilon_{i-1}) = -\frac{1-p}{\theta_n} - p\left(\epsilon_{i-1}\cdot g_n(p,\epsilon_{i-1}) - \int_0^{\epsilon_{i-1}} g_n(p,q)dq\right). \tag{13}$$

We proceed by induction. The base case $i = 1$ can be verified easily. Assuming that (13) holds for some $i < k$, by the fundamental theorem of calculus, we have

$$\begin{aligned}
g_n(p,\epsilon_{i+1}) - g_n(p,\epsilon_i) &= \int_{\epsilon_i}^{\epsilon_{i+1}} g_n'(p,q)dq = \int_{\epsilon_{i-1}}^{\epsilon_i} g_n'(p,q)(1-pq)dq && \text{(Equation (9))}\\
&= \int_{\epsilon_{i-1}}^{\epsilon_i} g_n'(p,q)dq - p\int_{\epsilon_{i-1}}^{\epsilon_i} g_n'(p,q)qdq\\
&= g_n(p,\epsilon_i) - g_n(p,\epsilon_{i-1}) - p\epsilon_i\cdot g_n(p,\epsilon_i) + p\epsilon_{i-1}\cdot g_n(p,\epsilon_{i-1}) + p\int_{\epsilon_{i-1}}^{\epsilon_i} g_n(p,q)dq
\end{aligned}$$

$$= -\frac{1-p}{\theta_n} + p\int_0^{\epsilon_i} g_n(p,q)dq - p\epsilon_i \cdot g_n(p,\epsilon_i)$$

where the last equality comes from plugging in the induction hypothesis.

Note that from (13), we obtain that $\varepsilon_i < \varepsilon_{i+1}$ and $g_n'(p,\varepsilon_{i+1})\varepsilon_{i+1}' = \frac{1-p}{\theta_n^2} + g_n'(p,\varepsilon_i)(1-p\varepsilon_i)\varepsilon_i'$, where the derivative is with respect to θ_n. From here, we see that all ε_i's are decreasing in θ_n. Thus, by making θ_n sufficiently large, we can sequentially define $\varepsilon_{k+1}, \varepsilon_{k+2}, \cdots \leq 1$ until we reach $\varepsilon_n = 1$. This concludes the proof of Lemma 5.1. $\square$

Asymptotic Analysis. We now show that $\liminf_{n\to\infty} \gamma_{n,p}^{\mathrm{AD}} \geq \theta^* \approx 0.745$. Let $f_n(p,\lambda) := g_n(p,\lambda/n)/(1-p)n$. Then, (13) becomes

$$\frac{f_n(p,\lambda_i) - f_n(p,\lambda_{i-1})}{1/n} = -\frac{1}{\theta_n} - p\left(\lambda_{i-1}\cdot f_n(p,\lambda_i) - \int_0^{\lambda_{i-1}} f_n(p,w)dw\right), \tag{14}$$

where $\lambda_i = n\varepsilon_i$ for all i. Note that $\theta_n \in [0,1]$. Then, there exists a subsequence that converges to some $\hat{\theta} \in [0,1]$. For simplicity, we abuse notation and denote this subsequence by θ_n so that $\theta_n \to \hat{\theta} \in [0,1]$. Now, doing a linear piece-wise approximation of λ_i via a function $\lambda_n(x)$'s such that $\lambda_i = \lambda_n(i/n)$ and taking the limit in n of (14), we obtain

$$\begin{aligned} f(p,\lambda(x))' &= -\frac{1}{\hat{\theta}} - p\left(\lambda(x)\cdot f(p,\lambda(x)) - \int_0^{\lambda(x)} f(p,w)dw\right) \\ &= 1 - \frac{1}{\hat{\theta}} - p\lambda(x)e^{-\lambda(x)p} - e^{-\lambda(x)p}, \end{aligned}$$

where $f(p,\lambda) = \lim_{n\to\infty} f_n(p,\lambda) = e^{-\lambda p}$. Furthermore, we have conditions $\lambda(0) = 0$ and $\lim_{x\to 1}\lambda(1) = +\infty$. Performing the change of variable $y(x) = e^{-\lambda(x)p}$, we obtain the following system:

$$\begin{aligned} y'(x) &= 1 - \frac{1}{\hat{\theta}} + (\ln y(x) - 1)y(x), \\ y(0) &= 1, \quad \lim_{x\to 1} y(x) = 0. \end{aligned}$$

This is exactly the Hill and Kertz equation presented in §3, which has a solution if and only if $\hat{\theta} = \theta^* \approx 0.745$, which yields that $\liminf_n \theta_n = \theta^*$. Finally per Theorem 5.2, for any $p \in (0,1)$, we have

$$\liminf_{n\to\infty} \gamma_{n,p}^{\mathrm{AD}} \geq \liminf_{n\to\infty} \theta_n \cdot \left(1 - \frac{(1-p)^{n-1}pn}{1-(1-p)^n}\right) = \theta^*.$$

5.2 Upper Bound

In this subsection, we prove Theorem 1.3 to establish that the lower bound derived in the previous section is optimal. Fix $p, \epsilon \in (0,1)$ and n sufficiently large so that $n \geq \left\lceil -\frac{\log(y(1-\epsilon))}{p} \right\rceil$. We define the following distribution through its inverse, where $y(\cdot)$ is the solution to the Hill and Kertz equation:

$$F^{-1}(1-u) = \frac{\theta^*}{(1-p)n} \cdot \delta_{\{0\}}(u) - \frac{p}{1-p} \int_{y^{-1}(e^{-pnu})}^{1-\epsilon} \frac{1}{y'(s)} ds \mathbb{I}_{(0,1]}(u).$$

Given $\epsilon > 0$, denote the value of the optimal algorithm by $v^\epsilon(\texttt{OPT})$ under $F^{-1}(1-u)$ that is defined above. The following proposition provides a characterization of $v^\epsilon(\texttt{OPT})$.

Proposition 5.3. $v^\epsilon(\texttt{OPT}) = \theta^* - \int_0^{1-\epsilon} \frac{1}{y'(s)} \left(1 - \left(1 + \frac{\log(y(s))}{n}\right)^n\right) ds.$

We then characterize the maximum value that can be obtained by any online algorithm. Consider the following dynamic program, where for any $n \geq 1$ and $p \in (0,1)$, D_i^ϵ is the value at observing X_i, $i \in [n+1]$, with the convention $X_{n+1} = 0$:

$$D_i^\epsilon = \sup_{q \in [0,1]} \left\{ (1-p) \int_0^q F^{-1}(1-u) du + (1-pq) D_{i+1}^\epsilon \right\}, \forall i \in [n] \quad \text{and} \quad D_{n+1}^\epsilon = 0. \tag{15}$$

In particular, we are interested in analyzing the dynamic program solution D_i^0. for all $i \in [n]$. We now examine its continuous approximate counterpart (denoted by $d(x)$) through Lemma 5.4 and link it back to discrete valued D_i^0 via Lemma 5.5. In order to find $d(0)$, we rewrite (15) from an ODE perspective. Consider the following Bellman equation:

$$-d'(x) = \sup_{\mu \in [0,\infty]} \left\{ \int_0^\mu h(u) du - p\mu d(x) \right\}, \tag{16}$$

where $h(u) := \theta^* \cdot \delta_{\{0\}}(u) - p \int_{y^{-1}(e^{-pu})}^1 (1/y'(s))\, ds \mathbb{I}_{(0,1)}(u)$. In what follows, we show that for any $x \in (0,1]$, (16) above can be satisfied by

$$d(x) = \int_x^1 -\frac{1}{y'(s)} ds \quad \text{and} \quad \mu = -\frac{\log(y(x))}{p}. \tag{17}$$

Lemma 5.4. (17) *provides a unique solution to the Bellman equation* (16)*, when* $x \in (0,1]$.

Next, we connect the continuous dynamic program value (denoted by $d(i/n)$) with D_i^ϵ.

Lemma 5.5. *Fix* $\sigma \in (0,1)$ *and* $n > -\log(y(1-\sigma))$. *Let*

$$\eta_\sigma := \frac{n\sigma - \log(y(1-\sigma))(1-\sigma)}{(1-\sigma)\left(n + \log(y(1-\sigma))\right)}. \tag{18}$$

Let $\widetilde{D}_i := (1+\eta_\sigma) \cdot d((1-\sigma)i/n)$. *Then for any* $i \in [n]$, *we have* $D_i^\epsilon \leq D_i^0 \leq \widetilde{D}_i$.

Proof. The first inequality follows from observing that $F^{-1}(1-u)$ is a positive, non-increasing function in ϵ. Thus, per (15), it holds that $D_i^\epsilon \leq D_i^0$ for any $\epsilon > 0$. To prove the second inequality, we use the following claim.

Claim 5.6. *For any* $q \in [0,1]$, *we have* $(1-p)\int_0^q F^{-1}(1-u)du + (1-pq)\widetilde{D}_{i+1} \leq \widetilde{D}_i$.

Now we are ready to show the second inequality. We proceed by induction. Note that for some $q_i \in [0,1]$, the bellman equation for D_i^0 is satisfied with equality:

$$
\begin{aligned}
D_i^0 &= (1-p)\int_0^{q_i} F^{-1}(1-u)du + (1-pq_i)D_{i+1}^0 \\
&\leq (1-p)\int_0^{q_i} F^{-1}(1-u)du + (1-pq_i)\widetilde{D}_{i+1}^0 \leq \widetilde{D}_i \qquad \text{(Induction Hypothesis and Claim 5.6)} \square
\end{aligned}
$$

Finally, we compare the value of dynamic program and the optimal algorithm. Specifically, we consider the ratio $D_1^\epsilon / v^\epsilon(\texttt{OPT})$, and analyze its behavior in the asymptotic regime. Noting that $D_1^\varepsilon \leq D_1^0 \leq (1+\eta_\sigma)d((1-\sigma)/n)$ and $\eta_\sigma \to \sigma/(1-\sigma)$ when $n \to \infty$, we have

$$
\lim_{n\to\infty} \frac{D_1^\epsilon}{v^\epsilon(\texttt{OPT})} \leq \lim_{n\to\infty} \frac{(1+\eta_\sigma)\cdot d((1-\sigma)/n))}{v^\epsilon(\texttt{OPT})} = \frac{1}{1-\sigma} \frac{\int_0^1 -\frac{1}{y'(s)}ds}{\theta^* - \int_0^{1-\epsilon} \frac{1-y(s)}{y'(s)}ds},
$$

where in the last inequality we use Proposition 5.3 and Lebesque's dominated convergence theorem. This bound holds for any $\sigma \in (0,1)$. Thus,

$$
\gamma_p^{\mathrm{AD}} \leq \frac{\int_0^1 -\frac{1}{y'(s)}ds}{\theta^* - \int_0^{1-\epsilon} \frac{1-y(s)}{y'(s)}ds}
$$

for any $\varepsilon \in (0,1)$. Then letting $\varepsilon \to 0$, we see that the right-hand side tends to θ^* per Theorem 3.11 in Liu et al. (2020), which concludes our proof.

6 Final Remarks

We introduced the Online Selection with Uncertain Disruption (OS-UD) problem, which captures unexpected disruptions resulting from serving requests. We provided a non-adaptive single-threshold algorithm with a tight competitive ratio of $1-1/e$. We also analyzed the general class of adaptive threshold algorithms and showed that an asymptotic competitive ratio of $\theta^* \approx 0.745$ is attainable, and this is tight.

Even though in this work we focus on the case of fixed disruption probability p, we can use the techniques developed for the non-adaptive single-threshold algorithms to analyze a rare disruption regime, in particular the case when $p = \alpha/n$ with $\alpha \leq 1$. Indeed, letting $q = 1$ in Lemma 4.2 and using the fact that $1 - x \leq e^{-x}$ for all $x \in \mathbb{R}$, we obtain the following lower bound on the competitive ratio: $\left(\frac{1-e^{-\alpha}}{\alpha}\right) \min\left\{1, \frac{\alpha}{1-(1-\alpha/n)^n}\right\} = (1-e^{-\alpha})/\alpha$. This competitive ratio is larger than $1-1/e$ for $\alpha \in [0,1)$, improving the ratio from the non-adaptive case for fixed p. Moreover, this asymptotic competitive ratio converges to 1 as $\alpha \to 0$.

Our study leaves open the question of determining a constant lower bound for $\gamma_{n,p}^{\mathrm{AD}}$ for all $n \geq 1$. Using the linear programming approach in Perez-Salazar et al. (2025) (see also (Jiang et al., 2023)) we can approximate $\gamma_{n,p}^{\mathrm{AD}}$ numerically (see Figure 1). We empirically observe that as n grows, $\gamma_{n,p}^{\mathrm{AD}}$

decreases to θ^*. We also note that, for finite n, when the disruption probability is close to 1, the problem aligns closely with the classical single-selection i.i.d. prophet inequality Indeed, although OS-UD allows multiple selections, if $p = 1 - \varepsilon$ with $\varepsilon \approx 0$, the expected gain beyond the first selection is multiplied by $O(\varepsilon^2)$. Therefore, this regime yields limited new insight, as it reduces to the single-selection case.

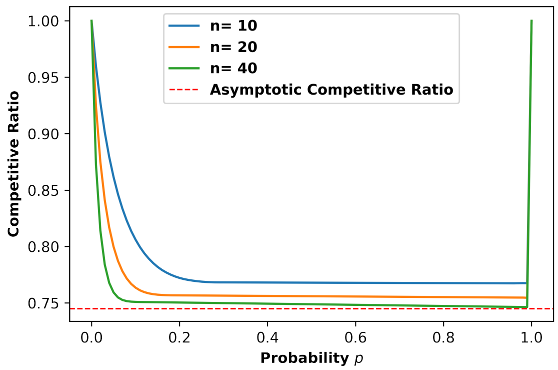

Figure 1: Estimated competitive ratios using linear programming formulation in Perez-Salazar et al. (2025) for $p \in [0, 1]$.

In this paper, we model the disruption as a memoryless process. A natural extension would be to consider broader classes of disruption processes. For example, the disruption probability might increase as more values are accepted, modeling a system that wears out over time; conversely, it could decrease, representing a system that becomes more reliable—such as a growing start-up stabilizing with each accepted value. However, analyzing such variants is non-trivial, as optimal algorithms must account for the number of selections made—a known challenge in multiple-selection problems and an active area of research (Alaei, 2014; Brustle et al., 2025; Jiang et al., 2023). Furthermore, the competitive ratio can depend on this disruption process. In Example 7.1, we show that the competitive ratio in the OS-UD problem can be zero for general disruption processes.

Lastly, OS-UD assumes a single disruption event that terminates service without collecting the last value. Another interesting direction is to incorporate penalties and partial recoveries on the last accepted value, akin to those studied in knapsack settings (Dean et al., 2008; Fu et al., 2018). This is particularly relevant in practical applications such as ticketing and reservations.

7 Missing Proofs from Section 3

Proof of Proposition 3.1. First, note that $\tau_n = 0$, since accepting always brings a non-negative value. Assume that an optimal algorithm, denoted by $\texttt{ALG}^{\text{OPT}}$, has $\tau_i < \tau_{i+1}$ for some $i \in [n-1]$. Consider the following alternative algorithm $\texttt{ALG}^{\text{ALT}}$ that swaps these two thresholds, i.e., the alternative algorithm accepts X_i with threshold τ_{i+1}, accepts X_{i+1} with threshold τ_i, and rest of the thresholds remain unchanged. Denote by $v(\texttt{ALG}^{\text{OPT}})$ and $v(\texttt{ALG}^{\text{ALT}})$ the expected total values collected under these two algorithms. Denote the expected total value obtained by both algorithms before observing the i^{th} value by $v_{[1,\dots,i-1]}$. Conditioned on observing the $(i+2)^{\text{th}}$ value, denote the expected total value obtained by both algorithms starting from observing $(i+2)^{\text{th}}$ value by $v_{[i+2,\dots,n]}$. Let $X^{\tau_i} := \mathbb{E}[X_i | X_i \ge \tau_i]$ and $p_i := \Pr[X_i \ge \tau_i]$. Finally, let p_r be the probability that the algorithm observes value X_i, and $c^i = v_{[1,\dots,i-1]} + p_r(1-p_{i+1}p)(1-p_i p)v_{[i+2,\dots,n]}$. Then the expected total values for both algorithms can be written as

$$
\begin{aligned}
v(\texttt{ALG}^{\text{OPT}}) &= c^i + p_r X^{\tau_i} p_i(1-p) + p_r(1-p_i p)X^{\tau_{i+1}} p_{i+1}(1-p), \\
v(\texttt{ALG}^{\text{ALT}}) &= c^i + p_r X^{\tau_{i+1}} p_{i+1}(1-p) + p_r(1-p_{i+1}p)X^{\tau_i} p_i(1-p),
\end{aligned}
$$

which yields $v(\texttt{ALG}^{\text{ALT}}) - v(\texttt{ALG}^{\text{OPT}}) = p_r(1-p)pp_i p_{i+1}(X_{\tau_{i+1}} - X_{\tau_i}) \ge 0$. Thus, one can construct an alternative algorithm that employs a non-increasing sequence of thresholds by iteratively swapping thresholds, which achieves optimality. □

Proof of Proposition 3.2. The first equality directly follows from the following claim, where its proof is given in the end of the current proof.

Claim 7.1. *For any $j \in [n]$, $\mathbb{E}[X_{(j)}] = \int_0^1 \Pr[Bin(n,v) \ge j] \cdot r(v)dv$, where $r(v) > 0$ satisfies $\int_u^1 r(v)dv = F^{-1}(1-u)$.*

To prove the second equality, we note that

$$
\begin{aligned}
v(\texttt{OPT}) = &\sum_{i=2}^{n} p \cdot (1-p)^{i-1} \sum_{j=1}^{i-1} \int_0^1 F^{-1}(1-q) \cdot (1-q)^{n-j} q^{j-1} \cdot j \cdot \binom{n}{j} dq \\
&\qquad + (1-p)^n \sum_{j=1}^{n} \int_0^1 F^{-1}(1-q) \cdot (1-q)^{n-j} q^{j-1} \cdot j \cdot \binom{n}{j} dq \\
= &\int_0^1 F^{-1}(1-q) \sum_{j=1}^{n-1} (1-q)^{n-j} q^{j-1} \cdot j \cdot \binom{n}{j} \sum_{i=j+1}^{n} p \cdot (1-p)^{i-1} dq \\
&\qquad + \int_0^1 F^{-1}(1-q) \sum_{j=1}^{n} (1-q)^{n-j} q^{j-1} \cdot j \cdot \binom{n}{j} (1-p)^n dq \\
= &\int_0^1 F^{-1}(1-q) \sum_{j=1}^{n} (1-q)^{n-j} q^{j-1} \cdot j \cdot \binom{n}{j} \cdot (1-p)^j dq = \int_0^1 F^{-1}(1-q) g_n(p,q) dq,
\end{aligned}
$$

where the last equality follows from the binomial theorem.

Proof of Claim 7.1 For any $j \in [n]$, let $b_{n,j} := j \cdot \binom{n}{j}$. Then we can write the expected value of the j^{th} top ordered statistics as follows:

$$
\begin{aligned}
\mathbb{E}[X_{(j)}] &= \int_0^\infty x f(x) F(x)^{n-j} (1-F(x))^{j-1} \cdot b_{n,j}\, dx = \int_0^1 F^{-1}(1-u) \cdot (1-u)^{n-j} u^{j-1} \cdot b_{n,j}\, du \\
&= \int_0^1 \int_u^1 r(v)dv \cdot (1-u)^{n-j} u^{j-1} \cdot b_{n,j}\, du = \int_0^1 \Pr[\text{Bin}(n,v) \ge j] \cdot r(v)\, dv, \qquad (19)
\end{aligned}
$$

where the second equality uses the substitution $1 - F(x) = u$, and the third equality comes from our assumption that F^{-1} is differentiable and strictly decreasing so that the existence of $r(v)$ is guaranteed, and the last equality changes the order of integration. $\square$

Example 7.1. *Here, we show that varying the disruption probability within the selection process can lead to a competitive ratio of* 0. *Let n and s be large constants, where $s \ll n$. Assume n is divisible by s for simplicity. Assume $X_i \sim \mathrm{Exp}(1)$. Consider the disruption process $\Pr(Y_i = 1) = 1$ if $i = 1, 1+s, \ldots, 1+(n/s-1)s$, and* 0 *otherwise. The optimal clairvoyant algorithm collects top n/s ordered statistics in expectation, where the exponential distribution's order statistics roughly follows harmonic numbers (e.g., see David and Nagaraja (2004)). Thus, $v(\texttt{OPT}) \approx \sum_{l=1}^{n/s} \log\left(\frac{n}{l}\right) \approx \left(\frac{n}{s}\log s + \frac{n}{s} - 1\right)$. On the other hand, the optimal online algorithm's value is n/s by accepting every value when $Y_i = 0$. From here, we see that $v(\texttt{ALG})/v(\texttt{OPT}) \to 0$ as n and s grow large.*

8 Missing Proofs from Section 4

Lemma 8.1. *Fix $p \in (0,1)$. Then $\frac{v}{\sum_{i=2}^{n} p\cdot(1-p)^{i-1}\sum_{j=1}^{i-1}\Pr[Bin(n,v)\geq j] + (1-p)^n\cdot nv}$ is a non-decreasing function in the interval $v \in [0,q]$.*

Proof. It is sufficient to show that the function $f(v) = v/\left(\sum_{j=1}^{i}\Pr[\mathrm{Bin}(n,v)\geq j]\right)$ is non-decreasing for fixed $i \in [n-1]$. Consider $g(v) = 1/f(v)$, and its derivative with respect to v:

$$\begin{aligned}
g'(v) &= \frac{1}{v^2}\left(\sum_{j=1}^{i}\Big(\sum_{k=j}^{n}\binom{n}{k}kv^k(1-v)^{n-k} - \binom{n}{k}(n-k)v^{k+1}(1-v)^{n-k-1}\Big) - \Pr[\mathrm{Bin}(n,v)\geq j]\right)\\
&= \frac{1}{v^2}\sum_{j=1}^{i}\frac{n!}{(n-j)!(j-1)!}v^j(1-v)^{n-j} - \frac{v^{n+1}}{1-v} - \Pr[\mathrm{Bin}(n,v)\geq j]\\
&= -\frac{iv^{n+1}}{v^2(1-v)} + \frac{1}{v^2}\sum_{j=1}^{i}(j\cdot\Pr[\mathrm{Bin}(n,v)=j] - \Pr[\mathrm{Bin}(n,v)\geq j]).
\end{aligned}$$

Observe that $iv^{n+1}/\left(v^2(1-v)\right)$ is non-negative, and for each $k \in [1,i]$, $k\cdot\Pr[\mathrm{Bin}(n,v)=k] - \sum_{j=1}^{k}\Pr[\mathrm{Bin}(n,v)=j] = 0$, as $\Pr[\mathrm{Bin}(n,v)\geq j] = \sum_{k=j}^{n}\Pr[\mathrm{Bin}(n,v)=k]$. Thus, $g'(v)\leq 0$, and $f(v)$ is non-decreasing in $v\in[0,q]$. $\square$

Proof of Lemma 4.4.

$$\begin{aligned}
v(\texttt{OPT}) &\geq \sum_{i=2}^{n} p\cdot(1-p)^{i-1}\sum_{j=1}^{i-1}\int_0^1 F^{-1}(1-u)\cdot(1-u)^{n-j}u^{j-1}\cdot j\cdot\binom{n}{j}du\\
&= \int_0^1\sum_{i=2}^{n} p\cdot(1-p)^{i-1}\sum_{j=1}^{i-1}\frac{a_1}{n}\delta_{\{0\}}(u)\cdot(1-u)^{n-j}u^{j-1}\cdot j\cdot\binom{n}{j}du\\
&\qquad + \int_0^{\beta/n}\sum_{i=2}^{n} p\cdot(1-p)^{i-1}\sum_{j=1}^{i-1}a_2\cdot(1-u)^{n-j}u^{j-1}\cdot j\cdot\binom{n}{j}du
\end{aligned}$$

$$= a_1 \sum_{i=2}^{n} p(1-p)^{i-1} + a_2 \sum_{i=2}^{n} \sum_{j=1}^{i-1} p \cdot (1-p)^{i-1} \Pr[\mathrm{Bin}(n, \beta/n) \geq j]$$

$$= a_1(1-p)(1-(1-p)^{n-1}) + a_2 \sum_{j=1}^{i-1} \Pr[\mathrm{Bin}(n, \beta/n) \geq j][(1-p)^j - (1-p)^n].$$

The second equality recognizes that $\Pr[\mathrm{Bin}(n, \beta/n) \geq j] = \int_0^{\beta/n} (1-u)^{n-j} u^{j-1} \cdot j \cdot \binom{n}{j} du$, and we recover the lemma by letting $n \to \infty$. $\square$

Proof of Lemma 4.5. We derive the performance of $\texttt{ALG}^q$ via direct calculation.

$$\begin{aligned}
v(\texttt{ALG}^q) &= A_n(q,p) \cdot \frac{\int_0^q F^{-1}(1-u)du}{q} = A_n(q,p) \cdot \frac{\int_0^q (a_1/n)\delta_{\{0\}}(u) + a_2 \mathbb{I}_{(0,\beta/n]}(u)du}{q} \quad \text{(Using (6))} \\
&= \frac{A_n(q,p)}{q} \cdot \left(\frac{a_1}{n} + a_2 \cdot \min\{q, \beta/n\}\right) \\
&= \max\left\{ \sup_{q \in [0,\beta/n]} \frac{A_n(q,p)}{q} \cdot \left(\frac{a_1}{n} + a_2 \cdot q\right), \sup_{q \in [\beta/n, 1]} \frac{A_n(q,p)}{q} \cdot \left(\frac{a_1}{n} + \frac{a_2 \beta}{n}\right) \right\} \\
&= \frac{1-p}{p} \cdot \max\left\{ \sup_{\lambda \in [0,\beta]} \frac{1}{\lambda}\left(1 - \left(1 - \frac{\lambda p}{n}\right)^n\right)(a_1 + a_2\lambda), \sup_{\lambda \in [\beta, n]} \frac{1}{\lambda}\left(1 - \left(1 - \frac{\lambda p}{n}\right)^n\right)(a_1 + a_2\beta) \right\},
\end{aligned} \tag{20}$$

where the last equality follows from the change of variable $q = \lambda/n$. We denote by $r_1(\lambda)$ the second term in the maximum argument in (20). We have $r_1'(\lambda) = \frac{1}{\lambda^2}\left(\lambda p \left(1 - \frac{\lambda p}{n}\right)^{n-1} + \left(1 - \frac{\lambda p}{n}\right)^n - 1\right)$, and for $\lambda = 2/p$, the derivative of r_1 becomes negative; hence, $\sup_{\lambda \in [\beta, n]} r_1(\lambda)$ can be restricted to $[\beta, \max\{\beta, 2/p\}]$. Next, letting $n \to \infty$ yields

$$\lim_{n \to \infty} v(\texttt{ALG}^q) = \frac{1-p}{p} \cdot \max\left\{ \sup_{\lambda \in [0,\beta]} \frac{1 - e^{-\lambda p}}{\lambda}(a_1 + a_2\lambda), \sup_{[\beta, \max\{\beta, 2/p\}]} \frac{1 - e^{-\lambda p}}{\lambda}(a_1 + a_2\beta) \right\}. \tag{21}$$

We denote by $r_2(\lambda)$ and $r_3(\lambda)$ the first and second terms in the maximum argument in (21), respectively. Note that $r_2'(\lambda) = [a_1(e^{-\lambda p} - 1) + p\lambda e^{-\lambda p}(a_1 + a_2\lambda)]/\lambda^2$. Thus, ${r_2}'(\lambda^*) = 0$ implies that $e^{-\lambda^* p}\left(a_1 + a_1 p \lambda^* + a_2 p (\lambda^*)^2\right) = a_1$; hence, $\sup_{\lambda \in [0,\beta]} r_2(\lambda) = \max\{r_2(\lambda^*), r_2(\beta), \lim_{\lambda \to 0} r_2(\lambda)\}$ where $0 < \lambda^* < \beta$, if it exists. Moreover, $r_3(\lambda)$ is a decreasing function in λ, thus the supremum is attained at $\lambda = \beta$. We combine both cases to conclude the proof. $\square$

9 Missing Proofs from Section 5

Proof of Proposition 5.3. By the second characterization of $v^\epsilon(\texttt{OPT})$ in (1), we have

$$v^\epsilon(\texttt{OPT}) = \int_0^1 F^{-1}(1-u) g_n(p,u) du$$

$$
\begin{aligned}
&= \int_0^1 \left(\frac{\theta^*}{(1-p)n} \cdot \delta_{\{0\}}(u) - \frac{p}{1-p} \int_{y^{-1}(e^{-pnu})}^{1-\epsilon} \frac{1}{y'(s)} ds \mathbb{I}_{(0,1)}(u) \right) (1-p) \cdot n(1-pu)^{n-1} du \\
&= \theta^* - \int_0^1 \int_{y^{-1}(e^{-pnu})}^{1-\epsilon} \frac{1}{y'(s)} ds (1-pu)^{n-1} pn du \\
&= \theta^* - \int_{e^{-np}}^1 \int_{y^{-1}(v)}^{1-\epsilon} \frac{1}{y'(s)} ds \left(1 + \frac{\log(v)}{n}\right)^{n-1} \frac{dv}{v} && \text{(Per } v = e^{-npu}\text{)} \\
&= \theta^* - \int_0^{1-\epsilon} \int_{\max\{y(s), e^{-np}\}}^1 \left(1 + \frac{\log(v)}{n}\right)^{n-1} \frac{dv}{v} \frac{1}{y'(s)} ds && \text{(Changing order of integration)} \\
&= \theta^* - \int_0^{1-\epsilon} \frac{1}{y'(s)} \left(1 - \left(1 + \frac{\log(y(s))}{n}\right)^n\right) ds. && \text{(Per } y(s) > e^{-np} \text{ when } s \in (0, 1-\epsilon)\text{)} \quad \square
\end{aligned}
$$

Proof of Lemma 5.4. We first show that (17) satisfies the Bellman equation (16).

$$
\begin{aligned}
&\int_0^\mu h(u) du - p\mu d(x) \\
&= \int_0^{-\log(y(x))/p} \left(\theta^* \cdot \delta_{\{0\}}(u) - p \int_{y^{-1}(e^{-pu})}^1 \frac{1}{y'(s)} ds \mathbb{I}_{(0,1)}(u) \right) du - p \cdot \frac{\log(y(x))}{p} \cdot \int_x^1 \frac{1}{y'(s)} ds \\
&= \theta^* - p \int_0^1 \int_0^{\min\{-\log(y(x))/p, -\log(y(s))/p\}} du \frac{1}{y'(s)} ds - \log(y(x)) \int_x^1 \frac{1}{y'(s)} ds \\
&= \theta^* + \int_0^1 \frac{\max\{\log(y(x)), \log(y(s))\}}{y'(s)} ds - \int_x^1 \frac{\log(y(x))}{y'(s)} ds \\
&= \theta^* + \int_0^x \frac{\log(y(s))}{y'(s)} ds + \int_x^1 \frac{\log(y(x))}{y'(s)} ds - \int_x^1 \frac{\log(y(x))}{y'(s)} ds && \text{(Monotonicity of } y(x)\text{)} \\
&= \theta^* + \int_0^x \frac{y''(s)}{(y'(s))^2} ds && \text{(Per } y''(x) = y'(x) \log(y(x))\text{)} \\
&= \theta^* - \frac{1}{y'(x)} + \frac{1}{y'(0)} = -\frac{1}{y'(x)} = -d'(x).
\end{aligned}
$$

Next, we show that μ is indeed the unique maximizer. We proceed in two steps: (i) the derivative with respect to μ is zero. (ii) the derivative is positive when $\mu = 0$, and the derivative is negative as $\mu \to \infty$. By Leibniz integral rule, we have $\frac{d}{d\mu}\left\{\int_0^\mu h(u)du - p\mu d(x)\right\} = h(\mu) - p \cdot d(x)$.
To verify (i), note that

$$
h(\mu) - p \cdot d(x) = \theta^* \cdot \delta_{\{0\}}(-\frac{\log(y(x))}{p}) - p \int_x^1 \frac{1}{y'(s)} ds \mathbb{I}_{(0,1)}(\mu) - p \cdot \int_x^1 -\frac{1}{y'(s)} ds = 0.
$$

To verify (ii), note that $\lim_{\mu \to \infty} h(\mu) - p \cdot d(x) = -p \cdot d(x) < 0$ and

$$
h(0) - p \cdot d(x) = \theta^* - p \int_0^1 \frac{1}{y'(s)} ds - p \cdot \int_x^1 -\frac{1}{y'(s)} ds = \theta^* - p \cdot \int_0^x \frac{1}{y'(s)} ds > 0. \qquad \square
$$

Proof of Claim 5.6. Let $d_{i,\sigma} := d((1-\sigma)i/n))$ and $y_{i,\sigma} := y((1-\sigma)i/n))$. Then, we have

$$
\begin{aligned}
&(1-p)\int_0^q F^{-1}(1-u)du + (1-pq)(1+\eta_\sigma)d_{i+1,\sigma} \\
&\leq \frac{1}{n}\int_0^\mu h(u)du + \left(1-\frac{p\mu}{n}\right)(1+\eta_\sigma)\left(d_{i,\sigma} + \frac{1-\sigma}{n}d'_{i,\sigma}\right) \qquad \text{(Rewrite } q = \tfrac{\mu}{n} \text{ \& concavity of } d_{i,\sigma}) \\
&= (1+\eta_\sigma)d_{i,\sigma} + \frac{1}{n}\left(\int_0^\mu h(u)du - (1+\eta_\sigma)p\mu d_{i,\sigma}\right) + \frac{1}{n}(1+\eta_\sigma)\left(d'_{i,\sigma}(1-\sigma) - \frac{p\mu}{n}d'_{i,\sigma}(1-\sigma)\right) \\
&\leq \widetilde{D}_i - \frac{d'_{i,\sigma}}{n} + \frac{\eta_\sigma \log(y_{i,\sigma})}{n}d_{i,\sigma} + \frac{d'_{i,\sigma}}{n} + \frac{\eta_\sigma d'_{i,\sigma}}{n} - \frac{\sigma d'_{i,\sigma}}{n} - \frac{\sigma\eta_\sigma d'_{i,\sigma}}{n} + \frac{\log(y_{i,\sigma})}{n^2}d'_{i,\sigma}(1-\sigma)(1+\eta_\sigma) \\
&= \widetilde{D}_i + \frac{\eta_\sigma \log(y_{i,\sigma})}{n}d_{i,\sigma} + \frac{\eta_\sigma d'_{i,\sigma}}{n} - \frac{\sigma d'_{i,\sigma}}{n} - \frac{\sigma\eta_\sigma d'_{i,\sigma}}{n} + \frac{\log(y_{i,\sigma})}{n^2}d'_{i,\sigma}(1-\sigma)(1+\eta_\sigma)
\end{aligned}
$$

Note that $\eta_\sigma \log(y_{i,\sigma})d_{i,\sigma} < 0$, and since $d'_{i,\sigma} < 0$, it is sufficient to show that $\eta_\sigma - \sigma - \sigma\eta_\sigma + \frac{\log(y_{i,\sigma})}{n}(1-\sigma)(1+\eta_\sigma) \geq 0$, which is guaranteed by (18). □

References

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

R. Alijani, S. Banerjee, S. Gollapudi, K. Munagala, and K. Wang. Predict and match: Prophet inequalities with uncertain supply. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 4(1):1–23, 2020.

A. Aouad and Ö. Saritaç. Dynamic stochastic matching under limited time. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 789–790, 2020.

N. Arnosti and W. Ma. Tight guarantees for static threshold policies in the prophet secretary problem. *Operations research*, 71(5):1777–1788, 2023.

J. Brustle, S. Perez-Salazar, and V. Verdugo. Splitting guarantees for prophet inequalities via nonlinear systems. *Mathematics of Operations Research*, 2025.

S. Chawla, J. D. Hartline, and R. Kleinberg. Algorithmic pricing via virtual valuations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 243–251, 2007.

S. Chawla, J. D. Hartline, D. L. Malec, and B. 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.

S. Chawla, N. Devanur, and T. Lykouris. Static pricing for multi-unit prophet inequalities. *Operations Research*, 72(4):1388–1399, 2024.

M. C. Cohen, P. W. Keller, V. Mirrokni, and M. Zadimoghaddam. Overcommitment in cloud services: Bin packing with chance constraints. *Management Science*, 65(7):3255–3271, 2019.

J. Correa, P. Foncea, D. Pizarro, and V. Verdugo. From pricing to prophets, and back! *Operations Research Letters*, 47(1):25–29, 2019.

J. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms and optimal threshold strategies for random arrivals. *Mathematics of operations research*, 46(4): 1452–1478, 2021.

H. A. David and H. N. Nagaraja. *Order statistics*. John Wiley & Sons, 2004.

B. C. Dean, M. X. Goemans, and J. Vondrák. Approximating the stochastic knapsack problem: The benefit of adaptivity. *Mathematics of Operations Research*, 33(4):945–964, 2008.

S. Delong, A. Farhadi, R. Niazadeh, B. Sivan, and R. Udwani. Online bipartite matching with reusable resources. *Mathematics of Operations Research*, 49(3):1825–1854, 2024.

B. Epstein and W. Ma. Selection and ordering policies for hiring pipelines via linear programming. *Operations Research*, 72(5):2000–2013, 2024.

H. Fu, J. Li, and P. Xu. A PTAS for a Class of Stochastic Dynamic Programs. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 56:1–56:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. ISBN 978-3-95977-076-7.

V. Goyal and R. Ravi. A ptas for the chance-constrained knapsack problem with random item sizes. *Operations Research Letters*, 38(3):161–164, 2010.

V. Gupta. Greedy algorithm for multiway matching with bounded regret. *Operations Research*, 72 (3):1139–1155, 2024.

M. T. Hajiaghayi, R. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 1*, AAAI'07, page 58–65, Vancouver, British Columbia, Canada, 2007. AAAI Press. ISBN 9781577353232.

S. He, Y. Wei, J. Xu, and S. H. Yu. Online resource allocation without re-solving: The effectiveness of primal-dual policies. *Available at SSRN*, 2025.

T. P. Hill and R. P. Kertz. Comparisons of Stop Rule and Supremum Expectations of I.I.D. Random Variables. *The Annals of Probability*, 10(2):336 – 345, 1982.

T. P. Hill and R. P. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.

T. P. Hill and U. Krengel. Minimax-optimal stop rules and distributions in secretary problems. *The Annals of Probability*, pages 342–353, 1991.

J. Jiang, W. Ma, and J. Zhang. Tightness without counterexamples: A new approach and new results for prophet inequalities. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, page 909, New York, NY, USA, 2023. Association for Computing Machinery. ISBN 9798400701047.

S. Kerimov, I. Ashlagi, and I. Gurvich. Dynamic matching: Characterizing and achieving constant regret. *Management Science*, 70(5):2799–2822, 2024.

R. 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.

U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745 – 747, 1977.

A. Liu, R. P. Leme, M. Pal, J. Schneider, and B. Sivan. Variable decomposition for prophet inequalities and optimal ordering, 2020.

W. Ma. Improvements and generalizations of stochastic knapsack and markovian bandits approximation algorithms. *Mathematics of Operations Research*, 43(3):789–812, 2018.

A. Mehta and D. Panigrahi. Online matching with stochastic rewards. In *Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, FOCS '12, page 728–737, USA, 2012. IEEE Computer Society. ISBN 9780769548746.

P. Naor. The regulation of queue size by levying tolls. *Econometrica: journal of the Econometric Society*, pages 15–24, 1969.

S. Perez-Salazar and V. Verdugo. Optimal guarantees for online selection over time, 2024.

S. Perez-Salazar, I. Menache, M. Singh, and A. Toriello. Dynamic resource allocation in the cloud with near-optimal efficiency. *Operations Research*, 70(4):2517–2537, 2022.

S. Perez-Salazar, M. Singh, and A. Toriello. The iid prophet inequality with limited flexibility. *Mathematics of Operations Research*, 2025.

E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *The Annals of Probability*, 12(4):1213–1216, 1984. ISSN 00911798, 2168894X.

E. Samuel-Cahn. Optimal stopping with random horizon with application to the full-information best-choice problem with random freeze. *Journal of the American Statistical Association*, 91 (433):357–364, 1996.

A. N. Shiryaev. *Optimal stopping rules*, volume 8. Springer Science & Business Media, Berlin, Germany, 2007.

J. A. Van Mieghem. Dynamic scheduling with convex delay costs: The generalized c— mu rule. *The Annals of Applied Probability*, pages 809–833, 1995.

Y. Wei, J. Xu, and S. H. Yu. Constant regret primal-dual policy for multi-way dynamic matching. In *Abstract Proceedings of the 2023 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 79–80, 2023.

J. Zhang and P. Jaillet. Secretary problems with random number of candidates: How prior distributional information helps, 2023.