OPTIMAL STOPPING PROBLEMS AND COMBINATORIAL OPTIMIZATION UNDER UNCERTAINTY

BY

VASILIS LIVANOS

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Computer Science
in the Graduate College of the
University of Illinois Urbana-Champaign, 2024

Urbana, Illinois

Doctoral Committee:

Associate Professor Ruta Mehta, Chair
Professor Chandra Chekuri
Professor Sariel Har-Peled
Assistant Professor Sahil Singla, Georgia Institute of Technology

ABSTRACT

This thesis studies optimal stopping and combinatorial optimization problems in an uncertain environment. Optimal stopping captures many natural scenarios in which decisions have to be made without knowledge of the future and cannot be amended later on. In combinatorial optimization settings the objective is to optimize a function over distinct elements subject to certain feasibility constraints.

Combinatorial optimization has been classically studied in the full-information setting where the entire input is known a priori. Combining combinatorial optimization with an uncertain environment leads to settings in which algorithms have only partial knowledge of the input, which is revealed element-by-element and all decisions of an algorithm have to be immediate and irrevocable. Such settings of optimization under uncertainty arise naturally in applications in which knowledge of the future cannot be obtained, either inherently or due to prohibiting costs or noise.

In this thesis we focus on two settings. The first is a classical model in optimal stopping theory, the prophet inequality, in which an algorithm has to pick one of many random variables whose realizations are observed sequentially and compares against a prophet who knows all realizations in advance. The second setting is rounding a solution to a linear program in an online manner via the use of an Online Contention Resolution Scheme (OCRS) which is very useful in settings of combinatorial optimization under uncertainty.

We initiate the study of prophet inequalities for independent and identically distributed (I.I.D.) random variables for cost minimization, showing distribution-dependent constant-factor guarantees for the competitive ratio that are qualitatively different from the maximization setting. In addition, we unify the maximization and minimization I.I.D. prophet inequalities via the theory of extreme values and show that the competitive ratio of both settings is governed by a single function that depends only on the extreme value index. We also obtain similar results for the objective of competition complexity, which captures how many more random variables an algorithm needs to observe in order to beat the prophet.

We then ask how our guarantees change if we allow our algorithms the ability to ask simple questions to an oracle that has knowledge of the future. Motivating this, we establish an equivalence between this setting and the top-1-of-m setting in which the algorithm can select m values but is judged only for the best one among them. For the oracle-augmented model, we obtain guarantees on the competitive ratio and the probability of selecting the maximum realization that are almost tight asymptotically with respect to the number of oracle calls, for both the I.I.D. case and the case of non-identical random variables whose arrival order is controlled by an adversary.

Afterwards, we turn to more general combinatorial optimization settings where multiple elements can be selected. We design optimal greedy OCRSs for special cases of matroids and provide matching upper bounds to show their optimality. We then use greedy OCRSs to obtain algorithms with significantly improved guarantees on the competitive ratio for prophet inequalities with a submodular objective function under several combinatorial feasibility constraints such as matroids, matchings and knapsacks.

To my friends and family, without whom life would not be nearly as incredible.

ACKNOWLEDGMENTS

Some say that a PhD is a marathon, and I tend to agree. So as I am approaching the end of this journey, I want to thank some people without whom I could not have made it. First of all, I would like to thank my advisor Ruta Mehta. She was the one who accepted me, believed in me when I first started out, and supported me greatly during this race. I have learned so much from her and I will be forever grateful for her guidance.

I also want to express my sincere thanks to my other advisor Chandra Chekuri. He worked with me during most of my early years as a PhD student and helped me develop tremendously as a researcher. From Chandra I learned so much about how to do research properly and he was always available to talk, offering heaps of wisdom and guidance. Continuing the marathon analogy, he was the one who helped me after the first couple miles, when the marathon starts getting tricky.

Near the end of the race I was fortunate enough to visit, meet and collaborate with plenty of extremely talented and incredibly kind people around the world. A special thanks has to go out to Sahil Singla, who is also part of my PhD committee, and who invited me to spend two months with him in Atlanta, patiently answering all of my, undoubtedly tedious, questions. I had a lot of fun working together while grabbing coffee and I am looking forward to doing it again as soon as possible. Working with Kristóf Bérczi in Budapest for a whole month was far and beyond the most enjoyable part of the entire race. He is someone that I consider my good friend and I feel extremely lucky to have met. Being able to visit José Correa and Victor Verdugo in Chile was a fantastic experience and I am extremely grateful for it. José, a one of a kind researcher who is always happy and smiling, instilled in me confidence and has taught me so much already. Victor is one of the nicest people in the world and he and I truly share the same passion for research. I am very happy and excited to continue working with both of them in Chile after my PhD. Sahil, Kristóf, José and Victor all found me near the end of the marathon, ran the last few miles and crossed the finish line with me.

Anyone who I've talked with in the last six years for more than five minutes knows how difficult it was for me, being used to the sun and warm Greek weather for most of the year, to live in Urbana. This thesis simply would not exist without the truly special group of people that made Urbana not just bearable but really fun. Grinding at Espresso Royale with my buddies Felix, Kesav, Cameron and Grigory and spending time together are some of my happiest moments and I will keep them in my memories forever. The fun times I had hanging out with Stathis, Nachiketa, Gayana and Manisha refreshed me every time the PhD and Urbana dehydrated me. Nowhere in the world is there a group of people so unique and supportive as the UIUC theory group. Manuel, Rucha, Pooja, Robert, Mitchell, Shant, Sahand, Xilin and Patrick all helped me survive initially, while Rhea, Farouk, Shubhang, Aniket, James, Phil, David, Justin and Bhaskar showed me how to be happy and have fun no matter where I live. I thank you all from the bottom of my heart for the fun and crazy times. A special thanks goes to Maria Paula for her unwavering support.

A few people have been with me long before this journey started. The smiles, encouragement and reassurance of Spyros, Kostas, Giannis, Tsapas, Bounias, Dongas, Sakis, Marina, Iro, Dionysia, Alexandra, Joakim and Konstantina have been an invaluable source of comfort whenever I thought I could not make it.

Nothing I have ever achieved would be possible without the unconditional love and support from my family: my cousins Sofia, Alexandros, Gogo, Dina, Andriana and Athanasia, my aunt Nia and uncle Achilleas. I am the luckiest person in the world to have such an incredible support system. A special thanks to my alter ego and soulmate Serafeim who, while not technically family, is as good as.

Mom and dad, there are no words to express how grateful I am for everything you have done and continue to do for me. Thank you.

CONTENTS

Chapter 1: INTRODUCTION

1.1 MOTIVATION

Consider the following scenario: you are the principal of a school and wish to sell the rights to operate the school cantina, aiming to maximize your revenue from the sale (Cantina-Sale). You do not know who will contact you with an interest to buy, or how much they would be willing to spend, but you have some estimate on how the cantina is valued by the general population. Whenever a potential buyer contacts you, they offer you a price for the rights. What should be your strategy? How should you decide to accept or decline the bids?

Continuing our scenario, let us assume instead that you are the town's mayor, in charge of selling the rights of all the schools' cantinas. To keep things simple, you have decided at a price for each of cantina, below which you will not sell its rights. The buyers contact you one after the other and declare their interest. In an effort to ensure a level of fairness to the process, you decide to guarantee that each interested buyer is selected at least with some probability (Fair-Cantina-Sale); the problem is that you have to decide whether to sell or not to an interested buyer before knowing who else is interested. What should be your strategy to maximize the probability that each interested buyer is selected?

Combinatorial Optimization under Uncertainty. These two problems are examples of *combinatorial optimization under uncertainty.* The term "combinatorial optimization" refers to our interest in selecting a number of elements from a larger set – called the *ground set* – in order to optimize a function over them. Uncertainty, meanwhile, is a natural part of many settings. This is due to a lack of knowledge about the future, prohibiting costs in estimating such knowledge or even potential noise in a problem's input.

Examples of such problems are extremely common in the real world, from grocers deciding how to set prices for their fruit and vegetables to state governments deciding how to auction off a public works project. The difficulty of such problems is amplified by their combinatorial nature when the constraints involved are more complex; just imagine the town mayor having to sell the rights to multiple schools' cantinas when each buyer is interested in only some of them, or trying to find the best possible matching of kidney donors to patients of a local hospital. Add in the element of uncertainty to the input, which is inherent in many real-world scenarios, and the problem seems almost intractable.

In this thesis, we study combinatorial problems in which the mechanism designer has partial knowledge about the input; specifically, knowledge of a probability distribution from which the input is drawn. Once the mechanism is decided, the input is revealed in a sequential manner, one element after another. At each step, we need to make *immediate* and *irrevocable* decisions. For example, in the Cantina-Sale scenario, every time a buyer contacts us, we need to decide whether to accept their bid and we cannot go back and accept a rejected bid. Since we need to make decisions without knowledge of the future and we cannot change them after the fact, we are bound to not be optimal in hindsight. Problems in which decisions have to be made immediately and irrevocably, without knowledge of the future, are called *optimal stopping* problems and fall within the field of online algorithms.

The Cantina-Sale and Fair-Cantina-Sale scenarios are applications of two classical and well-studied problems in the intersection of optimal stopping theory and combinatorial optimization under uncertainty: the *prophet inequality* and rounding via an *online contention resolution scheme.* We introduce these settings, and how to model stochastic combinatorial optimization problems in general, in §1.2. Afterwards, in §1.3, we provide a brief overview of our contributions as well as an outline for this thesis.

1.2 OPTIMAL STOPPING AND STOCHASTIC COMBINATORIAL OPTIMIZATION

Let N denote a ground set of n elements. A *feasibility constraint* is a set $\mathcal{F} \subseteq 2^N$, and a set $I \in \mathcal{F}$ is called a *feasible set*. Given an *objective function* $f : 2^N \to \mathbb{R}$, the problem of *combinatorial optimization* is to select a feasible set S that maximizes or minimizes f, depending on the setting. Access to $\mathcal{F}$ and f is usually provided via two oracles: a *feasibility oracle* that, given a set S, returns whether $S \in \mathcal{F}$ or not and a *value oracle* that, given a set S, returns the value $f(S)$. One can trivially solve this problem by computing the value of f for every feasible set I but, in general, this is inefficient since $\mathcal{F}$ may contain exponentially many sets with respect to n. When solving combinatorial optimization problems, we usually require our algorithms to run in time that is polynomial in n and make polynomially many calls to both oracles.

Returning to the Cantina-Sale scenario, N represents the buyers and $\mathcal{F}$ consists of all singletons of N, since we can only select one buyer to run the cantina. This is known as the *single-item* or *rank*-1 *matroid* setting. However, in this scenario, the value v_i of buyer i for the rights to the cantina is a random variable drawn from a known distribution $\mathcal{D}_i$, independently from the other buyers' values. Furthermore, we have to make a decision before we can observe all values. As we can see, uncertainty in the input is an inherent element of this setting.

There are several ways to combine combinatorial optimization and uncertainty that depend on the setting in question. For the simple Cantina-Sale scenario, we wish to select a single buyer i with the highest possible value and thus one can imagine that the value $f(\{i\}) \coloneqq v_i$, sometimes denoted by w_i, of each singleton $\{i\}$ is a random variable, drawn from a known distribution $\mathcal{D}_i$. The fact that the input is not deterministic but instead comes from a probability distribution makes this a *stochastic* combinatorial optimization setting. In addition, the random variables are revealed one after the other in an *online* manner, and every sell-or-continue decision needs to be made before we can observe future realizations. This special case of combinatorial optimization under uncertainty where we only need to select a single value is the characteristic problem in the field of optimal stopping theory. In general, the joint distribution of all singletons can be arbitrarily correlated. However, we usually assume that the elements realize independently and thus the joint distribution is simply the product distribution.

Objective Functions and Feasibility Constraints. In this thesis, we focus on two special families of feasibility constraints and two special classes of objective functions.

Definition 1.1 (Additive/Linear/Modular Function)**.** Given a vector $\boldsymbol{w} \in \mathbb{R}^n$, an *additive* objective function is defined as $f(S) = \sum_{i \in S} w_i$, for any $S \subseteq N$.

A generalization of additive functions is the class of submodular functions.

Definition 1.2 (Submodular Function)**.** An objective function f is called submodular if, for any $A, B \subseteq N$, we have $f(A) + f(B) \geq f(A \cup B) + f(A \cap B)$.

Notice that, for additive functions we have $f(A) + f(B) = f(A \cup B) + f(A \cap B)$, which is why they are also known as modular functions.

The two families of feasibility constraints that are most common and well-studied are packing and covering constraints.

Definition 1.3 (Packing/Downwards-Closed Constraint)**.** A feasibility constraint $\mathcal{F}$ is called a *packing* constraint if, for every $A \in \mathcal{F}$ and $B \subseteq A$, we have $B \in \mathcal{F}$.

Definition 1.4 (Covering/Upwards-Closed Constraint)**.** A feasibility constraint $\mathcal{F}$ is called a *covering* constraint if, for every $A \in \mathcal{F}$ and $B \supseteq A$, we have $B \in \mathcal{F}$.

Arrival Order. One way to classify optimal stopping problems is by the order in which the elements are observed; specifically, by who is in charge of deciding the *arrival order*. In one extreme, the arrival order is decided by an adversary, who aims to inflict the most harm to the algorithm; in such a setting the elements are said to arrive in *adversarial order*. The power that such an adversary has is not explicitly described here but one can consider multiple subclasses of adversarial arrival order settings based on how powerful or not the adversary is. When the arrival order is chosen by the universe, we obtain a more lenient setting. Specifically, in the case of a *random order* of arrival, the elements of N are observed in one of $n!$ permutations, drawn uniformly at random. An even better scenario is the *free order* setting, in which the algorithm is free to choose the next element it observes based on the distributions and the realizations observed so far. Perhaps the most ideal scenario is the case of *independently and identically distributed (I.I.D.)* elements, in which $\mathcal{D}_i = \mathcal{D}$ for all elements i.

1.2.1 Prophet Inequality

A classical setting in optimal stopping theory, first introduced in 1977 by Krengel, Sucheston and Garling [1], [2], is the *prophet inequality.* In this setup, one is presented with *take-it-or-leave-it* rewards $X_1, \ldots, X_n$ in an online manner, where each X_i is drawn from a known distribution $\mathcal{D}_i$, independently from the other rewards, and can stop at any point and collect the last reward seen.

From the perspective of stochastic combinatorial optimization, this is a special single-item case where the set of elements N in the prophet inequality is the set of random variables, $\mathcal{F}$ is the set of all singletons and $f(S) = \max_{i \in S} X_i$. Given that the distributions are known, the inequality ensures the existence of a *stopping strategy* S for any arrival order of the random variables, with expected reward at least half that of a *prophet* who can see the realizations of all the X_i's upfront and thus can always select the maximum, *i.e.,* $\mathbb{E}[S] \geq {}^1\!/_2\, \mathbb{E}[\max_i X_i]$. The *competitive ratio* $\frac{\mathbb{E}[S]}{\mathbb{E}[\max_i X_i]}$ is the usual objective in the prophet inequality setting and quantifies the loss that the algorithm incurs by observing the realizations sequentially instead of all at once. A simple example shows that the factor of ${}^1\!/_2$ is tight in this general setting.

The I.I.D. Setting. Part of this thesis focus on a special case of the prophet inequality; that of n I.I.D. random variables. For this setting, Hill and Kertz [3] showed the existence of a simple threshold strategy that is optimal. They showed it achieves a competitive ratio of $1 - {}^1\!/_e$, as well as providing a corresponding upper bound of ≈ 0.745. Subsequent work [4] improved the analysis of the optimal algorithm and showed a lower bound of 0.738, before Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld [5] obtained the tight factor of 0.745. Contrary to the adversarial order case, the optimal threshold strategy is not a single-threshold strategy and potentially uses a different threshold per random variable.

Applications. These results, and their variations and generalizations, have found extensive applications. Starting with the work of Hajiaghayi, Kleinberg and Sandholm [6] and later Chawla, Hartline, Malec and Sivan [7], it has been connected to the design of simple yet approximately optimal sequential posted-price mechanisms. As an example of such a mechanism, consider the Cantina-Sale scenario, and suppose we decide to offer the cantina rights for a pre-determined price T, selling them to the first buyer that is willing to purchase them at this price. The question of course becomes how to design T to obtain a guarantee on the competitive ratio. Due to the significance of such problems in Bayesian mechanism design, the study of prophet inequalities has seen a tremendous surge in the last decade [4], [5], [8]–[25].

1.2.2 Online Contention Resolution Schemes (OCRSs)

One approach to solving combinatorial optimization problems under uncertainty in general is to first relax the notion of containment in a set and consider a *fractional solution* in which we are allowed to select elements fractionally. This is typically done by designing a Linear Program (LP) that describes the problem and solving it to obtain a fractional solution. This solution can then be rounded to an *integral solution* in a manner which respects the feasibility constraints to obtain a final feasible set.

Contention Resolution Schemes. *Contention Resolution Schemes (CRSs)* are general rounding algorithms for certain constraint families, introduced by Chekuri, Vondrák and Zenklusen [26] for the purpose of maximizing a submodular function. Examples of such combinatorial constraints include selecting an independent set in a given matroid[1], selecting a feasible matching in a given graph in which the elements correspond to edges, or selecting a feasible set of elements subject to a knapsack constraint, where each element is associated with a size. For a given fractional solution $\boldsymbol{x}$, the main idea behind CRSs is to first obtain a random set R, drawn from the product distribution with marginals $\boldsymbol{x}$, hence called the set of *active elements*. Since R may be infeasible with respect to the constraint $\mathcal{C}$, the CRS proceeds to "drop" specific elements from R and obtain a new, feasible, set $R' \subseteq R$.

Online Contention Resolution Schemes. Feldman, Svensson and Zenklusen [27] introduced the notion of *Online Contention Resolution Schemes (OCRSs)*, which are able to provide guarantees even when one is required to round the elements in a given, potentially adversarial, order. Such rounding algorithms have been recently used to obtain several interesting results on both offline and online optimization in a multitude of combinatorial settings [20], [21], [26]–[30], and have more applications in online mechanism design and posted pricing mechanisms [6], [7]. An important subclass of OCRSs is the class of *greedy OCRSs*. Intuitively, a greedy OCRS π fixes a downward-closed subfamily of feasible sets $\mathcal{I}$ before the online process starts. During the online process, the greedy OCRS maintains a subset S of the elements which is feasible in $\mathcal{I}$, and then greedily accepts any active element i if $S \cup \{i\}$ is also feasible in $\mathcal{I}$, i.e. if i does not violate feasibility, with respect to $\mathcal{I}$, of the set maintained by the greedy OCRS. One can easily see that the final set $\pi(R)$ at the end of the online process is feasible by construction.

Selectability. For a greedy OCRS π, the quality of the approximation guarantee is governed by the notion of selectability. Informally, an OCRS is c-selectable if for any vector of marginals $\boldsymbol{x}$ and every element $i \in N$, it ensures $\Pr[i \in \pi(R) \mid i \in R] \geq c$, i.e. that the probability of every active element being included in the final set is at least c. Returning to the Fair-Cantina-Sale scenario, if we denote the probability that buyer i will be interested in purchasing the cantina rights by x_i, it is easy to see that using any c-selectable OCRS

[1] A matroid is a non-empty downward-closed feasibility constraint where if $A, B \subseteq N$ are both feasible and $|A| < |B|$, there exists an element $e \in B \setminus A$ such that $A \cup \{e\}$ is also feasible. The feasible sets of a matroid are called *independent sets*.

for $\mathcal{F}$ consisting of all singletons of N as a black-box, guarantees that every interested buyer will be selected with probability at least c, since the interested buyers are exactly the buyers in R. Therefore, providing a good guarantee for the Fair-Cantina-Sale scenario reduces to the design of OCRSs with high selectability.

OCRSs are strongly related to the prophet inequality setting. Feldman, Svensson and Zenklusen [27] observe that, given an OCRS π for a feasibility constraint $\mathcal{F}$, one can use π as a black-box to obtain guarantees for the prophet inequality setting in which multiple elements from $X_1, X_2, \ldots, X_n$ can be chosen and the chosen elements have to be feasible with respect to $\mathcal{F}$. The reverse, however, is not true; even for the simple case of selecting only one out of n I.I.D. elements, there exists an optimal stopping strategy that achieves a competitive ratio approximately equal to 0.745 [3], [5] as noted previously, whereas the best possible CRS is only $1 - 1/e \approx 0.632$-selectable [26].

1.3 THESIS CONTRIBUTIONS

Thesis Objectives. Our main motivation in this thesis is to provide a better understanding of the single-item setting, both for prophet inequalities and OCRSs. The technical parts of this thesis are contained in four chapters. The first two chapters study prophet inequalities in the single-item setting. We are interested in providing a unified analysis of the I.I.D. prophet inequality that is distribution-optimal for both maximization and minimization, study how many more random variables an algorithm needs to beat the prophet in each case, as well as how our guarantees improve if the algorithm has access to an oracle that notifies whether it should stop or continue. The next two chapters study greedy OCRSs for additive and submodular objective functions and several feasibility constraints. Our aim is to provide the first optimal greedy OCRSs for the single-item setting, or other simple constraints, and additive objective functions, as well as significantly improve the guarantees of prophet inequalities for combinatorial settings via corresponding greedy OCRSs for submodular objective functions.

We briefly highlight our results and techniques, in §1.3.1 for problems in the similar to Cantina-Sale and the prophet inequality setting, and in §1.3.2 for problems similar to Fair-Cantina-Sale and the OCRS setting. Afterwards, in §1.4, we summarize past literature that is relevant to this thesis. Finally, in §1.5, we present a roadmap of the chapters in this thesis and how they are organized.

1.3.1 Results and Techniques for Prophet Inequalities

In the first part of this thesis, we focus on the prophet inequality setting. The running example for this section is Cantina-Sale. Specifically, in Chapter 2, we provide a unified analysis of the prophet inequality with I.I.D. arrivals for both packing and covering constraints and provide optimal guarantees, while in Chapter 3 we equip our algorithm with the ability to make simple oracle calls about the future realizations and connect this setting to another well-studied prophet inequality setting in which the algorithm is allowed to select up to k realizations but is judged only according to the best among them.

Distribution-Optimal Prophet Inequalities. Recall the example of Cantina-Sale and consider a slight variant of this scenario, where the principal has decided to subsidize the cantina costs, and the buyers bid instead for how low they can keep the costs. The objective now becomes to select the buyer that minimizes the cost to run the school cantina, while having to make decisions without knowledge of the future. This variant incentivizes us to define the *I.I.D. Min-Prophet Inequality.*

In the I.I.D. Min-Prophet Inequality, we are given a probability distribution $\mathcal{D}$ and observe the realizations of n random variables $X_1, \ldots, X_n \sim \mathcal{D}$. At any point after observing an X_i, we can choose to select or discard it but, in contrast with the classical prophet inequality setting, we *must* stop at some point and take the last realization we observed. If we discard X_i, then this realization is lost forever and the process continues. An all-knowing prophet, who can see the realizations of all X_i's upfront can always select the minimum realized cost and hence their expected cost is $\mathbb{E}[\min_i X_i]$.

In this problem, the goal is to design a *stopping strategy* that minimizes the expected cost. We say that algorithm ALG is α-competitive if $\mathbb{E}[\text{ALG}] \leq \alpha \cdot \mathbb{E}[\min_i X_i]$. Our initial motivation for this problem were the following questions.

Question 1.1

(a) Is the I.I.D. Min-Prophet Inequality equivalent to the maximization setting?

(b) Does there exist a constant-competitive algorithm for the minimization setting?

What makes this setting interesting is that the answer to the first question turns out to be negative. The simple fact that the feasibility constraint in the minimization case is upwards-closed instead of downwards-closed is enough of a difference to yield qualitatively different results, compared to the maximization case. In Chapter 2 we show that, unlike the maximization setting, there is no constant $c > 1$ for which there exists a uniform c-competitive algorithm for the I.I.D. Min-Prophet Inequality. However, we also observe that one can obtain constant guarantees on the competitive ratio for every fixed distribution $\mathcal{D}$, with the constant being distribution-dependent.

Question 1.2

(a) Can we provide guarantees to the competitive ratio, for both the maximization and the minimization settings, that are distribution-sensitive?

(b) If so, which parameters of the distribution affect the competitive ratio?

Our main contribution here is a unified analysis of both the maximization and minimization settings, via the lens of Extreme Value Theory – for an introduction to Extreme Value Theory see [31]. Establishing a theory of Extreme Values was an achievement of several mathematicians of the past century, from Fisher, Tippett and Gnedenko to von Mises and de Haan, who were interested in the behaviour of the largest or smallest of an independent sample of n values from a common distribution, as n grows large. Their theory can be seen as an analogue of the classical Central Limit Theorem (CLT); just as the CLT predicts that the average of n samples from a distribution with finite variance converges in distribution to a standard Gaussian, Extreme Value Theory predicts that if the maxima and minima of n samples converge in distribution, the resulting distribution has to be one of three types, parametrized by the positive, zero and negative values, respectively, of a parameter γ called the *extreme value index*. For an introduction to Extreme Value Theory see [31].

In Chapter 2, we give a closed form for the asymptotic competitive ratio of both the I.I.D. Max and Min-Prophet Inequalities, as n goes to infinity, that depends only on γ. Perhaps surprisingly, We show the existence of a *unique* function $\Lambda(\gamma)$ that captures the competitive ratio for both maximization and minimization; when the asymptotic competitive ratio is not 1, it is equal to $\Lambda(\gamma)$ for $\gamma \geq 0$ in the I.I.D. Max-Prophet Inequality and for $\gamma \leq 0$ in the I.I.D. Min-Prophet Inequality.

One difference between our results and the results of [1], [2], [8], [32] for the classical prophet inequality is that in the latter setting, the optimal constant competitive ratio can be achieved via an algorithm that fixes a single threshold T based on the given distributions, and selects the first realization above T. Our results, however, are obtained by studying the optimal (multiple-threshold) algorithm. This motivates us to also investigate single-threshold algorithms for the I.I.D. Min-Prophet Inequality setting.

Question 1.3

(a) Do simple, single-threshold algorithms achieve a constant competitive ratio in the I.I.D. Min-Prophet Inequality as they do in the maximization setting?

(b) Is it possible to provide distribution-sensitive guarantees for single-threshold algorithms as well?

Another surprising outcome of our analysis is that, in contrast to all classical results in the prophet inequality setting, where the optimal single-threshold and optimal multiple-threshold algorithms are only a constant factor apart, the answer to the first question is negative. In particular, we analyze the optimal single-threshold for the I.I.D. Min-Prophet Inequality via the lens of Extreme Value Theory and show that it achieves a competitive ratio that is poly-logarithmic with respect to n, where the exponent is exactly $-\gamma$.

Competition Complexity. In [33], the authors introduced the study of *competition complexity* for optimal stopping problems. Informally, for an I.I.D. prophet inequality instance, the competition complexity is the smallest $c > 1$ such that the algorithm that sequentially observes $c \cdot n$ samples from the given distribution is able to "beat" the expected value of a prophet that observes n samples form the same distribution. In [33], the authors show that, for the I.I.D. Max-Prophet Inequality, the competition complexity can be unbounded. However, their constructed distribution crucially depends on n to achieve this worst-case behaviour. Therefore, their results raise some natural questions.

Question 1.4

(a) Is the competition complexity of the I.I.D. Max-Prophet Inequality bounded by a constant for distributions that are independent of n?

(b) Is the competition complexity of the I.I.D. Prophet Inequality for a large number of random variables the same for the maximization and the minimization settings?

(c) Can we quantify the competition complexity per distribution and provide a closed form based on the distributions parameters?

In Chapter 2, we answer all questions above positively. In contrast to the competitive ratio, we show that for the competition complexity objective, the maximization and minimization settings behave in *exactly* the same way. We quantify the asymptotic competition complexity, as n goes to infinity, of both settings by a single function that depends only on γ and show that it is always bounded above by $e \approx 2.718$. Since for distributions that depend on n, the competition complexity can be unbounded [33], our result establishes that this dependence on n is crucial for this worst-case behaviour.

Prophet Inequality with Oracle Access. Returning to the classical prophet inequality setting, we are interested in ways to improve upon the tight bounds of $1/2$ and ≈ 0.745 for the adversarial order and I.I.D. settings, respectively. Motivated by the idea of enhancing algorithms via the use of machine-learned *predictions*, in order to go beyond worst-case analysis [34]–[38], we provide the algorithm with access to an oracle $\mathcal{O}$ that has complete knowledge of all future realizations. After observing a realization X_i, the algorithm can decide to query the oracle and ask a simple yes or no answer: is X_i greater than all future realizations, or is an even larger realization waiting in the future? Given the oracle's answer, the optimal behaviour of the algorithm is obvious; if it is yes, select X_i and stop, otherwise continue to find the larger realization. We call this the *k-oracle setting*.

In this setting, aside from the competitive ratio, we can also consider a different objective, that of *maximizing the probability of selecting the highest* realization. For our running example of the Cantina-Sale scenario, suppose that you are only satisfied with selling the operation rights to the buyer that can offer the most money, and every other outcome is considered a failure. What is the largest probability of selecting this buyer that you can guarantee? This question has many connections to the well-known *secretary problem* [39] but is slightly different, as on one hand we possess distributional knowledge of the values of the arriving buyers but on the other they arrive in adversarial order. For this setting, Gilbert and Mosteller [40] showed that one can achieve a probability equal to ≈ 0.58 for the I.I.D. case, before Nuti [41] extended this result to show the same probability guarantee even for non-identical distributions and random arrival order. The above lead us to the following questions.

Question 1.5

(a) Do the guarantees on the competitive ratio improve when the algorithm has access to k oracle calls that instruct whether to stop or continue?

(b) Is it possible to obtain similar guarantees when the objective is maximizing the probability of selecting the highest realization, instead of the competitive ratio?

In Chapter 3, we answer both questions positively. In particular, we provide asymptotically tight upper and lower bounds on the competitive ratio of the k-oracle setting, for both the I.I.D. and the adversarial order cases. We also study the objective of maximizing the probability of selecting the highest realization, again showing asymptotically tight upper and lower bounds for the I.I.D. case. In all cases, our guarantees approach 1 as $k \to \infty$. Perhaps most importantly, we establish a connection between the k-oracle setting and the top-1-of-k setting, in which the algorithm is allowed to select k realizations but it is judged only by the highest one, regardless of whether the objective is the competitive ratio or the probability of selecting the highest. This setting was introduced for non-identical distributions by Assaf and Samuel-Cahn [42], in which the authors show a bound of $1 - {}^{1}\!/_{k+1}$ on the competitive ratio. Subsequently, Ezra, Feldman and Nehama [43] improved this to $1 - \mathrm{O}\left(e^{-k}\right)$. Via the connection we establish, our results imply further improvements upon these bounds for both the I.I.D. case as well as non-identical distributions.

1.3.2 Results and Techniques for OCRSs

In the second part of this thesis, we investigate online contention resolution schemes introduced informally in §1.2.2. Our motivation is to obtain optimal greedy OCRSs for simple feasibility constraints such as matroids, matchings and knapsacks. The running example for this section is Fair-Cantina-Sale. Specifically, in Chapter 4, we study provably optimal greedy OCRSs for additive objective functions and simple constraints, while in Chapter 5 we use greedy OCRSs for submodular objective functions to solve a combinatorial variant of prophet inequalities for several feasibility constraints of interest.

With every feasibility constraint $\mathcal{F}$, we can associate a set $\mathcal{P} \subseteq [0,1]^N$, called the *polyhedral relaxation* of $\mathcal{F}$, where $\mathcal{P}$ is a polyhedron and a characteristic vector $\mathbb{1}_S$ of a set S is in $\mathcal{P}$ if and only if $S \subseteq \mathcal{F}$. In other words, $\mathcal{P}$ is the convex hull of the characteristic vectors of the sets in $\mathcal{F}$. For downwards-closed and upwards-closed constraints, $\mathcal{P}$ is a bounded polyhedron, also known as a *polytope*.

A *Contention Resolution Scheme (CRS)* π for $\mathcal{P}$ is a procedure that, given as input a vector $\boldsymbol{x} \in \mathcal{P}$ and a set $A \subseteq N$, returns a random set $\pi_{\boldsymbol{x}}(A) \subseteq A \cap \operatorname{supp}(\boldsymbol{x})$ such that $\pi_{\boldsymbol{x}}(A) \in \mathcal{F}$ with probability 1. Let $R(\boldsymbol{x})$ denote a random set drawn from the product distribution with marginals $\boldsymbol{x}$. *Online Contention Resolution Schemes (OCRSs)* were introduced in [27] to handle online settings such as the Fair-Cantina-Sale scenario, where the elements of N reveal whether they are active, i.e., whether or not they are in $R(\boldsymbol{x})$, in adversarial order and the OCRS must decide whether to include an element in $\pi_{\boldsymbol{x}}(A)$ or not before it observes the remaining active elements. An OCRS π for a polytope $\mathcal{P}$ is an online algorithm that selects a subset $\pi(R(\boldsymbol{x})) \subseteq R(\boldsymbol{x})$ such that the characteristic vector of $\pi(R(\boldsymbol{x}))$ is in $\mathcal{P}$.

We focus on the special case of *greedy OCRSs*. A greedy OCRS π for $\mathcal{P}$ is an OCRS that, for any $\boldsymbol{x} \in \mathcal{P}$, defines a downward-closed subfamily of feasible sets $\mathcal{F}_{\pi,\boldsymbol{x}} \subseteq \mathcal{F}$. Then, an element i is selected when it arrives if, together with the already selected elements, the obtained set is in $\mathcal{F}_{\pi,\boldsymbol{x}}$. If the choice of $\mathcal{F}_{\pi,\boldsymbol{x}}$ given $\boldsymbol{x}$ is randomized, the greedy OCRS is called randomized; otherwise, it is called deterministic. A greedy OCRS for $\mathcal{P}$ is c-selectable if for any $\boldsymbol{x} \in \mathcal{P}$, we have

$$\Pr\left[I + i \in \mathcal{F}_{\pi,\boldsymbol{x}} \quad \forall I \subseteq R(\boldsymbol{x}),\, I \in \mathcal{F}_{\pi,\boldsymbol{x}}\right] \geq c, \quad \forall i \in N. \tag{1.1}$$

Optimal Greedy-OCRSs. Lee and Singla [44] designed an optimal OCRS for matroids that is $1/2$-selectable. However, it is not a greedy OCRS. Since greedy OCRSs provide guarantees against stronger adversaries than regular non-greedy OCRSs, the following question arises naturally.

Question 1.6

(a) Are greedy OCRSs as powerful as general non-greedy OCRSs?

(b) If not, what is the optimal selectability for simple feasibility constraints, e.g. matroids?

In Chapter 4, we provide the first separation between greedy and non-greedy OCRSs by answering the first question negatively. In particular, we show that there does not exist a c-selectable greedy OCRS for any $c > 1/e$ even for rank-1 matroids. Aiming to provide an answer to the second question, Feldman, Svensson and Zenklusen [27] presented a $1/4$-selectable greedy OCRSs for matroid polytopes. However, it is currently open whether this OCRS is optimal, or one can achieve a selectability higher than $1/4$. We make progress on this question as well, showing that there exists a $1/e$-selectable greedy OCRS for rank-1[2], partition and transversal matroids, thus making our impossibility result tight.

Submodular Prophet Inequalities. In several settings with externalities between the selected elements, the value of a solution cannot be captured by a simple additive function. For example, in the Fair-Cantina-Sale scenario, suppose that you are allowed to select multiple buyers to operate the school's cantina, hoping that each of them provides more value to the diversity of the menu and helps the cantina run more smoothly. However, there exists a point after which any additional operator does not really add value and potentially hinders the operating process. Thus, in the process of selecting multiple buyers for a single cantina, we experience *diminishing returns*.

[2] In a rank-1 matroid, every feasible set is a singleton. This is also known as the *single-item* setting.

Such instances can be captured by combinatorial optimization problems in which the objective is a *submodular function*. Greedy OCRSs were first applied to prophet inequality instances with submodular objectives by Rubinstein and Singla [20], where they showed the existence of constant-competitive submodular prophet inequalities for matroid polytopes. While impressive, their approach suffer from two important issues; the guarantee on the competitive ratio is several orders of magnitude small and, in addition, their approach is computationally inefficient.

Question 1.7

(a) Can we obtain computationally efficient submodular prophet inequalities with good guarantees for more general objective functions, e.g. submodular functions?

(b) If so, is it possible to obtain such guarantees for other feasibility constraints as well, e.g. matchings, knapsacks, etc?

In Chapter 5, we answer both questions positively. First, we improve upon the OCRS of [20] via different techniques and design a computationally efficient algorithm for submodular prophet inequalities using greedy OCRS for matroid polytopes with a guarantee that is several order of magnitudes higher than the one of [20]. In particular, we provide a simplified framework which directly generalizes any OCRS for additive functions to a greedy OCRS for submodular functions. Via this reduction, we obtain computationally efficient algorithms with good guarantees on the competitive ratio for submodular prophet inequalities with matching and knapsack feasibility constraints.

Continuous Extensions of Set Functions. Since polytopes arise naturally in the study of combinatorial optimization problems, one can obtain new insights for a problem by relaxing the notion of the objective function being a set function. Just as we relaxed the concept of containment in a set and allowed for fractional solutions earlier, we can consider different *continuous extensions* of the objective function, in order to assign a value to every point $\boldsymbol{x} \in \mathcal{P}$.

To start, recall that the general assumption in OCRSs is that the set $R(\boldsymbol{x})$ of active elements in drawn from the product distribution with marginals $\boldsymbol{x}$. Therefore, a natural first candidate for a continuous extension of a function f on a fractional point $\boldsymbol{x}$ is the expected value of the function for a random set S drawn from the product distribution with marginals $\boldsymbol{x}$. This is usually denoted as $F(\boldsymbol{x}) \triangleq \mathbb{E}_{S\sim\mathrm{Prod}(\boldsymbol{x})}\left[f(S)\right]$ and is known as the *multilinear relaxation*, introduced by Calinescu, Chekuri, Pal and Vondrák [45], and is very useful in maximization settings.

Another continuous extension that is useful in maximization settings is known as the *concave closure*, denoted by $f^+(\boldsymbol{x})$. Intuitively, the value of the concave closure at a fractional point $\boldsymbol{x}$ is equal to the maximum expected value of f for a random set S drawn from any distribution with marginals $\boldsymbol{x}$. In other words, the concave closure is similar to the multilinear relaxation, but allows for correlated distributions subject to the marginals being $\boldsymbol{x}$ and equals the highest expected value among all such distributions. Since the concave closure provides an upper bound to the objective of many combinatorial optimization settings, and the multilinear relaxation is easy to maximize for most polytopes, it is natural to ask whether how they are related.

Correlation Gap. For any function f, the smallest possible value of the ratio $\frac{F(\boldsymbol{x})}{f^+(\boldsymbol{x})}$ across all $\boldsymbol{x} \in [0,1]^n$ is known as the *correlation gap*. Introduced by Yan [46], the correlation gap intuitively measures how much the expected value of f changes, with respect to a random set, if only the marginal probabilities of the set are fixed. One can equivalently define the correlation gap of a polytope $\mathcal{P}$ defined by a feasibility constraint $\mathcal{F}$, by looking at the weighted rank function $r(\boldsymbol{w}) = \max_{S \subseteq N, S \in \mathcal{F}} \sum_{i \in S} w_i$ corresponding to the polytope. The, the correlation gap of $\mathcal{P}$ is the smallest out of all correlation gaps for $r(\boldsymbol{w})$, for every $\boldsymbol{w} \in \mathbb{R}^n_{\geq 0}$.

As it turns out, the correlation gap is deeply connected to CRSs. In [26], the authors use LP duality to show that the correlation gap of every polytope $\mathcal{P}$ is equal to the largest c for which there exists a c-selectable (offline) CRS. Of particular interest is the correlation gap of submodular functions, as these arise in many real-world problems. For monotone[3] submodular functions, it is known that the correlation gap is $1 - 1/e$ [47], [48]. However, for general, not necessarily monotone, submodular functions, a simple example where $N = \{u, v\}$ and $f(S) = 1 \iff S = \{u\}$ and 0 otherwise shows that the correlation gap can be arbitrarily small, since for $\boldsymbol{x} = (\varepsilon, 1 - \varepsilon)$, we have $F(\boldsymbol{x}) = \varepsilon^2$ and $f^+(\boldsymbol{x}) = \varepsilon$, and thus the correlation gap goes to 0 as $\varepsilon \to 0$. Due to this, it is natural to consider the following questions.

Question 1.8

(a) Can we obtain a fine-grained version of the correlation gap for general submodular functions?

(b) What parameters affect the behaviour of the correlation gap?

In Chapter 5, we provide a positive answer to the first question and identify one of the parameters that affect the behaviour of the correlation gap. Specifically, for a given point $\boldsymbol{x} \in [0, 1]^n$, parametrizing with respect to $p = \max_i x_i$ yields a nice interpolation between the correlation gap of $1 - 1/e$ for monotone submodular functions and 0 for general submodular functions as we vary p from 0 to 1. Therefore, one can obtain a non-trivial correlation gap when $\boldsymbol{x} \in [0, p]^n$, a setting which comes up in many practical applications of CRSs.

1.4 RELATED WORK

1.4.1 Optimal Stopping Problems

Since Krengel, Sucheston and Garling's [1] seminal result on the prophet inequality, several extensions and variants of the problem have been studied. For information on the early results in the literature, see an excellent survey by Correa [49]. As described in §1.2.1, a natural special case of significant importance the I.I.D. prophet inequality, where all distributions D_i are all the same. This setting was originally studied by Hill and Kertz [3], where they showed the best-possible competitive ratio is at least $1 - 1/e \approx 0.632$. Kertz [50] showed, via a recursive approach, that the competitive ratio approaches ≈ 0.745 as the number of random variables grows. A simpler proof of this can be found in [51]. We refer the reader to a survey by Hill and Kertz [52] for the initial results on this setting. Interestingly, the optimal algorithm is surprisingly simple and has been known for a long time; the challenging part is to analyze its performance. Abolhassani, Ehsani, Esfandiari, Hajiaghayi, Kleinberg and Lucier [4] showed that one can achieve a 0.738-competitive ratio, before the problem was finally resolved by Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld [5], who showed that one can achieve the upper bound of 0.745 due to Hill and Kertz. The proofs of both the upper and lower bounds were recently simplified, by [22] and [15] respectively, but perhaps the simplest proof of the optimal ≈ 0.745 competitive ratio can be found in Singla's thesis [53].

[3] A set function $f : 2^N \to \mathbb{R}$ is called monotone if for all $A \subseteq B \subseteq N$ we have $f(A) \leq f(B)$.

When the arrival order of the values is chosen uniformly at random, the model is known as the *prophet secretary* and was originally introduced by Esfandiari, Hajiaghayi, Liaghat and Monemizadeh [12]. They gave an adaptive-threshold algorithm that achieves a $(1 - 1/e)$-competitive ratio and showed no algorithm can achieve a factor better than 0.75. The factor of $1 - 1/e$ was beaten in subsequent work [13], [14], and the current best lower bound is ≈ 0.672 due to Harb [54]. An upper bound of 0.7254 on the competitive ratio was recently shown by Bubna and Chiplunkar [55].

Interestingly, when the algorithm can (adaptively) select the arrival order, based on the distributions given in the input and any values seen so far, strictly better competitive ratios are possible. This is called the *free order* prophet inequality setting, originally introduced by Hill [56]. This setting has been extensively studied, with incremental improvements on the competitive ratio [7], [13], [57]. Recently, Bubna and Chiplunkar [55] showed that one can achieve a 0.7258-competitive ratio in the free order setting, thereby separating it from prophet secretary. An upper bound of 0.745 trivially follows from the I.I.D. setting, which is a special case of the free order setting, and settling whether one can achieve the bound of ≈ 0.745 in the free-order setting as well is a major open problem.

The $1/2$-competitive factor guaranteed by the classical prophet inequality for adversarial arrival order has been shown to hold for more general classes of downwards-closed constraints, all the way up to matroids [8]. A general framework for combinatorial constraints was introduced in [10] via the technique of balanced prices, one of the two main techniques used to generalize prophet inequalities to combinatorial settings; the other being OCRSs. For the prophet secretary setting, Ehsani, Hajiaghayi, Kesselheim and Singla [9] extend the $1 - 1/e$-competitive ratio to matroid constraints.

The combinatorial auctions setting where a seller wants to sell distinct items to buyers that have combinatorial valuation functions for the items and aims to maximize either the social welfare or the revenue, has been studied a lot [10], [11], [58], [59]. Dütting, Feldman, Kesselheim and Lucier [10] obtained a 2-prophet inequality for submodular functions, while Dütting, Kesselheim and Lucier [11] obtained a $\mathrm{O}(\log \log m)$ prophet inequality for subadditive functions. For the latter, the authors also show that achieving a constant factor prophet inequality for subadditive valuation functions is impossible via their techniques and requires a different approach. This question was recently settled by Correa and Cristi in [60], where they showed the existence of a constant-factor prophet inequality for online combinatorial auctions, via a non-constructive fixed-point argument. Interestingly, their result holds even against an almighty adversary and can also be implemented in an incentive compatible way.

Secretary problems are closely related to prophet inequalities. In the classical version, an online algorithm sees n adversarially chosen values in a random order and has to pick one item irrevocably, aiming to maximize the probability of picking the highest value. A classical result of Dynkin [39] shows an optimal competitive ratio of $1/e$. This model has been extensively studied for combinatorial settings as well, where algorithms that achieve constant approximations to the best-possible offline selection are known for several special classes of matroids [61]–[71]. A survey on the secretary problem and variants is due to Dinitz [72].

1.4.2 Online Contention Resolution Schemes

Apart from balanced prices, the other technique that generalizes prophet inequalities to the combinatorial setting is rounding via an OCRS. For this reason, all the results we describe here immediately imply a corresponding prophet inequality for the corresponding setting. This connection between prophet inequalities and OCRSs actually goes deeper, as Lee and Singla [44] designed optimal OCRSs for several settings using ex-ante prophet inequalities.

Contention Resolution Schemes (CRSs) were originally introduced by Chekuri, Vondrák and Zenklusen [26] for the offline case. They gave a $1 - 1/e$-selectable CRS for matroids and used LP duality to show that the best-possible selectability of a feasibility constraint is exactly equal to the correlation gap of its polytope. Later, CRSs were generalized to OCRSs that work for online arrivals of the active elements.

Starting with the special case of k-uniform matroids, where one can select up to k values, Alaei [73] gave an OCRS that is $\left(1 - \frac{1}{\sqrt{k+3}}\right)$-selectable. This was recently improved for small k via the use of a static threshold [74] and later made tight for all k [75]. Perhaps the simplest asymptotically optimal OCRS for k-uniform matroids is due to Dinev and Weinberg [76], where the authors also show that greedy OCRSs for this setting are necessarily suboptimal.

For matching constraints, Ezra, Feldman, Gravin and Tang [29] showed a 0.337-selectable OCRS. Rubinstein [77] considered general downwards-closed feasibility constraints and obtained logarithmic approximations. The best-possible selectability for matchings remains an open question. A more general approach was taken by Feldman, Svensson and Zenklusen [27] in their seminal paper that introduced Online Contention Resolution Schemes. They presented greedy OCRSs for matroids, matchings and knapsack constraints, although whether their greedy OCRSs for matchings and knapsacks are optimal is unknown.

For the case of a uniformly random arrival order and k-uniform matroids, Arnosti and Ma [19] recently gave a surprising and quite beautiful single-threshold *Random Order Contention Resolution Scheme (ROCRS)* that achieves the best-possible competitive ratio of $1 - e^{-k}\frac{k^k}{k!}$; the upper bound is due to Alaei [73]. More general feasibility constraints have also been studied in the random arrival order case, i.e. for matroids [78] and matchings [79]–[81]. Adamczyk and Włodarczyk presented a unified framework for ROCRS in [28]. For more information on the random order model in several other settings, we refer the reader to a survey by Gupta and Singla [82]. ROCRSs have also found several applications to the stochastic probing model [83]–[88] and a good starting point for more information on stochastic probing is Singla's thesis [53].

1.4.3 Applications

The main application of prophet inequalities is on simple and truthful auctions. Given the intractability of the optimal (revenue-maximizing) mechanisms for selling items [89]–[92], the focus of the community turned to designing approximately optimal yet simple mechanisms, where prophet inequalities have proven extremely useful. The works of Hajiaghayi, Kleinberg and Sandholm [6] and Chawla, Hartline, Malec and Sivan [7] pioneered the use of prophet inequalities to analyze (sequential) posted price mechanisms for selling items. Specifically, [6] observed that the problem of designing posted price mechanisms that maximize welfare can be reduced to an appropriate optimal stopping theory problem, and this was extended to revenue-maximizing posted price mechanisms in [7].

This result led to a significant effort to understand how the expected revenue of an optimal posted price mechanism compares to that of the optimal auction [10], [11], [46], [58], [59], [73], [93]–[99]. In a surprising result, Correa, Foncea, Pizarro and Verdugo [17] showed that the reverse direction also holds, establishing an equivalence between finding stopping rules in an optimal stopping problem and designing optimal posted price mechanisms – for more information on these applications see a survey by Lucier [100].

1.5 ORGANIZATION

We present the roadmap for the remainder of this thesis.

Chapter 2: We introduce the I.I.D. Min-Prophet Inequality setting, and provide a unified analysis of both the I.I.D. Max-Prophet and Min-Prophet Inequality from the lens of Extreme Value Theory. We also study single-threshold algorithms for the minimization setting, as well as the competition complexity of both the maximization and minimization settings. This chapter is based on joint work with Ruta Mehta [101], on [102] and on joint work with Victor Verdugo [103].

Chapter 3: We introduce the k-Oracle Prophet Inequality, and show it is equivalent to the top-1-of-k problem. We study the I.I.D. and non-I.I.D. settings, when the objective is maximizing the competitive ratio or the probability of selecting the highest realization. This chapter is based on joint work with Sariel Har-Peled and Farouk Harb [104].

Chapter 4: We investigate the selectability of optimal greedy OCRSs for special classes of matroids, showing a separation between greedy and non-greedy OCRSs. This chapter is based on [105].

Chapter 5: We apply greedy OCRSs to obtain good guarantees on the competitive ratio for submodular objective functions and matroid, matching and knapsack constraints. We also study the correlation gap for general, not necessarily monotone, submodular functions. This chapter is based on joint work with Chandra Chekuri [21].

Chapter 6: We conclude with several open problems and potential further directions in combinatorial optimization under uncertainty.

Chapter 2: DISTRIBUTION-OPTIMAL PROPHET INEQUALITIES

2.1 OVERVIEW

In Chapter 1, we discussed the prophet inequality problem for reward maximization. This chapter introduces a natural counterpart, the prophet inequality for *cost minimization*, and shows that it is a remarkably different problem in terms of the guarantees on the competitive ratio. Our first attempt at obtaining optimal guarantees is via single-threshold algorithms, since these are indeed optimal for the rewards maximization case. When these fail to achieve a constant competitive ratio, we switch to studying the optimal threshold strategy. Our main result is a *unified analysis* of the rewards maximization and cost minimization cases, via the lens of *extreme value theory*.

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

$$\mathbb{E}\left[\text{ALG}\right] \leq \alpha \cdot \mathbb{E}\left[\min_i X_i\right]. \tag{2.1}$$

Prophet inequalities for cost minimization can find application in many settings, like their rewards maximization counterparts. For example, consider a house buyer trying to decide when to buy a house in a sellers' market, where houses are selling fast. When a house arrives with its cost listed, the buyer may have to decide the same day whether to buy it or not. Given that the buyer may have only distributional knowledge of future house prices, the goal is to devise a buying strategy so that the price paid is minimized.

For the rewards maximization setting, the competitive ratio of $1/2$ in the classical prophet inequality is achievable through simple single-threshold algorithms [8], [32] of the form "accept the first $X_i \geq \tau$ for some threshold τ", and is known to be tight. Furthermore, there exist simple online algorithms that achieve constant-factor approximations even for general multi-dimensional settings with complicated constraints (e.g. matroids, matchings, etc) [8], [29], [73], [75], [106]. Motivated by these works, our objective is to obtain similar results to the rewards setting and identify the competitive ratios of the optimal single-threshold algorithm and the optimal (multiple) threshold algorithm.

In this chapter, we study the above questions for the case of independent and identically distributed (I.I.D.) random variables. At first glance, one may wonder why minimization is not equivalent to maximization with *negative* X_i's. The reason is that a strategy for maximization is allowed to not pick any of the X_i's and hence will pick nothing if X_i's are negative. In contrast, in the minimization setting, one of the X_i's has to be selected, and hence such a reduction is impossible.

A first observation is that if the expected value of the given distribution is infinite, then there exists a family of instances, one for every $n \geq 2$, for which the competitive ratio is infinite (Proposition 2.1 due to Lucier [107]). This is because, for any algorithm, regardless of which realization it selects, their expected value will be infinite, whereas the expected minimum of the realizations is finite. This fact prevents any bounded factor approximation for all distributions. In the maximization setting this is impossible as the equivalent would be a distribution supported on $[0,+\infty)$ which has expected value equal to 0, but for which the expected maximum is strictly positive. Since such a distribution cannot exist by definition, the maximization setting avoids such pathological cases.

We thus turn our attention to distributions with finite expected value. In the maximization setting, the worst-case instance is achieved for n going to infinity. Motivated by this, we study fixed distributions that are independent of n and investigate the competitive ratio's behaviour as n goes to infinity. We call this the *Asymptotic Competitive Ratio (ACR)*. Our main contribution is in identifying that the behaviour of the ACR is governed, for any distribution, by a single quantity, the *extreme value index* from the field of extreme value theory.

2.1.1 Extreme Value Theory

The celebrated Extreme Value Theorem, also known as the Fisher-Tippett-Gnedenko Theorem, characterizes the class of distributions for which the maxima and minima of a sample converge in distribution. Our analysis assumes that the given distribution $\mathcal{D}$ belongs to this class, which is a mild assumption as this class is dense in the space of all probability distributions [108].

Recall that for a distribution $\mathcal{D}$ with cdf $F(x)$, the cdf of the maximum of n samples is $F^n(x)$ and the cdf of the minimum of n samples is $1-(1-F(x))^n$.

Theorem 2.1 (Extreme Value Theorem ([109], [110]))**.** *Let $X_1,\ldots,X_n$ be a sequence of I.I.D. random variables with cumulative distribution function F. Suppose that there exist two sequences $a_n>0, b_n\in\mathbb{R}$ such that the following limit converges to a non-degenerate distribution function*

$$\lim_{n\to\infty} F^n(a_n x+b_n) = G_\gamma(x).$$

Then, $G_\gamma(x)$ will be of the form

$$G_\gamma(x) = \begin{cases} \exp\left(-\left(1+\gamma x\right)^{-1/\gamma}\right), & \text{if } \gamma\neq 0,\\ \exp\left(-\exp\left(-x\right)\right), & \text{if } \gamma = 0,\end{cases}$$

for all $1+\gamma x>0$.

Similarly, suppose that there exist two sequences $a_n'>0, b_n'\in\mathbb{R}$ such that the following limit converges to a non-degenerate distribution function

$$\lim_{n\to\infty} \left(1-F(a_n x+b_n)\right)^n = G^*_\gamma(x).$$

Then $G^*_\gamma(x)$ will be of the form

$$G^*_\gamma(x) = \begin{cases} \exp\left(-\left(1-\gamma x\right)^{-1/\gamma}\right), & \text{if } \gamma\neq 0,\\ \exp\left(-\exp\left(x\right)\right), & \text{if } \gamma = 0.\end{cases}$$

for all $1-\gamma x>0$.

In the theorem above, γ is called the *extreme value index*, and it partitions the space of all distributions into equivalence classes based on which G_γ their maxima and minima converge to.

- For $\gamma > 0$, we obtain the *Fréchet class* of distributions.
- For $\gamma = 0$, we obtain the *Gumbel class* of distributions.
- For $\gamma < 0$, we obtain the *Reverse Weibull class* of distributions.

Definition 2.1 (Domain of Attraction)**.** When considering the I.I.D. Max-Prophet Inequality, we say that a distribution $\mathcal{D}$ with cdf F belongs in the domain of attraction of D_γ, indicated as $\mathcal{D} \in D_\gamma$, if there exist two sequences $a_n > 0, b_n \in \mathbb{R}$ such that the following limit converges to a non-degenerate distribution function

$$\lim_{n\to\infty} F^n(a_n x + b_n) = G_\gamma(x).$$

Similarly, when considering the I.I.D. Min-Prophet Inequality, we say that a distribution $\mathcal{D}$ with cdf F belongs in the domain of attraction of D_γ, indicated as $\mathcal{D} \in D_\gamma$, if there exist two sequences $a'_n > 0, b'_n \in \mathbb{R}$ such that the following limit converges to a non-degenerate distribution function

$$\lim_{n\to\infty} (1 - F(a'_n x + b'_n))^n = G^*_\gamma(x).$$

For simplicity, we sometimes write $F \in D_\gamma$ to indicate that the distribution $\mathcal{D}$ which has cdf F belongs in the domain of attraction of D_γ. Whether we refer to the I.I.D. Max-Prophet Inequality or the I.I.D. Min-Prophet Inequality will be clear from context.

One can think of the Extreme Value Theorem as the analogue of the Central Limit Theorem (CLT) for maxima and minima instead of averages; in the same way that the CLT ensures that the (properly scaled) average of n random variables drawn independently from the same distribution $\mathcal{D}$ converges in distribution to a standard Gaussian distribution, as long as their variance is finite, the Fisher-Tippett-Gnedenko Theorem ensures that, if the (properly scaled) maxima and minima of n random variables drawn independently from $\mathcal{D}$ converge in distribution, then they can only converge to a distribution of the form of G_γ for maxima G^*_γ for minima.

In our analysis, the choice of domain for $\mathcal{D}$ is done carefully and without loss of generality. Let $\text{supp}(\mathcal{D}) = [x_*, x^*)$, where $x_* \geq 0$ and $x^* \leq +\infty$. Notice that if $x_* \neq 0$, then in the I.I.D. Min-Prophet Inequality, we can obtain a $(1+\varepsilon)$-competitive ratio, for any ε and any distribution, by setting a threshold equal to $x_* + \varepsilon$. As $n \to \infty$, the probability that there exists a realization below $x_* + \varepsilon$ goes to 1. For this reason, we assume that $\text{supp}(\mathcal{D}) = [0, x^*)$ without loss of generality. In the I.I.D. Max-Prophet Inequality, for distributions with bounded domain, if $x^* < +\infty$, we can obtain a $(1-\varepsilon)$-competitive ratio, for any ε and any distribution, by setting a threshold equal to $x^* - \varepsilon$. Again, as $n \to \infty$, the probability that there exists a realization above $x^* - \varepsilon$ goes to 1. This fact will show up in our analysis. Also, notice that the Extreme Value Theorem for minima shows that $\gamma > 0$ implies that the left-most endpoint of the support of $\mathcal{D}$ has to be $-\infty$. Since these distributions are not studied in the prophet inequality setting, for the I.I.D. Min-Prophet Inequality it must be that $\gamma \leq 0$.

2.1.2 Competition Complexity

Apart from the competitive ratio, another notion of interest in optimal stopping problems such as the prophet inequality is the *competition complexity*. Informally, the competition complexity of a distribution is defined as how many more samples than the prophet, as a multiplicative ratio, an algorithm needs to have in order for the algorithm to outperform the prophet. More formally,

Definition 2.2 (Competition Complexity)**.** Let $\mathcal{D}$ be a distribution. For every $n \in \mathbb{N}$, the *competition complexity* of $\mathcal{D}$

- in the I.I.D. Max-Prophet Inequality, is defined as

$$CC_{M,\mathcal{D}}(n) \triangleq \inf \left\{ c \in \mathbb{N} \,\middle|\, \mathbb{E}\left[\text{ALG}(c \cdot n)\right] \geq \mathbb{E}\left[\max_{i=1}^{n} X_i\right] \right\}.$$

- in the I.I.D. Min-Prophet Inequality, is defined as

$$CC_{m,\mathcal{D}}(n) \triangleq \inf \left\{ c \in \mathbb{N} \,\middle|\, \mathbb{E}\left[\text{ALG}(c \cdot n)\right] \leq \mathbb{E}\left[\min_{i=1}^{n} X_i\right] \right\}.$$

The study of competition complexity in the context of prophet inequalities was introduced by [33], in which the authors show that the competition complexity for maxima can be unbounded in the worst-case, when $\mathcal{D}$ can depend on n. However, using the properties of the expected value of the optimal threshold algorithm and the prophet that we show in our analysis of the competitive ratio, we show that when $\mathcal{D}$ is independent of n, the competition complexity can be bounded by a small constant. We study the *asymptotic competition complexity (ACC)*, defined as $ACC_{M,\mathcal{D}} = \lim_{n\to\infty} CC_{M,\mathcal{D}}(n)$ and $ACC_{m,\mathcal{D}} = \lim_{n\to\infty} CC_{m,\mathcal{D}}(n)$ for maxima and minima respectively, for distributions which are in the domain of D_γ for some γ.

2.1.3 Our Contributions

Single-Threshold Algorithms. Since single-threshold algorithms have proven very successful in providing constant-factor approximations in the maximization setting, we start by asking whether we can achieve similar results for the minimization setting as well. The intuition behind this is that if n is very large, one could set a single threshold close to $\mathbb{E}[\min_i X_i]$ and with good probability there will be at least one realization below the threshold. Unfortunately, this intuition turns out to be wrong, even for simple distributions, like the exponential. We present an example in Section 2.3 explaining why this intuition fails.

For our first positive result, we show that single-threshold algorithms for the I.I.D. Min-Prophet Inequality achieve a poly-logarithmic competitive ratio for distributions which belong in the domain of attraction of D_γ for some γ.

Theorem 2.2. *For the I.I.D. Min-Prophet Inequality and any distribution $F \in D_\gamma$, where $\gamma < 0$, there exists a single-threshold algorithm that achieves a competitive ratio of $\mathrm{O}\left((\log n)^{-\gamma}\right)$. Furthermore, this factor is tight; no single-threshold algorithm can achieve a $\mathrm{o}\left((\log n)^{-\gamma}\right)$ competitive ratio for all instances.*

Optimal Threshold Algorithm. As the theorem above shows, single-threshold algorithms for the I.I.D. Min-Prophet Inequality fail to yield a constant competitive ratio, in contrast with the I.I.D. Max-Prophet Inequality. In our efforts to overcome this, we turn our attention to the optimal (multiple) threshold algorithm. Here, our contribution is twofold:

- We obtain a complete characterization of the I.I.D. Min-Prophet Inequality setting for every distribution $F \in D_\gamma$ for some γ, giving a closed form for the ACR that depends only on γ,

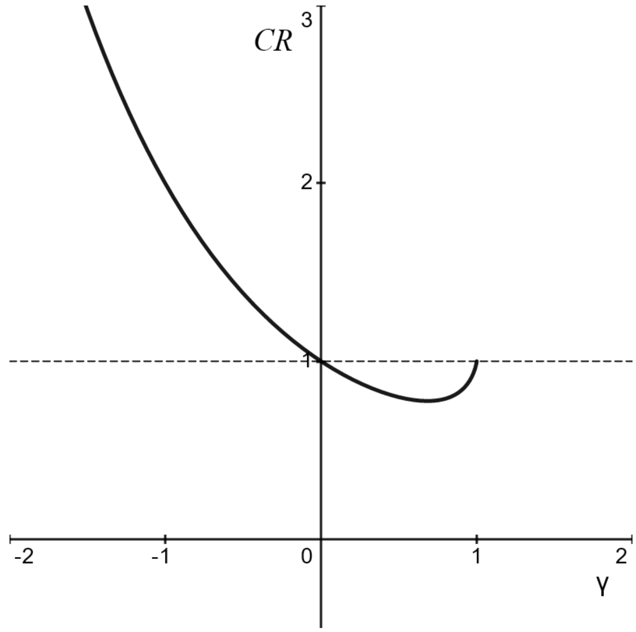

Figure 2.1: $\Lambda(\gamma)$

- We do this in a unified way for both the I.I.D. Max-Prophet Inequality and I.I.D. Min-Prophet Inequality, showing that the ACR for both settings is characterized by a *single* function of γ, where the positive domain of the function corresponds to the ACR for the maximization setting and the negative domain corresponds to the ACR for the minimization setting.

Let $\Gamma(x)$ denote the Gamma function[4]. What follows is our main theorem.

Theorem 2.3. *Let $F \in D_\gamma$. Then,*

- *for the I.I.D. Max-Prophet Inequality, $\gamma \in \mathbb{R}$ and the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is*
$$\lambda_M = \min\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}.$$

- *for the I.I.D. Min-Prophet Inequality, $\gamma \leq 0$ and the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is*
$$\lambda_m = \max\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}.$$

The proof relies on a unified analysis for both maxima and minima and, while relatively simple and easy to follow, makes heavy use of tools from extreme value theory and the theory of regularly varying functions.

We briefly discuss the unique function that characterizes the ACR. Let $\Lambda(\gamma) \triangleq \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}$. Λ can be seen in Figure 2.1.

To understand how $\Lambda(\gamma)$ grows as $\gamma \to -\infty$, consider Stirling's approximation for the Gamma function,

[4]For the definition of the Gamma function and more information, see Section 2.2.4

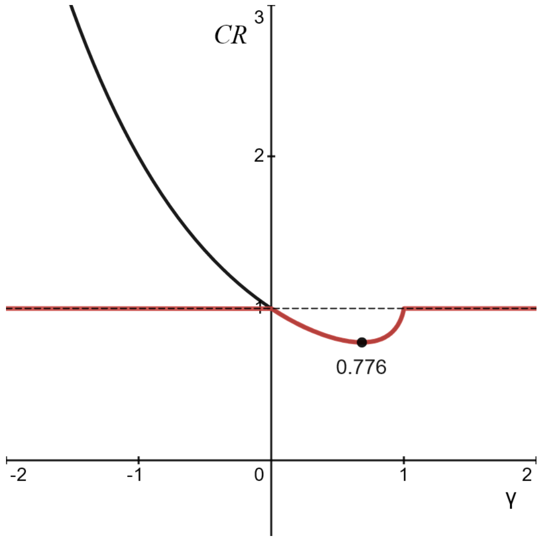

Figure 2.2: ACR(γ) for Max

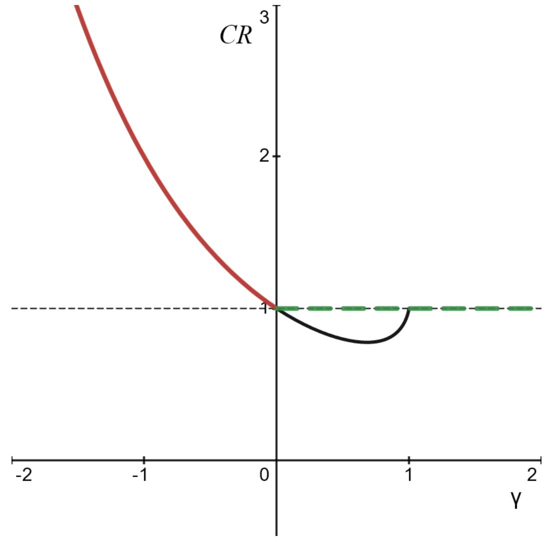

Figure 2.3: ACR(γ) for Min

$\Gamma(z) \approx \frac{\sqrt{2\pi}}{z} \left(\frac{z}{e}\right)^z$. Replacing this in the expression of $\Lambda(\gamma)$, we have

$$\Lambda(\gamma) = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)} \approx \frac{(1-\gamma)^{-\gamma}}{\frac{\sqrt{2\pi}}{1-\gamma}\left(\frac{1-\gamma}{e}\right)^{1-\gamma}} = \Theta\left(e^{-\gamma}\right). \tag{2.2}$$

Thus, the dependence of Λ on γ is (approximately) exponential, for $\gamma \to -\infty$.

Our main theorem essentially states that the interesting cases of the ACR for the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality, in which it is not equal to 1, are captured by the positive and negative domains of Λ, respectively. i.e. $\lambda_M = \Lambda(\gamma)$ for $\gamma \geq 0$ in the I.I.D. Max-Prophet Inequality and $\lambda_m = \Lambda(\gamma)$ for $\gamma \leq 0$ in the I.I.D. Min-Prophet Inequality.

MHR Distributions. Distributions with monotonically increasing hazard rate have been extensively studied in the mechanism design literature due to their sought after properties and applications (e.g., see [111]–[118]). These are known as *monotone hazard rate (MHR)* (also known as *increasing failure rate (IFR)*) distributions. For our third result, we show that for the special case of MHR distributions, $\gamma \leq 0$ for maxima and $\gamma \geq -1$ for minima. Due to the former, we recover the result of [96] which states that the competitive ratio of the I.I.D. Max-Prophet Inequality goes to 1 as n goes to infinity and the distribution is MHR. Due to the latter and the fact that $\Lambda(-1) = 2$, Theorem 2.3 implies the following.

Theorem 2.4. *In the I.I.D. Max-Prophet Inequality, the optimal threshold strategy is* 1*-competitive for every MHR distribution in the domain of attraction of* D_γ *for some* γ.

In the I.I.D. Min-Prophet Inequality, the optimal threshold strategy is 2*-competitive for every MHR distribution in the domain of attraction of* D_γ *for some* γ*. Furthermore, the factor of* 2 *is tight, since there is no* $(2-\varepsilon)$*-competitive algorithm for any* $\varepsilon > 0$ *for the exponential distribution, which is MHR.*

Competition Complexity. At first glance, one could think that, since the asymptotic competitive ratio is 1 for maxima and $\gamma \leq 0$ and for minima and $\gamma = 0$, this implies that the competition complexity for these cases is 1. However, this is not the case; as an example, consider the exponential distribution with

$F(x) = 1 - e^{-x}$, for which $\gamma = 0$ for maxima. We have $\mu_{n:n} \geq \log n + \gamma^*$ for every $n \geq 1$, where $\gamma^* \approx 0.577$ is the Euler-Mascheroni constant, and one can easily verify via induction that $G_M(n) \leq \log n + 1/6$ for $n \geq 4$. Assume that the competition complexity for some $n \geq 4$ is at most $3/2$; then, it must be that $\log(3/2 \cdot n) + 1/6 \geq \log n + \gamma^*$, which is not true for any $n \in \mathbb{N}$. In fact, we will later see that the competition complexity for the exponential distribution (and all distributions in the domain of attraction of D_0) is $e^{\gamma^*} \approx 1.781$.

Our main result for the asymptotic competition complexity is its characterization by a single function that depends only on γ. In fact, this function is the same for both maxima and minima, yielding a qualitatively different result than our result for the competitive ratio. Thus, we obtain that the ACC is the same for both the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality. Via our analysis of the (optimal) algorithm and prophet's values, we obtain the following result.

Theorem 2.5. *For every $F \in D_\gamma$, we have that*

$$ACC_{M,\mathcal{D}} = ACC_{m,\mathcal{D}} = (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}.$$

Recall that $\Lambda(\gamma) = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}$ is the function characterizing the competitive ratio for distributions in D_γ, and notice that the ACC for both maxima and minima is equal to $\Lambda(\gamma)^{-1/\gamma}$.

As a corollary of Theorem 2.5, we can upper bound the ACC for every class of distributions described by the Extreme Value Theorem.

Corollary 2.1. *Let $\mathcal{D}$ be a distribution in the domain of attraction of D_γ for some $\gamma \in \mathbb{R}$.*

- *If $\mathcal{D}$ is in the Fréchet family ($\gamma > 0$), then $ACC_{M,\mathcal{D}} \leq e^{\gamma^*} \approx 1.781$, where $\gamma^* \approx 0.577$ is the Euler-Mascheroni constant, obtained as $\gamma \to 0$.*
- *If $\mathcal{D}$ is in the Gumbel family ($\gamma = 0$), then $ACC_{M,\mathcal{D}}, ACC_{m,\mathcal{D}} \leq e^{\gamma^*} \approx 1.781$, where $\gamma^* \approx 0.577$ is the Euler-Mascheroni constant.*
- *If $\mathcal{D}$ is in the Reverse Weibull family ($\gamma < 0$), then $ACC_{M,\mathcal{D}}, ACC_{m,\mathcal{D}} \leq e$, obtained as $\gamma \to -\infty$.*

An interesting observation about Theorem 2.5 and Corollary 2.1 is that the Fréchet family has better competition complexity than the Gumbel family which in turn has better competition complexity than the Weibull family. This is in stark contrast with the competitive ratio in the maximization setting, where the ratio is worse for distributions in the Fréchet family.

2.1.4 Related Work

Our work is most closely related to the long line of work that considers the case of I.I.D. random variables drawn from a known distribution, which dates back to Hill and Kertz [3]. As stated previously, Kertz [50] showed that the competitive ratio in the I.I.D. case approaches ≈ 0.745 as n goes to infinity and conjectured its tightness. A simpler proof of this can be found in [51]. The bound of ≈ 0.745 was shown to be tight by Correa, Foncea, Hoeksma, Oosterwijk and Vredeveld [5]. The proofs of both the upper and lower bounds were recently simplified, by [22] and [15] respectively.

For minimization in the more general setting where the distributions are not necessarily identical, Esfandiari, Hajiaghayi, Liaghat and Monemizadeh [12] observed that no algorithm can achieve any bounded competitive factor. Recently, [119] initiated the study of buy-and-sell prophet inequalities, named trading prophets, a model that is similar yet orthogonal to the minimization setting. They obtained constant factor guarantees on the competitive ratio using single-threshold algorithms.

For more information on other related work, see §1.4.

Organization. Section 2.2 contains all relevant definitions and technical background needed for our results. In Section 2.3, we show that simple-threshold algorithms achieve a poly-logarithmic competitive ratio. Section 2.4 characterizes the optimal (multiple) threshold algorithm and contains our unified analysis for both the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality. Finally, in Section 2.5, we study the competition complexity of both the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality for large n.

2.2 PRELIMINARIES

Let $X_1, \ldots, X_n$ denote random variables drawn independently from a known distribution $\mathcal{D}$ supported on $[0, x^*)$, where $x^* \leq +\infty$. In the prophet inequality setting, we are presented with the realizations of X_i sequentially and at each step i we must make an immediate and irrevocable decision to accept or reject X_i. The process ends when we accept a realization and, once rejected, a realization cannot be obtained in the future. The benchmark is an all-knowing prophet who can see all realizations in advance and always select the optimal. In the I.I.D. Max-Prophet Inequality, the goal is to maximize the selected realization and the prophet's objective is $\mathbb{E}[\max_i X_i]$, whereas in the I.I.D. Min-Prophet Inequality, the goal is to minimize the selected realization and the prophet's objective is $\mathbb{E}[\max_i X_i]$ (but we are forced to select at least one realization).

Suppose we reorder the random variables such that $X_{(1)} \leq \cdots \leq X_{(n)}$. Then, $X_{(i)}$ is called the i-th *order statistic* of $\mathcal{D}$. We denote the expected value of the i-th order statistic of n samples from $\mathcal{D}$ by $\mu_{i:n} = \mathbb{E}_{\mathcal{D}}[X_{(i)}]$. Of special interest to us are the expectation of the *last* and *first* order statistic, i.e. the largest $\mu_{n:n}$ and smallest $\mu_{1:n}$ expected values, since they capture the prophet's objective in the I.I.D. Max-Prophet Inequality and I.I.D. Min-Prophet Inequality, respectively.

Let $F : [0, +\infty) \to [0, 1]$, where $F(x) = \Pr_{X \sim \mathcal{D}}[X \leq x]$, and $f : [0, +\infty) \to [0, 1]$ denote the *Cumulative Distribution Function (CDF)* and *Probability Density Function (PDF)* of $\mathcal{D}$, respectively.

Definition 2.3 (Left Continuous Inverse)**.** Let F be a non-decreasing function on $\mathbb{R}$. The *(left continuous) inverse* of F is defined as

$$F^{\leftarrow}(y) = \inf \{x \mid F(x) \geq y\}.$$

The above definition works with the convention that the infimum of an empty set is $+\infty$. For more information on the left continuous inverse see [120] (Section 0.2). In particular, if F denotes the CDF of a distribution $\mathcal{D}$, then $F^{\leftarrow}(y)$ denotes the quantile function of $\mathcal{D}$, i.e. $F^{\leftarrow}(y)$ is the smallest value τ for which $\Pr_{X \sim \mathcal{D}}[X \leq \tau] \geq y$.

We use $G_M(n)$ and $G_m(n)$ to denote the expected value of the optimal threshold policy for the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality settings with n random variables, respectively. To avoid confusion, we denote the asymptotic competitive ratio in the I.I.D. Max-Prophet Inequality by λ_M and in the I.I.D. Min-Prophet Inequality by λ_m.

Definition 2.4 (Asymptotic Competitive Ratio)**.** Consider an instance of the I.I.D. Max-Prophet Inequality (resp. I.I.D. Min-Prophet Inequality) with n random variables. Then, the *asymptotic competitive ratio* λ is

$$\lambda_M \triangleq \lim_{n\to\infty} \frac{G_M(n)}{\mu_{n:n}}, \qquad \text{in the I.I.D. Max-Prophet Inequality,}$$
$$\lambda_m \triangleq \lim_{n\to\infty} \frac{G_m(n)}{\mu_{1:n}}, \qquad \text{in the I.I.D. Min-Prophet Inequality.}$$

Since we analyze the asymptotic competitive ratio as $n \to \infty$, we use $\approx$ to denote asymptotic equality, for brevity and ease of presentation. In other words, whenever we write $a(n) \approx b(n)$ for two expressions $a(n), b(n)$ that depend on n, it implies that $\lim_{n\to\infty} \frac{a(n)}{b(n)} = 1$.

Next, we show that in the I.I.D. Min-Prophet Inequality, one cannot hope to obtain a bounded competitive factor for *all* distributions. The following counterexample is due to Lucier [107].

Proposition 2.1 ([107])**.** *For any $n \geq 2$, there exists an instance of the I.I.D. cost prophet inequality problem for which no algorithm is α-competitive for any $\alpha > 0$.*

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

$$\mathbb{E}[X] = \int_0^\infty (1 - F(x))\, dx = 1 + \int_1^\infty (1 - F(x))\, dx = 1 + \int_1^\infty \frac{1}{x}\, dx = +\infty. \tag{2.3}$$

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

$$\text{OPT} = \mu_{1:2} = \int_0^\infty (1 - F(x))^2\, dx = 1 + \int_1^\infty \frac{1}{x^2}\, dx = 2. \tag{2.4}$$

Therefore, no algorithm can achieve a finite competitive ratio.

Notice that the above counterexample can be easily extended to any $n > 2$. Due to the recursive nature of the optimal online algorithm, we have $\mathbb{E}[ALG] = +\infty$ regardless of which X_i the algorithm chooses to stop at. However, $\mu_{1:n}$ is finite for any $n \geq 2$. QED

2.2.1 Hazard Rate

Our results make use of the *hazard (failure) rate* of a distribution. We refer the reader to [121] for an extensive overview. Intuitively, for discrete distributions, the hazard rate for maxima (resp. minima) at a point t represents the probability that an event occurs at time t, given that the event has not occurred (resp. has occured) up to time t. For continuous distributions, the hazard rate instead quantifies the instantaneous rate of the event's occurrence at time t.

Definition 2.5 (Hazard Rate)**.** For a distribution $\mathcal{D}$ with cumulative distribution function F and probability density function f, the *maxima hazard rate* h and *minima hazard rate* r of $\mathcal{D}$ are defined as

$$h(x) \triangleq \frac{f(x)}{1 - F(x)} \qquad \text{and} \qquad r(x) \triangleq \frac{f(x)}{F(x)},$$

for all x in the support of $\mathcal{D}$.

Furthermore, let H and R denote the antiderivatives of h and r, which we call the *cumulative hazard rates* of $\mathcal{D}$ for maxima and minima respectively:

$$H(x) \triangleq \int_0^x h(u)\,du \qquad \text{and} \qquad R(x) \triangleq \int_0^x r(u)\,du.$$

Notice that,

$$H(x) = \int_0^x h(u)\,du = \int_0^x \frac{f(u)}{1-F(u)}\,du = -\int_0^x \left(\ln\left(1-F(u)\right)\right)'\,du = -\ln\left(1-F(x)\right), \tag{2.5}$$

which implies that $1 - F(x) = e^{-H(x)}$. Similarly,

$$R(x) = \int_0^x r(u)\,du = \int_0^x \frac{f(u)}{F(u)}\,du = \int_0^x \left(\ln\left(F(u)\right)\right)'\,du = \ln\left(F(x)\right), \tag{2.6}$$

which implies that $F(x) = e^{R(x)}$.

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

Definition 2.6 (Monotone Hazard Rate Distribution)**.** A distribution $\mathcal{D}$ is called a *Monotone Hazard Rate (MHR)* distribution if and only if the hazard rate function h of $\mathcal{D}$ is non-decreasing.

2.2.2 Regularly Varying Functions

Our analysis relies heavily on the theory of *regularly varying functions*, originally developed by Karamata. Regularly varying functions are, roughly speaking, functions that behave asymptotically like power functions. We present some basic definitions here; for more information on the topic see [120] and [122].

Definition 2.7. Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}$ be a Lebesgue measurable function that is eventually positive. We say that f is *regularly varying* (at infinity) if, for some $\alpha \in \mathbb{R}$ and every $x > 0$

$$\lim_{t\to\infty} \frac{f(tx)}{f(t)} = x^\alpha.$$

In this case, we indicate this as $f \in RV_\alpha$.

In the above definition, α is called the *index* of regular variation and whenever $\alpha = 0$, we say that f is *slowly varying*. Furthermore, we say that $f(x)$ is regularly varying at 0 if and only if $f(1/x)$ is regularly varying at infinity. If $f \in RV_\alpha$, then $L(x) = f(x)/x^\alpha \in RV_0$. In general, one can represent any $f \in RV_\alpha$ as $f(x) = x^\alpha L(x)$, where L is a slowly-varying function.

Next, we present a connection between distributions $F \in D_\gamma$ for some γ and regular variation. In particular, we show that the quantile function $F^{\leftarrow}$ of such a distribution is a regularly varying function. Let $U_M(x) \triangleq F^{\leftarrow}(1-x)$ and $U_m(x) \triangleq F^{\leftarrow}(x)$. Since $1 - F(x) = e^{-H(x)}$ and $F(x) = e^{R(x)}$, we have that $U_M(x) = H^{\leftarrow}(-\log x)$ and $U_m(x) = R^{\leftarrow}(\log x)$.

Lemma 2.1. *If $F \in D_\gamma$, then $U_m \in RV_{-\gamma}$ at 0. Furthermore, for $\gamma \in [0, 1)$, we have $U_M \in RV_{-\gamma}$ at 0, and for $\gamma \leq 0$, we have $x^* < +\infty$ and $x^* - U_M \in R_{-\gamma}$ at 0.*

Proof. Let $\gamma \in [0,1)$. Since $F \in D_\gamma$, by [31] (Corollary 1.2.10), we have that $U_M\left({}^1\!/\!_x\right) \in RV_\gamma$ which implies that

$$\lim_{t\to\infty} \frac{U_M\left({}^1\!/\!_{tx}\right)}{U_M\left({}^1\!/\!_t\right)} = x^\gamma. \tag{2.7}$$

Therefore,

$$\lim_{t\to 0^+} \frac{U_M(tx)}{U_M(t)} = x^{-\gamma}, \tag{2.8}$$

and $U_M \in RV_{-\gamma}$ at 0.

Next, let $\gamma < 0$. Since $F \in D_\gamma$, by [31] (Corollary 1.2.10), we have that $x^* - U_M\left({}^1\!/\!_x\right)$, $U_m\left({}^1\!/\!_x\right) \in RV_\gamma$, which implies that

$$\lim_{t\to\infty} \frac{x^* - U_M\left({}^1\!/\!_{tx}\right)}{x^* - U_M\left({}^1\!/\!_t\right)} = \lim_{t\to\infty} \frac{U_m\left({}^1\!/\!_{tx}\right)}{U_m\left({}^1\!/\!_t\right)} = x^\gamma. \tag{2.9}$$

Therefore,

$$\lim_{t\to 0^+} x^* - \frac{U_M(tx)}{x^* - U_M(t)} = \lim_{t\to 0^+} \frac{U_m(tx)}{U_m(t)} = x^{-\gamma}, \tag{2.10}$$

and $x^* - U_M, U_m \in RV_{-\gamma}$ at 0. QED

Next, we present a famous theorem in the theory of regularly varying functions, due to Karamata, that is very useful in all our results. For more details on this see [31] (Appendix B) and [120].

Lemma 2.2 (Karamata's Theorem)**.** *For $F \in D_\gamma$ and large enough n, we have*

$$\int_0^x F^{\leftarrow}(1-u)\,du \approx \frac{xF^{\leftarrow}(1-x)}{1-\gamma}, \quad \textit{for } \gamma \le 1,$$

and

$$\int_0^x F^{\leftarrow}(u)\,du \approx \frac{xF^{\leftarrow}(x)}{1-\gamma}, \quad \textit{for } \gamma \le 0.$$

In particular,

$$\int_0^{\exp(-H(G_M(n-1)))} H^{\leftarrow}(-\log u)\,du \approx \frac{G_M(n-1)\cdot e^{-H(G_M(n-1))}}{1-\gamma}, \quad \textit{for } \gamma \le 1,$$

and

$$\int_0^{\exp(R(G_m(n-1)))} R^{\leftarrow}(\log u)\,du \approx \frac{G_m(n-1)\cdot e^{R(G_m(n-1))}}{1-\gamma}, \quad \textit{for } \gamma \le 0.$$

Proof. The first two asymptotic equalities are due to Karamata ([31], Theorem B.1.5). Recall that $H^{\leftarrow}(-\log x), R^{\leftarrow}(\log x) \in RV_{-\gamma}$ at 0, by Lemma 2.1. Since, as $n \to \infty$, $e^{-H(G_M(n-1))}, e^{R(G_m(n-1))} \to 0$, the other two asymptotic equalities follow directly from the first two. QED

2.2.3 Expectations of Order Statistics

We now turn our attention to the expectation of the first and last order statistic of F, i.e. to $\mu_{1:n}$ and $\mu_{n:n}$. Here, we provide a few lemmas that will be very useful in our analysis of both single-threshold and multiple-threshold algorithms.

Lemma 2.3. *For $F \in D_\gamma$ and large enough n, we have*

$$\mu_{n:n} \approx \begin{cases} \Gamma(1-\gamma)F^{\leftarrow}\left(1-\frac{1}{n}\right), & \text{for } \gamma \in (0,1), \\ F^{\leftarrow}\left(1-\frac{e^{-\gamma^*}}{n}\right), & \text{for } \gamma = 0, \\ x^* - \Gamma(1-\gamma)\left(x^* - F^{\leftarrow}\left(1-\frac{1}{n}\right)\right), & \text{for } \gamma < 0, \end{cases}$$

and

$$\mu_{1:n} \approx \begin{cases} \Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n}\right), & \text{for } \gamma < 0, \\ F^{\leftarrow}\left(\frac{e^{-\gamma^*}}{n}\right), & \text{for } \gamma = 0. \end{cases}$$

where $\gamma^* \approx 0.577$ is the Euler-Mascheroni constant.

Proof. Since $F \in D_\gamma$, by Theorem 2.1, we know that there exist sequences $a_n > 0, b_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} F^n(a_n x + b_n) = G_\gamma(x). \tag{2.11}$$

Let $M_n = \max\{X_1, \dots, X_n\}$. If (2.11) is satisfied, it has to be satisfied for

- $a_n = U_M\left({}^1\!/n\right)$ and $b_n = 0$, if $\gamma \in (0,1)$,
- $b_n = U_M\left({}^1\!/n\right)$ and appropriately chosen a_n if $\gamma = 0$, and
- $a_n = x^* - U_M\left({}^1\!/n\right)$ and $b_n = x^*$, if $\gamma \in< 0$,

by [31] (Corollary 1.2.4). Let $Y_n = \max\left\{{}^{X_1 - b_n}\!/a_n, \dots, {}^{X_n - b_n}\!/a_n\right\} = \frac{M_n - b_n}{a_n}$. The above imply that $\lim_{n\to\infty} Y_n$ converges in distribution to a random variable Z distributed according to G_γ. Notice that, for $0 < \gamma < 1$, we have $\mathbb{E}[Z] = \Gamma(1-\gamma)$, for $\gamma = 0$, we have $\mathbb{E}[Z] = \gamma^*$ and for $\gamma < 0$, we have $\mathbb{E}[Z] = -\Gamma(1-\gamma)$ ([31], Theorem 5.3.1). This implies that, for large enough n and $\gamma \in (0,1)$

$$\mathbb{E}[Y_n] \approx \Gamma(1-\gamma) \iff \mu_{n:n} \approx \Gamma(1-\gamma)a_n = \Gamma(1-\gamma)F^{\leftarrow}\left(1-\frac{1}{n}\right). \tag{2.12}$$

For $\gamma = 0$, we have

$$\mathbb{E}[Y_n] \approx \gamma^* \iff \mu_{n:n} \approx a_n\gamma^* + b_n = a_n\gamma^* + U_M\left({}^1\!/n\right) = a_n\gamma^* + F^{\leftarrow}\left(1-\frac{1}{n}\right). \tag{2.13}$$

Finally, for $\gamma = 0$, by [31] (Theorem 1.1.6), we have that, for any $x > 1$

$$a_n \approx \frac{U_M\left({}^1\!/xn\right) - U_M\left({}^1\!/n\right)}{\log x} = \frac{F^{\leftarrow}\left(1-\frac{1}{xn}\right) - F^{\leftarrow}\left(1-\frac{1}{n}\right)}{\log x}. \tag{2.14}$$

Combining (2.13) and (2.14), we get

$$\mu_{n:n} \approx \frac{\gamma^*}{\log x}F^{\leftarrow}\left(1-\frac{1}{xn}\right) + \left(1-\frac{\gamma^*}{\log x}\right)F^{\leftarrow}\left(1-\frac{1}{n}\right). \tag{2.15}$$

Setting $x = e^{\gamma^*}$ yields

$$\mu_{n:n} \approx F^{\leftarrow}\left(1-\frac{e^{-\gamma^*}}{n}\right). \tag{2.16}$$

Finally, for $\gamma < 0$, we have

$$\mathbb{E}[Y_n] \approx -\Gamma(1-\gamma) \iff \mu_{n:n} \approx -\Gamma(1-\gamma)a_n + b_n = x^* - \Gamma(1-\gamma)\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right). \tag{2.17}$$

Similarly, by Theorem 2.1, we know that there exist sequences $a'_n > 0, b'_n \in \mathbb{R}$ such that

$$\lim_{n\to\infty} (1 - F(a_n x + b_n))^n = G^*_\gamma(x). \tag{2.18}$$

Let $m_n = \min\{X_1, \ldots, X_n\}$. If (2.18) is satisfied, it has to be satisfied for $a'_n = U_m\left({}^1\!/\!{}_n\right)$ and $b'_n = 0$, if $\gamma \in< 0$ and for $b'_n = U_m\left({}^1\!/\!{}_n\right)$ and appropriately chosen a'_n if $\gamma = 0$ ([31], Corollary 1.2.4). Let $Y'_n = \min\left\{{}^{X_1 - b'_n}\!/\!{}_{a'_n}, \ldots, {}^{X_n - b'_n}\!/\!{}_{a'_n}\right\} = \frac{m_n - b'_n}{a'_n}$. The above imply that $\lim_{n\to\infty} Y'_n$ converges in distribution to a random variable Z' distributed according to G^*_γ. Notice that, for $\gamma < 0$, we have $\mathbb{E}[Z'] = \Gamma(1-\gamma)$, whereas for $\gamma = 0$, we have $\mathbb{E}[Z'] = -\gamma^*$. This implies that, for large enough n,

$$\mathbb{E}[Y'_n] \approx \Gamma(1-\gamma) \iff \mu_{1:n} \approx \Gamma(1-\gamma)a'_n = \Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n}\right), \tag{2.19}$$

whereas for $\gamma = 0$

$$\mathbb{E}[Y'_n] \approx -\gamma^* \iff \mu_{1:n} \approx -a'_n\gamma^* + b'_n = -a'_n\gamma^* + U_m\left({}^1\!/\!{}_n\right) = -a'_n\gamma^* + F^{\leftarrow}\left(\frac{1}{n}\right). \tag{2.20}$$

Finally, for $\gamma = 0$, by [31] (Theorem 1.1.6), we have that, for any $x > 1$

$$a'_n \approx \frac{U_m\left({}^1\!/\!{}_{xn}\right) - U_m\left({}^1\!/\!{}_n\right)}{-\log x} = \frac{F^{\leftarrow}\left(\frac{1}{xn}\right) - F^{\leftarrow}\left(\frac{1}{n}\right)}{-\log x}. \tag{2.21}$$

Combining (2.20) and (2.21), we get

$$\mu_{1:n} \approx \frac{\gamma^*}{\log x}F^{\leftarrow}\left(\frac{1}{xn}\right) + \left(1 - \frac{\gamma^*}{\log x}\right)F^{\leftarrow}\left(\frac{1}{n}\right). \tag{2.22}$$

Setting $x = e^{\gamma^*}$ yields

$$\mu_{1:n} \approx F^{\leftarrow}\left(\frac{e^{-\gamma^*}}{n}\right). \tag{2.23}$$

QED

The asymptotic expression for $\mu_{n:n}$ and $\mu_{1:n}$ leads us to consider alternative representations of the tail quantiles of F.

Lemma 2.4. *For every $F \in D_\gamma$, $c > 0$ and large enough n, we have*

- *For $\gamma \in (0, 1)$,*
$$F^{\leftarrow}\left(1 - \frac{c}{n}\right) \approx c^{-\gamma}F^{\leftarrow}\left(1 - \frac{1}{n}\right).$$
- *For $\gamma < 0$,*
$$x^* - F^{\leftarrow}\left(1 - \frac{c}{n}\right) \approx c^{-\gamma}\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right),$$

and

$$F^{\leftarrow}\left(\frac{c}{n}\right) \approx c^{-\gamma} F^{\leftarrow}\left(\frac{1}{n}\right).$$

Proof. Let $\gamma \in (0,1)$. Since $F \in D_\gamma$, we know that $U_M\left({}^1\!/\!_x\right) \in RV_\gamma$ ([31], Corollary 1.2.10). By the definition of regular variation, we have

$$\lim_{n\to\infty} \frac{U_M\left({}^1\!/\!_{xn}\right)}{U_M\left({}^1\!/\!_n\right)} = x^\gamma, \tag{2.24}$$

for all $x > 0$. Recall that $U_M\left({}^1\!/\!_x\right) = F^{\leftarrow}\left(1 - \frac{1}{x}\right)$, and thus

$$\lim_{n\to\infty} \frac{F^{\leftarrow}\left(1 - \frac{1}{xn}\right)}{F^{\leftarrow}\left(1 - \frac{1}{n}\right)} = x^\gamma, \tag{2.25}$$

for all $x > 0$. For $x = \frac{1}{c}$, we obtain, for large enough n

$$F^{\leftarrow}\left(1 - \frac{c}{n}\right) \approx c^{-\gamma} F^{\leftarrow}\left(1 - \frac{1}{n}\right). \tag{2.26}$$

Next, let $\gamma < 0$. Since $F \in D_\gamma$, we know that $x^* - U_M\left({}^1\!/\!_x\right), U_m\left({}^1\!/\!_x\right) \in RV_\gamma$ ([31], Corollary 1.2.10). By the definition of regular variation, we have

$$\lim_{n\to\infty} \frac{x^* - U_M\left({}^1\!/\!_{xn}\right)}{x^* - U_M\left({}^1\!/\!_n\right)} = \lim_{n\to\infty} \frac{U_m\left({}^1\!/\!_{xn}\right)}{U_m\left({}^1\!/\!_n\right)} = x^\gamma, \tag{2.27}$$

for all $x > 0$. Recall that $x^* - U_M\left({}^1\!/\!_x\right) = x^* - F^{\leftarrow}\left(1 - \frac{1}{x}\right)$ and $U_m\left({}^1\!/\!_x\right) = F^{\leftarrow}\left(\frac{1}{x}\right)$. Therefore,

$$\lim_{n\to\infty} \frac{x^* - F^{\leftarrow}\left(1 - \frac{1}{xn}\right)}{x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)} = \lim_{n\to\infty} \frac{F^{\leftarrow}\left(\frac{1}{xn}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} = x^\gamma, \tag{2.28}$$

for all $x > 0$. For $x = \frac{1}{c}$, we obtain, for large enough n

$$x^* - F^{\leftarrow}\left(1 - \frac{c}{n}\right) \approx c^{-\gamma}\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right), \tag{2.29}$$

and

$$F^{\leftarrow}\left(\frac{c}{n}\right) \approx c^{-\gamma} F^{\leftarrow}\left(\frac{1}{n}\right). \tag{2.30}$$

QED

Next, we use Lemmas 2.3 and 2.4 to characterize exactly how the ratio of the prophet's expected value for $n-1$ and n behaves as $n \to \infty$.

Lemma 2.5. *For $F \in D_\gamma$ and large enough n, we have*

- *For $\gamma \in (0,1)$,*

$$\frac{\mu_{n-1:n-1}}{\mu_{n:n}} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right).$$

- *For $\gamma < 0$,*

$$\frac{x^* - \mu_{n-1:n-1}}{x^* - \mu_{n:n}} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right).$$

and
$$\frac{\mu_{1:n-1}}{\mu_{1:n}} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right).$$

Proof. Let $\gamma \in (0,1)$. Notice that, by Lemma 2.3, for large enough n, we have

$$\mu_{n-1:n-1} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{1}{n-1}\right) = \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{c}{n}\right), \tag{2.31}$$

for $c = 1 + \frac{1}{n-1}$. Using Lemma 2.4, we obtain

$$\mu_{n-1:n-1} \approx \Gamma(1-\gamma)\left(1 + \frac{1}{n-1}\right)^{-\gamma} F^{\leftarrow}\left(1 - \frac{1}{n}\right). \tag{2.32}$$

Also, again by Lemma 2.3, we have

$$\mu_{n:n} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{1}{n}\right), \tag{2.33}$$

and thus, for large enough n

$$\frac{\mu_{n-1:n-1}}{\mu_{n:n}} \approx \left(1 + \frac{1}{n-1}\right)^{-\gamma} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right). \tag{2.34}$$

Similarly, for $\gamma < 0$, by Lemma 2.3, for large enough n, we have

$$x^* - \mu_{n-1:n-1} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{1}{n-1}\right) = \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{c}{n}\right), \tag{2.35}$$

and
$$\mu_{1:n-1} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n-1}\right) = \Gamma(1-\gamma)F^{\leftarrow}\left(\frac{c}{n}\right), \tag{2.36}$$

for $c = 1 + \frac{1}{n-1}$. Using Lemma 2.4, we obtain

$$x^* - \mu_{n-1:n-1} \approx \Gamma(1-\gamma)\left(1 + \frac{1}{n-1}\right)^{-\gamma} F^{\leftarrow}\left(1 - \frac{1}{n}\right), \tag{2.37}$$

and
$$\mu_{1:n-1} \approx \Gamma(1-\gamma)\left(1 + \frac{1}{n-1}\right)^{-\gamma} F^{\leftarrow}\left(\frac{1}{n}\right). \tag{2.38}$$

Also, again by Lemma 2.3, we have

$$x^* - \mu_{n:n} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{1}{n}\right), \tag{2.39}$$

and
$$\mu_{1:n} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n}\right), \tag{2.40}$$

and thus, for large enough n

$$\frac{x^* - \mu_{n-1:n-1}}{x^* - \mu_{n:n}} \approx \left(1 + \frac{1}{n-1}\right)^{-\gamma} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right), \tag{2.41}$$

and

$$\frac{\mu_{1:n-1}}{\mu_{1:n}} \approx \left(1 + \frac{1}{n-1}\right)^{-\gamma} = 1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right). \tag{2.42}$$

QED

Successive applications of the lemma above yield the following corollary.

Corollary 2.2. *For $F \in D_\gamma$, large enough n and $m < n$, we have*

- *For $\gamma \in (0,1)$,*

$$\frac{\mu_{m:m}}{\mu_{n:n}} = \frac{\Gamma(m+1)}{\Gamma(n+1)} \cdot \frac{\Gamma(n+1-\gamma)}{\Gamma(m+1-\gamma)} + \mathrm{o}\left(\frac{1}{n}\right).$$

- *For $\gamma < 0$,*

$$\frac{x^* - \mu_{m:m}}{x^* - \mu_{n:n}} = \frac{\Gamma(m+1)}{\Gamma(n+1)} \cdot \frac{\Gamma(n+1-\gamma)}{\Gamma(m+1-\gamma)} + \mathrm{o}\left(\frac{1}{n}\right).$$

  and

$$\frac{\mu_{1:m}}{\mu_{1:n}} = \frac{\Gamma(m+1)}{\Gamma(n+1)} \cdot \frac{\Gamma(n+1-\gamma)}{\Gamma(m+1-\gamma)} + \mathrm{o}\left(\frac{1}{n}\right).$$

In particular, taking the series expansion of $\Gamma(\cdot)$ around infinity, it is easy to see that

$$\begin{aligned}
\frac{\Gamma(m+1)}{\Gamma(n+1)} \cdot \frac{\Gamma(n+1-\gamma)}{\Gamma(m+1-\gamma)} &= 1 - \frac{k\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right), \ \textit{for } m = n-k \\
\frac{\Gamma(m+1)}{\Gamma(n+1)} \cdot \frac{\Gamma(n+1-\gamma)}{\Gamma(m+1-\gamma)} &= \frac{1}{k^\gamma} - \frac{2\gamma(\gamma-1)}{k^\gamma n} + \mathrm{o}\left(\frac{1}{n}\right), \ \textit{for } m = {}^n\!/_k.
\end{aligned}$$

2.2.4 Gamma Function

The Gamma (Γ) function – which is an extension of the factorial function over the reals – and its relatives arise in our closed form of the ACR. For $x > 0$, it is defined as $\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}\,dt$. In particular, $\Gamma(n+1) = n!$ for every $n \in \mathbb{N}$. Like the factorial function, the Gamma function also satisfies the following recurrence $\Gamma(x+1) = x\Gamma(x)$. For a more extensive treatment along with many folklore results about the function, see [123].

2.3 SINGLE-THRESHOLD ALGORITHMS FOR THE I.I.D. MIN-PROPHET INEQUALITY

In this section, we investigate single-threshold algorithms for the I.I.D. Min-Prophet Inequality. We start by showing that the intuition from the I.I.D. Max-Prophet Inequality on how to select a single threshold fails in the I.I.D. Min-Prophet Inequality. Afterwards, for any distribution $F \in D_\gamma$, we design a single-threshold that achieves a poly-logarithmic competitive ratio, where the exponent is exactly $-\gamma$, and show that this dependence is optimal up to constants.

A natural first approach that is seemingly intuitive is to set a single threshold $T = c \cdot \mu_{1:n}$ for some $c \geq 1$ since, if n is large enough, with good probability there will be a realization below the threshold and one would achieve a good competitive ratio in this manner. The following example shows why this natural intuition fails.

Example 2.1. Consider the exponential distribution, for which $F(x) = 1 - e^{-x}$, $f(x) = e^{-x}$, $H(x) = x$, $E[X] = 1$ and

$$\mu_{1:n} = \int_0^\infty e^{-nx}\, dx = \frac{1}{n}. \tag{2.43}$$

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

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

$$\begin{aligned}
\mathbb{E}[\mathrm{ALG}_n] &= \left(1 - (1 - F(T))^{n-1}\right) \mathbb{E}\left[X \mid X \leq T\right] + (1 - F(T))^{n-1}\, E[X] && (2.44)\\
&= \left(1 - e^{-(n-1)T}\right) \mathbb{E}\left[X \mid X \leq T\right] + e^{-(n-1)T} \cdot 1 && (2.45)\\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \mathbb{E}\left[X \mid X \leq {}^{c}\!/_{n}\right] + e^{-c\frac{n-1}{n}} && (2.46)\\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{\int_0^{\frac{c}{n}} x f(x)\, dx}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} && (2.47)\\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{\int_0^{\frac{c}{n}} x e^{-x}\, dx}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} && (2.48)\\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \frac{1 - e^{-c/n} - \frac{c}{n}\, e^{-c/n}}{1 - e^{-c/n}} + e^{-c\frac{n-1}{n}} && (2.49)\\
&= \left(1 - e^{-c\frac{n-1}{n}}\right) \left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) + e^{-c\frac{n-1}{n}}. && (2.50)
\end{aligned}$$

Thus, the competitive ratio is

$$\lambda_m(n) = \frac{\mathbb{E}[\mathrm{ALG}_n]}{\mu_{1:n}} = n\left(\left(1 - e^{-c\frac{n-1}{n}}\right)\left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) + e^{-c\frac{n-1}{n}}\right). \tag{2.51}$$

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

$$\lim_{n \to +\infty} n\left(1 - e^{-c\frac{n-1}{n}}\right)\left(1 - \frac{c}{n} \cdot \frac{e^{-c/n}}{1 - e^{-c/n}}\right) = \frac{c\, e^{-c}\left(e^c - 1\right)}{2}, \tag{2.52}$$

but

$$\lim_{n \to +\infty} n\, e^{-c\frac{n-1}{n}} = +\infty, \tag{2.53}$$

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

2.3.1 Optimal Single Threshold

This section is dedicated to proving Theorem 2.2. We design an algorithm which sets a fixed threshold T and selects the first realization that is below T. If our algorithm ever reaches X_n and has not selected any value, it is forced to pick the realization of X_n regardless of its cost. Recall that $F \in D_\gamma$ for some $\gamma < 0$.

Our choice of T is

$$T = \Theta\left(\left(\frac{\log n}{n}\right)^{-\gamma}\right). \tag{2.54}$$

For clarity of presentation, we split the analysis of the upper and lower bounds on the competitive ratio into different sections. Theorem 2.2 follows by Theorems 2.6 and 2.7.

Remark 2.1. Theorem 2.2 holds for $\gamma < 0$. This is because for $\gamma = 0$, as we will see in the proof of Theorem 2.6, we have that $\mu_{1:1} \le c\mu_{1:n}$ for some constant $c > 1$ and all n. Thus, an algorithm that sets a single threshold $T \in (\mu_{1:n}, \mu_{1:1}]$ achieves a constant competitive ratio.

Upper Bound

Theorem 2.6. *For the I.I.D. Min-Prophet Inequality and any distribution $F \in D_\gamma$, there exists a single-threshold $T = T(n, \gamma, F)$ such that the algorithm that selects the first value $X_i \le T$ for $i < n$ and X_n otherwise, achieves a* $\mathrm{O}\left((\log n)^{-\gamma}\right)$*-competitive ratio, for large enough n.*

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

$$\begin{aligned}
\mathbb{E}[\mathrm{ALG}] &= \left(1 - (1 - F(T))^{n-1}\right) \mathbb{E}\left[X \mid X \le T\right] + (1 - F(T))^{n-1}\, \mathbb{E}[X] && (2.55)\\
&= \frac{1 - (1 - F(T))^{n-1}}{F(T)} \int_0^T (F(T) - F(x))\, dx + (1 - F(T))^{n-1}\, \mu_{1:1} && (2.56)\\
&= \frac{1 - (1 - F(T))^{n-1}}{F(T)} \left(TF(T) - \int_0^T F(x)\, dx\right) + (1 - F(T))^{n-1}\, \mu_{1:1}. && (2.57)
\end{aligned}$$

Let $t = F(x) \implies dx = (F^{\leftarrow}(t))'\, dt$. Thus,

$$\begin{aligned}
\mathbb{E}[\mathrm{ALG}] &= \frac{1 - (1 - F(T))^{n-1}}{F(T)} \left(TF(T) - \int_0^{F(T)} t\, (F^{\leftarrow}(t))'\, dt\right) + (1 - F(T))^{n-1}\, \mu_{1:1} && (2.58)\\
&= \frac{1 - (1 - F(T))^{n-1}}{F(T)} \left(TF(T) - TF(T) + \int_0^{F(T)} F^{\leftarrow}(t)\, dt\right) + (1 - F(T))^{n-1}\, \mu_{1:1} && (2.59)\\
&= \frac{1 - (1 - F(T))^{n-1}}{F(T)} \cdot \int_0^{F(T)} F^{\leftarrow}(t)\, dt + (1 - F(T))^{n-1}\, \mu_{1:1}. && (2.60)
\end{aligned}$$

However, notice that, since $F \in D_\gamma$, by [31] (Corollary 1.2.10), we have that $F^{\leftarrow} \in RV_{-\gamma}$. Thus, by Lemma 2.2, we have

$$\int_0^{F(T)} F^{\leftarrow}(t) \approx \frac{TF(T)}{1 - \gamma}, \tag{2.61}$$

as $T \to 0^+$, which for our choice of T corresponds to $n \to +\infty$.

Also, using Lemma 2.5 and ignoring lower order terms, we have

$$\mu_{1:1} \approx \prod_{j=2}^{n} \left(1 - \frac{\gamma}{j}\right) \mu_{1:n} = \frac{\Gamma(n + 1 - \gamma)}{(1 - \gamma)\Gamma(1 - \gamma)\Gamma(n + 1)} \mu_{1:n} = \left(\frac{n^{-\gamma}}{(1 - \gamma)\Gamma(1 - \gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right) \mu_{1:n}. \tag{2.62}$$

Therefore, for large enough n,

$$\mathbb{E}[\mathrm{ALG}] \approx \frac{1-(1-F(T))^{n-1}}{F(T)} \cdot \frac{TF(T)}{1-\gamma} + (1-F(T))^{n-1}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right)\mu_{1:n} \tag{2.63}$$

$$= \left(1-(1-F(T))^{n-1}\right)\frac{T}{1-\gamma} + (1-F(T))^{n-1}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right)\mu_{1:n}. \tag{2.64}$$

$$\tag{2.65}$$

Thus, the competitive ratio is

$$\lambda_m(n) \approx \left(1-(1-F(T))^{n-1}\right)\frac{T}{(1-\gamma)\mu_{1:n}} + (1-F(T))^{n-1}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right). \tag{2.66}$$

Next, let $T = F^{\leftarrow}\left(\frac{g(n,\gamma)}{n}\right)$, for some appropriate function g to be defined later. Then, we obtain

$$\lambda_m(n) \approx \left(1-\left(1-\frac{g(n,\gamma)}{n}\right)^{n-1}\right)\frac{F^{\leftarrow}\left(\frac{g(n,\gamma)}{n}\right)}{(1-\gamma)\mu_{1:n}} + \left(1-\frac{g(n,\gamma)}{n}\right)^{n-1}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right). \tag{2.67}$$

Using Lemmas 2.3 and 2.4, along with the fact that $\Gamma(2-\gamma) = (1-\gamma)\Gamma(1-\gamma)$, for large enough n we get

$$\lambda_m(n) \approx \left(1-\left(1-\frac{g(n,\gamma)}{n}\right)^{n-1}\right)\frac{F^{\leftarrow}\left(\frac{g(n,\gamma)}{n}\right)}{(1-\gamma)\Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n}\right)} \tag{2.68}$$

$$+ \left(1-\frac{g(n,\gamma)}{n}\right)^{n-1}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right) \tag{2.69}$$

$$\approx \left(1-e^{-g(n,\gamma)}\right)\frac{(g(n,\gamma))^{-\gamma}F^{\leftarrow}\left(\frac{1}{n}\right)}{(1-\gamma)\Gamma(1-\gamma)F^{\leftarrow}\left(\frac{1}{n}\right)} + e^{-g(n,\gamma)}\left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right) \tag{2.70}$$

$$\approx \frac{1}{\Gamma(2-\gamma)}\left(\left(1-e^{-g(n,\gamma)}\right)(g(n,\gamma))^{-\gamma} + e^{-g(n,\gamma)}\left(n^{-\gamma} + \mathrm{o}\left(n^{-\gamma}\right)\right)\right). \tag{2.71}$$

Finally, let $g(n,\gamma) = -\gamma\log\left(\frac{n}{\log n}\right)$. This implies that $e^{-g(n,\gamma)} = \left(\frac{\log n}{n}\right)^{-\gamma}$ and that $(g(n,\gamma))^{-\gamma} = (-\gamma)^{-\gamma}\left(\log\left(\frac{n}{\log n}\right)\right)^{-\gamma}$. Therefore, (2.71) becomes

$$\lambda_m(n) \approx \frac{1}{\Gamma(2-\gamma)}\left(\left(1-\left(\frac{\log n}{n}\right)^{-\gamma}\right)(-\gamma)^{-\gamma}\left(\log\left(\frac{n}{\log n}\right)\right)^{-\gamma} + \left(\frac{\log n}{n}\right)^{-\gamma}\left(n^{-\gamma} + \mathrm{o}\left(n^{-\gamma}\right)\right)\right) \tag{2.72}$$

$$\approx \frac{1}{\Gamma(2-\gamma)}\left(\left(1-\left(\frac{\log n}{n}\right)^{-\gamma}\right)(-\gamma)^{-\gamma}\left(\log\left(\frac{n}{\log n}\right)\right)^{-\gamma} + (\log n)^{-\gamma} + \mathrm{o}\left((\log n)^{-\gamma}\right)\right) \tag{2.73}$$

$$\approx \frac{\left((-\gamma)^{-\gamma}+1\right)}{\Gamma(2-\gamma)}(\log n)^{-\gamma} + \mathrm{o}\left((\log n)^{-\gamma}\right). \tag{2.74}$$

Thus, our choice of T achieves a competitive ratio of $\mathrm{O}\left((\log n)^{-\gamma}\right)$. QED

Lower Bound

Theorem 2.7. *For the I.I.D. Min-Prophet Inequality and any $\gamma \leq 0$, consider the distribution $\mathcal{D}$ for which $F(x) = 1 - e^{-x^{-1/\gamma}}$. For $\mathcal{D}$ and large enough n, no single-threshold algorithm is $\mathrm{o}\left((\log n)^{-\gamma}\right)$-competitive.*

Proof. Recall by (2.66) that for large n

$$\lambda_m(n) \approx \left(1 - (1 - F(T))^{n-1}\right) \frac{T}{(1-\gamma)\mu_{1:n}} + (1 - F(T))^{n-1} \left(\frac{n^{-\gamma}}{(1-\gamma)\Gamma(1-\gamma)} + \mathrm{o}\left(n^{-\gamma}\right)\right) \tag{2.75}$$

$$\approx \frac{1}{\Gamma(2-\gamma)} \left(\left(1 - e^{-(n-1)T^{-1/\gamma}}\right) \frac{T\Gamma(1-\gamma)}{\mu_{1:n}} + e^{-(n-1)T^{-1/\gamma}} \left(n^{-\gamma} + \mathrm{o}\left(n^{-\gamma}\right)\right)\right). \tag{2.76}$$

Next, notice that

$$\mu_{1:n} = \int_0^\infty e^{-nx^{-1/\gamma}}\, dx = \frac{\Gamma(1-\gamma)}{n^{-\gamma}}. \tag{2.77}$$

Thus, ignoring lower order terms, we have that

$$\lambda_m(n) \approx \frac{n^{-\gamma}}{\Gamma(2-\gamma)} \left(\left(1 - e^{-(n-1)T^{-1/\gamma}}\right) T + e^{-(n-1)T^{-1/\gamma}}\right). \tag{2.78}$$

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

$$\left(1 - e^{-(n-1)T^{-1/\gamma}}\right) T = \mathrm{o}\left(\left(\frac{\log n}{n}\right)^{-\gamma}\right), \tag{2.79}$$

and also that

$$e^{-(n-1)T^{-1/\gamma}} = \mathrm{o}\left(\left(\frac{\log n}{n}\right)^{-\gamma}\right). \tag{2.80}$$

By (2.80) and the definition of $\mathrm{o}\left(\cdot\right)$, we have that for every $\varepsilon > 0$, there must exist a $n_0 \geq 1$ such that for all $n \geq n_0$, we have

$$e^{-(n-1)T^{-1/\gamma}} \leq \varepsilon \left(\frac{\log n}{n}\right)^{-\gamma} \iff T \geq \left(-\frac{\log\left(\varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}\right)}{n-1}\right)^{-\gamma}. \tag{2.81}$$

However, by (2.79), we have that for every $\varepsilon' > 0$, there must exist a $n_1 \geq 1$ such that for all $n \geq n_1$, we have

$$\left(1 - e^{-(n-1)T^{-1/\gamma}}\right) T \leq \varepsilon' \left(\frac{\log n}{n}\right)^{-\gamma}. \tag{2.82}$$

From (2.81), we know that

$$1 - e^{-(n-1)T^{-1/\gamma}} \geq 1 - \varepsilon \left(\frac{\log n}{n}\right)^{-\gamma} \tag{2.83}$$

and thus, by (2.82) for $\varepsilon' = \varepsilon$, it must be the case that

$$\left(1 - \varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}\right) T \leq \varepsilon \left(\frac{\log n}{n}\right)^{-\gamma} \iff T \leq \frac{\varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}}{1 - \varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}}. \tag{2.84}$$

Thus, using (2.81) and (2.84), we have

$$\left(-\frac{\log\left(\varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}\right)}{n-1}\right)^{-\gamma} \leq \frac{\varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}}{1-\varepsilon\left(\frac{\log n}{n}\right)^{-\gamma}}. \tag{2.85}$$

Let $\varepsilon = \left(\frac{-\gamma}{2}\right)^{-\gamma} > 0$ for $\gamma > 0$. Also let $B_\gamma(n) = \left(\frac{-\gamma}{2}\frac{\log n}{n}\right)^{-\gamma}$. Then, (2.85) becomes

$$-\log B_\gamma(n) \leq (n-1)\left(\frac{B_\gamma(n)}{1-B_\gamma(n)}\right)^{-1/\gamma} \iff \tag{2.86}$$

$$-\gamma\log\left(\frac{2n}{(-\gamma)\log n}\right) \leq (n-1)\frac{\frac{(-\gamma)\log n}{2n}}{\left(1-\left(\frac{(-\gamma)\log n}{2n}\right)^{-\gamma}\right)^{-1/\gamma}}. \tag{2.87}$$

Taking the Taylor series of $\frac{\frac{(-\gamma)\log n}{2n}}{\left(1-\left(\frac{(-\gamma)\log n}{2n}\right)^{-\gamma}\right)^{-1/\gamma}}$ around infinity, we have that

$$\frac{\frac{(-\gamma)\log n}{2n}}{\left(1-\left(\frac{(-\gamma)\log n}{2n}\right)^{-\gamma}\right)^{-1/\gamma}} \approx \frac{-\gamma\log n}{2n}, \tag{2.88}$$

and thus (2.87) becomes

$$-\gamma\log\left(\frac{2n}{(-\gamma)\log n}\right) \leq \frac{-\gamma}{2}\cdot\frac{n-1}{n}\log n \iff \log\left(\frac{2n}{(-\gamma)\log n}\right) \leq \frac{1}{2}\cdot\frac{n-1}{n}\log n. \tag{2.89}$$

Since

$$\lim_{n\to\infty}\frac{\log\left(\frac{2n}{(-\gamma)\log n}\right)}{\log n} = 1, \tag{2.90}$$

we have that, for large enough n, (2.89) does not hold. Therefore, there exists $\varepsilon > 0$ such that for all $n_0 \in \mathbb{N}$ and $n \geq n_0$, (2.79) and (2.80) cannot simultaneously hold, and we arrive at a contradiction. QED

2.4 A UNIFIED APPROACH TO ACR

In this section, we prove our main theorem.

Theorem 2.3
Let $F \in D_\gamma$. Then,

- *for the I.I.D. Max-Prophet Inequality, $\gamma \in \mathbb{R}$ and the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is*

$$\lambda_M = \min\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}.$$

- *for the I.I.D. Min-Prophet Inequality, $\gamma \leq 0$ and the asymptotic competitive ratio of the optimal* *threshold policy as $n \to \infty$ is*

$$\lambda_m = \max\left\{\frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, 1\right\}.$$

We split our analysis into the following cases and show that

- For I.I.D. Max-Prophet Inequality,
  - $\lambda_M = 1$, for $\gamma \geq 1$,
  - $\lambda_M = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}$, for $\gamma \in (0,1)$,
  - $\lambda_M = 1$, for $\gamma = 0$,
  - $\lambda_M = 1$, for $\gamma < 0$.
- For I.I.D. Min-Prophet Inequality,
  - $\lambda_m = 1$, for $\gamma = 0$,
  - $\lambda_m = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}$, for $\gamma < 0$,

For the I.I.D. Max-Prophet Inequality setting, the case of $\gamma \geq 1$ is ease to see as it implies that $\mathbb{E}[X] = +\infty$, and thus accepting the first random variable trivially yields a competitive ratio of 1. Recall that the case $\gamma > 0$ for I.I.D. Min-Prophet Inequality is impossible as it would imply that the left-most endpoint of the domain is $-\infty$.

Recall that $G_M(n)$ and $G_m(n)$ denote the expected value of the optimal threshold policy for the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality settings, respectively. Also, let $\lambda_M(n)$ and $\lambda_m(n)$ denote the competitive ratios for the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality settings, for every $n \geq 1$. Moreover, assume that $\mathcal{D} = [0, x^*)$, where $x^* \leq +\infty$.

We start by getting a recursive form of $G_M(n)$ and $G_m(n)$ which will be more useful, in terms of the distribution's inverse hazard rate.

Lemma 2.6. *For any $n > 1$, we have*

$$G_M(n) = G_M(n-1) + \int_{G_M(n-1)}^{x^*} (1 - F(u))\, du,$$

and

$$G_m(n) = \int_0^{G_m(n-1)} (1 - F(u))\, du.$$

Proof. The optimal threshold policy sets a threshold τ_i when observing X_i equal to $G_M(n-i)$ or $G_m(n-i)$ for the max and min settings, respectively. Therefore, we have

$$G_M(n) = (1 - F(G_M(n-1)))\,\mathbb{E}\left[X | X \geq G_M(n-1)\right] + F(G_M(n-1))G_M(n-1) \tag{2.91}$$

$$= (1 - F(G_M(n-1)))\frac{\int_{G_M(n-1)}^{x^*} uf(u)\, du}{(1 - F(G_M(n-1)))} + F(G_M(n-1))G_M(n-1) \tag{2.92}$$

$$= \mathbb{E}[X] - \int_0^{G_M(n-1)} uf(u)\, du + F(G_M(n-1))G_M(n-1) \tag{2.93}$$

$$= \int_0^{x^*} (1 - F(u))\, du - \int_0^{G_M(n-1)} u\,(F(u))'\, du + F(G_M(n-1))G_M(n-1) \tag{2.94}$$

$$= \int_0^{x^*} (1 - F(u))\, du - [uF(u)]_0^{G_M(n-1)} + \int_0^{G_M(n-1)} F(u)\, du + F(G_M(n-1))G_M(n-1) \tag{2.95}$$

$$= G_M(n-1) + \int_{G_M(n-1)}^{x^*} (1 - F(u))\, du - F(G_M(n-1))G_M(n-1) + F(G_M(n-1))G_M(n-1) \tag{2.96}$$

$$= G_M(n-1) + \int_{G_M(n-1)}^{x^*} (1 - F(u))\, du. \tag{2.97}$$

Similarly, for $G_m(n)$, we obtain

$$G_m(n) = F(G_m(n-1))\, \mathbb{E}\left[X | X \le G_m(n-1)\right] + (1 - F(G_m(n-1)))G_m(n-1) \tag{2.98}$$

$$= F(G_m(n-1)) \frac{\int_0^{G_M(n-1)} uf(u)\, du}{F(G_m(n-1))} + (1 - F(G_m(n-1)))G_m(n-1) \tag{2.99}$$

$$= \int_0^{G_M(n-1)} u\, (F(u))'\, du + G_m(n-1) - G_m(n-1)F(G_m(n-1)) \tag{2.100}$$

$$= [uF(u)]_0^{G_m(n-1)} - \int_0^{G_M(n-1)} F(u)\, du + G_m(n-1) - G_m(n-1)F(G_m(n-1)) \tag{2.101}$$

$$= G_m(n-1)F(G_m(n-1)) + \int_0^{G_m(n-1)} 1\, du - \int_0^{G_M(n-1)} F(u)\, du - G_m(n-1)F(G_m(n-1)) \tag{2.102}$$

$$= \int_0^{G_m(n-1)} (1 - F(u))\, du. \tag{2.103}$$

QED

Recall that $H(x) = -\log(1 - F(x))$ and $R(x) = \log F(x)$.

Lemma 2.7. *For any $n > 1$, we have*

$$G_M(n) = G_M(n-1)\left(1 - e^{-H(G_M(n-1))}\right) + \int_0^{\exp(-H(G_M(n-1)))} H^{\leftarrow}(-\log u)\, du,$$

and

$$G_m(n) = G_m(n-1)\left(1 - e^{R(G_m(n-1))}\right) + \int_0^{\exp(R(G_m(n-1)))} R^{\leftarrow}(\log u)\, du.$$

Proof. Substituting the definitions of H and R into Lemma 2.6, we have

$$G_M(n) = G_M(n-1) + \int_{G_M(n-1)}^{x^*} e^{-H(u)}\, du, \tag{2.104}$$

and

$$G_m(n) = \int_0^{G_m(n-1)} 1 - e^{R(u)}\, du. \tag{2.105}$$

Let $x = e^{-H(u)}, t = e^{R(u)}$, which implies that $du = \left(H^{\leftarrow}(-\log x)\right)' dx = \left(R^{\leftarrow}(\log t)\right)' dt$. Thus,

$$G_M(n) = G_M(n-1) + \int_{\exp(-H(G_M(n-1)))}^{0} x \cdot \left(H^{\leftarrow}(-\log x)\right)' dx, \tag{2.106}$$

and

$$G_m(n) = \int_0^{\exp(R(G_m(n-1)))} (1-t)\left(R^{\leftarrow}(\log t)\right)' dt. \tag{2.107}$$

Integrating by parts, we obtain

$$G_M(n) = G_M(n-1) + \left[x \cdot H^{\leftarrow}(-\log x)\right]_{\exp(-H(G_M(n-1)))}^{0} - \int_{\exp(-H(G_M(n-1)))}^{0} H^{\leftarrow}(-\log x)\, dx \tag{2.108}$$

$$= G_M(n-1) - G_M(n-1) \cdot e^{-H(G_M(n-1))} + \int_0^{\exp(-H(G_M(n-1)))} H^{\leftarrow}(-\log x)\, dx \tag{2.109}$$

$$= G_M(n-1)\left(1 - e^{-H(G_M(n-1))}\right) + \int_0^{\exp(-H(G_M(n-1)))} H^{\leftarrow}(-\log x)\, dx, \tag{2.110}$$

and

$$G_m(n) = \left[(1-t)R^{\leftarrow}(\log t)\right]_0^{\exp(R(G_m(n-1)))} + \int_0^{\exp(R(G_m(n-1)))} R^{\leftarrow}(\log t)\, dt \tag{2.111}$$

$$= G_m(n-1)\left(1 - e^{R(G_m(n-1))}\right) + \int_0^{\exp(R(G_m(n-1)))} R^{\leftarrow}(\log t)\, dt. \tag{2.112}$$

QED

Now we can combine Lemmas 2.7 and 2.2 to obtain a simplified approximation to $G_M(n)$ and $G_m(n)$ for large n.

Lemma 2.8. *For every $F \in D_\gamma$, where $\gamma < 1$ for the I.I.D. Max-Prophet Inequality, and large enough n, we have*

$$G_M(n) \approx G_M(n-1)\left(1 - e^{-H(G_M(n-1))}\left(1 - \frac{1}{1-\gamma}\right)\right), \tag{2.113}$$

and

$$G_m(n) \approx G_m(n-1)\left(1 - e^{R(G_m(n-1))}\left(1 - \frac{1}{1-\gamma}\right)\right). \tag{2.114}$$

We are finally ready to prove our main theorem, for each case of γ.

Theorem 2.8. *Let $F \in D_\gamma$. Then,*

- *for the I.I.D. Max-Prophet Inequality, if $\gamma \in (0,1)$, the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is*

$$\lambda_M = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}.$$

- *for the I.I.D. Min-Prophet Inequality, if $\gamma < 0$, the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is*

$$\lambda_M = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}.$$

Proof. Recall that $\lambda_M(n)$ and $\lambda_m(n)$ denote the competitive ratio of the I.I.D. Max-Prophet Inequality and I.I.D. Min-Prophet Inequality settings respectively.

For large enough n, we have

$$\lambda_M(n) = \frac{G_M(n)}{\mu_{n:n}} \approx \frac{G_M(n-1)}{\mu_{n:n}}\left(1 - e^{-H(G_M(n-1))}\left(1 - \frac{1}{1-\gamma}\right)\right) \tag{2.115}$$

$$= \lambda_M(n-1)\frac{\mu_{n-1:n-1}}{\mu_{n:n}}\left(1 - e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right), \tag{2.116}$$

and

$$\lambda_m(n) = \frac{G_m(n)}{\mu_{1:n}} \approx \frac{G_m(n-1)}{\mu_{1:n}}\left(1 - e^{R(G_m(n-1))}\left(1 - \frac{1}{1-\gamma}\right)\right) \tag{2.117}$$

$$= \lambda_m(n-1)\frac{\mu_{1:n-1}}{\mu_{1:n}}\left(1 - e^{R(\lambda_m(n-1)\mu_{1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right), \tag{2.118}$$

by Lemma 2.8. Using Lemma 2.5, we obtain

$$\lambda_M(n) = \lambda_M(n-1)\left(1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right)\right)\left(1 - e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right), \tag{2.119}$$

and

$$\lambda_m(n) = \lambda_m(n-1)\left(1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right)\right)\left(1 - e^{R(\lambda_m(n-1)\mu_{1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right). \tag{2.120}$$

Next, by [8] and Theorem 2.2, we have that

$$\lambda_M(n) = \mathrm{O}\left(1\right) \text{ and } \lambda_M(n) = \mathrm{O}\left((\log n)^{-\gamma}\right), \tag{2.121}$$

and thus

$$\lambda_M(n) - \lambda_M(n-1) \approx 0 \text{ and } \lambda_m(n) - \lambda_m(n-1) \approx 0. \tag{2.122}$$

To see why the asymptotic equality for λ_m is true, notice that there exists a constant $c > 0$ such that

$$(\log(n+1))^{-\gamma} - (\log(n))^{-\gamma} \leq (\log(n+1))^{\lceil -\gamma \rceil} - (\log(n))^{\lceil -\gamma \rceil} \tag{2.123}$$

$$\leq (\log(n+1) - \log(n)) \sum_{j=1}^{\lceil -\gamma \rceil} (\log(n+1))^{\lceil -\gamma \rceil - j} (\log(n))^{j-1} \tag{2.124}$$

$$\leq \log\left(\frac{n+1}{n}\right) \cdot c\lceil -\gamma \rceil (\log n)^{\lceil -\gamma \rceil} \tag{2.125}$$

$$\leq \log\left(1 + \frac{1}{n}\right) \cdot c\lceil -\gamma \rceil (\log n)^{\lceil -\gamma \rceil} \tag{2.126}$$

$$\leq \left(\frac{1}{n} + \mathrm{o}\left(\frac{1}{n}\right)\right) \cdot c\lceil -\gamma \rceil (\log n)^{\lceil -\gamma \rceil} \tag{2.127}$$

$$\approx 0, \tag{2.128}$$

where the first inequality follows from the monotonicity of $\log(\cdot)$ and the last inequality follows from the series expansion of $\log(1+z)$ around $z = 0$. Therefore, by (2.119), (2.120), (2.121) and (2.122), we have

$$\lambda_M(n-1)\left(1 - \left(1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right)\right)\left(1 - e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right)\right) \approx 0, \tag{2.129}$$

and

$$\lambda_m(n-1)\left(1 - \left(1 - \frac{\gamma}{n} + \mathrm{o}\left(\frac{1}{n}\right)\right)\left(1 - e^{R(\lambda_m(n-1)\mu_{1:n-1})}\left(1 - \frac{1}{1-\gamma}\right)\right)\right) \approx 0. \tag{2.130}$$

Since $\lambda_M, \lambda_m \neq 0$,

$$\left(1-\frac{\gamma}{n}+\mathrm{o}\left(\frac{1}{n}\right)\right)\left(1-e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\right)\approx 1, \tag{2.131}$$

and

$$\left(1-\frac{\gamma}{n}+\mathrm{o}\left(\frac{1}{n}\right)\right)\left(1-e^{R(\lambda_m(n-1)\mu_{1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\right)\approx 1, \tag{2.132}$$

and by rearranging terms and ignoring lower-order terms, we get

$$e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\approx\frac{-\gamma}{n}, \tag{2.133}$$

and

$$e^{R(\lambda_m(n-1)\mu_{1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\approx\frac{-\gamma}{n}. \tag{2.134}$$

Therefore,

$$e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\approx\frac{1-\gamma}{n}\iff H(\lambda_M(n-1)\mu_{n-1:n-1})\approx-\log\left(\frac{1-\gamma}{n}\right), \tag{2.135}$$

and

$$e^{R(\lambda_m(n-1)\mu_{1:n-1})}\approx\frac{1-\gamma}{n}\iff R(\lambda_m(n-1)\mu_{1:n-1})\approx\log\left(\frac{1-\gamma}{n}\right). \tag{2.136}$$

Taking the inverses of H and R, we obtain

$$\lambda_M(n-1)\approx\frac{H^{\leftarrow\left(-\log\left(\frac{1-\gamma}{n}\right)\right)}}{\mu_{n-1:n-1}}, \tag{2.137}$$

and

$$\lambda_m(n-1)\approx\frac{R^{\leftarrow\left(\log\left(\frac{1-\gamma}{n}\right)\right)}}{\mu_{1:n-1}}. \tag{2.138}$$

Since $H^{\leftarrow}(-\log x)=F^{\leftarrow}(1-x)$ and $R^{\leftarrow}(\log x)=F^{\leftarrow}(x)$, we obtain

$$\lambda_M(n-1)\approx\frac{F^{\leftarrow}\left(1-\frac{1-\gamma}{n}\right)}{\mu_{n-1:n-1}}, \tag{2.139}$$

and

$$\lambda_m(n-1)\approx\frac{F^{\leftarrow}\left(\frac{1-\gamma}{n}\right)}{\mu_{1:n-1}}. \tag{2.140}$$

By Lemma 2.3, we get

$$\lambda_M(n-1)\approx\frac{F^{\leftarrow}\left(1-\frac{1-\gamma}{n}\right)}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}F^{\leftarrow}\left(1-\frac{1}{n}\right)}, \tag{2.141}$$

and

$$\lambda_m(n-1)\approx\frac{F^{\leftarrow}\left(\frac{1-\gamma}{n}\right)}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}F^{\leftarrow}\left(\frac{1}{n}\right)}. \tag{2.142}$$

Finally, using Lemma 2.4, we get

$$\lambda_M(n-1) \approx \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}} \cdot \frac{F^{\leftarrow}\left(1-\frac{1}{n}\right)}{F^{\leftarrow}\left(1-\frac{1}{n}\right)} = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}}, \tag{2.143}$$

and

$$\lambda_m(n-1) \approx \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}} \cdot \frac{F^{\leftarrow}\left(\frac{1}{n}\right)}{F^{\leftarrow}\left(\frac{1}{n}\right)} = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}}. \tag{2.144}$$

Since $\lambda_M(n-1) \to \lambda_M$ and $\lambda_m(n-1) \to \lambda_m$ as n goes to infinity, we have

$$\lambda_M = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, \text{ for } \gamma \in (0,1) \qquad \text{and} \qquad \lambda_m = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}, \text{ for } \gamma < 0. \tag{2.145}$$

QED

Theorem 2.9. *Let $F \in D_0$. Then, for both the I.I.D. Max-Prophet Inequality and the I.I.D. Min-Prophet Inequality, the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is* 1.

Proof. In this case, we will also need the following variant of Karamata's theorem.

Lemma 2.9. *For $F \in D_0$ and large enough n, we have*

$$\int_0^{\exp(-H(G_M(n-1)))} H^{\leftarrow}(-\log u)\, du \approx F^{\leftarrow}\left(1-\frac{1}{n}\right) e^{-H(G_M(n-1))},$$

and

$$\int_0^{\exp(R(G_m(n-1)))} R^{\leftarrow}(\log u)\, du \approx F^{\leftarrow}\left(\frac{1}{n}\right) e^{R(G_m(n-1))},$$

Proof. The proof follows directly from [31] (Corollary 1.2.15), for $t = \exp\{1 + H(G_M(n-1))\}$ and $t = \exp\{1 - R(G_m(n-1))\}$. QED

By Lemmas 2.9 and 2.2, we have that, for large n

$$G_M(n-1) \approx F^{\leftarrow}\left(1-\frac{1}{n}\right) \qquad \text{and} \qquad G_m(n-1) \approx F^{\leftarrow}\left(\frac{1}{n}\right) \tag{2.146}$$

Therefore, by Lemmas 2.8 and 2.3, we have

$$\lambda_M(n) = \frac{G_M(n)}{\mu_{n:n}} \approx \frac{F^{\leftarrow}\left(1-\frac{1}{n}\right)}{F^{\leftarrow}\left(1-\frac{e^{-\gamma^*}}{n}\right)}, \tag{2.147}$$

and

$$\lambda_m(n) = \frac{G_m(n)}{\mu_{1:n}} \approx \frac{F^{\leftarrow}\left(\frac{1}{n}\right)}{F^{\leftarrow}\left(\frac{e^{-\gamma^*}}{n}\right)}, \tag{2.148}$$

both of which go to 1 as n goes to infinity, by Lemma 2.1. QED

Theorem 2.10. *Let $F \in D_\gamma$, for $\gamma < 0$. Then, for the I.I.D. Max-Prophet Inequality, the asymptotic competitive ratio of the optimal threshold policy as $n \to \infty$ is* 1.

Proof. In this case, we necessarily have $x^* < +\infty$. The analysis follows in a similar manner to the $\gamma \in (0,1)$ case, but for the value $x^* - G_M(n)$ instead of $G_M(n)$.

Notice that, by Lemma 2.6, we have

$$G_M(n) = G_M(n-1) + \int_{G_M(n-1)}^{x^*} (1 - F(u)\, du) \tag{2.149}$$

$$= G_M(n-1) + x^* - G_M(n-1) - \int_{G_M(n-1)}^{x^*} (F(u)\, du) \iff \tag{2.150}$$

$$x^* - G_M(n) = \int_{G_M(n-1)}^{x^*} F(u)\, du. \tag{2.151}$$

Next, we analyze the integral on the right-hand side.

Lemma 2.10. *For $\gamma < 0$ and large n, we have*

$$\int_{G_M(n-1)}^{x^*} F(u)\, du \approx \left(1 - e^{-H(G_M(n-1))}\left(1 - \frac{1}{1-\gamma}\right)\right)(x^* - G_M(n-1))\,.$$

Proof. We have

$$\int_{G_M(n-1)}^{x^*} F(u)\, du = \int_{G_M(n-1)}^{x^*} 1 - e^{-H(u)}\, du. \tag{2.152}$$

Let $y = e^{-H(x)}$, which implies that $du = (H^{\leftarrow}(-\log y))'\, dy$. Then,

$$\int_{G_M(n-1)}^{x^*} 1 - e^{-H(u)}\, du = \int_{e^{-H(G_M(n-1))}}^{0} (1-y)\,(H^{\leftarrow}(-\log y))'\, dy \tag{2.153}$$

$$= [(1-y)H^{\leftarrow}(-\log y)]^0_{e^{-H(G_M(n-1))}} + \int_{e^{-H(G_M(n-1))}}^{0} H^{\leftarrow}(-\log y)\, dy \tag{2.154}$$

$$= x^* - G_M(n-1)\left(1 - e^{-H(G_M(n-1))}\right) - \int_0^{e^{-H(G_M(n-1))}} H^{\leftarrow}(-\log y)\, dy \tag{2.155}$$

$$= x^* - G_M(n-1)\left(1 - e^{-H(G_M(n-1))}\right) + \int_0^{e^{-H(G_M(n-1)}} x^* - x^*\, dy - \tag{2.156}$$

$$- \int_0^{e^{-H(G_M(n-1))}} H^{\leftarrow}(-\log y)\, dy \tag{2.157}$$

$$= \left(1 - e^{-H(G_M(n-1))}\right)(x^* - G_M(n-1)) + \int_0^{e^{-H(G_M(n-1))}} (x^* - H^{\leftarrow}(-\log y))\, dy. \tag{2.158}$$

Next, let $z = 1/y$, which implies $dy = -z^{-2}\, dz$, and thus

$$\int_{G_M(n-1)}^{x^*} 1 - e^{-H(u)}\, du = \left(1 - e^{-H(G_M(n-1))}\right)(x^* - G_M(n-1)) - \tag{2.159}$$

$$- \int_{e^{H(G_M(n-1))}}^{\infty} z^{-2}\,(x^* - H^{\leftarrow}(\log z))\, dz. \tag{2.160}$$

Finally, notice that, for $\gamma < 0$, $g(z) = x^* - H^{\leftarrow}(\log z) \in RV_\gamma$ by [31] (Corollary 1.2.10) and also that $1+\gamma < 1$

which, by the general form of Karamata's theorem ([122], Theorem 1.5.11) for $\sigma=-2$, implies that

$$\int_{e^{H(G_M(n-1))}}^{\infty} z^{-2}\left(x^*-H^{\leftarrow}(\log z)\right)dz \approx \frac{e^{-H(G_M(n-1)}\left(x^*-H^{\leftarrow}\left(\log e^{H(G_M(n-1))}\right)\right)}{1-\gamma} \tag{2.161}$$

$$= \frac{e^{-H(G_M(n-1)}}{1-\gamma}\left(x^*-G_M(n-1)\right). \tag{2.162}$$

Therefore,

$$\int_{G_M(n-1)}^{x^*} F(u)\,du = \left(1-e^{-H(G_M(n-1)}\left(1-\frac{1}{1-\gamma}\right)\right)\left(x^*-G_M(n-1)\right). \tag{2.163}$$

QED

Now, let $\lambda'_M(n)=\frac{x^*-G_M(n)}{x^*-\mu_{n:n}}$. We have

$$\lambda'_M(n) = \frac{x^*-G_M(n)}{x^*-\mu_{n:n}} \tag{2.164}$$

$$\approx \frac{x^*-G_M(n-1)}{x^*-\mu_{n-1:n-1}}\frac{x^*-\mu_{n-1:n-1}}{x^*-\mu_{n:n}}\left(1-e^{-H(G_M(n-1))}\left(1-\frac{1}{1-\gamma}\right)\right) \tag{2.165}$$

$$\approx \lambda'_M(n-1)\left(1-\frac{\gamma}{n}+\mathrm{o}\left(\frac{1}{n}\right)\right)\left(1-e^{-H(G_M(n-1))}\left(1-\frac{1}{1-\gamma}\right)\right) \tag{2.166}$$

$$\approx \lambda'_M(n-1)\left(1-\frac{\gamma}{n}+\mathrm{o}\left(\frac{1}{n}\right)\right)\left(1-e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\right), \tag{2.167}$$

where the second equality follows from Lemma 2.10 and the third equality from Lemma 2.5. Next, notice that, by [5], $G_M(n)\geq 0.745\mu_{n:n}$, and thus

$$\lambda'_M(n)-\lambda'_M(n-1) \leq \frac{x^*-0.745\mu_{n:n}}{x^*-\mu_{n:n}}-1 = \frac{0.255\mu_{n:n}}{x^*-\mu_{n:n}} \leq \frac{0.255x^*}{x^*-\mu_{1:1}}\prod_{j=2}^{n}\left(1-\frac{\gamma}{j}\right)\approx 0, \tag{2.168}$$

where the second inequality follows from Lemma 2.5 and the fact that $\mu_{n:n}\leq x^*$, and the asymptotic equality at the end follows from the fact that $\lim_{n\to\infty}\prod_{j=2}^{n}\left(1-\frac{\gamma}{j}\right)=0$, for any $\gamma\in(0,1)$. Thus,

$$\lambda'_M(n-1)\left(1-\left(1-\frac{\gamma}{n}+\mathrm{o}\left(\frac{1}{n}\right)\right)\left(1-e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\right)\right)\approx 0, \tag{2.169}$$

and thus

$$e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\left(1-\frac{1}{1-\gamma}\right)\approx\frac{-\gamma}{n}, \tag{2.170}$$

Therefore,

$$e^{-H(\lambda_M(n-1)\mu_{n-1:n-1})}\approx\frac{1-\gamma}{n} \iff H(\lambda_M(n-1)\mu_{n-1:n-1})\approx -\log\left(\frac{1-\gamma}{n}\right). \tag{2.171}$$

Taking the inverse of H, we obtain

$$\lambda_M(n-1)\mu_{n-1:n-1}\approx H^{\leftarrow\left(-\log\left(\frac{1-\gamma}{n}\right)\right)}. \tag{2.172}$$

Since $H^{\leftarrow}(-\log x) = F^{\leftarrow}(1-x)$, we obtain

$$G_M(n-1) \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{n}\right). \tag{2.173}$$

Therefore, by Lemma 2.3, we get

$$\frac{x^* - G_M(n)}{x^* - \mu_{n:n}} \approx \frac{x^* - F^{\leftarrow}\left(1 - \frac{1-\gamma}{n+1}\right)}{x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)}. \tag{2.174}$$

Finally, using Lemma 2.4, we get

$$\lambda'_M(n) \approx \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n}\right)^{-\gamma}} \cdot \frac{x^* - F^{\leftarrow}\left(1-\frac{1}{n}\right)}{x^* - F^{\leftarrow}\left(1-\frac{1}{n}\right)} = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)\left(1+\frac{1}{n-1}\right)^{-\gamma}}. \tag{2.175}$$

Since $\lambda'_M(n) \to \lambda'_M$ as n goes to infinity, we have

$$\lambda'_M = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}. \tag{2.176}$$

However, this implies that

$$\frac{x^*}{\mu_{n:n}} - \frac{G_M(n)}{\mu_{n:n}} \approx \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}\left(\frac{x^*}{\mu_{n:n}} - 1\right). \tag{2.177}$$

Notice that for large n, the right-hand side goes to 0, and thus it must be that

$$\lim_{n\to\infty} \frac{G_M(n)}{\mu_{n:n}} = 1. \tag{2.178}$$

Therefore, for $\gamma < 0$, we have $\lambda_M = 1$. QED

The proof of Theorem 2.3 now follows from Theorems 2.8, 2.9 and 2.10.

2.4.1 MHR Distributions

In this section, we show that for the special case of MHR distributions, $\gamma \leq 0$ for maxima and $\gamma \geq -1$ for minima. Combining this with Theorem 2.3, for the I.I.D. Max-Prophet Inequality with an MHR distribution, we recover that the asymptotic competitive ratio is 1, a result of [96] and for the I.I.D. Min-Prophet Inequality, since $\Lambda(-1) = 2$, we obtain the following uniform bound on the competitive ratio.

Lemma 2.11. *For every MHR distribution with CDF F,*

- *if $F \in D_\gamma$ for maxima, then $\gamma \leq 0$,*
- *if $F \in D_\gamma$ for minima, then $\gamma \geq -1$.*

Proof. First, let $F \in D_\gamma$ for maxima. Then, we have that $h(x)$ is monotonically non-decreasing, which implies that $H(x) = \Omega(x)$. Therefore, $H^{\leftarrow}(x) = \mathrm{O}(x)$ and in particular, $U_M(1/x) = H^{\leftarrow}(\log x) = \mathrm{O}(\log x)$.

Assume towards contradiction that $\gamma > 0$. Then, by Lemma 2.1, it must be that

$$\lim_{t\to\infty} \frac{U_M(1/tx)}{U_M(1/t)} = x^\gamma, \tag{2.179}$$

for all $x > 0$. However, we have that

$$\lim_{t\to\infty} \frac{U_M(1/tx)}{U_M(1/t)} \leq \lim_{t\to\infty} \frac{\log xt}{\log t} = 1, \tag{2.180}$$

and thus (2.179) does not hold and we arrive at a contradiction.

Next, let $F \in D_\gamma$ for minima. Then, we have that $H(x) = \Omega(x)$. Notice, however, that $e^{-H(x)} + e^{R(x)} = 1 - F(x) + F(x) = 1$, which implies that

$$R(x) = \log\left(1 - e^{-H(x)}\right), \tag{2.181}$$

and thus $R(x) = \log\left(1 - e^{-\Omega(x)}\right)$ for every MHR distribution. Therefore, we have

$$R^{\leftarrow}(x) = \mathrm{O}\left(-\log 1 - e^x\right) \implies U_m(x) = R^{\leftarrow}(\log x) = \mathrm{O}\left(-\log 1 - x\right). \tag{2.182}$$

Assume towards contradiction that $\gamma < -1 \implies -\gamma > 1$. Then, by Lemma 2.1, it must be that

$$\lim_{t\to 0^+} \frac{U_m(tx)}{U_m(t)} = x^{-\gamma}, \tag{2.183}$$

for all $x > 0$. However, we have that

$$\lim_{t\to 0^+} \frac{U_m(tx)}{U_m(t)} \leq \lim_{t\to 0^+} \frac{-\log 1 - xt}{-\log 1 - t} = x, \tag{2.184}$$

and thus (2.183) does not hold and we arrive at a contradiction. QED

Thus, as explained above, we immediately obtain the following theorem.

Theorem 2.4
In the I.I.D. Max-Prophet Inequality, the optimal threshold strategy is 1-competitive for every MHR distribution in the domain of attraction of D_γ for some γ.

In the I.I.D. Min-Prophet Inequality, the optimal threshold strategy is 2-competitive for every MHR distribution in the domain of attraction of D_γ for some γ. Furthermore, the factor of 2 is tight, since there is no $(2-\varepsilon)$-competitive algorithm for any $\varepsilon > 0$ for the exponential distribution, which is MHR.

Proof. The first statement of the theorem follows directly from Theorem 2.3 and Lemma 2.11. For the I.I.D. Min-Prophet Inequality, let $\Lambda(\gamma) = \frac{(1-\gamma)^{-\gamma}}{\Gamma(1-\gamma)}$. Since $\Lambda(-1) = 2$ and $\Lambda(\gamma) \leq \Lambda(-1)$ for all $\gamma \geq -1$, by Theorem 2.3 and Lemma 2.11 we have that the optimal threshold strategy is 2-competitive for every MHR distribution in the domain of attraction of D_γ for some γ.

Finally, let $\mathcal{D}_{\exp}$ be the exponential distribution, and notice that $F_{\exp} \in D_{-1}$ for minima. Therefore, by Theorem 2.3, we have that the asymptotic competitive ratio of the optimal threshold strategy is $\Lambda(-1) = 2$, and thus there is no $(2-\varepsilon)$-competitive algorithm for any $\varepsilon > 0$ for the exponential distribution. QED

2.5 COMPETITION COMPLEXITY

In this section, we show Theorem 2.5.

Theorem 2.5
For every $F \in D_\gamma$ and $\gamma \leq 1$, we have that

$$ACC_{M,\mathcal{D}} = ACC_{m,\mathcal{D}} = (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}.$$

Proof. First, let $\gamma > 0$. We have

$$G_M(c\,n) \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{c\,n+1}\right) \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{c\,n}\right) \approx \left(\frac{1-\gamma}{c}\right)^{-\gamma} F^{\leftarrow}\left(1 - \frac{1}{n}\right), \tag{2.185}$$

where the first asymptotic equality follows from (2.139), and the second one follows from Lemma 2.4. We also know, from Lemma 2.3 that

$$\mu_{n:n} \approx \Gamma(1-\gamma)F^{\leftarrow}\left(1 - \frac{1}{n}\right). \tag{2.186}$$

Notice that

$$\left(\frac{1-\gamma}{c}\right)^{-\gamma} \geq \Gamma(1-\gamma) \iff c \geq (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}, \tag{2.187}$$

and thus,

$$ACC_{M,\mathcal{D}} = \inf\left\{c \;\middle|\; \left(\frac{1-\gamma}{c}\right)^{-\gamma} \geq \Gamma(1-\gamma)\right\} = (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}. \tag{2.188}$$

Next, let $\gamma = 0$. We have that

$$G_M(c\,n) \approx F^{\leftarrow}\left(1 - \frac{1}{c\,n+1}\right) \approx F^{\leftarrow}\left(1 - \frac{1}{c\,n}\right), \tag{2.189}$$

and

$$G_m(c\,n) \approx F^{\leftarrow}\left(\frac{1}{c\,n+1}\right) \approx F^{\leftarrow}\left(\frac{1}{c\,n}\right), \tag{2.190}$$

where the first asymptotic equality follows from (2.146). We also know, from Lemma 2.3 that

$$\mu_{n:n} \approx F^{\leftarrow}\left(1 - \frac{e^{-\gamma^*}}{n}\right), \tag{2.191}$$

and

$$\mu_{1:n} \approx F^{\leftarrow}\left(\frac{e^{-\gamma^*}}{n}\right), \tag{2.192}$$

where $\gamma^* \approx 0.577$ is the Euler-Mascheroni constant. Notice that since F is a monotonically increasing function, we have

$$F^{\leftarrow}\left(1 - \frac{1}{c\,n}\right) \geq F^{\leftarrow}\left(1 - \frac{e^{-\gamma^*}}{n}\right) \iff 1 - \frac{1}{c\,n} \geq 1 - \frac{e^{-\gamma^*}}{n} \iff c \geq e^{\gamma^*}, \tag{2.193}$$

and

$$F^{\leftarrow}\left(\frac{1}{c\,n}\right) \leq F^{\leftarrow}\left(\frac{e^{-\gamma^*}}{n}\right) \iff \frac{1}{c\,n} \leq \frac{e^{-\gamma^*}}{n} \iff c \geq e^{\gamma^*}. \tag{2.194}$$

Thus,

$$ACC_{M,\mathcal{D}} = ACC_{m,\mathcal{D}} = e^{\gamma^*}. \tag{2.195}$$

Also notice that $\lim_{\gamma\to 0}(1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma} = e^{\gamma^*}$.

Finally, let $\gamma < 0$ and recall that, in this case, $x^* < +\infty$ for maxima (but not necessarily for minima). Again, we have

$$G_M(c\,n) \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{c\,n+1}\right) \approx F^{\leftarrow}\left(1 - \frac{1-\gamma}{c\,n}\right), \tag{2.196}$$

and

$$G_m(c\,n) \approx F^{\leftarrow}\left(\frac{1-\gamma}{c\,n+1}\right) \approx F^{\leftarrow}\left(\frac{1-\gamma}{c\,n}\right) \approx \left(\frac{1-\gamma}{c}\right)^{-\gamma} F^{\leftarrow}\left(\frac{1}{n}\right), \tag{2.197}$$

by (2.140) and (2.173). From Lemma 2.3, we know that

$$\mu_{n:n} \approx x^* - \Gamma(1-\gamma)\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right), \tag{2.198}$$

and

$$\mu_{1:n} \approx \Gamma(1-\gamma) F^{\leftarrow}\left(1 - \frac{1}{n}\right). \tag{2.199}$$

We have that

$$G_M(c\,n) \geq \mu_{n:n} \iff \tag{2.200}$$

$$x^* - G_M(c\,n) \leq x^* - \mu_{n:n} \iff \tag{2.201}$$

$$\left(\frac{1-\gamma}{c}\right)^{-\gamma}\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right) \leq \Gamma(1-\gamma)\left(x^* - F^{\leftarrow}\left(1 - \frac{1}{n}\right)\right), \tag{2.202}$$

where the third equivalence follows from Lemma 2.4. Also notice that

$$\left(\frac{1-\gamma}{c}\right)^{-\gamma} \leq \Gamma(1-\gamma) \iff c \geq (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}, \tag{2.203}$$

and thus,

$$ACC_{M,\mathcal{D}} = ACC_{m,\mathcal{D}} = \inf\left\{c \;\middle|\; \left(\frac{1-\gamma}{c}\right)^{-\gamma} \geq \Gamma(1-\gamma)\right\} = (1-\gamma)\left(\Gamma(1-\gamma)\right)^{1/\gamma}. \tag{2.204}$$

QED

Chapter 3: ORACLE-AUGMENTED PROPHET INEQUALITIES

3.1 OVERVIEW

In Chapter 2, we presented distribution-optimal prophet inequalities for the maximization and minimization settings with I.I.D. random variables. The goal of this chapter is to achieve better guarantees by augmenting the algorithm's capabilities. To this end, we allow the algorithm to query an oracle that has knowledge of all realizations, even the ones in the future, and study the classical single-item maximization prophet inequality setting under identical and non-identical distributions.

Motivation. Our oracle model is motivated by the idea of enhancing algorithms via the use of machine-learned *predictions*, in order to go beyond worst-case analysis [34]–[38]. This idea of using learning to improve the performance of algorithms has received significant attention recently, for example in designing auctions to maximize revenue [124], [125] or in matching problems [126], [127]. For more information on this line of work see the survey of Mitzenmacher and Vassilvitskii [128]. In real-world applications such as posted-pricing mechanisms for auctions, machine-learning models can capture behavioral patterns of buyers and accurately predict their future actions. This allows them to provide highly accurate predictions on future realizations in repeated prophet inequality settings, which makes studying prediction-enhanced models of prophet inequalities significantly important.

Prior work. The majority of past literature has focused on maximizing the competitive ratio which is also called the *Ratio of Expectations (RoE)*. A slightly different objective, introduced by Gilbert and Mosteller [40] for I.I.D. random variables, is that of maximizing the *Probability of selecting the Maximum realization (PbM)* realization. For this setting, Gilbert and Mosteller [40] gave an algorithm that achieves a probability of ≈ 0.58, which is the best possible. Later, Esfandiari, Hajiaghayi, Lucier and Mitzenmacher [12] studied the same objective for non-I.I.D. random variables, obtaining a tight probability equal to $1/e$ when the random variables arrive in adversarial order and 0.517 when the random variables arrive in random order. The latter case was recently improved to the tight ≈ 0.58 by Nuti [41], showing that the I.I.D. setting is not easier than the non-I.I.D. setting with random order.

A model that is quite related to ours is the *top* 1*-of-*m model, formally introduced by Assaf and Samuel-Cahn [42] for I.I.D. random variables, although it had been studied initially by Gilbert and Mosteller [40]. In this setting, the algorithm is allowed to select $m \geq 1$ values, but the value it gets judged by is the maximum selected value. Gilbert and Mosteller [40] gave numerical approximations of the PbM objective for $2 \leq m \leq 10$, using a simple single threshold algorithm. Later, Assaf and Samuel-Cahn [42] studied the RoE objective for non-identical distributions and gave a very elegant and simple $(1 - 1/m+1)$-competitive algorithm. The same authors, along with Larry Goldstein, later improved this [129] bounding the competitive ratio of the optimal algorithm by a recursive differential equation. They gave constants for $2 \leq m \leq 5$ numerically, but studying the asymptotic nature of the constants for large m turned out to be difficult. Ezra, Feldman, and Nehama [43] later revisited the problem and gave a new algorithm for large k that is $1 - \mathrm{O}\left(e^{-m/6}\right)$-competitive for the same problem. This improves the error term from [129] from linear in m to exponential in m.

Model. Our model generalizes the prophet inequality setting, allowing the algorithm some information about the future that is otherwise privy only to the prophet. Specifically, at any point in the process, upon seeing a reward X_i, the algorithm is allowed to query an oracle $\mathcal{O}$. The oracle $\mathcal{O}$ responds with a single bit answer: YES if the current realization is the largest of the remaining realizations, i.e $X_i \geq \max_{j=i+1}^{n} X_j$ and NO otherwise. In other words, the oracle $\mathcal{O}$ informs the algorithm on whether the latter should select X_i, or reject it because there is a higher reward coming in the future. Clearly, with no queries available, one recovers the classical prophet inequality setting, whereas with $n-1$ queries, the strategy of using a query on every X_i, for $i = 1, \ldots, n-1$, leads to the algorithm selecting the highest realization always. Thus, this model interpolates nicely between the two extremes of full information and no information about the future.

We use Z to denote $\max\{X_1, \ldots, X_n\}$. To help distinguish between the different settings, we denote each model as $\mathcal{M}(x, y, z)$, where

- x is either $Proph_m$ or $\mathcal{O}_m$ with $m \in \mathbb{N}$,
- y is either I.I.D. or Non-I.I.D. and
- z is either PbM or RoE.

$\mathcal{M}(\mathbf{Proph_m}, \mathbf{y}, \mathbf{z})$ This denotes the *top* 1-*of*-m setting, as described above. $X_1, \ldots, X_n$ may be identically distributed ($y =$ I.I.D.) or not ($y =$ Non-I.I.D.), and m denotes the total number of realizations that can be selected. At the end of the process, we have selected a set of realized values $S = \{X_{i_1}, X_{i_2}, \ldots, X_{i_k}\}$, where $k \leq m$. Let V denote the highest realized value in S. The objective depends on z; if $z = PbM$, then the payoff is 1 if and only if $V = Z$, otherwise it is 0. On the other hand, if $z = RoE$, we compare against the expected value of the prophet, i.e. $\mathbb{E}[Z]$, and our objective is to maximize the competitive ratio, $\frac{\mathbb{E}[V]}{\mathbb{E}[Z]}$.

$\mathcal{M}(\mathcal{O}_\mathbf{m}, \mathbf{y}, \mathbf{z})$ This denotes our oracle setting. As in the previous model, $X_1, \ldots, X_n$ may be identically distributed ($y =$ I.I.D.) or not ($y =$ Non-I.I.D.). However, there are two key differences between this setting and the previous one. On one hand, we are only allowed to select a single realization. On the other hand, we are allowed to up to m queries to a truthful oracle $\mathcal{O}$, as described above. Let V denote the value that we selected (or 0 if we reached the end of the process without selecting a value). Again, the objective depends on z in exactly the same way as in the $\mathcal{M}(Proph_m, y, z)$ model.

3.1.1 Our Contributions

In this chapter, we study the oracle model for identical and non-identical distributions with the PbM and RoE objectives and make the following contributions:

- We establish an equivalence between the oracle model and the top-1-of-m model for the PbM objective.
- We show that this equivalence fails to hold for the RoE objective, but guarantees for the competitive ratio in the oracle model translate to guarantees in the top-1-of-m model.
- We give a single-threshold algorithm for $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ that achieves a $1 - \mathrm{O}\left(m^{-m/5}\right)$ probability of selecting the maximum, as well as providing an upper bound that is asymptotically (almost) tight. To the best of our knowledge, this is the first result for the PbM objective and general m in the top-1-of-m model. Our algorithm achieves a probability of ≈ 0.797 even with $m = 1$ calls to the oracle, a significant improvement on the ≈ 0.58 achieved without oracle calls [40].

- We give a single-threshold algorithm for $\mathcal{M}(\mathcal{O}_m, Non-I.I.D., RoE)$ that improves upon the current best-known of [43] even for the weaker oracle model, as well as providing an upper bound that is asymptotically (almost) tight. We conjecture that our upper bound construction is the worst-case instance, as it generalizes standard counterexamples in prophet inequalities.

The main motivation behind our oracle model comes from our first two results which relate it to the top-1-of-m model.

Theorem 3.1. *The $\mathcal{M}(\mathcal{O}_m, y, PbM)$ model is equivalent to the $\mathcal{M}(Proph_{m+1}, y, PbM)$ model, where $y =$ I.I.D. or Non-I.I.D. In other words, for every prophet inequality instance, the probability achieved by the best-possible algorithm in the $\mathcal{M}(\mathcal{O}_m, y, PbM)$ model is the same as the one achieved by the best-possible algorithm in the $\mathcal{M}(Proph_{m+1}, y, PbM)$ model.*

Theorem 3.1 is perhaps not that surprising due to the apparent similarity of the two models. However, thinking about the top-1-of-m setting from the viewpoint of oracle calls allows for a different perspective that we exploit in our analysis. Perhaps more surprisingly, our oracle model and the top-1-of-m model stop being equivalent when one considers the RoE objective; as we show in our second result, the oracle model is strictly weaker.

Theorem 3.2. *There exists a prophet inequality instance and an algorithm $\mathcal{A}$ for $\mathcal{M}(Proph_2, y, PbM)$ on that instance for which no algorithm for $\mathcal{M}(\mathcal{O}_1, Non-I.I.D., RoE)$ can achieve the same competitive ratio as that of $\mathcal{A}$.*

However, for every instance of $\mathcal{M}(\mathcal{O}_m, y, RoE)$ where $y =$ I.I.D. or Non-I.I.D., there exists an algorithm $\mathcal{A}$ for the same instance of $\mathcal{M}(Proph_{m+1}, y, RoE)$ that achieves a competitive ratio that is at least as good as that of the optimal algorithm for $\mathcal{M}(\mathcal{O}_m, y, RoE)$.

After establishing the relationship between our oracle model and the top-1-of-m model, we turn our attention to upper and lower bounds for the oracle model. First, for the I.I.D. setting with m oracles calls and the PbM objective, we present a simple, single-threshold algorithm that selects the maximum realization with probability that approaches 1 in a super-exponential fashion. As a warm-up, we first present the analysis for $m = 1$ before generalizing it to all m.

Theorem 3.3. *For every instance of $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ and sufficiently large n, there exists an algorithm that selects the maximum realization with probability at least $1 - \mathrm{O}\left(m^{-m/5}\right)$.*

We also present an upper bound on the probability of success that is asymptotically tight, up to small multiplicative constants in the exponent. Because of Theorem 3.1, both upper and lower bounds on the probability of success carry over in the top 1-of-m setting as well.

Theorem 3.4. *There exists an instance of $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ for which no algorithm can select the maximum realization with probability greater than $1 - \mathrm{O}\left(m^{-m}\right)$.*

Next, we turn our attention to the Non-I.I.D. setting and the RoE objective. We first present an extremely simple single-threshold algorithm achieving a competitive ratio that approaches 1 exponentially in m. Even though our algorithm is for the oracle model, for which weaker guarantees are expected due to Theorem 3.2, it improves upon the best-known guarantee for the top-1-of-m setting, due to [43].

Theorem 3.5. *For every instance of $\mathcal{M}(\mathcal{O}_m, Non-I.I.D., RoE)$ and sufficiently large n, there exists an algorithm that achieves a competitive ratio of $1 - \mathrm{O}\left(e^{-m/5.178}\right)$.*

Despite the slight improvement in the exponent that our algorithm guarantees, our main result in the Non-I.I.D. setting is the construction of an instance that yields an almost matching upper bound on the competitive ratio. Specifically, we show that no algorithm can obtain a $1 - \frac{1}{2^{m+1}}$ competitive ratio for the Non-I.I.D. setting with m oracle calls. The same bound carries over for the PbM objective as well. Our instance generalizes standard counterexamples of prophet inequalities and for this reason we believe it to be the worst-case instance for the oracle model.

Theorem 3.6. *There exists an instance of the $\mathcal{M}(\mathcal{O}_m, Non-I.I.D., z)$ setting, where $z =$ RoE or PbM, in which no algorithm can achieve a $\left(1 - \frac{1}{2^{m+1}} + \delta\right)$-competitive ratio or select the maximum realization with probability $\left(1 - \frac{1}{2^{m+1}} + \delta\right)$, for any $\delta > 0$.*

3.1.2 Related Work

We have already mentioned the related work on algorithms with predictions, as well as the works of Gilbert and Mosteller [40], Esfandiari, Hajiaghayi, Lucier and Mitzenmacher [12] and Nuti [41] for the PbM objective. Related work includes the study of order-aware algorithms by Ezra, Feldman, Gravin and Tang [130], algorithms with fairness guarantees by Correa, Cristi, Dütting and Norouzi-Fard [131] and algorithms with a-priori information of some of the values by Correa, Cristi, Epstein and Soto [132]. In addition to these, Esfandiari, Hajiaghayi, Lucier and Mitzenmacher [12] study a related but distinct variant to ours. They relax the objective to allow the return of one out of the top k values, and show exponential upper and lower bounds. Their model is orthogonal to ours, and thus incomparable.

For more information on other related work, see §1.4.

Organization Section 3.2 contains some definitions and useful lemmas that we use in our analysis. In Section 3.3 we relate our model to top-1-of-m-selection model of Assaf and Samuel-Cahn [42]. Section 3.4 contains our (almost) asymptotically tight algorithms for the I.I.D. setting. In Section 3.5 we present our (almost) asymptotically tight algorithm for the non-I.I.D. setting.

3.2 PRELIMINARIES

We briefly present two lemmas that will be useful in the analysis of our algorithms; the standard Chernoff bound for binary random variables and Le Cam's theorem.

Lemma 3.1 ([133])**.** *Let $Y_1, \ldots, Y_n$ be independent indicator random variables with $p_i = \Pr Y_i = 1$ and $Y = \sum_i Y_i$. Let $\mu = \mathbb{E}[Y] = \sum_i p_i$. Then,*

1. *For $\delta \geq 0$,*
$$\Pr\left[Y \geq (1+\delta)\mu\right] \leq \left(\frac{e^{\delta}}{(1+\delta)^{(1+\delta)}}\right)^{\mu}.$$

2. *For $\delta \geq 0$,*
$$\Pr\left[Y \leq (1-\delta)\mu\right] \leq \left(\frac{e^{-\delta}}{(1-\delta)^{(1-\delta)}}\right)^{\mu}.$$

3. *For $\delta \in (0, 1]$,*
$$\Pr\left[Y \geq (1+\delta)\mu\right] \leq e^{-\mu\delta^2/3}.$$

4. *For* $\delta \in (0, 1]$
$$\Pr[Y \leq (1-\delta)\mu] \leq e^{-\mu\delta^2/2}.$$

5. *For* $\delta > e^2$,
$$\Pr[Y \geq (1+\delta)\mu] < e^{-\frac{\mu\delta\log\delta}{2}}.$$

Le Cam's theorem is useful in bounding the approximation error of a binomial distribution by a Poisson distribution. We will use a slightly tighter version [134].

Lemma 3.2 ([134], [135]). *For every* $n \in \mathbb{N}, p \in (0, 1)$, *we have*

$$\sum_{i=0}^{\infty} \left| \binom{n}{i} p^i (1-p)^{n-i} - e^{-np} \frac{(np)^i}{i!} \right| \leq \frac{2np^2}{\max\{1, np\}}.$$

3.3 REDUCTIONS

To motivate our oracle model, we start by establishing an equivalence between $\mathcal{M}(\mathcal{O}_m, y, PbM)$ and $\mathcal{M}(Prophet_{m+1}, y, PbM)$, for both the $y =$ I.I.D. and $y =$ Non-I.I.D. case (Theorem 3.1). We also show that, perhaps surprisingly, this equivalence does not hold for the RoE objective, but guarantees for $\mathcal{M}(\mathcal{O}_m, y, RoE)$ translate to guarantees for $\mathcal{M}(Prophet_{m+1}, y, RoE)$ (Theorem 3.2). Later, we will use this result to improve the best-known guarantees on $\mathcal{M}(Prophet_{m+1}, y, RoE)$.

3.3.1 The PbM objective

Theorem 3.1
The $\mathcal{M}(\mathcal{O}_m, y, PbM)$ model is equivalent to the $\mathcal{M}(Proph_{m+1}, y, PbM)$ model, where $y =$ I.I.D. or Non-I.I.D. In other words, for every prophet inequality instance, the probability achieved by the best-possible algorithm in the $\mathcal{M}(\mathcal{O}_m, y, PbM)$ model is the same as the one achieved by the best-possible algorithm in the $\mathcal{M}(Proph_{m+1}, y, PbM)$ model.

Theorem 3.1 follows from Lemmas 3.3 and 3.4.

Lemma 3.3. *Fix an instance of $\mathcal{M}(Proph_{m+1}, y, PbM)$ where $y =$ I.I.D. or Non-I.I.D., and let α denote the probability of selecting the maximum that an algorithm $\mathcal{A}$ for $\mathcal{M}(\mathcal{O}_m, y, PbM)$ achieves on this instance. Then, there exists an algorithm $\mathcal{B}$ for $\mathcal{M}(Proph_{m+1}, y, PbM)$ on this instance, with black-box access to $\mathcal{A}$ such that the probability that $\mathcal{B}$ selects the maximum realization is at least α.*

Proof. Again, the idea is that $\mathcal{B}$ can simulate $\mathcal{A}$'s behaviour by selecting each realization that $\mathcal{A}$ decides to query. Initially, $\mathcal{B}$ starts with an empty set S of selected values. Whenever $\mathcal{B}$ is presented with a realization X_i, it feeds it to $\mathcal{A}$. If $\mathcal{A}$ decides to select X_i or expend a query for X_i, regardless of the outcome of the query, $\mathcal{B}$ always selects X_i into S, otherwise $\mathcal{B}$ decides not to select X_i. By induction, S contains exactly all the realizations that were queried by $\mathcal{A}$ as well as at most one more realization that might have been selected by $\mathcal{A}$ if it run out of queries. Therefore, $|S| \leq m + 1$.

Now, notice that $\mathcal{A}$ succeeds if and only if it selects the maximum, and it only selects a realization X_i if (i) it chose to expend a query on X_i, or (ii) when it observed X_i it run out of queries. In both cases, by the description of $\mathcal{B}$, we know that $X_i \in S$, and thus the probability that $\mathcal{B}$ succeeds is at least α. QED

Next, we show that $\mathcal{M}(\mathcal{O}_m, y, PbM)$ can be reduced to $\mathcal{M}(Proph_{m+1}, y, PbM)$.

Lemma 3.4. *Fix an instance of $\mathcal{M}(\mathcal{O}_m, y, PbM)$ where y = I.I.D. or Non-I.I.D., and let α denote the probability of selecting the maximum that an algorithm $\mathcal{B}$ for $\mathcal{M}(Proph_{m+1}, y, PbM)$ achieves on this instance. Then, there exists an algorithm $\mathcal{A}$ for $\mathcal{M}(\mathcal{O}_m, y, PbM)$ on this instance, with black-box access to $\mathcal{B}$ such that the probability that $\mathcal{A}$ selects the maximum realization is at least α.*

Proof. The idea is that $\mathcal{A}$ can simulate $\mathcal{B}$'s behaviour using the oracle queries instead of storing the values like $\mathcal{B}$ does. Initially, $\mathcal{B}$ starts with an empty set S of selected values. Whenever $\mathcal{A}$ is presented with a realization X_i, it feeds it to $\mathcal{B}$. If $\mathcal{B}$ selects X_i into S, $\mathcal{A}$ chooses to expend a query and ask $\mathcal{O}$ whether $X_i \geq \max_{j=i+1}^{n} X_j$. Consider the first i where this happens. We distinguish between the two possible answers:

- If $\mathcal{O}$ answers YES, then we know that all future realizations are smaller than X_i. However, we also know that since the objective is PbM, any optimal algorithm for $Proph_{m+1}$ will only select a value X_i if it is larger than any previously observed value (otherwise it "wastes" a spot in S for a value that is definitely not the maximum). Therefore, if $\mathcal{B}$ selects X_i, we know that $X_i \geq \max_{j<i} X_j$. In this case, both $\mathcal{B}$ and $\mathcal{A}$ succeed in selecting the maximum realization.

- If $\mathcal{O}$ answers NO, then we know that there exists a future realization that is greater than X_i. In this case, the instance for $\mathcal{B}$ reduces to $\mathcal{M}(Proph_m, y, PbM)$ on $X_{i+1}, \ldots, X_n$, whereas the instance for $\mathcal{A}$ reduces to $\mathcal{M}(\mathcal{O}_{m-1}, y, PbM)$. Since we know that $\mathcal{M}(Proph_1, y, PbM) = \mathcal{M}(\mathcal{O}_0, y, PbM)$ by definition, we have that by induction, the probability that $\mathcal{A}$ succeeds is at least α.

QED

3.3.2 The RoE Objective

Given the apparent similarity of the two models, one may wonder whether the equivalence continues to hold even for the RoE objective. As we show in this section, this is not the case, but studying the oracle model for the RoE objective is still useful.

Theorem 3.2
There exists a prophet inequality instance and an algorithm $\mathcal{A}$ for $\mathcal{M}(Proph_2, y, PbM)$ on that instance for which no algorithm for $\mathcal{M}(\mathcal{O}_1, Non-I.I.D., RoE)$ can achieve the same competitive ratio as that of $\mathcal{A}$.

However, for every instance of $\mathcal{M}(\mathcal{O}_m, y, RoE)$ where y = I.I.D. or Non-I.I.D., there exists an algorithm $\mathcal{A}$ for the same instance of $\mathcal{M}(Proph_{m+1}, y, RoE)$ that achieves a competitive ratio that is at least as good as that of the optimal algorithm for $\mathcal{M}(\mathcal{O}_m, y, RoE)$.

We first present an example that shows the first part of the theorem.

Example 3.1. For a fixed $\varepsilon > 0$ and $m = 1$, consider the following instance:

$$X_1 = 1 \quad \text{w.p. } 1, X_2 = \begin{cases} 1+\varepsilon & \text{w.p. } \frac{1}{2}-\varepsilon \\ 0 & \text{w.p. } \frac{1}{2}+\varepsilon \end{cases}, X_3 = \begin{cases} \frac{1}{\varepsilon} & \text{w.p. } \varepsilon \\ 0 & \text{w.p. } 1-\varepsilon \end{cases}. \tag{3.1}$$

Let $Z = \max\{X_1, X_2, X_3\}$, with

$$\mathbb{E}[Z] = \frac{1}{\varepsilon} \cdot \varepsilon + (1+\varepsilon)(1-\varepsilon)\left(\frac{1}{2}-\varepsilon\right) + 1 \cdot (1-\varepsilon)\left(\frac{1}{2}+\varepsilon\right). \tag{3.2}$$

Notice that, for small ε, an algorithm $\mathcal{B}$ that is optimal for the $Proph_2$ model in this instance is to select X_1, ignore X_2 and then select X_3 if it is non-zero. This yields

$$\mathbb{E}[\mathcal{B}] = 1 \cdot (1-\varepsilon) + \frac{1}{\varepsilon} \cdot \varepsilon. \tag{3.3}$$

However, the optimal $\mathcal{A}$ will query $\mathcal{O}$ at X_1. With probability $(1-\varepsilon)(1/2+\varepsilon)$, it will stop and select X_1, getting a value of 1. Otherwise, it will continue, with no oracle calls left. It will ignore X_2 and select X_3. Thus,

$$\mathbb{E}[\mathcal{A}] = 1 \cdot \left(\frac{1}{2} + \varepsilon\right)(1-\varepsilon) + \frac{1}{\varepsilon} \cdot \varepsilon. \tag{3.4}$$

The competitive ratios of $\mathcal{A}$ and $\mathcal{B}$ respectively are

$$RoE_{\mathcal{A}} = \frac{\left(\frac{1}{2}+\varepsilon\right)(1-\varepsilon) + \frac{1}{\varepsilon}\cdot\varepsilon}{\frac{1}{\varepsilon}\cdot\varepsilon + (1+\varepsilon)(1-\varepsilon)\left(\frac{1}{2}-\varepsilon\right) + 1\cdot(1-\varepsilon)\left(\frac{1}{2}+\varepsilon\right)} \tag{3.5}$$

and

$$RoE_{\mathcal{B}} = \frac{(1-\varepsilon) + \frac{1}{\varepsilon}\cdot\varepsilon}{\frac{1}{\varepsilon}\cdot\varepsilon + (1+\varepsilon)(1-\varepsilon)\left(\frac{1}{2}-\varepsilon\right) + 1\cdot(1-\varepsilon)\left(\frac{1}{2}+\varepsilon\right)}, \tag{3.6}$$

and thus, as $\varepsilon \to 0$, we get

$$RoE_{\mathcal{A}} \to \frac{3/2}{2} = \frac{3}{4}, \qquad \text{and} \qquad RoE_{\mathcal{B}} \to \frac{2}{2} = 1. \tag{3.7}$$

The above example, appropriately generalized for $m > 1$ (for the exact instance see the proof of Theorem 3.6), yields the following corollary.

Corollary 3.1. *For every $m \geq 1$, there exists an instance such that*

$$\frac{\mathcal{M}(\mathcal{O}_m, Non-I.I.D., RoE)}{\mathcal{M}(Proph_{m+1}, Non-I.I.D., RoE)} \leq 1 - \frac{1}{2^{m+1}}.$$

Next, we present the proof of the second part of Theorem 3.2, showing that an algorithm for $\mathcal{M}(Prophet_{m+1}, y, RoE)$ that has access to an algorithm for $\mathcal{M}(\mathcal{O}_m, y, RoE)$ can always do at least as well. The theorem follows from Lemma 3.5, whose proof is essentially the same as the proof of Lemma 3.3.

Lemma 3.5. *Fix an instance of $\mathcal{M}(Proph_{m+1}, y, PbM)$ where $y =$ I.I.D. or Non-I.I.D., and let α denote the competitive ratio that an algorithm $\mathcal{A}$ for $\mathcal{M}(\mathcal{O}_m, y, PbM)$ achieves on this instance. Then, there exists an algorithm $\mathcal{B}$ for $\mathcal{M}(Proph_{m+1}, y, PbM)$ on this instance, with black-box access to $\mathcal{A}$, that achieves competitive ratio at least α.*

Proof. Again, the idea is that $\mathcal{B}$ can simulate $\mathcal{A}$'s behaviour by selecting each realization that $\mathcal{A}$ decides to query. Initially, $\mathcal{B}$ starts with an empty set S of selected values. Whenever $\mathcal{B}$ is presented with a realization X_i, it feeds it to $\mathcal{A}$. If $\mathcal{A}$ decides to select X_i or expend a query for X_i, regardless of the outcome of the query, $\mathcal{B}$ always selects X_i into S, otherwise $\mathcal{B}$ decides not to select X_i. By induction, S contains exactly all the realizations that were queried by $\mathcal{A}$ as well as at most one more realization that might have been selected by $\mathcal{A}$ if it run out of queries. Therefore, $|S| \leq m+1$.

Now, notice that, for every possible sequence of realizations, whatever value $\mathcal{A}$ has selected is also in S. Therefore, if $V_{\mathcal{A}}$ is the value selected by $\mathcal{A}$ and $V_{\mathcal{B}}$ is the value selected by $\mathcal{B}$, we have $\mathbb{E}[V_{\mathcal{B}}] \geq \mathbb{E}[V_{\mathcal{A}}]$, and thus the competitive ratio of $\mathcal{B}$ is at least α. QED

3.4 THE I.I.D. SETTING

After describing how the oracle model is related to the top-1-of-m model, we continue by providing algorithms for the oracle model. As a warm-up, we take a look at the I.I.D. setting with the PbM objective and the case of $m = 1$, providing a simple single-threshold algorithm as well as analyzing the optimal (multiple) threshold algorithm.

3.4.1 A Single-Threshold Algorithm for $m = 1$

Our single-threshold algorithm $\mathcal{A}_p$ for $\mathcal{M}(\mathcal{O}_1, I.I.D., PbM)$ selects a threshold τ equal to the p-th quantile of the given distribution $\mathcal{D}$, for some $p \in [0, 1]$. In other words, τ is set such that $p = \Pr[X_i \geq \tau]$. The first time the algorithm observes a realization above τ, it queries the oracle to see whether the realization should be selected or not. If it continues, it simply accepts the first value encountered above the observed realization on which it queried $\mathcal{O}$.

Lemma 3.6. *There exists $p \in [0, 1]$ such that $\mathcal{A}_p$ selects the maximum realization with probability at least 0.797 in the $\mathcal{M}(\mathcal{O}_1, I.I.D., PbM)$ model for large n.*

Proof. Let Y be the total number of realizations above τ, and $i_1 < i_2 < \cdots < i_Y$ be the indices of the random variables above τ, i.e. $X_{i_t} > \tau$, for $t = 1, \ldots, Y$. Furthermore, let r_t be the rank of X_{i_t} in $\mathcal{X} = \{X_{i_1}, \ldots, X_{i_Y}\}$, i.e. the number k such that X_{i_t} is the k-th largest number in $\mathcal{X}$, and Z be the maximum realization of $X_1, \ldots, X_n$.

X_{i_1} is the first realization we observe above τ. Notice that if $r_1 = 1$ or $r_1 = 2$ then the algorithm always selects the maximum realization Z. In other words, given that $Y = 1$ or $Y = 2$, the algorithm selects Z with probability 1. Consider the case $Y > 2$. Again, if $r_1 \leq 2$, the algorithm selects Z with probability 1. Otherwise, if $r_1 > 2$, the algorithm returns Z if and only if for all realizations above τ that appear after X_{i_1} and are also larger than X_{i_1}, the first to encounter is Z. In other words, for the algorithm to succeed in this case, it must be that among the $r_1 - 1$ values of rank smaller than r_1, the first one in the arrival order is the element of rank 1. Since the random variables are I.I.D., the probability of this event is exactly $1/r_1 - 1$.

Let j be the first index such that $X_{i_j} > X_{i_1}$, and $\alpha(Y) = \Pr\{\mathcal{A} \text{ selects } Z \mid Y\}$. Conditioned on $Y \geq 3$, the probability that the algorithm selects Z is

$$\alpha(Y \mid Y \geq 3) = \Pr[r_1 = 1] + \Pr[r_1 = 2] + \sum_{t=3}^{Y} \Pr[r_1 = t] \Pr\{r_j = 1 \mid r_1 = t\} \tag{3.8}$$

$$= \frac{2}{Y} + \sum_{t=3}^{Y} \frac{\Pr[r_z = 1 \mid r_1 = t]}{Y} \tag{3.9}$$

$$= \frac{1}{Y}\left(2 + \sum_{t=3}^{Y} \Pr[r_z = 1 \mid r_1 = t]\right) \tag{3.10}$$

$$= \frac{1}{Y}\left(2 + \sum_{t=3}^{Y} \frac{1}{t-1}\right) \tag{3.11}$$

$$= \frac{1}{Y}\left(1+\sum_{t=2}^{Y}\frac{1}{t-1}\right) \tag{3.12}$$

$$= \frac{1}{Y}\left(1+\sum_{t=1}^{Y-1}\frac{1}{t}\right) \tag{3.13}$$

$$= \frac{1}{Y}\left(1+H_{Y-1}\right), \tag{3.14}$$

where H_n denotes the n-th harmonic number. Recall also that $\alpha\left(Y \mid Y=1\right) = \alpha\left(Y \mid Y=2\right) = 1$.

Next, we estimate $\Pr[Y=i]$, by approximating Y with a Poisson distribution via Lemma 3.2 (Le Cam's theorem). Let

$$\delta_i = \left|\binom{n}{i}p^i(1-p)^{n-i} - e^{-np}\frac{(np)^i}{i!}\right|. \tag{3.15}$$

The idea is to set p such that $np = q$, where q is a fixed constant. We know that $\Pr[Y=i] = \binom{n}{i}p^i(1-p)^{n-i}$, and thus, by Lemma 3.2, we have

$$\sum_{i=0}^{\infty}\delta_i = \sum_{i=0}^{\infty}\left|\Pr[Y=i] - e^{-np}\frac{(np)^i}{i!}\right| = \sum_{i=0}^{\infty}\left|\Pr[Y=i] - e^{-q}\frac{(q)^i}{i!}\right| \le \frac{2qp}{\max\{1,q\}} \le 2p = \frac{2q}{n}. \tag{3.16}$$

Overall, the probability that $\mathcal{A}$ selects Z is

$$\alpha(Y) = \sum_{i=0}^{n}\Pr[Y=i]\cdot\alpha\left(Y \mid Y=i\right) \tag{3.17}$$

$$= \Pr[Y=1] + \sum_{i=2}^{n}\Pr[Y=i]\cdot\alpha\left(Y \mid Y=i\right) \tag{3.18}$$

$$\ge np(1-p)^{(n-1)} + \sum_{i=2}^{n}\left(e^{-q}\frac{q^i}{i!} - \delta_i\right)\cdot\alpha\left(Y \mid Y=i\right) \tag{3.19}$$

$$= q(1-q/n)^{(n-1)} + \sum_{i=2}^{n}e^{-q}\frac{q^i}{i!}\cdot\alpha\left(Y \mid Y=i\right) - \sum_{i=2}^{n}\delta_i\cdot\alpha\left(Y \mid Y=i\right) \tag{3.20}$$

$$\ge q(1-q/n)^{(n-1)} + \sum_{i=2}^{n}e^{-q}\frac{q^i}{i!}\frac{1+H_{i-1}}{i} - \sum_{i=2}^{n}\delta_i \tag{3.21}$$

$$\ge q(1-q/n)^{(n-1)} + e^{-q}\sum_{i=2}^{n}\frac{q^i\left(1+H_{i-1}\right)}{i!\cdot i} - \frac{2q}{n}. \tag{3.22}$$

It is easy to see that simply setting $q=2$, which corresponds to $p = 2/n$ and τ being the $2/n$-th quantile of $\mathcal{D}$, yields $\alpha(Y) > 0.5801$ for all $n \ge 20$. Thus, our simple single-threshold algorithm, augmented with a single oracle call, beats, even for small n, the optimal algorithm for the I.I.D. prophet inequality which uses different thresholds per distribution and achieves a probability of success approximately 0.5801 [40].

Since the worst-case probability of ≈ 0.5801 by [40] is achieved for $n \to \infty$, one might interested in the asymptotic behaviour of the probability of our algorithm, $\alpha(Y)$, for large n. It is not too difficult to see, after some calculations, that as $n \to \infty$ (3.22) is maximized for $q \approx 2.435$, yielding $\alpha(Y) \approx 0.798$. QED

3.4.2 A Single-Threshold Algorithm for General m

As we saw in the previous section, even for a simple, single-threshold algorithm, the analysis of the winning probability gets tedious very quickly. In this section, we generalize our single-threshold algorithm to the case of general m, and use the fact that the maximum of a uniformly random permutation of n values changes $\mathrm{O}(\log n)$ times with high probability to obtain a guarantee on the winning probability that is super-exponential with respect to m.

As before, our algorithm selects a threshold τ such that $p = \Pr[X \geq \tau]$ and every time the algorithm observes a realization above τ, it uses an oracle query and asks $\mathcal{O}$ if the realization should be selected or not. If the algorithm runs out of oracle calls, then it selects the first element above τ that is encounters, if any. In other words, the algorithm uses the oracle calls greedily for all realizations above τ.

Theorem 3.3
For every instance of $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ and sufficiently large n, there exists an algorithm that selects the maximum realization with probability at least $1 - \mathrm{O}\left(m^{-m/5}\right)$.

Proof. Let $L = e^{\sqrt{m}}$. The idea is to set τ so that $p = \Pr[X \geq \tau] = L/n$. As before, let Y be the number of realizations above τ. By Lemma 3.1, we have

$$\Pr[|Y - L| \geq \delta L] \leq 2e^{-\delta^2 L/3}. \tag{3.23}$$

Setting $\delta = 1$ yields that $1 \leq Y \leq 2L$ with probability at least $1 - 2e^{-L/3} = 1 - 2e^{-e^{\sqrt{m}}/3} \geq 1 - m^{-m/4}$ for all m.

Next, let $X'_1, \ldots, X'_Y$ be the subsequence of all realizations larger than τ, according to their arrival order, and let $Z_i = 1$ if $X'_i > \max_{j=1}^{i-1} X'_j$, in other words if X'_i is larger than all previous realizations, and $Z_i = 0$ otherwise. Observe that $\Pr[Z_i = 1] = 1/i$, and that the random variables $Z_1, \ldots, Z_n$ are independent. Furthermore, let $M = \sum_i Z_i$ be the number of times that the maximum realization changes in the sequence $X'_1, \ldots, X'_Y$. Observe that if $M \leq m + 1$, then m oracle queries are sufficient for the algorithm to always select the maximum realization. Therefore, our goal is to bound the probability that this event happens.

Conditioned on $1 \leq Y \leq 2L$, we have

$$\mathbb{E}[M] = \sum_{i=1}^{2L} \frac{1}{i} \leq \log(2L) + 1 \leq \sqrt{m} + 2. \tag{3.24}$$

For $\delta = {m+2}/{\mathbb{E}[M]} - 1$, we have

$$\Pr[M \geq m + 2] = \Pr[M \geq (1 + \delta)\, \mathbb{E}[M]]. \tag{3.25}$$

Notice that for $m \geq 98$, we have $\delta \geq e^2$, and thus, by Lemma 3.1, we obtain

$$\Pr[M \geq m + 2] \leq e^{-\mathbb{E}[M]\delta \log \delta/2} \leq e^{-\frac{(m-\sqrt{m})\left(\log\left(m-\sqrt{m}\right)-\log(m+2)/2\right)}{2}} \leq m^{-m/5}. \tag{3.26}$$

If we instead use the tight Chernoff bound in Lemma 3.1, we can show that $\Pr[M \geq m + 2] \leq m^{-m/4+\varepsilon}$ for all m and $\varepsilon > 0$.

Putting everything together, for our algorithm to succeed, it suffices to have $1 \leq Y \leq 2L$ and $M \leq m + 1$, both of which happen together with probability at least $1 - \mathrm{O}\left(m^{-m/5}\right)$. QED

3.4.3 An Asymptotically (Almost) Tight Upper Bound

Now that we have presented a simple, single-threshold algorithm for the $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ setting, a reasonable question to ask is how far it is from being optimal. As we show in this section, the algorithm is asymptotically almost optimal.

Theorem 3.4
There exists an instance of $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ for which no algorithm can select the maximum realization with probability greater than $1 - \mathrm{O}(m^{-m})$.

Proof. To construct an instance in which no algorithm can achieve a high probability, fix m and consider n random variables $X_1, \ldots, X_n$ drawn I.I.D. from the uniform distribution on $[0, 1]$, where n is a sufficiently large number. We first divide $[0, 1]$ into $k = n/m \log m$ intervals $B_1, \ldots, B_k$ of length $m \log m/n$ each, with $B_i = \left((i-1) \cdot m \log m/n, i \cdot m \log m/n\right]$. For each $i = 1, \ldots, n$, let Y_i denote the random variable that is equal to 1 if $X_i \in B_k$ and 0 otherwise, where B_k is the last interval. Also, let $Y = \sum_{i=1}^{n} Y_i$. Since the X_i's follow the uniform distribution, we have $\Pr Y_i = 1 = \frac{m \log m}{n}$ for all i, and $\mathbb{E}[Y] = m \log m$.

Next, consider an algorithm $\mathcal{A}$ for $\mathcal{M}(\mathcal{O}_m, I.I.D., PbM)$ on this instance, and assume that $Y \geq 1$, i.e. there exists at least one realization that falls in the last interval. Consider the moment that $\mathcal{A}$ observes a realization $X_i \in B_k$ that is larger than all previous realizations (including previous realizations in B_k). There are two cases:

- If $\mathcal{A}$ decides not to expend a query to $\mathcal{O}$ for this realization and skip it, there is a chance it fails to select the highest realization. This definitely happens if no other realization in the future is in B_k, which occurs with probability

$$\left(1 - \frac{m \log m}{n}\right)^{n-i} \geq \left(1 - \frac{m \log m}{n}\right)^{n} \geq e^{-m \log m - 1} = \Omega\left(m^{-m/1-\varepsilon}\right) \tag{3.27}$$

  for sufficiently large n, for any $\varepsilon > 0$.

- If $\mathcal{A}$ decides to expend a query to $\mathcal{O}$ for this realization, there is a chance it fails to select the highest realization by running out of queries, deciding to select the next realization in B_k that is higher than all previous ones, and missing out on a higher realization in the future. For this to happen, it must be that $Y \geq m + 2$. Let $\delta = 1 - 1 + 1/m/\log m$. By Lemma 3.1, this happens with probability

$$\Pr[Y > m+1] = 1 - \Pr[Y \leq m+1] = 1 - \Pr[Y \leq (1-\delta)\,\mathbb{E}[Y]] \geq 1 - e^{-\frac{m \log m \left(\log m - 1 - 1/m\right)^2}{2 \log m^2}} \geq 1 - m^{-m/4}. \tag{3.28}$$

  Given that $Y \geq m + 2$, the probability that the first $m + 2$ realizations arrive in increasing order is $1/(m+2)!$. Therefore, $\mathcal{A}$ misses out on the maximum realization in this case with probability at least

$$\frac{1 - m^{-m/4}}{(m+2)!} \geq m^{-m}. \tag{3.29}$$

Therefore, $\mathcal{A}$ must miss the maximum realization with probability at least $\Omega\left(m^{-m}\right)$. QED

3.5 THE NON-I.I.D. SETTING

Motivated by the early work of Gilbert and Mosteller [40] for the top-1-of-m model, we studied the I.I.D. setting and the PbM objective in the previous section. As we described, later works [42], [43], [129] studied the top-1-of-m model with the RoE objective. Using the reduction of Theorem 3.2, any guarantees we provide for the oracle model with the RoE objective can be directly translated to guarantees for the top-1-of-m model. For this reason, in this section we study our oracle model for not necessarily identical random variables and the RoE objective. We provide a simple, single-threshold algorithm that improves slightly upon the current best-known competitive ratio, but our main contribution is an asymptotically (almost) tight upper bound.

3.5.1 A Single-Threshold Algorithm

We first describe our algorithm that slightly improves the current best-known guarantee for the top-1-of-m model [43], even though it works for our oracle model, which is weaker under the RoE objective.

Theorem 3.5
For every instance of $\mathcal{M}(\mathcal{O}_m, \text{Non-I.I.D.}, \text{RoE})$ and sufficiently large n, there exists an algorithm that achieves a competitive ratio of $1 - \mathrm{O}\left(e^{-m/5.178}\right)$.

Proof. Just like our previous algorithms, this algorithm is quite simple; we select a threshold τ and greedily use all oracle calls on the first $m+1$ values above τ. As always, if $\mathcal{O}$ responds YES, the algorithm accepts the current realization and otherwise it continues. Once it run out of oracle calls, it accept the next value it sees that is above τ.

Let N be the number of realizations above τ. Let $p_i = \Pr[X_i \geq \tau]$, and notice that $\mathbb{E}[N] = \sum_i p_i$. We select τ such that $\mathbb{E}[N] = m/2 + 1$. For every i, let $Y_i = 1$ if $X_i \geq \tau$ and 0 otherwise. Recall that $Z = \max_i X_i$, and let $\mathcal{D}_{\max}$ denote the distribution of Z and V denote the value selected by our algorithm. Notice that when $1 \leq N \leq m+1$, we have $V = Z$, and when $N \geq m+2$, we can lower bound V by τ, since our algorithm will always select a value in this case. By Lemma 3.1, for $\delta = 1$, we have

$$\Pr[N \leq 0] = \Pr[N \leq (1-\delta)\,\mathbb{E}[N]] \leq e^{-(m/2+1)}, \tag{3.30}$$

and

$$\Pr[N \geq m+2] = \Pr[N \geq (1+\delta)\,\mathbb{E}[N]] \leq \left(\frac{e}{4}\right)^{-(m/2+1)} \leq e^{-2m/\log 4 - 1} = e^{-m/5.178}. \tag{3.31}$$

Thus,

$$\Pr[1 \leq N \leq m+1] \geq 1 - 2e^{-m/5.178}. \tag{3.32}$$

Now, for all $z \geq \tau$, we know that

$$\Pr_Z[Z \geq \tau \wedge V = Z \mid Z = z] \geq \Pr[1 \leq N \leq m+1] \geq 1 - 2e^{-m/5.178}, \tag{3.33}$$

since $1 \leq N \leq m+1$ immediately implies that our algorithm selects the maximum, regardless of its actual value. Putting everything together, the expected value of our algorithm is

$$\mathbb{E}[V] = \mathop{\mathbb{E}}_{z\sim\mathcal{D}_{\max}}\left[z \cdot \Pr_Z[Z \geq \tau \wedge V = Z \mid Z = z]\right] \tag{3.34}$$

$$= \mathop{\mathbb{E}}_{z\sim\mathcal{D}_{\max}}\left[z \cdot \Pr[1 \leq N \leq m+1]\right] \tag{3.35}$$

$$\geq \left(1 - 2e^{-m/5.178}\right) \underset{z \sim \mathcal{D}_{\max}}{\mathbb{E}} [z \mid Z = z] \tag{3.36}$$

$$= \left(1 - 2e^{-m/5.178}\right) \mathbb{E}[Z]. \tag{3.37}$$

QED

3.5.2 An Asymptotically (Almost) Tight Upper Bound

Next, we show that no algorithm for $\mathcal{M}(\mathcal{O}_m, \text{Non-I.I.D.}, RoE)$ can achieve a competitive ratio greater than $1 - \frac{1}{2^{m+1}}$. The same counterexample works also for the PbM objective, as can easily be seen from its proof, giving us the same upper bound for the $\mathcal{M}(\mathcal{O}_m, \text{Non-I.I.D.}, PbM)$ setting.

Theorem 3.6

There exists an instance of the $\mathcal{M}(\mathcal{O}_m, Non-I.I.D., z)$ setting, where $z = RoE$ or PbM, in which no algorithm can achieve a $\left(1 - \frac{1}{2^{m+1}} + \delta\right)$-competitive ratio or select the maximum realization with probability $\left(1 - \frac{1}{2^{m+1}} + \delta\right)$, for any $\delta > 0$.

Proof. Consider the following instance with $m + 2$ random variables: Let $\varepsilon > 0$ and

$$X_1 = 1 \quad \text{w.p. } 1, X_i = \begin{cases} (1+\varepsilon)^i & \text{w.p. } 1/2 \\ 0 & \text{w.p. } 1/2 \end{cases} \quad \text{for } i = 2, \dots, m+1, \text{ and } X_{m+2} = \begin{cases} \frac{1}{\varepsilon} & \text{w.p. } \varepsilon \\ 0 & \text{w.p. } 1 - \varepsilon \end{cases}. \tag{3.38}$$

We should note that the non-zero values of $X_2, \dots, X_{n-1}$ need not form a geometric series as shown here. It suffices that the non-zero value of X_i is $1 + \varepsilon_i$, where $\varepsilon_2 \leq \dots \leq \varepsilon_{n-1}$, but we used the construction above as it uses a single constant ε and it simplifies the calculations. The result remains the same in both cases, as in the end we take $\varepsilon \to 0$.

First notice that the prophet obtains a value of

$$\mathbb{E}[\max\{X_1, \dots, X_n\}] = \frac{1}{\varepsilon} \cdot \varepsilon + (1 - \varepsilon)\, \mathbb{E}[\max\{X_1, \dots, X_{n-1}\}], \tag{3.39}$$

where, for $\varepsilon \to 0$, the maximum of $X_1, \dots, X_{n-1}$ is always 1. Therefore, as $\varepsilon \to 0$, the prophet's expected value is 2.

Next, consider an algorithm for this setting. There are two cases: either the algorithm decides to query $\mathcal{O}$ at $X_1 = 1$ or not.

Assume that the algorithm decides not to query $\mathcal{O}$ at X_1. Then, it proceeds to observe $X_2, \dots, X_n$ with $m = n - 1$ oracle calls to $\mathcal{O}$ available. This implies that it can query $\mathcal{O}$ at every random variable from X_2 through to X_n, and obtain the optimal value of this residual instance. This value is $1/\varepsilon$ with probability ε, and with probability $1 - \varepsilon$ we obtain a value of 1, as $\varepsilon \to 0$, so long as at least one of $X_2, \dots, X_{n-1}$ is non-zero, which happens with probability $1 - 1/2^{n-2} = 1 - 1/2^m$. Overall, as $\varepsilon \to 0$, the expected value of the algorithm, if it decides not to query $\mathcal{O}$ at X_1, is $1 + 1 - 1/2^m = 2 - 1/2^m$. Since the prophet's value is 2, the algorithm obtains a competitive ratio of $2 - 1/2^m / 2 = 1 - \frac{1}{2^{m+1}}$.

Now assume that the algorithm decides to query $\mathcal{O}$ at X_1. We differentiate between two cases:

- First, assume that $X_2 = \dots = X_{n-1} = 0$. This happens with probability $1/2^m$. In this case, $\mathcal{O}$ will respond YES only if $X_n = 1/\varepsilon$, and thus the expected value of the algorithm is $\frac{1}{\varepsilon} \cdot \varepsilon + (1 - \varepsilon) \cdot 1 \to 2$ as $\varepsilon \to 0$.

- Next, assume that $X_2, \ldots, X_{n-1}$ are not all 0. Since the algorithm queried $\mathcal{O}$ at X_1, it will have $m-1$ remaining oracle calls for $X_2, \ldots, X_{n-1}$. This implies that no matter how the algorithm makes its decisions, there exists a (potentially randomly chosen) X_i on which the algorithm will not use a query to $\mathcal{O}$. Consider now the following realization of $X_2, \ldots, X_{n-1}$: $X_2, \ldots, X_i$ are all non-zero and $X_{i+1}, \ldots, X_{n-1}$ are all 0. This happens with probability exactly $1/2^m$. Assume, for simplicity, that in every other realization of $X_2, \ldots, X_{n-1}$, the algorithm obtains the optimal value of 2, as $\varepsilon \to 0$.

  In this realization, however, when the algorithm reaches X_i, it either selects it and gains an expected value of 1, for small ε, or it skips it without querying $\mathcal{O}$, but ends up regretting this decision as all the remaining random variables up to X_{n-1} are 0. In this latter case, the algorithm reaches X_n and obtains an expected value of $\frac{1}{\varepsilon} \cdot \varepsilon = 1$.

Overall, the algorithm's expected value if it decides to query X_1 is $2\left(1 - 1/2^m\right) + 1 \cdot 1/2^m = 2 - 1/2^m$. Therefore, as in the previous case, the algorithm obtains a competitive ratio of $2 - 1/2^m / 2 = 1 - \frac{1}{2^{m+1}}$. Notice that the example shown is sharp in that there exists an algorithm for such instances that achieves the $1 - \frac{1}{2^{m+1}}$ ratio.

As shown above, every algorithm for this instance fails to select the maximum realization with probability more than $1 - \frac{1}{2^{m+1}}$, and thus the same upper bound holds for the PbM objective as well. QED

Chapter 4: OPTIMAL GREEDY ONLINE CONTENTION RESOLUTION SCHEMES

4.1 OVERVIEW

In Chapter 1, we introduced *Online Contention Resolution Schemes (OCRSs)* to solve stochastic combinatorial optimization problems with adversarial arrival order such as Fair-Cantina-Sale. The general procedure to solve such problems is to model them as an Integer Linear Program (ILP), relax it to a Linear Program (LP), solve it in polynomial time to obtain a solution x, and then round it. Such rounding algorithms have recently been used to obtain several optimal and interesting results [20], [21], [26]–[30], and have more applications in online mechanism design and posted pricing mechanisms [6], [7].

For the rounding, we usually think of every variable x_i of the LP solution as a "probability" that element i is *active*. A first approach to round x is to create a selected set of elements in which every element appears independently with probability x_i. Notice that this corresponds to always selecting every active element, and can be done regardless of whether the active elements are revealed offline or online. For any additive objective function, by linearity of expectation, such an algorithm does not incur any loss. However, there is no guarantee that the resulting integral set is feasible.

Contention Resolution Schemes. To address this issue, one can decide to forego optimality with respect to the objective function, in order to ensure feasibility of the final solution. In other words, we can decide to select an active element with some probability $c \in (0, 1)$, with the hope that this leads to a feasible set at the end of the process. An algorithm that achieves this when it is able to see all the active elements offline is called a *Contention Resolution Scheme (CRS)* [26], and if it is able to do it even when the active elements are revealed in adversarial order, it is called an OCRS [27]. The highest c an OCRS can achieve, while maintaining that the final solution is always feasible, is called the *selectability* of the OCRS. From an application viewpoint, OCRSs are useful in a variety of online settings in Bayesian and stochastic online optimization, such as prophet inequalities [20], [21], [44], stochastic probing [83]–[86], and posted pricing mechanisms [6].

As it turns out, most known OCRSs are *greedy OCRSs*, in which the algorithm (randomly) commits to a subset S of the elements and greedily selects the active elements of S, as long as they don't violate the constraints. While greedy OCRSs constitute a special case, they have several advantages. Usually, to obtain an optimal non-greedy OCRS, one has to use LP duality [44]. This leads to a non-intuitive algorithm which, in many situations, can be difficult to implement. However, greedy OCRSs are inherently simpler than their non-greedy counterparts. Another advantage is that they are able to provide guarantees against stronger adversaries. In fact, for some settings, we are not aware of how to obtain good guarantees without greedy OCRSs. One such example is [136] in which the authors study the "delegation gap" of the generalized Pandora's box problem and in fact reduce the problem to the design of an OCRS which is necessarily greedy. For all these reasons, studying the potential loss in selectability one incurs by restricting to greedy OCRSs is an important problem in stochastic combinatorial optimization.

Adversaries. When considering online combinatorial optimization problems in which the arrival order is decided by an adversary, it is important to consider the power such an adversary has. On one end of the spectrum, an *offline* adversary has to commit to the arrival order of the elements before the process starts and without knowledge of any realizations. Clearly, this is a weak form of the adversary. An *online* adversary is endowed with more power, since at every step i, they can decide which element to reveal next based on the realizations of the elements on the previous steps $1, \ldots, i-1$ as well as the actions that the algorithm has taken so far. Perhaps the strongest adversary is the *almighty* adversary, who decides on the arrival order after observing all realizations and also all potential random coins used by the algorithm.

By their nature, the rounding guarantees that greedy OCRSs provide hold even against an almighty adversary. Therefore, if greedy OCRSs can achieve the same selectability as general OCRSs for a combinatorial optimization setting, there is no difference in power between an offline and an almighty adversary.

4.1.1 Our Contributions

We show that this is not the case. The simplest example of a feasibility constraint is a set of all singletons, i.e. when we only want to select a single element. We call this the *single item* or *rank*-1 *matroid* setting. For our first result, we show that for such a setting, no greedy OCRS can achieve a selectability greater than $1/e$. Since there exist (non-greedy) OCRSs that are $1/2$-selectable [44], [73], this result provides a separation between greedy and general OCRSs.

Theorem 4.1. *For every $\varepsilon > 0$, there exists no greedy OCRS for the single-item setting that selects an active element i with probability at least $1/e + \varepsilon$ for all $i \in N$.*

Feldman, Svensson and Zenklusen [27] designed an OCRS for the single-item setting that is $1/4$-selectable. However, their OCRS does not make use of the exact values of the fractional solution x, i.e. it is *oblivious*. A natural question then is whether this OCRS is optimal. Our second contribution is the design of a greedy OCRS for the single item setting that is $1/e$-selectable, thereby settling the selectability question for this setting. Our OCRS relies on the values of the fractional solution x and is thus not oblivious. Recently, [137] gave a beautifully simple OCRS for the single-item setting that is $1/e$-selectable, oblivious, but not greedy, and also showed that no oblivious OCRS can achieve a higher selectability. Determining the optimal selectability of OCRSs that are both oblivious and greedy is an interesting open question.

Theorem 4.2. *There exists a $1/e$-selectable (randomized) greedy OCRS for the single-item setting.*

The greedy OCRS we provide can be extended in a natural way to feasibility constraints beyond the single-item setting. In particular, by running k parallel instances of it, one can achieve a $1/e$-selectability for k-uniform matroids, which are constraints where every set of elements of size at most k is feasible. Similarly, we can extend our greedy OCRS to work for a partition matroid, where the constraint consists of given sets $S_j \subseteq N$ and a set of elements A is feasible if and only if $|A \cap S_i| \leq 1$ for every j. By decomposing the partition matroid into single-item instances, executing our greedy OCRS on every single-item instance and accepting an active element if and only if it was accepted by the greedy OCRS on the corresponding single-item instance, we obtain a $1/e$-selectable greedy OCRS for partition matroids. The details are omitted as the design of both extensions is straightforward.

Corollary 4.1. *There exists a $1/e$-selectable (randomized) greedy OCRS for uniform and partition matroids.*

Finally, we extend Theorem 4.2 to a more general class of matroids, namely transversal matroids. A transversal matroid is defined by a bipartite graph $G = (U \uplus V, E)$, where U is the ground set of the matroid, and a subset of U is independent if it can be matched to a subset of V. Transversal matroids form a well-studied class of matroids and have been used extensively to model matching markets [63], [138], [139]. The OCRS we design actually provides a stronger selectability guarantee for the special case in which every element's neighborhood has size at least 3.

Theorem 4.3. *Let $\mathcal{M} = (U, \mathcal{I})$ be a transversal matroid represented by a bipartite graph $G = (U \uplus V, E)$. There exists a $1/e$-selectable (randomized) greedy OCRS π for $\mathcal{M}$. Furthermore, if for every element $u \in U$ we have $|N(u)| \geq 3$, where $N(u)$ is the set of neighbours of u in G, then π is a $(1 - 1/e)$-selectable (randomized) greedy OCRS for $\mathcal{M}$.*

4.1.2 Related Work

Since their introduction [26], Contention Resolution Schemes (CRSs) have found several applications. As we discussed previously, applications of CRSs in Bayesian mechanism design and posted price mechanisms [7] can be found in [26]. Later, Yan [46] connected mechanism design with the notion of correlation gap [48]. OCRSs were developed [27] with applications to Bayesian mechanism design as one of the main motivations, since they directly translate to competitive ratios for the prophet inequality problem [20], [21], [27], [77]. In fact, Alaei's work on k-uniform matroids [73] precedes [27] and can be seen as an OCRS, even though it is formulated differently. Random order CRSs (ROCRSs) were introduced in [28] and yield improved bounds when the arrival order is random.

Expanding upon the previously mentioned related work, Feldman, Svensson and Zenklusen [27] gave the first greedy OCRS for matroids, which is $1/4$-selectable. Lee and Singla [44] showed a reverse connection between OCRSs and prophet inequalities, obtaining a $1/2$-selectable (non-greedy) OCRS and a $(1 - 1/e)$-selectable ROCRS for matroids. Adamczyk and Włodarczyk [28] obtained several results, including a $1/k+1$-selectable ROCRS for the intersection of k matroids. For matchings, apart from the work of [29], Bruggmann and Zenklusen [30] developed optimal monotone CRSs via a novel polyhedral approach.

Recently, Dughmi [140], [141] showed an equivalence between the existence of constant-factor OCRSs for specific correlated distributions and a constant-factor approximation to the famous matroid secretary problem [142]. Later, Feldman and Zenklusen showed that any guarantee for the matroid secretary problem with a modular objective can be translated to a guarantee for a submodular objective with only a constant factor loss. We discuss the matroid secretary problem more in Chapter 6.

For more information on other related work, see §1.4.

Organization. We begin in Section 4.2 with some background. Then, in Section 4.3, we present our first main result which shows that no greedy OCRS for the single-item case can achieve a selectability higher than $1/e$. Afterwards, in Section 4.4, we design an optimal greedy OCRS for the single-item setting, uniform matroids and partition matroids that is $1/e$-selectable. We proceed with our greedy OCRS for transversal matroids in Section 4.5, which also achieves the optimal $1/e$ selectability, and show it performs even better under some assumptions on the structure of the transversal matroid.

4.2 PRELIMINARIES

Before we proceed, we present the formal definitions of CRSs, OCRSs and greedy OCRSs and briefly describe the $1/4$-selectable single item OCRS by [27].

Let N be a finite ground set. Recall from Chapter 1 the definitions of a feasibility constraint, and its polyhedral relaxation. Given a polyhedral relaxation $\mathcal{P}_\mathcal{I}$ of a constraint $\mathcal{I}$ and a point $x \in \mathcal{P}_\mathcal{I}$, a natural question is whether we can round x in order to obtain a feasible set $S \in \mathcal{I}$. One way to achieve this is via Contention Resolution Schemes, which we define below. For the remainder of this thesis, for $b \in [0,1]$, let $b \cdot \mathcal{P}$ denote a "scaled-down" variant of $\mathcal{P}$, which contains every point $b \cdot x$ for all $x \in \mathcal{P}$.

Definition 4.1 (Contention Resolution Scheme [26])**.** Let $b, c \in [0,1]$. A (b,c)-balanced *Contention Resolution Scheme* π for $\mathcal{P}_\mathcal{I}$ is a procedure that for every $x \in b \cdot \mathcal{P}_\mathcal{I}$ and $A \subseteq N$, returns a random set $\pi_x(A) \subseteq A \cap \text{support}(x)$ and satisfies the following properties:

1. $\pi_x(A) \in \mathcal{I}$ with probability 1, $\quad \forall A \subseteq N, x \in b \cdot \mathcal{P}_\mathcal{I}$, and
2. for all $i \in \text{supp}\,(x)$, $\Pr\left[i \in \pi_x(R(x)) \mid i \in R(x)\right] \geq c, \quad \forall x \in b \cdot \mathcal{P}_\mathcal{I}$,

where $R(x) \subseteq N$ denotes a random set in which every element $i \in N$ appears independently with probability x_i. The scheme is said to be *monotone* if $\Pr\left[i \in \pi_x(A_1)\right] \geq \Pr\left[i \in \pi_x(A_2)\right]$ whenever $i \in A_1 \subseteq A_2$.

For the remainder of this chapter, we drop the subscript in $\mathcal{P}_\mathcal{I}$ and simply write $\mathcal{P}$ whenever the constraint is clear from context.

CRSs are offline rounding schemes. In the case where the arrival order of the elements is selected by an adversary, we can use the following notion of *Online Contention Resolution Schemes (OCRS)* to round x.

Definition 4.2 (Online Contention Resolution Scheme (OCRS) [27])**.** For an online selection setting where a point $x \in \mathcal{P}$ is given, we draw a random subset of the elements $R(x)$, in which each element i appears independently with probability x_i. We call $R(x)$ the set of *active* elements. Afterwards, we observe whether each element $e \in N$ is active ($e \in R(x)$), one by one, and have to immediately and irrevocably decide whether to select an element or not before the next element is revealed. An *Online Contention Resolution Scheme* π for $\mathcal{P}$ is an online algorithm which selects a subset $\pi_x(R(x)) \subseteq R(x)$ such that $\mathbb{1}_{\pi_x(R(x))} \in \mathcal{P}$.

We also define the notion of a *greedy* OCRS, which provide guarantees with respect to an almighty adversary.

Definition 4.3 (Greedy OCRS [27])**.** Let $\mathcal{P} \subseteq [0,1]^n$ be a relaxation of the feasible sets $\mathcal{I} \subseteq 2^N$. An OCRS π for $\mathcal{P}$ is called a *greedy OCRS* if, for any $x \in \mathcal{P}$, π defines a down-closed subfamily of feasible sets $\mathcal{F}_{\pi,x} \subseteq \mathcal{I}$, and it selects an active element e when it arrives if, together with the set of elements already selected, the resulting set is in $\mathcal{F}_{\pi,x}$. We say that π is a randomized greedy OCRS if, given x, the choice of $\mathcal{F}_{\pi,x}$ is randomized. Otherwise, we say that π is a deterministic greedy OCRS.

For the remainder of this paper, we drop the subscript in $\mathcal{F}_{\pi,x}$ and simply write $\mathcal{F}_x$ or $\mathcal{F}$, whenever π and x are clear from context.

Intuitively, we say a greedy OCRS is c-selectable if and only if an active element $e \in R(x)$ can be included in the currently selected elements $S \subseteq R(x)$ and maintain feasibility with probability at least c.

Definition 4.4 (c-selectability)**.** Let $c \in [0,1]$. A greedy OCRS for $\mathcal{P}$ is c-selectable if and only if for any $x \in P$ we have

$$\Pr\left[S \cup \{e\} \in \mathcal{F}_x\right] \geq c, \qquad \forall e \in N,\ S \subseteq R(x) \text{ such that } S \in \mathcal{F}_x.$$

Notice that a c-selectable greedy OCRS guarantees that each active element e is selected with probability at least c, even against the almighty adversary. We should note that the randomness in the above definition is with respect to both the randomness of $R(x)$ and also any potential randomness the greedy OCRS might use to decide $\mathcal{F}_x$.

Next, we briefly describe the $1/4$-selectable single item greedy OCRS by [27]. Given a fractional point x such that $\sum_{i=1}^n x_i \leq 1$, the greedy OCRS will, at step i, observe whether element i is active or not. If it is active, the greedy OCRS will choose to select with probability $1/2$ or discard it and move on to the next element. Since each element is active with probability x_i and is selected with probability $x_i/2$, the expected number of selected elements is at most half, and thus, by Markov's inequality, the probability the greedy OCRS selects no elements is at least $1/2$. Therefore, for every element i, we reach i without having selected an element with probability at least $1/2$ and we select i, given that it is active, with probability $1/2$, for an overall selectability of $1/4$.

We should note that for the single item setting there exists a $1/2$-selectable OCRS [73] but, crucially, it is not greedy. As we show in Section 4.3, there is no $1/2$-selectable greedy OCRS for the single item setting.

Other variants of online CRSs have been studied as well. One example is *Random Order Contention Resolution Schemes (ROCRS).* In an ROCRS, the arrival order of the elements is chosen uniformly at random, instead of being chosen by an adversary. Adamczyk and Wlodarczyk present several interesting results on ROCRSs in [28].

4.3 AN UPPER BOUND

In this section, we present the proof of Theorem 4.1. Consider the instance where $x_i = 1/n$ for all $e_i \in N$, where $n = |N|$, and let A denote the set of active elements. Any greedy OCRS π will select a subset S of $\mathcal{F} = \{e_i \mid e_i \in N\}$ with some probability α_S, and then accept the first element in S that comes up active. What is the worst-case probability that an element from S will be selected? This is minimized for the element in S which is last in the arrival order, which has a probability of being selected exactly equal to $(1 - 1/n)^{|S|-1}$, because the OCRS is greedy, and it would select an element from S which arrived earlier, if it came up active. Therefore, no greedy OCRS can guarantee, for any $S \subseteq \mathcal{F}$, that an element $e \in N$ will be selected, when $e \in A$, with probability greater than $(1 - 1/n)^{|S|-1}$. Thus, for any $e \in N$ and any greedy OCRS π, we have

$$\Pr\left[e \in \pi(A) \mid e \in A\right] \leq \sum_{\substack{S \subseteq N \\ e \in S}} \alpha_S \left(1 - \frac{1}{n}\right)^{|S|-1} \tag{4.1}$$

$$= \sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq N \,:\, |S|=k \\ e \in S}} \alpha_S. \tag{4.2}$$

Next, for a greedy OCRS π to be c-selectable, it needs to guarantee that $\min_{e \in N} \Pr\left[e \in \pi(A) \mid e \in A\right] \geq c$. Therefore, if we show that

$$\min_{e \in N} \left\{ \sum_{k=1}^{n} \left(1 - \frac{1}{n}\right)^{k-1} \sum_{\substack{S \subseteq N \,:\, |S|=k \\ e \in S}} \alpha_S \right\} \leq c, \tag{4.3}$$

by (4.2) it follows that π cannot be $(c + \varepsilon)$-selectable for any $\varepsilon > 0$.

Lemma 4.1.

$$\min_{e\in N}\left\{\sum_{k=1}^{n}\left(1-\frac{1}{n}\right)^{k-1}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right\}\leq\left(1-\frac{1}{n}\right)^{n-1}.$$

Proof. Assume towards contradiction, that

$$\min_{e\in N}\left\{\sum_{k=1}^{n}\left(1-\frac{1}{n}\right)^{k-1}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right\}>\left(1-\frac{1}{n}\right)^{n-1}. \tag{4.4}$$

The proof consists of a double counting argument. First, notice that, by the inequality above, we have

$$\sum_{e\in N}\left(\sum_{k=1}^{n}\left(1-\frac{1}{n}\right)^{k-1}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right)>n\left(1-\frac{1}{n}\right)^{n-1}. \tag{4.5}$$

For any $0\leq k\leq n$, let $\beta_k=\sum_{S\subseteq N\,:\,|S|=k}\alpha_S$ be the total probability mass assigned by the greedy OCRS to all sets of size k, and notice that $\sum_{k=0}^{n}\beta_k=1$. We can also compute the left-hand side of (4.5) as

$$\begin{aligned}\sum_{e\in N}\left(\sum_{k=1}^{n}\left(1-\frac{1}{n}\right)^{k-1}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right)&=\sum_{k=1}^{n}\left(\left(1-\frac{1}{n}\right)^{k-1}\sum_{e\in N}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right) &(4.6)\\ &=\sum_{k=1}^{n}\left(k\left(1-\frac{1}{n}\right)^{k-1}\sum_{S\subseteq N\,:\,|S|=k}\alpha_S\right) &(4.7)\\ &=\sum_{k=1}^{n}\left(\beta_k\cdot k\left(1-\frac{1}{n}\right)^{k-1}\right). &(4.8)\end{aligned}$$

where the second equality follows from the fact that in the double sum, for every S such that $|S|=k$, every coefficient a_S appears exactly k times, one for each element it contains. Under the constraint $\sum_{k=0}^{n}\beta_k=1$, we have that $\sum_{k=1}^{n}\left(\beta_k\cdot k\left(1-1/n\right)^{k-1}\right)$ is maximized for $\beta_n=1$ and $\beta_m=0$ for all $m<n$, as $k\left(1-1/n\right)^{k-1}$ is strictly increasing in k. Therefore,

$$\sum_{e\in N}\left(\sum_{k=1}^{n}\left(1-\frac{1}{n}\right)^{k-1}\sum_{\substack{S\subseteq N\,:\,|S|=k\\ e\in S}}\alpha_S\right)\leq n\left(1-\frac{1}{n}\right)^{n-1}. \tag{4.9}$$

Together, (4.5) and (4.9) yield a contradiction. QED

By Lemma 4.1, since $\lim_{n\to\infty}\left(1-1/n\right)^{n-1}=1/e$, it follows that there exists no greedy OCRS for $\mathcal{P}$ that selects an element e, when active, with probability at least $1/e+\varepsilon$ for all $e\in N$.

4.4 AN OPTIMAL GREEDY OCRS FOR RANK-1 MATROIDS

This section is dedicated to proving Theorem 4.2. Before we begin, we need the following lemma.

Lemma 4.2. *Let $a_1, \ldots, a_k \in [0,1]$. Then*

$$\ln\left(1 - \frac{a_k}{2}\right) + \sum_{j=1}^{k-1} \ln\left(1 - a_j + \frac{a_j^2}{2}\right) \geq -a_k - \sum_{j=1}^{k-1} a_j.$$

Proof. We split the statement into two separate parts.

Claim 4.1.

$$\ln\left(1 - \frac{a_k}{2}\right) \geq -a_k.$$

Proof. Consider the function $f : [0,1] \to \mathbb{R}_{\geq 0}$, where $f(x) = e^x\left(1 - x/2\right)$. Clearly, if $f(x) \geq 1$ for all $x \in [0,1]$, then the claim follows by taking the (natural) logarithm of each side of the inequality, and setting $x = a_k$.

We have $\frac{df(x)}{dx} = -e^x/2\,(x-1) \geq 0$ for all $x \in [0,1]$. Therefore, f is increasing in $[0,1]$, and thus attains its minimum for $x = 0$. Therefore, $f(x) \geq f(0) = 1$ for all $x \in [0,1]$ and the claim follows. QED

Claim 4.2. *For every $j \in \{1, 2, \ldots, k-1\}$, we have*

$$\ln\left(1 - a_j + \frac{a_j^2}{2}\right) \geq -a_j.$$

Proof. Fix an arbitrary a_j. Consider the function $g : [0,1] \to \mathbb{R}_{\geq 0}$, where $g(x) = e^x\left(1 - x + x^2/2\right)$. Clearly, if $g(x) \geq 1$ for all $x \in [0,1]$, then the claim follows by taking the (natural) logarithm of each side of the inequality, and setting $x = a_j$.

We have $\frac{dg(x)}{dx} = e^x x^2/2 \geq 0$ for all $x \in [0,1]$. Therefore, g is increasing in $[0,1]$, and thus attains its minimum for $x = 0$. Therefore, $g(x) \geq g(0) = 1$ for all $x \in [0,1]$ and the claim follows. QED

QED

Next, consider a ground set $N = \{e_1, e_2, \ldots, e_n\}$, and let $\mathcal{M} = (N, \mathcal{I})$ be the uniform matroid of rank 1 with respect to N, i.e. $\mathcal{I} = \{\{e_i\} \mid e_i \in N\}$. Let $\mathcal{P}$ be the following polyhedral relaxation of $\mathcal{M}$:

$$\mathcal{P} = \left\{ x \in [0,1]^n \;\middle|\; \sum_{i=1}^{n} x_i \leq 1 \right\}. \tag{4.10}$$

For a given $x \in \mathcal{P}$, let $\pi = \pi_x$ denote the OCRS we will create. π will draw a random set $R(q)$ where each element e_i appears in $R(q)$ independently with some probability q_i. The family of feasible subsets is

$$\mathcal{F}_{\pi,x} = \{\{e_i\} \mid e_i \in R(q)\}. \tag{4.11}$$

We set $q_i = 1 - x_i/2$ for all $e_i \in N$. Afterwards, π selects the first element e_i that is active and that $\{e_i\} \in \mathcal{F}$.

Lemma 4.3. *π is a randomized greedy OCRS.*

Proof. π is clearly a randomized OCRS because every time it sees an element, it makes an irrevocable decision to select it, if it is active, before it sees the next element, and also, by the choice of $\mathcal{F}_{\pi,x}$, it is easy to see that the set of elements it returns is always a singleton, and thus feasible in $\mathcal{I}$, since $\mathcal{F}_{\pi,x} \subseteq \mathcal{I}$. Furthermore, the choice of $\mathcal{F}_{\pi,x}$ is randomized, and thus π is a randomized OCRS.

Next, it is also easy to see that π is a greedy OCRS, because, given x, $\mathcal{F}_{\pi,x}$ is a down-closed subfamily of feasible sets and an active element e is always selected if $\{e\} \in \mathcal{F}_{\pi,x}$, since there are no previously selected elements. QED

Next, we quantify the probability that each element is selected by π, given that it is active.

Lemma 4.4. *π selects every element $e_i \in N$, given that e_i is active, with probability at least ${}^1\!/e$.*

Proof. We relabel the elements of N so that each e_i arrives in the i-th step. Consider an element $e_i \in N$. Given that e_i is active, since π is a greedy OCRS, π will select e_i if and only if it has not selected any elements before e_i and also $\{e_i\} \in \mathcal{F}_{\pi,x}$. Recall that we have $\{e_i\} \in \mathcal{F}_{\pi,x}$ with probability exactly $q_i = 1 - {}^{x_i}\!/2$. Furthermore, for every element e_j where $j < i$, it needs to be the case that we avoid having both $\{e_j\} \in \mathcal{F}_{\pi,x}$ and also e_j coming up active. This happens with probability $1 - x_j \cdot (1 - {}^{x_j}\!/2) = 1 - x_j + {}^{x_j^2}\!/2$ for every e_j where $j < i$. Overall, if we denote by r_i the probability that e_i is selected by π, given that it is active, we have

$$\ln r_i = \ln\left(\left(1 - \frac{x_i}{2}\right) \cdot \prod_{j=1}^{i-1}\left(1 - x_j + \frac{x_j^2}{2}\right)\right) = \ln\left(1 - \frac{x_i}{2}\right) + \sum_{j=1}^{i-1} \ln\left(1 - x_j + \frac{x_j^2}{2}\right) \tag{4.12}$$

$$\geq -x_i - \sum_{j=1}^{i-1} x_j \geq -1, \tag{4.13}$$

where the first inequality follows from Lemma 4.2 and the second inequality follows from $\sum_i x_i \leq 1$. Therefore $r_i \geq {}^1\!/e$, for all $i \in N$. QED

From Lemmas 4.3 and 4.4, it follows that π is a ${}^1\!/e$-selectable (randomized) greedy OCRS for $\mathcal{P}$.

4.4.1 An Alternative Scheme

After we notified him of our scheme, Jan Vondrák devised an alternate scheme for the problem. With his consent [143], we present his greedy OCRS below.

Let π denote the OCRS we will create. π will draw a random set R where each element e_i appears in R independently with some probability q_i. Afterwards, it will set

$$\mathcal{F}_{\pi,x} = \{\{e_i\} \mid e_i \in R\}\,. \tag{4.14}$$

We set $q_i = {}^{1-e^{-x_i}}\!/x_i$ for all $e_i \in N$. Afterwards, π selects the first element e_i that is active and that $\{e_i\} \in \mathcal{F}$.

The proof of the next lemma is identical to the proof of Lemma 4.3.

Lemma 4.5. *π is a randomized greedy OCRS.*

Next, we quantify the probability that each element is selected by π, given that it is active.

Lemma 4.6. *π selects every element $e_i \in N$, given that e_i is active, with probability at least ${}^1\!/e$.*

Proof. We relabel the elements of N so that each e_i arrives in the i-th step. Consider an element $e_i \in N$. Given that e_i is active, since π is a greedy OCRS, π will select e_i if and only if it has not selected any elements before e_i and also $\{e_i\} \in \mathcal{F}_{\pi,x}$. Recall that we have $\{e_i\} \in \mathcal{F}_{\pi,x}$ with probability exactly $q_i = {}^{1-e^{-x_i}}\!/x_i$.

Furthermore, for every element e_j where $j < i$, it needs to be the case that we avoid having both $\{e_j\} \in \mathcal{F}_{\pi,x}$ and also e_j coming up active. This happens with probability $1 - x_j \cdot {}^{1-e^{-x_j}}/_{x_j} = e^{-x_j}$ for every e_j where $j < i$. Overall, if we denote by r_i the probability that e_i is selected by π, given that it is active, we have

$$r_i = \frac{1 - e^{-x_i}}{x_i} \cdot \prod_{j<i} e^{-x_j} = \frac{1 - e^{-x_i}}{x_i} \cdot e^{-\sum_{j<i} x_j} \geq \frac{(1 - e^{-x_i})\, e^{x_i - 1}}{x_i} = \frac{e^{x_i-1} - e^{-1}}{x_i}, \tag{4.15}$$

where the inequality follows from $\sum_i x_i \leq 1$. This expression is minimized for $x_i \to 0$, and thus we get $r_i \geq {}^1/_e$, for all $i \in N$. QED

From Lemmas 4.5 and 4.6, it follows that π is a ${}^1/_e$-selectable (randomized) greedy OCRS for $\mathcal{P}$.

Remark 4.1. One can easily see that the difference between the two proofs is that, in our scheme, the probability of selection q_i of each element $i \in N$ is a linear approximation of the selection probability of Vondrák's scheme. The result then follows due to the convexity of the selection probability $q_i = {}^{1-e^{-x_i}}/_{x_i}$ of Vondrák's scheme.

4.5 EXTENSION TO TRANSVERSAL MATROIDS

In this section, we prove Theorem 4.3. Let $\mathcal{M} = (U, \mathcal{I})$ be a transversal matroid and $G = (U \uplus V, E)$ denote the underlying bipartite graph, where $|U| = n$. We know that a subset $S \subseteq U$ is independent if and only if there exists a matching in G that covers S. Let $\mathcal{P}$ be the natural polyhedral relaxation of $\mathcal{M}$. For a given $x \in \mathcal{P}$, let $\pi = \pi_x$ be the greedy OCRS we will create. Let $N(u)$ denote the set of neighbors of a vertex u in G. For each $v \in V$, π will draw a random set $R_v \subseteq N(v)$, in which each element $u \in U$ appears with probability q_u. Then, we set

$$q_u = 1 - \left(1 - \frac{1 - e^{-x_u}}{x_u}\right)^{\frac{1}{|N(u)|}}. \tag{4.16}$$

It is easy to see that $q_u \in [0, 1]$ for every $|N(u)| \geq 1$, and thus q_u is well-defined.

Next, we create a down-closed subfamily of feasible sets by taking all possible combinations of sets created by taking at most one element from each R_v and then taking the union of all such elements. Specifically,

$$\mathcal{F}_{\pi,x} = \left\{ S = \{u_1, \ldots, u_k\} \subseteq U \;\middle|\; \exists\, T = \{v_1, \ldots, v_k\} \subseteq V \text{ s.t. } u_j \in R_{v_j},\ \forall j \in \{1, \ldots, k\} \right\}. \tag{4.17}$$

Any set S in $\mathcal{F}$ is clearly an independent set of $\mathcal{M}$, as the constraints guarantee that there always exists a matching in G that covers S. During the online process, π starts with a set of selected elements $S = \emptyset$, and greedily selects an active element u if $S + u \in \mathcal{F}$.

The proof of the following lemma is identical to the proof of Lemma 4.3 and follows from the discussion above.

Lemma 4.7. *π is a randomized greedy OCRS.*

Next, we again lower bound the selection probability of an active element.

Lemma 4.8. *π selects every element $u \in U$, given that u is active, with probability at least ${}^1/_e$. Furthermore, if $|N(u)| \geq 3$ for all $u \in U$, π selects every element $u \in U$, given that it is active, with probability at least $1 - {}^1/_e$.*

Proof. Consider an active element $u \in U$. Since π is a greedy OCRS, it will select u if and only if there exists a neighbor v of u such that $u \in R_v$, and also, together with the set S of elements already selected by π, $S + u \in \mathcal{F}$. First, for every element $w \in U$, let $\mathcal{E}_w$ denote the event that there exists an element $v \in V$ such that $w \in R_v$. In other words, $\mathcal{E}_w$ is the event that w is in some set of $\mathcal{F}$. For $\mathcal{E}_u$, we have

$$\Pr[\mathcal{E}_u] = 1 - \prod_{v \in N(u)} (1 - q_u) = 1 - (1 - q_u)^{|N(u)|} = \frac{1 - e^{-x_u}}{x_u}. \tag{4.18}$$

Furthermore, the set S selected prior to seeing u has to be independent, thus $S \in \mathcal{F}$, and thus for $S + u \notin \mathcal{F}$, it has to be that for every $v \in N(u)$, we have $|S \cap R_v| \geq 1$. Therefore, the probability that $S + u \notin \mathcal{F}$ is

$$\begin{aligned}
\Pr\left[S + u \notin \mathcal{F} \mid S\right] &= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} \left(1 - x'_u \Pr[\mathcal{E}_{u'}]\right)\right) && (4.19)\\
&= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} \left(1 - x'_u \frac{1 - e^{-x_{u'}}}{x_{u'}}\right)\right) && (4.20)\\
&= \prod_{v \in N(u)} \left(1 - \prod_{\substack{u' \in N(v) \\ u' \neq u}} e^{-x_{u'}}\right) && (4.21)\\
&= \prod_{v \in N(u)} \left(1 - e^{-\sum_{u' \in N(v): u' \neq u} x_{u'}}\right) && (4.22)\\
&\leq \prod_{v \in N(u)} \left(1 - e^{-1 + x_u}\right) && (4.23)\\
&= \left(1 - e^{-1 + x_u}\right)^{|N(u)|}, && (4.24)
\end{aligned}$$

where the inequality follows from the fact that, for every $v \in V$, $\sum_{w \in N(v)} x_w \leq 1$ due to $x \in \mathcal{P}$. Therefore, we have

$$\begin{aligned}
\Pr[u \in \pi(R) | u \in R] &= \Pr[\mathcal{E}_u] \cdot \Pr\left[S + u \in \mathcal{F} \mid S\right] && (4.25)\\
&\geq \frac{1 - e^{-x_u}}{x_u} \left(1 - \left(1 - e^{-1 + x_u}\right)^{|N(u)|}\right). && (4.26)
\end{aligned}$$

Let $f_k(x) = {}^{1 - e^{-x}}/_{x} \left(1 - \left(1 - e^{-1+x}\right)^k\right)$, for $k \geq 1$ and $x \in [0, 1]$. It is easy to see that $f_k(x) \geq {}^1/_e$ for every $k \geq 1$ and $x \in [0, 1]$. Furthermore, we have that for $k \geq 3$, $f_k(x)$ is minimized in $[0, 1]$ for $x = 1$, and yields $f_k(1) = 1 - {}^1/_e$. QED

We conclude that π is a ${}^1/_e$-selectable greedy OCRS for $\mathcal{M}$ and that if $|N(u)| \geq 3$ for every $u \in U$, π is a $(1 - {}^1/_e)$-selectable greedy OCRS for $\mathcal{M}$.

Chapter 5: SUBMODULAR PROPHET INEQUALITIES

5.1 OVERVIEW

In Chapter 4, we studied greedy OCRSs for additive objective functions and simple feasibility constraints such as rank-1, uniform, partition and transversal matroids. However, more general objective functions have several applications. For these objectives, the value of a subset of items from N is specified by a set function $f : 2^N \to \mathbb{R}$. Prominent examples are submodular set functions, introduced in §1.2, and subadditive[5] set functions, as well as their special cases.

The application motivating this chapter is a model of combinatorial prophet inequalities introduced by Rubinstein and Singla [20]. Their model is a generalization of the standard prophet inequality, in which the objective is a set function. [20] reduce their problem to the design of greedy OCRSs for submodular objective functions which form a rich class and constitute the main object of study in this chapter.

Model description. In this model, called the *Submodular Prophet Inequality (SPI)*, the input consists of n independent random variables $X_1, X_2, \ldots, X_n$. Unlike the standard prophet inequality where X_i is a real-valued random variable, in this combinatorial setting, each X_i is a discrete-valued random variable over a finite set $\mathcal{U}_i$. Thus $\mathcal{D}_i$ is a discrete probability distribution over $\mathcal{U}_i$. For technical reasons one assumes that $\mathcal{U}_1, \mathcal{U}_2, \ldots, \mathcal{U}_n$ are mutually disjoint. Let $\mathcal{U} = \bigcup_i \mathcal{U}_i$. There exists a non-negative submodular function $f : 2^{\mathcal{U}} \to \mathbb{R}_+$ defined over the ground set $\mathcal{U}$ and given via a value oracle. As in the standard prophet inequality setting, the variables arrive in an adversarial order. After seeing the realization of a variable in the order, the algorithm has to make an irrevocable decision to accept it or not. Its goal is to maximize the value, $f(S)$, of the selected set $S \subseteq \mathcal{U}$. The input also specifies a feasibility constraint $\mathcal{I} \subseteq 2^N$, and the set S of chosen variables must belong to $\mathcal{I}$. The prophet is allowed to optimize offline after seeing all the realizations and obtains value equal to $\mathbb{E}[\max_{S \in \mathcal{I}} f(\bigcup_{i \in S}\{X_i\})]$. We say that an algorithm achieves an α competitive ratio, or is α-competitive, if the expected value of its choice is at least $\alpha \cdot \mathbb{E}[\max_{S \in \mathcal{I}} f(\bigcup_{i \in S}\{X_i\})]$. Observe that the SPI model generalizes the standard prophet inequality with additive functions and arbitrary downward-closed constraints.

Prior work and Limitations. Rubinstein and Singla [20] presented an $\mathrm{O}(1)$-competitive algorithm for SPI under a matroid constraint. However, the constant they obtained is very small (several orders of magnitude smaller than 1, although they did not try to optimize it) and they did not consider or emphasize the computational aspects of the online algorithm. Usually prophet inequalities in the standard setting of modular/additive objectives achieve large constant guarantees that are relatively close to 1. For instance, the well-known result of Kleinberg and Weinberg [8] showed a bound of $1/2$ even for arbitrary matroid constraints, and it is also known that the bound for a cardinality constraint with k items is $1 - \mathrm{O}(1/\sqrt{k})$ (hence it tends to 1 as $k \to \infty$) [73]. Moreover, [20] did not explicitly consider the case of monotone submodular functions, and did not generalize the constraint family beyond a single matroid.

Motivation. As a real-world application, recall the Fair-Cantina-Sale setting, and consider the following variation, in which the town's mayor has to choose the buyers of the k cantinas of the town schools. In this

[5] A real-valued set function $f : 2^N \to \mathbb{R}$ is subadditive if $f(A \cup B) \leq f(A) + f(B)$ for all $A, B \subseteq N$. A non-negative submodular function is subadditive.

setting, the mayor has to select k out of the n buyers, and the marginal value of selecting each buyer could be decreasing, as if a school's cantina does not open, the students could be served by the cantina of a nearby school, perhaps at a social or economic cost. This is captured by submodular functions. Since the model of [20] can be considered with arbitrary submodular functions over $\mathcal{U}$, it allows for substantial generality.

Our goal in this chapter is to study greedy OCRSs with submodular objective functions via a clean framework that applies to a wide variety of constraints. Via this framework, one can obtain significantly improved bounds for SPI as a corollary. This question was explicitly raised by Lucier in his survey on prophet inequalities [100]. Another motivation, related to the goal of obtaining improved guarantees, comes from a technical tool that Rubinstein and Singla relied upon, namely the notion of the *correlation gap*, which is the ratio of two important continuous extensions of a set function f – the multilinear relaxation F and the concave closure f^+ – and plays a crucial role in submodular optimization [26], [46].

5.1.1 Our Contributions

In this chapter we make two high-level contributions:

1. We address and provide improvements to three aspects of the SPI problem: (i) significantly improved constants for the prophet inequalities for monotone and non-monotone functions, (ii) a clean black-box reduction to greedy Online Contention Resolution Schemes that allows one to obtain prophet inequalities for various other constraints beyond a single matroid constraint and (iii) computational aspects of the prophet inequality that were not explicitly addressed in [20]. In essence, we answer the open question in [100] in the affirmative.

2. We consider the correlation gap for non-negative submodular functions. Rubinstein and Singla used a variant of the standard correlation gap. For both the original definition and the variant of [20], we obtain substantially improved bounds.

We now give a formal statement of our results. We refer the reader to Section 5.2.1 for some basic definitions and background on submodular functions, continuous extensions, the correlation gap and contention resolution schemes.

Submodular Prophet Inequality: For the SPI setting, we follow the high level framework of [20] via the correlation gap and greedy Online Contention Resolution Schemes (OCRSs) [27]. Our main contributions are several technical improvements and refinements that lead to significantly improved constants and clarity on the parameters that affect the constants. The competitive ratio that we achieve for a particular constraint family is dictated by the OCRS available for that family. As stated in Chapter 4, the approximation quality of the OCRS is governed by two parameters $b, c \in [0, 1]$ via the notion of (b, c)-selectability. Informally, we say that an OCRS π for a polytope $\mathcal{P}$ is (b, c)-selectable if it is c-selectable for any $x \in b \cdot \mathcal{P}$, where $b \cdot \mathcal{P}$ is the "scaled-down" variant of $\mathcal{P}$. Notice that an OCRS that is (b, c)-selectable is also bc-selectable in the sense of Definition 4.4, as one can multiply every x_i by b, losing a factor of b in the process, and then obtain a point $x' = bx \in b \cdot \mathcal{P}$. [27] provides $(b, c(b))$-selectable greedy OCRSs for several feasibility constraints for every $b \in [0, 1]$.

Theorem 5.1 (**Informal).** *For the SPI problem with a monotone submodular function f over a constraint family with a (b, c)-selectable OCRS, there is a $c \cdot (e^{-b} - \varepsilon)(1 - e^{-b})$-competitive algorithm for any fixed*

$\varepsilon > 0$. For non-negative submodular functions there is a $\frac{c}{4} \cdot (e^{-b} - \varepsilon)(1 - e^{-b})$-competitive algorithm for any fixed $\varepsilon > 0$. These competitive ratios can be achieved by an efficient randomized polynomial time algorithm, assuming value oracle access to f and efficiency of the corresponding OCRS.

Our results hold in the setting of an almighty adversary. We note that the competitive ratios we obtain are explicit and relatively close to 1. We summarize our concrete competitive ratios for several constraints of interest below. OCRSs for constraints can be composed nicely (similar to CRSs) and thus our black-box reduction is very useful.

Feasibility constraint	Competitive Ratio	
	Monotone Submodular	General Submodular
Uniform matroid of rank $k \to \infty$	$1/4.3$	$1/17.2$
Matroid	$1/7.4$	$1/30$
Matching	$1/9.5$	$1/38$
Knapsack	$1/17.5$	$1/70$
Intersection of k matroids	$\Omega(1/k)$	$\Omega(1/k)$

Table 5.1: A summary of our results for several feasibility constraints.

In subsequent work, Qiu and Singla [144] obtained improved bounds for SPI via the use of *submodular dominance*. In particular, they obtain an optimal $1-1/e$ bound for uniform matroids and monotone submodular functions. Obtaining optimal SPIs under different constraints remains an interesting open problem.

Correlation gap: For a non-negative submodular function, for any given $p \in [0, 1]$, there is a simple instance with $n = 2$ where $F(\boldsymbol{x}) \leq (1 - p)f^+(\boldsymbol{x})$, and this implies that, as p tends to 1, the correlation gap tends to 0. One way to overcome this is to restrict attention to settings where p is bounded away from 1. Nevertheless, there has been little work on precisely quantifying the correlation gap as a function of this parameter. Our first theorem addresses this.

Theorem 5.2. *Let $f : 2^N \to \mathbb{R}_{\geq 0}$ be a non-negative submodular function and let $\boldsymbol{x} \in [0, 1]^n$, where $n = |N|$. Let $p = \max_i x_i$. Then $F(\boldsymbol{x}) \geq (1 - p)(1 - 1/e)f^+(\boldsymbol{x})$. Given any $p \in [0, 1]$ there are instances such that $F(\boldsymbol{x}) \leq (1 - e^{-(1-p)})f^+(\boldsymbol{x})$.*

The upper bound of $(1 - p)(1 - 1/e)$ is optimal when p is close to 0 or when p is close to 1. The lower bound on the gap that we show, $1 - e^{-(1-p)}$, agrees nicely with the extremes, but we do not know whether it is the right bound for all ranges of p and leave it as an interesting open problem.

Rubinstein and Singla [20] instead use the correlation gap of a related function, $f_{\max}(S) = \max_{T \subseteq S} f(T)$. $f_{\max}$ is monotone, but in general is not submodular, even when f is. It is shown in [20] that for any non-negative submodular function f, $\inf_{\boldsymbol{x} \in [0,1]^n} \frac{F_{\max}(\boldsymbol{x})}{f^+(\boldsymbol{x})} \geq 1/200$, where $F_{\max}$ is the multilinear extension of $f_{\max}$. For this variant of the correlation gap, we observe that known results on the Measured Continuous Greedy (MCG) algorithm [28], [145] show that $F_{\max}(\boldsymbol{x}) \geq \frac{1}{e} f^+(\boldsymbol{x})$.

We strengthen this observation by considering the parameter p.

Theorem 5.3. *Let $f : 2^N \to \mathbb{R}_{\geq 0}$ be a non-negative submodular function and let $\boldsymbol{x} \in [0, 1]^n$, where $n = |N|$. Let $p = \max_i x_i$. There exists a point $\boldsymbol{y} \in [0, 1]^n$, where $\boldsymbol{y} \leq \boldsymbol{x}$ (coordinate wise), such that $F(\boldsymbol{y}) \geq \max\{\frac{1}{e}, (1 - p - \frac{1}{e}(1 + \ln(1 - p)))\} f^+(\boldsymbol{x})$.*

We obtain the preceding theorem as a corollary of the following.

Theorem 5.4. *Let $p \in [0,1)$, f be a non-negative submodular function with multilinear extension F and $\mathcal{P}$ be a downward-closed solvable polytope*[6] *on N, such that $\mathcal{P} \subseteq p \cdot [0,1]^N$ (that is, if $\boldsymbol{z} \in \mathcal{P}$ then $z_i \leq p$ for all $i \in N$). Then, the output of the Measured Continuous Greedy (MCG) algorithm on F and $\mathcal{P}$ at time $b \in [0,1]$ is a vector $\boldsymbol{x}(b) \in b \cdot \mathcal{P}$ such that*

$$F(\boldsymbol{x}(b)) \geq \begin{cases} b \cdot e^{-b} \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^{+}(\boldsymbol{z}), & 0 \leq b \leq \ln\left(\frac{1}{1-p}\right) \\ \left(1 - p - e^{-b}\left(1 + \ln\left(1-p\right)\right)\right) \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^{+}(\boldsymbol{z}), & \ln\left(\frac{1}{1-p}\right) \leq b \leq 1. \end{cases}$$

Theorems 5.2 and 5.4 are useful when p is small and we show later that this can indeed be achieved in some cases, such as the SPI problem, via a reduction.

5.1.2 Related Work

The work here is connected to submodular optimization, stochastic optimization, online algorithms, and mechanism design which have extensive literature. Singla's thesis [53] touches upon several of these themes and has several pointers. Contention Resolution Schemes (CRSs) have found many applications since their introduction [26]; in fact Bayesian mechanism design, posted price mechanisms [7] and subsequent work by Yan [46], connecting mechanism design with the correlation gap, played an important role in [26].

Submodular functions and constraints such as cardinality (k-uniform matroids), matroids and others provide generality and computational tractability. It is possible to consider more general objective functions such as subadditive and monotone XOS functions, as well as more complex and general feasibility constraints. In such settings one can ignore computational considerations and focus on the online competitive ratio or assume access to a demand oracle (even though a demand oracle may be NP-Hard in general). We refer to [20], [77], [144] for some recent work and pointers.

For more information on other related work, see §1.4.

Organization. Section 5.2 introduces our notation and provides background on submodular functions, constraint systems and contention resolution schemes. Section 5.3 describes the relaxation of the prophet's objective. Section 5.4 describes the algorithm and analysis for SPI. Section 5.5 presents our refined correlation gap.

5.2 PRELIMINARIES

5.2.1 Background and Definitions

Let N be a finite ground set. Recall the definitions of a feasibility constraint, a downward-closed constraint and a submodular set function from §1.2. Classical examples of downward-closed constraints include those induced by a matroid on N or intersections of several matroids on N, independent sets of graphs, matchings in graphs and hypergraphs, boolean vectors that satisfy packing constraints of the form $A\boldsymbol{x} \leq b$ for non-negative A, b, among many others. We will use the terminology $(N, \mathcal{I})$, or simply $\mathcal{I}$ when the ground set is clear form context, to indicate a constraint family.

[6] Informally, a polytope $\mathcal{P}$ is solvable if one can efficiently do linear optimization over it. A formal definition is given in Section 5.2.

The maximum weight independent set problem over a constraint system $(N, \mathcal{I})$ is the following: given $w : N \to \mathbb{R}$ solve $\max_{S \in \mathcal{I}} w(S)$ where $w(S) = \sum_{e \in N} w(e)$. Since many of these problems are NP-Hard, a common technique is to use polyhedral (or more generally convex) relaxations. We say $\mathcal{P} \subseteq [0,1]^N$ is a *polyhedral relaxation* of $(N, \mathcal{I})$ if $\mathcal{P}$ is a polyhedron and $\mathbb{1}_S \in \mathcal{P}$ for all $S \in \mathcal{I}$ (here $\mathbb{1}_S$ is the characteristic vector of S). We say that $\mathcal{P}$ is *solvable* if one can efficiently do linear optimization over $\mathcal{P}$, that is, given $w : N \to \mathbb{R}$, there is a polynomial time algorithm that computes $\max_{\boldsymbol{x} \in \mathcal{P}} \sum_i w_i x_i$.

A set function f is *monotone* if $f(A) \leq f(B)$ for all $A \subseteq B$. In the rest of this chapter we work with non-negative normalized functions that satisfy $f(\emptyset) = 0$ and $f(A) \geq 0$ for all $A \subseteq N$. We often equate N with $[n] = \{1, 2, \ldots, n\}$. We use the terminology $S + i$ and $S - i$ as shorthands for $S \cup \{i\}$ and $S \setminus \{i\}$ respectively. The following continuous extensions of submodular functions to $[0,1]^N$ play an important role in our discussion.

Definition 5.1 (Multilinear Extension)**.** Let $f : \{0,1\}^N \to \mathbb{R}_{\geq 0}$. For any $\boldsymbol{x} \in [0,1]^n$, let $S \sim \boldsymbol{x}$ denote a random set S that contains each element $i \in N$ independently w.p. x_i. The *multilinear extension* of f is defined as

$$F(\boldsymbol{x}) \coloneqq \mathop{\mathbb{E}}_{S \sim \boldsymbol{x}}[f(S)] = \sum_{S \subseteq N} f(S) \prod_{i \in S} x_i \prod_{i \notin S} (1 - x_i).$$

It should be noted that via the multilinear relaxation, the polyhedral approach to approximation has been extended successfully to submodular function maximization [26], [45], [146].

Definition 5.2 (Concave Closure)**.** Let $f : \{0,1\}^N \to \mathbb{R}_{\geq 0}$. Moreover, let $\mathbb{1}_S$ denote the characteristic vector of a set $S \subseteq N$ of length $n = |N|$. For any $\boldsymbol{x} \in [0,1]^n$, the *concave closure* of f is defined as

$$f^+(\boldsymbol{x}) \coloneqq \max_{\boldsymbol{a} \in [0,1]^{2^N}} \left\{ \sum_{S \subseteq N} a_S f(S) \;\middle|\; \sum_{S \subseteq N} a_S = 1, \sum_{S \subseteq N} a_S \mathbb{1}_S = \boldsymbol{x} \right\}.$$

$f^+(\boldsymbol{x})$ can be interpreted as the maximum expected value of $f(R)$ where R is generated by a distribution whose marginal values are given by $\boldsymbol{x}$. Since $F(\boldsymbol{x})$ corresponds to the product distribution defined by $\boldsymbol{x}$, which is a specific distribution, it follows that $F(\boldsymbol{x}) \leq f^+(\boldsymbol{x})$ for all $\boldsymbol{x}$. The *correlation gap*, introduced in the work of Agrawal, Ding, Saberi and Ye [48], provides an inequality in the opposite direction.

Definition 5.3 (Correlation Gap)**.** Let $f : \{0,1\}^N \to \mathbb{R}_{\geq 0}$ be a set function and F, f^+ denote its multilinear relaxation and concave closure respectively. The correlation gap of f is defined as

$$\inf_{\boldsymbol{x} \in [0,1]^{|N|}} \frac{F(\boldsymbol{x})}{f^+(\boldsymbol{x})}.$$

It is easy to see that the correlation gap of modular/additive functions is 1. An important result in submodular optimization is that the correlation gap is at most $1 - 1/e$ for any monotone submodular function [45], [47], [48]. However, it is known that the correlation gap for general non-negative submodular functions (which can be non-monotone) can be arbitrarily small.

Contention Resolution Schemes: These are rounding schemes introduced in [26] for submodular function maximization. Recall the definitions of a Contention Resolution Scheme (CRS), an Online Contention Resolution Scheme (OCRS), a greedy OCRS and a c-selectable OCRS from §4.2. For the remainder of this chapter, $R(\boldsymbol{x})$ denotes a random subset of N where each $i \in N$ appears independently with probability x_i.

For a greedy OCRS, the quality of the approximation guaranteed with respect to the multilinear relaxation is governed by the notion of (b, c)-selectability [27].

Definition 5.4 ((b, c)-selectability)**.** Let $b, c \in [0, 1]$. A greedy OCRS for $\mathcal{P}$ is (b, c)-selectable if for any $\boldsymbol{x} \in b \cdot \mathcal{P}$, we have

$$\Pr\left[I \cup \{e\} \in \mathcal{F}_{\pi,\boldsymbol{x}} \quad \forall I \subseteq R(\boldsymbol{x}),\ I \in \mathcal{F}_{\pi,\boldsymbol{x}}\right] \geq c, \quad \forall e \in N.$$

We introduce the following notation, which will be useful in our analysis when dealing with the input constraints.

Definition 5.5 (Blowup of a Ground Set)**.** Let N denote a finite set, and N' denote another finite set, to be defined, with $|N| \leq |N'|$. Suppose for each $e \in N$ there is an associated finite non-empty set $A_e \subseteq N'$ such that the sets $A_e, e \in N$ are mutually disjoint. Let $\mathcal{A} = \{A_e \mid e \in N\}$. We call $N' = \bigcup_{e \in N} A_e$ the *blowup* of N by $\mathcal{A}$.

Definition 5.6 (Partition Extension of a Constraint)**.** Let $\mathcal{I} = (N, \mathcal{I})$ be a downward-closed constraint family over N. Consider a blowup N' of N induced by sets $A_e, e \in N$. Consider the function $g : N' \to N$ where $g(e') = e$ if and only if $e' \in A_e$. The *partition extension* of $\mathcal{I}$, denoted by $\mathcal{I}'$, is a constraint family $(N', \mathcal{I}')$ where

$$\mathcal{I}'_A = \{S \subseteq N' \mid g(S) \in \mathcal{I} \text{ and } \forall e \in N,\ |A_e \cap S| \leq 1\}.$$

A natural question for the prophet inequality setting is whether one can obtain better prophet inequalities when the arrival order of the random variables is chosen uniformly at random or even chosen by the algorithm. In [12], the authors introduce the *prophet secretary* model, combining the best of both the secretary and prophet inequality worlds. There has been much work on this model and we refer to [9], [49] for several interesting results in this and related models. We can consider the *Submodular Prophet Secretary (SPS)* problem as a generalization of the standard prophet secretary problem. We note that one can obtain improved guarantees in the SPS problem by using a Random Order CRS instead of an OCRS, since our results utilize the given OCRS in a black-box manner.

5.2.2 Useful Lemmas

Below we state several two lemmas regarding sampling and submodular functions that we need.

Lemma 5.1 (Lemma 2.2 from [147])**.** *Let $f : 2^N \to \mathbb{R}_{\geq 0}$ be submodular. Denote by $A(p)$ a random subset of A where each element appears with probability at most p (not necessarily independently). Then,*

$$\mathbb{E}[f(A(p))] \geq (1 - p) \cdot f(\emptyset).$$

Lemma 5.2 (Lemma 2.2 from [47])**.** *Let $g : 2^N \to \mathbb{R}_{\geq 0}$ be submodular. Denote by $A(p)$ a random subset of A where each element appears with probability exactly p (not necessarily independently). Then*

$$\mathbb{E}[g(A(p))] \geq (1 - p) \cdot g(\emptyset) + p \cdot g(A).$$

5.3 SUBMODULAR PROPHET INEQUALITIES

In the *Submodular Prophet Inequality (SPI)* problem, we are given n random variables $X_1, \ldots, X_n$ following (known) distributions $D_1, \ldots, D_n$, along with a constraint $\mathcal{I}$ on $N = \{1, 2, \ldots, n\}$. The random variables are arranged in adversarial (worst-case) order. Let $\mathcal{U}_i$ denote the image (range) of X_i, and $\mathcal{I}$ denote the feasible sets of $\mathcal{I}$.

The online algorithm starts with a set $S = \emptyset$ of selected elements and a set $Z = \emptyset$ of selected indices from N. At the i-th time step, it is presented with the realization $e \in \mathcal{U}_i$ of X_i. At that moment, it has to decide irrevocably whether to include e in S (and hence i in Z) or not, subject to Z remaining feasible in $\mathcal{I}$. The algorithm is also given a non-negative submodular function $f : \mathcal{U} \to \mathbb{R}_{\geq 0}$, where $\mathcal{U} \triangleq \bigcup_{i=1}^{n} \mathcal{U}_i$. The algorithm's objective is to maximize $f(S)$, subject to Z being feasible in $\mathcal{I}$.

In this model, we are comparing against the almighty adversary who already knows all realizations and can adaptively change the order in which to reveal the random variables based on the algorithm's actions so far and also the random coins it uses (if the algorithm is randomized). The prophet/adversary will select the best possible set S^* according to the constraints with knowledge of the realizations. Thus, we compare the expected value of the online algorithm against the expected value of the prophet, which is

$$\mathrm{OPT} = \mathop{\mathbb{E}}_{\boldsymbol{X}}\left[\max_{T \in \mathcal{I}} f\left(\{X_i \mid i \in T\}\right)\right]. \tag{5.1}$$

Later, we will use an OCRS to round the fractional solution we obtain in this section. Since f is defined over $\mathcal{U}$ but the constraint given is over N, we cannot immediately apply an OCRS for rounding. To overcome this issue, we view $\mathcal{U}$ as the blowup of N with respect to $\{\mathcal{U}_i\}_{i=1}^{n}$. On each step i, only one element arrives. Therefore, our input constraint $\mathcal{I}$ is equivalent to a new constraint $\mathcal{I}'$ on $\mathcal{U}$, where we are allowed to pick only one element from each $\mathcal{U}_i$. Notice that this is exactly the partition extension $\mathcal{I}' = (\mathcal{U}, \mathcal{I}')$ of $\mathcal{I}$.

We also denote $\boldsymbol{X} = \{X_1, X_2, \ldots, X_n\}$ and $\boldsymbol{D} = \{D_1, D_2, \ldots, D_n\}$. For an element $e \in \mathcal{U}_i$ we let $D_i(e)$ denote the probability of e being realized; we also use the notation $\mathcal{D}(e)$ to denote the probability of $e \in \mathcal{U}$ when we do not need to specify the part it belongs to. Note that the elements within $\mathcal{U}_i$ are correlated and hence we do not have a product distribution on $\mathcal{U}$.

Algorithmic approach: Following the description in Section 5.1, we design an online algorithm following the general approach of [20] but with technical differences. First, we obtain a relaxation of the prophet's objective. Afterwards, to design an online algorithm, we obtain an *offline* fractional point $\boldsymbol{z}$ based on the input, and round it *online* using a greedy OCRS and other tools. In this section, we describe the relaxation of the prophet's objective and how to obtain an offline fractional point $\boldsymbol{z}$. The process of rounding $\boldsymbol{z}$ online using a greedy OCRS is presented in Section 5.4.

Before we proceed, we describe a simple but technically important reduction that allows us to use improved correlation gaps, as well as obtain better bounds in the rounding algorithm.

Observation 5.1 (Reduction to small probabilities)**.** *Let $I = (N, \mathcal{U}, \boldsymbol{D}, \boldsymbol{X}, f, \mathcal{I})$ be an instance of the Submodular Prophet Inequality problem. For every fixed $\varepsilon > 0$, there is a reduction of I to another instance $I' = (N, \mathcal{U}', \boldsymbol{D'}, \boldsymbol{Y}, g, \mathcal{I})$ of the SPI problem such that that (i) for all $e \in \mathcal{U}'$, $\boldsymbol{D'}(e) \leq \varepsilon$ and (ii) there exists an α-competitive algorithm for I if and only if there exists an α-competitive algorithm for I'.*

Proof Sketch. Consider the original instance I and recall that each $\mathcal{D}_i$ is a probability distribution over $\mathcal{U}_i$. Our goal is to ensure that $\mathcal{D}_i(e) \leq \varepsilon$ for every $e \in \mathcal{U}_i$. Suppose there is an element e such that $\mathcal{D}_i(e) > \varepsilon$. We obtain a new instance I' as follows. We replace $e \in \mathcal{U}_i$ by $h = \lceil 1/\varepsilon \rceil$ "copies" $e_1, e_2, \ldots, e_h$; let S_e denote this set of copies. Let $\mathcal{U}_i'$ be the new set of elements. We obtain a probability distribution $\mathcal{D}_i' : \mathcal{U}_i' \to [0, 1]$ as follows. If $e' \in \mathcal{U}_i$ such that $e' \neq e$ then $\mathcal{D}_i'(e') = \mathcal{D}_i(e')$ (nothing changes for e'). For each copy e_j

of e we set $\mathcal{D}'_i(e_j) = \mathcal{D}_i(e)/h$ and by our choice of h we have $\mathcal{D}'_i(e_j) \leq 1/h \leq \varepsilon$, for all $e_j \in S_e$. Thus, $\sum_{j=1}^h \mathcal{D}'_i(e_j) = \mathcal{D}_i(e)$. Since we replaced e by h copies of it, the ground set $\mathcal{U}$ changes to $\mathcal{U}'$ and we now define a new submodular function $g : \mathcal{U}' \to \mathbb{R}_+$ that is derived from f. The function g treats the copies of e as a "single" element and hence mimics f. More formally, for any $A \subseteq \mathcal{U}'$: $g(A) = f(A)$ if $A \cap S_e = \emptyset$, else $g(A) = f((A \setminus S_e) \cup \{e\})$. It is easy to verify that if f is non-negative and submodular, then g is also non-negative and submodular, and also inherits monotonicity from f. Let I' be the resulting modified instance. We observe that in I', the probability of an element from S_e being chosen is precisely equal to $\mathcal{D}_i(e)$ and hence the copies of e act as proxies for e and the submodular function g ensures that every copy behaves the same as e in f. Note that we crucially relied on the power of submodularity in this reduction. One can apply this reduction repeatedly to reduce all realization probabilities to at most ε. One also notices that the reduction is computationally efficient as a function of ε. For any fixed ε, the size of I' is at most $O(1/\varepsilon)$ times the size of I and a value oracle for f can be used to efficiently and easily obtain a value oracle for the new submodular function g. QED

Remark 5.1. The reduction's simplicity may make the reader wonder why it is useful in achieving improved bounds. The reason is a combination of the model and the power of submodularity. The fact that we can only pick a single element from each $\mathcal{U}_i$ allows us to make copies of the elements, and we can use a derived submodular function to treat the copies as a single element.

5.3.1 An Upper Bound on the Prophet's Value

Let $\mathcal{P}$ denote a solvable polyhedral relaxation of $\mathcal{I}$. Then one can easily develop a solvable polyhedral relaxation of $\mathcal{I}'$ as follows:

$$\mathcal{P}' = \left\{ \boldsymbol{y} \in [0,1]^{|\mathcal{U}|} \;\middle|\; \sum_{e \in \mathcal{U}_i} y_e = x_i, \;\; i \in [n], \boldsymbol{x} \in \mathcal{P} \right\}. \tag{5.2}$$

Consider any algorithm, including an offline algorithm, that computes a feasible output given the realizations of the random variables. For any fixed algorithm $\mathcal{A}$ (deterministic or randomized) we have a probability $p_{\mathcal{A}}(e)$ for each $e \in \mathcal{U}$ appearing in the output of $\mathcal{A}$. Since an element $e \in \mathcal{U}$ is realized with probability $\mathcal{D}(e)$, e cannot appear in the output of $\mathcal{A}$ with probability more than $\mathcal{D}(e)$. Moreover, for a given realization, each output of the algorithm is feasible. Putting these facts together we obtain the following observation.

Observation 5.2. *Let $\mathcal{A}$ be any online or offline algorithm for a given instance of the problem. Let $p_{\mathcal{A}}(e)$ denote the probability that e is in the output of $\mathcal{A}$. Then the vector $\boldsymbol{p}$ is in the polytope*

$$\mathcal{P}'' = \left\{ \boldsymbol{z} \in [0,1]^{|\mathcal{U}|} \;\middle|\; \boldsymbol{z} \in \mathcal{P}', z_e \leq \mathcal{D}(e) \; e \in \mathcal{U} \right\}.$$

We are now ready to proceed with the relaxation of the prophet's objective.

Claim 5.1. *Consider an instance of the Submodular Prophet Inequality problem. Then*

$$\max_{\boldsymbol{z} \in \mathcal{P}''} f^+(\boldsymbol{z}) \geq OPT.$$

Proof. Fix an optimal strategy for the prophet and let $\boldsymbol{y}^* \in [0,1]^{|\mathcal{U}|}$ denote the vector of probabilities of the elements appearing in the output of the prophet's strategy. We have $\boldsymbol{y}^* \in \mathcal{P}''$. By the definition of the concave closure of f, $f^+(\boldsymbol{y}^*)$ maximizes the value of f among all distributions with the marginals $\boldsymbol{y}^*$ (note that the distribution that achieves this may not be a feasible strategy for any algorithm). Therefore, $f^+(\boldsymbol{y}^*) \geq \mathrm{OPT}$, which also implies that $\max_{\boldsymbol{z}\in\mathcal{P}''} f^+(\boldsymbol{z}) \geq \mathrm{OPT}$. QED

5.3.2 Fractional Solution and Correlation Gap

From Claim 5.1, $\max_{\boldsymbol{z}\in\mathcal{P}''} f^+(\boldsymbol{z}) \geq \mathrm{OPT}$. Since OCRSs are designed to relate the quality of their output to that of the multilinear relaxation, we need to relate $F(\boldsymbol{z})$ to $f^+(\boldsymbol{z})$ and hence to OPT. We present two different ways to do this — via a direct correlation gap and via the Measured Continuous Greedy (MCG) algorithm — with the second yielding strictly better results than the first.

The direct correlation gap approach The first approach is not computationally efficient and relies on optimally solving the optimization problem $\max_{\boldsymbol{z}\in\mathcal{P}''} f^+(\boldsymbol{z})$. Let $\boldsymbol{z}^*$ be the optimum solution. We can then use the correlation gap to relate $F(\boldsymbol{z}^*)$ to OPT. For monotone functions we have $F(\boldsymbol{z}^*) \geq (1-1/e)f^+(\boldsymbol{z}^*) \geq (1-1/e)\mathrm{OPT}$. For non-negative functions we can use Theorem 5.2. Following the reduction that we described earlier, we can assume that $z_e^* \leq \max_e \boldsymbol{D}(e) \leq \varepsilon$ for all e and this implies, via Theorem 5.2 that $F(\boldsymbol{z}^*) \geq (1-\varepsilon)(1-1/e)f^+(\boldsymbol{z}^*) \geq (1-\varepsilon)(1-1/e)\mathrm{OPT}$. In rounding it is useful to have a solution $\boldsymbol{z} \in b\cdot\mathcal{P}''$ for some parameter $b \in (0,1)$. One can of course use $\boldsymbol{z} = b\boldsymbol{z}^*$ and in this case, we can use the concavity of f^+ to see that $f^+(b\boldsymbol{z}^*) \geq bf^+(\boldsymbol{z}^*)$, and then apply the correlation gap to $b\boldsymbol{z}^*$ to conclude that, in the monotone case, $F(b\boldsymbol{z}^*) \geq b(1-1/e)f^+(\boldsymbol{z}^*) \geq b(1-1/e)\mathrm{OPT}$ and, in the non-monotone case, $F(b\boldsymbol{z}^*) \geq b(1-\varepsilon)(1-1/e)f^+(\boldsymbol{z}^*) \geq b(1-\varepsilon)(1-1/e)\mathrm{OPT}$.

The measured continuous greedy approach The second approach is algorithmic and relies on the Measured Continuous Greedy (MCG) algorithm and its properties. We state two relevant known results about the algorithm. For these results as well as Theorem 5.4, we assume the submodular function f is given via a value oracle, and that the algorithms are randomized and run in polynomial time and are correct with high probability.

Lemma 5.3 (Lemma 4 of [83])**.** *Let f be a monotone submodular function with multilinear extension F, and let $\mathcal{P}$ be a solvable downward-closed polytope. Let $\boldsymbol{x}(b)$ be solution produced by the Continuous Greedy algorithm on F and $\mathcal{P}$ until time $b \in (0,1]$. Then (i) $\boldsymbol{x}(b) \in b\cdot\mathcal{P}$ and (ii) $F(\boldsymbol{x}(b)) \geq \big(1-e^{-b}-o(1)\big)\cdot\max_{\boldsymbol{y}\in\mathcal{P}} f^+(\boldsymbol{y})$.*

For a general non-negative submodular function, the MCG algorithm achieves the following bound.

Lemma 5.4 (Lemma 8.3 of [28])**.** *Let $b \in [0,1]$, f be a non-negative submodular function with multilinear extension F, and let $\mathcal{P}$ be a solvable downward-closed polytope. Then, the solution $\boldsymbol{x}(b) \in [0,1]^n$ produced by the MCG algorithm satisfies (i) $\boldsymbol{x}(b) \in b\cdot\mathcal{P}$ and (ii) $F(\boldsymbol{x}(b)) \geq \big(b\cdot e^{-b}-\varepsilon\big)\cdot\max_{\boldsymbol{y}\in\mathcal{P}} f^+(\boldsymbol{y})$, for any fixed $\varepsilon > 0$.*

The two preceding lemmas are algorithmic. If $\mathcal{P}$ is solvable then the underlying algorithms can be implemented efficiently. Based on our reduction to small probabilities it is useful to consider whether the preceding lemmas can take advantage of this. No advantage is possible in the monotone setting, however, we show below that one can indeed take advantage of the reduction when f is non-monotone. We provide a refined analysis of the standard bound of the MCG algorithm, which depends on a parameter p that quantifies the maximum value of any coordinate that is feasible in the polytope. For small enough p, Theorem 5.4 constitutes an improvement over Lemma 5.4, which comprises the main result of this section. Notice that Theorem 5.3 follows from Theorem 5.4 by setting $b = 1$.

Theorem 5.4
Let $p \in [0, 1)$, f be a non-negative submodular function with multilinear extension F and $\mathcal{P}$ be a downward-closed solvable polytope on N, such that $\mathcal{P} \subseteq p \cdot [0, 1]^N$ (that is, if $\boldsymbol{z} \in \mathcal{P}$ then $z_i \leq p$ for all $i \in N$). Then, the output of the Measured Continuous Greedy (MCG) algorithm on F and $\mathcal{P}$ at time $b \in [0, 1]$ is a vector $\boldsymbol{x}(b) \in b \cdot \mathcal{P}$ such that

$$F(\boldsymbol{x}(b)) \geq \begin{cases} b \cdot e^{-b} \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^+(\boldsymbol{z}), & 0 \leq b \leq \ln\left(\frac{1}{1-p}\right) \\ \left(1 - p - e^{-b}\left(1 + \ln(1-p)\right)\right) \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^+(\boldsymbol{z}), & \ln\left(\frac{1}{1-p}\right) \leq b \leq 1. \end{cases}$$

Remark 5.2. Notice that, for the SPI problem, due to our reduction, we can assume that all vectors $\boldsymbol{z} \in \mathcal{P}''$ have $z_i \leq \varepsilon'$ for all $i \in N$, for any fixed constant $\varepsilon' > 0$. Therefore, for any fixed constant $\varepsilon > 0$, there exists an ε' such that

$$F(\boldsymbol{x}(b)) \geq \left(1 - e^{-b} - \varepsilon\right) \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^+(\boldsymbol{z}),$$

where $\boldsymbol{x}(b) \in b \cdot \mathcal{P}''$ is the output of the MCG algorithm at time b.

Before we proceed with the proof of Theorem 5.4, we briefly sketch the idea implicit in prior work [28], [145] that implies $F_{\max}(\boldsymbol{x}) \geq \frac{1}{e} f^+(\boldsymbol{x})$. Consider a downward-closed polytope $\mathcal{P}$ defined by all points in $[0, 1]^n$ dominated by the given point $\boldsymbol{x}$: $\mathcal{P} \coloneqq \{\boldsymbol{y} \in [0, 1]^n \mid \forall\, 1 \leq i \leq n,\ y_i \leq x_i\}$. Suppose we run the MCG algorithm on $\mathcal{P}$. From Lemma 8.3 of [28] for $b = 1$, for any $\varepsilon > 0$, the algorithm can be used to find a point $\boldsymbol{z}_\varepsilon \in \mathcal{P}$ such that $F(\boldsymbol{z}_\varepsilon) \geq \left(\frac{1}{e} - \varepsilon\right) \max_{\boldsymbol{y} \in \mathcal{P}} f^+(\boldsymbol{y}) \geq \left(\frac{1}{e} - \varepsilon\right) f^+(\boldsymbol{x})$. Since such a point $\boldsymbol{z}_\varepsilon \in \mathcal{P}$ exists for any $\varepsilon > 0$, by the compactness of $\mathcal{P}$ and the continuity of F and f^+, it follows that there exists a point $\boldsymbol{y} \in \mathcal{P}$ such that $F(\boldsymbol{y}) \geq \frac{1}{e} \cdot f^+(\boldsymbol{x})$. Also notice that $\boldsymbol{x} \in \mathcal{P}$, and thus

$$F_{\max}(\boldsymbol{x}) = \max_{\boldsymbol{z} \in \mathcal{P}} F(\boldsymbol{z}) \geq F(\boldsymbol{y}) \geq \frac{1}{e} \cdot f^+(\boldsymbol{x}). \tag{5.3}$$

To prove Theorem 5.4, which generalizes Lemma 8.3 in [28], we use the same proof outline as above, but in the algorithm's analysis, we take advantage of the fact that $\|\boldsymbol{x}\|_\infty \leq p$.

Proof of Theorem 5.4. Let $\hat{\boldsymbol{x}} = \arg\max_{\boldsymbol{z} \in \mathcal{P}} f^+(\boldsymbol{z})$. Recall that there exists $\boldsymbol{\alpha} \in [0, 1]^{2^N}$ such that

$$f^+(\hat{\boldsymbol{x}}) = \sum_{S \subseteq N} \alpha_S f(S), \quad \sum_{S \subseteq N} \alpha_S = 1 \quad \text{and} \quad \sum_{S \subseteq N} \alpha_S \mathbb{1}_S = \hat{\boldsymbol{x}}. \tag{5.4}$$

From the analysis of Measured Continuous Greedy and the fact that $\boldsymbol{x}(b) \in \mathcal{P}$, we know that, at time b, for all $i \in N$ we have

$$x_i(b) \leq \min\{1 - e^{-b}, p\}. \tag{5.5}$$

Let $\boldsymbol{x} = \boldsymbol{x}(b)$, and, for $S \subseteq N$, consider a line of direction $\boldsymbol{d}_S = (\boldsymbol{x} \vee \mathbb{1}_S) - \boldsymbol{x} = (\mathbb{1}_S - \boldsymbol{x}) \vee \mathbf{0}$. Notice that $\mathbf{0} \leq \boldsymbol{d}_S \leq \mathbb{1}_S$ for all $S \subseteq N$. From Section 2.3 of [45], it follows that

$$\boldsymbol{d}_S \cdot \nabla F(\boldsymbol{x}) \geq F(\boldsymbol{x} \vee \mathbb{1}_S) - F(\boldsymbol{x}). \tag{5.6}$$

Since f may not be monotone, $\nabla F(\boldsymbol{x})$ may have negative entries. Let $\boldsymbol{d}'_S$ be a vector obtained from $\boldsymbol{d}_S$ as follows: $(\boldsymbol{d}'_S)_i = (\boldsymbol{d}_S)_i$ if $\nabla F(\boldsymbol{x})_i \geq 0$, otherwise $(\boldsymbol{d}'_S)_i = 0$. We have $\mathbf{0} \leq \boldsymbol{d}'_S \leq \boldsymbol{d}_S$ and,

$$\boldsymbol{d}'_S \cdot \nabla F(\boldsymbol{x}) \geq \max\{0, \boldsymbol{d}_S \cdot \nabla F(\boldsymbol{x})\} \geq \max\{0, F(\boldsymbol{x} \vee \mathbb{1}_S) - F(\boldsymbol{x})\}. \tag{5.7}$$

Since $\boldsymbol{x}(b)_i \leq \min\{1 - e^{-b}, p\}$ for all $i \in N$, by Lemma III.5 of [145], we have

$$F(\boldsymbol{x} \vee \mathbb{1}_S) \geq \left(1 - \min\{1 - e^{-b}, p\}\right) f(S). \tag{5.8}$$

Therefore,

$$\begin{aligned}
\boldsymbol{d}'_S \cdot \nabla F(\boldsymbol{x}) &\geq \max\{0, (1-p)f(S) - F(\boldsymbol{x}), e^{-b} f(S) - F(\boldsymbol{x})\} && (5.9)\\
&\geq \max\{1 - p, e^{-b}\} f(S) - F(\boldsymbol{x}). && (5.10)
\end{aligned}$$

Next, let $\hat{\boldsymbol{d}} = \sum_{S \subseteq N} \alpha_S \boldsymbol{d}'_S$. Since $\boldsymbol{d}_S \leq \mathbb{1}_S$ and $\boldsymbol{d}'_S \leq \boldsymbol{d}_S$, we have $\boldsymbol{d}'_S \leq \mathbb{1}_S$, and thus

$$\hat{\boldsymbol{d}} = \sum_{S \subseteq N} \alpha_S \boldsymbol{d}'_S \leq \sum_{S \subseteq N} \alpha_S \mathbb{1}_S = \hat{\boldsymbol{x}}. \tag{5.11}$$

Since $\mathcal{P}$ is downward-closed and $\hat{\boldsymbol{x}} \in \mathcal{P}$, we know that $\hat{\boldsymbol{d}} \in \mathcal{P}$. Therefore, from the above and the fact that $\boldsymbol{v}_{\max} = \arg\max_{\boldsymbol{v} \in \mathcal{P}} \boldsymbol{v} \cdot \nabla F(\boldsymbol{x})$, we have

$$\begin{aligned}
\frac{dF(\boldsymbol{x}(b))}{db} &= \boldsymbol{v}_{\max}(\boldsymbol{x}) \cdot \nabla F(\boldsymbol{x}) && (5.12)\\
&\geq \hat{\boldsymbol{d}}_S \cdot \nabla F(\boldsymbol{x}) && (5.13)\\
&= \sum_{S \subseteq N} \alpha_S \cdot \boldsymbol{d}'_S \cdot \nabla F(\boldsymbol{x}) && (5.14)\\
&\geq \sum_{S \subseteq N} \alpha_S \left(\max\{1 - p, e^{-b}\} f(S) - F(\boldsymbol{x})\right) && (5.15)\\
&\geq \max\{1 - p, e^{-b}\} \sum_{S \subseteq N} \alpha_S f(S) - \sum_{S \subseteq N} \alpha_S F(\boldsymbol{x}) && (5.16)\\
&\geq \max\{1 - p, e^{-b}\} f^+(\hat{\boldsymbol{x}}) - F(\boldsymbol{x}). && (5.17)
\end{aligned}$$

We proceed to solve the above differential inequality. For brevity, let $y = F(\boldsymbol{x})$. Then,

$$\begin{aligned}
dy + y\, db &\geq f^+(\hat{\boldsymbol{x}}) \max\{1 - p, e^{-b}\}\, db && (5.18)\\
e^b\, dy + y e^b\, db &\geq f^+(\hat{\boldsymbol{x}}) \max\{(1-p)e^b, 1\}\, db && (5.19)\\
d\left(y e^b\right) &\geq f^+(\hat{\boldsymbol{x}}) \max\{(1-p)e^b, 1\}\, db && (5.20)\\
y &\geq e^{-b} f^+(\hat{\boldsymbol{x}}) \int_0^b \max\{(1-p)e^u, 1\}\, du. && (5.21)
\end{aligned}$$

Notice that, for $0 \leq u \leq \ln\left(\frac{1}{1-p}\right)$, we have $(1-p)e^u \leq 1$, while for $\ln\left(\frac{1}{1-p}\right) \leq u \leq 1$, we have $1 \leq (1-p)e^u$.

Therefore, for $b \leq \ln\left(\frac{1}{1-p}\right)$, (5.21) becomes

$$y \geq e^{-b} f^{+}(\hat{\boldsymbol{x}}) \int_{0}^{b} 1\, du = b \cdot e^{-b} \cdot f^{+}(\hat{\boldsymbol{x}}), \tag{5.22}$$

whereas for $b \geq \ln\left(\frac{1}{1-p}\right)$, (5.21) becomes

$$\begin{aligned} y &\geq e^{-b} f^{+}(\hat{\boldsymbol{x}}) \left(\int_{0}^{\ln\left(\frac{1}{1-p}\right)} 1\, du + \int_{\ln\left(\frac{1}{1-p}\right)}^{b} (1-p) e^{u}\, du \right) & (5.23) \\ &= \left(1 - p - e^{-b}\left(1 + \ln\left(1-p\right)\right)\right) f^{+}(\hat{\boldsymbol{x}}). & (5.24) \end{aligned}$$

We conclude that

$$F(\boldsymbol{x}(b)) \geq \begin{cases} b \cdot e^{-b} \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^{+}(\boldsymbol{z}), & 0 \leq b \leq \ln\left(\frac{1}{1-p}\right) \\ \left(1 - p - e^{-b}\left(1 + \ln\left(1-p\right)\right)\right) \cdot \max_{\boldsymbol{z} \in \mathcal{P}} f^{+}(\boldsymbol{z}), & \ln\left(\frac{1}{1-p}\right) \leq b \leq 1. \end{cases} \tag{5.25}$$

QED

We summarize the results via both methods below. We observe that for both monotone and non-monotone functions the bounds are best when $p \to 0$, which we can ensure via the reduction. Once we make this assumption, the bounds provided by the correlation gap approach are essentially $(1 - 1/e)$ when $b = 1$ which is optimal. However, these bounds are matched by the Continuous Greedy approach. When $b < 1$, which will be the case when applying the rounding schemes, the bound in Lemma 5.3 and our new refined bound in Theorem 5.4 are superior and have the further advantage of being computable in polynomial time.

5.4 ROUNDING THE FRACTIONAL SOLUTION

In the preceding section we described ways to obtain a vector $\boldsymbol{z} \in b \cdot \mathcal{P}''$ for some $b \in [0, 1]$ such that $F(\boldsymbol{z}) \geq \alpha \cdot OPT$ for various constants α depending on the approach. In this section we show how to round $\boldsymbol{z}$ in an online fashion. We follow the high-level approach of [20] but refine it in several ways. We will use a greedy OCRS for $\mathcal{I}$ via the relaxation $\mathcal{P}$ as a black box. Recall that our rounding needs to produce a feasible set in $\mathcal{I}'$ with ground set $\mathcal{U}$, while the OCRS is for the constraint on the variables of N. Moreover the distribution $\boldsymbol{D}$ is *not* a product distribution on $\boldsymbol{\mathcal{U}}$. These are the technical challenges that need to be overcome in the algorithm and analysis. The quality of the output will depend on the properties of the OCRS for $\mathcal{P}$. We assume that the greedy OCRS for $\mathcal{P}$ is (b, c)-selectable, where c is some function of b. This depends on the specific constraint family $\mathcal{I}$ and the polyhedral relaxation $\mathcal{P}$. At the end of the section, we use known results to derive concrete competitive ratios for several constraint families of interest. We note that $\boldsymbol{z} \in \mathcal{P}''$, which also implies that $\boldsymbol{z} \in \mathcal{P}'$. For rounding purposes we only work with $\mathcal{P}'$ and $\mathcal{P}$; $\mathcal{P}''$ is only necessary to obtain an upper bound on OPT.

We rely on the certain parts of the analysis of OCRS for submodular function maximization from [27]. In the following, we will use π to denote the mapping function for the OCRS over the ground set N and the polytope $\mathcal{P}$. Technically the mapping π is a function of $\boldsymbol{x} \in \mathcal{P}$ and should be written as $\pi_{\boldsymbol{x}}$ but we omit $\boldsymbol{x}$ for notational simplicity. We also note that π can be randomized. An important definition from [27] in the analysis of OCRSs is the characteristic CRS of a greedy OCRS.

Definition 5.7 (Characteristic CRS of a greedy OCRS)**.** The *characteristic CRS* $\bar{\pi}$ of a greedy OCRS π for a polytope $\mathcal{P}$ is a CRS for the same polytope $\mathcal{P}$. It is defined for an input $x \in \mathcal{P}$ and a set $A \subseteq N$ by $\bar{\pi}(A) = \{e \in A \mid I \cup \{e\} \in \mathcal{F}_{\pi,\boldsymbol{x}},\ \forall I \subseteq A,\, I \in \mathcal{F}_{\pi,\boldsymbol{x}}\}$. Notice that, if π is randomized, then $\bar{\pi}$ is randomized as well.

We will also need the following known results from [27].

Observation 5.3 (Observation 3.3 of [27])**.** *For every set $A \subseteq N$ and a characteristic CRS $\bar{\pi}$ of a greedy OCRS π, the set $\bar{\pi}(A)$ is always a subset of the elements selected by π when the active elements are the elements of A.*

Lemma 5.5 (Lemma 3.4 of [27])**.** *The characteristic CRS $\bar{\pi}$ of a (b, c)-selectable greedy OCRS π is (b, c)-balanced and monotone.*

For any $S \subseteq \mathcal{U}$, we define $S_{\downarrow} \subseteq N$ to be the projection of S onto N, i.e.

$$S_{\downarrow} \coloneqq \{i \in N \mid S \cap \mathcal{U}_i \neq \emptyset\}\,. \tag{5.26}$$

Also, for a greedy OCRS π, we denote the characteristic CRS of π by $\bar{\pi}$. We now define a CRS π' for $\mathcal{P}'$ that we will need for our analysis later on. We define π' using the characteristic CRS $\bar{\pi}$ of π as follows. For any set $S \subseteq \mathcal{U}$,

$$\pi'(S) \coloneqq \bigcup_{\substack{i \in \bar{\pi}(S_{\downarrow}) \\ |S \cap \mathcal{U}_i| = 1}} (S \cap \mathcal{U}_i). \tag{5.27}$$

Intuitively, the characteristic CRS $\bar{\pi}$ of a greedy OCRS π returns, on input $A \subseteq N$ the set of all elements in A that are in $\pi(A)$ regardless of the arrival order chosen by the adversary. Given a set $S \subseteq \mathcal{U}$, $\pi'(S)$ is equal to the union of at most one element from each $\mathcal{U}_i$, for all such i that are in the projection $S_{\downarrow}$ of S and are selected by $\bar{\pi}$ on input $S_{\downarrow}$. The next lemma relates the balance guarantee provided by π' given a selectability guarantee on π.

Lemma 5.6. *For any (b, c)-selectable greedy OCRS π for $\mathcal{P}$ and $\boldsymbol{z} \in \mathcal{P}'$, the CRS π' is monotone and $(b, c \cdot \gamma)$-balanced, where $\gamma = \min_{i \in N} \prod_{e \in \mathcal{U}_i} (1 - z_e)$.*

Proof. First, notice that π' is a CRS, since $\pi'(S) \subseteq S$ for all $S \subseteq \mathcal{U}$. This follows immediately from the definition of π' as $S \cap \mathcal{U}_i \subseteq S$ for all $i \in N, S \subseteq \mathcal{U}$.

Next, we show that π' is monotone. Fix an element $e \in S_1 \subseteq S_2 \subseteq \mathcal{U}$, and an instantiation of $\mathcal{F}_{\pi,\boldsymbol{x}}$ (this is relevant if the OCRS is randomized). Let $e \in \mathcal{U}_i$ for some $i \in N$. Suppose $e \in \pi'(S_2)$. This implies that $|S_2 \cap \mathcal{U}_i| = 1$ and since $S_1 \subseteq S_2$ and $e \in S_1$, we have $|S_1 \cap \mathcal{U}_i| = 1$. Furthermore, we know that $i \in \bar{\pi}(S_{2\downarrow})$. Since $S_1 \subseteq S_2$, it follows that $S_{1\downarrow} \subseteq S_{2\downarrow}$. By Lemma 5.5, we know that $\bar{\pi}$ is monotone, and thus, since $i \in \bar{\pi}(S_{2\downarrow})$, it follows that $i \in \bar{\pi}(S_{1\downarrow})$. Therefore, we know that $e \in \pi'(S_1)$. Since $e \in \pi'(S_2)$ implies that $e \in \pi'(S_1)$, unconditioning over the instantiation of $\mathcal{F}_{\pi,\boldsymbol{x}}$ yields

$$\Pr\left[e \in \pi'(S_1)\right] \geq \Pr\left[e \in \pi'(S_2)\right]. \tag{5.28}$$

We now show that π' is $(b, c \cdot \gamma)$-balanced, for $\gamma = \min_{i \in N} \prod_{e \in \mathcal{U}_i} (1 - z_e)$. It suffices to show that, for any $e \in \mathcal{U}$

$$\Pr_{S \sim R(\boldsymbol{z})}\left[e \in \pi'(S) \mid e \in S\right] \geq c \cdot \gamma. \tag{5.29}$$

Notice that, for any realization S of $R(\boldsymbol{z})$, $e \in \pi'(S)$ if and only if $S \cap \mathcal{U}_i = \{e\}$ and $i \in \bar{\pi}(S_\downarrow)$. Thus,

$$\Pr_{S\sim R(\boldsymbol{z})}[e \in \pi'(S) \mid e \in S] = \Pr_{S\sim R(\boldsymbol{z})}[S \cap \mathcal{U}_i = \{e\} \wedge i \in \bar{\pi}(S_\downarrow) \mid e \in S] \tag{5.30}$$

$$= \Pr_{S\sim R(\boldsymbol{z})}[S \cap \mathcal{U}_i = \{e\} \mid e \in S] \tag{5.31}$$

$$\cdot \Pr_{S\sim R(\boldsymbol{z})}[i \in \bar{\pi}(S_\downarrow) \mid S \cap \mathcal{U}_i = \{e\}, e \in S] \tag{5.32}$$

$$= \Pr_{S\sim R(\boldsymbol{z})}[S \cap \mathcal{U}_i = \{e\} \mid e \in S] \tag{5.33}$$

$$\cdot \Pr_{S\sim R(\boldsymbol{z})}[i \in \bar{\pi}(S_\downarrow) \mid S \cap \mathcal{U}_i = \{e\}], \tag{5.34}$$

where the last equality follows from the fact that, if $S \cap \mathcal{U}_i = \{e\}$, then $e \in S$. We lower bound each probability in (5.34) separately, starting from

$$\Pr_{S\sim R(\boldsymbol{z})}[S \cap \mathcal{U}_i = \{e\} \mid e \in S] = \prod_{e'\neq e, e'\in\mathcal{U}_i}(1 - z_{e'}) \geq \prod_{e'\in\mathcal{U}_i}(1 - z_{e'}) \geq \gamma. \tag{5.35}$$

Also, notice that $\bar{\pi}$ is a CRS over N and does not depend on which $S \cap \mathcal{U}_i$ led to $i \in S_\downarrow$. Therefore,

$$\Pr[i \in \bar{\pi}(S_\downarrow) \mid i \in S_\downarrow] = \Pr[i \in \bar{\pi} \mid S \cap \mathcal{U}_i = T] \tag{5.36}$$

for all $T \subseteq \mathcal{U}_i$ such that $T \neq \emptyset$. Specifically, for $T = \{e\}$,

$$\Pr[i \in \bar{\pi} \mid S \cap \mathcal{U}_i = \{e\}] = \Pr[i \in \bar{\pi}(S_\downarrow) \mid i \in S_\downarrow] \geq c, \tag{5.37}$$

where the last inequality follows from the fact that fact that $\bar{\pi}$ is (b, c)-balanced, by Lemma 5.5.

Combining (5.34), (5.35) and (5.37), we obtain

$$\Pr_{S\sim R(\boldsymbol{z})}[e \in \pi'(S) \mid e \in S] \geq c \cdot \gamma. \tag{5.38}$$

QED

Remark 5.3. Notice that via Observation 5.1, we can assume without loss of generality that, for any fixed $\varepsilon' > 0$, $z_e \leq \varepsilon'$ for all $e \in \mathcal{U}$. By choosing ε' sufficiently small, for any fixed $\varepsilon > 0$ we have

$$\gamma = \min_{i\in N}\prod_{e\in\mathcal{U}_i}(1 - z_e) \geq \min_{i\in N}\left(\prod_{e\in\mathcal{U}_i} e^{-z_e}\right) - \varepsilon = \min_{i\in N}\left(e^{-\sum_{e\in\mathcal{U}_i} z_e}\right) - \varepsilon \geq e^{-b} - \varepsilon,$$

where the last inequality follows from the fact that $\boldsymbol{z} \in b \cdot \mathcal{P}'$. Thus, $c \cdot \gamma \geq c \cdot (e^{-b} - \varepsilon)$, and we obtain the following as corollary: For any (b, c)-selectable greedy OCRS π for $\mathcal{P}$ and fixed $\varepsilon > 0$, π' defined earlier is a $\left(b, c\left(e^{-b} - \varepsilon\right)\right)$-balanced monotone CRS for $\mathcal{P}'$.

Now we are ready to describe our online algorithm. We describe and analyze the algorithms for monotone and non-monotone cases separately, since there are technical differences. The algorithms are similar to the one in [20], however, the main technical difference is that we use the OCRS for N as a black box; in [20] the authors use an OCRS over $\mathcal{U}$ since they work in the special case of matroids.

5.4.1 Monotone Non-Negative Submodular Functions

We assume we have already computed a vector $\boldsymbol{z} \in b \cdot \mathcal{P}''$ for some $b \in [0,1]$ such that $F(\boldsymbol{z}) \geq \alpha \cdot \mathrm{OPT}$ for some α. Note that the adversary is almighty and can alter the order in which it feeds the variables to the algorithm based on knowledge of the full realizations of the variables and the actions of the algorithm so far.

Let $\boldsymbol{z_i}$ denote the product distribution on $\mathcal{U}_i$ defined by marginals $z_i(e), e \in \mathcal{U}_i$. We write $R \sim \boldsymbol{z_i}$ to denote a random set $R \subseteq \mathcal{U}_i$ realized according to this product distribution, and we denote $z_i(e)$ by z_e when i is clear from context or irrelevant. Furthermore, let $\boldsymbol{x} \in [0,1]^n$ be the vector where $x_i = \Pr_{R \sim \boldsymbol{z_i}}[R \neq \emptyset] = 1 - \prod_{e \in \mathcal{U}_i}(1 - z_e)$, for all $i \in N$. We assume that $\boldsymbol{x}$ is the input vector to our OCRS π_x for $\mathcal{P}$ and its characteristic CRS $\bar{\pi}_{\boldsymbol{x}}$. To simplify our notation, we denote $\pi_{\boldsymbol{x}}$ and $\bar{\pi}_{\boldsymbol{x}}$ by π and $\bar{\pi}$, respectively.

Algorithm 5.1: MONOTONE ROUNDING $(\mathcal{U}, f, \boldsymbol{D}, \mathcal{I}, \pi, \boldsymbol{z})$

$T_{\mathrm{ALG}} = \emptyset$
for $h \leftarrow 1$ **to** n **do**
  Let X_i be variable that arrives on step h
  Let $e \in \mathcal{U}_i$ be the realization of X_i
  With probability $\frac{\Pr_{R \sim \boldsymbol{z_i}}[R = \{e\}]}{\mathcal{D}_i(e)}$, set $T_i \leftarrow \{e\}$
  Otherwise, set T_i to be a random subset R of $\mathcal{U}_i$, drawn according to $\boldsymbol{z_i}$, conditioned on $|R| \neq 1$
  if $T_i \neq \emptyset$ **then**
    Feed i to OCRS π for $\mathcal{P}$
    if π *accepts* i and $T_i = \{e\}$ **then**
      $T_{\mathrm{ALG}} \leftarrow T_{\mathrm{ALG}} \cup \{e\}$
    end
  end
end
Return T_{ALG}

The online algorithm on the h-th step receives a random variable X_i decided by the almighty adversary, and once X_i is received the algorithm also sees the realization $e \in \mathcal{U}_i$ of X_i according to the distribution $\mathcal{D}_i$. The online algorithm generates a random set $T_i \subseteq \mathcal{U}_i$ *after* seeing the realization e. The idea is that if one does not see the realization e of X_i, the distribution of T_i appears identical to the product distribution generated by $\boldsymbol{z_i}$. Note that, for $S \subseteq \mathcal{U}_i$, $\Pr_{R \sim \boldsymbol{z_i}}[R = S] = \prod_{e \in S} z_e \prod_{e \in \mathcal{U} \setminus S}(1 - z_e)$.

Lemma 5.7. *For any $i \in N$ and $S \subseteq \mathcal{U}_i$,*

$$\Pr[T_i = S] = \Pr_{R \sim \boldsymbol{z_i}}[R = S].$$

Proof. Let $\mathcal{E}_e$ be the event that $e \in \mathcal{U}_i$ is the realization of X_i. Note that $\Pr[\mathcal{E}_e] = \mathcal{D}_i(e)$. Consider $S \subseteq \mathcal{U}_i$ such that $|S| \neq 1$. We see from the algorithm's description that

$$\Pr[T_i = S \mid \mathcal{E}_e] = \left(1 - \frac{\Pr_R[R(\boldsymbol{z}_i) = \{e\}]}{D_i(e)}\right) \cdot \frac{\Pr_R[R(\boldsymbol{z}_i) = S]}{1 - \Pr_R[|R(\boldsymbol{z}_i)| = 1]}. \tag{5.39}$$

Summing up over all realizations of X_i, we have that, for any S such that $|S| \neq 1$,

$$\Pr[T_i = S] = \sum_{e \in \mathcal{U}_i} D_i(e) \Pr[T_i = S \mid \mathcal{E}_e] \tag{5.40}$$

$$= \sum_{e \in \mathcal{U}_i} D_i(e) \left(1 - \frac{\Pr_R [R(\boldsymbol{z}_i) = \{e\}]}{D_i(e)}\right) \cdot \frac{\Pr_R [R(\boldsymbol{z}_i) = S]}{1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]} \tag{5.41}$$

$$= \frac{\Pr_R [R(\boldsymbol{z}_i) = S]}{1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]} \cdot \sum_{e \in \mathcal{U}_i} D_i(e) \left(1 - \frac{\Pr_R [R(\boldsymbol{z}_i) = \{e\}]}{D_i(e)}\right) \tag{5.42}$$

$$= \frac{\Pr_R [R(\boldsymbol{z}_i) = S]}{1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]} \cdot \left(\sum_{e \in \mathcal{U}_i} D_i(e) - \sum_{e \in U_i} \Pr_R [R(\boldsymbol{z}_i) = \{e\}]\right) \tag{5.43}$$

$$= \frac{\Pr_R [R(\boldsymbol{z}_i) = S]}{1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]} \cdot \left(1 - \sum_{e \in \mathcal{U}_i} \Pr_R [R(\boldsymbol{z}_i) = \{e\}]\right) \tag{5.44}$$

$$= \frac{\Pr_R [R(\boldsymbol{z}_i) = S]}{1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]} \cdot \left(1 - \Pr_R [|R(\boldsymbol{z}_i)| = 1]\right) \tag{5.45}$$

$$= \Pr_R [R(\boldsymbol{z}_i) = S]. \tag{5.46}$$

Next, consider any set S with $|S| = 1$ and, without loss of generality, assume $S = \{e\}$ for some $e \in \mathcal{U}_i$. It can be seen from the algorithm description that $T_i = \{e\}$ if and only if e is the realization of X_i and the algorithm succeeds in Line 5 in setting $T_i = \{e\}$ which happens with probability $\frac{\Pr_{R \sim \boldsymbol{z_i}}[R = \{e\}]}{\mathcal{D}_i(e)}$. Hence

$$\Pr[T_i = \{e\}] = \mathcal{D}_i(e) \cdot \frac{\Pr_{R \sim \boldsymbol{z_i}} [R = \{e\}]}{\mathcal{D}_i(e)} = \Pr_{R \sim \boldsymbol{z_i}} [R = \{e\}], \tag{5.47}$$

as desired. QED

We now analyze the expected value of $f(T_{\mathrm{ALG}})$ relying on the CRS π' that we set up (this is inspired by the use of characteristic CRS in [27]).

Lemma 5.8. *Given a (b, c)-selectable greedy OCRS π for $\mathcal{P}$, for any $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and fixed $\varepsilon > 0$, Algorithm 14 returns a set $T_{ALG} \subseteq \mathcal{U}$ such that*

$$\mathbb{E}[f(T_{ALG})] \geq c\left(e^{-b} - \varepsilon\right) \cdot F(\boldsymbol{z}).$$

Proof. It is easy to see from the algorithm's description that, for any X_i, only the actual realization of X_i can be potentially chosen to be added to T_{ALG}. Furthermore, the variables chosen by the algorithm are feasible in $\mathcal{I}$, since this is ensured by the OCRS.

Let T_i be the random set generated by the online algorithm for variable X_i. We see that T_i is independent of $T_{i'}$ for $i \neq i'$, due to independence of the realization of the random variables $X_1, \ldots, X_n$ and the independence of the coins used in the algorithm across all steps in N. From Lemma 5.7, the distribution of T_i is according to the product distribution $R \sim \boldsymbol{z_i}$ over $\mathcal{U}_i$. Let $Q = \bigcup_{i=1}^n T_i$. It follows that Q is a random set drawn from the product distribution induced by $\boldsymbol{z}$ over $\mathcal{U}$. Consider the distribution of the set $Q_\downarrow \in N$. Because of the product distribution of Q it can be see that the distribution of $Q_\downarrow$ is a product distribution on N where $i \in N$ appears in $Q_\downarrow$ with probability $x_i = 1 - \prod_{e \in \mathcal{U}_i}(1 - z_e) \leq b$ since $z \in b \cdot \mathcal{P}''$. Note that the algorithm feeds $Q_\downarrow$ to the OCRS π which is (b, c)-selectable. Let $\bar{\pi}$ be the characteristic CRS of π.

Fix a realization S of Q, along with an instantiation of $\mathcal{F}_{\pi, \boldsymbol{x}}$. Notice $e \in T_{\mathrm{ALG}} \cap \mathcal{U}_i$ if and only if

$|S \cap \mathcal{U}_i| = \{e\}$ and $i \in \pi(S_\downarrow)$. In fact,

$$T_{\mathrm{ALG}} = \bigcup_{\substack{i \in \pi(S_\downarrow) \\ |S \cap \mathcal{U}_i| = 1}} (S \cap \mathcal{U}_i), \tag{5.48}$$

by the description of Algorithm 14. By Observation 5.3, we have $\bar{\pi}(A) \subseteq \pi(A)$ for any $A \subseteq N$, and thus $\pi'(S) \subseteq T_{\mathrm{ALG}}$. Therefore, by the monotonicity of f, we have $f(T_{\mathrm{ALG}}) \geq f(\pi'(S))$, and by unconditioning

$$\mathbb{E}[f(T_{\mathrm{ALG}})] \geq \mathbb{E}[f(\pi'(Q))]. \tag{5.49}$$

Finally, by Lemma 5.6 and Remark 5.3, we have that for any $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and any fixed $\varepsilon > 0$,

$$\mathbb{E}[f(\pi'(Q))] \geq c\left(e^{-b} - \varepsilon\right) \cdot F(\boldsymbol{z}), \tag{5.50}$$

which yields

$$\mathbb{E}[f(T_{\mathrm{ALG}})] \geq c\left(e^{-b} - \varepsilon\right) \cdot F(\boldsymbol{z}). \tag{5.51}$$

QED

We are now ready for the main theorem of this section, which follows from Lemmas 5.8 and 5.3, and Claim 5.1.

Theorem 5.5. *Let $(N, \boldsymbol{D}, \mathcal{I}, f)$ be an instance of the Submodular Prophet Inequality model and let OPT denote the prophet's value. Given a (b, c)-selectable greedy OCRS π for $\mathcal{P}$, for a non-negative monotone submodular function f, $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and fixed $\varepsilon > 0$, Algorithm 14 returns a set T_{ALG} such that*

$$\mathbb{E}[f(T_{ALG})] \geq c\left(e^{-b} - \varepsilon\right)\left(1 - e^{-b}\right) \cdot OPT.$$

Next, we provide concrete results for several constraints, given an OCRS for these constraints. First, we summarize known greedy OCRSs for various constraints of interest below.

Lemma 5.9 (Theorem 1.1 from [27])**.** *There exist:*

- *For every $b \in [0, 1]$, a $(b, 1 - b)$-selectable deterministic greedy OCRS for matroid polytopes.*
- *For every $b \in [0, 1]$, a (b, e^{-2b})-selectable randomized greedy OCRS for matching polytopes.*
- *For every $b \in [0, \frac{1}{2}]$, a $(b, \frac{1-2b}{2-2b})$-selectable randomized greedy OCRS for the natural relaxation of a knapsack constraint.*

By combining Lemma 5.9 with Theorem 5.5, we obtain the following corollary.

Corollary 5.1. *Let $(N, \boldsymbol{D}, \mathcal{I}, f)$ be an instance of the Submodular Prophet Inequality model and let OPT denote the prophet's value. For a non-negative monotone submodular function f and any fixed $\varepsilon > 0$, Algorithm 14 returns a set T_{ALG} such that*

$$\mathbb{E}_{\boldsymbol{X}, \mathcal{T}}[f(T_{ALG})] \geq (1 - b)\left(e^{-b} - \varepsilon\right)\left(1 - e^{-b}\right) \cdot OPT, \qquad \forall b \in [0, 1], \tag{5.52}$$

if $\mathcal{I}$ is a matroid constraint

$$\mathbb{E}_{\boldsymbol{X}, \mathcal{T}}[f(T_{ALG})] \geq e^{-2b}\left(e^{-b} - \varepsilon\right)\left(1 - e^{-b}\right) \cdot OPT, \qquad \forall b \in [0, 1], \tag{5.53}$$

$$\mathop{\mathbb{E}}_{\boldsymbol{X},\mathcal{T}}[f(T_{ALG})] \geq \frac{1-2b}{2-2b}\left(e^{-b}-\varepsilon\right)\left(1-e^{-b}\right)\cdot OPT, \qquad \forall b \in \left[0, \frac{1}{2}\right], \tag{5.54}$$

if $\mathcal{I}$ is a matching constraint

if $\mathcal{I}$ is a knapsack constraint

where $\mathcal{T} = \{T^1, \dots, T^n\}$ denotes the set of random sets Algorithm 14 generates.

5.4.2 Non-Negative Submodular Functions

Below we describe the algorithm for non-negative functions. It is very similar to the monotone case except for a minor change in accepting an element e; in the final step, the algorithm tosses an additional random coin and accepts e with probability 1/2 (see Line 10 in the algorithm). This is inspired by the similar idea in [27] in handling non-monotone functions.

Algorithm 5.2: General Rounding $(\mathcal{U}, f, \boldsymbol{D}, \mathcal{I}, \pi, \boldsymbol{z})$

```
1  T_ALG = ∅
2  for h ← 1 to n do
3      Let X_i be variable that arrives on step h
4      Let e ∈ U_i be the realization of X_i
5      With probability Pr_{R~z_i}[R={e}] / D_i(e), set T_i ← {e}
6      Otherwise, set T_i to be a random subset R of U_i, drawn according to z_i, conditioned on |R| ≠ 1
7      if T_i ≠ ∅ then
8          Feed i to OCRS π for P
9          if π accepts i and T_i = {e} then
10             With probability 1/2, T_ALG ← T_ALG ∪ {e}
11         end
12     end
13 end
14 Return T_ALG
```

Notice that Lemmas 5.6 and 5.7 still hold, as they do not depend on the monotonicity of f. We present the following analogue of Lemma 5.8 for general submodular functions. The proof of the next lemma relies on an argument similar to that in [27].

Lemma 5.10. *Given a (b, c)-selectable greedy OCRS π for $\mathcal{P}$, for any $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and fixed $\varepsilon > 0$, Algorithm 14 returns a set $T_{ALG} \subseteq \mathcal{U}$ such that*

$$\mathbb{E}\left[f(T_{ALG})\right] \geq \frac{c\left(e^{-b}-\varepsilon\right)}{4} \cdot F(\boldsymbol{z}).$$

Proof. At every step i, Algorithm 14 draws a random set T_i according to the product distribution on $\mathcal{U}_i$ with probabilities $\boldsymbol{z_i}$, by Lemma 5.7. Let $Q = \bigcup_{i \in N} T_i$. Since the realizations between the steps are independent, Q is a random set that follows the product distribution on $\mathcal{U}$ with probabilities $\boldsymbol{z}$. Fix a realization S of Q and an instantiation of $\mathcal{F}_{\pi,\boldsymbol{x}}$. Notice that $e \in T_{\text{ALG}} \cap \mathcal{U}_i$ if and only if $|S \cap \mathcal{U}_i| = 1$, $i \in \pi(S_\downarrow)$ and the coin toss of Line 10 succeeds. In fact, if we denote

$$W = \bigcup_{\substack{i \in \pi(S_\downarrow) \\ |S \cap \mathcal{U}_i| = 1}} (S \cap \mathcal{U}_i), \tag{5.55}$$

we have that $\mathbb{E}\left[f(T_{\text{ALG}})\right] = \mathbb{E}[f(W(1/2))]$, by the description of Algorithm 14. By Observation 5.3, we have $\bar{\pi}(A) \subseteq \pi(A)$ for any $A \subseteq N$, and thus $\pi'(S) \subseteq W$. For ease of notation, we denote $\pi'(S)$ by L. For our fixed choice of S and $\mathcal{F}_{\pi,\boldsymbol{x}}$, L is deterministic. Therefore, we can think of $W(1/2)$ as obtained by first calculating a set $L(1/2)$ in which every element of L appears with probability $1/2$ independently, and then adding to it a random set $\Delta \subseteq \mathcal{U} \setminus L$. The almighty prophet can control the order in which the elements arrive, and thus can make the distribution of Δ depend on $L(1/2)$. However, Δ is guaranteed to contain every element with probability at most $1/2$, for every given realization of $L(1/2)$. Thus,

$$\begin{aligned}
\mathbb{E}\left[f(W(1/2)) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right] &= \mathbb{E}\left[f(L(1/2) \cup \Delta) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right] && (5.56)\\
&= \sum_{B \subseteq L} \Pr\left[L(1/2) = B \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right] \cdot \mathbb{E}\left[f(B \cup \Delta) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right] && (5.57)\\
&\geq \sum_{B \subseteq L} \Pr\left[L(1/2) = B \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right] \cdot \frac{\mathbb{E}\left[f(B) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right]}{2} && (5.58)\\
&= \frac{\mathbb{E}\left[f(L(1/2)) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right]}{2} && (5.59)\\
&= \frac{\mathbb{E}\left[f(L) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right]}{4}, && (5.60)
\end{aligned}$$

where the first inequality follows from Lemma 5.1 since the function $h_B(T) = h(B \cup T)$ is non-negative and submodular for all $B \subseteq \mathcal{U}$, and the second inequality follows from Lemma 5.2. Taking an expectation over all possible realizations of S and $\mathcal{F}_{\pi,\boldsymbol{x}}$, we obtain

$$\begin{aligned}
\mathbb{E}\left[f(W(1/2))\right] &= \mathop{\mathbb{E}}_{S, \mathcal{F}_{\pi,\boldsymbol{x}}}\left[\mathbb{E}\left[f(W(1/2)) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right]\right] \geq \mathop{\mathbb{E}}_{S, \mathcal{F}_{\pi,\boldsymbol{x}}}\left[\frac{\mathbb{E}\left[f(L) \mid S, \mathcal{F}_{\pi,\boldsymbol{x}}\right]}{4}\right] && (5.61)\\
&= \frac{\mathbb{E}\left[f(L)\right]}{4}. && (5.62)
\end{aligned}$$

Finally, by Lemma 5.6 and Remark 5.3, we have that for any $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and any fixed $\varepsilon > 0$,

$$\frac{\mathbb{E}\left[f(L)\right]}{4} \geq \frac{c\left(e^{-b} - \varepsilon\right)}{4} \cdot F(\boldsymbol{z}), \tag{5.63}$$

which implies

$$\mathbb{E}\left[f(T_{\text{ALG}})\right] = \mathbb{E}\left[f(W(1/2))\right] \geq \frac{c\left(e^{-b} - \varepsilon\right)}{4} \cdot F(\boldsymbol{z}). \tag{5.64}$$

QED

We are now ready to proceed with the main result for general submodular functions, which follows from Lemma 5.10, Theorem 5.4, and Claim 5.1.

Theorem 5.6. *Let $(N, \boldsymbol{D}, \mathcal{I}, f)$ be an instance of the Submodular Prophet Inequality model and let OPT denote the prophet's value. Given a (b, c)-selectable greedy OCRS π for $\mathcal{P}$, for a non-negative submodular* *function f, $\boldsymbol{z} \in b \cdot \mathcal{P}''$ and fixed $\varepsilon > 0$, Algorithm 14 returns a set T_{ALG} such that*

$$\mathbb{E}\left[f(T_{ALG})\right] \geq \frac{c\left(e^{-b} - \varepsilon\right)}{4} \cdot \left(1 - e^{-b} - \varepsilon\right) \cdot OPT.$$

By combining Lemma 5.9 with Theorem 5.6, we obtain the following corollary.

Corollary 5.2. *Let $(N, \boldsymbol{D}, \mathcal{I}, f)$ be an instance of the Submodular Prophet Inequality model and let OPT denote the prophet's value. For a non-negative submodular function f and any fixed $\varepsilon > 0$, Algorithm 14 returns a set T_{ALG} such that*

$$\mathbb{E}[f(T_{ALG})] \geq \frac{(1-b)\left(e^{-b} - \varepsilon\right)}{4} \cdot \left(1 - e^{-b} - \varepsilon\right) \cdot OPT, \qquad \forall b \in [0,1], \tag{5.65}$$

if $\mathcal{I}$ is a matroid constraint

$$\mathbb{E}[f(T_{ALG})] \geq \frac{e^{-2b}\left(e^{-b} - \varepsilon\right)}{4} \cdot \left(1 - e^{-b} - \varepsilon\right) \cdot OPT, \qquad \forall b \in [0,1], \tag{5.66}$$

if $\mathcal{I}$ is a matching constraint

$$\mathbb{E}[f(T_{ALG})] \geq \frac{(1-2b)\left(e^{-b} - \varepsilon\right)}{8-8b} \cdot \left(1 - e^{-b} - \varepsilon\right) \cdot OPT, \qquad \forall b \in \left[0, \frac{1}{2}\right], \tag{5.67}$$

if $\mathcal{I}$ is a knapsack constraint

where $\mathcal{T} = \{T^1, \ldots, T^n\}$ denotes the set of random sets Algorithm 14 generates.

5.5 THE CORRELATION GAP OF NON-NEGATIVE SUBMODULAR FUNCTIONS

In this section we prove Theorems 5.2 and 5.4 on the correlation gap for non-negative submodular functions.

The correlation gap for monotone functions [45], [47] used a continuous time argument by relating $F(\boldsymbol{x})$ and $f^+(\boldsymbol{x})$ via another continuous extension f^* and this was the same approach followed in [20]. We take a different approach. For the exact correlation gap in Theorem 5.2 we build on a proof for the monotone case from [26] which is less well-known; we adapt their proof for the non-negative case via the parameter p. The proof of Theorem 5.2 is qualitatively different from that of Theorem 5.4.

We split the proof into two parts, the upper bound and the lower bound, state them separately and give their proof. Before we begin, we present two lemmas that are useful in our analysis.

Lemma 5.11 (Lemma 4.3 from [47])**.** *Let $f : 2^N \to \mathbb{R}_{\geq 0}$ be a submodular function, let $A_1, A_2, \ldots, A_k \subseteq N$ be k (not necessarily disjoint) sets and let $A_1(p_1), A_2(p_2), \ldots, A_k(p_k)$ their independently sampled subsets, where each element of A_i appears in $A_i(p_i)$ with probability p_i, for all $1 \leq i \leq k$. Then*

$$\mathbb{E}\left[f\left(\bigcup_{i=1}^{k} A_i(p_i)\right)\right] \geq \sum_{I \subseteq [k]} \left(\prod_{j \in I} p_j \prod_{j \notin I} (1-p_j) f\left(\bigcup_{j \in I} A_j\right)\right).$$

The next Lemma appears in [26], but its proof is slightly obfuscated within Lemma B.2. For clarity, we present it here on its own.

Lemma 5.12 ([26]). *Let $a_1 \geq \cdots \geq a_m \in \mathbb{R}_{\geq 0}$, and $q_1, \ldots, q_m \in [0,1]$ such that $\sum_{k=1}^{m} q_k = 1$. Then*

$$\sum_{k=1}^{m} q_k \, a_k \prod_{j=1}^{k-1} (1 - q_j) \geq \left(1 - \frac{1}{e}\right) \cdot \sum_{j=1}^{m} q_j a_j.$$

Proof. Since the above inequality is linear in the parameters a_i, it suffices to prove it for the special case $a_1 = a_2 = \cdots = a_r = 1$ and $a_{r+1} = \cdots = a_m = 0$. (A general decreasing sequence of a_j can be obtained as a positive linear combination of such special cases). Hence, it remains to prove

$$\sum_{k=1}^{r} q_k \prod_{j=1}^{k-1} (1 - q_j) \geq \left(1 - \frac{1}{e}\right) \cdot \sum_{j=1}^{r} q_j. \tag{5.68}$$

We start from the left-hand side, which we expand to

$$\sum_{k=1}^{r} q_k \prod_{j=1}^{k-1} (1 - q_j) = 1 - \prod_{k=1}^{r} (1 - q_k) \geq 1 - \left(1 - \frac{1}{r} \sum_{k=1}^{r} q_k\right)^r, \tag{5.69}$$

where the inequality follows from the arithmetic-geometric mean inequality. Finally, we use the concavity of $\phi_r(x) \coloneqq 1 - \left(1 - \frac{x}{r}\right)^r$, and the fact that $\phi_r(0) = 0$, to get

$$\phi_r(x) \geq \phi_r(1) \cdot x = \left(1 - \left(1 - \frac{1}{r}\right)^r\right) \cdot x \tag{5.70}$$

for $x \in [0,1]$. Since $\left(1 - \left(1 - \frac{1}{r}\right)^r\right) \geq \left(1 - \frac{1}{e}\right)$ for all r, we get

$$\phi_r(x) \geq \left(1 - \frac{1}{e}\right) \cdot x. \tag{5.71}$$

which implies that

$$\phi_r\left(\sum_{k=1}^{r} q_k\right) = 1 - \left(1 - \frac{1}{r} \sum_{k=1}^{r} q_k\right)^r \geq \left(1 - \frac{1}{e}\right) \cdot \sum_{k=1}^{r} q_k. \tag{5.72}$$

QED

5.5.1 Upper Bound

The proof of this upper bound is inspired by the proof in [26] for the monotone case, which is different from the earlier one in [47].

Theorem 5.7. *Let $f : 2^N \rightarrow \mathbb{R}_{\geq 0}$ be a non-negative submodular function, where $n = |N|$. Let $\boldsymbol{x} \in [0,1]^n$, such that $\boldsymbol{x} \leq p \cdot \mathbf{1}_N$ for some $p \in [0,1]$. Then,*

$$F(\boldsymbol{x}) \geq (1 - p) \left(1 - \frac{1}{e}\right) f^+(\boldsymbol{x}).$$

Proof. Consider a basic feasible solution $(q_j, A_j)_{j \in [m]}$ to the linear program that defines $f^+(\boldsymbol{x})$. In other words, $f^+(\boldsymbol{x}) = \sum_{j=1}^{m} q_j f(A_j)$, where $\sum_{j=1}^{m} q_j = 1$, $\sum_{j : i \in A_j} q_j = x_i$, for all i, and $q_j \geq 0$ for all j. Notice that, since we chose a basic feasible solution and the LP that defines $f^+(\boldsymbol{x})$ has only $n+1$ constraints, apart from the non-negativity constraints, we have $m \leq n+1$.

Next, consider the following process to generate a subset of elements. For each $j \in [m]$ sample independently each element of A_j with probability q_j. An element $i \in N$ is not selected with probability equal to $\prod_{j:i\in A_j}(1-q_j)$, thus, i is selected with probability equal to $1-\prod_{j:i\in A_j}(1-q_j)$. Notice that we can assume without loss of generality that $q_j \neq 1$ for all j; if $q_j = 1$ for some j then that implies that $x_i = 1$ for every element $i \in A_j$, and $q_{j'} = 0$ for all $j' \neq j$, which then leads us to $F(\boldsymbol{x}) = f(A_j) = f^+(\boldsymbol{x})$.

However, we want to make each element i to be selected with probability exactly equal to $x_i = \sum_{j:i\in A_j} q_j$. To do this, we simply need to sample again each element i with probability r_i, where

$$1-(1-r_i)\cdot\prod_{j:i\in A_j}(1-q_j) = \sum_{j:i\in A_j} q_j. \tag{5.73}$$

It is easy to see that $0 \leq r_i \leq x_i \leq p$.

Consider the sampling scheme described above, and let R denote a random set created via this sampling scheme. Notice that in our sampling scheme, each element i is chosen independently with probability x_i, which implies that $\mathbb{E}_R[f(R)] = F(\boldsymbol{x})$.

We consider $m+n$ sets $B_1, B_2, \ldots, B_{m+n}$ where $B_j = A_j$ for $1 \leq j \leq m$, and $B_{m+i} = \{i\}$ for $1 \leq i \leq n$. Let $\mathcal{J}$ denote a random subset of $[m+n]$ obtained by including each $j \in \{1,2,\ldots,m\}$ independently with probability q_j and each $i \in \{m+1, m+2, \ldots, m+n\}$ independently with probability r_i. Also, let $R' \subseteq N$ denote a random set where

$$R' = \bigcup_{j\in\mathcal{J}} B_j. \tag{5.74}$$

The next claim is based on the submodularity of f.

Claim 5.2.

$$F(\boldsymbol{x}) \geq \mathop{\mathbb{E}}_{\mathcal{J}}\left[f(R')\right].$$

Proof. Since $F(\boldsymbol{x}) = \mathbb{E}_R[f(R)]$, it suffices to show that

$$\mathop{\mathbb{E}}_{R}\left[f(R)\right] \geq \mathop{\mathbb{E}}_{\mathcal{J}}\left[f(R')\right]. \tag{5.75}$$

We apply Lemma 5.11 for $k = m+n$, $A_j = B_j$ for $1 \leq j \leq m+n$, $p_j = q_j$ for $1 \leq j \leq m$, and $p_{m+i} = r_i$ for $1 \leq i \leq n$. Notice that

$$\mathop{\mathbb{E}}_{R}\left[f(R)\right] = \mathbb{E}\left[f\left(\bigcup_{i=1}^{k} B_i(p_i)\right)\right], \tag{5.76}$$

while

$$\mathop{\mathbb{E}}_{\mathcal{J}}\left[f(R')\right] = \sum_{I\subseteq[k]}\left(\prod_{j\in I} p_j \prod_{j\notin I}(1-p_j) f\left(\bigcup_{j\in I} B_j\right)\right), \tag{5.77}$$

and thus, by Lemma 5.11, we get

$$\mathop{\mathbb{E}}_{R}\left[f(R)\right] \geq \mathop{\mathbb{E}}_{\mathcal{J}}\left[f(R')\right]. \tag{5.78}$$

QED

Claim 5.3.

$$\mathop{\mathbb{E}}_{\mathcal{J}}\left[f(R')\right] \geq (1-p)\left(1-\frac{1}{e}\right)\cdot f^+(\boldsymbol{x}).$$

Proof. Assume, without loss of generality, that $f(A_1) \geq \cdots \geq f(A_m)$. We analyze $\mathbb{E}[f(R')]$ by conditioning on the minimum index j that belongs to $\mathcal{J}$. For $k \in [m]$, let

$$J_k = \{I \subseteq [m+n] \mid k \in I \text{ and } \ell \notin I, \forall \ell < k\}. \tag{5.79}$$

Furthermore, for $k \in [m]$ define the set function $g_k : 2^N \to \mathbb{R}_{\geq 0}$ where $g_k(S) = f(B_k \cup S)$ for all $S \subseteq N$. It is easy to verify that g_k is non-negative and submodular because f is non-negative and submodular. $\mathcal{J} \in J_k$ implies that $B_k \subseteq R'$, hence,

$$\mathop{\mathbb{E}}_{\mathcal{J}}[f(R') \mid \mathcal{J} \in J_k] = \mathop{\mathbb{E}}_{\mathcal{J}}[f(B_k \cup (R' \setminus B_k)) \mid \mathcal{J} \in J_k] \tag{5.80}$$

$$= \mathop{\mathbb{E}}_{\mathcal{J}}[g_k(R' \setminus B_k) \mid \mathcal{J} \in J_k]. \tag{5.81}$$

For any fixed $i \in N$ we analyze the probability that $i \in R' \setminus B_k$ conditioned on $\mathcal{J} \in J_k$. Using independence of the choice of each index in $\mathcal{J}$ we obtain the following.

$$\Pr_{\mathcal{J}}[i \in (R' \setminus B_k) \mid \mathcal{J} \in J_k] = 1 - (1 - r_i) \prod_{j : i \in A_j, k < j \leq m} (1 - q_j) \tag{5.82}$$

$$\leq 1 - (1 - r_i) \prod_{j : i \in A_j, j \in [m]} (1 - q_j) \tag{5.83}$$

$$= x_i \leq p. \tag{5.84}$$

Thus, applying Lemma 5.1 to g_k yields

$$\mathop{\mathbb{E}}_{\mathcal{J}}[g_k(R' \setminus B_k) \mid \mathcal{J} \in J_k] \geq (1 - p) g_k(\emptyset) = (1 - p) f(B_k). \tag{5.85}$$

Combining the above,

$$\mathop{\mathbb{E}}_{\mathcal{J}}[f(R') \mid \mathcal{J} \in J_k] \geq (1 - p) \cdot f(B_k). \tag{5.86}$$

Also notice that

$$\Pr_{\mathcal{J}}[\mathcal{J} \in J_k] = \Pr_{\mathcal{J}}[k \in \mathcal{J}] \cdot \prod_{j=1}^{k-1} \left(1 - \Pr_{\mathcal{J}}[j \in \mathcal{J}]\right) = q_k \cdot \prod_{j=1}^{k-1} (1 - q_j). \tag{5.87}$$

Therefore,

$$\mathop{\mathbb{E}}_{\mathcal{J}}[f(R')] = \sum_{k=1}^{m} \Pr_{\mathcal{J}}[\mathcal{J} \in J_k] \cdot \mathop{\mathbb{E}}_{\mathcal{J}}[f(R') \mid \mathcal{J} \in J_k] \tag{5.88}$$

$$+ \Pr_{\mathcal{J}}[\mathcal{J} \cap [m] = \emptyset] \cdot \mathop{\mathbb{E}}_{\mathcal{J}}[f(R') \mid \mathcal{J} \cap [m] = \emptyset] \tag{5.89}$$

$$\geq \sum_{j=1}^{m} \Pr_{\mathcal{J}}[\mathcal{J} \in J_k] \cdot \mathop{\mathbb{E}}_{\mathcal{J}}[f(R') \mid \mathcal{J} \in J_k] \tag{5.90}$$

$$\geq \sum_{k=1}^{m} \Pr_{\mathcal{J}}[\mathcal{J} \in J_k] \cdot (1 - p) \cdot f(B_k) \tag{5.91}$$

$$= (1-p) \sum_{k=1}^{m} q_k f(B_k) \prod_{j=1}^{k-1} (1-q_j), \tag{5.92}$$

where the first inequality follows from the non-negativity of f, the second inequality follows from (5.86) and the last equality from (5.87). However, for $1 \leq k \leq m$, we have $B_k = A_k$, and thus

$$\mathbb{E}_{\mathcal{J}}[f(R')] \geq (1-p) \sum_{k=1}^{m} q_k f(A_k) \prod_{j=1}^{k-1} (1-q_j). \tag{5.93}$$

Finally, utilizing Lemma 5.12 for $a_k = f(A_k)$, we get that

$$\sum_{k=1}^{m} q_k f(A_k) \prod_{j=1}^{k-1} (1-q_j) \geq \left(1 - \frac{1}{e}\right) \cdot \sum_{j=1}^{m} q_j f(A_j) = \left(1 - \frac{1}{e}\right) \cdot f^+(\boldsymbol{x}). \tag{5.94}$$

Combining (5.88) and (5.94),

$$\mathbb{E}_{\mathcal{J}}[f(R')] \geq (1-p) \left(1 - \frac{1}{e}\right) \cdot f^+(\boldsymbol{x}). \tag{5.95}$$

QED

Finally, combining Claims 5.2 and 5.3, we obtain

$$F(\boldsymbol{x}) \geq (1-p) \left(1 - \frac{1}{e}\right) \cdot f^+(\boldsymbol{x}), \tag{5.96}$$

which completes the proof. QED

5.5.2 Lower Bound

A simple example on $n = 2$ shows that $F(\boldsymbol{x}) \leq (1-p) f^+(\boldsymbol{x})$; the function is the cut function of a directed graph on two vertices. For monotone functions, a simple coverage example shows that $F(\boldsymbol{x}) \leq (1 - 1/e) f^+(\boldsymbol{x})$. We combine and generalize these two examples to create an instance for non-monotone functions and obtain the following theorem.

Theorem 5.8. *There exists a non-negative submodular function $f : 2^N \to \mathbb{R}_{\geq 0}$ such that, for any $0 \leq p \leq 1$, there exists an $\boldsymbol{x} \in [0,1]^n$ where $\|\boldsymbol{x}\|_\infty \leq p$ and*

$$F(\boldsymbol{x}) \leq \left(1 - e^{-(1-p)}\right) f^+(\boldsymbol{x}).$$

Proof. Consider the following graph $G = (V, E)$, where $V = \{u_1, \ldots, u_n, v\}$, and $E = \{(u_i, v) \mid 1 \leq i \leq n\}$. Let $x_{u_i} = \frac{1-p}{n}$ for all $i \in \{1, \ldots, n\}$ and $x_v = p$. We define a function $f : 2^V \to \mathbb{R}_{\geq 0}$ as follows

$$f(S) = \begin{cases} 1 & \text{if } v \notin S \text{ and } S \neq \emptyset, \\ 0 & \text{otherwise.} \end{cases} \tag{5.97}$$

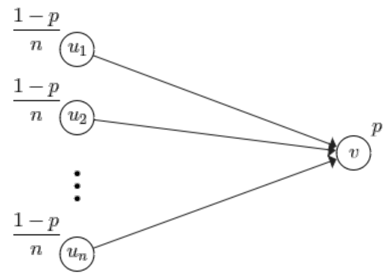

Figure 5.1: Graph G which yields the desired lower bound.

It is easy to see that f is submodular. Notice that

$$f^+(\boldsymbol{x}) = 1 - p, \tag{5.98}$$

as the coefficients that maximize $\sum_S a_S f(S)$ subject to the constraints are $a_{\{v\}} = p$, $a_{\{u_i\}} = \frac{1-p}{n}$ for all $i \in \{1, \ldots, n\}$ and $a_S = 0$, for $|S| \neq 1$. In other words, $a_{\{u\}} = x_u$ for all $u \in V$, and $a_S = 0$, if $|S| \neq 1$.

Next, notice that, if $R(\boldsymbol{x}) \subseteq V$ is a random set, where each element $u \in V$ is sampled with probability x_u, then $f(R(\boldsymbol{x})) = 1$ if and only if v is not selected in $R(\boldsymbol{x})$, but at least one element of $V \setminus \{v\}$ is selected. Therefore,

$$F(\boldsymbol{x}) = \mathbb{E}[f(R(\boldsymbol{x}))] = (1-p) \cdot \left(1 - \left(1 - \frac{1-p}{n}\right)^n\right), \tag{5.99}$$

which implies that

$$\frac{F(\boldsymbol{x})}{f^+(\boldsymbol{x})} = \frac{(1-p) \cdot \left(1 - \left(1 - \frac{1-p}{n}\right)^n\right)}{1-p} = 1 - \left(1 - \frac{1-p}{n}\right)^n. \tag{5.100}$$

As $n \to \infty$, we get

$$\lim_{n \to \infty} 1 - \left(1 - \frac{1-p}{n}\right)^n = 1 - e^{-(1-p)}. \tag{5.101}$$

We conclude that, for any $0 \leq p < 1$, when $x_i \leq p$ for all i,

$$F(\boldsymbol{x}) \leq \left(1 - e^{-(1-p)}\right) f^+(\boldsymbol{x}). \tag{5.102}$$

QED

Chapter 6: CONCLUSIONS AND FUTURE DIRECTIONS

6.1 DELEGATION OF ONLINE PROBLEMS

In the last chapter of this thesis, we discuss several further directions on optimal stopping and stochastic combinatorial optimization problems.

Starting with the work of Kleinberg and Kleinberg [148], prophet inequalities have recently been connected to delegation settings, in which a principal delegates a search problem to one or more agents, whose objectives may not be aligned with the principal's objective. Formally, there exists a set of potential solutions Ω, known to both the principal and the agents. The agents observe elements of Ω drawn I.I.D. from a common distribution. Out of all solutions that each agent observed, they propose one of them to the principal. There exist valuation functions $x : \Omega \to \mathbb{R}$ and $y_i : \Omega \to \mathbb{R}$ for each agent i such that f the principal selects a proposed solution $\omega \in \Omega$, they receive value $x(\omega)$ and the agent i who proposed the solution receives value $y_i(\omega)$.

A delegation mechanism is a set $R \subset \Omega$ that the principal announces to the agents, promising that the solution accepted by the principal has to lie in R. Kleinberg and Kleinberg show that a simple threshold mechanism in which the principal proposes $R = \{\omega \mid x(\omega) \geq \tau\}$ to a single agent achieves at least a $1/2$-approximation to the utility of the optimal mechanism, via an elegant reduction to the prophet inequality setting. Subsequent work generalizes this result to multiple agent delegation [149].

In the Kleinberg-Kleinberg model of delegation, the agents observe their realization offline, and decide what to propose to the principal after all their observations have concluded. This inspires the following question.

Question 6.1
Is there a $\mathrm{O}(1)$*-approximate delegation mechanism when the agent observes their realizations in an online manner?*

In other words, what guarantees can we achieve when the principal delegates an inherently online instance such as a prophet inequality instance to the agent? In ongoing work with our collaborators, we are able to show the following.

Theorem 6.1. *There exists a* $k \in \mathbb{N}$ *such that if* P *delegates a prophet inequality instance to* k *agents they achieve, on expectation, the same utility they would receive if they observed the realizations and made all the decisions on their own.*

Another interesting question in delegation concerns the effect of competition among multiple agents.

Question 6.2
Can we obtain a Bulow-Klemperer style result for delegation, relating the utility of the optimal mechanism when delegating to k *agents to the utility of the best threshold mechanism when delegating to* $c \cdot k$ *agents for some constant* c*? How large does* c *need to be to beat the optimal mechanism for* k *agents?*

6.2 GOING BEYOND THE INDEPENDENCE ASSUMPTION

A strong assumption in most works in the literature, including the results of this thesis, is independence between the realizations of different elements. For prophet inequalities, the case of correlated distributions was first studied only recently for linearly correlated distributions [150] and pairwise-independent distributions [151], and is relatively unexplored.

A graphical model that is quite general and is usually considered in probability theory and statistical physics is the Markov Random Field (MRF), in which the random variables correspond to vertices of a graph and the strength of correlation between them is captured by the maximum edge weight Δ. Cai and Oikonomou [152] gave a prophet inequality for the MRF model that is exponential in Δ, but left open the possibility of a competitive ratio that is polynomial in Δ, which would be optimal due to an upper bound they present. In ongoing work with our collaborators, we are able to show the following.

Theorem 6.2. *There exists a an algorithm for the prophet inequality setting under the MRF model that is* $\mathrm{O}\left(1/\Delta\right)$*-competitive.*

Studying correlated realizations is of significant importance even beyond the prophet inequality setting. Specifically for contention resolution, Duhgmi [140], [141] connected correlated contention resolution to the famous matroid secretary problem [142]. Recall that in the matroid secretary problem we are given a matroid and are presented with n elements in a uniformly random order. Whenever we observe an element, we get to see its weight and our objective is to select an independent set of the matroid that is a constant approximation to the maximum-weight basis of the matroid.

Returning to CRSs, consider offline CRSs for matroids and let $\mathcal{H}$ denote the set of all distributions $\mathcal{D}$ for which there exists a $\mathrm{O}(1)$-selectable CRS when the active elements are drawn according to $\mathcal{D}$. In a surprising turn of events, Dughmi [140], [141] showed that the matroid secretary problem is equivalent to showing that for every distribution $\mathcal{D} \in \mathcal{H}$, there exists a $\mathrm{O}(1)$-selectable Random Order CRS (ROCRS) for matroids when the active elements are drawn according to $\mathcal{D}$. In other words, to settle the matroid secretary problem it is necessary to develop an understanding of ROCRSs for correlated distributions.

Question 6.3
Are offline contention resolution and online contention resolution equivalent up to a constant?

Very recently, Dughmi, Kalayci and Patel [153] studied contention resolution for matroids and pairwise-independent distributions, giving hardness results for general matroids and constant-factor CRSs for special classes of matroids.

6.3 GENERALIZATIONS OF MINIMIZATION PROPHET INEQUALITIES

Our results on the I.I.D. Min-Prophet Inequality in Chapter 2 open up a series of questions. We presented a thorough examination of the single-item setting and I.I.D. random variables, but the non-I.I.D. case and more general combinatorial constraints are left open.

Question 6.4
Can we obtain (distribution-dependent) constant-competitive prophet inequalities for the minimization setting? Which parameters of the different distributions do the guarantees depend on?

For the I.I.D. Min-Prophet Inequality and k-uniform matroids the competitive ratio goes to 1 as $k \to \infty$, since both the algorithm and the prophet need to select all realizations.

Question 6.5
What is the dependence on k of the competitive ratio in the I.I.D. Min-Prophet Inequality for k-uniform matroids?

Interestingly, as we show in Chapter 2, no single-threshold algorithm is constant-competitive in the single-item I.I.D. Min-Prophet Inequality, even if the constant is distribution-dependent. However, the optimal algorithm, which uses a different threshold per random variable, is constant-competitive. This immediately raises the following question.

Question 6.6
How many different thresholds are needed for an algorithm to be constant-competitive in the I.I.D. Min-Prophet Inequality?

6.4 SEPARATING ADVERSARIES BASED ON THEIR POWER

Talk about counterexample to greedy OCRSs. Can we separate almighty adversary for general matroids? (Mention Matt's paper). In Chapter 4 we present the first separation between greedy and non-greedy OCRSs. However, our counterexample fails to separate between almighty and online adversaries, as technically we do not need to use a greedy OCRS to compete with an almighty adversary.

Question 6.7
Is there a $^1/_2$*-selectable OCRS for matroids against an almighty adversary, or are almighty adversaries more powerful than online adversaries?*

Recently, Dinev and Weinberg [76] showed that the best-possible selectability of an OCRS against the almighty adversary is asymptotically worse than the best-possible selectability one can obtain against the online adversary, providing some evidence that the answer to the question above might be negative. This motivates another, more general question.

Question 6.8
Are greedy OCRSs optimal with respect to an almighty adversary?

6.5 OPTIMAL GUARANTEES AND HARDNESS RESULTS

Next, we briefly mention some settings for which the optimal guarantees are unknown.

6.5.1 The Free-Order Prophet Inequality

As we mentioned earlier, in the single-item free-order prophet inequality the algorithm is allowed to select the order in which it observes the realizations, but the decision to select or reject a realization still needs to be immediate and irrevocable. As a model, it sits between the prophet secretary and the I.I.D. setting, which motivates the following question.

Question 6.9
Is the competitive ratio of ≈ 0.745 *for the I.I.D. setting achievable in the free order setting as well?*

6.5.2 Prophet Secretary

Similarly, the best-possible competitive ratio for the prophet secretary is still open, but due to the work of Peng and Tang [154] it is strictly smaller than that of the free-order setting.

Question 6.10
What is the best-possible competitive ratio achievable in the prophet secretary model?

6.5.3 Greedy OCRSs for Matroids and Matchings

The results we presented in Chapter 4 give the first optimal greedy OCRSs for special classes of matroids, but the optimal selectability of greedy OCRSs for most feasibility constraints is still unknown. Of particular interest are constraints that are commonly used in applications such as (general) matroids and matching constraints. For the latter, an upper bound of $2/5$ is known, but the current best OCRS is 0.344-selectable [81] and the current best greedy OCRS is $1/2e$-selectable [27].

Question 6.11
Does there exist a $1/e$-selectable greedy OCRS for all matroids?

Question 6.12
What is the best-possible selectability of greedy and non-greedy OCRSs for matching constraints?

6.5.4 Oracle-Augmented Prophet Inequalities

In Chapter 3 we presented an instance that upper bounds the competitive ratio of an oracle-augmented algorithm by $1 - \frac{1}{2^{m+1}}$, for m oracle calls. This instance generalizes the standard prophet inequality counterexample and for this reason we conjecture it is actually the worst-case instance for the oracle model.

Question 6.13
Is there an algorithm for the oracle-augmented prophet inequality model with m oracle calls that is $(1 - 1/2^{m+1})$-competitive?

6.5.5 Correlation Gap

Finally, closing the gap between the upper and lower bounds of the correlation gap for general submodular functions is an interesting open question. At the same time, for other applications it might be useful to investigate the correlation gap under a different parametrization.

Question 6.14
How does the correlation gap of general submodular functions depend on the largest coordinate of the given point? Also, how does it depend on other parameters of the submodular function?

REFERENCES

[1] U. Krengel and L. Sucheston, "Semiamarts and finite values," *Bull. Amer. Math. Soc.*, vol. 83, no. 4, pp. 745–747, Jul. 1977. [Online]. Available: https://projecteuclid.org:443/euclid.bams/1183538915.

[2] U. Krengel and L. Sucheston, "On semiamarts, amarts, and processes with finite value," *Probability on Banach spaces*, vol. 4, pp. 197–266, 1978.

[3] T. P. Hill and R. P. Kertz, "Comparisons of stop rule and supremum expectations of i.i.d. random variables," *Ann. Probab.*, vol. 10, no. 2, pp. 336–345, May 1982. DOI: 10.1214/aop/1176993861. [Online]. Available: https://doi.org/10.1214/aop/1176993861.

[4] M. Abolhassani, S. Ehsani, H. Esfandiari, M. Hajiaghayi, R. D. Kleinberg, and B. Lucier, "Beating 1-1/e for ordered prophets," in *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, H. Hatami, P. McKenzie, and V. King, Eds., ACM, 2017, pp. 61–71. DOI: 10.1145/3055399.3055479. [Online]. Available: https://doi.org/10.1145/3055399.3055479.

[5] J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld, "Posted price mechanisms and optimal threshold strategies for random arrivals," *Math. Oper. Res.*, vol. 46, no. 4, pp. 1452–1478, 2021. DOI: 10.1287/moor.2020.1105. [Online]. Available: https://doi.org/10.1287/moor.2020.1105.

[6] 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*, ser. AAAI'07, Vancouver, British Columbia, Canada: AAAI Press, 2007, pp. 58–65, ISBN: 978-1-57735-323-2. [Online]. Available: http://dl.acm.org/citation.cfm?id=1619645.1619656.

[7] 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*, ser. STOC '10, Cambridge, Massachusetts, USA: Association for Computing Machinery, 2010, pp. 311–320, ISBN: 9781450300506. DOI: 10.1145/1806689.1806733. [Online]. Available: https://doi.org/10.1145/1806689.1806733.

[8] R. Kleinberg and S. M. Weinberg, "Matroid prophet inequalities and applications to multi-dimensional mechanism design," *Games Econ. Behav.*, vol. 113, pp. 97–115, 2019. DOI: 10.1016/j.geb.2014.11.002. [Online]. Available: https://doi.org/10.1016/j.geb.2014.11.002.

[9] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. Singla, "Prophet secretary for combinatorial auctions and matroids," in *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, A. Czumaj, Ed., SIAM, 2018, pp. 700–714. DOI: 10.1137/1.9781611975031.46. [Online]. Available: https://doi.org/10.1137/1.9781611975031.46.

[10] P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier, "Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs," *SIAM Journal on Computing*, vol. 49, no. 3, pp. 540–582, 2020. DOI: 10.1137/20M1323850. eprint: https://doi.org/10.1137/20M1323850. [Online]. Available: https://doi.org/10.1137/20M1323850.

[11] P. Dütting, T. Kesselheim, and B. Lucier, "An $O(\log \log m)$ prophet inequality for subadditive combinatorial auctions," *SIAM Journal on Computing*, vol. 0, no. 0, FOCS20-239-FOCS20–275, 0. DOI: 10.1137/20M1382799. eprint: https://doi.org/10.1137/20M1382799. [Online]. Available: https://doi.org/10.1137/20M1382799.

[12] H. Esfandiari, M. Hajiaghayi, V. Liaghat, and M. Monemizadeh, "Prophet secretary," *SIAM Journal on Discrete Mathematics*, vol. 31, no. 3, pp. 1685–1701, 2017. DOI: 10.1137/15M1029394. eprint: https://doi.org/10.1137/15M1029394. [Online]. Available: https://doi.org/10.1137/15M1029394.

[13] Y. Azar, A. Chiplunkar, and H. Kaplan, "Prophet secretary: Surpassing the 1-1/e barrier," in *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, É. Tardos, E. Elkind, and R. Vohra, Eds., ACM, 2018, pp. 303–318. DOI: 10.1145/3219166.3219182. [Online]. Available: https://doi.org/10.1145/3219166.3219182.

[14] J. Correa, R. Saona, and B. Ziliotto, "Prophet secretary through blind strategies," *Mathematical Programming*, Aug. 2020. DOI: 10.1007/s10107-020-01544-8.

[15] A. Liu, R. P. Leme, M. Pál, J. Schneider, and B. Sivan, "Variable decomposition for prophet inequalities and optimal ordering," in *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18-23, 2021*, P. Biró, S. Chawla, and F. Echenique, Eds., ACM, 2021, p. 692. DOI: 10.1145/3465456.3467598. [Online]. Available: https://doi.org/10.1145/3465456.3467598.

[16] 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*, vol. 46, no. 4, pp. 1452–1478, 2021. DOI: 10.1287/moor.2020.1105. eprint: https://doi.org/10.1287/moor.2020.1105. [Online]. Available: https://doi.org/10.1287/moor.2020.1105.

[17] J. Correa, P. Foncea, D. Pizarro, and V. Verdugo, "From pricing to prophets, and back!" *Operations Research Letters*, vol. 47, no. 1, pp. 25–29, 2019, ISSN: 0167-6377. DOI: https://doi.org/10.1016/j.orl.2018.11.010. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0167637718302554.

[18] J. Jiang, W. Ma, and J. Zhang, "Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack," in *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022, pp. 1221–1246. DOI: 10.1137/1.9781611977073.51. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611977073.51. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611977073.51.

[19] N. Arnosti and W. Ma, "Tight guarantees for static threshold policies in the prophet secretary problem," in *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds., ACM, 2022, p. 242. DOI: 10.1145/3490486.3538377. [Online]. Available: https://doi.org/10.1145/3490486.3538377.

[20] A. Rubinstein and S. Singla, "Combinatorial prophet inequalities," in *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, Longer ArXiv version is at http://arxiv.org/abs/1611.00665, SIAM, 2017, pp. 1671–1687.

[21] C. Chekuri and V. Livanos, "On submodular prophet inequalities and correlation gap," in *Algorithmic Game Theory - 14th International Symposium, SAGT 2021, Aarhus, Denmark, September 21-24, 2021, Proceedings*, I. Caragiannis and K. A. Hansen, Eds., ser. Lecture Notes in Computer Science, vol. 12885, Springer, 2021, p. 410. [Online]. Available: https://link.springer.com/content/pdf/bbm%5C%3A978-3-030-85947-3%5C%2F1.pdf.

[22] J. Jiang, W. Ma, and J. Zhang, *Tightness without counterexamples: A new approach and new results for prophet inequalities*, 2022. DOI: 10.48550/ARXIV.2205.00588. [Online]. Available: https://arxiv.org/abs/2205.00588.

[23] P. D. Azar, R. Kleinberg, and S. M. Weinberg, "Prior independent mechanisms via prophet inequalities with limited information," *Games Econ. Behav.*, vol. 118, pp. 511–532, 2019. DOI: 10.1016/j.geb.2018.05.006. [Online]. Available: https://doi.org/10.1016/j.geb.2018.05.006.

[24] J. R. Correa, P. Dütting, F. A. Fischer, and K. Schewior, "Prophet inequalities for independent and identically distributed random variables from an unknown distribution," *Math. Oper. Res.*, vol. 47, no. 2, pp. 1287–1309, 2022. DOI: 10.1287/moor.2021.1167. [Online]. Available: https://doi.org/10.1287/moor.2021.1167.

[25] A. Rubinstein, J. Z. Wang, and S. M. Weinberg, "Optimal single-choice prophet inequalities from samples," in *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, T. Vidick, Ed., ser. LIPIcs, vol. 151, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 60:1–60:10. DOI: 10.4230/LIPIcs.ITCS.2020.60. [Online]. Available: https://doi.org/10.4230/LIPIcs.ITCS.2020.60.

[26] C. Chekuri, J. Vondrák, and R. Zenklusen, "Submodular function maximization via the multilinear relaxation and contention resolution schemes," *SIAM J. Comput.*, vol. 43, no. 6, pp. 1831–1879, 2014. DOI: 10.1137/110839655. [Online]. Available: https://doi.org/10.1137/110839655.

[27] M. Feldman, O. Svensson, and R. Zenklusen, "Online contention resolution schemes with applications to bayesian selection problems," *SIAM J. Comput.*, vol. 50, no. 2, pp. 255–300, 2021. DOI: 10.1137/18M1226130. [Online]. Available: https://doi.org/10.1137/18M1226130.

[28] M. Adamczyk and M. Wlodarczyk, "Random order contention resolution schemes," in *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, 2018, pp. 790–801. DOI: 10.1109/FOCS.2018.00080. [Online]. Available: https://doi.org/10.1109/FOCS.2018.00080.

[29] T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang, "Prophet matching with general arrivals," *Mathematics of Operations Research*, vol. 47, no. 2, pp. 878–898, 2022. DOI: 10.1287/moor.2021.1152. eprint: https://doi.org/10.1287/moor.2021.1152. [Online]. Available: https://doi.org/10.1287/moor.2021.1152.

[30] S. Bruggmann and R. Zenklusen, "An optimal monotone contention resolution scheme for bipartite matchings via a polyhedral viewpoint," *Mathematical Programming*, Aug. 2020, ISSN: 1436-4646. DOI: 10.1007/s10107-020-01570-6. [Online]. Available: https://doi.org/10.1007/s10107-020-01570-6.

[31] L. Haan and A. Ferreira, *Extreme Value Theory: An Introduction*. Springer, Jan. 2006, ISBN: 978-0-387-23946-0. DOI: 10.1007/0-387-34471-3.

[32] E. Samuel-Cahn, "Comparison of threshold stop rules and maximum for independent nonnegative random variables," *The Annals of Probability*, vol. 12, no. 4, pp. 1213–1216, 1984, ISSN: 00911798. [Online]. Available: http://www.jstor.org/stable/2243359.

[33] J. Brustle, J. Correa, P. Dütting, and V. Verdugo, "The competition complexity of dynamic pricing," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, ser. EC '22, Boulder, CO, USA: Association for Computing Machinery, 2022, pp. 303–320, ISBN: 9781450391504. DOI: 10.1145/3490486.3538366. [Online]. Available: https://doi.org/10.1145/3490486.3538366.

[34] K. Gatmiry, T. Kesselheim, S. Singla, and Y. Wang, "Bandit algorithms for prophet inequality and pandora's box," *CoRR*, vol. abs/2211.08586, 2022. DOI: 10.48550/arXiv.2211.08586. arXiv: 2211.08586. [Online]. Available: https://doi.org/10.48550/arXiv.2211.08586.

[35] K. Anand, R. Ge, A. Kumar, and D. Panigrahi, "Online algorithms with multiple predictions," in *Proceedings of the 39th International Conference on Machine Learning*, K. Chaudhuri, S. Jegelka, L. Song, C. Szepesvari, G. Niu, and S. Sabato, Eds., ser. Proceedings of Machine Learning Research, vol. 162, PMLR, Jul. 2022, pp. 582–598. [Online]. Available: https://proceedings.mlr.press/v162/anand22a.html.

[36] S. Banerjee, V. Cohen-Addad, A. Gupta, and Z. Li, *Graph searching with predictions*, 2022. DOI: 10.48550/ARXIV.2212.14220. [Online]. Available: https://arxiv.org/abs/2212.14220.

[37] A. Antoniadis, C. Coester, M. Elias, A. Polak, and B. Simon, "Learningaugmented dynamic power management with multiple states via new ski rental bounds," in *NIPS_2021*, M. Ranzato, A. Beygelzimer, Y. Dauphin, P. Liang, and J. W. Vaughan, Eds., vol. 34, Curran Associates, Inc., 2021, pp. 16 714–16 726. [Online]. Available: https://proceedings.neurips.cc/paper/2021/file/8b8388180314a337c9aa3c5aa8e2f37a-Paper.pdf.

[38] T. Lykouris and S. Vassilvitskii, "Competitive caching with machine learned advice," *J. ACM*, vol. 68, no. 4, Jul. 2021, ISSN: 0004-5411. DOI: 10.1145/3447579. [Online]. Available: https://doi.org/10.1145/3447579.

[39] E. B. Dynkin, "The optimum choice of the instant for stopping a Markov process," *Soviet Math. Dokl*, vol. 4, 1963.

[40] J. P. Gilbert and F. Mosteller, "Recognizing the maximum of a sequence," *Journal of the American Statistical Association*, vol. 61, no. 313, pp. 35–73, 1966, ISSN: 01621459. [Online]. Available: http://www.jstor.org/stable/2283044.

[41] P. Nuti, "The secretary problem with distributions," in *Integer Programming and Combinatorial Optimization: 23rd International Conference, IPCO 2022, Eindhoven, The Netherlands, June 27–29, 2022, Proceedings*, Eindhoven, The Netherlands: Springer-Verlag, 2022, pp. 429–439, ISBN: 978-3-031-06900-0. DOI: 10.1007/978-3-031-06901-7_32. [Online]. Available: https://doi.org/10.1007/978-3-031-06901-7_32.

[42] D. Assaf and E. Samuel-Cahn, "Simple ratio prophet inequalities for a mortal with multiple choices," *Journal of Applied Probability*, vol. 37, no. 4, pp. 1084–1091, 2000, ISSN: 00219002. [Online]. Available: http://www.jstor.org/stable/3215496 (visited on 04/27/2023).

[43] T. Ezra, M. Feldman, and I. Nehama, "Prophets and secretaries with overbooking," in *Proceedings of the 2018 ACM Conference on Economics and Computation*, ser. EC '18, Ithaca, NY, USA: Association for Computing Machinery, 2018, pp. 319–320, ISBN: 9781450358293. DOI: 10.1145/3219166.3219211. [Online]. Available: https://doi.org/10.1145/3219166.3219211.

[44] E. Lee and S. Singla, "Optimal online contention resolution schemes via ex-ante prophet inequalities," English (US), in *26th European Symposium on Algorithms, ESA 2018*, H. Bast, G. Herman, and Y. Azar, Eds., ser. Leibniz International Proceedings in Informatics, LIPIcs, Germany: Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, Aug. 2018, ISBN: 9783959770811. DOI: 10.4230/LIPIcs.ESA.2018.57.

[45] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák, "Maximizing a monotone submodular function subject to a matroid constraint," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1740–1766, 2011. DOI: 10.1137/080733991. eprint: https://doi.org/10.1137/080733991. [Online]. Available: https://doi.org/10.1137/080733991.

[46] Q. Yan, "Mechanism design via correlation gap," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2011, San Francisco, California, USA, January 23-25, 2011*, D. Randall, Ed., SIAM, 2011, pp. 710–719. DOI: 10.1137/1.9781611973082.56. [Online]. Available: https://doi.org/10.1137/1.9781611973082.56.

[47] J. Vondrák, "Submodularity in combinatorial optimization," Ph.D. dissertation, Univerzita Karlova, Matematicko-fyzikální fakulta, 2007.

[48] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye, "Price of correlations in stochastic optimization," *Operations Research*, vol. 60, no. 1, pp. 150–162, 2012, ISSN: 0030364X, 15265463. [Online]. Available: http://www.jstor.org/stable/41476345 (visited on 12/27/2022).

[49] J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld, "Recent developments in prophet inequalities," *SIGecom Exch.*, vol. 17, no. 1, pp. 61–70, 2018. DOI: 10.1145/3331033.3331039. [Online]. Available: https://doi.org/10.1145/3331033.3331039.

[50] R. P. Kertz, "Stop rule and supremum expectations of i.i.d. random variables: A complete comparison by conjugate duality," *Journal of Multivariate Analysis*, vol. 19, no. 1, pp. 88–112, 1986, ISSN: 0047-259X. DOI: https://doi.org/10.1016/0047-259X(86)90095-3. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0047259X86900953.

[51] U. Saint-Mont, "A simple derivation of a complicated prophet region," *J. Multivar. Anal.*, vol. 80, no. 1, pp. 67–72, Jan. 2002, ISSN: 0047-259X. DOI: 10.1006/jmva.2000.1972. [Online]. Available: https://doi.org/10.1006/jmva.2000.1972.

[52] T. P. Hill and R. P. Kertz, "A survey of prophet inequalities in optimal stopping theory," *Contemp. Math*, vol. 125, pp. 191–207, 1992.

[53] S. Singla, "Combinatorial optimization under uncertainty: Probing and stopping-time algorithms," http://reports-archive.adm.cs.cmu.edu/anon/2018/CMU-CS-18-111.pdf, Ph.D. dissertation, CMU, 2018.

[54] E. Harb, *Fishing for better constants: The prophet secretary via poissonization*, 2023. arXiv: 2307.00971 [cs.DS].

[55] A. Bubna and A. Chiplunkar, "Prophet inequality: Order selection beats random order," in *Proceedings of the 24th ACM Conference on Economics and Computation*, ser. EC '23, London, United Kingdom: Association for Computing Machinery, 2023, pp. 302–336. DOI: 10.1145/3580507.3597687. [Online]. Available: https://doi-org.proxy2.library.illinois.edu/10.1145/3580507.3597687.

[56] T. Hill, "Prophet inequalities and order selection in optimal stopping problems," *Proceedings of the American Mathematical Society*, vol. 88, no. 1, pp. 131–137, 1983.

[57] H. Beyhaghi, N. Golrezaei, R. P. Leme, M. Pál, and B. Sivan, "Improved revenue bounds for posted-price and second-price mechanisms," *Operations Research*, vol. 69, no. 6, pp. 1805–1822, 2021. DOI: 10.1287/opre.2021.2121. eprint: https://doi.org/10.1287/opre.2021.2121. [Online]. Available: https://doi.org/10.1287/opre.2021.2121.

[58] M. Feldman, N. Gravin, and B. Lucier, "Combinatorial auctions via posted prices," in *Proceedings of the 2015 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2015, pp. 123–135. DOI: 10.1137/1.9781611973730.10. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611973730.10. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611973730.10.

[59] S. Assadi, T. Kesselheim, and S. Singla, "Improved truthful mechanisms for subadditive combinatorial auctions: Breaking the logarithmic barrier," in *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2021, pp. 653–661. DOI: 10.1137/1.9781611976465.40. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611976465.40. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611976465.40.

[60] J. Correa and A. Cristi, "A constant factor prophet inequality for online combinatorial auctions," in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, ser. STOC 2023, Orlando, FL, USA, 2023, pp. 686–697, ISBN: 9781450399135. DOI: 10.1145/3564246.3585151. [Online]. Available: https://doi.org/10.1145/3564246.3585151.

[61] M. Babaioff, M. Dinitz, A. Gupta, N. Immorlica, and K. Talwar, "Secretary problems: Weights and discounts," in *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '09, New York, New York: Society for Industrial and Applied Mathematics, 2009, pp. 1245–1254.

[62] N. B. Dimitrov and C. G. Plaxton, "Competitive weighted matching in transversal matroids," *Algorithmica*, vol. 62, no. 1, pp. 333–348, Feb. 2012, ISSN: 1432-0541. DOI: 10.1007/s00453-010-9457-2. [Online]. Available: https://doi.org/10.1007/s00453-010-9457-2.

[63] M. Babaioff, N. Immorlica, and R. Kleinberg, "Matroids, secretary problems, and online mechanisms," in *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '07, New Orleans, Louisiana: Society for Industrial and Applied Mathematics, 2007, pp. 434–443, ISBN: 9780898716245.

[64] N. Korula and M. Pál, "Algorithms for secretary problems on graphs and hypergraphs," in *Automata, Languages and Programming*, S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikoletseas, and W. Thomas, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 508–520, ISBN: 978-3-642-02930-1.

[65] M. Dinitz and G. Kortsarz, "Matroid secretary for regular and decomposable matroids," *SIAM Journal on Computing*, vol. 43, no. 5, pp. 1807–1830, 2014. DOI: 10.1137/13094030X. eprint: https://doi.org/10.1137/13094030X. [Online]. Available: https://doi.org/10.1137/13094030X.

[66] S. Im and Y. Wang, "Secretary problems: Laminar matroid and interval scheduling," in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '11, San Francisco, California: Society for Industrial and Applied Mathematics, 2011, pp. 1265–1274.

[67] D. Abdolazimi, A. R. Karlin, N. Klein, and S. Oveis Gharan, "Matroid Partition Property and the Secretary Problem," in *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, Y. Tauman Kalai, Ed., ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 251, Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 2:1–2:9, ISBN: 978-3-95977-263-1. DOI: 10.4230/LIPIcs.ITCS.2023.2. [Online]. Available: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.ITCS.2023.2.

[68] P. Jaillet, J. A. Soto, and R. Zenklusen, "Advances on matroid secretary problems: Free order model and laminar case," in *Integer Programming and Combinatorial Optimization - 16th International Conference, IPCO 2013, Valparaíso, Chile, March 18-20, 2013. Proceedings*, M. X. Goemans and J. R. Correa, Eds., ser. Lecture Notes in Computer Science, vol. 7801, Springer, 2013, pp. 254–265. DOI: 10.1007/978-3-642-36694-9_22. [Online]. Available: https://doi.org/10.1007/978-3-642-36694-9%5C_22.

[69] T. Ma, B. Tang, and Y. Wang, "The Simulated Greedy Algorithm for Several Submodular Matroid Secretary Problems," in *30th International Symposium on Theoretical Aspects of Computer Science (STACS 2013)*, N. Portier and T. Wilke, Eds., ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 20, Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2013, pp. 478–489, ISBN: 978-3-939897-50-7. DOI: 10.4230/LIPIcs.STACS.2013.478. [Online]. Available: https://drops.dagstuhl.de/entities/document/10.4230/LIPIcs.STACS.2013.478.

[70] J. A. Soto, "Matroid secretary problem in the random-assignment model," *SIAM Journal on Computing*, vol. 42, no. 1, pp. 178–211, 2013. DOI: 10.1137/110852061. eprint: https://doi.org/10.1137/110852061. [Online]. Available: https://doi.org/10.1137/110852061.

[71] Z. Huang, Z. Parsaeian, and Z. Zhu, "Laminar matroid secretary: Greedy strikes back," *CoRR*, vol. abs/2308.09880, 2023. DOI: 10.48550/ARXIV.2308.09880. arXiv: 2308.09880. [Online]. Available: https://doi.org/10.48550/arXiv.2308.09880.

[72] M. Dinitz, "Recent advances on the matroid secretary problem," *SIGACT News*, vol. 44, no. 2, pp. 126–142, Jun. 2013, ISSN: 0163-5700. DOI: 10.1145/2491533.2491557. [Online]. Available: https://doi.org/10.1145/2491533.2491557.

[73] S. Alaei, "Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers," *SIAM Journal on Computing*, vol. 43, no. 2, pp. 930–972, 2014. DOI: 10.1137/120878422. eprint: https://doi.org/10.1137/120878422. [Online]. Available: https://doi.org/10.1137/120878422.

[74] S. Chawla, N. R. Devanur, and T. Lykouris, "Static pricing for multi-unit prophet inequalities (extended abstract)," in *Web and Internet Economics - 17th International Conference, WINE 2021, Potsdam, Germany, December 14-17, 2021, Proceedings*, M. Feldman, H. Fu, and I. Talgam-Cohen, Eds., ser. Lecture Notes in Computer Science, vol. 13112, Springer, 2021, pp. 545–546. [Online]. Available: https://link.springer.com/content/pdf/bbm%5C%3A978-3-030-94676-0%5C%2F1.

[75] J. Jiang, W. Ma, and J. Zhang, "Tight guarantees for multi-unit prophet inequalities and online stochastic knapsack," in *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 2022, pp. 1221–1246. DOI: 10.1137/1.9781611977073.51. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611977073.51. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611977073.51.

[76] A. Dinev and S. M. Weinberg, *Simple and optimal online contention resolution schemes for k-uniform matroids*, 2023. arXiv: 2309.10078 `[cs.DS]`.

[77] A. Rubinstein, "Beyond matroids: Secretary problem and prophet inequality with general constraints," in *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, ser. STOC '16, Cambridge, MA, USA: Association for Computing Machinery, 2016, pp. 324–332, ISBN: 9781450341325. DOI: 10.1145/2897518.2897540. [Online]. Available: https://doi.org/10.1145/2897518.2897540.

[78] M. Adamczyk and M. Wlodarczyk, “Multi-dimensional mechanism design via random order contention resolution schemes,” *SIGecom Exch.*, vol. 17, no. 2, pp. 46–53, Jan. 2020. DOI: 10.1145/3381329.3381334. [Online]. Available: https://doi.org/10.1145/3381329.3381334.

[79] B. Brubach, N. Grammel, W. Ma, and A. Srinivasan, “Improved guarantees for offline stochastic matching via new ordered contention resolution schemes,” in *NeurIPS*, 2021.

[80] T. Pollner, M. Roghani, A. Saberi, and D. Wajc, “Improved online contention resolution for matchings and applications to the gig economy,” in *EC ’22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds., ACM, 2022, pp. 321–322. DOI: 10.1145/3490486.3538295. [Online]. Available: https://doi.org/10.1145/3490486.3538295.

[81] C. MacRury, W. Ma, and N. Grammel, “On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs,” *CoRR*, vol. abs/2209.07520, 2022. DOI: 10.48550/arXiv.2209.07520. arXiv: 2209.07520. [Online]. Available: https://doi.org/10.48550/arXiv.2209.07520.

[82] A. Gupta and S. Singla, “Random-order models,” in *Beyond the Worst-Case Analysis of Algorithms*. Cambridge University Press, 2021, pp. 234–258. DOI: 10.1017/9781108637435.015.

[83] M. Adamczyk, M. Sviridenko, and J. Ward, “Submodular stochastic probing on matroids,” *Mathematics of Operations Research*, vol. 41, no. 3, pp. 1022–1038, 2016. DOI: 10.1287/moor.2015.0766. [Online]. Available: https://doi.org/10.1287/moor.2015.0766.

[84] A. Gupta and V. Nagarajan, “A stochastic probing problem with applications,” in *Integer Programming and Combinatorial Optimization*, M. Goemans and J. Correa, Eds., Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 205–216, ISBN: 978-3-642-36694-9.

[85] A. Gupta, V. Nagarajan, and S. Singla, “Algorithms and adaptivity gaps for stochastic probing,” in *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, R. Krauthgamer, Ed., SIAM, 2016, pp. 1731–1747. DOI: 10.1137/1.9781611974331.ch120. [Online]. Available: https://doi.org/10.1137/1.9781611974331.ch120.

[86] A. Gupta, V. Nagarajan, and S. Singla, “Adaptivity gaps for stochastic probing: Submodular and XOS functions,” in *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, P. N. Klein, Ed., SIAM, 2017, pp. 1688–1702. DOI: 10.1137/1.9781611974782.111. [Online]. Available: https://doi.org/10.1137/1.9781611974782.111.

[87] D. Bradac, S. Singla, and G. Zuzic, “(Near) Optimal Adaptivity Gaps for Stochastic Multi-Value Probing,” in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*, D. Achlioptas and L. A. Végh, Eds., ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 145, Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2019, 49:1–49:21, ISBN: 978-3-95977-125-2. DOI: 10.4230/LIPIcs.APPROX-RANDOM.2019.49. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2019/11264.

[88] A. Asadpour and H. Nazerzadeh, “Maximizing stochastic monotone submodular functions,” *Management Science*, vol. 62, no. 8, pp. 2374–2391, 2016. DOI: 10.1287/mnsc.2015.2254. eprint: https://doi.org/10.1287/mnsc.2015.2254. [Online]. Available: https://doi.org/10.1287/mnsc.2015.2254.

[89] P. Briest, S. Chawla, R. Kleinberg, and S. M. Weinberg, "Pricing lotteries," *Journal of Economic Theory*, vol. 156, pp. 144–174, 2015, Computer Science and Economic Theory, ISSN: 0022-0531. DOI: https://doi.org/10.1016/j.jet.2014.04.011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0022053114000659.

[90] S. Hart and N. Nisan, "Selling multiple correlated goods: Revenue maximization and menu-size complexity," *J. Econ. Theory*, vol. 183, pp. 991–1029, 2019. DOI: 10.1016/j.jet.2019.07.006. [Online]. Available: https://doi.org/10.1016/j.jet.2019.07.006.

[91] C. Daskalakis, A. Deckelbaum, and C. Tzamos, "Strong duality for a multiple-good monopolist," in *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15, Portland, OR, USA, June 15-19, 2015*, T. Roughgarden, M. Feldman, and M. Schwarz, Eds., ACM, 2015, pp. 449–450. DOI: 10.1145/2764468.2764539. [Online]. Available: https://doi.org/10.1145/2764468.2764539.

[92] C. Daskalakis, A. Deckelbaum, and C. Tzamos, "The complexity of optimal mechanism design," in *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, C. Chekuri, Ed., SIAM, 2014, pp. 1302–1318. DOI: 10.1137/1.9781611973402.96. [Online]. Available: https://doi.org/10.1137/1.9781611973402.96.

[93] S. Chawla, J. D. Hartline, and R. Kleinberg, "Algorithmic pricing via virtual valuations," in *Proceedings of the 8th ACM Conference on Electronic Commerce*, ser. EC '07, San Diego, California, USA: Association for Computing Machinery, 2007, pp. 243–251, ISBN: 9781595936530. DOI: 10.1145/1250910.1250946. [Online]. Available: https://doi.org/10.1145/1250910.1250946.

[94] L. Blumrosen and T. Holenstein, "Posted prices vs. negotiations: An asymptotic analysis," in *Proceedings of the 9th ACM Conference on Electronic Commerce*, ser. EC '08, Chicago, Il, USA: Association for Computing Machinery, 2008, p. 49, ISBN: 9781605581699. DOI: 10.1145/1386790.1386801. [Online]. Available: https://doi.org/10.1145/1386790.1386801.

[95] M. Adamczyk, A. Borodin, D. Ferraioli, B. D. Keijzer, and S. Leonardi, "Sequential posted-price mechanisms with correlated valuations," *ACM Trans. Econ. Comput.*, vol. 5, no. 4, Dec. 2017, ISSN: 2167-8375. DOI: 10.1145/3157085. [Online]. Available: https://doi.org/10.1145/3157085.

[96] A. Braun, M. Buttkus, and T. Kesselheim, "Asymptotically optimal welfare of posted pricing for multiple items with MHR distributions," in *29th Annual European Symposium on Algorithms, ESA 2021, September 6-8, 2021, Lisbon, Portugal (Virtual Conference)*, P. Mutzel, R. Pagh, and G. Herman, Eds., ser. LIPIcs, vol. 204, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 22:1–22:16. DOI: 10.4230/LIPIcs.ESA.2021.22. [Online]. Available: https://doi.org/10.4230/LIPIcs.ESA.2021.22.

[97] S. Dobzinski, "Breaking the logarithmic barrier for truthful combinatorial auctions with submodular bidders," *SIAM J. Comput.*, vol. 50, no. 3, 2021. DOI: 10.1137/16M1088594. [Online]. Available: https://doi.org/10.1137/16M1088594.

[98] S. Dobzinski and J. Vondrák, "Impossibility results for truthful combinatorial auctions with submodular valuations," *J. ACM*, vol. 63, no. 1, 5:1–5:19, 2016. DOI: 10.1145/2786754. [Online]. Available: https://doi.org/10.1145/2786754.

[99] S. Assadi and S. Singla, "Improved truthful mechanisms for combinatorial auctions with submodular bidders," *SIGecom Exch.*, vol. 18, no. 1, pp. 19–27, 2020. DOI: 10.1145/3440959.3440964. [Online]. Available: https://doi.org/10.1145/3440959.3440964.

[100] B. Lucier, "An economic view of prophet inequalities," *SIGecom Exch.*, vol. 16, no. 1, pp. 24–47, Sep. 2017. DOI: 10.1145/3144722.3144725. [Online]. Available: https://doi.org/10.1145/3144722.3144725.

[101] V. Livanos and R. Mehta, "Prophet inequalities for cost minimization," *CoRR*, vol. abs/2209.07988, 2022. DOI: 10.48550/arXiv.2209.07988. arXiv: 2209.07988. [Online]. Available: https://doi.org/10.48550/arXiv.2209.07988.

[102] V. Livanos, "A unified theory of i.i.d. prophet inequalities," *CoRR*, 2023.

[103] V. Livanos and V. Verdugo, "Competition complexity for large markets," *CoRR*, 2023.

[104] S. Har-Peled, E. Harb, and V. Livanos, "Prophets and oracles," *CoRR*, 2023.

[105] V. Livanos, "Simple and optimal greedy online contention resolution schemes," in *Advances in Neural Information Processing Systems*, S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, Eds., vol. 35, Curran Associates, Inc., 2022, pp. 9687–9699. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/file/3f45a1768ed452c1706e0a3a272e3bbc-Paper-Conference.pdf.

[106] N. Gravin and H. Wang, "Prophet inequality for bipartite matching: Merits of being simple and nonadaptive," *Mathematics of Operations Research*, vol. 48, no. 1, pp. 38–52, 2023. DOI: 10.1287/moor.2022.1251. eprint: https://doi.org/10.1287/moor.2022.1251. [Online]. Available: https://doi.org/10.1287/moor.2022.1251.

[107] B. Lucier, personal communication, Sep. 20, 2022.

[108] P. Leonetti and A. K. Chokami, "The maximum domain of attraction of multivariate extreme value distributions is small," *Electronic Communications in Probability*, vol. 27, no. none, pp. 1–8, 2022. DOI: 10.1214/22-ECP501. [Online]. Available: https://doi.org/10.1214/22-ECP501.

[109] R. A. Fisher and L. H. C. Tippett, "Limiting forms of the frequency distribution of the largest or smallest member of a sample," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 24, no. 2, pp. 180–190, 1928. DOI: 10.1017/S0305004100015681.

[110] B. Gnedenko, "Sur la distribution limite du terme maximum d'une serie aleatoire," *Annals of Mathematics*, vol. 44, no. 3, pp. 423–453, 1943, ISSN: 0003486X. [Online]. Available: http://www.jstor.org/stable/1968974 (visited on 11/13/2023).

[111] Y. Giannakopoulos, D. Poças, and K. Zhu, "Optimal pricing for mhr and λ-regular distributions," *ACM Trans. Econ. Comput.*, vol. 9, no. 1, Jan. 2021, ISSN: 2167-8375. DOI: 10.1145/3434423. [Online]. Available: https://doi.org/10.1145/3434423.

[112] M. Babaioff, L. Blumrosen, S. Dughmi, and Y. Singer, "Posting prices with unknown distributions," *ACM Trans. Econ. Comput.*, vol. 5, no. 2, Mar. 2017, ISSN: 2167-8375. DOI: 10.1145/3037382. [Online]. Available: https://doi.org/10.1145/3037382.

[113] S. Bhattacharya, G. Goel, S. Gollapudi, and K. Munagala, "Budget-constrained auctions with heterogeneous items," *Theory Comput.*, vol. 8, no. 1, pp. 429–460, 2012. DOI: 10.4086/toc.2012.v008a020. [Online]. Available: https://doi.org/10.4086/toc.2012.v008a020.

[114] Y. Cai and C. Daskalakis, "Extreme value theorems for optimal multidimensional pricing," *Games Econ. Behav.*, vol. 92, pp. 266–305, 2015. DOI: 10.1016/j.geb.2015.02.003. [Online]. Available: https://doi.org/10.1016/j.geb.2015.02.003.

[115] C. Daskalakis and S. M. Weinberg, "Symmetries and optimal multi-dimensional mechanism design," in *Proceedings of the 13th ACM Conference on Electronic Commerce*, ser. EC '12, Valencia, Spain: Association for Computing Machinery, 2012, pp. 370–387, ISBN: 9781450314152. DOI: 10.1145/2229012.2229042. [Online]. Available: https://doi.org/10.1145/2229012.2229042.

[116] P. Dhangwatnotai, T. Roughgarden, and Q. Yan, "Revenue maximization with a single sample," *Games and Economic Behavior*, vol. 91, pp. 318–333, 2015, ISSN: 0899-8256. DOI: https://doi.org/10.1016/j.geb.2014.03.011. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0899825614000645.

[117] Y. Giannakopoulos, E. Koutsoupias, and P. Lazos, "Online Market Intermediation," in *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, I. Chatzigiannakis, P. Indyk, F. Kuhn, and A. Muscholl, Eds., ser. Leibniz International Proceedings in Informatics (LIPIcs), vol. 80, Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, 47:1–47:14, ISBN: 978-3-95977-041-5. DOI: 10.4230/LIPIcs.ICALP.2017.47. [Online]. Available: http://drops.dagstuhl.de/opus/volltexte/2017/7481.

[118] J. D. Hartline and T. Roughgarden, "Simple versus optimal mechanisms," *SIGecom Exch.*, vol. 8, no. 1, 2009. DOI: 10.1145/1598780.1598785. [Online]. Available: https://doi.org/10.1145/1598780.1598785.

[119] J. Correa, A. Cristi, P. Dütting, M. Hajiaghayi, J. Olkowski, and K. Schewior, *Trading prophets*, 2023. arXiv: 2305.15566 `[cs.DS]`.

[120] S. Resnick, *Extreme Values, Regular Variation and Point Processes*. Springer, 2014, ISBN: 9781489908360. [Online]. Available: https://books.google.com/books?id=JzQjswEACAAJ.

[121] R. Barlow, F. Proschan, and L. Hunter, *Mathematical Theory of Reliability* (Classics in Applied Mathematics). Society for Industrial and Applied Mathematics, 1996, ISBN: 9780898713695. [Online]. Available: https://books.google.com/books?id=08PBufZtyWwC.

[122] N. H. Bingham, C. M. Goldie, and J. L. Teugels, *Regular Variation* (Encyclopedia of Mathematics and its Applications). Cambridge University Press, 1987. DOI: 10.1017/CBO9780511721434.

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

[124] R. Cole and T. Roughgarden, "The sample complexity of revenue maximization," in *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing*, ser. STOC '14, New York, New York: Association for Computing Machinery, 2014, pp. 243–252, ISBN: 9781450327107. DOI: 10.1145/2591796.2591867. [Online]. Available: https://doi.org/10.1145/2591796.2591867.

[125] A. M. Medina and S. Vassilvitskii, "Revenue optimization with approximate bid predictions," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, Long Beach, California, USA: Curran Associates Inc., 2017, pp. 1856–1864, ISBN: 9781510860964.

[126] N. R. Devanur and T. P. Hayes, "The adwords problem: Online keyword matching with budgeted bidders under random permutations," in *Proceedings of the 10th ACM Conference on Electronic Commerce*, ser. EC '09, Stanford, California, USA: Association for Computing Machinery, 2009, pp. 71–78, ISBN: 9781605584584. DOI: 10.1145/1566374.1566384. [Online]. Available: https://doi.org/10.1145/1566374.1566384.

[127] E. Vee, S. Vassilvitskii, and J. Shanmugasundaram, "Optimal online assignment with forecasts," in *Proceedings of the 11th ACM Conference on Electronic Commerce*, ser. EC '10, Cambridge, Massachusetts, USA: Association for Computing Machinery, 2010, pp. 109–118, ISBN: 9781605588223. DOI: 10.1145/1807342.1807360. [Online]. Available: https://doi.org/10.1145/1807342.1807360.

[128] M. Mitzenmacher and S. Vassilvitskii, "Algorithms with predictions," *Commun. ACM*, vol. 65, no. 7, pp. 33–35, Jun. 2022, ISSN: 0001-0782. DOI: 10.1145/3528087. [Online]. Available: https://doi.org/10.1145/3528087.

[129] D. Assaf, L. Goldstein, and E. Samuel-Cahn, "Ratio prophet inequalities when the mortal has several choices," *The Annals of Applied Probability*, vol. 12, no. 3, pp. 972–984, 2002.

[130] T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang, ""who is next in line?" on the significance of knowing the arrival order in bayesian online settings," in *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, N. Bansal and V. Nagarajan, Eds., SIAM, 2023, pp. 3759–3776. DOI: 10.1137/1.9781611977554.ch145. [Online]. Available: https://doi.org/10.1137/1.9781611977554.ch145.

[131] J. Correa, A. Cristi, P. Duetting, and A. Norouzi-Fard, "Fairness and bias in online selection," in *Proceedings of the 38th International Conference on Machine Learning*, M. Meila and T. Zhang, Eds., ser. Proceedings of Machine Learning Research, vol. 139, PMLR, Jul. 2021, pp. 2112–2121. [Online]. Available: https://proceedings.mlr.press/v139/correa21a.html.

[132] J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto, "The two-sided game of googol and sample-based prophet inequalities," in *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 2066–2081. DOI: 10.1137/1.9781611975994.127. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611975994.127. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611975994.127.

[133] D. P. Dubhashi and A. Panconesi, *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, 2009, ISBN: 978-0-521-88427-3. [Online]. Available: http://www.cambridge.org/gb/knowledge/isbn/item2327542/.

[134] F. Den Hollander, "Probability theory: The coupling method,"

[135] L. L. Cam, "An approximation theorem for the poisson binomial distribution.," *Pacific Journal of Mathematics*, vol. 10, pp. 1181–1197, 1960.

[136] C. Bechtel, S. Dughmi, and N. Patel, "Delegated pandora's box," in *EC '22: The 23rd ACM Conference on Economics and Computation, Boulder, CO, USA, July 11 - 15, 2022*, D. M. Pennock, I. Segal, and S. Seuken, Eds., ACM, 2022, pp. 666–693. DOI: 10.1145/3490486.3538267. [Online]. Available: https://doi.org/10.1145/3490486.3538267.

[137] H. Fu, P. Lu, Z. G. Tang, *et al.*, *Oblivious online contention resolution schemes*, 2021. arXiv: 2111.10607 [cs.DS].

[138] S. Dughmi, T. Roughgarden, and M. Sundararajan, "Revenue submodularity," in *Proceedings of the 10th ACM Conference on Electronic Commerce*, ser. EC '09, Stanford, California, USA: Association for Computing Machinery, 2009, pp. 243–252, ISBN: 9781605584584. DOI: 10.1145/1566374.1566409. [Online]. Available: https://doi.org/10.1145/1566374.1566409.

[139] F. Constantin, J. Feldman, S. Muthukrishnan, and M. Pál, "An online mechanism for ad slot reservations with cancellations," in *Proceedings of the 2009 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 1265–1274. DOI: 10.1137/1.9781611973068.137. eprint: https://epubs.siam.org/doi/pdf/10.1137/1.9781611973068.137. [Online]. Available: https://epubs.siam.org/doi/abs/10.1137/1.9781611973068.137.

[140] S. Dughmi, "The outer limits of contention resolution on matroids and connections to the secretary problem," in *47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference)*, A. Czumaj, A. Dawar, and E. Merelli, Eds., ser. LIPIcs, vol. 168, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 42:1–42:18. DOI: 10.4230/LIPIcs.ICALP.2020.42. [Online]. Available: https://doi.org/10.4230/LIPIcs.ICALP.2020.42.

[141] S. Dughmi, "Matroid secretary is equivalent to contention resolution," in *13th Innovations in Theoretical Computer Science Conference, ITCS 2022, January 31 - February 3, 2022, Berkeley, CA, USA*, M. Braverman, Ed., ser. LIPIcs, vol. 215, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 58:1–58:23. DOI: 10.4230/LIPIcs.ITCS.2022.58. [Online]. Available: https://doi.org/10.4230/LIPIcs.ITCS.2022.58.

[142] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg, "Matroid secretary problems," *J. ACM*, vol. 65, no. 6, Nov. 2018, ISSN: 0004-5411. DOI: 10.1145/3212512. [Online]. Available: https://doi.org/10.1145/3212512.

[143] J. Vondrák, personal communication, Jun. 30, 2021.

[144] F. Qiu and S. Singla, "Submodular dominance and applications," in *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2022, September 19-21, 2022, University of Illinois, Urbana-Champaign, USA (Virtual Conference)*, A. Chakrabarti and C. Swamy, Eds., ser. LIPIcs, vol. 245, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, 44:1–44:21. DOI: 10.4230/LIPIcs.APPROX/RANDOM.2022.44. [Online]. Available: https://doi.org/10.4230/LIPIcs.APPROX/RANDOM.2022.44.

[145] M. Feldman, J. (Naor, and R. Schwartz, "A unified continuous greedy algorithm for submodular maximization," in *Proceedings of the 2011 IEEE 52Nd Annual Symposium on Foundations of Computer Science*, ser. FOCS '11, Washington, DC, USA: IEEE Computer Society, 2011, pp. 570–579, ISBN: 978-0-7695-4571-4. DOI: 10.1109/FOCS.2011.46. [Online]. Available: http://dx.doi.org/10.1109/FOCS.2011.46.

[146] N. Buchbinder and M. Feldman, "Submodular functions maximization problems," in *Handbook of Approximation Algorithms and Metaheuristics: Methologies and Traditional Applications, Volume 1*, T. F. Gonzalez, Ed., 2nd, CRC Press, 2018. DOI: 10.1201/9781351236423. [Online]. Available: https://doi.org/10.1201/9781351236423.

[147] N. Buchbinder, M. Feldman, J. (Naor, and R. Schwartz, "Submodular maximization with cardinality constraints," in *Proceedings of the Twenty-fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '14, Portland, Oregon: Society for Industrial and Applied Mathematics, 2014, pp. 1433–1452, ISBN: 978-1-611973-38-9. [Online]. Available: http://dl.acm.org/citation.cfm?id=2634074.2634180.

[148] J. Kleinberg and R. Kleinberg, "Delegated search approximates efficient search," in *Proceedings of the 2018 ACM Conference on Economics and Computation*, ser. EC '18, Ithaca, NY, USA: Association for Computing Machinery, 2018, pp. 287–302, ISBN: 9781450358293. DOI: 10.1145/3219166.3219205. [Online]. Available: https://doi.org/10.1145/3219166.3219205.

[149] M. Hajiaghayi, K. Rezaei, and S. Shin, *Delegating to multiple agents*, 2023. arXiv: 2305.03203 [cs.GT].

[150] N. Immorlica, S. Singla, and B. Waggoner, "Prophet inequalities with linear correlations and augmentations," in *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, P. Biró, J. D. Hartline, M. Ostrovsky, and A. D. Procaccia, Eds., ACM, 2020, pp. 159–185. DOI: 10.1145/3391403.3399452. [Online]. Available: https://doi.org/10.1145/3391403.3399452.

[151] I. Caragiannis, N. Gravin, P. Lu, and Z. Wang, "Relaxing the independence assumption in sequential posted pricing, prophet inequality, and random bipartite matching," in *Web and Internet Economics: 17th International Conference, WINE 2021, Potsdam, Germany, December 14–17, 2021, Proceedings*, Potsdam, Germany: Springer-Verlag, 2021, pp. 131–148, ISBN: 978-3-030-94675-3. DOI: 10.1007/978-3-030-94676-0_8. [Online]. Available: https://doi.org/10.1007/978-3-030-94676-0_8.

[152] Y. Cai and A. Oikonomou, "On simple mechanisms for dependent items," in *Proceedings of the 22nd ACM Conference on Economics and Computation*, ser. EC '21, Budapest, Hungary: Association for Computing Machinery, 2021, pp. 242–262, ISBN: 9781450385541. DOI: 10.1145/3465456.3467643. [Online]. Available: https://doi.org/10.1145/3465456.3467643.

[153] S. Dughmi, Y. H. Kalayci, and N. Patel, *Limitations of stochastic selection problems with pairwise independent priors*, 2023. arXiv: 2310.05240 [cs.DS].

[154] B. Peng and Z. Tang, "Order selection prophet inequality: From threshold optimization to arrival time design," in *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, Los Alamitos, CA, USA: IEEE Computer Society, Nov. 2022, pp. 171–178. DOI: 10.1109/FOCS54457.2022.00023. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/FOCS54457.2022.00023.