



Prophet Inequalities with Cancellation Costs

Farbod Ekbatani

University of Chicago
Chicago, USA
fekbatan@chicagobooth.edu

Rad Niazadeh

University of Chicago
Chicago, USA
rad.niazadeh@chicagobooth.edu

Pranav Nuti

Stanford University
Stanford, USA
pranavn@stanford.edu

Jan Vondrák

Stanford University
Stanford, USA
jvondrak@stanford.edu

ABSTRACT

Most of the literature on online algorithms and sequential decision-making focuses on settings with “irrevocable decisions” where the algorithm’s decision upon arrival of the new input is set in stone and can never change in the future. One canonical example is the classic prophet inequality problem, where realizations of a sequence of independent random variables X_1, X_2, \dots with known distributions are drawn one by one and a decision maker decides when to stop and accept the arriving random variable, with the goal of maximizing the expected value of their pick. We consider “prophet inequalities with recourse” in the linear buyback cost setting, where after accepting a variable X_i , we can still discard X_i later and accept another variable X_j , at a *buyback cost* of $f \times X_i$. The goal is to maximize the expected net reward, which is the value of the final accepted variable minus the total buyback cost. Our first main result is an optimal prophet inequality in the regime of $f \geq 1$, where we prove that we can achieve an expected reward $\frac{1+f}{1+2f}$ times the expected offline optimum. The problem is still open for $0 < f < 1$ and we give some partial results in this regime. In particular, as our second main result, we characterize the asymptotic behavior of the competitive ratio for small f and provide almost matching upper and lower bounds that show a factor of $1 - \Theta\left(f \log\left(\frac{1}{f}\right)\right)$. Our results are obtained by two fundamentally different approaches: One is inspired by various proofs of the classical prophet inequality, while the second is based on combinatorial optimization techniques involving LP duality, flows, and cuts.*

CCS CONCEPTS

• **Theory of computation** → **Online algorithms; Algorithmic game theory and mechanism design.**

KEYWORDS

Online Algorithms, Prophet Inequalities, Buyback, Online Resource Allocation

*The full online version of this paper is available at [arXiv:2404.00527](https://arxiv.org/abs/2404.00527) [24].



This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649786>

ACM Reference Format:

Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. 2024. Prophet Inequalities with Cancellation Costs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, June 24–28, 2024, Vancouver, BC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3618260.3649786>

1 INTRODUCTION

Consider a monopolist seller tasked with selling a single item to a sequence of arriving buyers. A quintessential problem in online algorithms and mechanism design — with a wide range of applications from selling seats in a concert hall to the multibillion dollar online display advertising industry — is how to allocate this item in a sequential fashion to the buyers so as to maximize social welfare, that is, to maximize the willingness to pay, or the value of the allocated buyer.¹ Given distributional knowledge about the sequence of arriving values, a common approach to this problem is to design online algorithms that attain the so-called *prophet inequalities*; the goal there is to evaluate the performance of the online algorithm relative to an “omniscient prophet”, who knows the entire sequence of values and simply maximizes social welfare by allocating the item to the highest value buyer. This fundamental algorithm design question, which can alternatively be described as an optimal stopping problem, has its roots in the classic work of Krengel and Sucheston [43] in the 70s and has since been studied quite extensively in computer science, mathematics, and operations research.

This significant line of work on prophet inequalities shows that in many settings of Bayesian allocations and mechanism design (even beyond single parameter environments), optimal or near-optimal prophet inequalities can be obtained by using simple and elegant take-it-or-leave-it pricing rules. For example, the seminal work of [53] shows that a single threshold algorithm can achieve the optimal 0.5 competitive prophet inequality in a single item setting when the values are independent of each other. Other examples include, but are not limited to, obtaining more powerful prophet inequalities in special cases such as i.i.d. values or independent values arriving in a random order [1, 19, 26, 38], and extensions to more general feasibility environments such as multi-units [4, 35, 41, 55], matroids [6, 16, 42], matchings [5, 49], general downward closed

¹By standard reductions in Bayesian mechanism design, in particular replacing values with Myerson’s virtual values [47], one can reduce the problem of maximizing expected revenue to social welfare in single-parameter settings. Hence, alternatively one could think of maximizing revenue in these applications.

environments [51], combinatorial valuations [52], and combinatorial auctions [22, 32] (see [46] and [20] for comprehensive surveys). Importantly, one common aspect in all of this existing work in the literature on prophet inequalities is the assumption that the underlying algorithms should only make *irrevocable* decisions — once the decision is made for an arriving buyer/random variable, that decision is set in stone and can never be modified in the future. Relaxing this assumption gives more power to an online algorithm, as it can now hedge against possible early commitments to low-valued buyers arriving at the beginning of the sequence.

Motivated by some modern applications of Bayesian online allocation in electronic marketplaces and platforms — in particular, the criticality of maintaining supply efficiency in cloud spot markets and online hotel reservation systems through overbooking — we initiate the study of the single-item prophet inequality problem with (costly) *recourse*. More formally, we consider the setting in which the decision maker is presented with a sequence of independent random variables X_1, X_2, \dots one by one (think of them as the sequence of values of arriving buyers), where each X_i is drawn independently from a known distribution F_i . Once X_i is presented to the decision maker, she has the option of accepting this variable or going to the next round to observe the next variable. In the vanilla prophet inequality setting, the decision maker stops once she accepts the first variable. We diverge from this setting by allowing her to accept X_i even after the first allocation, but if the item has already been allocated to a buyer $j < i$ with value X_j , the decision maker should revoke this previous allocation by paying a cost in addition to losing the original reward X_j .

We focus on a simple, yet practical and fundamental model for costly recourse known as the *linear buyback setting*, which has been introduced and studied in the literature on online allocations under adversarial arrivals [9, 11, 23]. In this model, the additional cost of discarding X_j is equal to $f \cdot X_j$, where $f \geq 0$ is called the *buyback parameter*. The goal is to maximize the expected net reward, defined as the value of the final accepted variable minus the total buyback cost.² From a technical perspective, this is a natural model, as the resulting problem is invariant under scaling of the values. The model is also natural from the perspective of applications, as it captures scenarios where X_i 's are the willingness-to-pay of the arriving users of a service, and the decision maker uses “compensation fees” in the form of a fixed percentage of the original transfer paid by the user to take back the resource. One might also imagine scenarios in which the X_i 's are a sequence of demand requests (received by a platform during a decision-making horizon) declaring how long they would like to use a service in the future (e.g., renting an item for X_i hours next week). Now, the platform can serve only one request (e.g., because it has only one rental item) and its service is offered at a price of 1 per unit of time (so giving the service to demand X_i generates a revenue of $1 \times X_i$). The platform also has the option of “outsourcing” the request by offering the user the same service from a third party after paying a price of $1 + f > 1$ per unit of time to that third party (e.g., can

borrow the same rental item from a third party to rent to the user for X_i hours). Hence, if it outsources X_j , it needs to pay $(1 + f) \times X_j$. The goal of the platform is to decide to whom it should allocate its own service and to whom it should offer an outsourced service to maximize its net profit.

To establish prophet inequalities in the above model, we compare the expected net reward of our algorithms with the prophet benchmark $\mathbb{E}[\max_i X_i]$. Let $\alpha(f)$ be defined as the optimal competitive ratio that can be achieved by an online algorithm in our setting. Knowing the sequence of distributions, one can use polynomial-time dynamic programming/backward induction to devise an *optimal online policy* for our problem that maximizes the expected net reward amongst all online algorithms (note that the state of such a dynamic programming is essentially the pair of time and the value of the previous allocation, and hence is polynomial-time solvable with proper discretization; we detail this later in Section 2). As a result, $\alpha(f)$ can be defined as the worst-case ratio of the online optimum over the offline optimum. In the special case of $f = +\infty$ (or alternatively, the version in which no buyback is allowed), our problem is essentially the classic prophet inequality problem — therefore, $\alpha(f) \in [0.5, 1]$ for any choice of f and $\lim_{f \rightarrow +\infty} \alpha(f) = 0.5$. Intuitively speaking, as f decreases, the performance of the best online algorithm should improve. Furthermore, $\lim_{f \rightarrow 0} \alpha(f) = 1$, as in the extreme case of zero buyback cost with $f = 0$, a simple greedy algorithm, performing buyback whenever it has a positive gain, achieves the exact offline optimum. We now ask the following research question:

Can we characterize the optimal competitive ratio $\alpha(f)$, that is, the worst-case performance ratio of the optimal online algorithm over the offline optimum benchmark, as a function of the buyback parameter $f \geq 0$? In particular, (i) can we always obtain a constant competitive ratio $\alpha(f) > 0.5$ for every fixed $f > 0$? (ii) can we understand the behavior of $\alpha(f)$ in the asymptotic regime when $f \rightarrow 0$?

1.1 Our Results

Our first main result answers the above question affirmatively when the buyback parameter is in the range $f \in [1, +\infty)$. More specifically, we show that:

(informal) Theorem I: For any buyback parameter $f \in [1, +\infty)$, there exists a polynomial-time online algorithm that achieves a competitive ratio $\alpha(f) = \frac{1+f}{1+2f}$ against the optimum offline benchmark. Furthermore, there exists an instance of $n = 2$ two-point random variables such that no online algorithm can obtain a competitive ratio better than $\frac{1+f}{1+2f}$.

Note that as f becomes smaller, the competitive ratio of the optimal online policy can only improve; Therefore, as a simple corollary of the above result, we obtain a constant improvement over 0.5 for any fixed f . Our bad example in the above result is actually quite intuitive (and a natural generalization of the bad example for the classic prophet inequality): Consider two random variables where $X_1 = 1$, and $X_2 = 1 + f$ with probability $\frac{1}{1+f}$ and 0

²While it seems discarding decisions are made in an online fashion in this model, one could alternatively frame the problem as an online allocation problem with overbooking, but at the end only the highest allocated value receives the item and the decision maker must pay the required cancellation cost to discard all the other overbooked buyers.

otherwise. It is easy to verify that every online algorithm obtains value at most 1: If it selects $X_1 = 1$, it is indifferent between ignoring X_2 or performing a buyback, obtaining an expected reward of 1 in either case. If we ignore X_1 , we also obtain at most 1, since $\mathbb{E}[X_2] = 1$. On the other hand, $\mathbb{E}[\max\{X_1, X_2\}] = \frac{1}{f+1}(f+1) + \frac{f}{f+1} = \frac{2f+1}{f+1}$. Based on this example, it is quite tempting to conjecture that $\frac{1+f}{1+2f}$ is the optimal answer for every $f \geq 0$. Our result above shows that this is indeed the case for $f \geq 1$. We also show (in the full version of the paper) that the same result holds for every instance with $n = 2$, that is, with only two variables (regardless of whether $f \geq 1$ or $0 \leq f < 1$). Furthermore, we show that if the random variable $X_{\max} \triangleq \max_i X_i$ is bounded within a multiplicative range of $1+f$, then a simple thresholding algorithm with no buyback achieves the ratio $\frac{1+f}{1+2f}$ for all $f \geq 0$.

However, we show that the bound $\frac{1+f}{1+2f}$ is actually *not possible to achieve* for $0 < f < 1$ for the general problem. We establish this result by constructing an instance of the problem with $n = 3$ variables for which no online algorithm can obtain a competitive ratio better than $\frac{(1+f)(\sqrt{f(2-f)+1})}{(1+f)\sqrt{f(2-f)+3f+1}} < \frac{1+f}{1+2f}$ for $0 < f < 1$; we also conjecture that this bound is tight for $n = 3$ variables. With the help of a computer-aided search, we also tried to find the worst-case instance for $n = 4$, and we observed a continuation of the pattern: There seems to exist an instance where, for any choice of $0 < f < \frac{1}{3}$, the competitive ratio is strictly worse than the ratio obtained by $\frac{(1+f)(\sqrt{f(2-f)+1})}{(1+f)\sqrt{f(2-f)+3f+1}}$. We leave finishing the investigation of identifying the worst-case instance for any n as an interesting open problem.

This investigation suggests that the problem presents new challenges for smaller values of f , since the worst-case instance seems to employ more and more variables as f decreases. Based on this evidence, it may not be feasible to obtain a closed form for the competitive ratio as a function of the buyback parameter f in the entire $0 < f < 1$ regime. Despite the fact that we could establish a complete characterization of the optimal competitive ratio for $f \in [1, +\infty)$, the problem is still open for $0 < f < 1$ due to this issue. Therefore, we attempt only to answer weaker questions in this regime. In particular, we ask whether it is possible to characterize the asymptotic behavior of $\alpha(f)$ when $f \rightarrow 0$ by establishing (almost) matching upper and lower bounds on the performance of the online optimum policy. This brings us to our second main result. We show that:

(informal) Theorem II: For some constant $c_0 > 0$, for any buyback parameter $f \in [0, c_0]$, there is a polynomial-time online algorithm that achieves a competitive ratio $\alpha(f) = 1 - O\left(f \log \frac{1}{f}\right)$ against the optimum offline benchmark. Furthermore, for every $f \in [0, c_0]$, there is an instance of $n = \Theta\left(\log \frac{1}{f}\right)$ two-point random variables such that no online algorithm can achieve a competitive ratio better than $1 - \Omega\left(f \log \frac{1}{f}\right)$.

Our online policy in the result above is a simple threshold-greedy algorithm that works as follows: depending on the value of f and the set of distributions, it chooses a threshold T . Then it performs the first selection the first time an arriving random variable X_i exceeds the threshold T . After that, it accepts any arriving random variable X_i as long as it generates a positive marginal gain, that is $X_i > (1+f)X_j$, where X_j is the previously accepted random variable at time i . We show that this simple algorithm achieves a competitive

$$\text{ratio } \gamma(f) = \max \left\{ 1 / \left(\frac{f}{f+1} + \left(2 + \frac{1}{f} \right)^{\frac{f}{1+f}} \right), \frac{1}{2} + \frac{1}{50(1+f)} \right\}, \text{ which is}$$

asymptotically $1 - O\left(f \log \frac{1}{f}\right)$ for small values of $f > 0$, and is always strictly larger than 0.5 for all $f \geq 0$. We note that this algorithm is *order oblivious*, which means that it only needs to know the set of arriving distributions – similar to the single threshold algorithm in the classical prophet inequality [53] – and not the order in which they arrive. Another important feature of this algorithm is that it can potentially perform an unbounded number of buybacks, which seems necessary to achieve the asymptotic optimal factor as $f \rightarrow 0$.

It is also interesting to compare our asymptotic bound of $1 - \Theta\left(f \log \frac{1}{f}\right)$ with the best known competitive ratio for small values of f in the adversarial setting, without prior stochastic information. In this setting, [9] established the optimal competitive ratio of $1/(1 + 2f + 2\sqrt{f(1+f)})$ for deterministic algorithms and the follow-up work of [11] established the optimal competitive ratio of $1/W_{-1}(-1/e(1+f))$, where W_{-1} is the non-principal branch of the Lambert W function.³ Both of these bounds are asymptotically equal to $1 - \Theta\left(\sqrt{f}\right)$, which converges to 1 much slower than $1 - \Theta\left(f \log \frac{1}{f}\right)$. This gap can be interpreted as the “value of information” about the arriving distributions in our setting, which turns out to be critical in obtaining the asymptotically optimal competitive ratio.

Our results are summarized in Figure 1. Due to constraints on length, some proofs are omitted. All omitted proofs can be found in the full version of the paper [25].

1.2 Overview of Our Techniques

Broadly speaking, our results in this paper are obtained using two fundamentally different approaches: Our first approach is based on combinatorial optimization techniques involving LP duality, flows, and cuts, while the second approach is inspired by various proofs of the classical prophet inequality.

Flow Approach. While it is always possible to set up a nonlinear mathematical program to characterize the worst-case performance of the optimal online policy against the prophet benchmark, studying such a program is not always tractable. We get around this issue by formulating instead a parametric linear program that characterizes the worst-case performance of the optimal online policy, and studying the program under the worst-case choice of these parameters. This turns out to involve using LP duality to encode

³The Lambert W function, also called the product logarithm, is the inverse of the function $y(x) = x \cdot e^x$. The non-principal branch $x = W_{-1}(y)$ is the inverse when $x \leq -1$.




Figure 1: Comparison of different lower and upper bounds on the competitive ratio. Blue and green curves refer to the optimal online policy (for $f \rightarrow 0$, we depict an exaggerated version of the $1 - \Theta(f \log \frac{1}{f})$ curve to emphasize the qualitative behavior), the red curve is our threshold-greedy algorithm, and the black curves are optimal bounds in the adversarial setting [9, 11].

the problem as a certain *generalized flow* problem. The generalized flow problem has a special structure which is amenable to explicit solutions; however, discovering this explicit solution (for $f \geq 1$) seems to require a rather circuitous route.

In order to set up our parametric linear program, we start by proving that for every instance of our problem, there must be a harder instance which is made up of two-point random variables $X_i \in \{0, v_i\}$, with $v_1 \leq v_2 \leq \dots \leq v_n$. This kind of a reduction to a discrete, monotonic setting is intuitive because it ensures that information about the maximum value of the random variables we are observing is revealed as slowly as possible, thereby making the problem hard for an online decision-maker who gets to see these variables sequentially.

Formally achieving this reduction depends on being able to carefully analyze the optimal online policy using a Bellman recurrence that describes how much value the policy can obtain starting from a certain point in the middle of the sequence. We demonstrate that the reduction to the discrete and monotone setting formally depends only upon verifying a few special properties of the recurrence (Lemmas 3 and 4).

Once we have an instance in this monotonic form, we can set up a factor-revealing LP to understand the worst-case competitive ratio; however, the probabilities q_t of the two-point variables must remain external parameters, since the constraints that determine the competitive ratio are nonlinear in q_t . LP duality reveals another linear program, which we dub “contention resolution with recourse”. This is somewhat analogous to contention resolution LPs, where each variable encodes the probability of being at a certain state during the process and we would like to guarantee a certain amount of *allocation coverage* to each state; however it has some new features: think of each state (i, j) , $i < j$ as a directed edge from node i to j , where this state encodes that the decision maker is at time j and the last variable accepted was at time i . The LP has the form of a *generalized flow* in this digraph (with “leakage” along each edge related to the buyback parameter). We also have capacities on each edge, related to the final policy being implementable in an

online fashion, and the goal of the flow problem is find a generalized flow that satisfies certain demands at each node i . The task that remains is to find an explicit solution to this LP for every choice of the probabilities q_t .

It turns out that in the worst-case, it is safe to assume that all the probabilities q_t are small. This allows us to reduce (using standard Riemann sum arguments) the problem of finding a solution to the LP for potentially many different values of q_t to solving a single continuous flow problem. This part of the reduction in our analysis is inspired by known “variable splitting” methods for other prophet-inequality-type problems (e.g., [45]) to reduce the problem to the case with small random variables, but the particular splitting we employ to make the instance harder is unique to our setting.

Explicitly solving continuous linear programs, such as our continuous generalized flow problem, is typically a challenging task. A key insight in solving our continuous flow problem is that if we believe the optimal competitive ratio in a certain range of f is obtained on a particular small instance (for example on 2 variables), then the structure of the general solution should mimic the solution for the small example. More precisely, by grouping variables in a larger instance (or by splitting a continuous interval) into two blocks, we can view the larger instance as a blow-up of the hard instance on 2 variables. We therefore hypothesize that the flow has a block structure, which allows us to make some natural assumptions that reduce the space of possible solutions that we need to examine to find a flow.

After we make these assumptions, one way to prove the existence of the desired flow is to use the max-flow min-cut theorem. We omit this proof (our first) of being able to obtain a factor of $\frac{1+f}{1+2f}$ for $f \geq 1$. Instead, we exploit the intuition obtained from the study of minimum cuts to find a feasible flow explicitly. In the construction of our flows, inspired by the structure obtained at a block level, we posit that certain capacity constraints related to pairs of random variables that are far enough from each other are tight. We translate this tightness restriction to certain differential equations that can aid in partially constructing the continuous flow, and then use the remaining constraints in our continuous program to construct the final flow.

It turns out, however, that our 2-block solution cannot extend to the regime of $f < 1$. Indeed, we are able to use the Bellman recurrence and its properties to construct examples for which it is impossible to achieve the factor of $\frac{1+f}{1+2f}$. In all the examples we construct, the random variables satisfy an important property—if the optimal algorithm accepts a certain variable X_i , then it is indifferent between accepting or rejecting X_{i+1} . While we cannot prove that the worst-case instance on a given number of variables must always have this structure, it seems plausible that this is true. Our examples suggest a sequence of intervals I_k where the optimal factor $\alpha(f)$ is determined by an instance of this type on k variables; however, this remains as a conjecture.

Threshold-Greedy Approach. Our second approach is a natural algorithm motivated by existing proofs of the classical prophet inequality: We set a threshold T , wait for the first variable above T ; then, if we are holding a variable X_i and encounter another variable $X_j > (1+f)X_i$, we buyback and swap X_i for X_j . While this is certainly not the optimal online policy, it has the advantage

of being simple and practical, and it requires less prior information than the LP-based algorithm discussed above. Here, we only need to know the distribution of X_{\max} (in fact certain special quantiles of X_{\max}), and we don't even need to know the individual distributions \mathcal{D}_i , their ordering, or even the number of variables n .

The analysis of the threshold-greedy algorithm can be viewed as an extension of the classical prophet inequality; however, with some significant complications. One of our ideas here is that in the usual proof of the prophet inequality, we really only count gains from events where exactly one variable is above the threshold. Sometimes surprisingly, events where multiple variables appear above T cannot be used to obtain a better prophet inequality; however, these events provide extremely useful gains in our problem with buyback. Roughly speaking, if the number of variables above T is k , we reach the actual maximum of these variables after at most $k - 1$ buybacks, and the reward we obtain is at least $\frac{X_{\max}}{(1+f)^{k-1}}$. The number of variables above T can be approximated by a Poisson random variable, which can be decoupled from the actual maximum X_{\max} using the FKG inequality. Finally, we can optimize the performance of the algorithm by carefully choosing the threshold T .

All of these ideas taken together lead to an almost-optimal lower bound of $1 - \Theta(f \log \frac{1}{f})$ for small $f > 0$ (we choose T so that $\Pr(X_{\max} < T) = \frac{f}{1+2f}$). Furthermore, a separate analysis shows that this simple algorithm achieves a factor of $1/2 + \Omega(1/f)$ for any fixed $f > 1$. This is obviously weaker than the optimal factor of $\frac{1+f}{1+2f}$, but it shows that even limited prior information can be used to obtain an improvement over $\frac{1}{2}$ for any buyback parameter.

1.3 Further Related Work

Buyback and Recourse in Online Allocations. The study of the single item buyback setting under adversarial arrivals originates from the work [9], in which the optimal competitive ratio for deterministic integral algorithms is determined using a modification of the greedy algorithm. The follow-up work of [11] shows that the optimal competitive ratio can be achieved by a randomized algorithm that combines an elegant correlated randomization step with the greedy algorithm. The recent work of [23] shows an optimal competitive primal-dual fractional algorithm for this problem, which together with a lossless randomized rounding for the single item buyback setting, yields another optimal competitive randomized algorithm under adversarial arrivals. Beyond the single item setting, the buyback setting has also been studied for general matroids [9, 11], online matchings [23], and general matroid intersections [10]. The buyback setting is also the special case of an unconstrained online allocation with a combinatorial valuation $\mathcal{F}(S) = (1+f)\max_{i \in S} X_i - f \sum_{i \in S} X_i$, which is nonmonotone but submodular. Prophet inequalities with submodular combinatorial valuations are studied in [52], in which they devise an $\mathcal{O}(1)$ competitive algorithm (weaker than our bound) for this general problem. Another related problem is prophet inequalities with overbooking [7, 29], where the decision maker is allowed to accept more random variables than their capacity. We diverge from this setting by considering a capacity (single item) and a linear cost for cancellations. Finally, the role of “recourse” (or preemption)

in resource allocation has been studied in other settings. Examples include, but are not limited to, general online algorithms with rejection power [8], online submodular maximization with preemption [13], scheduling [14], online packing with removal cost [36], online matching with augmentations [15], and designing online BIC ad mechanisms with cancellations [18]. We diverge from all of this work by considering the different setting of the classic prophet inequality with buyback.

Online Contention Resolution and LP Duality. Our generalized flow problem is the dual of a specific factor-revealing LP for the worst-case competitive ratio of the optimal online policy against the prophet benchmark. As mentioned earlier, this flow problem is intimately connected to a generalization of the online contention resolution scheme (CRS) problem when (costly) recourse is allowed. The OCRS problem is rooted in the seminal work of [17, 30] (in the offline setting) and [33] on performing contention resolution for various combinatorial optimization problems, and has since been studied extensively in the CS and OR literature. For example, the work of [4] obtains an (almost) optimal OCRS for uniform matroids, which has been further optimized or simplified [21, 40] and extended to reusable resources [34]. In a different direction, the beautiful work of [44] establishes a duality between prophet inequalities with respect to ex-ante relaxation and OCRS, which is, at a high level, also connected to our LP approach and other LP approaches in the literature on obtaining prophet inequalities or competitive ratios against fluid-approximations in Bayesian online allocation, e.g., [3–5, 41]. OCRS-based approaches have several applications in prophet inequalities and Bayesian mechanism design. Some classic examples are sequential posted pricing [16, 37, 55], combinatorial auctions [22, 31], adding correlations to prophet inequalities [39], prophet inequalities for matchings [28, 50], and the prophet secretary setting [2, 26].

Optimum Online Policy. A major part of the literature on prophet inequalities focuses on understanding simple algorithms, e.g., a single threshold, and developing competitive ratios against the prophet benchmark for these algorithms. However, some recent work tries to understand the optimum online policy (mostly from a computational perspective). Some examples are the work of [6] that establishes a PTAS in the laminar matroid environment, or the work of [54] that shows how to obtain a PTAS for the optimum online in a wide range of problems including prophet inequalities and stochastic probing. On the computational complexity side of the problem, finding the optimum online in a combinatorial setting such as matchings is known to be PSPACE-hard [49]. Moving beyond computational aspects, there is a line of work that aims to understand the advantages of knowing the order, e.g. [48] and the more recent work of [27]. In our setting, fortunately, the state space of the DP is not exponentially-sized and hence the optimum online policy can be computed exactly using polynomial-time backward induction.

2 DYNAMIC PROGRAMMING FORMULATION OF THE OPTIMAL POLICY

In this section, in order to develop a more general understanding of the optimal online policy, we first focus on formulating this policy

as a simple dynamic program given the knowledge of the sequence of distributions of arriving random variables. We provide some basic properties of the value function of this DP, which turn out to be helpful in later parts of our analysis.

To obtain the DP formulation, we can derive a recurrence, also known as the *optimality equation* [12], that describes what the optimal expected reward is at each point in the process. Suppose that the instance consists of a sequence of random variables X_1, X_2, \dots, X_n , with known distributions. For $t = 0, \dots, n - 1$, we define

$$\Phi_t(x) = \mathbb{E}[\text{optimal reward that an algorithm can obtain, given that before observing } X_{t+1}, \text{ it holds value } x].$$

In this notation, we ignore any buyback cost that the algorithm paid prior to observing X_{t+1} . We define $\Phi_n(x) = x$, which is consistent with the above since there is no variable X_{n+1} to observe and the algorithm simply keeps value x .

The key property is that the function $\Phi_t(x)$ satisfies the following recurrence.

LEMMA 1.

$$\Phi_{t-1}(x) = \mathbb{E}[\max\{\Phi_t(x), \Phi_t(X_t) - fx\}].$$

PROOF. Upon observing X_t , we have two options: either we keep the value x that we held from before, and in the future obtain expected reward $\Phi_t(x)$, or we pay the buyback cost fx , and accept value X_t , in which case we obtain expected reward $\Phi_t(X_t) - fx$. Given the value of X_t , the optimal algorithm chooses the better of the two options, and in expectation obtains $\mathbb{E}[\max\{\Phi_t(x), \Phi_t(X_t) - fx\}]$. \square

Since we know $\Phi_n(x) = x$, we can use backward induction to compute the function Φ_t for every $t = 0, \dots, n - 1$, to determine the optimal algorithm's behavior. Note that whether the algorithm accepts or rejects a given value of X_t is determined exactly by which of the two arguments of $\max\{\Phi_t(x), \Phi_t(X_t) - fx\}$ is larger. To define the optimal online policy uniquely, let us assume that it buys back x and accepts X_t only if $\Phi_t(X_t) - fx > \Phi_t(x)$ (but this choice is not essentially important).

We have the following properties.

LEMMA 2. *The function Φ_t for each $t = 0, \dots, n$ satisfies:*

- Φ_t is convex,
- $\Phi_t(x) + fx$ is non-decreasing,
- $\Phi_t(x) - x$ is non-increasing.

The proof proceed by backward induction on t , and is postponed to the full version of the paper. Particularly important for the subsequent analysis will be the following:

LEMMA 3. *If at time t , while holding a value x , the optimal policy buys back x and swaps for a variable $X_t = v$, then it will also buyback x and swap for X_t if $X_t > v$. Formally, if $\Phi_t(v) - fx > \Phi_t(x)$, then we also have $\Phi_t(v') - fx \geq \Phi_t(v) - fx > \Phi_t(x)$ for $v' > v$.*

(Note that this is not completely obvious, since holding a higher value is not necessarily more profitable — Φ_t is not always monotone. Nevertheless the statement above is true.)

PROOF. If the optimal policy swaps x for $X_t = v$, this means that $\Phi_t(v) - fx > \Phi_t(x)$. This implies that $\Phi_t(v) > \Phi_t(x)$, and

by convexity $\Phi_t(y)$ must be increasing for $y \geq v$. Hence, for any $v' > v$, we have $\Phi_t(v') > \Phi_t(v)$. \square

Finally, we mention one more important lemma:

LEMMA 4. *If $x < y$, then $\Phi_{t-1}(y) - \Phi_{t-1}(x) \leq \Phi_t(y) - \Phi_t(x)$.*

Given the DP formulation, our ultimate goal is to compare the performance of the online optimum, that is, $\Phi_0(0)$ with that of the offline optimum, that is, $\mathbb{E}[\max_i X_i]$. In the next section, we switch to a linear programming formulation to analyze the worst-case performance of this DP policy with respect to the prophet benchmark.

3 COMPETITIVE RATIO ANALYSIS VIA LINEAR PROGRAMMING

In general, one can set up a nonlinear exponential-size mathematical program to characterize the worst-case performance of the optimal online policy against the prophet benchmark — in the same spirit as characterizing the worst-case competitive ratio as the equilibrium of a zero-sum game played between an algorithm who picks an online policy and an adversary who picks the sequence of distributions. In this section, instead of taking this naive approach, we show how to develop a parametric linear program, where the parameters somehow encode a lower-dimensional “sufficient statistic” of the worst-case instance. We then show how this LP characterizes the worst-case ratio of the optimal online policy and the offline optimum under the worst-case choices of these parameters.

To develop this parametric LP, first we need to describe a reduction that converts any instance to a restricted form, such that the worst-case competitive ratio does not change when restricting to these instances.

3.1 Reduction to a Monotonic Sequence of 2-Point Distributions

We claim that without loss of generality, we can consider a sequence of random variables X_i with 2-point distributions, i.e, $X_i \in \{0, v_i\}$, where $v_1 \leq v_2 \leq \dots \leq v_n$. Moreover the probability of each random variable being the last nonzero one can be made arbitrarily small. More precisely, we prove the following.

THEOREM 5. *For any $\epsilon > 0$, if a competitive factor $\alpha(f)$ can be achieved for any instance of the buyback problem with random variables $X_i = v_i \cdot \text{Be}(q_i)$, where $\text{Be}(q_i)$ is a Bernoulli $0/1$ random variable with expectation q_i , $0 \leq v_1 \leq \dots \leq v_n$, $q_i \prod_{j=i+1}^n (1 - q_j) \leq \epsilon$, and $q_1 = 1$, then the same competitive factor $\alpha(f)$ can be achieved for any instance with nonnegative random variables of finite expectation.*

We proceed in a sequence of simple reductions. The first step is to discretize the random variables.

LEMMA 6. *If a competitive factor $\alpha(f)$ can be achieved for any instance with discrete nonnegative random variables (with finitely many possible values), then the same competitive factor $\alpha(f)$ can be achieved for any instance with nonnegative random variables of finite expectation.*

The second step is to reduce the random variables by “splitting” to 2-point distributions.

LEMMA 7. *If a competitive factor $\alpha(f)$ can be achieved for any instance with scaled Bernoulli random variables (with values in $\{0, v_i\}$ for some $v_i > 0$), then the same competitive factor $\alpha(f)$ can be achieved for any instance with discrete nonnegative random variables.*

The third step is to rearrange the sequence of 2-point distributions, so that the values v_i are non-decreasing. We prove that this can only make the instance more difficult.

LEMMA 8. *If a competitive factor $\alpha(f)$ can be achieved for any instance with scaled Bernoulli random variables (with values in $\{0, v_i\}$ for some $v_i > 0$), and $0 \leq v_1 \leq \dots \leq v_n$, then the same competitive factor $\alpha(f)$ can be achieved for any instance with scaled Bernoulli random variables.*

PROOF. Consider an instance I_1 with variables $X_i = v_i Be(q_i)$ and $X_{i+1} = v_{i+1} Be(q_{i+1})$, $v_i < v_{i+1}$, and the same instance with the ordering of variables X_i, X_{i+1} swapped, which we call I_2 . Suppose we hold value x before observing the variables X_i, X_{i+1} , and consider 2 cases:

Case 1. The optimal policy \mathcal{P}_1 for I_1 does not accept X_i even if nonzero. Then, the value of this variable is irrelevant, and we can define \mathcal{P}_2 to accept X_{i+1} based on whether \mathcal{P}_1 accepts X_{i+1} , which yields an equivalent policy.

Case 2. The optimal policy \mathcal{P}_1 for I_1 accepts X_i if $X_i = v_i$; i.e. $\Phi_i(v_i) - fx > \Phi_i(x)$. This also implies that $\Phi_{i+1}(v_i) - fx > \Phi_{i+1}(x)$, by Lemma 4, and $\Phi_{i+1}(v_{i+1}) - fx \geq \Phi_{i+1}(v_i) - fx$, by Lemma 3. We define \mathcal{P}_2 to accept X_{i+1} if $X_{i+1} = v_{i+1}$, and if not, then accept X_i if $X_i = v_i$. This yields a policy which performs at least as well as \mathcal{P}_1 in each of the 4 possible pairs of values of the variables X_i, X_{i+1} . The only non-trivial case is when $X_i = v_i, X_{i+1} = v_{i+1}$, and \mathcal{P}_1 accepts X_i and then rejects X_{i+1} . In this case, \mathcal{P}_1 obtains a value of $\Phi_{i+1}(v_i) - fx$ which is less than or equal to $\Phi_{i+1}(v_{i+1}) - fx$, the value \mathcal{P}_2 obtains. If \mathcal{P}_1 accepts X_i and then accepts X_{i+1} , \mathcal{P}_2 performs better because it pays lesser in buyback costs. In the other 3 when at least of X_i or X_{i+1} is zero, both policies perform the same. Therefore, the expected reward of \mathcal{P}_2 is better than that of \mathcal{P}_1 .

By this argument, if there is any adjacent pair of variables where $v_i > v_{i+1}$, we can swap them and obtain an instance where the optimal reward can only go down. By repeating this swapping operation, we can sort the original instance and obtain an instance where $v_1 \leq v_2 \leq \dots \leq v_n$. The argument above shows that the monotonic instance is at least as difficult as the original instance. \square

The last reduction step is not necessary for the LP formulation, but it will be useful later: we can split the random variables further, to obtain a sequence such that the probability of any random variable being the last nonzero variable is arbitrarily small, and we can also assume that $q_1 = 1$.

LEMMA 9. *For any $\epsilon > 0$, if a competitive factor $\alpha(f)$ can be achieved for any instance with scaled Bernoulli random variables $X_i = v_i Be(q_i)$, $0 \leq v_1 \leq v_2 \leq \dots \leq v_n$, $q_i \prod_{j=i+1}^n (1 - q_j) \leq \epsilon$, and $q_1 = 1$, then the same competitive factor $\alpha(f)$ can be achieved for any instance with scaled Bernoulli variables $X_i = v_i Be(q_i)$, $0 \leq v_1 \leq v_2 \leq \dots \leq v_n$.*

Combining Lemmas 6,7,8,9 implies Theorem 5.

3.2 LP Formulations

Let us now formulate a parametric linear program that captures the worst-case ratio of the optimal online policy to that of the prophet benchmark by using the recurrence for the optimal online policy (Lemma 1). We assume that the instance is in the form provided by Theorem 5: a sequence of 2-point distributions, $X_i = v_i Be(q_i)$, and $0 \leq v_1 \leq v_2 \leq \dots \leq v_n$. Lemma 1 then takes the following form:

$$\Phi_{t-1}(x) = \max\{\Phi_t(x), (1 - q_t)\Phi_t(x) + q_t(\Phi_t(v_t) - fx)\}. \quad (1)$$

We will formulate an LP with the following variables: $\Phi_{i,t}$, representing the value of $\Phi_t(v_i)$, and v_i representing the value v_i itself. We also use $\Phi_{0,t}$ to represent $\Phi_t(0)$ (and it is convenient to assume that $v_0 = 0$). Note that $\Phi_{i,t} = \Phi_t(v_i)$ obviously depends on the v_i 's, but in the LP formulation these will be independent variables; the connection between them will be provided by the constraints. The LP constraints are chosen to be valid for the optimal online policy: $\Phi_{i,t-1} \geq \Phi_{i,t}$, and $\Phi_{i,t-1} \geq (1 - q_t)\Phi_{i,t} + q_t(\Phi_{i,t} - fv_i)$. Another constraint is that $\Phi_{i,n} = v_i$: at the end of the process, the reward is the value that we are holding. The values should satisfy $0 \leq v_1 \leq \dots \leq v_n$. And finally, we have a constraint normalizing the offline optimum to be 1:

LEMMA 10. *Given $q_1, \dots, q_n \in [0, 1]$, the following LP gives the best possible factor that an online policy can achieve on an instance in the form $X_i = v_i Be(q_i)$, $0 \leq v_1 \leq \dots \leq v_n$:*

$$\begin{aligned} \min \quad & \Phi_{0,0} \quad \text{s.t.} \\ \forall 0 \leq i < t \leq n; \quad & \Phi_{i,t-1} \geq \Phi_{i,t} \\ \forall 0 \leq i < t \leq n; \quad & \Phi_{i,t-1} \geq (1 - q_t)\Phi_{i,t} + q_t(\Phi_{i,t} - fv_i) \\ & \Phi_{0,n} = 0 \\ \forall 1 \leq i \leq n; \quad & \Phi_{i,n} = v_i \\ \forall 1 \leq i < n; \quad & v_{i+1} \geq v_i \\ \forall 1 \leq i \leq n; \quad & v_i \geq 0 \\ \sum_{i=1}^n v_i q_i \prod_{j=i+1}^n (1 - q_j) & = 1 \end{aligned}$$

The proof follows from the properties of the $\Phi_t(x)$; we omit the details.

Now we appeal to LP duality to obtain another LP which will be much more useful for us. We use the following notation:

$$\hat{q}_t = q_t \prod_{j=t+1}^n (1 - q_j). \quad (2)$$


Note that this is exactly the probability that $X_t = X_{\max}$ (in the monotonic instance). The dual LP can be written after some simplifications as follows (again we omit the details).

$$\begin{aligned} \max \quad & \Theta \quad \text{s.t.} \\ \forall 1 \leq t \leq n; \quad & x_{0,t} \leq q_t \left(1 - \sum_{i=1}^{t-1} x_{0,i}\right) \\ \forall 1 \leq s < t \leq n; \quad & x_{s,t} \leq q_t \left(\sum_{i=0}^{s-1} x_{i,s} - \sum_{j=s+1}^{t-1} x_{s,j}\right) \\ \forall 1 < t \leq n; \quad & \Theta \hat{q}_t \leq \sum_{i=0}^{t-1} x_{i,t} - (1 + f) \sum_{j=t+1}^n x_{t,j} \\ \forall 0 \leq s < t \leq n; \quad & x_{s,t} \geq 0 \end{aligned}$$

This LP has a natural interpretation: $x_{0,t}$ represents the probability that X_t is the first variable that we accept, and $x_{s,t}$ for $1 \leq s < t$ represents the probability that at some point we swap X_s for X_t . The first constraint expresses the fact that by the time we get to X_t , there is probability $1 - \sum_{i=1}^{t-1} x_{0,i}$ that we have not accepted any variable yet; and this is independent of the probability that $X_t = v_t$,

which is q_t . Therefore, $q_t(1 - \sum_{i=1}^{t-1} x_{0,i})$ is an upper bound on the probability that we accept X_t . Similarly, $\sum_{i=0}^{s-1} x_{i,s} - \sum_{j=s+1}^{t-1} x_{s,j}$ is the probability that we accepted X_s and haven't discarded it yet by the time we get to X_t . Therefore, $q_t(\sum_{i=0}^{s-1} x_{i,s} - \sum_{j=s+1}^{t-1} x_{s,j})$ is an upper bound on the probability that we swap X_s for X_t . We call these constraints *online implementability* constraints.

The third constraint expresses the fact that the total probability of accepting variable X_t minus the buyback cost of possibly discarding it should be at least Θ times the probability that $X_t = X_{\max}$. If we satisfy this constraint, then we recover at least a Θ -fraction of the contribution of X_t to the offline optimum.



This LP can be viewed as a generalized flow problem, with $x_{s,t}$ representing the flow from node s to node t , each node extracting a certain demand proportional to Θ , and a certain fraction of the flow along each edge being lost due to the buyback cost. In the following, our goal is to prove that this LP has a solution for $f \geq 1$ and $\Theta = \frac{1+f}{1+2f}$, and thereby obtain a prophet inequality for the buyback problem.

3.3 Solving the LP: Intuition


We describe in this section intuition which will lead us to a feasible solution for the LP.

If we want to match the hardness factor $\frac{1+f}{1+2f}$ given by the worst-case 2-variable example, our solution certainly has to solve this example optimally. The probabilities in this instance are $\hat{q}_1 = \frac{f}{1+f}$ and $\hat{q}_2 = \frac{1}{1+f}$, and we can find an optimal solution of the corresponding dual LP, which gives an algorithm to solve 2-variable instances in this form. However, this still does not tell us how to solve a general instance.

If we want to solve every instance in a manner consistent with this example, we hazard a guess that the solution should have a 2-block structure where the two blocks correspond to variables of total probability mass $\ell_1 = \sum_{s \in A} \hat{q}_s = \frac{f}{1+f}$, and $\ell_2 = \sum_{t \in B} \hat{q}_t = \frac{1}{1+f}$. (As a special case, the 2 blocks could be the 2 variables in the hardness instance.) Crucially, the flow at the level of these blocks should be consistent with the optimal solution for the hard instance on 2 variables. Furthermore, there should be no flow inside each block, since this would cause us to incur an additional buyback cost which we cannot afford. Therefore, buyback edges should go only from A to B , and not within each block. This leads to constraints that imply the following aggregate flow between the initial node and the two blocks:

$$x_{0,A} = \frac{1+f}{1+2f}, \quad x_{0,B} = \frac{f}{(1+f)(1+2f)}, \quad x_{A,B} = \frac{1}{(1+f)(1+2f)}.$$

Note that the probabilities \hat{q}_t in a given instance do not necessarily allow us to form blocks of exactly the desired sizes ℓ_1 and ℓ_2 . However, we can always split variables into identical independent copies with smaller probabilities, in a way similar to Lemma 7, in order to obtain the desired block size, so this is not an issue.



Once we postulate that the solution should have this block structure, it remains to find the specific flow from the initial node 0 ($x_{0,s}$ for $s \in A \cup B$), and the flow between the two blocks ($x_{s,t}$ for $s \in A$ and $t \in B$). The interpretation of $x_{0,s}$ is that it describes the probability of X_s being selected as our first pick. The interpretation of $x_{s,t}$ is that it describes the probability of swapping from X_s in block A to X_t in block B . Since at the beginning of the process, we can select variables without the online implementability constraint imposing a significant restriction, $x_{0,s}$ for s in the first block is not very restricted. However, we need to be careful not to make $x_{0,s}$ too large because that would hamper our ability to select variables in the second block.

No buyback happens within the first block, and hence $x_{0,s}$ is the only source of acceptance probability in the first block. This means that we should have $x_{0,s} \geq \Theta \hat{q}_s$ in the first block. On the other hand, a variable in the second block will never be discarded, and it can be either selected as our first pick or via buyback. Therefore, we have $x_{0,t} \leq \Theta \hat{q}_t$ in the second block. It is natural to assume that $x_{0,s}$ should be decreasing in s . We also hypothesize that in the second block, we should accept any nonzero variable if we have not accepted anything so far; i.e., the online implementability should be tight: $x_{0,t} = q_t(1 - \sum_{s=1}^{t-1} x_{0,s})$. This, together with the block flow $x_{0,B}$, determines $x_{0,t}$ for $t \in B$. For $s \in A$, we choose $x_{0,s}$ to depend linearly on s in a way that produces the desired block flow $x_{0,A}$.

Finding the interblock flow $x_{s,t}$ is the trickiest part of our solution. We have several constraints to satisfy here, and the more restrictive constraints appear towards the end of the second block. Here, we have the maximum deficit $\Theta - x_{0,t}$, while the online implementability constraint simultaneously takes its most restrictive form. Therefore, we should try to send as much flow as possible to the end of the second block, from every point $s \in A$. Given these constraints, we send a flow $x_{s,t} = \sigma(s) \hat{q}_s \hat{q}_t$ to a certain interval $\tau(s) \leq t \leq n$, where $\sigma(s)$ is determined by the online implementability constraint and $\tau(s)$ is chosen so that $\sum_{\tau(s) \leq t \leq n} x_{st} = \frac{1}{1+f}(x_{0,s} - \Theta) \hat{q}_s$ is exactly the surplus that we allowed to send from $s \in A$. It remains to verify that the choices described here form a feasible solution of the dual LP.

3.4 Solving the LP: Technical Details

Here we show rigorously how to describe the LP solution. Most details are postponed to the full version of the paper. First, note from our earlier discussion that it is safe to assume that $\hat{q}_i \leq \epsilon$ for any fixed $\epsilon > 0$. The first step in solving the LP is to reduce the problem to a continuous version. Intuitively, we can imagine that we have a continuum of random variables, each of which attains a nonzero value with infinitesimal probability. To begin the reduction, let us write down

$$x_{0,s} = H(s) \hat{q}_s, \quad \text{and} \quad x_{s,t} = G(s, t) \hat{q}_s \hat{q}_t$$

for some positive functions H and G . Note that the LP we're interested in now has the form:

$$\begin{aligned} \max \quad & \Theta \quad \text{s.t.} \\ \forall t \geq 1; \quad & H(t) \frac{\hat{q}_t}{q_t} \leq 1 - \sum_{i=1}^{t-1} H(i) \hat{q}_i \\ \forall t > s; \quad & G(s, t) \frac{\hat{q}_t}{q_t} \leq H(s) + \sum_{i=1}^{s-1} G(i, s) \hat{q}_i - \sum_{j=s+1}^{t-1} G(s, j) \hat{q}_j \\ \forall t \geq 1; \quad & \Theta - H(t) \leq \sum_{i=1}^{t-1} G(i, t) \hat{q}_i - (1+f) \sum_{j=t+1}^n G(t, j) \hat{q}_j \end{aligned}$$

The next lemma demonstrates that it suffices to analyze the natural continuous version of the above LP:

LEMMA 11. *If there exist non-negative, continuous functions h and g with domains $[0, 1]$ and $[0, 1] \times [0, 1]$, and $h(x) > 0$ such that*

$$\begin{aligned} \forall 0 \leq t \leq 1; \quad & th(t) \leq 1 - \int_0^t h(r) dr \\ \forall 0 \leq s \leq t \leq 1; \quad & tg(s, t) \leq h(s) + \int_0^s g(r, s) dr - \int_s^t g(s, r) dr \\ \forall 0 \leq t \leq 1; \quad & \Theta - h(t) \leq \int_0^t g(r, t) dr - (1+f) \int_t^1 g(t, r) dr. \end{aligned}$$

then $\alpha(f) \geq \Theta$.

The proof of this lemma uses a standard Riemann sum argument, and the fact that we can assume all the \hat{q}_i are less than ϵ . We postpone the proof to the full version.

Now we explicitly define such functions h and g for $f \geq 1$. Further motivation for our choices, and a complete demonstration that they satisfy the continuous version of the LP, is provided in the full version of the paper. In the following, we denote by $\Theta = \frac{1+f}{1+2f}$ the desired approximation factor. Let us define

$$h(t) := \begin{cases} \Theta \cdot \left(1 + \frac{2}{f} \left(1 - \frac{1+f}{f} t\right)\right) & \text{if } t \leq \frac{f}{1+f} \\ \Theta \cdot \frac{f^2}{(1+f)^2 t^2} & \text{if } t > \frac{f}{1+f} \end{cases}$$

$$g(s, t) := \begin{cases} \frac{h(s)+\Theta}{1+f} & \text{if } s \leq \frac{f}{1+f}, t \geq \frac{2\Theta}{h(s)+\Theta} \left(\geq \frac{f}{1+f}\right) \\ 0 & \text{otherwise.} \end{cases}$$

For $f \geq 1$, we can show that h and g satisfy the continuous form of our LP with $\Theta = \frac{1+f}{1+2f}$, by letting $A = \{s \mid s \leq \frac{f}{1+f}\}$ and $B = \{s \mid h(s) < \Theta\} = \{s \mid s > \frac{f}{1+f}\}$, and checking that h is online implementable (in fact tightly for $t \in B$), and furthermore that:

$$\begin{aligned} \int_A \frac{h(s)-\Theta}{1+f} ds &= \int_B (\Theta - h(t)) dt = \frac{1}{(1+f)(1+2f)}, \\ \int_B g(s, t) dt &= \frac{h(s)-\Theta}{1+f}, \forall s \in A, \\ \int_A g(s, t) ds &= \Theta - h(t), \forall t \in B, \\ g(s, t) &\leq h(s) - \int_B g(s, r) dr = \frac{fh(s)+\Theta}{1+f}, \forall s \in A, t \in B. \end{aligned}$$

We conclude that we can obtain a competitive ratio of $\frac{1+f}{1+2f}$ for the prophet inequality with buyback if $f \geq 1$.

REMARK 1. *As mentioned earlier in the introduction, there is a simple example with only 2 variables, such that no online policy can obtain better than $\frac{1+f}{1+2f}$ fraction of the prophet benchmark. So the competitive ratio bound obtained in this section for $f \in [1, +\infty)$ is tight.*

4 COMPETITIVE RATIO UPPER-BOUNDS FOR SMALL BUYBACK FACTORS

In the previous section, we showed that the competitive ratio of $\frac{1+f}{1+2f}$ can be achieved by the optimal online policy when $f \in [1, +\infty)$. Also, we already saw an example with 2 random variables showing that a factor better than $\frac{1+f}{1+2f}$ cannot be achieved (Remark 1). Therefore, it is tempting to conjecture that this bound is the best possible for all values of the buyback parameter f . However, as we show next, even with $n = 3$ variables, we can construct hard instances where any algorithm obtains a factor that is strictly worse than $\frac{1+f}{1+2f}$ when $0 < f < 1$.

Hard Instance with 3 Variables. Building on the intuition for the worst-case instance with two variables, we consider the following instance with three variables for $x > 1 + f$:

$$X_1 = 1 \quad X_2 = \begin{cases} x & \text{w.p. } \frac{1}{x} \\ 0 & \text{otherwise} \end{cases} \quad X_3 = \begin{cases} x(1+f) & \text{w.p. } \frac{x-1-f}{(1+f)(x-1)} \\ 0 & \text{otherwise} \end{cases}$$

Later, we pick x to minimize the competitive ratio.⁴ The expected value of the prophet benchmark is easy to calculate:

$$\mathbb{E}[X_{\max}] = \frac{f}{1+f} + \frac{xf}{(1+f)(x-1)} + \frac{x(x-1-f)}{(x-1)} = x + f \frac{x-1-xf}{(x-1)(1+f)}.$$

To find the expected reward for the optimal online algorithm we can solve the recurrence described in Lemma 1 to calculate $\Phi_0(0)$:

$$\begin{aligned} \Phi_0(0) &= x - f & \Phi_1(0) &= x - f & \Phi_2(0) &= x - \frac{fx}{x-1} \\ \Phi_1(1) &= x - f & \Phi_2(1) &= x - f & \Phi_2(x) &= x \end{aligned}$$

Hence, no online algorithm can get a competitive ratio better than⁵

$$\min_x \frac{x-f}{x+f \frac{x-1-xf}{(x-1)(1+f)}} = \frac{(1+f)(\sqrt{f(2-f)}+1)}{(1+f)\sqrt{f(2-f)}+3f+1}.$$

The minimum competitive ratio is obtained when $x = \frac{f+2+\sqrt{f(2-f)}}{2}$, and the instance is:

$$X_1 = 1, \quad X_2 = \begin{cases} \frac{f+2+\sqrt{f(2-f)}}{2} & \text{w.p. } \frac{2}{f+2+\sqrt{f(2-f)}} \\ 0 & \text{otherwise} \end{cases}$$

$$X_3 = \begin{cases} \frac{f+2+\sqrt{f(2-f)}}{2} (1+f) & \text{w.p. } \frac{1-\sqrt{f(2-f)}}{(1+f)(1-f)} \\ 0 & \text{otherwise} \end{cases}$$

Notice that the instance only makes sense when $\frac{1-\sqrt{f(2-f)}}{(1+f)(1-f)} \leq 1$, i.e., when $f^4 \leq f(2-f)$. This happens precisely when $f \leq 1$. (Furthermore, $x > 1 + f$ only when $f < 1$.)

The above instance suggests that as f becomes smaller, the adversary may require a higher number of random variables to push the optimal online policy to achieve its worst-case competitive ratio. We conjecture that this instance is indeed the worst-case instance with $n = 3$ variables for the competitive ratio of the optimal online

⁴As it turns out, with this choice of x , this instance has the property that having accepted X_i , the optimal online algorithm is indifferent between accepting or rejecting X_{i+1} . However, we do not rely on this property in our calculations.

⁵For all $0 < f < 1$: $\frac{(1+f)(\sqrt{f(2-f)}+1)}{(1+f)\sqrt{f(2-f)}+3f+1} < \frac{1+f}{1+2f}$.

policy against the prophet benchmark, and leave it without a proof. With the help of a computer-aided search, we also tried to find the worst case instance for $n = 4$, and we observed that there exists an instance where, for any choice of $0 < f < \frac{1}{3}$, the competitive ratio is strictly worse than the ratio obtained by $\frac{(1+f)(\sqrt{f(2-f)+1})}{(1+f)(\sqrt{f(2-f)+1})+3f+1}$.

Hard Instances with Multiple Variables. More generally, we construct an example with n Bernoulli variables $X_i = v_i \text{Be}(p_i)$ which has the property that the optimal policy, while holding a value v_i , is indifferent to accepting v_{i+1} or rejecting it. In terms of $\Phi_t(x)$, this means the condition

$$\Phi_t(v_{t-1}) = \Phi_t(v_t) - f v_{t-1}.$$

We also set $p_i = 1/2$, which is not the optimal choice, but simplifies the analysis. This example leads to the following hardness result in the regime of small $f > 0$ (we defer the proof to the full version).

THEOREM 12. *Given a buyback factor $0 < f < \frac{1}{16}$, there is no algorithm for the buyback problem which achieves a competitive ratio better than $1 - \frac{1}{2}f \log_2 \frac{1}{16f}$.*

This demonstrates an interesting phenomenon: Recall that for $f = 0$, we can achieve a competitive ratio of 1, simply by performing buyback whenever possible. For small $f > 0$, we might aim to achieve a competitive ratio of $1 - cf$ for some fixed constant $c > 0$. However, the counterexample shows that this is impossible, and the optimal competitive ratio $\alpha(f)$ as a function of the buyback factor f satisfies $\alpha'(0) = -\infty$. As we prove in Theorem 16, the behavior of the optimal competitive ratio for small $f > 0$ is indeed $1 - \Theta(f \log \frac{1}{f})$.

5 THE THRESHOLD-GREEDY ALGORITHM

In this section, we present a simple algorithm which achieves a competitive ratio of $1 - \Theta(f \log \frac{1}{f})$ for small $f > 0$, matching Theorem 12 up to constant factors. This algorithm also gives a constant improvement over $1/2$, for any fixed buyback factor $f > 0$, but we omit the proof. The algorithm is a simple extension of the standard threshold algorithm: After setting an initial threshold, and accepting some variable above the threshold, we buy back and replace the current variable if and only if this brings positive profit.

The Algorithm.

- Set an initial threshold T (to be determined).
- If a variable X_i arrives and $X_i \geq T$, then select X_i .
- Given a currently selected value X_i , if a variable X_j arrives and $X_j > (1+f)X_i$, then discard X_i and select X_j .

An advantage of this algorithm is that is order oblivious, and as will see T will depend only on the distribution of X_{\max} .

5.1 General Analysis Approach

Let us first develop some general formulae to analyze this algorithm.

LEMMA 13. *Suppose that the algorithm selects at least k variables and the k -th selected variable is $X^{(k)}$. Then the profit of the algorithm is at least $X^{(k)}/(1+f)^{k-1}$.*

PROOF. By induction: If the first selected variable is $X^{(0)}$, then the profit at that point is clearly $X^{(0)}$. Observe also that profit never goes down after future buybacks, so the profit at the end is at least $X^{(0)}$.

Assuming that the k -th selected variable is $X^{(k)}$ and the variable selected just before that was $X^{(k-1)}$, by induction our profit after selecting $X^{(k-1)}$ was $p_{k-1} \geq X^{(k-1)}/(1+f)^{k-2}$. After the last buyback, since we select $X^{(k)}$ only if $X^{(k)} > (1+f)X^{(k-1)}$, and we pay $f \cdot X^{(k-1)}$ for discarding $X^{(k-1)}$, our profit is

$$\begin{aligned} p_k &= p_{k-1} + (X^{(k)} - (1+f)X^{(k-1)}) \\ &\geq \frac{X^{(k-1)}}{(1+f)^{k-2}} + \frac{1}{(1+f)^{k-1}}(X^{(k)} - (1+f)X^{(k-1)}) \\ &= \frac{X^{(k)}}{(1+f)^{k-1}}, \end{aligned}$$

as desired. \square

Next, consider the classical analysis of the prophet inequality, which uses the fact that our algorithm achieves expected value at least $T \cdot \Pr[X_{\max} \geq T] + \mathbb{E}[(X_{\max} - T)_+] \cdot \Pr[X_{\max} < T]$. We claim the following extension of this bound.

LEMMA 14. *The expected profit of our algorithm is*

$$\mathbb{E}[ALG] \geq T \cdot \Pr[X_{\max} \geq T] + \mathbb{E}\left[\left(\frac{X_{\max}}{(1+f)^{S'}} - T\right)_+\right],$$

where $S' = \sum_j \mathbb{1}(X'_j \geq T)$ and X'_j is an independent copy of X_j .

PROOF. First, let us note that if the maximum is X_i , and the algorithm makes θ_i picks before X_i , Lemma 4 implies that in the case that the maximum is above T , the algorithm obtains at least

$$\max\left(T, \frac{X_i}{(1+f)^{\theta_i}}\right).$$

It follows that

$$\begin{aligned} \mathbb{E}[ALG] &\geq T \cdot \Pr[X_{\max} \geq T] \\ &+ \sum_{i=1}^n \mathbb{E}\left[\left(\frac{X_i}{(1+f)^{\theta_i}} - T\right)_+ \mathbb{1}(X_i = X_{\max})\right] \\ &\geq T \cdot \Pr[X_{\max} \geq T] \\ &+ \sum_{i=1}^n \mathbb{E}\left[\left(\frac{X_i}{(1+f)^{\sum_{j \neq i} \mathbb{1}(X_j \geq T)}} - T\right)_+ \mathbb{1}(X_i = X_{\max})\right], \end{aligned}$$

since clearly, $\theta_i \leq \sum_{j \neq i} \mathbb{1}(X_j \geq T)$. Observe that conditioned on X_i , $\left(\frac{X_i}{(1+f)^{\sum_{j \neq i} \mathbb{1}(X_j \geq T)}} - T\right)_+$ is a decreasing function of $(X_j : j \neq i)$, and $\mathbb{1}(X_i = X_{\max})$ is also a decreasing function of $(X_j : j \neq i)$. So we can apply the FKG inequality on the product space of $(X_j : j \neq i)$, for every fixed value of X_i , and replace the variables $(X_j : j \neq i)$ in the first expression by independent copies $(X'_j : j \neq i)$. Using the notation $S' = \sum_j \mathbb{1}(X'_j \geq T)$, we obtain that $\mathbb{E}[ALG]$ is at least:

$$\begin{aligned} &T \Pr[X_{\max} \geq T] \\ &+ \sum_{i=1}^n \mathbb{E}\left[\left(\frac{X_i}{(1+f)^{\sum_{j \neq i} \mathbb{1}(X'_j \geq T)}} - T\right)_+ \mathbb{1}(X_i = X_{\max})\right] \\ &\geq T \Pr[X_{\max} \geq T] + \sum_{i=1}^n \mathbb{E}\left[\left(\frac{X_i}{(1+f)^{S'}} - T\right)_+ \mathbb{1}(X_i = X_{\max})\right] \\ &= T \Pr[X_{\max} \geq T] + \mathbb{E}\left[\left(\frac{X_{\max}}{(1+f)^{S'}} - T\right)_+\right], \end{aligned}$$

as desired. \square

We will also use the following useful comparison between summations of independent Bernoulli variables and a Poisson variable.

LEMMA 15. If $S' = \sum_{j=1}^n \mathbb{1}(X'_j \geq T)$, for independent random variables X'_1, \dots, X'_n , and $\Pr[\forall j; X'_j < T] = e^{-\lambda}$, then S' is stochastically dominated by a Poisson random variable with mean λ . In particular, for any $\ell \geq 1$,

$$\Pr[S' \leq \ell] \geq e^{-\lambda} \sum_{k=0}^{\ell} \frac{\lambda^k}{k!}.$$

PROOF. Let $p_i = \Pr[X'_i \geq T]$, and let Y_i be a Poisson random variable with mean $\lambda_i = -\log(1 - p_i)$, where $Y_i = 0$ whenever $X'_i < T$. (Note that $\Pr[Y_i = 0] = e^{-\lambda_i} = 1 - p_i = \Pr[X'_i < T]$.) Hence, $Y_i \geq \mathbb{1}(X'_i \geq T)$, and $S' = \sum_{j=1}^n \mathbb{1}(X'_j \geq T) \leq \sum_{j=1}^n Y_j$ with probability 1. Finally, note that $Y = \sum_{j=1}^n Y_j$ is also a Poisson random variable, with mean $\lambda = -\sum_{j=1}^n \log(1 - p_j) = -\log \prod_{j=1}^n (1 - p_j) = -\log \Pr[\forall j; X'_j < T]$. \square

5.2 Analysis for Small Buyback Factors

Consider now the case of small f (e.g. $f \in (0, 1/2)$). We can use Lemmas 14 and 15 to get a good bound on the algorithm's performance in this regime. We note that this bound beats the optimal randomized algorithm for the buyback problem with no information, for all f [11].⁶

THEOREM 16. There is an algorithm which achieves a competitive ratio of at least

$$\max_{0 \leq x \leq 1} \frac{(1-x) \cdot x^{\frac{f}{1+f}}}{1-x+x\frac{2f+1}{f+1}} \geq \frac{1}{\frac{f}{1+f} + \left(2 + \frac{1}{f}\right)^{\frac{f}{1+f}}}.$$

We remark that the last expression in the theorem follows from plugging in $x = \frac{f}{1+2f}$. For $f \in (0, \frac{1}{2})$, it is easy to see that this is $1 - \Theta\left(f \log\left(\frac{1}{f}\right)\right)$, indeed, $(2 + \frac{1}{f})^{\frac{f}{1+f}} < (\frac{2}{f})^f = e^{f \ln \frac{2}{f}} < 1 + 2f \ln \frac{2}{f}$, so the competitive ratio is at least $\frac{1}{\frac{f}{1+f} + 1 + 2f \ln \frac{2}{f}} > 1 - 2f \ln \frac{4}{f}$.

PROOF. From Lemmas 14 and 15, we get: If Z is a Poisson random variable with mean λ (with $e^{-\lambda} = \Pr[X_{\max} < T]$), and $0 \leq A \leq 1$ is any random variable (possibly correlated with X_1, \dots, X_n), then

$$\begin{aligned} \mathbb{E}[ALG] &\geq T \Pr[X_{\max} \geq T] + \mathbb{E}\left[\left(\frac{X_{\max}}{(1+f)^{S'}} - T\right)_+\right] \\ &\geq T \Pr[X_{\max} \geq T] + \mathbb{E}\left[A \cdot \left(\frac{X_{\max}}{(1+f)^Z} - T\right)\right] \\ &= T(\Pr[X_{\max} \geq T] - \mathbb{E}[A]) + \mathbb{E}\left[A \cdot \frac{X_{\max}}{(1+f)^Z}\right]. \end{aligned}$$

Recall that $\Pr[X_{\max} \geq T] = 1 - e^{-\lambda}$. Taking $A = c \mathbb{1}(X_{\max} \geq T)$:

$$\begin{aligned} \mathbb{E}[ALG] &\geq T(\Pr[X_{\max} \geq T] - \mathbb{E}[A]) + \mathbb{E}\left[A \cdot \frac{X_{\max}}{(1+f)^Z}\right] \\ &= T(1 - e^{-\lambda})(1 - c) \\ &\quad + c \cdot \mathbb{E}\left[\frac{1}{(1+f)^Z}\right] \cdot \mathbb{E}[X_{\max} \mathbb{1}(X_{\max} \geq T)] \\ &= T(1 - e^{-\lambda})(1 - c) + ce^{-\frac{\lambda f}{1+f}} \mathbb{E}[X_{\max} \mathbb{1}(X_{\max} \geq T)]. \end{aligned}$$

⁶The bound presented here is worse than the classical prophet inequality when $f \rightarrow \infty$ (in which case our bound tends to $1/3$) as we are aiming primarily to match the asymptotic upper bound we obtained of $1 - \Theta\left(f \log\left(\frac{1}{f}\right)\right)$.

where $\mathbb{E}\left[\frac{1}{(1+f)^Z}\right] = e^{-\frac{\lambda f}{1+f}}$ follows from an elementary computation for the Poisson random variable Z . Since

$$\begin{aligned} \mathbb{E}[X_{\max}] &\leq T \Pr[X_{\max} < T] + \mathbb{E}[X_{\max} \mathbb{1}(X_{\max} \geq T)] \\ &= T e^{-\lambda} + \mathbb{E}[X_{\max} \mathbb{1}(X_{\max} \geq T)], \end{aligned}$$

setting

$$(1 - e^{-\lambda})(1 - c) = ce^{-\frac{\lambda f}{1+f}} e^{-\lambda}$$

guarantees a competitive ratio of $ce^{-\frac{\lambda f}{1+f}}$. Letting $x = e^{-\lambda}$, and solving for c , we get

$$c = \frac{1-x}{1-x+x\frac{2f+1}{f+1}},$$

which means we are guaranteed a competitive ratio of

$$\max_{0 \leq x \leq 1} \frac{(1-x) \cdot x^{\frac{f}{1+f}}}{1-x+x\frac{2f+1}{f+1}}. \quad \square$$

We remark that with $A = \mathbb{1}(X_{\max} \geq T) \mathbb{1}(Z \leq 1)$, we could have obtained another, different lower bound on the competitive ratio, better for medium values of f . Another reasonable option is $A = c_1 \mathbb{1}(X_{\max} \geq T) c_2^Z$, which can recover approximately the correct result both near $f = 0$ and $f = \infty$.

6 CONCLUSION AND OPEN PROBLEMS

In this paper, we studied a version of the prophet inequality where recourse is possible by paying a penalty f times the value of the discarded item. For $f \geq 1$ we showed that the optimal online algorithm attains a competitive ratio of $\frac{1+f}{1+2f}$ against the optimal offline solution. We also designed an algorithm which is asymptotically optimal in the regime of small $f \geq 0$, and studied a simple threshold-greedy algorithm.

Open problems: The exploration of optimal competitive ratios $\alpha(f)$ for $f < 1$ remains an intriguing avenue for future research. We conjecture that the 3 variable example presented in section 4 is the worst-case instance on 3 variables, and that we can use this example to construct a solution to the continuous flow problem, using 3 blocks, in the regime where $\frac{1}{3} \leq f < 1$. More generally, we conjecture that there is a sequence of intervals I_k (where $I_2 = [1, \infty)$, $I_3 = [1/3, 1)$), and the optimal factor for I_k is determined by the worst-case instance on k variables, and the optimal flow has a k -block structure. The problem of identifying the worst case instance for $k \geq 3$ remains open.

ACKNOWLEDGMENTS

Rad Niazadeh was partially supported by Asness Junior Faculty Fellowship at The University of Chicago Booth School of Business. Jan Vondrák was supported by NSF award 2127781.

REFERENCES

- [1] Melika Abolhassani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi Hajiaghayi, Robert Kleinberg, and Brendan Lucier. 2017. Beating $1-1/e$ for ordered prophets. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 61–71.
- [2] Marek Adamczyk and Michał Włodarczyk. 2018. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 790–801.

- [3] Daniel Adelman. 2004. A price-directed approach to stochastic inventory/routing. *Operations Research* 52, 4 (2004), 499–514.
- [4] Saeed Alaei. 2014. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM J. Comput.* 43, 2 (2014), 930–972.
- [5] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. 2012. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. 18–35.
- [6] Nima Anari, Rad Niazadeh, Amin Saberi, and Ali Shamel. 2019. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 91–92.
- [7] David Assaf and Ester Samuel-Cahn. 2000. Simple ratio prophet inequalities for a mortal with multiple choices. *Journal of applied probability* 37, 4 (2000), 1084–1091.
- [8] Yossi Azar, Avrim Blum, and Yishay Mansour. 2003. Combining online algorithms for rejection and acceptance. In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*. 159–163.
- [9] Moshe Babaioff, Jason Hartline, and Robert Kleinberg. 2008. Selling banner ads: Online algorithms with buyback. In *Fourth workshop on ad auctions*.
- [10] Ashwinkumar Badanidiyuru. 2011. Buyback problem-approximate matroid intersection with cancellation costs. In *International Colloquium on Automata, Languages, and Programming*. Springer, 379–390.
- [11] Ashwinkumar Badanidiyuru and Robert Kleinberg. 2009. Randomized online algorithms for the buyback problem. In *International Workshop on Internet and Network Economics*. Springer, 529–536.
- [12] Richard Bellman. 1954. The theory of dynamic programming. *Bull. Amer. Math. Soc.* 6, 6 (1954), 503–515.
- [13] Niv Buchbinder, Moran Feldman, and Roy Schwartz. 2014. Online submodular maximization with preemption. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1202–1216.
- [14] Ran Canetti and Sandy Irani. 1995. Bounding the power of preemption in randomized scheduling. In *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*. 606–615.
- [15] Kamalika Chaudhuri, Constantinos Daskalakis, Robert D Kleinberg, and Henry Lin. 2009. Online bipartite perfect matching with augmentations. In *IEEE INFOCOM 2009*. IEEE, 1044–1052.
- [16] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. 2010. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*. 311–320.
- [17] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. 2011. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *Proceedings of the forty-third annual ACM symposium on Theory of computing*. 783–792.
- [18] Florin Constantin, Jon Feldman, S Muthukrishnan, and Martin Pál. 2009. An online mechanism for ad slot reservations with cancellations. In *Proceedings of the twentieth annual ACM-SIAM symposium on discrete algorithms*. SIAM, 1265–1274.
- [19] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2017. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. 169–186.
- [20] Jose Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. 2019. Recent developments in prophet inequalities. *ACM SIGecom Exchanges* 17, 1 (2019), 61–70.
- [21] Atanas Dinev and S Matthew Weinberg. 2023. Simple and Optimal Online Contention Resolution Schemes for k -Uniform Matroids. *arXiv preprint arXiv:2309.10078* (2023).
- [22] Paul Dutting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. 2020. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.* 49, 3 (2020), 540–582.
- [23] Farbod Ekbatani, Yiding Feng, and Rad Niazadeh. 2022. Online Matching with Cancellation Costs. *arXiv preprint arXiv:2210.11570* (2022).
- [24] Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. 2024. Prophet Inequalities with Cancellation Costs. [arXiv:2404.00527](https://arxiv.org/abs/2404.00527) [cs.DS]
- [25] Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. 2024. Prophet Inequalities with Cancellation Costs. (2024). https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4779633
- [26] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. 2017. Prophet secretary. *SIAM Journal on Discrete Mathematics* 31, 3 (2017), 1685–1701.
- [27] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. 2022. On the Significance of Knowing the Arrival Order in Prophet Inequality. *arXiv preprint arXiv:2202.09215* (2022).
- [28] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. 2022. Prophet matching with general arrivals. *Mathematics of Operations Research* 47, 2 (2022), 878–898.
- [29] Tomer Ezra, Michal Feldman, and Ilan Nehama. 2018. Prophets and Secretaries with Overbooking. In *Proceedings of the 2018 ACM Conference on Economics and Computation*. 319–320.
- [30] Uriel Feige. 2006. On maximizing welfare when utility functions are subadditive. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*. 41–50.
- [31] Uriel Feige, Michal Feldman, Nicole Immorlica, Rani Izsak, Brendan Lucier, and Vasilis Syrgkanis. 2015. A unifying hierarchy of valuations with complements and substitutes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.
- [32] Michal Feldman, Hu Fu, Nick Gravin, and Brendan Lucier. 2013. Simultaneous auctions are (almost) efficient. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*. 201–210.
- [33] Moran Feldman, Ola Svensson, and Rico Zenklusen. 2016. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*. SIAM, 1014–1033.
- [34] Yiding Feng, Rad Niazadeh, and Amin Saberi. 2022. Near-optimal bayesian online assortment of reusable resources. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 964–965.
- [35] Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. 2007. Automated online mechanism design and prophet inequalities. In *AAAI*, Vol. 7. 58–65.
- [36] Xin Han, Yasushi Kawase, and Kazuhisa Makino. 2014. Online unweighted knapsack problem with removal cost. *Algorithmica* 70, 1 (2014), 76–91.
- [37] Jason D Hartline and Tim Roughgarden. 2009. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. 225–234.
- [38] Theodore P Hill and Robert P Kertz. 1982. Comparisons of stop rule and supremum expectations of iid random variables. *The Annals of Probability* (1982), 336–345.
- [39] Nicole Immorlica, Sahil Singla, and Bo Waggoner. 2020. Prophet inequalities with linear correlations and augmentations. In *Proceedings of the 21st ACM Conference on Economics and Computation*. 159–185.
- [40] Jiashuo Jiang, Will Ma, and Jiawei Zhang. 2022. 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, 1221–1246.
- [41] Jiashuo Jiang, Will Ma, and Jiawei Zhang. 2023. Tightness without Counterexamples: A New Approach and New Results for Prophet Inequalities. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 909–909.
- [42] Robert Kleinberg and Seth Matthew Weinberg. 2012. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. 123–136.
- [43] Ulrich Krengel and Louis Sucheston. 1978. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces* 4 (1978), 197–266.
- [44] Euiwoong Lee and Sahil Singla. 2018. Optimal online contention resolution schemes via ex-ante prophet inequalities. *arXiv preprint arXiv:1806.09251* (2018).
- [45] Allen Liu, Renato Paes Leme, Martin Pál, Jon Schneider, and Balasubramanian Sivan. 2021. Variable Decomposition for Prophet Inequalities and Optimal Ordering. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 692–692.
- [46] Brendan Lucier. 2017. An economic view of prophet inequalities. *ACM SIGecom Exchanges* 16, 1 (2017), 24–47.
- [47] Roger B Myerson. 1981. Optimal auction design. *Mathematics of operations research* 6, 1 (1981), 58–73.
- [48] Rad Niazadeh, Amin Saberi, and Ali Shamel. 2018. Prophet inequalities vs. approximating optimum online. In *Web and Internet Economics: 14th International Conference, WINE 2018*. Springer, 356–374.
- [49] Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. 2021. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 763–764.
- [50] Tristan Pollner, Mohammad Roghani, Amin Saberi, and David Wajc. 2022. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 321–322.
- [51] Aviad Rubinfeld. 2016. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 324–332.
- [52] Aviad Rubinfeld and Sahil Singla. 2017. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 1671–1687.
- [53] Ester Samuel-Cahn. 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability* (1984), 1213–1216.
- [54] Danny Segev and Sahil Singla. 2021. Efficient approximation schemes for stochastic probing and prophet problems. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 793–794.
- [55] Qiqi Yan. 2011. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*. SIAM, 710–719.

Received 13-NOV-2023; accepted 2024-02-11