

Prophet Inequalities with Linear Correlations and Augmentations

Nicole Immorlica^{*}Sahil Singla[†]Bo Waggoner[‡]

May 26, 2020

Abstract

In a classical online decision problem, a decision-maker who is trying to maximize her value inspects a sequence of arriving items to learn their values (drawn from known distributions), and decides when to stop the process by taking the current item. The goal is to prove a “prophet inequality”: that she can do approximately as well as a prophet with foreknowledge of all the values. In this work, we investigate this problem when the values are allowed to be correlated. Since non-trivial guarantees are impossible for arbitrary correlations, we consider a natural “linear” correlation structure introduced by Bateni et al. [BDHS15] as a generalization of the common-base value model of Chawla et al. [CMS15].

A key challenge is that threshold-based algorithms, which are commonly used for prophet inequalities, no longer guarantee good performance for linear correlations. We relate this road-block to another “augmentations” challenge that might be of independent interest: many existing prophet inequality algorithms are not robust to slight increase in the values of the arriving items. We leverage this intuition to prove bounds (matching up to constant factors) that decay gracefully with the amount of correlation of the arriving items. We extend these results to the case of selecting multiple items by designing a new $(1 + o(1))$ approximation ratio algorithm that is robust to augmentations.

^{*}(nicimm@microsoft.com) Microsoft Research.

[†](singla@cs.princeton.edu) Princeton University and Institute for Advanced Study. Supported in part by the Schmidt Foundation.

[‡](bwag@colorado.edu) University of Colorado.

Contents

1	Introduction	3
1.1	Linear Correlations Model	3
1.2	Techniques and the Augmented Prophets Problem	5
1.3	Related Work	6
2	Model and Fixed-Threshold Algorithms	7
2.1	Model and Notation	7
2.2	Fixed-Threshold Algorithms	7
3	Our Approach of Handling Correlations via Augmentations	8
4	Selecting a Single Item	10
4.1	Bounded Column Sparsity	11
4.2	Bounded Row Sparsity	11
4.3	Lower Bounds	14
5	Selecting Multiple Items	16
5.1	Bounded Column Sparsity	16
5.2	Bounded Row Sparsity	17
5.3	Multiple-Items Augmentation Lemma	18
A	Unweighted Linear Correlations	22
A.1	Notation and Proof Overview	23
A.2	Proof	24
B	Negatively Associated Values	25
C	Bounded s_{col} and Small Cardinality Constraint	26
D	Missing Proofs	27
D.1	Missing Proofs from Section 2	27
D.2	Missing Proofs from Section 5.3	28

1 Introduction

In classic optimal-stopping problems, a decision-maker wishes to select one of a set $[n] = \{1, \dots, n\}$ of options whose values are distributed according to a known joint distribution. Option i materializes at time i , revealing its value X_i . The decision-maker must then either select option i , receiving a value of X_i , or permanently reject it and continue. Her goal is to choose an option whose value, in expectation over her selection algorithm and the randomness in the problem instance, obtains at least a $1/\alpha$ -fraction of the expected maximum value for some approximation ratio $\alpha \geq 1$. Such an approximation is referred to as a “prophet inequality” as it compares the decision-maker’s performance to that of a prophet who knows the realizations of all values in advance and can always stop at the maximum. Examples of optimal-stopping problems include hiring, in which an employer interviews a sequence of candidates and must make a hiring decision on the spot; or house-buying in a sellers’ market, in which a buyer must make an offer at the open house. These optimal-stopping problems became more popular in the last 15 years particularly because of their applications in mechanism design. E.g., an α -prophet inequality often implies a posted-pricing mechanism that gets a $1/\alpha$ -fraction of the maximum welfare for a sequence of bidders arriving online (see related work in Section 1.3).

The approximability of optimal stopping problems depends heavily on the correlations of the option values. In the case where all the values are independent, a tight 2 approximation ratio was shown by Krengel and Sucheston [KS78]. In 1984, Samuel-Cahn [SC84] presented a simple *median-of-maximum* “threshold-based” rule with the same performance: compute τ as the median of the distribution of the maximum value, and stop at the first X_i exceeding τ . Other threshold rules are also known to obtain 2 approximation, e.g., Kleinberg and Weinberg [KW12] showed this for $\tau = \frac{1}{2} \mathbb{E} [\max_i X_i]$. When the values are negatively correlated, the problem intuitively becomes even easier than the independent case: observing and rejecting a low X_i increases the chances of seeing a high $X_{i'}$ in the future (and vice versa accepting a large X_i decreases the chances of having missed a high $X_{i'}$ in the future). E.g., a simple implication of threshold-based algorithms is a 2 approximation ratio for negatively associated (a form of negative correlation) values; see Appendix B. Rinott and Samuel-Cahn [RSC⁺87, RSC91] indeed show that the value of the optimal stopping algorithm for negative correlations exceeds that of the independent case, holding the marginals fixed. On the other hand, with general and positive correlation structures, no algorithm can guarantee better than $\Omega(n)$ approximation ratio (as is known from Hill and Kertz [HK92], and will also be a special case of our lower bounds). Therefore, the question is how to impose a structure on the correlations that both models interesting applications and allows for better bounds.

1.1 Linear Correlations Model

We consider a *linear correlations model* in which there exists a set of m *independent* variables Y_1, \dots, Y_m , with each option value X_i being a positive linear combination of some subset:

$$\mathbf{X} = A \cdot \mathbf{Y}$$

where A is a nonnegative matrix. The algorithm is given the matrix A and the distributions of all the Y_j s, but when X_i arrives, it only finds X_i and not any of the realizations of the Y_j s. This model was introduced in an auctions context by Bateni et al. [BDHS15], where it was inspired by

the common-base value model of Chawla et al. [CMS15]. It has two natural parameters capturing the degree of correlation of an instance. If each row of A has at most s_{row} nonzero entries (*row sparsity* s_{row}), this implies that each X_i only depends on at most s_{row} different Y_j s. On the other hand, if each column of A has at most s_{col} nonzero entries (*column sparsity* s_{col}), this implies that each Y_j is only relevant to at most s_{col} different X_i s.

General applications. Linear correlations occur in applications when each option i is defined by the degree A_{ij} to which it exhibits each feature $j \in [m]$. The value $\sum_j A_{ij}Y_j$ of the option is then determined by the values Y_j of its features, which is unknown to the decisionmaker. Of particular interest in our setting are applications with many features, such as the hiring and house-buying examples often used to motivate prophet problems. Other relevant applications include selecting hotel rooms, restaurants, and movies. Here we elaborate on the hiring and house-buying examples, noting how they naturally exhibit column or row sparsity.¹

In a hiring application, the features of a candidate might include where he received his education, his major, his work experience in each relevant industry, aspects of his personality, etc. When the employer interviews a candidate, she learns how much she likes him, but not how to attribute her value for the candidate to particular features (every school/major/industry is a different feature). If the candidate pool is diverse, so that candidates come from many different schools/majors/industries, we might expect the instance to have a low column sparsity. Similarly, in house-buying, the features of a house might include the commute time, the square footage, and various bells-and-whistles like the existence of a patio, a hot-tub, a roof-deck, the number of parking spaces, if any. Again, the value of a house is a linear combination of the value of its features, but when seeing a house, the buyer may only be able to access and articulate an overall valuation. If each house has a limited number of bells-and-whistles, we expect the instance to have low row-sparsity.

Mechanism design applications: Welfare for linearly correlated values. Prophet inequalities can directly imply social welfare and revenue guarantees for sequential posted-price mechanisms [CHMS10, KW12]. In the simplest model, a single item is for sale to a sequence of arriving bidders with values X_1, \dots, X_n , drawn from distributions known to the seller. A threshold- τ stopping rule immediately translates to a posted price τ . The item is purchased by the first bidder whose value satisfies $X_i \geq \tau$. In particular, social welfare is the value of the bidder who purchases the item, so a prophet inequality directly translates to a social welfare guarantee.² While we show that threshold-based policies fail for linearly correlated values, we obtain positive results with *inclusion-threshold* policies. These correspond to offering a fixed posted price to a predetermined subset of buyers, while the others are automatically rejected.

For linearly correlated bidder values, our positive results immediately imply social welfare guarantees using such inclusion-posted-price mechanisms. Here, linear correlations capture some component of common values in bidder preferences. Namely, there are different features Y_1, \dots, Y_m of the object, with bidder i placing weight A_{ij} on feature j . In the mechanism-design setting, it is

¹Note however that our bounds are expressed in terms of the minimum of row and column sparsity of an instance, and hence apply even to instances with high row/column sparsity.

²Revenue guarantees, at least in the classical independent- X_i model, can be obtained using a threshold in virtual value space.

particularly natural to make our assumption that the decisionmaker (here, the seller) is not able to access Y_j when value X_i arrives.

Results. We start from the observation that threshold-based algorithms cannot guarantee good approximation ratios as soon as *any* correlations are introduced. Therefore, we define *inclusion-threshold* algorithms that probabilistically include a subset of the arrivals for consideration and take the first arrival *in this subset* to exceed a threshold.

We first design an inclusion-threshold algorithm to obtain an $O(s_{\text{col}})$ approximation ratio, i.e., a guarantee that degrades gracefully as the amount of correlation increases, from the known $O(1)$ bound for the independent case to the known $\Theta(n)$ worst-case bound. Then, we design a more complex inclusion-threshold algorithm to obtain an $O(s_{\text{row}})$ approximation ratio, i.e., another gracefully degrading guarantee. Together, these prove an $O(\min\{s_{\text{row}}, s_{\text{col}}\})$ approximation guarantee for the linear correlations model. We then design a lower bound instance and prove that this is tight up to constants, i.e., no algorithm can guarantee better than an $\Omega(\min\{s_{\text{row}}, s_{\text{col}}\})$ approximation.

Main Result 1 (Informal Theorem 4.1). For the linearly correlated prophet problem, there exists an $O(\min\{s_{\text{col}}, s_{\text{row}}\})$ approximation ratio algorithm.

Finally, we extend these results to the case of selecting a subset of up to r of the arriving options with a goal of maximizing their expected sum (also known as an *r -uniform matroid* constraint). It is known that for independent distributions, $1 + o(1)$ approximation ratio prophet inequality algorithms are possible for the case of large r [HKS07, Ala11]. We show a similar result for linearly correlated instances with bounded column sparsity s_{col} .

Main Result 2 (Informal Theorem 5.1). For the linearly correlated prophet problem where we select r options, there exists a $1 + o(1)$ approximation ratio algorithm when $r \gg s_{\text{col}}$.

The case of bounded row sparsity, however, turns out to be harder: regardless of r , no algorithm can guarantee better than an $\Omega(s_{\text{row}})$ approximation ratio for unbounded s_{col} , as in the $r = 1$ case.

1.2 Techniques and the Augmented Prophets Problem

A crucial technique for our results is to introduce and solve the *augmented prophets* problem. The idea is to suppose we have an instance with independent random variables Z_1, \dots, Z_n and an algorithm, say a threshold rule, guaranteeing some approximation ratio. Now suppose we “augment” the instance by sending instead the values $X_1 := Z_1 + W_1, \dots, X_n := Z_n + W_n$ where the W_i s are nonnegative. Does the algorithm (which does not know W_i s) continue to guarantee its original approximation ratio (measured against the maximum Z_i)? One would hope so, as each arriving option has only increased, while the benchmark has not. However, this turns out not to be true for the median-of-maximum threshold rule. E.g., if $Z_i \sim \text{Bernoulli}(p)$ i.i.d. for $p \ll \frac{1}{n}$, the median is zero, and augmenting the first arrival to a miniscule positive value causes the (strict) median threshold rule to always take it, resulting in an arbitrarily poor approximation. Luckily, we show that the half-of-expected-maximum threshold algorithm retains its approximation guarantees, even when the W_i s are chosen by an adversary depending on the past $X_{i'}s$, i.e., $i' < i$.

Armed with this “augmentation lemma”, we use subsampling to obtain inclusion sets of arrivals $\{X_i\}$ with significant independent portions $\{Z_i\}$. In the bounded s_{col} case, direct subsampling of arrivals succeeds. The case where s_{row} is bounded but Y_j can appear in any number of arrivals is more challenging. We show it suffices to obtain a contention-resolution style subsampling of the arrivals such that each Y_j is well-represented, but only with its maximum coefficient A_{ij} . We then use a graph-theoretic argument to construct such a scheme.

The augmented prophets problem is also our key technique for the linearly correlated prophets problem with an r -uniform matroid constraint. In this case, however, we notice that none of the existing $1 + o(1)$ algorithms are robust to augmentations. Hence we design a new $1 + o(1)$ algorithm and prove its robustness using a much more sophisticated analysis involving a sequence of thresholds and “buckets” with different cardinality constraints. By combining this augmentation result with random partitioning of the input, we obtain the $1 + o(1)$ approximation for the r -uniform matroid problem with fixed s_{col} .

Main Result 3 (Informal Lemma 3.2 and 3.4). For the augmented prophets problem, there exists a 2 approximation ratio algorithm when selecting a single option and a $1 + o(1)$ approximation ratio algorithm when selecting $r \gg 1$ options.

One can also view the augmented prophets problem as capturing correlations induced by a mischievous wish-granting genie who awards bonuses $W_i \geq 0$ at each step, but tries to choose them so as to worsen the algorithm’s performance. We think this problem is of independent interest and will find further applications in designing robust prophet inequality algorithms.

1.3 Related Work

The last decade has seen significant interest in prophet inequalities motivated by their applications in mechanism design. Many works focus on multiple-choice prophet inequality problems. This includes prophet inequalities for uniform matroids in [HKS07, Ala11], for general matroids in [CHMS10, Yan11, KW12, FSZ16, LS18], for matchings and combinatorial auctions in [AHL12, FGL15, DFKL17, EHKS18], and for arbitrary packing constraints in [Rub16, RS17]. There has also been a lot of work on variants of prophet inequalities: the prophet secretary problem where the values arrive in a random order [EHL17, AEE⁺17, CFH⁺17, EHKS18, ACK18, LS18, CSZ19], and the limited information setting where we only have sample-access to distributions [AKW14, CDFS19, RWW20]. All these works assume mutually independent values, whereas capturing correlations and designing robust algorithms is the main challenge in our work.

Rinott and Samuel-Cahn [RSC⁺87, RSC91, RSC92] study correlated prophet inequalities. However, their techniques are not applicable to our work because their positive results hold only for negatively correlated values. Furthermore, their benchmark is the expected maximum of independent values having the same marginal distributions. This benchmark could be a factor n larger than the expected maximum for positively correlated values.

Our approach via the augmented prophets problem is also related to the line of work on designing robust stochastic optimization algorithms. Since algorithms that assume *known* input distributions tend to over-fit, here the goal is to design algorithms that are robust to adversarial

noise (see [Dia18, Moi18, DKK⁺16, LRV16, CSV17, DKK⁺18, EKM18, LMPL18, BGSZ20] and references therein). Our single-item and multiple-items augmentation algorithms can be seen as robust prophet inequality algorithms that retain their guarantees even when the input distributions are augmented by an adversary. Another relevant reference is that of Dütting and Kesselheim [DK19], which gives prophet inequalities assuming only probability distributions that are ϵ -close (in some metric) to the true distributions. Their technical results, however, are not useful here because augmented distributions can be very far from the original distributions.

2 Model and Fixed-Threshold Algorithms

2.1 Model and Notation

In the *linear correlations* model, there are n random variables X_1, X_2, \dots, X_n that linearly depend on m independent nonnegative random variables (sometimes called *features*) Y_1, Y_2, \dots, Y_m as

$$\mathbf{X} = A \cdot \mathbf{Y},$$

where matrix A only contains non-negative entries. Let s_{row} denote the row-sparsity of A (maximum number of nonzero entries in any row) and s_{col} denote the column-sparsity of A (maximum number of nonzero entries in any column).

An online algorithm is initially given A and the distributions of Y_1, \dots, Y_m . Then, it observes the realizations of X_1, \dots, X_n one at a time. After observing X_i , the algorithm decides either to stop and take the reward X_i , ending the process, or to reject X_i and continue to X_{i+1} . Given such an algorithm ALG, we abuse notation by writing ALG for the reward of the algorithm, a random variable. The algorithm has an *approximation ratio* of α for some $\alpha(n, s_{\text{row}}, s_{\text{col}})$ if for all $n, s_{\text{row}}, s_{\text{col}}$ and all instances of the problem with these parameters,

$$\mathbb{E}[\text{ALG}] \geq \frac{1}{\alpha(n, s_{\text{row}}, s_{\text{col}})} \cdot \mathbb{E}[\max_i X_i].$$

Such a guarantee is often called a *prophet inequality* because it compares the algorithm to a “prophet” that can predict the realizations of all X_i in advance and take $\max_i \{X_i\}$ every time.

We use the notation $(\cdot)^+$ to mean $\max\{\cdot, 0\}$.

Our examples frequently use random variables that are either zero or some fixed positive value. We say the variable is *active* if it takes its positive value. We sometimes say that X_i “includes” Y_j if $A_{ij} > 0$.

2.2 Fixed-Threshold Algorithms

A *fixed-threshold* algorithm selects a single threshold τ and takes the first arrival X_i that exceeds τ . We refer to such an algorithm as ALG_τ . Fixed-threshold algorithms have been very successful in prophet inequality design. However, our first result shows their severe limitation for even mildly correlated prophet inequalities.

Lemma 2.1. *In the linear correlations model, even for $s_{\text{row}} = s_{\text{col}} = 2$ there exist instances where every fixed threshold τ algorithm ALG_τ has an approximation ratio at least $\Omega(n)$.*

The full proof is deferred to Appendix D, but the instance is important and described next. Intuitively, the problem is that cases where an arrival X_i just crosses the threshold may be correlated with some later $X_{i'}$ being very large. Taking X_i prevents the algorithm from ever obtaining the gains from $X_{i'}$. Our proof uses the following “tower” variables, which will also be useful later.

Definition 2.2 (Tower Y variables). Given $\epsilon > 0$, define the *tower Y variables* as $Y_i = \frac{1}{\epsilon^i}$ with probability ϵ^i and $Y_i = 0$ otherwise for each $i \in \{1, \dots, m\}$.

Example 2.3 ($s_{\text{row}} = s_{\text{col}} = 2$ tower instance). Take the tower Y variables. Let A be an $n \times n$ matrix with entry $A_{i,i} = 1$ for all i and $A_{i,i+1} = \epsilon$ for $i \in [1, n-1]$. All other entries are 0. Visually,

$$X_1 = Y_1 + \epsilon Y_2, \quad X_2 = Y_2 + \epsilon Y_3, \quad \dots, \quad X_n = Y_n.$$

We have $s_{\text{col}} = s_{\text{row}} = 2$. The point is that if X_i is nonzero, then almost certainly $X_i = \frac{1}{\epsilon^i}$. But in this case, a threshold algorithm cannot distinguish between the more likely case that $Y_i = \frac{1}{\epsilon^i}$, in which case it should stop and take X_i , and the unlikely case that $Y_i = 0$ and $Y_{i+1} = \frac{1}{\epsilon^{i+1}}$, in which case it should wait and take X_{i+1} .

Indeed, these coefficients of matrix A play an important role, and in Appendix A we show that when entries of A are restricted to being only 0 or 1, there exists a constant-factor approximation fixed-threshold algorithm. Roughly this happens because each Y_j has “limited influence” on any X_i : either Y_j appears with coefficient 0 and has no influence on X_i , or it appears with coefficient 1 and has the same influence on every such X_i .

Theorem 2.4. *The unweighted linear correlations problem has a fixed threshold constant-factor approximation algorithm.*

This raises the question whether for a general matrix A any policy can achieve a better approximation, let alone a simple policy. Our positive results will show that relatively simple *inclusion-threshold* algorithms can achieve tight prophet inequalities.

Definition 2.5. An *inclusion-threshold* algorithm selects a subset $S \subseteq \{1, \dots, n\}$ and threshold τ , possibly both at random, and selects the first X_i such that $i \in S$ and $X_i \geq \tau$. In other words, it commits to a subset S of arrivals and applies a threshold policy to those X_i , ignoring the others.

3 Our Approach of Handling Correlations via Augmentations

In analysis of prophet inequalities, the problem is to upper-bound the expected maximum of the variables X_i as compared to one’s algorithm. An important and common approach is to use the fact that for any threshold τ ,

$$\mathbb{E}[\max_i X_i] \leq \tau + \mathbb{E}[\max_i (X_i - \tau)^+] \leq \tau + \sum_i \mathbb{E}[(X_i - \tau)^+]. \quad (1)$$

When all the arrivals X_i are independent, it is known that one can always select τ such that the left and right sides of (1) differ by at most a constant factor, i.e., $\frac{e}{e-1} \approx 1.6$ (this is related to the correlation gap [ADSY12]). In fact, the prototypical prophets analysis shows that setting some threshold τ allows ALG_τ to approximate the right hand side up to a constant factor. However, when $\{X_i\}$ are correlated, this could be a very loose upper bound. E.g., consider $X_1 = \dots = X_n = Y_1 \sim \text{Bernoulli}(p)$. The left side equals p while the right side equals $\tau + np(1 - \tau) = np + \tau(1 - np) \geq np$ for

$p < \frac{1}{n}$. So the right side can be a factor n larger than the left, and we cannot hope to approximate the right side with any algorithm.

One approach to correlated prophets *could be* a direct analysis of the right-hand side of (1) in cases of limited correlation. Here, we take a different approach. The first key idea is to use inclusion-threshold algorithms. To see why, consider a first attempt: discard certain X_i such that we are only left with a subset that are all independent of each other. Now a standard prophet algorithm that only considers these X_i would obtain a constant factor of the maximum in this subset. One could then hope to argue that this subset's maximum approximates the original maximum up to a factor depending on the amount of correlation. Indeed, one can show that this approach succeeds on the tower instance in Example 2.3 with $s_{\text{row}} = s_{\text{col}} = 2$, e.g., by including every other X_i . But in general this approach cannot give tight bounds. This is because each X_i contains s_{row} variables, each of which can appear in up to $s_{\text{col}} - 1$ other $X_{i'}$, so including X_i requires eliminating $\approx s_{\text{row}} s_{\text{col}}$ other variables. Our goal is to achieve approximations to within $\min\{s_{\text{row}}, s_{\text{col}}\}$ factors even if $\max\{s_{\text{row}}, s_{\text{col}}\} = n$.

So in addition to including only a subset of X_i , we will use a second key idea: decompose each variable as $X_i = Z_i + W_i$, where the Z_i s satisfy independence requirements and W_i s are viewed as “bonus” augmentations. We will show that $\mathbb{E}[\max_i Z_i]$ is an approximation to $\mathbb{E}[\max_i X_i]$. Then we will compete with $\mathbb{E}[\max_i Z_i]$. However, the augmentations add an additional challenge, requiring us to solve the following problem.

Definition 3.1 (Single-Item Augmented Prophets Problem). The algorithm is given the distributions of a set of independent nonnegative random variables Z_1, \dots, Z_n . Then, it observes one at a time the realizations of $X_i = Z_i + W_i$ for $i \in [1, n]$, where each W_i is nonnegative and Z_i is independent of X_1, \dots, X_{i-1} for each i (but W_i could depend on X_1, \dots, X_i). The algorithm chooses at each step whether to continue or stop and obtain value X_i . It must compete with $\mathbb{E}[\max_i Z_i]$.

One can view the augmented prophets problem as capturing correlations induced by a mischievous genie who awards bonuses $W_i \geq 0$ at each step so as to worsen the algorithm's performance. We note that the genie cannot base her choices on the future, i.e., W_i is a random variable that may be correlated with $Z_{i'}$ if $i' \leq i$ but not if $i' > i$.

Intuitively, it might seem like algorithms for prophets problems should continue to perform well, as the genie can only increase the rewards at each step. However, this is not true for the classical median stopping rule, i.e, for $\tau = \text{the median of } \max_i Z_i$ (see an example in Section 1.2). Luckily, a different threshold rule is robust:

Lemma 3.2 (Single-item Augmentation lemma). *For the augmented prophets problem, a fixed threshold algorithm with $\tau = \frac{1}{2}\mathbb{E}[\max_i Z_i]$ guarantees $\frac{\mathbb{E}[\text{ALG}_\tau]}{\mathbb{E}[\max_i Z_i]} \geq \frac{1}{2}$.*

Proof. We “augment” a standard prophet inequality proof. Let $P = \Pr[\max_i X_i \geq \tau]$. Now,

$$\begin{aligned} \mathbb{E}[\text{ALG}_\tau] &= P \cdot \tau + \sum_i \Pr[X_{i'} < \tau \ (\forall i' < i)] \cdot \mathbb{E}[(X_i - \tau)^+ \mid X_{i'} < \tau \ (\forall i' < i)] \\ &\geq P \cdot \tau + \sum_i (1 - P) \cdot \mathbb{E}[(X_i - \tau)^+ \mid X_{i'} < \tau \ (\forall i' < i)] \end{aligned}$$

because ALG_τ selects no element with probability $1 - P$. Nonnegativity of W_i implies $X_i \geq Z_i$, so

$$\mathbb{E}[\text{ALG}_\tau] \geq P \cdot \tau + \sum_i (1 - P) \cdot \mathbb{E}[(Z_i - \tau)^+ \mid X_{i'} < \tau \ (\forall i' < i)]$$

$$= P \cdot \tau + (1 - P) \cdot \mathbb{E} \left[\sum_i (Z_i - \tau)^+ \right]$$

because Z_i is independent of the event $\{X_{i'} < \tau \ (\forall i' < i)\}$. Since $\sum_i (Z_i - \tau)^+ \geq \max_i (Z_i - \tau)^+$,

$$\begin{aligned} \mathbb{E}[\text{ALG}_\tau] &\geq P \cdot \tau + (1 - P) \cdot \mathbb{E} \left[\max_i (Z_i - \tau)^+ \right] \\ &\geq P \cdot \tau + (1 - P) \cdot \mathbb{E} \left[\max_i Z_i - \tau \right] = P \cdot \tau + (1 - P)\tau = \tau. \end{aligned}$$

This proves that $\mathbb{E}[\text{ALG}_\tau] \geq \frac{1}{2} \mathbb{E}[\max_i Z_i]$, as claimed. \square

Multiple Items. The key idea in proving our $1 + o(1)$ approximation ratio result for selecting multiple items for bounded s_{col} is to extend the augmentation lemma to cardinality constraints.

Definition 3.3 (Multiple-Items Augmented Prophets Problem). In the *augmented prophets problem with cardinality constraint r* , the algorithm is given the distributions of a set of independent nonnegative random variables Z_1, \dots, Z_n . Then, it observes one at a time the realizations of $X_i = Z_i + W_i$ for $i \in [1, n]$, where each W_i is nonnegative and satisfies that each Z_i is independent of X_1, \dots, X_{i-1} . The algorithm chooses at each step to reject or accept X_i subject to taking at most r variables total. It must compete with $\mathbb{E} \left[\sum_{i=1}^r Z^{(i)} \right]$, where $Z^{(i)}$ is the i th-largest $\{Z_1, \dots, Z_n\}$.

Since none of the prior $1 + o(1)$ approximation ratio algorithms for multiple-items is robust to augmentations, in Section 5.3 we design a new algorithm to prove the following multiple-items augmentation lemma.

Lemma 3.4. (Multiple-Items Augmentation Lemma). *There is an algorithm for the augmented prophets problem with cardinality constraint r achieving a $\left(1 + O\left(\frac{(\log r)^{3/2}}{r^{1/4}}\right)\right)$ approximation ratio.*

Next, we utilize the single-item augmentation lemma, along with careful decompositions of $\{X_i\}$, to separately attack the single-item problem for the cases of bounded s_{row} and s_{col} .

4 Selecting a Single Item

In this section we prove our main theorem. Later, we will also address cases where the algorithm can take multiple items.

Theorem 4.1. *There exists an inclusion-threshold algorithm for the linearly correlated prophet problem with approximation ratio $O(\min\{s_{\text{col}}, s_{\text{row}}\})$.*

We will first show in Proposition 4.2 that an inclusion-threshold algorithm guarantees $O(s_{\text{col}})$; then in Proposition 4.4 that an inclusion-threshold algorithm achieves $O(s_{\text{row}})$. The algorithm that runs one of these according to which of $s_{\text{col}}, s_{\text{row}}$ is smaller is an inclusion-threshold algorithm achieving the claimed performance.

We will see that bounded column sparsity is the easier case, requiring a simpler algorithm and analysis. For the case of bounded row sparsity, we will need much more careful reasoning about dependencies and correlations between Y_j . This difficulty will also manifest quantitatively when we move to the cardinality-constraint setting in Section 5, where better bounds will be achievable only in the case of bounded column sparsity.

4.1 Bounded Column Sparsity

Recall that column sparsity s_{col} is the maximum number of times a given feature Y_j may appear with nonzero coefficient. We now give a relatively straightforward algorithm for achieving $\Omega\left(\frac{1}{s_{\text{col}}}\right)$ fraction of $\mathbb{E}[\max_i X_i]$. The idea is similar to the “first attempt” described in Section 3, using the single-item augmentation lemma (Lemma 3.2) to overcome the challenges discussed there.

Proposition 4.2. *There exists an inclusion-threshold algorithm for the linearly correlated prophet problem with approximation ratio $2e \cdot s_{\text{col}}$.*

Proof. Choose $S \subseteq [n]$ by including each $i \in [n]$ independently with probability $\frac{1}{s_{\text{col}}}$. This gives the inclusion subset; now we define the threshold τ . Assign each Y_j to the first surviving X_i that includes it, i.e., for each $i \in S$, construct a set $T_i := \{j : A_{ij} > 0 \text{ and } A_{i'j} = 0 \ (\forall i' \in S \text{ where } i' < i)\}$. Let $Z_i = \sum_{j \in T_i} A_{ij} Y_j$ and set $\tau = \frac{1}{2} \mathbb{E}[\max\{Z_1, \dots, Z_n\}]$. If $i \notin S$, then $T_i = \emptyset$ and $Z_i = 0$.

By definition of an inclusion-threshold algorithm (Definition 2.5), we automatically reject any X_i such that $i \notin S$, and we select the first arriving X_i such that $i \in S$ and $X_i \geq \tau$.

Now, by construction, we can write $X_i = Z_i + W_i$ where each Z_i contains only variables Y_j not appearing in any prior $X_{i'}$ for $i' \in S$ and $i' < i$. So Z_i is independent of the preceding $X_{i'}$ under consideration. Hence by the single-item augmentation lemma (Lemma 3.2), $\mathbb{E}[\text{ALG}] \geq \frac{1}{2} \mathbb{E}[\max_i Z_i]$.

Next, we show that $\mathbb{E}[\max_i Z_i]$ is comparable to $\mathbb{E}[\max_i X_i]$.

Claim 4.3. $\mathbb{E}[\max\{Z_1, \dots, Z_n\}] \geq \frac{1}{e \cdot s_{\text{col}}} \mathbb{E}[\max\{X_1, \dots, X_n\}]$, where the expectation is over the construction of S as well as Y_1, \dots, Y_n .

Proof of Claim 4.3. We prove that for every fixed realization of Y_1, \dots, Y_n , the inequality holds in expectation over S . Let $X_{i^*} = \max_i X_i$. For each Y_j with $A_{i^*j} > 0$, we claim $\Pr[j \in T_i] \geq \frac{1}{e \cdot s_{\text{col}}}$ because X_i survives with probability $\frac{1}{s_{\text{col}}}$ and independently, the other at most $s_{\text{col}} - 1$ variables $X_{i'}$ with $A_{i'j} > 0$ all fail to survive with probability at least³ $\left(1 - \frac{1}{s_{\text{col}}}\right)^{s_{\text{col}}-1} \geq \frac{1}{e}$. (If $s_{\text{col}} = 1$, then this probability is 1.) So with probability only over the construction of S ,

$$\mathbb{E}_S[Z_{i^*}] = \sum_j \Pr[j \in T_{i^*}] A_{i^*j} Y_j \geq \frac{1}{e \cdot s_{\text{col}}} \sum_j A_{i^*j} Y_j = \frac{1}{e \cdot s_{\text{col}}} X_{i^*}.$$

So we have $\mathbb{E}_S[Z_{i^*}] \geq \frac{1}{e \cdot s_{\text{col}}} \max_i X_i$, so $\mathbb{E}_S[\max_i Z_i] \geq \frac{1}{e \cdot s_{\text{col}}} \max_i X_i$. This holds for each fixed realization of Y_1, \dots, Y_n , so it holds in expectation. \square

Claim 4.3 completes the proof of Proposition 4.2, as we have

$$\mathbb{E}[\text{ALG}] \geq \frac{1}{2} \mathbb{E}[\max_i Z_i] \geq \frac{1}{2e \cdot s_{\text{col}}} \mathbb{E}[\max_i X_i]. \quad \square$$

4.2 Bounded Row Sparsity

Recall that row sparsity s_{row} implies that each X_i only depends on at most s_{row} different features Y_j ; however, a given Y_j may appear in arbitrarily many X_i s. In this section, for notational convenience,

³We often use the inequality $(1 - \frac{1}{N})^{N-1} = (\frac{N-1}{N})^{N-1} \geq \frac{1}{e}$, which follows from $\left(\frac{N}{N-1}\right)^{N-1} = \left(1 + \frac{1}{N-1}\right)^{N-1} \leq e$.

we assume without loss of generality that $\max_i \{A_{ij}\} = 1$ for all j . (If this is not the case, one can renormalize each column and redefine a scaled version of Y_j .) We prove the following:

Proposition 4.4. *There exists an inclusion-threshold algorithm for the linearly correlated prophets problem achieving approximation ratio $2e^3 \cdot s_{\text{row}}$.*

This case requires more care. There does not seem to be an analogous approach to randomly excluding X_i , as for bounded column sparsity. Moreover, an important observation is that the X_i cannot be treated identically in a manner oblivious to the structure of A . For every “important” row that ought to be included, there can be many unimportant rows. Indeed, we can take any instance and prepend it with arbitrarily many variables of the form $X_i = Y_1$ without changing the row sparsity s_{row} . An oblivious inclusion-threshold algorithm would essentially keep only variables from this prefix, ignoring the actual instance.

Before the formal proof of Proposition 4.4, we develop a tool to address this challenge. The key idea is to design an inclusion scheme that, for any instance structure, allows each Y_j to be both represented and “independent” with a reasonably high probability. Here independence refers to not sharing an X_i with any other included $Y_{j'}$. Inspired by contention-resolution schemes, which have a similar flavor, we define a *representative* construction of a subset of the Y_j and corresponding X_i with $A_{ij} = 1$, where X_i is matched to $Y_{j(i)}$.

Definition 4.5. Consider a randomized selection of $S \subseteq \{1, \dots, n\}$ and $T \subseteq \{1, \dots, m\}$ of equal size with a perfect matching $j(i)$ satisfying $A_{ij(i)} = 1$. Call this construction α -*representative* or α -*rep.* if

- (i) for all $j \in \{1, \dots, m\}$, we have $\Pr[j \in T] \geq \alpha$, and
- (ii) for all $i, i' \in S$, $i \neq i'$, we have $A_{i'j(i)} = 0$.

Note that we cannot hope for better than an $O(\frac{1}{s_{\text{row}}})$ -rep. construction, as any inclusion of some Y_j can rule out $s_{\text{row}} - 1$ other features. This raises the question of whether one can achieve $\Omega(\frac{1}{s_{\text{row}}})$ -rep.

Lemma 4.6. *For any linearly correlated instance there exists a $\frac{1}{e^2 \cdot s_{\text{row}}}$ rep. construction.*

Proof. For each Y_j , define its *primary* $i(j)$ by picking any i such that $A_{ij} = 1$ (by our renormalization assumption, there is at least one). Consider a directed graph G where the nodes are $\{1, \dots, m\}$ representing the independent variables Y_j . There is a *directed* edge (j, j') if $A_{i(j)j'} > 0$, i.e., j points to all other j' who are included in its primary variable $X_{i(j)}$. We note that both edges (j, j') and (j', j) might be present.

The key property is that each vertex in G has out degree $\leq s_{\text{row}} - 1$, as Y_j has only one primary $X_{i(j)}$ and at most $s_{\text{row}} - 1$ other variables j' have $A_{i(j)j'} > 0$. Because average in-degree equals average out-degree, this implies there exists a vertex with in-degree at most $s_{\text{row}} - 1$. Applying this argument recursively, we get the following claim.

Claim 4.7. *There exists an order π of the vertices of G such that for every j , the induced subgraph on $\pi(1), \dots, \pi(j)$ satisfies that the in-degree of $\pi(j)$ is at most $s_{\text{row}} - 1$.*

Proof. As shown, there exists some j with in-degree at most $s_{\text{row}} - 1$. Set $\pi(m) = j$. Now delete j from the graph, including all edges to and from j . In this graph again all out-degrees are at most $s_{\text{row}} - 1$, so we can recursively construct $\pi(m - 1), \dots, \pi(1)$. \square

Consider all the Y_j variables in the order π given by Claim 4.7. Initialize $S, T = \emptyset$. On considering j , if j does not have an edge with any vertex $j' \in T$ (neither incoming nor outgoing), then independently with probability $\frac{1}{s_{\text{row}}}$, add j to T and add $j(i)$ (its primary variable) to S . With the remaining probability, ignore j and continue. We show that this randomized construction of S, T satisfies the two properties of a $\frac{1}{e^2 \cdot s_{\text{row}}}$ rep. construction.

For Property (i), note each j has at most $2(s_{\text{row}} - 1)$ total edges (both incoming and outgoing) to nodes j' appearing prior to j in the permutation: j has at most $s_{\text{row}} - 1$ outgoing edges in total, and by construction of π , has at most $s_{\text{row}} - 1$ incoming edges from nodes prior to j in π . So when we reach j in the permutation, we consider it with probability at least the probability that all these $2(s_{\text{row}} - 1)$ neighbors have been rejected, which is at least $\left(1 - \frac{1}{s_{\text{row}}}\right)^{2(s_{\text{row}} - 1)} \geq \frac{1}{e^2}$; and then we include it with probability $\frac{1}{s_{\text{row}}}$. This shows that each j is included with probability at least $\frac{1}{e^2 \cdot s_{\text{row}}}$.

Next, for Property (ii), consider any $i, i' \in S$ with respective partners $j, j' \in T$. We must show $A_{ij'} = 0$. Note that by construction, $i = i(j)$ and $i' = i(j')$, i.e., they are the primary variables for j, j' . Either j was selected into T before or after j' . In either case, the second variable could only be selected if edge (j, j') did not exist in the graph, which implies $A_{ij'} = 0$.

This completes the proof of Lemma 4.6. \square

Given our representative construction, we are ready to complete the algorithm and proof.

Proof of Proposition 4.4. The algorithm is an inclusion-threshold algorithm. Its inclusion set S is obtained by calling our representative construction of Lemma 4.6, which also produces a choice $j(i)$ for each $i \in S$. Define $Z_i = Y_{j(i)}$ for each $i \in S$ and $Z_i = 0$ if $i \notin S$. Set $\tau = \frac{1}{2}\mathbb{E}[\max_i Z_i]$.

By the second property of representative constructions, Z_i is independent of $X_{i'}$ for all $i' \in S, i' \neq i$. Therefore, by the Augmentation Lemma (Lemma 3.2),

$$\mathbb{E}[\text{ALG}] \geq \frac{1}{2}\mathbb{E}[\max_i Z_i]. \quad (2)$$

Combining this with the following Lemma 4.8 will prove Proposition 4.4.

Lemma 4.8. $\mathbb{E}[\max_i Z_i] \geq \frac{1}{e^3 \cdot s_{\text{row}}}\mathbb{E}[\max_i X_i]$.

Before proving Lemma 4.8, we will need one more idea. Let $R_i = \{j : A_{ij} > 0\}$, the variables included in X_i . Notice that we may have $|R_i \cap T| \geq 2$, i.e., multiple variables Y_j are members of X_i and appear in the construction T . This can occur when X_i is not primary for any of them.⁴ It will help to lower-bound the probability that Y_j is the unique member of $R_i \cap T$.

Claim 4.9. *Under the construction of Lemma 4.6, for each $j \in R_i$, $\Pr[R_i \cap T = \{j\}] \geq \frac{1}{e^3 \cdot s_{\text{row}}}$.*

Proof. We have $\Pr[j \in T] \geq \frac{1}{e^2 \cdot s_{\text{row}}}$ by the representative construction. Meanwhile, conditioned on any other decisions, each j' is included in T with probability at most $\frac{1}{s_{\text{row}}}$, because it is considered with probability at most 1 and included independently with probability $\frac{1}{s_{\text{row}}}$ conditioned on being considered. So $\Pr[R_i \cap T = \{j\} | j \in T] \geq \left(1 - \frac{1}{s_{\text{row}}}\right)^{s_{\text{row}} - 1} \geq \frac{1}{e}$. \square

Proof of Lemma 4.8. Fix the realizations of all Y_j . Let $i^* = \arg\max_i X_i$. First, notice that

$$\mathbb{E}[\max_i Z_i] = \mathbb{E}[\max_{j \in T} \{Y_j\}] \geq \mathbb{E}[\max_{j \in T \cap R_{i^*}} \{A_{i^*j} \cdot Y_j\}].$$

⁴An illuminating instance is: $X_i = Y_i$ for all $i \leq s_{\text{row}}$ and $X_{s_{\text{row}}+1} = 0.99(Y_1 + \dots + Y_{s_{\text{row}}})$.

Now the expected maximum of elements in $T \cap R_{i^*}$ is at least the sum over the elements of each's contribution to the max, which is at least the chance it is the unique survivor times its value. This allows us to relate a max to a sum, and it relies on the fact that the representative construction's randomness is independent of the realizations of the $\{Y_j\}$. Thus,

$$\mathbb{E}[\max_i Z_i] \geq \mathbb{E}\left[\max_{j \in T \cap R_{i^*}} \{A_{i^*j} \cdot Y_j\}\right] \geq \sum_j \Pr[T \cap R_{i^*} = \{j\}] \cdot A_{i^*j} Y_j \geq \sum_j \frac{1}{e^3 \cdot s_{\text{row}}} A_{i^*j} Y_j,$$

where the last inequality uses Claim 4.9. Since $\sum_j A_{i^*j} Y_j = X_{i^*}$, we have $\mathbb{E}[\max_i Z_i] \geq \frac{1}{e^3 \cdot s_{\text{row}}} X_{i^*}$. Taking expectations over Y_1, \dots, Y_n completes the proof. \square

Finally, combining (2) with Lemma 4.8 completes the proof of Proposition 4.4. \square

4.3 Lower Bounds

We now give a matching hardness result, showing that no algorithm can do better than our results in the previous section up to constant factors.

Example 4.10 (General tower instance). Take the tower \mathbf{Y} variables (recall $Y_i = \frac{1}{\epsilon^i}$ with probability ϵ^i , and 0 otherwise). Given input integer c , set $n = s_{\text{row}} = s_{\text{col}} = c$. Let A be an $n \times n$ matrix with entry $A_{i,j} = 0$ for $j < i$ and $A_{i,j} = \epsilon^{j-i}$ for $j \geq i$. Visually,

$$X_1 = Y_1 + \epsilon Y_2 + \dots + \epsilon^n Y_n, \quad X_2 = Y_2 + \epsilon Y_3 + \dots + \epsilon^{n-1} Y_n, \quad \dots, \quad X_n = Y_n.$$

The difficulty here amplifies that of Example 2.3. If any of Y_i, Y_{i+1}, \dots, Y_n are active, then this will cause X_i to be nonzero. Assuming only one of these variables is active (by far the most likely case), it is impossible for the algorithm to tell whether to stop or continue.

It will turn out that this instance is hard even if the algorithm is given additional power.

Definition 4.11. In the *fractional variant of the prophet problem*, at each arrival i , the algorithm may choose to take a fraction p_i of the current arrival X_i subject to always taking at most one unit in total, i.e., $\sum_{i=1}^n p_i \leq 1$ with probability 1. Its reward is $\sum_{i=1}^n p_i X_i$.

One strategy available in the fractional prophet problem is to spend the entire budget on a single arrival, which is an algorithm for the standard prophets problem. So a lower bound for the fractional problem immediately implies a lower bound for the prophets problem.

Theorem 4.12. *In the linearly correlated prophet problem, even if fractional, no online algorithm can guarantee a smaller approximation ratio than $\min\{s_{\text{col}}, s_{\text{row}}\}$.*

Proof. We consider a family of instances of Example 4.10 where $n = s_{\text{col}} = s_{\text{row}}$. Since for every i we have $X_i \geq Y_i$, we get that for sufficiently small ϵ ,

$$\mathbb{E}[\max_i \{X_i\}] \geq \mathbb{E}[\max_j \{Y_j\}] = \sum_{i=1}^n \epsilon^i \prod_{j>i} (1 - \epsilon^j) \cdot \frac{1}{\epsilon^i} \geq n(1 - \epsilon)^n \geq n(1 - n\epsilon).$$

On the other hand, we show every online algorithm that may even fractionally select elements has value at most $1/(1 - \epsilon)^2$. This implies the approximation ratio can be $n(1 - n\epsilon)(1 - \epsilon)^2$, which tends to n as $\epsilon \rightarrow 0$.

Lemma 4.13. Suppose at arrival i we have $p = \sum_{i' < i} p_{i'}$ and $X_i = 1/\epsilon^i$. Conditioned on this event, any online algorithm obtains expected value from elements i, \dots, n at most $(1-p)/\epsilon^i$.

Before proving Lemma 4.13, we use it to prove that every algorithm has $O(1)$ expected value. Notice that if $X_i > 1/\epsilon^i$ then the online algorithm should never accept any fraction of the element X_i as X_{i+1} is guaranteed to be larger. Hence by Lemma 4.13, the optimal algorithm ALG takes the smallest i for which $X_i = 1/\epsilon^i$, which means

$$\begin{aligned} \mathbb{E}[\text{ALG}] &\leq \sum_{i \geq 1} \Pr[(X_j > 1/\epsilon^j \text{ for all } j < i) \ \& \ (X_i = 1/\epsilon^i)] \cdot \frac{1}{\epsilon^i} \\ &= \sum_{i \geq 1} \Pr[Y_{i-1} = 1/\epsilon^{i-1}] \cdot (\Pr[X_i = 1/\epsilon^i]) \cdot \frac{1}{\epsilon^i} \\ &\leq \sum_{i \geq 1} \Pr[Y_{i-1} = 1/\epsilon^{i-1}] \cdot \left(\sum_{j \geq i} \Pr[Y_j = 1/\epsilon^j] \right) \cdot \frac{1}{\epsilon^i} \\ &= \sum_{i \geq 1} \epsilon^{i-1} \sum_{j \geq i} \epsilon^j \cdot \frac{1}{\epsilon^i} \leq \sum_{i \geq 1} \frac{\epsilon^{i-1}}{1-\epsilon} \leq \frac{1}{(1-\epsilon)^2}. \quad \square \end{aligned}$$

Now we prove the missing lemma.

Proof of Lemma 4.13. We prove this lemma by reverse induction on i . It is immediately true for $i = n$, as spending the entire remaining budget $1-p$ on acquiring $X_i = \frac{1}{\epsilon^i}$ is optimal. To prove the inductive step, notice the optimal online algorithm can be written as a convex combination of the following two algorithms: one that spends the entire remaining budget of $(1-p)$ on X_i and another one that spends no budget on X_i and plays optimally afterwards. We argue that the first algorithm is better, which means its expected value is $(1-p)/\epsilon^i$.

Observe that the second algorithm obtains nonnegative reward only if one of Y_j s for $j > i$ is active. In this case $X_{i+1} = 1/\epsilon^{i+1}$ (note it cannot be larger because $X_i = 1/\epsilon^i$), and hence by induction hypothesis the optimal online algorithm gets value $X_{i+1} = (1-p)/\epsilon^{i+1}$. Thus, the expected value of the algorithm is

$$\Pr[\exists j > i \text{ s.t. } Y_j = 1/\epsilon^j \mid X_i = 1/\epsilon^i] \cdot (1-p)/\epsilon^{i+1}.$$

We show $\Pr[\exists j > i \text{ s.t. } Y_j = 1/\epsilon^j \mid X_i = 1/\epsilon^i] \leq \epsilon$, which implies the first algorithm is always better. To see this, notice

$$\begin{aligned} \Pr[\exists j > i \text{ s.t. } Y_j = 1/\epsilon^j \mid X_i = 1/\epsilon^i] &= \frac{\Pr[(\exists j > i \text{ s.t. } Y_j = 1/\epsilon^j) \ \& \ (X_i = 1/\epsilon^i)]}{\Pr[X_i = 1/\epsilon^i]} \\ &= \frac{\sum_{j > i} \Pr[(Y_j = 1/\epsilon^j) \ \& \ (X_i = 1/\epsilon^i)]}{\sum_{j \geq i} \Pr[(Y_j = 1/\epsilon^j) \ \& \ (X_i = 1/\epsilon^i)]} \\ &= \frac{\sum_{j > i} \epsilon^j / (1 - \epsilon^j) \cdot \prod_{j' \geq i} (1 - \epsilon^{j'})}{\sum_{j \geq i} \epsilon^j / (1 - \epsilon^j) \cdot \prod_{j' \geq i} (1 - \epsilon^{j'})} = \frac{\sum_{j > i} \epsilon^j / (1 - \epsilon^j)}{\sum_{j \geq i} \epsilon^j / (1 - \epsilon^j)}. \end{aligned}$$

Now using $1+x \leq \frac{1}{1-x} \leq 1+2x$ for $0 \leq x < 0.5$, we get

$$\begin{aligned} \Pr[\exists j > i \text{ s.t. } Y_j = 1/\epsilon^j \mid X_i = 1/\epsilon^i] &\leq \frac{\sum_{j \geq i+1} \epsilon^j \cdot (1+2\epsilon^j)}{\sum_{j \geq i} \epsilon^j \cdot (1+\epsilon^j)} \\ &= \frac{\epsilon^{i+1} - \epsilon^{n+1} + (2\epsilon^{2i+2} - 2\epsilon^{2n+2})/(1+\epsilon)}{\epsilon^i - \epsilon^{n+1} + (\epsilon^{2i} - \epsilon^{2n+2})/(1+\epsilon)} \leq \epsilon, \end{aligned}$$

where the last inequality uses $\epsilon < 1/2$ and $\epsilon^n < 2$. \square

5 Selecting Multiple Items

In this section, we show that our approach via augmentations extends to a variant of the problem in which one may take up to $r \in \mathbb{N}$ of the arriving variables X_1, \dots, X_n . Let $Q \subseteq \{1, \dots, n\}$ be the indices chosen by ALG with $|Q| \leq r$. Let $X^{(i)}$ denote the i th-largest realized variable (we later use notation $Z^{(i)}$ for the i th-largest among $\{Z_1, \dots, Z_n\}$ as well). We have

$$\text{ALG} = \sum_{i \in Q} X_i \quad \text{while} \quad \text{OPT} = \sum_{i=1}^r X^{(i)}.$$

We refer to this problem as a cardinality constraint of r . It is also referred to as selecting an independent set of a rank- r matroid in the special case of r -uniform matroids.

In this setting, there will be significant differences between row and column sparsity assumptions. We will show that for bounded column sparsity s_{col} , one can design $(1 + o(1))$ -approximation algorithms for cardinalities $r \rightarrow \infty$, while this does *not* hold for bounded row sparsity s_{row} .

5.1 Bounded Column Sparsity

As $r \rightarrow \infty$, we will show for bounded column sparsity an approximation ratio approaching 1.

Theorem 5.1. *For a fixed s_{col} , the linearly correlated prophets problem with cardinality constraint r admits a $\left(1 + O\left(\left(\frac{s_{\text{col}}}{r}\right)^{1/5} (\log r)^{6/5}\right)\right)$ -approximation.*

The key idea is to prove an augmentation lemma for selecting multiple items (restated below).

Lemma 3.4. (Multiple-Items Augmentation Lemma). *There is an algorithm for the augmented prophets problem with cardinality constraint r achieving a $\left(1 + O\left(\frac{(\log r)^{3/2}}{r^{1/4}}\right)\right)$ approximation ratio.*

In Section 5.3 we prove this augmentation lemma, but before we use it to prove Theorem 5.1.

The idea of the reduction is that by randomly partitioning the variables into s_{col}/ϵ' “groups” gives us multiple independent Augmented Prophets problems. We think of each group as a subproblem of selecting $\epsilon' r/s_{\text{col}}$ elements and use the Augmentation Lemma to approximately solve it.

Proof of Theorem 5.1. Formally, let there be $c = s_{\text{col}}/\epsilon'$ sets, which we call *groups*, B_1, \dots, B_c . For each X_i , place it in a group $j \in \{1, \dots, c\}$ chosen uniformly at random. For each X_i , let

$$X'_i = \sum A_{ij} Y_j \cdot \mathbf{1}[Y_j \text{ only appears once in the group containing } X_i]$$

denote the sum of X_i ’s components Y_j that do not appear with any other variable in the group containing X_i . Let OPT'_j denote the sum of the largest $\epsilon' r/s_{\text{col}}$ elements X'_i in group B_j .

Claim 5.2. $\mathbb{E}[\sum_j \text{OPT}'_j] \geq (1 - \epsilon') \cdot \mathbb{E}[\text{OPT}]$.

Proof. Consider a fixed X_i . Condition on X_i landing in a group. Notice that each of its Y_j s have at least $1 - \epsilon'$ chance of appearing only with X_i in this group, and hence it contributes to OPT'_j . \square

Since each group B_j forms a separate instance of the Augmented Prophets problems, we can apply the Augmentation Lemma 3.4 on each of them. Let ALG_j denote the algorithm's performance in group j . By selecting ϵ' less than $O(\frac{(\log r)^{3/2}}{(\epsilon' r / s_{\text{col}})^{1/4}})$, i.e., choosing $\epsilon' = \Theta((\frac{s_{\text{col}}}{r})^{1/5} (\log r)^{6/5})$, we get

$$\sum_j \mathbb{E}[\text{ALG}_j] \geq \sum_j (1 - O(\epsilon')) \cdot \mathbb{E}[\text{OPT}'_j] \geq (1 - O(\epsilon')) \cdot \mathbb{E}[\text{OPT}],$$

where the last inequality uses Claim 5.2. \square

5.2 Bounded Row Sparsity

For cardinality constraints, the symmetry between bounds for row and column sparsity breaks: One cannot guarantee better than a $\Theta(s_{\text{row}})$ approximation in general even as $r \rightarrow \infty$. In fact, this follows by reducing to our previous hardness result for fractional prophets.

Theorem 5.3. *No algorithm for linearly correlated prophets with cardinality constraint can guarantee better than an $\Omega(s_{\text{row}})$ -approximation, even as $r \rightarrow \infty$.*

Finally, we show that the $O(s_{\text{row}})$ -approximation upper bound for single item can be extended to the cardinality constraint setting. The intuition is straightforward, as we can simply instantiate r parallel versions of our previous single item algorithm and assign arriving variables to each at random. The analysis needs to show that no more than a constant factor is lost due to cases where members of OPT are sent to the same bucket.

Theorem 5.4. *For all r , there is an $O(s_{\text{row}})$ -approximation for the linearly correlated prophet problem with cardinality constraint r .*

We now present formal proofs of the last two theorems.

Proof of Theorem 5.3. For any r, s_{row} , we construct the general tower instance of Example 4.10 with parameter s_{row} , but we simply make r copies of each variable X_i (not independent, but exact copies). The row sparsity is unchanged. Now we will make the problem easier in two steps.

(1) Allow variables to arrive in batches of size r . The algorithm may select any subset of the variables (until fulfilling its cardinality constraint), then reject the rest and receive the next batch. This is a strictly easier problem, so an algorithm's performance can only improve. Of course this will not help on this instance, since each batch of r are all identical, and it will turn out to be optimal to either take them all or none.

(2) Now instead we send the original tower instance (with no duplication), and we give the algorithm a cardinality constraint of 1, but we allow it to pick a fractional amount of each variable X_i . In other words, we return exactly to the setting of Theorem 4.12. In each case, the algorithm sees the same information before making each decision, i.e., the value of the current X_i . In the fractional problem, the algorithm can pick any fraction p_i of X_i , so long as the total amount picked is at most one. In problem (1) above, the algorithm can pick any fraction $\frac{c}{r}$ of the variables that equal the original X_i , as long as it has does not exceed a total of $\frac{r}{r}$. The fractional problem allows more choice and the benchmarks (normalized) are the same, since the maximum of the duplicated instance will take all r copies of the largest X_i . So the fractional problem is only easier.

We now invoke Theorem 4.12, which gives an lower bound of $\Omega(s_{\text{row}})$ -approximation on the fractional problem. (Note that s_{row} did not change during the above reduction, although s_{col} did.) \square

Proof of Theorem 5.4. Given an instance, we create r buckets B_1, \dots, B_r and place each variable X_i in a bucket B_j uniformly at random. We then run our algorithm for the linearly correlated prophet problem with bounded s_{row} in each bucket j (call it ALG_j), selecting one item. Given a fixed assignment of variables to buckets, in each bucket B_j we have by the algorithm's guarantee that, with expectation over realizations of $\{X_i : i \in B_j\}$, by Proposition 4.4

$$\mathbb{E}_{\mathbf{X}}[\text{ALG}_j] \geq \Omega\left(\frac{1}{s_{\text{row}}}\right) \cdot \mathbb{E}_{\mathbf{X}}\left[\max_{X_i \in B_j} X_i\right].$$

Taking expectations over both buckets and variables, we have

$$\mathbb{E}[\text{ALG}] = \sum_{j=1}^r \mathbb{E}_B \mathbb{E}_{\mathbf{X}}[\text{ALG}_j] \geq \Omega\left(\frac{1}{s_{\text{row}}}\right) \sum_{j=1}^r \mathbb{E}_B \mathbb{E}_{\mathbf{X}}\left[\max_{X_i \in B_j} X_i\right] = r \cdot \Omega\left(\frac{1}{s_{\text{row}}}\right) \mathbb{E}_{\mathbf{X}} \mathbb{E}_B\left[\max_{X_i \in B_1} X_i\right],$$

where the last equality is by symmetry of the buckets. Now since $\text{OPT} = X^{(1)} + \dots + X^{(r)}$ where $X^{(i)}$ is the i th largest variable, and since for fixed variable realizations $\Pr[X^{(i)} = \max_{X_{i'} \in B_1} X_{i'}] = \frac{1}{r} \prod_{i'=1}^{i-1} (1 - \frac{1}{r}) \geq \frac{1}{e \cdot r}$, we get

$$\begin{aligned} \mathbb{E}[\text{ALG}] &\geq r \cdot \Omega\left(\frac{1}{s_{\text{row}}}\right) \mathbb{E}_{\mathbf{X}}\left[\sum_{i=1}^r \Pr[X^{(i)} = \max_{X_{i'} \in B_1} X_{i'}] \cdot X^{(i)}\right] \\ &\geq r \cdot \Omega\left(\frac{1}{s_{\text{row}}}\right) \mathbb{E}_{\mathbf{X}}\left[\sum_{i=1}^r \frac{1}{e \cdot r} X^{(i)}\right] = \Omega\left(\frac{1}{s_{\text{row}}}\right). \end{aligned} \quad \square$$

5.3 Multiple-Items Augmentation Lemma

In Section 3 we showed a 2-approximation single-item augmentation lemma using the half of expected-maximum as a threshold. In this section, we prove a $1+o(1)$ approximation multiple-items augmentation lemma (Lemma 3.4), assuming the cardinality constraint r is sufficiently large. We first give a surrogate benchmark OPT' that competes with OPT . We then define the algorithm and show that it competes with OPT' .

Surrogate benchmark.

The analysis hinges on threshold $\tau_0 := \frac{\mathbb{E}[\text{OPT}]}{\epsilon}$, where we call variables Z_i and X_i that fall above the threshold *heavy* and below *light*. For light variables, we exclude the scenario where they are very small, below some threshold $\tau_c \leq \epsilon \frac{\mathbb{E}[\text{OPT}]}{r}$. We will generally use the prime symbol $'$ to denote a version of a variable that is zeroed out if it's too heavy (or too light). Let

$$Z'_i = Z_i \cdot \mathbf{1}[\tau_c \leq Z_i < \tau_0] \quad \text{and} \quad Z^{(i)'} = Z^{(i)} \cdot \mathbf{1}[\tau_c \leq Z_i < \tau_0].$$

Let

$$\begin{aligned} \text{OPT}' &= \text{OPT}'_1 + \text{OPT}'_2 \quad \text{where} \\ \text{OPT}'_1 &= \max_i Z_i \cdot \mathbf{1}[Z_i \geq \tau_0] \quad \text{and} \quad \text{OPT}'_2 = \sum_{j=1}^r Z^{(j)'} \end{aligned}$$

Note that OPT' only considers one heavy variable and throws all other heavy variables away. Nevertheless, in Appendix D we will show the following claim that it competes with OPT .

Claim 5.5. $\mathbb{E}[\text{OPT}'] \geq (1 - 2\epsilon)\mathbb{E}[\text{OPT}]$.

Algorithm overview.

For intuition, the problem with setting any particular fixed threshold is that variables with insignificant values of Z_i can be boosted by the adversary to some X_i just above the threshold. The algorithm would use up its r slots and be unable to take the later, larger arrivals that contribute to OPT.

Therefore, we will define a sequence of thresholds, each a factor of $1 - \epsilon$ apart. The algorithm will have a certain number of slots \tilde{r}_j for the “bucket” j of arrivals between any two thresholds. For this fixed bucket, such a boosting strategy by the adversary can only cost the algorithm a factor of $1 - \epsilon$.

Roughly, this strategy will cover the case where OPT is concentrated. To allow for cases where most of OPT comes from very rare, very large variables, we will also reserve a slot for such variables and analyze it separately.

Algorithm definition.

We define a sequence of thresholds. Recall that $\text{OPT} = \sum_{j=1}^r Z^{(j)}$ where $Z^{(j)}$ is the j th-largest of Z_1, \dots, Z_n . Let $c = \lceil \frac{1}{\epsilon} \ln \frac{r}{\epsilon^2} \rceil$ and define thresholds

$$\tau_j = (1 - \epsilon)^j \frac{\mathbb{E}[\text{OPT}]}{\epsilon} \quad (j = 0, \dots, c)$$

Observation 5.6. The largest threshold is $\tau_0 = \frac{\mathbb{E}[\text{OPT}]}{\epsilon}$ and the smallest threshold $\tau_c \leq \frac{\epsilon \mathbb{E}[\text{OPT}]}{r}$.

Proof. τ_0 is immediate, and we have $(1 - \epsilon)^c \leq e^{-\epsilon c} \leq \exp\left(-\epsilon \frac{1}{\epsilon} \ln \frac{r}{\epsilon^2}\right) = \frac{\epsilon^2}{r}$. \square

Now we define the size of each bucket. Recall that we abuse notation by writing $Z_i \in \text{OPT}$ if Z_i is one of the r variables included in the OPT solution. Let

$$r_j = \mathbb{E}|\{i : Z_i \in \text{OPT}, \tau_j \leq Z_i \leq \tau_{j-1}\}| \quad (j = 1, \dots, c)$$

$$\beta = 3\sqrt{r \ln(c/\epsilon)}$$

$$\tilde{r}_j = r_j + \beta \quad (j = 1, \dots, c)$$

$$\tilde{r}_0 = 1.$$

We first define an algorithm ALG' that does not quite achieve the cardinality constraint r . We will then modify it to obtain ALG with only a small loss in performance. ALG' initializes $b_j = 0$ for $j = 0, \dots, c$ and proceeds as follows when a variable X_i arrives.

1. If $X_i < \tau_c$, we discard X_i and continue.
2. Otherwise, let $j = \min\{j' : X_i \geq \tau_{j'}\}$.
3. If $b_j < \tilde{r}_j$, we take X_i and increment b_j . Otherwise (bucket j is full), increment j and repeat this step. If $j > c$, stop and discard X_i .

In other words, we attempt to assign X_i to its original bucket, but if that is full, we allow it to fall into buckets reserved for smaller variables (higher indices j).

Now, the final algorithm ALG is defined as follows: run ALG', but each time ALG' takes an arrival X_i , discard it independently with probability ϵ . If a variable is not discarded, but the cardinality constraint r is reached, then discard it anyways.

Analysis.

We show in Appendix D that ALG approximates ALG'.

Lemma 5.7. *For all $\epsilon \geq \frac{9(\ln r)^{3/2}}{r^{1/4}}$, we have $\mathbb{E}[ALG] \geq (1 - 2\epsilon)\mathbb{E}[ALG']$.*

Now, we analyze ALG'.

Let us define the contributions of ALG' and OPT' bucket-by-bucket. The top bucket of ALG' is split into cases where the corresponding Z_i is heavy or light. The following are random sets:

$$\begin{aligned} O_j &= \{i \in \text{OPT}' : \tau_j \leq Z_i < \tau_{j-1}\} & (j = 1, \dots, c) \\ O_0 &= \{i \in \text{OPT}' : \tau_0 \leq Z_i\} \\ B_j &= \{i \in \text{ALG}' : \tau_j \leq X_i < \tau_{j-1}\} & (j = 1, \dots, c) \\ B_0^{\text{light}} &= \{i \in \text{ALG}' : Z_i < \tau_0 \leq X_i\} \\ B_0^{\text{heavy}} &= \{i \in \text{ALG}' : \tau_0 \leq Z_i\}. \end{aligned}$$

We use the notation $i \in \text{OPT}'$ to denote that i contributes to OPT', i.e. either Z_i is the largest among $\{Z_j\}$ or Z'_i is among the r largest of $\{Z'_j\}$. Similarly, we write $i \in \text{ALG}'$ to mean that the algorithm takes X_i .

Now, we break down ALG' as follows.

$$\begin{aligned} \text{ALG}' &= \text{ALG}'_1 + \text{ALG}'_2 \quad \text{where} \\ \text{ALG}'_1 &= \sum_{i \in B_0^{\text{heavy}}} X_i \quad \text{and} \quad \text{ALG}'_2 = \sum_{i \in B_0^{\text{light}}} X_i + \sum_{j=1}^c \sum_{i \in B_j} X_i. \end{aligned}$$

In other words, ALG'_1 tracks the contribution of the special ‘‘heavy’’ bucket, but only in the case where the underlying variable Z_i is heavy. ALG'_2 tracks the remaining case and all other buckets. Notice these definitions are only for the purpose of analysis, as Z_i is not observable to the algorithm.

Finally, we define

$$P := \Pr[\max_i X_i \geq \tau_0].$$

A key point will be that if P is large, then the algorithm will often get some variable larger than τ_0 , which is good enough to compete with OPT.

Lemma 5.8. $\mathbb{E}[ALG'] \geq \frac{P}{\epsilon} \mathbb{E}[\text{OPT}]$.

Proof. With probability P , some X_i exceeds $\tau_0 = \frac{\mathbb{E}[\text{OPT}]}{\epsilon}$. Since Bucket 0 is reserved for such arrivals with budget $b_0 = 1$, the algorithm gets such X_i if this occurs, so its expectation is $\geq P\tau_0$. \square

Thus, if $P \geq \epsilon$, we are already done. The rest of the analysis will leverage cases where P is small.

Lemma 5.9. $\mathbb{E}[ALG'_1] \geq (1 - P)\mathbb{E}[\text{OPT}'_1]$.

Proof.

$$\begin{aligned}
\mathbb{E}[\text{ALG}'_1] &= \mathbb{E}\left[\sum_i X_i \cdot \mathbf{1}[i \in B_0^{\text{heavy}}]\right] \geq \mathbb{E}\left[\sum_i Z_i \cdot \mathbf{1}[i \in B_0^{\text{heavy}}]\right] \\
&= \sum_i \Pr[b_0 = 0 \text{ when } i \text{ arrives}] \cdot \mathbb{E}[Z_i \cdot \mathbf{1}[Z_i \geq \tau_0]] \quad (\text{using independence}) \\
&\geq (1 - P) \sum_i \mathbb{E}[Z_i \cdot \mathbf{1}[Z_i \geq \tau_0]] \\
&\geq (1 - P) \mathbb{E}[\max_i \{Z_i \cdot \mathbf{1}[Z_i \geq \tau_0]\}] = (1 - P) \cdot \mathbb{E}[\text{OPT}'_1]. \quad \square
\end{aligned}$$

The final piece of the argument is to show that the “buckets” strategy works, i.e., it cannot be disrupted by augmentations. The idea is that we have reserved an accurate number of slots in each bucket for the case where there is no augmentation. An augmented variable X_i can take away a bucket slot from some $Z_{i'}$, with $Z_{i'} \gg Z_i$, but then it will contribute about as much to ALG' as $Z_{i'}$ did to OPT'_2 . In this case, we should be concerned that OPT'_2 gets both Z_i and $Z_{i'}$ while ALG' only gets X_i , with $X_{i'}$ disappearing thanks to the bucket being full. However, the algorithm allows such an $X_{i'}$ to “trickle down” into a lower-tier bucket, in particular, the slot that is not being used by X_i . And if this slot is full as well, then in any case ALG' is competing with OPT'_2 .

Lemma 5.10. *For $\epsilon \geq \frac{9(\ln r)^{3/2}}{r^{1/4}}$, we have $\mathbb{E}[\text{ALG}'_2] \geq (1 - \epsilon)^2 \mathbb{E}[\text{OPT}'_2]$.*

Proof. First, let C denote the event that none of the OPT'_2 buckets are filled to the \tilde{r}_j capacities, i.e. C is the event that $|O_j| \leq \tilde{r}_j$ for all $j = 1, \dots, c$.

Claim 5.11. $\Pr[C] \geq 1 - \epsilon$.

Proof. Recall that $r_j = \mathbb{E}|O_j|$. Because $|O_j|$ is a sum of independent Bernoulli random variables, we have by a standard Bernstein bound that

$$\Pr[|O_j| - r_j \geq \sqrt{2r_j t} + t] \leq e^{-t}.$$

Setting $t = \ln \frac{\epsilon}{c}$, we get a bound of $\frac{\epsilon}{c}$; a union bound over the buckets will complete the proof. We just need to show that $\sqrt{2r_j \ln \frac{\epsilon}{c}} + \ln \frac{\epsilon}{c} \leq \beta$, as then the probability of exceeding $r_j + \beta$ is only smaller. As in the proof of Claim D.1, for this choice of ϵ , we have $\ln \frac{\epsilon}{c} \leq \ln(r)$, and $\ln(r) \leq r$, so

$$\sqrt{2r_j \ln \frac{\epsilon}{c}} + \ln \frac{\epsilon}{c} \leq \sqrt{2r \ln \frac{\epsilon}{c}} + \sqrt{r \ln \frac{\epsilon}{c}} \leq \beta. \quad \square$$

Next we argue that at each tier of thresholds, ALG' is getting just as many variables as OPT'_2 , even if their identities are different. First, a helpful property:

Claim 5.12. *Conditioned on C , suppose $|B_j| < \tilde{r}_j$ for some $j \geq 1$. Then $|B_0^{\text{light}}| + \sum_{j'=1}^j |B_{j'}| \geq \sum_{j'=1}^j |O_{j'}|$.*

Proof. Consider any arrival i that contributes to the right side. We claim it is also counted on the left. We know $X_i \geq Z_i \geq \tau_j$, because the variable contributes to the right side. So the algorithm will attempt to place X_i in some assigned bucket $j' \leq j$. If it does not succeed because the bucket is full, it will proceed to $j' + 1, \dots$, and possibly eventually j . Because $|B_j| < \tilde{r}_j$, we know there is space for i in bucket j , so the algorithm definitely takes i in bucket j or earlier. By definition, $Z_i < \tau_0$, so i cannot be a member of B_0^{heavy} . Therefore, it is counted by the left side. \square

Now we can show the key fact.

Claim 5.13. *Conditioned on C , we have for all $j = 1, \dots, c$ that*

$$\left| B_0^{\text{light}} \right| + \sum_{j'=1}^j |B_{j'}| \geq \sum_{j'=1}^j |O_{j'}|.$$

Proof. By induction on j . For $j = 1$, we must show $|B_0^{\text{light}}| + |B_1| \geq |O_1|$. Recall that an arrival i is a member of O_1 if $\tau_1 \leq Z_i < \tau_0$. There are two cases. If $|B_1| = \tilde{r}_1$, i.e. the bucket is full, then the case is proven as we have assumed event C , which implies $|O_1| \leq \tilde{r}_1$. Otherwise, the case follows by Claim 5.12.

Now consider $j > 1$. If $|B_j| < \tilde{r}_j$, i.e. the bucket is not full, then the case follows by Claim 5.12. Otherwise, i.e. bucket j is full, then we have $|B_j| \geq |O_j|$ because of property C . Combining this with the induction hypothesis proves that $|B_0^{\text{light}}| + \sum_{j'=1}^j |B_{j'}| \geq \sum_{j'=1}^j |O_{j'}|$. \square

We are now ready to complete the proof of Lemma 5.10. From Claim 5.13, given event C , we can make a one-to-one mapping from contributions Z'_i of OPT'_2 to contributions X_j of ALG'_2 , such that X_j is in the same bucket or a higher bucket than Z'_i . (I.e. map all elements of O_1 to elements of B_1 or B_0^{light} ; map all elements of O_2 to remaining elements of these or to elements of B_2 ; and so on.) For each such pair, we have $X_j \geq (1 - \epsilon)Z'_i$ because, at worst, both are in the same bucket.⁵ In total, this implies that, conditioned on C , we always have $\text{ALG}'_2 \geq (1 - \epsilon)\text{OPT}'_2$; and C occurs with probability at least $1 - \epsilon$. \square

Corollary 5.14. *For $\epsilon \geq \frac{9(\ln r)^{3/2}}{r^{1/4}}$, we have $\mathbb{E}[\text{ALG}'] \geq (1 - \epsilon)^2 \cdot \mathbb{E}[\text{OPT}']$.*

Proof. If $P \geq \epsilon$, then by Lemma 5.8, we have $\mathbb{E}[\text{ALG}'] \geq \mathbb{E}[\text{OPT}] \geq \mathbb{E}[\text{OPT}']$, proving the claim. Otherwise, by Lemma 5.9, $\mathbb{E}[\text{ALG}'_1] \geq (1 - P)\mathbb{E}[\text{OPT}'_1] \geq (1 - \epsilon)\mathbb{E}[\text{OPT}'_1]$; and by Lemma 5.10, $\mathbb{E}[\text{ALG}'_2] \geq (1 - \epsilon)^2\mathbb{E}[\text{OPT}'_2]$. So in the case $P < \epsilon$, we have

$$\begin{aligned} \mathbb{E}[\text{ALG}'] &= \mathbb{E}[\text{ALG}'_1] + \mathbb{E}[\text{ALG}'_2] &>> (1 - \epsilon)\mathbb{E}[\text{OPT}'_1] + (1 - \epsilon)^2\mathbb{E}[\text{OPT}'_2] \\ &>> (1 - \epsilon)^2\mathbb{E}[\text{OPT}']. \end{aligned} \quad \square$$

Proof of Augmentation Lemma 3.4. For $\epsilon \leq \frac{1}{2}$, $\epsilon \geq \frac{9(\ln r)^{3/2}}{r^{1/4}}$, we have:

$$\mathbb{E}[\text{ALG}] \stackrel{\text{Lemma 5.7}}{\geq} (1 - 2\epsilon)\mathbb{E}[\text{ALG}'] \stackrel{\text{Corollary 5.14}}{\geq} (1 - 2\epsilon)^3\mathbb{E}[\text{OPT}'] \stackrel{\text{Claim 5.5}}{\geq} (1 - 2\epsilon)^4\mathbb{E}[\text{OPT}].$$

So we obtain a $(1 + O(\epsilon))$ -approximation. \square

Acknowledgments. We are thankful to the anonymous reviewers of EC 2020 for helpful comments on improving the presentation of this paper.

A Unweighted Linear Correlations

In this section we consider a linear correlations model where the nonnegative matrix A from $\mathbf{X} = A \cdot \mathbf{Y}$ is unweighted, i.e., each of its entry is either 0 or 1. Alternately, for $i \in [n]$ there are known sets $S_1, S_2, \dots, S_n \subseteq [m]$ such that $X_i = \sum_{j \in S_i} Y_j$. Our lower bound tower instance from Section 4.3 no longer holds as it crucially exploits that matrix A has entries that decrease exponentially in ϵ . Can we do better than a $\Theta(\min\{s_{\text{col}}, s_{\text{row}}\})$ approximation ratio? One might wonder if there exists

⁵For example, we may have $Z'_i \approx \tau_4$ while $X_j = \tau_3 = (1 - \epsilon)\tau_4$.

an alternate hardness instance that only has 0 – 1 entries in A . We show that this is not the case. In fact, there exist simple threshold-based constant approximation algorithms.

Theorem 2.4. *The unweighted linear correlations problem has a fixed threshold constant-factor approximation algorithm.*

The main intuition in the proof of Theorem 2.4 is that for the unweighted problem each *independent* Y_j has limited “influence” on the X_i s. This is because either Y_j appears with coefficient 0, in which case it has no influence on the value X_i , or it appears with coefficient 1, in which case it has the same influence in the value of every such X_i . A threshold algorithm is therefore difficult to fool because unlike the tower instance, it is not possible to have a scenario where a Y_j is very large but our algorithm selects it within an X_i where it appears with a small coefficient ϵ .

For readers familiar with the revenue-maximization result of Babaioff et al. [BILW14], i.e., the best of selling items individually and selling all the items together in a single bundle is a constant factor approximation to optimal revenue, our result has a similar flavor, although the technical details are quite different. We decompose our problem instance into a “core” and a “tail” part. The tail consists of cases where any Y_j exceeds a boundary τ ; the core, the rest. For the tail case we show that approximating $\mathbb{E}[\max_j \{Y_j\}]$ (the best individual item) suffices, and in the core case we can approximate the best bundle X_i .

This argument will show that there is one fixed threshold τ_{core} such that the algorithm taking the first arrival above τ_{core} achieves a constant approximation to the optimal core contribution to $\max\{X_i\}$; and similarly for τ_{tail} and the tail part. There remains a corner case, where we show a fixed threshold equal to the boundary τ gives a constant factor. Thus, for any given instance, one can select among $\tau, \tau_{\text{core}}, \tau_{\text{tail}}$ to get a fixed-threshold constant-factor approximation algorithm.

A.1 Notation and Proof Overview

We first choose a real number τ representing a boundary. Let $p_j = \Pr[Y_j > \tau]$. We set τ such that $\prod_j (1 - p_j) = 1/2$, i.e., with half probability all Y_j are below τ . We let the set A be all “heavy” Y_j variables: $A := \{j : Y_j > \tau\}$.

Recall that for each X_i , the set of active indices is S_i , so we have $X_i = \sum_{j \in S_i} Y_j$. We first upper-bound OPT by contributions from the *core* event that $A = \emptyset$ (all Y_j are small) and from the remaining *tail* event.

Claim A.1.

$$\mathbb{E}[\text{OPT}] \leq \mathbb{E}[\max_i X_i \mid A = \emptyset] + \sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau].$$

Proof. For any outcome of Y_j s, we separately count those in A (larger than τ) and the rest, and relax the objective to always take the large ones:

$$\max_i X_i = \max_i \left\{ \sum_{j \in S_i} Y_j \right\} \leq \max_i \left\{ \sum_{j \in S_i \cap \bar{A}} Y_j \right\} + \sum_{j \in A} Y_j.$$

Now taking expectations on both sides,

$$\mathbb{E}[\text{OPT}] = \mathbb{E}[\max_i X_i] \leq \mathbb{E}[\max_i \left\{ \sum_{j \in S_i \cap \bar{A}} Y_j \right\}] + \mathbb{E}[\sum_{j \in A} Y_j]$$

$$\leq \mathbb{E}[\max_i X_i \mid A = \emptyset] + \sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau].$$

To justify that $\mathbb{E}[\max_i \{\sum_{j \in S_i \cap \bar{A}} Y_j\}] \leq \mathbb{E}[\max_i X_i \mid A = \emptyset]$, use a coupling argument: First draw all Y_j from their initial distributions and consider the value of $\max_i \{\sum_{j \in S_i \cap \bar{A}} Y_j\}$. Now take any variables $Y_j > \tau$, and redraw them until they fall below τ . The value of the inner sum can only increase, but now we are exactly obtaining $\mathbb{E}[\max_i X_i \mid A = \emptyset]$. \square

Given that Claim A.1 upper-bounds OPT by the sum of two terms, our proof goes in two steps. First, we approximate the *tail* contributions $\sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau]$. In Lemma A.2 we use the Augmentation Lemma 3.2 to give a simple fixed threshold- τ_{tail} algorithm with expected value $\Omega(\max\{Y_j\})$. This is a bit surprising because the lower bound in Section 4.3 actually proves such a result is not possible for *weighted* linear correlations. In Claim A.3, we show that $\Omega(\max\{Y_j\})$ suffices to capture the tail term.

To capture the *core* contributions $\mathbb{E}[\max_i X_i \mid A = \emptyset]$, in Claim A.5 we argue that for all instances where Y_j s are bounded by τ , we can use concentration of XOS functions to argue that $\Pr[\max_i X_i > 1/2 \cdot \mathbb{E}[\max_i X_i]]$ is at least a constant, and hence a simple fixed threshold- τ_{core} algorithm suffices. This second step also holds for prophets with weighted linear correlations. There is also a corner case where τ is too large to apply concentration. But in this case, setting a threshold τ will directly achieve a constant factor.

A.2 Proof

Lemma A.2. *For the prophet inequality problem with unweighted linear correlations, there exists a fixed threshold algorithm with expected value $\Omega(\max\{Y_j\})$.*

Proof. Define Z_i to be the sum of Y_j s that appear in X_i and have not appeared in any $X_{i'}$ for $i' < i$. Since every Y_j appears in some Z_i , we know $\max\{Z_i\} \geq \max\{Y_j\}$. The Augmentation Lemma 3.2 now completes the proof. \square

Now we argue that $\mathbb{E}[\max_j Y_j]$ takes care of the second term in Claim A.1.

Claim A.3.

$$\mathbb{E}[\max_j Y_j] \geq \frac{1}{2} \cdot \sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau].$$

Proof. We have

$$\mathbb{E}[\max_j Y_j] \geq \sum_j \Pr[A = \{j\}] \cdot \mathbb{E}[\max_j Y_j \mid A = \{j\}] \geq \sum_j \frac{p_j}{2} \cdot \mathbb{E}[Y_j \mid A = \{j\}],$$

where the second inequality uses $\Pr[A = \{j\}] = \Pr[Y_j > \tau] \Pr[Y_{j'} \leq \tau (\forall j' \neq j)] \geq (p_j) (\frac{1}{2})$ and that $\max_j Y_j$ given that $A = \{j\}$ is the same as Y_j . \square

Corollary A.4. *For any instance with unweighted linear correlations, there exists τ_{tail} such that the algorithm setting a fixed threshold of τ_{tail} obtains $\mathbb{E}[ALG] \geq \Omega\left(\sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau]\right)$.*

We now turn to the core portion of contributions to OPT.

Claim A.5. *Let $V = \mathbb{E}[\max_i X_i \mid A = \emptyset]$. If the boundary satisfies $\tau \leq V/10$, there exists τ_{core} such that the algorithm setting a fixed threshold of τ_{core} obtains $\mathbb{E}[ALG] \geq \Omega(V)$.*

Proof. Let Y'_j be a copy of Y_j conditioned on falling into the range $[0, \tau]$. Let $X'_i = \sum_{j \in S_i} Y'_j$ and let $W = \max_i X'_i$. We have $\mathbb{E}[\max_i X_i \mid A = \emptyset] = \mathbb{E}[W] = V$.

Now, note that W is an XOS function of the independent Y'_i variables (meaning is a maximum of weighted combinations). Thus we can apply the concentration of XOS functions (more generally, for self-bounding functions, see e.g. [Von10]) to get

$$\Pr[W < (1 - \delta) \cdot \mathbb{E}[W]] \leq \exp\left(-\delta^2 \cdot \frac{\mathbb{E}[W]}{2\tau}\right).$$

In particular, for $\delta = 1/2$ we get

$$\Pr\left[W < \frac{1}{2}\mathbb{E}[W]\right] \leq \exp\left(-\frac{\mathbb{E}[W]}{8\tau}\right) \leq \gamma,$$

for some constant $\gamma < 1$, using that $\tau \leq \mathbb{E}[W]/10$. Hence the expected value of an algorithm that sets a threshold of $\frac{1}{2}\mathbb{E}[W]$ is at least

$$\Pr\left[W \geq \frac{1}{2}\mathbb{E}[W]\right] \cdot \frac{1}{2}\mathbb{E}[W] \geq \Omega(\mathbb{E}[W]). \quad \square$$

Now we have all the tools to prove our main theorem.

Proof of Theorem 2.4. By Claim A.1, one of the following is at least a 2-approximation to the prophet: $V := \mathbb{E}[\max_i X_i \mid A = \emptyset]$, and $\mathbb{E}[\sum_j p_j \cdot \mathbb{E}[Y_j \mid Y_j > \tau]]$. Suppose it is the latter. Then by Claim A.3, setting a fixed threshold of τ_{tail} gives a constant-factor approximation.

So suppose we have $V \geq \mathbb{E}[\text{OPT}]/2$. If $\tau > V/10$, then we can set a fixed threshold of τ : with probability at least $\frac{1}{2}$, some $X_i \geq \tau$ (because some $Y_j \geq \tau$), so we obtain performance at least $\frac{1}{2}\tau \geq \frac{1}{40}\mathbb{E}[\text{OPT}]$.

Finally, if $V \geq \mathbb{E}[\text{OPT}]/2$ and $\tau \leq V/10$, then by Claim A.5, setting a fixed threshold of τ_{core} gives a constant-factor approximation. \square

B Negatively Associated Values

In this section we show a 2 approximation ratio for *negatively associated* random values, a property that is known to imply negative correlation⁶ [JDP83]. Formally, we say $\{X_j\}$ are negatively associated if for all monotone increasing functions f, g and disjoint subsets $S, S' \subseteq \{1, \dots, n\}$, and for all $a, b \in \mathbb{R}$, we have

$$\Pr[f(X_j : j \in S) \geq a] \leq \Pr[f(X_j : j \in S) \geq a \mid g(X_j : j \in S') \leq b].$$

Let $\tau = \frac{1}{2}\mathbb{E}[\max_i X_i]$ and let $P = \Pr[\max_i X_i \geq \tau]$. We can write down precisely the usual prophet proof with just one line requiring additional justification.

$$\begin{aligned} \mathbb{E}[\text{ALG}_\tau] &= P \cdot \tau + \sum_{i=1}^n \Pr[X_{i'} < \tau(\forall i' < i)] \cdot \mathbb{E}[(X_i - \tau)^+ \mid X_{i'} < \tau(\forall i' < i)] \\ &\geq P \cdot \tau + (1 - P) \cdot \sum_{i=1}^n \mathbb{E}[(X_i - \tau)^+ \mid X_{i'} < \tau(\forall i' < i)] \\ &\geq P \cdot \tau + (1 - P) \cdot \sum_{i=1}^n \mathbb{E}[(X_i - \tau)^+], \end{aligned}$$

where the last inequality uses negative association as $(X_i - \tau)^+$ is a monotone function as is $\max_{i' < i} X_{i'}$. We can also use weaker notions of negative correlation for the last inequality, e.g.,

⁶We say $\{X_j\}$ are negatively correlated if for all i, i' we have $\text{Cov}(X_i, X_{i'}) \leq 0$.

NLODS as in Section 2 of [RSC92]. Now repeating the old prophet inequality analysis,

$$\sum_{i=1}^n \mathbb{E}[(X_i - \tau)^+] \geq \mathbb{E}\left[\sum_{i=1}^n (X_i - \tau)^+\right] \geq \mathbb{E}\left[\max_i X_i - \tau\right] = \tau.$$

Thus, $\mathbb{E}[\text{ALG}_\tau] \geq P \cdot \tau + (1 - P) \cdot \tau = \tau$.

C Bounded s_{col} and Small Cardinality Constraint

For the r -uniform matroid problem with bounded s_{col} , we have shown in Section 5.1 a $1 + o(1)$ approximation for large r tending to infinity. Here, we complement that result with a gracefully-improving approximation ratio for all r that smoothly interpolates between $O(s_{\text{col}})$ for $r = 1$ (the classic result) and $O(1)$ for $r \geq s_{\text{col}}$.

The approach is an extension of our algorithm for bounded column sparsity in the $r = 1$ case.

Theorem C.1. *For any s_{col}, r , the linearly correlated prophets problem with column sparsity s_{col} and cardinality constraint r admits an approximation ratio of $2e^2 \cdot \max\{1, \frac{s_{\text{col}}}{r}\}$.*

In other words, as $r = 2, 3, \dots, s_{\text{col}}$, the guarantee improves to a constant factor times $\frac{2}{s_{\text{col}}}, \frac{3}{s_{\text{col}}}, \dots, 1$.

Proof. Let there be r sets (“buckets”) B_1, \dots, B_r . If $r < s_{\text{col}}$, let there also be a “discard pile” B_0 . Let $c = \max\{r, s_{\text{col}}\}$.

For each X_i , place it in a bucket $j \in \{1, \dots, r\}$ each chosen with probability $\frac{1}{c}$. If $r < s_{\text{col}}$, then with the remaining probability of $\frac{s_{\text{col}} - r}{s_{\text{col}}}$, place X_i in the discard bucket B_0 .

For each bucket $j = 1, \dots, r$, give the bucket a cardinality constraint of 1 item and run the following algorithm (based on the $r = 1$ case). Let $S_j = \{i : X_i \in B_j\}$ and note that they are disjoint for different j . When a variable X_i arrives, send it to the algorithm for its bucket, or if it is in B_0 , discard X_i and continue. In bucket j , we run an inclusion-threshold algorithm with S_j and with τ_j to be determined next. Assign each $Y_{j'}$ to the first $X_i \in B_j$ that includes it, i.e., let $T_i = \{j' : A_{ij'} > 0 \text{ and } A_{i'j'} = 0 (\forall i' < i, i' \in S_j)\}$. Let $Z_i = \sum_{j' \in T_i} A_{ij'} Y_{j'}$, and let $\tau_j = \frac{1}{2} \mathbb{E}[\max_{i \in S_j} Z_i]$.

Claim C.2. *For each bucket $j \in \{1, \dots, r\}$, the expected value selected by the algorithm is at least $\frac{1}{2e} \cdot \mathbb{E}[\max_{i \in S_j} X_i]$.*

Proof. By construction each Z_i is independent of all previous X_i in the same bucket, so by the Augmentation Lemma 3.2, bucket j obtains expected reward at least $\frac{1}{2} \mathbb{E}[\max_{i \in S_j} Z_i]$ with randomness over the variables.

For each $Y_{j'}$ with $A_{ij'} > 0$, we claim $\Pr[j \in T_i] \geq \frac{1}{e}$ because there are at most $s_{\text{col}} - 1$ other variables $X_{i'}$ that include $Y_{j'}$, and each misses bucket j with probability at least $1 - \frac{1}{s_{\text{col}}}$, so they all miss bucket j with probability at least $(1 - \frac{1}{s_{\text{col}}})^{s_{\text{col}} - 1} \geq \frac{1}{e}$. In this case, we must have $j \in T_i$. So for each fixed X_i , we have with probability only over the bucket assignments and construction of S_j ,

$$\mathbb{E}[Z_i] = \sum_{j'} \Pr[j' \in T_i] \cdot A_{ij'} Y_{j'} \geq \frac{1}{e} \sum_{j'} A_{ij'} Y_{j'} = \frac{1}{e} X_i.$$

Combining these facts, each bucket j obtains expected reward at least $\frac{1}{2e} \mathbb{E}[\max_{i \in S_j} X_i]$. \square

Write \mathbb{E}_B for an expectation taken over the bucketing and $\mathbb{E}_{\mathbf{X}}$ for expectation over the realizations of the variables. The above gives that for every set of variable realizations, $\mathbb{E}_B[\text{ALG}] \geq \frac{1}{2e} \sum_{j=1}^r \mathbb{E}_B[\max_{i \in S_j} X_i]$. By linearity of expectation and symmetry of the buckets, we have

$$\mathbb{E}_{\mathbf{X},B}[\text{ALG}] \geq \frac{r}{2e} \mathbb{E}_{\mathbf{X},B}[\max_{i \in S_1} X_i]. \quad (3)$$

Claim C.3. $\mathbb{E}_{\mathbf{X},B}[\max_{i \in S_1} X_i] \geq \frac{1}{e \cdot c} \cdot \mathbb{E}_{\mathbf{X}}[\text{OPT}]$, where $c = \max\{r, s_{\text{col}}\}$.

Proof. Let the random variable $X^{(i)}$ equal the i th-largest realized variable, i.e., in particular

$$\text{OPT} = \sum_{i=1}^r X^{(i)}.$$

Recall that any fixed variable falls into bucket B_1 with probability $\frac{1}{c}$. Fixing realizations of $X^{(1)}, \dots, X^{(r)}$, we have with probability taken only over the buckets,

$$\begin{aligned} \mathbb{E}_B[\max_{i \in S_1} X_i] &\geq \sum_{i=1}^r \Pr[X^{(i)} \in B_1 \text{ and } X^{(i')} \notin B_1 (\forall i' < i)] X^{(i)} \\ &= \sum_{i=1}^r \frac{1}{c} \left(1 - \frac{1}{c}\right)^{i-1} X^{(i)} \geq \frac{1}{e \cdot c} \sum_{i=1}^r X^{(i)} = \frac{1}{e \cdot c} \text{OPT}. \end{aligned}$$

Now taking an expectation on both sides over the realizations of \mathbf{X} proves the claim. \square

Finally, combine Claim C.3 with Inequality (3) to get

$$\mathbb{E}[\text{ALG}] \geq \frac{r}{2e} \mathbb{E}_{\mathbf{X}} \mathbb{E}_B[\max_{i \in S_1} X_i] \geq \frac{r}{2e^2 c} \mathbb{E}_{\mathbf{X}}[\text{OPT}],$$

which proves Theorem C.1. \square

D Missing Proofs

D.1 Missing Proofs from Section 2

Lemma 2.1. *In the linear correlations model, even for $s_{\text{row}} = s_{\text{col}} = 2$ there exist instances where every fixed threshold τ algorithm ALG_τ has an approximation ratio at least $\Omega(n)$.*

Proof of Lemma 2.1. We consider the $s_{\text{row}} = s_{\text{col}} = 2$ tower instance (Example 2.3), with sufficiently small ϵ chosen later. We claim that $\mathbb{E}[\max_i X_i] = \Omega(n)$ while any fixed threshold- τ algorithm has $\mathbb{E}[\text{ALG}_\tau] \leq 3$.

First, we bound $\mathbb{E}[\text{ALG}_\tau]$. Let $p_j = \Pr[\text{ALG}_\tau \text{ takes } X_j \text{ and } Y_j \text{ is active}]$. In this case the algorithm's reward includes $Y_j = \frac{1}{\epsilon^j}$ with coefficient 1. Let $p'_j = \Pr[\text{ALG}_\tau \text{ takes } X_{j-1} \text{ and } Y_j \text{ is active}]$. In this case its reward includes $Y_j = \frac{1}{\epsilon^j}$ with coefficient ϵ . We note that $p_j, p'_j \leq \Pr[Y_j \text{ is active}] = \epsilon^j$. By summing over Y_1, \dots, Y_n , we have

$$\mathbb{E}[\text{ALG}_\tau] = \sum_{j=1}^n \left(p_j \frac{1}{\epsilon^j} + (p'_j)(\epsilon) \frac{1}{\epsilon^j} \right) \leq \sum_{j=1}^n \left(p_j \frac{1}{\epsilon^j} + \epsilon \right) = n\epsilon + \sum_{j=1}^n p_j \frac{1}{\epsilon^j}.$$

We argue that $p_j = 0$ for all but at most two terms. Let j^* satisfy $\epsilon^{-(j^*+1)} < \tau \leq \epsilon^{-j^*}$, or $j^* = 1$ if $\tau \leq \frac{1}{\epsilon}$. We claim $p_j = 0$ if $j \leq j^* - 2$: assuming $\epsilon < \frac{1}{2}$, we have

$$X_j \leq \frac{2}{\epsilon^j} \leq \frac{1}{\epsilon^{j+1}} < \tau,$$

so X_j is never taken. We also claim $p_j = 0$ if $j \geq j^* + 1$: If Y_j is active, then

$$X_{j-1} \geq \epsilon \frac{1}{\epsilon^j} \geq \frac{1}{\epsilon^{j-1}} \geq \tau,$$

so X_{j-1} is taken and X_j is not. So we have $p_j = 0$ unless $j \in \{j^* - 1, j^*\}$, in which case $p_j \leq \epsilon^j$. So

$$\mathbb{E}[\text{ALG}_\tau] \leq n\epsilon + \sum_{j=1}^n p_j \frac{1}{\epsilon^j} \leq n\epsilon + 2 \leq 3$$

for $\epsilon \leq \frac{1}{n}$.

For the benchmark, since $\max_i \{X_i\} \geq \max_j \{Y_j\}$, it suffices to show $\mathbb{E}[\max_j Y_j] = \Omega(n)$. Now,

$$\begin{aligned} \mathbb{E}[\max_j \{Y_j\}] &= \sum_{i=1}^n \Pr[Y_j = 0 \ (\forall j > i)] \cdot \Pr[Y_i \neq 0] \cdot \left(\frac{1}{\epsilon^i}\right) \\ &\geq \Pr[Y_j = 0 \ (\forall j)] \cdot \sum_{i=1}^n \Pr[Y_i \neq 0] \cdot \left(\frac{1}{\epsilon^i}\right). \end{aligned}$$

Since $\Pr[Y_i \neq 0] = \epsilon^i$, we get

$$\begin{aligned} \mathbb{E}[\max_j \{Y_j\}] &\geq \Pr[Y_j = 0 \ (\forall j)] \cdot n \\ &\geq (1 - \sum_j \Pr[Y_j \neq 0]) \cdot n \geq (1 - n\epsilon) \cdot n \geq n/2 \end{aligned}$$

for any choice of $\epsilon \leq \frac{1}{2n}$. This gives an approximation ratio of at least $\frac{n/2}{3} = \frac{n}{6}$ for $\epsilon \leq \frac{1}{2n}$. \square

D.2 Missing Proofs from Section 5.3

Proof of Claim 5.5. First, consider supplementing OPT' by including tiny elements below τ_c when there is room, i.e. let $Z^{(i)''} = Z^{(i)} \mathbf{1}[Z_i < \tau_0]$ and consider $\text{OPT}_2'' = \sum_{j=1}^r Z^{(j)''}$. Let $\text{OPT}'' = \text{OPT}'_1 + \text{OPT}_2''$. On a case-by-case basis, OPT'' differs from OPT' by at most r elements, each at most $\tau_c \leq \frac{\epsilon}{r} \mathbb{E}[\text{OPT}]$. This proves that $\mathbb{E}[\text{OPT}'] \geq \mathbb{E}[\text{OPT}''] - \epsilon \mathbb{E}[\text{OPT}]$. We next prove that $\mathbb{E}[\text{OPT}'] \geq (1 - \epsilon) \mathbb{E}[\text{OPT}]$, which completes the proof of the claim.

Let H be the event there exists a heavy element, i.e. $\max_i Z_i \geq \tau_0 = \mathbb{E}[\text{OPT}]/\epsilon$. Let $p = \Pr[H]$. So,

$$\begin{aligned} \mathbb{E}[\text{OPT}'] &= p \cdot \mathbb{E}[\text{OPT}' | H] + (1 - p) \cdot \mathbb{E}[\text{OPT}' | \neg H] \\ &= p \cdot \mathbb{E}[\text{OPT}'' | H] + (1 - p) \cdot \mathbb{E}[\text{OPT} | \neg H] \\ &\geq p \cdot \mathbb{E}[Z^{(1)} | H] + (1 - p) \cdot \mathbb{E}[\text{OPT} | \neg H]. \end{aligned} \tag{4}$$

Now, we claim $\mathbb{E}[\text{OPT} | H] \leq \mathbb{E}[Z^{(1)} | H] + \mathbb{E}[\text{OPT}]$. Proof: let M_i be the event that $i = \arg \max_{i'} Z_{i'}$ and $Z_i \geq \tau_0$. Let OPT_{-i} be the sum of the largest $r - 1$ elements excluding Z_i . Note that conditioning on all others lying below Z_i , for any fixed Z_i , only decreases OPT_{-i} , as the variables are independent.

$$\begin{aligned} \mathbb{E}[\text{OPT} | H] &= \mathbb{E}[Z^{(1)} | H] + \sum_{i=1}^n \Pr[M_i | H] \cdot \mathbb{E}[\text{OPT}_{-i} | M_i] \\ &\leq \mathbb{E}[Z^{(1)} | H] + \sum_{i=1}^n \Pr[M_i | H] \cdot \mathbb{E}[\text{OPT}_{-i}] \\ &\leq \mathbb{E}[Z^{(1)} | H] + \sum_{i=1}^n \Pr[M_i | H] \cdot \mathbb{E}[\text{OPT}] = \mathbb{E}[Z^{(1)} | H] + \mathbb{E}[\text{OPT}]. \end{aligned}$$

Using this,

$$\begin{aligned} \mathbb{E}[\text{OPT}] &= p \cdot \mathbb{E}[\text{OPT} | H] + (1 - p) \cdot \mathbb{E}[\text{OPT} | \neg H] \\ &\leq p \cdot \mathbb{E}[Z^{(1)} | H] + p \cdot \mathbb{E}[\text{OPT}] + (1 - p) \cdot \mathbb{E}[\text{OPT} | \neg H]. \end{aligned}$$

This implies

$$(1-p)\mathbb{E}[\text{OPT}] \leq p \cdot \mathbb{E}[Z^{(1)} \mid H] + (1-p) \cdot \mathbb{E}[\text{OPT} \mid \neg H].$$

Combining with Inequality (4) gives $\mathbb{E}[\text{OPT}'] \geq (1-p)\mathbb{E}[\text{OPT}]$. We have $p \leq \epsilon$ by Markov's inequality: $p = \Pr[\max_i Z_i \geq \tau_0] \leq \mathbb{E}[\max_i Z_i]/\tau_0 \leq \mathbb{E}[\text{OPT}]/\tau_0 = \epsilon$. This completes the proof. \square

Proof of Lemma 5.7. Suppose r is large enough that $\epsilon \leq 0.5$, otherwise the lemma is immediate. Let $\delta := \frac{1+c\beta}{r}$.

Claim D.1. $\epsilon \geq 2\delta$.

Proof. Using that $\epsilon \geq r^{-1/4}$, we have $c := \lceil \frac{1}{\epsilon} \ln \frac{r}{\epsilon^2} \rceil \leq \frac{3}{\epsilon} \ln(r)$. Further using that $\epsilon \geq 3r^{-1/4}$, this implies $\ln \frac{r}{\epsilon} \leq \ln(r^{1/2} \ln(r)) \leq \ln(r)$. Therefore, $\beta := 3\sqrt{r \ln \frac{r}{\epsilon}} \leq 3\sqrt{r \ln(r)}$. Then $c \cdot \beta \leq \left(\frac{3}{\epsilon} \ln(r)\right) \left(3\sqrt{r \ln(r)}\right) \leq \frac{9\sqrt{r}(\ln r)^{3/2}}{\epsilon}$. Using that $\epsilon \geq 9r^{-1/4}(\ln r)^{3/2}$, this gives $c \cdot \beta \leq r^{3/4}$, so $1 + c \cdot \beta \leq 2r^{3/4}$, so $\delta \leq 2r^{-1/4} \leq \epsilon/2$. \square

Let K be the number of arrivals taken by ALG, a random variable.

Claim D.2. *With at least $1 - \epsilon$ probability, $K < r$ (i.e. ALG does not reach its cardinality constraint).*

Proof. The number of arrivals taken by ALG' is at most $\sum_{j=0}^c \tilde{r}_j = 1 + c \cdot \beta + \sum_{j=1}^c r_j$. Because OPT takes at most r arrivals pointwise, and thus in expectation, we have $\sum_{j=1}^c r_j \leq r$. So ALG' takes at most, in the worst case, $K' = r + 1 + c \cdot \beta = r(1 + \delta)$ arrivals. Because ALG keeps each independently with probability $1 - \epsilon \leq 1 - 2\delta < (1 - \delta)^2$, the chance it reaches r is upper-bounded by the chance that a $\text{Binomial}(K', (1 - \delta)^2)$ variable exceeds r . This is upper-bounded by the chance it exceeds $K'(1 - \delta)^2(1 + \delta) = r(1 - \delta)^2(1 + \delta)^2 = r(1 - \delta^2)^2 < r$. So by a Chernoff bound,

$$\Pr[K \geq r] \leq \Pr[\text{Binom}(K', (1 - \delta)^2) \geq K'(1 - \delta)^2(1 + \delta)] \leq \exp\left(\frac{-\delta^2 K'(1 - \delta)^2}{3}\right).$$

We have $\delta \leq \epsilon/2 \leq 0.25$, and $K' \geq r$, so $K'(1 - \delta)^2 \geq \frac{r}{2}$. Also, $\delta \geq \frac{c\beta}{r} \geq \frac{\beta}{r}$.

$$\Pr[K \geq r] \leq \exp\left(\frac{-\delta^2 r}{6}\right) \leq \exp\left(\frac{-\beta^2}{6r}\right) \leq \exp\left(-\ln \frac{c}{\epsilon}\right) \leq \epsilon. \quad \square$$

Now, each time ALG' obtains some variable X_i , ALG also obtains it unless either: it has reached its cardinality constraint; or it independently discards X_i (with probability ϵ). By a union bound over these two events, when ALG' obtains X_i , ALG also obtains it except with probability $1 - 2\epsilon$. So $\mathbb{E}[\text{ALG}] \geq (1 - 2\epsilon)\mathbb{E}[\text{ALG}']$, which proves Lemma 5.7. \square

References

- [ACK18] Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the $1-1/e$ barrier. In *Proceedings of EC*, 2018.
- [ADSY12] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Price of correlations in stochastic optimization. *Operations Research*, 60(1):150–162, 2012.
- [AEE⁺17] Melika Abolhasani, Soheil Ehsani, Hossein Esfandiari, MohammadTaghi HajiAghayi, Robert Kleinberg, and Brendan Lucier. Beating $1-1/e$ for ordered prophets. In *Proceedings of STOC*, pages 61–71. ACM, 2017.

- [AHL12] Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of EC*, pages 18–35, 2012.
- [AKW14] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of SODA*, 2014.
- [Ala11] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *Proceedings of FOCS*, 2011.
- [BDHS15] MohammadHossein Bateni, Sina Dehghani, MohammadTaghi Hajiaghayi, and Saeed Seddighin. Revenue maximization for selling multiple correlated items. In *Algorithms-ESA 2015*, pages 95–105. Springer, 2015.
- [BGSZ20] Domagoj Bradac, Anupam Gupta, Sahil Singla, and Goran Zuzic. Robust algorithms for the secretary problem. In *11th Innovations in Theoretical Computer Science Conference*, 2020.
- [BILW14] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *Proceedings of FOCS*, 2014.
- [CDFS19] José R. Correa, Paul Dütting, Felix A. Fischer, and Kevin Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In *Proceedings of EC*, pages 3–17, 2019.
- [CFH⁺17] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of EC*, pages 169–186, 2017.
- [CHMS10] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of STOC*, 2010.
- [CMS15] Shuchi Chawla, David Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- [CSV17] Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *Proceedings of STOC*, 2017.
- [CSZ19] José R. Correa, Raimundo Saona, and Bruno Ziliotto. Prophet secretary through blind strategies. In *Proceedings of SODA*, pages 1946–1961, 2019.
- [DFKL17] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing non-stochastic inputs. In *Proceedings of FOCS*, pages 540–551, 2017.
- [Dia18] Ilias Diakonikolas. Algorithmic high-dimensional robust statistics. Webpage <http://www.iliasdiakonikolas.org/simons-tutorial-robust.html>, 2018. Tutorial at Foundations of Data Science bootcamp.
- [DK19] Paul Dütting and Thomas Kesselheim. Posted pricing and prophet inequalities with inaccurate priors. In *Proceedings of EC*, pages 111–129, 2019.
- [DKK⁺16] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high dimensions without the computational intractability. In *Proceedings of FOCS*, 2016.
- [DKK⁺18] Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robustly learning a gaussian: Getting optimal error, efficiently. In *Proceedings of SODA*, pages 2683–2702, 2018.
- [EHKS18] Soheil Ehsani, Mohammad Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of SODA*, 2018.

- [EHL17] Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- [EKM18] Hossein Esfandiari, Nitish Korula, and Vahab Mirrokni. Allocation with traffic spikes: Mixing adversarial and stochastic models. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):14, 2018.
- [FGL15] Michal Feldman, Nick Gravin, and Brendan Lucier. Combinatorial auctions via posted prices. In *Proceedings of SODA*, pages 123–135, 2015.
- [FSZ16] Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of SODA*, pages 1014–1033, 2016.
- [HK92] Theodore P Hill and Robert P Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.
- [HKS07] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of AAAI*, 2007.
- [JDP83] Kumar Joag-Dev and Frank Proschan. Negative association of random variables with applications. *The Annals of Statistics*, pages 286–295, 1983.
- [KS78] Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Prob*, 4:197–266, 1978.
- [KW12] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of STOC*, pages 123–136, 2012.
- [LMPL18] Thodoris Lykouris, Vahab S. Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *Proceedings of STOC*, pages 114–122, 2018.
- [LRV16] Kevin A. Lai, Anup B. Rao, and Santosh Vempala. Agnostic estimation of mean and covariance. In *IEEE 57th Annual Symposium on Foundations of Computer Science*, 2016.
- [LS18] Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms, ESA*, 2018.
- [Moi18] Ankur Moitra. Robustness meets algorithms (invited talk). In *16th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2018*, pages 3:1–3:1, 2018.
- [RS17] Aviad Rubinstein and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of SODA*, pages 1671–1687. SIAM, 2017.
- [RSC⁺87] Yosef Rinott, Ester Samuel-Cahn, et al. Comparisons of optimal stopping values and prophet inequalities for negatively dependent random variables. *The Annals of Statistics*, 15(4), 1987.
- [RSC91] Yosef Rinott and Ester Samuel-Cahn. Orderings of optimal stopping values and prophet inequalities for certain multivariate distributions. *Journal of multivariate analysis*, 37(1), 1991.
- [RSC92] Yosef Rinott and Ester Samuel-Cahn. Optimal stopping values and prophet inequalities for some dependent random variables. *Lecture Notes-Monograph Series*, pages 343–358, 1992.
- [Rub16] Aviad Rubinstein. Beyond matroids: secretary problem and prophet inequality with general constraints. In *Proceedings of STOC*, pages 324–332, 2016.
- [RWW20] Aviad Rubinstein, Jack Z. Wang, and S. Matthew Weinberg. Optimal single-choice prophet inequalities from samples. In *Proceedings of ITCS*, 2020.
- [SC84] Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.
- [Von10] Jan Vondrák. A note on concentration of submodular functions. *arXiv:1005.2791*, 2010.
- [Yan11] Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of SODA*, 2011.