Optimal Prophet Inequality for Unions of Easy Systems

Submission 245

Abstract

We give an $O(\log k / \log \log k)$ prophet inequality when the system of feasible sets is the union of k "easy" systems (e.g., matroids and knapsacks) each admitting a constant-factor prophet inequality. This implies an $O(\log k / \log \log k)$ prophet inequality for downward-closed systems with k maximal feasible sets, matching the lower bound by Babaioff, Immorlica, and Kleinberg [BIK07]. As a result, under the parametrization by k, our upper bound is asymptotically optimal for downward-closed systems, and by implication, unions of k easy systems.

1 Introduction

Prophet inequalities are algorithms that, given distributional information, select items online subject to feasibility, so as to maximize the total value of items selected. More specifically, a total of n items have independent values, whose distributions are known to the algorithm upfront. Items arrive online one by one in an adversarial order. Upon the arrival of each item, the algorithm observes the realized value of the item, and must decide immediately and irrevocably whether to select the item. Throughout the procedure, the items selected must always be feasible, i.e., the set of items selected must be a member of a prescribed system of feasible sets. Different choices of the system of feasible sets lead to natural special cases of the problem, including single-choice prophet (the system consists of all singletons), multi-unit prophet (the system consists of all sets below a certain size), matroid prophet (the system corresponds to independent sets of a given matroid), etc.

As a result of persistent interest and effort in prophet inequalities, most important variants have now been largely resolved, in the sense that we know at least the asymptotically tight competitive ratio (i.e., value achievable online as a fraction of the offline optimum), and often even the precise numerical value thereof. Super-constant gaps remain in very few cases, among which perhaps the cleanest one is prophet inequalities with *general downward-closed systems* of feasible sets: Rubinstein [Rub16] presents an $O(\log^2 n)$ upper bound,[1] while Rubinstein and Singla [RS26] give a lower bound of $\Omega((\log n/\log\log n)^2)$. The two bounds almost match each other, but not exactly. This gap, while small, motivates our work in a curious sense discussed below — in fact, as a corollary of our main result, we will "close" the gap by looking at it in a different way.

The above bounds are parametrized by n, the number of items. We take another perspective, which, in the case of downward-closed systems, involves parametrizing the problem by k, the number of maximal feasible sets. While less popular, this is not unreasonable — for example, Feldman et al. [FGGS25] parametrize their bound using the same parameter in the context of clock auctions. More importantly, this parametrization enables a new, more general interpretation of downward-closed systems: A downward-closed system with k maximal sets is the union of k power sets (or free matroids), each generated by one maximal set. Under this interpretation, downward-closed systems generated by k maximal sets become a *special case* of something highly general: For example, one could consider systems that are unions of k matroids, or k sub-systems each defined by a knapsack constraint, or even mixtures thereof. In fact, our main result applies to any system that is the union of k *easy systems*, each admitting a (for example) constant-factor prophet inequality on its own. In other words, we measure the complexity of a problem instance by not how many items, but how many disjunctive easy systems / constraints it involves.

1.1 Results and Techniques

Our main result is a generic algorithmic reduction, from a problem instance defined on the union of k easy systems, to those defined on each of these easy systems:

Theorem 1 (informal version of Theorem 2)**.** *Consider k systems $\mathcal{F}_1, \ldots, \mathcal{F}_k$ defined over a common ground set of independent items. Suppose there is some $\alpha \geq 1$ such that for each $i \in [k]$, the problem instance defined by $\mathcal{F}_i$ alone admits an α-competitive prophet inequality. Then the problem*

[1] To be more precise, the bound is $O(\log n \cdot \log r)$, where r is the size of the largest feasible set.

instance defined by $\mathcal{F} = \bigcup_{i\in[k]} \mathcal{F}_i$ *admits an* $O(\alpha \cdot \log k/\log\log k)$*-competitive prophet inequality, where* $O(\cdot)$ *hides an absolute constant.*

As a direct corollary, there exists an $O(\log k/\log\log k)$-competitive prophet inequality for any problem instance defined by the union of k systems each admitting a constant-factor prophet inequality, including matroids, knapsacks, etc. This in particular subsumes downward-closed systems as a special case, which are unions of free matroids. As a result, we obtain the following bound for downward-closed systems generated by k maximal sets:

Corollary 1 (informal version of Corollary 2)**.** *Any problem instance defined by a downward-closed system containing* k *maximal sets admits an* $O(\log k/\log\log k)$*-competitive prophet inequality.*

Note that this bound is *asymptotically tight*: The hard instance by Babaioff, Immorlica, and Kleinberg [BIK07] (in which $k = \text{poly}(n)$) rules out $o(\log k/\log\log k)$-competitive prophet inequalities even restricted to the case of disjoint maximal sets and iid items with binary values. By implication, our main result for unions of k easy systems is also tight under the k-parametrization.

Proposition 1 (informal, corollary of [BIK07])**.** *The above upper bounds for downward-closed systems and unions of easy systems are asymptotically optimal under the* k*-parametrization.*

In fact, our bound for downward-closed systems is the first asymptotically tight bound, albeit under a less canonical parametrization. While this bound does not seem to easily imply an improved bound under the n-parametrization, it (as well as the underlying technical insights explained below) does hint at the possibility of such improvements.[2] At the same time, it narrows down the space of potential stronger lower bound constructions: We know that any improved lower bound construction would necessarily involve a large number of maximal feasible sets. On top of that, it must be structurally very rich, in the sense that it cannot be written as the union of polynomially many easy systems.

Overview of algorithm and analysis. The reduction itself is incredibly clean. We first rule out two easy cases:

- With constant probability, a single item contributes a significant (i.e., $\Omega(\log\log k/\log k)$) fraction of the offline optimum benchmark (call it opt).
- Items with large realized values (we will define "large" appropriately so that large realized values are rare) contribute more than half of opt in expectation.

In both cases, we simply take the first item that is large enough. In the remaining nontrivial case, we essentially pick the "most promising" sub-system out of the k, which is most likely to contain a set whose value is at least a constant fraction of opt, and disregard all the other possibilities. Restricted to the most promising sub-system, we run the existing α-competitive algorithm as a blackbox subroutine. This works because of concentration[3], as we explain below.

[2] Superficially, to obtain a better bound than $O(\log^2 n)$, one would need $k = \exp(o(\log^2 n \log\log n))$, whereas in general k could be as large as $\exp(\Omega(n))$ (e.g., a system containing all sets of size precisely $n/2$ and all subsets thereof). Certain more complex ideas, e.g., trying to cover a system using a few matroids (which in particular handles the above "hard" instance), do not easily give better bounds either.

[3] Here, concentration requires opt to be distributed somewhat evenly among sufficiently many small items. This can possibly be violated only in the easy cases above, which are handled separately.

Let $v_1, \ldots, v_n$ be the values of the n items, which are independent random variables. Consider the maximum value, say X, generated by the best set in any given system $\mathcal{F}$ (which could be either the union system or a sub-system, depending on the context), i.e.,

$$X = \max_{S \in \mathcal{F}} \sum_{i \in S} v_i.$$

First, imagine $\mathcal{F}$ is the union system. Rubinstein [Rub16] already observes that X is reasonably likely to be large enough, i.e., with constant probability, X is at least a constant fraction of its expectation (opt in this case). Given this, applying a simple union bound over the k sub-systems, we see that with probability $\Omega(1/k)$, the most promising sub-system generates at least a constant fraction of opt. That is, the most promising system is at least *somewhat likely to be very good*. Below we argue that this implies that the most promising system must be *somewhat good in expectation*, i.e., it generates a substantial fraction (i.e., $\Omega(\log\log k/\log k)$) of opt in expectation.

Now imagine $\mathcal{F}$ is our most promising system. We observe (as a corollary of well-established statistical facts) that X (which is now the value generated by the most promising system) is rarely much larger than its expectation. In fact, as long as $\mathbb{E}[X]$ is not too small (which, as we argue, must be the case), the tail probability $\Pr[X \geq t \cdot \mathbb{E}[X]]$ for $t \geq 1$ decays *super-exponentially* in t:

$$\Pr[X \geq t \cdot \mathbb{E}[X]] \leq \exp(-\Omega(t \log t)).$$

So, the probability that X is more than $\mathbb{E}[X] \cdot C \log k/\log\log k$ for some sufficiently large constant C must be sufficiently small (i.e., $o(1/k)$). On the other hand, recall that the probability that X is a constant fraction of opt is $\Omega(1/k)$, which means $\mathsf{opt} \leq C \log k/\log\log k \cdot \mathbb{E}[X]$. Now if we run an α-competitive algorithm for the most promising sub-system, the value we get would be $\mathbb{E}[X]/\alpha = \mathsf{opt} \cdot \log\log k/(\alpha \cdot C \log k)$. In other words, in addition to the factor of α, we lose a factor of $O(\log k/\log\log k)$ in the entire process.

1.2 Further Related Work

Prophet inequalities have been extremely well-studied in different models and contexts. In the classical single-choice prophet inequality, Krengel and Sucheston [KS77] show that the online algorithm can achieve at least half of the expected offline optimum. Subsequent research has explored the prophet problem in a variety of settings, including matroids [KW12; FSZ16; LS18], knapsack constraints [FSZ15; DFKL17], polymatroids [DK15], k-choice (uniform matroids) [HKS07; Ala11; JMZ22], matching constraints [GW19; EFGT20], independent sets in graphs [GHK+14], combinatorial allocation [FGL15; DFKL20; DKL20; Zha22; CC23], single-sample algorithms [AKW14; RWW20; CDF+22; DKL+24], etc. For comprehensive coverage of the problem and its variants, we refer interested readers to surveys [Luc17; CFH+18].

Having set the broader context, we now turn to the lines of research most closely related to the present work. For prophet inequalities on general downward-closed set systems, in addition to the aforementioned papers [BIK07; Rub16; RS26], Rubinstein and Singla [RS17] extend the problem to combinatorial valuation functions. This research is also relevant to intersections and unions of matroid systems. Kleinberg and Weinberg [KW12] show a competitive ratio of $(4k-2)$ for the intersection of k matroids and a lower bound of $\Omega(\sqrt{k})$. Feldman, Svensson, and Zenklusen [FSZ16] demonstrate an improved upper bound of $e(k+1)$, and Saxena, Velusamy, and Weinberg [SVW23] establish an improved lower bound of $k^{1/2+\Omega(1/\log\log k)}$. Alon et al. [AGP+25] consider "k-fold matroid unions", and present a prophet inequality which achieves at least $(1 - O(\sqrt{(\log k)/k}))$ of the expected optimum. This is different from unions of k matroids that we consider: The k-fold union of matroids is defined by taking one independent set from each of the k matroids, putting the union of these independent sets into the new system, and enumerating all combinations of independent sets. As a result, the problem defined by the k-fold union of matroids becomes *easier*. To the best of our knowledge, no prior work has studied prophet inequalities for unions of easy systems, or even special cases thereof, e.g., unions of matroid independent systems. Our main result extends the problem to a broader class of constraints.

2 Preliminaries

The downward-closed prophet problem. Consider a collection of n items $[n]$. Each item i has a value v_i, drawn independently from a known distribution D_i over $\mathbb{R}_{\geq 0}$. Let $\mathcal{F}$ be an arbitrary downward-closed set system on the item set. The algorithm has prior knowledge of n, $\mathcal{F}$, and all distributions $\{D_i\}$, and maintains a set ALG of selected elements.[4] An (adaptive online) adversary determines the arrival order: At each time, some unseen item i arrives, and its value $v_i \sim D_i$ realizes. Upon observing the realized value $v_i \sim D_i$, the algorithm must make an immediate and irreversible decision about whether to include the corresponding item in ALG, while keeping ALG feasible (i.e., $\mathsf{ALG} \in \mathcal{F}$). The objective is to maximize (the expectation of) the total selected value, i.e.,

$$\mathsf{alg} = \mathbb{E}\left[\sum_{i \in \mathsf{ALG}} v_i\right].$$

Let $\mathsf{OPT} := \operatorname{argmax}_{S \in \mathcal{F}} \sum_{i \in S} v_i$ denote the optimal feasible set in hindsight (ties broken arbitrarily), and

$$\mathsf{opt} := \mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i\right] = \mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i\right]$$

denote the expected offline optimum. We say that an online algorithm induces a prophet inequality with competitive ratio $\beta \geq 1$, if for any distributions $\{D_i\}_{i=1}^n$,

$$\beta \cdot \mathsf{alg} = \beta \cdot \mathbb{E}\left[\sum_{i \in \mathsf{ALG}} v_i\right] \geq \mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i\right] = \mathsf{opt},$$

where the expectations are taken over the randomness in the realization process (and possibly the randomness of the algorithm).

The system of feasible sets as the union of k easy systems. We say that a set system $\mathcal{F}$ is the union of k independence systems, iff there exist independence systems $\mathcal{F}_1, \cdots, \mathcal{F}_k \subseteq 2^{[n]}$ such that $\mathcal{F} = \cup_{j \in [k]} \mathcal{F}_j$. In other words, a set S is feasible in $\mathcal{F}$ iff there exists some $j \in [k]$ such that $S \in \mathcal{F}_j$.

As briefly discussed earlier, considering the union of k easy systems is more general than that of k maximal feasible sets. Since each system $\mathcal{F}_i \subseteq 2^{[n]}$ is a collection of maximal sets, it may, for example, be a matroid or a knapsack constraint. Also note that in principle, the k sub-systems are

[4] Slightly abusing notation, we will use ALG to denote both the algorithm and the set of items it selects.

not necessarily part of the problem instance, but something the algorithm is free to choose. For example, one could in principle view the collection of independent sets of a matroid as the union of the power sets generated by all its maximal members, where each maximal member induces its own easy sub-system. However, treating the entire collection of independent sets as one easy system, in our framework, normally leads to a better overall competitive ratio.

3 Algorithm for Unions of Easy Systems

In this section, we assume that for each $j \in [k]$, there exists a prophet algorithm ALG_j for the system $\mathcal{F}_j$ that gives a competitive ratio of α. For ease of presentation, we assume items arrive from 1 to n (our algorithm does not utilize the order of arrival). We begin with a verbal description of the algorithm:

1. Initialize $\mathsf{ALG} \leftarrow \emptyset$. Set t as the largest number such that $\Pr[\max_i v_i \geq t] \geq \frac{1}{2}$.

2. If $t > (\mathsf{opt} \cdot \log\log k)/(200000 \log k)$, take the first item with a value greater than or equal to t.

3. Otherwise, let $\mathsf{opt}_{>2t}$ denote the expected contribution from values greater than $2t$, i.e.,

$$\mathsf{opt}_{>2t} := \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i > 2t} v_i\right].$$

   If $\mathsf{opt}_{>2t} \geq \frac{\mathsf{opt}}{2}$, then take the first item with a value $> 2t$.

4. Otherwise, ignore values greater than $2t$, and let $\mathsf{opt}_{\leq 2t}$ denote the expected offline optimum of the resulting instance. That is,

$$\mathsf{opt}_{\leq 2t} := \mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i^{\leq 2t}\right], \quad \text{where} \quad v_i^{\leq 2t} = \begin{cases} v_i, & \text{if } v_i \leq 2t \\ 0, & \text{otherwise} \end{cases}.$$

   Choose the independence system $\mathcal{F}_j$ that maximizes $\Pr[\max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i \geq \frac{\mathsf{opt}_{\leq 2t}}{2}]$. Choose the set ALG by applying the existing α-competitive algorithm ALG_j to $\mathcal{F}_j$. (This guarantees an expected value no less than $\frac{1}{\alpha} \cdot \mathbb{E}[\max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i]$.)

The formal description is given as Algorithm 1. We argue that this algorithm admits a competitive ratio of $O(\alpha \cdot \log k / \log\log k)$.

Theorem 2. *Consider k independence systems $\mathcal{F}_1, \cdots, \mathcal{F}_k$ defined over $[n]$. Suppose that for each $j \in [k]$, the problem instance defined by $\mathcal{F}_j$ alone admits an α-competitive prophet inequality. Then Algorithm 1 is $O(\alpha \cdot \log k / \log\log k)$-competitive for the instance defined by $\mathcal{F} = \cup_{j \in [k]} \mathcal{F}_j$.*

Theorem 2 is proved in Sections 3.1–3.4. Before proceeding to the proof, we state the following corollary of Theorem 2.

Corollary 2. *For any downward-closed prophet instance containing k maximal feasible sets, Algorithm 1 obtains a competitive ratio of $O(\log k / \log\log k)$.*

Algorithm 1: Algorithm for unions of k easy systems.

Input: n, downward-closed set system $\mathcal{F}$, and distributions $\{D_i\}_{i=1}^n$
Output: Set of items ALG

/* Assume for each $j \in [k]$, there exists an α-competitive algorithm ALG_j for the system $\mathcal{F}_j$. */

Initialize $\mathsf{ALG} \leftarrow \emptyset$;
$t \leftarrow \max\{x : \Pr[\max_i v_i \geq x] \geq \frac{1}{2}\}$;
if $t > \frac{\mathsf{opt} \cdot \log\log k}{200000 \log k}$ **then**
  $i^* \leftarrow \min\{i : v_i \geq t\}$;
  $R \leftarrow \{i^*\}$;
else if $\mathsf{opt}_{>2t} \geq \frac{\mathsf{opt}}{2}$ **then**
  $i^* \leftarrow \min\{i : v_i > 2t\}$;
  $R \leftarrow \{i^*\}$;
else
  Ignore values $> 2t$;
  $\mathcal{F}_{j^*} \leftarrow \arg\max_{\mathcal{F}_j \subseteq \mathcal{F}} \Pr\left[\max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i \geq \frac{\mathsf{opt}_{\leq 2t}}{2}\right]$;
  $R \leftarrow \mathsf{ALG}_{j^*}(\mathcal{F}_{j^*})$;
$\mathsf{ALG} \leftarrow \mathsf{ALG} \cup R$;
return ALG;

Proof of Corollary 2 given Theorem 2. Each maximal feasible set is a free matroid. For each $j \in [k]$, take $\mathcal{F}_j$ as the power set of the j-th maximal feasible set $M_j \in \mathcal{F}$, we have $\mathcal{F} = \cup_{j \in [k]} \mathcal{F}_j$. For each prophet instance defined by $\mathcal{F}_j$, selecting all items in the set (i.e., $\mathsf{ALG}_j = M_j$ for each $j \in [k]$) gives a 1-competitive prophet inequality. By Theorem 2, Algorithm 1 yields a competitive ratio of $O(\log k / \log\log k)$. □

3.1 The Easy Cases

As introduced earlier, there are two easy cases where the algorithm takes the first item that is large enough. We summarize the results in the following lemma.

Lemma 1. *Consider any downward-closed prophet instance defined as the union of k independence systems. Algorithm 1 attains a competitive ratio of $O(\log k / \log\log k)$ for the following two cases:*

1. $t > \frac{\mathsf{opt} \cdot \log\log k}{200000 \log k}$.
2. $\mathsf{opt}_{>2t}$ *is no less than* $\frac{\mathsf{opt}}{2}$.

Proof. Recall that t is the largest number such that $\Pr[\max_i v_i \geq t] \geq \frac{1}{2}$. For case *1.*, the probability of taking an item with a value $\geq t$ is at least $\frac{1}{2}$. Therefore, the algorithm obtains an expected value of at least $(\mathsf{opt} \cdot \log\log k)/(400000 \log k)$.

For case *2.*, the algorithm takes the first item with a value $> 2t$. The probability of taking an item is less than $\frac{1}{2}$. In other words, for any item i, the probability that the algorithm reaches i without selecting any item is greater than $\frac{1}{2}$. We have:

$$\begin{aligned}\mathsf{alg} \geq & \sum_{i=1}^{n} \Pr[\text{reaching item } i] \cdot \Pr[v_i > 2t] \cdot \mathbb{E}[v_i | v_i > 2t] \\ > & \sum_{i=1}^{n} \frac{1}{2} \cdot \mathbb{E}[v_i \cdot \mathbf{1}_{v_i > 2t}] \\ = & \frac{1}{2} \sum_{i=1}^{n} \mathbb{E}[v_i \cdot \mathbf{1}_{v_i > 2t}].\end{aligned}$$

$\mathsf{opt}_{>2t}$ is no greater than $\sum_{i=1}^{n} \mathbb{E}[v_i \cdot \mathbf{1}_{v_i > 2t}]$. Therefore, the algorithm achieves at least a constant fraction of the expected contribution from values $> 2t$, and thus a constant fraction of opt. □

3.2 Tail Behavior of the Maximum Value

Our proof relies on the following tail bound by [Led97]:

Theorem 3 (Theorem 2.4 in [Led97])**.** *Let Y_i's be independent random variables with values in some space S; let $\mathcal{C}$ be a countable class of measurable functions $f : S \to [0, 1]$; and let $Z = \sup_{f \in \mathcal{C}} \sum_{i=1}^{n} f(Y_i)$. Then, for every $t \geq 0$,*

$$\Pr\left[Z \geq \mathbb{E}[Z] + t\right] \leq \exp\left(-\frac{t}{K} \log \frac{\mathbb{E}[Z] + t}{\mathbb{E}[Z]}\right). \tag{1}$$

For some numerical constant $K > 0$.

We interpret this bound similarly to [Rub16]. Let Y_i be the vector in $[0, 1]^{|\mathcal{F}_j|}$ whose S-th coordinate is v_i if $i \in S$, and 0 otherwise. Let $f_S(Y_i) := [Y_i]_S$ (i.e., the S-th coordinate of Y_i), so $\sum_{i=1}^{n} f_S(Y_i)$ is the value of set S. Let $\mathcal{C} := \{f_S\}_{S \in \mathcal{F}_j}$. And $Z = \sup_{f_S \in \mathcal{C}} \sum_{i=1}^{n} f_S(Y_i) = \max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i$. We derive the following lemma.

Lemma 2. *Consider an independence system $\mathcal{F}_j$ and independent variables $v_i \in [0, 1]$. Let $Z = \max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i$ with $\mathbb{E}[Z] = \lambda$. For $x \geq 2$, it holds that:*

$$\Pr\left[Z \geq x \cdot \lambda\right] \leq \exp\left(-\frac{1}{6590} \lambda \cdot x \log x\right).$$

We defer the proof of Lemma 2 to Appendix A.

3.3 The Main Reduction

We next present a concentration bound that is almost identical to Lemma 1 in [Rub16] — essentially a generalization of the bound by Rubinstein to random variables in $[0, 1]$. Let $V := \max_{S \in \mathcal{F}} \sum_{i \in S} v_i$ denote the value of optimal offline solution (which is a random variable).

Lemma 3. *For* $\mathsf{opt} \geq \Omega(\frac{\log k}{\log \log k})$,

$$\Pr\left[V \geq \frac{\mathsf{opt}}{2}\right] \geq \frac{1}{4}.$$

Proof. We have,

$$\begin{aligned}
\mathsf{opt} &= \int_0^\infty \Pr\left[V \geq u\right] \mathrm{d}u \\
&= \int_{-\mathsf{opt}}^\infty \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t \\
&= \int_{-\mathsf{opt}}^{-\mathsf{opt}/2} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t + \int_{-\mathsf{opt}/2}^{\mathsf{opt}} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t + \int_{\mathsf{opt}}^{\infty} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t.
\end{aligned}$$

Setting $t = u - \mathsf{opt}$ gives the second equation. For the first two integrals,

$$\begin{aligned}
&\int_{-\mathsf{opt}}^{-\mathsf{opt}/2} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t \leq \int_{-\mathsf{opt}}^{-\mathsf{opt}/2} 1 \mathrm{d}t \leq \frac{\mathsf{opt}}{2}, \\
&\int_{-\mathsf{opt}/2}^{\mathsf{opt}} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t \leq \int_{-\mathsf{opt}/2}^{\mathsf{opt}} \Pr\left[V \geq \frac{\mathsf{opt}}{2}\right] \mathrm{d}t \leq \frac{3\mathsf{opt}}{2} \Pr\left[V \geq \frac{\mathsf{opt}}{2}\right].
\end{aligned}$$

For the third part, we apply the concentration bound by [Led97] in Theorem 3. The interpretation is slightly adjusted. Let Y_i be the vector in $[0,1]^{|\mathcal{F}|}$ whose S-th coordinate is v_i if $i \in S$, and 0 otherwise. Let $f_S(Y_i) := [Y_i]_S$ (i.e., the S-th coordinate of Y_i), so $\sum_{i=1}^n f_S(Y_i)$ is the value of set S. Let $\mathcal{C} := \{f_S\}_{S \in \mathcal{F}}$. And $Z = \sup_{f_S \in \mathcal{C}} \sum_{i=1}^n f_S(Y_i) = \max_{S \in \mathcal{F}} \sum_{i \in S} v_i = V$. The concentration inequality can be rewritten as:

$$\Pr\left[V \geq \mathsf{opt} + t\right] \leq \exp\left(-\frac{t}{K} \log\left(1 + \frac{t}{\mathsf{opt}}\right)\right).$$

Therefore, we have,

$$\begin{aligned}
\int_{\mathsf{opt}}^{\infty} \Pr\left[V \geq \mathsf{opt} + t\right] \mathrm{d}t &\leq \int_{\mathsf{opt}}^{\infty} \exp\left(-\frac{t}{K} \log\left(1 + \frac{t}{\mathsf{opt}}\right)\right) \mathrm{d}t \\
&\leq \int_{\mathsf{opt}}^{\infty} \exp\left(-\frac{t}{K}\right) \mathrm{d}t = K \cdot \exp\left(-\frac{\mathsf{opt}}{K}\right).
\end{aligned}$$

$K \cdot \exp\left(-\frac{\mathsf{opt}}{K}\right)$ is negligible since $\mathsf{opt} \geq \Omega(\frac{\log k}{\log\log k})$. Collecting the three parts, we get,

$$\mathsf{opt} \leq \frac{\mathsf{opt}}{2} + \frac{3\mathsf{opt}}{2} \Pr\left[V \geq \frac{\mathsf{opt}}{2}\right] + o(1).$$

It follows that,

$$\Pr\left[V \geq \frac{\mathsf{opt}}{2}\right] \geq \frac{1}{3} - o(1).$$

□

The next lemma establishes that for sufficiently large opt, we can get a constant lower bound on the expected value of the selected sub-system.

Lemma 4. *When $v_i \in [0,1]$ for any item i and $\mathsf{opt} \geq \frac{50000 \log k}{\log\log k}$, for the system $\mathcal{F}_j$ that maximizes $\Pr[\max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i \geq \frac{\mathsf{opt}}{2}]$, it holds that $\mathbb{E}[\max_{S \in \mathcal{F}_j} \sum_{i \in S} v_i] \geq 1$.*

Proof. Assume, for the sake of contradiction, that $\mathbb{E}[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i] < 1$. Consider a more promising imaginary set $S' \in \mathcal{F}'_j$ that dominates any $S \in \mathcal{F}_j$: let $\{v'_i\}_{i\in S'}$ denote the values in set S': for all $i \in S'$, $v'_i \in [0,1]$ and stochastically dominates v_i; and $\mathbb{E}[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i] = 1$. Then,

$$\Pr\left[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i \geq \frac{\mathsf{opt}}{2}\right] \geq \Pr\left[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i \geq \frac{\mathsf{opt}}{2}\right].$$

Apply Lemma 2 and $\mathbb{E}[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i] = 1$, we have:

$$\begin{aligned}
\Pr\left[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i \geq \frac{\mathsf{opt}}{2}\right] &\leq \Pr\left[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i \geq \frac{25000\log k}{\log\log k}\right] \\
&\leq \exp\left(-\frac{1}{6590}\cdot\frac{25000\log k}{\log\log k}\log\left(\frac{25000\log k}{\log\log k}\right)\right) \\
&= \exp\left(-\frac{2500\log k}{659\log\log k}\left(\log 25000 + \log\log k - \log\log\log k\right)\right) \\
&= k^{\left(-\frac{2500}{659}\left(\frac{\log 25000}{\log\log k}+1-\frac{\log\log\log k}{\log\log k}\right)\right)} \\
&< k^{\frac{2500}{659}\left(\frac{\log\log\log k}{\log\log k}-1\right)} \\
&\leq k^{\frac{2500}{659}\left(\frac{1}{e}-1\right)} \\
&< k^{-2.3}.
\end{aligned}$$

The second last inequality is due to $\frac{\log\log\log k}{\log\log k} \leq \frac{1}{e}$.

By Lemma 3, $\Pr[V \geq \frac{\mathsf{opt}}{2}] \geq \frac{1}{4}$. According to the union bound, for the system $\mathcal{F}_j$ that maximizes $\Pr[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i \geq \frac{\mathsf{opt}}{2}]$, we have,

$$\Pr[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i \geq \frac{\mathsf{opt}}{2}] \geq \frac{1}{4k}.$$

That implies,

$$\Pr[\max_{S'\in\mathcal{F}_j}\sum_{i\in S'} v'_i \geq \frac{\mathsf{opt}}{2}] \geq \Pr[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i \geq \frac{\mathsf{opt}}{2}] \geq \frac{1}{4k}.$$

This contradicts $\Pr\left[\max_{S'\in\mathcal{F}'_j}\sum_{i\in S'} v'_i \geq \frac{\mathsf{opt}}{2}\right] < k^{-2.3}$. Therefore, the initial assumption is false, and we conclude that $\mathbb{E}[\max_{S\in\mathcal{F}_j}\sum_{i\in S} v_i] \geq 1$. ☐

It remains to consider the case that $t \leq (\mathsf{opt}\cdot\log\log k)/(200000\log k)$ and $\mathsf{opt}_{>2t} < \frac{\mathsf{opt}}{2}$. The following lemma shows the algorithm incurs an additional factor of $O(\log k/\log\log k)$ in the reduction process.

Lemma 5. *Suppose that for each $j \in [k]$, the problem instance defined by $\mathcal{F}_j$ alone admits an α-competitive prophet inequality. When $t \leq (\mathsf{opt}\cdot\log\log k)/(200000\log k)$ and $\mathsf{opt}_{>2t} < \frac{\mathsf{opt}}{2}$, Algorithm 1 achieves a competitive ratio of $O(\alpha\cdot\frac{\log k}{\log\log k})$.*

Proof. We first show that $\mathsf{opt}_{\leq 2t} > \frac{\mathsf{opt}}{2}$.

Let $\mathbf{V}$ be the vector of all random variables. For any fixed realization $\mathbf{V} = \mathbf{v} = (v_1, \cdots, v_n)$, the set OPT is fixed:

$$\begin{aligned}
\mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i \middle| \mathbf{V} = \mathbf{v}\right] &= \sum_{i \in \mathsf{OPT}: v_i > 2t} v_i + \sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i \\
&= \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i > 2t} v_i \middle| \mathbf{V} = \mathbf{v}\right] + \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i \middle| \mathbf{V} = \mathbf{v}\right].
\end{aligned}$$

Since this inequality holds for any realization $\mathbf{v}$, we have:

$$\mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i \middle| \mathbf{V}\right] = \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i > 2t} v_i \middle| \mathbf{V}\right] + \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i \middle| \mathbf{V}\right].$$

Take the unconditional expectation $\mathbb{E}[\cdot]$ on both sides. By linearity of expectation and the law of total expectation,

$$\mathsf{opt} = \mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i\right] = \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i > 2t} v_i\right] + \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i\right].$$

We have $\mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i\right] = \mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i^{\leq 2t}\right]$. Since $\mathsf{OPT} \in \mathcal{F}$, for any fixed realization $\mathbf{v}$:

$$\mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i^{\leq 2t} \middle| \mathbf{V} = \mathbf{v}\right] \geq \mathbb{E}\left[\sum_{i \in \mathsf{OPT}} v_i^{\leq 2t} \middle| \mathbf{V} = \mathbf{v}\right].$$

Similarly, apply linearity of expectation and the law of total expectation, it follows that:

$$\mathsf{opt}_{\leq 2t} = \mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v_i^{\leq 2t}\right] \geq \mathbb{E}\left[\sum_{i \in \mathsf{OPT}: v_i \leq 2t} v_i\right] > \frac{\mathsf{opt}}{2}.$$

Recall that we ignore all values greater than $2t$. All items have values in $[0, 2t]$ (i.e., $[0, (\mathsf{opt} \cdot \log\log k)/(100000 \log k)]$). Scale the valuation support such that all items have values in $[0, 1]$. Let $\{v'_i\}_{i \in [n]}$ denote the values after scaling; and let opt' denote the expected offline optimum after scaling:

$$v'_i = \begin{cases} v_i / \frac{\mathsf{opt} \cdot \log\log k}{100000 \log k}, & \text{if } v_i \leq 2t \\ 0, & \text{otherwise} \end{cases} \qquad \text{and} \qquad \mathsf{opt}' = \mathbb{E}\left[\max_{S \in \mathcal{F}} \sum_{i \in S} v'_i\right].$$

We get that:

$$\mathsf{opt}' = \frac{\mathsf{opt}_{\leq 2t} \cdot 100000 \log k}{\mathsf{opt} \cdot \log\log k} > \frac{50000 \log k}{\log\log k}.$$

To show that the algorithm obtains an $\Omega(\frac{1}{\alpha} \cdot \frac{\log\log k}{\log k})$-fraction of opt, it is equivalent to show that after scaling, it gets an $\Omega(\frac{1}{\alpha} \cdot \frac{\log\log k}{\log k})$-fraction of opt'.

Let $\mathcal{F}_{j^*}$ be the independence system selected by the algorithm that maximizes $\Pr[\max_{S\in\mathcal{F}_j} \sum_{i\in S} v'_i \geq \frac{\mathsf{opt}'}{2}]$. Let γ denote the expected optimum value in $\mathcal{F}_{j^*}$, i.e., $\gamma := \mathbb{E}[\max_{S\in\mathcal{F}_{j^*}} \sum_{i\in S} v'_i]$. By Lemma 4, we get $\gamma \geq 1$.

If $\frac{\mathsf{opt}'}{2\gamma} < 2$, we would directly get a constant upper bound on $\frac{\mathsf{opt}'}{\gamma}$, which implies a competitive ratio of $O(\alpha)$ for the algorithm. For any $\frac{\mathsf{opt}'}{2\gamma} \geq 2$, Lemma 2 implies that,

$$\begin{aligned}\Pr\left[\max_{S\in\mathcal{F}_{j^*}} \sum_{i\in S} v'_i \geq \frac{\mathsf{opt}'}{2\gamma}\cdot\gamma\right] &\leq \exp\left(-\frac{\gamma}{6590}\cdot\frac{\mathsf{opt}'}{2\gamma}\log\left(\frac{\mathsf{opt}'}{2\gamma}\right)\right)\\ &\leq \exp\left(-\frac{1}{6590}\cdot\frac{\mathsf{opt}'}{2\gamma}\log\left(\frac{\mathsf{opt}'}{2\gamma}\right)\right).\end{aligned}$$

Let V' denote $\left(\max_{S\in\mathcal{F}}\sum_{i\in S} v'_i\right)$. Applying Lemma 3 yields $\Pr[V' \geq \frac{\mathsf{opt}'}{2}] \geq \frac{1}{4}$. By the union bound, for system $\mathcal{F}_{j^*}$, we have $\Pr[\max_{S\in\mathcal{F}_{j^*}} \sum_{i\in S} v'_i \geq \frac{\mathsf{opt}'}{2}] \geq \frac{1}{4k}$. Thus,

$$\frac{1}{4k} \leq \Pr\left[\max_{S\in\mathcal{F}_{j^*}} \sum_{i\in S} v'_i \geq \frac{\mathsf{opt}'}{2}\right] \leq \exp\left(-\frac{1}{6590}\frac{\mathsf{opt}'}{2\gamma}\log\frac{\mathsf{opt}'}{2\gamma}\right).$$

This implies,

$$\frac{\mathsf{opt}'}{2\gamma} = O\left(\frac{\log k}{\log\log k}\right).$$

Applying Algorithm ALG_{j^*} to system $\mathcal{F}_{j^*}$ achieves at least $\frac{\gamma}{\alpha}$. Therefore, Algorithm 1 is $O(\alpha\cdot\frac{\log k}{\log\log k})$-competitive. □

Note that we use the α-competitive algorithm ALG_{j^*} on the selected sub-system $\mathcal{F}_{j^*}$ as a subroutine. Given the existence of constant-factor prophet inequalities on e.g., matroids and knapsacks ([KW12; DFKL17]), it follows directly that our algorithm provides $O(\frac{\log k}{\log\log k})$-competitive prophet inequalities for unions of matroids, knapsack constraints, or mixtures.

3.4 Putting Everything Together

Having established the necessary results, we now prove the main theorem.

Proof of Theorem 2. By Lemma 1, we get that the algorithm is $O(\log k/\log\log k)$-competitive for the cases corresponding to line 3 and line 6 of Algorithm 1.

Lemma 5 shows that the algorithm is $O(\alpha\cdot\log k/\log\log k)$-competitive for the case corresponding to line 9 of Algorithm 1. The proof of the theorem is complete. □

4 A Matching Lower Bound

To be self-contained, we restate the lower bound construction by Babaioff, Immorlica, and Kleinberg [BIK07] and show that it gives a lower bound asymptotically matching our upper bound, under the parameterization by k. We include the construction stated in Appendix B of [Rub16].

The idea is the following: Partition n items into disjoint sets of size $\log n/\log\log n$, each set corresponds to a maximal feasible set in $\mathcal{F}$. Assign values $\{0,1\}$ independently to the items such that for each item i, $v_i = 1$ with probability $\frac{\log\log n}{\log n}$ and $v_i = 0$ with probability $1 - \frac{\log\log n}{\log n}$.

Theorem 4 (adapted from [BIK07; Rub16])**.** *No online algorithm can achieve a competitive ratio of* $o(\frac{\log k}{\log\log k})$ *for downward-closed prophet.*

Proof. The instance above contains $k = \frac{n \cdot \log\log n}{\log n}$ independence systems. The Theorem in Appendix B of [Rub16] establishes the following conclusion that is relevant to our analysis: opt is approximately $(\log n / \log\log n)$, and any online algorithm cannot achieve an expected value of more than 2. It is impossible for any online algorithm to achieve a ratio better than $o(\frac{\log n}{\log\log n})$.

We next show that for $\log n > 4$,

$$\frac{\log n}{\log\log n} \geq \frac{\log k}{2\log\log k}.$$

Let m denote $\log n$. We have,

$$\log k = m + \log\log m - \log m \leq m.$$

For $m > 4$, $m - \log m > m/2$, we have $\log k > m - \log m > \frac{m}{2}$. Therefore,

$$\log\log k > \log m - \log 2.$$

Thus, $\frac{\log k}{\log\log k} \leq \frac{m}{\log m - \log 2}$. For any $m > 4$,

$$\frac{\log n / \log\log n}{\log k / \log\log k} \geq \frac{m}{\log m} \cdot \frac{\log m - \log 2}{m} = \frac{\log m - \log 2}{\log m} \geq \frac{1}{2}.$$ □

References

[AGP+25] Noga Alon, Nick Gravin, Tristan Pollner, Aviad Rubinstein, Hongao Wang, S. Matthew Weinberg, and Qianfan Zhang. "A Bicriterion Concentration Inequality and Prophet Inequalities for k-Fold Matroid Unions". In: *16th Innovations in Theoretical Computer Science Conference*. Ed. by Raghu Meka. LIPIcs. 2025, 4:1–4:22.

[AKW14] Pablo Daniel Azar, Robert Kleinberg, and S. Matthew Weinberg. "Prophet Inequalities with Limited Information". In: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Chandra Chekuri. 2014, pp. 1358–1377.

[Ala11] Saeed Alaei. "Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers". In: *IEEE 52nd Annual Symposium on Foundations of Computer Science*. Ed. by Rafail Ostrovsky. 2011, pp. 512–521.

[BIK07] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. "Matroids, secretary problems, and online mechanisms". In: *Symposium on Discrete Algorithms (SODA'07)*. 2007, pp. 434–443.

[CC23] José Correa and Andrés Cristi. "A Constant Factor Prophet Inequality for Online Combinatorial Auctions". In: *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*. Ed. by Barna Saha and Rocco A. Servedio. 2023, pp. 686–697.

[CDF+22] Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhuser. "Single-Sample Prophet Inequalities via Greedy-Ordered Selection". In: *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2022), pp. 1298–1325.

[CFH+18] José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. "Recent developments in prophet inequalities". In: *SIGecom Exch.* 17.1 (2018), pp. 61–70.

[DFKL17] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. "Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Non-Stochastic Inputs". In: *58th IEEE Annual Symposium on Foundations of Computer Science*. Ed. by Chris Umans. 2017, pp. 540–551.

[DFKL20] Paul Dütting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. "Prophet Inequalities Made Easy: Stochastic Optimization by Pricing Nonstochastic Inputs". In: *SIAM J. Comput.* 49.3 (2020), pp. 540–582.

[DK15] Paul Dütting and Robert Kleinberg. "Polymatroid Prophet Inequalities". In: *Algorithms - ESA 2015 - 23rd Annual European Symposium*. Ed. by Nikhil Bansal and Irene Finocchi. Lecture Notes in Computer Science. 2015, pp. 437–449.

[DKL+24] Paul Dütting, Thomas Kesselheim, Brendan Lucier, Rebecca Reiffenhäuser, and Sahil Singla. "Online Combinatorial Allocations and Auctions with Few Samples". In: *65th IEEE Annual Symposium on Foundations of Computer Science*. 2024, pp. 1231–1250.

[DKL20] Paul Dütting, Thomas Kesselheim, and Brendan Lucier. "An O(log log m) Prophet Inequality for Subadditive Combinatorial Auctions". In: *61st IEEE Annual Symposium on Foundations of Computer Science*. Ed. by Sandy Irani. 2020, pp. 306–317.

[EFGT20] Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. "Online Stochastic Max-Weight Matching: Prophet Inequality for Vertex and Edge Arrival Models". In: *EC '20: The 21st ACM Conference on Economics and Computation*. Ed. by Péter Biró, Jason D. Hartline, Michael Ostrovsky, and Ariel D. Procaccia. 2020, pp. 769–787.

[FGGS25] Michal Feldman, Vasilis Gkatzelis, Nick Gravin, and Daniel Schoepflin. "Bayesian and Randomized Clock Auctions". In: *Operations Research* 73.4 (2025), pp. 1965–1982.

[FGL15] Michal Feldman, Nick Gravin, and Brendan Lucier. "Combinatorial Auctions via Posted Prices". In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Piotr Indyk. 2015, pp. 123–135.

[FSZ15] Moran Feldman, Ola Svensson, and Rico Zenklusen. "A Simple $O(\log\log(\text{rank}))$-Competitive Algorithm for the Matroid Secretary Problem". In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Piotr Indyk. 2015, pp. 1189–1201.

[FSZ16] Moran Feldman, Ola Svensson, and Rico Zenklusen. "Online Contention Resolution Schemes". In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Robert Krauthgamer. 2016, pp. 1014–1033.

[GHK+14] Oliver Göbel, Martin Hoefer, Thomas Kesselheim, Thomas Schleiden, and Berthold Vöcking. "Online Independent Set Beyond the Worst-Case: Secretaries, Prophets, and Periods". In: *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014*. Ed. by Javier Esparza, Pierre Fraigniaud, Thore Husfeldt, and Elias Koutsoupias. Lecture Notes in Computer Science. 2014, pp. 508–519.

[GW19] Nikolai Gravin and Hongao Wang. "Prophet Inequality for Bipartite Matching: Merits of Being Simple and Non Adaptive". In: *Proceedings of the 2019 ACM Conference on Economics and Computation*. Ed. by Anna R. Karlin, Nicole Immorlica, and Ramesh Johari. 2019, pp. 93–109.

[HKS07] Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. "Automated Online Mechanism Design and Prophet Inequalities". In: *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence*. 2007, pp. 58–65.

[JMZ22] Jiashuo Jiang, Will Ma, and Jiawei Zhang. "Tight Guarantees for Multi-unit Prophet Inequalities and Online Stochastic Knapsack". In: *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Joseph (Seffi) Naor and Niv Buchbinder. 2022, pp. 1221–1246.

[KS77] Ulrich Krengel and Louis Sucheston. "Semiamarts and finite values". In: *Bulletin of the American Mathematical Society* 83.1977 (1977), pp. 745–747.

[KW12] Robert Kleinberg and S. Matthew Weinberg. "Matroid prophet inequalities". In: *Proceedings of the 44th Symposium on Theory of Computing Conference*. Ed. by Howard J. Karloff and Toniann Pitassi. 2012, pp. 123–136.

[Led97] Michel Ledoux. "On Talagrand's deviation inequalities for product measures". In: *ESAIM: Probability and Statistics* 1 (1997), pp. 63–87.

[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*. Ed. by Yossi Azar, Hannah Bast, and Grzegorz Herman. LIPIcs. 2018, 57:1–57:14.

[Luc17] Brendan Lucier. "An economic view of prophet inequalities". In: *SIGecom Exch.* 16.1 (2017), pp. 24–47.

[RS17] Aviad Rubinstein and Sahil Singla. "Combinatorial Prophet Inequalities". In: *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. Ed. by Philip N. Klein. 2017, pp. 1671–1687.

[RS26] Aviad Rubinstein and Sahil Singla. "Secretary, Prophet, and Stochastic Probing via Big-Decisions-First". In: *arXiv preprint arXiv:2604.00437* (2026).

[Rub16] Aviad Rubinstein. "Beyond matroids: Secretary problem and prophet inequality with general constraints". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 324–332.

[RWW20] Aviad Rubinstein, Jack Z. Wang, and S. Matthew Weinberg. "Optimal Single-Choice Prophet Inequalities from Samples". In: *11th Innovations in Theoretical Computer Science Conference*. Ed. by Thomas Vidick. LIPIcs. 2020, 60:1–60:10.

[SVW23] Raghuvansh R. Saxena, Santhoshini Velusamy, and S. Matthew Weinberg. "An Improved Lower Bound for Matroid Intersection Prophet Inequalities". In: *14th Innovations in Theoretical Computer Science Conference*. Ed. by Yael Tauman Kalai. LIPIcs. 2023, 95:1–95:20.

[Zha22] Hanrui Zhang. "Improved prophet inequalities for combinatorial welfare maximization with (approximately) subadditive agents". In: *J. Comput. Syst. Sci.* 123 (2022), pp. 143–156.

A Omitted Proofs

Proof of Lemma 2. As shown in the proof of Theorem 2.4 in [Led97]:

$$\Pr\left[Z \geq \mathbb{E}[Z] + t\right] \leq \exp\left(-\frac{t^2}{32\mathbb{E}[Z]}\right) \qquad \forall\, t \leq 4\mathbb{E}[Z] \tag{i}$$

$$\Pr\left[Z \geq \mathbb{E}[Z] + t\right] \leq \exp\left(-\frac{t}{8}\right) \qquad \forall\, t > 4\mathbb{E}[Z] \tag{ii}$$

$$\Pr\left[Z \geq t\right] \leq \exp\left(-\frac{t}{8}\log\frac{t}{\mathbb{E}[Z]} + t\right) \qquad \forall\, t \geq 8\mathbb{E}[Z]. \tag{iii}$$

To establish an upper bound of the form $\exp\left(-\frac{t}{K}\log\frac{\mathbb{E}[Z]+t}{\mathbb{E}[Z]}\right)$ for some constant K, we consider the following constraints, depending on the range of t [5]:

$$\exp\left(-\frac{t^2}{32\mathbb{E}[Z]}\right) \leq \exp\left(-\frac{t}{K}\log\frac{\mathbb{E}[Z]+t}{\mathbb{E}[Z]}\right) \qquad \forall\, t \leq 4\mathbb{E}[Z] \tag{a}$$

$$\exp\left(-\frac{t}{8}\right) \leq \exp\left(-\frac{t}{K}\log\frac{\mathbb{E}[Z]+t}{\mathbb{E}[Z]}\right) \qquad \forall\, 4\mathbb{E}[Z] < t < K\mathbb{E}[Z] \tag{b}$$

$$\exp\left(-\frac{t}{8}\log\frac{t}{\mathbb{E}[Z]} + t\right) \leq \exp\left(-\frac{t-\mathbb{E}[Z]}{K}\log\frac{t}{\mathbb{E}[Z]}\right) \qquad \forall\, t \geq K\mathbb{E}[Z]. \tag{c}$$

We check the three constraints to give a constant lower bound for K. For condition (a): take the natural logarithm of both sides and divide by t. Substituting $\frac{t}{\mathbb{E}[Z]} = u$, the constraint can be rewritten as:

$$\frac{u}{32} \geq \frac{1}{K}\log(u+1) \qquad \forall\, u \in (0,4].$$

That is $K \geq \frac{32\log(u+1)}{u}$ for $u \in (0,4]$. Let $f(u) = \frac{32\log(u+1)}{u}$. Differentiate,

$$f'(u) = \frac{32\left(\frac{u}{u+1} - \log(u+1)\right)}{u^2}.$$

Let $g(u) = \frac{u}{u+1} - \log(u+1)$. Note that $g(0) = 0$.

$$g'(u) = \frac{1}{(u+1)^2} - \frac{1}{u+1} < 0 \qquad \forall\, u > 0.$$

Therefore, $f(u)$ is decreasing in u. The constraint becomes,

$$K \geq \lim_{u\to 0^+} f(u) = \lim_{u\to 0^+} \frac{32\cdot\frac{1}{u+1}}{1} = 32.$$

The first equality follows from L'Hôpital's rule.

For constraint (c), take the natural logarithm of both sides and divide by t. Setting $u = \frac{t}{\mathbb{E}[Z]}$, the constraint can be rewritten as:

$$u\left(\log u - 8\right) \geq \frac{8(u-1)}{K}\log u \qquad \forall\, u \geq K.$$

[5] Condition (c) is derived from from inequality (iii), whose left-hand side is $\Pr\left[Z \geq t\right]$; replacing t by $(\mathbb{E}[Z] + t)$ yields the desired form.

Let $h(u) = Ku(\log u - 8) - 8(u-1)\log u$. We need $h(u) \geq 0$ for all $u \geq K$. Take the derivative,

$$\begin{aligned} h'(u) &= K\log u + K - 8K - 8\log u - 8 + \frac{8}{u} \\ &> (K-8)\log u - 7K - 8. \end{aligned}$$

Take $K = e^{8.1}$: for $u \geq K$, $h'(u) > 0$, $h(u)$ is increasing in u. $h(K) = 0.1e^{16.2} - 64.8e^{8.1} + 64.8 \approx 8.72 \times 10^5 > 0$. Therefore, taking $K = e^{8.1}$ suffices for the third condition.

We next consider condition (b) for $t \in (4\mathbb{E}[Z], K\mathbb{E}[Z])$. Take the natural logarithm of both sides and divide by t. Let $u = \frac{t}{\mathbb{E}[Z]}$, the constraint can be rewritten as:

$$K \geq 8\log(u+1) \qquad \forall\, u \in (4, K).$$

As $8\log(u+1)$ is increasing in u, $K \geq 8\log(K+1)$. For $K = e^{8.1}$, $K \geq 8\log(e^{8.1}+1) \approx 64.8$.

Combining these cases, taking $K = e^{8.1} \approx 3295$ suffices for the upper bound.

Rewrite (1) in terms of λ and x, we get:

$$\Pr\left[Z \geq x \cdot \lambda\right] \leq \exp\left(-\frac{\lambda}{K} \cdot (x-1)\log x\right).$$

Take $K = 3295$, for any $x > 1$,

$$\Pr\left[Z \geq x \cdot \lambda\right] \leq \exp\left(-\frac{\lambda}{3295} \cdot (x-1)\log x\right).$$

For $x \geq 2$, $\frac{1}{2}\lambda x \log x \leq \lambda(x-1)\log x$. Therefore,

$$\Pr\left[Z \geq x \cdot \lambda\right] \leq \exp\left(-\frac{1}{6590}\lambda \cdot x\log x\right). \qquad \square$$