

Prophet Inequalities with Limited Information

Pablo D. Azar

MIT, CSAIL

Cambridge, MA 02139, USA

azar@csail.mit.edu

Robert Kleinberg

Cornell University

Ithaca, NY 14853, USA

rdk@cs.cornell.edu

S. Matthew Weinberg

MIT, CSAIL

Cambridge, MA 02139, USA

smweinberg@csail.mit.edu

October 9, 2013

Abstract

In the classical prophet inequality, a gambler observes a sequence of stochastic rewards V_1, \dots, V_n and must decide, for each reward V_i , whether to keep it and stop the game or to forfeit the reward forever and reveal the next value V_i . The gambler's goal is to obtain a constant fraction of the expected reward that the optimal offline algorithm would get. Recently, prophet inequalities have been generalized to settings where the gambler can choose k items, and, more generally, where he can choose any independent set in a matroid. However, all the existing algorithms require the gambler to know the distribution from which the rewards V_1, \dots, V_n are drawn.

The assumption that the gambler knows the distribution from which V_1, \dots, V_n are drawn is very strong. Instead, we work with the much simpler assumption that the gambler only knows a few samples from this distribution. We construct the first single-sample prophet inequalities for many settings of interest, whose guarantees all match the best possible asymptotically, *even with full knowledge of the distribution*. Specifically, we provide a novel single-sample algorithm when the gambler can choose any k elements whose analysis is based on random walks with limited correlation. In addition, we provide a black-box method for converting specific types of solutions to the related *secretary problem* to single-

sample prophet inequalities, and apply it to several existing algorithms. Finally, we provide a constant-sample prophet inequality for constant-degree bipartite matchings.

In addition, we apply these results to design the first posted-price and multi-dimensional auction mechanisms with limited information in settings with asymmetric bidders. Connections between prophet inequalities and posted-price mechanisms are already known, but applying the existing framework requires knowledge of the underlying distributions, as well as the so-called “virtual values” even when the underlying prophet inequalities do not. We therefore provide an extension of this framework that bypasses virtual values altogether, allowing our mechanisms to take full advantage of the limited information required by our new prophet inequalities.

1 Introduction

Prophet inequalities are a fundamental tool in optimal stopping theory. In the classical prophet inequality, a gambler observes a sequence V_1, \dots, V_n of n rewards sampled independently from known distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$. After seeing the i^{th} reward, the gambler has two options: he can stop the game and keep reward V_i , or he can continue the game. If he chooses to continue the game, he forfeits reward V_i forever, and is shown the next reward V_{i+1} . The gam-

bler’s goal is to obtain an expected reward competitive with the best offline algorithm, represented by a *prophet* who can observe the values of all the variables V_1, \dots, V_n before making her selection. A seminal result of Krengel, Sucheston and Garling [20, 21] states a strategy for the gambler exists that guarantees an expected reward of at least half of the prophet’s. Recently there has been a renewed interest in prophet inequalities, generalizing the problem to settings where the prophet and gambler can choose any k out of the n presented items [1, 6], and more generally to settings where the prophet and gambler can choose any independent set in a matroid or matroid intersection environment [18]. However, all existing results require the gambler to know $\mathcal{D}_1, \dots, \mathcal{D}_n$.

We improve on the existing literature by giving the first prophet inequalities with limited information. More concretely, we show how the gambler can obtain a constant factor of the prophet’s expected reward, even when he only knows a single sample from each \mathcal{D}_i .¹ This approach is robust, and guarantees—in expectation over the observed sample and the realized state of the world—a simultaneous approximation to the prophet’s reward for all possible distributions \mathcal{D} . Our work is inspired by recent literature on mechanism design [11, 14] and on ad auctions [9, 10] which explores how to obtain approximately optimal revenue with limited information about an existing distribution of bidders’ values. Our work applies this limited information framework beyond auctions. Indeed, while our work has applications in online and multi-dimensional mechanism design, it also applies to the setting of optimal stopping problems.

1.1 Our Results. In the list below, we summarize our new prophet inequalities. We remark that, for all the results below, the weights of the items we are choosing online are revealed in an adversarial order (where the adversary observes the values in advance before deciding how to order the elements) and where the online algorithm has no knowledge of the distribution \mathcal{D} from which the values are drawn except for a single sample. The only exception is our result for constant degree bipartite matching environments, where the online algorithm requires a constant number samples from the distribution \mathcal{D} .

- **k-Uniform Matroids.** A $1 - O(\frac{1}{\sqrt{k}})$ -competitive single-sample prophet inequality for k -uniform matroids. This competitive ratio is

asymptotically optimal as a function of k .

- **Transversal Matroids.** A $\frac{1}{16}$ -competitive single-sample prophet inequality.
- **Graphic Matroids.** A $\frac{1}{8}$ -competitive single-sample prophet inequality.
- **Laminar Matroids.** A $\frac{1}{12\sqrt{3}}$ -competitive single-sample prophet inequality.
- **Constant Degree Bipartite Matchings.** A $\frac{1}{6.75}$ -competitive constant-sample prophet inequality.

1.2 New Results in Mechanism Design. Myerson’s seminal paper [22] shows how to construct the revenue-optimal single-item auction when each buyer’s valuation is drawn independently from a known distribution. Starting with work by Hartline and Roughgarden [14] and by Dhangwatnotai, Roughgarden and Yan [11], some recent attention has been focused on designing auctions that guarantee a constant-factor approximation to Myerson’s optimal auction, even when the seller has limited information about these distributions. However, prior to this work, progress on this front has been mostly limited to single-dimensional settings.

We apply our new prophet inequalities to construct the first truthful and approximately optimal auctions for certain multi-dimensional settings that use limited information. It is worth noting that we cannot simply plug our new prophet inequalities into the existing machinery of Chawla, Hartline, Malec and Sivan [6] to obtain these results, as their machinery requires full knowledge of the distributions, as well as the ability to compute “virtual values.”² Our main contribution on this front is an extension of their framework that allows us to analyze the expected virtual surplus of our mechanisms without ever learning the virtual values.

It is also worth noting that our results apply whenever the buyers’ valuations are drawn either from identical regular distributions, or from *distinct* monotone hazard rate (MHR) distributions. In contrast, all existing multi-dimensional mechanisms with limited information work only when bidders have identical distributions [8, 23]. More concretely, our results will apply to the following settings:

¹As described below, one of our results requires a constant number of samples.

²Virtual values were introduced in Myerson’s seminal paper and are known to have strong connections to revenue maximization. The virtual value of a bidder with value v sampled from distribution D_i with CDF F and PDF f is $v - \frac{1-F(v)}{f(v)}$.

- **Sequential Posted Price Mechanisms (SPMs).**

In this setting, a seller offers a service to buyers who arrive online, in an order chosen by the seller. Each buyer i has a value v_i for receiving service, and is offered a take-it-or-leave-it price p_i . The seller may face constraints on which buyers can be served simultaneously, such as matroid constraints (that is, a set S of buyers can be simultaneously allocated service if and only if S is an independent set in a matroid). We show a new approximately optimal single-sample SPM for *all* matroid settings. This improves over previously known SPMs, which applied to k -uniform settings and required bidder distributions to be identical [25].

- **Order-Oblivious Posted Price Mechanisms (OPMs) for multi-dimensional environments.**

Order-Oblivious Posted Price mechanisms are approximately optimal SPMs, whose revenue guarantee holds regardless of the order in which bidders arrive (i.e. the seller may no longer choose the order in which bidders arrive), and are known to imply truthful mechanisms for corresponding multi-dimensional settings [6, 18]. We construct single-sample OPMs for all environments for which we construct single-sample prophet inequalities, including graphic, laminar, transversal and partition matroids, as well as (constant-sample OPMs for) constant-degree bipartite matching settings. To the best of our knowledge, our mechanisms are the first OPMs that do not require full knowledge of the distribution or the ability to compute virtual values.

- **Multi-Dimensional Matching environments.**

In these environments, there are n buyers and m goods, and no buyer can be allocated more than one good, or good be allocated to more than one buyer. This induces a bipartite graph between buyers and goods, with an edge (i, j) present if $v_{ij} > 0$. When this graph has maximum degree d (no buyer has value for more than d goods, and no good is valued by more than d buyers), we give a mechanism that uses $d^2 + 1$ samples. We note this is the first limited-sample mechanism for matchings when bidders are asymmetric. In the case of i.i.d. regular distributions, Roughgarden, Talgam-Cohen and Yan [23] and Devanur, Hartline, Karlin and Nguyen [8] give limited-information mechanisms for general matching settings.

1.3 Our techniques.

1. **Reduction from existing secretary problems.**

In section 3, we give a black-box reduction that obtains single-sample prophet inequalities from existing order-oblivious³ algorithms for the secretary problem.⁴ This allows us to obtain prophet inequalities for transversal, graphic and laminar matroids based on corresponding secretary algorithms given by Dimitrov and Plaxton [12], Korula and Pal [19] and Jaillet, Zoto and Zenklusen [16]. However, not all algorithms for the secretary problem are order-oblivious. In particular, Kleinberg’s algorithm for k -uniform matroids [17] is not order-oblivious, and neither is Korula and Pal’s algorithm for matchings [19].

2. **Sufficient thresholds with limited samples.**

In section 5, we give a constant-sample prophet inequality for constant-degree bipartite matching settings. A prophet would accept element i only if it were above a certain threshold, determined by the values of all other items. Since the elements arrive one by one, we cannot compute these thresholds, and with a constant number of samples, we cannot even estimate them accurately. Instead, we use our samples to set sufficient thresholds that do not necessarily bear any relation to the prophet’s thresholds.

3. **Analysis of correlated random walks**

The best known secretary algorithms [17] and full-information prophet inequalities [1] for k -uniform matroids both guarantee a $1 - O(\frac{1}{\sqrt{k}})$ competitive ratio. In order to asymptotically match this competitive ratio, we give a new algorithm in section 4, whose analysis models the drawing of “samples” or “values” as positive and negative steps in a random walk. This random walk is correlated because for every “sample” s_i that we observe (which makes the walk move upward), there is a corresponding “value” v_i which will make the walk move “downward”. By estimating the expected height of this correlated random walk, we are able to guarantee that each of the top k values are selected by our online algorithm with probability $1 - O(\frac{1}{\sqrt{k}})$.

³We define order-oblivious algorithms in Section 3.

⁴In the secretary problem, the value of weights can be arbitrary, but the elements are revealed in a random order. In the prophet inequality problem, the value of weights come from distributions, but the order in which items are presented can be arbitrary.

There are many settings (arbitrary matroids, the intersection of any k arbitrary matroids) for which full-information prophet inequalities exist but limited-information prophet inequalities don't. We hope that our techniques can help develop such new limited-information algorithms in the future.

2 Preliminaries

Environments and Offline Selection Problems. An environment $\mathcal{I} = (\mathcal{U}, \mathcal{J})$ is given by a universe of elements $\mathcal{U} = \{1, \dots, n\}$ and a collection $\mathcal{J} \subset 2^{\mathcal{U}}$ of *feasible subsets* of \mathcal{U} . An algorithm \mathcal{A} for the *offline selection problem on \mathcal{I}* takes as input a vector of positive weights $v = (v_1, \dots, v_n)$ for elements of \mathcal{U} and outputs the independent set $MAX(v) = \operatorname{argmax}_{S \in \mathcal{J}} \sum_{i \in S} v_i$ with the maximum weight. We denote by $OPT(v) = \sum_{i \in MAX(v)} v_i$ the weight of this maximum independent set.

Online Selection Problems. Given an environment $\mathcal{I} = (\mathcal{U}, \mathcal{J})$, an algorithm \mathcal{A} for the *online selection problem* takes as *online* input a vector of values $v = (v_1, \dots, v_n)$ in some order $(v_{i_1}, \dots, v_{i_n})$ (this order will be specified below). The algorithm must maintain a set A of accepted elements, and element $i_j \in \mathcal{U}$ must be either accepted when its value v_{i_j} is revealed, or rejected forever before moving on to the next item i_{j+1} . At all times, the set A of accepted items must be an independent set (that is, $A \in \mathcal{J}$). For convenience of notation, we define $A^*(v) = A(v_{i_1}, \dots, v_{i_n})$ to be the final set of items accepted by \mathcal{A} , and note that $A^*(v)$ *depends on the order in which the items v_{i_1}, \dots, v_{i_n} are revealed*.

Prophet Inequalities. Given an environment \mathcal{I} with universe set $\mathcal{U} = \{1, \dots, n\}$, let $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$ be a product distribution over $\mathbb{R}_{\geq 0}^n$.⁵ Let $v = (v_1, \dots, v_n)$ be drawn from \mathcal{D} . We say that an algorithm \mathcal{A} for the online selection problem induces a *prophet inequality* with competitive ratio α for environment \mathcal{I} if

$$\mathbb{E}_{v \leftarrow \mathcal{D}} \left[\sum_{i \in A^*(v)} v_i \right] \geq \alpha \cdot \mathbb{E}_{v \leftarrow \mathcal{D}} [OPT(v)]$$

where the expectations are taken with respect to the random choice of v and the random coin tosses of \mathcal{A} . The above inequality holds *regardless* of the order in which the elements v_{i_1}, \dots, v_{i_n} are revealed. We

remark that this is a stronger property than that guaranteed by the prophet inequalities in previous papers [18], where the adversary had to choose which element i_j to reveal at time j using only knowledge of the items and values $(i_1, v_{i_1}), \dots, (i_{j-1}, v_{i_{j-1}})$ revealed up to time $j - 1$.

Limited-Information Prophet Inequalities.

In order to guarantee a prophet inequality with a constant competitive ratio, the online algorithm \mathcal{A} must have some information about the distributions $\mathcal{D}_1, \dots, \mathcal{D}_n$ from which the values are drawn. We say that \mathcal{A} is a constant-sample prophet inequality if it has access only to a constant number of samples $s^1 = (s_1^1, \dots, s_n^1), \dots, s^d = (s_1^d, \dots, s_n^d)$, each drawn from the joint distribution \mathcal{D} . When \mathcal{A} is constant-sample, its expected reward $\mathbb{E}_{v, s^1, \dots, s^d} [\sum_{i \in A^*(s^1, \dots, s^d; v)} v_i]$ is computed over the randomness in the vector of values v , the random samples s^1, \dots, s^d and the random coin tosses of the algorithm. We remark that, except for our results for matching environments, all our limited-information prophet inequalities use only one sample $s = (s_1, \dots, s_n)$ from the joint distribution \mathcal{D} .

Our Constraints. We can give different feasibility constraints by placing different structure on \mathcal{J} . We consider constraints that are matroids, specific types of matroids, or bipartite matchings. We refer the reader who is not familiar with these constraints to Appendix A for a formal definition of each setting we consider.

Secretary Problems. The secretary problem for an environment $(\mathcal{U}, \mathcal{J})$ [5] is an online selection problem where the item values v_1, \dots, v_n can be adversarially chosen, and they are revealed to the online algorithm in a *random order*. This is incomparable in terms of hardness with the prophet inequality setting described above, where the values are random variables, and they are presented in an adversarial order. We remark that there exist competitive algorithms for the secretary problem when \mathcal{J} is a uniform matroid [17], a laminar matroid [16], graphic matroid [19], a transversal matroid [12], or a bipartite matching [19]. If the online algorithm can choose the order in which the weights are revealed, then there exists a competitive algorithm for general matroids [16]. If the weight for item i is not completely adversarial, but is instead chosen randomly without replacement from a list (w_1, \dots, w_n) , then there also exists a competitive algorithm for matroids [24], even when the order in which the items is revealed is adversarially chosen [13].

⁵We remark that the assumption that the rewards V_1, \dots, V_n are independent is somewhat necessary if we want a constant competitive ratio. Hill and Kertz [15] show that if we allow arbitrary correlation between the rewards, then the gambler cannot obtain more than a $\frac{1}{n}$ fraction of the gambler's expected reward.

3 Prophet Inequalities from Secretary Algorithms

In this section, we provide a formal black-box method to convert specific kinds of solutions to the secretary problem to single-sample prophet inequalities. More formally, our reduction will work for *order-oblivious algorithms*, which we define as follows.

DEFINITION 1. *We say that an algorithm \mathcal{S} for the secretary problem (together with its corresponding analysis) is **order-oblivious** if, on a randomly ordered input vector $(v_{i_1}, \dots, v_{i_n})$:*

1. (algorithm) \mathcal{S} sets a (possibly random) number k , observes without accepting the first k values $S = \{v_{i_1}, \dots, v_{i_k}\}$, and uses information from S to choose elements from $V = \{v_{i_{k+1}}, \dots, v_{i_n}\}$.
2. (analysis) \mathcal{S} maintains its competitive ratio even if the elements from V are revealed in any (possibly adversarial) order. In other words, the analysis does not fully exploit the randomness in the arrival of elements, it just requires that the elements from S arrive before the elements of V , and that the elements of S are the first k items in a random permutation of values.

We argue in appendix C of the full version of the paper that existing algorithms for graphic, transversal and laminar matroids are order-oblivious [4]. Furthermore, Oveis Gharan and Vondrak [13]’s matroid secretary algorithm for the random assignment model is also order-oblivious (a fact that they claim in their paper). Combined with Theorem 3.1 below, this gives us single-sample prophet inequalities for graphic, transversal and laminar matroids, as well as arbitrary matroids when each \mathcal{D}_i is identical. This is stated formally in Corollary 3.1.

We now show how to construct an algorithm \mathcal{P} for the limited-information prophet problem given an order-oblivious algorithm \mathcal{S} for the secretary problem. Recall that the algorithm \mathcal{P} takes as offline input a vector $s = (s_1, \dots, s_n)$ of samples drawn from a distribution \mathcal{D} , and takes as online input a vector v also drawn from \mathcal{D} , and whose individual components are provided in an adversarial order.

$\mathcal{P}_{\mathcal{S}}(s_1, \dots, s_n; v_{i_1}, \dots, v_{i_n})$

Offline Stage

1. Let k be the number of elements that \mathcal{S} observes before it starts accepting elements (i.e., $k = |S|$).
2. Let s_{j_1}, \dots, s_{j_n} be a random permutation of $s = (s_1, \dots, s_n)$. Pass s_{j_1}, \dots, s_{j_k} as the first k inputs to \mathcal{S} .

Online Stage

3. For each index $i \in \{i_1, \dots, i_n\}$:
 - a. If $i \in \{j_1, \dots, j_k\}$, then index i has already been processed as a “sample”. Ignore it and continue.
 - b. If $i \in \{j_{j+1}, \dots, j_n\}$, then pass the value v_i to algorithm \mathcal{S} , and accept i if and only if \mathcal{S} accepts i .

THEOREM 3.1. *If \mathcal{S} is an order-oblivious algorithm for the secretary problem with competitive ratio α , then $\mathcal{P}_{\mathcal{S}}$ is a single-sample prophet inequality with competitive ratio α .*

We give the proof for Theorem 3.1 in appendix C. The proof that $\mathcal{P}_{\mathcal{S}}$ inherits the competitive ratio of \mathcal{S} uses the fact that the joint distribution of values associated to the items in our simulation of \mathcal{S} is exactly the same as the true value distribution \mathcal{D} . Note that our single-sample algorithm $\mathcal{P}_{\mathcal{S}}$ does not use any sampled values for elements in the set V . This is important, as we can then reuse the samples for items in V for other purposes, such as setting reserve prices in auctions, as we will see in Section 6.

COROLLARY 3.1.

1. *For graphic matroids, there exists a $\frac{1}{8}$ -competitive single-sample prophet inequality based on the secretary algorithm of Korula and Pal [19]*
2. *For transversal matroids, there exists a $\frac{1}{16}$ -competitive single-sample prophet inequality based on the secretary algorithm of Dimitrov and Plaxton [12].*
3. *For laminar matroids, there exists a $\frac{1}{12\sqrt{3}}$ -competitive single-sample prophet inequality based on the secretary algorithm of Jaillet, Soto, and Zenklusen [16].*
4. *For general matroid settings, when weights are drawn from identical and independent distributions, there exists a $\frac{1-\frac{1}{e}}{20}$ -competitive single-sample prophet inequality based on the secretary*

algorithm of Oveis Gharan and Vondrak for matroids in the random assignment model [13].⁶

4 Single-Sample Prophet Inequalities for k -Uniform Matroids

Recently, Alaei [1] gave a full-information prophet inequality that is $(1 - \frac{1}{\sqrt{k+3}})$ -competitive, which is asymptotically optimal. This raises the question of whether there also exists a $1 - O(\frac{1}{\sqrt{k}})$ competitive single-sample prophet inequality for k -uniform matroids. Since the corresponding algorithm (of Kleinberg, which obtains a competitive ratio of $1 - O(\frac{1}{\sqrt{k}})$) for the secretary problem is *not order-oblivious*, we cannot use our reduction from the previous section. Instead, we develop a new algorithm, and show that we can guarantee a $1 - O(\frac{1}{\sqrt{k}})$ competitive ratio by giving a new analysis for prophet inequalities based on correlated random walks. We note also that our algorithm is comparatively simpler than previous algorithms.

4.1 The Rehearsal Algorithm. We now describe our algorithm, which we call the *Rehearsal Algorithm*. The algorithm needs to fill k slots, and each slot i is associated with a threshold T_i (which is defined below). Each slot i can only be filled by a value that is above the threshold T_i , and can only be filled once. Each observed value can only fill a single slot. When we see an element that can fill at least one available slot, we fill the slot with the highest threshold. When we see an element that cannot fill any available slots, we reject it.

Intuitively, one might try to set the i^{th} threshold T_i to the i^{th} largest sample. This algorithm doesn't quite work, but a small modification suffices: instead, we set the first $k - 2\sqrt{k}$ thresholds equal to the top $k - 2\sqrt{k}$ samples, then set the remaining $2\sqrt{k}$ thresholds equal to the $k - 2\sqrt{k}^{\text{th}}$ highest sample (essentially repeating this sample $2\sqrt{k}$ times as a threshold). This is necessary in order for the probability of selecting the highest-value items to be sufficiently close to 1. (See Lemmas 10 and 11 in appendix G of the full version of the paper [4].)

In appendix G, we prove the following theorem. As we mentioned above, the proof may be interesting in its own right for its use of correlated random walks to analyze prophet inequalities. We defer part of the

proof to the last appendix, and the remainder to the complete version of the paper [4].

THEOREM 4.1. *Let $\mathcal{I} = (\mathcal{U}, \mathcal{J})$ be a k -uniform matroid. The rehearsal algorithm is a single-sample prophet inequality with a competitive ratio of $1 - O(\frac{1}{\sqrt{k}})$.*

Rehearsal($s_1, \dots, s_n; v_{i_1}, \dots, v_{i_n}$)

1. Offline Phase

- 1.a Let $s^{(1)} > \dots > s^{(n)}$ be the observed samples in decreasing order.
- 1.b For $j \in \{1, \dots, k - 2\sqrt{k}\}$ set $T_j = s^{(j)}$.
- 1.c For $k - 2\sqrt{k} < j \leq k$, set $T_j = T_{k-2\sqrt{k}} = s^{(k-2\sqrt{k})}$.

2. Online Phase

Initialize $S = \{1, \dots, k\}$ as the set of available slots. For $j \in \{1, \dots, n\}$:

- 2.a Let v_{i_j} be the value of the j^{th} revealed item. Let α be an index such that $T_{\alpha-1} > v_{i_j} > T_\alpha$.
- 2.b Let $S \cap \{\alpha, \alpha + 1, \dots, k\}$ be the set of slots that have not been filled, and that could be filled by v_{i_j} . Let $m = \min S \cap \{\alpha, \dots, k\}$. This is the first slot that could be occupied by v_{i_j} .
- 2.c If $S \cap \{\alpha, \dots, k\}$ is empty, reject v_{i_j} .
- 2.d If $S \cap \{\alpha, \dots, k\}$ is not empty, accept v_{i_j} and update $S \leftarrow S - m$.

5 Bipartite Matching Environments

Before we give our algorithm, we establish some notation to make our exposition clearer.

Edge Indices. Let $G = (L \cup R, E)$ be a degree- d bipartite graph, and let $e = (\ell, r)$ be an edge in this graph. There are at most d edges incident to ℓ , and we can assign them an arbitrary order $\{0, 1, \dots, d-1\}$. Analogously, we can assign the edges incident to r an order $\{0, 1, \dots, d-1\}$. Without loss of generality, assume that e is the j^{th} edge incident to ℓ , and the k^{th} edge incident to r . Define $\text{Index}(e) = 1 + j + d \cdot k$. This index function has two key properties

1. $\text{Index}(e) \in \{1, \dots, d^2\}$
2. If e, e' share a vertex, then $\text{Index}(e) \neq \text{Index}(e')$.

Edge Thresholds. Given an vector of values $v = (v_1, \dots, v_{|E|})$ and an edge $e \in E$ define $x_e(v)$ to be 1 if e is in the maximum weight matching

⁶We note that a similar result for general matroids under i.i.d. distributions was already proved by two of the authors [18]. Their result did not emphasize the single-sample nature of the algorithm.

when the weights are given by v , and 0 if e is not in this maximum weight matching.⁷ Note that x_e is a deterministic increasing function of v_e when all the other weights v_{-e} are fixed. Thus, there exists a threshold function that takes as input the weight v_{-e} of all the other edges, and outputs the lowest weight that edge e needs to have to be in the maximum weight matching.

$$T_e(v_{-e}) = \inf\{v_e : x_e(v_e, v_{-e}) = 1\}.$$

Our algorithm. We construct an algorithm $\mathcal{P}_{Matching}$ that takes as offline input a collection $s^1 = (s_1^1, \dots, s_n^1), \dots, s^{d^2} = (s_1^{d^2}, \dots, s_n^{d^2})$ of samples, and as online input a vector v of values $(v_{i_1}, \dots, v_{i_n})$. It proceeds as follows:

$\mathcal{P}_{Matching}(s^1, \dots, s^{d^2}; v_{i_1}, \dots, v_{i_{|E|}})$

Offline Phase:

- 1 For each edge e , compute $i = \text{Index}(e)$.
- 2 For each edge e , set its corresponding sample to be s^i . Set its price to be $p_e = T_e(s_{-e}^i)$.

Online Phase:

- 3 Initialize a set A of accepted items to \emptyset .
- 4 For $e \in \{i_1, \dots, i_{|E|}\}$:
 - 4.a Flip a coin $c_e = \begin{cases} 1 & \text{with probability } \frac{1}{3} \\ 0 & \text{with probability } \frac{2}{3} \end{cases}$
 - 4.b If $c_e = 0$, discard edge e and move on to the next edge.
 - 4.c If $c_e = 1$, accept edge e if and only if $v_e > p_e$ and $A \cup \{e\}$ is a matching in the bipartite graph G .

THEOREM 5.1. *The algorithm $\mathcal{P}_{Matching}$ guarantees a $\frac{1}{6.75}$ competitive ratio for environments \mathcal{I} that are degree- d bipartite matchings.*

We present the proof of this theorem in appendix D. We remark that, for general bipartite matchings (and, more generally, for intersections of two partition matroids), an analogous algorithm with n samples obtains the same competitive ratio.

⁷We can set a tie-breaking rule so the maximum weight matching is unique.

Even though our algorithm is not an auction, it is inspired by an approximately optimal auction for bipartite matching environments given by Chawla, Hartline, Malec and Sivan [6]. Their auction requires knowledge of the distribution from which edge weights are drawn, and requires knowledge of the virtual values associated with these distributions, which can be estimated in their paper with $n^4 \log n$ samples. In contrast, our algorithm only requires a constant number of samples and approximately maximizes the weight of the matching (as opposed to its virtual weight).

6 Mechanism Design with Limited Information

In this section, we give new limited-information auctions for online and multi-dimensional mechanism design. In particular, we improve over existing literature as follows

- **Single-Dimensional SPMs with Non-Identical Distributions** We give the first limited-information sequential posted price mechanisms (SPMs) for matroids and constant-degree bipartite matching settings. Our results guarantee a constant approximation to revenue when distributions are identical and regular, or when distributions are distinct and MHR. The best previously known limited-information SPM [25] applies only to k -uniform matroids and requires distributions to be i.i.d.
- **OPMs for Multidimensional Unit-Demand Mechanism Design** We give the first limited-information OPMs for partition, graphic, laminar, and transversal matroid settings, as well as constant-degree bipartite matchings. For bipartite matchings, there exist limited-information auctions that approximately maximize revenue when bidders have identical distributions [8] [23]. Our auction is the first that is approximately optimal for bidders with distinct distributions satisfying the monotone hazard rate condition.
- **A new reduction from welfare to revenue maximization** We give a new reduction from approximate welfare maximization to approximate revenue maximization for single-dimensional environments when buyers' preferences are identical and regular. This reduction generalizes the well know fact that the Vickrey Clarke Groves (VCG) auction with appropriate

reserves is approximately optimal for matroid environments [14, 11] to show that *any mechanism that approximately maximizes welfare* (not necessarily VCG) also approximately maximizes revenue when valuations are regular and i.i.d.

Before stating our results more formally, we establish some preliminaries and recall prior work on mechanism design.

6.1 Mechanism Design Preliminaries. To improve readability, some details are deferred to the appendix. Contained in Appendix B is a formal definition of a mechanism, posted-price mechanism, as well as the specific mechanism design problems we solve (called Bayesian Single-Dimensional Mechanism Design (BSMD) and Bayesian Multi-Dimensional Unit-Demand Mechanism Design (BMUMD) in [6]). Contained also is a brief list of facts related to mechanism design (such as the connection between revenue and virtual valuations). We include here the relevant related work necessary to understand our approach.

Mechanisms with Reserves. The idea of combining simple, welfare-optimizing mechanisms with revenue-optimizing reserve prices originated in [14]. In [14], the authors first remove every bidder who does not meet their reserve, and then run the welfare maximizing mechanism. This process was later dubbed an “eager” combination of mechanisms with reserves. The authors of [11] introduce a “lazy” combination of mechanisms with reserves that first runs the mechanism, and then removes all bidders who do not meet their reserve. In this work, we concern ourselves primarily with lazy reserves. When we refer to *monopoly reserves*, we mean setting the reserve price $\phi_i^{-1}(0)$ for each bidder i . When we refer to *sample reserves*, we mean setting a random reserve price $r_i \leftarrow \mathcal{D}_i$ for bidder i , that is drawn from the same distribution as \mathcal{D}_i .

A reduction from OPMs to multi-dimensional mechanism design. Chawla, Hartline, Malec and Sivan [6] show how to reduce designing (approximately) optimal multi-dimensional mechanisms to (approximately) solving a related single-dimensional problem in a specific way. Given an instance \mathcal{I} of a multi-dimensional mechanism design problem with n items and m buyers, they construct an analogous single-dimensional instance $\mathcal{I}^{\text{copies}}$ with nm buyers. That is, each buyer i in the original setting gets split into m buyers in $\mathcal{I}^{\text{copies}}$. The $(i, j)^{\text{th}}$ buyer in $\mathcal{I}^{\text{copies}}$ only values the $(i, j)^{\text{th}}$ good, and her valuation v_{ij} is drawn from the same distribution \mathcal{D}_{ij} as in the original setting. We use the following result

from [6]:

LEMMA 6.1. ([6]) *Let \mathcal{I} be an instance of the BMUMD, and let $\mathcal{I}^{\text{copies}}$ be its analogous single-dimensional environment. If there exists an OPM for $\mathcal{I}^{\text{copies}}$ that achieves an α -approximation to the optimal revenue, then there exists a truthful mechanism for \mathcal{I} that achieves an α -approximation to the optimal revenue.*⁸

6.2 From Prophet Inequalities to Mechanisms. Let $\mathcal{P}(v_{i_1}, \dots, v_{i_n})$ be a limited-information prophet inequality with a competitive ratio of α . All of the limited-information algorithms that we gave in the previous sections are monotonic in v , meaning that the higher a value v_i is, the higher the probability that our algorithms accept item i . This means that any of our limited-information algorithms induces a *limited-information online allocation rule* $x(v)$, and this allocation rule is monotonic. When each value corresponds to a different bidder (single-dimensional setting), this monotonic allocation rule implies a pricing rule $p(v)$ which makes the mechanism (x, p) truthful. This means that all our limited-information algorithms can be used to give truthful online mechanisms to maximize welfare. Furthermore, our mechanisms are posted price mechanisms. This is because when we need to decide whether to accept bidder i or not, the decision to accept depends only on the set A of already accepted bidders and on the samples that we have from \mathcal{D} . If \mathcal{P} obtains a competitive ratio of α , we have $\mathbb{E}_v[x_i(v) \cdot v] \geq \alpha \mathbb{E}_v[\text{OPT}(v)]$. Thus, our prophet inequalities give sequential posted price mechanisms that approximately maximize welfare in single-dimensional settings.

6.3 From Welfare to Revenue: The I.I.D. Case. At this point, we have proven prophet inequalities and turned them into posted-price mechanisms with good welfare guarantees, but have said nothing about revenue. We show in this section how to guarantee a good revenue approximation given a guarantee for a good approximation to welfare. We again

⁸Formally, they show that there exists a truthful mechanism for \mathcal{I} that obtains an α -approximation to the optimal revenue achievable by any deterministic mechanism. It is shown in [7] that the optimal revenue achievable by any (possibly randomized) mechanism is at most five times larger than that of the optimal deterministic mechanism. So an OPM for $\mathcal{I}^{\text{copies}}$ that achieves an α -approximation to the optimal revenue implies the existence of a truthful mechanism for \mathcal{I} that achieves an $\alpha/5$ approximation to the optimal revenue of any (possibly randomized) mechanism.

note that this process is novel and cannot be replaced by simply plugging our prophet inequalities into the machinery of [6], which requires full knowledge of the distributions to apply, even if our prophet inequalities do not.

Comparison Based Mechanisms. Our reduction from welfare to revenue when distributions are i.i.d. requires the mechanism \mathbb{M} to be comparison-based. We define below what it means for a mechanism to be comparison based when it uses samples.

DEFINITION 2. Let $\mathbb{M}(v; s^1, \dots, s^d)$ be a mechanism for single-dimensional settings which depends on a vector of bids $v = (v_1, \dots, v_n) \leftarrow \mathcal{D}$ and also on a collection of samples $s^1 = (s_1^1, \dots, s_n^1), \dots, s^d = (s_1^d, \dots, s_n^d)$, each drawn from \mathcal{D} . Let x be the allocation rule associated with \mathbb{M} . We say that \mathbb{M} is comparison-based if the allocation rule $x(v_1, \dots, v_n, s_1^1, \dots, s_n^d)$ only depends on the relative order of its arguments, and not on their respective values.

The rehearsal algorithm and the algorithms derived from our black-box reduction in corollary 3.1 are all comparison-based. The only algorithm which is not comparison-based is our matching algorithm $\mathcal{P}_{\text{Matching}}$, which uses an algorithm for computing maximum weight matchings as a black-box to set a threshold price $p_e = \inf\{v_e : e \text{ is in a maximum weight matching when all other weights are } s_{-e}^{\text{Index}(e)}\}$. Since p_e cannot necessarily be computed by comparisons between the samples in $s^{\text{Index}(e)}$, $\mathcal{P}_{\text{Matching}}$ is not comparison-based. If we use the Greedy algorithm (which is comparison-based) instead of an optimal bipartite matching algorithm, then $\mathcal{P}_{\text{Matching}}$ becomes comparison-based but loses a factor of 2 in its competitive ratio.

THEOREM 6.1. Let \mathcal{J} be any downwards-closed set system, and let each \mathcal{D}_i be identical and regular. Let also \mathbb{M} be any single-dimensional comparison-based mechanism whose expected welfare competitive ratio is α . Then the mechanism that combines (either eagerly or lazily) \mathbb{M} with monopoly reserves has expected revenue competitive ratio α .

Of course, computing the monopoly reserves requires knowledge of the distributions. These reserves can be replaced by samples, using a result (stated in Appendix E) from Azar, Daskalakis, Micali and Weinberg [3].

COROLLARY 6.1. If \mathbb{M} is a single-dimensional mechanism that guarantees an α approximation to welfare

when distributions are i.i.d. and regular then \mathbb{M} combined with lazy sample reserves guarantees an $\frac{\alpha}{2}$ approximation to revenue and an $\frac{\alpha}{2}$ approximation to welfare.

6.4 From Welfare to Revenue: the MHR case. Since we want mechanisms that guarantee good revenue for asymmetric bidders, we also need a reduction from welfare maximization to revenue maximization when distributions are not identical. It is well known (and stated in Appendix E) that, when bidders' distributions have a monotone hazard rate, a single-dimensional mechanism that approximates welfare combined with lazy monopoly reserves gives a good approximation to revenue [11]. We emphasize that an analogous result is not known for multi-dimensional settings.⁹ Combining this with lemma E.1, we obtain the following corollary.

COROLLARY 6.2. If \mathbb{M} guarantees an α approximation to welfare and distributions are MHR then \mathbb{M} combined with lazy sample reserves guarantees an $\frac{\alpha}{2e}$ approximation to revenue and an $\frac{\alpha}{2}$ approximation to welfare.

6.5 Our mechanisms. Since our limited-information prophet inequalities guarantee a good approximation to welfare, we are now ready to give our approximately optimal multi-dimensional OPMs. Given an environment \mathcal{J} for which we have a limited-information online algorithm \mathcal{P} , our online mechanism for \mathcal{J} will behave as follows

1. Use \mathcal{P} to choose a set $W \in \mathcal{J}$ of winners that approximately maximizes welfare.
2. Use a sample $r \leftarrow \mathcal{D}$ as a vector of lazy reserves. Keep only winners $i \in W$ that satisfy $v_i \geq r_i$.

We note that for all the limited-information algorithms that we obtain from our black-box reduction in section 3, we only use the samples s_i corresponding to items i that are never chosen by our algorithms. The samples s_i corresponding to items i that are chosen by the algorithm (that is, corresponding to auction winners) are never used, and hence can be used to set reserve prices.

In Appendix E, we state two theorems for OPMs, one when distributions are i.i.d. and regular, and

⁹If such a result existed, then the VCG auction together with appropriate reserves would be a very simple, approximately optimal multidimensional mechanism when distributions are MHR.

the other one when distributions have a monotone hazard rate, but are not necessarily identical. We remark, as described above, that to apply our algorithm $\mathcal{P}_{\text{Matching}}$ in the i.i.d. regular setting, we need to modify it so it uses the greedy matching algorithm as a black-box. Theorems E.1 and E.2 are direct applications of Corollaries 6.1 and 6.2. Essentially, they state that we can obtain limited-information multi-dimensional for in any unit-demand setting for which we have a limited-information prophet inequality. If we start with a limited-information prophet inequality with competitive ratio α , then the corresponding mechanism for i.i.d. regular environments has revenue and welfare competitive ratio $\alpha/2$, and the corresponding mechanism for non-i.i.d. MHR environments has revenue competitive ratio $\alpha/2e$ and welfare competitive ratio $\alpha/2$. We separately state below our theorems as they apply to bipartite matching, which models settings where goods are matched to buyers.

THEOREM 6.2. *For the BMUMD problem on constant-degree bipartite matching settings, there exists a $\frac{1}{13.5e}$ -competitive auction using a constant number of samples when buyers' valuations are drawn from MHR distributions. A modification of this algorithm gives a $\frac{1}{27}$ -competitive limited-information auction when buyers' valuations are drawn from i.i.d. regular distributions.*

Finally, even for settings where we do not have limited-information prophet inequalities, we can leverage existing results to obtain improved mechanism design results. Jaillet, Soto and Zenklusen [16] give an algorithm for the matroid secretary problem in the *free order model*, where the algorithm gets to choose the order in which values are revealed. This model corresponds to a Sequential Posted Price Mechanism. We give in appendix F an improved analysis of Jaillet, Soto and Zenklusen, improving their competitive ratio from $\frac{1}{9}$ to $\frac{1}{4}$. We use this improved analysis to give the following SPM.

THEOREM 6.3. *Let \mathcal{J} be any matroid and let each \mathcal{D}_i be MHR. Then there exists a truthful SPM requiring only a single sample from \mathcal{D} that guarantees a revenue competitive ratio of $\frac{1}{8e}$ and a welfare competitive ratio of $\frac{1}{8}$. When the distributions \mathcal{D}_i are independent and regular, this algorithm obtains a revenue competitive ratio of $\frac{1}{8}$.*

References

- [1] Saeed Alaei. Bayesian Combinatorial Auctions: Expanding Single Buyer Mechanisms to Many Buyers. In *FOCS'11*, 2011.
- [2] A. Archer and É. Tardos. Truthful mechanisms for one-parameter agents. In *Foundations of Computer Science, 2001. Proceedings. 42nd IEEE Symposium on*, pages 482–491. IEEE, 2001.
- [3] Pablo Azar, Constantinos Daskalakis, Silvio Micali, and Matt Weinberg. Optimal and Efficient Parametric Auctions. In *SODA'13*, 2013.
- [4] Pablo D. Azar, Robert Kleinberg, and S. Matthew Weinberg. Prophet Inequalities with Limited Information. *ArXiv Report*, <http://arxiv.org/abs/0812.2277>, 20013.
- [5] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *SODA'07*, pages 434–443, 2007.
- [6] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-Parameter Mechanism Design and Sequential Posted Pricing. In *STOC'10*, 2010.
- [7] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The Power of Randomness in Bayesian Optimal Mechanism Design. In *EC'10*, 2010.
- [8] Nikhil Devanur, Jason Hartline, Anna Karlin, and Thach Nguyen. Prior-independent multi-parameter mechanism design. In *Internet and Network Economics*, pages 122–133. Springer, 2011.
- [9] Nikhil R Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher A Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 29–38. ACM, 2011.
- [10] Nikhil R Devanur, Balasubramanian Sivan, and Yossi Azar. Asymptotically optimal algorithm for stochastic adwords. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 388–404. ACM, 2012.
- [11] P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. In *EC'10*, pages 129–138. ACM, 2010.
- [12] Nedialko B. Dimitrov and C. Greg Plaxton. Competitive weighted matching in transversal matroids. *Algorithmica*, 62(1-2):333–348, 2012.
- [13] Shayan Oveis Gharan and Jan Vondrák. On variants of the matroid secretary problem. In *ESA'11*, pages 335–346, 2011.
- [14] J.D. Hartline and T. Roughgarden. Simple versus optimal mechanisms. In *EC'09*, pages 225–234, 2009.
- [15] Theodore P. Hill and Robert P. Kertz. Stop rule inequalities for uniformly bounded sequences of random variables. *Transactions of the American Mathematical Society*, 278(1):197–207, 1983.
- [16] Patrick Jaillet, José A. Soto, and Rico Zenklusen. Advances on matroid secretary prob-

- lems: Free order model and laminar case. *CoRR*, <http://arxiv.org/abs/1207.1333>, 2012.
- [17] R. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA'05*, pages 630–631, 2005.
 - [18] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *STOC'12*, pages 123–136, 2012.
 - [19] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *ICALP'09*, pages 508–520, 2009.
 - [20] U. Krengel and L. Sucheston. Semiamarts and finite values. *Bull. Amer. Math. Soc.*, 83:745–747, 1977.
 - [21] U. Krengel and L. Sucheston. On semiamarts, amarts and processes with finite value. *Advances in Prob.*, 4:197–266, 1978.
 - [22] Roger B. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
 - [23] Tim Roughgarden, Inbal Talgam-Cohen, and Qiqi Yan. Supply-limiting mechanisms. In *EC'12*, pages 844–861, 2012.
 - [24] José A. Soto. Matroid secretary problem in the random assignment model. In Dana Randall, editor, *SODA*, pages 1275–1284. SIAM, 2011.
 - [25] Qiqi Yan. *Prior-Independence: A new lens for mechanism design*. PhD thesis, Stanford Univ., 2012.

Appendix

A Matroids and Feasibility Constraints

- **Matroids.** \mathcal{J} is a matroid if and only if \mathcal{J} is downwards-closed¹⁰, contains \emptyset , and satisfies the augmentation property: for all $S, S' \in \mathcal{J}$ with $|S| > |S'|$, there exists some $x \in S - S'$ such that $S' \cup \{x\} \in \mathcal{J}$.
- **Uniform matroids of rank k .** A set $S \subset \mathcal{U}$ is in \mathcal{J} if and only if $|S| \leq k$.
- **Partition matroids.** Let B_1, \dots, B_ℓ be disjoint subsets of \mathcal{U} such that $\mathcal{U} = B_1 \cup \dots \cup B_\ell$. Associate a positive integer capacity c_i with each block B_i . A set $S \subset \mathcal{U}$ is in \mathcal{J} if and only if $|S \cap B_i| \leq c_i$ for every $i \in \{1, \dots, \ell\}$.
- **Laminar matroids.** Let $\mathcal{F} \in 2^{\mathcal{U}}$ be a *laminar family* of subsets of \mathcal{U} . \mathcal{F} is a laminar family iff for all $A, B \in \mathcal{F}$, we have $A \subseteq B$, $B \subseteq A$, or $A \cap B = \emptyset$. Associate also, for every set $A \in \mathcal{F}$,

¹⁰ \mathcal{J} is downward-closed if for any $S \in \mathcal{J}$ and any $T \subset S$, we have $T \in \mathcal{J}$.

a positive integer capacity c_A . A set $S \in \mathcal{J}$ if and only if $|S \cap A| \leq c_A$ for all $A \in \mathcal{F}$.

- **Graphic Matroids.** Let $G = (V, E)$ be a graph with vertex set V and edge set E . The universe \mathcal{U} of the set system is given by the set of edges E . A subset $S \subset E$ is in \mathcal{J} if and only if E induces no cycles in the graph G . In other words, a subset of edges is feasible if and only if it is a forest.
- **Transversal Matroids.** Let $G = (L \cup R, E)$ be a bipartite graph, with left-vertex set L and right-vertex set R . The universe \mathcal{U} of the set system is L , and a subset $S \subset L$ is in \mathcal{J} if and only if there is a matching in the graph G that matches every vertex of S to some vertex in R .
- **Bipartite Matchings.** Let $G = (L \cup R, E)$ be a bipartite graph and let $\mathcal{U} = E$. A set $S \subset E$ is independent if and only if it induces a matching in G . The bipartite matching has degree d if at most d edges are incident to any given vertex.

B Omitted Details From Section 6.1

Mechanisms. An instance of the Bayesian Single-Dimensional Mechanism Design problem (BSMD) is specified by a set system $(\mathcal{U}, \mathcal{J})$ and a product distribution $\mathcal{D} = \mathcal{D}_1 \times \dots \times \mathcal{D}_n$, where $n = |\mathcal{U}|$. Each element of \mathcal{U} represents a buyer, interested in obtaining a service. The collection $\mathcal{J} \subset 2^{\mathcal{U}}$ represents constraints on which buyers can receive service simultaneously. Each buyer i 's value for receiving service is a random variable v_i drawn from the distribution \mathcal{D}_i . A mechanism is said to be *dominant strategy truthful* if it is in each bidder's interest to report truthfully their value for each item, no matter what values are reported by the other bidders.

Formally, a mechanism is a pair of vector-valued functions (x, p) where, given a vector of bids $b = (b_1, \dots, b_n)$, $x_i(b)$ is player i 's probability of receiving service and $p_i(b)$ is player i 's expected payment. If bidder i 's true preferences are given by v_i , then her expected utility when the profile of reported bids is b is $U(v_i, b_i, b_{-i}) = x_i(b) \cdot v_i - p_i(b)$. A mechanism is dominant strategy truthful if for all v_i, b_i, b_{-i} , we have $U(v_i, v_i, b_{-i}) \geq U(v_i, b_i, b_{-i})$. We also require mechanisms to be individually rational. That is, $U(v_i, v_i, b_{-i}) \geq 0$ for all v_i, b_{-i} .

Allocation Rules Determine Prices [22, 2]. If $\mathbb{M} = (x, p)$ is a single-dimensional mechanism, then \mathbb{M} is truthful if and only if $x_i(b_i, b_{-i})$ is a monotonically increasing function of b_i (regardless of

the vector of other bids b_{-i}) and the price function satisfies

$$p_i(b_i) = b_i x_i(b_i) - \int_0^{b_i} x_i(z) dz$$

where the dependence on b_{-i} has been omitted. Thus, a monotonic allocation rule immediately specifies a truthful mechanism for single-dimensional settings.

Monotone Hazard Rate. The hazard rate function $h(v)$ of a distribution with cumulative distribution function $F(v)$ and probability density function $f(v)$ is defined as $h(v) = \frac{f(v)}{1-F(v)}$. The distribution has a monotone hazard rate (MHR) if $h(v)$ is increasing in v .

Virtual Valuations and Revenue. The virtual value of a bidder with value v sampled from a distribution with CDF F and PDF f is usually denoted by $\phi(v)$, and is equal to $v - \frac{1-F(v)}{f(v)}$. The distribution is called regular if $\phi(v)$ is monotonically increasing in v . It is immediate that all MHR distributions are regular. Myerson's famous theorem shows that in all single dimensional settings, the expected revenue of a truthful mechanism is exactly its expected virtual welfare. That is $\mathbb{E}_v[\sum_{i=1}^n p_i(v)] = \mathbb{E}_v[\sum_i x_i(v) \phi_i(v_i)]$.

Posted Price Mechanisms. A single-dimensional *sequential posted price mechanism* (SPM) serves bidders one at a time, offering each a price upon arrival that depends only on the previously observed bids and the underlying distributions. The mechanism maintains a set S of bidders who have been assigned service, initialized to be \emptyset , and adds each bidder to S iff their reported bid exceeds the price offered. An *order-oblivious posted price mechanism* (OPM) is a sequential posted price mechanism that maintains its approximation guarantee when the order is chosen by an adversary instead of the mechanism.¹¹

Bayesian Multi-parameter Unit-demand Mechanism Design (BMUMD). In a Bayesian multidimensional mechanism design problem, there are n buyers interested in m items for sale. Each buyer i has a value v_{ij} for receiving item j . Let $\mathcal{U} = [n] \times [m]$, with the element (i, j) denoting the event that bidder i receives item j . Further denote by \mathcal{J} the subsets of \mathcal{U} corresponding to feasible allocations. That is, a set $S \in \mathcal{J}$ iff it is feasible to simultaneously allocate item j to bidder i for all $(i, j) \in S$.

¹¹We remark that our definition matches that of [18], which extends the one given in [6].

A setting is said to be *unit-demand* if for all $S \in \mathcal{J}$, $(i, j) \in S \Rightarrow (i, j') \notin S$ for all $j \neq j'$ (i.e. it is infeasible to allocate any bidder more than one item). As in [6], we also assume that each v_{ij} is sampled independently from a known distribution \mathcal{D}_{ij} . As in the single dimensional setting, we seek to devise a truthful mechanism whose expected revenue is (approximately) optimal with respect to the maximum over all truthful mechanisms.

C Omitted Proofs and Algorithms from section 3

We now give a proof of theorem 3.1.

THEOREM 1. (THEOREM 3.1) *If \mathcal{S} is an order-oblivious algorithm for the secretary problem with competitive ratio α , then $\mathcal{P}_{\mathcal{S}}$ is a single-sample algorithm for the prophet problem with competitive ratio α .*

Proof. The algorithm $\mathcal{P}_{\mathcal{S}}$ first permutes the vector s of samples into a random permutation s_{j_1}, \dots, s_{j_n} and takes the first k elements s_{j_1}, \dots, s_{j_k} of this permutation and passes them as inputs to the secretary algorithm \mathcal{S} . After that, the secretary algorithm \mathcal{S} is passed all the inputs v_i where $i \notin \{j_1, \dots, j_k\}$ in an arbitrary order. Since \mathcal{S} is order-oblivious, the set it selects has a weight of at least $\alpha \cdot \text{OPT}(v)$, where $\text{OPT}(v) = \max_{A \in \mathcal{J}} \sum_{i \in A} v_i$. So if we let $f(v)$ denote the probability density function associated with the joint distribution \mathcal{D} , we have that our algorithm $\mathcal{P}_{\mathcal{S}}$ obtains expected reward of at least

$$\int_v f(v) \alpha \cdot \text{OPT}(v) dv$$

The prophet's expected reward is

$$\text{OPT} = \int_v f(v) \cdot \text{OPT}(v) dv$$

which immediately says that $\mathcal{P}_{\mathcal{S}}$ obtains competitive ratio α , completing the proof.

C.1 Existing order-oblivious secretary algorithms. In the complete version of the paper, we overview the existing secretary algorithms used in this paper and argue why they are order-oblivious. The arguments are omitted in this version for space considerations [4].

D Omitted Proofs from Section 5

THEOREM 2. (THEOREM 5.1) *The algorithm $\mathcal{P}_{\text{Matching}}$ guarantees a $\frac{1}{6.75}$ competitive ratio for environments \mathcal{I} that are degree- d bipartite matchings.*

Proof. Let $v = (v_1, \dots, v_{|E|})$ be drawn from a joint distribution $\mathcal{D}_1 \times \dots \times \mathcal{D}_{|E|}$. Recall that $T_e(v_{-e}) = \inf\{v_e : e \text{ is in the maximum weight matching, given all other weights are } v_{-e}\}$. Thus, the optimal offline algorithm selects a matching that has an expected weight of

$$\sum_{e=1}^{|E|} Pr_{v \leftarrow \mathcal{D}}[v_e \geq T_e(v_{-e})] \cdot \mathbb{E}_{v \leftarrow \mathcal{D}}[v_e | v_e \geq T_e(v_{-e})]$$

Let $q_e = Pr_{v \leftarrow \mathcal{D}}[v_e \geq T_e(v_{-e})]$ and recall that $p_e = T_e(s_{-e}^{Index(e)})$. Since $s^{Index(e)}$ is a sample drawn from the same distribution that v is drawn, we have that $Pr[v_e \geq p_e] = q_e$. We also have $\mathbb{E}[v_e | v_e \geq p_e] = \mathbb{E}_{v \leftarrow \mathcal{D}}[v_e | v_e \geq T_e(v_{-e})]$. So we can write the optimal reward as

$$OPT = \sum_e Pr[v_e \geq p_e] \mathbb{E}[v_e | v_e \geq p_e].$$

What is the reward obtained by our algorithm $\mathcal{P}_{Matching}$? Recall that $\mathcal{P}_{Matching}$ first sets a price p_e for each edge e . When the value v_e is revealed, the algorithm flips a coin c_e that is equal to one with probability $\frac{1}{3}$, and accepts e if and only if $c_e = 1$ and $v_e \geq p_e$ and $A \cup \{e\}$ is an independent set (i.e. a matching in the given bipartite graph). For each edge $e \in E$, define the following three random events

1. $c_e = 1$,
2. $v_e \geq p_e$,
3. $A \cup \{e\}$ is an independent set.

Call these events X_e, Y_e and Z_e , respectively.

Thus, the expected reward obtained by $\mathcal{P}_{Matching}$ is

$$W = \sum_e Pr[X_e \text{ and } Y_e \text{ and } Z_e] \cdot \mathbb{E}[v_e | X_e, Y_e, Z_e]$$

Clearly, X_e is independent from Y_e, Z_e and v_e . This means we can write

$$W = \sum_e \frac{1}{3} Pr[Y_e \text{ and } Z_e] \cdot \mathbb{E}[v_e | Y_e, Z_e].$$

However, Y_e and Z_e are not necessarily independent. Recall that $Z_e = "A \cup \{e\} \text{ is an independent set}"$, where A is the set of items accepted before e , and $Y_e = "v_e \geq p_e"$. The price p_e depends on a sample $s^{Index(e)}$ that may have been used to price an edge e' arriving before e , and hence to influence the set A .

For any edge $e = (\ell, r)$, we can define the following two events E_1, E_2 , stating that no other edge e' incident to ℓ and no edge e' incident to r get chosen by \mathcal{P}

$$E_1 = \{e' = (\ell, r') : e' \neq e \wedge v_{e'} \geq p_{e'} \wedge c_{e'} = 1\} = \emptyset$$

$$E_2 = \{e' = (\ell', r) : e' \neq e \wedge v_{e'} \geq p_{e'} \wedge c_{e'} = 1\} = \emptyset$$

If both events E_1 and E_2 hold, then $A \cup \{e\}$ will always be an independent set. Recall that edge e 's contribution to the $\mathcal{P}_{Matching}$'s expected reward is $\frac{1}{3} Pr[Y_e \text{ and } Z_e] \cdot \mathbb{E}[v_e | Y_e \text{ and } Z_e]$. Since Z_e always holds whenever both E_1, E_2 hold, we have

$$Pr[Y_e \wedge Z_e] \cdot \mathbb{E}[v_e | Y_e \wedge Z_e] \geq Pr[Y_e \wedge E_1 \wedge E_2] \cdot \mathbb{E}[v_e | Y_e \wedge E_1 \wedge E_2].$$

Note that events E_1, E_2 only depend on values $v_{e'}$ and prices $p_{e'}$ for $e' \neq e$. Since \mathcal{D} is a product distribution, v_e is independent of $v_{e'}$. Also, since e, e' share a vertex, we have that the prices $p_e, p_{e'}$ are determined using different samples $s^{Index(e)}, s^{Index(e')}$. Thus Y_e is independent of E_1 and of E_2 . This means that we can write

$$Pr[Y_e \text{ and } E_1 \text{ and } E_2] \cdot \mathbb{E}[v_e | Y_e \text{ and } E_1 \text{ and } E_2] = Pr[E_1 \text{ and } E_2] \cdot Pr[Y_e] \cdot \mathbb{E}[v_e | Y_e].$$

Thus, it suffices to give a constant lower bound on $Pr[E_1 \text{ and } E_2]$ in order to guarantee a constant factor competitive ratio for $\mathcal{P}_{Matching}$.

We now follow a line of argument from Chawla, Hartline, Malec and Sivan [6]. Since the edges in a maximum matching form an independent set, and the probability of any edge e being present in a maximum matching is $Pr[v_e \geq p_e] = Pr[Y_e]$, we have

$$\sum_{e' : e' = (\ell, r')} Pr[Y_{e'}] \leq 1$$

$$\sum_{e' : e' = (\ell', r)} Pr[Y_{e'}] \leq 1.$$

Now, the probability of $\mathcal{P}_{Matching}$ choosing an element i is $Pr[X_e \text{ and } Y_e \text{ and } Z_e] \leq Pr[X_e] \cdot Pr[Y_e] = \frac{1}{3} Pr[Y_e]$, so we have

$$\sum_{e' : e' = (\ell, r')} Pr[X_e \text{ and } Y_e \text{ and } Z_e] \leq \frac{1}{3}$$

$$\sum_{e' : e' = (\ell', r)} Pr[X_e \text{ and } Y_e \text{ and } Z_e] \leq \frac{1}{3}$$

This means that the probability that event E_1 does not happen is at most $\frac{1}{3}$, and analogously for event E_2 . Thus, $Pr[E_1] \geq \frac{2}{3}$, $Pr[E_2] \geq \frac{2}{3}$. Since events E_1 is more likely to happen when event E_2 happens, we have

$$Pr[E_1 \text{ and } E_2] \geq Pr[E_1] \cdot Pr[E_1|E_2] \geq \frac{2}{3} \cdot \frac{2}{3} = \frac{4}{9}.$$

We can conclude that

$$\begin{aligned} W &= \sum_{i=1}^n Pr[X_i \text{ and } Y_i \text{ and } Z_i] \cdot \mathbb{E}[v_i|X_i, Y_i, Z_i] \\ &= \sum_{i=1}^n \frac{1}{3} Pr[Y_i \text{ and } Z_i] \cdot \mathbb{E}[v_i|Y_i, Z_i] \\ &\geq \sum_{i=1}^n \frac{1}{3} Pr[Y_i] \cdot Pr[E_1 \text{ and } E_2] \cdot \mathbb{E}[v_i|Y_i] \\ &\geq \sum_{i=1}^n \frac{1}{6.75} Pr[Y_i] \cdot \mathbb{E}[v_i|Y_i] \\ &= \frac{1}{6.75} OPT \end{aligned}$$

We remark that the only place where we needed d^2 samples was to argue that any two incident edges e, e' have independent prices $p_e, p_{e'}$. For general bipartite matchings, if we have $|E|$ samples $s^1, \dots, s^{|E|}$, we can use sample s^e to compute p_e , and then all prices are independent. Thus, our algorithm can be used for general matchings if we have access to $|E|$ samples from \mathcal{D} .

E Omitted Proofs from Section 6

LEMMA E.1. ([3]¹²) *Let \mathcal{J} be any downwards-closed set system and let each \mathcal{D}_i be regular (not necessarily identical). Let \mathbb{M} be a mechanism such that the lazy combination of \mathbb{M} with monopoly reserves has an expected revenue competitive ratio of α . Then the lazy combination of \mathbb{M} with single sample reserves¹³ obtains an expected revenue competitive ratio of $\frac{\alpha}{2}$.¹⁴*

¹²This result was stated for VCG auctions, but it applies without modifying the proof to any auction that approximately maximizes welfare. We note that Dhangwatnotai, Roughgarden and Yan proved this result for VCG auctions with sample reserves. [11]. We also note that the result depends on the fact, proved in [11], that when there is only a single-buyer with distribution \mathcal{D} , the mechanism that offers a posted price equal to a sample from \mathcal{D} obtains $\frac{1}{2}$ of the optimal revenue.

¹³Sample each bidder's reserve r_i independently from \mathcal{D}_i

¹⁴We could also replace the median with the p^{th} quantile and get a competitive ratio of $\alpha \cdot \min\{p, 1-p\}$. Any error in approximating the median (or quantile) is directly absorbed into the competitive ratio as well.

Furthermore, if \mathbb{M} obtains expected welfare competitive ratio of β , then the lazy combination of \mathbb{M} with single sample reserves or median reserves obtains expected welfare competitive ratio of $\frac{\beta}{2}$.

PROPOSITION E.1. ([11]) *Let \mathcal{J} be any downwards-closed set system, and let each \mathcal{D}_i be MHR. Let also \mathbb{M} be any single-dimensional universally truthful mechanism¹⁵ whose expected welfare competitive ratio is α . Then the mechanism \mathbb{M}' that combines (lazily) \mathbb{M} with monopoly reserves has a revenue competitive ratio of $\frac{\alpha}{e}$.*

In order to prove Proposition E.1, we need to borrow a lemma from Yan [25].

LEMMA E.2. ([25]) *Let \mathcal{D} be an MHR distribution with Myerson reserve r^* . Let also $V(t)$ denote the expected welfare of the single bidder mechanism that sets price t , and $R(t)$ denote the expected revenue of the single bidder mechanism that sets price t (when the bidder's value is drawn from \mathcal{D}). Then:*

$$R(\max\{t, r^*\}) \geq \frac{1}{e} V(t)$$

The proof of Proposition E.1 parallels that of Theorem 4.9 from [25], but replaces VCG with an arbitrary truthful mechanism. We again note that it is observed in [11] that their proof for VCG applies to any approximation algorithm, but as their setting and claim is slightly different, we repeat it here for clarity.

Proof of Proposition E.1: Observe first that if we prove the claim for deterministic mechanisms, then the claim immediately follows for universally truthful mechanisms as well. So we can fix bidder i and v_{-i} for the remaining bids and look at the conditional expected revenue from bidder i in this case. For deterministic mechanisms \mathbb{M} , there is some threshold t such that bidder i wins the item if and only if his value is above t . So the conditional contribution to the expected welfare of \mathbb{M} is $V(t)$, and the conditional contribution to the expected revenue of the lazy combination of \mathbb{M} with Myerson reserves is $R(\max\{t, r_i^*\})$. By Lemma E.2, this is at least $\frac{1}{e} V(t)$. So in all cases, the conditional contribution to the expected revenue of the lazy combination of \mathbb{M} with Myerson reserves is at least a $\frac{1}{e}$ fraction of the conditional contribution to the expected welfare of \mathbb{M} , and therefore the

¹⁵A mechanism is universally truthful if it is a distribution over deterministic truthful mechanisms. All posted-price mechanisms are universally truthful.

expected revenue of \mathbb{M} combined lazily with Myerson reserves is at least a $\frac{1}{e}$ fraction of the expected welfare of \mathbb{M} . As the optimal expected welfare upper bounds the optimal expected revenue, this completes the proof. \square

To prove Theorem 6.1 for the lazy combination with Myerson reserves, we need a technical lemma regarding properties of comparison-based algorithms. Lemma E.3 below says that in order for a comparison-based mechanism to achieve good welfare, it must accept a good fraction of the highest bidders in expectation (where “good fraction” means relative to the best possible).

LEMMA E.3. *Let \mathbb{M} be any comparison-based mechanism for feasibility constraints \mathcal{J} whose expected welfare competitive ratio is α . Fix an ordering of bidders x_1, \dots, x_n and relative ordering of values $v_1 > \dots > v_n$ (but not the values themselves). Let also $J(i) = \max_{S \in \mathcal{J}} |S \cap \{1, \dots, i\}|$, and q_j denote the probability that \mathbb{M} selects x_j . Then for all i , we have:*

$$\sum_{j \leq i} q_j \geq \alpha J(i)$$

Proof. Observe first that q_j is well-defined: As \mathbb{M} is a comparison-based mechanism, once we fix the bidders and their relative ordering of values, the behavior of the mechanism is also fixed, independent of what the actual values are. So assume for contradiction that the lemma is false, and let i be an index for which $\sum_{j \leq i} q_j < \alpha J(i)$. Then set $v_j = 1$ for all $j \leq i$ and $v_k = 0$ for all $k > i$. Then \mathbb{M} obtains expected welfare $\sum_{j \leq i} q_j < \alpha J(i)$, and the optimal mechanism obtains expected welfare $J(i)$. So \mathbb{M} does not have expected welfare competitive ratio α .

THEOREM 3. (THEOREM 6.1) *Let \mathcal{J} be any downwards-closed set system, and let each \mathcal{D}_i be identical and regular. Let also \mathbb{M} be any single-dimensional comparison-based mechanism whose expected welfare competitive ratio is α . Then the mechanism that combines (either eagerly or lazily) \mathbb{M} with monopoly reserves has expected revenue competitive ratio α .*

Proof. We first recall Myerson’s lemma that expected revenue (for all truthful mechanisms) is exactly expected virtual welfare [22]. We now make the same observation as [6]: if we run a good welfare mechanism on the *virtual values* instead of the values, then the welfare guarantee of the original mechanism immediately gives us a virtual welfare (i.e. revenue)

guarantee. As the original mechanism was truthful, its allocation rule must have been monotone, and therefore whenever the virtual valuation function, ϕ_i , is monotone, the resulting mechanism is also truthful. ϕ_i is monotone exactly when \mathcal{D}_i is regular.

So the mechanism we would like to implement is \mathbb{M} on the virtual values (which we will denote by $\phi(\mathbb{M})$), but we want to implement $\phi(\mathbb{M})$ without knowing the virtual values. Because each \mathcal{D}_i is identical and regular, whenever $\phi(\mathbb{M})$ wants to compare two virtual values, we can just compare the values instead. This is because the comparison will yield the same result. So all that’s left is to handle negative virtual values.

We could just remove all negative virtual values first, and then run $\phi(\mathbb{M})$ on the remaining bidders. This is exactly the same as removing all bidders who don’t meet their Myerson reserve first, and running \mathbb{M} on the remaining bidders by the observation in the previous paragraph. As \mathbb{M} obtains expected welfare competitive ratio α when all values are positive, we get that $\phi(\mathbb{M})$ obtains expected virtual welfare (revenue) competitive ratio α when run only on bidders with positive virtual values. Therefore, the eager combination of \mathbb{M} with Myerson reserves gives a revenue competitive ratio of α .

We also could just run $\phi(\mathbb{M})$ first, and remove the negative virtual values after. However, it’s not obvious that this mechanism succeeds, as we are no longer directly running $\phi(\mathbb{M})$ on bidders with positive virtual value. Nevertheless, we can use Lemma E.3 to argue that we still get good revenue with lazy removal of negative virtual values. For any fixed bids, relabel the bidders so that $v_1 > \dots > v_n$. Let m denote the largest index such that $v_m \geq 0$, and q_j denote the probability that \mathbb{M} selects bidder x_j , and $Q_i = \sum_{j=1}^i q_j$. Then we can write the expected virtual welfare of $\phi(\mathbb{M})$ with lazy removal of negative virtual values as:

$$\begin{aligned} \sum_{j=1}^m q_j \cdot \phi(v_j) &= Q_m \cdot \phi(v_m) + \\ &\sum_{i=1}^{m-1} Q_i \cdot (\phi(v_i) - \phi(v_{i+1})) \end{aligned}$$

We can also let $p_j = 1$ if Myerson’s auction selects x_j and 0 otherwise, and $P_i = \sum_{j=1}^i p_j$. Then the expected revenue of Myerson’s auction is just:

$$P_m \cdot \phi(v_m) + \sum_{i=1}^{m-1} P_i \cdot (\phi(v_i) - \phi(v_{i+1}))$$

Again let $J(i)$ denote the maximum size of a feasible set in \mathcal{J} using only bidders in $\{x_1, \dots, x_i\}$. Then we clearly have $P_i \leq J(i)$. By Lemma E.3, we also have $Q_i \geq \alpha \cdot J(i)$. Putting this together with the above work we get:

$$\begin{aligned} & Q_m \cdot \phi(v_m) + \sum_{i=1}^{m-1} Q_i \cdot (\phi(v_i) - \phi(v_{i+1})) \\ & \geq \alpha \cdot J(m) \cdot \phi(v_m) + \sum_{i=1}^{m-1} \alpha \cdot J(i) \cdot (\phi(v_i) - \phi(v_{i+1})) \end{aligned}$$

and

$$\begin{aligned} & P_m \cdot \phi(v_m) + \sum_{i=1}^{m-1} P_i \cdot (\phi(v_i) - \phi(v_{i+1})) \\ & \leq J(m) \cdot \phi(v_m) + \sum_{i=1}^{m-1} J(i) \cdot (\phi(v_i) - \phi(v_{i+1})) \end{aligned}$$

which exactly says that the expected virtual welfare competitive ratio of $\phi(\mathbb{M})$ with lazy removal of negative virtual values is α . Again, we observe that this is exactly the same mechanism as \mathbb{M} combined lazily with Myerson reserves and complete the proof of the Theorem.

THEOREM E.1. *Let \mathcal{J} be a downwards-closed set system and let each \mathcal{D}_i be identical and regular. Then there exist truthful OPMs with the following guarantees:*

1. *When \mathcal{J} is a k -uniform matroid, a revenue competitive ratio of $\frac{1}{2} - O(\frac{1}{\sqrt{k}})$ and a welfare competitive ratio of $\frac{1}{2} - O(\frac{1}{\sqrt{k}})$ using two samples from \mathcal{D} .¹⁶*
2. *When \mathcal{J} is a graphic matroid, a revenue competitive ratio of $\frac{1}{16}$, and a welfare competitive ratio of $\frac{1}{16}$ using one sample from \mathcal{D} .*
3. *When \mathcal{J} is a transversal matroid, a revenue competitive ratio of $\frac{1}{32}$ and a welfare competitive ratio of $\frac{1}{32}$ using one sample from \mathcal{D} .*

¹⁶Alternatively, instead of using the rehearsal algorithm, we can use a simpler single-sample algorithm which guarantees a competitive ratio of $\frac{1}{4}$ for the prophet problem. Recall that our motivation for the rehearsal algorithm was purely algorithmic: we want to obtain a single-sample prophet inequality whose competitive ratio of $1 - O(\frac{1}{\sqrt{k}})$ is asymptotically optimal in k . While this motivation still holds from an algorithmic point of view, its not very strong in a mechanism design setting since our use of reserves reduces the competitive ratio by a factor of at least $\frac{1}{2}$.

4. *When \mathcal{J} is a laminar matroid, a revenue competitive ratio of $\frac{1}{24\sqrt{3}}$ and a welfare competitive ratio of $\frac{1}{24\sqrt{3}}$ using one sample from \mathcal{D} .*
5. *When \mathcal{J} is a general matroid, a revenue competitive ratio of $\frac{1-\frac{1}{e}}{40}$ and a welfare competitive ratio of $\frac{1-\frac{1}{e}}{40}$ using one sample from \mathcal{D} .*
6. *When \mathcal{J} is a degree d -bipartite matching, a revenue competitive ratio of $\frac{1}{27}$ and a $\frac{1}{27}$ welfare competitive ratio using $d^2 + 1$ samples from \mathcal{D} .*

Our results for MHR distributions are very similar, with the exception that for the MHR case, our $\mathcal{P}_{\text{Matching}}$ algorithm is the same one as the one described in section 5.

THEOREM E.2. *Let \mathcal{J} be a downwards-closed set system and let each \mathcal{D}_i be MHR (not necessarily identical). Then there exist truthful OPMs with the following guarantees:*

1. *When \mathcal{J} is a k -uniform matroid, a revenue competitive ratio of $\frac{1}{2e} - O(\frac{1}{\sqrt{k}})$ and a welfare competitive ratio of $\frac{1}{2} - O(\frac{1}{\sqrt{k}})$ using two samples from \mathcal{D} .*
2. *When \mathcal{J} is a graphic matroid, a revenue competitive ratio of $\frac{1}{16e}$, and a welfare competitive ratio of $\frac{1}{16}$ using one sample from \mathcal{D} .*
3. *When \mathcal{J} is a transversal matroid, a revenue competitive ratio of $\frac{1}{32e}$ and a welfare competitive ratio of $\frac{1}{32}$ using one sample from \mathcal{D} .*
4. *When \mathcal{J} is a laminar matroid, a revenue competitive ratio of $\frac{1}{24e\sqrt{3}}$ and a welfare competitive ratio of $\frac{1}{24\sqrt{3}}$ using one sample from \mathcal{D} .*
5. *When \mathcal{J} is a degree d -bipartite matching, a revenue competitive ratio of $\frac{1}{13.5e}$ and a $\frac{1}{13.5}$ welfare competitive ratio using $d^2 + 1$ samples.*

F The Free-Order Model

In this section, we provide an improved and simplified analysis of the secretary algorithm in the free-order model proposed by Jaillet, Soto, and Zenklusen [16]. It is easy to see that their algorithm satisfies a modified definition of “order-oblivious” from Section 3 appropriate for the free-order model (the algorithm can choose the order of P instead of having them come in adversarial order), meaning that their algorithm implies a single-sample prophet inequality for the free-order model as well. Let’s first recall their algorithm:

1. Initialize the set of accepted elements, A , to \emptyset .
2. Sample $k = \text{Binomial}(n, 1/2)$ elements uniformly at random from \mathcal{U} and call these the sample set, S . Call the remaining elements P .
3. Find the max-weight basis of S under \mathcal{J} . Label these elements in decreasing order of weight, X_1, \dots, X_k .
4. Set $i = 1$.
5. Draw one at a time in any order each element $y \in P \cap (\text{span}(\{X_1, \dots, X_i\}) - \text{span}(\{X_1, \dots, X_{i-1}\}))$. Add y to A iff $A \cup \{y\} \in \mathcal{J}$ and $v_y > v_{X_i}$.
6. Increment i by one and return to step 5. If $i = k$, and there are any elements not spanned by $\{X_1, \dots, X_m\}$, process them as in step 5.

We first recall a lemma from [16]:

LEMMA F.1. ([16]) *If y is in the max-weight basis of \mathcal{U} under \mathcal{J} , and $y \in P$, then we will always have $v_y > v_{X_i}$ when it is processed in step 5. The only way the algorithm will not accept y is if A already spans y .*

Proof. By definition, we know that $y \in \text{span}(\{X_1, \dots, X_i\})$, and $v_{X_1} > \dots > v_{X_i}$. So if $v_y < v_{X_i}$, greedy would not select y , and y cannot possibly be in the max-weight basis of \mathcal{U} under \mathcal{J} .

DEFINITION 3. *Let Z_1, \dots, Z_m list elements of S in decreasing order of weight for any $S \subseteq \mathcal{U}$. Let $i(y)$ be the minimum i such that $y \in \text{span}(\{Z_1, \dots, Z_i\})$ (if one exists). Then we say the cost of y with respect to S is $v(Z_{i(y)})$ (or 0 if no $i(y)$ exists). Denote this by $C(y, S)$.*

LEMMA F.2. *For all $y \in \mathcal{U}$, if $y \in P$ and $C(y, S) > C(y, P - \{y\})$, A will not span y when it is processed by the algorithm in step 5.*

Proof. First, we observe by the definition of the algorithm that when y is processed, the only elements that could possibly be added to A are of weight at least v_{X_i} . So if y is already spanned, it must be spanned by a subset of $P - \{y\}$ whose elements all have weight at least v_{X_i} . However, it is obvious that $C(y, S) = v_{X_i}$. It is also obvious that if y is spanned by a subset of $P - \{y\}$ whose elements all have weight at least v_{X_i} , that $C(y, P - \{y\})$ is at least v_{X_i} . Therefore, if A spans y at the time the algorithm processes y , it must be the case that $C(y, P - \{y\}) > C(y, S)$, proving the lemma.

THEOREM F.1. *The algorithm of [16] obtains a competitive ratio of $\frac{1}{4}$ whenever \mathcal{J} is a matroid.*

Proof. Clearly, for all $y, y' \in P$ with probability $1/2$. Conditioned on this, it is also clear that $C(y, S) > C(y, P - \{y\})$ with probability $1/2$. This is because whenever we sample $P - \{y\}$ and S , they are switched with probability $1/2$ and the costs are flipped as well. By Lemma F.1 and F.2, every element in the max-weight basis of \mathcal{U} under \mathcal{J} , y , is accepted whenever $y \in P$ and $C(y, S) > C(y, P - \{y\})$. As this happens with probability $1/4$, every element of the max-weight basis is accepted with probability $1/4$, so the algorithm obtains a competitive ratio of $1/4$.

G Analysis of the Rehearsal Algorithm

In this appendix we prove Theorem 4.1

THEOREM 4. (THEOREM 2) *Let $\mathcal{I} = (\mathcal{U}, \mathcal{J})$ be a k -uniform matroid. The rehearsal algorithm is a single-sample algorithm for the prophet problem with a competitive ratio of $1 - O(\frac{1}{\sqrt{k}})$.*

G.1 Part I: The worst adversarial ordering and defining the random walk RW.

Here, we provide the first step in analyzing the rehearsal algorithm, reducing the analysis to answering a question about correlated random walks. We first state a convenient property of the rehearsal algorithm. (In fact, it holds no matter how the thresholds T_1, \dots, T_k are set.)

LEMMA G.1. *For any vector of values $v = (v_1, v_2, \dots, v_n)$, and any thresholds T_1, \dots, T_k , the worst-case order for the rehearsal algorithm is when the values v_i are revealed in increasing order.*

Proof. Consider any fixed v_1, \dots, v_n and T_1, \dots, T_n and assume w.l.o.g. that $v_1 < \dots < v_n$. Also, say there exists some j, j' such that v_j is revealed right before $v_{j'}$ and $v_j > v_{j'}$. Clearly, such j, j' exist whenever the values are not revealed in increasing order. We now want to consider the behavior of the rehearsal algorithm if we swap the order in which v_j and $v_{j'}$ are revealed.

First, observe that whether v_i is accepted or not depends *only* on what slots are available when v_i is revealed and *not* on what elements already filled the slots that are not available. So let S denote the set of available slots right before v_j is revealed. Let S_j denote the subset of S of slots whose threshold is below v_j , and $S_{j'}$ the subset whose threshold is below $v_{j'}$. Since $v_{j'} < v_j$, we have that $S_{j'} \subseteq S_j$. Now we consider a few cases:

First, maybe $S_j = \emptyset$. Then no matter what order v_j and $v_{j'}$ are revealed in, the rehearsal algorithm will reject them both and the same set of thresholds will be available to the remaining elements. So the set of accepted elements will be exactly the same regardless of the order of v_j and $v_{j'}$.

Second, maybe $S_{j'} = \emptyset$, $S_j \neq \emptyset$. Then no matter what order v_j and $v_{j'}$ are revealed in, the rehearsal algorithm will reject $v_{j'}$ and accept v_j to fill the lowest available slot in S_j . So the same set of thresholds will be available to the remaining elements and the set of accepted elements will be exactly the same regardless of the order of v_j and $v_{j'}$.

Third, maybe $S_j = S_{j'}$ and $|S_j| \geq 2$. Then no matter what order v_j and $v_{j'}$ are revealed, the rehearsal algorithm will accept both v_j and $v_{j'}$ and fill the two lowest slots of S_j . So the same set of thresholds will be available to the remaining elements and the set of accepted elements will be exactly the same regardless of the order of v_j and $v_{j'}$.

Fourth, maybe $|S_j| > |S_{j'}| > 0$. Then no matter what order v_j and $v_{j'}$ are revealed, v_j will fill the slot of S_j with the highest threshold value (which is necessarily not in $S_{j'}$), and $v_{j'}$ will fill the slot in $S_{j'}$ with the highest threshold value. So the same slots will be available to the remaining elements and set of accepted elements will be exactly the same regardless of the order of v_j and $v_{j'}$.

Finally, maybe $S_j = S_{j'}$ and $|S_j| = 1$. Then whichever of v_j and $v_{j'}$ is revealed first will fill the single available slot. The second will be rejected. However, the same slots will be available to the remaining elements regardless of their order, so the exact same set of remaining elements will be accepted. The only difference is whether v_j or $v_{j'}$ was accepted. This is the only case where the set of accepted elements will differ, and it differs exactly by replacing v_j with $v_{j'}$, which strictly increases the value of accepted elements.

So we can start from any ordering of the v_i 's and swapping elements a finite number of times until the v_i 's are sorted so that the values are revealed in increasing order. By the above argument, we did not improve the value of accepted elements at any swapping step. Therefore, revealing the v_i 's in order of increasing values is indeed the worst-case order for the rehearsal algorithm.

Using Lemma G.1, we may assume w.l.o.g. that all elements are revealed so that the values are in increasing order. Using this, we will now reduce the problem of analyzing the rehearsal algorithm to answering a question about correlated random

walks. When we run the rehearsal algorithm, the following experiment happens. First, a sample vector $s = (s_1, \dots, s_n)$ is drawn from \mathcal{D} and thresholds T_1, \dots, T_k are set. Then, values v_1, \dots, v_n are revealed in increasing order and accepted/rejected according to the algorithm. Instead, imagine the following equivalent experiment. First, *two* samples are taken from each \mathcal{D}_i , y_i and y'_i . Then, independently for all i , we permute the pair (y_i, y'_i) to determine which element is a “sample” and which one is a “value.” That is, we set $v_i = y_i$ and $s_i = y'_i$ with probability $\frac{1}{2}$, or $v_i = y'_i$ and $s_i = y_i$ with probability $\frac{1}{2}$. We will show that, for *any* $y_1, y'_1, \dots, y_n, y'_n$, the rehearsal algorithm obtains good reward in expectation, where the expectation is taken over the coin tosses that determine which of (y_i, y'_i) is a “value” and which one is a “sample.”

Fix the list $y_1, y'_1, \dots, y_n, y'_n$ and let Y_j denote the j^{th} highest value of this list. Let p_j denote the probability, over the randomness of the coin flips, that the prophet selects Y_j (i.e. the probability that Y_j is one of the k largest “values”). Let's observe a simple upper bound on the expected value the prophet attains with samples Y_1, \dots, Y_{2n} :

OBSERVATION 1. $\sum_{j=1}^{2n} p_j \cdot Y_j \leq \sum_{j=1}^{2k} \frac{1}{2} \cdot Y_j$.

Proof. The prophet chooses element Y_j with probability p_j . Thus $OPT = \sum_{j=1}^{2n} p_j Y_j$. Since the prophet cannot select more than k items, we must have $\sum_{j=1}^{2n} p_j \leq k$. Furthermore, each Y_j has a $\frac{1}{2}$ chance of being a “sample” and thus the prophet will never choose it. Thus $p_j \leq \frac{1}{2}$ for all j . Since $Y_1 \geq \dots \geq Y_{2n}$, these constraints imply that $\sum_{j=1}^{2n} p_j Y_j \leq \sum_{j=1}^{2k} \frac{1}{2} Y_j$.

Our goal is to show that the gambler can guarantee a reward of $(1 - O(\frac{1}{\sqrt{k}})) \cdot OPT$ by using the rehearsal algorithm. Let q_j denote the probability that the rehearsal algorithm selects Y_j . By Observation 1, it suffices to show that $\sum_{j=1}^{2k} q_j Y_j \geq \frac{c}{2} \sum_{j=1}^{2k} Y_j$ for $c = 1 - O(\frac{1}{\sqrt{k}})$. In fact, a sufficient condition for this is that $\sum_{j=1}^i q_j \geq ci/2$ for all $i \leq 2k$.¹⁷

The rest of this section is spent proving this claim. We do this by defining a random walk RW associated with the performance of the rehearsal algorithm. The random walk starts at 0 and goes up or down depending on whether Y_j is a “sample” or a “value”. A formal definition of RW is on the following page.

¹⁷It is easy to see that minimizing $\sum_j q_j Y_j$ subject to this condition will set $q_j = c/2$ for all $j \leq 2k$.

To clarify, if Y_j is a “value,” the walk moves down by 1 at step j . If Y_j is a “sample” and would have set a threshold, the walk moves up by 1 at step j . If Y_j is a “sample” and would have set the threshold that is repeated $2\sqrt{k} + 1$ times, then the walk moves up by $2\sqrt{k} + 1$ at step j . If Y_j is a “sample” and would not have set a threshold, the walk does not move at step j . Now we state some facts that relate the performance of the rehearsal algorithm to facts about this random walk. Still assuming that all x_i are revealed so that the values are in increasing order, we show how to figure out, just by looking at this random walk, which elements are selected by the rehearsal algorithm. We first need a definition and some facts. Figure G.1 illustrates these facts, assigning different colors to accepted and rejected values, as well as filled and unfilled thresholds.

Random Walk RW

- 1 Define $RW(0) = 0$.
- 2 For $j > 0$, given the value $RW(j - 1)$ of the random walk at time $j - 1$, define the value $RW(j)$ of the random walk at time j as:
 - 2.a $RW(j) = RW(j - 1) - 1$ if Y_j is a “value”.
 - 2.b $RW(j) = RW(j - 1) + 1$ if Y_j is a “sample,” and there are at most $k - 2\sqrt{k} - 2$ different $i < j$ that are also “samples.”
 - 2.c $RW(j) = RW(j - 1) + 2\sqrt{k} + 1$ if Y_j is a “sample,” and there are exactly $k - 2\sqrt{k} - 1$ different $i < j$ that are also “samples.”
 - 2.d $RW(j) = RW(j - 1)$ if Y_j is a “sample,” and there are at least $k - 2\sqrt{k}$ different $i < j$ that are also “samples.”

DEFINITION 4. For any j , $H_j^R(RW)$ is the height of RW to the right of j . Or formally, $H_j^R(RW) = \max_{i \geq j} \{RW(i) - RW(j)\}$. Similarly, $H_j^L(RW)$ is the height of RW to the left of j . Formally, $H_j^L(RW) = \max_{i \leq j} \{RW(i) - RW(j)\}$.

If it is clear from context, we will just write H_j^L instead of $H_j^L(RW)$. We can now prove two facts about this random walk and its relation to the rehearsal algorithm when values are revealed by the adversary in increasing order.

FACT G.1. Assuming that the v_i are revealed so that the values are in increasing order, for all j , Y_j is




Figure 1: An illustration of our random walk. The steps in blue correspond to selected values (since the random walk returns to these values eventually), the values in red correspond to rejected values. The samples in black are unfilled thresholds, the samples in green are filled thresholds. The samples in yellow are samples that do not determine a threshold. Notice that there’s a threshold that produces a large jump in the random walk. We also highlight a point i , together with its corresponding left and right heights. The value is accepted because its right height is greater than zero. The number of values to the left that are *not* accepted is exactly $H_i^L - H_i^R$.

chosen by the rehearsal algorithm if and only if Y_j is a “value” and $H_j^R > 0$.

Proof. If $H_j^R > 0$, then there is some $i > j$ with $RW(i) > RW(j)$. RW increases every time it sees a threshold, and decreases every time it sees a value. So that means that there are more thresholds than “values” in the list (Y_{j+1}, \dots, Y_i) . This necessarily means that the first “value” revealed that is at least Y_j will be selected, because there will be at least one available threshold between Y_i and Y_j . Because we are assuming that the values are revealed in increasing order, Y_j is exactly the first value revealed that is at least Y_j , and is therefore selected.

If $RW(i) \leq RW(j)$, then there are at least as many “values” as there are thresholds in the list (Y_{j+1}, \dots, Y_i) . Because the values are revealed in increasing order, this means that the slot using threshold Y_i will certainly be filled before Y_j is revealed. If $H_j^R = 0$, then it is true that $RW(i) \leq RW(j)$ for all $i > j$, which means that all possible slots that Y_j could use will be filled before Y_j is revealed, and therefore Y_j will not be selected by the rehearsal algorithm.

FACT G.2. For all i , the number of “values” in $\{Y_1, \dots, Y_i\}$ that are not selected by the rehearsal algorithm is $\max\{H_i^L - H_i^R, 0\}$.

Proof. Let j_1, \dots, j_h denote the indices of the “values” in (Y_1, \dots, Y_i) that are not selected by the rehearsal algorithm in increasing order. We show that $H_i^L - H_i^R = h$ by first showing that $H_i^L - H_i^R \geq h$, and then showing that $H_i^L - H_i^R \leq h$.

For any index k in $\{1, \dots, h\}$, Y_{j_k} is not selected. Thus, Fact G.1 tells us that it must be the case that $RW(z) \leq RW(j_k)$ for all $z \geq j_k$. In particular, this must hold for $z = j_{k+1} - 1$. Because $Y_{j_{k+1}}$ is a “value”, we know that $RW(j_{k+1}) = RW(j_{k+1} - 1) - 1$, and therefore $RW(j_{k+1}) \leq RW(j_k) - 1$. Chaining this together for all k in $\{1, \dots, h\}$, we get that $RW(j_h) \leq RW(j_1) - (h - 1)$. Because j_1 is a “value”, $RW(j_1) = RW(j_1 - 1) - 1$, which means that we get $RW(j_h) \leq RW(j_1 - 1) - h$.

Since j_h is the index of a “value” that was not selected by the rehearsal algorithm, we know from fact G.1 that $RW(z) \leq RW(j_h)$ for all indices $z \geq j_h$ (which includes all $z \geq i$, since $j_h \in \{1, \dots, i\}$). Let $m = RW(j_h) - RW(i)$ and note that $H_i^L \geq RW(j_1) - RW(i) \geq h + RW(j_h) - RW(i) = h + m$. Furthermore, since $RW(z) \leq RW(j_h)$ for all $z \geq i$, we have $H_i^R \leq RW(j_h) - RW(i) = m$. We conclude that $H_i^L - H_i^R \geq h + m - m = h$.

Let $H = H_i^L - H_i^R$. We will show that $H \leq h$, thus concluding the proof. Since $H_i^L = H_i^R + H$, there exists an index $j \in \{1, \dots, i\}$ such that $RW(j) = RW(i) + H_i^R + H$. So, for every k in $\{1, \dots, H\}$, choose j_k to be the largest index in $\{1, \dots, i\}$ such that $RW(j_k - 1) \geq RW(i) + H_i^R + k$. By this definition, we have $RW(j_k) < RW(i) + H_i^R + k \leq RW(j_k - 1)$, and thus the random walk goes down at step j_k . This means that Y_{j_k} is a “value”. Furthermore, the value Y_{j_k} is not selected by the rehearsal algorithm because $H_{j_k}^R = 0$. To see this, note that for any index j between j_k and i , we have $RW(j) \leq RW(j_k)$ by the definition of j_k (otherwise j_k would not be the largest index satisfying $RW(j_k - 1) \geq RW(i) + H_i^R + k$). Furthermore, for every index $j \geq i$, we have $RW(j) \leq RW(i) + H_i^R < RW(i) + H_i^R + k \leq RW(j_k - 1) = RW(j_k) + 1$. Thus, we have $RW(j) \leq RW(j_k)$ for every $j > j_k$. By Fact G.1 this implies that Y_{j_k} is a value that does not get selected by the rehearsal algorithm. We showed in this paragraph that there are at least $H = H_i^L - H_i^R$ such values. In the previous paragraph we show that there are at most H such values. Thus, we conclude that the number of values in $\{1, \dots, i\}$ that are not selected by the rehearsal algorithm is exactly $H_i^L - H_i^R$.

The expected number of “values” in $\{Y_1, \dots, Y_i\}$ is $\frac{i}{2}$. By Fact G.2, we have that the expected number of values in $\{Y_1, \dots, Y_i\}$ selected by the rehearsal

algorithm is $\frac{i}{2} - \mathbb{E}[\max\{H_i^L - H_i^R, 0\}]$, where the expectation is taken with respect to the coin tosses of the random walk. Thus, to show that $\sum_{j=1}^i q_j \geq \frac{ci}{2}$ for $c = 1 - \frac{d}{\sqrt{k}}$ (where we have made explicit the constant d in $O(\frac{1}{\sqrt{k}})$), it suffices to show that

$$\mathbb{E}[\max\{H_i^L - H_i^R, 0\}] \leq \frac{d \cdot i}{2\sqrt{k}}.$$

Due to space restrictions, a proof of this inequality can be found in the complete version of the paper [4].