Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Prior independent mechanisms via prophet inequalities with limited information

Pablo D. Azar^{a,1}, Robert Kleinberg^{b,2}, S. Matthew Weinberg^{c,*,3}

^a Department of Economics, MIT, United States

^b Department of Computer Science, Cornell University, United States

^c Department of Computer Science, Princeton University, United States

A R T I C L E I N F O

Article history: Received 4 September 2016 Available online 28 June 2018

JEL classification: C790

Keywords: Game Theory Bargaining Theory

ABSTRACT

Prophet inequalities have recently become a fundamental tool in the design of sequential and multi-dimensional mechanisms in Bayesian settings. However, existing mechanisms— as well as the underlying prophet inequalities behind their analysis—require sophisticated information about the distribution from which inputs are drawn.

Our goal in this work is to design prior-independent sequential and multi-dimensional mechanisms. To this end, we first design prophet inequalities that require knowing only a single sample from the input distribution. These results come in two forms: the first is via a reduction from single-sample prophet inequalities to secretary algorithms. The second is via novel single-sample prophet inequalities for k-uniform matroids.

Leveraging our new prophet inequalities, we construct the first prior-independent sequential mechanisms where the seller does not know the order in which buyers arrive, and buyers may have asymmetric value distributions. We also construct the first priorindependent multi-dimensional mechanism where buyers may have asymmetric value distributions.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Myerson's seminal paper (Myerson, 1981) shows how to construct a revenue-optimal incentive-compatible mechanism when buyers have single-dimensional independent private values, act simultaneously, and the mechanism designer knows the exact distributions from which these values are drawn. However, when these conditions are not met, characterizing the revenue-optimal mechanism is a notoriously more difficult task. For example, Manelli and Vincent (2007) show that, even when there are only two goods for sale and the valuation distributions are uniform over [0, 1], the optimal mechanism can be prohibitively complex. Numerous recent works identify further prohibitive complexities, such as non-monotonicity, unbounded menu complexity, and computational intractability (Briest et al., 2015; Hart and Nisan, 2013; Daskalakis et al., 2014, 2015; Hart and Reny, 2015).

In light of these difficulties, recent literature has focused on finding approximately revenue-optimal mechanisms for multi-dimensional settings, or those where the seller lacks an accurate prior. Chawla et al. (2010) construct the first mecha-

* Corresponding author.

- ¹ Supported by Virtual Digital Trust, Grant # 6926059 and the MIT Laboratory for Financial Engineering.
- ² Supported by NSF awards CCF-0643934 and AF-0910940, an Alfred P. Sloan Foundation Fellowship, and a Microsoft Research New Faculty Fellowship.

³ Work done in part while the author was supported by a Microsoft Research Graduate Fellowship, and in part while the author was visiting the Simons Institute for the Theory of Computing.





E-mail addresses: pazar@mit.edu (P.D. Azar), robert.kleinberg@cornell.edu (R. Kleinberg), smweinberg@princeton.edu (S.M. Weinberg).

nisms that guarantee a constant fraction of the optimal revenue when buyers are unit-demand, and have different valuations for different items. They obtain these mechanisms by first finding approximately optimal mechanisms for settings where buyers have single-dimensional valuations, but arrive sequentially, and then by showing how to reduce multi-dimensional mechanism design problems to sequential mechanism design problems. However, their mechanisms still require precise knowledge of the distributions from which buyers' valuations are known.

In parallel to the multi-dimensional mechanism design literature, many recent results starting with the work of Hartline and Roughgarden (2009) and Dhangwatnotai et al. (2010) focus on prior-independent mechanisms for single-dimensional settings. These mechanisms do not require that the seller know the entire distribution of bidder valuations. Instead, they assume that the seller can only observe a single sample from these distributions.

Prior work has indeed considered both lines of work together, although results are limited. For example, Roughgarden et al. (2012) and Devanur et al. (2011) provide prior-independent mechanisms in settings with many unit-demand bidders and many items. More recent work of Goldner and Karlin (2016) provides prior-independent mechanisms in settings with many additive bidders and many items. Still, all prior work requires bidders to be identically and independently distributed. In comparison to these works, we construct the first prior-independent, approximately-optimal mechanisms in multi-dimensional settings with asymmetric bidders.

The main technical tool in our constructions are new prophet inequalities with limited information, which are informationally robust variants of the classical prophet inequality. In the classical setting, a gambler observes a sequence V_1, \ldots, V_n of *n* rewards sampled independently from known distributions $\mathcal{D}_1, \ldots, \mathcal{D}_n$. 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 gambler's goal is to obtain an expected reward that guarantees a constant-fraction of the expected reward obtained by the best offline algorithm, represented by a *prophet* who can observe the values of all the variables V_1, \ldots, V_n before making her selection. A seminal result of Krengel et al. (1977, 1978) states that there is a strategy for the gambler so that his expected reward is at least half of the prophet's expected reward. Owing to their applications in Bayesian mechanism design, there has recently 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 (Alaei, 2011; Chawla et al., 2010), and more generally to settings where the prophet and gambler can choose any independent set in a matroid or matroid intersection (Kleinberg and Weinberg, 2012). However, all existing results require the gambler to know the distributions $\mathcal{D}_1, \ldots, \mathcal{D}_n$ outright (or at least be able to accurately compute expectations of various functions over the joint prior).

We show how the gambler can obtain a constant factor of the prophet's expected reward, even when he only learns a single sample from each D_i . This approach guarantees, for all distributions D, in expectation over the observed sample and the realized state of the world, a constant-factor approximation to the prophet's expected reward.

1.1. Sequential prior-independent mechanism design

The first step in designing prior-independent multi-dimensional mechanisms is of its own independent interest: the construction of *single-dimensional* prior-independent mechanisms when bidders arrive *sequentially*, and valuations may be drawn from independent but not necessarily identical distributions.

In a sequential single-dimensional mechanism design problem, there are *n* buyers interested in receiving some service. Each buyer *i* has value v_i for receiving service, which is drawn from a distribution \mathcal{D}_i . The buyers arrive one at a time, and the seller must make each buyer a "take-it-or-leave-it" offer for service at a price p_i , that may depend on previous offers and previous buyers' choices.

The problem is most interesting when the seller cannot offer service to all buyers simultaneously. That is, there is a collection $\mathcal{J} \subset 2^{\{1,...,n\}}$ of subsets such that a set *S* of buyers can be served simultaneously if and only if $S \in \mathcal{J}$. For example, if the seller has *k* copies of a good for sale, then only *k* buyers can be served, implying in this case that $\mathcal{J} = \{S \subset \{1,...,n\} : |S| \leq k\}$.

Following the lead of Chawla et al. (2010), we consider two variants of this problem. The first is when the seller can choose the order in which buyers arrive – such mechanisms are referred to as Sequential Posted Price Mechanisms (SPM's). The second is when the seller does not know the order in which buyers will arrive – such mechanisms are referred to as Order-Oblivious Posted Price Mechanisms (OPM's).

We show a new approximately optimal *prior-independent* SPM whenever the collection \mathcal{J} of feasibility constraints has a matroid structure.⁴ Prior work covers matroid settings when the designer knows the prior (or at least enough information to compute virtual values) (Yan, 2011). SPMs are not known to imply multi-dimensional mechanisms, so these results are of interest purely as posted-price mechanisms in single-dimensional settings.

We further construct *prior-independent* OPMs for all collections \mathcal{J} of feasibility constraints for which prior-independent prophet inequalities exist. To this end, we also design new single-sample prophet inequalities (discussed further in Section 1.3).

⁴ That is, \mathcal{J} is downwards-closed, 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 $S' \cup \{x\} \in \mathcal{J}$.

OPMs are known to imply multi-dimensional mechanisms via the reduction of Chawla et al. (2010) (covered more clearly in Section 3), so these results are also of interest for the multi-dimensional mechanisms they imply. In particular, when \mathcal{J} is a bipartite matching, our prior-independent OPM implies a prior-independent multi-dimensional mechanism described below.

1.2. Multi-dimensional prior-independent mechanism design

The second contribution of this paper to the mechanism design literature is the construction of the first multidimensional prior-independent mechanisms when bidders' valuations may be drawn from independent but asymmetric distributions.

Consider the following 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*, drawn from a distribution \mathcal{D}_{ij} . The buyers are *unit-demand*, in that their value for a set of items is simply their value for their favorite item in that set. We construct a mechanism that guarantees a constant-factor approximation to the expected revenue of the optimal mechanism for multi-dimensional unit-demand settings, where the seller only needs to know **one** sample from each \mathcal{D}_{ij} .

We briefly overview how this result is obtained, as it demonstrates the full chain of our tools. We start by treating each (bidder, item) pair as an element, and observing that because bidders are unit-demand, it is feasible to simultaneously award item *j* to bidder *i* (and have bidder *i* obtain non-zero value for it) for all (i, j) in some set *S* if and only if *S* is a matching in the bipartite graph where left nodes are bidders and right nodes are items. A recent algorithm of Feldman et al. (2018) achieves a $\frac{1}{256}$ approximation for the *secretary problem* in bipartite matching environments (described in further detail in the following section). We apply our new reduction (Section 5) to convert their algorithm into a *single-sample prophet inequality* for the same setting that obtains the same competitive ratio.

Ideally, the next step would be to use the machinery of Chawla et al. (2010), and simply run our new prophet inequality on the distribution of virtual values, and obtain an OPM for a related single-dimensional *copies* setting (which essentially treats all (i, j) as a distinct bidder rather than bidder *i*'s value for item *j*). To do this, we would only require one sample from the virtual value distribution. Unfortunately, even obtaining a single sample from the virtual value distribution requires extensive information about the underlying value distribution, so this approach is infeasible.

Instead, we develop machinery similar to Hartline and Roughgarden (2009), Dhangwatnotai et al. (2010), which takes our new prophet inequality with competitive ratio $\frac{1}{256}$ and turns it into a truthful OPM that guarantees a $\frac{1}{256e}$ approximation for the related copies setting, whenever each D_{ij} has a Monotone Hazard Rate (but are not necessarily identical), or a $\frac{1}{512}$ approximation whenever the distributions are i.i.d. and regular.

Finally, portions of machinery from Chawla et al. (2010) do still apply and allow us to conclude that any OPM for the related copies setting immediately provides the same guarantee for the original multi-dimensional setting (when compared against the revenue optimal Dominant-Strategy Incentive Compatible mechanism). The competitive ratio against the optimal Bayesian Incentive Compatible mechanism for the same setting is at most an additional factor of 4 smaller (Chawla et al., 2015; Cai et al., 2016).

So the full approach is to start with the secretary algorithm of Feldman et al. (2018), plug it into our reduction from Section 5 to get a new single-sample prophet inequality. Then plug our single-sample prophet inequality into new machinery based on Hartline and Roughgarden (2009), Dhangwatnotai et al. (2010) to get an OPM for this single-dimensional copies setting. Finally, machinery of Chawla et al. (2010) turns this into a truthful mechanism for the original Bayesian multi-dimensional setting.

1.3. New prophet inequalities

The key technical contribution that allows us to construct new prior-independent OPMs (and, by the reduction of Chawla et al., 2010, referenced above, prior-independent multi-dimensional mechanisms) is our development of *limited-information* prophet inequalities. As described above, these are analogous to traditional prophet inequalities, where a gambler must choose a feasible set of items that arrive one by one, with the twist that the only information that the gambler has about the value of future items is a single sample from their distributions. We derive our limited-information prophet inequalities via two different approaches.

- 1. **Reduction from existing secretary problems.** In Section 5, we give a black-box reduction that obtains single-sample prophet inequalities from existing order-oblivious⁵ algorithms for the secretary problem.⁶ The ratio obtained after our reduction is exactly the same as the corresponding secretary algorithm. Many existing secretary algorithms are order-oblivious, listed below:
 - *k*-uniform matroids: A modification of the 1/e-approximation of Babaioff et al. (2007a, 2007b) is order-oblivious and achieves a 1/4-approximation. The $1 O(1/\sqrt{k})$ -approximation of Kleinberg (2005) is not order-oblivious.

⁵ We define what order-oblivious algorithms are in Section 5.

⁶ In 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.

- Graphic matroids: A modification of the 1/(2e)-approximation of Korula and Pál (2009) is order-oblivious and achieves an 1/8-approximation.
- **Co-graphic matroids:** A modification of the 1/(3e)-approximation of Soto (2011) is order-oblivious and achieves a 1/12-approximation.
- Transversal matroids: Dimitrov and Plaxton's (2012) 1/16-approximation is order-oblivious. Kesselheim et al. (2013) 1/e-approximation is not order-oblivious (and it does not appear that a simple modification of it is).
- Laminar matroids: Ma et al. (2016) 1/9.6-approximation is order-oblivious.
- **Regular and max-flow-min-cut matroids:** A modification of the 1/(9e)-approximation of Dinitz and Kortsarz (2014)
- is order-oblivious and achieves a 1/36-approximation. **General matroids:** Oveis Gharan and Vondrák's $\frac{1-1/e}{40}$ -approximation in the random assignment model is order-oblivious. Lachish's (2014), and Feldman et al. (2015) $1/O(\log \log k)$ -approximation in the standard model is orderoblivious.
- **Bipartite matchings:** Feldman et al. (2018) 1/256 approximation is order-oblivious.

The astute reader might notice that in the above list, whenever a modification is necessary to make an existing algorithm order-oblivious, the competitive ratio degrades by a factor of e/4. This is essentially due to replacing a subroutine that runs Dynkin's 1/e-approximation for the original secretary problem (Dynkin, 1963) with a sub-optimal algorithm and simpler competitive analysis that only proves a 1/4-approximation. More details on this appear in Section 5.

2. Analysis of correlated random walks. The best known secretary algorithms (Kleinberg, 2005) and full-information prophet inequalities (Alaei, 2011) for k-uniform matroids both guarantee a $1 - O(\frac{1}{\sqrt{k}})$ competitive ratio, but Kleinberg's secretary algorithm is not order-oblivious, and Alaei's prophet inequality requires heavy knowledge of the distributions. In order to asymptotically match this competitive ratio, we give a new algorithm in Section 6, whose analysis models the drawing of "samples" or "values" as positive and negative steps in a random walk. By estimating the expected height of this correlated random walk, we are able to guarantee that each of the top k values (that is, the values that are accepted by the optimal offline algorithm) are selected by our online algorithm with probability $1 + O(\frac{1}{\mu})$.

1.4. Organization

In Section 2 below, we cover related work on prophet inequalities and prior independent mechanism design. In Section 3, we provide notation and definitions for the following theorems. We begin our technical Sections by showing how to obtain prior independent/limited-information mechanisms from limited-information prophet inequalities in Section 4. Sections 5 and 6 contain our new prophet inequalities. The results in Section 5 are via a reduction to the secretary problem. Section 6 contains our prophet inequality for k-uniform matroids. Since our main results concern OPMs and their implied multi-dimensional mechanisms, we delay the treatment of our SPM for matroids to Appendix A.

2. Related work

2.1. Prophet inequalities

The literature on prophet inequalities is large, and we will not try to summarize it all here. Following Krengel, Sucheston, and Garling's seminal work (Krengel and Sucheston, 1977, 1978), Samuel-Cahn (1984) designed an extremely elegant singlechoice prophet inequality guaranteeing the optimal competitive ratio of 2: set a threshold T such that $Pr[max_i V_i > T] = 1/2$, and accept any element that exceeds this threshold. Most related to our work is that on multiple-choice prophet inequalities, where the gambler's reward is equal to the sum of rewards from selected elements. Here, the first results are from Chawla et al. (2010) who design constant-factor approximations for uniform matroids, graphic matroids, and bipartite matchings. Alaei obtained an asymptotically optimal ratio of $1 - 1/\sqrt{k+3}$ for k-uniform matroids (Alaei, 2011). For general matroids, Yan (2011) obtained a (1 - 1/e)-approximation when the gambler gets to choose the order, and Kleinberg and Weinberg (2012) obtained a 1/2-approximation in the general case. In comparison, our new prophet inequalities, in settings where they apply, match the asymptotic guarantees of prior work while only requiring a single sample from each distribution instead of outright knowledge.

Our work also provides a formal connection between the secretary problem and prophet inequalities. Independently, Göbel, Hoefer, Kesselheim, Schleiden, and Vöcking also discover a formal connection (Göbel et al., 2014). In our work, we show that certain kinds of secretary algorithms imply single-sample prophet inequalities. Göbel et al. instead pose a more general problem, such that both prophet inequalities and secretary problems are special cases. In other words, any solution to their problem immediately imply both prophet inequalities and secretary problems. Aside from the fact that both works provide the first formal connections between secretary problems and prophet inequalities, it is unclear that these results are further related.

2.2. Prior independent mechanism design

Myerson's seminal work (Myerson, 1981) shows how to find the revenue-optimal auction in any single-dimensional setting, so long as the designer has enough information about the prior to compute "virtual values." Another seminal work of Bulow and Klemperer (1996) showed the following: in single-item settings, if all bidders are i.i.d. and have values drawn from regular distributions, the revenue of a second price auction (with no reserve) with an additional bidder is always larger than the revenue of Myerson's optimal auction in the original setting. More recent seminal work of Hartline and Roughgarden (2009) provided approximate Bulow–Klemperer type Theorems for more general single-dimensional settings with asymmetric bidders – in many cases, running the Vickrey–Clarke–Groves mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973) with additional bidders gets a constant factor of the revenue of Myerson's optimal auction in the original setting. Dhangwatnotai et al. (2010) and Azar et al. (2013) also design approximately optimal prior-independent mechanisms in more general single-dimensional settings.

In the realm of posted-price mechanisms, Chawla et al. (2010) provided a generic reduction from designing multidimensional mechanisms in certain settings (formalized in Section 3) to designing OPMs in a related single-dimensional setting. They also identified SPMs as being interesting in their own right as simpler mechanisms in single-dimensional settings, and designed SPMs and OPMs in numerous settings. Yan (2011) designed a tight (1 - 1/e)-approximate SPM in all matroid settings. Kleinberg and Weinberg designed 1/(4p - 2)-approximate OPMs in all settings that can be written as the intersection of p matroids. In comparison to these works, our results, in settings where they apply, match the asymptotic guarantees of prior work while only requiring a single sample from each distribution. However, our results only apply when the distributions are i.i.d. and regular or asymmetric and MHR. It is also still open to extend our single-sample results to general matroids and matroid intersections.

3. Preliminaries

Environments and offline selection problems. An environment $\mathcal{I} = (\mathcal{U}, \mathcal{J})$ is given by a universe of elements $\mathcal{U} = \{1, ..., n\}$ 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 $v = (v_1, ..., v_n)$ 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 $v = (v_1, \ldots, v_n)$ in some order $(v_{i_1}, \ldots, v_{i_n})$ (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 $A^*(v) = A(v_{i_1}, \ldots, v_{i_n})$ 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}, \ldots, v_{i_n} are revealed.

Prophet inequalities. Given an environment \mathcal{I} with universe set $\mathcal{U} = \{1, ..., n\}$, let $\mathcal{D} = \mathcal{D}_1 \times ... \times \mathcal{D}_n$ be a product distribution over $\mathbb{R}^n_{\geq 0}$.⁷ Let $v = (v_1, ..., v_n)$ 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

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

where the expectations are taken with respect to the random choice of v and the random coin tosses of A. The above inequality holds *regardless* of the order in which the elements v_{i_1}, \ldots, v_{i_n} are revealed. We remark that this is a stronger property than that guaranteed by the prophet inequalities in previous papers (Kleinberg and Weinberg, 2012), where the adversary had to choose which element i_j to reveal at time j using only knowledge of the items and values $(i_1, v_{i_1}), \ldots, (i_{j-1}, v_{i_{j-1}})$ 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 $\mathcal{D}_1, \ldots, \mathcal{D}_n$ from which the values are drawn. We say that \mathcal{A} is a single-sample prophet inequality if it has access only to a single sample $s^1 = (s_1, \ldots, s_n)$, drawn from the joint distribution \mathcal{D} . When \mathcal{A} is single-sample, its expected reward $\mathbb{E}_{v,s}[\sum_{i \in \mathcal{A}^*(s;v)} v_i]$ is computed over the randomness in the vector of values v, the random sample s and the random coin tosses of the algorithm.

Our constraints. We can give different feasibility constraints by placing different structure on \mathcal{J} . Some readers may choose not to parse definitions for some of the more advanced matroid classes below. It is not necessary to understand the proofs of our results, only to which domains they apply. We are not advocating that advanced matroid classes correspond to reasonable auction design settings, but as our results extend to these settings for free, we wish to at least state them (with the exception of max-flow-min-cut matroids, for which a formal definition would require excessive matroid-specific language that is likely outside the interest of most readers).

⁷ We remark that the assumption that the rewards V_1, \ldots, V_n are independent is somewhat necessary if we want a constant competitive ratio. Hill and Kertz (1983) 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.

- **Matroids.** \mathcal{J} is a matroid if and only if \mathcal{J} is downwards-closed⁸, 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 $S \subset U$ is in \mathcal{J} if and only if $|S| \leq k$.
- **Partition matroids.** Let B_1, \ldots, B_ℓ be disjoint subsets of \mathcal{U} such that $\mathcal{U} = B_1 \cup \ldots \cup B_\ell$. 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 $|S \cap B_i| \le c_i$ for every $i \in \{1, \ldots, \ell\}$.
- **Graphic matroids.** Let G = (V, E) be a graph with vertex set V and edge set E. The universe 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 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 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.
- 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 $A \in \mathcal{F}$, a positive integer capacity c_A . A set $S \in \mathcal{J}$ if and only if $|S \cap A| \leq c_A$ for all $A \in \mathcal{F}$.
- **Co-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 $S \subseteq E$ is in \mathcal{J} if and only if E S contains a spanning forest.⁹
- **Regular matroids.** A matroid is \mathbb{F} -representable if its universe \mathcal{U} can be interpreted as elements in a vector space over \mathbb{F} , and a subset *S* of elements is in \mathcal{J} if and only if the vectors in *S* are linearly independent. A matroid is regular if it is \mathbb{F} -representable for all fields \mathbb{F} .

Secretary problems. The secretary problem for an environment $(\mathcal{U}, \mathcal{J})$ (Babaioff et al., 2007b) is an online selection problem where the item values v_1, \ldots, 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. However, our reduction in Section 5 shows that designing certain kinds of solutions for the secretary problem is strictly harder than designing prophet inequalities for the same setting.

3.1. Mechanism design preliminaries

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 $\mathcal{D} = \mathcal{D}_1 \times ... \times \mathcal{D}_n$, 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 $b = (b_1, ..., b_n)$, $x_i(b)$ 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 $U(v_i, v_i, b_{-i}) \ge U(v_i, b_i, b_{-i})$. We also require mechanisms to be individually rational. That is, $U(v_i, v_i, b_{-i}) \ge 0$ for all v_i, b_{-i} .

Allocation rules determine prices. (Myerson, 1981; Archer and Tardos, 2001) 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.

⁸ \mathcal{J} is downward-closed if for any $S \in \mathcal{J}$ and any $T \subset S$, we have $T \in \mathcal{J}$.

⁹ Co-graphic matroids are *duals* to graphic matroids, but it is not necessary for this work to understand matroid duals.

Welfare. For a given instance of BSMD, the optimum welfare is $\mathbb{E}[\max_{s \in \mathcal{J}} \{\sum_{i \in S} v_i\}]$. We refer to this as OPT_W . We also refer to the expected welfare of a *mechanism* as the expected sum of values for allocated bidders. We say that a mechanism gets a welfare competitive ratio of α if its expected welfare exceeds $\alpha \cdot OPT_W$ for all \mathcal{D} . We may also discuss a mechanism's welfare competitive ratio for a subclass of distributions, such as i.i.d. or regular.

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 seminal result (Myerson, 1981) shows that in all single dimensional settings, the expected revenue of a truthful mechanism is exactly its expected virtual welfare. That is $\mathbb{E}_{v}[\sum_{i=1}^{n} p_{i}(v)] = \mathbb{E}_{v}[\sum_{i} x_{i}(v)\phi_{i}(v_{i})]$. Therefore, the optimum revenue is $\mathbb{E}[\max_{S \in \mathcal{J}} \{\sum_{i \in S} \phi_{i}(v_{i})\}]$. We refer to this as OPT_{R} . We say that a mechanism gets a revenue competitive ratio of α if its expected revenue exceeds $\alpha \cdot OPT_{R}$ for all \mathcal{D} .

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.¹⁰

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 $(i, j) \in S$. 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 Chawla et al. (2010), 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.

Mechanisms with reserves. The idea of combining simple, welfare-optimizing mechanisms with revenue-optimizing reserve prices originated in Hartline and Roughgarden (2009). In Hartline and Roughgarden (2009), 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. Dhangwatnotai et al. (2010) 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 $r_i \leftarrow D_i$ for bidder *i*, that is drawn from the same distribution as D_i .

A reduction from OPMs to multi-dimensional mechanism design. Chawla et al. (2010) 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)th buyer in $\mathcal{I}^{\text{copies}}$ only values the (i, j)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 Chawla et al. (2010):

Theorem 1. (Chawla et al., 2010) 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 O PM 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 of any deterministic DSIC mechanism, and an $\alpha/4$ -approximation to the optimal randomized mechanism.¹¹

4. Mechanism design with limited information

4.1. From prophet inequalities to mechanisms

Consider a limited-information prophet inequality for any setting \mathcal{J} with a competitive ratio of α . All reasonable prophet inequalities (certainly every prophet inequality referenced in this work, and all those that the authors are aware of) are monotonic in v, meaning that the higher a value v_i is (fixing all other v_j), the higher the probability i is selected. All such prophet inequalities therefore induce an allocation rule $x(\cdot)$, and this allocation rule is monotonic. When each value

¹⁰ We remark that our definition matches that of Kleinberg and Weinberg (2012), which extends the one given in Chawla et al. (2010).

¹¹ The extension to randomized mechanisms holds even when compared to the optimal Bayesian IC mechanism, and is due to Chawla et al. (2015) (who proved a bound of α /5). The improvement to α /4 is due to Cai et al. (2016). A mechanism is Bayesian IC (BIC) if it is in every bidder's best interest to tell the truth assuming that all other bidders' reports are drawn from \mathcal{D}_{-i} .

corresponds to a different bidder (single-dimensional setting), this monotonic allocation rule implies a pricing rule $p(\cdot)$ which makes the mechanism (x, p) truthful. This implies that every monotone prophet inequality immediately implies a truthful online mechanism whose expected welfare is at least α OPT_W. Furthermore, these mechanisms are all posted-price mechanisms, because they accept *i* iff v_i exceeds some threshold T_i .

So the challenge is figuring out how to get a guarantee on the revenue obtained by such mechanisms. Chawla, Hartline, Malec, and Sivan provide a framework with which to do this by considering the bidder's virtual values instead of values (Chawla et al., 2010). If we run a prophet inequality with competitive ratio α on the distributions of *virtual values*, we select a set of bidders whose virtual values in expectation are at least an α fraction of the optimal expected virtual welfare. Therefore, by Myerson's Theorem (Myerson, 1981), such an allocation rule combined with appropriate payments obtains at least an α fraction of the optimal expected revenue.

One issue with this approach is that even getting one sample from a bidder's virtual value distribution (or even knowing the virtual values of the actual bidders) requires extensive knowledge of the distributions. So single-sample prophet inequalities do not immediately imply prior independent mechanisms through this framework. Instead, we take an approach more similar to Hartline and Roughgarden (2009), who show that welfare-maximizing auctions combined with appropriately chosen reserves are also revenue-maximizing in certain settings.

4.2. From welfare to revenue: the i.i.d. case

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, and will emphasize in all of our prophet inequalities whether or not they are comparison-based.

Definition 1. Let $\mathbb{M}(v; s^1, \ldots, s^d)$ be a mechanism for single-dimensional settings which depends on a vector of bids $v = (v_1, \ldots, v_n) \leftarrow \mathcal{D}$ and also on a collection of samples $s^1 = (s_1^1, \ldots, s_n^1), \ldots, s^d = (s_1^d, \ldots, s_n^d)$, 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 $x(v_1, \ldots, v_n, s_1^1, \ldots, s_n^d)$ only depends on the relative order of its arguments, and not on their respective values.

Theorem 2. 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 .

To prove Theorem 2 for the lazy combination with Myerson reserves, we need a Lemma regarding properties of comparison-based algorithms. Lemma 1 below says that in order for a comparison-based mechanism to achieve good welfare on all product distributions, it must accept a good fraction of the highest bidders in expectation (where "good fraction" means relative to the best possible).

Lemma 1. Let \mathbb{M} be any comparison-based mechanism for downwards-closed feasibility constraints \mathcal{J} whose expected welfare competitive ratio is α . Fix an ordering of bidders x_1, \ldots, x_n and relative ordering of values $v_1 > \ldots > v_n$ (but not the values themselves). Let also $J(i) = \max_{S \in \mathcal{J}} \{|S \cap \{1, \ldots, i\}|\}$, and q_i denote the probability that \mathbb{M} selects x_i . 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 $\sum_{j \le i} q_j < \alpha J(i)$. Then consider the distribution for which \mathcal{D}_j is a point mass at 1 (and therefore $v_j = 1$ with probability 1) for all $j \le i$ and \mathcal{D}_k is a point mass as 0 (and therefore $v_k = 0$ with probability 1) for all k > i. Then \mathbb{M} obtains expected welfare $\sum_{j \le i} q_j < \alpha J(i)$, and the optimal mechanism obtains expected welfare J(i). So \mathbb{M} does not have expected welfare competitive ratio α .

Proof of Theorem 2. We first observe that whenever all D_i are regular and i.i.d., the relative ordering by values is *exactly the same* as the relative ordering by virtual values. So if values are sorted in decreasing order as v_1, \ldots, v_n , the following are also sorted in decreasing order: $\max\{\phi(v_1), 0\}, \ldots, \max\{\phi(v_n), 0\}$. As everything in the latter list is positive, immediately because \mathbb{M} obtains a welfare competitive ratio of α , we get that the expected virtual surplus of \mathbb{M} with eager monopoly reserves (replace all bidders who fall below their monopoly reserve with a dummy bidder of virtual value 0) is at least $\mathbb{E}[\alpha \max_{S \in \mathcal{J}} \{\sum_{i \in S} \max\{\phi(v_i), 0\}]$, exactly an α -fraction of the optimal expected virtual surplus. By Myerson's Lemma, this implies that such a mechanism also obtains an expected revenue competitive ratio of α .

For lazy reserves, we need to be more careful because of negative virtual values. The problem is that no comparisonbased mechanism can guarantee a non-trivial welfare competitive ratio if values might be negative (because it can't distinguish positive from negative values). So we cannot simply invoke the algorithm's competitive ratio while some virtual values are negative. Fortunately, Lemma 1 captures exactly what is necessary to show that removing negative virtual values ex-post still works. Let *m* denote the largest index such that $\phi(v_m) \ge 0$, 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:

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

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 virtual surplus from Myerson's auction is just:

$$P_m \cdot \phi(\mathbf{v}_m) + \sum_{i=1}^{m-1} P_i \cdot (\phi(\mathbf{v}_i) - \phi(\mathbf{v}_{i+1}))$$

Again let J(i) denote the maximum size of a feasible set in \mathcal{J} using only bidders in $\{x_1, \ldots, x_i\}$. Then we clearly have $P_i \leq J(i)$. By Lemma 1, we also have $Q_i \geq \alpha \cdot J(i)$. Putting this together with the above work we get the following two inequalities:

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

which exactly says that on every profile, running \mathbb{M} with lazy removal of negative virtual values gets an α fraction of the optimal virtual surplus. Taking an expectation over all profiles, we see that the expected revenue of \mathbb{M} with lazy monopoly reserves has 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 from Azar et al. (2013), Dhangwatnotai et al. (2010).

Lemma 2. (*Azar et al.*, 2013; Dhangwatnotai et al., 2010) 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 reserves obtains an expected revenue competitive ratio of $\frac{\alpha}{2}$. Furthermore, if \mathbb{M} obtains expected welfare competitive ratio of β , then the lazy combination of \mathbb{M} with single sample reserves obtains expected welfare competitive ratio of β .

We highlight that we could also replace the single sample with the median of D_i , or more generally with the *p*th quantile of D_i 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.

As a corollary of Theorem 2 and Lemma 2, we conclude that any mechanism which guarantees an constant-fraction approximation to welfare can be combined with lazy sample reserves to guarantee a constant-fraction approximation to revenue when the valuations are drawn from i.i.d. regular distributions.

Corollary 1. Let \mathbb{M} be a comparison-based, single-dimensional mechanism that guarantees an α approximation to welfare for all product distributions. Then when \mathcal{D} is i.i.d. and regular, \mathbb{M} combined with lazy sample reserves guarantees an $\frac{\alpha}{2}$ approximation to revenue and an $\frac{\alpha}{2}$ approximation to welfare.

The proof of Corollary 1 follows by first plugging \mathbb{M} into Theorem 2. This guarantees that \mathbb{M} combined with lazy monopoly reserves yields a revenue competitive ratio of α . Now, the hypotheses of Lemma 2 are satisfied and we can further conclude that \mathbb{M} combined with lazy sample reserves yields the stated competitive ratios.

4.3. From welfare to revenue: the MHR case

Here, we can fortunately just use prior work directly. It now known 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 (Dhangwatnotai et al., 2010).

Lemma 3. (Dhangwatnotai et al., 2010) 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 mechanism¹² 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 $\frac{\alpha}{\sigma}$.

¹² A mechanism is universally truthful if it is a distribution over deterministic truthful mechanisms. All posted-price mechanisms are universally truthful.

Combining this with Lemma 2, we obtain the following corollary.

Corollary 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.

5. Single-sample prophet inequalities from existing 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 2. We say that an algorithm S for the secretary problem (together with its corresponding analysis) is **order-oblivious** if, on a randomly ordered input vector (v_{i_1}, \ldots, v_{i_n}):

- 1. (algorithm) S sets a (possibly random) number k, observes without accepting the first k values $S = \{v_{i_1}, \ldots, v_{i_k}\}$, and uses information from S to choose elements from $V = \{v_{i_{k+1}}, \ldots, v_{i_n}\}$.
- 2. (analysis) 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 now show how to construct a prophet inequality \mathcal{P} given an order-oblivious algorithm \mathcal{S} for the secretary problem. Recall that the algorithm \mathcal{P} takes as offline input a vector $s = (s_1, \ldots, s_n)$ 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.

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

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

Online Stage

3. For each index i ∈ {i₁,..., i_n}:
a. If i ∈ {j₁,..., j_k}, then index i has already been processed as a "sample". Ignore it and continue.
b. If i ∈ {j_{j+1},..., j_n}, pass the value v_i to algorithm S, and accept i if and only if S accepts i.

Theorem 3. If S is an order-oblivious algorithm for the secretary problem with competitive ratio α , then \mathcal{P}_S is a single-sample prophet inequality with competitive ratio α .

Proof. The algorithm \mathcal{P}_{S} first permutes the vector *s* of samples into a random permutation s_{j_1}, \ldots, s_{j_n} and takes the first *k* elements s_{j_1}, \ldots, s_{j_k} of this permutation and passes them as inputs to the secretary algorithm S. After that, the secretary algorithm S is passed all the inputs v_i where $i \notin \{j_1, \ldots, j_k\}$ in an arbitrary order. Since S is order-oblivious, the set it selects has a weight of at least $\alpha \cdot OPT(v)$, where $OPT(v) = \max_{A \in \mathcal{J}} \sum_{i \in A} v_i$. 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}_S obtains expected reward of at least

$$\int_{v} f(v)\alpha \cdot OPT(v)dv,$$

because whether or not element *i* is a sample is independent of s_i , v_i . So the distribution of values passed into the secretary algorithm is indeed consistent with $f(\cdot)$. The prophet's expected reward is exactly $\int_{v} f(v) \cdot OPT(v)dv$, which immediately says that \mathcal{P}_{S} obtains competitive ratio α .

Note that our single-sample algorithm \mathcal{P}_{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.

5.1. Existing order-oblivious secretary algorithms

In this Section, we sketch existing secretary algorithms and briefly argue why they are order-oblivious. We conclude the Section with a formal statement of the prophet inequalities implied by these results, as well as the prior-independent mechanisms.

General matroids in the random assignment model (Gharan and Vondrák, 2011). If the rank of the matroid given by \mathcal{J} is less than 12, this algorithm runs the rank-1 matroid algorithm. Otherwise it observes a set the first half of its input and sets a threshold T equal to the $\lfloor \frac{r}{4} \rfloor$ + 1st largest value it observes, where r is the rank of the matroid. For the second half of the input, it accepts all items above the threshold T, as long as accepting them does not violate the matroid constraints. In their proof, the authors directly lower bound the probability that a certain weight is selected by the algorithm for the worst-case ordering of the second half, so both the algorithm and analysis are order-oblivious.

Transversal matroids (Dimitrov and Plaxton, 2012). The algorithm begins by assigning an ranking to the set R of right vertices. It then chooses a set S of "samples" consisting of the first $k = Binom(n, \frac{1}{2})$ values seen. All the values in S are discarded, but they are used to construct an auxiliary matching $M_0(S)$, where each item in S is matched to the highest ranking right-node that is still available. The algorithm then constructs the "real matching" M_1 using elements from V =L-S. As each of the remaining left-vertices $\ell \in L-S$ arrives, ℓ is matched with the highest ranked right vertex r that is not matched in $M_0(S)$, as long as r is not already matched in M_1 . Dimitrov and Plaxton show that this is a $\frac{1}{16}$ competitive algorithm, and that this competitive ratio holds regardless of the order in which elements from V are revealed. Thus, the algorithm is order-oblivious.

Rank-1 matroids. Many of the remaining algorithms involve a subroutine for 1-uniform matroids. Dynkin's algorithm is not order-oblivious, so we first give a very simple $\frac{1}{4}$ -competitive algorithm for the classical secretary problem that is.

 $S_{Rank-1}(v_{i_1},\ldots,v_{i_n})$

- 1 Let $k = Binomial(n, \frac{1}{2})$.
- 2 Let $T = \max\{v_{i_1}, \ldots, v_{i_k}\}.$
- 3 Accept the first element in $v_{i_{k+1}}, \ldots, v_{i_n}$ satisfying $v_i > T$.

With probability 1/4, the highest element is somewhere in $v_{i_{k+1}}, \ldots, v_{i_n}$ and the second-highest is a "sample" in v_{i_1},\ldots,v_{i_k} . In this case, the highest element is accepted no matter what order the elements in V are revealed. Thus S_{Rank-1} is order-oblivious.

Graphic matroids (Korula and Pál, 2009), co-graphic matroids (Soto, 2011), laminar matroids (Ma et al., 2016), regular and max-flow-min-cut matroids (Dinitz and Kortsarz, 2014), general matroids (Lachish, 2014; Feldman et al., 2015). At a high level, all of these algorithms do the following: process a constant fraction of the input. Then, possibly randomly, partition the remaining elements (without seeing them) into disjoint sets \mathcal{U}_1,\ldots and put restricted feasibility constraints \mathcal{J}_i on \mathcal{U}_i such that for any sets $S_i \in \mathcal{J}_i$ we have $\cup_i S_i \in \mathcal{J}$. For some algorithms, the decomposition is chosen ahead of time without even needing to see the samples, and for some the decomposition depends on the ignored elements. In some, no decomposition is necessary. In each case, each \mathcal{J}_i is simple enough that some greedy-like algorithm (accept every element that is in the max-weight basis of elements revealed so far) gets a constant factor approximation, no matter the order. The analysis looks similar to that of S_{Rank-1} , but is obviously more technical. In some cases, the \mathcal{J}_i are actually rank-1 matroids, in which case S_{Rank-1} is exactly the greedy-like algorithm used.

Corollary 3. The following single-sample prophet inequalities exist:

- Bipartite matchings: a $\frac{1}{256}$ -approximation based on Feldman et al. (2018).
- Graphic matroids: a $\frac{1}{8}$ -approximation based on Korula and Pál (2009).
- Transversal matroids: a 1/16-approximation based on Dimitrov and Plaxton (2012).
 Co-graphic matroids: a 1/12-approximation based on Soto (2011).
- Laminar matroids: a $\frac{1}{9.6}$ -approximation based on Ma et al. (2016).
- Regular and max-flow-min-cut matroids: a ¹/₃₆-approximation based on Dinitz and Kortsarz (2014).
 General matroids: a ¹/_{0 (log log k)}-approximation based on Lachish (2014), Feldman et al. (2015). If all weights are i.i.d., then the ratio improves to $\frac{1-\frac{1}{e}}{40}$ (Gharan and Vondrák, 2011).

Finally, We can make use of Corollary 3 above combined with Corollaries 1 and 2 we get the following. The mechanisms are obtained by running the prophet inequalities guaranteed by Corollary 3 above, and then applying a lazy sample reserve. Remember that all of the above prophet inequalities only use samples from bidders who are unallocated, so we can use the remaining samples from \mathcal{D} to act as the sample reserves.

Corollary 4. For each of the settings below, when all bidders' values are MHR, there exists a single-sample OPM with the following guarantees:

- Bipartite matchings: a revenue competitive ratio of $\frac{1}{256e}$ and welfare competitive ratio of $\frac{1}{256}$.
- Graphic matroids: a revenue competitive ratio of $\frac{1}{16e}$ and welfare competitive ratio of $\frac{1}{16}$.
- Transversal matroids: a revenue competitive ratio of $\frac{1}{32e}$ and welfare competitive ratio of $\frac{1}{32}$.
- Co-graphic matroids: a revenue competitive ratio of $\frac{1}{24e}$ and welfare competitive ratio of $\frac{1}{24}$. Laminar matroids: a revenue competitive ratio of $\frac{1}{19.2e}$ and welfare competitive ratio of $\frac{1}{19.2e}$.
- Regular and max-flow-min-cut matroids: a revenue competitive ratio of $\frac{1}{72e}$ and welfare competitive ratio of $\frac{1}{72}$.
- General matroids: a revenue competitive ratio and welfare competitive ratio of $\frac{1}{O(\log \log k)}$

If instead bidders' values are i.i.d. and regular, there exists a single-sample OPM for all matroids with revenue competitive ratio $\frac{1-1/e}{80a}$ and welfare competitive ratio $\frac{1-1/e}{80}$.

Note that by Theorem 1, this immediately implies mechanisms for BMUMD in the same settings with the same guarantees against the optimal deterministic DSIC mechanism, and with all guarantees degraded by a factor of four against the optimal randomized BIC mechanism.

6. Single-sample prophet inequalities for k-uniform matroids

In this Section we provide an asymptotically optimal single-sample prophet inequality for k-uniform matroids. Note that Kleinberg's algorithm (Kleinberg, 2005) for k-uniform matroids is not order-oblivious (and seems quite far from modifications that would make it so), so we cannot simply plug into the reduction from Section 5. Note again that Alaei's prophet inequality already obtains the optimal asymptotic competitive ratio (Alaei, 2011), so the contribution of this Section is the algorithm's simplicity and need for only a single sample from \mathcal{D} .

6.1. 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 $k - 2\sqrt{k}$ 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 7 and 8.)

We describe the algorithm formally below.

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

1. Offline Phase

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

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

2.a Let v_{i_i} be the value of the *j*th revealed item. Let α be an index such that $T_{\alpha-1} > v_{i_i} > T_{\alpha}$.

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_i} . Let m = $\min\{S \cap \{\alpha, \ldots, k\}\}$. This is the first slot that could be occupied by v_{i_i} .

2.c If $S \cap \{\alpha, \ldots, k\}$ is empty, reject v_{i_i}

2.d If $S \cap \{\alpha, \ldots, k\}$ is not empty, accept v_{i_i} and update $S \leftarrow S - m$.

Theorem 4. The rehearsal algorithm obtains a competitive ratio of $1 - O(\frac{1}{\sqrt{k}})$ for k-uniform matroids.

The proof of Theorem 4 is more involved than those of the previous Sections, so we split the analysis into parts below. In Section 6.2, we reduce the analysis of the rehearsal algorithm to a question about mildly correlated random walks. Section 6.3 is purely an analysis of this random walk, showing it has the desired properties. Before getting into the proof, we state the Theorems implications for mechanism design:

Corollary 5. Let \mathcal{J} be a k-uniform matroid. Then there exists a two-sample OPM guaranteeing a revenue competitive ratio of $\frac{1-O(1/\sqrt{k})}{2e}$ and welfare competitive ratio of $\frac{1-O(1/\sqrt{k})}{2}$ when all \mathcal{D}_i are MHR. If all \mathcal{D}_i are i.i.d. and regular, then the OPM guarantees a revenue and welfare competitive ratio of $\frac{1-O(1/\sqrt{k})}{2}$.

6.2. Part I: the worst adversarial ordering and connection to random walks

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, \ldots, T_k are set.)

Lemma 4. For any vector of values $v = (v_1, v_2, ..., v_n)$, and any thresholds $T_1, ..., T_k$, 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, \ldots, v_n and T_1, \ldots, T_n and assume w.l.o.g. that $v_1 < \ldots < v_n$. Also, say there exist 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 $|S_j| \ge 2$. 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 4, 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 $s = (s_1, ..., s_n)$ is drawn from \mathcal{D} and thresholds $T_1, ..., T_k$ are set. Then, values $v_1, ..., v_n$ 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 $y_1, y'_1, \ldots, y_n, y'_n$, 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 $\sum_{j=1}^{2n} p_j \le k$. Furthermore, each Y_j has a $\frac{1}{2}$ chance of being a "sample" and thus the prophet will never choose it. Thus $p_j \le \frac{1}{2}$ for all j. Since $Y_1 \ge ... \ge Y_{2n}$, these constraints imply that $\sum_{j=1}^{2n} p_j Y_j \le \sum_{j=1}^{2k} \frac{1}{2} Y_j$.

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 $\sum_{j=1}^{2k} q_j Y_j \ge \frac{c}{2} \sum_{j=1}^{2k} Y_j$ for $c = 1 - O(\frac{1}{\sqrt{k}})$. In fact, a sufficient condition for this is that $\sum_{j=1}^{i} q_j \ge ci/2$ for all $i \le 2k$.¹³

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". More formally, RW's definition is as follows:

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."

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. Fig. 1 illustrates these facts, assigning different colors to accepted and rejected values, as well as filled and unfilled thresholds.

Definition 3. For any *j*, $H_j^R(RW)$ is the height of RW to the right of *j*. Or formally, $H_j^R(RW) = \max_{i \ge j} \{RW(i) - RW(j)\}$. Similarly, $H_i^L(RW)$ is the height of RW to the left of *j*. Formally, $H_i^L(RW) = \max_{i \le j} \{RW(i) - RW(j)\}$.

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 1. Assuming that the v_i are revealed so that the values are in increasing order, for all j, Y_j is chosen by the rehearsal algorithm if and only if Y_j is a "value" and $H_i^R > 0$.

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 $(Y_{j+1}, ..., Y_i)$. This

¹³ It is easy to see that minimizing $\sum_{i} q_{j} Y_{j}$ subject to this condition will set $q_{j} = c/2$ for all $j \le 2k$.



Fig. 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 $H_i^L - H_i^R$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

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_i , and is therefore selected.

If $RW(i) \le RW(j)$, then there are at least as many "values" as there are thresholds in the list (Y_{j+1}, \ldots, 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 $RW(i) \le RW(j)$ 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 2. For all *i*, the number of "values" in $\{Y_1, \ldots, Y_i\}$ that are *not* selected by the rehearsal algorithm is max $\{H_i^L - H_i^R, 0\}$.

Proof. Let j_1, \ldots, j_h denote the indices of the "values" in (Y_1, \ldots, 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 $H_i^L - H_i^R \ge h$, and then showing that $H_i^L - H_i^R \le h$.

For any index k in $\{1, ..., h\}$, Y_{j_k} is not selected. Thus, Fact 1 tells us that it must be the case that $RW(z) \le RW(j_k)$ for all $z \ge j_k$. 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 $RW(j_{k+1}) \le RW(j_k) - 1$. Chaining this together for all k in $\{1, ..., h\}$, we get that $RW(j_h) \le RW(j_1) - (h-1)$. Because j_1 is a "value", $RW(j_1) = RW(j_1 - 1) - 1$, which means that we get $RW(j_h) \le RW(j_1 - 1) - h$.

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

Let $H = H_i^L - H_i^R$. We will show that $H \le h$, thus concluding the proof. Since $H_i^L = H_i^R + H$, there exists an index $j \in \{1, ..., i\}$ such that $RW(j) = RW(i) + H_i^R + H$. So, for every k in $\{1, ..., H\}$, choose j_k to be the largest index in $\{1, ..., i\}$ such that $RW(j_k - 1) \ge RW(i) + H_i^R + k$. By this definition, we have $RW(j_k) < RW(i) + H_i^R + k \le RW(j_k - 1)$, 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 $RW(j) \le RW(j_k)$ by the definition of j_k (otherwise j_k would not be the largest index satisfying $RW(j_k - 1) \ge RW(i) + H_i^R + k$). Furthermore, for every index $j \ge i$, we have $RW(j) \le RW(i) + H_i^R < RW(i) + H_i^R + k \le RW(j_k - 1) = RW(j_k) + 1$. Thus, we have $RW(j) \le RW(j_k)$ for every $j > j_k$. By Fact 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 $\{1, ..., i\}$ that are not selected by the rehearsal algorithm is exactly $H_i^L - H_i^R$.

The expected number of "values" in $\{Y_1, \ldots, Y_i\}$ is $\frac{i}{2}$. By Fact 2, we have that the expected number of values in $\{Y_1, \ldots, Y_i\}$ 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 $\sum_{j=1}^{i} q_j \ge \frac{ci}{2}$ 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

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

Our next subsection is dedicated to proving this inequality.

6.3. Rehearsal algorithm analysis. Part II: bounding the height of the random walk

In light of the previous Section, we have reduced the analysis of the rehearsal algorithm to proving the following Theorem.

Theorem 5. $\mathbb{E}[\max\{H_i^L - H_i^R, 0\}] \le O(\frac{i}{\sqrt{k}}) \forall i \le 2k$, where the constant implicit in the $O(\cdot)$ notation is the same for all *i*.

Recall that our random walk is non-traditional in two ways. First, after $k - \sqrt{2k}$ positive steps, the random walk jumps an additional $2\sqrt{k} + 1$ units. Second, the steps of the random walk are slightly correlated. In each pair $y_i, y'_i \leftarrow D_i$, exactly one induces a non-negative step (by being a "sample") and the other one must induce a negative step (by being a "value"). Thus, the steps in the random walk are correlated. Our proof of Theorem 5 accounts for these obstacles using the following steps.

- 1. We show that for large *i* we in fact have $\mathbb{E}[H_i^L] \leq O(i/\sqrt{k})$. It is clear that $\mathbb{E}[H_i^L] \geq \mathbb{E}[\max\{H_i^L H_i^R, 0\}]$, so this is enough. We prove this by first observing that if there were no correlation between steps and no jump, then this is a well-known fact about the expected height of random walks. Then we show that the jump and correlation can only decrease $\mathbb{E}[H_i^L]$.
- 2. The analysis is made difficult by the fact that RW jumps up at a random location. To circumvent this difficulty, we will describe a new random walk RW' that jumps up at a fixed index instead of after the $(k 2\sqrt{k})$ th threshold seen. For all small *i*, it will be clear that $H_i^L(RW) = H_i^L(RW')$, and we will show that $H_i^R(RW') \le H_i^R(RW)$ with very high probability. (The probability that $H_i^R(RW') > H_i^R(RW)$ is inversely exponential in *k*.) As $H_i^R(RW)$ is clearly at most *k*, this means that for small *i*, we only have to bound $\mathbb{E}[\max\{H_i^L(RW') H_i^R(RW'), 0\}]$, which is still challenging but much cleaner.
- 3. We show in *RW'* that for small *i* and j < i, $H_j^R = 0$ with low probability. We first prove that this is true if there was no correlation, and show that correlation can only decrease the probability that $H_j^R = 0$. By Facts 1 and 2, this exactly says that $\mathbb{E}[\max\{H_i^L H_i^R, 0\}]$ is small.

We now proceed to show step 1, that for all $i \ge k/2$, $\mathbb{E}[H_i^L] \le O(i/\sqrt{k})$. First, it is clear that the jump cannot possibly increase $\mathbb{E}[H_i^L]$, because for all j < i, either the jump does not affect RW(j) - RW(i), or it decreases RW(j) - RW(i) by $2\sqrt{k} + 1$. So we may ignore the jump as doing so only increases $\mathbb{E}[H_i^L]$. Next, it is clear that if there is no correlation between steps to the left of *i*, then H_i^L is just the height of a truly random walk starting at *i* going back to 0. It is a well-known consequence of the reflection principle that the expected height of a random walk on *i* steps is $O(\sqrt{i})$, see e.g. Feller (1968). Because $i \ge k/2$, this would exactly say that $\mathbb{E}[H_i^L] \le O(i/\sqrt{k})$. Now we just have to show that the same bound holds even if there are correlated pairs before *i*. To do this, we show that for any pair of correlated steps, decorrelating them only increases $\mathbb{E}[H_i^L]$, regardless of any other correlation. We can then apply this argument a finite number of times, decorrelating every pair of correlated steps to increase $\mathbb{E}[H_i^L]$ to a value that is $O(i/\sqrt{k})$ by our previous observation. Therefore, it must be the case that $\mathbb{E}[H_i^L] \le O(i/\sqrt{k})$.

Lemma 5. Let RW be any random walk of n steps where steps x and y are negatively correlated random variables, each uniformly distributed in $\{\pm 1\}$. Consider modifying RW by replacing steps x, y with i.i.d. uniform samples from $\{\pm 1\}$ that are independent of the other steps in RW. This modification cannot decrease the expected height of RW, even if there are other correlated steps in RW.

Proof. Imagine that the random walk is fixed except for what happens at x and y. Then this random walk has a height. And we can consider how the height is expected to change by filling in what happens at x and y if they are correlated and decorrelated respectively. We just need to show that the expected change is greater when x and y are decorrelated.

Imagine in this fixed random walk that we have removed the step at *x* and at *y*. Or in other words, the random walk stays level at these steps. Then let *a* denote the height of the peak before *x*, *b* the height of the peak between *x* and *y*, and *c* the height of the peak after *y*. If there are no steps in the walk in any of these positions, then the value of the appropriate variable is $-\infty$. We then consider adding in steps at *x* and *y* (i.e. changing the fixed walk from staying level at these two points to taking a genuine step). We consider what happens when the two steps are correlated and uncorrelated, showing that no matter what relations are satisfied by *a*, *b*, *c* that if *x* and *y* are uncorrelated, the expected height is always greater. There are several different cases to consider, but they are all simple.

Case 1: a = b = c. If x and y are correlated, we change b to b - 1 and b + 1 each with probability 1/2, and don't change c. So with probability 1/2 we increase the height by 1, with probability 1/2 it is unchanged. If x and y are uncorrelated, with probability 1/4 we increase c by 2 and b by 1. With probability 1/4 we leave c unchanged and increase b by 1. With probability 1/4 we decrease b by 1 and leave c unchanged, and with probability 1/4 we decrease b by 1 and c by 2. So with probability 1/4 we increase the height by 1, with probability 1/4 we increase it by 2, and with probability 1/2 we leave it unchanged.

Case 2: $a \le b < c$. If x and y are correlated, they cannot change the height ever. If x and y are uncorrelated, we increase the height by 2 with probability 1/4 and decrease the height by 2 (or 1 if c = b + 1) with probability 1/4.

Case 3: $b \le a < c$. Same as above.

Case 4: a > b, a > c. If x and y are correlated, we do not change the height ever. If x and y are uncorrelated, we never decrease the height, and sometimes may increase the height if a = c + 1.

Case 5: b > a, b > c. Whether or not x and y are correlated, we increase the height by 1 with probability 1/2 and decrease it by 1 with probability 1/2.

Case 6: a = b > c. Whether or not x and y are correlated, we increase the height by 1 with probability 1/2 and never decrease it.

Case 7: a = c > b. If x and y are correlated, we never change the height. If x and y are uncorrelated, we sometimes increase height by 2, and sometimes don't change it.

Case 8: b = c > a. If x and y are correlated, we never decrease c and increase b by 1 with probability 1/2. So the expected increase is 1/2. When x and y are uncorrelated, we increase c by 2 with probability 1/4, increase b by 1 without changing c with probability 1/4, and decrease b by 1 without changing c with probability 1/4, and decrease b by 1 and c by 2 with probability 1/4. So the expected increase is 1/2 + 1/4 - 1/4 = 1/2.

In all cases, it is easy to see that the expected increase in height when x and y are uncorrelated is at least as large as the expected increase in height when x and y are correlated. This covers all cases and does not depend on any other existing correlations in *RW*. Therefore, decorrelating steps x and y can only increase the expected height of *RW*.

Using Lemma 5 and the reasoning above, we complete step 1 of the proof with the following corollary:

Corollary 6. $\forall i \geq k/2$, $\mathbb{E}[H_i^L] \leq O(i/\sqrt{k})$.

We now shift to step 2 of the proof. First, define the following random walk RW'

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:
 - RW'(j) = RW'(j-1) 1 if Y_j is a "value" and $1 \le j < 2k 4\sqrt{k} + 2k^{2/3}$.
 - RW'(j) = RW'(j-1) + 1 if Y_j is a "sample" and $1 \le j < 2k 4\sqrt{k} + 2k^{2/3}$.
 - $RW'(j) = RW'(j-1) + \sqrt{k}$ when $j = 2k 4\sqrt{k} + 2k^{2/3}$.
 - RW'(j) = RW'(j-1) for $j > 2k 4\sqrt{k} + 2k^{2/3}$.

We can prove the following Lemma about RW'.

Lemma 6. $H_i^R(RW') \le H_i^R(RW)$ for all $i \le k/2$ with probability $1 - e^{-\Omega(k)}$.

Proof. Let i^* denote the index where *RW* shoots up by $2\sqrt{k} + 1$. We first show that with high probability both of the following events hold:

1. $2k - 4\sqrt{k} - 2k^{2/3} \le i^* \le 2k - 4\sqrt{k} + 2k^{2/3}$. 2. For all $i, j \in [2k - 4\sqrt{k} - 2k^{2/3}, 2 - 4\sqrt{k} + 2k^{2/3}]$, $RW'(i) - RW'(j) \le \sqrt{k}$.

Part 1 is a simple application of the Chernoff bound. If we are to have $i^* < T = 2k - 4\sqrt{k} - 2k^{2/3}$, then we must have seen $k - 2\sqrt{k}$ rehearsal elements by then. If we let k' denote the number of indices before T whose correlated partner also comes before T, then clearly there will be exactly k'/2 rehearsal elements from such indices. For the remaining indices, whether that element is rehearsed or real is independent of all other indices before T. The expected number of rehearsal elements from the remaining indices is exactly (T - k')/2. So in order to see at least $k - 2\sqrt{k}$, this value must deviate from its expectation by at least $k^{2/3}$. Using the additive Chernoff bound we get that:

 $Pr[more than k - 2\sqrt{k} rehearsals before T] \le 2e^{-k^{4/3}/(2T-2k')} \le 2e^{-k^{1/3}/4}$

An analogous argument holds to show that $i^* < 2k - 4\sqrt{k} + 2k^{2/3}$ with high probability by showing that the probability that we see fewer than $k - 2\sqrt{k}$ rehearsals by then is equally tiny. Therefore, using a union bound, part 1 holds with probability at least $1 - 4e^{-k^{1/3}/4}$.

Part 2 is also an application of the Chernoff bound. For any fixed i, j, the expected value of RW'(i) - RW'(j) is 0. There are some steps between *i* and *j* that are correlated, and will always cancel each other out. The remaining steps are all independent and there are at most $4k^{2/3}$ of them. So RW'(i) - RW'(j) must deviate from its expectation by at least \sqrt{k} and we can apply the Chernoff bound again to say that:

$$Pr[|RW'(i) - RW'(j)| \ge \sqrt{k}] \le 2e^{-k^{1/3}/8}$$

We can now take a union bound over all $O(k^{4/3})$ ordered pairs of i, j to get that with probability at least 1 – $8k^{4/3}e^{-k^{1/3}/8}$, $RW'(i) - RW'(j) \le \sqrt{k}$ for all *i*, *j*. So taking a final union bound gives us that with high probability parts 1 and 2 both hold.

Now let's couple RW and RW' to use the same coin flips. In other words, when Y_i is determined to be real or rehearsal, it is the same for both walks. Also assume that parts 1 and 2 hold for RW and RW' respectively. We now show that as long as these two assumptions hold, then for any $i \le k/2$, $H_i^R(RW') \le H_i^R(RW)$.

Because $i \le k/2$, it must be the case that $i < i^*$, so RW(i) = RW'(i). Let $j \ge i$ be the index maximizing RW'(j) - RW'(i). Then $H_i^R(RW') = RW'(j) - RW'(i)$. There are two cases to consider. Say $j < i^*$. Then RW'(j) = RW(j), and therefore RW(j) - RW(i) = RW'(j) - RW'(i), so we immediately get that $H_i^R(RW) \ge H_i^R(RW')$. Otherwise, $i^* \le j \le 2k - 4\sqrt{k} + 2k^{2/3}$. The $RW'(j) - RW'(i) \le 2\sqrt{k} + RW'(i^*) - RW(i)$ by our two assumptions. By the definition of RW, we also have that $RW(i^*) - RW(i) = RW'(i^*) + 2\sqrt{k} - RW(i)$, so this exactly says that $RW'(j) - RW'(i) \le RW(i^*) - RW(i)$, also giving us that $H_i^R(RW) \ge H_i^R(RW')$. It cannot be the case that $j > 2k - 4\sqrt{k} + 2k^{2/3}$ because we defined RW' to stop changing after this. So this covers every possible case, and in all cases $H_i^R(RW) \ge H_i^R(RW')$. Because our assumptions hold with high probability, so does the result.

We now finish by showing that for all $j \le k/2$, $H_i^R(RW') = 0$ with probability $O(1/\sqrt{k})$. We prove this claim in two steps. First, we show that if RW' had no correlated steps, then $H_i^R(RW') = 0$ with probability $O(1/\sqrt{k})$ for all j. Then we show that *removing* a specific correlated pair only increases the probability that $H_i^R(RW') = 0$, regardless of any other correlation in RW'. We can apply this argument a finite number of times to remove all correlated pairs without decreasing the probability that $H_j^R(RW') = 0$. Therefore, because this probability is now $O(1/\sqrt{k})$, it must be the case that $Pr[H_i^R(RW') = 0] \le O(1/\sqrt{k})$ to begin with.

We now take the first step. Let RW'' denote RW' without the \sqrt{k} jump at the end. Then in order for $H_i^R(RW') = 0$, we must have RW''(j) < RW''(i) for all j > i and $RW''(2k - 4\sqrt{k} + 2k^{2/3}) < RW''(i) - \sqrt{k}$. We show that if RW'' has no correlated steps, then both of these occur with low probability.

Lemma 7. Let RW'' be a random walk with n truly independent steps. Then for all n, the probability that H(RW'') = 0 and RW''(n) < 0 $-\sqrt{k}$ is $O(1/\sqrt{k})$.

Proof. We first compute the probability that H(RW'') > 0 and $RW''(n) \le -\sqrt{k}$ using the reflection principle. For any fixed walk with H(RW'') > 0 and $RW''(n) \le -\sqrt{k}$, let *i* be the last index with RW''(i) = 1. Consider the mapping that sets RW''(j) = 2 - RW''(j) for all j > i. This mapping is clearly injective and always has $RW''(n) \ge \sqrt{k} + 2$. In fact, the same mapping takes any fixed random walk with $RW''(n) \ge \sqrt{k} + 2$ and turns it into a random walk with H(RW'') > 0 and $RW''(n) \leq -\sqrt{k}$, thereby creating a bijection. In other words, this mapping bears evidence that $Pr[H(RW'') > 0 \land RW''(n) \leq 1$ $-\sqrt{k} = Pr[RW''(n) > 2 + \sqrt{k}].$

Furthermore, we can write $Pr[H(RW'') = 0 \land RW''(n) \le -\sqrt{k}] = Pr[RW''(n) \le -\sqrt{k}] - Pr[H(RW'') > 0 \land RW''(n) \le -\sqrt{k}]$ $-\sqrt{k}$], which by the above work is exactly $Pr[RW''(n) \ge \sqrt{k}] - Pr[RW''(n) \ge 2 + \sqrt{k}] = Pr[RW''(n) \in \{\sqrt{k}, \sqrt{k} + 1\}] \approx 1$ $\binom{n}{n/2+\sqrt{k}/2}/2^n$. So now we just want to bound this value. We observe first that for all *n* that:

$$\frac{\binom{n+2}{n/2+1+\sqrt{k}/2}}{2^{n+2}} = \frac{\binom{n}{n/2+\sqrt{k}/2}}{2^n} \times \frac{(n+2)(n+1)}{4(n/2-\sqrt{k}/2+1)(n/2+\sqrt{k}/2+1)} = \frac{\binom{n}{n/2+\sqrt{k}/2}}{2^n} \times \frac{n^2+3n+2}{n^2+4n+4-k}$$

In other words, for n < k - 2, the value increases when we increase n by 2. For n > k - 2, the value decreases when we increase *n* by 2. Therefore, the value is maximized around n = k, where it is obvious that $\binom{k}{k/2+\sqrt{k}/2}/2^k \le O(1/\sqrt{k})$. Therefore, for all *n*, the probability that H(RW'') = 0 and $RW''(n) \le -\sqrt{k}$ is $O(1/\sqrt{k})$.

Finally, we prove that removing the correlated pairs in RW' only increases the probability that $H_i^R = 0$:

Lemma 8. Let RW'' be a random walk on n steps where some pairs of steps $(x_1, y_1), \ldots, (x_z, y_z)$ are negatively correlated. Let $x_i < y_i$ for all i and $y_1 < \ldots < y_z$. Then removing x_1, y_1 from RW'' only increases the probability that H(RW'') = 0 and $RW''(n) \le -m$, for all n, m.

Proof. Observe first that we are not claiming that removing any correlated pair can only increase this probability, but that there is always a "correct" pair that we can remove without decreasing the probability. For a fixed random walk, imagine removing steps x_1 and y_1 (i.e. don't move at these steps). Then let *a* denote the height of the highest peak before x_1 , *b* denote the height of the highest peak between x_1 and y_1 , *c* denote the height of the highest peak after y_1 , and *d* the value of RW''(n). Also let S(a, b, c, d) denote the set of all instances of RW'' that respect the correlation between the pairs of steps (x_2, y_2) through (x_z, y_z) with respective peak heights *a*, *b*, *c* and also satisfy RW''(n) = d. Then every instance of RW'' is in exactly one set, and whether or not H(RW'') = 0 and $RW''(n) \le -m$ depends only on which S(a, b, c, d) the instance is in. We now want to look at which sets will satisfy this regardless of how steps x_1 and y_1 are set, and which sets may or may not satisfy it depending on how x_1 and y_1 are set.

We observe that setting x_1 and y_1 can never change a, c, or d, but may increase or decrease b by 1. So if a > 0, b > 1, c > 0, or d > -m, then we will never have H(RW'') = 0 and $RW''(n) \le -m$ no matter how x_1, y_1 are set. Likewise, if we have $a \le 0, b < 0, c \le 0$, and $d \le -m$, then we will always have H(RW'') = 0 and $RW''(n) \le -m$ no matter how x_1, y_1 are set. The interesting cases are when we have $a \le 0, c \le 0, d \le -m$ and $b \in \{0, 1\}$. If we remove x_1 and y_1 , then all of these cases with b = 1 will not have H(RW'') = 0, and those with b = 0 will. If we keep x_1 and y_1 , then exactly half of both cases will have H(RW'') = 0. We show that there are more of the latter case than the former. In other words, if we removed x_1 and y_1 , instead of splitting these cases 50–50, more of them would yield H(RW'') = 0 and $RW''(n) \le -m$. Therefore removing x_1 and y_1 only increases the probability that H(RW'') = 0 and $RW''(n) \le -m$. We prove this by giving an injective map from the former case to the latter.

Consider any instance of RW'' in S(a, 1, c, d) with $a \le 0$. Let *i* denote the first index after x_1 with RW''(i) = 1. Then it must be the case (because $a \le 0$) that RW''(i - 1) = 0. So consider changing RW'' to take a step down at *i* instead of up (i.e. set RW''(i) = -1). If *i* was part of a correlated pair, then also change RW'' to take a step up at its partner, *j*. It is clear that we have not changed *a*. We might have decreased *c* by 2, 1, or 0, depending on if *i* was part of a correlated pair and where its partner was located, and we might have decreased *d* by 2 or 0, depending on if *i* was part of a correlated pair. Furthermore, this map is injective. Observe first that we can determine the index *i* of the instance of RW'' where the flip happened by looking at its image under the map. A priori, *i* could be any index between x_1 and y_1 with RW''(i - 1) = 0 and RW''(i) = -1 in the image. Then the pre-image would have taken a step up at *i* instead of down, and we would have had RW''(i' - 1) = 2 in the pre-image, meaning that the instance was not in S(a, 1, c, d). Even if *i* was part of a correlated step, by our choice of x_1, y_1 , its partner *necessarily occurs after* y_1 , and therefore will not cancel out the change from switching RW''(i) by the time we take step i' - 1. Since we can determine the index *i* from the image, and it is obvious that if two instances of RW'' have the same image and had the same step switched they must be the same, this map is injective. Finally, the map only decreases *c* and *d*. So in particular, if:

$$S_1 = \bigcup_{a \le 0, c \le 0, d \le -m} S(a, 1, c, d)$$

$$S_0 = \bigcup_{a < 0, c < 0, d < -m} S(a, 0, c, d)$$

then we have shown an injective map from S_1 to S_0 . Also denote by S_2 all other instances of RW'' with H(RW'') = 0 and $RW''(n) \le -m$, and S_3 the remaining instances of RW''. Then the probability that H(RW'') = 0 and $RW''(n) \le -m$ when we have removed x_1 and y_1 is exactly:

$$\frac{|S_0| + |S_2|}{|S_0| + |S_1| + |S_2| + |S_3|}$$

And the probability that H(RW'') = 0 and $RW''(n) \le -m$ when we keep x_1 and y_1 is exactly:

$$\frac{|S_0|/2 + |S_1|/2 + |S_2|}{|S_0| + |S_1| + |S_2| + |S_3|}$$

By showing an injective map from S_1 to S_0 , we have shown that the first probability is greater. Namely, removing x_1 and y_1 can only increase the probability that H(RW'') = 0 and $RW''(n) \le -m$.

Now by Lemma 8, we can continue removing the earliest-ending correlated pair from RW' until we get a random walk with truly independent steps (and \sqrt{k} jump at the end) whose probability of having $H(RW') \ge 0$ has only increased. By Lemma 7, we know that this value is $O(1/\sqrt{k})$. So together, this says that $Pr[H_j^R(RW') = 0] \le O(1/\sqrt{k})$ for all $j \le k/2$. Finally, by Lemma 6 and the fact that $H_i^R(RW) \le k$ always, we get that $Pr[H_j^R(RW) = 0] \le O(1/\sqrt{k})$. This exactly says that the expected number of $j \le i$ with $H_j^R(RW) = 0$ is $O(i/\sqrt{k})$ for all $i \le k/2$. By Facts 1 and 2 we now have that $\mathbb{E}[\max\{H_i^L(RW) - H_i^R(RW), 0\}] \le O(i/\sqrt{k})$.

So now we have shown that for all $i \le 2k$, $\mathbb{E}[\max\{H_i^L(RW) - H_i^R(RW), 0\}] \le O(i/\sqrt{k})$, completing the proof of Theorem 5, and proving that the rehearsal algorithm obtains a competitive ratio of $1 - O(1/\sqrt{k})$.

7. Conclusion

This paper proposes the first prior-independent mechanisms for multi-dimensional settings with asymmetric bidders. Our results also yield new prior-independent mechanisms when valuation distributions are regular, independent and identical. Our approach is quite general, and builds off of previously developed tools of Chawla et al. (2010), Hartline and Roughgarden (2009), Dhangwatnotai et al. (2010). Our main technical workhorse is the development of limited information prophet inequalities.

There are two exciting directions for future work. The first is to extend our mechanism design framework to accommodate further generalizations (for instance, asymmetric regular distributions), or to identify formal barriers that would prevent these results from being generalized. The second is to design new single-sample prophet inequalities in settings where full-information prophet inequalities exist, but where algorithms for the secretary problem are unknown, such as arbitrary matroids.

Appendix A. The free-order model and SPMs

In this Section, we provide an improved and simplified analysis of the secretary algorithm in the free-order model proposed by Jaillet et al. (2012), and show that it implies a single-sample SPM for all matroids. Note that Yan (2011) has already shown a (1 - 1/e)-approximate SPM for all matroids, so the contribution of the work below is that the SPM requires just a single sample. Let's first recall their algorithm:

- 1. Initialize the set of accepted elements, A, to \emptyset .
- 2. Sample k = 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, \ldots, 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 $\{X_1, \ldots, X_m\}$, process them as in step 5.

We first recall a Lemma from Jaillet et al. (2012):

Lemma 9. (Jaillet et al., 2012) 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 4. Let $Z_1, \ldots, Z_{m'}$ list elements of *S* in decreasing order of weight for any $S \subseteq U$. Let i(y) be the minimum *i* such that $y \in \text{span}(\{Z_1, \ldots, Z_i\})$ (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 10. For all $y \in U$, 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 6. The algorithm of Jaillet et al. (2012) 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 9 and 10, every element in the max-weight basis of 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.

It is easy to see that this algorithm implies a single-sample prophet inequality for all matroids where the gambler gets to choose the order, and that no samples from elements that can possibly be accepted are used. Therefore, we immediately get the following SPM:

Theorem 7. Let \mathcal{J} be any matroid and let each \mathcal{D}_i be MHR. The 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 and welfare competitive ratio of $\frac{1}{8}$.

References

Alaei, Saeed, 2011. Bayesian combinatorial auctions: expanding single buyer mechanisms to many buyers. In: FOCS'11.

Archer, A., Tardos, É., 2001. Truthful mechanisms for one-parameter agents. In: 42nd IEEE Symposium on Foundations of Computer Science, 2001, Proceedings. IEEE, pp. 482–491.

Azar, Pablo, Daskalakis, Constantinos, Micali, Silvio, Weinberg, Matt, 2013. Optimal and efficient parametric auctions. In: SODA'13.

- Babaioff, Moshe, Immorlica, Nicole, Kempe, David, Kleinberg, Robert, 2007a. A knapsack secretary problem with applications. In: Charikar, Moses, Jansen, Klaus, Reingold, Omer, Rolim, José D.P. (Eds.), Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, 10th International Workshop, APPROX 2007, and 11th International Workshop, RANDOM 2007, Princeton, NJ, USA, August 20–22, 2007, Proceedings. In: Lect. Notes Comput. Sci., vol. 4627. Springer, Princeton, NJ, USA, pp. 16–28.
- Babaioff, Moshe, Immorlica, Nicole, Kleinberg, Robert, 2007b. Matroids, secretary problems, and online mechanisms. In: SODA'07, pp. 434-443.

Briest, Patrick, Chawla, Shuchi, Kleinberg, Robert, Weinberg, S. Matthew, 2015. Pricing lotteries. J. Econ. Theory 156, 144–174.

Bulow, Jeremy, Klemperer, Paul, 1996. Auctions versus negotiations. Am. Econ. Rev. 86 (1), 180-194.

- Cai, Yang, Devanur, Nikhil R., Weinberg, Matthew S., 2016. A duality based unified approach to bayesian mechanism design. In: Wichs, Daniel, Mansour, Yishay (Eds.), Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing. STOC 2016, Cambridge, MA, USA, June 18–21, 2016. ACM, pp. 926–939.
- Chawla, Shuchi, Hartline, Jason D., Malec, David L., Sivan, Balasubramanian, 2010. Multi-parameter mechanism design and sequential posted pricing. In: STOC'10.
- Chawla, Shuchi, Malec, David L., Sivan, Balasubramanian, 2015. The power of randomness in bayesian optimal mechanism design. Games Econ. Behav. 91, 297–317.

Clarke, Edward H., 1971. Multipart pricing of public goods. Public Choice 11 (1), 17-33.

Daskalakis, Constantinos, Deckelbaum, Alan, Tzamos, Christos, 2014. The complexity of optimal mechanism design. In: 25th ACM–SIAM Symposium on Discrete Algorithms (SODA).

- Daskalakis, Constantinos, Deckelbaum, Alan, Tzamos, Christos, 2015. Strong duality for a multiple-good monopolist. In: 16th ACM Conference on Economics and Computation (EC).
- Devanur, Nikhil, Hartline, Jason, Karlin, Anna, Nguyen, Thach, 2011. Prior-independent multi-parameter mechanism design. In: Internet and Network Economics. Springer, pp. 122–133.
- Dhangwatnotai, P., Roughgarden, T., Yan, Q., 2010. Revenue maximization with a single sample. In: EC'10. ACM, pp. 129-138.
- Dimitrov, Nedialko B., Plaxton, C. Greg, 2012. Competitive weighted matching in transversal matroids. Algorithmica 62 (1-2), 333-348.

Dinitz, Michael, Kortsarz, Guy, 2014. Matroid secretary for regular and decomposable matroids. SIAM J. Comput. 43 (5), 1807–1830.

Dynkin, E.B., 1963. The optimum choice of the instant for stopping a Markov process. Sov. Math. Dokl. 4, 627–629.

- Feldman, Moran, Svensson, Ola, Zenklusen, Rico, 2015. A simple O(log log(rank))-competitive algorithm for the matroid secretary problem. In: Indyk, Piotr (Ed.), Proceedings of the Twenty-Sixth Annual ACM–SIAM Symposium on Discrete Algorithms. SODA 2015, San Diego, CA, USA, January 4–6, 2015. SIAM, pp. 1189–1201.
- Feldman, Moran, Svensson, Ola, Zenklusen, Rico, 2018. A framework for the secretary problem on intersections of matroids. In: Indyk, Piotr (Ed.), Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA.
- Feller, William, 1968. An Introduction to Probability Theory and Its Applications, vol. 1. Wiley.

Gharan, Shayan Oveis, Vondrák, Jan, 2011. On variants of the matroid secretary problem. In: ESA'11, pp. 335-346.

- Göbel, Oliver, Hoefer, Martin, Kesselheim, Thomas, Schleiden, Thomas, Vöcking, Berthold, 2014. Online independent set beyond the worst-case: secretaries, prophets, and periods. In: Esparza, Javier, Fraigniaud, Pierre, Husfeldt, Thore, Koutsoupias, Elias (Eds.), Automata, Languages, and Programming 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8–11, 2014, Proceedings, Part II. In: Lect. Notes Comput. Sci., vol. 8573. Springer, pp. 508–519.
- Goldner, Kira, Karlin, Anna, 2016. A prior-independent revenue-maximizing auction for multiple additive bidders. In: 12th Conference on Web and Internet Economics (WINE).
- Groves, Theodore, 1973. Incentives in teams. Econometrica 41 (4), 617-631.
- Hart, Sergiu, Nisan, Noam, 2013. The menu-size complexity of auctions. In: 14th ACM Conference on Electronic Commerce (EC).
- Hart, Sergiu, Reny, Philip J., 2015. Maximizing revenue with multiple goods: nonmonotonicity and other observations. Theor. Econ. 10 (3), 893-922.

Hartline, J.D., Roughgarden, T., 2009. Simple versus optimal mechanisms. In: EC'09, pp. 225-234.

- Hill, Theodore P., Kertz, Robert P., 1983. Stop rule inequalities for uniformly bounded sequences of random variables. Trans. Am. Math. Soc. 278 (1), 197–207. Jaillet, Patrick, Soto, José A., Zenklusen, Rico, 2012. Advances on matroid secretary problems: free order model and laminar case. CoRR. http://arxiv.org/abs/ 1207.1333.
- Kesselheim, Thomas, Radke, Klaus, Tönnis, Andreas, Vöcking, Berthold, 2013. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In: Bodlaender, Hans L., Italiano, Giuseppe F. (Eds.), Algorithms – ESA 2013–21st Annual European Symposium, Sophia Antipolis. France, September 2–4, 2013. Proceedings. In: Lect. Notes Comput. Sci., vol. 8125. Springer, pp. 589–600.
- Kleinberg, R., 2005. A multiple-choice secretary algorithm with applications to online auctions. In: SODA'05, pp. 630-631.
- Kleinberg, Robert, Weinberg, S. Matthew, 2012. Matroid prophet inequalities. In: STOC'12, pp. 123-136.
- Korula, Nitish, Pál, Martin, 2009. Algorithms for secretary problems on graphs and hypergraphs. In: ICALP'09, pp. 508-520.
- Krengel, U., Sucheston, L., 1977. Semiamarts and finite values. Bull. Am. Math. Soc. 83, 745–747.
- Krengel, U., Sucheston, L., 1978. On semiamarts, amarts and processes with finite value. Adv. Probab. 4, 197-266.
- Lachish, Oded, 2014. O(loglog rank) competitive ratio for the matroid secretary problem. In: 55th IEEE Annual Symposium on Foundations of Computer Science. FOCS 2014, Philadelphia, PA, USA, October 18–21, 2014. IEEE Computer Society, pp. 326–335.
- Ma, Tengyu, Tang, Bo, Wang, Yajun, 2016. The simulated greedy algorithm for several submodular matroid secretary problems. Theory Comput. Syst. 58 (4), 681–706.

Manelli, A.M., Vincent, D.R., 2007. Multidimensional mechanism design: revenue maximization and the multiple-good monopoly. J. Econ. Theory 137 (1), 153-185.

Roughgarden, Tim, Talgam-Cohen, Inbal, Yan, Qiqi, 2012. Supply-limiting mechanisms. In: EC'12, pp. 844-861.

Samuel-Cahn, E., 1984. Comparison of threshold stop rules and maximum for independent nonnegative random variables. Ann. Probab. 12 (4), 1213–1216. Soto, José A., 2011. Matroid secretary problem in the random assignment model. In: Randall, Dana (Ed.), SODA. SIAM, pp. 1275–1284.

Vickrey, William, 1961. Counterspeculations, auctions, and competitive sealed tenders. J. Finance 16 (1), 8-37.

Yan, Qiqi, 2011. Mechanism design via correlation gap. In: Randall, Dana (Ed.), Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms. SODA 2011, San Francisco, California, USA, January 23–25, 2011. SIAM, pp. 710–719.

Myerson, Roger B., 1981. Optimal auction design. Math. Oper. Res. 6 (1), 58-73.