Polymatroid Prophet Inequalities

Paul Dütting[1] and Robert Kleinberg[2]

[1] London School of Economics, London, UK, p.d.duetting@lse.ac.uk
[2] Cornell University, Ithaca, NY, USA, rdk@cs.cornell.edu

Abstract. Prophet inequalities bound the reward of an online algorithm—or gambler—relative to the optimum offline algorithm—the prophet—in settings that involve making selections from a sequence of elements whose order is chosen adversarially but whose weights are random. The goal is to maximize total weight.

We consider the problem of choosing quantities of each element subject to polymatroid constraints when the weights are arbitrary concave functions. We present an online algorithm for this problem that does at least half as well as the optimum offline algorithm. This is best possible, as even in the case where a single number has to be picked no online algorithm can do better.

An important application of our result is in algorithmic mechanism design, where it leads to novel, truthful mechanisms that, under a monotone hazard rate (MHR) assumption on the conditional distributions of marginal weights, achieve a constant-factor approximation to the optimal revenue for this multi-parameter setting. Problems to which this result applies arise, for example, in the context of Video-on-Demand, Sponsored Search, or Bandwidth Markets.

1 Introduction

Prophet inequalities compare the performance of an online algorithm to the optimum offline algorithm in settings that involve making selections from a sequence of random elements. The online algorithm knows the distribution from which the elements will be sampled, while the optimum offline algorithm knows the sequence of sampled elements. Prophet inequalities thus bound the relative power of online and offline algorithms in Bayesian settings. Not surprisingly, they play an important role in the analysis of online and offline algorithms in these settings. Less obviously, but no less importantly, they have a growing number of applications in algorithmic mechanism design. Specifically, they have been used to design simple yet approximately optimal (=revenue maximizing) mechanisms for multi-parameter settings, in which Myerson [17]'s seminal characterization of optimal mechanisms does not apply. Revenue maximization in multi-parameter settings is considered one of the biggest challenges in this field.

A classic result of Krengel and Sucheston [15, 16] shows that when both the online algorithm and the offline algorithm get to pick exactly one element, then the online algorithm can do at least half as well as the offline algorithm. More formally, if $w_1, \dots, w_n$ is a sequence of independent, non-negative, real-valued random variables satisfying $\mathbb{E}[\max_i w_i] < \infty$, then there exists a stopping rule τ such that

$$\mathbb{E}[w_\tau] \geq \frac{1}{2} \cdot \mathbb{E}[\max_i w_i].$$

The bound is achieved, for example, by an elegant algorithm of Samuel-Cahn [18]. This algorithm chooses a threshold T such that $\Pr(\max_i w_i > T) = \frac{1}{2}$, and selects the first element whose weight exceeds this threshold. Alternatively, as described by Kleinberg and Weinberg [14], this bound can be obtained by choosing threshold $T = \mathbb{E}[\max_i X_i]/2$ and picking the first element whose weight exceeds the threshold.

Kleinberg and Weinberg [14] recently extended this result to matroid settings. In a matroid setting we are given a ground set $\mathcal{U}$ and a non-empty downward-closed family of independent sets $\mathcal{I} \subseteq 2^{|\mathcal{U}|}$ satisfying the exchange axiom: for all pairs of sets $I, J \in \mathcal{I}$ and $|I| < |J|$ there exists an element $j \in J$ such that $I \cup \{j\} \in \mathcal{I}$; a maximal element of $\mathcal{I}$ is called a *basis*. For these settings they prove that if both the online and the offline algorithm have to pick an independent set of elements, then the online algorithm again can do at least half as well as the offline algorithm. More formally, if $w_1, \ldots, w_n$ is a sequence of independent, non-negative, real-valued random variables satisfying $\mathbb{E}[\max_i w_i] < \infty$, then there is a way to pick $A \in \mathcal{I}$ in an online fashion such that

$$\mathbb{E}\left[\sum_{i \in A} w_i\right] \geq \frac{1}{2} \cdot \mathbb{E}\left[\max_{B \in \mathcal{I}} \sum_{i \in B} w_i\right].$$

A common restriction of the original result of Krengel and Sucheston and the Kleinberg and Weinberg result is that they only apply to settings with binary decisions (i.e., an element can either be picked or not).

A Prophet Inequality for Polymatroids. Our main technical contribution is a prophet inequality for settings in which the gambler and the prophet have to choose quantities of each element subject to polymatroid constraints and the weights are arbitrary concave functions. That is, we consider settings in which we are given a ground set $\mathcal{U}$ and a submodular[3] set function $f : 2^{\mathcal{U}} \to \mathbb{R}$ and a vector of quantities $z \in \mathbb{N}^{|\mathcal{U}|}$ is feasible if $z \in P_f = \{q \in \mathbb{N}^{|\mathcal{U}|} \mid \sum_{u \in S} q(u) \leq f(S) \text{ for all } S \subseteq \mathcal{U}\}$. We will restrict ourselves to integer quantities and integer-valued set functions for ease of exposition; our results trivially extend to rational quantities and rational-valued functions by scaling. For this setting we prove that if the goal of the online and the offline algorithm is to maximize $\sum_{u \in \mathcal{U}} w(u, z(u))$ over feasible z, and if the w's are random concave weights chosen independently for each element, then the online algorithm can do at least half as well as the offline algorithm.

More formally, we show that if $w_1, \ldots, w_n$ is a sequence of independent, non-negative, real-valued concave weight functions for elements $u_1, \ldots, u_n$, then there exists a way to choose a feasible $z = (z_1, \ldots, z_n)$ in an online fashion (i.e.,,

[3] A set function f is submodular if for all $X \subset Y \subseteq \mathcal{U}$, $f(X \cup Y) + f(X \cap Y) \leq f(X) + f(Y)$.

choosing z_i when $w_1, \ldots, w_i$ have been revealed but $w_{i+1}, \ldots, w_n$ have not yet been revealed) such that

$$\mathbb{E}\left[\sum_{i=1}^{n} w(u_i, z_i)\right] \geq \frac{1}{2} \cdot \mathbb{E}\left[\max_{q \in P_f} \sum_{i=1}^{n} w(u_i, q_i)\right].$$

Our result contains the previous results as a special case, and is best possible as even in the case where a single element has to be picked no online algorithm can do better.

To prove this result we apply a known reduction from polymatroids to matroids (see, e.g., Section 44.6b of [19]). Applying this reduction, we transform an input sequence to the polymatroid problem to an input sequence of the matroid problem by repeating the (element, weight) pairs in the input sequence to the polymatroid problem. While this construction turns inputs to the polymatroid problem into inputs to the matroid problem, it violates the independence of weights assumption. Different matroid elements corresponding to the same polymatroid element will have identical (and hence dependent) weights.

A second potential difficulty that arises when reducing the polymatroid problem to the matroid problem in this manner is that the canonical way of doing so (by repeating elements of the ground set of the polymatroid and assigning the j-th copy of an element in the resulting matroid problem the marginal weight of the j-th unit of the corresponding element in the polymatroid setting) only leads to a meaningful interpretation if the matroid algorithm always picks contiguous elements from the beginning of each sequence of matroid elements corresponding to the same polymatroid element.

The Kleinberg-Weinberg algorithm does not apply to dependent weights and it also does not necessarily pick consecutive matroid elements. Our main technical workhorse is therefore a novel algorithm for the matroid setting that is capable of handling the dependencies resulting from the reduction, and that ensures that a solution to the matroid problem can be meaningfully translated back to the polymatroid setting. To ensure the latter our algorithm sets increasing thresholds within each block of elements corresponding to the same polymatroid element, and accepts an element precisely if the weight of that element passes the threshold. Once an element fails to pass the threshold it "freezes" the threshold at the current niveau. It thereby ensures that subsequent elements will not be selected as their weight can only be lower. We control for the former, i.e., the potential dependencies across weights of matroid elements corresponding to the same polymatroid element, by introducing the notion of surrogate thresholds and performing large parts of the analysis using these surrogate thresholds as proxies.

Truthful Mechanisms with Near-Optimal Revenue. The most important implication of our prophet inequality result are novel, truthful mechanisms that achieve constant-factor approximations to the optimal revenue for a multi-parameter mechanism design problem. The problem to which our mechanisms apply is multi-parameter as each agent can receive multiple units, and can have arbitrary concave valuations. The requirement that the valuations are concave corresponds to the standard economic assumption that valuations have decreasing marginals. Like prior results our mechanisms are posted-price mechanisms; that is, they approach the agents in turn and present them with a price that the agents can either accept or not. However, prior results that have used prophet inequalities to devise posted-price mechanisms were restricted to unit-demand settings (e.g., [5, 1, 14]). To the best of our knowledge, our posted price mechanisms are the first such mechanisms for a multi-unit demand setting, and yield the first constant-factor revenue guarantees for problems with polymatroid structure and valuations with decreasing marginals.

In a Bayesian mechanism design problem with polymatroid structure we are given a set N of n agents. Each agent i has a private, concave valuation function $v_i : \mathbb{N} \to \mathbb{R}_+$, drawn independently from not necessarily identical distributions F_i with support V_i that are common knowledge. A mechanism (x, p) consists of an outcome rule $x : \prod_i V_i \to \mathbb{R}^n_+$, where x_i specifies how much service agent i gets, and a payment rule $p : \prod_i V_i \to \mathbb{R}^n_+$, where p_i specifies the payment of agent i. An outcome is feasible if $\sum_{i \in S} x_i \leq f(S)$ for all $S \subseteq N$, where f is an integer-valued submodular function. Agent i's utility is $u_i(b, v_i) = v_i(x_i(b)) - p_i(b)$, where b denotes the bids of the agents. The welfare of a mechanism is $\sum_{i \in N} v_i(x_i(b))$ and its revenue is $\sum_{i \in N} p_i(b)$. A mechanism is dominant strategy incentive compatible (DSIC) (or truthful) if for every agent i, value v_i, bid b_i and bids $b_{-i} = (b_1, \ldots, b_{i-1}, b_{i+1}, \ldots, b_n)$, $u_i((v_i, b_{-i}), v_i) \geq u_i((b_i, b_{-i}), v_i)$.

Practical mechanism design problems with polymatroid structure include:

(1) **Video-on-Demand** [4]: Consider a collection of spatially dispersed user groups, each of which wants to watch various movies using a streaming service. We can model this via a graph $G = (\bigcup_{i \in N} T_i \cup \{s\}, E)$ in which $T_i \cap T_j = \emptyset$ for all $i, j \in N$ and each edge $e \in E$ has a capacity c_e. The seller is identical with source node s. Each agent $i \in N$ is identified with a number of demand nodes T_i corresponding to the members of user group i. The allocation to each agent $i \in N$ is $\sum_{t \in T_i} x_t$, where x_t is the flow into t. An allocation x is feasible if and only if $\sum_{t \in S} x_t \leq f(S)$ for all $S \subseteq \bigcup_{i \in N} T_i$, where f is the submodular function giving the value of a minimum s-S-cut.

(2) **Local Purchasing Collectives** [3]: Consider a group of buyers which is interested in purchasing a certain good from local providers. We can model this via a bipartite graph with vertices on the left side representing providers and those on the right side representing buyers (elements of N). An edge represents a buyer having access to a particular provider. Suppose that each provider j has a positive supply $s(j)$. A vector of quantities purchased is feasible if each buyer's quantity can be fulfilled by one or more of the adjacent providers without exceeding any provider's supply. More formally, for a set S of buyers let $\Gamma(S)$ denote the set of providers adjacent to at least one element of S, and let $f(S) = \sum_{j \in \Gamma(S)} s(j)$ which is a submodular function. A vector x of quantities purchased is feasible if and only if for every set of buyers S, we have $\sum_{i \in S} x_i \leq f(S)$.

(3) **Sponsored Search** [10]: In sponsored search a set of advertisers seeks to be assigned clicks on ad slots. Denote the set of advertisers by N and the

set of ad slots by M. Sort the ad slots $j \in M$ by non-increasing number of clicks $\alpha_j \in \mathbb{N}_{\geq 0}$. An allocation x of clicks to advertisers is feasible if and only if $\sum_{i \in S} x_i \leq f(S)$ for all $S \subseteq N$, where $f(S) = \sum_{j=1}^{|S|} \alpha_j$ is a submodular function.

(4) **Bandwidth Markets** [4]: In wireless communication settings agent $i \in N$ seeks to maximize its transmission rate x_i. In a Gaussian multiple-access channel the set of feasible transmission rates x—the so-called Cover-Wyner region—forms a polymatroid (see [20] for details).

We present two DSIC posted-price mechanisms for these problems. The first combines the thresholds of our algorithm with "eager" reserves, the second combines them with "lazy" reserves [7]. The difference between eager and lazy reserves is that the former are applied during the computation of the allocation, while the latter are applied only after the fact. In our case, however, both can be implemented in an online fashion. We prove that these mechanisms achieve at least a $1/2e^2$ resp. $1/2e$ fraction of the optimal revenue by proving a lower bound in terms of the optimal welfare. For "eager" reserves we use Chebyshev's Integral Inequality and inductively apply a single-sample argument of [7]. For "lazy" reserves we only need the single-sample result.

Related Work. We have already described the result by Krengel and Sucheston [15, 16] for the case in which both the online algorithm and the offline algorithm are allowed to pick one number, showing that the online algorithm can do at least half as well as the offline algorithm. This bound is tight. The result has been extended to the case where both the online algorithm and the offline algorithm can pick k numbers by Alaei [1], showing that the online-to-offline ratio is at least $1 - 1/(\sqrt{k+3})$. This matches the aforementioned tight bound when $k = 1$, and it remains nearly tight for $k > 1$, in the sense that a ratio of $1 - o(1/\sqrt{k})$ is known to be unattainable. Kleinberg and Weinberg [14] have extended the bound of two to settings where the elements picked must be an independent set in a matroid. This bound is also tight, as it subsumes the case where both the online and offline algorithm have to pick one number.

Hajiaghayi et al. [13] observed the following relationship between prophet inequalities and algorithmic mechanism design: algorithms used to prove prophet inequalities can be interpreted as truthful online auction mechanisms, and the prophet inequality in turn can be interpreted as the mechanism's approximation guarantee. Chawla et al. [5] observed an even subtler relationship between the two topics: questions about the approximability of offline Bayesian optimal mechanisms by sequential posted-price mechanisms could be translated into questions about prophet inequalities, via the use of virtual valuation functions. Alaei [1] and Kleinberg and Weinberg [14], armed with stronger prophet inequalities, deepen this relationship even further. More recently, and in parallel to this work, Feldman et al. [9] have designed posted-price mechanisms for combinatorial auctions, which are not based on prophet inequalities.

Another related line of literature is work on secretary problems, which also concerns relations between optimal offline selection rules and suboptimal online stopping rules, but under the assumption of a randomly ordered input rather than independent random numbers in a fixed order. While the polymatroid prophet inequality that we solve here contains the matroid prophet inequality problem as a special case, the matroid secretary problem introduced by Babaioff et al. [2] remains largely unsolved despite recent progress.

A final related direction is work on exponential-sized Markov decision processes (MDP's) [6, 11, 12]. The connection here is that algorithms for prophet inequalities can be formulated as exponential-sized MDP's, whose state reflects the entire set of decisions made prior to a specified point during the algorithm's execution. Most of the algorithms with provable approximation guarantees for exponential-sized MDP's are LP-based, while our algorithm is combinatorial.

2 Preliminaries

In a *Bayesian online selection problem* we are given a ground set $\mathcal{U}$ and for each $x \in \mathcal{U}$ a probability distribution F_x with support $\mathbb{R}^d_{\geq} = \{m \in \mathbb{R}^d_+ : i \leq i' \Rightarrow m_i \geq m_{i'}\}$ of finite dimension $d \in \mathbb{N}_+$. This induces a probability distribution over functions $w : \mathcal{U} \times \{1, \ldots, d\} \to \mathbb{R}_+$ in which the multivariate random variables $\{(w(x,1), \ldots, w(x,d)) : x \in \mathcal{U}\}$ are independent and the marginals $(w(x,1) - w(x,0), \ldots, w(x,d) - w(x,d-1))$ where we set $w(x,0) = 0$ have distribution F_x. We refer to $w(x,k)$ as the *weight* of k units of x, and to $w(x,k) - w(x,k-1)$ as the *marginal weight* of the k-th unit of x. By our assumption regarding the distributions F_x for $x \in \mathcal{U}$, the marginal weights $w(x,k) - w(x,k-1)$ for all $x \in \mathcal{U}$ and $k \geq 1$ are decreasing and the weight $w(x,k)$ of any given $x \in \mathcal{U}$ is a concave function in k.

The goal is to choose a vector $z \in \mathbb{R}^{|\mathcal{U}|}$ that maximizes $\sum_{x \in \mathcal{U}} w(x, z(x))$. For a given weight function w we use $\text{OPT}(w)$, or simply OPT, to denote the optimal value. The vector z will typically be restricted to come from a space of feasible vectors $\mathcal{F} \subseteq \mathbb{R}^{|\mathcal{U}|}$. One common restriction is $\mathcal{F} \subseteq \{0,1\}^{|\mathcal{U}|}$ in which case $z_i \in \{0,1\}$ can be thought of as encoding membership in a subset $A \subseteq \mathcal{U}$. Two further restrictions, matroids and polymatroids, were discussed already in Section 1. For matroids the distribution F_x for $x \in \mathcal{U}$ has dimension 1; for polymatroids defined by f taking values in $\{1, \ldots, M\}$ it suffices to consider distributions F_x for $x \in \mathcal{U}$ of dimension M.

An *input sequence* is a sequence σ of ordered pairs $(x_i, w_i)_{i=1,\ldots,|\mathcal{U}|}$ consisting of an element $x_i \in \mathcal{U}$ and marginals $w_i \in \mathbb{R}^d_{\geq}$ such that every $x_i \in \mathcal{U}$ occurs exactly once in the sequence. A *deterministic online selection algorithm* is a function z mapping every input sequence σ to a vector $z(\sigma) \in \mathcal{F}$ such that for any pair of input sequences σ, σ' that match on the first i pairs $(x_1, w_1), \ldots, (x_i, w_i)$ we have $z_j(\sigma) = z_j(\sigma')$ for all $1 \leq j \leq i$. An *online weight-adaptive adversary* that has chosen $x_1 \ldots, x_{i-1}$ and has learned about $w_1, \ldots, w_{i-1}$ chooses x_i without knowing w_i.

Notation. For a real number z, we use z^+ to denote $\max\{z, 0\}$. We use the shorthand $w(S)$ to denote the weight of a feasible (multi-)set S of elements $x \in \mathcal{U}$.

3 Algorithm for Polymatroids

We derive our algorithm for the polymatroid prophet inequality by reducing to the matroid case. We begin by defining *block-structured matroids*, *block-restricted* *weight distributions*, and *block-restricted adversaries.* Although a prophet inequality for the resulting matroid problem would translate into a prophet inequality for the polymatroid problem, we cannot simply apply the Kleinberg-Weinberg algorithm for matroids to derive it. The reason is twofold. First, the reduction from polymatroids to matroids leads to weight distributions that are no longer independent. Second, for the weights in the matroid setting to be in one-to-one correspondence to the weights in the polymatroid setting we need to ensure that the matroid algorithm chooses consecutive elements of each block. The crux of our analysis is therefore a novel algorithm for the matroid setting that can handle the dependencies that result from the reduction and that guarantees that weights are consistent.

Due to space limitations, most proofs in this section and the following one are deferred to the full version of the paper [8].

3.1 Block-structured matroids

We first define block-structured matroids and show that to every polymatroid defined by an integer-valued submodular function there is an associated block-structured matroid.

Definition 1. *A* block-structured matroid *is one whose ground set is partitioned into blocks $B_1, ..., B_n$ such that the independence relation is preserved under permutations of the ground set that preserve the pieces of the partition.*

For a set $S \subseteq B_1 \cup \cdots \cup B_n$ we define its cardinality vector $\mathbf{q}(S) = (q_1(S), q_2(S), \ldots, q_n(S))$ by setting $q_i(S) = |S \cap B_i|$ for $i = 1, \ldots, n$.

Lemma 1 (cf. Chapter 44.6b of [19]). *Suppose f is a submodular function on ground set $\mathcal{U} = \{u_1, \ldots, u_n\}$, taking values in $\{0, 1, \ldots, M\}$. There is a block-structured matroid $\mathcal{M}_f$ on ground set $\mathcal{U} \times [M]$ with blocks $B_i = \{u_i\} \times [M]$ $(i = 1, \ldots, n)$, whose independent sets are those S satisfying $\mathbf{q}(S) \in P_f$.*

Next we define block-restricted weight distributions and block-restricted adversaries to capture the type of input sequences generated by our reduction. We also define a property of algorithms for block-restricted matroids that ensures that the weights in the matroid setting can be translated back into the polymatroid setting.

Definition 2. *A* block-restricted weight distribution *on a block-structured matroid is a joint distribution of weights for its elements, such that the elements of a block receive non-increasing weights, the weights within each block can be arbitrarily correlated, but the weight assignments to different blocks are mutually independent.*

Definition 3. *A* block-restricted adversary *is one who is restricted to choose an ordering of the input sequence in which the elements of each block appear consecutively, and after any proper subset of the blocks have been presented, the choice of which block is presented next may only depend on the weights of elements that have already been presented.*

Definition 4. *A deterministic algorithm for block-restricted matroids is* consistent *if whenever it picks the j-th element of a block if it has already picked elements* $1, \dots, j-1$ *from that block. In other words a consistent algorithm for block-restricted matroids may only choose consecutive blocks of items at the beginning of each block in the matroid.*

Note that when all blocks have size 1, a block-structured matroid is simply a matroid, and a block-restricted distribution is simply an independent distribution. Furthermore, a block-restricted adversary is exactly the same as the notion of *online weight-adaptive adversary* defined in [14]. Thus, the special case in which all blocks have size 1 is precisely the setting of the matroid prophet inequality of [14].

3.2 Prophet Inequality for Block-Structured Matroids

Next we describe our algorithm for the matroid problem. The algorithm is similar in spirit to the algorithm of Kleinberg and Weinberg in that it sets a threshold for each element and accepts the element if and only if its weight exceeds the threshold. However, it significantly differs from the Kleinberg-Weinberg algorithm in the way it chooses the thresholds.

Consider a block-restricted matroid $\mathcal{M} = (\mathcal{U}, \mathcal{I})$. Let $w, w' : \mathcal{U} \to \mathbb{R}_+$ denote two assignments of weights to the elements of $\mathcal{U}$ sampled independently from a block-restricted weight distribution. For a given input sequence $\sigma = (x_1, w(x_1)), \dots, (x_n, w(x_n))$ we compare the set $A = A(\sigma)$ selected by the algorithm to the basis B that maximizes $w'(B)$. The matroid exchange axiom guarantees the existence of a partition of B into disjoint subsets C, R such that $A \cup R$ is also a basis of $\mathcal{M}$. Among all such partitions, let $C(A), R(A)$ denote the one that maximizes $w'(R)$. Let $g(A) = w'(R(A))$.

The selection algorithm when faced with element x_i proceeds as follows: Denote the (possibly empty) set of elements already selected by A_{i-1} and denote the (possibly empty) set of indices of elements belonging to the same block as x_i and that precede x_i in the input sequence by $Pred(x_i)$. Element x_i is accepted if and only if $w(x_i) \geq T_i$ where the threshold T_i is determined as follows. If $A_{i-1} \cup \{x_i\} \notin \mathcal{I}$ then $T_i = \infty$. Otherwise,

$$\begin{aligned} T_i &= \max\{\max_{j \in Pred(x_i)} T_j, \frac{1}{2} \cdot \mathbb{E}[g(A_{i-1}) - g(A_{i-1} \cup \{x_i\})]\} \\ &= \max\{\max_{j \in Pred(x_i)} T_j, \frac{1}{2} \cdot \mathbb{E}[w'(R(A_{i-1})) - w'(R(A_{i-1} \cup \{x_i\}))]\} \qquad (1) \\ &= \max\{\max_{j \in Pred(x_i)} T_j, \frac{1}{2} \cdot \mathbb{E}[w'(C(A_{i-1} \cup \{x_i\})) - w'(C(A_{i-1}))]\}. \qquad (2) \end{aligned}$$

Note that (1) and (2) define the same quantity: Let B be the maximum weight basis of $\mathcal{M}$ with weights w'. Then, $w'(B) = w'(C(A_{i-1})) + w'(R(A_{i-1}))$ and $w'(B) = w'(C(A_{i-1} \cup \{x_i\})) + w'(R(A_{i-1} \cup \{x_i\}))$. Equalizing and rearranging gives

$$w'(R(A_{i-1})) - w'(R(A_{i-1} \cup \{x_i\})) = w'(C(A_{i-1} \cup \{x_i\})) - w'(C(A_{i-1})).$$

Theorem 1. *For every block-restricted matroid* $(\mathcal{U}, \mathcal{I})$ *with block-restricted weight distribution there is a deterministic, consistent online selection algorithm that achieves the following performance guarantee against block-restricted adversaries:*

$$\mathbb{E}[w(A)] \geq \frac{1}{2} \cdot OPT.$$

Before we outline how this theorem can be proved, we use it to derive a prophet inequality for polymatroids.

3.3 Prophet Inequality for Polymatroids

The algorithm that achieves the prophet inequality in the polymatroid setting (with integer-valued submodular function f taking values in $\{1, \ldots, M\}$) does so by reducing the problem to the block-structured matroid setting with the matroid $\mathcal{M}_f$ defined in Lemma 1 as follows. If in the polymatroid setting the elements are presented in order $u_1, \ldots, u_n$, then the reduction constructs an input sequence in the matroid setting by presenting the elements in order $(u_1, 1), (u_1, 2), \ldots, (u_2, 1), (u_2, 2), \ldots$ (lexicographic order, $\mathcal{U}$ coordinate first). If in the polymatroid setting the marginal weights of element u_i are $w_i = (w_{i,1}, w_{i,2}, \ldots, w_{i,M})$ then element (u_i, j) is presented in the matroid setting with weight $w_{i,j}$. If the matroid algorithm, while processing elements $(u_i, 1), (u_i, 2), \ldots, (u_i, M)$, selects a subset $\{u_i\} \times S_i$, then the polymatroid algorithm when processing u_i sets $z_i = |S_i|$.

Theorem 2. *For every polymatroid* P_f *defined by a rational-valued submodular function* f *and concave weights there exists a deterministic online selection algorithm that satisfies the following performance guarantee against online weight-adaptive adversaries:*

$$\mathbb{E}\left[\sum_{i=1}^{n} w(u_i, z(u_i))\right] \geq \frac{1}{2} \cdot OPT.$$

3.4 Proof of the Block-Restricted Matroid Prophet Inequality

We start with a proposition that provides a lower bound on the sum of the thresholds of the elements that are selected by the algorithm. The proof of this proposition exploits the definition of the thresholds and, in addition, linearity of expectation and a telescoping sum.

Proposition 1. *For every input sequence* σ, *if* $A = A(\sigma)$, *then* $\sum_{x_i \in A} T_i \geq \frac{1}{2} \cdot \mathbb{E}[w'(C(A))]$.

Next we describe our main technical insight. Namely, that the thresholds within a given block have a specific form (see Figure 1 for an illustration). Specifically, consider any block consisting of elements $u_{i_0}, u_{i_0+1}, \ldots, u_{i_1-1}$. For all $i_0 \leq i \leq i_1$ define $A^i = A_{i_0-1} \cup \{x_{i_0}, \ldots, x_i\}$, and

$$t_i = \frac{1}{2} \cdot \mathbb{E}_{w'}\left[g(A^{i-1}) - g(A^i)\right], \tag{3}$$

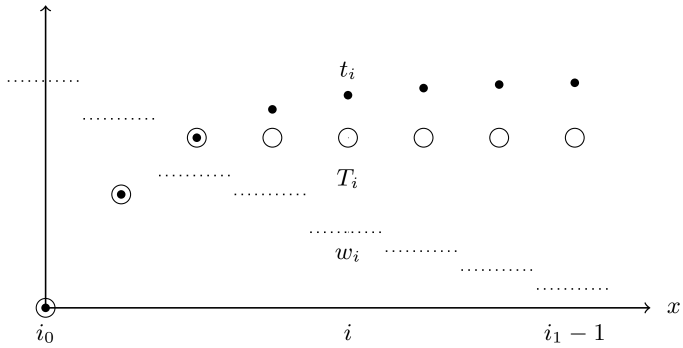

Fig. 1. Visualization of the thresholds set by the algorithm

where for convenience we also set $A^{i_0-1} = A_{i_0-1}$. We will show that the sequence of numbers defined by (3) forms a non-decreasing sequence depending only on the weights associated with previous elements $w_1, w_2, \ldots, w_{i_0-1}$, and that for $i_0 \leq i \leq i_1 - 1$ the algorithm sets threshold $T_i = t_i$ if $t_i \leq w_i$ and $T_i > w_i$ otherwise.

Lemma 2. *Consider a block-structured matroid $(\mathcal{U}, \mathcal{I})$ with blocks $B_1, \ldots, B_n$. For any input sequence σ generated by a block-restricted adversary, and any block B_j, let $i_0, i_0+1, \ldots, i_1$ denote the times when the elements of B_j are presented in σ. The sequence of numbers $t_{i_0}, \ldots, t_{i_1}$ defined by (3) satisfies $t_{i_0} \leq t_{i_0+1} \leq \cdots \leq t_{i_1}$ and depends only on the subsequence of σ preceding time i_0. Moreover, the algorithm is consistent and sets $T_i = t_i$ for all $i_0 \leq i \leq i_1$ such that $t_i \leq w_i$, and $T_i > w_i$ otherwise.*

An important corollary of the preceding structural result regarding the thresholds is the following assertion for two weight assignments w, w' drawn independently from a block-restricted weight distribution.

Corollary 1. *Let w, w' be two weight assignments drawn independently from a block-restricted weight distribution. For any input sequence σ generated by a block-restricted adversary, and any block B_j, let $i_0, i_0+1, \ldots, i_1$ denote the times when the elements of B_j are presented in σ. Then, for all $i_0 \leq i < i_1$, $(w_i - T_i)^+ = (w_i - t_i)^+$, and $w_i, t_i, w'(x_i)$ are mutually independent, so*

$$\mathbb{E}[(w_i - T_i)^+] = \mathbb{E}[(w_i - t_i)^+] = \mathbb{E}[(w'(x_i) - t_i)^+].$$

The final ingredient is an upper bound on the sum of the surrogate thresholds t_i for the elements x_i in $R(A)$, where A is the set of elements accepted by the algorithm on a given input sequence σ.

Proposition 2. *For every input sequence σ generated by a block-restricted adversary, let $A = A(\sigma)$. Then $\sum_{x_i \in R(A)} t_i \leq \frac{1}{2} \cdot \mathbb{E}[w'(R(A))]$.*

The proof of Theorem 1 uses our structural insight regarding the thresholds to lift the proof from the actual thresholds to the surrogate thresholds. It then uses the upper and lower bounds on the surrogate thresholds from this section to establish the claimed bound.

4 Application to Mechanism Design

We conclude by showing how our prophet inequality algorithm can be used to derive dominant strategy incentive compatible (DSIC), constant factor-approximations to the optimal revenue for a multi-parameter setting in which Myerson's analysis of the revenue-maximizing auction does not apply. Our result applies to concave weights whose distribution satisfies a conditional analog of the monotone hazard rate (MHR) condition. Specifically, we will assume that for each element $u_i \in \mathcal{U}$ the conditional distribution of the marginal weight $w_{i,j}$ of the j-th unit given the marginal weights $w_{i,1}, \ldots, w_{i,j-1}$ of the preceding units is MHR. That is, $\frac{f(w_{i,j}|w_{i,j_0},\ldots,w_{i,j-1})}{1-F(w_{i,j}|w_{i,j_0},\ldots,w_{i,j-1})}$ is non-decreasing in $w_{i,j}$. One example of a distribution satisfying this assumption is obtained by first drawing $w_{i,1} \sim U[0,1]$, then drawing $w_{i,2} \sim U[0, w_{i,1}]$, and so on.

We obtain posted-price mechanisms by combining the algorithm for polymatroids with "eager" or "lazy" monopoly reserves [7]. The monopoly reserve r^* for a given distribution F over valuations v with density f is $r^* = \phi^{-1}(0)$ where $\phi(v) = v - \frac{1-F(v)}{f(v)}$ is the virtual valuation. In the case of "eager" reserves, we modify the algorithm so that it only awards element x_i if its weight $w(x_i)$ exceeds the threshold T_i and the monopoly reserve r_i^* of the conditional distribution of $w(x_i)$. In the case of "lazy" reserves, we first run the algorithm to determine a tentative allocation, but then we only allocate elements whose weight also exceeds the reserve. Note that this can be done in an online fashion by computing thresholds as if all tentative assignments were made, but only actually awarding an element if it also exceeds the reserve.

Both mechanisms are DSIC as they are posted price. To prove the revenue bounds we need the following single-sample result.

Lemma 3 (Lemma 3.10 of Dhangwatnotai et al. [7]). *Let F be an MHR distribution with monopoly price r^* and revenue function $\hat{R}$. Let $V(t)$ denote the expected welfare of a single-item auction with a posted price of t and a single bidder with valuation drawn from F. For every nonnegative number $t \geq 0$, $\hat{R}(\max\{t, r^*\}) \geq \frac{1}{e} \cdot V(t)$.*

Theorem 3. *For polymatroids P_f defined by rational-valued submodular function f and concave weights that satisfy the conditional analog of the MHR condition, combining the polymatroid prophet inequality algorithm with "eager" or "lazy" reserves yields a DSIC mechanism whose revenue R_{EAGER} or R_{LAZY} on any input sequence σ generated by an online weight-adaptive adversary satisfies*

$$R_{EAGER}(\sigma) \geq \frac{1}{2e^2} \cdot R_{OPT}(\sigma) \text{ or } R_{LAZY} \geq \frac{1}{2e} \cdot R_{OPT}(\sigma),$$

where R_{OPT} denotes the optimal revenue.

Corollary 2. *For MHR valuations with decreasing marginals, there is a truthful $1/2e$ approximation to revenue for video-on-demand, bandwidth markets, sponsored search, and local purchasing collectives.*

References

[1] Alaei, S.: Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In: Proc. of 52nd FOCS. pp. 512–521 (2011)
[2] Babaioff, M., Immorlica, N., Kleinberg, R.: Matroids, secretary problems, and online mechanisms. In: Proc. of 18th SODA. pp. 434–443 (2007)
[3] Babaioff, M., Lucier, B., Nisan, N.: Bertrand networks. In: Proc. of 14th EC. pp. 33–34 (2013)
[4] Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R.V.: An ascending vickrey auction for selling bases of a matroid. Operations Research 59(2), 400–413 (2011)
[5] Chawla, S., Hartline, J.D., Malec, D.L., Sivan, B.: Multi-parameter mechanism design and sequential posted pricing. In: Proc. of 41th STOC. pp. 311–320 (2010)
[6] Dean, B.C., Goemans, M.X., Vondrák, J.: Approximating the stochastic knapsack problem: The benefit of adaptivity. In: Proc. of 45th FOCS. pp. 208–217 (2004)
[7] Dhangwatnotai, P., Roughgarden, T., Yan, Q.: Revenue maximization with a single sample. Games and Economic Behavior in press (2014)
[8] Dütting, P., Kleinberg, R.: Polymatroid prophet inequalities (2013), http://arxiv.org/abs/1307.5299
[9] Feldman, M., Gravin, N., Lucier, B.: Combinatorial auctions via posted prices. In: Proc. of 26th SODA. pp. 123–135 (2015)
[10] Goel, G., Mirrokni, V.S., Paes Leme, R.: Polyhedral clinching auctions and the adwords polytope. In: Proc. of 44th STOC. pp. 107–122 (2012)
[11] Guha, S., Munagala, K.: Multi-armed bandits with metric switching costs. In: Proc. of 36th ICALP. pp. 496–507 (2009)
[12] Guha, S., Munagala, K., Shi, P.: Approximation algorithms for restless bandit problems. Journal of the ACM 58, 3:1–3:50 (2010)
[13] Hajiaghayi, M., Kleinberg, R., Sandholm, T.W.: Automated mechanism design and prophet inequalities. In: Proc. of 22nd AAAI. pp. 58–65 (2007)
[14] Kleinberg, R., Weinberg, S.M.: Matroid prophet inequalities. In: Proc. of 44th STOC. pp. 123–136 (2012)
[15] Krengel, U., Sucheston, L.: Semiamarts and finite values. Bulletin of the American Mathematical Society 83, 745–747 (1977)
[16] Krengel, U., Sucheston, L.: On semiamarts, amarts, and processes with finite value. Advances in Probability and Related Topics 4, 197–266 (1978)
[17] Myerson, R.: Optimal auction design. Mathematics of Operations Research 6, 58–73 (1981)
[18] Samuel-Cahn, E.: Comparison of threshold stop rules and maximum for independent nonnegative random variables. Annals of Probability 12, 1213–1216 (1984)
[19] Schrijver, A.: Combinatorial Optimization: Polyhedra and Efficiency. Springer Verlag (2002)
[20] Tse, D.N.C., Hanly, S.V.: Multiaccess fading channels-part I: Polymatroid structure, optimal resource allocation and throughput capacities. Transactions on Information Theory 44, 2796–2815 (1998)