

Optimal Online Contention Resolution Schemes via Ex-Ante Prophet Inequalities

Euiwoong Lee

Courant Institute of Mathematical Sciences, New York University, New York City, USA

euiwoong@cims.nyu.edu

Sahil Singla

Computer Science Department, Carnegie Mellon University, Pittsburgh, USA

ssingla@cmu.edu

Abstract

Online contention resolution schemes (OCRSs) were proposed by Feldman, Svensson, and Zenklusen [11] as a generic technique to round a fractional solution in the matroid polytope in an online fashion. It has found applications in several stochastic combinatorial problems where there is a *commitment* constraint: on seeing the value of a stochastic element, the algorithm has to immediately and irrevocably decide whether to select it while always maintaining an independent set in the matroid. Although OCRSs immediately lead to prophet inequalities, these prophet inequalities are not optimal. Can we instead use prophet inequalities to design optimal OCRSs?

We design the first optimal $1/2$ -OCRS for matroids by reducing the problem to designing a matroid prophet inequality where we compare to the stronger benchmark of an ex-ante relaxation. We also introduce and design optimal $(1-1/e)$ -random order CRSs for matroids, which are similar to OCRSs but the arrival order is chosen uniformly at random.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis

Keywords and phrases Prophets, Contention Resolution, Stochastic Optimization, Matroids

Digital Object Identifier 10.4230/LIPIcs.ESA.2018.57

Related Version A full version of the paper is available at <https://arxiv.org/abs/1806.09251>.

Acknowledgements We are thankful to Ravishankar Krishnaswamy and Deeparnab Chakrabarty for useful discussions in early part of this project. Part of this work was done while the authors were visiting the Simons Institute for the Theory of Computing. The second author was supported in part by NSF awards CCF-1319811, CCF-1536002, and CCF-1617790.

1 Introduction

Given a combinatorial optimization problem, a common algorithmic approach is to first solve a convex relaxation of the problem and to then round the obtained fractional solution \mathbf{x} into a *feasible integral* solution while (approximately) preserving the objective. Contention resolution schemes (CRSs), introduced in [8], is a way to perform this rounding given a fractional solution $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$. For $c > 0$, intuitively a c -CRS is a rounding algorithm that guarantees every element i is selected into the final feasible solution w.p. at least $c \cdot x_i$. For a *maximization* problem with a linear objective, by linearity of expectation such a c -CRS directly implies a c -approximation algorithm.



© Euiwoong Lee and Sahil Singla;
licensed under Creative Commons License CC-BY
26th Annual European Symposium on Algorithms (ESA 2018).

Editors: Yossi Azar, Hannah Bast, and Grzegorz Herman; Article No. 57; pp. 57:1–57:14



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In a recent work, Feldman et al. [11] introduced an Online CRS (OCRS), which is a CRS with an additional property that it performs the rounding in an “online fashion”. This property is crucial for the *prophet inequality* problem (or any stochastic combinatorial problem with a *commitment* constraint; see §1.3).

► **Definition 1** (Prophet inequality). Suppose each element $i \in N$ takes a value $v_i \in \mathbb{R}_{\geq 0}$ independently from some known distribution \mathcal{D}_i . These values are presented one-by-one to an online algorithm in an adversarial order. Given a packing feasibility constraint $\mathcal{F} \subseteq 2^N$, the problem is to immediately and irrevocably decide whether to select the next element i , while always maintaining a feasible solution and maximizing the sum of the selected values.

A c -*approximation prophet inequality* for $0 \leq c \leq 1$ means there exists an online algorithm with expected value at least c times the expected value of an offline algorithm that knows all values from the beginning. As shown in [11], a c -OCRS immediately implies a c -approximation prophet inequality. Some other applications are oblivious posted pricing mechanisms and stochastic probing.

Although powerful, the above approach of using OCRSs to design prophet inequalities does not give us *optimal* prophet inequalities. For example, while we know a $1/2$ -approximation prophet inequality over matroids [20], we only know a $1/4$ -OCRS over matroids [11]. This indicates that the currently known OCRSs may not be optimal. Can we design better OCRSs? The main contribution of this work is to design an *optimal* OCRS over matroid constraints using the following idea:

Not only can we design prophet inequalities from OCRSs, we can also design OCRSs from prophet inequalities.

More specifically, our OCRS is based on an *ex-ante* prophet inequality: we compare the online algorithm to the stronger benchmark of a convex relaxation. We modify existing prophet inequalities to obtain ex-ante prophet inequalities while *preserving* the approximation factors. As a corollary, this gives the first optimal $1/2$ -OCRS over matroids.

Since for many applications the arrival order is not chosen by an adversary, some recent works have also studied prophet secretary inequalities where the arrival order is chosen uniformly at random [10, 9, 5]. Motivated by these works, we introduce *random order contention resolution schemes* (RCRS), which is an OCRS for uniformly random arrival¹. Again by designing the corresponding random order ex-ante prophet inequalities, we obtain optimal $(1 - 1/e)$ -RCRS over matroids.

In §1.1 we formally define an OCRS/RCRS and an ex-ante prophet inequality. In §1.2 we describe our results and proof techniques. See §1.3 for further related work.

1.1 Model

CRSs are a powerful tool for offline and stochastic optimization problems [8, 15]. For a given $\mathbf{x} \in [0, 1]^N$, let $R(\mathbf{x})$ denote a random set containing each element $i \in N$ independently w.p. x_i . We say an element i is *active* if it belongs to $R(\mathbf{x})$.

► **Definition 2** (Contention resolution scheme). Given a finite ground set N with $n = |N|$ and a packing (downward-closed) family of feasible subsets $\mathcal{F} \subseteq 2^N$, let $P_{\mathcal{F}} \subseteq [0, 1]^N$ be the convex hull of all characteristic vectors of feasible sets. For a given $\mathbf{x} \in P_{\mathcal{F}}$, a c -*selectable*

¹ A parallel independent work has also introduced RCRS [1]; however, their technical results are very different.

CRS (or simply, c -CRS) is a (randomized) mapping $\pi : 2^N \rightarrow 2^N$ satisfying the following three properties:

- (i) $\pi(S) \subseteq S$ for all $S \subseteq N$.
- (ii) $\pi(S) \in \mathcal{F}$ for all $S \subseteq N$.
- (iii) $\Pr_{R(\mathbf{x}), \pi}[i \in \pi(R(\mathbf{x}))] \geq c \cdot x_i$ for all $i \in N$.

Notice, if f is a monotone linear function then $\mathbb{E}[f(\pi(R(\mathbf{x})))] \geq c \cdot \mathbb{E}[f(R(\mathbf{x}))]$. By constructing CRSs for various constraint families of \mathcal{F} , Chekuri et al. [8] give improved approximation algorithms for linear and submodular maximization problems under knapsack, matroid, matchoid constraints, and their intersections².

In the above applications to offline optimization problems, the algorithm first flips all the random coins to sample $R(\mathbf{x})$, and then obtains $\pi(R(\mathbf{x})) \subseteq R(\mathbf{x})$. For various online problems such as the prophet inequality, this randomness is an inherent part of the problem. Feldman et al. [11] therefore introduce an OCRS where the random set $R(\mathbf{x})$ is sampled in the same manner, but whether $i \in R(\mathbf{x})$ (or not) is only revealed one-by-one to the algorithm in an adversarial order³. After each revelation (arrival), the OCRS has to irrevocably decide whether to include $i \in R(\mathbf{x})$ into $\pi(R(\mathbf{x}))$ (if possible). A c -selectable OCRS (or simply, c -OCRS) is an OCRS satisfying the above properties (i) to (iii) of a c -CRS.

In this work, we also study RCRS which is an OCRS with the arrival order chosen uniformly at random. A c -selectable RCRS (or simply, c -RCRS) is an RCRS satisfying the above properties (i) to (iii) of a c -CRS, where in Property (iii) we also take expectation over the arrival order.

While prophet inequalities have been designed using OCRSs, our main result in this paper is to show a deeper *reverse* connection between OCRSs and prophet inequalities. We first define an *ex-ante* prophet inequality. Given a prophet inequality problem instance with packing constraints \mathcal{F} and r.v.s $v_i \sim \mathcal{D}_i$ for $i \in N$, the following *ex-ante relaxation* gives an upper bound on the expected offline optimum:

$$\max_{\mathbf{x}} \sum_i x_i \cdot \mathbb{E}_{v_i \sim \mathcal{D}_i}[v_i \mid v_i \text{ takes value in its top } x_i \text{ quantile}] \quad \text{s.t.} \quad \mathbf{x} \in P_{\mathcal{F}}. \quad (1)$$

To prove that (1) is an upper bound, we interpret x_i as the probability that i is in the offline optimum. It is also known that (1) is a convex program and can be solved efficiently; see [11] for more details.

► **Definition 3** (Ex-ante prophet inequality). For $0 \leq c \leq 1$, a c -approximation ex-ante prophet inequality for packing constraints \mathcal{F} is a prophet inequality algorithm with expected value at least c times (1).

Before describing our results, to build some intuition for the above definitions we discuss the special case of a rank 1 matroid, i.e., where we can only select one of the n elements.

Example: Rank 1 matroid

For simplicity, in this section we assume that all random variables are Bernoulli, i.e., v_i takes value y_i independently w.p. p_i , and is 0 otherwise. We first show why a c -OCRS implies a c -approximation prophet inequality for rank 1 matroids.

² Some “greedy” properties are also required from the CRS for the guarantees to hold for a submodular function f [8].

³ For adversarial arrival order, we assume that this order is known to the OCRS algorithm in advance. This offline adversary is weaker than the almighty adversary considered in [11], but is common in the prophet inequality literature [24, 25]. We need this assumption in §2 to define our exponential sized linear program.

Consider the optimum solution \mathbf{x} to the ex-ante relaxation (1) for the above Bernoulli instance. Its objective value is $\sum_i x_i y_i$ where \mathbf{x} satisfies $\sum_i x_i \leq 1$. Moreover, $x_i \leq p_i$ for all i because selecting i beyond p_i does not increase (1). To see why (1) gives an upper bound on the expected offline maximum, observe that if we interpret x_i as the probability that v_i is the offline maximum, this gives a feasible solution to $\sum_i x_i \leq 1$ and with value at most $\sum_i x_i y_i$. Thus, to prove a c -approximation prophet inequality, it suffices to design an online algorithm with value at least $c \cdot \sum_i x_i y_i$. Consider an algorithm that runs a c -OCRS on \mathbf{x} , where i is considered active independently w.p. x_i/p_i whenever v_i takes value y_i . This ensures element i is active w.p. exactly x_i . Since a c -OCRS guarantees each element is selected w.p. $\geq c$ when it is active, by linearity of expectation such an algorithm has expected value at least $c \cdot \sum_i x_i y_i$.

We now discuss a simple $1/4$ -OCRS for a rank 1 matroid. Given \mathbf{x} satisfying $\sum_i x_i \leq 1$, consider an algorithm that ignores each element i independently w.p. $1/2$, and otherwise selects i only if it is active. Since this algorithm selects any element i w.p. at most $x_i/2$ (when i is not ignored and is active), by Markov's inequality the algorithm selects no element till the end w.p. at least $1 - \sum_i x_i/2 \geq 1/2$. Hence the algorithm reaches each element i w.p. at least $1/2$ without selecting any of the previous elements. Moreover, it does not ignore i w.p. $1/2$, which implies it considers each element w.p. at least $1/4$. The OCRS due to Feldman et al. [11] can be thought of generalizing this approach to a general matroid.

An interesting result of Alaei [4] shows that the above $1/4$ -OCRS can be improved to a $1/2$ -OCRS over a rank 1 matroid by “greedily” maximizing the probability of ignoring the next element i , but considering i w.p. $1/2$ on average. In the full version of the paper we present Alaei's proof for completeness, and also show how to obtain a simple $(1 - 1/e)$ -RCRS for a rank 1 matroid. This raises the question whether one can obtain a $1/2$ -OCRS and a $(1 - 1/e)$ -RCRS for general matroids.

1.2 Results and Techniques

Our first theorem gives an approximation factor preserving reduction from OCRSs to ex-ante prophet inequalities.

► **Theorem 4.** *For $0 \leq c \leq 1$, a c -approximation ex-ante prophet inequality for adversarial (random) arrival order over a packing constraint \mathcal{F} implies a c -OCRS (c -RCRS) over \mathcal{F} .*

We complement the above theorem by designing ex-ante prophet inequalities over matroids.

► **Theorem 5.** *For matroids, there exists a $1/2$ -approximation ex-ante prophet inequality for adversarial arrival order and a $(1 - 1/e)$ -approximation ex-ante prophet inequality for uniformly random arrival order.*

As a corollary, the above two theorems give optimal OCRS and RCRS over matroids. This generalizes the rank 1 results discussed in the previous section to general matroids; although the proof techniques are very different.


► **Corollary 6.** *For matroids, there exists a $1/2$ -OCRS and a $(1 - 1/e)$ -RCRS.*

Our $1/2$ -OCRS above assumes that the arrival order is known to the algorithm. It is an interesting open question to find a $1/2$ -OCRS for an almighty/online adversary as in [11].

We first prove that both the factors $1/2$ and $(1 - 1/e)$ in Corollary 6 are *optimal*.

Optimality of $1/2$ -OCRS and $(1 - 1/e)$ -RCRS

We argue that the factors $1/2$ and $(1 - 1/e)$ in Corollary 6 are optimal even in the special case of a rank 1 matroid. For adversarial arrival, consider just two elements, i.e., $n = 2$, with



■ **Figure 1** The Hat example on $n+2$ vertices. The following \mathbf{x} belongs to the graphic matroid: $x_e = 1/2$ for $e = (u_i, v_j)$ where $i \in \{1, 2\}$ and $j \in \{1, \dots, n\}$, and $x_e = 1$ for $e = (u_1, u_2)$.

$x_1 = 1 - \epsilon$ and $x_2 = \epsilon$ for some $\epsilon \rightarrow 0$. Since the OCRS algorithm has to select the first element at least $1/2$ fraction of the times, it can attempt to select the second element at most $1/2 + \epsilon/2$ fraction of the times.

For random arrival order, consider the feasible solution \mathbf{x} with $x_i = 1/n$ for every $i \in N$. We show that no online RCRS algorithm can guarantee each element is selected w.p. greater than $\frac{(1-1/e)}{n}$. This is because for the product distribution, w.p. $1/e$ none of the n elements is active (more precisely, w.p. $(1 - 1/n)^n$). Hence the RCRS algorithm, which only selects active elements, selects some element w.p. $1 - 1/e$. This implies on average it cannot pick every element w.p. greater than $\frac{(1-1/e)}{n}$. This example, originally shown in [8], also proves that offline CRS cannot be better than $(1 - 1/e)$ -selectable.

Our techniques

We first see the difficulty in extending Alaei's greedy approach from a rank 1 matroid to a general matroid. Consider the graphic matroid for the *Hat* example (see Figure 1). Suppose the base edge (u_1, u_2) appears in the end of an adversarial order. Notice that any algorithm which ignores the structure of the matroid is very likely to select some pair of edges (u_1, v_i) and (v_i, u_2) for some i . Since this pair spans the base edge (u_1, u_2) , such an OCRS algorithm will not satisfy c -selectability for (u_1, u_2) . To overcome this, Feldman et al. [11] decompose the matroid into "simpler" matroids using \mathbf{x} . However, it is not clear how to extend their approach beyond a $1/4$ -OCRS.

In this paper we take an alternate LP based approach to design OCRSs, which was first used by Chekuri et al. [8] to design offline CRSs. The idea is to define an exponential sized linear program where each variable denotes a deterministic OCRS algorithm. The objective of this linear program is to maximize c s.t. each element is selected at least c fraction of the times (c -selectability). Thus to show existence of a $1/2$ -OCRS, it suffices to prove this linear program has value $c \geq 1/2$. In §2 we prove this by showing that the dual LP has value at least $1/2$ because it can be interpreted as an ex-ante prophet inequality.

Next, to show there exists a $1/2$ approximation ex-ante prophet inequality, our approach is inspired from the matroid prophet inequality of Kleinberg and Weinberg [20]. They give an online algorithm that gets at least half of the expected offline optimum for the product distribution (independent r.v.s). Unfortunately, their techniques do not directly extend because the ex-ante relaxation objective could be significantly higher than for the product distribution (this is known as the *correlation gap*, which can be $e/(e-1)$ [2, 6]). Our primary technique is to view the ex-ante relaxation solution as a "special kind" of a correlated value distribution. Although prophet inequalities are not possible for general correlated distributions [19], we show that in this special case the original proof of the matroid prophet inequality algorithm retains its $1/2$ approximation after some modifications.

1.3 Further Related Work

Krengel and Sucheston gave the first tight $1/2$ -single item prophet inequality [22, 21]. The connection between multiple-choice prophet inequalities and mechanism design was recognized in [18]; they proved a prophet inequality for uniform matroids. This bound was later improved by Alaei [3] using the *Magician's problem*, which is an OCRS in disguise. Chawla et al. [7] further developed the connection between prophet inequalities and mechanism design, and showed how to be $O(1)$ -prophet inequality for general matroids in a variant where the algorithm may choose the element order. Yan [26] improved this result to $e/(e-1)$ -competitive using the *correlation gap* for submodular functions, first studied in [2, 6]. Chekuri et al. [8] adapted correlation gaps to a polytope to design CRSs. Improved correlation gaps were presented in [26, 17]. The matroid prophet inequality was first explicitly formulated in [20]. Feldman et al. [11] gave an alternate proof, and extended to Bernoulli submodular functions, using OCRSs. Finally, information theoretic $O(\text{poly log}(n))$ -prophet inequalities are also known for general downward-closed constraints [24, 25].

The prophet secretary notion was first introduced in [10], where the elements arrive in a uniformly random order and draw their values from known independent distributions. Their results have been recently improved [9, 5]. There is a long line of work on studying the *commitment constraints* for combinatorial probing problems, e.g., see [13, 15, 16, 14]. In these models the algorithm starts with some stochastic knowledge about the input and on probing an element has to irrevocably commit if the element is to be included in the final solution. A common approach to handle such a constraint is using a prophet inequality/OCRS.

2 OCRS Assuming an Ex-Ante Prophet Inequality

In this section we prove Theorem 4, showing how to reduce the problem of designing an OCRS to a prophet inequality where we compare ourself to the ex-ante relaxation instead of the expected offline maximum.

2.1 Using LP Duality

Given a finite ground set N with $n = |N|$ and a downward-closed family of feasible subsets $\mathcal{F} \subseteq 2^N$, let $P_{\mathcal{F}} \subseteq [0, 1]^N$ be the convex hull of all characteristic vectors of feasible sets. Let $\mathbf{x} \in P_{\mathcal{F}}$ and $R(\mathbf{x})$ denote a random set containing each element $i \in N$ independently w.p. x_i . For offline CRSs, let Φ^* be the set of valid offline *deterministic* mappings; i.e., $\phi : 2^N \rightarrow \mathcal{F}$ is in Φ^* iff $\phi(A) \subseteq A$ and $\phi(A) \in \mathcal{F}$ for all $A \subseteq N$. For $\phi \in \Phi^*$ and $i \in N$, let $q_{i,\phi} := \Pr_{R(\mathbf{x})}[i \in \phi(R(\mathbf{x}))]$ denote the probability of selecting i if the CRS executes ϕ . The following LP relaxation, introduced by Chekuri et al. [8], finds a c -selectable *randomized* CRS. It has variables $\{\lambda_{\phi}\}_{\phi \in \Phi^*}$ and c .

$$\begin{aligned} \max_{\lambda, c} \quad & c \\ \text{s.t.} \quad & \sum_{\phi \in \Phi^*} q_{i,\phi} \lambda_{\phi} \geq x_i \cdot c && i \in N \\ & \sum_{\phi \in \Phi^*} \lambda_{\phi} = 1 \\ & \lambda_{\phi} \geq 0 && \forall \phi \in \Phi^* \end{aligned}$$

Observe that if the above LP has value c , there exists a randomized c -CRS. This is because we can randomly select one of the ϕ 's w.p. λ_{ϕ} , and the constraint $\sum_{\phi \in \Phi^*} q_{i,\phi} \lambda_{\phi} \geq x_i \cdot c$

ensures c -selectability for every $i \in N$. Chekuri et al. noticed that by strong duality, to prove the above LP has value at least c , it suffices to show that the following dual program has value at least c . It has variables $\{y_i\}_{i \in N}$ and μ .

$$\begin{aligned}
& \min_{\mathbf{y}, \mu} \mu \\
& \text{s.t.} \quad \sum_{i \in N} q_{i, \phi} y_i \leq \mu && \phi \in \Phi^* \\
& \quad \sum_{i \in N} x_i y_i = 1 \\
& \quad y_i \geq 0 && \forall i \in N
\end{aligned}$$

To design OCRSs (RCRSs), we take a similar approach as Chekuri et al and let Φ^* be the set of all *deterministic online algorithms*. Formally, $\phi : 2^N \times 2^N \times N \rightarrow \{0, 1\}$ belongs to Φ^* iff $\phi(A, B, i) = 1$ only for $B \subseteq A$, $i \notin A$, and $B \cup \{i\} \in \mathcal{F}$. Intuitively, $\phi(A, B, i) = 1$ indicates that the online algorithm selects element i in the current iteration after processing elements in A and selecting elements in B . Let $q_{i, \phi}$ denote the probability of selecting i if the OCRS (RCRS) executes ϕ , where for RCRS we also take probability over the random order. By the above duality argument, to show existence of a c -OCRS (c -RCRS) it suffices to prove the dual LP has value at least c . We prove this by showing that for any $\mathbf{y} \geq 0$ s.t. $\sum_{i \in N} x_i y_i = 1$, there exists $\phi \in \Phi^*$ such that $\sum_{i \in N} q_{i, \phi} y_i \geq c$.

Consider a Bernoulli prophet inequality instance where each element $i \in N$ has value y_i with probability x_i , and 0 otherwise. Since $\mathbf{x} \in P_{\mathcal{F}}$, notice that $\sum_{i \in N} x_i y_i = 1$ is exactly the value of the ex-ante relaxation (1) for this instance. Thus, a c -approximation ex-ante prophet inequality implies there exists a $\phi \in \Phi^*$ with value at least c . By linearity of expectation, the value of ϕ is $\sum_{i \in N} q_{i, \phi} y_i$, which proves $\sum_{i \in N} q_{i, \phi} y_i \geq c$.

2.2 Solving the LP Efficiently

While the original primal LP has an exponential number of variables, we can compute an OCRS (or RCRS) that achieves value at least c as follows. In the dual program, given \mathbf{y} s.t. $\sum_i x_i y_i = 1$, we can use the ex-ante prophet inequality to find $\phi \in \Phi^*$ with value $\sum_i q_{i, \phi} y_i \geq c$ in polynomial time. (Notice $q_{i, \phi}$ can be computed in polynomial time because the adversarial order is known to the OCRS algorithm.) This implies for any $\epsilon > 0$, the polytope $Q_{c-\epsilon} := \{\mathbf{y} : \mathbf{y} \geq 0, \sum_i x_i y_i = 1, \sum_i q_{i, \phi} y_i \leq c - \epsilon \text{ for all } \phi \in \Phi^*\}$ is empty.

Since we have an efficient *separation oracle* (for any \mathbf{y} , we can find a violated constraint in polynomial time) for $Q_{c-\epsilon}$, by running the ellipsoid algorithm [12] we can find a subset $\Phi' \subseteq \Phi^*$ with $|\Phi'| = \text{poly}(n)$ in polynomial time such that $Q'_{c-\epsilon} := \{\mathbf{y} : \mathbf{y} \geq 0, \sum_i x_i y_i = 1, \sum_i q_{i, \phi} y_i \leq c - \epsilon \text{ for all } \phi \in \Phi'\}$ is empty. Now the following linear program, which has a polynomial number of variables and constraints, with optimal value at least $c - \epsilon$ can be solved efficiently.

$$\begin{aligned}
& \max_{\lambda, c} \quad c \\
& \text{s.t.} \quad \sum_{\phi \in \Phi'} q_{i, \phi} \lambda_{\phi} \geq x_i \cdot c && i \in N \\
& \quad \sum_{\phi \in \Phi'} \lambda_{\phi} = 1 \\
& \quad \lambda_{\phi} \geq 0 && \forall \phi \in \Phi'
\end{aligned}$$

3 Ex-Ante Prophet Inequalities for a Matroid

This section proves Theorem 5 by designing for a matroid a $1/2$ -ex-ante prophet inequality under adversarial arrival and a $(1 - 1/e)$ -ex-ante prophet inequality under random arrival.

3.1 Notation

Let $\mathbf{v} \sim \mathcal{D}$ be a set of random element values $\{v_1, \dots, v_n\}$ where each v_i is independently drawn from \mathcal{D}_i . Let \mathbf{x} be the optimal solution to the ex-ante relaxation in (1) for a given matroid $\mathcal{M} = (N, \mathcal{I})$. For $i \in N$, denote

$$y_i := \mathbb{E}_{v_i \sim \mathcal{D}_i}[v_i \mid v_i \text{ takes value in its top } x_i \text{ quantile}]. \quad (2)$$

Since $\mathbf{x} \in P_{\mathcal{M}}$, we can write it as a convex combination of independent sets in the matroid. In particular, this gives a correlated distribution $\hat{\mathcal{D}}$ over independent sets of \mathcal{M} such that for each $i \in N$, we have $\Pr_{I \sim \hat{\mathcal{D}}}[i \in I] = x_i$. Let $\hat{\mathbf{v}} = \{\hat{v}_1, \dots, \hat{v}_n\}$ be a set of random values obtained by sampling $I \sim \hat{\mathcal{D}}$ and setting $\hat{v}_i = y_i$ for $i \in I$, and $\hat{v}_i = 0$ otherwise. Notice the optimal value of (1) is $\sum_i x_i y_i$ and for each $i \in N$, we have $\mathbb{E}[\hat{v}_i] = x_i y_i$.

We need the following notation to describe our algorithms.

► **Definition 7.** For any vector $\hat{\mathbf{v}}$ denoting values of elements of N and any $A \subseteq N$, we define:

- Let $\text{Opt}(\hat{\mathbf{v}} \mid A) \subseteq N \setminus A$ denote the maximum value independent set in the contracted matroid \mathcal{M}/A .
- Let $R(A, \hat{\mathbf{v}}) := \sum_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A)} \hat{v}_i$ denote the *remaining value* after selecting set A .

We next define a *base price* of for every element i .

► **Definition 8.** For $A \in \mathcal{I}$ denoting an independent set of elements accepted by our algorithm, we define

- Let $b_i(A, \hat{\mathbf{v}}) := R(A, \hat{\mathbf{v}}) - R(A \cup \{i\}, \hat{\mathbf{v}})$ denote a threshold for element i .
- Let $b_i(A) := \mathbb{E}_{\hat{\mathbf{v}} \sim \hat{\mathcal{D}}}[b_i(A, \hat{\mathbf{v}})]$ denote the *base price* for element i .

3.2 Reducing to Bernoulli Distributions

In this section we show that it suffices to only prove Theorem 5 for Bernoulli distributions.

► **Lemma 9.** *If there exists an α -approximation ex-ante prophet inequality for Bernoulli distributed independent random values then there exists an α -approximation ex-ante prophet inequality for general distributed independent random values.*

Proof. Given a prophet inequality instance where $\mathbf{v} \sim \mathcal{D}$ for a general distribution \mathcal{D} , consider a new Bernoulli prophet inequality instance $\mathbf{v}' \sim \mathcal{D}'$ where for each $i \in N$, r.v. $v'_i \sim \mathcal{D}'_i$ independently takes value y_i (defined in (2)) w.p. x_i , and is 0 otherwise. Since the optimal ex-ante fractional value for both the general and Bernoulli instance is the same, to prove this theorem we use an ex-ante prophet inequality for the Bernoulli instance to design an ex-ante prophet inequality for the general instance with the same expected value.

On arrival of an element i , consider an algorithm for the general distribution that treats i is active iff v_i takes value in its top x_i quantile. If active, the algorithm asks the ex-ante prophet inequality of the Bernoulli instance to decide whether to select i . We claim that the expected value of this algorithm is $\alpha \cdot \sum_i x_i y_i$, which will prove this theorem. The claim is true because for the above algorithm each element i is active independently w.p. exactly x_i , and conditioned on being active its expected value is exactly y_i . Thus by linearity of expectation, the expected value is the same as the Bernoulli instance, which is $\alpha \cdot \sum_i x_i y_i$. ◀

3.3 Adversarial Order

We prove the optimal ex-ante prophet inequality for a matroid under the adversarial arrival.

► **Theorem 10.** *For matroids, there exists a 1/2-approximation ex-ante prophet inequality for adversarial arrival order.*

Given the notation and definitions in §3.1, the proof of Theorem 10 is similar to the proof of the matroid prophet inequality in [20].

By Lemma 9, we know it suffices to prove this theorem only for Bernoulli distributions. Consider $\mathbf{v} \sim \mathcal{D}$ as the input to our online algorithm, where v_i takes value y_i w.p. x_i and is 0 otherwise. Given \mathbf{v} , our algorithm is deterministic and let $A := A(\mathbf{v})$ denote the set of elements that it selects. Relabel the elements such that the arrival order of the elements is $1, \dots, n$. Let $A_i = A \cap \{1, \dots, i\}$.

Our algorithm selects the next element i iff both $v_i > \mathcal{T}_i := \alpha \cdot b_i(A_{i-1})$ and selecting i is feasible in \mathcal{M} , where $\alpha = \frac{1}{2}$. Thus, the total value of algorithm Alg := $\sum_{i \in A} v_i = \text{Revenue} + \text{Utility}$, where

$$\text{Revenue} := \sum_{i \in A} \mathcal{T}_i \quad \text{and} \quad \text{Utility} := \sum_{i \in A} (v_i - \mathcal{T}_i)^+.$$

Since $\sum_{i \in N} x_i y_i$ is the optimal value of (1), to prove Theorem 10 it suffices to show $\mathbb{E}[\text{Alg}] = \mathbb{E}[\text{Revenue}] + \mathbb{E}[\text{Utility}] \geq \alpha \cdot \sum_{i \in N} x_i y_i$.

We keep track of the algorithm's progress using the following *residual* function:

$$r(i) := \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}} [R(A_{i-1}, \hat{\mathbf{v}})].$$

Clearly, $r(0) = \sum_{i \in N} x_i y_i$. In the following Lemma 11 and Lemma 12, we use the residual function to lower bound $\mathbb{E}[\text{Revenue}]$ and $\mathbb{E}[\text{Utility}]$.

► **Lemma 11.** $\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\text{Revenue}] = \alpha \cdot (r(0) - r(n))$.

Proof. From the definition of Revenue, we get

$$\begin{aligned} \text{Revenue} &= \alpha \cdot \sum_{i \in A} b_i(A_{i-1}) = \alpha \cdot \sum_{i \in A} \left(\mathbb{E}_{\hat{\mathbf{v}}} [R(A_{i-1}, \hat{\mathbf{v}})] - \mathbb{E}_{\hat{\mathbf{v}}} [R(A_{i-1} \cup \{i\}, \hat{\mathbf{v}})] \right) \\ &= \alpha \cdot \sum_{i \in A} \left(\mathbb{E}_{\hat{\mathbf{v}}} [R(A_{i-1}, \hat{\mathbf{v}})] - \mathbb{E}_{\hat{\mathbf{v}}} [R(A_i, \hat{\mathbf{v}})] \right) \\ &= \alpha \cdot \left(\mathbb{E}_{\hat{\mathbf{v}}} [R(A_0, \hat{\mathbf{v}})] - \mathbb{E}_{\hat{\mathbf{v}}} [R(A, \hat{\mathbf{v}})] \right). \end{aligned}$$

Taking expectation over $\mathbf{v} \sim \mathcal{D}$ and using definitions of $r(0)$ and $r(n)$, the lemma follows. ◀

► **Lemma 12.** $\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\text{Utility}] \geq (1 - \alpha) \cdot r(n)$.

Proof. We prove the following two inequalities:

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\text{Utility}] \geq \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} (\hat{v}_i - \mathcal{T}_i)^+ \right] \quad (3)$$

and

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} (\hat{v}_i - \mathcal{T}_i)^+ \right] \geq (1 - \alpha) \cdot \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}} [R(A, \hat{\mathbf{v}})]. \quad (4)$$

Lemma 12 now follows by summing (3) and (4), and using $r(n) = \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}} [R(A, \hat{\mathbf{v}})]$.

To prove (3), notice that for any i not selected by the algorithm, $v_i \leq \mathcal{T}_i$. This implies

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\text{Utility}] = \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in A} (v_i - \mathcal{T}_i)^+ \right] = \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} (v_i - \mathcal{T}_i)^+ \right].$$

Now observe that for any fixed i and v_1, \dots, v_{i-1} , the threshold \mathcal{T}_i is determined. Since v_i and \hat{v}_i are independent random variables with the same distribution, we get

$$\mathbb{E}_{\mathbf{v}}[(v_i - \mathcal{T}_i)^+ | v_1, \dots, v_{i-1}] = \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}}[(\hat{v}_i - \mathcal{T}_i)^+ | v_1, \dots, v_{i-1}].$$

This implies

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}}[\text{Utility}] = \mathbb{E}_{\mathbf{v}} \left[\sum_{i \in N} (v_i - \mathcal{T}_i)^+ \right] = \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in N} (\hat{v}_i - \mathcal{T}_i)^+ \right] \geq \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} (\hat{v}_i - \mathcal{T}_i)^+ \right].$$

Finally, to prove (4), we have

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} [R(A, \hat{\mathbf{v}})] &= \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} \hat{v}_i \right] \leq \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} \mathcal{T}_i \right] + \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} (\hat{v}_i - \mathcal{T}_i)^+ \right] \\ &\leq \alpha \cdot \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} [R(A, \hat{\mathbf{v}})] + \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}} \left[\sum_{i \in \text{Opt}(\hat{\mathbf{v}}|A)} (\hat{v}_i - \mathcal{T}_i)^+ \right], \end{aligned}$$

where the first inequality uses $\hat{v}_i \leq \mathcal{T}_i + (\hat{v}_i - \mathcal{T}_i)^+$ and the second inequality uses Claim 13 for $S = \text{Opt}(\hat{\mathbf{v}} | A)$. After rearranging, this implies (4). ◀

We need the following Claim 13 in the proof of Lemma 12.

► **Claim 13.** For every pair of disjoint sets A, S such that $A \cup S \in \mathcal{M}$,

$$\alpha \cdot \mathbb{E}_{\hat{\mathbf{v}} \sim \hat{\mathcal{D}}} \left[\sum_{i \in S} R(A_{i-1}, \hat{\mathbf{v}}) - R(A_{i-1} \cup \{i\}, \hat{\mathbf{v}}) \right] = \sum_{i \in S} \mathcal{T}_i \leq \alpha \cdot \mathbb{E}_{\hat{\mathbf{v}} \sim \hat{\mathcal{D}}} [R(A, \hat{\mathbf{v}})]. \quad (5)$$

Proof. This directly follows from [20], as they proved it for every fixed $\hat{\mathbf{v}}$. The proof is similar to Claim 18 in the next section. ◀

Proof of Theorem 10. Using Lemma 11 and Lemma 12, and substituting $\alpha = \frac{1}{2}$, we get

$$\mathbb{E}[\text{Alg}] = \mathbb{E}[\text{Utility}] + \mathbb{E}[\text{Revenue}] \geq \frac{1}{2} \cdot r(0) = \frac{1}{2} \cdot \sum_{i \in N} x_i y_i. \quad \blacktriangleleft$$

3.4 Random Order

We prove the optimal ex-ante prophet inequality for a matroid for random arrival.

► **Theorem 14.** For matroids, there exists a $(1 - 1/e)$ -approximation ex-ante prophet inequality for uniformly random arrival order.

The proof of Theorem 14 is similar to the matroid prophet secretary inequality in [9]. We consider the model where each item chooses the arrival time from $[0, 1]$ uniformly and independently, which is equivalent to the random permutation model. Starting with $A_0 = \emptyset$, let A_t denote the set of accepted elements by our algorithm *before* time t . This is a random variable that depends on the values \mathbf{v} and arrival times \mathbf{T} . For $t \in [0, 1]$, let

$$\alpha(t) := 1 - \exp(t - 1).$$

Suppose an element i arrives at time t , then our algorithm selects i iff both $v_i > \alpha(t) \cdot b_i(A_t)$ and selecting i is feasible in \mathcal{M} .

Similar to §3.3, we keep track of the algorithm's progress using the *residual* function

$$r(t) := \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}, \mathbf{T}}[R(A_t, \hat{\mathbf{v}})],$$

where A_t is a function of \mathbf{v} and \mathbf{T} . Clearly, $r(0) = \sum_{i \in N} x_i y_i$.

► **Claim 15.** $\mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \mathbf{T}}[\text{Revenue}] = - \int_{t=0}^1 \alpha(t) \cdot r'(t) dt.$

Proof. This follows directly from the definition of Revenue. See [9] for details. ◀

► **Lemma 16.** $\mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \mathbf{T}}[\text{Utility}] \geq \int_{t=0}^1 (1 - \alpha(t)) \cdot r(t) dt.$

Proof. The utility for element i arriving at time t is given by

$$\mathbb{E}_{\mathbf{v}, \mathbf{T}}[u_i \mid T_i = t] = \mathbb{E}_{\mathbf{v}, \mathbf{T}_{-i}} \left[(v_i - \alpha(t) \cdot b_i(A_t))^+ \cdot \mathbf{1}_{i \notin \text{Span}(A_t)} \mid T_i = t \right].$$

Observe that A_t does not depend on v_i if $T_i = t$ because it includes only the acceptances *before* t . It does not depend on \hat{v}_i either, as \hat{v}_i is only used for analysis purposes and not known to the algorithm. Since v_i and \hat{v}_i are identically distributed, we can also write

$$\mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \mathbf{T}}[u_i \mid T_i = t] = \mathbb{E}_{\mathbf{v} \sim \mathcal{D}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}, \mathbf{T}_{-i}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t))^+ \cdot \mathbf{1}_{i \notin \text{Span}(A_t)} \mid T_i = t \right]. \quad (6)$$

Now observe that element i can belong to $\text{Opt}(\hat{\mathbf{v}} \mid A_t)$ only if it's not already in $\text{Span}(A_t)$, which implies $\mathbf{1}_{i \notin \text{Span}(A_t)} \geq \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)}$. Using this and removing non-negativity, we get

$$\mathbb{E}_{\mathbf{v}, \mathbf{T}}[u_i \mid T_i = t] \geq \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}, \mathbf{T}_{-i}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)} \mid T_i = t \right].$$

Now we use Lemma 17 to remove the conditioning on element i arriving at time t as this gives a valid lower bound on expected utility,

$$\mathbb{E}_{\mathbf{v}, \mathbf{T}}[u_i \mid T_i = t] \geq \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}, \mathbf{T}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)} \right]. \quad (7)$$

We can now lower bound sum of all the utilities using Eq. (7) to get

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, \mathbf{T}}[\text{Utility}] &= \sum_i \int_{t=0}^1 \mathbb{E}_{\mathbf{v}, \mathbf{T}}[u_i \mid T_i = t] \cdot dt \\ &\geq \sum_i \int_{t=0}^1 \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}} \sim \hat{\mathcal{D}}, \mathbf{T}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)} \right] \cdot dt. \end{aligned}$$

By moving the sum over elements inside the integrals, we get

$$\begin{aligned} \mathbb{E}_{\mathbf{v}, \mathbf{T}}[\text{Utility}] &\geq \int_{t=0}^1 \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}, \mathbf{T}} \left[\sum_i (\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)} \right] \cdot dt \\ &= \int_{t=0}^1 \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}, \mathbf{T}} \left[R(A_t, \hat{\mathbf{v}}) - \alpha(t) \cdot \sum_{i \in \text{Opt}(\hat{\mathbf{v}} \mid A_t)} b_i(A_t) \right] \cdot dt. \end{aligned}$$

Finally, using Claim 18 for $S = \text{Opt}(\hat{\mathbf{v}} \mid A_t)$, we get

$$\mathbb{E}_{\mathbf{v}, \mathbf{T}}[\text{Utility}] \geq \int_{t=0}^1 \mathbb{E}_{\mathbf{v}, \hat{\mathbf{v}}, \mathbf{T}} \left[(1 - \alpha(t)) \cdot R(A_t, \hat{\mathbf{v}}) \right] \cdot dt. \quad \blacktriangleleft$$

Proof of Theorem 14. Using Lemma 16 and Claim 15, we get

$$\begin{aligned}\mathbb{E}[\text{Alg}] &= \mathbb{E}[\text{Revenue}] + \mathbb{E}[\text{Utility}] \\ &\geq - \int_{t=0}^1 \alpha(t) \cdot r'(t) \cdot dt + \int_{t=0}^1 (1 - \alpha(t)) \cdot r(t) \cdot dt \\ &= \int_{t=0}^1 r(t) \cdot (1 - \alpha(t) + \alpha'(t)) \cdot dt - [r(t) \cdot \alpha(t)]_{t=0}^1.\end{aligned}$$

Notice that for $\alpha(t) = 1 - e^{t-1}$, we have $1 - \alpha(t) + \alpha'(t) = 0$. Hence, we get

$$\mathbb{E}[\text{Alg}] \geq -[r(t) \cdot \alpha(t)]_{t=0}^1 = \left(1 - \frac{1}{e}\right) \cdot r(0) = \left(1 - \frac{1}{e}\right) \cdot \sum_{i \in N} x_i y_i. \quad \blacktriangleleft$$

Finally, we prove the missing Lemma 17 that removes the conditioning on i arriving at t .

► **Lemma 17.** *For any i , any time t , and any fixed $\mathbf{v}, \hat{\mathbf{v}}$, we have*

$$\mathbb{E}_{\mathbf{T}_{-i}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}}|A_t)} \mid T_i = t \right] \geq \mathbb{E}_{\mathbf{T}} \left[(\hat{v}_i - \alpha(t) \cdot b_i(A_t)) \cdot \mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}}|A_t)} \right].$$

Proof. We prove the lemma for any fixed \mathbf{T}_{-i} . Suppose we draw a uniformly random $T_i \in [0, 1]$. Observe that if $T_i \geq t$ then we have equality in the above equation because set A_t is the same both with and without i . This is also the case when $T_i < t$ but i is not selected into A_t . Finally, when $T_i < t$ and $i \in A_t$ we have $\mathbf{1}_{i \in \text{Opt}(\hat{\mathbf{v}}|A_t)} = 0$ in the presence of element i (i.e., RHS of lemma), making the inequality trivially true. ◀

► **Claim 18.** *For any fixed \mathbf{v}, \mathbf{T} , time t , and set of elements $S \subseteq N$ that is independent in the matroid \mathcal{M}/A_t , we have*

$$\sum_{i \in S} b_i(A_t) \leq \mathbb{E}_{\hat{\mathbf{v}}} [R(A_t, \hat{\mathbf{v}})].$$

Proof. By definition

$$\sum_{i \in S} b_i(A_t) = \mathbb{E}_{\hat{\mathbf{v}}} \left[\sum_{i \in S} (R(A_t, \hat{\mathbf{v}}) - R(A_t \cup \{i\}, \hat{\mathbf{v}})) \right].$$

Fix the values $\hat{\mathbf{v}}$ arbitrarily, we also have

$$\sum_{i \in S} (R(A_t, \hat{\mathbf{v}}) - R(A_t \cup \{i\}, \hat{\mathbf{v}})) \leq R(A_t, \hat{\mathbf{v}}).$$

This follows from the fact that $R(A_t, \hat{\mathbf{v}}) - R(A_t \cup \{i\}, \hat{\mathbf{v}})$ are the respective critical values of the greedy algorithm on \mathcal{M}/A_t with values $\hat{\mathbf{v}}$. Therefore, the bound follows from Lemma 3.2 in [23]. An alternative proof is given as Proposition 2 in [20] while in our case the first inequality can be skipped and the remaining steps can be followed replacing A by A_t .

Taking the expectation over $\hat{\mathbf{v}}$, the claim follows. ◀

References

- 1 Marek Adamczyk and Michal Włodarczyk. Random order contention resolution schemes. *CoRR*, abs/1804.02584, 2018.
- 2 Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Price of correlations in stochastic optimization. *Operations Research*, 60(1):150–162, 2012. Preliminary version in SODA 2010.

- 3 Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 512–521, 2011.
- 4 Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- 5 Yossi Azar, Ashish Chiplunkar, and Haim Kaplan. Prophet secretary: Surpassing the $1-1/e$ barrier. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 303–318, 2018.
- 6 Gruia Călinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM J. Comput.*, 40(6):1740–1766, 2011.
- 7 Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320, 2010.
- 8 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM J. Comput.*, 43(6):1831–1879, 2014.
- 9 Soheil Ehsani, Mohammad Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2018.
- 10 Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- 11 Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1014–1033, 2016.
- 12 Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1(2):169–197, 1981.
- 13 Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *STOC*, pages 104–113. ACM, 2007. Full version as: *Approximation Algorithms for Bayesian Multi-Armed Bandit Problems*, <http://arxiv.org/abs/1306.3525>.
- 14 Anupam Gupta, Haotian Jiang, Ziv Scully, and Sahil Singla. The markovian price of information, 2018.
- 15 Anupam Gupta and Viswanath Nagarajan. A stochastic probing problem with applications. In *Integer Programming and Combinatorial Optimization - 16th International Conference, IPCO 2013, Valparaíso, Chile, March 18-20, 2013. Proceedings*, pages 205–216, 2013.
- 16 Anupam Gupta, Viswanath Nagarajan, and Sahil Singla. Algorithms and adaptivity gaps for stochastic probing. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1731–1747. SIAM, 2016.
- 17 Guru Guruganesh and Euiwoong Lee. Understanding the Correlation Gap For Matchings. In *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2017)*, volume 93 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 32:1–32:15, 2018.
- 18 Mohammad Taghi Hajiaghayi, Robert D. Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*, pages 58–65, 2007.
- 19 Theodore P Hill and Robert P Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.

- 20 Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136, 2012.
- 21 Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bull. Am. Math. Soc*, 1977.
- 22 Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Prob*, 4:197–266, 1978.
- 23 Brendan Lucier and Allan Borodin. Price of anarchy for greedy auctions. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 537–553. Society for Industrial and Applied Mathematics, 2010.
- 24 Aviad Rubinfeld. Beyond matroids: secretary problem and prophet inequality with general constraints. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*, pages 324–332, 2016.
- 25 Aviad Rubinfeld and Sahil Singla. Combinatorial prophet inequalities. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1671–1687. SIAM, 2017.
- 26 Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. Society for Industrial and Applied Mathematics, 2011.