

The Outer Limits of Contention Resolution on Matroids and Connections to the Secretary Problem

Shaddin Dughmi

Department of Computer Science, University of Southern California, Los Angeles, CA, USA
shaddin@usc.edu

Abstract

Contention resolution schemes have proven to be a useful and unifying abstraction for a variety of constrained optimization problems, in both offline and online arrival models. Much of prior work restricts attention to product distributions for the input set of elements, and studies contention resolution for increasingly general packing constraints, both offline and online. In this paper, we instead focus on generalizing the input distribution, restricting attention to matroid constraints in both the offline and online random arrival models. In particular, we study contention resolution when the input set is arbitrarily distributed, and may exhibit positive and/or negative correlations between elements. We characterize the distributions for which offline contention resolution is possible, and establish some of their basic closure properties. Our characterization can be interpreted as a distributional generalization of the matroid covering theorem. For the online random arrival model, we show that contention resolution is intimately tied to the secretary problem via two results. First, we show that a competitive algorithm for the matroid secretary problem implies that online contention resolution is essentially as powerful as offline contention resolution for matroids, so long as the algorithm is given the input distribution. Second, we reduce the matroid secretary problem to the design of an online contention resolution scheme of a particular form.

2012 ACM Subject Classification Theory of computation → Design and analysis of algorithms

Keywords and phrases Contention Resolution, Secretary Problems, Matroids

Digital Object Identifier 10.4230/LIPIcs.ICALP.2020.42

Category Track A: Algorithms, Complexity and Games

Related Version <https://arxiv.org/abs/1909.04268>

Funding *Shaddin Dughmi*: This work was supported by NSF CAREER Award CCF-1350900.

1 Introduction

The notion of a contention resolution scheme (CRS) abstracts a familiar task in constrained optimization: converting a (random) set-valued solution which is *ex-ante* (i.e., on average) feasible for a packing problem to one which is *ex-post* (i.e., always) feasible. Unlike randomized rounding algorithms more broadly, which in general may be catered to both the constraint and objective function at hand, a contention resolution scheme is specific only to the constraints of the problem, and preserves solution quality in a manner which is largely agnostic to the objective function¹ – element by element. Since they were formalized by Chekuri et al [10], CRSs have been connected to a variety of online and offline computational tasks, including rounding the solutions of mathematical programs [10], online mechanism design and stochastic probing [17, 1], and prophet inequalities [17, 25].

¹ In its most general form, a CRS approximately preserves all linear objective functions simultaneously, whereas a *monotone* CRS approximately preserves all submodular objectives [10].



© Shaddin Dughmi;
licensed under Creative Commons License CC-BY

47th International Colloquium on Automata, Languages, and Programming (ICALP 2020).

Editors: Artur Czumaj, Anuj Dawar, and Emanuela Merelli; Article No. 42; pp. 42:1–42:18

Leibniz International Proceedings in Informatics



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



Starting with [10], prior work defines an (offline) contention resolution scheme for a set system $(\mathcal{E}, \mathcal{I})$ – where \mathcal{E} is a *ground set of elements* and $\mathcal{I} \subseteq 2^{\mathcal{E}}$ is a downwards-closed family of *feasible sets* – as an algorithm which takes as input the marginal probabilities $x \in [0, 1]^{\mathcal{E}}$ of a product distribution \mathcal{D} supported on $2^{\mathcal{E}}$ as well as a random set $R \sim \mathcal{D}$ of *active elements*, and must output a feasible subset S of R . The contention resolution scheme is α -competitive if $\Pr[i \in S] \geq \frac{1}{\alpha} \Pr[i \in R]$ holds for all product distributions of interest – typically those with marginals x in the convex hull of indicator vectors of \mathcal{I} (ex-ante feasibility). In *online* contention resolution schemes, first explored by Feldman et al [17] and subsequently by Adamczyk and Włodarczyk [1] and Lee and Singla [25], the active elements R arrive sequentially and the decision to include an element in S must irrevocably be made online.

The existing literature on (offline and online) contention resolution has mostly restricted attention to ex-ante-feasible and given product distributions, and varied the set system (e.g. matroids, knapsacks, and their intersections), all the while drawing connections to applications such as stochastic online problems, approximation algorithms, mechanism design, and prophet inequalities. In this paper, we restrict our attention to matroid constraints,² and instead focus on generalizing the class of input distributions. Our main goal is to understand the power and limitations of contention resolution, offline and online, in the presence of correlations in the input distribution and without regard to ex-ante feasibility. A secondary goal is to understand how knowledge of the distribution influences contention resolution. In pursuit of both goals, we draw connections between contention resolution and the secretary problem on matroids, shedding light on challenges posed by the matroid secretary conjecture in the process.

Results

Our first set of results develops an understanding of offline contention resolution on matroids. We begin with a characterization of the class of α -*uncontentious distributions*: those distributions $\mathcal{D} \in \Delta(2^{\mathcal{E}})$ permitting α -competitive offline contention resolution for a given matroid. Most notably, we show that a distribution is α -uncontentious if and only if it satisfies a family of $2^{|\mathcal{E}|}$ inequalities, one for each subset of the ground set. Moreover, we observe that our inequality characterization is the natural generalization of the *matroid base covering theorem* (see e.g. [31]) from covering a set of elements to covering a distribution over sets of elements. In other words, we show that *contention resolution is the natural distributional generalization of base covering*. Leveraging our characterization, we establish some basic closure properties of the class of uncontentious distributions, and present some examples of uncontentious distributions exhibiting negative and positive correlation between elements. Finally, we examine whether knowledge of the distribution \mathcal{D} is essential to contention resolution, and exhibit an impossibility result: any contention resolution scheme which has nontrivial guarantees for all α -uncontentious distributions cannot be *prior-independent*, in that it cannot make do with a finite number of samples from the distribution, even for very simple matroids.

Our second set of results concerns online contention resolution on matroids in the random arrival model, and in particular its connection to the matroid secretary problem. First, we show that a competitive secretary algorithm for a matroid implies that online contention resolution is essentially as powerful as offline contention resolution for that same matroid: a γ -competitive secretary algorithm implies that any α -uncontentious distribution permits $\gamma\alpha$ -competitive online contention resolution.

² Though some of our results hold beyond matroids; we discuss those in the conclusion section.

Second, we provide evidence that contention resolution might hold the key to resolving the matroid secretary conjecture. As our most technically-involved result, we leverage our characterization of uncontentious distributions to show that the random set of *improving elements* in a weighted matroid – as originally defined by Karger [20] – is $O(1)$ -uncontentious. Since the improving elements can be recognized online, and moreover hold a constant fraction of the weighted rank of the matroid in expectation, our result can be loosely interpreted as a reduction from the matroid secretary problem to online contention resolution for a particular uncontentious distribution. There is one major caveat to this interpretation of our result, however: not only does the set of active (improving) elements arrive online, but so does the description of the uncontentious distribution from which that set is drawn. Though we present our proof of this result in an elementary form, the underlying arguments are reminiscent of – and inspired by – those often encountered in the analysis of martingales: we condition on carefully-chosen random variables, and employ a delicate charging argument between different probability events.

Third, in response to feedback on the previous version of this manuscript, we show that our aforementioned result – that improving elements are uncontentious – cannot be derived as a consequence of prior work.

Additional Discussion of Related Work

Contention Resolution Schemes

Contention resolution schemes were introduced by Chekuri et al [10], motivated by the problem of maximizing a submodular function subject to packing constraints. In particular, offline CRS were used to transform a randomized rounding algorithm which respects the packing constraints ex-ante to one which respects them ex-post, at the cost of the competitive ratio of the CRS. Their focus – like that of all related work prior to ours – was on product input distributions, in which case the optimal competitive ratio of an offline CRS was shown to equal the worst-case *correlation gap* (first studied by [2, 7]) of the weighted rank function associated with the packing constraint. The characterization result of [10] result forms the basis for ours.

Online contention resolution was first studied by Feldman et al [17], and applied to a number of online selection problems. They show that simple packing constraints – such as matroids, knapsacks, and matchings – permit constant competitive online contention resolution schemes even when elements arrive in an unknown and adversarial order. Moreover, they show how to combine competitive online schemes for different constraints in order to yield competitive online schemes for their intersection. Lee and Singla [25] obtain optimal online CRS in both the known adversarial-order model as well as the random-arrival model. Adamczyk and Włodarczyk [1] restrict attention to the random-arrival model, and obtain a particularly elegant algorithm and argument based on martingales, as well as improved competitive ratios for intersections of matroids and knapsacks.

Prophet Inequalities

Contention resolution is intimately tied to *prophet inequality* problems, also known as *Bayesian online selection* problems. In the traditional model for these problems, independent real-valued random variables with known distributions arrive online in a known but adversarial order, and the goal is to select a subset of the variables with maximum sum, subject to a packing constraint. An α -competitive algorithm for a Bayesian online selection problem is also referred to as a prophet inequality with ratio α , for historical reasons. Krenkel, Sucheston,

and Garling [22, 23] proved the first (classical) single-choice prophet inequality with ratio $1/2$ for selecting a single variable (i.e., a 1-uniform matroid packing constraint). Motivated by applications in algorithmic mechanism design, more recent work (e.g. [18, 3, 8, 32]) pursued prophet inequalities for more general packing constraints. Of particular note is the work of Kleinberg and Weinberg [21], who proved an optimal prophet inequality with ratio $1/2$ for matroids. Also notable is polylogarithmic prophet inequality for general packing constraints due to Rubinfeld [28]. The (easier) variant of Bayesian online selection problems in which the variables arrive in a uniformly random order has also received recent interest, resulting in improved prophet inequalities for various packing constraints [15, 14, 5].

It was shown by Feldman et al [17] that an online CRS yields a prophet inequality with the same competitive ratio, and in the same arrival model. A weak converse is also true, as shown by Lee and Singla [25]: a stronger form of prophet inequality – in particular one which competes against the *ex-ante relaxation* of the Bayesian online selection problem – yields an online CRS with the same competitive ratio and in the same arrival model.

Beyond Known Product Distributions

The vast majority of work on contention resolution or prophet inequalities, and all such work discussed thus far, restricts attention to known product distributions, and crucially exploits the product structure and knowledge of the distribution. We note the few exceptions next.

Rinott et al [27] and Samuel-Cahn [29] show that the single-choice prophet inequality, and some slight generalizations, continue to hold for negatively dependent random variables. It is known [19] that there is no single-choice prophet inequality with ratio better than the number of variables in the presence of arbitrary positive correlation. Moreover, we are unaware of any nontrivial positive results for a class of distributions exhibiting positive correlation, in either prophet inequality or contention resolution models. We note that whereas [25] and [1] use specially-crafted correlated distributions as benchmarks, their results and techniques do not appear to shed light on contention resolution or prophet inequalities in the presence of correlation more generally.

Some work has relaxed the requirement that the distributions be known in prophet inequality problems. Azar et al [4] study prophet inequality problems when only a single sample is given from each distribution, and obtain constant competitive ratios for a variety of constraints. Wang [30] obtains an optimal algorithm for the single-choice prophet inequality, with ratio $1/2$, in the same single-sample model. Correa et al [11] study the single-choice prophet inequality with i.i.d. variables drawn from an unknown distribution, and characterize the relationship between the competitive ratio and the number of samples available from the distribution.

Secretary Problems

In a *generalized secretary problem*, a set of adversarially chosen variables arrive online in a random order, and the goal is to select a subset of the variables with maximum sum subject to a packing constraint. The (classical) single-choice *secretary problem*, corresponding to a 1-uniform matroid constraint, was introduced and solved by Dynkin [13]. The *matroid secretary problem* was introduced by Babaioff et al [6], and has since spawned a long line of work. Constant-competitive algorithms have been discovered for most natural matroids and for some alternative models – see [12] for a semi-recent survey – though the general conjecture remains open. The state of the art for the general matroid secretary problem is a $O(\log \log \mathbf{rank})$ -competitive algorithm due to Lachish [24], which was henceforth simplified by Feldman et al [16].

2 Preliminaries

2.1 Matroid Theory Basics

We use standard definitions from matroid theory; for details see [26, 31]. A matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ consists of a *ground set* \mathcal{E} of *elements*, and a family $\mathcal{I} \subseteq 2^{\mathcal{E}}$ of *independent sets*, satisfying the three *matroid axioms*. A *weighted matroid* (\mathcal{M}, w) consists of a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ together with weights $w \in \mathbb{R}^{\mathcal{E}}$ on the elements. We use the standard notions of a *dependent set*, *circuit*, *flat*, and *minor* in a matroid. We denote the *rank* of a matroid \mathcal{M} as $\mathbf{rank}(\mathcal{M})$, and the rank of a set of elements A in \mathcal{M} as $\mathbf{rank}_{\mathcal{M}}(A)$, or $\mathbf{rank}(A)$ when \mathcal{M} is clear from context. Overloading notation, we use $\mathbf{rank}_w^{\mathcal{M}}(A)$ to denote the *weighted rank* of a set A – the maximum weight of an independent subset of A – in the weighted matroid (\mathcal{M}, w) , though we omit the superscript \mathcal{M} when the matroid is clear from context. We note that both rank and weighted rank are submodular set functions on the ground set of the matroid. For $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and $A \subseteq \mathcal{E}$, we denote the *restriction* of \mathcal{M} to A as $\mathcal{M}|_A$, *deletion* of A as $\mathcal{M} \setminus A$, and *contraction* by A as \mathcal{M}/A .

When \mathcal{E} is clear from context, and $S \subseteq \mathcal{E}$, we use $\mathbf{1}_S \in \{0, 1\}^{\mathcal{E}}$ to denote the vector indicating membership in S . We often reference the *matroid polytope* $\mathcal{P}(\mathcal{M})$ of a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, defined as the convex hull of $\{\mathbf{1}_S : S \in \mathcal{I}\}$, or equivalently as the family of $x \in [0, 1]^{\mathcal{E}}$ satisfying $\sum_{i \in S} x_i \leq \mathbf{rank}_{\mathcal{M}}(S)$ for all $S \subseteq \mathcal{E}$.

Throughout this paper we assume that any weighted matroid has distinct weights. This assumption is made merely to simplify some of our proofs, and – using standard tie-breaking arguments – can be shown to be without loss of generality in as much as our results are concerned. Under this assumption, we define $\mathbf{OPT}_w^{\mathcal{M}}(A)$ as the (unique) maximum-weight independent subset of A of minimum cardinality (excluding zero-weight elements), and we omit the superscript when the matroid is clear from context.

2.2 The Matroid Secretary Problem

In the *matroid secretary problem*, originally defined by [6] there is matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with nonnegative weights $w : \mathcal{E} \rightarrow \mathbb{R}_+$ on the elements. Elements \mathcal{E} arrive online in a uniformly random order Π , and an online algorithm must irrevocably accept or reject an element when it arrives, subject to accepting an independent set of \mathcal{M} . The algorithm is given \mathcal{M} at the outset (as an independence oracle), but the weights w are chosen adversarially before the order Π is drawn and then are revealed online. The goal of the online algorithm is to maximize the expected weight of the accepted set of elements. We say that an algorithm is α -*competitive* for a class of matroids if for every matroid \mathcal{M} in that class and every adversarial choice of w , the expected weight of the accepted set (over the random choice of Π and any internal randomness of the algorithm) is at least an α fraction of the maximum weight of an independent set of (\mathcal{M}, w) .

The *matroid secretary conjecture*, posed by [6], postulates that the matroid secretary problem admits an (online) algorithm which is constant-competitive for all matroids.

2.3 Miscellaneous Notation and Terminology

We denote the natural numbers by \mathbb{N} , and the nonnegative real numbers by \mathbb{R}_+ . Given a set \mathcal{A} with weights $w \in \mathbb{R}^{\mathcal{A}}$, and a subset $B \subseteq \mathcal{A}$, we use the shorthand $w(B) = \sum_{i \in B} w_i$. We use $[n]$ as shorthand for the set $1, \dots, n$. For a set A , we use $\Delta(A)$ to denote the family of distributions over A , and 2^A to denote the family of subsets of A .

Let \mathcal{A} be a finite ground set. For a distribution \mathcal{D} supported on $2^{\mathcal{A}}$, we define the vector $x \in [0, 1]^{\mathcal{A}}$ of *marginals* of \mathcal{D} by $x_i = \Pr_{B \sim \mathcal{D}}[i \in B]$, and refer to x_i as the *marginal probability* of i in \mathcal{D} . When marginals $x \in [0, 1]^{\mathcal{A}}$ are given, we use $\mathbf{Ind}(x)$ to denote the distribution of the random set $B \subseteq \mathcal{A}$ which includes each element $i \in \mathcal{A}$ independently with probability x_i . We also use $\mathbf{Ind}_p(\mathcal{A})$ as shorthand for $\mathbf{Ind}(x)$ when $x_i = p$ for all $i \in \mathcal{A}$.

3 Understanding Contention Resolution

3.1 The Basics of Contention Resolution

The definitions below are parametrized by a given matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$.

► **Definition 3.1.** A contention resolution map (CRM) ϕ is a randomized function from $2^{\mathcal{E}}$ to \mathcal{I} with the property that $\phi(R) \subseteq R$ for all $R \subseteq \mathcal{E}$. Such a map is α -competitive for a distribution $\mathcal{D} \in \Delta(2^{\mathcal{E}})$ if, for $R \sim \mathcal{D}$, we have $\Pr[i \in R] \leq \alpha \Pr[i \in \phi(R)]$ for all $i \in \mathcal{E}$.

The following is known from Chekuri et al [10].

► **Theorem 3.2** ([10]). Every product distribution with marginals in $\mathcal{P}(\mathcal{M})$ admits an $\frac{e}{e-1}$ -competitive CRM.

► **Definition 3.3.** An online random-order contention resolution map (henceforth online CRM for short) is a contention resolution map ϕ which can be implemented as an algorithm in the online random-arrival model. In the online random-arrival model, \mathcal{E} is presented to the algorithm in a uniformly random order (e_1, \dots, e_n) , and at the i th step the algorithm learns whether e_i is active – i.e., $e_i \in R \sim \mathcal{D}$ – and if so must make an irrevocable decision on whether to include e_i in $\phi(R)$.

The following is known from Lee and Singla [25].

► **Theorem 3.4** ([25]). Every product distribution with marginals in $\mathcal{P}(\mathcal{M})$ admits a $\frac{e}{e-1}$ -competitive online CRM.

3.2 Uncontentious Distributions and their Characterization

As shorthand, we refer to distributions which permit competitive (offline) CRMs as *uncontentious*.

► **Definition 3.5.** Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$. For $\alpha \geq 1$, we say that a distribution $\mathcal{D} \in \Delta(2^{\mathcal{E}})$ is α -uncontentious if it admits an α -competitive contention resolution map.

For convenience, we also refer to a random set $R \sim \mathcal{D}$ as α -uncontentious if its distribution \mathcal{D} is α -uncontentious. We prove the following characterization of uncontentious distributions.

► **Theorem 3.6.** Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let \mathcal{D} be a distribution supported on $2^{\mathcal{E}}$. The following are equivalent for every $\alpha \geq 1$.

- (a) \mathcal{D} is α -uncontentious (i.e., admits an α -competitive contention resolution map).
- (b) For every weight vector $w \in \mathbb{R}_+^{\mathcal{E}}$, the following holds for $R \sim \mathcal{D}$: $\mathbf{E}[\mathbf{rank}_w(R)] \geq \frac{\mathbf{E}[w(R)]}{\alpha}$.
- (c) For every $\mathcal{F} \subseteq \mathcal{E}$, the following holds for $R \sim \mathcal{D}$: $\mathbf{E}[|R \cap \mathcal{F}|] \leq \alpha \mathbf{E}[\mathbf{rank}(R \cap \mathcal{F})]$.

Proof. Property (a) implies property (c) by applying an α -CRM ϕ to R , noting that $\phi(R) \cap \mathcal{F}$ is necessarily an independent subset of $R \cap \mathcal{F}$.

$$\begin{aligned} \mathbf{E}[\mathbf{rank}(R \cap \mathcal{F})] &\geq \mathbf{E}[|\phi(R) \cap \mathcal{F}|] \\ &= \sum_{i \in \mathcal{F}} \Pr[i \in \phi(R)] \\ &\geq \frac{1}{\alpha} \sum_{i \in \mathcal{F}} \Pr[i \in R] \\ &= \frac{1}{\alpha} \mathbf{E}[|R \cap \mathcal{F}|]. \end{aligned}$$

Property (c) implies property (b) by a summation argument. Sort and number the elements $\mathcal{E} = (e_1, \dots, e_n)$ in decreasing order of weights $w_1 \geq w_2 \geq \dots \geq w_n \geq 0$, where w_i denotes the weight of e_i . Denote $\mathcal{E}_i = \{e_1, \dots, e_i\}$, and let $\mathcal{E}_0 = \emptyset$, and $w_{n+1} = 0$. Recalling that the greedy algorithm computes the maximum weight independent subset of a matroid:

$$\begin{aligned} \mathbf{E}[\mathbf{rank}_w(R)] &= \mathbf{E} \left[\sum_{i=1}^n w_i (\mathbf{rank}(R \cap \mathcal{E}_i) - \mathbf{rank}(R \cap \mathcal{E}_{i-1})) \right] && \text{(Greedy alg. on } \mathcal{M}|R) \\ &= \mathbf{E} \left[\sum_{i=1}^n (w_i - w_{i+1}) \mathbf{rank}(R \cap \mathcal{E}_i) \right] && \text{(Reverse summations)} \\ &\geq \frac{1}{\alpha} \mathbf{E} \left[\sum_{i=1}^n (w_i - w_{i+1}) |R \cap \mathcal{E}_i| \right] && \text{((c) and linearity of exp.)} \\ &= \frac{1}{\alpha} \mathbf{E} \left[\sum_{i=1}^n w_i (|R \cap \mathcal{E}_i| - |R \cap \mathcal{E}_{i-1}|) \right] && \text{(Reverse summations)} \\ &= \frac{1}{\alpha} \mathbf{E}[w(R)]. \end{aligned}$$

Property (b) implies property (a) by a duality argument identical to that presented in [10]. We present a self-contained proof here. Let $x = x(\mathcal{D}) \in [0, 1]^\mathcal{E}$ denote the marginals of \mathcal{D} . The distribution \mathcal{D} is α -uncontentious if the optimal value of the following LP, with variables β and λ_ϕ for each deterministic CRM ϕ , is at least $\frac{1}{\alpha}$.

$$\begin{aligned} &\text{maximize } \beta \\ &\text{subject to } \sum_{\phi} \lambda_{\phi} \Pr_{R \sim \mathcal{D}}[i \in \phi(R)] \geq \beta x_i, \quad \text{for } i \in \mathcal{E}. \\ &\quad \sum_{\phi} \lambda_{\phi} = 1 \\ &\quad \lambda \succeq 0 \end{aligned}$$

The dual of the preceding LP is the following

$$\begin{aligned} &\text{minimize } \mu \\ &\text{subject to } \sum_{i \in \mathcal{E}} w_i \Pr_{R \sim \mathcal{D}}[i \in \phi(R)] \leq \mu, \quad \text{for all CRM } \phi. \\ &\quad \sum_{i \in \mathcal{E}} w_i x_i = 1 \\ &\quad w \succeq 0 \end{aligned}$$

It is not hard to see that, at optimality, the binding constraint on μ corresponds to the CRM ϕ which maps each set R to its maximum weight independent subset according to weights w . It follows that the optimal value of the dual, and hence the primal, equals the minimum over all weight vectors $w \succeq 0$ of the ratio $\frac{\mathbf{E}[\mathbf{rank}_w(R)]}{\sum_i w_i x_i}$. (b) implies that this quantity is at least $\frac{1}{\alpha}$, as needed. \blacktriangleleft

We note that the equivalence between (a) and (b) is essentially implicit in the arguments of [10]. Condition (c) is the most notable part of Theorem 3.6, in no small part because it is reminiscent of the *matroid base covering theorem* (see e.g., [31]). This theorem can equivalently be stated as follows: a (deterministic) set $T \subseteq \mathcal{E}$ in a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ can be *covered by* (i.e., expressed as a union of) $\alpha \in \mathbb{N}$ independent sets if and only if $|S| \leq \alpha \mathbf{rank}_{\mathcal{M}}(S)$ for all $S \subseteq T$. In light of part (c) of Theorem 3.6, a set T of elements can be covered by α independent sets if and only if the point distribution on T is α -uncontentious. Therefore, *we can interpret contention resolution as a distributional generalization of base covering.*

3.3 Elementary Properties of Uncontentious Distributions

We state some elementary, yet quite useful, properties of uncontentious distributions.

► **Proposition 3.7.** *Fix a matroid \mathcal{M} . Every α -uncontentious distribution \mathcal{D} for $\alpha \geq 1$ has marginals $x(\mathcal{D}) \in \alpha\mathcal{P}(\mathcal{M})$.*

Proof. Let $x = x(\mathcal{D})$ and $R \sim \mathcal{D}$. From Theorem 3.6 (c), for every set of ground set elements \mathcal{F} we have

$$\sum_{i \in \mathcal{F}} x_i = \mathbf{E}[|R \cap \mathcal{F}|] \leq \alpha \mathbf{E}[\mathbf{rank}_{\mathcal{M}}(R \cap \mathcal{F})] \leq \alpha \mathbf{rank}_{\mathcal{M}}(\mathcal{F}).$$

These are the inequalities describing $\alpha\mathcal{P}(\mathcal{M})$. ◀

► **Proposition 3.8.** *Fix a matroid. A mixture of α -uncontentious distributions is α -uncontentious.*

Proof. Follows directly from Theorem 3.6 (b) and linearity of expectations. ◀

► **Proposition 3.9.** *Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let $\mathcal{M}' = (\mathcal{E}', \mathcal{I}')$ be a minor of \mathcal{M} , with $\mathcal{E}' \subseteq \mathcal{E}$. If a random set $R \subseteq \mathcal{E}'$ is α -uncontentious in \mathcal{M}' , then R is also α -uncontentious in \mathcal{M} .*

Proof. An independent set of \mathcal{M}' is also independent in \mathcal{M} . Therefore, the proposition follows by simply applying the same CRM in the context of the larger matroid \mathcal{M} . ◀

► **Proposition 3.10.** *Fix a matroid. Let R be an α -uncontentious random set, and let $R' \sim \mathbf{Ind}_p(R)$ for some $p \in [0, 1]$. The random set R' is α -uncontentious as well.*

Proof. We use Theorem 3.6 (b). For any weight vector w , submodularity of the weighted rank function implies that $\mathbf{E}[\mathbf{rank}_w(R')] \geq p \mathbf{E}[\mathbf{rank}_w(R)]$. It follows that $\mathbf{E}[w(R')] = p \mathbf{E}[w(R)] \leq p\alpha \mathbf{E}[\mathbf{rank}_w(R)] \leq \alpha \mathbf{E}[\mathbf{rank}_w(R')]$. ◀

We note that Proposition 3.10 is tight when both p and α are absolute constants. In particular, the random set R' cannot be guaranteed to be α' -uncontentious for a constant $\alpha' < \alpha$, even if p is a very small constant. To see this, consider the a 1-uniform matroid with elements $[n]$, and the following 2-uncontentious random set R : For every singleton $i \in [n]$ we have $\Pr[R = \{i\}] = \frac{1}{n+1}$, and $\Pr[R = [n]] = \frac{1}{n+1}$.

3.4 Examples of Uncontentious Distributions

We now present some examples of uncontentious distributions in order to develop a feel for them. As mentioned previously, and shown in [10], every product distribution with marginals in the matroid polytope is $\frac{e}{e-1}$ -uncontentious. Combined with Proposition 3.8, this extends to mixtures of product distributions.

► **Proposition 3.11.** *Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let $\mathcal{D} \in \Delta(2^{\mathcal{E}})$ be a mixture of product distributions, each with marginals in $\mathcal{P}(\mathcal{M})$. It follows that \mathcal{D} is $\frac{e}{e-1}$ -uncontentious.*

Going beyond product distributions and their mixtures, if a distribution satisfies a certain strong notion of negative correlation, defined in [9], then it also is $\frac{e}{e-1}$ -uncontentious.

► **Proposition 3.12.** *Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let $\mathcal{D} \in \Delta(2^{\mathcal{E}})$ be a distribution with marginals $x = x(\mathcal{D}) \in \mathcal{P}(\mathcal{M})$. Assume that \mathcal{D} satisfies the property of increasing submodular expectations: for every submodular function f we have $\mathbf{E}_{R \sim \mathcal{D}}[f(R)] \geq \mathbf{E}_{S \sim \text{Ind}(x)}[f(S)]$.³ It follows that \mathcal{D} is $\frac{e}{e-1}$ -uncontentious.*

Proof. This is immediate by combining Theorem 3.6 (b) with the property of increasing submodular expectations and the fact that $\text{Ind}(x)$ is $\frac{e}{e-1}$ -uncontentious. ◀

As shown in [9], the property of increasing submodular expectations is stronger than the following standard notion of negative correlation for $R \sim \mathcal{D}$: For all sets T , $\Pr[T \subseteq R] \leq \prod_{i \in T} \Pr[i \in R]$ and $\Pr[T \subseteq \bar{R}] \leq \prod_{i \in T} (1 - \Pr[i \in R])$.⁴ However, we can show that there are distributions exhibiting positive correlation which are also uncontentious for specific matroids. We now list some examples of uncontentious distributions exhibiting positive correlation.

► **Example 3.13.** Let \mathcal{M} be a k -uniform matroid with n elements where $2 \leq k \leq n$. Let the random set R be empty with probability $1/2$, and a uniformly random base of \mathcal{M} otherwise.

It is clear that R is 1-uncontentious, since it is supported on the family of independent sets. However, for each distinct pair of elements i and j , we have $\Pr[i \in R] = \Pr[j \in R] = \frac{k}{2n}$, yet $\Pr[i \in R | j \in R] = \Pr[j \in R | i \in R] = \frac{k-1}{n-1} > \frac{k}{2n}$.

The next example will feature repeatedly in this paper, since it is the random set of improving elements for the rank 1 matroid.

► **Example 3.14.** Consider the 1-uniform matroid with elements $[n] = \{1, \dots, n\}$. For $k = 0, 1, \dots, n-1$, let $R = \{1, \dots, k\}$ with probability $2^{-(k+1)}$, and let $R = [n]$ with remaining probability 2^{-n} . The random set R is 2-uncontentious, as evidenced by the CRM ϕ with $\phi(\{1, \dots, k\}) = \{k\}$ and $\phi(\emptyset) = \emptyset$, and a simple calculation. Note the positive correlation between elements $i < j$: $\Pr[j \in R] = 2^{-j}$, and $\Pr[j \in R | i \in R] = 2^{i-j} > \Pr[i \in R]$.

As a generalization of the previous example, we get the following.

► **Example 3.15.** Let \mathcal{M} be a matroid with m pairwise-disjoint bases B_1, \dots, B_m . For each $k = 1, \dots, m-1$, let $R = \cup_{i=1}^k B_i$ with probability 2^{-k} , and let $R = \cup_{i=1}^m B_i$ with the remaining probability 2^{1-m} . The set R is 2-uncontentious, as evidenced by the CRM $\phi(\cup_{i=1}^k B_i) = B_k$. However, for $e_i \in B_i$ and $e_j \in B_j$ with $i < j$, we have $\Pr[e_j \in R] = 2^{1-j}$ and $\Pr[e_j \in R | e_i \in R] = 2^{i-j} > \Pr[e_j \in R]$.

³ In fact, it suffices for \mathcal{D} to satisfy the (weaker) property of increasing expectations for matroid rank functions (or, equivalently, their weighted sums).

⁴ A natural question is whether negative correlation suffices for the distribution to be $\frac{e}{e-1}$ -uncontentious. This is open as far as we know.

3.5 Contention Resolution Schemes, Universality, and Prior Dependence

A *contention resolution scheme (CRS)* Φ for a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and class of distributions $\mathbb{D} \subseteq \Delta(2^{\mathcal{E}})$ is an algorithm which takes as input a (possibly partial) description of a distribution $\mathcal{D} \in \mathbb{D}$ and a sample $R \sim \mathcal{D}$, and outputs $S \in \mathcal{I}$ satisfying $S \subseteq R$. In effect, Φ is a collection of contention resolution maps $\phi_{\mathcal{D}}$, one for each $\mathcal{D} \in \mathbb{D}$. In much of the prior work on contention resolution, \mathbb{D} was taken to be the class of product distributions with marginals in $\mathcal{P}(\mathcal{M})$, and each $\mathcal{D} \in \mathbb{D}$ is described completely via its marginals $x \in \mathcal{P}(\mathcal{M})$. In such a setting, the notion of a CRS offers little beyond the notion of a CRM, as each distribution gets its own dedicated CRM. More generally, however, we allow \mathbb{D} to be an arbitrary class of distributions, and we allow the description to be partial and/or random; for example, \mathcal{D} may be described by m independent samples from \mathcal{D} .

Next, we set the stage by defining some desirable contention resolution schemes, and establish some limitations on their existence.

► **Definition 3.16.** *Fix a matroid. For $\beta \geq \alpha > 1$, an α -universal β -competitive CRS is a CRS which is β -competitive for the class of α -uncontentious distributions.*

By definition, there exists an (offline) α -universal α -competitive CRS for every α and every matroid. The notion of a universal scheme becomes more interesting when we restrict dependence on the prior, as per the following definitions.

► **Definition 3.17.** *Fix a matroid. A contention resolution scheme Φ is said to be prior-independent if it is not given a complete description of \mathcal{D} as input, but rather is given a set of independent samples from \mathcal{D} . When the number of samples is m , we say Φ is a prior-independent m -sample scheme. The number of samples may be function of the size of the matroid. If $m = 0$, we say the scheme is oblivious: the scheme consists of a single contention resolution map.*

We now show that, if a scheme is universal, it cannot be prior-independent with any finite number of samples, even for very simple matroids.

► **Theorem 3.18.** *Let \mathcal{M} be the 1-uniform matroid on n elements. For every finite m , and every $1 < \alpha \leq \beta < n$, there does not exist a β -competitive α -universal CRS for \mathcal{M} which is prior independent with m samples.*

To prove Theorem 3.18, we first show that a prior-independent universal scheme implies the existence of an oblivious universal scheme; then we show that an oblivious universal scheme does not exist for the uniform matroid. This is captured in the two following lemmas.

► **Lemma 3.19.** *Fix a matroid \mathcal{M} . If there exists a β -competitive α -universal CRS Φ which is prior-independent with m samples, then there exists an oblivious β -competitive α -universal scheme Φ' .*

Proof. Let \mathcal{D} be any α -uncontentious distribution. Let $\epsilon \in (0, 1)$, and let $\mathcal{D}' = \mathcal{D}'(\epsilon)$ be the mixture of \mathcal{D} with the point distribution on the empty set with proportions ϵ and $1 - \epsilon$ respectively. By Proposition 3.8 and the fact that the point distribution on the empty set is 1-uncontentious, it follows that \mathcal{D}' is α -uncontentious.

The CRS Φ induces a CRM $\phi_{\mathcal{D}'}$ on the distribution \mathcal{D}' , and by assumption $\phi_{\mathcal{D}'}$ is β -competitive for \mathcal{D}' . Since Φ is prior-independent with m samples, its induced CRM $\phi_{\mathcal{D}'}$ is a mixture over CRMs ϕ_S , where $S = (S_1, \dots, S_m)$ is a random vector of m samples from \mathcal{D}' .

With probability at least $(1-\epsilon)^m$, we have $S = \emptyset^m := (\emptyset, \dots, \emptyset)$. For $\phi_{\mathcal{D}'}$ to be β -competitive, in particular when with probability ϵ it is queried with a draw $R \sim D$, a simple calculation shows that ϕ_{\emptyset^m} must be β' -competitive for \mathcal{D} for $\beta' = \frac{(1-\epsilon)^m}{1/\beta + (1-\epsilon)^m - 1}$. As ϵ tends to 0, β' tends to β , and a basic analytic argument implies that ϕ_{\emptyset^m} is β -competitive for \mathcal{D} . Since \mathcal{D} was chosen arbitrarily among α -uncontentious distributions, and ϕ_{\emptyset^m} does not depend on \mathcal{D} , it follows that the oblivious scheme Φ' with $\phi'_{\mathcal{D}} = \phi_{\emptyset^m}$ for every \mathcal{D} is β -competitive and α -universal. \blacktriangleleft

► **Lemma 3.20.** *The 1-uniform matroid with n elements does not admit an oblivious β -competitive α -universal CRS for any $1 < \alpha \leq \beta < n$.*

Proof. Let $[n]$ be the ground set of the matroid, and fix α such that $1 < \alpha < n$. An oblivious CRS consists of a single CRM ϕ . There exists at least one element $i \in [n]$ such that $\Pr[i \in \phi([n])] \leq 1/n$. Let $\epsilon = \alpha - 1 > 0$, and consider the following random set R : For each $j \in [n] \setminus i$ we have $R = \{j\}$ with probability $\frac{1}{n-1+\epsilon}$, and $R = [n]$ with the remaining probability $\frac{\epsilon}{n-1+\epsilon}$. The random set R is α -uncontentious: consider the CRM ϕ' with $\phi'(\{j\}) = j$ for $j \neq i$, and $\phi'([n]) = i$. However, our original CRM ϕ is no better than n -competitive for R , since its probability of selecting i is no more than $\frac{1}{n} \Pr[R = [n]] = \frac{1}{n} \Pr[i \in R]$. \blacktriangleleft

4 An Online Universal CRS from a Secretary Algorithm

We show that competitive matroid secretary algorithms imply that every contention resolution scheme can be made online, in the random arrival model, without much loss.

► **Theorem 4.1.** *Suppose that there is a γ -competitive online algorithm for the secretary problem on matroid \mathcal{M} . It follows that every α -uncontentious distribution admits an online $\gamma\alpha$ -competitive contention resolution map. In other words, for every α there exists an online $\gamma\alpha$ -competitive α -universal contention resolution scheme for \mathcal{M} .*

We interpret the above theorem as follows: the design of competitive universal online schemes is a necessary technical hurdle towards resolving the matroid secretary conjecture.

We now proceed with proving Theorem 4.1. Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let $\mathcal{D} \in \Delta(2^{\mathcal{E}})$. Recall that an online CRM operates in the following model: a set of active elements $R \sim \mathcal{D}$ and a random permutation Π are (independently) sampled by nature, then \mathcal{E} arrive online in order Π . When $i \in \mathcal{E}$ arrives, it is revealed whether $i \in R$, and if so the online CRS must determine whether to select i . The online CRM must only select an independent subset of R .

Suppose we are given a secretary algorithm \mathcal{A} for \mathcal{M} with competitive ratio γ . Without loss of generality, we assume that \mathcal{A} selects only non-zero weight elements. Consider the following online CRM ϕ_w for \mathcal{M} , parametrized by a weight vector $w \in \mathbb{R}_+^{\mathcal{E}}$. When element i arrives, if $i \in R$ then it is presented to \mathcal{A} with weight w_i , and if $i \notin R$ then it is presented to \mathcal{A} with weight 0. ϕ_w selects precisely the elements selected by \mathcal{A} .

► **Lemma 4.2.** *For every distribution \mathcal{D} , we have $\mathbf{E}_{R \sim \mathcal{D}}[w(\phi_w(R))] \geq \frac{1}{\gamma} \mathbf{E}_{R \sim \mathcal{D}}[\mathbf{rank}_w(R)]$.*

Proof. Condition on the choice of R , and let $w'_i = w_i$ if $i \in R$ and $w'_i = 0$ otherwise. \mathcal{E} are presented to \mathcal{A} in a uniformly random order, with weights w'_i , and $\phi_w(R) \subseteq R$ is the set of elements selected by \mathcal{A} . Since \mathcal{A} is γ -competitive, it follows that $\mathbf{E}[w'(\phi_w(R))] \geq \frac{1}{\gamma} \mathbf{rank}_{w'}(\mathcal{M})$. Since $w'(\phi_w(R)) = w(\phi_w(R))$ and $\mathbf{rank}_{w'}(\mathcal{M}) = \mathbf{rank}_w(R)$, we are done. \blacktriangleleft

► **Lemma 4.3.** *If \mathcal{D} is α -uncontentious, then $\mathbf{E}_{R \sim \mathcal{D}}[w(\phi_w(R))] \geq \frac{1}{\gamma\alpha} \mathbf{E}_{R \sim \mathcal{D}}[w(R)]$.*

Proof. Combining the previous lemma with Theorem 3.6 (b). \blacktriangleleft

42:12 The Outer Limits of Contention Resolution

Recall that we are assuming for now that we know the α -uncontentious distribution \mathcal{D} , and we can design an online CRM $\phi_{\mathcal{D}}$ accordingly. $\phi_{\mathcal{D}}$ will be a random mixture of the maps ϕ_w described above; in particular, we will show that there exists a distribution $\mathcal{W} = \mathcal{W}(\mathcal{D})$ over weight vectors such that the (randomized) online CRM $\phi_{\mathcal{W}}$ which samples $w \sim \mathcal{W}$ upfront then invokes ϕ_w is an online $\gamma\alpha$ -CRM for \mathcal{D} .

For each element $i \in \mathcal{E}$, let $x_i = \Pr_{R \sim \mathcal{D}}[i \in R]$. For each weight vector w and $i \in \mathcal{E}$, let $y_i(w) = \Pr_{R \sim \mathcal{D}}[i \in \phi_w(R)]$. For each distribution \mathcal{W} over weight vectors and element $i \in \mathcal{E}$, let $y_i(\mathcal{W}) = \Pr_{R \sim \mathcal{D}}[i \in \phi_{\mathcal{W}}(R)] = \Pr_{R \sim \mathcal{D}, w \sim \mathcal{W}}[i \in \phi_w(R)]$. Let $\mathcal{Y} = \{y(\mathcal{W}) : \mathcal{W} \in \Delta(\mathbb{R}_+^{\mathcal{E}})\} \subseteq [0, 1]^{\mathcal{E}}$ be the family of all inclusion probabilities achievable by some online CRM of the form $\phi_{\mathcal{W}}$. It is immediate that $\mathcal{Y} = \mathbf{convexhull}(\{y(w) : w \in \mathbb{R}_+^{\mathcal{E}}\})$, and hence \mathcal{Y} is a convex subset of $[0, 1]^{\mathcal{E}}$.

An online $\alpha\gamma$ -CRM for \mathcal{D} of the form $\phi_{\mathcal{W}}$ exists if and only if \mathcal{Y} intersects with the upwards closed convex set $\frac{x}{\alpha\gamma} + \mathbb{R}_+^{\mathcal{E}}$. Suppose for a contradiction that this intersection is empty; by the separating hyperplane theorem, this implies that there exists $w \in \mathbb{R}_+^{\mathcal{E}}$ such that $\frac{1}{\alpha\gamma} \sum_i w_i x_i > \sum_i w_i y_i$ for all $y \in \mathcal{Y}$. In particular, $\frac{1}{\alpha\gamma} \sum_i w_i x_i > \sum_i w_i y_i(w)$. Since $\sum_i w_i x_i = \mathbf{E}_{R \sim \mathcal{D}} w(R)$ and $\sum_i w_i y_i(w) = \mathbf{E}_{R \sim \mathcal{D}} [w(\phi_w(R))]$, we get a contradiction with Lemma 4.3. This concludes the proof of the theorem.

5 From Contention Resolution to a Secretary Algorithm?

One might hope that online contention resolution is equivalent to the secretary problem on matroids. In particular, does a competitive universal online CRS imply a competitive secretary algorithm? We make partial progress towards this question. In particular, we reduce the secretary problem to online contention resolution on a particular uncontentious distribution derived from the matroid and sample of its elements: the distribution of “improving elements”, as originally defined by Karger [20] for purposes different from ours.

► **Definition 5.1.** Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with weights $w \in \mathbb{R}_+^{\mathcal{E}}$, and let $p \in (0, 1)$. The random set R of improving elements with parameter p is sampled as follows: Let $S \sim \mathbf{Ind}_p(\mathcal{E})$, and let $R = R(S) = \{i \in \mathcal{E} : \mathbf{rank}_w(S \cup i) > \mathbf{rank}_w(S)\}$. Equivalently, R is the set of elements in $\mathcal{E} \setminus S$ which are not spanned by higher weight elements in S . Another equivalent definition is $R = \{i \in \mathcal{E} \setminus S : i \in \mathbf{OPT}_w(S \cup i)\}$.

The maximum-weight independent subset of the improving elements is $(1 - p)$ -approximately optimal in expectation:

► **Fact 5.2.** Fix a weighted matroid (\mathcal{M}, w) , and let R be the random set of improving elements with parameter p . Each element of $\mathbf{OPT}_w(\mathcal{M})$ is in R with probability $1 - p$. It follows that $\mathbf{E}[w(R)] \geq \mathbf{E}[\mathbf{rank}_w(R)] \geq (1 - p)\mathbf{rank}_w(\mathcal{M})$.

Note that the random set R of improving elements does not follow a product distribution. In fact, elements are (weakly) positively correlated in general. This is illustrated by the special case of the 1-uniform matroid on $[n]$ with weights $w_1 > w_2 > \dots > w_n$, and $p = 1/2$: the distribution of R is as described in Example 3.14. As our main result in this section, we nevertheless show that the random set of improving elements is uncontentious.

► **Theorem 5.3.** Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a matroid with weights $w \in \mathbb{R}_+^{\mathcal{E}}$, and let $p \in (0, 1)$. The random set of improving elements with parameter p is $\frac{1}{p}$ -uncontentious.

Theorem 5.3 and Fact 5.2, taken together, essentially reduce the matroid secretary problem to online contention resolution for the distribution of the random set of improving elements, with one caveat we will discuss shortly. In particular, consider the following blueprint for a secretary algorithm:

1. Let S be the first $\text{Binom}(|\mathcal{E}|, p)$ elements arriving online.
2. Let $R = R(S) \subseteq \mathcal{E} \setminus S$ be a sample of the set of improving elements with parameter p .
3. After observing S , the elements of $\mathcal{E} \setminus S$ arrive online in random order and are presented as such to an online contention resolution algorithm, along with their membership status in R . Note that membership in R can be determined “on the spot” as required for online contention resolution.⁵

Now given a β -competitive α -universal online CRS, we set $p = \frac{1}{\alpha}$ and obtain a $\frac{\beta}{1-p}$ -competitive secretary algorithm. However, the following caveat prevents us from proving a formal theorem of this form: we cannot provide the online CRS with a complete description of the prior distribution. In particular, the distribution \mathcal{D} of improving elements – while fully described by the weighted matroid (\mathcal{M}, w) and the parameter p – can not be fully described to the contention resolution algorithm prior to its invocation, since entries of w are revealed online. As such, we learn both the sample $R \sim \mathcal{D}$ and the distribution \mathcal{D} gradually as elements arrive. An oblivious universal online CRS would resolve this difficulty, but unfortunately we proved in Theorem 3.18 that such a CRS can not exist even for simple matroids and even offline. A reduction from the matroid secretary problem to contention resolution must therefore require a CRS which can make do with only partial knowledge of the prior. We leave exploration of these possibilities for future work, and discuss them further in the Conclusion section.

5.1 Proof of Theorem 5.3

Let p , S , and R be as in Definition 5.1. We prove that R is uncontentious by leveraging (c) from Theorem 3.6. In particular we will show that, for arbitrary $F \subseteq \mathcal{E}$, we have $\mathbf{E}[\text{rank}(R \cap F)] \geq p \mathbf{E}[|R \cap F|]$. We break this up into the following three lemmas.

► **Lemma 5.4.** $\mathbf{E}[\text{rank}(R \cap F)] \geq (1-p) \mathbf{E}[|F \cap \text{OPT}_w(S \cup F)|]$

Proof. Let $T = S \setminus F$, and note that $S \cup F = T \uplus F$. We condition on the random variable T and show conditionally that $\mathbf{E}[\text{rank}(R \cap F)] \geq (1-p) |F \cap \text{OPT}_w(T \uplus F)|$.

Take $i \in F \cap \text{OPT}_w(T \uplus F)$. We will show that i is in R , and hence is in $R \cap F$, with probability $1-p$. Since $i \in S \cup i \subseteq T \uplus F$ and $i \in \text{OPT}_w(T \uplus F)$, it follows from the matroid axioms that $i \in \text{OPT}_w(S \cup i)$. With probability $1-p$ we also have $i \notin S$, in which case $i \in R$ by definition. Since $F \cap \text{OPT}_w(T \uplus F)$ is an independent set, the claim follows. ◀

► **Lemma 5.5.** $|F \cap \text{OPT}_w(S \cup F)| \geq |F \cap \text{OPT}_w(S)|$

Proof. We prove this by induction on a set T with $S \subseteq T \subseteq S \cup F$, initialized to $T = S$ at the base case. Consider how the value of $|F \cap \text{OPT}_w(T)|$ changes as we add elements of $F \setminus S$ to T one by one. When adding an element $i \in F \setminus T$ to T , there are three cases:

- $i \notin \text{OPT}_w(T \cup i)$: In this case, $\text{OPT}_w(T \cup i) = \text{OPT}_w(T)$ and $|F \cap \text{OPT}_w(T \cup i)| = |F \cap \text{OPT}_w(T)|$.
- i is not spanned by T , and $i \in \text{OPT}_w(T \cup i)$: In this case, $\text{OPT}_w(T \cup i) = \text{OPT}_w(T) \cup \{i\}$, and therefore $|F \cap \text{OPT}_w(T \cup i)| = 1 + |F \cap \text{OPT}_w(T)|$.
- i is spanned by T , and $i \in \text{OPT}_w(T \cup i)$: In this case, elementary application of the matroid axioms implies that $\text{OPT}_w(T \cup i) = \text{OPT}_w(T) \cup \{i\} \setminus \{j\}$ for some $j \in T$. Since $i \in F$, it follows that $|F \cap \text{OPT}_w(T \cup i)|$ is either equal to $|F \cap \text{OPT}_w(T)|$ or exceeds it by 1, depending on whether $j \in F$. ◀

⁵ Technically, a CRM requires that elements of \mathcal{E} – rather than merely $\mathcal{E} \setminus S$ – be presented in uniform random order along with their membership status in R . This is easily accomplished by appropriately interleaving the elements of S – none of which are in R – among the elements of $\mathcal{E} \setminus S$.

► **Lemma 5.6.** $\mathbf{E}[\mathbf{OPT}_w(S) \cap F] \geq \frac{p}{1-p} \mathbf{E}[R \cap F]$

Proof. For each $i \in F$, we will show that $\Pr[i \in \mathbf{OPT}_w(S)] \geq \frac{p}{1-p} \Pr[i \in R]$, which suffices.

Take $i \in F$, and let $S_{>i} = \{j \in S : w_j > w_i\}$. Conditioning on $S_{>i}$, there are two cases:

- $i \in \mathbf{span}(S_{>i})$: It follows that $i \notin \mathbf{OPT}_w(S)$ and $i \notin R$, with certainty.
- $i \notin \mathbf{span}(S_{>i})$: With probability p we have $i \in S$ and therefore $i \in \mathbf{OPT}_w(S)$ and $i \notin R$. With the remaining probability $(1-p)$ we have $i \notin S$ and therefore $i \in R$ and $i \notin \mathbf{OPT}_w(S)$.

In both cases, the conditional probability that $i \in \mathbf{OPT}_w(S)$ is at least $\frac{p}{1-p}$ times the conditional probability that $i \in R$. The lemma follows. ◀

5.2 Where Prior Work Fails

There has been speculation in the community that contention resolution for improving element distributions can be accomplished online using the ideas of Feldman et al [17]. If this were true, then a stronger (online) form of our Theorem 5.3 would follow. We show that such conjecture is fatally flawed: there exists no $o(n)$ -competitive online CRS in the worst-case arrival model, even when both the order and the distribution of improving elements are known to the algorithm. In other words, any competitive online CRS for improving element distributions must make and exploit assumptions on the arrival order. This rules out direct application of the arguments and techniques of Feldman et al [17], which – in holding for an (unknown) worst-case arrivals – cannot exploit the uniform arrival order. The same can be said for the work of Lee and Singla [25], which operates in the known worst-case arrival model.

We prove the following theorem, then elaborate on how algorithms from prior work tend to fail on simple examples.

► **Theorem 5.7.** *Let \mathcal{M} be a matroid on n elements. There is no $o(n)$ -competitive online CRS for (known) improving element distributions on \mathcal{M} in the worst-case arrival model. This holds even for the 1-uniform matroid, for every constant parameter p of the distribution of improving elements, and even when the arrival order is known to the algorithm.*

Proof. Let $\{1, \dots, n\}$ denote the ground set of of a 1-uniform matroid, listed in decreasing order of weight. Let R be the random set of improving elements with parameter p . Note that R is supported on sets of the form $\{1, \dots, k\}$ for $k = 0, \dots, n$. In the special case of $p = 1/2$, the distribution of R is as described in Example 3.14. In general, $\Pr[R = \{1, \dots, k\}] = p(1-p)^k$. The random set R is $1/p$ uncontentious, as shown by Theorem 5.3. Concretely, the offline CRM $\phi(\{1, \dots, k\}) = k$ is $\frac{1}{p}$ -competitive.

Now suppose that elements are known to arrive online in the order $1, 2, 3, \dots, n$, and consider an α -competitive online CRM for some $\alpha \geq 1$. Let $T \subseteq R$ be the (random) set of elements selected by the CRM. Conditioned on $i \in R$, the CRM must select i with probability at least $\frac{1}{\alpha}$. Formally, $\Pr[i \in T | i \in R] \geq \frac{1}{\alpha}$.

When element i arrives, the CRM learns whether $i \in R$, and if so must decide whether to select i . Since the online CRM has only observed elements $1, \dots, i$, and must make its decision on the spot, it cannot distinguish between different sets of the form $R = \{1, \dots, k\}$ for $k \geq i$. In other words, it cannot distinguish between the different realizations of R which include i , and must therefore select i with probability at least $\frac{1}{\alpha}$ in every realization of R which includes i . Formally, $\frac{1}{\alpha} \leq \Pr[i \in T | i \in R] = \Pr[i \in T | R = \{1, \dots, k\}]$ for every $k \geq i$.

Since i was chosen arbitrarily, we can take $k = n$ and conclude that $\Pr[i \in T | R = \{1, \dots, n\}] \geq \frac{1}{\alpha}$ for all i . Feasibility requires that $\sum_{i=1}^n \Pr[i \in T | R = \{1, \dots, n\}] \leq 1$. Therefore, $\alpha \geq n$. ◀

It is instructive to examine where the algorithm of Feldman et al [17] fails in the special case of the 1-uniform matroid on n elements, even when improving elements are presented in a uniformly random order. Indeed, we will argue that no “simple tricks” seem to save the day. Recall that the algorithm of [17] defines a sequence of nested flats $\emptyset \subset F_1 \subset F_2 \subset \dots \subset F_k$, and runs the greedy online algorithm on each contracted submatroid F_i/F_{i-1} . The 1-uniform matroid contains only a single non-empty flat, containing all elements. Therefore, the algorithm of [17] reduces merely to the naive greedy online algorithm which simply selects the first element it encounters, which in the case of a uniform arrival order is a uniformly random improving element.

Now, let $[n] = \{1, \dots, n\}$ denote the ground set of the 1-uniform matroid listed in decreasing order of weight, and consider the distribution of improving elements R with parameter $p = 1/2$ as described in Example 3.14. Element k is improving with probability 2^{-k} , yet is selected by the algorithm with probability $\sum_{i=k}^{n-1} 2^{-(i+1)} \cdot \frac{1}{i} + 2^{-n} \cdot \frac{1}{n} < \frac{2^{-k}}{k} = \frac{\Pr[k \in R]}{k}$. Intuitively, when k is improving, so are elements $1, \dots, k-1$, which easily span k and are not distinguished from k by the algorithm. It is easy to show that the algorithm suffers the same fate for any other choice of p .

One might be tempted to employ other tricks, such as for example “canceling” each element in R with independent constant probability in order to reduce contention and place the marginal probability vector deep in the matroid polytope. Such tricks are doomed to fail all the same: the algorithm groups all elements into the same (unique) flat, and in doing so does not distinguish between the “uncanceled” elements of R , so cannot select element k with probability exceeding $\frac{\Pr[k \in R]}{k}$.

It is hopefully now clear that any online CRS for improving element distributions must make and exploit assumptions on the arrival order. Whereas this rules out obvious extensions of Feldman et al [17] and Lee and Singla [25], one might hope that the algorithm of Adamczyk and Włodarczyk [1] might fare better, since they do exploit the random ordering assumption. Sadly, their algorithm also fails for the 1-uniform matroid: it also does not distinguish between different improving elements in this special case, and therefore also selects element k with probability no more than $\frac{\Pr[k \in R]}{k}$. That being said, we are more hopeful that the techniques of [1], if combined with significant new ideas, might yield progress on online contention resolution for positively correlated distributions.

6 Conclusions and Open Problems

In this paper, we begin an exploration of the power and limitations of contention resolution beyond known product distributions, as well as its connections to secretary problems. We hope that our results are a first step towards broader application of the techniques behind contention resolution and online selection. Most notably, our results highlight approaches to resolving the matroid secretary conjecture. We identify several intriguing open questions in pursuit of these agendas.

- Can the result of Theorem 4.1 be shown unconditionally; i.e., can we show a competitive universal online CRS for matroids without assuming the matroid secretary conjecture? We believe this to be a reasonable first step towards proving the matroid secretary conjecture. As we show in Section 5.2, prior work on online contention resolution fails in the presence of even the modest positive correlation exhibited by (uncontentious) improving element distributions on simple matroids. Therefore, we believe significant new ideas are required.

- Recalling the caveat to our results from Section 5, can a tighter connection be made between the secretary problem and contention resolution? Is there a natural model of contention resolution on matroids which permits a reduction both from and to the matroid secretary problem? The knee-jerk approaches using duality-like arguments fail to establish such an equivalence, so new ideas appear to be required.
- The caveat to our results from Section 5 suggests that resolving contention with limited knowledge of the prior is closely related to the matroid secretary conjecture. Recalling our impossibility result of Theorem 3.18, we can start by examining prior-independent contention resolution for interesting classes of distributions. For example, is there a competitive prior-independent (or even oblivious) CRS for ex-ante-feasible product distributions?
- Can Theorem 4.1 be made computationally efficient? Given only oracle access to an arbitrary uncontentious distribution and an arbitrary algorithm for the matroid secretary problem, this is unclear.
- Is there an analogue of our characterization of uncontentious distributions for prophet inequality problems? In particular, can we characterize joint distributions of random variables which permit competitive prophet inequalities with respect to a given matroid?
- Do more general set systems permit a characterization of uncontentious distributions with a finite set of inequalities, a-la Theorem 3.6?

We restricted our attention to matroids in the paper, though some notes are in order on extensions of our results to more general constraints. In the characterization of Theorem 3.6, the equivalence of (a) and (b) holds for a general downwards-closed set systems, and is implicit in the arguments of [10]. The equivalence with (c) exploits the matroid structure, however. Theorem 4.1 also holds for general downwards-closed set systems, and our proof does not invoke the matroid assumption. The results and arguments of Section 5, in particular Theorem 5.3, heavily rely on the matroid structure and do not appear to be easily extensible beyond matroids. We leave further extensions of our results beyond matroids for future work.

References

- 1 Marek Adamczyk and Michał Włodarczyk. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 790–801. IEEE, 2018.
- 2 Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Price of correlations in stochastic optimization. *Operations Research*, 60(1):150–162, 2012.
- 3 Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43(2):930–972, 2014.
- 4 Pablo D Azar, Robert Kleinberg, and S Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1358–1377. Society for Industrial and Applied Mathematics, 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*, pages 303–318. ACM, 2018.
- 6 Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 434–443. Society for Industrial and Applied Mathematics, 2007.
- 7 Grigori Calinescu, Chandra Chekuri, Martin Pál, and Jan Vondrák. Maximizing a monotone submodular function subject to a matroid constraint. *SIAM Journal on Computing*, 40(6):1740–1766, 2011.

- 8 Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320. ACM, 2010.
- 9 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 575–584. IEEE, 2010.
- 10 Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.
- 11 José Correa, Paul Dütting, Felix Fischer, and Kevin Schewior. Prophet inequalities for iid random variables from an unknown distribution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 3–17. ACM, 2019.
- 12 Michael Dinitz. Recent advances on the matroid secretary problem. *ACM SIGACT News*, 44(2):126–142, 2013.
- 13 Evgenii Borisovich Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics*, 4:627–629, 1963.
- 14 Soheil Ehsani, MohammadTaghi 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*, pages 700–714. SIAM, 2018.
- 15 Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.
- 16 Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $o(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 1189–1201. SIAM, 2014.
- 17 Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. Society for Industrial and Applied Mathematics, 2016.
- 18 Mohammad Taghi Hajiaghayi, Robert Kleinberg, and Tuomas Sandholm. Automated online mechanism design and prophet inequalities. In *AAAI*, volume 7, 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 David R Karger. Random sampling and greedy sparsification for matroid optimization problems. *Mathematical Programming*, 82(1-2):41–81, 1998.
- 21 Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136. ACM, 2012.
- 22 Ulrich Krengel and Louis Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745–747, 1977.
- 23 Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.
- 24 Oded Lachish. $O(\log \log \text{rank})$ competitive ratio for the matroid secretary problem. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 326–335. IEEE, 2014.
- 25 Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2018.
- 26 J. G. Oxley. *Matroid Theory*. Oxford University Press, 1992.
- 27 Yosef Rinott, Ester Samuel-Cahn, et al. Comparisons of optimal stopping values and prophet inequalities for negatively dependent random variables. *The Annals of Statistics*, 15(4):1482–1490, 1987.

42:18 The Outer Limits of Contention Resolution

- 28 Aviad Rubinfeld. Beyond matroids: Secretary problem and prophet inequality with general constraints. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 324–332. ACM, 2016.
- 29 Ester Samuel-Cahn. Prophet inequalities for bounded negatively dependent random variables. *Statistics & probability letters*, 12(3):213–216, 1991.
- 30 Jack Wang. The prophet inequality can be solved optimally with a single set of samples. *arXiv preprint*, 2018. [arXiv:1812.10563](https://arxiv.org/abs/1812.10563).
- 31 Dominic JA Welsh. *Matroid theory*. Courier Corporation, 2010.
- 32 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.