From Contention Resolution to Matroid Secretary and Back*

Shaddin Dughmi†
Department of Computer Science
University of Southern California
shaddin@usc.edu

December 9, 2024

Abstract

We show that the matroid secretary problem is equivalent to correlated contention resolution in the online random-order model. Specifically, the matroid secretary conjecture is true if and only if every matroid admits an online random-order contention resolution scheme which, given an arbitrary (possibly correlated) prior distribution over subsets of the ground set, matches the balance ratio of the best offline scheme for that distribution up to a constant.

Integral to our result is a polyhedral characterization of the (correlated) distributions permitting offline contention resolution over a given matroid — we refer to these distributions as uncontentious. Our characterization can be viewed as a distributional generalization of the matroid covering theorem, and isolates the kind and degree of positive correlation that is benign for offline contention resolution. Using this characterization, we are able to show that the set of *improving elements* for a subsample of a weighted matroid is uncontentious — a fact that serves as a key technical component of our result.

One direction of our equivalence is relatively straightforward: a competitive secretary algorithm yields a random-order contention resolution scheme — one which approximately matches the best possible offline balance ratio — by providing an approximate solution to its dual. The other direction is more technical, and involves a composition of three reductions which each isolates a technical hurdle: from the secretary problem to the (correlated) prophet secretary problem, then from that to a labeled generalization of (random-order) contention resolution, and finally from labeled contention resolution to its unlabeled counterpart. The uncontentiousness of the set of improving elements implies that the resulting contention resolution problem features an (offline) uncontentious distribution, which therefore implies our main result.

One interpretation of our result is that handling the positive correlation inherent to uncontentious distributions is the key technical barrier to resolving the matroid secretary conjecture.

1 Introduction

This paper follows in the hallowed tradition in theoretical computer science of reducing the number of questions without providing any answers. We establish an equivalence between one of the central open problems in online algorithm design, the *matroid secretary conjecture*, and the increasingly rich and fruitful framework of *contention resolution*. Specifically, we show that the matroid secretary problem admits a constant-competitive algorithm if and only if matroid contention resolution for general (correlated) distributions is approximately as powerful (up to a constant) in the online

*This journal paper unifies and expands upon results that span two conference publications: [13] and [14].
†This work was supported by NSF CAREER Award CCF-1350900 and NSF Grant CCF-2009060.

random-order model as it is in the offline model. Our result paves the way for application of the many recent advances in contention resolution, and in stochastic decision-making problems more generally, to resolving the conjecture.

The classical (single-choice) secretary problem [19], and its many subsequent combinatorial generalizations, capture the essence of online decision making when adversarial datapoints arrive in a non-adversarial order. The paradigmatic such generalization is the *matroid secretary problem*, originally proposed by Babaioff et al. [4]. Here, elements of a known matroid arrive online in a uniformly random order, each equipped with a nonnegative weight chosen at the outset by an adversary. An algorithm for this problem must decide online whether to accept or reject each element, knowing only the weights of the elements which have arrived thus far, subject to accepting an independent set of the matroid. The goal is to maximize the total weight of accepted elements. The *matroid secretary conjecture* of [4] postulates the existence of an (online) algorithm for this problem which is constant competitive, as compared to the offline optimal, for all matroids. Though much prior work has designed competitive algorithms for specific classes of matroids, the general conjecture has remained open.

Recent years have seen an explosion of interest in a variety of online decision-making problems of a similar flavor, albeit distinguished from secretary problems in that the uncertainty in the data is stochastic, with known distribution, rather than adversarial. Such models include variants and generalizations of the classical *prophet inequality*, adaptive stochastic optimization models such as *stochastic probing*, and what is increasingly emerging as the technical core of such problems: *contention resolution*.

Contention Resolution, as originally formalized in Chekuri et al. [11], 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 function[1] — element by element. The offline model of contention resolution was introduced by Chekuri et al. [11], motivated by applications to approximation algorithm design. It has since been extended to various online settings (e.g. [24, 1]), and moreover has emerged as the basic technical building block of a number of important results for stochastic decision-making (see e.g. [24, 1, 36, 5, 6, 15]).

In contention resolution, elements of a set system — for our purposes, a matroid — are each equipped with a single-bit stochastic datapoint indicating whether that element is *active* or *inactive*. The joint distribution of these datapoints, henceforth referred to as the *prior distribution*, is assumed to be ex-ante feasible (perhaps approximately) and given upfront. An algorithm for this problem — which we often refer to as a *contention resolution scheme (CRS)* — is tasked with accepting an independent set of active elements with the goal of maximizing the *balance ratio*: the minimum, over all elements, of the ratio of the probability the element is accepted to the probability the element is active. In the original *offline* setting of contention resolution, the algorithm observes all datapoints before choosing which elements to accept. In the various *online* settings, elements and their datapoints are presented sequentially to an algorithm, which must then irrevocably decide whether or not to accept each element subject to feasibility. The most pertinent online model to this paper is that in which elements arrive in a uniformly *random order*, in which case we refer to the algorithm as a *random-order contention resolution scheme (RO-CRS)*.

[1] In its most general form, a CRS approximately preserves all linear objective functions simultaneously, whereas a *monotone* CRS approximately preserves all submodular objectives [11].

Most work on contention resolution prior to ours has restricted attention to product prior distributions: elements are active independently, with given probabilities. Sweeping positive results hold for product priors, and those results tend to extend to negative correlation between elements. In the absence of positive correlation, approximate ex-ante feasibility is both necessary and sufficient for balanced contention resolution, and this holds for matroids and most other "simple" set systems (see e.g. [11, 24, 36]). In contrast, it is easy to see that not much is possible in the presence of unrestrained positive correlation, even offline and for the simplest of set systems. Understanding the power of contention resolution more thoroughly, as a function of the kind and degree of positive correlation, the underlying set system, and the arrival model, forms the starting point of our investigations. As this paper shows, such understanding will turn out to be intimately related to the matroid secretary conjecture.

Results and Technical Approach

Characterizing Contention Resolution. Our first set of results provides a structural understanding of the class of *α-uncontentious distributions* for a given matroid: those distributions permitting α-balanced offline contention resolution. We provide a polyhedral characterization of α-uncontentious distributions consisting of an exponential-sized system of linear inequalities, one for each subset of the ground set. We point out that this characterization can be viewed as a natural distributional generalization of the *matroid covering theorem* of Edmonds [20] (see also [44]). We then leverage our characterization to 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.

Improving Elements are Uncontentious. Given a weighted matroid and a random sample consisting of a constant fraction of its elements, we call an unsampled element *improving* if it increases the omniscient optimal value for the sample. This (random) set of *improving elements*, originally defined by Karger [30], plays an important role in our main result, enabled by the following key technical lemma. We show that the set of improving elements is $O(1)$-uncontentious by showing that it satisfies all inequalities in our polyhedral characterization. This is particularly noteworthy in light of the fact that the improving elements may exhibit significant positive correlation in general. Our proof involves induction, as well some delicate charging arguments between different probability events that are reminiscent of — and inspired by — the analysis of martingales. In defense of our somewhat technical proof, we also show that the contention resolution schemes from prior work do not provide an alternative route to our lemma.

Main Result. We then proceed to proving our main result: the equivalence, up to a constant factor, between the matroid secretary problem and contention resolution. Given a matroid, we say an (online) random-order contention resolution scheme is *c-universal* if, for every prior distribution, it c-approximates the balance ratio of the best offline contention resolution scheme. In other words, a c-universal scheme must be $c\alpha$-balanced when the input is drawn from an α-uncontentious distribution. Our main result is that the matroid secretary conjecture is equivalent to the existence of a c-universal scheme for all matroids, with c a universal constant that is independent of the matroid. Since negative correlation is typically benign for contention resolution, one interpretation of this result is that handling the positive correlation inherent to uncontentious distributions is the key technical barrier to resolving the matroid secretary conjecture.

One direction of this equivalence is relatively straightforward: A c-competitive secretary algorithm solves what is essentially a dual of contention resolution, yielding a $(c-\epsilon)$-universal RO-CRS for arbitrarily small positive ϵ.

The other direction, a reduction from the secretary problem to universal contention resolution, ends up being much more technically involved: a c-universal RO-CRS yields a $\frac{c-\epsilon}{4096}$-competitive secretary algorithm for an arbitrarily small positive ϵ. However, it emanates from the following simple — though ultimately unsuccessful — approach: (a) Sample a constant fraction of the elements as they arrive online and observe their weights, then (b) resolve contention for the elements that improve the sample. The improving elements hold a constant fraction of the omniscient optimal value, and can be recognized online. Moreover, as shown in aforementioned technical lemma, the improving elements are uncontentious. Therefore, this appears at first glance to be a reduction from the secretary problem to contention resolution for a particular uncontentious distribution, and therefore to universal contention resolution. However, this approach fails for the following subtle reason: the prior distribution of improving elements, being defined by the matroid and all of its (a-priori unknown) weights, is unknown to the contention resolution algorithm. In other words, both the set of improving elements and its distribution are revealed gradually online, whereas online contention resolution requires that the distribution be given upfront.

This failure turns out to be the first of a series of three technical obstacles, which we isolate by expressing our reduction from the secretary problem to contention resolution as the composition of three component reductions. This takes us through two "bridge problems" along the way. The first of these bridge problems is the correlated version of the familiar *matroid prophet secretary* problem of Ehsani et al. [21], which relaxes the matroid secretary problem by assuming that weights are drawn from a known distribution rather than adversarially.[2] The second bridge problem is a generalization of contention resolution — in particular on matroids, in the online random-order model — which we define and term *labeled contention resolution*. Here, each active element comes with a stochastic label, and balance is evaluated with respect to (element,label) pairs rather than merely with respect to elements. Figure 1 summarizes the cycle of reductions between all four problems, which we conclude are all equivalent up to constant factors in their competitive and balance ratios.

Our first component reduction, motivated by the distribution of improving elements being unknown, is from the secretary problem to the prophet secretary problem on matroids. Fairly standard duality arguments allow us to replace the adversarial weight vector in the secretary problem with a stochastic one of known distribution. Modulo some simple normalization and discretization of the weights, at the expense of a constant in the competitive ratio, this yields an instance of the prophet secretary problem. With a stochastic weight vector drawn from a known distribution, we now face a *known* mixture of improving element distributions. Moreover, since our characterization of uncontentious distributions implies their closure under mixing, it would appear that we have now overcome the first obstacle.

Unfortunately, shifting to a stochastic weight vector introduces a second obstacle. With the set of improving elements now correlated with the vector of element weights, balanced contention resolution no longer guarantees extracting a constant fraction of the expected weight of the set. This is because a contention resolution scheme may preferentially accept an improving element when it has low weight, and reject it when it has high weight, while still satisfying the balance requirement in the aggregate. In fact, we show by way of a simple example that egregious instantiations of this phenomenon are not difficult to come by. This motivates our reduction from the matroid prophet secretary problem to *labeled* contention resolution, also in the online random-order model. By labeling each element with its weight, and requiring balance with respect to (element,label)

[2] An alternate, equivalent, description of the prophet secretary problem is as a relaxation of the prophet inequality problem to random order arrivals.

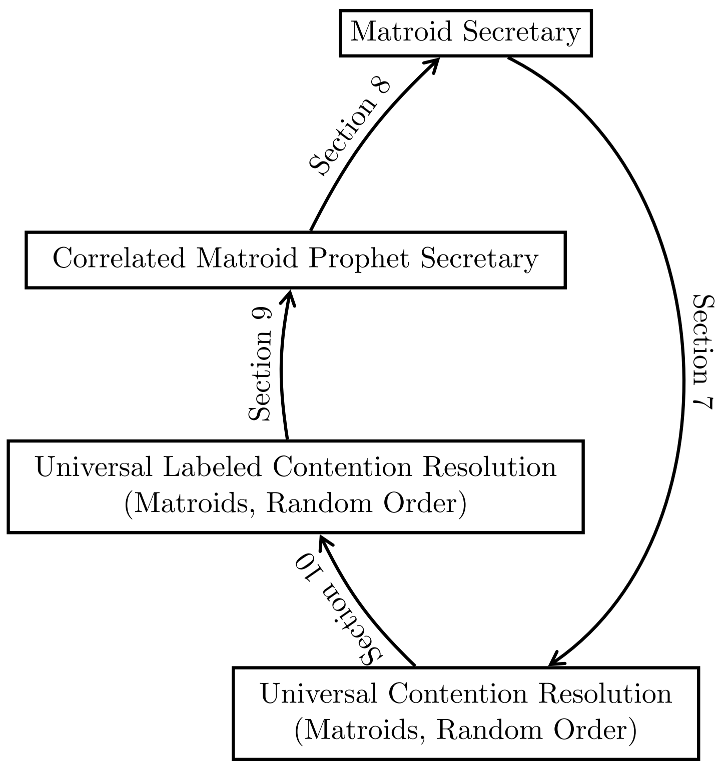

Figure 1: Reductions between the four relevant problems. An arrow $A \to B$ indicates using an algorithm for problem A to solve problem B, i.e., a reduction from B to A. All reductions preserve the competitive or balance ratio up to a constant.

pairs, we exclude contention resolution policies which favor low-weight elements. This overcomes our second technical obstacle.

Our final, and most technically involved, component reduction is from labeled to unlabeled contention resolution, for matroids in the random-order model. Such a reduction would be trivial in the offline setting: by thinking of each (element,label) pair as a distinct parallel copy of the element, we obtain an equivalent instance of unlabeled contention resolution, albeit on a larger matroid. An online version of this reduction is more difficulty to come by, however: in the labeled version of the problem only the active (labeled) copy of an element arrives online, whereas all copies arrive in a uniformly random order in the corresponding unlabeled version. Any reduction must somehow "hallucinate" some of the inactive copies online without knowing the future, and therefore runs the risk of "collision" with a future active copy of the same element. This forms our third and final technical obstacle. We overcome this obstacle by "blowing up" the matroid even further, creating a large number of identical duplicates of each label. As the number of duplicates grows large, a random order of (element,label) pairs converges in distribution to a deterministic order (modulo the equivalence relation between duplicates). The required "hallucinations" are now approximately independent — in a statistical distance sense — of the future identities of active elements and their labels.

Additional Discussion of Related Work

Contention resolution in the offline setting was formalized by Chekuri et al. [11], motivated by applications to approximation algorithm design via randomized rounding. For product priors and a given packing set system, [11] shows that the optimal offline balance ratio equals the worst-case *correlation gap*, as defined by Agrawal et al. [2], of the set system's weighted rank function. Starting with the work of Feldman et al. [24], contention resolution was extended to online settings and applied to a variety of problems in mechanism design and adaptive stochastic optimization (see also [1, 36]). Regardless of the set system, balanced contention resolution is obviously only possible for priors that are (approximately) *ex-ante feasible*: the random set is feasible on average, in the sense that the per-element marginal probabilities lie in the polytope associated with the set system. One message of the aforementioned prior work is that — for product priors and many natural set systems such as matroids, knapsacks, and their intersections — approximate ex-ante feasibility is also sufficient for balanced contention resolution, whether offline or online.

The (single-choice) secretary problem is due to Dynkin [19]. It was subsequently generalized to a uniform matroid constraint by Kleinberg [32], and to a general matroid constraint by Babaioff et al. [4]. A long line of work has designed constant-competitive algorithms for special cases of the matroid secretary problem, and we refer the reader to the semi-recent survey by Dinitz [12]. The current state-of-the art for general matroids is an $O(\log \log \mathbf{rank})$-competitive algorithm due to Lachish [35], which was since simplified by Feldman et al. [23]. Beyond matroids, the secretary problem with general packing constraints was recently studied by Rubinstein [41].

Closely related to the secretary problem are the prophet inequality problem and the prophet secretary problem, which analogously admit combinatorial generalizations to matroids and other packing set systems. Whereas a secretary problem features adversarial data (i.e., element weights) arriving online in a random order, a prophet inequality problem features stochastic data (typically assumed to be independent) arriving online in an adversarial order. A prophet secretary problem is a relaxation of both, featuring stochastic data arriving online in a random order. The original (single-choice) prophet inequality is due to Krengel, Sucheston, and Garling [33, 34], and was generalized to matroids by Kleinberg and Weinberg [31]. Generalizations beyond matroids have also received much study; see for example [18, 17, 41]. The (single-choice) prophet secretary problem was introduced by Esfandiari et al. [22], and further studied in [3]. Generalizations to combinatorial constraints, including matroids, were studied by Ehsani et al. [21].

One take-away from our work is that stochastic decision making in the presence of correlation, and in particular positive correlation, is deserving of more attention. Most prior work on aforementioned stochastic decision-making models restricts attention to product priors, with a few exceptions which we now mention. For contention resolution, we are not aware of any work prior to ours which explicitly studies positive correlation. The classical (single-choice) prophet inequality was extended to negatively correlated variables by Rinott and Samuel Cahn [40, 42], whereas no nontrivial prophet inequality holds in the presence of unrestrained positive correlation [26]. To our knowledge, prior to the first conference publication of this work the only nontrivial prophet inequalities in the presence of positive correlation were from the work of Immorlica et al. [29], who posed a particular linear model of correlated distributions. Since then, several works have examined other models which allow for some positive correlation, notably pairwise (and more generally k-wise) independence [8, 16] as well as probabilistic graphical models [7, 37].

Paper Outline

Section 2 presents requisite technical preliminaries. Section 3 develops a structural understanding of offline contention resolution in the presence of correlation, which will serve as a foundation for our results. Section 4 establishes a technical lemma that is key to our main result: that the random set of improving elements for a matroid permits offline contention resolution. Section 5 presents a generalization of contention resolution to include labels on the elements, which will serve as an useful waypoint in our reduction from the matroid secretary problem to contention resolution. Section 6 formally states our main theorem, namely the equivalence between the matroid secretary problem and correlated contention resolution in the random order model, and outlines the high-level structure of its proof. The remaining technical sections present said proof: Section 7 reduces contention resolution to the secretary problem for matroids and more generally independence systems, while Sections 8, 9, and 10 present our three-step reduction from the matroid secretary problem to contention resolution. We end with concluding thoughts and a discussion of open questions in Section 11.

2 Preliminaries

2.1 Miscellaneous Notation and Terminology

We denote the natural numbers by $\mathbb{N}$, the real numbers by $\mathbb{R}$, and the nonnegative real numbers by $\mathbb{R}_+$. We also use $[n]$ as shorthand for the set of integers $\{1, \ldots, n\}$.

For a set $\mathcal{E}$, we use $\Delta(\mathcal{E})$ to denote the family of distributions over $\mathcal{E}$, use $2^{\mathcal{E}}$ to denote the family of subsets of $\mathcal{E}$, and use $\mathcal{E}^*$ to denote finite strings with alphabet $\mathcal{E}$. When $\mathcal{E}$ is finite, we use $e \sim \mathcal{E}$ to denote uniformly sampling e from $\mathcal{E}$. We also use $\mathcal{E}!$ to denote the family of permutations of $\mathcal{E}$, where we think of $\pi \in \mathcal{E}!$ as a bijection from positions $\{1, \ldots, |\mathcal{E}|\}$ to $\mathcal{E}$. When $\mathcal{E}$ is equipped with weights $w \in \mathbb{R}^{\mathcal{E}}$, and $A \subseteq \mathcal{E}$, we use the shorthand $w(A) = \sum_{i \in A} w_i$.

Let ρ be a distribution supported on $2^{\mathcal{E}}$ for some set $\mathcal{E}$. We say ρ is a *product distribution* if, for $A \sim \rho$, the events $\{i \in A\}_{i \in \mathcal{E}}$ are jointly independent. Otherwise, we say ρ is *correlated*. We refer to the vector $x = x(\rho) \in [0,1]^{\mathcal{E}}$ of *marginals* of ρ, where $x_i = \mathbf{Pr}_{A \sim \rho}[i \in A]$ is the *marginal probability* of i in ρ. There is a unique product distribution with any given set of marginals $x \in [0,1]^{\mathcal{E}}$, but many different correlated distributions with those marginals.

2.2 Set Systems and Matroids

We briefly review some of the basics of set systems, and in particular matroid theory. For more background, we refer the reader to Oxley [38] and Welsh [44].

A *set system* $\mathcal{M}$ consists of a *ground set* $\mathcal{E}$ of *elements*, and a family $\mathcal{I} \subseteq 2^{\mathcal{E}}$ of *feasible sets*. A *weighted* set system is a pair $(\mathcal{M}, w)$ where $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ is a set system and $w \in \mathbb{R}^{\mathcal{E}}$ are *weights* associated with its elements. We say a set system $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ is *non-empty* if $\mathcal{I} \neq \emptyset$, and we say it is *downwards-closed* if whenever $T \in \mathcal{I}$ and $S \subseteq T$ we have $S \in I$ as well. Set systems which are non-empty and downwards-closed are often referred to as *independence systems* or *packing constraints*.

A *matroid* $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ is a set system which is non-empty, downwards-closed, and satisfies the so-called *exchange property*: If $S, T \in \mathcal{I}$ and $|S| < |T|$, then there is $i \in T \setminus S$ such that $S \cup \{i\} \in \mathcal{I}$. The feasible sets of a matroid, and more generally an independence system, are traditionally called its *independent sets*. A set of elements which is not independent is called *dependent*, and a minimal dependent set is called a *circuit*. In this paper we are primarily interested in matroids, though we sometimes point out some of our results that hold more generally, such as for independence systems or arbitrary set systems.

Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a set system. The *restriction* of $\mathcal{M}$ to some set of elements $A \subseteq \mathcal{E}$, denoted $\mathcal{M}|A$, is the set system with elements A and feasible sets $\mathcal{I} \cap 2^A$. If $\mathcal{M}$ is an independence system then so is $M|A$, and if $\mathcal{M}$ is a matroid then so is $\mathcal{M}|A$. The *rank* of $\mathcal{M}$, which we denote by $\mathbf{rank}(\mathcal{M})$, is the maximum cardinality of a feasible set of $\mathcal{M}$. When $\mathcal{M}$ is weighted with w, we also use $\mathbf{rank}_w(\mathcal{M})$ to denote its *weighted rank* — i.e. the maximum total weight of a feasible set. Overloading notation, we use $\mathbf{rank}^{\mathcal{M}}(A)$ to denote the rank of $\mathcal{M}|A$, and $\mathbf{rank}^{\mathcal{M}}_w(A)$ to denote the weighted rank of $\mathcal{M}|A$ with weights $\{w_e\}_{e \in A}$, and we refer to these as the *rank function* and *weighted rank function* respectively. The rank and weighted rank functions of a set system $\mathcal{M}$ are nondecreasing set functions. If $\mathcal{M}$ is a matroid then they are also submodular.

When $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ is a matroid, we further use the following standard notions. The *contraction* of $\mathcal{M}$ by $A \subseteq \mathcal{E}$, denoted by $\mathcal{M}/A$, is the matroid $(\mathcal{E} \setminus A, \mathcal{I}')$ where $S \in \mathcal{I}'$ if $A \cup S \in \mathcal{I}$. The *truncation* of $\mathcal{M}$ to rank $k \in \mathbb{N}$ is the matroid $(\mathcal{E}, \mathcal{I}')$ where $S \in \mathcal{I}'$ if $S \in \mathcal{I}$ and $|S| \leq k$. The *span* of a set $S \subseteq \mathcal{E}$ is defined by $\mathbf{span}(S) = \{i \in \mathcal{E} : \mathbf{rank}_{\mathcal{M}}(S \cup i) = \mathbf{rank}_{\mathcal{M}}(S)\}$. If $i \in \mathbf{span}(S)$ we say i is *spanned* by S. A *flat* is a set S that is closed under spanning — i.e., with $\mathbf{span}(S) = S$.

We restrict attention without loss of generality to matroids with no loops (i.e., each singleton is independent). In parts of this paper, we also restrict attention to weighted matroids where all non-zero weights are distinct. This assumption is made merely to simplify some of our proofs, and — using standard tie-breaking arguments — can be shown to hold without loss of generality in as much as our results are concerned. Under this assumption, we define $\mathbf{OPT}^{\mathcal{M}}_w(A)$ 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. We also use $\mathbf{OPT}_w(\mathcal{M}) = \mathbf{OPT}^{\mathcal{M}}_w(\mathcal{E})$ as shorthand for the maximum-weight independent set of $\mathcal{M}$ of minimum cardinality.

For a set system $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, we define its polytope $\mathcal{P}(\mathcal{M})$ to be the convex hull of indicator vectors of feasible sets. When $\mathcal{M}$ is an independence system, $\mathcal{P}(\mathcal{M})$ is a non-empty *packing polytope*: whenever $y \in \mathcal{P}(\mathcal{M})$ and $0 \preceq x \preceq y$ (element-wise), we have $x \in \mathcal{P}(\mathcal{M})$ as well. When $\mathcal{M}$ is a matroid, $\mathcal{P}(\mathcal{M})$ is called a *matroid polytope*.

2.3 The Matroid Secretary Problem

In the *matroid secretary problem* there is matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with nonnegative weights $\{w_e\}_{e \in \mathcal{E}}$ on the elements. Elements $\mathcal{E}$ arrive online in a uniformly random order $\pi \sim \mathcal{E}!$, and an online algorithm must irrevocably accept or reject an element when it arrives, subject to accepting an independent set of $\mathcal{M}$. Only the matroid $\mathcal{M}$ is given to the algorithm at the outset — say, as an independence oracle. The weights w, on the other hand, are chosen adversarially, and without knowledge of the random order π. The elements then arrive online, along with their weights, in the random order π.

The goal of the online algorithm is to maximize the expected weight of the accepted set of elements. Given $c \in [0, 1]$, we say that an algorithm for the secretary problem is *c-competitive* for a class of matroids, in the worst-case, 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 order π and any internal randomness of the algorithm) is at least a c fraction of the offline optimal — i.e., at least $c \cdot \mathbf{rank}_w(\mathcal{M})$.

The matroid secretary problem was defined by Babaioff et al. [4], generalizing the classical (single-choice) secretary problem of Dynkin [19] which can be viewed as the special case where $\mathcal{M}$ is the rank one matroid. The *matroid secretary conjecture*, posed by Babaioff et al. [4], can be stated as follows.

Conjecture 1 (Matroid Secretary Conjecture [4])**.** *There exists an absolute constant $c > 0$ such that the matroid secretary problem admits an (online) algorithm which is c-competitive for all matroids.*

We note that we are considering the *known matroid* model of the secretary problem, which is the original model defined by Babaioff et al. [4]. A potentially more challenging variant, where only the size of the ground set is known at the outset, but the structure of the matroid is revealed online, has also been considered (see e.g. [25]). We are unaware of any evidence of a separation between the two models, and in fact most algorithms in the matroid secretary literature work for both models. Nonetheless, the known setting lends itself best to our reduction.

Whereas we are primarily interested in matroid constraints, one direction of our equivalence holds for more general set systems. For a class of set systems of interest, this is referred to as a *generalized secretary problem.* When a particular set system $\mathcal{M}$ is given, we refer to the *secretary problem on $\mathcal{M}$.*

2.4 The Matroid Prophet Secretary Problem

The matroid prophet secretary problem relaxes the matroid secretary problem by assuming that the weights $w \in \mathbb{R}_+^{\mathcal{E}}$ are drawn from a known *prior distribution* μ, independent of the random order π, rather than being chosen adversarially. Both $\mathcal{M}$ and μ are given at the outset, whereas the random order π and the realized weight vector w are revealed online as elements arrive. The single-choice prophet secretary problem was introduced by Esfandiari et al. [22], and later studied for matroids and other set systems by Ehsani et al. [21]. To our knowledge, all prior work on the prophet secretary problem has considered independent weights — i.e., μ is a product distribution. We make no such assumption here, allowing the weights to be correlated arbitrarily.

Given $c \in [0, 1]$, we say that an algorithm for the secretary problem is *c-competitive* for a class of matroids and prior distributions if for every matroid $\mathcal{M}$ and distribution μ in that class, the expected weight of the accepted set (over the random order π, the weight vector $w \sim \mu$, and any internal randomness of the algorithm) is at least a c fraction of the expected offline optimal — i.e., at least $c \cdot \mathbf{E}[\mathbf{rank}_w(\mathcal{M})]$.

The matroid prophet secretary problem also relaxes the *matroid prophet inequality* problem of Kleinberg and Weinberg [31], in particular by assuming that the arrival order is uniformly random rather than adversarial. It follows that the competitive ratio of $\frac{1}{2}$ for the matroid prophet inequality from [31] generalizes to the matroid prophet secretary problem when weights are independent. This was improved to $1 - \frac{1}{e}$ by [21]. No constant is known for the matroid prophet secretary problem with general correlated priors, though one would immediately follow from the matroid secretary conjecture. In fact, along the way to our results we show that the existence of a constant competitive algorithm for the matroid prophet secretary problem, with arbitrary matroids and arbitrary correlated priors, is equivalent to the matroid secretary conjecture.

3 Understanding Contention Resolution

3.1 Background and Definitions

Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a set system. We define a *contention resolution map (CRM)* ϕ for $\mathcal{M}$ as a randomized function from $2^{\mathcal{E}}$ to $\mathcal{I}$ with the property that $\phi(R) \subseteq R$ for all $R \subseteq \mathcal{E}$. We refer to the input R to a CRM ϕ as the set of *active* elements. If an active element $i \in R$ is in the output $\phi(R)$, we say it has been *accepted* (or *selected*), otherwise we say it has been *rejected*. Given a *prior distribution* ρ supported on $2^{\mathcal{E}}$, we say ϕ is *α-balanced* for ρ if, for $R \sim \rho$, we have

$$\mathbf{Pr}[i \in \phi(R)] \geq \alpha \, \mathbf{Pr}[i \in R]$$

for all $i \in \mathcal{E}$. In other words, an α-balanced CRM selects each element with probability at least α conditioned on it being active, in expectation over the prior.

Every CRM can be implemented by some algorithm in the *offline model*, where the set R is provided to the algorithm at the outset; when we emphasize this we sometimes say it is an *offline CRM*. A *random-order contention resolution map (RO-CRM)* is a CRM ϕ which can be implemented as an algorithm in the *online random-order model*. Here, $\mathcal{E}$ is presented to the algorithm in a uniformly random order $\pi = (e_1, \ldots, e_n)$, and at the ith step the algorithm learns whether e_i is active — i.e., whether $e_i \in R$ — and if so must make an irrevocable decision on whether to accept or reject e_i — i.e., whether or not to include it in $\phi(R)$. In this random order model, we can conveniently think of the randomness in the arrival order π as endogenous to the randomized function ϕ, and the balance ratio of ϕ is evaluated in expectation over π. Whereas this paper is primarily concerned with the offline and random-order models, we will briefly encounter other online models in which the arrival order π is determined either entirely or partially by an adversary; in such settings, we think of the CRM ϕ^π as being parametrized by the arrival order π, and any balance guarantee must hold for all choices of the adversary.

A *contention resolution scheme (CRS)* Φ for a set system $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and class of distributions $\Delta \subseteq \Delta(2^{\mathcal{E}})$ is an algorithm which takes as input a description of a prior distribution $\rho \in \Delta$ and implements a CRM ϕ_ρ for $\mathcal{M}$ that is catered to the prior distribution ρ. Every CRS can be implemented offline, and we call it an *offline CRS* when we wish to emphasize this. If each ϕ_ρ is an RO-CRM, we say that Φ is a *random order CRS (RO-CRS)*. More generally, if each ϕ_ρ is an online CRM in some adversarial or semi-adversarial arrival model, then Φ is an online CRS in that same arrival model. Regardless of the input model, if ϕ_ρ is an α-balanced CRM for each prior $\rho \in \Delta$, we say that the Φ is an *α-balanced CRS* for the class of distributions Δ.

A prior distribution $\rho \in \Delta(2^{\mathcal{E}})$ is *ex-ante feasible* for set system $\mathcal{M} = (\mathcal{E}, I)$ if its marginals $x = x(\rho)$ are in the polytope $\mathcal{P}(\mathcal{M})$ associated with the set system. Previous work on contention resolution has often restricted attention to the class of ex-ante feasible product distributions —

i.e., $x \in \mathcal{P}(\mathcal{M})$ and each element $i \in \mathcal{E}$ is active independently with probability x_i. Under this restriction, positive results have followed for most "tractable" set systems, whether offline or online (see e.g. [11, 24, 36, 1]). In particular, for matroids there is a $(1 - 1/e)$-balanced CRS both offline [11] and in the random order model [36].

3.2 Uncontentious Distributions and their Characterization

As shorthand, we refer to distributions which permit balanced (offline) contention resolution as *uncontentious*.

Definition 3.1. *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a set system, and let $\rho \in \Delta(2^{\mathcal{E}})$ be a prior distribution. For $\alpha \in [0, 1]$, we say ρ is* α-uncontentious *with respect to $\mathcal{M}$ if there is a CRM for $\mathcal{M}$ which is α-balanced for ρ.*

When $R \sim \rho$ and ρ is α-uncontentious, we often abuse terminology and also say that the random set R is α-uncontentious. We omit reference to $\mathcal{M}$ in these definitions when the set system is clear from context.

We prove the following characterization of uncontentious distributions.

Theorem 3.2. *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a set system, and let $\rho \in \Delta(2^{\mathcal{E}})$. The following are equivalent for every $\alpha \in [0, 1]$:*

(a) ρ is α-uncontentious.

(b) $\mathop{\mathbf{E}}_{R\sim\rho}[\mathbf{rank}_w(R)] \geq \alpha \mathop{\mathbf{E}}_{R\sim\rho}[w(R)]$ for every weight vector $w \in \mathbb{R}_+^{\mathcal{E}}$.

Moreover, if $\mathcal{M}$ is a matroid then both are also equivalent to:

(c) $\mathop{\mathbf{E}}_{R\sim\rho}[\mathbf{rank}(R \cap F)] \geq \alpha \mathop{\mathbf{E}}_{R\sim\rho}[|R \cap F|]$ for every $F \subseteq \mathcal{E}$.

Proof. Property (a) implies property (b) by applying an α-balanced CRM ϕ to $R \sim \rho$, using the fact that $\phi(R)$ is a feasible subset of R.

$$\begin{aligned}
\mathbf{E}[\mathbf{rank}_w(R)] &\geq \mathbf{E}[w(\phi(R))] \\
&= \sum_{i \in \mathcal{E}} w_i \cdot \mathbf{Pr}[i \in \phi(R)] \\
&\geq \alpha \sum_{i \in \mathcal{E}} w_i \cdot \mathbf{Pr}[i \in R] \\
&= \alpha \, \mathbf{E}[w(R)].
\end{aligned}$$

Property (b) implies property (a) by a duality argument identical to that presented in [11]. We present a self-contained proof here. Let $x = x(\rho) \in [0, 1]^{\mathcal{E}}$ denote the marginals of ρ. The distribution ρ 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 α.

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

The dual is the following LP with variables μ and w_i for each element $i \in \mathcal{E}$.

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

At optimality, the binding constraint on μ corresponds to the CRM ϕ which maps each set R to its maximum weight feasible subset according to weights w. It follows that the optimal value of the dual, and hence the primal, equals $\mathbf{E}[\mathbf{rank}_w(R)]$. Property (b) implies that this is at least $\alpha \, \mathbf{E}[w(R)]$. Since $\mathbf{E}[w(R)] = \sum_i w_i x_i = 1$ by linearity of expectations, the optimal value is at least α as needed.

Property (c) is the special case of (b) for binary weight vectors. When $\mathcal{M}$ is a matroid, (c) also implies (b) by the following standard summation argument. Sort and number the elements $\mathcal{E} = (e_1, \ldots, e_n)$ in decreasing order of weights $w_1 \geq w_2 \geq \ldots \geq w_n \geq 0$, where w_i denotes the weight of e_i. Denote $\mathcal{E}_i = \{e_1, \ldots, e_i\}$, 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:

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

□

We note that the equivalence between (a) and (b) for general (correlated) distributions is implicit in the arguments of [11]. Property (c) is reminiscent of the classical *matroid covering theorem*, which can be stated as follows.

Theorem 3.3 (Edmonds [20]). *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a matroid. A set $T \subseteq \mathcal{E}$ can be* covered by *(i.e., expressed as a union of) k independent sets if and only if*

$$|S| \leq k \; \mathbf{rank}_{\mathcal{M}}(S) \textit{ for all } S \subseteq T. \tag{1}$$

Whether T is k-coverable, in the sense of Theorem 3.3, turns out to be equivalent to the point distribution on T being $\frac{1}{k}$-uncontentious. One direction is immediate: If T is covered by the k independent sets $B_1, \ldots, B_k \subseteq T$, then choosing a uniformly random B_i gives a $\frac{1}{k}$-balanced CRS for the point distribution on T. The converse is not immediate: a $\frac{1}{k}$-balanced CRS may output an arbitrary distribution over independent subsets of T. Comparing the characterizations in Theorem 3.3 and Theorem 3.2 eliminates this integrality gap. Observe that (1) can be rewritten as

$$\mathbf{rank}_{\mathcal{M}}(F \cap T) \geq \frac{1}{k} |F \cap T| \text{ for every } F \subseteq \mathcal{E},$$

which in turn is equivalent to the characterization in part (c) of Theorem 3.2 when applied to the point distribution on T. It follows that a set T of elements can be covered by k independent sets if and only if the point distribution on T is $\frac{1}{k}$-uncontentious. Generally speaking, contention resolution generalizes matroid covering beyond point distributions, and removes any integrality requirements on the distribution of outputs. Therefore, *we can interpret contention resolution on matroids as a distributional generalization and relaxation of matroid covering.*

3.3 Elementary Properties of Uncontentious Distributions

We now state two elementary, yet quite useful, properties of uncontentious distributions. First, we observe that an uncontentious distribution is, by necessity, approximately ex-ante feasible up to a scaling factor.

Proposition 3.4. *Let $\mathcal{M}$ be an independence system, and let $\mathcal{P}(\mathcal{M})$ be its polytope. If ρ is α-uncontentious and $x = x(\rho)$ are its marginals, then $\alpha x \in \mathcal{P}(\mathcal{M})$.*

Proof. Let ϕ be an α-balanced CRM for ρ. Let $R \sim \rho$, and let y be the marginals of the (random) set $\phi(R)$. Since $\phi(R)$ is independent, $y \in \mathcal{P}(\mathcal{M})$. The balance guarantee implies that $y \geq \alpha x$. Since $\mathcal{P}(\mathcal{M})$ is a packing polytope, it follows that $\alpha x \in \mathcal{P}(\mathcal{M})$. □

Next, we observe that uncontentious distributions are closed under mixing

Proposition 3.5. *Fix a set system. A mixture of α-uncontentious distributions is α-uncontentious.*

Proof. Follows directly from Theorem 3.2 (b) and linearity of expectations. □

3.4 Examples of Uncontentious Distributions

We now restrict attention to matroids, and present some instructive examples of uncontentious distributions. As mentioned previously, and shown in [11], every product distribution which is ex-ante feasible is $(1 - \frac{1}{e})$-uncontentious. Combined with Proposition 3.5, this extends to mixtures of product distributions.

Proposition 3.6. *Fix a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and let $\rho \in \Delta(2^{\mathcal{E}})$ be a mixture of ex-ante-feasible product distributions. It follows that ρ is $(1 - \frac{1}{e})$-uncontentious.*

More generally, if a distribution satisfies a certain strong notion of negative correlation, defined in [10] and further explored in [39], then it also is $(1 - \frac{1}{e})$-uncontentious.

Proposition 3.7. *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a matroid, and let $\rho \in \Delta(2^{\mathcal{E}})$ be a prior distribution which is ex-ante feasible for $\mathcal{M}$. Let ρ' denote the product distribution with the same marginals as ρ. Suppose that ρ satisfies the property of* submodular dominance*: for every submodular function f we have $\mathbf{E}_{R\sim\rho}[f(R)] \geq \mathbf{E}_{R'\sim\rho'} \mathbf{E}[f(R')]$.*[3] *It follows that ρ is $(1 - \frac{1}{e})$-uncontentious.*

Proof. Since ρ' is an ex-ante feasible product distribution, it is $(1 - \frac{1}{e})$-uncontentious as shown in [11]. Combined with Theorem 3.2 (b), the submodular dominance property, and the fact that matroid rank functions are submodular, we conclude that ρ is also $(1 - \frac{1}{e})$-uncontentious. □

[3] In fact, it suffices for ρ to satisfy the (weaker) property of dominance with respect to matroid rank functions (or, equivalently, their weighted sums).

As shown in [10], submodular dominance is stronger than the following standard notion of negative correlation for $R \sim \rho$: For all sets T, $\mathbf{Pr}[T \subseteq R] \leq \prod_{i \in T} \mathbf{Pr}[i \in R]$ and $\mathbf{Pr}[T \subseteq \overline{R}] \leq \prod_{i \in T}(1 - \mathbf{Pr}[i \in R])$.[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 3.8. *Let $\mathcal{M}$ be a k-uniform matroid with n elements where $2 \leq k \leq n$. Let the random set R be a uniformly random base with some small constant probability ϵ, and the empty set 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 $\mathbf{Pr}[i \in R] = \mathbf{Pr}[j \in R] = \epsilon \frac{k}{n}$, yet $\mathbf{Pr}[i \in R | j \in R] = \mathbf{Pr}[j \in R | i \in R] = \frac{k-1}{n-1} \gg \epsilon \frac{k}{n}$.*

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

Example 3.9. *Consider the rank one matroid with elements $[n] = \{1, \ldots, n\}$. For $k = 0, 1, \ldots, n-1$, let $R = \{1, \ldots, k\}$ with probability $2^{-(k+1)}$, and let $R = [n]$ with remaining probability 2^{-n}. The random set R is $\frac{1}{2}$-uncontentious, as evidenced by the CRM ϕ with $\phi(\{1, \ldots, k\}) = \{k\}$ and $\phi(\emptyset) = \emptyset$, and a simple calculation. Note the positive correlation between elements $i < j$: $\mathbf{Pr}[j \in R] = 2^{-j}$, and $\mathbf{Pr}[j \in R | i \in R] = 2^{i-j} \gg \mathbf{Pr}[j \in R]$.*

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

Example 3.10. *Let $\mathcal{M}$ be a matroid with m pairwise-disjoint bases $B_1, \ldots, B_m$. For each $k = 1, \ldots, m-1$, let $R = \cup_{i=1}^{k} B_i$ with probability 2^{-k}, and let $R = \cup_{i=1}^{m} B_m$ with the remaining probability 2^{1-m}. The set R is $\frac{1}{2}$-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 $\mathbf{Pr}[e_j \in R] = 2^{1-j}$ and $\mathbf{Pr}[e_j \in R | e_i \in R] = 2^{i-j} \gg \mathbf{Pr}[e_j \in R]$.*

4 The Set of Improving Elements

We now examine the "improving elements" for a subsample of a weighted matroid, as originally defined by Karger [30] for purposes different from ours. We recall the well-known fact that the improving elements contribute a constant fraction of the omniscient optimal value, and point out that they are positively correlated in general. As the primary contribution of this section, we leverage our characterization in Theorem 3.2 to prove that the improving elements are uncontentious. This serves as a key technical lemma for our main result, and may also be of independent interest. Our proof is somewhat technical, so we also explain in Appendix A why prior work does not appear to provide an alternate route to the same result.

We begin with the formal definition of improving elements.

Definition 4.1. *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 $S \subseteq \mathcal{E}$ be a* sample *which includes each element $i \in \mathcal{E}$ independently with probability p.*
- *Let $R = R(S) = \{i \in \mathcal{E} : \mathbf{rank}_w(S \cup i) > \mathbf{rank}_w(S)\}$.*

We use $\mathbf{Imp}(\mathcal{M}, p, w)$ to denote the distribution of the set R.

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

Equivalently, R is the set of elements which are not spanned by elements of greater or equal weight in the sample S. We assume — without loss of generality by standard tie-breaking arguments — that non-zero weights are distinct, in which case a convenient equivalent definition of the improving elements is

$$R = R(S) = \{i \in \mathcal{E} \setminus S : i \in \mathbf{OPT}_w(S \cup i)\}.$$

It is easy to see that, for matroids, the improving elements contribute at least a $(1-p)$ fraction of the weighted rank.

Fact 4.2. *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*

$$\mathbf{E}[w(R)] \geq \mathbf{E}[\mathbf{rank}_w(R)] \geq (1-p)\mathbf{rank}_w(\mathcal{M}).$$

We note that the improving elements do not follow a product distribution. This is illustrated by the special case of the rank one matroid on $[n]$ with weights $w_1 > w_2 > \ldots > w_n$, and $p = 1/2$: the distribution of the improving elements is that described in Example 3.9, which we showed exhibits positive correlation. In fact, the improving elements are positively correlated in general.

Proposition 4.3. *Let R be the set of improving elements for a weighted matroid $(\mathcal{M}, w)$. For any pair of elements i and j, membership in R is (weakly) positively correlated; formally*

$$Pr[\{i,j\} \in R] \geq \mathbf{Pr}[i \in R]\,\mathbf{Pr}[j \in R].$$

Moreover, the correlation inequality is strict if there is a circuit containing both i and j.

Proof. Let S and $R = R(S)$ be as in the definition of improving elements. If $i \in R(S)$, then $i \in R(S')$ for every $S' \subseteq S$, and the same for j. Thus, the events $i \in R$ and $j \in R$ are both decreasing in S. Since S is drawn from a product distribution, weak positive correlation follows by invoking the FKG inequality.

For strict correlation, let $C \cup \{i,j\}$ be a circuit, and suppose without loss of generality that $w_i > w_j$. Conditioned on $S \setminus \{i,j\} = C$, the events $i \in R$ and $j \in R$ are strictly positively correlated: $j \in R$ implies $i \in R$ with certainty, whereas $\mathbf{Pr}[i \in R] \leq 1 - p < 1$. Conditioned on $S \setminus \{i,j\} = B$ for any other B, however, the events $i \in R$ and $j \in R$ are still decreasing in $S \cap \{i,j\}$, and therefore are still (weakly) positively correlated. It follows that $i \in R$ and $j \in R$ are strictly positively correlated overall. □

Despite being positively correlated, we show that the set of improving elements is uncontentious.

Lemma 4.4. *Given a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with weights w, the random set of improving elements with parameter $p \in (0,1)$ is p-uncontentious.*

We prove this lemma by leveraging our polyhedral characterization of uncontentious distributions. Let R be the set of improving elements with parameter p, and let S be the associated sample in Definition 4.1. We use (c) from Theorem 3.2: for arbitrary $F \subseteq \mathcal{E}$, we show that $\mathbf{E}[\mathbf{rank}(R \cap F)] \geq p\,\mathbf{E}[|R \cap F|]$. We break this up into the following three sublemmas.

Sublemma 4.5. $\mathbf{E}[\mathbf{rank}(R \cap F)] \geq (1-p)\,\mathbf{E}[|F \cap \mathbf{OPT}_w(S \cup F)|]$

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 $\mathbf{E}[\mathbf{rank}(R \cap F)] \geq (1-p)|F \cap \mathbf{OPT}_w(T \uplus F)|]$.

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

Sublemma 4.6. $|F \cap \mathbf{OPT}_w(S \cup F)| \geq |F \cap \mathbf{OPT}_w(S)|$

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 $|F \cap \mathbf{OPT}_w(T)|$ 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:

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

□

Sublemma 4.7. $\mathbf{E}[|\mathbf{OPT}_w(S) \cap F|] = \frac{p}{1-p} \mathbf{E}[|R \cap F|]$

Proof. We will show that $\mathbf{Pr}[i \in \mathbf{OPT}_w(S)] = \frac{p}{1-p} \mathbf{Pr}[i \in R]$ for arbitrary $i \in F$, which suffices.

Fix $i \in F$, and let $S_{>i} = \{j \in S : w_j > w_i\}$. Observe that $i \in \mathbf{OPT}_w(S)$ if and only if $i \notin \mathbf{span}(S_{>i})$ and $i \in S$. Furthermore, $i \in R$ if and only if $i \notin \mathbf{span}(S_{>i})$ and $i \notin S$. Since membership of i in S is independent of $S_{>i}$, we get

$$\mathbf{Pr}[i \in \mathbf{OPT}_w(S)] = p\,\mathbf{Pr}[i \notin \mathbf{span}(S_{>i})], \text{ and} \tag{2}$$

$$\mathbf{Pr}[i \in R] = (1-p)\,\mathbf{Pr}[i \notin \mathbf{span}(S_{>i})]. \tag{3}$$

Combining (2) and (3) yields $\mathbf{Pr}[i \in \mathbf{OPT}_w(S)] = \frac{p}{1-p} \mathbf{Pr}[i \in R]$, as needed. □

5 Labeled Contention Resolution

In this section, we introduce a generalization of contention resolution which associates labels with the active elements, and requires balance not only across elements but also across the different labels. This generalization will serve as an important waypoint in our reduction from the matroid secretary problem to contention resolution.

Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a set system, and let $\mathcal{L}$ be a finite set of *labels*. A *labeled set* for $(\mathcal{M}, \mathcal{L})$ is a pair (R, L) where $R \subseteq \mathcal{E}$ and $L : R \to \mathcal{L}$ is a *labeling* of R with $\mathcal{L}$. A *labeled contention resolution map (LCRM)* ϕ for $(\mathcal{M}, \mathcal{L})$ takes as input such a labeled set (R, L), where R is again referred to as the set of *active* elements, and outputs $T \in \mathcal{I}$ with the property $T \subseteq R$. If an element is in the output of ϕ, we again say it has been *accepted* (or *selected*), otherwise we say it has been *rejected*. We say ϕ is *α-balanced* for a *prior distribution* ρ over labeled sets if, when the input (R, L) is drawn from ρ, we have

$$\mathbf{Pr}[e \in \phi(R, L) \wedge L(e) = \ell] \geq \alpha\,\mathbf{Pr}[e \in R \wedge L(e) = \ell]$$

for every $e \in \mathcal{E}$ and $\ell \in \mathcal{L}$.

When an (offline) α-balanced LCRM exists for a distribution ρ over labeled sets, we again say that ρ is *α-uncontentious* for $\mathcal{M}$. When $(R, L) \sim \rho$ and ρ is α-uncontentious, we often abuse terminology and also say that the random labeled set (R, L) is α-uncontentious.

In the online random order setting, elements of $\mathcal{E}$ arrive in a uniformly random order $(e_1, \ldots, e_n)$, and at the ith step the algorithm learns whether e_i is active — i.e., whether $e_i \in R$ — and if so the algorithm also learns its label $L(e_i)$. The algorithm must then make an irrevocable decision on whether to accept e_i. We refer to an LCRM that can be implemented in this model as a *random-order LCRM (RO-LCRM).*

Remaining notions and terms from unlabeled contention resolution generalize naturally to the labeled setting: A *labeled contention resolutions scheme (LCRS)* Φ for set system $\mathcal{M}$ takes as input a set $\mathcal{L}$ of labels and a description of a prior distribution ρ over labeled sets for $(\mathcal{M}, \mathcal{L})$, and implements an LCRM ϕ_ρ for $(\mathcal{M}, \mathcal{L})$. As before, an LCRS Φ may be offline or online, and is α-balanced for a class of prior distributions if, for every ρ in that class, ϕ_ρ is α-balanced for ρ. When an LCRS operates in the random order model, we refer to it as an RO-LCRS for short. We omit reference to $\mathcal{M}$ and/or $\mathcal{L}$ when they are clear from context.

Note that classical contention resolution is the special case of labeled contention resolution in which each element of the ground set is associated with a single label. We also note that labeled contention resolution offers little beyond classical contention resolution in the offline model for matroids: if we think of labeled copies of an element as parallel elements in a new matroid, we obtain an equivalent unlabeled contention resolution problem.[5] Formally, for a matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and set $\mathcal{L}$ of labels, we define their "tensor product" $\mathcal{M} \otimes \mathcal{L} = (\mathcal{E} \times \mathcal{L}, \mathcal{I} \otimes \mathcal{L})$, where $\mathcal{I} \otimes \mathcal{L}$ includes $S \odot L = \{(e, L(e)) : e \in S\}$ for each $S \in \mathcal{I}$ and each $L : S \to \mathcal{L}$.[6] It is easy to verify that $\mathcal{M} \otimes \mathcal{L}$ is a matroid: each element of $\mathcal{M}$ was just replaced with $|\mathcal{L}|$ parallel copies of itself, one for each label. In the offline setting, a labeled contention resolution problem on $\mathcal{M}$ and $\mathcal{L}$ is equivalent to an unlabeled one on $\mathcal{M} \otimes \mathcal{L}$. In particular, we can think of a labeled set (R, L) for $\mathcal{M}$ as an unlabeled set $R \odot L = \{(e, L(e)) : e \in R\}$ for $\mathcal{M} \otimes \mathcal{L}$, where at most one copy of each original element $e \in \mathcal{E}$ — that corresponding to label $L(e)$, in the event $e \in R$ — is ever active. It follows that a random labeled set (R, L) is α-uncontentious (in the labeled sense, for $\mathcal{M}$) if and only if the corresponding unlabeled set $R \odot L$ is α-uncontentious (in the unlabeled sense, for $M \otimes \mathcal{L}$). Given this equivalence, the following labeled analogue of Proposition 3.5, which will be useful in Section 9, is immediate.

Proposition 5.1. *Fix a matroid and a set of labels. A mixture of α-uncontentious distributions over labeled sets is α-uncontentious.*

Our main concern will be labeled contention resolution in the online random order model. Unlike in the offline model, the reduction from the labeled to the unlabeled problem is nontrivial, as will be shown in Section 10.[7]

6 Overview of Main Result

As our main result, we show that the matroid secretary conjecture is equivalent to the existence of random-order contention resolution schemes for matroids which are competitive with the best offline scheme, universally for all prior distributions. We formalize this through the notion of a *universal* contention resolution scheme, which can be defined in both the labeled and unlabeled settings.

[5] More generally, this is also the case for any family of set systems closed under duplication of elements.

[6] We note that concurrent independent work [9] defined an essentially identical operation, which they termed a "partition extension" of a constraint.

[7] Though not a concern of this paper, the relationship between the labeled and unlabeled problems is interesting to contemplate in other online order models. In the adversarial order model, it is not too hard to see that the two problems are again equivalent. In the free order model, however, no such equivalence is immediately obvious.

Definition 6.1. *Given a set system, universal contention resolution schemes are defined as follows:*

- *For $\alpha, \beta \in [0,1]$ with $\beta \leq \alpha$, a (β,α)-*universal CRS [LCRS] *is a CRS [LCRS] which is β-balanced for the class of α-uncontentious distributions.*
- *For $c \in [0,1]$, a c-*universal CRS [LCRS] *is one which is $(c\alpha,\alpha)$-universal, simultaneously for all $\alpha \in [0,1]$.*

Notice that we distinguish two notions of universality, one seemingly stronger than the other. The *strong* notion of c-universality requires competing with the best offline scheme for *every* distribution. The *weak* notion of (β,α)-universality only requires such a guarantee for α-uncontentious distributions where α is fixed.

The question of the existence of universal schemes is only interesting in non-offline models: By the definition of uncontentious distributions, there is a 1-universal offline CRS for every set system. The only nontrivial settings of the parameters in Definition 6.1 are $c > 0$ and $\beta, \alpha \notin \{0,1\}$, since every CRS is 0-universal, while the identity CRS is $(1,1)$-universal. We will be concerned with the existence of universal RO-CRSs with $0 < \beta \leq \alpha < 1$ or $c > 0$, and matroid set systems.

We are now ready to state our main result.

Theorem 6.2. *The following three statements are equivalent*

(i) The matroid secretary conjecture (Conjecture 1).

(ii) There exists a constant $c > 0$ such that every matroid admits a c-universal RO-CRS.

(iii) There exist constants $\beta, \alpha \in (0,1)$ such that every matroid admits a (β,α)-universal RO-CRS.

We emphasize that Theorem 6.2 refers only to the classical, unlabeled, form of contention resolution. Note that the strong notion of universal contention resolution (ii) clearly implies the weaker notion (iii).[8]

Our proof of Theorem 6.2 proceeds as follows. First, we reduce the strong form of universal contention resolution to the secretary problem. This shows that (i) implies (ii). We note that this reduction holds beyond matroids, for any independence system.

Lemma 6.3. *If there is a c-competitive algorithm for the secretary problem on an independence system $\mathcal{M}$, then for each $\widetilde{c} < c$ there is a $\widetilde{c}$-universal RO-CRS for $\mathcal{M}$.*

Second, we reduce the matroid secretary problem to the weak form of universal contention resolution. This shows that (iii) implies (i), thereby completing the proof of Theorem 6.2. We break this up into three component reductions: one from the matroid secretary problem to the (correlated) matroid prophet secretary problem, one from the matroid prophet secretary problem to universal labeled contention resolution, and finally one from labeled to unlabeled contention resolution, all in the online random order model. Recall that each of these reductions serves to tackle a distinct technical obstacle, as discussed in Section 1.

Lemma 6.4. *If there is a c-competitive algorithm for the matroid prophet secretary problem with finitely-supported arbitrarily-correlated priors, then there is a $\frac{c}{4096}$-competitive algorithm for the matroid secretary problem.*

[8] A notable, and perhaps surprising, consequence of Theorem 6.2 is that the existence of an $(\Omega(1),\alpha)$-universal RO-CRS on matroids for some $\alpha \in (0,1)$ implies the same for all other $\alpha' \in (0,1)$.

Lemma 6.5. *If there is a (β,α)-universal RO-LCRS for a matroid $\mathcal{M}$, then there is a $\beta(1-\alpha)$-competitive algorithm for the matroid prophet secretary problem on $\mathcal{M}$ with finitely-supported arbitrarily-correlated priors.*

Lemma 6.6. *If every matroid admits a (β,α)-universal RO-CRS, then for each $\widetilde{\beta} < \beta$, every matroid admits a $(\widetilde{\beta},\alpha)$-universal RO-LCRS.*

We prove Lemmas 6.3, 6.4, 6.5, and 6.6 in Sections 7, 8, 9, and 10, respectively. Recall Figure 1, as well as the accompanying discussion in Section 1 of the challenges and technical approaches associated with our four reductions.

7 Reducing Contention Resolution to Secretary

In this section we reduce universal contention resolution in the random order model to the secretary problem on the same independence system, thereby proving Lemma 6.3.

Fix an independence system $\mathcal{M} = (\mathcal{E},\mathcal{I})$, and let A be a c-competitive algorithm for the secretary problem on $\mathcal{M}$. We assume without loss of generality that A only accepts non-zero weight elements — any secretary algorithm for a downwards-closed set system can be easily modified to reject zero-weight elements without affecting its competitive ratio.

For a given α-uncontentious prior $\rho \in \Delta(2^{\mathcal{E}})$, it suffices to construct an RO-CRM ϕ which is $\widetilde{c}\alpha$-balanced for ρ, for an arbitrary $\widetilde{c} < c$. We exhibit such an RO-CRM ϕ which simply runs the secretary algorithm A in parallel, using a particular randomized assignment of "dual" weights to the elements, and accepts an element precisely when A accepts it. The weight assigned to an element i depends on whether that element is active as well as on some internal randomness in ϕ, but not on the position of i in the online order nor on the identities of the other active elements. This allows the weights to be computed online as elements arrive, and guarantees that the weights are independent of the arrival order as stipulated in secretary problems. We note that we only show existence, for each prior ρ and secretary algorithm A, of a randomized weight assignment which can be computed online and yields our balance guarantee. We do not show that the associated computations can be performed efficiently, and anticipate that tractability of these computations hinges on the form of the algorithm A, as well as on how the prior ρ is represented as input.

Our RO-CRM ϕ will be a mixture of RO-CRMs ϕ_w, where $w \in \mathbb{R}_+^{\mathcal{W}}$ is a (random) weight vector we will define later. The map ϕ_w is implemented online as follows: When element $i \in \mathcal{E}$ arrives, if i is active then it is presented to A with weight w_i, otherwise it is presented to A with weight 0. The element is accepted by ϕ_w if and only if it is accepted by A. By our assumption that A discards zero-weight elements, the accepted elements are an independent subset of the active elements, as required of a contention resolution map. As described in Section 3.1, we think of an RO-CRM mathematically as a randomized function mapping the set R of active elements to an independent subset of R.

Sublemma 7.1. *Let $R \sim \rho$ be the set of active elements, where ρ is α-uncontentious. For every weight vector w, we have*

$$\mathbf{E}[w(\phi_w(R))] \geq c\alpha\, \mathbf{E}[w(R)],$$

where expectation is over R, the random arrival order, and any internal randomness in the algorithm A.

Proof. Condition on the set R of active elements, and let $w'_i = w_i$ if $i \in R$ and $w'_i = 0$ otherwise. The elements $i \in \mathcal{E}$ are presented to A in a uniformly random order with weights w'_i, and $\phi_w(R) \subseteq R$

is the set of elements accepted by A. Since A is c-competitive, it follows that

$$\mathbf{E}[w'(\phi_w(R))] \geq c\, \mathbf{rank}_{w'}(\mathcal{M}).$$

Since $w'(\phi_w(R)) = w(\phi_w(R))$, and $\mathbf{rank}_{w'}(\mathcal{M}) = \mathbf{rank}_w(R)$ by downwards closure, we get

$$\mathbf{E}[w(\phi_w(R))] \geq c\, \mathbf{rank}_w(R) \tag{4}$$

for every fixed R, where expectation is over the random order and any internal randomness in A.

For $R \sim \rho$, Theorem 3.2 (b) gives

$$\mathbf{E}[\mathbf{rank}_w(R)] \geq \alpha\, \mathbf{E}[w(R)]. \tag{5}$$

Taking expectation with respect to R in Equation (4), and combining with Equation (5), completes the proof. □

Our RO-CRM ϕ for the α-uncontentious distribution ρ will be a mixture of the ϕ_w described above. Using what is essentially a duality argument, we will show that there exists a distribution $\sigma = \sigma(\rho)$ over weight vectors such that the RO-CRM ϕ_σ which samples $w \sim \sigma$ upfront then invokes ϕ_w is $\widetilde{c}\alpha$-balanced for ρ. This duality argument is most easily formalized by direct application of the separating hyperplane theorem.

Let $R \sim \rho$ be the set of active elements. For each element $i \in \mathcal{E}$, let $x_i = \mathbf{Pr}[i \in R]$ be the probability i is active, and $y_i(w) = \mathbf{Pr}[i \in \phi_w(R)]$ be the probability — taken over R as well as the random arrival order — that i is selected by ϕ_w. We assume without loss of generality that $x_i > 0$ for all i — elements that are never active can be discarded. For each distribution σ over weight vectors w and element $i \in \mathcal{E}$, let $y_i(\sigma) = \mathbf{Pr}[i \in \phi_\sigma(R)] = \mathbf{Pr}[i \in \phi_w(R)]$, where the probability is taken over the active elements R, the random arrival order, and the random weight vector $w \sim \sigma$. Now let

$$\mathcal{Y} = \left\{ y(\sigma) : \sigma \in \Delta(\mathbb{R}_+^{\mathcal{E}}) \right\} \subseteq [0,1]^{\mathcal{E}}$$

be the family of all selection probabilities achievable by some RO-CRM of the form ϕ_σ. It is immediate that $\mathcal{Y}$ is the convex hull of $\left\{ y(w) : w \in \mathbb{R}_+^{\mathcal{E}} \right\}$, and hence $\mathcal{Y}$ is a convex subset of $[0,1]^{\mathcal{E}}$.[9]

A $\widetilde{c}\alpha$-balanced scheme for ρ of the form ϕ_σ exists if and only if $\mathcal{Y}$ intersects with the upwards-closed convex set $\left\{ z \in \mathbb{R}_+^{\mathcal{E}} : z_i \geq \widetilde{c}\alpha x_i \text{ for all } i \right\}$. Suppose for a contradiction that this intersection is empty. By the separating hyperplane theorem, there exists non-zero $w \in \mathbb{R}_+^{\mathcal{E}}$ such that $w_i y_i \leq \widetilde{c}\alpha \sum_i w_i x_i$ for all $y \in \mathcal{Y}$. In particular,

$$\sum_i w_i y_i(w) \leq \widetilde{c}\alpha \sum_i w_i x_i < c\alpha \sum_i w_i x_i.$$

Since $\sum_i w_i y_i(w) = \mathbf{E}[w(\phi_w(R))]$ and $\sum_i w_i x_i = \mathbf{E}[w(R)]$, we obtain a contradiction with Sublemma 7.1. This concludes the proof of Lemma 6.3.

8 Reducing Secretary to Prophet Secretary

In this section we reduce the secretary problem to the prophet secretary problem on the same matroid with an arbitrarily-correlated and finitely-supported distribution on weight vectors, thereby proving Lemma 6.4. Our reduction loses a constant factor in the competitive ratio. The purpose

[9] However, one can easily describe (perhaps contrived) secretary algorithms A which result in $\mathcal{Y}$ not being closed. This precludes application of the strict separating hyperplane theorem in general.

of this reduction is to remedy the first technical obstacle to reducing the secretary problem to contention resolution, as outlined in Section 1: The distribution of improving elements — being a function of an adversarial weight vector — is a-priori unknown. Reducing to the prophet secretary problem replaces the adversarial weight vector with a stochastic one, thereby inducing a contention resolution problem on a known — and uncontentious by Lemma 4.4 and Proposition 3.5 — mixture of improving element distributions.

First, we observe that we can restrict attention to instances of the matroid secretary problem which are *normalized*, in that the offline optimal value is roughly 1, and *discretized*, in that weights are contained in a known finite set. The following Sublemma is shown using standard arguments, and its proof is therefore deferred to Appendix B. We note that we make no attempt to optimize the constants here.

Sublemma 8.1. *The matroid secretary problem reduces, at a cost of a factor of* 256 *in the competitive ratio, to its special case where the matroid* $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ *and weights* w *are guaranteed to satisfy the following:*

- *Normalized:* $\mathbf{rank}_w(\mathcal{M}) \in \left[\frac{1}{16}, 1\right]$.
- *Discretized: The weight* w_e *of each element* $e \in \mathcal{E}$ *is either zero, or is an integer power of* 2 *contained in* $\left[\frac{1}{256\ \mathbf{rank}(\mathcal{M})}, 1\right]$.

We now fix the matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$, and reduce the normalized and discretized matroid secretary problem on $\mathcal{M}$, in the sense of Sublemma 8.1, to the prophet secretary problem on the same matroid $\mathcal{M}$, losing a constant factor in the reduction. To keep the proof generic, we use $a = \frac{1}{16}$ to denote the (known) constant such that offline optimal value is guaranteed to lie in $[a, 1]$, and use $W = \{0\} \cup \left\{2^{-i} : i \in \mathbb{N}, i \leq \log_2(256\ \mathbf{rank}(\mathcal{M}))\right\}$ to denote the (known) finite set of *permissible weights* for $\mathcal{M}$. We also use $\mathcal{W} = \left\{w \in W^{\mathcal{E}} : \mathbf{rank}_w(\mathcal{M}) \in [a, 1]\right\}$ to denote the (known) finite set of *permissible weight vectors* for $\mathcal{M}$, yielding a normalized and discretized instance.

Our reduction invokes minimax duality to replace the adversarially-chosen weight vector w, as in the secretary problem, with a weight vector drawn from a known and arbitrarily-correlated distribution μ, as in the prophet secretary problem. Discretization is needed so that we can invoke the minimax theorem for finite games. However, straightforward application of the minimax theorem produces a variant of the prophet secretary problem where the goal is to maximize the expected ratio between the online and offline optimal values, rather than the (usual) goal of maximizing the ratio of the two expectations. Normalization serves to obviate the distinction between these two goals.

An algorithm A for a normalized and discretized secretary problem on $\mathcal{M}$ maps a permissible weight vector $w \in \mathcal{W}$ and an order $\pi \in \mathcal{E}!$ on the elements to an independent set $A(\pi, w) \in \mathcal{I}$. When A is deterministic, we can think of it as a function from $\mathcal{W} \times \mathcal{E}!$ to $\mathcal{I}$. Since $\mathcal{W}$, $\mathcal{E}!$, and $\mathcal{I}$ are all finite sets, there are finitely many such functions that are computable online. A randomized algorithm can be thought of as simply a distribution over these functions. For an algorithm A, be it deterministic or randomized, we use $\mathbf{val}(A, w) = \mathbf{E}_\pi[w(A(\pi, w))]$ to denote the expected weight of the independent set chosen by algorithm A for weight vector w, where expectation is over the uniformly random order $\pi \sim \mathcal{E}!$. Note that $\mathbf{val}(A, w)$ is a random variable when A is randomized.

Consider the following finite two-player zero-sum game played between an algorithm player and an adversary. The pure strategies of the algorithm player are deterministic algorithms for the secretary problem on $\mathcal{M}$, which we think of as functions from $\mathcal{W} \times \mathcal{E}!$ to $\mathcal{I}$, and mixed strategies are randomized algorithms. The pure strategies for the adversary are the permissible weight vectors

$\mathcal{W}$. The algorithm player's utility if he plays a deterministic algorithm A and the adversary plays w is simply the competitive ratio of A on w, given by $\frac{\mathbf{val}(A,w)}{\mathbf{rank}_w(\mathcal{M})}$.

For a randomized algorithm A for the secretary problem, its competitive ratio on a weight vector w is given by $\frac{\mathbf{E}[\mathbf{val}(A,w)]}{\mathbf{rank}_w(\mathcal{M})} = \mathbf{E}\left[\frac{\mathbf{val}(A,w)}{\mathbf{rank}_w(\mathcal{M})}\right]$, where expectation is over any internal randomness in A. The worst-case competitive ratio of A is at least d if

$$\forall w \in \mathcal{W} : \mathop{\mathbf{E}}_{A}\left[\frac{\mathbf{val}(A,w)}{\mathbf{rank}_w(\mathcal{M})}\right] \geq d \tag{6}$$

Inequality (6) can be equivalently interpreted as follows: if the algorithm player moves first by playing mixed strategy A, he guarantees an expected utility of at least d regardless of the response w of the adversary. By the minimax theorem for finite two-player zero-sum games, and through the associated dual pair of linear programs, the design of an algorithm A satisfying Inequality (6) reduces to the following (dual) problem faced by an algorithm player who moves second: for each $\mu \in \Delta(\mathcal{W})$ (a mixed strategy of the adversary), design an algorithm $B = B(\mu)$ for the secretary problem on $\mathcal{M}$ which satisfies:

$$\mathop{\mathbf{E}}_{B}\mathop{\mathbf{E}}_{w\sim\mu}\left[\frac{\mathbf{val}(B,w)}{\mathbf{rank}_w(\mathcal{M})}\right] \geq d \tag{7}$$

We note that our minimax reduction is not necessarily efficient, as both players in our zero-sum game have exponentially many strategies in the size of the ground set of the matroid. An efficient reduction is not necessary, however, for our (information theoretic) result. We also note that there is no benefit to randomization in B when computational efficiency is not a concern: a randomized algorithm B satisfying inequality (7) can be derandomized, albeit perhaps inefficiently, by appropriately choosing a deterministic algorithm in its support. Nevertheless, we permit randomization in B for our reduction to be as general as possible.[10]

Finally, we claim that a c-competitive algorithm B for the prophet secretary problem on $\mathcal{M}$ and μ satisfies inequality (7) with $d = a \cdot c$. By definition, the assumption that B is a c-competitive prophet secretary algorithm for $\mathcal{M}$ and μ can be written as

$$\frac{\mathbf{E}_B\,\mathbf{E}_{w\sim\mu}[\mathbf{val}(B,w)]}{\mathbf{E}_{w\sim\mu}[\mathbf{rank}_w(\mathcal{M})]} \geq c.$$

It follows that

$$\begin{aligned}
\mathop{\mathbf{E}}_{B}\mathop{\mathbf{E}}_{w\sim\mu}\left[\frac{\mathbf{val}(B,w)}{\mathbf{rank}_w(\mathcal{M})}\right] &\geq \mathop{\mathbf{E}}_{B}\mathop{\mathbf{E}}_{w\sim\mu}[\mathbf{val}(B,w)] && (\mathbf{rank}_w(\mathcal{M}) \leq 1 \text{ for all } w \in \mathcal{W})\\
&= a \cdot \frac{\mathbf{E}_B\,\mathbf{E}_{w\sim\mu}[\mathbf{val}(B,w)]}{a}\\
&\geq a \cdot \frac{\mathbf{E}_B\,\mathbf{E}_{w\sim\mu}[\mathbf{val}(B,w)]}{\mathbf{E}_{w\sim\mu}[\mathbf{rank}_w(\mathcal{M})]} && (\mathbf{rank}_w(\mathcal{M}) \geq a \text{ for all } w \in \mathcal{W})\\
&\geq a \cdot c
\end{aligned}$$

Since our reduction lost a factor of 256 in the normalization and discretization step (Sublemma 8.1), and a factor of $1/a = 16$ due to the discrepancy between the objective of the matroid prophet secretary problem and the dual of the matroid secretary problem, this completes the proof of Lemma 6.4 with the claimed loss in the competitive ratio of $256 \times 16 = 4096$.

[10] This is convenient since the reduction from the prophet secretary problem to contention resolution in Sections 9 and 10 will, in general, produce a randomized algorithm, as contention resolution schemes are typically randomized.

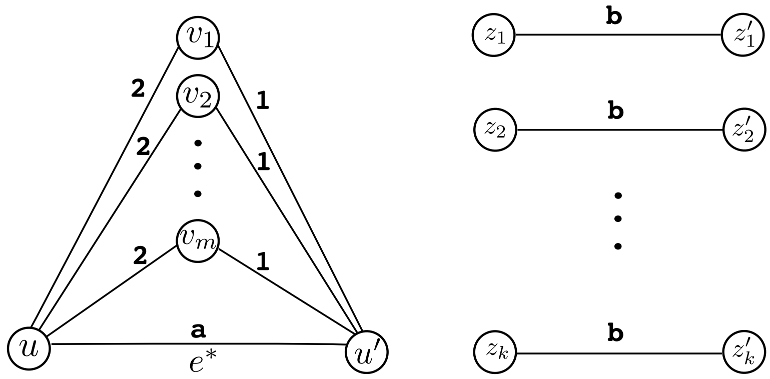

Figure 2: Modified Hat Example. This graphical matroid is truncated to rank m.

9 Reducing Prophet Secretary to Labeled Contention Resolution

In this section we reduce the prophet secretary problem to labeled contention resolution on the same matroid, thereby proving Lemma 6.5. The purpose of this reduction is to remedy the second technical obstacle outlined in Section 1: With a stochastic weight vector exhibiting correlations, contention resolution applied to the set of improving elements is no longer guaranteed to recover a constant fraction of the expected weighted rank. This is illustrated by the following example.

Example 9.1. *Consider the truncated graphical matroid in Figure 2, with the weights labeling the edges and $k \gg 2m$. We can guarantee that weights are distinct by introducing small perturbations. The graph on the left is the classical "hat example" often employed in the literature on the matroid secretary problem. We take the disjoint union of the hat example with the free matroid on k elements (represented by the k isolated edges on the right), and truncate the resulting matroid to a rank of m. We fix the sampling parameter $p = \frac{1}{2}$, and examine the set of improving elements for two settings of the weights a and b.*

For the first scenario, let $a = 4$ and $b = 3$. With high probability as k grows large, the set R of improving elements does not include any of the "hat" edges with weights 1 or 2. Moreover, $\mathbf{Pr}[e^* \in R] = \frac{1}{2}$. The following simple scheme is $(\frac{1}{2}-o(1))$-balanced: Discard the edges $\{(z_i, z_i') : i \in [k]\}$ with probability $\frac{1}{2}$ (and otherwise discard nothing),[11] then run greedy random-order contention resolution on the remaining edges.

For the second scenario we let $a = \infty$ (or a very large constant) and $b = 0$. Setting $b = 0$ effectively takes the edges $\{(z_i, z_i') : i \in [k]\}$ "out of the running", leaving only the hat example. The set R of improving elements, though $\frac{1}{2}$-uncontentious, is now less amenable to greedy contention resolution: when $e^* \in R$, there are typically many "hats" in R as well: for a constant fraction of $i \in [m]$, both edges (u, v_i) and (v_i, u') are in R. It follows that the above-described discard-then-greedy scheme is no longer $\Omega(1)$-balanced. In particular, it selects e^* with probability $O(\frac{1}{m}) = o(1)$, despite the fact that $\mathbf{Pr}[e^* \in R] = \frac{1}{2}$. A slightly more involved contention resolution scheme is needed for a constant balance ratio.

Suppose we randomize between the above scenarios, with each scenario equally likely. Let R be the resulting set of improving elements, and note that R is $\frac{1}{2}$-uncontentious. It is easy to verify that the discard-then-greedy scheme is $(\frac{1}{4} - o(1))$-balanced here. However, e^* is accepted with probability $\frac{1}{2} - o(1)$ when it is active with weight 4 (in the first scenario), but with probability $o(1)$ when it is

[11] Discarding these edges serves solely to guarantee balance for the "hat edges", in the low probability event that any of the hat edges are improving.

active with weight ∞ (in the second scenario). Therefore, the discard-then-greedy scheme does not recover a constant fraction of the expected weighted rank of R, despite being $\Omega(1)$-balanced.

A similar situation arises for any nontrivial randomization between the two scenarios, even if we make the second scenario exceedingly unlikely.

Example 9.1 suggests that we must constrain contention resolution to not "favor" improving elements that have low weight. We accomplish this by labeling each improving element with its weight, and requiring contention resolution in the (stronger) labeled sense. We assume without loss of generality that non-zero weights are distinct, and use the following labeled notion of improving elements.

Definition 9.2. *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a matroid, let $p \in (0, 1)$ be a parameter, and let $w \in \mathbb{R}_+^{\mathcal{E}}$ be a weight vector. The random* labeled set of improving elements *for $(\mathcal{M}, p, w)$ is the pair (R, L), where $R \sim \mathbf{Imp}(\mathcal{M}, p, w)$ is the (random) set of improving elements, and $L : R \to \mathbb{R}_+$ is the labeling with $L(e) = w(e)$ for all $e \in R$. We use $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, w)$ to denote the distribution of the labeled set (R, L).*

When w is fixed, each element e is associated with a single label $w(e)$, so labeled contention resolution for $(R, L) \sim \mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, w)$ is equivalent to unlabeled contention resolution for $R \sim \mathbf{Imp}(\mathcal{M}, p, w)$, and by Lemma 4.4 it follows that $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, w)$ is p-uncontentious in the labeled sense. When w is a drawn from a known prior μ with finite support, the labeled set of improving elements (R, L) is drawn from a mixture of the p-uncontentious distributions $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, w)$, for the finitely-many realizations of $w \in \mathbf{supp}(\mu)$. When $w \sim \mu$ and $(R, L) \sim \mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, w)$, we refer to (R, L) as the *labeled set of improving elements for* $(\mathcal{M}, p, \mu)$, and denote its distribution by $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$. The following is then a direct consequence of Proposition 5.1.

Sublemma 9.3 (Follows from Lemma 4.4 and Proposition 5.1)**.** *Let $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ be a matroid, let $p \in (0, 1)$, and let $\mu \in \Delta(\mathbb{R}_+^{\mathcal{E}})$ be a distribution over weight vectors with finite support. The distribution $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$ is p-uncontentious (in the labeled sense) for $\mathcal{M}$.*

Fixing matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and $p \in (0, 1)$, we reduce the prophet secretary problem on $\mathcal{M}$ to (β, α)-universal random-order labeled contention resolution with $\alpha = p$, and arbitrary β. The reduction is shown in Algorithm 1 for the prophet secretary problem on $\mathcal{M}$, which takes as an offline input a prior μ on weight vectors, and as its online inputs a sequence of weighted elements of $\mathcal{M}$. We assume that the online inputs to the Algorithm are distributed as specified in the prophet secretary problem, namely with $w \sim \mu$ and $\pi \sim \mathcal{E}!$ drawn independently, and analyze the algorithm's competitive ratio. In particular, we will see that the algorithm achieves its competitive ratio by resolving contention, in the random order model, for a labeled set drawn from the p-uncontentious distribution $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$.

Let R denote the elements improving S, as determined in Step (11), and let $L(e) = w(e)$ be the label of $e \in R$ determined in Step (12). Note that the algorithm interleaves elements S with those of $\mathcal{E} \setminus S$, as needed to produce a uniformly random permutation independent of (R, L). We let $\pi' \in \mathcal{E}!$ denote the resulting list of elements (whether active or inactive) fed to ϕ by Algorithm 1, in that order. First, we show that the inputs to ϕ are as stipulated in random-order contention resolution for $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$, and that ϕ is β-balanced for that distribution.

Sublemma 9.4. *The labeled set (R, L) follows the distribution $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$. Moreover, π' is a uniformly random order on $\mathcal{E}$ independent of (R, L).*

Proof. Since π is a uniformly random permutation of $\mathcal{E}$, and $k \sim \mathbf{Binom}(n, p)$, it follows that S includes each element of $\mathcal{E}$ independently with probability p. The set R consists of all elements

Algorithm 1 Reduction from matroid prophet secretary to labeled contention resolution.

Parameter: Matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with n elements.

Parameter: (β, α)-universal RO-LCRS Φ for matroid $\mathcal{M}$

Input: Finitely-supported prior distribution $\mu \in \Delta(\mathbb{R}_+^{\mathcal{E}})$.

Input: Online string $(e_1, w(e_1)), \ldots, (e_n, w(e_n))$, where $\pi = (e_1, \ldots, e_n) \in \mathcal{E}!$, and $w \in \mathbf{supp}(\mu)$.

```
Let p = α
Instantiate Φ with prior distribution Imp_lbl(M, p, μ), and let φ denote the resulting RO-LCRM
    for matroid M and finite set of labels L = {w'(e) : w' ∈ supp(μ), e ∈ E}.
Sample k ~ Binom(n, p).
Observe first k online inputs (e_1, w(e_1)), ... (e_k, w(e_k)) without accepting any.
Let S = {e_1, ..., e_k}.
Let i = 1 and j = k + 1                        ▷ Indexes elements e_i ∈ S and e_j ∈ E \ S
while i ≤ k or j ≤ n do                        ▷ While not all elements in E have been fed to φ
    Flip a biased coin with heads probability (n−(j−1)) / (n−(j−1)+k−(i−1))
    if Coin came up heads then                 ▷ Feed next element in E \ S to φ
        Read the next online input (e_j, w(e_j))
        if rank_w^M(S ∪ e_j) > rank_w^M(S) (i.e., e_j improves S) then
            Feed e_j as active to φ, with label w(e_j). Accept e_j if φ accepts it, otherwise Reject e_j.
        else
            Feed e_j as inactive to φ.
        end if
        Increment j
    else                                       ▷ Coin came up tails. Feed an element from S to φ
        Feed e_i as inactive to φ
        Increment i
    end if
end while
```

improving S with respect to weight vector w, so $R \sim \mathbf{Imp}(\mathcal{M}, p, w)$ by Definition 4.1. Since $L(e) = w(e)$ and $w \sim \mu$, it follows that $(R, L) \sim \mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$.

We now condition on S and w, which in turn fixes (R, L), and show that π' is a uniformly random permutation of $\mathcal{E}$. Since each iteration of the while loop feeds one of e_i or e_j to ϕ, and increments the corresponding counter (i or j), it follows that $\pi' = (e'_1, \ldots, e'_n)$ is a permutation of $\mathcal{E}$. Now consider the tth iteration of the while loop, let $S_t = S \setminus \{e'_1, \ldots, e'_{t-1}\}$ and $\overline{S}_t = (\mathcal{E} \setminus S) \setminus \{e'_1, \ldots, e'_{t-1}\}$, and notice that $S_t \cup \overline{S}_t = \mathcal{E} \setminus \{e'_1, \ldots, e'_{t-1}\}$ is the set of elements not yet fed to ϕ. It is easy to see inductively that $S_t = \{e_i, \ldots, e_k\}$ and $\overline{S}_t = \{e_j, \ldots, e_n\}$, where i and j are as in iteration t. Since π is uniformly random, e_i is a uniformly random element of S_t, and e_j is a uniformly random element of $\overline{S}_t$. The bias of the coin in Step (8) is such that $e'_t = e_j$ with probability $\frac{|\overline{S}_t|}{|S_t \cup \overline{S}_t|}$, and $e'_t = e_i$ with probability $\frac{|S_t|}{|S_t \cup \overline{S}_t|}$. Therefore, e'_t is a uniformly-random sample, without replacement, from $\overline{S}_t \cup S_t = \mathcal{E} \setminus \{e'_1, \ldots, e'_{t-1}\}$. It follows inductively that π' is a uniformly random permutation of $\mathcal{E}$. □

Sublemma 9.5. *The RO-LCRM ϕ instantiated in Step (2) is β-balanced for $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$.*

Proof. Follows directly from the fact that Φ is (β, p)-universal, and the fact that $\mathbf{Imp}_{\mathtt{lbl}}(\mathcal{M}, p, \mu)$ is p-uncontentious as shown in Sublemma 9.3. □

Let $T \subseteq R$ denote the set of elements accepted by Algorithm 1, as determined in Step (12). We can bound the expected weight of these elements as follows, where expectations are with respect to $w \sim \mu$, $\pi \sim \mathcal{E}!$, the internal randomness in Algorithm 1, and any randomness in the instantiated contention resolution map ϕ.

$$
\begin{aligned}
\mathbf{E}[w(T)] &= \sum_{e \in \mathcal{E}} \sum_{w_0 \in \mathcal{L}} w_0 \cdot \mathbf{Pr}[e \in T \wedge w(e) = w_0] \\
&= \sum_{e \in \mathcal{E}} \sum_{w_0 \in \mathcal{L}} w_0 \cdot \mathbf{Pr}[e \in T \wedge L(e) = w_0] \\
&\geq \beta \sum_{e \in \mathcal{E}} \sum_{w_0 \in \mathcal{L}} w_0 \cdot \mathbf{Pr}[e \in R \wedge L(e) = w_0] && \text{(Sublemmata 9.4 and 9.5)} \\
&= \beta \sum_{e \in \mathcal{E}} \sum_{w_0 \in \mathcal{L}} w_0 \cdot \mathbf{Pr}[e \in R \wedge w(e) = w_0] \\
&= \beta \, \mathbf{E}[w(R)] \\
&\geq \beta(1-p) \, \mathbf{E}[\mathbf{rank}_w(\mathcal{M})] && \text{(Fact 4.2 and Sublemma 9.4)} \\
&= \beta(1-\alpha) \, \mathbf{E}[\mathbf{rank}_w(\mathcal{M})]
\end{aligned}
$$

We conclude that Algorithm 1 is $\beta(1-\alpha)$ competitive for the prophet secretary problem on $\mathcal{M}$ with a finitely-supported prior. This concludes the proof of Lemma 6.5.

10 Reducing Labeled to Unlabeled Contention Resolution, Online

In this section we reduce labeled contention resolution to its unlabeled counterpart, in the random order model, thereby proving Lemma 6.6. This overcomes the final technical obstacle outlined in Section 1. We elaborate on this obstacle next.

Consider labeled contention resolution for matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and labels $\mathcal{L}$ in the online random-arrival model, and denote $n = |\mathcal{E}|$ and $m = |\mathcal{L}|$. Here, a labeled set (R, L) drawn from a known distribution is presented online to an LCRM for $\mathcal{M}$ and $\mathcal{L}$ as the string

$$x = x(R, L, \pi) = (e_1, \ell_1), (e_2, \ell_2), \ldots, (e_n, \ell_n), \tag{8}$$

where $\pi = (e_1, \ldots, e_n)$ is a uniformly random permutation of $\mathcal{E}$, and $\ell_i \in \mathcal{L} \cup \{\bot\}$ is the label $L(e_i)$ if $e_i \in R$ (i.e. e_i is active) and is $\bot$ otherwise. Entries of x are revealed online, with iteration i revealing (e_i, ℓ_i), at which point the LCRM must immediately decide whether to accept e_i in the event it is active.

Recall from Section 5 that, in the offline setting, labeled contention resolution on $\mathcal{M}$ and $\mathcal{L}$ reduces to unlabeled contention resolution on $\mathcal{M} \otimes \mathcal{L}$, via the map $(R, L) \to R \odot L$. It is therefore tempting to attempt a similar reduction in the random order model. When the unlabeled problem on $\mathcal{M} \otimes \mathcal{L}$ is considered in the random order model, the (unlabeled) active set $R \odot L$ is presented online to an (unlabeled) CRM for $\mathcal{M} \otimes \mathcal{L}$ as the string

$$y = y(R \odot L, \pi') = ((e'_1, \ell'_1), a_1), ((e'_2, \ell'_2), a_2), \ldots, ((e'_{nm}, \ell'_{nm}), a_{nm}), \tag{9}$$

where $\pi' = (e'_1, \ell'_1), \ldots, (e'_{nm}, \ell'_{nm})$ is a uniformly random permutation of $\mathcal{E} \times \mathcal{L}$, and $a_i \in \{\top, \bot\}$ designates whether $(e'_i, \ell'_i) \in R \odot L$. The string y is revealed online, with iteration i revealing $((e'_i, \ell'_i), a_i)$, at which point the CRM must immediately decide whether to accept (e'_i, ℓ'_i) in the event that $a_i = \top$. We emphasize that the string y is longer than x: whereas an element $e \in \mathcal{E}$ appears exactly once in x, it appears m times in y (once for each possible label, with at most one of these appearances active).

In attempting an online reduction from the labeled problem to its unlabeled counterpart, the problem we face at this point, intuitively, is the following: Given x, how do we "interleave" the "missing" (element,label) pairs to form the string y. This interleaving must be done online, before we know exactly which elements are active and what their labels are. Moreover, it must be such that the resulting order of (element,label) pairs in y is uniformly distributed, at least approximately, in order to make use of any guarantee on the balance ratio of the (unlabeled) RO-CRM. This, it so happens, is nontrivial.

The reader might understandably furrow their brow at this point: *Surely, any "reasonable" random-order contention resolution algorithm need only exploit the relative ordering of active elements. This is already uniformly random in x, so an arbitrary interleaving of the missing (element,label) pairs should suffice! Certainly, this additional difficulty is an artifact of the precise technical definition of the random order model, rather than a conceptually interesting distinction!* The reader would be justified in expressing such skepticism. However, intuitive as it may seem, this knee-jerk reaction is flawed in a formal sense. We exhibit in Appendix C a simple distribution which is uncontentious for the rank one matroid, though balanced contention resolution is impossible absent any guarantees on the order of inactive elements. Therefore, for online contention resolution to plausibly encode the matroid secretary problem, it needs to exploit randomness in the arrival order of both active and inactive elements!

10.1 Difficulties with Direct Approaches

We begin by explaining how the direct approach, namely reducing the online labeled contention resolution for $(\mathcal{M}, \mathcal{L})$ to online unlabeled contention resolution for $\mathcal{M} \otimes \mathcal{L}$, appears unlikely to succeed. Let $x = x(R, L, \pi)$ be the online input string to the labeled problem, as in Equation (8). All online reductions to the corresponding unlabeled problem which are conceivable to us fit the following template, which has oracle access to an online CRM ϕ' for $\mathcal{M} \otimes \mathcal{L}$, and produces an online LCRM ϕ for $(\mathcal{M}, \mathcal{L})$.

- While not all (element,label) pairs have been fed to ϕ', do one of the following:
  - (i) Read the next active (element,label) pair (e, ℓ) in x (if any), skipping inactive elements as needed. If (e, ℓ) has not previously been fed to ϕ', then feed $((e, \ell), \top)$ to ϕ', and accept e iff ϕ' accepts (e, ℓ).
  - (ii) "Hallucinate" an (element,label) pair (e, ℓ) which has not yet been fed to ϕ', and feed $((e, \ell), \bot)$ to ϕ'.

Notice that, in each iteration, the choice to do (i) or (ii), and the choice of "hallucination" (e, ℓ) in (ii), can depend on previously observed entries of x, on previous acceptance/rejection decisions of ϕ', and on previous "hallucinations". These choices may also be randomized. Let y denote the string fed to ϕ' through the course of the reduction, and let π' denote the sequence of (element,label) pairs appearing in y.

For an instantiation of the above template to serve as an approximation preserving reduction (up to a constant) from the labeled problem to its unlabeled counterpart in the online random order model, the following properties appear needed.

(a) Condition on the labeled set (R, L), and assume that the order $\pi \in \mathcal{E}!$ of elements in x is uniformly distributed (as is guaranteed by the random order model for the labeled problem). The order $\pi' \in (\mathcal{E} \times \mathcal{L})!$ of (element,label) pairs in y should be uniformly distributed (as is required by the random order model for the unlabeled problem) or approximately so (say, in terms of total variation distance).

(b) In the event that (e, ℓ) is an entry of x (i.e., e is active with label ℓ), it should hold with constant probability that $((e, \ell), \top)$ is an entry of y (i.e., (e, ℓ) is active in the corresponding unlabeled instance). This requires that (e, ℓ) is not "hallucinated" before it arrives in x.

Trivial instantiations of our template satisfy one of (a) or (b), but satisfying both (a) and (b) simultaneously appears impossible. To illustrate the difficulty, consider the special case where the number of active elements $|R|$ is known in advance. Arguably the most natural instantiation of our template in this special case, and one which at first glance appears promising, is as follows. In each iteration, with r active entries of x remaining and k (element,label) pairs not yet fed to ϕ', we choose (i) with probability $p = p(r, k) = \frac{r}{k}$ and choose (ii) otherwise. When (ii) is chosen, we let (e, ℓ) be a uniformly random draw from the k remaining (element,label) pairs. The probability p is chosen to reflect the proportion of active to inactive (element,label) pairs.

It is not too difficult to verify that (b) is satisfied for this reduction. However, it can be shown that the permutation π' is not uniformly distributed after conditioning on (R, L). To see this, consider an element $e \in R$ with $L(e) = \ell$. The probability that (e, ℓ) is the first (element,label) pair appearing in y is given by

$$\frac{|R|}{mn} \cdot \frac{1}{|R|} + \frac{mn - |R|}{mn} \cdot \frac{1}{mn},$$

where the first term corresponds to the event that (i) is chosen and (e, ℓ) is the first active (element,label) pair in x, and the second term corresponds to the event that (ii) is chosen and (e, ℓ) is hallucinated. Since $|R| \leq n$, this expression is at least $(2 - \frac{1}{m}) \cdot \frac{1}{mn}$. When the number of labels

m is large, this is almost twice the probability that (e, ℓ) would appear first in a uniformly random permutation on (element,label) pairs! In other words, an active (element,label) pair is almost twice as likely to appear early in π' than an inactive (element,label) pair, rendering π' far from uniformly distributed. In fact, we can show that the total variation distance between π' and the uniform distribution tends to 1 as m grows large, violating (a).

One might hope that different choices of $p(r, k)$, coupled with a different rule for choosing the hallucinated (element,label) pair in (ii), might remedy this failure. However, some examination suggests that such approaches are unlikely to succeed. The difficulty, intuitively, is the following: when hallucinating inactive (element,label) pairs early in the sequence y, we must do so without knowledge of which active (element,label) pairs appear later in x, and this is due to the online nature of the reduction. This gives active (element,label) pairs in x a "greater than fair" shot at appearing early in the sequence y (violating (a)), unless one is content with "ignoring" entries of x with high probability (which results in violating (b)). Therefore, there is a tension between requirements (a) and (b).

These difficulties appear intrinsic to online reductions from the labeled problem on $(\mathcal{M}, \mathcal{L})$ to the unlabeled problem on $\mathcal{M} \otimes \mathcal{L}$, leaving little hope for preserving the balance ratio with such a direct approach. A new idea appears to be needed.

10.2 An Indirect Approach: Duplicating the labels

We overcome these difficulties by reducing labeled contention resolution on $\mathcal{M}$ and $\mathcal{L}$ to unlabeled contention resolution on a much larger matroid than $\mathcal{M} \otimes \mathcal{L}$. Specifically, we "duplicate" each label a large number of times, creating many "identical copies" of each (element,label) pair. We associate an active (element,label) pair in x with one of its copies uniformly at random, leaving all other copies inactive. Roughly speaking, a random permutation π' of the duplicated (element,label) pairs converges in probability to a limiting permutation as the number of copies grows large, modulo the symmetry between copies. An active (element,label) pair from x is now merely a drop in a sea of its inactive brethren, and therefore interleaving x into π' has little influence on the probability distribution of π'.

Formally, we duplicate each label in $\mathcal{L}$ a large number K of times to form an expanded set of labels $\mathcal{L} \times \mathcal{C}$, where $\mathcal{C}$ is an abstract set for indexing copies with $|\mathcal{C}| = K$. We then reduce labeled contention resolution on $\mathcal{M}$ and $\mathcal{L}$ to unlabeled contention resolution on the matroid $\mathcal{M} \otimes (\mathcal{L} \times \mathcal{C}) = \mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$. For $\ell \in \mathcal{L}$ and $c \in \mathcal{C}$, we say the pair (ℓ, c) is a *copy* of label ℓ. We also say that an element $(e, \ell, c) \in \mathcal{E} \times \mathcal{L} \times \mathcal{C}$ of $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ is a *copy* of (e, ℓ).

An offline version of our reduction maps an active set in $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with labels in $\mathcal{L}$ to an (unlabeled) active set of $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ by selecting a copy of each label uniformly at random. Specifically, a labeled set of active elements (R, L) is mapped to the (unlabeled) set $R \odot L \odot C = \{(e, L(e), C(e)) : e \in R\}$ of elements of the matroid $\mathcal{M} \otimes \mathcal{L} \otimes C$, where $C(e) \in \mathcal{C}$ is chosen independently and uniformly at random for each $e \in R$.[12] It is easy to verify that if the random labeled set (R, L) is α-uncontentious (in the offline sense, of course), so is the random unlabeled set $R \odot L \odot C$.

Observation 10.1. *If (R, L) is an α-uncontentious labeled set for $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and $\mathcal{L}$, and $C : \mathcal{E} \to \mathcal{C}$ is chosen uniformly at random independently of (R, L), then $R \odot L \odot C$ is an α-uncontentious (unlabeled) set for $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$.*

[12] For convenience, we sometimes think of C as a function from $\mathcal{E}$ to $\mathcal{C}$, with the understanding that the restriction of C to R, which we denote by $C|_R$, is all that is relevant for defining $R \odot L \odot C$.

Proof. First, we can interpret an α-balanced offline LCRM for (R, L) as an offline CRM for $R \odot L$ in the matroid $\mathcal{M} \otimes \mathcal{L}$. Second, we can interpret the latter as an offline CRM for $R \odot L \odot C$ in the matroid $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ — in particular, one which ignores the index $C(e)$ for each element $(e, L(e), C(e))$. Since the indices $\{C(e)\}_{e \in \mathcal{E}}$ are independent of R and L, it follows that that the balance ratio is preserved. $\square$

In the online random order model, our reduction approximates the above-described map $(R, L) \to R \odot L \odot C$. Even more importantly, if the elements $e \in \mathcal{E}$ are presented to our reduction in uniformly random order (each tagged with its label $L(e)$ if $e \in R$, or $\perp$ otherwise), then its output is (approximately) a uniformly random permutation of $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$, with each (e, ℓ, c) tagged with $\top$ if $e \in R$, $L(e) = \ell$, and $C(e) = c$, and with $\perp$ otherwise. The error in both these approximations (the active set itself and the permutation), as measured in total variation distance, tends to 0 as the number of copies K of each label grows large. The reduction is summarized in Algorithm 2.

Algorithm 2 is an online LCRM for the matroid $\mathcal{M}$ with labels $\mathcal{L}$, which uses an online CRM ϕ' for $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ as a subroutine. At iteration i, the algorithm is presented with $x_i = (e_i, \ell_i)$, where ℓ_i is either a label (if e_i is active) or $\perp$ (if e_i is inactive), and in the former case must decide "on the spot" whether to accept e_i. To guide these decisions, the algorithm runs a parallel execution of the CRM ϕ', and feeds the elements of $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ (with each labeled as active or inactive) to ϕ' in a uniformly random order π'. For each active e_i in the input string x, the algorithm (tries to) activate the k_ith copy of (e_i, ℓ_i) in the order of appearance in π', where $\vec{k} = (k_1, \ldots, k_n)$ are n i.i.d. uniform samples from $[K]$ ordered in non-decreasing order. The algorithm accepts e_i if the corresponding activated copy of (e_i, ℓ_i) is accepted by ϕ'. To enable online acceptance/rejection decisions, we do the following: In each iteration i where e_i is active, the algorithm "skips ahead" in π' — feeding skipped over elements as inactive to ϕ' — until the desired k_ith copy of (e_i, ℓ_i) is reached, at which point this copy of (e_i, ℓ_i) is fed to ϕ' as active. The algorithm can fail when it "skips over" an element of $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$ which we later realize should have been activated.

The following sequence of sublemmata lead to a proof of Lemma 6.6.

Sublemma 10.2. *For each input string x, the probability that Algorithm 2* **FAIL***s tends to* 0 *as* $K \to \infty$.

Proof. In iteration i, the algorithm skips through π' until it finds the k_ith copy of (e_i, ℓ_i) in π'. It fails when that copy has already been passed over, in an earlier iteration $j < i$ while searching for the k_jth copy of (e_j, ℓ_j). In particular, for the algorithm to fail it must be that there are $j < i$ such that at least k_i copies of (e_i, ℓ_i) precede the k_jth copy (e_j, ℓ_j). We will show that this is a low-probability event.

First, we show that $k_i - k_j \geq K^{0.75}$, simultaneously for all $1 \leq j < i \leq n$, with high probability at least $1 - \frac{n^2}{K^{0.25}}$. By definition, this is equivalent to showing that n i.i.d. samples from the uniform distribution on $[K]$ are pairwise separated by at least $K^{0.75}$ with the claimed probability. The probability that the $(j+1)$st sample is at least $K^{0.75}$ away from the first j samples is at least $1 - 2j\frac{K^{0.75}}{K} = 1 - 2jK^{-0.25}$, so we get

$$\begin{aligned}
\mathbf{Pr}[\forall i \; k_{i+1} - k_i \geq K^{0.75}] &\geq \prod_{j=0}^{n-1} (1 - 2jK^{-0.25}) \\
&\geq 1 - \sum_{j=0}^{n-1} 2jK^{-0.25}
\end{aligned}$$

Algorithm 2 Reduction from labeled to unlabeled online contention resolution.

Parameter: Matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ with n elements, and a set $\mathcal{L}$ of m labels.

Parameter: Abstract index set $\mathcal{C}$ with $|\mathcal{C}| = K$

Parameter: Oracle access to an online (unlabeled) CRM ϕ' for $M \otimes \mathcal{L} \otimes \mathcal{C}$

Input: String $x = (e_1, \ell_1), \ldots, (e_n, \ell_n)$ given online, where $\pi = (e_1, \ldots, e_n)$ is a permutation of $\mathcal{E}$, and $\ell_i \in \mathcal{L} \cup \{\bot\}$.

Let $\pi' = \pi'(1), \ldots, \pi'(nmk)$ be a uniformly-random permutation of $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$

Draw n integers i.i.d. from the uniform distribution on $[K] = \{1, \ldots, K\}$. Sort these integers in non-decreasing order $k_1 \leq k_2 \leq \ldots \leq k_n$.

Let $i' = 1$ be the current position in π'

for $i = 1$ **to** n **do** $\triangleright$ Receive and process the ith online input $x_i = (e_i, \ell_i)$

  Read the next online input $x_i = (e_i, \ell_i)$

  if $\ell_i = \bot$ **then** $\triangleright$ e_i is inactive

    Do nothing

  else if There are at least k_i copies of (e_i, ℓ_i) among $\pi'(1), \ldots, \pi'(i'-1)$ **then**

    FAIL $\triangleright$ We already "missed" the k_ith copy of (e_i, ℓ_i)

  else $\triangleright$ Skip ahead to the k_ith copy of (e_i, ℓ_i)

    while $\pi'(i')$ is not the k_ith copy of (e_i, ℓ_i) seen so far in π' **do**

      Feed $(\pi'(i'), \bot)$ to ϕ' as its next online input.

      Increment i'

    end while $\triangleright$ $\pi'(i')$ is k_ith copy of (e_i, ℓ_i) in the ordered list π'

    Feed $(\pi'(i'), \top)$ to ϕ' as its next online input, and **ACCEPT** e_i if ϕ' accepts $\pi'(i')$, otherwise **REJECT** e_i.

    Increment i'

  end if

end for

while $i' \leq nmk$ **do** $\triangleright$ Complete the execution of ϕ' (may be omitted)

  Feed $(\pi'(i'), \bot)$ to ϕ' as its next online input.

  Increment i'

end while

$$= 1 - 2K^{-0.25} \sum_{j=0}^{n-1} j$$
$$\geq 1 - n^2 K^{-0.25}$$

Next, for $j < i$ we bound the probability that at least $k_j + K^{0.75}$ — with high probability a lower bound on k_i — copies of (e_i, ℓ_i) precede the k_jth copy of (e_j, ℓ_j). Consider K red balls and K blue balls ordered uniformly at random, with red balls corresponding to copies of (e_i, ℓ_i) and blue balls corresponding to copies of (e_j, ℓ_j). It suffices to upperbound the probability that, for any prefix of the randomly ordered balls, the number of red balls exceeds the number of blue balls by more than $K^{0.75}$. For the first b balls in the random order, we use the Hoeffding bound for sampling without replacement (see [28]) to get a probability upperbound of $\exp(-\frac{2K^{1.5}}{b}) \leq \exp(-\frac{2K^{1.5}}{2K}) = \exp(-\sqrt{K})$. Taking the union bound over all $b = 1, \ldots, 2K$, we get a bound of $\frac{2K}{e^{\sqrt{k}}}$.

Using the union bound, we conclude that the probability of failure is at most

$$\frac{n^2}{K^{0.25}} + \sum_{j<i} \frac{2K}{e^{\sqrt{k}}} \leq n^2 \left(\frac{1}{K^{0.25}} + \frac{2K}{e^{\sqrt{K}}} \right),$$

which tends to 0 as $K \to \infty$. □

Sublemma 10.3. *Let (R, L) be a random labeled set for an n-element matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and labels $\mathcal{L}$, let $\widetilde{C} : \mathcal{E} \to \mathcal{C}$ be chosen uniformly at random, let π be a uniformly random permutation of $\mathcal{E}$, and let $\widetilde{\pi}$ be a uniformly random permutation of $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$, with all four mutually independent. Consider running Algorithm 2 on the (random) input string $x = x(R, L, \pi)$ (see Equation (8)), and let y' be the string of inputs passed to ϕ'. After conditioning on Algorithm 2 not* **FAIL***ing, the total variation distance between y' and $y(R \odot L \odot \widetilde{C}, \widetilde{\pi})$ (see Equation 9) tends to 0 as $K \to \infty$.*

Proof. Let $x = x(R, L, \pi) = (e_1, \ell_1), \ldots (e_n, \ell_n)$, and recall that $\ell_i = \perp$ if $e_i \notin R$, and $\ell_i = L(e) \in \mathcal{L}$ if $e_i \in R$. When the algorithm succeeds, it feeds the elements $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$ to ϕ' in the order π', and for each $e \in R$ it designates precisely one copy of $(e, L(e))$ as active — namely, the $k_{\pi^{-1}(e)}$th copy of $(e, L(e))$ appearing in π'. We use $C'(e) \in \mathcal{C}$ to denote the index of this $k_{\pi^{-1}(e)}$th copy of $(e, L(e))$, and note that $C' : R \to \mathcal{C}$ is a function that depends on (R, L), π', π, and $\vec{k}$. In summary, when the algorithm succeeds we have $y' = y(R \odot L \odot C', \pi')$.

We now condition on (R, L) and π', and show (conditionally) that C' is a uniformly random function from R to $\mathcal{C}$. Since π is a uniformly random order on $\mathcal{E}$, it follows that the map $e \to k_{\pi^{-1}(e)}$ is a uniformly random perfect matching of $\mathcal{E}$ to $\{k_i\}_{i=1}^n$. Since $\{k_i\}_{i=1}^n$ consists of n i.i.d. draws from $[K]$, we conclude that $(k_{\pi^{-1}(e)})_{e \in \mathcal{E}}$ are i.i.d. draws uniformly from $[K]$. In other words, for each element $e \in R$ we independently activate a copy of $(e, L(e))$ uniformly at random — in particular the $k_{\pi^{-1}(e)}$th copy in order of appearance in π', where $k_{\pi^{-1}(e)} \sim [K]$. It follows that $(C'(e))_{e \in R}$ are i.i.d. uniform draws from $\mathcal{C}$, as needed.

Since π' is a uniformly random permutation of $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$ independent of (R, L), and $C' : R \to \mathcal{C}$ is uniformly random for each realization of π' and (R, L), it follows that $(R, L, C', \pi') \sim (R, L, \widetilde{C}|_R, \widetilde{\pi})$. Recall that $y' = y(R \odot L \odot C', \pi')$ when the algorithm succeeds. Since the probability of failure tends to 0 as $K \to \infty$ (Sublemma 10.2), we conclude that the total variation distance between y' and $y(R \odot L \odot \widetilde{C}|_R, \widetilde{\pi}) = y(R \odot L \odot \widetilde{C}, \widetilde{\pi})$ tends to 0 as $K \to \infty$, as needed. □

Sublemma 10.4. *Let (R, L) be a random labeled set for matroid $\mathcal{M}$ and labels $\mathcal{L}$, and let $C : \mathcal{E} \to \mathcal{C}$ be chosen uniformly at random independent of (R, L). If ϕ' is a β-balanced random-order CRM for the random set $R \odot L \odot C$ of elements of the matroid $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$, then Algorithm 2 instantiated with ϕ' is a $\widetilde{\beta}(K)$-balanced random-order LCRM for (R, L), where $\widetilde{\beta}(K)$ converges to β as $K \to \infty$.*

Proof. In the random order model, Algorithm 2 applied to the random labeled set (R, L) receives the string $x = x(R, L, \pi) = (e_1, \ell_1), \ldots, (e_n, \ell_n)$ as input, where π is a uniform random order independent of (R, L). Sublemmata 10.2 and 10.3 imply that the input to the parallel execution of ϕ' tends to $y = y(R \odot L \odot C, \widetilde{\pi})$ as $K \to \infty$, where $\widetilde{\pi}$ is a uniformly random permutation of $\mathcal{E} \times \mathcal{L} \times \mathcal{C}$ independent of (R, L) and C.

Recall that e_i is accepted by the algorithm if and only if a copy of (e_i, ℓ_i) is accepted by the parallel execution of ϕ'. Let $S \subseteq R$ be the set of elements accepted by the algorithm. Similarly, let $S' \subseteq \mathcal{E} \times \mathcal{L} \times \mathcal{C}$ be the set of elements accepted by the parallel execution of ϕ'. It follows that $e \in S$ and $L(e) = \ell$ if and only if $(e, \ell, c) \in S'$ for some $c \in \mathcal{C}$. Since the input string to ϕ' tends to $y = y(R \odot L \odot C, \widetilde{\pi})$ (in the sense of their total variation distance tending to 0), and ϕ' is β-balanced for $R \odot L \odot C$ in the random order model, it follows that there is $\widetilde{\beta}$ converging to β such that

$$\begin{aligned}
\mathbf{Pr}[(e, \ell, c) \in S'] &\geq \widetilde{\beta}\, \mathbf{Pr}[(e, \ell, c) \in R \odot L \odot C] \\
&= \widetilde{\beta}\, \mathbf{Pr}[e \in R \wedge L(e) = \ell \wedge C(e) = c] \\
&= \widetilde{\beta}\, \mathbf{Pr}[e \in R \wedge L(e) = \ell] \cdot \frac{1}{K}
\end{aligned}$$

simultaneously for all $(e, \ell, c) \in \mathcal{E} \times \mathcal{L} \times \mathcal{C}$. Now fix $e \in \mathcal{E}$ and $\ell \in \mathcal{L}$. Since different copies of (e, ℓ) are parallel in $\mathcal{M} \otimes \mathcal{L} \otimes \mathcal{C}$, and ϕ' accepts an independent set, it follows that the events $(e, \ell, c) \in S'$ are mutually exclusive. Therefore,

$$\begin{aligned}
\mathbf{Pr}[e \in S \wedge L(e) = \ell] &= \mathbf{Pr}[\exists c \in \mathcal{C} \text{ s.t. } (e, \ell, c) \in S'] \\
&= \sum_{c \in \mathcal{C}} \mathbf{Pr}[(e, \ell, c) \in S'] \\
&\geq \sum_{c \in \mathcal{C}} \widetilde{\beta}\, \mathbf{Pr}[e \in R \wedge L(e) = \ell] \cdot \frac{1}{K} \\
&= \widetilde{\beta}\, \mathbf{Pr}[e \in R \wedge L(e) = \ell],
\end{aligned}$$

as needed to show that Algorithm 2, instantiated with ϕ', is a $\widetilde{\beta}$-balanced LCRM for (R, L) in the random order model. □

Lemma 6.6 follows directly from Sublemma 10.4 and Observation 10.1.

11 Conclusion

This paper takes a journey from contention resolution to the matroid secretary problem and back. We obtain structural understanding of the distributions permitting offline contention resolution, and then exploit that structure to prove an equivalence between random-order contention resolution for correlated distributions and the matroid secretary problem. It is worth noting that our results are primarily information theoretic, pertaining to the power of online algorithms; i.e., we did not concern ourselves with computational efficiency of our reductions.

Our results indicate that the main challenge of the matroid secretary conjecture is resolving contention in the presence of a particular form of positive correlation. Specifically, it suffices to resolve contention, in the online random-order model, for the class of uncontentious distributions. These distributions are quite structured, as captured by the polyhedral characterization in Theorem 3.2. Said structure might lend just enough tractability to enable progress on the conjecture. More generally, identifying classes of distributions which exhibit nontrivial positive correlation, and yet permit online contention resolution, might stimulate the development of the needed contention resolution techniques.

Another conceptual takeaway from our result pertains to the importance of *cardinal* information in the matroid secretary problem, as compared to just *ordinal* information about the relative ordering of the weights. Ordinal algorithms for secretary problems were explored by [27, 43], though whether the *ordinal matroid secretary problem* is fundamentally more difficult than its classical (cardinal) counterpart remains open. Whereas our result does not definitively answer this question, it does indicate that the "hard part" of the matroid secretary problem is fundamentally ordinal in nature. Indeed, contention resolution involves no weights at all, and the set of improving elements can be determined online using just ordinal information. That said, our reduction from the secretary problem to contention resolution does use cardinal information,[13] leaving open the possibility that the cardinal matroid secretary problem is strictly easier than its ordinal counterpart. In other words, it is conceivable that some cardinal information is necessary to identify the pertinent ordinal information. Determining whether this is indeed the case is an intriguing open problem.

[13] The minimax reduction in Section 8 from secretary to prophet secretary heavily employs cardinal information. Furthermore, the reduction to labeled contention resolution in Section 9 requires associating each element with a discretization of its weight.

References

[1] M. Adamczyk and M. Włodarczyk. Random order contention resolution schemes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 790–801. IEEE, 2018.

[2] S. Agrawal, Y. Ding, A. Saberi, and Y. Ye. Price of correlations in stochastic optimization. *Operations Research*, 60(1):150–162, 2012.

[3] Y. Azar, A. Chiplunkar, and H. 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.

[4] M. Babaioff, N. Immorlica, and R. 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.

[5] C. Bechtel and S. Dughmi. Delegated stochastic probing. In *Proceedings of the Conference on Innovations in Theoretical Computer Science (ITCS)*, 2021.

[6] C. Bechtel, S. Dughmi, and N. Patel. Delegated pandora's box. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 666–693, 2022.

[7] Y. Cai and A. Oikonomou. On Simple Mechanisms for Dependent Items. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 242–262, New York, NY, USA, July 2021. Association for Computing Machinery.

[8] I. Caragiannis, N. Gravin, P. Lu, and Z. Wang. Relaxing the independence assumption in sequential posted pricing, prophet inequality, and random bipartite matching. In *International Conference on Web and Internet Economics*, pages 131–148. Springer, 2021.

[9] C. Chekuri and V. Livanos. On submodular prophet inequalities and correlation gap. *Theoretical Computer Science*, 1019:114814, 2024.

[10] C. Chekuri, J. Vondrák, and R. 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.

[11] C. Chekuri, J. Vondrák, and R. Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6): 1831–1879, 2014.

[12] M. Dinitz. Recent advances on the matroid secretary problem. *ACM SIGACT News*, 44(2): 126–142, 2013.

[13] S. Dughmi. The outer limits of contention resolution on matroids and connections to the secretary problem. In *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.

[14] S. Dughmi. Matroid secretary is equivalent to contention resolution. In *Proceedings of the Conference on Innovations in Theoretical Computer Science (ITCS)*, 2022.

[15] S. Dughmi, Y. H. Kalayci, and N. Patel. On sparsification of stochastic packing problems. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*, 2023.

[16] S. Dughmi, Y. H. Kalayci, and N. Patel. Limitations of Stochastic Selection with Pairwise Independent Priors. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC)*, 2024.

[17] P. Dütting and R. Kleinberg. Polymatroid prophet inequalities. In *Algorithms-ESA 2015*, pages 437–449. Springer, 2015.

[18] P. Dutting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3): 540–582, 2020.

[19] E. B. Dynkin. The optimum choice of the instant for stopping a markov process. *Soviet Mathematics*, 4:627–629, 1963.

[20] J. Edmonds. Minimum partition of a matroid into independent subsets. *J. Res. Nat. Bur. Standards Sect. B*, 69:67–72, 1965.

[21] S. Ehsani, M. Hajiaghayi, T. Kesselheim, and S. 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.

[22] H. Esfandiari, M. Hajiaghayi, V. Liaghat, and M. Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.

[23] M. Feldman, O. Svensson, and R. Zenklusen. A simple o (log log (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.

[24] M. Feldman, O. Svensson, and R. 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.

[25] S. O. Gharan and J. Vondrák. On variants of the matroid secretary problem. *Algorithmica*, 67(4):472–497, 2013.

[26] T. P. Hill and R. P. Kertz. A survey of prophet inequalities in optimal stopping theory. *Contemp. Math*, 125:191–207, 1992.

[27] M. Hoefer and B. Kodric. Combinatorial secretary problems with ordinal information. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[28] W. Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

[29] N. Immorlica, S. Singla, and B. Waggoner. Prophet inequalities with linear correlations and augmentations. In *Proceedings of the ACM Conference on Economics and Computation (EC)*. ACM, 2020.

[30] D. R. Karger. Random sampling and greedy sparsification for matroid optimization problems. *Mathematical Programming*, 82(1-2):41–81, 1998.

[31] R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136. ACM, 2012.

[32] R. D. Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, volume 5, pages 630–631, 2005.

[33] U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83(4):745–747, 1977.

[34] U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4:197–266, 1978.

[35] O. Lachish. O (log log rank) competitive ratio for the matroid secretary problem. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 326–335. IEEE, 2014.

[36] E. Lee and S. 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.

[37] V. Livanos, K. Patton, and S. Singla. Improved Mechanisms and Prophet Inequalities for Graphical Dependencies. In *Proceedings of the 25th ACM Conference on Economics and Computation*. Association for Computing Machinery, 2024.

[38] J. G. Oxley. *Matroid Theory*. Oxford University Press, 1992.

[39] F. Qiu and S. Singla. Submodular Dominance and Applications. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022)*, pages 44:1–44:21. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022.

[40] Y. Rinott, E. 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.

[41] A. Rubinstein. 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.

[42] E. Samuel-Cahn. Prophet inequalities for bounded negatively dependent random variables. *Statistics & probability letters*, 12(3):213–216, 1991.

[43] J. A. Soto, A. Turkieltaub, and V. Verdugo. Strong algorithms for the ordinal matroid secretary problem. *Mathematics of Operations Research*, 2021.

[44] D. J. Welsh. *Matroid theory*. Courier Corporation, 2010.

A Relation of Lemma 4.4 to Prior Work

It is tempting to conjecture that contention resolution algorithms from prior work can be shown to succeed for improving elements, and in doing so provide an alternate route to Lemma 4.4. We now present evidence that this is not the case.

First, we show that there exists no $\omega(\frac{1}{n})$-balanced online CRS in the worst-case arrival model, even when both the order and the distribution of improving elements are known to the algorithm. Therefore, 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. [24], which hold for an (unknown) worst-case arrival order.

Proposition A.1. *Fix a rank one matroid on n elements, equipped with distinct weights on the elements. For every $p \in (0, 1)$ and $\alpha > \frac{1}{n}$, there is no α-balanced online CRM for the improving elements with parameter p in the worst-case arrivals model. This holds even when the arrival order is known to the algorithm.*

Proof. Let $\{1, \ldots, n\}$ denote the ground set of of a rank one 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 $\{1, \ldots, k\}$ for $k = 0, \ldots, n$. In the special case of $p = 1/2$, the distribution of R is as described in Example 3.9. In general, $\mathbf{Pr}[R = \{1, \ldots, k\}] = p(1 - p)^k$. The random set R is p-uncontentious, as shown by Lemma 4.4. Concretely, the offline CRM $\phi(\{1, \ldots, k\}) = k$ is p-balanced.

Now suppose that the elements are known to arrive online in the order $1, 2, 3, \ldots, n$, and consider an α-balanced online CRM for some $\alpha \in [0, 1]$. 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 α. Formally, $\mathbf{Pr}[i \in T | i \in R] \geq \alpha$.

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 $1, \ldots, i$, and must make its decision on the spot, it cannot distinguish between different sets of the form $R = \{1, \ldots, k\}$ 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 α in every realization of R which includes i. Formally, $\alpha \leq \mathbf{Pr}[i \in T | i \in R] = \mathbf{Pr}[i \in T | R = \{1, \ldots, k\}]$ for every $k \geq i$.

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

It is instructive to examine how the algorithm of Feldman et al. [24] fails in the special case of the rank one matroid, 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 [24] defines a sequence of nested flats $\emptyset \subset F_1 \subset F_2 \subset \ldots \subset F_k$, and runs the greedy online algorithm on each contracted submatroid F_i/F_{i-1}. The rank one matroid contains only a single non-empty flat, containing all elements. Therefore, the algorithm of [24] reduces merely to the naive greedy online algorithm which simply selects the first active element it encounters, which in the case of a uniform arrival order is a uniformly random active element.

Now, let $[n] = \{1, \ldots, n\}$ denote the ground set of the rank one matroid listed in decreasing order of weight, and consider the distribution of improving elements R with parameter $p = 1/2$ as described in Example 3.9. When element k is improving, so are elements $1, \ldots, k - 1$. Since the algorithm selects a uniformly random active element, k is selected with probability at most $\frac{\mathbf{Pr}[k \in R]}{k}$. It is also easy to show that the algorithm suffers a similar fate for any other choice of p.

One might be tempted to employ other tricks, such as for example "canceling" each active element 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. Therefore, it cannot select element k with probability exceeding $\frac{\mathbf{Pr}[k \in R]}{k}$.

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. [24], 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 rank one 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 $\frac{\mathbf{Pr}[k \in R]}{k}$. 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.

B Proof of Sublemma 8.1

Consider the matroid secretary problem on matroid $\mathcal{M} = (\mathcal{E}, \mathcal{I})$ and arbitrary (unknown) weights $w \in \mathbb{R}_+^{\mathcal{E}}$. Denote $r^* = \mathbf{rank}(\mathcal{M})$, $n = |\mathcal{E}|$, $v^* = \mathbf{rank}_w(\mathcal{M})$, and let $T^* \in \mathcal{I}$ be a maximum-weight independent set (i.e., with $w(T^*) = v^*$). Consider the algorithm which, with probability $\frac{1}{2}$, runs the $\frac{1}{e}$-competitive algorithm for the single-choice secretary problem with weights w, and otherwise runs the following reduction to a normalized and discretized instance on a restriction of $\mathcal{M}$.

- Sample roughly half the elements: Let $k \sim \mathbf{Binom}\left(n, \frac{1}{2}\right)$, and observe the weights of the first k elements S in the arrival order π, without accepting any.
- Let $r = \mathbf{rank}^{\mathcal{M}}(S)$ and $v = \mathbf{rank}_w^{\mathcal{M}}(S)$ be rank and weighted rank, respectively, of the sample.
- Let $\overline{S} = \mathcal{E} \setminus S$ be the remaining (unsampled) elements.
- Define transformed weights for the unsampled elements $e \in \overline{S}$ as follows: $\widehat{w}_e = 0$ if $w_e < \frac{v}{32r}$, otherwise $\widehat{w}_e$ is the result of rounding down $\frac{w_e}{v}$ to the nearest power of 2.
- To select an independent subset of the remaining elements $\overline{S}$, invoke a matroid secretary algorithm for the remaining matroid $\mathcal{M}|\overline{S}$ with weights $\widehat{w}$.

It is clear that the elements in $\overline{S}$ arrive in uniformly random order after S. It is also clear that the transformed weights $\{\widehat{w}_e\}_{e \in \overline{S}}$ can be computed online from the original weights $\{w_e\}_{e \in \overline{S}}$, as well as the rank r and weighted rank v of the sample. It follows that, for each realization of the random sample S, this is indeed a valid reduction to the matroid secretary problem on $\mathcal{M}|\overline{S}$ and $\widehat{w}$. The following relationship between the original and transformed weights is easy to see, and will be useful for the remainder of this proof.

$$\frac{w_e}{2v} - \frac{1}{64r} = \frac{w_e - v/32r}{2v} \leq \widehat{w}_e \leq \frac{w_e}{v} \text{ for all } e \in \overline{S} \tag{10}$$

Observe that if there is an element with weight exceeding $\frac{v^*}{16}$, then running single-choice secretary algorithm with probability $\frac{1}{2}$ guarantees that we obtain a competitive ratio of at least $\frac{1}{2} \cdot \frac{1}{e} \cdot \frac{1}{16} = \frac{1}{32e} > \frac{1}{256}$. Therefore, we henceforth assume that $w_e \leq \frac{v^*}{16}$ for all $e \in \mathcal{E}$ and analyze the above reduction.

We first show that v is within a constant of v^* with constant probability. It is immediate that v is upper-bounded by v^*. The lower-bound follows from a series of elementary calculations, using the fact that each element of T^*, with total weight $w(T^*) = v^*$, is in S independently with probability $\frac{1}{2}$.

$$\begin{aligned}
\mathbf{Pr}\left[v < \frac{v^*}{4}\right] &\leq \mathbf{Pr}\left[w(S \cap T^*) < \frac{v^*}{4}\right] && \text{(Since } v > w(S \cap T^*) \text{)} \\
&\leq \exp\left(-\frac{2(v^*/4)^2}{\sum_{e \in T^*} w_e^2}\right) && \text{(Hoeffing's Inequality)} \\
&= \exp\left(-\frac{(v^*)^2}{8\sum_{e \in T^*} w_e^2}\right) \\
&\leq \exp\left(-\frac{(v^*)^2}{8(\max_{e \in T^*} w_e)(\sum_{e \in T^*} w_e)}\right) && \text{(Holder's inequality)} \\
&= \exp\left(-\frac{(v^*)^2}{8(\max_{e \in T^*} w_e) \cdot v^*}\right) \\
&\leq \exp(-2) && \text{(Since } \max_e w_e \leq v^*/16\text{)}
\end{aligned}$$

It follows that v is in $[\frac{v^*}{4}, v^*]$ with probability at least $1 - \frac{1}{e^2}$.

An even simpler argument shows that r is within a constant of r^*. By our assumption that $w_e \leq v^*/16$, it follows that $r^* \geq 16$. A simple application of the Hoeffding bound, akin to that above, implies that r is in $[\frac{r^*}{4}, r^*]$ with probability at least $1 - \exp\left(-\frac{2(r^*/4)^2}{r^*}\right) = 1 - \frac{1}{e^2}$.

Now denote $\overline{v} = \mathbf{rank}_w^{\mathcal{M}}(\overline{S})$ and $\overline{r} = \mathbf{rank}^{\mathcal{M}}(\overline{S})$, and observe that v and $\overline{v}$ are identically distributed, and the same is true for r and $\overline{r}$. It follows from the union bound that v and $\overline{v}$ are in $[\frac{v^*}{4}, v^*]$, and moreover r and $\overline{r}$ are in $[\frac{r^*}{4}, r^*]$, with probability at least $1 - \frac{4}{e^2} > \frac{1}{4}$. In this event, symmetry implies that $v \geq \overline{v}$ with probability at least $\frac{1}{2}$. Therefore, the following hold with probability at least $\frac{1}{8}$:

$$\frac{v^*}{4} \leq \overline{v} \leq v \leq v^* \tag{11}$$

and

$$r, \overline{r} \in \left[\frac{r^*}{4}, r^*\right]. \tag{12}$$

We now condition on (11) and (12), which hold with probability at least $\frac{1}{8}$, and show that the matroid secretary instance $(\mathcal{M}|\overline{S}, \widehat{w})$ is normalized and discretized, and moreover that our reduction to this instance is approximation preserving up to a constant.

Normalization follows easily from (10), (11) and (12):

$$\mathbf{rank}_{\widehat{w}}(\mathcal{M}|\overline{S}) \leq \frac{1}{v}\mathbf{rank}_w(\mathcal{M}|\overline{S}) = \frac{\overline{v}}{v} \leq 1.$$

and

$$\begin{aligned}
\mathbf{rank}_{\widehat{w}}(\mathcal{M}|\overline{S}) &\geq \frac{1}{2v}\mathbf{rank}_w(\mathcal{M}|\overline{S}) - \frac{1}{64r}\mathbf{rank}(\mathcal{M}|\overline{S}) \\
&= \frac{\overline{v}}{2v} - \frac{\overline{r}}{64r} \\
&\geq \frac{v^*/4}{2v^*} - \frac{r^*}{64r^*/4} = \frac{1}{16}
\end{aligned}$$

For discretization, recall that by definition each transformed weight $\widehat{w}_e$ for $e \in \overline{S}$ is either zero or the result of rounding down w_e/v to a power of 2, for $v/32r \leq w_e \leq \mathbf{rank}_w(\mathcal{M}|\overline{S}) = \overline{v} \leq v$. It follows that a non-zero $\widehat{w}_e$ is a power of 2 between $1/64r$ and 1. Since $r \leq r^* \leq 4\overline{r}$, a non-zero $\widehat{w}_e$ is a power of 2 between $\frac{1}{256\overline{r}}$ and 1.

For the approximation, consider any c-competitive solution $\widehat{T}$ for the instance $(\mathcal{M}|\overline{S}, \widehat{w})$, and let T' be an optimal solution for the instance $(\mathcal{M}|\overline{S}, w)$. We can show that $\widehat{T}$ is $\frac{c}{16}$-competitive for the original instance $(\mathcal{M}, w)$, using (10), (11), and (12):

$$\begin{aligned} w(\widehat{T}) &\geq v \cdot \widehat{w}(\widehat{T}) \\ &\geq v \cdot c \cdot \widehat{w}(T') \\ &\geq v \cdot c \cdot \left(\frac{w(T')}{2v} - \frac{|T'|}{64r}\right) \\ &\geq v \cdot c \cdot \left(\frac{w(T')}{2v} - \frac{\overline{r}}{64r}\right) \\ &\geq v \cdot c \cdot \left(\frac{w(T')}{2v} - \frac{1}{16}\right) \\ &= c \cdot \left(\frac{\overline{v}}{2} - \frac{v}{16}\right) \\ &\geq c \cdot \left(\frac{v^*}{8} - \frac{v^*}{16}\right) \\ &= \frac{c}{16} v^* \end{aligned}$$

Recall that we run the reduction (rather than the single-choice secretary algorithm) with probability $\frac{1}{2}$. Also recall that we conditioned on (11) and (12), an event which holds with probability at least $\frac{1}{8}$. Therefore, the loss in the approximation ratio is no worse than $\frac{1}{16} \times \frac{1}{8} \times \frac{1}{2} = \frac{1}{256}$.

C Only Active Elements Arrive in Uniformly Random Order

We now consider a semi-random model of online arrivals, where the relative order of active elements is uniformly random, but the order is otherwise arbitrary. We will show that there exists a β-uncontentious distribution for the rank one matroid, where β can be made arbitrarily close to 1, admitting no constant-balanced CRM in this semi-random arrival model. In fact, we will show this to be true even when active elements arrive first (in uniformly random order), followed by all inactive elements in an arbitrary order.

Let ϵ be such that $0 < \epsilon \leq \frac{1}{2}$, and let n and m be positive integers. We will later choose these parameters to enable our impossibility result. Let $\mathcal{E}_i$ be a class of $N_i = n^{m-i}$ elements for each $i = 0, \dots m$, and let $\mathcal{M}$ be the rank one matroid on $\mathcal{E} = \cup_{i=0}^m \mathcal{E}_i$. Denote $\delta = \frac{\epsilon^2}{n^m}$ and draw the set $R \subseteq \mathcal{E}$ of active elements as follows:

- Let k be a draw from the geometric distribution with parameter $1-\delta$, and let $p = p(k) = \frac{\epsilon}{n^{m-k}}$
- Let R include each element of $\cup_{i \leq k} \mathcal{E}_i$ independently with probability p.

In the subsequent analysis, for a quantity $x = x(\epsilon)$ we say $x \to y$, if $\lim_{\epsilon \to 0} x = y$. We also say a probability $p = p(\epsilon)$ approaches q if $\lim_{\epsilon \to 0} p \geq q$. We say an event holds with high probability if its probability approaches 1.

We argue that R is β-uncontentious, for $\beta \to 1$, by considering the following offline CRM: For k in the above sampling procedure, accept an arbitrary element in $R \cap \mathcal{E}_k$, if any.[14] Observe that each element $e \in \mathcal{E}_i$ is accepted by the CRM with probability at least $\mathbf{Pr}[k=i] \cdot \frac{\epsilon}{N_i} \cdot (1-\epsilon/N_i)^{N_i - 1} \geq \mathbf{Pr}[k=i] \frac{\epsilon}{e^\epsilon N_i}$, and is active (i.e., in R) with probability at most $\mathbf{Pr}[k=i]\frac{\epsilon}{N_i} + \mathbf{Pr}[k>i]$. By definition of the geometric distribution, and using the fact that $\delta \leq \frac{1}{2}$, we get $\mathbf{Pr}[k>i] = \frac{\delta}{1-\delta}\,\mathbf{Pr}[k=i] \leq 2\delta\,\mathbf{Pr}[k=i]$. Bounding $N_i \leq n^m$, this yields $\beta = \frac{1}{e^\epsilon(1+2\epsilon)}$ as needed.

We also argue that any CRM which is α-balanced for R must, in the event that $k=i$ and at least one element of $\mathcal{E}_i$ is active, accept an element of $\mathcal{E}_i$ with conditional probability approaching α. First, for an element $e \in \mathcal{E}_i$, we show that $k=i$ with high conditional probability given e is active.

$$\begin{aligned}
\mathbf{Pr}[k=i|e \in R] &= \frac{\mathbf{Pr}[k=i]\,\mathbf{Pr}[e \in R|k=i]}{\mathbf{Pr}[e \in R]} \\
&\geq \frac{\epsilon\,\mathbf{Pr}[k=i]/N_i}{\epsilon\,\mathbf{Pr}[k=i]/N_i + \mathbf{Pr}[k>i]} \\
&\geq \frac{\epsilon\,\mathbf{Pr}[k=i]/N_i}{\epsilon\,\mathbf{Pr}[k=i]/N_i + 2\delta\,\mathbf{Pr}[k=i]} \\
&= \frac{\epsilon}{\epsilon + 2\delta N_i} \\
&\geq \frac{\epsilon}{\epsilon + 2\delta n^m} \\
&= 1/(1+2\epsilon) \to 1
\end{aligned}$$

It follows that $\mathbf{Pr}[e \text{ accepted}|e \in R, k=i] \geq \alpha'$ for some $\alpha' \to \alpha$. Note also that, when e is active and $k=i$, there are no other active elements in $\mathcal{E}_i$ with high probability. In other words, the events $f \in R | k=i$ for elements $f \in \mathcal{E}_i$ tend to disjointness as $\epsilon \to 0$. The initial claim follows.

Now fix an online CRM ϕ with balance ratio α for R in our semi-random arrival model. Suppose that the active elements $R = R(k)$ are presented to ϕ in a uniformly random random order $\pi = (e_1, \ldots, e_{|R|})$, followed by all the inactive elements in an arbitrary order. Note that ϕ does not know k a-priori, but can only glean information about it from observing R. In the case that $k=0$, R is empty with probability $(1-\frac{\epsilon}{n^m})^{n^m} \approx 1-\epsilon$, and consists of a single element in $\mathcal{E}_0$ with probability $n^m \frac{\epsilon}{n^m}(1-\frac{\epsilon}{n^m})^{n^m-1} \approx \epsilon(1-\epsilon)$. The argument in the previous paragraph implies, therefore, that ϕ must accept the first active element (in $\mathcal{E}_0$, if any) with probability approaching α when $k=0$. Now consider the case of $k=1$: R consists of $\mathbf{Binom}(n^m, \frac{\epsilon}{n^{m-1}}) \approx \frac{\epsilon}{n^{m-1}} n^m = \epsilon n$ elements of $\mathcal{E}_0$, and $\mathbf{Binom}(n^{m-1}, \frac{\epsilon}{n^{m-1}}) \approx O(1)$ elements of $\mathcal{E}_1$. More formally, if we choose $n = \omega(\frac{1}{\epsilon})$, Chernoff bounds imply that R consists of $\Omega(\epsilon n)$ elements of $\mathcal{E}_0$ and $O(1)$ elements of $\mathcal{E}_1$ with high probability. Therefore, with high probability the first element e_1 in the sequence will be in $\mathcal{E}_0$, and by our previous argument for the case of $k=0$ — since the ϕ cannot distinguish between $k=0$ and $k=1$ at the beginning of the sequence — it must be accepted with probability approaching α. Moreover, by our previous paragraph if there is an active element in $\mathcal{E}_1$ then the first such element must be accepted by ϕ with probability approaching α.

This pattern continues inductively. Consider the case $k=i$, for an arbitrary i. Let e'_j be the first active element in $\mathcal{E}_j$ appearing in the online order, if any. With high probability, e'_j exists for all $j<i$, though e'_i may not (in the event there are no active elements in $\mathcal{E}_i$). Notice that the relative proportion of $\mathcal{E}_j \cap R$ to $\mathcal{E}_{j+1} \cap R$ is $\Omega(n)$ with high probability (i.e., with probability approaching 1 as ϵ approaches 0). We can therefore choose m as an increasing function of $\frac{1}{\epsilon}$ such that, with high probability, e'_j precedes e'_{j+1} in π simultaneously for all $j = 1, \ldots, m-1$. Also notice

[14] Note that, in the offline model, we can assume without loss of generality that the CRM has access to k: it can simply sample the distribution $k|R$.

that, for $j, \ell \leq i-1$, the (distribution of) the relative size of $\mathcal{E}_j \cap R$ to that of $\mathcal{E}_\ell \cap R$ is the same whether $k = i$ or $k = i-1$; the principle of deferred decisions then implies that ϕ cannot distinguish $k = i$ from $k = i-1$ until it first encounters e'_i. It follows that ϕ accepts each of $e'_1, \ldots, e'_{i-1}$ with probability approaching α by induction. Moreover, as previously argued it must accept e'_i, in the event it exists, with probability approaching α. Taking $i = m$, it follows that $\alpha = O(1/m)$. Since m grows without bound, this proves that no absolute constant balance ratio is possible.