Stationary Online Contention Resolution Schemes:

Theory and Applications to Bayesian Online Resource Allocation

Mohammad Reza Aminian
The University of Chicago, Booth School of Business, Chicago, IL, maminian@chicagobooth.edu

Rad Niazadeh
The University of Chicago, Booth School of Business, Chicago, IL, rad.niazadeh@chicagobooth.edu

Pranav Nuti
The University of Chicago, Booth School of Business, Chicago, IL, pranav.nuti@chicagobooth.edu

Abstract.

In several modern applications involving online platforms and digital marketplaces, the data-driven allocation of supply to demand involves new difficulties such as accommodating combinatorial feasibility environments or reusable resources (or both). Modeling these settings as instances of a Bayesian online allocation problem raises a natural question: Is there a universal technical framework for designing algorithmic solutions across all these settings? Online contention resolution schemes (OCRSs) aim to provide such solutions by converting ex-ante relaxations into approximately optimal online feasible policies. Despite their importance—and relative universality—existing constructions of (near-)optimal OCRSs are often technically challenging. Many constructions rely on indirect reductions to prophet inequalities and LP duality, resulting in algorithms that are ad hoc and difficult to interpret or implement. Moreover, it is unclear whether such vanilla OCRSs can accommodate new paradigms such as resource reusability.

In this paper, we introduce *stationary online contention resolution schemes (S-OCRSs)*, a permutation-invariant class of OCRSs in which the distribution of the selected feasible set is independent of the arrival order. We establish a distributional characterization of S-OCRSs and provide a universal online implementation, thereby reducing the design of online policies to the construction of suitable distributions over feasible sets. This reduction applies to arbitrary, potentially complex feasibility environments. We further show that S-OCRSs are precisely the tool needed to address reusable resource allocation problems. We then develop a general 'maximum-entropy-based' approach for constructing and analyzing S-OCRSs across a variety of settings. Together, these results provide a universal and principled technical framework for designing simple and potentially improved OCRSs, yielding approximately optimal online allocation algorithms for a wide range of modern Bayesian online allocation problems.

We demonstrate the power of our new framework across several canonical feasibility environments. For bipartite matchings, we obtain an improved $(3-\sqrt{5})/2$-selectable OCRS, attaining the independence benchmark conjectured to be optimal for OCRSs. For knapsack constraints, we obtain a $1/4$-selectable S-OCRS. We also obtain a $1-\sqrt{2/(\pi k)}+O(1/k)$-selectable S-OCRS for k-uniform matroids; a simple, explicit $1/2$-selectable S-OCRS for weakly Rayleigh matroids, including all $\mathbb{C}$-representable matroids such as graphic and laminar matroids; a $1-\sqrt{2\log L/k}-O(1/\sqrt{k})$-selectable S-OCRS for rank-L hypergraph matchings with capacity k (also known as the network revenue management setting); and a $\Theta(1/\sqrt{N})$-selectable S-OCRS for general downward-closed constraints on sets of size N. These latter guarantees match the best-known bounds for OCRSs via simple and principled algorithms, in contrast to existing algorithms that are either ad-hoc, complex, or non-explicit. All of the aforementioned S-OCRS guarantees for different feasibility environments extend to approximation factors for the Bayesian online allocation problems of reusable products in the same feasibility environments.

1. Introduction

In several modern applications involving online platforms and digital marketplaces, the data-driven, real-time allocation of supply to demand gives rise to new difficulties that require novel technical treatments. These challenges often stem from constraints on feasible allocations—typically imposed by marketplace rules or by compatibility restrictions between the demand and supply sides of the platform—or from specific properties of the resources being allocated. One class of examples comes from *combinatorial feasibility environments*, including matching, knapsack, network revenue management, and complex matroid constraints. Such environments arise in applications such as two-sided matching platforms, crowdsourcing marketplaces, and revenue management over complex networks. Another class of examples involves simple—and, in some cases, complex and combinatorial—feasibility environments with *reusable* (or *recurring*) resources, where each resource unit is occupied for some duration before becoming available again to serve future requests. Models with reusable resources arise prominently in online revenue management and assortment optimization (Rusmevichientong et al., 2023, 2020; Feng et al., 2024; Goyal et al., 2025; Feng et al., 2025; Ekbatani et al., 2025), and are motivated by applications including cloud and edge computing platforms that allocate compute capacity or GPU memory to jobs, digital rental marketplaces, volunteer matching platforms, and, more generally, service systems in which capacity cycles between busy and idle states.

Modeling all these settings as instances of a Bayesian online resource allocation problem raises a natural question: Is there a *universal* technical framework for designing algorithmic solutions across all these settings? To address this question, a recurring algorithmic theme in operations research and related work on Bayesian online decision making—particularly Bayesian online resource allocation—is to approach problems such as those described above through the following (potentially universal) design recipe. First, one computes an *ex-ante* feasible plan that satisfies the relevant feasibility constraints—arising from the combinatorial structure of the allocations, the nature of the resources, or both—only in expectation. One then converts this plan into an online policy that maintains *ex-post* feasibility in every realization while incurring only a small loss, relative to the ex-ante plan, in the expected reward of the allocation. This ex-ante-to-ex-post conversion is central to problems in dynamic resource allocation, revenue management, and mechanism design, where decisions must be made sequentially subject to hard feasibility constraints.[1]

Online contention resolution schemes (OCRSs)—studied by Alaei (2014) for the special case of k-uniform matroids and later introduced for general feasibility environments by Feldman et al. (2016, 2021)—provide a principled framework for carrying out this conversion. Fix a downward-closed feasibility family $\mathcal{F} \subseteq 2^E$ over a ground set E, and let $\mathcal{P} := \mathrm{conv}\{\mathbf{1}_S \in \mathcal{F}\}$ denote its associated polytope. Given a fractional point $\boldsymbol{x} \in \mathcal{P}$, each element $e \in E$ is independently "active" with probability $\mathbb{P}[e \text{ active}] = x_e$, and the activations are

[1] It also underlies several well-studied problem classes at the intersection of operations research, economics, and computer science, including prophet inequalities (Krengel and Sucheston, 1978; Samuel-Cahn, 1984; Kleinberg and Weinberg, 2012; Dutting et al., 2020; Ezra et al., 2022) (see Lucier (2017) for a comprehensive survey), Bayesian online resource allocation (Alaei et al., 2012; Alaei, 2014; Feng et al., 2022; Anari et al., 2019), sequential posted-price mechanisms (Chawla et al., 2010; Correa et al., 2017; Yan, 2011; Pollner et al., 2022), and, more broadly, dynamic resource-constrained reward collection models (Balseiro et al., 2024).

revealed online in an arrival order that may be adversarial. An OCRS is an online policy that irrevocably selects a feasible set $S \in \mathcal{F}$ of active elements while seeking to guarantee sufficient *selectability* for every element. Specifically, the scheme is α-selectable if, for every arrival order and every element e,

$$\mathbb{P}[e \in S \mid e \text{ is active}] \;\geq\; \alpha, \quad \text{or, equivalently,} \quad \mathbb{P}[e \in S] \;\geq\; \alpha x_e.$$

An OCRS can therefore be viewed as a black-box online rounding primitive that preserves each marginal probability up to a common factor. This is precisely the type of guarantee needed to convert an ex-ante solution—for example, an LP-based allocation rule for a resource allocation problem that is feasible only in expectation—into an implementable online allocation.[2] Owing to their versatility, OCRSs have served as a modular tool in Bayesian online resource allocation and mechanism design since their introduction.

Designing OCRSs, however, is often technically delicate, and two major issues arise when one attempts to use them as the algorithmic engine of the universal framework envisioned above:

(I) The first concerns the design process itself. A fundamental obstruction is the intrinsic asymmetry of online information: elements arriving early must compete against an unknown future, whereas elements arriving late face a system that has already committed some of its capacity. Consequently, many OCRS designs and analyses depend essentially on the arrival order, and the resulting algorithms can be complex, ad hoc, and specific to the underlying feasibility environment—rather than following a universal design template. Moreover, for certain feasibility environments, such as matroids, the best-known guarantees are obtained only indirectly through computationally intensive implicit algorithms derived via reductions to prophet inequalities with respect to ex-ante relaxations (Lee and Singla, 2018).[3] Although such reductions are powerful, they can obscure the eventual rounding rule and make it difficult to characterize the distribution over feasible outcomes that the algorithm implicitly implements. From an operations research perspective, this lack of transparency can limit interpretability when the OCRS is used as a building block within a larger online control policy.

(II) The second issue concerns the reach of the OCRS abstraction itself. Standard OCRSs are inherently *one-shot*: each element arrives exactly once, in some possibly adversarial order, and, once selected, occupies its share of the resources until the end of the horizon. In the reusable environments highlighted at the outset, by contrast, allocated capacity is released over time as new requests continue to arrive. Thus, there is no single arrival order to speak of: the pattern of contention among requests keeps changing as resource units are occupied and released. It is therefore unclear whether vanilla OCRSs, whose selectability guarantees are stated with respect to a single arrival of each element, can accommodate resource reusability.

[2] This perspective also explains the close connection between OCRSs and prophet-inequality-style guarantees obtained through ex-ante relaxations (Lee and Singla, 2018).

[3] These reductions and implicit constructions rely on duality and minimax arguments, together with algorithms for solving the resulting minimax problems. They typically produce a randomized online policy by applying first-order optimization procedures, such as multiplicative-weights updates, over the *space of online policies*; see Lee and Singla (2018) for details.

1.1. Our conceptual contribution: "Stationary OCRS"

This paper explores a novel design principle for OCRSs by introducing and studying a symmetry requirement that directly addresses the "first vs. last" asymmetry (and, as we will see, the reusability challenge): *permutation invariance* of the output distribution, that is, the distribution of the selected set should be independent of the online arrival order of the elements. Formally, for every fixed $\boldsymbol{x} \in \mathcal{P}$, we require the existence of a single distribution $\mu^{(\boldsymbol{x})}$ over feasible sets $S \in \mathcal{F}$ such that, under every arrival order, the selected set S is distributed according to $\mu^{(\boldsymbol{x})}$. We call an OCRS satisfying this symmetry requirement a *stationary OCRS (S-OCRS)*. Stationarity can be interpreted as requiring the scheme to "treat every element as if it were the last".

A natural question is why stationary OCRSs merit study. Beyond their conceptual appeal, stationarity directly addresses the two issues discussed above. First, we posit that even though imposing symmetry through stationarity restricts the class of admissible policies, it introduces additional structure that facilitates the design of simple—and, in some cases, improved or (near-)optimal—OCRS algorithms that arise from a universal design recipe. Second, stationarity is precisely the property needed to address the reusable-resource settings introduced at the outset. Recent work uses OCRSs as admission-control subroutines to convert solutions to ex-ante LP benchmarks into ex-post feasible and approximately optimal online policies in these settings (Feng et al., 2022, 2024; Baek and Ma, 2022). In such models, as noted earlier, a single resource may repeatedly encounter dynamically reordered contention processes, making permutation-invariant admission-control primitives a natural choice; see Figure 1 for an illustration of this recurring pattern of contention.[4]

Our perspective gives rise to a set of fundamental technical questions, which we address in this paper:

(i) Can we design stationary OCRSs with constant-factor selectability for natural feasibility environments, such as matching, matroid, and knapsack constraints? In important special cases, does imposing stationarity lead to simpler or even improved OCRSs?

(ii) Is there a universal recipe for designing stationary OCRSs across a broad class of feasibility environments? Can the resulting schemes be made explicit—that is, can they yield concrete, implementable online algorithms rather than merely indirect existence proofs?

(iii) Can stationary OCRSs help address resource reusability? Specifically, is there a universal recipe for designing policies for the Bayesian online allocation of reusable products (under stationary request arrivals) in any feasibility environment (e.g., k-uniform matroids, matchings, or general matroids)—for instance, by reducing the problem to the design of stationary OCRSs in the same feasibility environment?

[4] Besides the main reasons we study S-OCRS, we anticipate our framework has further applications. For instance, another consideration that recurs in the theory and applications of Bayesian online resource allocation is fairness and robustness when agents can influence *timing*. Order-dependent rules can create incentives for groups of friendly agents to manipulate their arrival times, because their joint distribution of allocations is a function of the arrival order. In contrast, a stationary OCRS removes any advantage that is purely due to the arrival times. This is in spirit related to time-strategy-proofness or order incentive compatibility notions studied, e.g., in the context of secretary problems and online auctions (Buchbinder et al., 2009; Hajiaghayi et al., 2004).

1.2. Our main technical contributions

We introduce a new conceptual and technical framework for designing and analyzing S-OCRSs and OCRSs. Our results provide affirmative answers to all the questions posed above, with a particular emphasis on explicit distributions and explicit (near-)optimal algorithms—culminating in a formal reduction from Bayesian online allocation with reusable products to the design of S-OCRSs.

Distributional LP characterization & meta-algorithm. We begin by giving a precise characterization of stationary OCRSs via a linear feasibility problem over distributions on feasible sets, which we call the *stationary OCRS LP*. In addition to the usual marginal (selectability) constraints—requiring $\mathbb{P}_{S\sim\mu^{(\boldsymbol{x})}}[e \in S] \geq \alpha x_e$ for all e—this LP imposes a strengthened *stationary implementability* condition. Unlike standard OCRSs, where implementability is defined by conditioning on an arrival prefix, stationarity requires conditioning on the selection outcomes of *all other* elements. This provides a mathematical formalization of the principle of "treating every element as if it might be the last." We show that the stationary OCRS LP is exact: a distribution is feasible in this LP if and only if it can be implemented by a stationary OCRS.

A key algorithmic consequence of our LP characterization is a universal implementation procedure. Given any feasible witness distribution $\mu^{(\boldsymbol{x})}$ for the stationary OCRS LP, we provide a simple *simulate-then-replace* meta-algorithm (sharing similarities with a Gibbs sampling update or single-site Glauber dynamics, see Levin and Peres (2017) for instance) that implements $\mu^{(\boldsymbol{x})}$ for *every* arrival order.

The algorithm maintains a random set $\widehat{S}$ (initialized as a draw from $\mu^{(\boldsymbol{x})}$ at time 0) whose distribution is preserved throughout the execution. $\widehat{S}$ consists of a "real" part containing the elements selected so far and an "imaginary" part representing a random subset of future elements. When an element arrives, its membership in $\widehat{S}$ is resampled using the appropriate conditional probability under $\mu^{(\boldsymbol{x})}$ and in consistency with its activation.

As a result, once a distribution $\mu^{(\boldsymbol{x})}$ is specified, the online policy is essentially fixed, and the design of a stationary OCRS reduces to finding a "good" feasible solution to the stationary OCRS LP. Accordingly, most of the paper focuses on constructing explicit witness distributions with constant selectability factors.

Maximum-entropy template & the addability analysis. Our main distributional construction is based on a maximum-entropy principle. Fixing target marginals $\boldsymbol{p} = \alpha\boldsymbol{x}$, where $\alpha \in (0, 1)$ is the selectability factor of interest, we consider the maximum-entropy distribution μ^* over feasible sets with these marginals. For downward-closed systems, this choice leads to a crucial structural simplification: the resulting distribution over feasible sets is a *Gibbs distribution* (defined later; cf. Singh and Vishnoi (2014); Gharan et al. (2011)), and has a product form with weights. In particular, the conditional inclusion probability of an element depends only on its weight and on whether the element is *addable* to the current set. As a result, feasibility of μ^* for the stationary OCRS LP with parameter α reduces to verifying a single, intuitive condition—namely, lower bounding the probability that a random feasible set drawn from μ^* allows for the addition of any given element with probability α. This reduction converts the analysis of a stationary OCRS into a static probabilistic statement, and it is the main technical engine behind our results for matchings and uniform matroids.

Comparison with greedy OCRSs. The greedy OCRSs of Feldman et al. (2016, 2021) are also analyzed using a related notion of addability. They fix a random downward-closed subfamily of feasible sets, start from the empty set, and accept an active element whenever it can be added to the set selected so far. Our approach uses the same feasibility test, but changes the state on which it is applied—the simulate-then-replace policy starts from a simulated imaginary state $\widehat{S}$ drawn from a target distribution rather than the empty set. Furthermore, unlike that paper, we allow the probability with which we discard feasible elements to depend on our past decisions. For Gibbs distributions, our algorithm can be viewed as a greedy policy (accepting elements with a fixed probability when they can be added to $\widehat{S}_{-e}$) initialized at a good imaginary state.[5] Our results will explain why greedy behavior performs significantly better when started from the correct stationary state.

Stationary OCRSs for matchings & uniform matroids. We instantiate the maximum-entropy framework in several canonical feasibility environments and obtain simple, explicit stationary OCRSs. For (non-bipartite) matchings, we obtain a clean $\frac{1}{3}$-selectable S-OCRS (and extend it to rank-L hypergraph matchings to obtain a $\frac{1}{L+1}$-selectable S-OCRS, matching the usual benchmark for this environment; see Ma et al. (2026); Correa et al. (2024)). For bipartite matchings, we obtain an explicit selectability constant

$$\alpha^* = \frac{3-\sqrt{5}}{2},$$

and we show that this constant is optimal within the stationary class. This result provides an improved OCRS for this environment, and therefore settles an open problem in Ma et al. (2024). The factor is also optimal for a natural class of OCRSs (satisfying a concentration property), and is conjectured to be optimal in general. This result also provides the best known competitive ratio for the bipartite matching prophet inequality problem, improving on Gravin and Wang (2023); Ezra et al. (2022); MacRury et al. (2025). Our proof highlights the advantage of the distributional viewpoint over previous approaches. In our analysis, the addability of an edge is the same as both endpoints being unmatched, and a classical correlation inequality for Gibbs measures over matchings gives a short analysis.

For the k-uniform matroid (selecting at most k elements from E), also known as the *magician's problem* (Alaei, 2014), we also provide an explicit Gibbs distribution achieving the optimal stationary selectability. The optimal constant has a simple description using Poisson CDFs:

$$\alpha_k := \mathbb{P}[Q<k \mid Q \le k] = \frac{\mathbb{P}[Q<k]}{\mathbb{P}[Q\le k]} \approx 1-\sqrt{\frac{2}{\pi k}}+\mathcal{O}\left(\frac{1}{k}\right), \qquad \text{where } Q \sim \mathrm{Poisson}(k).$$

The associated simulate-then-replace algorithm fixes the greedy algorithm (with independent rejection w.p. $\mathcal{O}(\sqrt{\log k}/k)$) by initializing $\widehat{S}$ at a Gibbs distribution. This greedy algorithm is known to achieve a sub-optimal selectability of $1-\mathcal{O}(\sqrt{\log k/k})$.

S-OCRSs for hypergraph matchings with capacity (or network revenue management). We further develop the previous insight by designing an S-OCRS distribution for rank-L hypergraph matchings in which

[5] We will always assume that the order in which we observe elements is not influenced by this initially sampled imaginary state.

vertices have capacity k. Rather than using the solution of a maximum-entropy convex program, we use an explicit Gibbs distribution where every hyperedge is selected independently with probability γx_e for $\gamma = 1 - \mathcal{O}(\sqrt{\log L/k})$, after which we condition on obtaining a feasible set. This leads to a greedy-with-random-discarding algorithm started from a stationary distribution that is a (near-optimal) $1 - \sqrt{2\log L/k} - \mathcal{O}(1/\sqrt{k})$-selectable OCRS.[6] The dependence on L and k in this result matches large-capacity results in several related online resource allocation models. See for instance the work of Gupta and Molinaro (2016); Kesselheim et al. (2018), who obtain an analogous result for random-order online packing LPs, and Ghuge et al. (2025) for single sample prophet inequalities. We also provide a different OCRS for this setting with a slightly improved selectability, which uses our earlier k-uniform S-OCRS as a black box and has a provable polynomial-time implementation.

Further extensions beyond the maximum-entropy approach. Finally, we delineate the limits of the maximum-entropy approach and explore how the approach can be modified to yield S-OCRSs in other feasibility environments. While these approaches are different from our standard maximum-entropy approach, they rely on similar ideas or are slightly modified versions of this approach.

Weakly Rayleigh matroids: relative entropy maximization & thinning. For general matroids, maximum-entropy distributions over independent sets can concentrate on configurations with poor addability, and therefore fail to yield strong S-OCRSs. To overcome this barrier, we develop a modified template for a broad class of matroids with negative dependence structure, which we call *weakly Rayleigh matroids*. Our construction replaces entropy maximization with relative entropy maximization with respect to a Rayleigh base measure, followed by a simple thinning step. This yields an explicit $\frac{1}{2}$-selectable S-OCRS, matching the implicit OCRS in Lee and Singla (2018). Note that $\frac{1}{2}$ is the best factor possible for this class.

Knapsack constraints: Heavy & light items. For knapsack constraints, as with matroids, there is a barrier to applying the maximum entropy approach directly (also faced by Feldman et al. (2016, 2021) in designing deterministic greedy OCRSs). However, there is still a $1/4$-selectable S-OCRS using a modified maximum entropy approach. We obtain this guarantee by using the technique of splitting the items into light and heavy types (Feldman et al., 2016, 2021; Dutting et al., 2020), and then using a mixture of two different maximum entropy S-OCRSs. We note that our OCRS for the knapsack constraint is not optimal, since Jiang et al. (2025a) give an OCRS with selectability $1/(3+e^{-2})$. This is not just a limitation of our analysis, since in their tight instance, we can show that no S-OCRS can obtain a selectability larger than $3/10 < 1/(3+e^{-2})$.

General downward-closed feasibility environments. For general downward closed feasibility environments, we show that there is a $\Theta(1/\sqrt{|E|})$-selectable S-OCRS, and no (even ordinary) OCRS can do better. Our results in this general setting are thus sharp. To the best of our knowledge, these are the first results for OCRSs in general downward closed constraints, though the impossibility result is closely related to the Kleinberg

[6] It is well known that the factor of $1 - \mathcal{O}\left(\sqrt{\frac{\log L}{k}}\right)$ cannot be beaten, even for offline contention resolution schemes on uniform multi-hypergraphs.

and Weinberg (2012) proof of impossibility of obtaining a competitive ratio better than $O\left(\frac{1}{\sqrt{q}}\right)$ for prophet inequalities for the intersection of q matroids.[7]

Extension to $\boldsymbol{x} \in b\boldsymbol{\mathcal{P}}$. Our results extend to the case where $\boldsymbol{x}$ belongs to $b\mathcal{P}$, a version of $\mathcal{P}$ scaled by a parameter $b \geq 0$ (we must still assume that $x_e \in [0,1]$). For rank-L hypergraph matchings, we obtain a $\frac{1}{1+bL}$-selectable S-OCRS. For bipartite matchings, we obtain a $\frac{2b+1-\sqrt{4b+1}}{2b^2}$-selectable S-OCRS. For k-uniform matroids, we obtain an $\alpha_k(b)$-selectable S-OCRS, where

$$\alpha_k(b) \;:=\; \mathbb{P}[Q < k \mid Q \leq k], \qquad \text{and } Q \sim \mathrm{Poisson}(bk).$$

For weakly Rayleigh matroids, we obtain a $\frac{1}{1+b}$-selectable S-OCRS. Finally, for knapsack constraints, we obtain a $\frac{1}{2+2b}$-selectable S-OCRS.

Algorithmic tractability. Across all environments, we complement the distributional results with a systematic discussion of polynomial-time implementation. This emphasis on implementability is particularly important for applications in operations research. In particular, in Appendix EC.6.1, we discuss how to sample (exactly or approximately) from these distributions in polynomial time and how to compute the required conditional expectations, enabling a computationally efficient implementation of our algorithms in several feasibility environments. Our S-OCRSs are fully explicit and implementable for matchings, k-uniform matroids, knapsack constraints, $\mathbb{C}$-representable matroids, and, more generally, whenever we can efficiently find a maximum-entropy distribution or KL-projected distribution. In addition, we also provide an OCRS for the network revenue management setting that runs in polynomial time.

Reducing Bayesian online allocation of reusable products to S-OCRS design. Finally, we return to our second motivation for studying S-OCRSs and show that they are exactly the tool needed to handle reusable resources. We consider a Bayesian online allocation model with reusable products over an arbitrary feasibility environment (and with stationary request arrivals). We let E be a ground set of all *products*, and $\mathcal{N}$ a family of multisets of products that can be feasibly allocated at any particular time (so multiple units of the same product can also be simultaneously allocated; $\mathcal{F} \subseteq \mathcal{N}$). Requests of different types $i \in \mathcal{I}$ arrive according to an independent homogeneous Poisson process with constant rate λ_i. When a type-i request arrives, the algorithm may reject it or allocate one product $e \in E$. Such an allocation earns reward $r_{ie} \geq 0$ and occupies product e for an exponentially distributed duration with mean $\ell_{ie} > 0$, after which the product becomes available again.[8]

For this general model, we establish a formal reduction from the Bayesian online allocation of reusable products in any feasibility environment to the design of S-OCRSs in the same environment: any α-selectable S-OCRS yields an online policy that obtains an α-fraction of the long-run average reward of the optimal omniscient policy (which knows all future arrivals). The reduction proceeds in two steps.

[7] Note that while logarithmic competitive ratios are known for prophet inequalities in this general setting, OCRSs are dual to *ex-ante* prophet inequalities rather than the usual prophet inequalities, which explains why our results do not contradict the existing literature.

[8] The exponential duration distribution assumption is not necessary, and is only made for expositional clarity.

First, starting from an ex-ante fluid LP relaxation of the reusable product model, we use a propose-then-admit argument to reduce the construction of good policies to the construction of admission-control primitives. Second, we show that an S-OCRS for a particular feasibility environment, through its witness distribution as a feasible solution to the S-OCRS LP, can be used to design admission control primitives. Combined with the S-OCRSs in this paper, this recipe yields approximately optimal online allocation algorithms for reusable products across all of our feasibility environments—thereby answering question (iii) in the affirmative.

In a companion paper (Aminian et al., 2026b), we focus on the special feasibility environment of k-uniform matroids, and allow for non-homogeneous Poisson process arrivals of requests, together with a finite horizon. By leveraging the ideas related to stationarity developed in the current paper, we solve a long-standing open problem in revenue management with reusable resources. More precisely, it turns out that in the vanishing regime for k-uniform matroids, where each of the x_e is tiny, there is an alternative description of the solution to the maximum-entropy convex program. This description leads to a simpler optimal S-OCRS construction with selectability factor α_k that can be applied to non-homogeneous Poisson arrivals. This yields an α_k-selectable policy for Bayesian online allocation with reusable resources under heterogeneous arrivals, improving upon the previously best known bound of $1-\Theta(\sqrt{\log k/k})$ in Feng et al. (2022) and matching the best explicit bound known for the non-reusable setting Wang et al. (2018).

Table 1 summarizes our selectability guarantees for various feasibility environments, together with the principal contribution of each result. Our work is also related to several lines of research in operations research and computer science, which we briefly summarize in Section EC.1.

Table 1 Summary of our explicit S-OCRS guarantees and their main contributions.

Feasibility environment	**Selectability**	**Main contribution**	**Section**
Bipartite matchings	$\frac{3-\sqrt{5}}{2} \approx 0.382$	Improved OCRS; optimal S-OCRS.	5
Rank-L hypergraph matchings	$\frac{1}{L+1}$	Stationary scheme matching basic OCRS benchmark.	5
k-uniform matroids	$1-\sqrt{\frac{2}{\pi k}}+\mathcal{O}\left(\frac{1}{k}\right).$	Optimal S-OCRS; matches best explicit OCRS upper bound asymptotically.	5
Rank-L hypergraphs with capacity k (NRM)	$1-\mathcal{O}\left(\sqrt{\frac{\log L}{k}}\right)$	Optimal even among offline CRSs for $k \gg \log L$.	6
Weakly Rayleigh matroids	$\frac{1}{2}$	Optimal even among OCRSs.	EC.2
Knapsack constraints	$\frac{1}{4}$	Stationary scheme matching basic OCRS benchmark.	EC.3
General downward-closed	$\Theta\left(1/\sqrt{\lvert E\rvert}\right)$	First OCRS guarantee; optimal among OCRSs.	EC.4

Consequence for reusable resources. For every environment above, an α-selectable S-OCRS yields an α-approximation for the Bayesian online allocation of reusable products against the long-run omniscient benchmark (Section 7).

2. Preliminaries and Notation

We begin by formalizing the general setting studied in this paper and the notion of an online contention resolution scheme (OCRS), introduced by Feldman et al. (2016, 2021). We also introduce the notation and provide some details for the specific feasibility environments we focus on.

General setting. Throughout this paper, we work with a finite ground set of elements E and a downward-closed family of feasible sets $\mathcal{F} \subseteq 2^E$, i.e., for every $S \in \mathcal{F}$ and every $T \subseteq S$ we have $T \in \mathcal{F}$. We refer to $(E, \mathcal{F})$ as a *feasibility environment*. Let

$$\mathcal{P} := \mathrm{conv}\{\mathbf{1}_S : S \in \mathcal{F}\} \subseteq [0,1]^E$$

be the corresponding feasibility polytope, where $\mathbf{1}_S \in \{0,1\}^E$ is the incidence vector of S. For a (deterministic or random) set $S \subseteq E$ and an element $e \in E$, we write S_{-e} for the set obtained by removing e, i.e., $S_{-e} = S \setminus \{e\}$. Note that if $S \in \mathcal{F}$ then $S_{-e} \in \mathcal{F}$ as well.

Given a feasibility environment $(E, \mathcal{F})$, we fix a point $\boldsymbol{x} \in \mathcal{P}$ and interpret x_e as the marginal *activation* probability of element e. The *active set* $A \subseteq E$ is then a random set generated by including each element $e \in E$ independently with probability x_e. We will always assume $x_e > 0$ for all $e \in E$ (in all our arguments, if x_e were ever equal to 0, we could always replace E with $E \setminus \{e\}$). We then consider an online setting in which elements arrive sequentially and are revealed to an online algorithm that makes irrevocable decisions. An *arrival order* is a permutation $\pi \in \mathrm{Sym}(E)$, that is, a bijection $\pi : 1, \ldots, |E| \to E$. With this setup[9], we recall the following standard definitions (cf. Feldman et al. (2016, 2021)).

DEFINITION 1 (OCRS). An online contention resolution scheme for $(E, \mathcal{F})$ is an online algorithm that, for every $\boldsymbol{x} \in \mathcal{P}$ (known by the online algorithm upfront) and every permutation π, processes elements sequentially in the order $\pi(1), \pi(2), \ldots, \pi(|E|)$. When an element e is revealed, the algorithm learns whether $e \in A$ and must irrevocably decide whether to accept or reject e. In the end, the algorithm must output a random set S satisfying $S \subseteq A$ and $S \in \mathcal{F}$.

DEFINITION 2 (α-SELECTABILITY). An OCRS is *α-selectable* (also called *α-balanced*) if for every $\boldsymbol{x} \in \mathcal{P}$, for every arrival order π and every $e \in E$,

$$\mathbb{P}[e \in S] \geq \alpha x_e.$$

Equivalently, since $e \in S$ implies $e \in A$, we can also write this as $\mathbb{P}[e \in S \mid e \in A] \geq \alpha$.[10]

Central to our technical framework in this paper, we also rely on a distributional view of OCRSs. Fixing $\boldsymbol{x}$ and an arrival order π, any OCRS induces a distribution over feasible sets; we denote by $\mu^\pi \in \Delta(\mathcal{F})$ the distribution of the final selected set S. In general, μ^π may depend on π.

[9] In our positive results, we will allow for an 'online' (or 'adaptive') adversary who can change the order in which future elements are revealed based on the algorithm's decisions in the past, but does not have access to the activations of unrevealed elements, or the random bits the algorithm uses.

[10] Unlike in (Feldman et al., 2016, 2021), we consider the *selectability* of schemes that are not necessarily greedy, and our schemes are not tailored to an almighty adversary.

Feasibility environments studied in our paper. We will instantiate $(E,\mathcal{F})$ with the following concrete feasibility environments (and use the corresponding notation throughout the paper):

(i) **Rank-L hypergraph matchings (Section 5.1, Section 6, Section EC.5.3, and Section EC.5.8):** E is the edge set of a hypergraph $H=(V,E)$ of rank at most L (each $e\in E$ satisfies $|e|\le L$), and $S\subseteq E$ is feasible iff it is a matching, i.e., its hyperedges are pairwise vertex-disjoint, and hence

$$\mathcal{F}:=\Big\{S\subseteq E:\forall v\in V,\big|\{e\in S:v\in e\}\big|\le 1\Big\}.$$

We allow parallel edges. Note that for every $\boldsymbol{x}\in\mathcal{P}$, we must necessarily have:

$$\boldsymbol{x}\in[0,1]^E \text{ and } \sum_{e\in\delta(v)} x_e\le 1 \qquad \forall v\in V. \tag{1}$$

The special case $L=2$ corresponds to (generally non-bipartite) graph matchings. We will also be interested in the case where each vertex has a capacity k. In this case,

$$\mathcal{F}:=\Big\{S\subseteq E:\forall v\in V,\big|\{e\in S:v\in e\}\big|\le k\Big\},$$

and for every $\boldsymbol{x}\in\mathcal{P}$, we have

$$\boldsymbol{x}\in[0,1]^E \text{ and } \sum_{e\in\delta(v)} x_e\le k \qquad \forall v\in V. \tag{2}$$

This model captures *network revenue management (NRM)* with products that consume up to L resources and each resource has k available units: vertices represent resources, hyperedges represent products (bundles of resources), and a feasible allocation corresponds to choosing hyperedges that are collectively incident to each vertex at most k times; see, e.g., Ma et al. (2026).

(ii) **Bipartite matchings (Section 5.2):** E is the edge set of a bipartite graph $G=(U\cup V,E)$, and $S\subseteq E$ is feasible iff it is a (not-necessarily-perfect) matching. We allow parallel edges. Every $\boldsymbol{x}\in\mathcal{P}$ is a *fractional matching*, that is, it belongs to the following integral (and totally unimodular) polytope:

$$\boldsymbol{x}\in[0,1]^E: \quad \sum_{e\in\delta(z)} x_e\le 1 \qquad \forall z\in U\cup V. \tag{3}$$

(iii) ***k*-selection or *k*-uniform matroid (Section 5.3):** E is an arbitrary ground set and $S\subseteq E$ is feasible iff $|S|\le k$. Therefore $\mathcal{F}:=\big\{S\subseteq E:|S|\le k\big\}$ and $\boldsymbol{x}\in\mathcal{P}$ is characterized by a simple *capacity constraint*:

$$\boldsymbol{x}\in[0,1]^E: \quad \sum_{e\in E} x_e\le k. \tag{4}$$

(iv) **(Weakly Rayleigh) matroids (Section EC.2):** A *matroid* is a pair $\mathcal{M}=(E,\mathcal{I})$ where $\mathcal{I}\subseteq 2^E$ is a hereditary family of *independent sets* satisfying: (i) $\emptyset\in\mathcal{I}$; (ii) if $A\in\mathcal{I}$ and $B\subseteq A$ then $B\in\mathcal{I}$; (iii) if $A,B\in\mathcal{I}$ and $|A|<|B|$, then there exists $e\in B\setminus A$ such that $A\cup\{e\}\in\mathcal{I}$. A set S is feasible iff it is independent, i.e., $\mathcal{F}=\mathcal{I}$. A matroid can equivalently be described using its set of *bases* $\mathcal{B}$, consisting of all maximal independent sets in $\mathcal{I}$. A matroid has *rank* r iff all bases have size r. Given a matroid $\mathcal{M}$, the *rank function* (also known as the *rank oracle*) of the matroid is defined as $\mathrm{rank}_{\mathcal{M}}(S) := \max_{I \in \mathcal{I}: I \subseteq S} |I|$. Given the rank function, the feasibility polytope $\mathcal{P} = \mathrm{conv}\{\mathbf{1}_S : S \in \mathcal{F}\}$, called the *independence polytope*, is given by:

$$\mathcal{P} = \Big\{\boldsymbol{x} \in [0,1]^E : \sum_{e \in T} x_e \leq \mathrm{rank}_{\mathcal{M}}(T), \quad \forall T \subseteq E\Big\}. \tag{5}$$

It will also be convenient to define the *base polytope* of a rank-r matroid:

$$\mathcal{P}_{\mathcal{B}} := \Big\{\boldsymbol{q} \in [0,1]^E : \sum_{e \in T} q_e \leq \mathrm{rank}_{\mathcal{M}}(T) \forall T \subseteq E, \sum_{e \in E} q_e = r\Big\}. \tag{6}$$

It is a standard result that there is a polynomial-time equivalence between the rank and separation oracles for the independence (or base) polytope; see, e.g., Oxley (2011); Schrijver et al. (2003). We further focus on the subclass of *weakly Rayleigh* matroids that admit distributions over bases with a certain type of negative dependence (see Definitions EC.1 and EC.2). Many commonly encountered classes of matroids are weakly Rayleigh, including linear matroids over $\mathbb{C}$ (hence, graphic, laminar, and regular matroids).[11]

(v) **Knapsack constraints (Section EC.3):** E is a set of items with sizes $a_e \in (0,1]$, and $S \subseteq E$ is feasible iff $\sum_{e \in S} a_e \leq 1$. For every $\boldsymbol{x} \in \mathcal{P}$, we necessarily have that:

$$\boldsymbol{x} \in [0,1]^E \text{ and } \sum_{e \in E} a_e x_e \leq 1. \tag{7}$$

We note that our positive results will also work when we only assume that $\boldsymbol{x}$ belongs to a standard relaxation of the feasibility polytope $\mathcal{P}$, rather than the actual convex hull of feasibility sets. In particular, for hypergraph matchings, we only need to assume:

$$\boldsymbol{x} \in \widetilde{\mathcal{P}} = \left\{\boldsymbol{x} \in [0,1]^E : \sum_{e \in \delta(v)} x_e \leq k \quad \forall v \in V\right\}.$$

and for knapsack constraints, we only need to assume:

$$\boldsymbol{x} \in \widetilde{\mathcal{P}} = \left\{\boldsymbol{x} \in [0,1]^E : \sum_{e \in E} a_e x_e \leq 1\right\}.$$

3. Stationary OCRS: Definitions & LP Formulation

This section formalizes the main mathematical object studied in this paper: the *Stationary Online Contention Resolution Scheme (S-OCRS)*. We start by defining S-OCRS and providing intuition for this notion in Section 3.1. We then characterize the problem of designing S-OCRS in Section 3.2 via an LP over the space of distributions, which we call *the stationary OCRS LP*. This reduction is constructive and yields a meta S-OCRS algorithm that can implement any feasible solution of this LP—and hence plays the role of a template for all S-OCRS algorithms in the rest of the paper.

[11] Weakly Rayleigh matroids are a common weakening of Weak Half-Plane Property (WHPP) matroids and Rayleigh matroids. See Choe et al. (2004); Choe and Wagner (2006) for more details on the differences between these matroid classes.

3.1. Definition and motivation

One motivation to study stationarity (to be formalized soon) comes from a variant of the standard OCRS problem that naturally appears in Bayesian online allocation models with *reusable resources* (Feng et al., 2024, 2022; Rusmevichientong et al., 2020). In such models, one often reduces the design of an online policy to an OCRS-style *admission* subroutine, abstracting away the arrival process and rewards from the admission decisions. In that reduction, the input to the admission subroutine is stochastic requests generated in an online fashion by sampling the solution to an ex-ante offline fluid linear program for the original problem. The generated requests can be feasibly fulfilled in expectation. The goal of the subroutine is then to maintain the feasibility of the set of admitted requests, while ensuring that each request is selected with a large enough probability, à la OCRS selectability in Definition 2.

In reusable settings, an allocated agent (or request) uses our resources for a duration of time and releases the resource afterwards. A stylized abstraction of the reduced model that arises is the following recurring variant of the OCRS problem, which is a generalization of the standard OCRS problem when the requests (or elements) come back repeatedly.

DEFINITION 3 (THE RECURRING OCRS PROBLEM). Fix a downward-closed feasibility environment $(E, \mathcal{F})$ with corresponding feasibility polytope $\mathcal{P} = \mathrm{conv}(\mathcal{F})$ and a point $\boldsymbol{x} \in \mathcal{P}$ (known by the algorithm upfront). An adversary fixes an infinite sequence $e_1, e_2, \dots$ with $e_t \in E$ that is unknown to the algorithm. At time t, the algorithm observes e_t, and a fresh independent coin determines whether e_t is *active*, with probability x_{e_t}. An α-selectable recurring OCRS is an online algorithm that starts with $S_0 = \emptyset$ and, after each observation, maintains a random set S_t with the following properties:

(i) If e_t is inactive, then $S_t = S_{t-1} \setminus \{e_t\}$; if e_t is active, then $S_t \in \{S_{t-1} \setminus \{e_t\}, S_{t-1} \cup \{e_t\}\}$.

(ii) $S_t \in \mathcal{F}$ at all times t.

(iii) $\mathbb{P}[e_t \in S_t] \geq \alpha x_{e_t}$, where the probability is over the activity coins and the algorithm's randomness.

As in standard OCRSs, the goal is to maximize the selectability factor α.

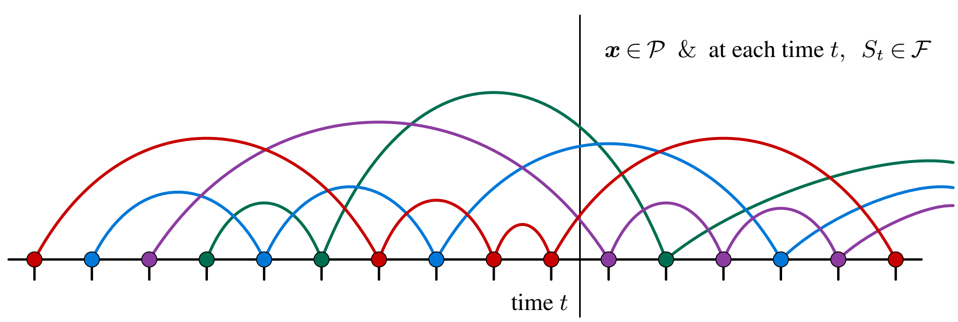

Figure 1 A recurring OCRS instance. Each color corresponds to one element $e \in E$. Consecutive arcs depict successive "epochs" of the same element; at the beginning of each epoch, e is independently active with probability x_e. If an active epoch is accepted, e remains selected until the end of that arc.

Informally, we may consider, for each element $e \in E$, *epochs* which are the intervals between its two consecutive appearances in the sequence. A *renewal event* occurs when an epoch begins. At each renewal of element

e, a fresh independent coin determines whether e is *active* in the next epoch. A recurring OCRS must decide at every renewal whether to select an active element e. If the algorithm selects e at the beginning of an active epoch, then e remains selected for the entire epoch and becomes unselected when it ends. The algorithm must maintain feasibility at all times, and the goal is to be α*-selectable* at *every* renewal of *every* element.

Unlike a standard OCRS instance—where each element appears once and the adversary chooses a single arrival permutation—in the recurring OCRS problem the algorithm repeatedly faces a new contention pattern as requests depart and re-arrive. Since epoch lengths can be arbitrary, the stream of renewal events can interleave in a highly irregular way. From the viewpoint of the a policy, this means that the relative order in which elements are encountered keeps reshuffling over time (Figure 1), and still the algorithm must remain feasible and maintain α-selectability at each renewal.

The above feature suggests that, in reusable settings, the exact permutation in which renewals are processed should be treated as incidental rather than as a meaningful part of the model. This motivates focusing on OCRSs whose output distribution is invariant to the arrival permutation: for each fixed $\boldsymbol{x} \in \mathcal{P}$, all arrival orders should induce the same distribution over the selected feasible set. Formally, we have the following definition.

DEFINITION 4 (STATIONARY OCRS (S-OCRS)). An OCRS is *stationary* if, for every $\boldsymbol{x} \in \mathcal{P}$, the distribution of the output set is independent of the arrival order. Formally, for every two permutations π, π' of E:

$$\mu^{(\boldsymbol{x},\pi)} = \mu^{(\boldsymbol{x},\pi')} \equiv \mu^{(\boldsymbol{x})},$$

where $\mu^{(\boldsymbol{x},\pi)} \in \Delta(\mathcal{F})$ (resp. $\mu^{(\boldsymbol{x},\pi')} \in \Delta(\mathcal{F})$) is the distribution of the output set of the OCRS algorithm for $\boldsymbol{x} \in \mathcal{P}$ under arrival order π (resp. π'). We write $\mu^{(\boldsymbol{x})} \in \Delta(\mathcal{F})$ for the common distribution of S (and drop $\boldsymbol{x}$ when it is clear from the context).

The following proposition formalizes the conceptual reason for why stationarity is the "right" abstraction for the recurring OCRS problem:

PROPOSITION 1 **(S-OCRS $\Rightarrow$ Recurring OCRS).** *Fix $\boldsymbol{x} \in \mathcal{P}$. If there exists an α-selectable stationary OCRS for $(E, \mathcal{F})$ at $\boldsymbol{x}$ with common output distribution $\mu \in \Delta(\mathcal{F})$, then there exists an online policy for the recurring OCRS problem with activation probabilities $\boldsymbol{x}$ that is α-selectable.*

We postpone the proof to Section EC.5.1, since it depends on concepts developed in the next section, where we identify an alternative characterization of S-OCRS via a certain linear program. This characterization is accompanied by a meta algorithm, which can be used to prove the reduction from recurring OCRS to S-OCRS.

3.2. Characterization via LP, and a "Simulate-then-Replace" meta algorithm

We start by introducing the following LP, where the decision variables are distributions over $\mathcal{F}$.

DEFINITION 5 (THE STATIONARY OCRS LP). Fix $\boldsymbol{x} \in \mathcal{P}$. In the stationary OCRS LP, we seek the largest $\alpha \in [0, 1]$ such that there exists a distribution $\mu \in \Delta(\mathcal{F})$ satisfying

$$\mathbb{P}_{S\sim\mu}[e \in S] \geq \alpha\, x_e \qquad \forall e \in E, \qquad \text{(SELECTABILITY)}$$

$$\mathbb{P}_{S\sim\mu}\big[e \in S \mid S_{-e} = T\big] \leq x_e \qquad \forall e \in E,\ \forall T \in \mathcal{F} \text{ with } e \notin T. \qquad \text{(STATIONARY-IMPLEMENTABILITY)}$$

A few comments are in order. Constraint (SELECTABILITY) is exactly α-selectability in the distributional view. Constraint (STATIONARY-IMPLEMENTABILITY) is a *stationary implementability* condition: even after conditioning on the selection status of *all* other elements, the conditional probability of selecting e cannot exceed its activation probability x_e. Note that when we condition on an event of the type $S_{-e} = T$, we always implicitly assume $\mathbb{P}_{S\sim\mu}\left[S_{-e} = T\right] > 0$. We now have the following proposition, demonstrating an equivalence between stationary OCRS algorithms and feasible solutions to the LP:

PROPOSITION 2 **(S-OCRS $\Longleftrightarrow$ stationary OCRS LP).** *Fix $\boldsymbol{x} \in \mathcal{P},\ \alpha \in [0,1]$. There exists an α-selectable stationary OCRS with activation probabilities $\boldsymbol{x}$ if and only if the stationary OCRS LP (Definition 5) is feasible.*

The full proof of the proposition is straightforward, and we postpone to Section EC.5.2. Instead, it is convenient for our purposes to go over an alternative constructive proof of *(LP $\Rightarrow$ S-OCRS)* using a randomized policy we call *"simulate-then-replace."* Assume we are given $\mu \in \Delta(\mathcal{F})$ that satisfies (STATIONARY-IMPLEMENTABILITY). The simulate-then-replace policy—formalized in Algorithm 1—begins by simulating a feasible *simulated set* $\widehat{S} \sim \mu$ of all elements. Then, once each element arrives, the policy *replaces* the simulated membership of that element with an activation-feasible decision (i.e., accepting the active element e with a target conditional probability $q_e/x_e = \mathbb{P}_{S\sim\mu}[e \in S \mid S_{-e} = T]/x_e$), updating S_e, while keeping the remaining allocation $T = \widehat{S}_{-e}$ intact. See Figure 2 for an illustration of how set $\widehat{S}$ evolves as we run Algorithm 1.[12]

To see why this algorithm is an α-selectable S-OCRS, note that due to Constraint STATIONARY-IMPLEMENTABILITY, we have $q_e \le x_e$, which is exactly what makes the acceptance probability q_e/x_e well-defined. Moreover, if $T \cup \{e\} \notin \mathcal{F}$ then $q_e = 0$; therefore, the algorithm maintains the feasibility of the selected set of elements throughout (as T is a superset of actually selected elements). Furthermore, if we consider the arrival time of element e and condition on $T = \widehat{S}_{-e}$, the updated membership of e is $\mathrm{Bernoulli}(q_e)$, and the rest of the simulated set remains equal to T. Therefore, the one-step update leaves the joint distribution of $\widehat{S}$ equal to μ. Since this holds for every processed element, after the full arrival sequence we still have $\mathrm{Law}(\widehat{S}) = \mu$, and hence $\mathrm{Law}(S) = \mu$. Finally, because of Constraint SELECTABILITY, we have $\mathbb{P}_{S\sim\mu}[e \in S] \ge \alpha\, x_e$, which is exactly the α-selectability.

$$\widehat{S} = \Big(\underbrace{\boxed{1}\boxed{0}\boxed{1}\boxed{0}\boxed{1}}_{\text{actual (already decided)}}\ \underbrace{\boxed{0}\boxed{1}\boxed{0}\boxed{1}}_{\text{simulated}} \Big) \sim \mu$$

coordinates $1, \ldots, i$ coordinates $i+1, \ldots, n$

Figure 2 Evolution of set $\widehat{S}$ throughout the execution of Algorithm 1.

[12] Our policy shares similarities with single-site Glauber dynamics (see for instance Chapter 3 in Levin and Peres (2017), or Vigoda (1999)), but we start at (and therefore stay at) the stationary distribution, the order in which we update sites is arbitrary, and updates are 'gated' by the activation.

Algorithm 1: Simulate-then-Replace policy for distribution μ

Input: Activation vector $\boldsymbol{x} \in \mathcal{P}$ and target distribution $\mu \in \Delta(\mathcal{F})$ satisfying (Stationary-Implementability)

Sample a feasible simulated set $\widehat{S} \sim \mu$; `// initialize simulated state with target distribution`

foreach $e \in E$ *revealed online (in the given order)* **do**

  $T \leftarrow \widehat{S}_{-e}$; `// forget the membership of e in the simulated set`

  if $e \in A$ **then**

    $q_e \leftarrow \mathbb{P}_{S\sim\mu}\left[e \in S \mid S_{-e} = T\right]$; `// calculate target conditional probability under μ`

    $T \leftarrow T \cup \{e\}$ with probability q_e / x_e; `// randomly replace the membership of e if active`

    Accept element e if $e \in T$

  else

    Reject element e ; `// reject element e if inactive`

  $\widehat{S} \leftarrow T$

return $S \leftarrow \widehat{S}$; `// final set has law μ`

We end this section by noting that although the reduction established in Proposition 2 conceptually simplifies the design of S-OCRS algorithms and reduces the design problem to solving the stationary OCRS LP, it is highly unclear how to solve this (exponential-size) LP or even find an appropriate feasible solution for this program. We address this point in the next section.

4. S-OCRS via the Maximum-Entropy Principle

In order to find good feasible solutions for the stationary OCRS LP, in this section we introduce a general framework that utilizes the maximum-entropy method (Gharan et al., 2011; Singh and Vishnoi, 2014; Straszak and Vishnoi, 2019). More specifically, for the design and analysis of our S-OCRS algorithms, we follow a template with the following four steps:

(i) Pick a target selectability parameter $\alpha \in (0,1)$ and define scaled marginals $\boldsymbol{p} := \alpha\boldsymbol{x}$,

(ii) Consider the *maximum-entropy distribution* μ^* over feasible sets $\mathcal{F}$ with marginals $\boldsymbol{p}$, which is the solution of a particular concave program (formally defined later),

(iii) Prove that under μ^*, the stationary implementability constraint is always satisfied,

(iv) Run the simulate-then-replace policy (Algorithm 1) with μ^* as input.

We formalize our template in Section 4.1. We then show in Section 4.2 that step (iii) is equivalent to proving a certain combinatorial statement related to adding an element to a given set. In the following section (Section 5), we will show how to apply this method to three feasibility environments: matchings in (hyper-)graphs, bipartite matchings, and k-uniform matroids.

4.1. Maximum-entropy with prescribed marginals

Fix any target marginal vector $\boldsymbol{p}$ that is in the relative interior of $\mathcal{P}$, which is the case in our applications after scaling by $\alpha < 1$. Consider the entropy maximization problem:

$$\max_{\mu \in \Delta(\mathcal{F})} \quad \mathsf{H}(\mu) := -\sum_{S \in \mathcal{F}} \mu(S) \log \mu(S) \qquad \text{(MAX-ENTROPY)}$$
$$\text{s.t.} \quad \mathbb{P}_{S\sim\mu}[e \in S] = p_e \qquad \forall e \in E. \qquad (8)$$

Note that this optimization problem is indeed a concave optimization over the space $\Delta(\mathcal{F})$. As Slater's condition holds (cf. Boyd and Vandenberghe (2004)), the optimizer is unique. Moreover, it is well known in combinatorics and statistical physics that the solution to this problem has a product form, which is often referred to as a *Gibbs distribution* (Jaynes, 1957; Georgii, 2011).

DEFINITION 6 (GIBBS DISTRIBUTION ON A FEASIBILITY SET). Fix a feasibility environment $(E, \mathcal{F})$ and a vector of weights $\boldsymbol{w} = \{w_e\}_{e\in E}$ with $w_e > 0$ for all e. The *Gibbs distribution* on $\mathcal{F}$ with parameters $\boldsymbol{w}$ is the probability distribution $\mu_{\boldsymbol{w}} \in \Delta(\mathcal{F})$ defined by

$$\mu_{\boldsymbol{w}}(S) := \frac{1}{\mathsf{Z}(\boldsymbol{w})} \Big(\prod_{e\in S} w_e\Big) \mathbf{1}\{S \in \mathcal{F}\}, \qquad \mathsf{Z}(\boldsymbol{w}) := \sum_{T\in\mathcal{F}} \prod_{e\in T} w_e. \qquad (9)$$

We now have the following standard lemma, for which we include a proof for completeness.

LEMMA 1 (**Vishnoi (2021), for instance**). *Let $\boldsymbol{p}$ be in the relative interior of $\mathcal{P}$ and let μ^* be an optimal solution of* (MAX-ENTROPY). *Then there exists a parameter vector $\boldsymbol{\theta} \in \mathbb{R}^E$ such that μ^* is a Gibbs distribution on $\mathcal{F}$ with parameters $w_e = e^{\theta_e}$. Moreover, for each $e \in E$, define*

$$\rho_e := \frac{w_e}{1 + w_e} \in (0, 1). \qquad (10)$$

Then for every feasible $T \in \mathcal{F}$ with $e \notin T$,

$$\mathbb{P}_{S\sim\mu^*}\big[e \in S \mid S_{-e} = T\big] = \frac{\mu^*(T \cup \{e\})}{\mu^*(T) + \mu^*(T\cup\{e\})} = \begin{cases} \rho_e, & \textit{if } T \cup \{e\} \in \mathcal{F}, \\ 0, & \textit{otherwise.} \end{cases} \qquad (11)$$

Proof. Introduce a Lagrange multiplier $\gamma \in \mathbb{R}$ for the normalization constraint and Lagrange multipliers $\{\theta_e\}_{e\in E}$ for the marginal constraints. The Lagrangian of (MAX-ENTROPY) is

$$\mathcal{L}(\mu, \gamma, \boldsymbol{\theta}) = -\sum_{S\in\mathcal{F}} \mu(S) \log \mu(S) + \gamma\Big(\sum_{S\in\mathcal{F}} \mu(S) - 1\Big) + \sum_{e\in E} \theta_e \Big(\sum_{S\ni e} \mu(S) - p_e\Big).$$

Since $\boldsymbol{p}$ lies in the relative interior of $\mathcal{P}$, Slater's condition holds (Boyd and Vandenberghe, 2004). Therefore the optimum is unique, attained in the interior, and satisfies the KKT conditions (strong duality). In particular, letting $(\boldsymbol{\theta}, \gamma)$ be the optimal duals of the concave program, for each $S \in \mathcal{F}$, the first-order condition gives

$$0 = \frac{\partial \mathcal{L}}{\partial \mu(S)} = -(\log \mu(S) + 1) + \gamma + \sum_{e\in S} \theta_e \implies \mu^*(S) = \exp(\gamma - 1) \exp\Big(\sum_{e\in S} \theta_e\Big).$$

Normalizing over $S \in \mathcal{F}$ yields the Gibbs form in (9), with $w_e := e^{\theta_e} > 0$.

Fix any $e \in E$ and condition on $S_{-e} = T$. As $\mathcal{F}$ is downward-closed, the only feasible sets compatible with this conditioning are T and (if feasible) $T \cup \{e\}$. Using the multiplicative form $\mu^*(S) \propto \prod_{f \in S} w_f$, we have

$$\frac{\mu^*(T \cup \{e\})}{\mu^*(T)} = w_e \qquad \text{whenever } T \cup \{e\} \in \mathcal{F}.$$

Therefore the conditional probability $\mathbb{P}_{S \sim \mu^*}\left[e \in S \mid S_{-e} = T\right]$ is $\mu^*(T \cup \{e\})/(\mu^*(T) + \mu^*(T \cup \{e\})) = w_e/(1 + w_e) = \rho_e$ when $T \cup \{e\} \in \mathcal{F}$, and 0 otherwise, proving (11). ☐

Using the maximum-entropy solution. Consider setting the target marginals to $\boldsymbol{p} = \alpha \boldsymbol{x}$ in the concave program (MAX-ENTROPY), where $\alpha \in (0, 1)$ is a target selectability parameter. Given the Gibbs distribution μ^* on $\mathcal{F}$ with parameters $\boldsymbol{w} = (w_e)_{e \in E}$—as the solution to the concave program—Lemma 1 shows that the simulate-then-replace meta-algorithm for μ^* takes a particularly simple form. Once an initial sample $\widehat{S} \sim \mu^*$ is drawn, each arriving element e, if active, can be processed by (i) testing whether $T \cup \{e\} \in \mathcal{F}$ and (ii), if $T \cup \{e\}$ is feasible, sampling a Bernoulli random variable with parameter ρ_e/x_e to decide whether to accept or reject e. See Algorithm 2 for a detailed description.[13]

By the guarantee of the simulate-then-replace meta-algorithm (Proposition 2), if μ^* is feasible for the stationary OCRS LP, then Algorithm 2 implements μ^* for every arrival order, and therefore yields an α-selectable S-OCRS. Checking constraint STATIONARY-IMPLEMENTABILITY for a general distribution μ is not always easy, but Lemma 1 is precisely what makes this step tractable for the special case of Gibbs distribution: For a Gibbs witness μ^*, the conditional probability $\mathbb{P}_{S \sim \mu^*}\left[e \in S \mid S_{-e} = T\right]$ depends only on whether e is *addable* to T, and in that case equals a constant ρ_e that does not depend on the particular choice of T. We formalize this statement in the next subsection.

Algorithm 2: SIMULATE-THEN-REPLACE for a Gibbs distribution μ^*

Input: Activation vector $\boldsymbol{x} \in \mathcal{P}$ and Gibbs distribution $\mu^* \in \Delta(\mathcal{F})$ with parameters $\boldsymbol{w} \in \mathbb{R}^E_{>0}$

Sample a feasible simulated set $\widehat{S} \sim \mu^*$

foreach $e \in E$ *revealed online (in the given order)* **do**

- $T \leftarrow \widehat{S}_{-e}$ // remove e from the simulated set
- **if** $e \in A$ *and* $T \cup \{e\} \in \mathcal{F}$ **then**
  - Accept element e with probability $\frac{\rho_e}{x_e} := \frac{w_e}{x_e(1+w_e)}$, and set $T \leftarrow T \cup \{e\}$ // randomly accept e if active and addable
- **else**
  - Reject element e // reject element e if inactive or adding it creates infeasibility
- $\widehat{S} \leftarrow T$

4.2. Addability as the sufficient condition for stationarity

For a feasible set $S \in \mathcal{F}$ and an element $e \in E$, define the *addability event*

$$\mathsf{Add}(e) := \{S_{-e} \cup \{e\} \in \mathcal{F}\}. \tag{12}$$

[13] Once again, we note the similarity to Glauber dynamics (Vigoda, 1999; Levin and Peres, 2017)

Under μ^*, combining Lemma 1 with the law of iterated expectations yields the factorization

$$p_e = \mathbb{P}[e \in S] = \mathbb{E}\big[\mathbb{P}_{S\sim\mu^*}\big[e \in S \mid S_{-e}\big]\big] = \mathbb{P}[\mathsf{Add}(e)] \cdot \rho_e. \tag{13}$$

Thus, once we set marginal probabilities $p_e = \alpha x_e$, (STATIONARY-IMPLEMENTABILITY)—which is equivalent to the constraint $\rho_e \le x_e$—reduces to proving the following lower bound on addability:

$$\mathbb{P}_{S\sim\mu^*}[\mathsf{Add}(e)] \ \ge\ \alpha \qquad \forall e \in E. \tag{ADDABILITY}$$

Therefore, if (ADDABILITY) holds, μ^* provides a feasible solution to the stationary OCRS LP satisfying Constraints (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY) with selectability α.

Analysis via addability. Putting everything together, we have reduced the construction of a stationary OCRS with selectability α to establishing the addability lower bound (ADDABILITY) for the maximum-entropy distribution with marginals $\alpha\boldsymbol{x}$. Once established for an environment $(E, \mathcal{F})$, Algorithm 2 would be a α-selectable S-OCRS for that environment.

5. Maximum-Entropy for Matchings & Uniform Matroids

We now investigate each of the three promised feasibility environments. We begin with a warm-up in Section 5.1 for non-bipartite matchings (easily extendable to hypergraph matchings). We then turn to bipartite matchings in Section 5.2, where we prove an optimal S-OCRS for this environment. Finally, in Section 5.3, we study k-uniform matroids and obtain an optimal S-OCRS for this setting.

The underlying approach is common across all three environments and follows the four-step maximum-entropy template introduced in Section 4. In particular, all algorithms are instantiations of Algorithm 2 for the corresponding Gibbs distribution. The only place where problem-specific structure is used is in establishing the addability lower bound in each case.

Finally, we note that a polynomial-time implementation of these algorithms requires two ingredients: (i) computing the Gibbs parameters $\boldsymbol{w}$ by solving (MAX-ENTROPY) (or its dual), and (ii) sampling (exactly or approximately) an initial set $\widehat{S} \sim \mu^*$ in polynomial time. We show that both steps admit polynomial-time routines for (bipartite and general) matchings and uniform matroid feasibility environments studied in the current section, and we defer the corresponding results and technical details to Section EC.6.

5.1. Warm-up: matchings in non-bipartite graphs

Recall from Section 2 the (non-bipartite) matching environment defined on a graph $G = (V, E)$. Our first application of the maximum-entropy template gives a simple S-OCRS for this setting. Our result matches the most basic analysis of the algorithm of Ezra et al. (2022).

THEOREM 1 **(S-OCRS for matchings).** *For the matching environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a Gibbs distribution μ^* such that Algorithm 2, when instantiated with μ^*, is a $\frac{1}{3}$-selectable S-OCRS.*

Proof. Fix any $\boldsymbol{x} \in \mathcal{P}$ and set $\alpha := \frac{1}{3}$ and $\boldsymbol{p} := \alpha\,\boldsymbol{x}$. Let μ^* be the maximum-entropy distribution—that is, solution to (MAX-ENTROPY)—with marginals $\boldsymbol{p}$. As we discussed in Section 4.2, we only need to verify the condition in (ADDABILITY). Fix an edge $e = \{u, v\} \in E$. The addability event

$$\mathsf{Add}(e) \;=\; \{S_{-e} \cup \{e\} \in \mathcal{F}\}$$

is exactly the event that neither endpoint of e is matched by S_{-e}. Now, for each vertex $v \in V$, we define the following "matched" event:

$$\mathsf{Matched}(v) \;:=\; \Big\{\exists f \in S \text{ with } v \in f\Big\}.$$

Since S is always a matching, for any fixed v the events $\{f \in S\}$ over distinct $f \in \delta(v)$ are disjoint. Therefore,

$$\mathbb{P}_{S\sim\mu^*}\big[\mathsf{Matched}(v)\big] \;=\; \sum_{f\in\delta(v)} \mathbb{P}_{S\sim\mu^*}[f \in S] \;=\; \sum_{f\in\delta(v)} p_f. \tag{14}$$

Using $p_f = \alpha x_f$ and (1), we get

$$\mathbb{P}_{S\sim\mu^*}\big[\mathsf{Matched}(v)\big] \;=\; \alpha \sum_{f\in\delta(v)} x_f \;\le\; \alpha. \tag{15}$$

Finally, $\mathsf{Add}(e)$ fails only if $\mathsf{Matched}(u)$ or $\mathsf{Matched}(v)$ occurs, so by a union bound and (15),

$$\mathbb{P}_{S\sim\mu^*}[\mathsf{Add}(e)] \;\ge\; 1 - \mathbb{P}[\mathsf{Matched}(u)] - \mathbb{P}[\mathsf{Matched}(v)] \;\ge\; 1 - 2\alpha \;=\; \alpha.$$

This is exactly (ADDABILITY). The conclusion follows. □

We should highlight that we use the same style of the above simple analysis to extend our results to other feasibility environments in the following subsection. Also, an almost identical argument to the proof of Theorem 1 shows the existence of $\frac{1}{L+1}$-selectable S-OCRS for rank-L hypergraph matching. We postpone the details to Section EC.5.3. This bound matches the common benchmark in this environment (and the known upper-bound on the best possible factor achievable by slightly restricted OCRS algorithms, c.f. Ma et al. (2026)).

5.2. Bipartite matchings

Recall from Section 2 the bipartite matching environment defined on a graph $G = (U \cup V, E)$. Our next application of the maximum-entropy template yields an optimal S-OCRS for this setting.

THEOREM 2 (**Optimal S-OCRS for bipartite matchings**). *For the bipartite matching environment $(E, \mathcal{F})$ and every fractional matching $\boldsymbol{x} \in \mathcal{P}$, there exists a Gibbs distribution μ^* such that Algorithm 2, when instantiated with μ^*, is an α^*-selectable S-OCRS with $\alpha^* = \frac{3-\sqrt{5}}{2} \approx 0.382$. Moreover, α^* is the best possible selectability achievable by any S-OCRS algorithm.*

Before proceeding to the proof, a few remarks are in order. As we show in Section EC.6, our algorithm runs in polynomial time. Consequently, it resolves an open problem in Ma et al. (2024) by providing an improved OCRS for bipartite matchings in the non-vanishing regime, and also results in a prophet inequality with the same competitive ratio. Prior to our work, Gravin and Wang (2023) obtained a $1/3$-prophet inequality via a simple pricing-based policy. This factor was subsequently improved by Ezra et al. (2022) to 0.337, and then by MacRury et al. (2025) to 0.349. Our work improves all of these results.

Not only is $\alpha^* = (3-\sqrt{5})/2$ the best bound achievable by any S-OCRS algorithm, but it is also the best bound attainable by a broad class of OCRS policies satisfying the so-called *concentration property*, and is conjectured to be optimal for general matching OCRSs (Ma et al., 2024).

Positive result. The proof of the positive result follows the same addability argument as in Section 5.1, but it relies on the following well-known *positive correlation* lemma for Gibbs measures over matchings in bipartite graphs (see Theorem 6.3 in Heilmann and Lieb (1972) or Lemma 22.15 in Molloy and Reed (2002)), which has been extensively studied in combinatorics and statistical physics, sometimes under the name "monomer-dimer systems." To state this lemma, for $z \in U \cup V$, define the following event (similar to Section 5.1)

$$\mathsf{Matched}(z) \;:=\; \Big\{\exists f \in S \text{ with } z \in f\Big\}.$$

LEMMA 2 **(Gibbs positive correlation for bipartite matchings).** *When $S \subseteq E$ is sampled from a Gibbs distribution on matchings of a bipartite graph $G = (U \cup V, E)$, the events* $\mathsf{Matched}(u)$ *and* $\mathsf{Matched}(v)$ *(equivalently, their complements) for $u \in U$ and $v \in V$ are positively correlated.*

We first prove Theorem 2 assuming Lemma 2, and then provide a combinatorial proof of the lemma for completeness in Section EC.5.4.

Proof of Theorem 2. Fix any $\boldsymbol{x} \in \mathcal{P}$ and set $\alpha := \alpha^*$ and $\boldsymbol{p} := \alpha\boldsymbol{x}$. Let μ^* be the maximum-entropy distribution—that is, the solution to (MAX-ENTROPY)—with marginals $\boldsymbol{p}$. Based on the discussion in Section 4.2, we only need to verify the condition in (ADDABILITY). Fix an edge $e = \{u, v\} \in E$ with $u \in U$ and $v \in V$. The addability event (12) is

$$\mathsf{Add}(e) \;=\; \big\{S_{-e} \text{ contains no edge incident to } u \text{ or } v\big\}.$$

As before, for any $z \in U \cup V$ the events $\{f \in S\}$ over distinct $f \in \delta(z)$ are disjoint. Therefore, we have that

$$\mathbb{P}_{S\sim\mu^*}\big[\mathsf{Matched}(z)\big] \;=\; \sum_{f\in\delta(z)} \mathbb{P}_{S\sim\mu^*}[f \in S] \;=\; \sum_{f\in\delta(z)} p_f.$$

Using $p_f = \alpha x_f$ and feasibility in expectation in (3), we obtain

$$\mathbb{P}_{S\sim\mu^*}\big[\mathsf{Matched}(z)\big] \;=\; \alpha \sum_{f\in\delta(z)} x_f \;\le\; \alpha. \tag{16}$$

If u and v are simultaneously unmatched, then $\mathsf{Add}(e)$ occurs. Hence

$$\begin{aligned}
\mathbb{P}_{S\sim\mu^*}[\mathsf{Add}(e)] &\ge \mathbb{P}_{S\sim\mu^*}\Big[\overline{\mathsf{Matched}(u)} \cap \overline{\mathsf{Matched}(v)}\Big] \\
&\ge \mathbb{P}\Big[\overline{\mathsf{Matched}(u)}\Big]\,\mathbb{P}\Big[\overline{\mathsf{Matched}(v)}\Big] \qquad \text{(by Lemma 2)} \\
&= (1-\mathbb{P}[\mathsf{Matched}(u)])(1-\mathbb{P}[\mathsf{Matched}(v)]) \\
&\ge (1-\alpha)^2 \;=\; \alpha. \qquad \text{(by the choice of } \alpha = \alpha^* = \tfrac{3-\sqrt{5}}{2}\text{)}
\end{aligned}$$

This is exactly (ADDABILITY) and finishes the proof of α^*-selectability.[14] □

Negative results. We first establish the impossibility result underlying Theorem 2. To show that $\alpha^* = \frac{3-\sqrt{5}}{2}$ is the best possible selectability constant, we prove the following lemma—making essential use of our LP characterization in Proposition 2—by constructing a counterexample for the stationary OCRS LP. We defer the proof to Section EC.5.5.

LEMMA 3 (**Bipartite matching upper-bound**). *For every $\varepsilon \in (0, 1/2)$ there exists a bipartite graph G and a fractional matching $\boldsymbol{x} \in \mathcal{P}$ such that any distribution $\mu \in \Delta(\mathcal{F})$ satisfying the stationary constraints* (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY) *has selectability at most $\alpha^* + O(\varepsilon)$, where $\alpha^* = \frac{3-\sqrt{5}}{2}$. In particular, letting $\varepsilon \to 0$ implies that every S-OCRS has selectability at most α^*.*

We conclude this subsection by highlighting a limitation of the maximum-entropy approach. We have shown that this approach achieves selectability $\alpha^* = (3-\sqrt{5})/2$ for *bipartite* matchings. A natural question is whether the same approach can yield an α^*-selectable S-OCRS for general (non-bipartite) matchings.[15] Unfortunately, we show that the maximum-entropy approach *does not* extend to general matchings to obtain α^* (recall that we also showed this approach yields a $\frac{1}{3}$-selectable S-OCRS in Theorem 1). In particular, already on a four-vertex example (also studied in section 2.3 of MacRury et al. (2025)), the maximum-entropy distribution fails to provide selectability better than $(\sqrt{3}-1)/2 < \alpha^*$. We conjecture that the maximum-entropy distribution always provides an S-OCRS with selectability $(\sqrt{3}-1)/2$, which would beat the best known OCRS bound for this environment; we leave this as an open question. (We have formalized the impossibility result in Section EC.5.6) Evidence suggests that, in fact, *no S-OCRS* can obtain a selectability of α^* for general matchings.[16]

5.3. The $\mathbf{k}$-selection problem ($\mathbf{k}$-uniform matroid)

Recall from Section 2 the k-uniform matroid environment on a finite ground set E, where a set $S \subseteq E$ is feasible iff $|S| \le k$. Our final application of the maximum-entropy template in this section yields an optimal S-OCRS with selectability factor α_k for this environment, where

$$\alpha_k \;:=\; \mathbb{P}[Q < k \mid Q \le k] = \frac{\mathbb{P}[Q < k]}{\mathbb{P}[Q \le k]} \approx 1 - \sqrt{\frac{2}{\pi k}} + \mathcal{O}\left(\frac{1}{k}\right), \qquad \text{where } Q \sim \text{Poisson}(k). \tag{17}$$

THEOREM 3 (**Optimal S-OCRS for the $\mathbf{k}$-uniform matroid**). *For the k-uniform matroid environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a Gibbs distribution μ^* such that Algorithm 2, when instantiated with μ^*, is an α_k-selectable S-OCRS. Moreover, α_k is the best possible selectability achievable by any S-OCRS.*

[14] We note that our proof should also extend to the setting of non-bipartite matchings in the vanishing regime studied in Ma et al. (2024) due to correlation decay arguments from the literature implying that Lemma 2 holds approximately.

[15] Ma et al. (2024) conjecture that in the general matching environment no OCRS can be better than α^*-selectable, even in the vanishing regime with tree graphs, and that α^* should be achievable by some (ideally polynomial-time) OCRS algorithm.

[16] Let $V = \{a, b, c, u, v, h\}$ and $G = K_6$. Define the fractional matching x^ε by $x_e^\varepsilon = \frac{1}{3} - \frac{2\varepsilon}{3}$ for $e \in \{ab, ac, bc\}$, $\frac{1}{6} + \frac{\varepsilon}{6}$, for $e \in \{au, av, bu, bv, cu, cv\}$, $\frac{1}{2} - \frac{3\varepsilon}{2}$, for $e \in \{hu, hv\}$, and ε for $e \in \{uv, ah, bh, ch\}$. Numerical evidence with the S-OCRS LP as $\varepsilon \to 0$ suggests that we cannot obtain a selectability of $(3-\sqrt{5})/2$ on this instance.

Before turning to the proof, we briefly place Theorem 3 in context. The k-uniform matroid (selecting at most k active elements) is one of the most extensively studied OCRS environments in both computer science and operations research. As mentioned in the introduction, the seminal *magician's algorithm* of Alaei (2014) is optimal for this setting; it is obtained by solving an appropriate dynamic program for the best selectability via LP duality. The analysis in Alaei et al. (2012) established the explicit guarantee $1-\sqrt{1/(k+3)}$, and Jiang et al. (2025b) later made this analysis tight, proving optimality and characterizing the optimal selectability through a factor-revealing LP. The best explicit asymptotic bound currently available is due to Wang et al. (2018), who analyze the DP-based optimal policy in the vanishing regime (infinitesimal activation probabilities) and obtain the explicit factor $1-\sqrt{2/(\pi k)}+\mathcal{O}(1/k)$; since the vanishing regime is worst-case for the standard OCRS problem, this bound also applies to general instances. Finally, Dinev and Weinberg (2024) propose a simpler adaptive, combinatorial scheme that achieves a near-optimal factor of $1-\mathcal{O}(\sqrt{1/k})$, but this algorithm is ad-hoc and somewhat difficult to analyze.

In contrast to these DP/LP-based approaches, our policy is simple and highly interpretable. Algorithm 2 is almost static: it only maintains a simulated set $\widehat{S}$ as an *imaginary state* of the system, and makes local acceptance decisions so as to preserve the target distribution of $\widehat{S}$. For the uniform matroid, this can be viewed as a minor modification of the *greedy* rule: we run greedy on the imaginary state $\widehat{S}$, but accept an active element e only with a fixed probability ρ_e/x_e (rather than with probability 1). This yields the explicit selectability α_k for all k (and, in particular, recovers the optimal factor $1/2$ when $k=1$). Moreover, as we show in our companion paper (Aminian et al., 2026b), the stationarity property of our S-OCRS approach can be leveraged to obtain the same factor α_k for the *reusable* variant of the OCRS problem (and the reusable Bayesian online allocation problem), thereby improving upon the previously best known bound of $1-\mathcal{O}\left(\sqrt{\log(k)/k}\right)$ due to Feng et al. (2022). In that companion paper, we also further simplify the algorithm and its analysis.

Positive result. The proof of selectability in Theorem 3 again follows the addability argument from Section 5.1. The main additional ingredient is the following technical lemma (more general results are known in the literature, see Marsiglietti and Melbourne (2026), and the citations therein), comparing truncated Poisson and truncated Poisson–binomial distributions and establishing a consequence of a second-order stochastic dominance relationship between them.

LEMMA 4 **(Poisson comparison).** *Let $T=\sum_{e\in E}Y_e$ be a sum of independent Bernoulli random variables (with arbitrary success probabilities), and let Q be a Poisson random variable (with arbitrary mean). If $\mathbb{E}[T\mid T\le k]\le\mathbb{E}[Q\mid Q\le k]$, then*

$$\mathbb{P}[T<k\mid T\le k]\;\ge\;\mathbb{P}[Q<k\mid Q\le k]. \tag{18}$$

We first prove Theorem 3 assuming Lemma 4, and then give a self-contained proof of the lemma in Section EC.5.7 for the sake of completeness, using the *ultra-log-concavity* of the Poisson–binomial distribution (Pemantle, 2000).

Proof of Theorem 3. Fix any $\boldsymbol{x} \in \mathcal{P}$ and set $\alpha := \alpha_k$ and $\boldsymbol{p} := \alpha\, \boldsymbol{x}$. Once again, let μ^* be the maximum-entropy distribution—that is, the solution to (MAX-ENTROPY)—with marginals $\boldsymbol{p}$, and draw $S \sim \mu^*$. Based on the discussion in Section 4.2, it suffices to verify the addability lower bound (ADDABILITY) for $\alpha = \alpha_k$. For the k-uniform matroid, the addability event (12) is

$$\mathsf{Add}(e) \;=\; \big\{\, |S_{-e}| < k \,\big\}. \tag{19}$$

In particular, $\mathsf{Add}(e)$ holds whenever $|S| < k$. Since μ^* is a Gibbs distribution, we may write $\mu^*(S) \propto \prod_{e\in S} w_e$ over all S with $|S| \le k$, for some weights $w_e > 0$. Define $\rho_e := w_e/(1+w_e)$, and let $(Y_e)_{e\in E}$ be independent Bernoulli random variables with $\mathbb{P}[Y_e = 1] = \rho_e$. Then, for every $S \subseteq E$,

$$\mathbb{P}\big[\{e : Y_e = 1\} = S\big] = \prod_{e\notin S}(1-\rho_e)\cdot\prod_{e\in S}\rho_e = \Big(\prod_{e\in E}(1-\rho_e)\Big)\cdot\prod_{e\in S} w_e,$$

so conditioning on the event $\{T \le k\}$, where $T := \sum_{e\in E} Y_e$, yields the Gibbs distribution with parameters $\boldsymbol{w}$ on feasible sets (sets of size at most k). Equivalently, $S \sim \mu^*$ has the same distribution as $\{e : Y_e = 1\}$ conditioned on $\{T \le k\}$.

We now verify the mean condition in Lemma 4. On the one hand,

$$\mathbb{E}[T \mid T \le k] = \mathbb{E}[|S|] = \sum_{f\in E}\mathbb{P}[f\in S] = \sum_{f\in E} p_f = \alpha_k \sum_{f\in E} x_f \le k\alpha_k\ ,$$

where the last inequality uses feasibility of $\boldsymbol{x}$. On the other hand, for $Q \sim \mathrm{Poisson}(k)$[17],

$$\mathbb{E}[Q\mathbf{1}\{Q\le k\}] = \sum_{t=1}^{k} t\,\mathbb{P}[Q=t] = \sum_{t=1}^{k} k\,\mathbb{P}[Q=t-1] = k\,\mathbb{P}[Q<k],$$

and hence

$$\mathbb{E}[Q \mid Q\le k] = \frac{\mathbb{E}[Q\mathbf{1}\{Q\le k\}]}{\mathbb{P}[Q\le k]} = k\,\mathbb{P}[Q<k \mid Q\le k] = k\alpha_k.$$

Therefore $\mathbb{E}[T \mid T\le k] \le \mathbb{E}[Q \mid Q \le k]$. Applying Lemma 4 gives

$$\mathbb{P}_{S\sim\mu^*}\big[\,|S| < k\,\big] = \mathbb{P}[T < k \mid T \le k] \ge \alpha_k.$$

Since $\mathsf{Add}(e)$ holds whenever $|S| < k$, the above inequality implies $\mathbb{P}[\mathsf{Add}(e)] \ge \alpha_k$ for every e, which is exactly (ADDABILITY) for $\alpha = \alpha_k$—hence finishing the proof. □

Negative results. We next show that α_k is the best possible selectability constant in Theorem 3.

LEMMA 5 **(Upper bound approaching α_k).** *Fix $k \ge 1$. For every $n \ge k$, consider the symmetric instance $|E| = n$ and $x_e = q := k/n$ for all $e \in E$. Any distribution $\mu \in \Delta(\mathcal{F})$ satisfying the stationary OCRS LP constraints* (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY) *has selectability at most*

$$\frac{\mathbb{P}\big[\mathrm{Bin}(n-1,q) < k\big]}{\mathbb{P}\big[\mathrm{Bin}(n,q) \le k\big]}. \tag{20}$$

In particular, letting $n\to\infty$ implies that every S-OCRS has selectability at most α_k.

We defer the proof of Lemma 5 to Section EC.5.9.

[17] If $\boldsymbol{x} \in b\mathcal{P}$, we may let $Q \sim \mathrm{Poisson}(bk)$.

6. Beyond Maximum-Entropy: Network Revenue Management

In this section, we study a feasibility environment that shares similarities with both hypergraph matchings and k-uniform matroids, being a common generalization of both. As mentioned earlier, this class is commonly known as the *network revenue management (NRM)* environment, which is mathematically equivalent to a capacitated version of the hypergraph matching environment. While, as in the previous section, we will use a Gibbs distribution to solve the S-OCRS LP, we will not solve a maximum-entropy convex program.

Recall from Section 2 the NRM environment, i.e., the rank-L hypergraph matching environment with vertices with capacity k, defined on a hypergraph $H = (V, E)$. In this case, we have:

$$\mathcal{F} := \Big\{ S \subseteq E : \forall v \in V, \big| \{ e \in S : v \in e \} \big| \leq k \Big\},$$

and for every $\boldsymbol{x} \in \mathcal{P}$, we have

$$\boldsymbol{x} \in [0,1]^E \text{ and } \sum_{e \in \delta(v)} x_e \leq k \qquad \forall v \in V.$$

A simpler algorithm without maximum-entropy. We show that for this feasibility environment, the Gibbs distribution with $\rho_e = \gamma x_e$ with $\gamma = 1 - O(\sqrt{\log L/k})$, yields a selectability of $1 - O(\sqrt{\log L/k})$. While this choice of ρ_e does not arise as a solution to a max-entropy program[18], we can still run Algorithm 2. The advantage of this choice is that we accept each active element with the same constant probability γ, and the resulting algorithm is an even simpler greedy-with-homogeneous-randomized-discarding algorithm (started from the stationary distribution). Furthermore, (STATIONARY-IMPLEMENTABILITY) is satisfied automatically.

To formally state the result, fix $\varepsilon \in [0, 1)$. Let

$$Q \sim \text{Poisson}((1-\varepsilon)k), \quad \beta := \mathbb{P}[Q = k \mid Q \leq k], \quad \gamma := (1-\varepsilon)(1-\beta).$$

Let μ be the Gibbs distribution on $\mathcal{F}$ with parameters $\rho_e = \gamma x_e$.

THEOREM 4 (**S-OCRS for rank-L hypergraph matchings with capacity k**). *For the rank-L hypergraph matchings with capacity k environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, the Gibbs distribution μ with $\rho_e = \gamma x_e$ and γ as defined above is such that Algorithm 2, when instantiated with μ, is a $\gamma(1 - L\beta)$-selectable S-OCRS. In particular, for $k > 2\log L$, taking $\varepsilon = \sqrt{2\log L/k}$ gives selectability*

$$1 - \sqrt{\frac{2\log L}{k}} - O\left(\sqrt{\frac{1}{k}}\right).$$

The full proof of Theorem 4 is deferred to Section EC.5.8. The main idea is that to establish $\mathbb{P}_{S \sim \mu}[e \in S] \geq \gamma(1 - L\beta)x_e$, it suffices to show that $\mathbb{P}[\text{Add}(e)] \geq 1 - L\beta$. We prove this by using Lemma 4 to show that each individual vertex is at capacity with probability at most β, and then applying a union bound over the L vertices that e includes.

[18] We do not claim that the max-entropy approach does not work in this feasibility environment; our choice of Gibbs distribution gives a simple analysis and therefore we only consider studying this choice.

We note that there is also an ordinary OCRS with a slightly stronger selectability guarantee. First, independently thin each active edge with probability $1-\varepsilon$. Then, at every vertex v, run a separate virtual k-uniform-matroid S-OCRS on the thinned edges in $\delta(v)$, and accept an edge e precisely when all vertices $v \in e$ accept it in their virtual schemes. Since $\sum_{e\in\delta(v)}(1-\varepsilon)x_e \le (1-\varepsilon)k$, the proof of Theorem 3 shows that each virtual scheme is $(1-\beta)$-selectable, and hence rejects a thinned edge with probability at most β. By a union bound,

$$\mathbb{P}[e \text{ is accepted} \mid e \text{ is active}] \ge (1-\varepsilon)(1-|e|\beta) \ge (1-\varepsilon)(1-L\beta).$$

We show in Section EC.6 that the uniform-matroid S-OCRS is polynomial-time implementable, so this OCRS is polynomial-time implementable as well. Although every virtual scheme is stationary, their intersection need not be, so this construction is not an S-OCRS.

7. Application to Bayesian Online Allocation of Reusable Products

We now demonstrate the application of S-OCRS to Bayesian online allocation problems with reusable products under homogeneous Poisson arrivals of types. This reduction, at a high level, follows a propose-then-admit framework from our companion paper Aminian et al. (2026b)—inspired by earlier related work in the revenue management literature (Feng et al., 2024, 2022).

7.1. A Bayesian reusable-product model

Fix a set of *products* E, and $\mathcal{N} \subseteq \mathbb{Z}^E_{\ge 0}$ a family of multisets of products that can be feasibly allocated at any particular time. Note that we allow for multiple units of the same product to be simultaneously allocated, but we will always assume that at most N units can be allocated, so $\mathcal{N} \subseteq [N]^E$. We will also assume that $\mathcal{N}$ is downward closed, i.e., if $n \in \mathcal{N}$ and $n' \le n$ coordinatewise, then $n' \in \mathcal{N}$. Let $\mathcal{P}$ be the convex hull of $\mathcal{N}$.

Let $\mathcal{I}$ be a finite set of request types. Requests of type $i \in \mathcal{I}$ arrive according to an independent homogeneous Poisson process with constant rate λ_i. When a type-i request arrives, the algorithm may reject it or allocate one unit of a product $e \in E$. Such an allocation earns reward $r_{ie} \ge 0$ and occupies the unit of e for a duration that is independently exponentially distributed with mean $\ell_{ie} > 0$, after which the unit becomes available again. The rates, rewards, duration means, and feasibility environment are known to the algorithm. The allocation decision is immediate and irrevocable. If $N_e(t)$ is the number of units of product e occupied at time t, the feasibility requirement is

$$(N_e(t))_{e\in E} \in \mathcal{N} \quad \text{for every } t \ge 0.$$

For an online policy π, let $N^{\pi}_{ie}(T)$ denote the number of type-i requests allocated a unit of product e during $[0,T]$. The long-run average reward of the policy is:

$$\mathrm{R}(\pi) := \liminf_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i\in\mathcal{I}}\sum_{e\in E} r_{ie} N^{\pi}_{ie}(T)\right]. \tag{21}$$

We write OPT for the supremum of (21) over all (online and offline) policies, including those policies that have knowledge of all future arrivals at the time of making an allocation decision.[19] Thinking of OPT as the

[19] We assume all policies are ignorant of the time at which an allocated unit will return to the system, until the unit actually returns.

benchmark, we would like to identify an online (and ideally polynomial-time implementable) policy π for which $\mathrm{R}(\pi) \geq \alpha\, \mathrm{OPT}$ holds for all problem instances, for a constant $\alpha \in (0,1]$ as large as possible—referred to as the *competitive ratio* of the online policy π.

7.2. The static fluid relaxation and proposal step

Consider the following static fluid relaxation LP for the problem:

$$
\begin{aligned}
\mathrm{FLP} := \max_{y_{ie} \geq 0} \quad & \sum_{i \in \mathcal{I}} \sum_{e \in E} r_{ie} y_{ie} && && (22)\\
\text{s.t.} \quad & \sum_{e \in E} y_{ie} \leq \lambda_i && \forall i \in \mathcal{I}, && (23)\\
& x_e = \sum_{i \in \mathcal{I}} \ell_{ie} y_{ie} && \forall e \in E, && (24)\\
& \boldsymbol{x} \in \mathcal{P}. && && (25)
\end{aligned}
$$

Here, we interpret y_{ie} as the long-run average rate at which type-i requests are allocated a unit of product e by a policy, and constraint (24) gives the corresponding long-run average load on product e. The long-run reward then equals (22). The constraint (23) is a consequence of the fact that request of type i can only be allocated one unit each, while (25) follows from averaging the fact that at each time t, the set of allocated goods must be feasible. We therefore have the following proposition:

PROPOSITION 3. *For the reusable-product allocation problem,* $\mathrm{OPT} \leq \mathrm{FLP}$. *In particular, to identify a policy with* $\mathrm{R}(\pi) \geq \alpha\, \mathrm{OPT}$*, it suffices to identify a policy with* $\mathrm{R}(\pi) \geq \alpha\, \mathrm{FLP}$.

In order to use the static fluid relaxation LP to design an online policy, let $\{y^*_{ie}\}$ be an optimal solution to this fluid LP. Consider an online policy that upon the arrival of a type-i request, *proposes* each product e (mutually exclusively) with probability y^*_{ie}/λ_i, and makes no proposal with the remaining probability. Then, if the proposal is *admitted* by the product e, the request is allocated a unit of product e; else the request is rejected with no product allocated.

The proposals for every pair (i,e) form independent homogeneous Poisson processes of rate y^*_{ie}. If every proposal were admitted, the reward of the policy would be FLP, with each product receiving a load of x_e. Unfortunately, it is not possible to admit every request since we must have $(N_e(t))_{e \in E} \in \mathcal{N}$. However, if we can show that every request in the long run can be admitted with probability α, this provides a way to turn the ex-ante feasible proposals each product receives into an ex-post feasible allocation, at a cost of only a factor α for the rate which allocations are made. Thus, if we can find an admission rule that ensures that every proposal in the long run is accepted with probability α, we obtain a policy such that $\mathrm{R}(\pi) \geq \alpha\, \mathrm{OPT}$.

7.3. S-OCRSs are sufficient for admission control

For the purpose of the analysis, it is more convenient to work in a setting where the rate of proposals for every pair (i,e) is small; this is because to apply recurring OCRSs, we need consecutive requests for a product not to overlap. To achieve this goal, we can split each unit of a product into virtual clones, and decide when a proposal arrives which of the virtual clones it actually proposes to. This splitting also allows us to apply recurring OCRSs designed for feasibility families of *sets* to the setting with the multiset $\mathcal{N}$.

Formally speaking, an n*-split* of $(E,\mathcal{N})$ replaces every product e by a finite nonempty set $C_e^{(n)} := \{e^1,\dots,e^n\}$ of virtual clones. Writing $E_n = \bigsqcup_{e\in E} C_e^{(n)}$, the feasible family of sets associated with E_n is:

$$\mathcal{F}_n := \left\{ S \subseteq E_n : (|S\cap C_e^{(n)}|)_e \in \mathcal{N} \right\},$$

and we let $\mathcal{P}_n := \mathrm{conv}(\mathcal{F}_n)$. For a sufficiently large n where each product is split into n virtual clones, it turns out that an S-OCRS may be used as an approximate admission rule. We formalize this in the following lemma:

LEMMA 6. *Suppose that for every n-split of $(E,\mathcal{N})$, there is an α-selectable S-OCRS. Then, for every $\varepsilon>0$, there is an admission rule which ensures that* $\lim_{T\to\infty}\frac{1}{T}\mathbb{E}\left[N^\pi_{ie}(T)\right] \geq (\alpha-\varepsilon)y_{ie}$.

Proof. Let $\boldsymbol{x}\in\mathcal{P}$ be the vector of loads for the products in E for which we desire an admission rule, so:

$$x_e = \sum_{i\in\mathcal{I}} \ell_{ie} y_{ie},$$

where y_{ie} are rates of proposals. Note that $\|\boldsymbol{x}\|_\infty \leq N$. For $n \geq N\cdot\lceil\frac{1}{\varepsilon}\rceil$, consider the n-split $(E_n,\mathcal{F}_n)$. For $e\in E$, and $j\in[n]$, let $x^{(n)}_{ej} := \frac{x_e}{n+x_e}$, and note that $\mathbf{x}^{(n)}\in\mathcal{P}_n$ (because $\frac{\mathbf{x}}{n}\in\mathcal{P}_n$, and $\frac{x_e}{n+x_e}\leq\frac{x_e}{n}$). From our assumption that there is an α-selectable S-OCRS for $(E_n,\mathcal{F}_n)$ and $\mathbf{x}^{(n)}$, we know that there is a distribution μ_n such that for every clone e^j and every $T\in\mathcal{F}_n$ with $e^j\notin T$:

$$\frac{\mu_n(T\cup\{e^j\})}{\mu_n(T)+\mu_n(T\cup\{e^j\})} \leq \frac{x_e}{n+x_e} \implies \mu_n(T\cup\{e^j\}) \leq \frac{x_e}{n}\mu_n(T).$$

Furthermore, note also that:

$$\mathbb{P}_{S\sim\mu_n}[e^j\in S] \geq \alpha\frac{x_e}{n+x_e} \geq (\alpha-\varepsilon)\frac{x_e}{n}.$$

We now define the admission rule we consider. For each arriving proposal for a pair (i,e), assign it uniformly at random to one of the n clones $e^1,e^2,\dots,e^n$. For every time t, let $S(t)\in\mathcal{F}_n$ be the set of currently admitted proposals just before time t with $S(0)=\emptyset$. Suppose that a type-(i,e) proposal arriving at time t is assigned to clone e^j. If $e_j\in S(t)$ or $S(t)\cup\{e^j\}\notin\mathcal{F}_n$, reject the proposal. Otherwise, admit it with probability

$$f(S(t),e^j) := \frac{n\mu_n(S(t)\cup\{e^j\})}{x_e\mu_n(S(t))} \leq 1,$$

and then remove e^j from $S(t)$ after an exponentially distributed time with mean ℓ_{ie} passes. Note that since $S(t)$ is always in $\mathcal{F}_n$, the vector of the number of admitted proposals of each product $N_e(t) = |S(t)\cap C_e^{(n)}|$ is always in $\mathcal{N}$, so this is a valid admission rule. We claim that with our admission rule, $S(t)$ converges in distribution to μ_n. We postpone the proof of this fact to Section EC.5.17.

Now fix a type i, a product e, and a clone e^j. The type-(i,e) proposals assigned to e^j form a homogeneous Poisson process of rate y_{ie}/n. Let $N^\pi_{ie^j}(T)$ denote the number of type (i,e) proposals assigned to e^j that are admitted during $[0,T]$. Then, by the Campbell-Mecke formula (see equation (17) in Coeurjolly et al. (2017)):

$$\mathbb{E}\left[N^\pi_{ic}(T)\right] = \frac{y_{ie}}{n}\int_0^T \mathbb{E}\left[f(S(t),e^j)\right]dt.$$

Since $S(t)$ converges to μ_n, we know that: $\mathbb{E}[f(S(t),e^j)] \to \sum_{T\in\mathcal{F}_n} \mu_n(T) f(T,e^j)$. Therefore,

$$\begin{aligned}\lim_{T\to\infty} \frac{1}{T}\mathbb{E}\left[N^{\pi}_{ie^j}(T)\right] &= \frac{y_{ie}}{n} \sum_{T\in\mathcal{F}_n} \mu_n(T) f(T,e^j) \\ &= \frac{y_{ie}}{n}\frac{n}{x_e} \sum_{\substack{T\in\mathcal{F}_n \\ e^j\notin T}} \mu_n(T\cup\{e^j\}) \\ &= \frac{y_{ie}}{x_e}\mathbb{P}_{S\sim\mu_n}[e^j \in S] \\ &\geq (\alpha-\varepsilon)\frac{y_{ie}}{n}.\end{aligned}$$

The first equality is sometimes referred to as PASTA, or Poisson arrivals see time averages; see Wolff (1982). Summing over all n clones of product e gives $\liminf_{T\to\infty} \frac{1}{T}\mathbb{E}\left[N^{\pi}_{ie}(T)\right] \geq (\alpha-\varepsilon) y_{ie}$, as desired. □

Note that we may remove the ε loss if we wish by projecting the distribution μ_n onto a distribution on $\mathcal{N}$, and finding a subsequence of the projections that converges to a distribution μ. For the concrete Gibbs distribution we work with, this limit can sometimes be found directly. Furthermore, the assumption that the distribution of the duration for which a unit is occupied is exponential is not essential, and is made only for the sake of expositional ease in section EC.5.17.

7.4. Applications to environments in this paper

In case the constraints on the set of products E that can be allocated are weakly Rayleigh matroids, bipartite matchings, or general graph matchings, we may take $\mathcal{N} \subseteq \{0,1\}^E$ to be the family of feasible sets $\mathcal{F}$ in these environments, as defined in section 2. In these cases, an n-split of $(E,\mathcal{N})$ results in a feasibility environment $(E_n, \mathcal{F}_n)$ that is once again a weakly Rayleigh matroid, bipartite matching, or a matching constraint, respectively, so we obtain algorithms for the reusable product problem with competitive ratios $1/2, (3-\sqrt{5})/2$, and $1/3$, respectively. As discussed in Section EC.6, these can all be implemented in polynomial time.

Suppose that we now have resources in a set V, each of which have k identical units, and requests are for products $e \in E$, each of which consume at most one unit of at most L different resources. This is the classic NRM problem, and in this case:

$$\mathcal{N} = \left\{ m \in \mathbb{Z}^E_{\geq 0} : \sum_{e\ni v} m_e \leq k \quad \forall v \right\}.$$

Every n-split on $(E,\mathcal{N})$ results in a feasibility environment that is a rank-L hypergraph matching with vertices with capacity k, so we obtain an algorithm for the reusable production allocation problem with competitive ratio $1-\sqrt{2\log L/k} - \mathcal{O}(1/\sqrt{k})$. In case $L=1$, we get a competitive ratio of $\alpha_k = \mathbb{P}[Q<k \mid Q\leq k]$, where $Q \sim \mathrm{Poisson}(k)$, and our algorithm runs in polynomial time (as discussed in section EC.6). In case $L>1$, we may apply the trick used in section 6 where we scale the rate of proposals down by $1-\mathcal{O}(\sqrt{\log L/k})$, and each individual resource uses its own admission control to obtain an alternative polynomial time algorithm.

Finally, suppose we have products $e \in E$ with many units, each with size a_e, and at any point of time, we can allocate units with total size at most 1. In this case,

$$\mathcal{N} = \left\{ m \in \mathbb{Z}^E_{\geq 0} : \sum_{e\in E} a_e m_e \leq 1 \right\}.$$

Any n-split gives an ordinary knapsack constraint, so our $1/4$-selectable knapsack S-OCRS yields a $\frac{1}{4}$-competitive algorithm for the reusable products problem. Since our knapsack S-OCRS can be implemented in polynomial time (once again, see section EC.6), as long as

$$N = \max\{m_e : m \in \mathcal{N}, e \in E\} = \max_{e \in E} \left\lfloor \frac{1}{a_e} \right\rfloor$$

is polynomially bounded, then the admission rule can be implemented in polynomial time.

8. Conclusion & Future Directions

Motivated by online Bayesian selection problems, we introduced the stationary OCRS (S-OCRS) framework which gives us OCRSs that satisfy an arrival order invariance property. This structural restriction enabled the design of explicit OCRS policies across various feasibility environments, including matchings, hypergraph matchings, bipartite matchings, and k-uniform matroids in a unified manner using a maximum-entropy approach. Furthermore, for environments where maximum entropy is insufficient—specifically, knapsack constraints and weakly Rayleigh matroids—we developed alternative approaches that modify the maximum entropy approach. We studied the computational efficiency of all our policies and showed that for many common feasibility environments, they admit polynomial-time execution.

We have also provided a list of open problems and future directions for this work in Section EC.7.

References

Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43 (2):930–972, 2014.

Saeed Alaei, MohammadTaghi Hajiaghayi, and Vahid Liaghat. Online prophet-inequality matching with applications to ad allocation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 18–35, 2012.

Mohammad Reza Aminian, Will Ma, and Linwei Xin. Online job selection: Reward rate vs. remaining value. *Available at SSRN 4644495*, 2026a.

Mohammad Reza Aminian, Rad Niazadeh, and Pranav Nuti. Optimal bayesian online allocation of reusable resources. *Available at SSRN 6384158*, 2026b.

Nima Anari, Rad Niazadeh, Amin Saberi, and Ali Shameli. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 91–92, 2019.

Jackie Baek and Will Ma. Bifurcating constraints to improve approximation ratios for network revenue management with reusable resources. *Operations Research*, 70(4):2226–2236, 2022.

Santiago R Balseiro, Omar Besbes, and Dana Pizarro. Survey of dynamic resource-constrained reward collection problems: Unified model and analysis. *Operations Research*, 72(5):2168–2189, 2024.

Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

Niv Buchbinder, Kamal Jain, and Mohit Singh. Secretary problems and incentives via linear programming. *ACM SIGecom Exchanges*, 8(2):1–5, 2009.

Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.

Young-Bin Choe, James G Oxley, Alan D Sokal, and David G Wagner. Homogeneous multivariate polynomials with the half-plane property. *Advances in Applied Mathematics*, 32(1-2):88–187, 2004.

Youngbin Choe and David G Wagner. Rayleigh matroids. *Combinatorics, Probability and Computing*, 15(5):765–781, 2006.

Jean-François Coeurjolly, Jesper Møller, and Rasmus Waagepetersen. A tutorial on palm distributions for spatial point processes. *International Statistical Review*, 85(3):404–420, 2017.

José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 169–186, 2017.

José Correa, Andrés Cristi, Andrés Fielbaum, Tristan Pollner, and S Matthew Weinberg. Optimal item pricing in online combinatorial auctions. *Mathematical Programming*, 206, 2024.

Atanas Dinev and S Matthew Weinberg. Simple and optimal online contention resolution schemes for k-uniform matroids. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 39–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

Paul Dutting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582, 2020.

Farbod Ekbatani, Yiding Feng, Ian Kash, and Rad Niazadeh. Online job assignment. *arXiv preprint arXiv:2506.06893*, 2025.

Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet inequality for vertex and edge arrival models. *Math. Oper. Res*, 47(2):878–898, 2022.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM Journal on Computing*, 50(2):255–300, 2021.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Near-optimal bayesian online assortment of reusable resources. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 964–965, 2022.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Near-optimal bayesian online assortment of reusable resources. *Operations Research*, 72(5), 2024.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Robustness of online inventory balancing to inventory shocks. *arXiv preprint arXiv:2511.16044*, 2025.

Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.

Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE, 2011.

Rohan Ghuge, Sahil Singla, and Yifan Wang. Single-sample and robust online resource allocation. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1442–1453, 2025.

Vineet Goyal, Garud Iyengar, and Rajan Udwani. Asymptotically optimal competitive ratio for online allocation of reusable resources. *Operations Research*, 73(4):1897–1915, 2025.

Nick Gravin and Hongao Wang. Prophet inequality for bipartite matching: Merits of being simple and nonadaptive. *Mathematics of Operations Research*, 48(1):38–52, 2023.

Anupam Gupta and Marco Molinaro. How the experts algorithm can help solve lps online. *Mathematics of Operations Research*, 41 (4):1404–1431, 2016.

Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 71–80, 2004.

Ole J Heilmann and Elliott H Lieb. Theory of monomer-dimer systems. *Communications in mathematical Physics*, 25(3):190–232, 1972.

Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tight guarantees for multiunit prophet inequalities and online stochastic knapsack. *Operations Research*, 73(3):1703–1721, 2025a.

Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tightness without counterexamples: A new approach and new results for prophet inequalities. *Mathematics of Operations Research*, 2025b.

Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. *SIAM Journal on Computing*, 47(5):1939–1964, 2018.

Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.

Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4(197-266):1–2, 1978.

Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112, page 57. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.

David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.

Will Ma, Calum MacRury, and Pranav Nuti. Online matching and contention resolution for edge arrivals with vanishing probabilities. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 159–159, 2024.

Will Ma, Calum MacRury, and Jingwei Zhang. Online contention resolution schemes for network revenue management and combinatorial auctions. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, pages 100–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026.

Calum MacRury, Will Ma, and Nathaniel Grammel. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. *Operations Research*, 73(2):689–703, 2025.

Arnaud Marsiglietti and James Melbourne. Concentration inequalities for log-concave sequences. *International Mathematics Research Notices*, 2026(4):rnag023, 2026.

Michael Molloy and Bruce Reed. Hard-core distributions on matchings. In *Graph Colouring and the Probabilistic Method*, pages 247–264. Springer, 2002.

James Oxley. *Matroid Theory*. Oxford University Press, 2 edition, 2011.

Robin Pemantle. Towards a theory of negative dependence. *Journal of Mathematical Physics*, 41(3):1371–1390, 2000.

Tristan Pollner, Mohammad Roghani, Amin Saberi, and David Wajc. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 321–322, 2022.

Paat Rusmevichientong, Mika Sumida, and Huseyin Topaloglu. Dynamic assortment optimization for reusable products with random usage durations. *Management Science*, 66(7):2820–2844, 2020.

Paat Rusmevichientong, Mika Sumida, Huseyin Topaloglu, and Yicheng Bai. Revenue management with heterogeneous resources: Unit resource capacities, advance bookings, and itineraries over time intervals. *Operations Research*, 71(6):2196–2216, 2023.

Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.

Alexander Schrijver et al. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

Mohit Singh and Nisheeth K Vishnoi. Entropy, optimization and counting. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 50–59, 2014.

Damian Straszak and Nisheeth K Vishnoi. Maximum entropy distributions: Bit complexity and stability. In *Conference on Learning Theory*, pages 2861–2891. PMLR, 2019.

Eric Joseph Vigoda. *Sampling from Gibbs Distributions*. PhD thesis, EECS Department, University of California, Berkeley, 1999. URL https://www2.eecs.berkeley.edu/Pubs/TechRpts/1999/CSD-99-1088.pdf. Ph.D. dissertation; Technical Report No. UCB/CSD-99-1088.

Nisheeth K Vishnoi. *Algorithms for convex optimization*. Cambridge University Press, 2021.

Xinshang Wang, Van-Anh Truong, and David Bank. Online advance admission scheduling for services with customer preferences. *arXiv preprint arXiv:1805.10412*, 2018.

Ronald W Wolff. Poisson arrivals see time averages. *Operations research*, 30(2):223–231, 1982.

Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. SIAM, 2011.

This page is intentionally blank. The e-companion (EC) starts next page.

EC.1. Further related work

Our work is related to several lines of research in operations research and computer science, which we briefly summarize below.

Contention resolution schemes and order of arrivals. Contention resolution schemes (CRSs) were introduced as offline rounding primitives for the multilinear relaxation in submodular maximization; see the framework of Chekuri et al. (2014). OCRSs extend this primitive to online settings with independently active elements. The notion of OCRS for general downward-closed systems was first introduced by Feldman et al. (2016, 2021) and has since become a standard tool for converting ex-ante solutions into online policies. Since their introduction, OCRSs have also been studied in alternative arrival models, e.g., random-order (Adamczyk and Włodarczyk, 2018) and free-order (Ma et al., 2024) variants. Our work is complementary to these models: we impose a stronger symmetry requirement—permutation invariance of the *output distribution*—and study the resulting stationary OCRS policies, in contrast to the order-based standard OCRS policies.

CRSs for matchings, and applications. The contention resolution scheme problem for the matching feasibility system has been extensively studied. In the offline setting, initial results include those of Guruganesh and Lee (2018); Bruggmann and Zenklusen (2022), and the best known results may be found in Nuti and Vondrák (2025). The online setting (under edge and vertex arrivals) was first studied by Ezra et al. (2022), and their results were subsequently improved by MacRury et al. (2025). Random order arrival models have also been studied in Pollner et al. (2022) and MacRury et al. (2025). The first of these papers also discusses applications to pricing. In addition, vertex arrival models have also been explored: Ezra et al. (2022); MacRury and Ma (2024); Fu et al. (2021). A special regime of interest is when activation probabilities are tiny; this case has been studied by Nuti and Vondrák (2025); Ma et al. (2024). These results are connected to applications in revenue-management. Rank-L hypergraph matchings provide a canonical abstraction for network revenue management, and Ma et al. (2026) develop OCRS-type guarantees tailored to that model.

Bayesian online allocation of reusable resources. Motivated by the applications highlighted at the outset (e.g., cloud computing, rental marketplaces, and service systems), a growing line of work studies Bayesian online allocation and assortment optimization with *reusable* resources, allowing for possibly random usage durations. Earlier OR and CS literature on reusable capacity includes LP-based approaches for revenue management with reusable resources and reservations (Levi and Radovanović, 2010; Chen et al., 2017; Owen and Simchi-Levi, 2018), online Bayesian bipartite matching with reusable resources (Dickerson et al., 2018), and dynamic assortment optimization for rental products with random usage durations (Rusmevichientong et al., 2020); see also (Rusmevichientong et al., 2023; Ao et al., 2024) for richer models with advance bookings and reservations, (Sumita et al., 2022; Shanks et al., 2023; Dong et al., 2021) for stochastic task assignment and matching with reusable agents, and (Huang et al., 2024) for joint inventory and assortment planning. Related Bayesian/LP-guided policies appear for flexible products and network revenue management (Gallego et al., 2016; Ma et al., 2020; Baek and Ma, 2022; Anari et al., 2019), for admission control in overloaded loss networks with logarithmic regret (Xie et al., 2025), and for variants with volunteer nudging or recovering rewards (Manshadi and Rodilitz, 2022; Sumida, 2025). Closest to our work, Feng et al. (2022, 2024) study Bayesian online assortment/allocation of reusable resources under non-stationary arrivals and obtain a $1 - O(\sqrt{\log k/k})$ guarantee in the large-inventory regime via an LP-guided random-greedy (admission-control) policy. Our reduction in Section 7 complements this literature by showing that, for any downward-closed feasibility environment, designing an α-selectable S-OCRS immediately yields an α-approximate policy for the Bayesian online allocation of reusable products in the same environment but under stationary arrivals (i.e., homogeneous Poisson processes) for each type; our companion paper (Aminian et al., 2026b) leverages this viewpoint to tackle the problem under non-stationary arrivals (i.e., non-homogeneous Poisson processes). The main result of this companion paper solves an open problem in the k-uniform case, obtaining an α_k-selectable policy that removes the logarithmic loss in Feng et al. (2022) and matches the known lower-bound guarantee for the multi-unit, non-reusable counterpart (Alaei, 2014; Wang et al., 2018; Jiang et al., 2025a).

Reusable resources under adversarial arrivals & online learning. A complementary line of work studies online allocation of reusable resources when the request sequence is adversarial, under both matching and assortment models, generalizing classical adversarial models such as single-leg revenue management (Ball and Queyranne, 2009). In a semi-stationary model—where rewards are time-invariant and usage durations are i.i.d. or fixed—the greedy policy is $1/2$-competitive even for the more general assortment problem (Gong et al., 2022), its ratio provably improves with the degree of reusability (Baek and Wang, 2023), and LP/fluid-guided designs achieve asymptotically optimal competitive ratios as capacities grow (Feng et al., 2025; Goyal et al., 2025); see also (Simchi-Levi et al., 2025) for decaying rewards and (Delong et al., 2024) for online bipartite matching with small capacities, motivated by ride-sharing applications. Beyond the semi-stationary model, recent work studies multi-class arrivals and customer-dependent usage durations (Huo and Cheung, 2023b,a), online job assignment and its preemptive variant (Ekbatani et al., 2025; Aminian et al., 2026a), and—closest in spirit to our framework—online rounding schemes for k-rental problems, which separate the ex-ante fractional plan from an online feasibility-preserving rounding step (Nekouyan et al., 2025). Finally, when arrival or usage statistics are unknown, online learning and blind policies have been developed for nonstationary environments (Zhang and Cheung, 2025), assortment planning (Wu et al., 2025), and fair allocation of reusable resources (Liu and Hajiesmaili, 2025).

CRSs for multiple selection (or the uniform matroid). The k-uniform matroid (a.k.a. "multiple selection" or the "magician" setting) has a long history in Bayesian online allocation and mechanism design. The seminal OCRS constructions of Alaei (2014) characterize the optimal policy via a dynamic-programming/LP-duality viewpoint. Subsequent work sharpened the guarantees and refined the optimality story; see Wang et al. (2018) for the analysis in the vanishing regime, and the tight analysis in Jiang et al. (2025b,a) using factor-revealing and LP duality. More recently, Dinev and Weinberg (2024) propose a simpler adaptive and combinatorial algorithm with near-optimal asymptotic performance. Still, our algorithm for k-uniform matroid has (arguably) a simpler form and the analysis is quite short (with the right setup). Also, our stationary viewpoint is particularly compatible with online allocation of reusable resources, which is a natural extension of k-uniform setting and is extensively studied in the revenue management literature under both Bayesian and adversarial arrivals (see the dedicated discussion of reusable resources above).

The CRS problem for uniform matroids has also been studied in other related settings. In the offline setting, Yan (2011) established the optimal correlation gap for k-uniform matroids of $1 - \frac{e^{-k}k^k}{k!}$, which through an LP duality argument gives an optimal CRS. A simpler and more explicit scheme was given by Kashaev and Santiago (2023). In the random order setting, Arnosti and Ma (2023) obtain the optimal prophet inequality, and through a duality argument of Lee and Singla (2018) give an OCRS in the random order. We note the best selectability guarantees in the offline and random order online settings are the same.

CRSs for general matroids. The original papers introducing contention resolution schemes were particularly interested in the case of matroids, and Chekuri et al. (2014) already obtained the optimal selectability of $1-1/e$ in the offline case. The optimal contention resolution schemes in the online setting are obtained by Lee and Singla (2018) through a reduction to ex-ante prophet inequalities and an LP duality argument; their analysis is a refinement of the optimal algorithm in Kleinberg and Weinberg (2012) for the matroid prophet inequality problem. A recent result has established a connection between contention resolution schemes and the matroid secretary problem: Dughmi (2025). This has inspired further interest in sample-based contention resolution schemes for matroids: Fu et al. (2024); Feldman et al. (2026).

CRSs for knapsack constraints. The original OCRS framework includes an OCRS for knapsack constraints (Feldman et al., 2016, 2021) with selectability $3/2 - \sqrt{2}$. This OCRS works even against an almighty adversary. Jiang et al. (2025a) later introduced a tight OCRS for points x in the standard knapsack relaxation, with selectability $1/(3 + e^{-2})$. They allow for items to have stochastic sizes. Finally, Yoshinaga and Kawase (2023) consider an OCRS in a related setting where sizes are unknown to the algorithm.

Prophet inequalities and ex-ante relaxations. The prophet inequality literature originates from the classical work of Krengel and Sucheston (1978) and Samuel-Cahn (1984), and has since been developed for a variety of feasibility constraints, including matroids (Kleinberg and Weinberg, 2012), matchings (Ezra et al., 2022), knapsack constraints (Dutting et al., 2020), and more general downward-closed systems (Rubinstein, 2016; Rubinstein and Singla, 2026). In addition, various variants of this problem have been studied in the literature, for example, in random order Esfandiari et al. (2017); Ehsani et al. (2018), with access to samples Azar et al. (2014); Rubinstein et al. (2020); Cristi and Ziliotto (2024); Ghuge et al. (2025), or with the power of recourse/ cancellation Ekbatani et al. (2024). See Lucier (2017) for a comprehensive survey. Prophet inequalities are also tightly connected to sequential posted-price mechanisms and ex-ante relaxations in Bayesian mechanism design; see, e.g., Hartline and Roughgarden (2009); Chawla et al. (2010); Yan (2011); Correa et al. (2017); Lucier (2017); Alaei et al. (2019, 2022) and the nearly-optimal or improved posted-price/profit bounds versus optimal online policies (the so called "philosopher's inequality") in Anari et al. (2019); Braverman et al. (2025); Niazadeh et al. (2018); Papadimitriou et al. (2021). The link between OCRSs and ex-ante prophet inequalities is made explicit in Lee and Singla (2018), and underlies multiple optimal OCRS constructions (including the $1/2$-OCRS for matroids). Related to our motivation for permutation invariance, Kessel et al. (2022) study "stationary" prophet inequalities in an infinite-horizon setting; we emphasize that this notion is *different* from our stationary (permutation-invariant) OCRS. Finally, timing/order incentives have been studied in closely related online-selection models, e.g., in the secretary problem (Buchbinder et al., 2009; Kesselheim et al., 2015) and online auctions (Hajiaghayi et al., 2004).

Maximum entropy, Gibbs measures, and negative dependence. Maximum-entropy distributions go back to the classical maximum-entropy principle of Jaynes (1957). For a textbook treatment of Gibbs measures, see Georgii (2011). Algorithmically, maximum-entropy distributions and their associated exponential-family structure have been used as rounding primitives in combinatorial optimization; see, e.g., Gharan et al. (2011); Singh and Vishnoi (2014); Straszak and Vishnoi (2019). Our matroid results use negative dependence and Rayleigh-type properties. For background on negative dependence and correlation inequalities, see the survey of Pemantle (2000). The strong Rayleigh framework (and its polynomial viewpoint) was developed in Borcea et al. (2009), and Rayleigh matroids were studied in, e.g., Choe and Wagner (2006); Wagner (2008); Kahn and Neiman (2010). These tools motivate the KL-projection-to-a-Rayleigh-base-measure approach developed in Section EC.2. Our bipartite addability analysis uses classical correlation inequalities in the monomer–dimer model, going back to Heilmann and Lieb (1972).

EC.2. Beyond Maximum-Entropy: "Weakly Rayleigh" Matroids & Modified Maximum-Entropy

In this section, we develop a modified maximum-entropy approach for matroids. We first outline the new template in Section EC.2.1, and then show how it can be instantiated and analyzed in Section EC.2.2. Although the construction itself is conceptually simple, implementing the resulting S-OCRSs in polynomial time requires several ingredients, including solving relative-entropy (KL-divergence) optimization, sampling from certain distributions over bases, and computing conditional expectations for these distributions using counting oracles. We defer these algorithmic details to Section EC.6.

EC.2.1. Modified template: Rayleigh measures, minimizing KL divergence and thinning

Recall the matroid feasibility environment defined in Section 2. Let $\mathcal{M} = (E, \mathcal{I})$ be a rank-r matroid on E with collection of independent sets $\mathcal{I}$ and collection of bases $\mathcal{B}$. In this section, we focus on a broad class of matroids that admit a certain kind of negatively dependent distribution on bases, which we call *weakly Rayleigh*. We begin by defining the underlying notion of a *Rayleigh measure*.

DEFINITION EC.1 (RAYLEIGH DISTRIBUTION). We say that a probability measure μ_0 on $\{0,1\}^E$ is *Rayleigh* if for every weight vector $\boldsymbol{w} \in \mathbb{R}^E_{>0}$, the "*tilted measure with respect to* $\boldsymbol{w}$," defined as

$$\mu_{\boldsymbol{w}}(A) \;\propto\; \mu_0(A) \cdot \prod_{e \in A} w_e \tag{EC.1}$$

satisfies the *Rayleigh inequality:* for every $e \in E$ and every $T \subseteq E \setminus \{e\}$,

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T \subseteq B] \cdot \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e \in B] \ \geq\ \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T \cup \{e\} \subseteq B]. \tag{EC.2}$$

REMARK EC.1. It is straightforward to show that μ_0 is a Rayleigh measure[20] iff it satisfies the pairwise negative correlation under all tilts: for every $e, f \in E$ and every $\boldsymbol{w}$,

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e \in B] \cdot \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[f \in B] \ \geq\ \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[\{e, f\} \subseteq B].$$

(see Section EC.5.10 for a proof of this well-known result.)

DEFINITION EC.2 (WEAKLY RAYLEIGH MATROID). We say that $\mathcal{M} = (E, \mathcal{I})$ is *weakly Rayleigh* if there exists a Rayleigh probability measure μ_0 on $\mathcal{B}$ with full support (i.e., $\mu_0(B) > 0$ for all $B \in \mathcal{B}$).

We adopt the terminology of "weakly Rayleigh" from (Wagner, 2008), who studied weakly Rayleigh set systems. Many commonly encountered classes of matroids are weakly Rayleigh. For example, all matroids representable over $\mathbb{C}$ (e.g., the class of linear matroids over $\mathbb{R}$, which includes graphic and laminar matroids) are weakly Rayleigh: if $\mathcal{M}$ is represented by a matrix A, then the determinantal measure $\mu_0(B) \propto \det(A_B^* A_B)$ is Rayleigh. In addition, all matroids representable over $\mathbb{F}_2$ that do not have S_8 as a minor, all matroids of rank ≤ 3, and all matroids on at most 7 elements are weakly Rayleigh. We highlight that not all matroids are weakly Rayleigh. We prove this in Section EC.5.12 via showing that S_8 is a counterexample; this might be of independent interest.

We are now ready to present the main result of this section.

THEOREM EC.1 **(Optimal S-OCRS for weakly Rayleigh matroids).** *For the weakly Rayleigh matroid environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a distribution $\mu^* \in \Delta(\mathcal{F})$ such that Algorithm 1 (simulate-then-replace), when instantiated with μ^*, is a $\frac{1}{2}$-selectable S-OCRS. Moreover, $\frac{1}{2}$ is the best possible universal selectability constant for this class.*

Note that factor $\frac{1}{2}$ is already the best possible selectability constant for standard OCRS and rank-1 uniform matroids—therefore, our S-OCRS should be optimal.

Why we need a different method. The method in Section 4 constructs μ^* as the maximum-entropy distribution on feasible sets with marginals $\boldsymbol{p} = \alpha \boldsymbol{x}$ and then proves addability bounds under μ^*. For matroids, this "direct" maximum-entropy approach can fail: the maximum-entropy distribution on independent sets may place most of its mass on configurations that make certain elements rarely addable. In Section EC.5.11, we give an explicit counterexample for a simple graphic matroid (known as the "hat" instance in the literature).

A KL-divergence template on bases. For weakly Rayleigh matroids we instead work with distributions on *bases*. Starting from a Rayleigh base measure μ_0, we compute a new base distribution with certain marginals,

[20] Such measures μ_0 are also referred to as h-NLC+, S-MRR$_2$, or NC+ (Pemantle, 2000) in the literature. Choe and Wagner (2006) show that this notion can also be referred to as BLC[2] or BLC[3], and Kahn and Neiman (2010) later show that this is also equivalent to BLC[4] and BLC[5].

and then obtain an independent-set distribution by a simple *thinning procedure*. In the counterexample from Section EC.5.11, for instance, this thinning step deletes every u–v path with probability at least $1/2$, and is essential for the new approach to work.

Fix an activation vector $\boldsymbol{x} \in \mathcal{P}$ and any Rayleigh base measure μ_0. At a high-level, our new framework for design and analysis of S-OCRS algorithms follows a *modified* five-step template:

(i) Find $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ such that $\boldsymbol{q} \geq \boldsymbol{x}$ coordinatewise,

(ii) Build a Rayleigh distribution μ on bases with marginals $\boldsymbol{q}$ by minimizing *KL divergence* from μ_0 subject to prescribed marginals $\boldsymbol{q}$ (a "relative-entropy" analogue of maximum entropy)[21],

(iii) $1/2$-thin inside μ (independently across elements) and show that the resulting independent-set distribution satisfies Constraint STATIONARY-IMPLEMENTABILITY with caps $\boldsymbol{q}$,

(iv) Thin further—independently coordinatewise—to match the desired activation vector $\boldsymbol{x}$ (i.e., to have marginals exactly equal to $\boldsymbol{x}/2$) and obtain a distribution μ^* on independent sets,

(v) Run the simulate-then-replace policy (Algorithm 1) with μ^* as input.

To show the correctness of our S-OCRS, we need to show that the final witness distribution μ^* is feasible for the stationary OCRS LP with $\alpha = \frac{1}{2}$. Note that μ^* has marginals exactly equal to $\frac{1}{2}\boldsymbol{x}$. Therefore, it only remains to show that μ^* satisfies (STATIONARY-IMPLEMENTABILITY).[22]

EC.2.2. Implementing the modified template: proof of Theorem EC.1

We now explain each step of the template and why μ^* satisfies (STATIONARY-IMPLEMENTABILITY).

Step (i): Finding a dominating base point. We first note that every point in the independent-set polytope can be dominated by a point in the base polytope. This is a standard fact in matroid theory; see, e.g., Schrijver et al. (2003, Chapter 44). We still include a proof for completeness in Section EC.5.13.

LEMMA EC.1 **(Dominating base point).** *For every point $\boldsymbol{x} \in \mathcal{P}$, there exists a dominating base point $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ such that $\forall e \in E:\ q_e \geq x_e$.*

Fix such a $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ for the remainder of this subsection.

Step (ii): KL projection onto a Rayleigh base measure. We now build a Rayleigh distribution μ on bases with marginals $\boldsymbol{q}$. When $\boldsymbol{q}$ lies in the relative interior of $\mathcal{P}_{\mathcal{B}}$, this can be done via a KL-divergence projection. Concretely, consider this convex program over distributions ν on $\mathcal{B}$:

$$\min_{\nu \in \Delta(\mathcal{B})} \quad \mathrm{KL}(\nu \| \mu_0) \tag{MIN-KL}$$

$$\text{s.t.} \quad \mathbb{P}_{B\sim\nu}[e \in B] = q_e \qquad \forall e \in E\ , \tag{EC.3}$$

[21] For several classes of matroids (graphic matroids, for instance), the uniform distribution over bases is already Rayleigh, so this step is the same as finding the maximum entropy distribution with marginals $\boldsymbol{q}$.

[22] If $\boldsymbol{x} \in b\mathcal{P}$, we may choose $\boldsymbol{q} \geq \frac{\boldsymbol{x}}{b}$, $b/(1+b)$-thin inside μ, and then thin further coordinatewise by $\frac{x_e}{bq_e}$ to match the desired activation vector $\boldsymbol{x}$.

where $\mathrm{KL}(\nu\|\mu_0) := \sum_{B\in\mathcal{B}} \nu(B)\log\big(\nu(B)/\mu_0(B)\big)$. By strict convexity and Slater's condition, the minimizer is unique; moreover, the KKT conditions (as in the proof of Lemma 1) imply that the unique minimizer ν^* has the tilted form $\nu^* = \mu_{\boldsymbol{w}}$ for some $\boldsymbol{w} \in \mathbb{R}^E_{>0}$ as in (EC.1).[23] In particular, we obtain a tilted measure $\mu_{\boldsymbol{w}}$ satisfying

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e\in B] = q_e \qquad \forall e\in E\ ,$$

and, since μ_0 is Rayleigh, we also have the Rayleigh inequality (EC.2) under $\mu_{\boldsymbol{w}}$.

If $\boldsymbol{q}$ lies on the boundary of $\mathcal{P}_{\mathcal{B}}$, we approximate it by a sequence of interior points and pass to a limit. Since 2^E is finite, this limiting procedure produces a probability measure on $\mathcal{B}$ with marginals $\boldsymbol{q}$; Rayleigh-ness is preserved under limits. Let μ denote the resulting base measure with the desired marginals and the Rayleigh property (regardless of whether $\boldsymbol{q}$ is on the boundary or interior of $\mathcal{P}_{\mathcal{B}}$).

Step (iii): $1/2$-thinning inside bases. Sample a base $B\sim\mu$ and then include each element of B independently with probability $\frac{1}{2}$. Let S denote this random set. Then $S\subseteq B$ is independent, and for each $e\in E$,

$$\mathbb{P}[e\in S] = \tfrac{1}{2}\,\mathbb{P}_{B\sim\mu}[e\in B] = \tfrac{1}{2}\,q_e. \tag{EC.4}$$

We claim that S satisfies (STATIONARY-IMPLEMENTABILITY) with caps $\boldsymbol{q}$, meaning that for every $e\in E$ and every $T\subseteq E\setminus\{e\}$ with $\mathbb{P}[S_{-e}=T]>0$, we have:

$$\mathbb{P}[e\in S \mid S_{-e}=T] \le q_e. \tag{EC.5}$$

Indeed, conditioning on $S_{-e}=T$ leaves only two possibilities for S: either $S=T$ or $S=T\cup\{e\}$. Moreover, since the thinning is independent inside a base of fixed size r, for any $U\subseteq E$ we have $\mathbb{P}[S=U] = 2^{-r}\cdot \mathbb{P}_{B\sim\mu}[U\subseteq B]$, and therefore

$$\begin{aligned}
\mathbb{P}[e\in S\mid S_{-e}=T] &= \frac{\mathbb{P}[S=T\cup\{e\}]}{\mathbb{P}[S=T]+\mathbb{P}[S=T\cup\{e\}]}\\
&= \frac{\mathbb{P}_{B\sim\mu}[T\cup\{e\}\subseteq B]}{\mathbb{P}_{B\sim\mu}[T\subseteq B]+\mathbb{P}_{B\sim\mu}[T\cup\{e\}\subseteq B]}\\
&\le \frac{\mathbb{P}_{B\sim\mu}[T\cup\{e\}\subseteq B]}{\mathbb{P}_{B\sim\mu}[T\subseteq B]} \ \le\ \mathbb{P}_{B\sim\mu}[e\in B]\ =\ q_e,
\end{aligned}$$

where the penultimate inequality is the Rayleigh inequality (EC.2) applied to (T,e). This proves (EC.5).

Step (iv): Coordinatewise thinning to match $\boldsymbol{x}$. The caps in (EC.5) are stated in terms of q_e, which may exceed the desired activations x_e. Define thinning probabilities $s_e\in[0,1]$ by

$$s_e := \begin{cases} x_e/q_e, & q_e>0,\\ 0, & q_e=0.\end{cases}$$

Let $(C_e)_{e\in E}$ be independent Bernoulli variables with $\mathbb{P}[C_e=1]=s_e$, independent of S, and set $R:=\{e\in S: C_e=1\}$. Then $R\subseteq S$ is independent and, using (EC.4),

$$\mathbb{P}[e\in R] = \mathbb{P}[e\in S]\mathbb{P}[C_e=1] = \tfrac{1}{2}\,q_e\,s_e = \tfrac{1}{2}\,x_e,$$

[23] Note that a probability distribution of this tilted form is not, in general, a Gibbs distribution.

which is the selectability constraint with $\alpha = \frac{1}{2}$.

For stationary implementability, fix e and T with $\mathbb{P}[R_{-e} = T] > 0$. Since C_e is independent of the event $\{R_{-e} = T\}$, we have

$$\mathbb{P}[e \in R \mid R_{-e} = T] = s_e \cdot \mathbb{P}[e \in S \mid R_{-e} = T]. \tag{EC.6}$$

Let $U := S_{-e}$. Conditioned on U, the event $\{R_{-e} = T\}$ depends only on the coins $\{C_f\}_{f \neq e}$, and is therefore conditionally independent of the event $\{e \in S\}$. Hence

$$\mathbb{P}[e \in S \mid U,\ R_{-e} = T] = \mathbb{P}[e \in S \mid U] \le q_e,$$

where the inequality is (EC.5) (applied to the realized value of U). Taking conditional expectations given $\{R_{-e} = T\}$ yields $\mathbb{P}[e \in S \mid R_{-e} = T] \le q_e$. Plugging into (EC.6) gives

$$\mathbb{P}[e \in R \mid R_{-e} = T] \le s_e q_e = x_e,$$

which is exactly the stationary implementability constraint.

Step (v): Conclusion. Let $\mu^* := \mathrm{Law}(R)$. We have shown that μ^* satisfies the constraints of stationary OCRS LP, i.e., (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY), with $\alpha = \frac{1}{2}$. The theorem now follows from Proposition 2: for every arrival order, the simulate-then-replace procedure (Algorithm 1) implements μ^* online, yielding a $\frac{1}{2}$-selectable S-OCRS.

EC.3. Beyond Maximum-Entropy: Knapsack Constraints & Modified Maximum-Entropy

In this section, we develop a modified maximum-entropy approach for knapsack constraints. Recall from Section 2 the knapsack constraint environment defined on a set of items E with sizes $a_e \in (0, 1]$. For knapsack constraints, like for matroids, the direct maximum-entropy approach can fail, and in Section EC.5.14, we give an explicit counterexample following Feldman et al. (2016, 2021). The main issue is that maximum entropy distributions can make it very hard for very heavy items to be addable.

We therefore need a different strategy. Similar to Feldman et al. (2016, 2021) and Dutting et al. (2020), our idea is to split E into *heavy* items and *light* items. We then use two different stationary OCRS witness distributions for the heavy items and light items, and mix the two. For the heavy items, we select at most one item. For the light items, we use the maximum-entropy approach. We defer the algorithmic details of implementing our scheme in polynomial time to Section EC.6.

THEOREM EC.2 **(S-OCRS for knapsack constraints).** *For the knapsack constraint environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a distribution $\mu^* \in \Delta(\mathcal{F})$ such that Algorithm 1 (simulate-then-replace), when instantiated with μ^*, is a $\frac{1}{4}$-selectable S-OCRS.*

Before presenting the proof, we make a few remarks regarding the context of our result. Following Feldman et al. (2016, 2021), we work in a setting in which items have deterministic sizes. Note, however, that our schemes, unlike theirs, do not work against almighty adversaries. Jiang et al. (2025a) work with more general stochastically sized items. It is possible to extend our result to stochastic sizes and obtain a $1/5$-selectable S-OCRS. This then yields a $1/5$-competitive prophet inequality, matching the results of Dutting et al. (2020).

Proof. Partition the ground set into heavy and light items:

$$E_H := \{e \in E : a_e > 1/2\}, \qquad E_L := \{e \in E : a_e \le 1/2\},$$

and define

$$X := \sum_{e \in E_H} x_e, \qquad B := \sum_{e \in E_L} a_e x_e.$$

Step (i): The heavy-item case. Let $c_H := \frac{1}{1+X}$, and define $\mu^H \in \Delta(\mathcal{F})$ by

$$\mu^H(\emptyset) = c_H, \qquad \mu^H(\{e\}) = c_H x_e \quad (e \in E_H),$$

and $\mu^H(T) = 0$ for all other feasible sets T. This is clearly a probability distribution over feasible sets, with each heavy item getting selectability c_H. Furthermore,

$$\mathbb{P}_{S\sim\mu^H}[e \in S \mid S_{-e} = \emptyset] = \frac{c_H x_e}{c_H + c_H x_e} \le x_e,$$

and all remaining constraints are trivial. Hence (STATIONARY-IMPLEMENTABILITY) holds. Note that this measure is also in fact a solution to a maximum entropy program.

Step (ii): The light-item case. Let

$$\mathcal{F}_L := \{T \subseteq E_L : a(T) \le 1\}$$

be the light-item knapsack family. The projection $x_L = (x_e)_{e\in E_L}$ belongs to $\mathrm{conv}\{\mathbf{1}_T : T \in \mathcal{F}_L\}$. We may assume without loss of generality that $B > 0$. Define $c_L := \frac{1}{1+2B}$, and let $p_e := c_L x_e$ for $e \in E_L$, and let $\mu^L \in \Delta(\mathcal{F}_L)$ be the maximum-entropy distribution with marginals p:

$$\mathbb{P}_{S\sim\mu^L}[e \in S] = p_e = c_L x_e.$$

For $S \sim \mu^L$, define the addability event $\mathsf{Add}(e) := \{S_{-e} \cup \{e\} \in \mathcal{F}_L\}$. Based on the discussion in Section 4.2, we only need to verify the condition in (ADDABILITY), so it suffices to show

$$\mathbb{P}_{S\sim\mu^L}[\mathsf{Add}(e)] \ge c_L.$$

If $\mathsf{Add}(e)$ fails, then $a(S_{-e}) > 1 - a_e$. Since $e \in E_L$, we have $a_e \le 1/2$, and therefore $1 - a_e \ge 1/2$. By Markov's inequality,

$$\mathbb{P}_{S\sim\mu^L}[\mathsf{Add}(e)^c] \le \frac{\mathbb{E}[a(S_{-e})]}{1-a_e} \le \frac{\mathbb{E}[a(S)]}{1-a_e} \le 2\mathbb{E}[a(S)] = 2\sum_{f\in E_L} a_f p_f = 2c_L \sum_{f\in E_L} a_f x_f = 2c_L B.$$

Therefore,

$$\mathbb{P}_{S\sim\mu^L}[\mathsf{Add}(e)] \ge 1 - 2c_L B = \frac{1}{1+2B} = c_L,$$

as needed.

Step (iii): Mixing the two cases. We now combine the two cases, viewing μ^L as a distribution in $\Delta(\mathcal{F})$. Let

$$\theta := \frac{c_L}{c_H + c_L}, \qquad \mu := \theta\mu^H + (1-\theta)\mu^L.$$

An equivalent form of (STATIONARY-IMPLEMENTABILITY) is

$$(1-x_e)\mu(T \cup \{e\}) \le x_e \mu(T),$$

which is linear in μ. Since both μ^H and μ^L satisfy (STATIONARY-IMPLEMENTABILITY), the mixture μ also satisfies it. The marginals of heavy and light items under μ are

$$\mathbb{P}_{S\sim\mu}[e \in S] = \theta c_H x_e \text{ for } e \in E_H \quad \text{ and } \quad \mathbb{P}_{S\sim\mu}[e \in S] = (1-\theta)c_L x_e \text{ for } e \in E_L.$$

By the choice of θ,

$$\theta c_H = (1-\theta)c_L = \frac{c_H c_L}{c_H + c_L} = \frac{1}{c_H^{-1} + c_L^{-1}}.$$

Since $c_H^{-1} = 1 + X$ and $c_L^{-1} = 1 + 2B$, we obtain the common selectability factor $\alpha = \frac{1}{2+X+2B}$. Finally,

$$X + 2B = \sum_{e \in E_H} x_e + 2\sum_{e \in E_L} a_e x_e \le 2\sum_{e \in E} a_e x_e \le 2.$$

Thus $\alpha \ge \frac{1}{4}$, as desired.[24] □

To obtain an S-OCRS in the more general case where items may have stochastic sizes (that are revealed before a decision must be made on whether to accept them), we associate each item with a set of elements E_i with each element corresponding to different possible size of the item. The elements within each E_i are active mutually exclusively rather than independently. Let x_e and a_e be the probability and size of each element e, and assume $\sum_{e\in E_i} x_e \le 1$ for every i and $\sum_e a_e x_e \le 1$. Feasible sets contain at most one outcome from each E_i, with stationarity defined by conditioning on the selections outside E_i. Then, we may obtain a $1/5$-selectable S-OCRS by leaving the heavy-item distribution unchanged and for the light items, taking the maximum-entropy distribution on

$$\{S \subseteq E_L : a(S) \le 1,\ |S \cap E_i| \le 1\ \forall i\}$$

with marginals $c_L x_e$, where $c_L = (2+2B)^{-1}$.

Negative results. We also establish an impossibility result, showing that in the tight instance in Jiang et al. (2025a), no S-OCRS can obtain a selectability better than $3/10$. Note that this is strictly less than the selectability of the OCRS of Jiang et al. (2025a). Recall the standard knapsack relaxation:

$$\widetilde{\mathcal{P}} = \left\{ \boldsymbol{x} \in [0,1]^E : \sum_{e\in E} a_e x_e \le 1 \right\}.$$

[24] If $\boldsymbol{x} \in b\mathcal{P}$, then $X + 2B \le 2b$.

LEMMA EC.2 **(Knapsack upper-bound).** *For every $\varepsilon \in (0, 1/4)$ there exists a knapsack instance, and a point $x \in \widetilde{\mathcal{P}}$ such that any distribution $\mu \in \Delta(\mathcal{F})$ satisfying the stationary constraints (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY) has selectability at most $3/10 + O(\varepsilon)$. In particular, letting $\varepsilon \to 0$ implies that every S-OCRS has selectability at most $3/10$.*

The proof is postponed to Section EC.5.15. We note that there is also a knapsack instance and a point $x \in \mathcal{P}$ (i.e., in the actual convex hull rather than the relaxation) where the selectability is at most $1/3$: Consider $m = 1/\varepsilon^2$ large items with size 1 and probability of activation $(1-\varepsilon)/m$, and m small items with size $1/m$ and probability of activation ε, with ε close to 0.

EC.4. Beyond Maximum-Entropy: General Downward Closed Feasibility Environments

In this section, we let $(E, \mathcal{F})$ be any finite downward-closed feasibility environment with $|E| = N$. Our goal is to establish the following result, which, to the best of our knowledge, is the first on OCRSs for general downward-closed feasibility environments.

THEOREM EC.3 **(Optimal S-OCRS for general downward-closed feasibility environments).** *For a general downward closed feasibility environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a distribution $\mu^* \in \Delta(\mathcal{F})$ such that Algorithm 1 (simulate-then-replace), when instantiated with μ^*, is a $1/(2\sqrt{|E|}+1)$-selectable S-OCRS. Moreover, for every positive integer m, there is a downward-closed feasibility environment on $N = m^2$ elements and a point in $\mathcal{P}$ for which every OCRS has selectability at most $2/\sqrt{N}$.*

While previously, it was known that there exists an $O((\log N)^2)$ prophet inequality in this setting (Rubinstein, 2016; Rubinstein and Singla, 2026), these results do not extend to OCRSs, which are dual to *ex-ante* prophet inequalities.

Positive result. To prove the positive results, we let $q := \frac{1}{\sqrt{N}}$, and partition the ground set into more likely and less likely items:

$$E_H := \{e \in E : x_e > q\}, \qquad E_L := \{e \in E : x_e \le q\}.$$

Since $x \in \mathcal{P}$, there exists a distribution $\nu \in \Delta(\mathcal{F})$ such that

$$\mathbb{P}_{I \sim \nu}[e \in I] = x_e \qquad \forall e \in E.$$

We will use two different stationary OCRS witness distributions for the more likely items and less likely items, and mix the two. For the more likely items, we will sample a feasible set from ν, and then keep each element with probability q. For the light items, we will select at most one item.

Proof of Theorem EC.3. We start by working with the more likely item case.

Step (i): The more likely item case. Define μ^H as follows. Sample $I \sim \nu$, discard all items in E_L, and independently keep each item of $I \cap E_H$ with probability q. Let R^H be the resulting set, and set $\mu^H := \mathrm{Law}(R^H)$. Since $R^H \subseteq I$ and $\mathcal{F}$ is downward closed, μ^H is supported on $\mathcal{F}$. Moreover,

$$\mathbb{P}_{S \sim \mu^H}[e \in S] = \begin{cases} q x_e, & e \in E_H, \\ 0, & e \in E_L. \end{cases}$$

We claim that μ^H satisfies (STATIONARY-IMPLEMENTABILITY). If $e \in E_L$, this is trivial, since e is never selected. If $e \in E_H$, let D_e be the independent retention coin of e. Conditioning on an event of the form $(R^H)_{-e} = T$ does not reveal D_e. Hence, whenever the conditioning event has positive probability,

$$\mathbb{P}[e \in R^H \mid (R^H)_{-e} = T] = q\mathbb{P}[e \in I \mid (R^H)_{-e} = T] \le q \le x_e,$$

where the last inequality uses $e \in E_H$. Thus μ^H is stationary-implementable with respect to x.

Step (ii): The less likely item case. Define μ^L by

$$\mu^L(\{e\}) = \frac{qx_e}{1+q} \quad (e \in E_L), \qquad \mu^L(\emptyset) = 1 - \frac{q\sum_{e \in E_L} x_e}{1+q},$$

and $\mu^L(T) = 0$ for all other feasible sets T. This is a probability distribution, because

$$\frac{q\sum_{e \in E_L} x_e}{1+q} \le \frac{q^2}{1+q} \cdot |E_L| \le \frac{1}{1+q}.$$

The marginals of μ^L are

$$\mathbb{P}_{S\sim\mu^L}[e \in S] = \begin{cases} \dfrac{q}{1+q} x_e, & e \in E_L, \\ 0, & e \in E_H. \end{cases}$$

We now check (STATIONARY-IMPLEMENTABILITY). The only important case is when $e \in E_L$ and $T = \emptyset$. Since $\mu^L(\emptyset) \ge \frac{q}{1+q}$, we have

$$\mathbb{P}_{S\sim\mu^L}[e \in S \mid S_{-e} = \emptyset] = \frac{qx_e}{(1+q)\mu^L(\emptyset) + qx_e} \le \frac{qx_e}{(1+q)\mu^L(\emptyset)} \le x_e.$$

Thus μ^L is stationary-implementable with respect to x.

Step (iii): Mixing the two cases. Let

$$\mu := \frac{1}{2+q}\mu^H + \frac{1+q}{2+q}\mu^L.$$

Since (STATIONARY-IMPLEMENTABILITY) is linear when written in the form:

$$(1 - x_e)\mu(T \cup \{e\}) \le x_e \mu(T),$$

and since both μ^H and μ^L satisfy (STATIONARY-IMPLEMENTABILITY), the mixture μ also satisfies (STATIONARY-IMPLEMENTABILITY). For $e \in E_H$,

$$\mathbb{P}_{S\sim\mu}[e \in S] = \frac{1}{2+q} qx_e = \frac{x_e}{2\sqrt{N}+1}.$$

For $e \in E_L$,

$$\mathbb{P}_{S\sim\mu}[e \in S] = \frac{1+q}{2+q} \cdot \frac{q}{1+q} x_e = \frac{x_e}{2\sqrt{N}+1}.$$

Thus the construction gives a $1/(2\sqrt{N}+1)$-selectable stationary OCRS. $\square$

Negative result. The next result shows that the previous lower bound has the right order of magnitude even for usual OCRSs. This is a stronger impossibility statement than an upper bound for just stationary OCRSs. The proof is closely related to the Kleinberg and Weinberg (2012) proof of impossibility of beating $O\left(\frac{1}{\sqrt{q}}\right)$ for prophet inequalities for the intersection of q matroids.

LEMMA EC.3 **(General downward closed feasibility environments upper-bound).** *Let m be a positive integer. There is a downward-closed feasibility environment $(E,\mathcal{F})$ with $|E| = N = m^2$ and a point $x \in \mathcal{P}$ such that every OCRS has selectability at most $2/\sqrt{N}$ on this instance.*

EC.5. Deferred Proofs and Technical Details

EC.5.1. Proof of Proposition 1 (S-OCRS $\Longrightarrow$ recurring OCRS)

Proof. Maintain a feasible set $\widehat{S} \in \mathcal{F}$ as the state of the system, representing the set of elements that are currently selected, together with simulations for all elements that have never appeared in the sequence. We aim to keep this state distributed as μ, a solution to the S-OCRS LP, at all times. Initialize by sampling $\widehat{S} \sim \mu$. Whenever element e renews (its current epoch ends and the next epoch starts), first "forget" its previous membership by setting $T := \widehat{S}_{-e}$. Next, if the new epoch of e is inactive, keep $\widehat{S} \leftarrow T$. If the new epoch is active, accept it with probability $\mathbb{P}_{S\sim\mu}[e \in S \mid S_{-e} = T]/x_e$, and set $\widehat{S} \leftarrow T \cup \{e\}$ upon acceptance. This is exactly the simulate-then-replace update that preserves the law of $\widehat{S}$. Since the imaginary state distribution stays equal to μ after every renewal (regardless of the renewal order), each active epoch of e is accepted with probability $\mathbb{P}_{S\sim\mu}[e \in S]/x_e \geq \alpha$. Moreover, the algorithm is feasible in every realization. □

EC.5.2. Proof of Proposition 2 (S-OCRS $\Longleftrightarrow$ S-OCRS LP)

Proof. The proof consists of two steps:

(Stationary OCRS $\Rightarrow$ LP). Let a stationary OCRS at $\boldsymbol{x}$ be given, and let μ denote the common output distribution (Definition 4). Selectability of the OCRS is exactly (SELECTABILITY). To see why (STATIONARY-IMPLEMENTABILITY) holds, fix $e \in E$ and $T \in \mathcal{F}$ with $e \notin T$ and consider running the same OCRS under an arrival order in which e arrives *last*. By stationarity, the final selected set still has law μ. Moreover, in this run the event $\{S_{-e} = T\}$ is fully determined before e is revealed, hence it is independent of the activation of e. Since the OCRS can select e only when e is active (an event of probability x_e), we obtain $\mathbb{P}[e \in S \mid S_{-e} = T] \leq x_e$, establishing (STATIONARY-IMPLEMENTABILITY).

(LP $\Rightarrow$ Stationary OCRS). Assume $\mu \in \Delta(\mathcal{F})$ satisfies (STATIONARY-IMPLEMENTABILITY). We describe an online procedure that (for *any* arrival order π) produces an output set S with $\mathrm{Law}(S) = \mu$, while never selecting inactive elements. Process the elements in the arrival order π, and let $\mathcal{H}$ denote the (random) history of past selection decisions. When element e arrives, define the target conditional inclusion probability

$$q_e := \mathbb{P}_{S\sim\mu}[e \in S \mid \mathcal{H}],$$

where the conditioning event $\mathcal{H}$ is interpreted as the event that S agrees with the already-realized past selections on the elements that have arrived so far. q_e is always bounded by x_e; indeed, $\mathbb{P}[e \in S \mid \mathcal{H}]$ is an average over configurations of S_{-e}, each of which has conditional probability at most x_e by (STATIONARY-IMPLEMENTABILITY).

Now, upon observing whether e is active, act as follows: if $e \notin A$, reject it; if $e \in A$, accept it with probability q_e/x_e (using fresh internal randomness). Since $q_e \le x_e$, this is well-defined, and it ensures that $\mathbb{P}[e \in S \mid \mathcal{H}] = q_e$, i.e., the online process matches the conditionals of μ. Therefore the final output set has distribution μ for every π, hence the procedure is stationary. Finally, (SELECTABILITY) guarantees α-selectability of the resulting stationary OCRS. □

EC.5.3. Maximum entropy S-OCRS for rank-L hypergraph matchings

Recall from Section 2 the rank-L hypergraph matching environment defined on a hypergraph $H = (V, E)$ of rank at most L. Our next application of the maximum-entropy template gives a simple S-OCRS for this setting. For completeness, recall that a set $S \subseteq E$ is feasible if it is a (hypergraph) matching, i.e., its hyperedges are pairwise vertex-disjoint. Accordingly,

$$\mathcal{F} := \Big\{ S \subseteq E : \ \forall v \in V, \ \big|\{e \in S : \ v \in e\}\big| \le 1 \Big\}.$$

Note that for every $\boldsymbol{x}$ in the feasibility polytope $\mathcal{P} = \mathrm{conv}\{\mathbf{1}_S : S \in \mathcal{F}\}$, we must have:

$$\sum_{e \in \delta(v)} x_e \le 1 \qquad \forall v \in V, \tag{EC.7}$$

where $\delta(v)$ denotes the set of hyperedges incident to v.

THEOREM EC.4 (**S-OCRS for rank-L hypergraph matchings**). *For the rank-L hypergraph matching environment $(E, \mathcal{F})$ and every $\boldsymbol{x} \in \mathcal{P}$, there exists a Gibbs distribution μ^* such that Algorithm 2, when instantiated with μ^*, is a $\frac{1}{L+1}$-selectable S-OCRS.*

Proof. Fix any $\boldsymbol{x} \in \mathcal{P}$ and set $\alpha := \frac{1}{L+1}$ and $\boldsymbol{p} := \alpha \boldsymbol{x}$. Let μ^* be the maximum-entropy distribution—that is, the solution to (MAX-ENTROPY)—with marginals $\boldsymbol{p}$. Based on the discussion in Section 4.2, we only need to verify the condition in (ADDABILITY). Fix a hyperedge $e \in E$. The addability event

$$\mathsf{Add}(e) = \{S_{-e} \cup \{e\} \in \mathcal{F}\}$$

is exactly the event that no selected hyperedge in S_{-e} intersects e.

For each vertex $v \in V$, define

$$\mathsf{Matched}(v) := \Big\{ \exists f \in S \text{ with } v \in f \Big\}.$$

Since S is always a matching, for any fixed vertex v the events $\{f \in S\}$ over distinct $f \in \delta(v)$ are disjoint. Therefore,

$$\mathbb{P}_{S \sim \mu^*}\big[\mathsf{Matched}(v)\big] = \sum_{f \in \delta(v)} \mathbb{P}_{S \sim \mu^*}[f \in S] = \sum_{f \in \delta(v)} p_f. \tag{EC.8}$$

Using $p_f = \alpha x_f$ and (1), we get

$$\mathbb{P}_{S\sim\mu^*}\big[\mathsf{Matched}(v)\big] \;=\; \alpha \sum_{f\in\delta(v)} x_f \;\le\; \alpha. \tag{EC.9}$$

Finally, $\mathsf{Add}(e)$ fails only if at least one vertex of e is occupied by some hyperedge in S_{-e}. Thus, by a union bound and (EC.9),

$$\mathbb{P}_{S\sim\mu^*}[\mathsf{Add}(e)] \;\ge\; 1-\sum_{v\in e}\mathbb{P}\big[\mathsf{Matched}(v)\big] \;\ge\; 1-|e|\alpha \;\ge\; 1-L\alpha \;=\; \alpha.$$

This is exactly (ADDABILITY). The conclusion follows. □

EC.5.4. Proof of Lemma 2 (Gibbs positive correlation for bipartite matchings)

Fix the weight vector $\boldsymbol{w} = (w_e)_{e\in E}$ induced by a Gibbs distribution μ. For any subgraph H of G, define the partition function

$$\mathsf{Z}(H) \;:=\; \sum_{M\in\mathcal{M}(H)} \prod_{e\in M} w_e,$$

where $\mathcal{M}(H)$ denotes the set of matchings in H. For a vertex $z \in U\cup V$, write $G-z$ for the graph obtained by deleting z and all incident edges, and similarly define $G-u-v$.

Under the Gibbs distribution μ, the probability that a vertex is unmatched is a ratio of partition functions:

$$\mathbb{P}_{S\sim\mu}[\overline{\mathsf{Matched}(z)}] \;=\; \frac{\mathsf{Z}(G-z)}{\mathsf{Z}(G)}. \tag{EC.10}$$

Indeed, z is unmatched if and only if S uses no edge incident to z, i.e., if and only if S is a matching of $G-z$. Similarly,

$$\mathbb{P}_{S\sim\mu}[\overline{\mathsf{Matched}(u)}\cap\overline{\mathsf{Matched}(v)}] \;=\; \frac{\mathsf{Z}(G-u-v)}{\mathsf{Z}(G)}. \tag{EC.11}$$

Thus, the lemma reduces to showing that for any bipartite $G = (U\cup V, E)$, any $u\in U$, $v\in V$, and any weights $w_e > 0$, we have

$$\mathsf{Z}(G)\,\mathsf{Z}(G-u-v) \;\ge\; \mathsf{Z}(G-u)\,\mathsf{Z}(G-v). \tag{EC.12}$$

This inequality is well known for monomer–dimer systems (see for instance Lemma 22.15 in Molloy and Reed (2002)); for completeness, we give a simple combinatorial proof.

Proof of inequality (EC.12). Let $w(M) := \prod_{e\in M} w_e$ denote the weight of a matching. We construct an injective map

$$\Phi:\; \mathcal{M}(G-u)\times\mathcal{M}(G-v) \;\longrightarrow\; \mathcal{M}(G)\times\mathcal{M}(G-u-v),$$

with the property that for every $(M_u, M_v) \in \mathcal{M}(G-u)\times\mathcal{M}(G-v)$, we have:

$$w(M_u)w(M_v) = w(\Phi_1(M_u,M_v))\,w(\Phi_2(M_u,M_v)).$$

Summing over all pairs then gives

$$\begin{aligned}\mathsf{Z}(G-u)\,\mathsf{Z}(G-v) = \sum_{(M_u,M_v)} w(M_u)w(M_v) &= \sum_{(M_u,M_v)} w(\Phi_1(M_u,M_v))\,w(\Phi_2(M_u,M_v)) \\ &\le \mathsf{Z}(G)\,\mathsf{Z}(G-u-v),\end{aligned}$$

which is exactly (EC.12).

To define Φ, consider the symmetric difference

$$H := M_u \triangle M_v = (M_u \setminus M_v) \cup (M_v \setminus M_u).$$

Every vertex has degree at most 2 in H, so each connected component of H is either a cycle or a path. Let C_v be the (possibly empty) connected component of H containing the vertex v. Since M_v is a matching in $G - v$, the vertex v has degree at most 1 in C_v, and hence is an endpoint of a (possibly empty) path. This path, starting at v, alternates between edges from M_u and M_v. Moreover, because G is bipartite and $u \in U$, $v \in V$, this alternating path cannot contain u.

Define $\Phi(M_u, M_v) = (M, N)$ by swapping the matchings on C_v:

$$M := (M_v \setminus C_v) \cup (M_u \cap C_v),$$
$$N := (M_u \setminus C_v) \cup (M_v \cap C_v).$$

Since we swap the two matchings on the entirety of C_v, both M and N are matchings. By construction, N contains no edge incident to v. In addition, because C_v avoids u, the matching N contains no edge incident to u either. Therefore $M \in \mathcal{M}(G)$ and $N \in \mathcal{M}(G - u - v)$. We now need to show the two following properties of Φ to finish the proof:

(i) Weight preservation: On components other than C_v we do not change the matchings; on C_v we swap which matching receives each edge. Thus $w(M_u)w(M_v) = w(M)w(N)$.

(ii) Injectivity: Given an image pair (M, N), let C_v be the connected component containing v in the symmetric difference $M \triangle N$. Swapping back on C_v recovers a unique preimage (M_u, M_v), so Φ is injective. □

EC.5.5. Proof of Lemma 3 (Bipartite matching upper bound)

Proof. Fix $\varepsilon \in (0, 1/2)$ and let $n := 1/\varepsilon$ (assume $1/\varepsilon$ is an integer for simplicity). Construct the bipartite graph $G = (U \cup V, E)$ with

$$U = \{a\} \cup \{u_1, \ldots, u_n\}, \quad V = \{v_1, \ldots, v_n\},$$

and $2n$ edges $e_i := \{a, v_i\}$, and $g_i := \{u_i, v_i\}$. Define the activation vector $\boldsymbol{x}$ by

$$x_{e_i} = \varepsilon, \quad x_{g_i} = 1 - \varepsilon.$$

Then $\boldsymbol{x}$ is a fractional matching (and hence in the bipartite matching polytope $\mathcal{P}$): the degree constraint at a is tight since $\sum_i x_{e_i} = n\varepsilon = 1$, and each v_i has $x_{e_i} + x_{g_i} = 1$.

Let μ be any distribution on matchings satisfying stationary implementability, i.e., constraint (STATIONARY-IMPLEMENTABILITY), for $\boldsymbol{x}$. Moreover, suppose it attains selectability α, i.e.,

$$\mathbb{P}_{S \sim \mu}[e \in S] \geq \alpha x_e \quad \forall e \in E.$$

Let $X_i := \mathbf{1}\{e_i \in S\}$ and $A_i := 1 - \mathbf{1}\{g_i \in S\}$. By selectability, we know that $\mathbb{E}[X_i] \geq \alpha\varepsilon$, and $\mathbb{E}[A_i] \leq 1-\alpha(1-\varepsilon)$. Since at most one of the edges $\{e_i\}$ can be chosen,

$$\mathbb{P}[a \text{ is matched in } S] = \sum_{i=1}^{n} \mathbb{E}[X_i] \;\geq\; n\alpha\varepsilon \;=\; \alpha. \tag{EC.13}$$

Let $N := \sum_{i=1}^{n} A_i$. Then

$$\mathbb{E}[N] = \sum_{i=1}^{n} \mathbb{E}[A_i] \;\leq\; n\big(1-\alpha(1-\varepsilon)\big). \tag{EC.14}$$

Our strategy is to show, using the stationarity constraint, that

$$\mathbb{P}[a \text{ is matched in } S] \;\leq\; \frac{\rho\mathbb{E}[N]}{1+\rho\mathbb{E}[N]}, \text{ where } \rho := \frac{\varepsilon}{1-\varepsilon}.$$

Combined with (EC.13) and (EC.14), this immediately implies

$$\alpha \;\leq\; \frac{1-\alpha(1-\varepsilon)}{2-\varepsilon-\alpha(1-\varepsilon)},$$

and letting $\varepsilon \to 0$ yields $\alpha^2 - 3\alpha + 1 \geq 0$, hence $\alpha \leq (3-\sqrt{5})/2 = \alpha^*$.

To prove the inequality above, fix $i \in [n]$ and condition on the selection status of all edges other than e_i. The stationary constraint (STATIONARY-IMPLEMENTABILITY) gives

$$\mathbb{P}[X_i = 1 \mid S_{-e_i}] \leq \varepsilon.$$

Since $X_i = 1$ implies $A_i = 1$, we may write

$$\mathbb{P}[X_i = 1 \mid S_{-e_i}] \leq \rho A_i \mathbb{P}[a \text{ unmatched} \mid S_{-e_i}]. \tag{EC.15}$$

Here, the inequality above follows from noting that if any e_j with $j \neq i$ is selected, then both sides are 0, and if no e_j with $j \neq i$ is selected, then a unmatched is exactly the event $X_i = 0$.

Now condition only on $A = (A_1, \ldots, A_n)$ and take expectations in (EC.15) to obtain

$$\mathbb{E}[X_i \mid A] \;\leq\; \rho A_i \mathbb{P}[a \text{ unmatched} \mid A]. \tag{EC.16}$$

Summing (EC.16) over i and using disjointness of the events $\{X_i = 1\}$ gives, for $p(A) := \mathbb{P}[a \text{ matched} \mid A]$,

$$p(A) = \sum_{i=1}^{n} \mathbb{E}[X_i \mid A] \;\leq\; \rho N(1-p(A)).$$

Solving yields $p(A) \leq f(N) := \rho N/(1+\rho N)$. Since f is concave in $N \geq 0$, Jensen's inequality gives

$$\mathbb{P}[a \text{ matched}] = \mathbb{E}[p(A)] \leq \mathbb{E}[f(N)] \leq f(\mathbb{E}[N]),$$

which completes the proof. $\square$

EC.5.6. Maximum-entropy barrier on K_4

PROPOSITION EC.1 (**Maximum-entropy barrier on** K_4). *Let G be the complete graph on vertex set $\{1,2,3,4\}$. Write the edges as a 4-cycle ("outer" edges) $E_{\text{out}} := \{12,23,34,41\}$ together with the two diagonals $E_{\text{diag}} := \{13,24\}$. Fix a parameter $\varepsilon \in (0,1/2)$ and define an activation vector $\boldsymbol{x} \in [0,1]^E$ by*

$$x_e = \frac{1-\varepsilon}{2} \quad \forall e \in E_{\text{out}} \quad , \quad x_e = \varepsilon \quad \forall e \in E_{\text{diag}} \;, \tag{EC.17}$$

where $E := E_{\text{out}} \cup E_{\text{diag}}$. Then $\boldsymbol{x}$ lies in the (general) matching polytope of G. For $\alpha \in (0,1)$, let $\boldsymbol{p} := \alpha \boldsymbol{x}$ and let μ^* be the unique maximum entropy distribution (MAX-ENTROPY) *over matchings with marginals $\mathbf{p}$. If μ^* satisfies stationary implementability* (STATIONARY-IMPLEMENTABILITY) *with respect to $\boldsymbol{x}$, then*

$$\limsup_{\varepsilon \to 0} \alpha \le \frac{\sqrt{3}-1}{2} \approx 0.366.$$

Proof. Fix $\varepsilon \in (0,1/2)$ and $\alpha \in (0,1)$, set $\boldsymbol{p} = \alpha \boldsymbol{x}$, and let μ^* be the maximum-entropy distribution with marginals $\boldsymbol{p}$. Since $\mu^*(S) \propto \prod_{e \in S} w_e$ over matchings S, symmetry implies that all outer edges share a common weight $t > 0$ and both diagonals share a common weight $s > 0$:

$$w_e = t \; \forall e \in E_{\text{out}}, \quad w_e = s \; \forall e \in E_{\text{diag}}.$$

Every matching in K_4 has size 0, 1, or 2. There is one empty matching, six single-edge matchings (four outer, two diagonal), and three perfect matchings: two consisting of opposite outer edges and one consisting of both diagonals. Therefore the partition function is

$$\mathsf{Z}(t,s) = 1 + 4t + 2s + 2t^2 + s^2. \tag{EC.18}$$

For any fixed outer edge $e \in E_{\text{out}}$, the matchings containing e are $\{e\}$ and $\{e, e'\}$ where e' is the opposite outer edge, hence

$$\mathbb{P}_{S \sim \mu^*}[e \in S] = \frac{t + t^2}{\mathsf{Z}(t,s)}$$

Similarly, for any diagonal $d \in E_{\text{diag}}$,

$$\mathbb{P}_{S \sim \mu^*}[d \in S] = \frac{s + s^2}{\mathsf{Z}(t,s)}$$

By construction of μ^*,

$$\mathbb{P}_{S \sim \mu^*}[e \in S] = p_e = \alpha(1-\varepsilon)/2 \quad \forall e \in E_{\text{out}}, \quad \mathbb{P}_{S \sim \mu^*}[d \in S] = p_d = \alpha \varepsilon \quad \forall d \in E_{\text{diag}}.$$

Simple algebra therefore gives the identities:

$$\frac{s(1+s)}{t(1+t)} = \frac{\alpha \varepsilon}{\alpha(1-\varepsilon)/2} = \frac{2\varepsilon}{1-\varepsilon}, \qquad \alpha = \frac{2}{1-\varepsilon} \cdot \frac{t(1+t)}{\mathsf{Z}(t,s)}. \tag{EC.19}$$

Fix a diagonal $d \in E_{\text{diag}}$. By Lemma 1 and the stationary implementability constraint in the LP, i.e., (STATIONARY-IMPLEMENTABILITY), whenever d is addable to S_{-d},

$$\mathbb{P}_{S \sim \mu^*}\left[d \in S \mid S_{-d} = T\right] = \rho_d = \frac{s}{1+s} \le \varepsilon.$$

Equivalently, $s \le \varepsilon/(1-\varepsilon)$, and hence

$$s(1+s) \le \frac{\varepsilon}{1-\varepsilon} \cdot \frac{1}{1-\varepsilon} = \frac{\varepsilon}{(1-\varepsilon)^2}.$$

Combining this bound with (EC.19) gives

$$t(1+t) \le \frac{1}{2(1-\varepsilon)}. \tag{EC.20}$$

Letting $\varepsilon \to 0$ in (EC.20) yields $t(1+t) \le 1/2$. Moreover, when $\varepsilon \to 0$ we have $s \to 0$, so $\mathsf{Z}(t,s) \to 1+4t+2t^2$ and also

$$\alpha \to \frac{2t(1+t)}{1+4t+2t^2} := f(t).$$

A direct calculation shows that $f(t)$ is increasing for $t \ge 0$, so the largest possible α occurs at the largest feasible t. Let t_0 be the positive root of $t^2+t = 1/2$, namely $t_0 = (\sqrt{3}-1)/2$. Therefore, $\alpha \le f(t_0)$. Evaluating f at t_0 gives:

$$f(t_0) = \frac{2t_0(1+t_0)}{1+4t_0+2t_0^2} = \frac{1}{1+\sqrt{3}} = \frac{\sqrt{3}-1}{2}.$$

Therefore $\limsup_{\varepsilon\to 0} \alpha \le (\sqrt{3}-1)/2$, as claimed. □

EC.5.7. Proof of Lemma 4 (Poisson comparison)

Proof. Write $a_t := \mathbb{P}[T=t]$ and $b_t := \mathbb{P}[Q=t]$, and define the corresponding truncated distributions on $\{0,1,\dots,k\}$:

$$\widehat{a}_t \;:=\; \mathbb{P}[T=t \mid T \le k] \;=\; \frac{a_t}{\mathbb{P}[T\le k]}, \qquad \widehat{b}_t \;:=\; \mathbb{P}[Q=t \mid Q \le k] \;=\; \frac{b_t}{\mathbb{P}[Q\le k]}.$$

It is well known in the combinatorics and probability literature that the coefficient sequence (a_t) is *ultra log-concave* (Liggett, 1997; Pemantle, 2000), meaning that for all $t \ge 1$,

$$a_t^2 \;\ge\; a_{t-1}a_{t+1} \cdot \frac{t+1}{t}. \tag{EC.21}$$

Clearly, the same inequality holds with a_t replaced by $\widehat{a}_t$ for $t = 1,\dots,k-1$. For a Poisson distribution, a direct calculation gives equality:

$$\widehat{b}_t^2 \;=\; \widehat{b}_{t-1}\widehat{b}_{t+1} \cdot \frac{t+1}{t} \qquad \text{for all } t = 1,\dots,k-1. \tag{EC.22}$$

Dividing (EC.21) (for $\widehat{a}_t$) by (EC.22) yields that the ratio

$$r_t \;:=\; \frac{\widehat{a}_t}{\widehat{b}_t} \qquad (t = 0,1,\dots,k)$$

is log-concave: $r_t^2 \ge r_{t-1}r_{t+1}$ for $t = 1,\dots,k-1$. In particular, $(r_t)_{t=0}^k$ is unimodal. Assume now for the sake of contradiction that $\widehat{a}_k > \widehat{b}_k$. Then $r_k > 1$. Since (r_t) is unimodal, the set $\{t : r_t > 1\}$ is an interval containing

k, so there exists an index $r \in \{0,1,\dots,k\}$ such that $r_t \le 1$ for $t < r$ and $r_t > 1$ for $t \ge r$. Define $\delta_t := \widehat{a}_t - \widehat{b}_t$. Then $\delta_t \le 0$ for $t < r$, $\delta_t > 0$ for $t \ge r$, and $\sum_{t=0}^{k} \delta_t = 0$. Therefore,

$$\sum_{t=0}^{k} t\,\delta_t \;\ge\; r\sum_{t=r}^{k}\delta_t + (r-1)\sum_{t=0}^{r-1}\delta_t \;=\; \sum_{t=r}^{k}\delta_t \;>\; 0,$$

This shows

$$\mathbb{E}[T \mid T \le k] \;=\; \sum_{t=0}^{k} t\,\widehat{a}_t \;>\; \sum_{t=0}^{k} t\,\widehat{b}_t \;=\; \mathbb{E}[Q \mid Q \le k],$$

contradicting the assumption in the statement of Lemma 4. Thus we must have $\widehat{a}_k \le \widehat{b}_k$. □

EC.5.8. Proof of Theorem 4 (S-OCRS for rank-L hypergraph matchings with capacity k)

Proof. Recall that we fix $\varepsilon \in [0,1]$, and let:

$$Q \sim \text{Poisson}((1-\varepsilon)k), \quad \beta := \mathbb{P}[Q = k \mid Q \le k], \quad \gamma := (1-\varepsilon)(1-\beta).$$

Let μ be the Gibbs distribution on $\mathcal{F}$ with parameters $\rho_e = \gamma x_e$. Note that $\mathbb{P}_{S\sim\mu}[e \in S \mid S_{-e} = T]$ is equal to either ρ_e or 0. Since we have explicitly chosen $\rho_e = \gamma x_e$ with $\gamma = (1-\varepsilon)(1-\beta) \le 1$, STATIONARY-IMPLEMENTABILITY is satisfied automatically, so we concentrate on checking SELECTABILITY with $\alpha = \gamma(1-L\beta)$, i.e.,

$$\mathbb{P}_{S\sim\mu}[e \in S] \ge \gamma(1-L\beta)x_e.$$

Now, we know that

$$\mathbb{P}[e \in S] = \mathbb{P}[\mathsf{Add}(e)] \cdot \rho_e = \mathbb{P}[\mathsf{Add}(e)] \cdot \gamma x_e.$$

Following the notation of Theorem 3, let $(Y_e)_{e\in E}$ be independent Bernoulli random variables with $\mathbb{P}[Y_e = 1] = \rho_e = \gamma x_e$. Letting $T_v := \sum_{e\in\delta(v)} Y_e$, we know e is addable if none of the vertices that are a part of e are at capacity. So by the union bound it follows that:

$$\mathbb{P}[\mathsf{Add}(e)] \ge 1 - \sum_{v\in e} \mathbb{P}[T_v = k | (Y_e)_{e\in E} \in \mathcal{F}].$$

So it suffices to show that $\mathbb{P}[T_v = k | (Y_e)_{e\in E} \in \mathcal{F}] \le \beta$, since then

$$\mathbb{P}[\mathsf{Add}(e)] \ge 1 - |e|\beta \ge 1 - L\beta,$$

and the scheme is $\gamma(1-L\beta)$-selectable, as desired. We will show in Proposition EC.2 that: $\mathbb{P}[T_v = k | (Y_e)_{e\in E} \in \mathcal{F}] \le \mathbb{P}[T_v = k | T_v \le k]$. Assuming this result, note that T_v is Poisson–binomial, and

$$\mathbb{E}[T_v \mid T_v \le k] \le \mathbb{E}[T_v] = \sum_{f\in\delta(v)} \rho_f = \gamma \sum_{f\in\delta(v)} x_f \le \gamma k = (1-\varepsilon)(1-\beta)k.$$

On the other hand, for $Q \sim \text{Poisson}((1-\varepsilon)k)$ we have the identity

$$\mathbb{E}[Q \mid Q \le k] = (1-\varepsilon)k\mathbb{P}[Q < k \mid Q \le k] = (1-\varepsilon)(1-\beta)k.$$

Applying Lemma 4 therefore gives

$$\mathbb{P}[T_v = k \mid T_v \le k] \le \mathbb{P}[Q = k \mid Q \le k] = \beta,$$

as needed. For the concrete bound, note that standard Chernoff and Stirling bounds give

$$\mathbb{P}[Q \ge k] \le e^{-\frac{k\varepsilon^2}{2}}, \quad \mathbb{P}[Q = k] = e^{-(1-\varepsilon)k}\frac{((1-\varepsilon)k)^k}{k!} \le \frac{e^{k(\varepsilon+\log(1-\varepsilon))}}{\sqrt{2\pi k}} \le \frac{e^{-\frac{k\varepsilon^2}{2}}}{\sqrt{2\pi k}}.$$

Hence for $L \ge 2$ and $k > 2\log L$, taking $\varepsilon = \sqrt{2\log L/k}$, for which $e^{-k\varepsilon^2/2} = 1/L$ therefore gives:

$$\beta \le \frac{1}{\sqrt{2\pi k}} \cdot \frac{1/L}{1-1/L} = \frac{1}{\sqrt{2\pi k}} \cdot \frac{1}{L-1}, \qquad L\beta \le \frac{L}{L-1} \cdot \frac{1}{\sqrt{2\pi k}} \le \sqrt{\frac{2}{\pi k}}.$$

With this choice, the selectability lower bound becomes

$$\gamma(1-L\beta) = (1-\varepsilon)(1-\beta)(1-L\beta) \ge 1 - \sqrt{\frac{2\log L}{k}} - O\left(\sqrt{\frac{1}{k}}\right),$$

as claimed. (The result also works for $L = 1$, since with $\varepsilon = 0$, we have $\beta = \sqrt{\frac{2}{\pi k}} + O\left(\frac{1}{k}\right)$.) □

It remains to show that $\mathbb{P}[T_v = k|(Y_e)_{e\in E} \in \mathcal{F}] \le \mathbb{P}[T_v = k|T_v \le k]$. Intuitively speaking, this ought to be true because conditioning on the stricter condition $(Y_e)_{e\in E} \in \mathcal{F}$ ought to decrease the values of the Y variables, since vertices v' other than v will also need $T_{v'} \le k$. This in turn should reduce the chance that $T_v = k$. Formally, we establish the result in the following proposition using Efron monotonicity (Efron, 1965):

PROPOSITION EC.2. *Let $(Y_e)_{e\in E}$ be independent Bernoulli random variables with $\mathbb{P}[Y_e = 1] = \rho_e = \gamma x_e$. Fix a vertex $v \in V$, and let $T_v = \sum_{e\in\delta(v)} Y_e$. Then,*

$$\mathbb{P}[T_v = k|(Y_e)_{e\in E} \in \mathcal{F}] \le \mathbb{P}[T_v = k|T_v \le k].$$

Proof. Consider for the vertex v the edges which are neighbors, and the rest:

$$N := \{f \in \delta(v) : Y_f = 1\} \quad \text{and} \quad R := \{f \in E \setminus \delta(v) : Y_f = 1\},$$

so that $|N| = T_v$. Condition on a realization $R = r$. Then

$$\begin{aligned}\mathbb{P}\left[T_v = k \mid (Y_f)_{f\in E} \in \mathcal{F},\ R = r\right] &= \frac{\mathbb{P}[T_v = k,\ N \cup r \in \mathcal{F} \mid R = r]}{\mathbb{P}[N \cup r \in \mathcal{F} \mid R = r]} \\ &= \frac{\mathbb{P}[N \cup r \in \mathcal{F} \mid T_v = k, R = r]\mathbb{P}[T_v = k]}{\sum_{t=0}^{k} \mathbb{P}[N \cup r \in \mathcal{F} \mid T_v = t, R = r]\mathbb{P}[T_v = t]}.\end{aligned}$$

Here we have used the independence of R and N. The discrete version of Efron's monotonicity theorem states that, for independent Bernoulli random variables, the conditional expectation of any coordinatewise increasing function, conditional on their sum being t, is increasing in t (Goldschmidt et al., 2008, Proposition 2.3); see also Efron (1965) for the original theorem. Given $R = r$, the indicator of the event $N \cup r \in \mathcal{F}$ is coordinatewise decreasing in $(Y_f)_{f\in\delta(v)}$, since $\mathcal{F}$ is downward-closed. Applying the theorem to the complementary indicator therefore gives, for every $t \le k$,

$$\mathbb{P}[N\cup r\in\mathcal{F}\mid T_v=t,R=r]\ge\mathbb{P}[N\cup r\in\mathcal{F}\mid T_v=k,R=r].$$

Substituting this inequality into the denominator above gives

$$\begin{aligned}\mathbb{P}\left[T_v=k\mid (Y_f)_{f\in E}\in\mathcal{F},\ R=r\right] &\le \frac{\mathbb{P}[N\cup r\in\mathcal{F}\mid T_v=k,R=r]\mathbb{P}[T_v=k]}{\displaystyle\sum_{t=0}^{k}\mathbb{P}[N\cup r\in\mathcal{F}\mid T_v=k,R=r]\mathbb{P}[T_v=t]}\\ &=\frac{\mathbb{P}[T_v=k]}{\sum_{t=0}^{k}\mathbb{P}[T_v=t]}=\mathbb{P}[T_v=k\mid T_v\le k].\end{aligned}$$

Averaging over r under the conditional distribution of R given $(Y_f)_{f\in E}\in\mathcal{F}$ gives the desired result. □

EC.5.9. Proof of Lemma 5 (Uniform matroid upper bound)

Proof. Let $\mu\in\Delta(\mathcal{F})$ be any feasible distribution for the symmetric instance satisfying stationary implementability (STATIONARY-IMPLEMENTABILITY), and suppose it attains some selectability α. Averaging μ over all permutations of E preserves these properties. Thus, we may assume μ is exchangeable: $\mu(S)$ depends only on $|S|$. Define

$$\pi_\ell \;:=\; \mathbb{P}_{S\sim\mu}[\,|S|=\ell\,] \qquad (\ell=0,1,\dots,k).$$

Then every subset of size ℓ has probability $\pi_\ell/\binom{n}{\ell}$.

Fix $e\in E$ and fix any $T\subseteq E\setminus\{e\}$ with $|T|=\ell\le k-1$. Conditioning on $S_{-e}=T$ leaves only the possibilities $S=T$ and $S=T\cup\{e\}$, so

$$\mathbb{P}[e\in S\mid S_{-e}=T]=\frac{\pi_{\ell+1}/\binom{n}{\ell+1}}{\pi_\ell/\binom{n}{\ell}+\pi_{\ell+1}/\binom{n}{\ell+1}}=\frac{(\ell+1)\pi_{\ell+1}}{(n-\ell)\pi_\ell+(\ell+1)\pi_{\ell+1}}.$$

The constraint from the stationary OCRS LP—that is, (STATIONARY-IMPLEMENTABILITY) (with $x_e=q$)—therefore implies, for all $\ell=0,\dots,k-1$,

$$(1-q)(\ell+1)\pi_{\ell+1}\;\le\;q(n-\ell)\pi_\ell. \tag{EC.23}$$

Since the instance is symmetric, $\mathbb{P}[e\in S]=\mathbb{E}[|S|]/n$ for every e and hence the selectability α satisfies

$$\alpha\;\le\;\frac{\mathbb{P}[e\in S]}{q}\;=\;\frac{\mathbb{E}[|S|]}{nq}\;=\;\frac{\mathbb{E}[|S|]}{k}.$$

Thus, upper-bounding α reduces to upper-bounding $\mathbb{E}[|S|]=\sum_{\ell=0}^{k}\ell\pi_\ell$ subject to (EC.23) and $\sum_{\ell=0}^{k}\pi_\ell=1$. At an optimum, all constraints (EC.23) must be tight: if for some ℓ the inequality is strict, then shifting an infinitesimal mass from π_ℓ to $\pi_{\ell+1}$ increases $\mathbb{E}[|S|]$ and preserves feasibility for sufficiently small shifts. Hence, for $\ell=0,\dots,k-1$,

$$\frac{\pi_{\ell+1}}{\pi_\ell}\;=\;\frac{q}{1-q}\cdot\frac{n-\ell}{\ell+1}. \tag{EC.24}$$

The right-hand side equals the ratio of successive binomial probabilities for $\mathrm{Bin}(n,q)$, so $|S|$ has the distribution of $B\sim\mathrm{Bin}(n,q)$ conditioned on $\{B\le k\}$. Therefore

$$\mathbb{E}[|S|] \;=\; \mathbb{E}[B \mid B\le k] \;=\; \frac{\mathbb{E}[B\mathbf{1}\{B\le k\}]}{\mathbb{P}[B\le k]}.$$

Using the standard binomial identity $\mathbb{E}[B\mathbf{1}\{B\le k\}] = nq\cdot\mathbb{P}[\mathrm{Bin}(n-1,q)<k]$ and $nq=k$, we obtain

$$\alpha \;\le\; \frac{\mathbb{E}[|S|]}{k} \;=\; \frac{\mathbb{P}[\mathrm{Bin}(n-1,q)\le k-1]}{\mathbb{P}[\mathrm{Bin}(n,q)\le k]},$$

which is (20). Finally, as $n\to\infty$ with $q=k/n$, the Poisson limit theorem gives $\mathrm{Bin}(n,q)\Rightarrow\mathrm{Poisson}(k)$ and $\mathrm{Bin}(n-1,q)\Rightarrow\mathrm{Poisson}(k)$, so the ratio (20) converges to

$$\frac{\mathbb{P}[\mathrm{Poisson}(k)<k]}{\mathbb{P}[\mathrm{Poisson}(k)\le k]} = \alpha_k,$$

as desired. $\square$

EC.5.10. Proof of Remark EC.1

Proof. The implication (EC.2) $\Rightarrow$ (pairwise form) is immediate by taking $T=\{f\}$. For the converse, assume the pairwise inequality holds for every tilt: for all $e,f\in E$ and all $\boldsymbol{w}$,

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e\in B]\cdot\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[f\in B] \;\ge\; \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[\{e,f\}\subseteq B].$$

Fix $\boldsymbol{w}$, $e\in E$, and a set $T=\{f_1,\dots,f_m\}\subseteq E\setminus\{e\}$. If $\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T\subseteq B]=0$, then $\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T\cup\{e\}\subseteq B]=0$ as well and (EC.2) holds trivially. Hence assume $\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T\subseteq B]>0$.

For any $A\subseteq E$, let $\mu_{\boldsymbol{w}}\mid A$ denote the conditional law of $B\sim\mu_{\boldsymbol{w}}$ given $A\subseteq B$. We first note that $\mu_{\boldsymbol{w}}\mid A$ is a limit of tilts of $\mu_{\boldsymbol{w}}$. Indeed, for $\lambda>0$ define $\boldsymbol{w}^{(\lambda)}$ by

$$w_g^{(\lambda)} \;:=\; \begin{cases}\lambda\, w_g, & g\in A,\\ w_g, & g\notin A.\end{cases}$$

Since $\mu_{\boldsymbol{w}^{(\lambda)}}(B)\propto\mu_{\boldsymbol{w}}(B)\cdot\lambda^{|A\cap B|}$, as $\lambda\to\infty$ the mass concentrates on sets with $A\subseteq B$, and hence $\mu_{\boldsymbol{w}^{(\lambda)}}$ converges to $\mu_{\boldsymbol{w}}\mid A$. Because the state space is finite, this convergence implies convergence of all event probabilities.

Now fix $j\in\{1,\dots,m\}$ and write $A_j:=\{f_1,\dots,f_j\}$ for $j=1,\dots,m$. Applying the assumed pairwise inequality to the tilt $\mu_{\boldsymbol{w}^{(\lambda)}}$ and then taking $\lambda\to\infty$ yields that the conditional measure $\mu_{\boldsymbol{w}}\mid A_{j-1}$ also satisfies pairwise negative correlation, in particular

$$\mathbb{P}[e\in B\mid A_{j-1}\subseteq B]\cdot\mathbb{P}[f_j\in B\mid A_{j-1}\subseteq B] \;\ge\; \mathbb{P}[\{e,f_j\}\subseteq B\mid A_{j-1}\subseteq B].$$

Whenever $\mathbb{P}[f_j\in B\mid A_{j-1}\subseteq B]>0$, dividing both sides by this quantity gives

$$\mathbb{P}[e\in B\mid A_j\subseteq B] \;\le\; \mathbb{P}[e\in B\mid A_{j-1}\subseteq B].$$

(If $\mathbb{P}[f_j \in B \mid A_{j-1} \subseteq B] = 0$, then $\mathbb{P}[A_j \subseteq B] = 0$ and the above inequality is again trivial.)

Iterating over $j = 1, \dots, m$ yields

$$\mathbb{P}[e \in B \mid T \subseteq B] \;\le\; \mathbb{P}[e \in B].$$

Multiplying by $\mathbb{P}[T \subseteq B]$ and rewriting gives exactly (EC.2):

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T \subseteq B] \cdot \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e \in B] \;\ge\; \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[T \cup \{e\} \subseteq B].$$

□

EC.5.11. A counterexample to the standard maximum-entropy method on matroids

To see a counterexample to applying the maximum-entropy approach in the graphic matroid case, let G_n be the graph with terminals u, v and intermediate vertices $m_1, \dots, m_n$, with edges

$$E(G_n) = \{um_i, m_i v : i = 1, \dots, n\}.$$

Let $\mathcal{I}$ be the family of forests of G_n (i.e., the independent sets of the graphic matroid of G_n), and let $x_e = 1/2$ for every edge e. Consider the *unweighted* maximum-entropy distribution on $\mathcal{I}$, namely the uniform distribution over forests. A forest connects u to v iff it contains *both* edges of at least one length-2 path $u - m_i - v$. We can count the forests as follows:

- If u and v are disconnected, then for each i we may choose $\emptyset$, $\{um_i\}$, or $\{m_i v\}$, yielding 3^n forests.
- If u and v are connected, then there is a unique index i whose full path is present (choose i in n ways), and for each $j \ne i$ we again choose among $\emptyset, \{um_j\}, \{m_j v\}$, yielding $n \cdot 3^{n-1}$ forests.

Hence, under the uniform (maximum-entropy) forest,

$$\mathbb{P}[u \text{ and } v \text{ are disconnected}] = \frac{3^n}{3^n + n \cdot 3^{n-1}} = \frac{3}{n+3}$$

If we add a further element $e^* = (u, v)$ with x_{e^*} tiny, then e^* is addable precisely when u and v are disconnected, so $\mathbb{P}[\mathsf{Add}(e^*)] = o(1)$ under this max-entropy distribution. This illustrates the obstruction: the high-entropy forest distribution overwhelmingly prefers configurations that realize *some* u–v connection (there are n choices), and this drives the addability probability of e^* to 0.

The same phenomenon persists for weighted Gibbs measures with any fixed edge-weight $w > 0$ on every edge: the weight of forests with no u–v path is $(1 + 2w)^n$, while forests with a (unique) u–v path have total weight $nw^2(1 + 2w)^{n-1}$. If we fix some constant selectability that we aim for, say α, and take $n \to \infty$, it's straightforward to show that for the maximum entropy distribution, $w \to \frac{\alpha}{2(1-\alpha)}$. Thus, the maximum entropy approach is incompatible with proving (or implementing) any constant-factor stationary OCRS.

EC.5.12. Not all matroids are weakly Rayleigh

This follows from a simple variation on the results in Brändén and D'León (2010). We start by noting that Theorem 2.1 in Brändén and D'León (2010) also holds when the polynomial in question is Rayleigh rather than stable:

LEMMA EC.4. *Let $\mathcal{M} = (E, \mathcal{I})$ be a weakly Rayleigh matroid with set of bases $\mathcal{B}$ and base measure μ. Let*

$$(B_1, B_2, B_3, B_4) = \big(A \cup \{i,k\}, A \cup \{i,\ell\}, A \cup \{j,\ell\}, A \cup \{j,k\}\big)$$

be a degenerate quadrangle of bases of $\mathcal{M}$, i.e., $i,j,k,l \notin A$, and at most one of $A \cup \{i,j\}$ and $A \cup \{k,\ell\}$ is a basis. Then

$$\mu(B_1)\mu(B_3) = \mu(B_2)\mu(B_4).$$

Proof. Since μ is Rayleigh, we have that for every distinct $e, f \in E$ and every $\boldsymbol{w} \in \mathbb{R}^E_{>0}$,

$$\mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[e \in B] \cdot \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[f \in B] \;\geq\; \mathbb{P}_{B\sim\mu_{\boldsymbol{w}}}[\{e,f\} \subseteq B].$$

Equivalently, writing

$$g_\mu(\boldsymbol{w}) = \sum_{B\in\mathcal{B}} \mu(B) \prod_{e\in B} w_e,$$

we have for every distinct $e, f \in E$ and every $\boldsymbol{w} \in \mathbb{R}^E_{>0}$,

$$\Delta_{ef} g_\mu(\boldsymbol{w}) := \frac{\partial g_\mu}{\partial w_e}(\boldsymbol{w}) \, \frac{\partial g_\mu}{\partial w_f}(\boldsymbol{w}) - g_\mu(\boldsymbol{w}) \, \frac{\partial^2 g_\mu}{\partial w_e \partial w_f}(\boldsymbol{w}) \;\geq\; 0,$$

in which case we say the polynomial g_μ is Rayleigh. By symmetry, we may assume that $A \cup \{i,j\} \notin \mathcal{B}$. Fix $\lambda > 0$, $\varepsilon > 0$, and define a weight vector $\boldsymbol{w}^{(\lambda,\varepsilon)} \in \mathbb{R}^E_{>0}$ by

$$w_e^{(\lambda,\varepsilon)} := \begin{cases} \lambda, & e \in A, \\ w_e, & e = i,j,k,\ell \\ \varepsilon, & e \in E \setminus (A \cup \{i,j,k,\ell\}). \end{cases}$$

Consider

$$h_{\lambda,\varepsilon}(w_i, w_j, w_k, w_\ell) := \lambda^{-|A|} \, g_\mu\big(\boldsymbol{w}^{(\lambda,\varepsilon)}\big)\,.$$

Since $h_{\lambda,\varepsilon}$ is obtained from g_μ by rescaling and setting values for some variables, it is Rayleigh.

Now let $\lambda \to \infty$ and $\varepsilon \to 0$. The only surviving terms are those corresponding to bases containing A and no element of $E \setminus (A \cup \{i,j,k,\ell\})$. Because $A \cup \{i,j\} \notin \mathcal{B}$, the limit polynomial is

$$h(w_i, w_j, w_k, w_\ell) = \mu(B_1)\, w_i w_k + \mu(B_2)\, w_i w_\ell + \mu(B_3)\, w_j w_\ell + \mu(B_4)\, w_j w_k + \eta\, w_k w_\ell,$$

where

$$\eta := \begin{cases} \mu(A \cup \{k,\ell\}), & \text{if } A \cup \{k,\ell\} \in \mathcal{B}, \\ 0, & \text{otherwise.} \end{cases}$$

Since each $h_{\lambda,\varepsilon}$ is Rayleigh, and Δ_{ef} depends continuously on the coefficients, the limit polynomial h is also Rayleigh.

A direct calculation gives

$$\Delta_{ik}h = w_\ell\Big(\left(\mu(B_2)\mu(B_4) - \mu(B_1)\mu(B_3)\right)w_j + \mu(B_2)\eta\, w_\ell\Big),$$

and

$$\Delta_{i\ell}h = w_k\Big(\left(\mu(B_1)\mu(B_3) - \mu(B_2)\mu(B_4)\right)w_j + \mu(B_1)\eta\, w_k\Big).$$

Since h is Rayleigh, both expressions are nonnegative for all $w_i, w_j, w_k, w_\ell > 0$. Fix $w_j = 1$. From $\Delta_{ik}h \geq 0$, divide by $w_\ell > 0$ and let $w_\ell \to 0$ to obtain

$$\mu(B_2)\mu(B_4) - \mu(B_1)\mu(B_3) \geq 0.$$

Similarly, from $\Delta_{i\ell}h \geq 0$, divide by $w_k > 0$ and let $w_k \to 0$ to obtain

$$\mu(B_1)\mu(B_3) - \mu(B_2)\mu(B_4) \geq 0.$$

Hence

$$\mu(B_1)\mu(B_3) = \mu(B_2)\mu(B_4),$$

as desired. □

The spaces $V_{\mathcal{M}}$ and $W_{\mathcal{M}}$. Following Brändén and D'León (2010), let $V_{\mathcal{M}} \subseteq \mathbb{R}^{\mathcal{B}}$ be the linear subspace of all functions $\nu : \mathcal{B} \to \mathbb{R}$ satisfying

$$\nu(B_1) + \nu(B_3) - \nu(B_2) - \nu(B_4) = 0$$

for every degenerate quadrangle (B_1, B_2, B_3, B_4) of bases of $\mathcal{M}$. Let $W_{\mathcal{M}} \subseteq V_{\mathcal{M}}$ be the subspace of functions of the form

$$\nu(B) = \sum_{e \in B} v_e.$$

for some vector $(v_e)_{e \in E} \in \mathbb{R}^E$. We have the following analog of Theorem 2.3 in Brändén and D'León (2010). The proof is exactly the same as Brändén and D'León (2010).

THEOREM EC.5. *Let $\mathcal{M} = (E, \mathcal{I})$ be a matroid. Assume that*

$$\dim(W_{\mathcal{M}}) = \dim(V_{\mathcal{M}}).$$

Then $\mathcal{M}$ is weakly Rayleigh if and only if the uniform measure on $\mathcal{B}$ is Rayleigh.

To finish, note that by Corollary 2.7 and Table 1 in Brändén and D'León (2010),

$$\dim(W_{S_8}) = \dim(V_{S_8}) = 8.$$

Moreover, S_8 is not Rayleigh; see Choe and Wagner (2006). Therefore, S_8 is not weakly Rayleigh.

EC.5.13. Proof of Lemma EC.1 (Dominating base point)

Proof. Consider the nonempty polyhedron

$$Q := \Big\{ \boldsymbol{y} \in [0,1]^E : \sum_{e\in T} y_e \le \mathrm{rank}_{\mathcal{M}}(T) \forall T \subseteq E, \boldsymbol{y} \ge \boldsymbol{x} \Big\}.$$

Let $\boldsymbol{q}$ maximize $\sum_{e\in E} y_e$ over Q. If $\sum_e q_e = r$, then $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ and $\boldsymbol{q} \ge \boldsymbol{x}$, as desired.

Suppose by way of contradiction that $\sum_e q_e < r = \mathrm{rank}_{\mathcal{M}}(E)$. Let $\mathcal{T}$ be the family of tight sets at $\boldsymbol{q}$: $\mathcal{T} := \big\{ T \subseteq E : \sum_{e\in T} q_e = \mathrm{rank}_{\mathcal{M}}(T) \big\}$. By submodularity of $\mathrm{rank}_{\mathcal{M}}(\cdot)$, $\mathcal{T}$ is closed under unions, so $T^* := \bigcup_{T\in\mathcal{T}} T$ is also tight. Since $\sum_e q_e < \mathrm{rank}_{\mathcal{M}}(E)$, we must have $T^* \ne E$; pick $e \in E \setminus T^*$.

For sufficiently small $\varepsilon > 0$, increasing q_e to $q_e + \varepsilon$ keeps all inequalities feasible: any violated rank inequality would have to correspond to a set T with $e \in T$ that is tight at $\boldsymbol{q}$, contradicting $e \notin T^*$. This produces a feasible point in Q with a larger objective value, contradicting the maximality of $\boldsymbol{q}$. Therefore, $\sum_e q_e = r$ and $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$. □

EC.5.14. A counterexample to the standard maximum-entropy method on knapsack constraints

To see a counterexample to applying the maximum-entropy approach for knapsack constraints, consider a unit-capacity knapsack with ground set $E = [n]$, and assume the sizes are all equal to $1/n$. Let $x_i = 1$ for every item. The feasible sets are all subsets of these items.

Therefore, under any weighted Gibbs measures with fixed item-weight $w > 0$, the probability we pick none of these items is $\frac{1}{(1+w)^n}$. If we fix some constant selectability that we aim for, say α, and take $n \to \infty$, it's straightforward to show that for the maximum entropy distribution, $w \to \frac{\alpha}{1-\alpha}$.

Now, if we add a further item b with size 1 and x_b tiny, then b is addable precisely when none of the original items are selected, so $\mathbb{P}[\mathsf{Add}(b)] = o(1)$ under this max-entropy distribution. This illustrates the obstruction: the high-entropy distribution overwhelmingly prefers configurations that pick *some* light item (there are n choices), and this drives the addability probability of the heavy item to 0.

EC.5.15. Proof of Lemma EC.2 (Knapsack upper bound)

Proof. Fix $\varepsilon \in (0, 1/4)$, let $d = 1/2 + \varepsilon$, and let m be large. Consider an instance with an always-active small item s with $(a_s, x_s) = (\varepsilon, 1)$, a rarely active large item b with $(a_b, x_b) = (1, \varepsilon)$, and a family $H = \{h_1, \dots, h_m\}$ of m items with size d and common activation probability

$$p := \frac{1-2\varepsilon}{md}.$$

Note that $a_s x_s + a_b x_b + \sum_i a_{h_i} x_{h_i} = \varepsilon + \varepsilon + (1-2\varepsilon) = 1$, so $x \in \widetilde{\mathcal{P}}$. The only feasible sets are $\emptyset$, $\{s\}$, $\{b\}$, $\{h_i\}$, and $\{s, h_i\}$. Let μ be an α-selectable stationary witness. For b, selectability and stationary implementability give $\mu(\{b\}) \ge \alpha\varepsilon$ and $(1-\varepsilon)\mu(\{b\}) \le \varepsilon\mu(\emptyset)$, so:

$$\mu(\emptyset) + \mu(\{b\}) \ge \alpha.$$

Next, summing the selectability constraint over the items in H gives

$$\sum_i (\mu(\{h_i\}) + \mu(\{s,h_i\})) \geq \alpha mp.$$

Finally, stationary implementability between $\{s\}$ and $\{s,h_i\}$ gives $(1-p)\mu(\{s,h_i\}) \leq p\mu(\{s\})$. Summing these inequalities and using selectability for the always-active item s gives

$$\alpha \leq \mu(\{s\}) + \sum_i \mu(\{s,h_i\}) \leq \left(1 + \frac{mp}{1-p}\right)\mu(\{s\}),$$

and hence:

$$\mu(\{s\}) \geq \alpha/(1 + mp/(1-p)).$$

We conclude that

$$1 = \mu(\emptyset) + \mu(\{b\}) + \mu(\{s\}) + \sum_i (\mu(\{h_i\}) + \mu(\{s,h_i\})) \geq \alpha\left(1 + mp + \frac{1}{1 + \frac{mp}{1-p}}\right).$$

Letting $m \to \infty$ and then $\varepsilon \to 0$, we have $p \to 0$ and $mp = (1-2\varepsilon)/(1/2+\varepsilon) \to 2$. Thus the upper bound on α converges to $(1+2+1/3)^{-1} = 3/10$. □

EC.5.16. Proof of Lemma EC.3 (Downward closed constraints upper-bound)

Proof. Partition the ground set into m blocks

$$E = E_1 \sqcup E_2 \sqcup \cdots \sqcup E_m, \quad |E_j| = m.$$

Let

$$\mathcal{F} := \{S \subseteq E : S \subseteq E_j \text{ for some } 1 \leq j \leq m\}.$$

Thus a feasible set may contain any subset of one block, but it may not mix elements from different blocks. This family is downward closed. Set $x_e = \frac{1}{m}\ \forall e \in E$. Choosing a block uniformly and taking it whole has marginals x, so $x \in \mathcal{P}$. Now fix an arbitrary OCRS and consider revealing the elements of E in any fixed order. Let A denote the final selected set of the OCRS under this order. For each block j, define

$$Y_j := |A \cap E_j|, \quad D_j := \{Y_j > 0\}.$$

By α-selectability under this particular order,

$$\mathbb{E}[Y_j] = \sum_{e \in E_j} \mathbb{P}[e \in A] \geq \sum_{e \in E_j} \alpha x_e = \alpha. \tag{EC.25}$$

We next upper-bound the expected number of selected elements from a block, conditional on selecting at least one element from that block. On the event D_j, let $e_j^\star$ be the first selected element of E_j in the within-block arrival order. After the arrival of $e_j^\star$, any additional selected element from E_j must be an active element among the remaining at most $m-1$ elements of E_j. Conditional on the history up to and including $e_j^\star$, these future activations are still independent Bernoulli random variables with success probability $1/m$. Hence the expected number of additional selected elements from E_j is at most 1. Therefore

$$\mathbb{E}[Y_j \mid D_j] \leq 2. \tag{EC.26}$$

Combining (EC.25) and (EC.26), we get that $\mathbb{P}[D_j] \geq \frac{\mathbb{E}[Y_j]}{2} \geq \frac{\alpha}{2}$. Finally, the events $D_1, \ldots, D_m$ are disjoint, because a feasible selected set cannot contain elements from two different blocks. Thus

$$1 \geq \sum_{j=1}^{m} \mathbb{P}[D_j] \geq m\alpha/2,$$

so $\alpha \leq 2/m = 2/\sqrt{N}$, as needed. □

EC.5.17. Proof of convergence of set of admitted proposals (S-OCRS to admission control)

In this section, we prove the convergence claim used in Lemma 6. Recall that we fix a distribution over clones μ_n, and type-i proposals assigned to a clone e^j arrive at rate y_{ie}/n. For every time t, we let $S(t) \in \mathcal{F}_n$ be the set of currently admitted proposals just before time t with $S(0) = \emptyset$. If $e_j \in S(t)$ or $S(t) \cup \{e^j\} \notin \mathcal{F}_n$, the proposal is rejected. Otherwise, the proposal is admitted with probability

$$f(S(t), e^j) = \frac{n\mu_n(S(t) \cup \{e^j\})}{x_e \mu_n(S(t))} \leq 1,$$

where $x_e = \sum_{i \in \mathcal{I}} \ell_{ie} y_{ie}$, and e^j is removed from $S(t)$ after time ℓ_{ie} passes. The goal is to show that $S(t)$ converges in distribution to μ_n.

The set-valued process $S(t)$ is not Markov on its own, because the durations for which clones are occupied depends on the type of the proposal occupying it. Therefore, for the purpose of analysis, we label the type of the proposal occupying each clone and write

$$Z(t) := \left(S(t), (I_c(t))_{c \in S(t)}\right),$$

where $I_c(t) \in \mathcal{I}$ is the type of request using the clone c. For every product e and type i, define

$$\omega_{ie} := \frac{y_{ie} \ell_{ie}}{x_e}.$$

For a state (A, τ_A), where $A \in \mathcal{F}_n$ and $\tau_c \in \mathcal{I}$ denotes the type occupying $c \in A$, define

$$\Pi(A, \tau_A) := \mu_n(A) \prod_{c = e^j \in A} \omega_{\tau_c e}.$$

The occupied-set marginal of Π is μ_n. Now $Z(t)$ is a continuous-time Markov chain on $\mathrm{supp}(\Pi)$, with finite rates of transition between states. For any state in $\mathrm{supp}(\Pi)$, there is a positive rate at which it transitions to the empty state (if all occupied units return before there are any new arrivals of requests). Furthermore, for any state in $\mathrm{supp}(\Pi)$, there is a positive rate at which we transition to that state from the empty state (if we have arrivals of the appropriate types of requests and we choose the appropriate clones; note that $\mu_n(A) > 0$ and

$c \in A$, then $\mu_n(A \setminus \{c\}) > 0$). We conclude that $Z(t)$ is a non-explosive positive recurrent CTMC on $\mathrm{supp}(\Pi)$, and so converges to a unique stationary distribution (see Chapter 3 of (Norris, 1998)). We will show that Π is an invariant distribution for $Z(t)$, and therefore that $S(t)$ converges in distribution to μ_n, as needed.

Now fix two adjacent states in $\mathrm{supp}(\Pi)$ of the form $z = (B, \tau_B)$ and $z' = (B \cup \{c\}, (\tau_B, i))$, where $c = e^j \notin B$. The transition from z to z' occurs when a type-i proposal is assigned to clone c and is admitted. Its transition rate is therefore

$$q(z, z') = \frac{y_{ie}}{n} f(B, c).$$

The reverse transition occurs when the type-i occupation of c ends. Since its duration is exponential with mean ℓ_{ie}, the reverse transition rate is $q(z', z) = \frac{1}{\ell_{ie}}$. Using the definitions of f, ω_{ie}, and Π, we obtain

$$\begin{aligned}
\Pi(z) q(z, z') &= \mu_n(B) \left(\prod_{b = g^k \in B} \omega_{\tau_b g} \right) \frac{y_{ie}}{n} f(B, c) \\
&= \mu_n(B) \left(\prod_{b = g^k \in B} \omega_{\tau_b g} \right) \frac{y_{ie}}{n} \frac{n \mu_n(B \cup \{c\})}{x_e \mu_n(B)} \\
&= \mu_n(B \cup \{c\}) \left(\prod_{b = g^k \in B} \omega_{\tau_b g} \right) \frac{y_{ie}}{x_e} \\
&= \mu_n(B \cup \{c\}) \left(\prod_{b = g^k \in B} \omega_{\tau_b g} \right) \omega_{ie} \frac{1}{\ell_{ie}} \\
&= \Pi(z') q(z', z).
\end{aligned}$$

Thus every pair of adjacent states satisfies detailed balance. It follows that Π is invariant for $Z(t)$ as required.

EC.6. Discussion of Time Complexity of Algorithms

This appendix discusses how to implement (in polynomial time, up to a prescribed accuracy) the stationary OCRS policies constructed in the paper.

Computational requirements for polynomial-time running of a stationary OCRS. Fix $\boldsymbol{x} \in \mathcal{P}$ and let μ be any distribution on $\mathcal{F}$ that satisfies the stationary constraints (SELECTABILITY)–(STATIONARY-IMPLEMENTABILITY). The discussion after Proposition 2 gives an explicit implementation via a simulate-then-replace policy. This discussion showed that to run the stationary OCRS corresponding to μ, it suffices to have a polynomial-time procedure that samples (or approximately samples) from μ, and is able to (exactly or approximately) compute:

$$\mathbb{P}_{S \sim \mu}\left[e \in S \mid S_{-e} = T \right],$$

for any e and any set T that does not contain e. In our constructions, μ has additional structure (maximum-entropy for the first three environments, a Gibbs distribution for hypergraph matchings with vertex capacities, a mixture of maximum-entropy distributions for the knapsack constraint, and a KL-divergence based construction for weakly Rayleigh matroids), which reduces these requirements to simpler primitives. We spell this out in the following sections.

EC.6.1. The maximum-entropy constructions

In the three feasibility environments whose proofs follow the maximum-entropy method in Section 4 (matchings (Theorem 1), bipartite matchings (Theorem 2), and the rank-k uniform matroid (Theorem 3)), μ is the maximum-entropy distribution μ^* with prescribed marginals $\boldsymbol{p} = \alpha \boldsymbol{x}$ (Section 4.1). By Lemma 1, μ^* is a Gibbs distribution with parameters $\boldsymbol{w} = (w_e)_{e\in E}$ and constants $\rho_e = w_e/(1+w_e)$. In this case, simulate-then-replace takes the particularly simple form given by Algorithm 2.

The only online computations are (a) determining whether $T \cup \{e\} \in \mathcal{F}$ and (b) sampling a Bernoulli draw with parameter ρ_e/x_e. Thus, once the constants $\{\rho_e\}$ and an initial sample $\widehat{S} \sim \mu^*$ are available, each arrival can be processed in polynomial time. We thus need to explain:

(i) how we compute (approximately) the Gibbs parameters $\boldsymbol{w}$ and hence $\rho_e = w_e/(1+w_e)$ in polynomial time, so that μ^* has marginals $\boldsymbol{p} = \alpha\boldsymbol{x}$; and

(ii) how we can sample (approximately) an initial set $\widehat{S} \sim \mu^*$ in polynomial time.

Formally, by constructing polynomial time implementations for the above two steps, we will prove the following theorem:

THEOREM EC.6 **(Polynomial-time S-OCRS for matchings, k-uniform matroids).** *For every $\varepsilon > 0$:*

1. *There is a $(\frac{1}{3} - \varepsilon)$-selectable S-OCRS for the matching feasibility environment $(E, \mathcal{F})$ that runs in time* $\mathrm{poly}(|E|, 1/\varepsilon)$.
2. *There is a $(\frac{3-\sqrt{5}}{2} - \varepsilon)$-selectable S-OCRS for the bipartite matching feasibility environment $(E, \mathcal{F})$ that runs in time* $\mathrm{poly}(|E|, 1/\varepsilon)$.
3. *There is a $(\alpha_k - \varepsilon)$-selectable S-OCRS for the k-uniform matroid feasibility environment $(E, \mathcal{F})$ that runs in time* $\mathrm{poly}(|E|, \log(1/\varepsilon))$.

To compute Gibbs parameters as in step (i) above, we need to solve a maximum-entropy program. Before proceeding to the proof of Theorem EC.6, we summarize the approach of Singh and Vishnoi (2014); Straszak and Vishnoi (2019) to achieve this goal in a convenient form in the next subsection.

EC.6.2. Solving max-entropy/ minimum KL divergence programs

An initial observation is that if we fix $\boldsymbol{p}$ in the relative interior of $\mathcal{P}$, the maximum-entropy problem (MAX-ENTROPY) has a convex dual in variables $\boldsymbol{\theta} \in \mathbb{R}^E$:

$$\min_{\boldsymbol{\theta}\in\mathbb{R}^E} h(\boldsymbol{\theta}) \;:=\; \log \mathsf{Z}(\boldsymbol{\theta}) \;-\; \langle \boldsymbol{\theta}, \boldsymbol{p}\rangle, \qquad \mathsf{Z}(\boldsymbol{\theta}) = \sum_{S\in\mathcal{F}} \exp\Big(\sum_{e\in S}\theta_e\Big). \tag{EC.27}$$

The gradient of h is the marginal mismatch:

$$(\nabla h(\boldsymbol{\theta}))_e = \mathbb{P}_{S\sim\mu_{\boldsymbol{\theta}}}[e\in S] \;-\; p_e,$$

where $\mu_{\boldsymbol{\theta}}$ is the Gibbs distribution proportional to $\exp(\sum_{e\in S}\theta_e)\mathbf{1}\{S\in\mathcal{F}\}$. Thus, in any environment where we can evaluate or approximate the Gibbs marginals $\mathbb{P}_{S\sim\mu_{\boldsymbol{\theta}}}[e\in S]$ we should be able to use standard convex optimization algorithms (ellipsoid methods or first-order methods) and compute an ε-approximate minimizer $\boldsymbol{\theta}$ in time polynomial in $|E|$ and $\log(1/\varepsilon)$, given a good bound on the search region for $\boldsymbol{\theta}$.[25] From $\boldsymbol{\theta}$ we recover the optimal weights $\boldsymbol{w}$ and ρ_e via $w_e = e^{\theta_e}$ and $\rho_e = w_e/(1+w_e)$.

However, establishing a good bound on the search region for $\boldsymbol{\theta}$ is not always easy. This is the principal contribution of Singh and Vishnoi (2014). Unfortunately, their running time is still polynomial in the inverse of the distance of $\boldsymbol{p}$ from the boundary of $\mathcal{P}$ (which can happen, for instance, if $\boldsymbol{x}$ has a particularly small coordinate). Furthermore, their results are not stated for the case of minimizing KL divergence, which is useful for us when working with weakly Rayleigh matroids. Straszak and Vishnoi (2019) remove this dependence on the distance from the boundary, and generalize to minimizing KL divergence.

Formally, fix a reference distribution $\mu_0 \in \Delta(\mathcal{F})$ with full support, i.e. $\mu_0(S) > 0$ for all $S \in \mathcal{F}$, and fix a target marginal vector $\boldsymbol{p} \in \mathcal{P}$. Consider the convex program

$$\min\Big\{\mathrm{KL}(\nu\|\mu_0):\ \nu \in \Delta(\mathcal{F}),\ \mathbb{P}_{S\sim\nu}[e \in S] = p_e\ \ \forall e \in E\Big\}, \tag{EC.28}$$

where $\mathrm{KL}(\nu\|\mu_0) := \sum_{S\in\mathcal{F}} \nu(S)\log\big(\nu(S)/\mu_0(S)\big)$. Let μ^* denote an optimal solution. When μ_0 is uniform on $\mathcal{F}$, the objective differs from $\mathsf{H}(\mu)$ by an additive constant, so this is equivalent to the usual maximum-entropy program (MAX-ENTROPY).

For $\boldsymbol{\theta} \in \mathbb{R}^E$ define the tilt of μ_0 by

$$\mu_{\boldsymbol{\theta}}(S) \ :=\ \frac{\mu_0(S)\exp\big(\sum_{e\in S}\theta_e\big)}{\sum_{T\in\mathcal{F}}\mu_0(T)\exp\big(\sum_{e\in T}\theta_e\big).} \qquad (S \in \mathcal{F}),$$

and let $g_{\mu_0}(\boldsymbol{w}) = \sum_{S\in\mathcal{F}} \mu_0(S)\prod_{f\in S} w_f$, the weighted base generating polynomial. The work of Straszak and Vishnoi (2019) shows that access to a counting oracle for g_{μ_0} suffices to solve the convex program.

DEFINITION EC.3 (APPROXIMATE GENERALIZED COUNTING ORACLE WITH COST $\mathsf{T}(\eta)$). For $e \in E$ define the marginal sum

$$g_{\mu_0,e}(\boldsymbol{w}) \ :=\ \sum_{\substack{S\in\mathcal{F}:\\ S\ni e}} \mu_0(S)\prod_{f\in S} w_f.$$

An oracle $\mathcal{O}$ takes as input $(\boldsymbol{w},\eta)$ with $\boldsymbol{w} \in \mathbb{R}^E_{>0}$ and $\eta \in (0,1/2)$, and returns numbers $\widehat{g}$ and $\{\widehat{g}_e\}_{e\in E}$ such that

$$(1-\eta)\,g_{\mu_0}(\boldsymbol{w}) \ \le\ \widehat{g}\ \le\ (1+\eta)\,g_{\mu_0}(\boldsymbol{w}), \qquad (1-\eta)\,g_{\mu_0,e}(\boldsymbol{w}) \ \le\ \widehat{g}_e\ \le\ (1+\eta)\,g_{\mu_0,e}(\boldsymbol{w}) \quad \forall e \in E,$$

and runs in time $\mathsf{T}(\eta)$ (suppressing polynomial factors in $|E|$ and the bit-complexity of $\boldsymbol{w}$).

We need two further technical definitions to state the formal result.

DEFINITION EC.4 (UNARY FACET COMPLEXITY). Let $\mathcal{P} \subseteq \mathbb{R}^E$ be a polytope with integer vertices. Its *unary facet complexity*, denoted $\mathrm{fc}(\mathcal{P})$, is the smallest $M \in \mathbb{N}$ for which there exist finitely many vectors $a^{(j)} \in \mathbb{Z}^E$ with $\|a^{(j)}\|_\infty \le M$, scalars $b_j \in \mathbb{R}$, and a linear subspace $H \subseteq \mathbb{R}^E$ such that

$$\mathcal{P} = \Big(\{x \in \mathbb{R}^E:\ \langle a^{(j)}, x\rangle \le b_j \text{ for all } j\}\Big) \cap H.$$

[25] In our applications we take $\boldsymbol{p} = \alpha\boldsymbol{x}$ with a fixed $\alpha < 1$, which places $\boldsymbol{p}$ in the relative interior of $\mathcal{P}$ and ensures that an optimal dual solution exists with finite coordinates.

DEFINITION EC.5 (LOG-MASS PARAMETER). Define

$$L_{\mu_0} \;:=\; \max_{S\in\mathcal{F}} \;|\log \mu_0(S)|.$$

THEOREM EC.7. *Let $\mathcal{F}\subseteq 2^E$ and let $\mathcal{P}=\operatorname{conv}\{\mathbf{1}_S : S\in\mathcal{F}\}$. Let $M:=\operatorname{fc}(\mathcal{P})$. Fix any $\mu_0\in\Delta(\mathcal{F})$ with full support and define L_{μ_0} as above. There exists an algorithm which, given: (i) $\boldsymbol{p}\in\mathcal{P}$, (ii) $\varepsilon\in(0,1)$, and (iii) oracle access to $\mathcal{O}$ for μ_0, outputs a vector $\boldsymbol{\theta}\in\mathbb{R}^E$ such that*

$$\|\mu_{\boldsymbol{\theta}}-\mu^*\|_1 \;<\; \varepsilon,$$

where $\mu_{\boldsymbol{\theta}}$ is the tilted Gibbs distribution above. Moreover, the algorithm guarantees a norm bound of the form

$$\|\boldsymbol{\theta}\|_2 \;\le\; \operatorname{poly}\Big(|E|,\; M,\; L_{\mu_0},\; \log(1/\varepsilon)\Big).$$

The algorithm makes $\operatorname{poly}(|E|,M,L_{\mu_0},\log(1/\varepsilon))$ *oracle calls, all with a common accuracy parameter*

$$\eta \;=\; \frac{1}{\operatorname{poly}(|E|,1/\varepsilon)},$$

and hence runs in total time

$$\operatorname{poly}\Big(|E|,\; M,\; L_{\mu_0},\; \log(1/\varepsilon)\Big)\cdot \mathsf{T}(\eta).$$

Note that while the results in Straszak and Vishnoi (2019) are stated for an evaluation oracle that can compute $g_{\mu_0}(\boldsymbol{w})$ exactly, examination of the argument reveals that it also works for an approximate oracle, as long as we can also access values of $g_{\mu_0,e}(\boldsymbol{w})$. This type of approximate counting oracle is what is used in Singh and Vishnoi (2014).

Theorem EC.7 makes $\operatorname{poly}(|E|,M,L_{\mu_0},\log(1/\varepsilon))$ oracle calls, all at a common accuracy $\eta = 1/\operatorname{poly}(|E|,1/\varepsilon)$. Consequently:

- If the generalized counting oracle can be answered to relative error η in time $\operatorname{poly}(|E|,\log(1/\eta))$ (a high-precision, or evaluation, oracle model), then the Gibbs parameters can be computed in time $\operatorname{poly}(|E|,\log(1/\varepsilon))$. We note that in this case, it actually suffices to have an oracle for g_{μ_0} alone, since $g_{\mu_0,e}$ can be calculated by interpolation, see Straszak and Vishnoi (2019) for details.
- If the best available oracle is only an FPRAS-type primitive with running time $\operatorname{poly}(|E|,1/\eta)$, then the same reduction yields running time polynomial in $1/\varepsilon$.

EC.6.3. Proof of Theorem EC.6 for matchings and uniform matroids, and discussion of hypergraph matchings with vertex capacities

We now discuss for matchings and uniform matroids the (non-)existence of the needed approximate counting oracles, and our ability to sample from the max-entropy distribution. We will also indicate which oracle regime applies in each feasibility environment.

Matchings, bipartite matchings, and hypergraph matchings. For general matchings (and hence bipartite matchings) classical MCMC methods give an FPRAS for the matching partition function, and hence for the functions g_{μ_0} and $g_{\mu_0,e}$ (see chapter 12, Jerrum and Sinclair (1996)). The running time is polynomial in the problem size and $1/\eta$ for relative error η. Plugging this into Theorem EC.7 yields a randomized polynomial-time procedure for computing the Gibbs parameters needed by Algorithm 2, with overall running time polynomial in $1/\varepsilon$ and $\log(1/\delta)$ where δ is the allowed probability of error (since the unary facet complexity of the matching polytope is just 1, and the log-mass parameter is polynomially bounded). Note that the counting oracle being randomized is not an issue since we only make polynomially many calls to the oracle (the success probability can always be boosted appropriately).

Once we have access to the Gibbs parameters, there is a rapidly mixing Markov chain whose stationary distribution is the Gibbs distribution; see Chapter 12 in Jerrum and Sinclair (1996). Thus we can sample $\widehat{S} \sim \mu^*$ up to total-variation error ε in time polynomial in the instance size and $\log(1/\varepsilon)$.

Putting these pieces together yields an OCRS for matchings that obtains a selectability of $1/3 - \varepsilon$ (for general matchings) or $(3-\sqrt{5})/2 - \varepsilon$ (for bipartite matchings), that runs in time polynomial in the problem size and $1/\varepsilon$. Note that most of the time is spent 'preprocessing' and calculating the Gibbs parameters, and both sampling the Gibbs distribution, and actually running the algorithm are fast.

In the general case of hypergraph matchings, there is no efficient counting oracle, nor is the unary facet complexity necessarily polynomially bounded, so we do not claim a polynomial time algorithm. In special subclasses of hypergraphs, there are sometimes evaluation oracles and bounds on the facet complexity, in which case solving the maximum entropy program in polynomial time is possible. For such classes, we can also utilize a standard counting to sampling reduction of Jerrum et al. (1986) to sample from the Gibbs distribution and then run our algorithm.

***k*-uniform matroids.** In this case, a simple DP based method provides an evaluation oracle, and we can then apply Theorem EC.7. But Chen et al. (1994) also provide a direct method to compute the Gibbs parameters needed by Algorithm 2 in polynomial time to any needed precision, i.e., to an error of ε in time polynomial in $\log(1/\varepsilon)$. They also show how to sample from the resulting Gibbs distribution in polynomial time. Thus, we have an OCRS for k-uniform matroids with a selectability of $\alpha_k - \varepsilon$ that runs in time polynomial in the problem parameters and $\log(1/\varepsilon)$.

Hypergraph matchings with vertex capacities. For our S-OCRS for hypergraph matchings with vertex capacities discussed in Section 6, μ is a Gibbs distribution with constants $\rho_e = \gamma x_e$, where γ was an explicit constant. Like in the previous feasibility environments, simulate-then-replace takes the particularly simple form given by Algorithm 2. There are no online computations, since ρ_e/x_e is constant. Thus, the only requirement to implement the algorithm efficiently is that we sample (approximately) an initial set $\widehat{S} \sim \mu$ in polynomial time. Unfortunately, even knowing ρ_e there is no general efficient way to sample from this Gibbs distribution, so we do not claim a polynomial time algorithm (though we give a different poly-time OCRS in Section 6). Once again, if we have access to an appropriate counting oracle, and bounds on the facet complexity, we can utilize the standard counting to sampling reduction (Jerrum et al., 1986) to sample from the Gibbs distribution and then run our algorithm.

EC.6.4. Knapsack constraints

For the knapsack constraint, we have the following theorem:

THEOREM EC.8 **(Polynomial-time S-OCRS for knapsack constraint).** *For every $\varepsilon > 0$, there is a $(\frac{1}{4} - \varepsilon)$-selectable S-OCRS for the knapsack feasibility environment $(E, \mathcal{F})$ that runs in time $\mathrm{poly}(|E|, 1/\varepsilon)$.*

The S-OCRS for knapsack constraints discussed in section EC.3 is a mixture of two different maximum-entropy distributions, μ^H and μ^L. Accordingly, simulate-then-replace involves tossing a coin, and choosing to run Algorithm 2 with either the large item S-OCRS, or the small item S-OCRS. For each of these S-OCRS, we need to be able to sample from an initial distribution, and compute appropriate conditional probabilities.

Unfortunately for the knapsack polytope, we cannot invoke Theorem EC.7 directly, since there is no polynomial-sized bound on the unary facet complexity. Instead, we apply Theorem 2.8 from Singh and Vishnoi (2014). This result is analogous to Theorem EC.7, but does not require a bound on the facet complexity; instead there is an assumption on the marginals lying in the relative interior of the polytope, and the runtime of the algorithm is dependent on the distance of the marginals from the boundary.

To get around this dependence, we start by discarding each active item with probability ε; this effectively scales the x_e down by a factor of $1-\varepsilon$ and impacts the selectability by a factor of at most $1-\varepsilon$, but lets us assume that $x_e \le 1-\varepsilon$. We will also add the items with $x_e \le \frac{\varepsilon}{|E|}$ to the heavy items; this does not change any of the calculations in section EC.3, except for the bound on $X + 2B$, which is now bounded by $2+\varepsilon$ rather than 2.

For the heavy item S-OCRS, we select at most one item, and we calculated in Section EC.3 that: $\mathbb{P}_{S\sim\mu^H}[e \in S \mid S_{-e} = \emptyset] = \frac{x_e}{1+x_e}$. It is thus straightforward to sample an initial state from this distribution, and to process each arrival in polynomial time. For the light-item S-OCRS, we have:

$$E_L := \{e \in E : a_e \le 1/2, x_e > \varepsilon/|E|\}, \quad \mathcal{F}_L := \{T \subseteq E_L : a(T) \le 1\}.$$

The construction in Section EC.3 chose $c_L := \frac{1}{1+2\sum_{e\in E_L} a_e x_e}$, let $p_e = c_L x_e$ for $e \in E_L$, and $\mu^L \in \Delta(\mathcal{F}_L)$ was the maximum-entropy distribution with marginals p. Once the constants $\{\rho_e\}$ associated with this maximum-entropy distribution and an initial sample $\widehat{S} \sim \mu^L$ are available, each arrival can be processed in polynomial time, just as discussed in Section EC.6.1. Therefore it remains to show how to compute these parameters and sample the initial state in polynomial time.

Now we know that $p_e = c_L x_e$, and $c_L \le 1$. Moreover, $B \le 1$ implies $c_L \ge 1/3$. So we conclude that:

$$\frac{\varepsilon}{3|E|} \le p_e \le x_e \le 1-\varepsilon.$$

Thus, $\boldsymbol{p}$ lies in the relative interior of the knapsack polytope, and the distance from the boundary is polynomial in $|E|$ and $\frac{1}{\varepsilon}$. Next, we must find approximate counting oracles for the knapsack polytope. For weights $\boldsymbol{w} = (w_e)_{e\in E_L}$, define

$$g(\boldsymbol{w}) := \sum_{S\in\mathcal{F}_L} \prod_{e\in S} w_e, \quad g_e(\boldsymbol{w}) := \sum_{\substack{S\in\mathcal{F}_L\\ e\in S}} \prod_{f\in S} w_f.$$

Let $(Y_e)_{e\in E_L}$ be independent random variables with $\mathbb{P}[Y_e = 1] = \frac{w_e}{1+w_e}$, and write $R = \{e : Y_e = 1\}$. Then

$$g(\boldsymbol{w}) = \left(\prod_{e\in E_L}(1+w_e)\right)\mathbb{P}[a(R)\le 1], \quad g_e(\boldsymbol{w}) = \left(\prod_{f\in E_L}(1+w_f)\right)\mathbb{P}[e\in R,\ a(R)\le 1].$$

Theorem 3.6 of Gopalan et al. (2011) computes an approximation to $\mathbb{P}[a(R)\le 1]$ in time polynomial in $|E|$, the input encoding length, and the inverse relative error. The same theorem computes an approximation for g_e after fixing $Y_e = 1$ and reducing the capacity by a_e.

The approximate counting oracle, together with bound on the distance of $\boldsymbol{p}$ from the boundary of the polytope mean we can apply Theorem 2.8 in Singh and Vishnoi (2014) to compute the needed Gibbs distribution to accuracy ε in time polynomial in $|E|$ and $1/\varepsilon$. Once the approximate Gibbs parameters $\boldsymbol{w}$ are available, we can approximately sample from the corresponding Gibbs distribution using the standard counting-to-sampling self-reduction (Jerrum et al., 1986).

EC.6.5. Weakly Rayleigh matroids

The weakly Rayleigh matroid construction differs from the first three environments: the distribution μ^* is *not* the maximum-entropy distribution on independent sets. Instead, the construction proceeds by (i) finding a distribution on bases with prescribed marginals (a KL projection of a Rayleigh base measure), and (ii) applying independent thinning inside a sampled base. We therefore discuss its algorithmic implementation separately.

Implementing our algorithm requires certain natural assumptions—satisfied by a broad subclass of weakly Rayleigh matroids. The following theorem consolidates the assumptions we need to obtain a polynomial-time computational guarantee for the modified template that we use for weakly Rayleigh matroids.

THEOREM EC.9 **(Polynomial-time S-OCRS for weakly Rayleigh matroids).** *Let $\mathcal{M} = (E, \mathcal{I})$ be a weakly Rayleigh matroid (Definition EC.2) of rank r on $n := |E|$ elements, with Rayleigh base measure μ_0. Suppose the following conditions hold:*

(i) ***Rank oracle:*** *There is an oracle that, given $T \subseteq E$, returns $\text{rank}_{\mathcal{M}}(T)$ in polynomial time.*

(ii) ***High precision/ evaluation oracle:*** *There is an oracle that, given $\boldsymbol{w} \in \mathbb{R}^E_{>0}$ and accuracy parameter $\eta \in (0, 1/2)$, returns a multiplicative $(1\pm\eta)$-approximation to the weighted base generating polynomial, that is,*

$$g_{\mu_0}(\boldsymbol{w}) := \sum_{B\in\mathcal{B}} \mu_0(B) \prod_{f\in B} w_f\ ,$$

in time polynomial in $n, \log(1/\eta)$, and the bit complexity of $\boldsymbol{w}$.

(iii) ***Bounded log-mass:*** *The log-mass parameter $L_{\mu_0} := \max_{B\in\mathcal{B}} |\log \mu_0(B)|$ is polynomially bounded in n.*

Then, for every $\boldsymbol{x} \in \mathcal{P}$ and every $\varepsilon > 0$, there is a $(\frac{1}{2}-\varepsilon)$-selectable S-OCRS for the matroid feasibility environment $(E, \mathcal{F})$ that runs in time $\text{poly}(n, \log(1/\varepsilon))$.[26]

[26] Note the assumption (i) is not strictly speaking necessary, since a rank oracle can be implemented using an evaluation oracle under the assumption of bounded log-mass.

REMARK EC.2 ($\mathbb{C}$-REPRESENTABLE MATROIDS SATISFY THE ORACLE ASSUMPTIONS). Any subclass of weakly Rayleigh matroids for which an efficient evaluation oracle exists and for which the log-mass parameter is polynomially bounded falls within the scope of Theorem EC.9. The conditions of Theorem EC.9 are satisfied in particular by $\mathbb{C}$-representable matroids.

If $\mathcal{M}$ is represented by a matrix $A \in \mathbb{C}^{r\times n}$, there is a natural base measure that is *determinantal* (see Lyons (2003) for a discussion): $\mu_0(B) \propto \det(A_B^* A_B)$, and this measure is Rayleigh. The Cauchy–Binet formula gives that for every $\boldsymbol{w} \in \mathbb{R}^n_{>0}$, we have $g_{\mu_0}(\boldsymbol{w}) = \det(A\,\mathrm{diag}(\boldsymbol{w})\,A^*)/\det(AA^*)$, computable in $\mathrm{poly}(n, \log(1/\eta))$ time via standard determinant computation. The log-mass bound follows from Hadamard's inequality. This class includes graphic matroids and laminar matroids.

For graphic matroids in particular, we can let μ_0 be uniform on spanning trees. The evaluation oracle reduces to computing the weighted spanning tree polynomial, which equals a determinant of a weighted Laplacian by Kirchhoff's matrix-tree theorem.

The remainder of this subsection is devoted to the proof of Theorem EC.9.

Proof of Theorem EC.9. We start by recalling the construction of μ^*. Fix $\boldsymbol{x} \in \mathcal{P}$. First, compute a dominating base point $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ with $\boldsymbol{q} \geq \boldsymbol{x}$ (Lemma EC.1). Next, construct a probability measure μ on bases $\mathcal{B}$ such that

$$\mathbb{P}_{B\sim\mu}[e \in B] = q_e \quad \forall e \in E, \qquad \text{and} \qquad \mu \text{ is Rayleigh (Definition EC.2).}$$

When $\boldsymbol{q}$ lies in the relative interior of $\mathcal{P}_{\mathcal{B}}$, such a measure μ can be obtained as the unique minimizer of the KL projection program (MIN-KL), and it has the exponential-family form $\mu = \mu_{\boldsymbol{w}}$ for some $\boldsymbol{w} \in \mathbb{R}^E_{>0}$. When $\boldsymbol{q}$ lies on the boundary of $\mathcal{P}_{\mathcal{B}}$, an exact exponential tilt need not exist; in this case we take any sequence $\boldsymbol{q}^{(t)}$ in the relative interior of $\mathcal{P}_{\mathcal{B}}$ with $\boldsymbol{q}^{(t)} \to \boldsymbol{q}$, form the corresponding tilted measures $\mu^{(t)} = \mu_{\boldsymbol{w}^{(t)}}$, and pass to a limit along a convergent subsequence. Since $\mathcal{B}$ is finite, this produces a well-defined limiting base measure μ with marginals $\boldsymbol{q}$; moreover the Rayleigh inequalities are preserved under this limit, so μ is Rayleigh. (Note that when we try to compute $\boldsymbol{w}$, we will not take this limit, but just solve the program approximately even when $\boldsymbol{q}$ lies on the boundary.)

Finally, define the retention probabilities

$$\tau_e := \begin{cases} \dfrac{x_e}{2q_e}, & q_e > 0, \\ 0, & q_e = 0, \end{cases}$$

and generate $S \sim \mu^*$ by drawing $B \sim \mu$ and then keeping each element $e \in B$ independently with probability τ_e.

Simulate-then-replace for $\boldsymbol{\mu^*}$. Pre-sample $\widehat{S} \sim \mu^*$. When e arrives, let $T = \widehat{S}_{-e}$ and define the target conditional inclusion probability

$$\pi_e(T) := \mathbb{P}_{S\sim\mu^*}\big[e \in S \mid S_{-e} = T\big] = \frac{\mu^*(T \cup \{e\})}{\mu^*(T) + \mu^*(T \cup \{e\})}.$$

Since μ^* is supported on $\mathcal{I}$, we have $\pi_e(T) = 0$ whenever $T \cup \{e\} \notin \mathcal{I}$. Now:

- If $e \notin A$, set $\widehat{S} \leftarrow T$.
- If $e \in A$ and $T \cup \{e\} \notin \mathcal{I}$, set $\widehat{S} \leftarrow T$.
- If $e \in A$ and $T \cup \{e\} \in \mathcal{I}$, accept e with probability $\pi_e(T)/x_e$ and update

$$\widehat{S} \leftarrow \begin{cases} T \cup \{e\}, & \text{if accepted,} \\ T, & \text{if rejected.} \end{cases}$$

Clearly, the online computations are (a) determining whether $T \cup \{e\} \in \mathcal{I}$ (an independence test in $\mathcal{M}$) and (b) sampling a Bernoulli draw with parameter $\pi_e(T)/x_e$. Thus, once an initial sample $\widehat{S} \sim \mu^*$ is available and we can evaluate (or approximate) the conditionals $\pi_e(T)$ for the encountered sets T, each arrival can be processed in polynomial time.

The remainder of the proof is devoted to explaining:

(i) how we compute a dominating base point $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ and how we approximate a Rayleigh base measure μ on $\mathcal{B}$ with marginals $\boldsymbol{q}$; and

(ii) how we can approximately sample from μ (and hence from μ^*) and evaluate the ratios $\mu^*(T \cup \{e\})/\mu^*(T)$ needed to compute $\pi_e(T)$.

Step (i): Finding a dominating base point. Let us start by noting that the first step of calculating an ε-approximate base point $\boldsymbol{q} \in \mathcal{P}_{\mathcal{B}}$ which dominates $\boldsymbol{x}$ can be done in time polynomial in $\log(1/\varepsilon)$. Indeed, it suffices to obtain a maximizer $\sum_e y_e$ over the polytope:

$$Q := \Big\{ \boldsymbol{y} \in [0,1]^E : \sum_{e \in T} y_e \le \text{rank}_{\mathcal{M}}(T) \ \forall T \subseteq E, \ \ \boldsymbol{y} \ge \boldsymbol{x} \Big\},$$

and this is possible using the rank oracle, via the ellipsoid method (see for instance Grötschel et al. (1988)). Note that we have already shown in Section EC.5.13 that the optimizer necessarily satisfies $\sum_e y_e = r$.

Step (ii): KL projection onto a Rayleigh base measure. Here, we use access to our evaluation oracle, and Theorem EC.7 in a black box manner, to find an ε-approximate Rayleigh measure μ with the (approximately) correct marginals in running time polynomial in $\log(1/\varepsilon)$. (Note that the matroid base polytope has unary facet complexity 1, Theorem EC.7 accepts points on the boundary of the polytope as input, and the discussion after Theorem EC.7 clarifies that a high-precision or evaluation oracle for g_{μ_0} suffices to evaluate $g_{\mu_0,e}$ using interpolation.) At the end of this step, we have access to approximate Gibbs parameters $\boldsymbol{w}$.

Step (iii): Sampling $\boldsymbol{B \sim \mu}$ and sampling from $\boldsymbol{\mu^*}$. To sample from μ, we follow a standard counting-to-sampling self-reduction template (see Jerrum et al. (1986)). For disjoint sets $I, J \subseteq E$, define the constrained function

$$g_{\mu_0,J,I}(\boldsymbol{w}) := \sum_{B \in \mathcal{B}:\ I \subseteq B,\ B \cap J = \emptyset} \mu_0(B) \prod_{f \in B} w_f.$$

Consider the bivariate polynomial

$$F_{I,J}(t,u) := g_{\mu_0}(\boldsymbol{w}(t,u)), \quad \text{where} \quad w_i(t,u) = \begin{cases} t w_i, & i \in I, \\ u w_i, & i \in J, \\ w_i, & \text{otherwise.} \end{cases}$$

Since g is multi-affine, $F_{I,J}(t,u)$ has $\deg_t \le |I|$ and $\deg_u \le |J|$, we can check directly that

$$g_{\mu_0,J,I} \;=\; [t^{|I|}u^0]\, F_{I,J}(t,u).$$

Let $d := |I|$ and $m := |J|$. Using univariate interpolation at the positive points $\{1,2,\ldots,m+1\}$ in the u-variable and a d-th forward-difference identity in the t-variable, one may write $g_{\mu_0,J,I}$ as the following explicit linear combination of evaluations of $g_{\mu_0}(\cdot)$ at positive points:

$$g_{\mu_0,J,I}(\boldsymbol{w}) \;=\; \frac{1}{d!}\sum_{a=0}^{d}(-1)^{d-a}\binom{d}{a}\sum_{b=1}^{m+1}(-1)^{b-1}\binom{m+1}{b}\, g_{\mu_0}\big(\boldsymbol{w}(1+a,b)\big), \tag{EC.29}$$

Thus, with access to an oracle that can evaluate $g_{\mu_0}(\cdot)$ to relative error η in time polynomial in the problem parameters and $\log(1/\eta)$, each constrained value $g_{\mu_0,J,I}(\boldsymbol{w})$ can also be computed to relative error η in the same polynomial time.

To achieve this, first of all, note that we can detect whether the polynomial is 0 using the rank oracle for the matroid. If it is non-zero, the bound on the log-mass guarantees that the value cannot be less than exponentially small. We can therefore choose the evaluation accuracy used for the terms $g_{\mu_0}(\boldsymbol{w}(1+a,b))$ sufficiently small so that the resulting error in the linear combination translates to the desired relative accuracy for $g_{\mu_0,J,I}(\boldsymbol{w})$. Requesting this higher precision increases the running time only polynomially, due to the fact that the sum of the coefficients of $g_{\mu_0,J,I}(\boldsymbol{w})$ is bounded exponentially in the size of the matroid, and we use a high-precision oracle.

Note that a similar argument is used in Straszak and Vishnoi (2019) to calculate marginals.

Now fix any ordering $1,2,\ldots,n$ and run the following procedure: initialize $I \leftarrow \emptyset$ and $J \leftarrow \emptyset$. For $k = 1,2,\ldots,n$, compute

$$Z_1 \;:=\; g_{\mu_0,J,I\cup\{k\}}(\boldsymbol{w}), \qquad Z_0 \;:=\; g_{\mu_0,J\cup\{k\},I}(\boldsymbol{w}),$$

(using only evaluations of $g_{\mu_0}(\cdot)$ via (EC.29)). Define

$$\pi_k \;:=\; \mathbb{P}_{B\sim\mu}\big[k\in B \mid I\subseteq B,\ B\cap J=\emptyset\big] \;=\; \frac{Z_1}{Z_1+Z_0} \;=\; \frac{Z_1}{g_{\mu_0,J,I}(\boldsymbol{w})}.$$

Include k in I with probability π_k and otherwise add it to J. Lastly, output $B := I$. If, at each step k, we use an approximate conditionals $\tilde{\pi}_k$ satisfying $|\tilde{\pi}_k - \pi_k| \;\le\; \varepsilon/n$, the distribution of the output base $\tilde{B}$ is within ε in total variation distance of μ by a standard step-by-step coupling argument. Thus we obtain an approximate sample of μ in time polynomial in $\log(1/\varepsilon)$.[27]

Then, sample from μ^* by keeping each element of the base with probability $\tau_e = \frac{x_e}{2q_e}$.

Step (iv): Compute $\boldsymbol{\mu}^*$(T) for any set: Fix $T \subseteq [n]$. Clearly we have:

$$\mu^*(T) = \sum_{B\supseteq T}\mu(B)\Big(\prod_{i\in T}\tau_i\Big)\Big(\prod_{i\in B\setminus T}(1-\tau_i)\Big).$$

[27] We note that in some special cases, we can also sample using a random walk on the bases, see Anari et al. (2021), for instance.

(Recall that $\tau_i = \frac{x_i}{2q_i}$.) To calculate $\mu^*(T)$, we follow a similar interpolation strategy as above. Define a weight vector $\boldsymbol{w}^{(T)}(t) \in \mathbb{R}^n_{>0}$ by

$$w_i^{(T)}(t) \;:=\; \begin{cases} t\, w_i, & i \in T, \\ w_i(1-\tau_i), & i \notin T, \end{cases} \qquad (t > 0).$$

The polynomial $F_T(t) := g_{\mu_0}(\boldsymbol{w}^{(T)}(t))$ has degree at most $|T|$ and

$$[t^{|T|}]\, g_{\mu_0}(\boldsymbol{w}^{(T)}(t)) \;=\; \sum_{B \supseteq T} \mu_0(B) \Big(\prod_{i \in T} w_i \Big) \Big(\prod_{i \in B \setminus T} w_i(1-\tau_i) \Big),$$

which means

$$\mu^*(T) \;=\; \frac{\prod_{i \in T} \tau_i}{g_{\mu_0}(\boldsymbol{w})} \cdot [t^{|T|}]\, g_{\mu_0}(\boldsymbol{w}^{(T)}(t)).$$

Let $d := |T|$. Since $F_T(t) = g_{\mu_0}(\boldsymbol{w}^{(T)}(t))$ has degree at most d, we can calculate its coefficients from positive evaluations using the same interpolation trick as before:

$$[t^d]\, F_T(t) \;=\; \frac{1}{d!} \sum_{a=0}^{d} (-1)^{d-a} \binom{d}{a} F_T(1+a),$$

so computing $\mu^*(T)$ reduces to evaluations of g_{μ_0}. We can therefore compute $\mu^*(T)$ (and the ratio $\frac{\mu^*(T \cup \{e\})}{\mu^*(T) + \mu^*(T \cup \{e\})}$ to accuracy ε in time polynomial in $\log(1/\varepsilon)$). This is exactly the acceptance probability used by the online simulate-then-replace policy. Combining Steps (i)–(iv) completes the proof of Theorem EC.9: under the stated oracle assumptions, we obtain a $(\frac{1}{2} - \varepsilon)$-selectable S-OCRS for weakly Rayleigh matroids, running in time $\mathrm{poly}(n, \log(1/\varepsilon))$. □

EC.7. Future directions

A significant open problem is to determine whether it is possible to find a $1/2$-selectable S-OCRS for general matroids. We conjecture that the answer is yes. As discussed in the introduction, while the existence of a $1/2$-selectable OCRS is known via indirect duality arguments, it would be interesting to find a more explicit OCRS. The KL-divergence minimization based approach is not the only way to sample a distribution over bases of a matroid with the correct marginals; perhaps it is possible to use other methods and thin the sampled base in a correlated way to obtain a witness distribution for the stationary OCRS LP.

Another open problem is to the optimal selectability of an S-OCRS for matchings, or at the very least, determine whether the maximum-entropy approach gives $(\sqrt{3}-1)/2$-selectable S-OCRS. Once again, we conjecture that the answer to this question is yes. Also open is the optimal selectability for an S-OCRS for knapsack constraints.

One problem arising from our work is to understand the gap between OCRS and S-OCRS better. We should clarify that although one may conjecture that S-OCRSs can obtain the same selectability as OCRSs, that would be false—for instance, we can show that the best S-OCRS for 2-uniform matroids has a selectability of $3/5$, while the best OCRS has a selectability > 0.61, and the best S-OCRS for the knapsack relaxation has a selectability at most $3/10$, while the best OCRS has selectability $1/(3+e^{-2})$ (Jiang et al., 2025a). Understanding this gap better, and possibly characterizing it in different feasibility environments is a compelling avenue for future research.

We showed in Proposition 1 that for any feasibility environment a stationary OCRS with a selectability of α necessarily yields a recurring OCRS with a selectability of α. We conjecture, but are unable to prove that the converse is true as well.

Finally, OCRSs have seen significant applications in revenue management. In a companion paper Aminian et al. (2026b), we develop an application of S-OCRSs to Bayesian online allocation of reusable resources, originally studied in Rusmevichientong et al. (2020) and further developed in Feng et al. (2022). It would be useful to find other applications of S-OCRSs beyond the reusable setting. We leave the development of connections between the concept of S-OCRS and applications in revenue management, as well as queueing theory, as interesting directions for future research.

References

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

Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing*, 43 (2):930–972, 2014.

Saeed Alaei, Jason Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. *Games and Economic Behavior*, 118:494–510, 2019.

Saeed Alaei, Ali Makhdoumi, Azarakhsh Malekian, and Rad Niazadeh. Descending price auctions with bounded number of price levels and batched prophet inequality. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 246–246, 2022.

Mohammad Reza Aminian, Will Ma, and Linwei Xin. Online job selection: Reward rate vs. remaining value. *Available at SSRN 4644495*, 2026a.

Mohammad Reza Aminian, Rad Niazadeh, and Pranav Nuti. Optimal bayesian online allocation of reusable resources. *Available at SSRN 6384158*, 2026b.

Nima Anari, Rad Niazadeh, Amin Saberi, and Ali Shameli. Nearly optimal pricing algorithms for production constrained and laminar bayesian selection. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 91–92, 2019.

Nima Anari, Kuikui Liu, Shayan Oveis Gharan, Cynthia Vinzant, and Thuy-Duong Vuong. Log-concave polynomials iv: approximate exchange, tight mixing times, and near-optimal sampling of forests. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 408–420, 2021.

Ruicheng Ao, Hengyu Fu, and David Simchi-Levi. Two-stage online reusable resource allocation: Reservation, overbooking and confirmation call. *arXiv preprint arXiv:2410.15245*, 2024.

Nick Arnosti and Will Ma. Tight guarantees for static threshold policies in the prophet secretary problem. *Operations research*, 71(5): 1777–1788, 2023.

Pablo D Azar, Robert Kleinberg, and S Matthew Weinberg. Prophet inequalities with limited information. In *Proceedings of the twenty-fifth annual ACM-SIAM symposium on Discrete algorithms*, pages 1358–1377. SIAM, 2014.

Jackie Baek and Will Ma. Bifurcating constraints to improve approximation ratios for network revenue management with reusable resources. *Operations Research*, 70(4):2226–2236, 2022.

Jackie Baek and Shixin Wang. Leveraging reusability: Improved competitive ratio of greedy for reusable resources. *arXiv preprint arXiv:2304.03377*, 2023.

Michael O Ball and Maurice Queyranne. Toward robust revenue management: Competitive analysis of online booking. *Operations Research*, 57(4):950–963, 2009.

Julius Borcea, Petter Brändén, and Thomas M. Liggett. Negative dependence and the geometry of polynomials. *Journal of the American Mathematical Society*, 22(2):521–567, 2009.

Petter Brändén and Rafael S González D'León. On the half-plane property and the tutte group of a matroid. *Journal of Combinatorial Theory, Series B*, 100(5):485–492, 2010.

Mark Braverman, Mahsa Derakhshan, Tristan Pollner, Amin Saberi, and David Wajc. New philosopher inequalities for online bayesian matching, via pivotal sampling. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3029–3068. SIAM, 2025.

Simon Bruggmann and Rico Zenklusen. An optimal monotone contention resolution scheme for bipartite matchings via a polyhedral viewpoint. *Mathematical Programming*, 191(2):795–845, 2022.

Niv Buchbinder, Kamal Jain, and Mohit Singh. Secretary problems and incentives via linear programming. *ACM SIGecom Exchanges*, 8(2):1–5, 2009.

Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 311–320, 2010.

Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. *SIAM Journal on Computing*, 43(6):1831–1879, 2014.

Xiang-Hui Chen, Arthur P Dempster, and Jun S Liu. Weighted finite population sampling to maximize entropy. *Biometrika*, 81(3): 457–469, 1994.

Yiwei Chen, Retsef Levi, and Cong Shi. Revenue management of reusable resources with advanced reservations. *Production and Operations Management*, 26(5):836–859, 2017.

Youngbin Choe and David G Wagner. Rayleigh matroids. *Combinatorics, Probability and Computing*, 15(5):765–781, 2006.

José Correa, Patricio Foncea, Ruben Hoeksma, Tim Oosterwijk, and Tjark Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 169–186, 2017.

Andrés Cristi and Bruno Ziliotto. Prophet inequalities require only a constant number of samples. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 491–502, 2024.

Steven Delong, Alireza Farhadi, Rad Niazadeh, Balasubramanian Sivan, and Rajan Udwani. Online bipartite matching with reusable resources. *Mathematics of Operations Research*, 49(3), 2024.

John Dickerson, Karthik Sankararaman, Aravind Srinivasan, and Pan Xu. Allocation problems in ride-sharing platforms: Online matching with offline reusable resources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Atanas Dinev and S Matthew Weinberg. Simple and optimal online contention resolution schemes for k-uniform matroids. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*, pages 39–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024.

Zehao Dong, Sanmay Das, Patrick Fowler, and Chien-Ju Ho. Efficient nonmyopic online allocation of scarce reusable resources. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems*, 2021.

Shaddin Dughmi. From contention resolution to matroid secretary and back. *SIAM Journal on Computing*, 54(3):585–624, 2025.

Paul Dutting, Michal Feldman, Thomas Kesselheim, and Brendan Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM Journal on Computing*, 49(3):540–582, 2020.

Bradley Efron. Increasing properties of polya frequency function. *The Annals of Mathematical Statistics*, pages 272–279, 1965.

Soheil Ehsani, MohammadTaghi Hajiaghayi, Thomas Kesselheim, and Sahil Singla. Prophet secretary for combinatorial auctions and matroids. In *Proceedings of the twenty-ninth annual acm-siam symposium on discrete algorithms*, pages 700–714. SIAM, 2018.

Farbod Ekbatani, Rad Niazadeh, Pranav Nuti, and Jan Vondrák. Prophet inequalities with cancellation costs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1247–1258, 2024.

Farbod Ekbatani, Yiding Feng, Ian Kash, and Rad Niazadeh. Online job assignment. *arXiv preprint arXiv:2506.06893*, 2025.

Hossein Esfandiari, MohammadTaghi Hajiaghayi, Vahid Liaghat, and Morteza Monemizadeh. Prophet secretary. *SIAM Journal on Discrete Mathematics*, 31(3):1685–1701, 2017.

Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Prophet inequality for vertex and edge arrival models. *Math. Oper. Res*, 47(2):878–898, 2022.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 1014–1033. SIAM, 2016.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Online contention resolution schemes with applications to bayesian selection problems. *SIAM Journal on Computing*, 50(2):255–300, 2021.

Moran Feldman, Ola Svensson, and Rico Zenklusen. Nearly tight sample complexity for matroid online contention resolution. In *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 4692–4711. SIAM, 2026.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Near-optimal bayesian online assortment of reusable resources. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 964–965, 2022.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Near-optimal bayesian online assortment of reusable resources. *Operations Research*, 72(5), 2024.

Yiding Feng, Rad Niazadeh, and Amin Saberi. Robustness of online inventory balancing to inventory shocks. *arXiv preprint arXiv:2511.16044*, 2025.

Hu Fu, Zhihao Gavin Tang, Hongxun Wu, Jinzhao Wu, and Qianfan Zhang. Random order vertex arrival contention resolution schemes for matching, with applications. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, pages 68–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021.

Hu Fu, Pinyan Lu, Zhihao Gavin Tang, Hongxun Wu, Jinzhao Wu, and Qianfan Zhang. Sample-based matroid prophet inequalities. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 781–781, 2024.

Guillermo Gallego, Anran Li, Van-Anh Truong, and Xinshang Wang. Online personalized resource allocation with customer choice, 2016. Working paper.

Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9. Walter de Gruyter, 2011.

Shayan Oveis Gharan, Amin Saberi, and Mohit Singh. A randomized rounding approach to the traveling salesman problem. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 550–559. IEEE, 2011.

Rohan Ghuge, Sahil Singla, and Yifan Wang. Single-sample and robust online resource allocation. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pages 1442–1453, 2025.

Christina Goldschmidt, James B Martin, and Dario Spano. Fragmenting random permutations. *Electronic Communications in Probability [electronic only]*, 13:461–474, 2008.

Xiao-Yue Gong, Vineet Goyal, Garud N Iyengar, David Simchi-Levi, Rajan Udwani, and Shuangyu Wang. Online assortment optimization with reusable resources. *Management Science*, 68(7):4772–4785, 2022.

Parikshit Gopalan, Adam Klivans, Raghu Meka, Daniel Štefankovic, Santosh Vempala, and Eric Vigoda. An fptas for# knapsack and related counting problems. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science*, pages 817–826. IEEE, 2011.

Vineet Goyal, Garud Iyengar, and Rajan Udwani. Asymptotically optimal competitive ratio for online allocation of reusable resources. *Operations Research*, 73(4):1897–1915, 2025.

Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1988.

Guru Guruganesh and Euiwoong Lee. Understanding the correlation gap for matchings. In *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, 2018.

Mohammad Taghi Hajiaghayi, Robert Kleinberg, and David C Parkes. Adaptive limited-supply online auctions. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 71–80, 2004.

Jason D Hartline and Tim Roughgarden. Simple versus optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 225–234, 2009.

Ole J Heilmann and Elliott H Lieb. Theory of monomer-dimer systems. *Communications in mathematical Physics*, 25(3):190–232, 1972.

Yukai Huang, Jacob Feldman, and Heng Zhang. Joint inventory and assortment optimization with reusable resources. *Available at SSRN 5015696*, 2024.

Tianming Huo and Wang Chi Cheung. Online reusable resource assortment planning with customer-dependent usage durations. *Available at SSRN 4504713*, 2023a.

Tianming Huo and Wang Chi Cheung. Online reusable resource allocations with multi-class arrivals. *Available at SSRN 4320423*, 2023b.

Edwin T Jaynes. Information theory and statistical mechanics. *Physical review*, 106(4):620, 1957.

Mark Jerrum and Alistair Sinclair. The markov chain monte carlo method: an approach to approximate counting and integration. In Dorit S. Hochbaum, editor, *Approximation Algorithms for NP-hard Problems*. PWS Publishing, 1996.

Mark R. Jerrum, Leslie G. Valiant, and Vijay V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986. doi: 10.1016/0304-3975(86)90174-X.

Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tight guarantees for multiunit prophet inequalities and online stochastic knapsack. *Operations Research*, 73(3):1703–1721, 2025a.

Jiashuo Jiang, Will Ma, and Jiawei Zhang. Tightness without counterexamples: A new approach and new results for prophet inequalities. *Mathematics of Operations Research*, 2025b.

Jeff Kahn and Michael Neiman. Negative correlation and log-concavity. *Random Structures & Algorithms*, 37(3):367–388, 2010.

Danish Kashaev and Richard Santiago. A simple optimal contention resolution scheme for uniform matroids. *Theoretical Computer Science*, 940:81–96, 2023.

Kristen Kessel, Ali Shameli, Amin Saberi, and David Wajc. The stationary prophet inequality problem. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 243–244, 2022.

Thomas Kesselheim, Robert Kleinberg, and Rad Niazadeh. Secretary problems with non-uniform arrival order. In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, pages 879–888, 2015.

Robert Kleinberg and Seth Matthew Weinberg. Matroid prophet inequalities. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 123–136, 2012.

Ulrich Krengel and Louis Sucheston. On semiamarts, amarts, and processes with finite value. *Probability on Banach spaces*, 4(197-266):1–2, 1978.

Euiwoong Lee and Sahil Singla. Optimal online contention resolution schemes via ex-ante prophet inequalities. In *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112, page 57. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2018.

Retsef Levi and Ana Radovanović. Provably near-optimal LP-based policies for revenue management in systems with reusable resources. *Operations Research*, 58(2):503–507, 2010.

Thomas M Liggett. Ultra logconcave sequences and negative dependence. *Journal of Combinatorial Theory, Series A*, 79(2):315–325, 1997.

Qingsong Liu and Mohammad Hajiesmaili. Online fair allocation of reusable resources. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 9(2), 2025.

Brendan Lucier. An economic view of prophet inequalities. *ACM SIGecom Exchanges*, 16(1):24–47, 2017.

Russell Lyons. Determinantal probability measures. *Publications Mathématiques de l'IHÉS*, 98:167–212, 2003. doi: 10.1007/s10240-003-0016-0.

Will Ma, Calum MacRury, and Pranav Nuti. Online matching and contention resolution for edge arrivals with vanishing probabilities. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 159–159, 2024.

Will Ma, Calum MacRury, and Jingwei Zhang. Online contention resolution schemes for network revenue management and combinatorial auctions. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, pages 100–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2026.

Yuhang Ma, Paat Rusmevichientong, Mika Sumida, and Huseyin Topaloglu. An approximation algorithm for network revenue management under nonstationary arrivals. *Operations Research*, 68(3):834–855, 2020.

Calum MacRury and Will Ma. Random-order contention resolution via continuous induction: Tightness for bipartite matching under vertex arrivals. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 1629–1640, 2024.

Calum MacRury, Will Ma, and Nathaniel Grammel. On (random-order) online contention resolution schemes for the matching polytope of (bipartite) graphs. *Operations Research*, 73(2):689–703, 2025.

Vahideh Manshadi and Scott Rodilitz. Online policies for efficient volunteer crowdsourcing. *Management Science*, 68(9):6572–6590, 2022.

Michael Molloy and Bruce Reed. Hard-core distributions on matchings. In *Graph Colouring and the Probabilistic Method*, pages 247–264. Springer, 2002.

Hossein Nekouyan, Bo Sun, Raouf Boutaba, and Xiaoqi Tan. Online rounding schemes for k-rental problems. *arXiv preprint arXiv:2507.19649*, 2025.

Rad Niazadeh, Amin Saberi, and Ali Shameli. Prophet inequalities vs. approximating optimum online. In *International Conference on Web and Internet Economics*, pages 356–374. Springer, 2018.

James R Norris. *Markov chains*. Number 2. Cambridge university press, 1998.

Pranav Nuti and Jan Vondrák. Towards an optimal contention resolution scheme for matchings: Towards an optimal contention resolution scheme... *Mathematical Programming*, 210(1):761–792, 2025.

Zachary Owen and David Simchi-Levi. Price and assortment optimization for reusable resources. *Available at SSRN 3070625*, 2018.

Christos Papadimitriou, Tristan Pollner, Amin Saberi, and David Wajc. Online stochastic max-weight bipartite matching: Beyond prophet inequalities. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 763–764, 2021.

Robin Pemantle. Towards a theory of negative dependence. *Journal of Mathematical Physics*, 41(3):1371–1390, 2000.

Tristan Pollner, Mohammad Roghani, Amin Saberi, and David Wajc. Improved online contention resolution for matchings and applications to the gig economy. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 321–322, 2022.

Aviad 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, 2016.

Aviad Rubinstein and Sahil Singla. Secretary, prophet, and stochastic probing via big-decisions-first. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, pages 1242–1253, 2026.

Aviad Rubinstein, Jack Z Wang, and S Matthew Weinberg. Optimal single-choice prophet inequalities from samples. *Innovations in Theoretical Computer Science*, 2020.

Paat Rusmevichientong, Mika Sumida, and Huseyin Topaloglu. Dynamic assortment optimization for reusable products with random usage durations. *Management Science*, 66(7):2820–2844, 2020.

Paat Rusmevichientong, Mika Sumida, Huseyin Topaloglu, and Yicheng Bai. Revenue management with heterogeneous resources: Unit resource capacities, advance bookings, and itineraries over time intervals. *Operations Research*, 71(6):2196–2216, 2023.

Ester Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *the Annals of Probability*, pages 1213–1216, 1984.

Alexander Schrijver et al. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

Meghan Shanks, Ge Yu, and Sheldon H Jacobson. Approximation algorithms for stochastic online matching with reusable resources. *Mathematical Methods of Operations Research*, 98(1):43–56, 2023.

David Simchi-Levi, Zeyu Zheng, and Feng Zhu. On greedy-like policies in online matching with reusable network resources and decaying rewards. *Management Science*, 71(10):8908–8926, 2025.

Mohit Singh and Nisheeth K Vishnoi. Entropy, optimization and counting. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 50–59, 2014.

Damian Straszak and Nisheeth K Vishnoi. Maximum entropy distributions: Bit complexity and stability. In *Conference on Learning Theory*, pages 2861–2891. PMLR, 2019.

Mika Sumida. Dynamic resource allocation with recovering rewards under non-stationary arrivals. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pages 159–159, 2025.

Hanna Sumita, Shinji Ito, Kei Takemura, Daisuke Hatano, Takuro Fukunaga, Naonori Kakimura, and Ken-ichi Kawarabayashi. Online task assignment problems with reusable resources. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 2022.

David G Wagner. Negatively correlated random variables and mason's conjecture for independent sets in matroids. *Annals of Combinatorics*, 12(2):211–239, 2008.

Xinshang Wang, Van-Anh Truong, and David Bank. Online advance admission scheduling for services with customer preferences. *arXiv preprint arXiv:1805.10412*, 2018.

Zerui Wu, Ran Liu, and Xu Sun. Blind and near-optimal dynamic assortment optimization with reusable resources. *Available at SSRN 5340627*, 2025.

Xinchang Xie, Itai Gurvich, and Simge Küçükyavuz. Dynamic allocation of reusable resources: Logarithmic regret in overloaded networks. *Operations Research*, 73(4):2097–2124, 2025.

Qiqi Yan. Mechanism design via correlation gap. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 710–719. SIAM, 2011.

Toru Yoshinaga and Yasushi Kawase. Size-stochastic knapsack online contention resolution schemes. *arXiv preprint arXiv:2305.08622*, 2023.

Xilin Zhang and Wang Chi Cheung. Online allocation of reusable resources in nonstationary environments. *Mathematics of Operations Research*, 2025.