The Optimal $1/e$ Guarantee for Rank-Two Matroid Secretary

Shiva Kintali*

September 15, 2026

Abstract

We prove that a sequential record algorithm obtains expected weight at least OPT/e on every matroid of rank at most two, with arbitrary nonnegative weights and a uniformly random arrival order. The algorithm knows only the total number of elements and has independence-oracle access to subsets of already arrived elements. It uses only comparisons between weights.

1 Introduction and main result

The classical secretary problem asks how much of an offline optimum can be recovered when observations arrive in random order and every decision is irrevocable. Sampling an initial segment and then accepting the next record gives the familiar asymptotic success probability $1/e$; see Ferguson [1]. The matroid secretary problem replaces the single available position by a matroid independence constraint and asks for a maximum-weight independent set online. It was introduced by Babaioff, Immorlica, and Kleinberg [2]; see also the journal treatment of Babaioff, Immorlica, Kempe, and Kleinberg [3].

Even at rank two, independence matters. After loops are deleted, a rank-two matroid consists of parallel classes, and an independent pair must use two different classes. Several of the heaviest elements may therefore compete for the same position. We prove that the classical fraction $1/e$ is nevertheless attainable in expected weight without knowing the parallel classes or any representation of the matroid in advance.

For a matroid $M = (E, \mathcal{I})$ and a nonnegative weight function $w : E \to \mathbb{R}_{\geq 0}$, write

$$\mathrm{OPT}(M, w) = \max_{I \in \mathcal{I}} \sum_{x \in I} w(x).$$

Theorem 1.1 (Main guarantee)**.** *Suppose that the elements of a matroid $M = (E, \mathcal{I})$ of rank at most two arrive in a uniformly random order and reveal fixed nonnegative weights on arrival. There is a randomized online algorithm that always returns an independent set $\mathcal{A}$ and satisfies*

$$\mathbb{E}\left[\sum_{x \in \mathcal{A}} w(x)\right] \geq \frac{1}{e}\,\mathrm{OPT}(M, w).$$

The expectation is over the arrival order and the algorithm's internal randomness. The algorithm is told only $N = |E|$; it uses comparisons between revealed weights and independence queries involving already arrived elements.

*Email: shiva.kintali@gmail.com Homepage: shivakintali.github.io

The theorem permits loops, zero weights, and ties. Ties are resolved by an arrival-independent tie-breaking order, as specified in Section 2. The conclusion is an expected-weight, or utility, guarantee. It does not assert that every element of a designated optimal basis is selected with probability at least $1/e$. Example 6.1 shows that the particular algorithm analyzed here can violate this stronger individual-inclusion statement even when the optimum is unique.

1.1 Related work and comparison

We use the convention that a c-competitive utility guarantee means expected weight at least OPT/c. For general rank r, the first matroid-secretary guarantee was $O(\log r)$-competitive [2]. It was improved to $O(\sqrt{\log r})$ [4] and then to $O(\log\log r)$ [5]; Feldman, Svensson, and Zenklusen [6] later gave a simpler algorithm with the same asymptotic ratio. Singla [7] obtains a $1/4$ inclusion probability for every element of the canonical greedy optimum on an arbitrary matroid. The algorithm knows only the total input size and has independence-oracle access to arrived elements. By linearity of expectation, it also gives an expected-weight fraction $1/4$. Our result increases that expected-weight fraction to $1/e$ for rank at most two in the same arrived-element oracle model. It neither gives an arbitrary-rank guarantee nor strengthens Singla's individual-inclusion conclusion.

Several special matroid classes admit stronger guarantees. Kleinberg's multiple-choice secretary algorithm gives an expected fraction $1 - O(r^{-1/2})$ for uniform matroids [8]. Random-order weighted bipartite matching yields an e-competitive utility algorithm for transversal matroids [9]. Soto, Turkieltaub, and Verdugo [10] prove the stronger statement that every element of the canonical optimum is selected with probability at least $1/e$. These algorithms require a bipartite presentation, revealed either in advance or through the neighbors of each arriving element.

This does not cover every rank-two matroid. Consider the rank-two truncation of the partition matroid with capacity one on each of three two-element parts. In any transversal presentation, the union of the neighborhoods of two parallel nonloops must have size one; otherwise Hall's condition would give a matching of size two. Thus the two members of each parallel class have the same singleton neighbor. Different classes must have different neighbors because every cross-class pair is independent. Choosing one element from each of the three classes would then give a matching of size three, contradicting rank two. Hence this matroid is not transversal.

After loops are deleted, every rank-two matroid is laminar: impose capacity one on each parallel class and capacity two on their union. Constant guarantees for laminar matroids were developed through several methods [11, 12, 13, 10]. Huang, Parsaeian, and Zhu [14] proved that a greedy rule selects every optimal element with probability at least $1/4.75$. Bérczi, Livanos, Soto, and Verdugo [15] proved that the greedy-improving algorithm has the tight worst-case guarantee $1 - \ln 2$ on the full laminar class. Their separate rank-two algorithm selects every optimal element with probability at least 0.3462, and therefore obtains the same fraction in expected weight.

Against this last guarantee, Theorem 1.1 increases the expected-weight fraction to $1/e \approx 0.367879$ while using only the arrived-element oracle. The comparison is between the stated guarantees: the earlier result supplies an element-by-element guarantee, whereas ours supplies the larger expected-weight fraction. The counterexample in Section 6 confirms that our algorithm does not attain a $1/e$ individual-inclusion guarantee.

1.2 Contributions and proof strategy

Fix an arrival-independent strict priority order refining nonincreasing weight. The algorithm draws $K \sim \mathrm{Bin}(N, 1/e)$ and rejects the first K ground-set arrivals, including loops. It then accepts the first later nonloop whose priority is better than that of every previously observed nonloop; call it

x. Finally, it accepts the first later nonloop outside the parallel class of x whose priority is better than that of every earlier observation outside this class. Both record tests include the sampled observations, and the second also includes observations made before x was accepted. Section 2 gives the precise online implementation and proves feasibility using $O(N)$ independence queries and $O(N)$ storage.

The proof has two principal ingredients. First, let P_j contain the first j nonloops in priority order. The priority-prefix identity reduces the expected-weight theorem to

$$\mathbb{E}|\mathcal{A} \cap P_j| \geq p\, r(P_j) \qquad \text{for every } j, \qquad p = 1/e.$$

Rank-one prefixes follow from the finite classical secretary calculation. For rank two, it suffices to analyze the first prefix of rank two: it contains m elements of one parallel class and the highest-priority element of a second class. Every later rank-two prefix contains this one, so its selected count can only be larger.

Second, represent arrivals by independent uniform times and use the cutoff p. This is distributionally identical to the binomial sample used by the online algorithm. Conditioning on the highest-priority sampled nonloop isolates at most one exceptional class in the analysis of the second selection, namely the sampled element's parallel class. Averaging the resulting conditional bounds gives a master lower bound $L_{m,n}$, where n is the number of nonloops. We prove, entirely analytically, that

$$L_{m,n} \geq L_{m,\infty} \geq L_{1,\infty} \geq \frac{p}{60}(55 + 228p - 129p^2 - 28p^3) > 2p, \qquad p = 1/e.$$

The first comparison charges the omitted finite tail to the empty-sample event. The second is a monotonicity statement for the lower-bound expressions as the target prefix grows; it is not a monotonicity claim about the algorithm over different instances. The final comparison uses exact values for two initial terms and a uniform conditional envelope. The detailed envelope proof is in Appendix A. No numerical minimization, executable certificate, or external supplement is required.

The technical contributions are the rank-two conditional decomposition, the master expression, and its uniform analytical lower bound. The initial sampling device and the priority-prefix identity are established tools. Section 2 gives the model and algorithm; Section 3 proves the prefix reduction; Section 4 derives the master expression; Section 5 proves its uniform bound; and Section 6 completes the theorem and gives the individual-inclusion counterexample.

2 Online model and sequential record algorithm

Definition 2.1 (Matroid secretary model)**.** A matroid $M = (E, \mathcal{I})$ consists of a finite ground set E and a nonempty family $\mathcal{I} \subseteq 2^E$ satisfying heredity and augmentation. Its members are the independent sets. Heredity means that $A \subseteq B \in \mathcal{I}$ implies $A \in \mathcal{I}$. Augmentation means that, for $A, B \in \mathcal{I}$ with $|A| < |B|$, some $b \in B \setminus A$ satisfies $A \cup \{b\} \in \mathcal{I}$. Its rank function is

$$r(X) = \max\{|I| : I \subseteq X,\ I \in \mathcal{I}\}.$$

Each element x has a fixed nonnegative weight $w(x)$, and $w(X) := \sum_{x \in X} w(x)$ for $X \subseteq E$. The elements arrive in a uniformly random permutation, independent of the algorithm's randomness. When an element arrives, its identity and weight are revealed, and the algorithm must accept or reject it irrevocably before the next arrival. The accepted elements must form an independent set at every time. Write $\mathcal{A}$ for the random set of accepted elements. The benchmark is

$$\mathrm{OPT}(M, w) = \max_{I \in \mathcal{I}} \sum_{x \in I} w(x).$$

The algorithm is told $N = |E|$. It receives no advance representation of the matroid, its rank, its number of nonloops, or its parallel classes. It has an independence oracle: for any set consisting entirely of arrived elements, the oracle reports whether that set belongs to $\mathcal{I}$. Queries involving an unarrived element are forbidden. The input is promised to have rank at most two, but the algorithm need not know its exact rank.

A loop is an element x for which $\{x\} \notin \mathcal{I}$; every other element is a nonloop. Loops are always rejected and excluded from priority comparisons, but they still count toward the N arrivals and toward the initial sample length. Two distinct nonloops are parallel if their pair is dependent.

Fix a strict priority order that refines nonincreasing weight and is independent of the arrival permutation. Concretely, assign independent continuous tie-breaking labels to the elements and order equal-weight elements by these labels. A label may be generated and revealed when its element arrives; this produces the same joint distribution as assigning all labels in advance. In the analysis we condition on the resulting strict order. Because the labels are independent of the arrival permutation, this conditioning preserves the uniform arrival law. Index the nonloops in this order, with smaller indices denoting better priority. Thus ties in weight, including ties at zero, cause no ambiguity.

2.1 Rank-two structure

Lemma 2.2 (Rank-two structure)**.** *On the nonloops of a matroid of rank at most two, equality or parallelism is an equivalence relation. Its equivalence classes are called parallel classes. A set of nonloops is independent if and only if it has at most two elements and contains at most one element from each parallel class.*

Proof. Reflexivity follows from adjoining equality to parallelism, and symmetry follows because unordered pairs are symmetric. For transitivity, take three distinct nonloops x, y, z such that $\{x, y\}$ and $\{y, z\}$ are dependent. If $\{x, z\}$ were independent, then augmentation applied to the independent sets $\{y\}$ and $\{x, z\}$ would give an element $b \in \{x, z\}$ for which $\{y, b\}$ is independent. This contradicts the dependence of both $\{x, y\}$ and $\{y, z\}$. Hence x and z are parallel.

The empty set is independent by nonemptiness and heredity. Every singleton nonloop is independent, and a pair of distinct nonloops is independent exactly when its members are in different parallel classes. Every set of more than two elements is dependent because the rank is at most two. These observations prove both directions of the final characterization. □

Set throughout

$$p = e^{-1}, \qquad q = 1 - p.$$

Relative to any specified collection of observations, a newly arriving element is a *record* if its priority is better than that of every earlier element in the collection. If the collection has no earlier element, the new arrival is a record.

2.2 Sequential record algorithm

Loops count toward the sample length but never enter priority comparisons or selections. Until the first acceptance, retain the identity and priority of every observed nonloop, including nonloops rejected after the sample.

1. Draw $K \sim \mathrm{Bin}(N, p)$ independently of the arrival order and the tie-breaking labels. Reject the first K ground-set arrivals; loops are included in this count.

2. After the sample, accept the first nonloop that is a record among *all* nonloops seen so far. Denote it by x, and its parallel class by $C(x)$. If no such element arrives, the algorithm returns the empty set.

3. After accepting x, accept the first nonloop outside $C(x)$ that is a record relative to *all earlier nonloop observations outside* $C(x)$. These observations include sampled elements and elements rejected before or after x. After this second acceptance, make no further selections. If no qualifying element arrives, return $\{x\}$.

Lemma 2.3 (Implementation and feasibility)**.** *The algorithm is feasible and implementable with the information in Theorem 1.1.*

Proof. A singleton query determines on arrival whether an element is a loop. A loop is rejected immediately. During the sample, store every observed nonloop and maintain the best priority seen so far. After the sample and before the first acceptance, continue storing every rejected nonloop. A new nonloop is a global record exactly when its priority is better than the stored best priority; this condition holds automatically if no nonloop has yet appeared. Thus Step 2 is implementable using only revealed information.

Suppose that Step 2 accepts x. For every previously observed nonloop $y \neq x$, query $\{x, y\}$. By Lemma 2.2, the answer is independent exactly when $y \notin C(x)$. Among all such elements, store the one with best priority, or an empty sentinel if there is none. Now consider a later arrival y. After the singleton query has established that y is a nonloop, query $\{x, y\}$. If the pair is dependent, then $y \in C(x)$ and is irrelevant to the second record process. If the pair is independent, then $y \notin C(x)$; it is the record required by Step 3 exactly when the sentinel is empty or the priority of y is better than that of the stored outside-class observation. In that case accept y and stop. Otherwise reject it; the stored best outside-class observation remains unchanged.

All oracle queries contain only arrived elements. The first accepted element is a nonloop, and a second element is accepted only after the oracle has certified its pair with x to be independent. Hence the output is always independent. Rank zero produces no first acceptance, and in rank one every pair of nonloops is dependent, so no second acceptance occurs.

There is at most one singleton query per ground-set element. Apart from x, each nonloop is involved in at most one pair query with x, either when x is accepted or upon that nonloop's later arrival. Thus the total number of independence queries is at most $2N$. The algorithm stores at most all previously observed nonloops, so its storage is $O(N)$. The draw of K and the tie-breaking labels are internal randomness. The continuous arrival times introduced in Section 3.1 are used only for the analysis. $\square$

3 Reduction to priority-prefix counts

3.1 Arrival-time representation and record estimates

Lemma 3.1 (Arrival-time coupling)**.** *On the same probability space as the independent tie-breaking labels, assign independent $U_y \sim \mathrm{Unif}[0, 1]$ to the elements $y \in E$. Process the elements in increasing order of U_y and reject every element with $U_y \leq p$. The resulting output has exactly the same distribution as the output of the online algorithm. After the fixed set of loops is removed, the remaining arrival times are still independent uniform random variables. Write n for the number of these nonloops in the analysis.*

Proof. Let π be the permutation of element identities obtained by sorting the arrival times, and let

$$K' = |\{y \in E : U_y \leq p\}|.$$

Because the times are independent, identically distributed, and continuous, π is uniform and is independent of the ordered time values. In particular, $K' \sim \mathrm{Bin}(N, p)$ and K' is independent of π. Both are also independent of the tie-breaking labels. Moreover, with probability one, the elements with time at most p are exactly the first K' elements of π. Thus the joint distribution of the arrival permutation, sample length, and priority order in this construction agrees with that of the online algorithm.

Once these three objects are fixed, the two procedures make the same decisions: after the sample, the algorithm uses only arrival order, priority comparisons, and independence-oracle answers, all of which agree in the two descriptions. Hence their outputs agree in distribution. Finally, the set of loops is determined by the fixed matroid. Deleting those coordinates from a family of independent uniform variables leaves the remaining coordinates independent and uniform. The analysis may therefore use the fixed number n of nonloops without supplying n to the algorithm. $\square$

Lemma 3.2 (No record after a cutoff)**.** *Consider $L \geq 0$ fixed, strictly ranked elements with independent uniform arrival times in $[0, 1]$. Call an element a record if, when it arrives, no higher-ranked element among these L elements has arrived earlier. For $0 < t \leq u \leq 1$, the probability that no record has arrival time in (t, u) is*

$$S_L(t, u) = \frac{t}{u} + \left(1 - \frac{t}{u}\right)(1 - u)^L \ \geq \ \frac{t}{u}. \tag{1}$$

We use $(1 - u)^0 = 1$, including at $u = 1$. Arrival at a specified endpoint has probability zero and does not affect the formula.

Proof. Let A be the set of elements whose arrival times are at most u. With probability $(1 - u)^L$, the set A is empty, in which case the event holds. Condition instead on any particular nonempty value of A. Conditional on this event, the arrival times of the elements in A are independent and uniform on $[0, u]$. Let y_A be the highest-ranked element of A. No record occurs in (t, u) if and only if $U_{y_A} \leq t$. Indeed, if $U_{y_A} \leq t$, then y_A precedes and outranks every arrival in (t, u); if $U_{y_A} > t$, then y_A itself is a record in that interval. Consequently, for every nonempty value of A, the conditional probability is t/u. Combining the empty and nonempty cases gives $(1 - u)^L + (1 - (1 - u)^L)t/u$, which is (1). The argument also covers $L = 0$, when this expression equals one. $\square$

Lemma 3.3 (Finite classical secretary bound)**.** *For $d \geq 1$ fixed ranked elements with independent uniform arrival times, consider the rule that rejects arrivals at times at most $t \in (0, 1)$ and accepts the first subsequent record. Its probability of selecting the highest-ranked element is*

$$\int_t^1 \left[\frac{t}{u} + \left(1 - \frac{t}{u}\right)(1 - u)^{d-1}\right] du \ \geq \ -t \log t. \tag{2}$$

Proof. Condition on the highest-ranked element's arrival time u. If $u \leq t$, the rule rejects it. If $u > t$, the rule selects it if and only if the other $d - 1$ elements produce no record in (t, u): any such record would be accepted first, whereas in the absence of one the highest-ranked element is the first record after the cutoff. Its arrival time is independent of all other times and has density one. Lemma 3.2 therefore gives the displayed integral. Dropping its nonnegative second term gives $\int_t^1 (t/u)\, du = -t \log t$. $\square$

3.2 Prefix criterion and the canonical rank-two prefix

No independent set contains a loop, and the algorithm never accepts one. Deleting the loops therefore changes neither the algorithm's weight nor $\mathrm{OPT}(M, w)$. If $n = 0$, both quantities are zero and the theorem is immediate. For the remainder of the proof assume $n \geq 1$.

Fix a realization of the tie-breaking labels and hence of the strict priority order. Until the final averaging step, all probabilities and expectations are conditional on this order. The argument is uniform over the realized order, so averaging over the labels will recover the unconditional guarantee in Theorem 1.1. Index the nonloops $1, \dots, n$ from highest to lowest priority, and write w_i for their weights. Define $P_j = \{1, \dots, j\}$, $P_0 = \varnothing$, and $w_{n+1} = 0$. We use the standard prefix criterion for ordinal matroid secretary algorithms (see Soto, Turkieltaub, and Verdugo [10]) and include its proof.

Lemma 3.4 (Prefix criterion)**.** *For every random subset Z of nonloops,*

$$\mathbb{E}w(Z) = \sum_{j=1}^{n} (w_j - w_{j+1}) \mathbb{E}|Z \cap P_j|. \tag{3}$$

Moreover,

$$\mathrm{OPT}(M, w) = \sum_{j=1}^{n} (w_j - w_{j+1}) r(P_j). \tag{4}$$

Consequently, $\mathbb{E}|\mathcal{A} \cap P_j| \geq p\, r(P_j)$ for every j implies the conditional expected-weight bound $\mathbb{E}w(\mathcal{A}) \geq p\, \mathrm{OPT}(M, w)$. If these prefix inequalities hold for every realized priority order, then Theorem 1.1 follows by averaging over the tie-breaking labels.

Proof. For each i, telescoping gives $w_i = \sum_{j=i}^{n} (w_j - w_{j+1})$. Sum this identity over $i \in Z$, interchange the finite sums, and take expectations to obtain (3).

Run the offline matroid greedy algorithm in priority order, and let G_j be its accepted set after processing P_j. The set G_j is maximal independent in P_j. To see this, suppose that $e \in P_j \setminus G_j$ was rejected when the current accepted set was S. Then $S \cup \{e\}$ was dependent and $S \subseteq G_j$. If $G_j \cup \{e\}$ were independent, heredity would imply that $S \cup \{e\}$ was independent, a contradiction. Thus no rejected element can augment G_j. In a matroid every maximal independent subset of P_j has size $r(P_j)$, so $|G_j| = r(P_j)$.

Let $G = G_n$. Applying (3) deterministically to G gives

$$w(G) = \sum_{j=1}^{n} (w_j - w_{j+1}) r(P_j).$$

For every independent set I, the inequalities $|I \cap P_j| \leq r(P_j)$ and $w_j - w_{j+1} \geq 0$ show, by the same layer identity, that $w(I) \leq w(G)$. Hence G is maximum-weight and the last display is (4). Finally, applying (3) to $Z = \mathcal{A}$ and using the assumed prefix inequalities proves the conditional expected-weight bound. Its average over the tie-breaking labels is the assertion of Theorem 1.1. □

Apply Lemma 3.3 with $d = n$ and $t = p$ to the algorithm's first-selection rule. Element 1 is the first selection with probability at least $-p \log p = p$, where the last equality uses $p = e^{-1}$. Whenever $r(P_j) = 1$, the prefix P_j contains element 1, and therefore

$$\mathbb{E}|\mathcal{A} \cap P_j| \geq \mathbb{P}(1 \in \mathcal{A}) \geq p = p\, r(P_j).$$

Thus the required inequality holds for every rank-one prefix, regardless of the rank of the ambient matroid. Together with the case $n = 0$, this proves the theorem for matroids of rank at most one.

Now suppose that the matroid has rank two. Let $m + 1$ be the first index for which $r(P_{m+1}) = 2$. Then $m \geq 1$. By Lemma 2.2, the first m elements lie in a single parallel class C_A, while element $m + 1$ lies in a different class C_B. Denote the canonical rank-two prefix by

$$P = P_{m+1}, \qquad n \geq m + 1.$$

The rank-one prefix inequalities have already been proved. For the remaining prefixes it suffices to prove the stronger inequality

$$\mathbb{E}|\mathcal{A} \cap P| > \frac{2}{e}. \tag{5}$$

Indeed, every prefix P_j of rank two contains P, and hence, for every outcome,

$$|\mathcal{A} \cap P_j| \geq |\mathcal{A} \cap P|.$$

Since the ambient matroid has rank two, no prefix has larger rank. Thus (5) implies $\mathbb{E}|\mathcal{A} \cap P_j| > 2/e = p\, r(P_j)$ for every rank-two prefix and completes the reduction.

4 Conditional analysis and the master expression

Throughout this section the strict priority order fixed in Section 3 remains fixed; probabilities and expectations are over the independent arrival times only. We condition on the best sampled nonloop, analyze the two selections under that conditioning, and then sum the resulting bounds.

Let J be the priority index of the highest-priority sampled nonloop. Set $J = \infty$ when no nonloop is sampled. For $0 \leq s \leq n-1$,

$$\mathbb{P}(J = s+1) = pq^s, \qquad \mathbb{P}(J = \infty) = q^n. \tag{6}$$

Indeed, the event $J = s+1$ is

$$\{U_1 > p, \ldots, U_s > p,\ U_{s+1} \leq p\},$$

with no restriction on the coordinates of index greater than $s+1$. Independence gives probability $q^s p$. The event $J = \infty$ is $\{U_1 > p, \ldots, U_n > p\}$ and has probability q^n.

Lemma 4.1 (Conditional first selection)**.** *Condition on $J = s+1$ with $1 \leq s \leq n-1$. The elements $1, \ldots, s$, called* eligible elements *in this conditional analysis, have independent uniform times in $(p, 1)$. Every element with priority index greater than $s+1$ retains its independent uniform time in $(0, 1)$. The first selection is the earliest of the s eligible elements. Its identity I is uniform on $\{1, \ldots, s\}$ and independent of its time T_s, whose density is*

$$f_s(t) = \frac{s(1-t)^{s-1}}{q^s}, \qquad p < t < 1. \tag{7}$$

In the sense of conditional densities, conditional on $I = i$ and $T_s = t$, the other eligible times are independent uniform variables in $(t, 1)$, and the lower-than-J times remain unrestricted and independent of them. Conditioning only on the first selected class also leaves the law of T_s unchanged.

Proof. The conditioning event restricts the coordinates separately. It therefore preserves their independence: $U_1, \ldots, U_s$ become independent uniform variables on $(p, 1)$, U_{s+1} becomes uniform on $(0, p)$, and every lower-priority coordinate retains its original uniform law on $(0, 1)$.

At the end of the sample, element $s+1$ has better priority than every element of index greater than $s+1$. Consequently, no such lower-priority element can be a global record before one of $1, \ldots, s$ arrives. The earliest of $1, \ldots, s$ has better priority than the sampled element and than every intervening arrival, so it is the first post-sample global record and is selected. Thus $T_s = \min_{1 \leq i \leq s} U_i$. For $p < t < 1$,

$$\mathbb{P}(T_s > t \mid J = s+1) = \left(\frac{1-t}{q}\right)^s;$$

taking the negative derivative gives (7).

The joint density of the eligible times is invariant under permutation of their coordinates. Hence the index I of the minimum is uniform on $\{1, \dots, s\}$ and independent of the minimum value T_s. More explicitly, conditional on $I = i$ and $T_s = t$ in the regular-density sense, each remaining eligible coordinate has the uniform density $1/(1-t)$ on $(t, 1)$, and their joint density factors. The lower-than-J coordinates are independent of the eligible coordinates and are not involved in this conditioning. Finally, the first selected parallel class is determined by I. Since I and T_s are independent, conditioning on that class does not change the law of T_s. □

The density (7) is defined for every positive integer s, regardless of the finite instance size. When using auxiliary quantities indexed by arbitrary $s \geq 1$, we take T_s to be the minimum of s independent uniform variables on $(p, 1)$, with that density. This convention does not require the event $J = s + 1$ to exist in a given instance.

If $s = 0$, element 1, the highest-priority nonloop, was sampled. It suppresses every later nonloop, so no first selection is possible and the conditional target count is zero.

4.1 The highest sampled element is inside the target prefix

For every positive integer s, define

$$d_s := \mathbb{E}[-T_s \log T_s]. \tag{8}$$

Lemma 4.2 (Early conditional cases)**.** *For $1 \leq s \leq m - 1$, conditional on $J = s + 1$, the expected target count is at least $1 + d_s$. For $s = m$, the conditional target count is exactly one. For every positive integer s, we have the identity*

$$d_s = s \sum_{j=1}^{\infty} \frac{q^j}{j} \left(\frac{1}{s+j} - \frac{q}{s+j+1} \right). \tag{9}$$

Proof. When $s < m$, all eligible elements and J belong to C_A. The first selection is one of the first m elements, contributing one target. Every element outside C_A is lower than J and hence has an unrestricted uniform time, independently of T_s. The highest-priority element outside C_A is target element $m + 1$. There is at least one such outside element because the matroid has rank two. Conditional on $T_s = t$, the second stage of the algorithm is precisely the classical record rule on the elements outside C_A: all outside arrivals at times at most t form the observed prefix, and the algorithm accepts the first subsequent outside record. Lemma 3.3, applied to this fixed outside set with cutoff t, gives probability at least $-t \log t$ of selecting element $m + 1$. Averaging over T_s gives the conditional lower bound $1 + d_s$.

When $s = m$, the sampled element $J = m + 1$ is the highest-priority element outside C_A. The first selection is one of the m target elements in C_A. Every later element outside C_A has lower priority than the already observed element $m + 1$, so no second acceptance occurs. Thus the target count is exactly one.

For the series identity, put $x = 1 - t$ in (7) and expand $-\log(1 - x) = \sum_{j \geq 1} x^j / j$. Every term is nonnegative, so termwise integration is justified by monotone convergence:

$$\begin{aligned} d_s &= \frac{s}{q^s} \int_0^q x^{s-1}(1-x) \sum_{j \geq 1} \frac{x^j}{j} \, dx \\ &= s \sum_{j \geq 1} \frac{q^j}{j} \left(\frac{1}{s+j} - \frac{q}{s+j+1} \right). \end{aligned}$$

This is (9). □

4.2 The entire target prefix is unsampled

For the rest of this section suppose $s \geq m+1$ and condition on $J = s+1$. For each parallel class C, let v_C be its number of eligible elements and $z_C = |P \cap C|$. Write

$$a = v_{C_A} \geq m, \qquad b = v_{C_B} \geq 1, \qquad a + b \leq s.$$

Then $z_{C_A} = m$, $z_{C_B} = 1$, and every other class has $z_C = 0$. The first selection's expected target contribution is exactly

$$\mathbb{P}(I \in P \mid J = s+1) = \frac{m+1}{s}. \tag{10}$$

This does not assume that every selected member of C_A or C_B is a target: there may be lower-priority eligible members of either class.

Lemma 4.3 (Second selection outside the first class)**.** *Condition further on the event that the first selected element belongs to parallel class C and that $T_s = t$. Let $v = v_C$ and $k = s - v$. Then $k \geq 1$.*

(i) If $J \notin C$, each specified eligible element outside C is the second accepted element with probability exactly $1/k$.

(ii) If $J \in C$, each specified eligible element outside C is the second accepted element with probability at least

$$h_k(t) = \int_t^1 \frac{t}{u} \frac{(1-u)^{k-1}}{(1-t)^k}\, du. \tag{11}$$

Furthermore $0 < h_k(t) \leq 1/k$.

Proof. Because $s \geq m+1$, every member of the rank-two prefix P is eligible. Thus the eligible set meets at least two parallel classes, and no first class C contains all eligible elements; hence $k \geq 1$. By Lemma 4.1, conditional on the first class and on $T_s = t$, the k eligible elements outside C have independent uniform times in $(t, 1)$.

Suppose first that $J \notin C$. It was observed by time $p < t$ and belongs to the outside-C observation set. It has better priority than every element of index greater than $s+1$. Thus it suppresses all lower-than-J outside-C arrivals. Every eligible outside-C element has index at most s and hence has better priority than J. Consequently, the earliest of the k eligible outside-C elements is exactly the next outside-C record and is accepted second. The k eligible times are exchangeable, so each specified element is earliest with probability $1/k$.

If $J \in C$, consider the L lower-than-J elements outside C. Their times are independent uniform variables in $(0, 1)$ and remain independent after conditioning on the first class and on $T_s = t$: that conditioning involves only the eligible coordinates, while the event $J = s+1$ restricts no lower-than-J coordinate. Until the earliest eligible outside-C element arrives, every outside-C observation is one of these L lower-priority elements. A lower-element record after time t would be accepted immediately and consume the second slot.

Fix a specified eligible element outside C. For it to be accepted at time u, it is necessary and sufficient that it be the earliest of the k eligible outside-C elements and that the L lower elements produce no record in (t, u). Under these two conditions no second selection has yet occurred, and the specified eligible element outranks every lower outside observation, so it is the required outside record. Its density for arriving at u before the other $k-1$ eligible outside elements is

$$\frac{1}{1-t}\left(\frac{1-u}{1-t}\right)^{k-1}.$$

This event is independent of the record process of the lower elements. Lemma 3.2 therefore gives the exact per-element selection probability

$$G_{k,L}(t) := \int_t^1 S_L(t,u)\frac{(1-u)^{k-1}}{(1-t)^k}\,du.$$

In particular,

$$G_{k,L}(t) - h_k(t) = \int_t^1 \left(1-\frac{t}{u}\right)(1-u)^L\frac{(1-u)^{k-1}}{(1-t)^k}\,du \geq 0.$$

This proves (11) and displays explicitly the nonnegative contribution being discarded. When $L=0$, the exact probability is $1/k$. Conditioning on the first class rather than on its particular first identity does not change the calculation: for every possible first identity in C, the outside eligible coordinates have the same product distribution on $(t,1)$. Finally, $t/u \leq 1$ gives

$$h_k(t) \leq \int_t^1 \frac{(1-u)^{k-1}}{(1-t)^k}\,du = \frac{1}{k},$$

and positivity follows from the positive integrand. □

Lemma 4.4 (Averaged second-selection probability). *Define $H_{s,k} = \mathbb{E}[h_k(T_s)]$ for positive integers s,k. Then*

$$H_{s,k} = s\sum_{\ell=0}^{\infty}\frac{q^\ell}{k+\ell}\left(\frac{1}{s+\ell}-\frac{q}{s+\ell+1}\right), \qquad 0 < H_{s,k} \leq \frac{1}{k}. \tag{12}$$

Proof. With $x = 1-t$ and $y = 1-u$, (11) becomes

$$h_k(t) = \frac{1-x}{x^k}\int_0^x \frac{y^{k-1}}{1-y}\,dy = (1-x)\sum_{\ell\geq 0}\frac{x^\ell}{k+\ell}.$$

The geometric-series expansion of $(1-y)^{-1}$ has nonnegative terms, so monotone convergence justifies the last equality. The variable $x = 1-T_s$ has density sx^{s-1}/q^s on $(0,q)$. A second application of monotone convergence yields

$$\begin{aligned} H_{s,k} &= \frac{s}{q^s}\sum_{\ell\geq 0}\frac{1}{k+\ell}\int_0^q (1-x)x^{s+\ell-1}\,dx \\ &= s\sum_{\ell\geq 0}\frac{q^\ell}{k+\ell}\left(\frac{1}{s+\ell}-\frac{q}{s+\ell+1}\right), \end{aligned}$$

which is (12). The pointwise bounds $0 < h_k(t) \leq 1/k$ from Lemma 4.3 remain valid after averaging. □

Let D be the parallel class of J. Summing first over the first selected class gives the conditional lower bound

$$\mathbb{E}[|\mathcal{A}\cap P| \mid J = s+1] \geq \frac{m+1}{s} + \sum_{C:v_C>0}\frac{v_C}{s}(m+1-z_C)\,\eta_C, \tag{13}$$

where $\eta_C = 1/(s-v_C)$ for $C \neq D$ and $\eta_D = H_{s,s-v_D}$ if $v_D > 0$. To justify the formula, condition on the first selected class C. Lemma 4.1 gives this event probability v_C/s, and there are exactly $m+1-z_C$ target elements outside C. If $C \neq D$, part (i) of Lemma 4.3 assigns each of them second-selection probability $1/(s-v_C)$. If $C = D$, part (ii) gives the pointwise lower bound $h_{s-v_C}(T_s)$.

The law of T_s is unchanged by conditioning on the first class, so its average is $H_{s,s-v_C}$. Adding the first-selection contribution (10) proves (13). This use of linearity of expectation requires no independence between the first and second target contributions.

We call D *penalized* when $v_D > 0$, because its summand uses H in place of $1/(s - v_D)$. Every other class is *unpenalized*. If $v_D = 0$, the class D cannot be selected first and no summand is penalized.

4.3 Uniform expressions for arbitrary parallel classes

An unpenalized class other than C_A, C_B, with $c \geq 1$ eligible elements, contributes

$$\frac{c(m+1)}{s(s-c)} \geq \frac{c(m+1)}{s(s-1)} \tag{14}$$

to the second-target part of (13). The exact expression follows from $v_C = c$, $z_C = 0$, and $\eta_C = 1/(s-c)$, and the inequality follows from $s - c \leq s - 1$. Thus its contribution can be replaced in a lower bound by that of c singleton classes. This is an arithmetic relaxation of (13), not a claim that physically splitting a class preserves the algorithm's law.

There is at most one penalized class, namely D. The following three expressions correspond respectively to $D = C_A$, $D = C_B$, and $D \notin \{C_A, C_B\}$. If $D = C_A$, apply (14) to all other classes except C_B, obtaining

$$V_A(m, s; a, b) = \frac{m+1}{s} + \frac{a}{s} H_{s,s-a} + \frac{bm}{s(s-b)} + \frac{(s-a-b)(m+1)}{s(s-1)}. \tag{15}$$

If D is C_B, the same procedure gives

$$V_B(m, s; a, b) = \frac{m+1}{s} + \frac{a}{s(s-a)} + \frac{bm}{s} H_{s,s-b} + \frac{(s-a-b)(m+1)}{s(s-1)}. \tag{16}$$

If D is another class with $c \geq 1$ eligible elements, keep it intact and split all remaining other classes. This yields

$$\begin{aligned} V_C(m, s; a, b, c) = {} & \frac{m+1}{s} + \frac{a}{s(s-a)} + \frac{bm}{s(s-b)} \\ & + \frac{(s-a-b-c)(m+1)}{s(s-1)} + \frac{c(m+1)}{s} H_{s,s-c}. \end{aligned} \tag{17}$$

Define the nonempty finite set

$$\mathcal{D}_{m,s} = \{(a, b) \in \mathbb{Z}^2 : a \geq m,\ b \geq 1,\ a + b \leq s\}$$

for $m \geq 1$ and $s \geq m + 1$. All denominators above are positive: the pair $(m, 1)$ shows that $\mathcal{D}_{m,s}$ is nonempty; $a, b \leq s - 1$; and an admissible c satisfies $c \leq s - a - b \leq s - 2$. Let

$$\begin{aligned} v_{m,s} = \min\Big(& \{V_A(m, s; a, b), V_B(m, s; a, b) : (a, b) \in \mathcal{D}_{m,s}\} \\ & \cup \{V_C(m, s; a, b, c) : (a, b) \in \mathcal{D}_{m,s},\ c \in \mathbb{Z},\ 1 \leq c \leq s - a - b\}\Big). \end{aligned} \tag{18}$$

Proposition 4.5 (Uniform conditional bound)**.** *For every instance and $m + 1 \leq s \leq n - 1$,*

$$\mathbb{E}[|\mathcal{A} \cap P| \mid J = s + 1] \geq v_{m,s}.$$

Proof. Suppose first that $v_D > 0$. If $D = C_A$, the actual counts a, b lie in $\mathcal{D}_{m,s}$ and the relaxation (14) turns (13) into the lower bound $V_A(m, s; a, b)$. The case $D = C_B$ gives $V_B(m, s; a, b)$. If D is another class, put $c = v_D$; then $1 \le c \le s - a - b$, and the same relaxation gives $V_C(m, s; a, b, c)$. Each applicable expression occurs among the candidates in (18), so it is at least $v_{m,s}$.

Now suppose that $v_D = 0$. Then D cannot be the first class, so every summand in (13) is unpenalized. Apply (14) to all classes except C_A, C_B. The resulting lower bound contains the unpenalized C_A term $a/[s(s-a)]$. Replacing it by $(a/s)H_{s,s-a}$ can only decrease the expression, because $H_{s,s-a} \le 1/(s-a)$. What remains is exactly $V_A(m, s; a, b)$ for the actual pair $(a, b) \in \mathcal{D}_{m,s}$, and is therefore at least $v_{m,s}$.

The minimizing counts may differ for different s; no simultaneous realizability of all minimizers is used or required. Each conditional expectation is bounded separately before averaging. □

4.4 The empty sample and the master inequality

Lemma 4.6 (Empty nonloop sample)**.** *Conditional on $J = \infty$, the expected target count is at least*

$$u_{m,n} = \min_{(a,b)\in\mathcal{D}_{m,n}} \left\{ \frac{m+1}{n} + \frac{a}{n(n-a)} + \frac{bm}{n(n-b)} + \frac{(n-a-b)(m+1)}{n(n-1)} \right\}. \tag{19}$$

Proof. The event $J = \infty$ restricts every nonloop coordinate separately to $(p, 1)$. Conditional on it, the nonloop times remain independent and uniform on that interval, so their relative arrival order is uniform. No nonloop was observed in the sample. The algorithm therefore accepts the first arriving nonloop and then the first arriving element outside its parallel class; the latter is automatically an outside-class record.

The first accepted identity is uniform among the n nonloops, so its expected target contribution is $(m+1)/n$. Let a and b now denote the total numbers of nonloops in C_A and C_B. Conditional on a first class C containing v nonloops, every specified element outside C is the earliest such element, and hence is selected second, with probability $1/(n - v)$. Consequently, the second-target contributions from first classes C_A and C_B are respectively $a/[n(n-a)]$ and $bm/[n(n-b)]$.

For every other first class of size c, the contribution is $c(m+1)/[n(n-c)] \ge c(m+1)/[n(n-1)]$. Summing this relaxation over the other classes gives the last term in (19). The actual pair (a, b) lies in $\mathcal{D}_{m,n}$; minimizing over that set yields the stated uniform lower bound. The event is an empty *nonloop* sample. Sampled loops, if any, do not participate in either record process and have no effect. □

Proposition 4.7 (Master inequality)**.** *For all $m \ge 1$ and $n \ge m + 1$,*

$$\begin{aligned} \mathbb{E}|\mathcal{A} \cap P| \ge L_{m,n} := &\sum_{s=1}^{m-1} pq^s(1 + d_s) + pq^m \\ &+ \sum_{s=m+1}^{n-1} pq^s v_{m,s} + q^n u_{m,n}. \end{aligned} \tag{20}$$

An empty sum is zero. Every term on the right is nonnegative.

Proof. The events $J = 1, \ldots, n, \infty$ are disjoint and exhaustive. By (6), the finite event $J = s + 1$ has probability pq^s, while $J = \infty$ has probability q^n. The case $s = 0$ has conditional target count zero. For $1 \le s < m$, Lemma 4.2 gives conditional lower bound $1 + d_s$; for $s = m$, it gives the exact value one. For $m + 1 \le s \le n - 1$, Proposition 4.5 gives $v_{m,s}$, and Lemma 4.6 gives $u_{m,n}$ on the final event. Multiplying each conditional bound by its event probability and applying the law of total expectation gives (20). The stated empty-sum convention covers the boundary cases $m = 1$ and $n = m + 1$. $\square$

5 Uniform analytical lower bound

Theorem 5.1 (Analytical master bound). *For $p = 1/e$, $q = 1 - p$, and every pair of integers $m \geq 1$, $n \geq m + 1$,*

$$L_{m,n} \geq \frac{p}{60}\left(55 + 228p - 129p^2 - 28p^3\right) > 2p.$$

We prove this bound by reducing to one infinite series and then estimating its terms uniformly. All estimates below are explicit analytical inequalities; no numerical minimization is required.

We will use the elementary enclosure

$$2 < e < \frac{8}{3} + \frac{1/4!}{1 - 1/5} = \frac{87}{32} < \frac{68}{25} < \frac{11}{4}. \tag{21}$$

Indeed, $8/3$ is the exponential series through $1/3!$. From $1/4!$ onward, the ratio of each succeeding term to its predecessor is at most $1/5$, with strict inequality after the first such ratio. The geometric tail therefore gives the displayed strict upper bound. Inverting the positive inequalities gives $32/87 < p < 1/2$; the weaker bound $e < 3$ also gives $p > 1/3$.

5.1 Passing to an infinite population bound

Define

$$L_{m,\infty} := \sum_{s=1}^{m-1} pq^s(1 + d_s) + pq^m + \sum_{s=m+1}^{\infty} pq^s v_{m,s}. \tag{22}$$

Lemma 5.2 (Finite populations dominate the infinite bound). *For every $m \geq 1$ and $n \geq m + 1$, the series in (22) converges and $L_{m,n} \geq L_{m,\infty}$.*

Proof. Write the expression minimized in (19) as

$$F_{m,n}(a, b) = \frac{m + 1}{n} + \frac{a}{n(n - a)} + \frac{bm}{n(n - b)} + \frac{(n - a - b)(m + 1)}{n(n - 1)}.$$

For fixed admissible a, b, its first and last terms sum to $(m + 1)(2n - a - b - 1)/(n(n - 1))$. The difference of this sum at n and $n + 1$ is

$$\frac{2(m + 1)(n - a - b)}{n(n - 1)(n + 1)} \geq 0.$$

Its remaining terms also decrease with n, since their positive denominators increase. A minimizing pair at n is still admissible at $n + 1$. Thus $u_{m,n+1} \leq u_{m,n}$.

For any admissible a, b, the bound $H_{s,s-a} \leq 1/(s - a)$ gives $v_{m,s} \leq V_A(m, s; a, b) \leq F_{m,s}(a, b)$. Minimizing the last expression yields $0 \leq v_{m,s} \leq u_{m,s}$. Here nonnegativity follows directly from the defining expressions (15)–(17). To bound $u_{m,s}$ from above, the admissible pair $(a, b) = (s - 1, 1)$ gives

$$u_{m,s} \leq F_{m,s}(s - 1, 1) = 1 + \frac{m}{s - 1} \leq 2.$$

Thus the infinite series converges absolutely, by comparison with a geometric series. Finally,

$$p\sum_{s=n}^{\infty} q^s v_{m,s} \le p\sum_{s=n}^{\infty} q^s u_{m,s} \le p\sum_{s=n}^{\infty} q^s u_{m,n} = q^n u_{m,n}.$$

The second inequality uses the monotonicity of $u_{m,s}$ just proved, and the last equality uses $1-q=p$. The left side is the tail added in (22); the right side is the empty-sample term of (20). All earlier terms agree. $\square$

5.2 Reducing to a prefix with two targets

Lemma 5.3 (Monotonicity in the target-prefix length)**.** *For every $m \ge 1$, $L_{m+1,\infty} > L_{m,\infty}$.*

Proof. Fix $s \ge m+2$. The candidate domain defining $v_{m+1,s}$ is a subset of the domain defining $v_{m,s}$. For fixed class counts, the three candidates are affine functions of m, with respective slopes

$$\begin{aligned}
\partial_m V_A &= \frac{1}{s} + \frac{b}{s(s-b)} + \frac{s-a-b}{s(s-1)},\\
\partial_m V_B &= \frac{1}{s} + \frac{b}{s}H_{s,s-b} + \frac{s-a-b}{s(s-1)},\\
\partial_m V_C &= \frac{1}{s} + \frac{b}{s(s-b)} + \frac{s-a-b-c}{s(s-1)} + \frac{c}{s}H_{s,s-c}.
\end{aligned}$$

Every summand in (12) is positive, because

$$\frac{1}{s+\ell} - \frac{q}{s+\ell+1} = \frac{1+p(s+\ell)}{(s+\ell)(s+\ell+1)} > 0.$$

The only k-dependent factor is $1/(k+\ell)$, so $H_{s,k}$ decreases as k increases. Since $b \ge 1$ and $H_{s,s-1} \le 1/(s-1)$, we have

$$\frac{b}{s-b} \ge \frac{1}{s-1} \ge H_{s,s-1}, \qquad bH_{s,s-b} \ge H_{s,s-1}.$$

For the V_A and V_C slopes, the first inequality supplies the required $H_{s,s-1}/s$ beyond the common term $1/s$; for the V_B slope, the second inequality does so. Every omitted slope term is nonnegative. Consequently, each slope is at least $(1+H_{s,s-1})/s$. Keeping only the $\ell=0$ summand in (12) gives

$$H_{s,s-1} \ge \frac{s}{s-1}\left(\frac{1}{s} - \frac{q}{s+1}\right) = \frac{1+sp}{s^2-1}, \qquad \frac{1+H_{s,s-1}}{s} \ge \frac{s+p}{s^2-1}.$$

Every candidate at $m+1$ is therefore at least its value at m plus $(s+p)/(s^2-1)$, and its value at m is at least $v_{m,s}$. Taking the minimum over the smaller domain gives

$$v_{m+1,s} - v_{m,s} \ge \frac{s+p}{s^2-1}. \tag{23}$$

The changing boundary term cancels using the exact identity

$$d_m = \frac{mq}{m+1}H_{m+1,1}. \tag{24}$$

To see this, direct integration in (11) gives $h_1(t) = -t\log t/(1-t)$, and the densities (7) satisfy $(1-t)f_m(t) = (mq/(m+1))f_{m+1}(t)$. Therefore

$$d_m = \int_p^1 h_1(t)(1-t)f_m(t)\,dt = \frac{mq}{m+1}\int_p^1 h_1(t)f_{m+1}(t)\,dt,$$

which is (24). At $s = m+1$, the admissible candidate $a = m, b = 1$ in V_A gives

$$v_{m,m+1} \le 1 + \frac{1}{m+1} + \frac{m}{m+1}H_{m+1,1}.$$

Writing out (22) for $m+1$ and for m, and subtracting the absolutely convergent series term by term, gives the exact identity

$$L_{m+1,\infty} - L_{m,\infty} = pq^m d_m + pq^{m+1}(1 - v_{m,m+1}) + \sum_{s=m+2}^{\infty} pq^s(v_{m+1,s} - v_{m,s}).$$

The identity for d_m and the upper bound on $v_{m,m+1}$ cancel the $H_{m+1,1}$ terms, leaving a lower bound $-pq^{m+1}/(m+1)$ for the first two terms. Applying (23) to the remaining sum and putting $\mu = m+1$, we obtain

$$L_{m+1,\infty} - L_{m,\infty} \ge pq^\mu \left[-\frac{1}{\mu} + \sum_{j=1}^{\infty} q^j \frac{\mu+j+p}{(\mu+j)^2-1}\right]. \tag{25}$$

For real $\mu \ge 2$ and integers $j \ge 1$, set $g_j(\mu) = \mu(\mu+j+p)/((\mu+j)^2-1)$. Writing $z = \mu+j$, differentiation gives

$$g_j'(\mu) = \frac{(1-p)(z-1)^2 + (j-1)(z^2+2pz+1)}{(z^2-1)^2} > 0.$$

Thus $g_j(\mu) \ge g_j(2)$ and $\sum_{j\ge1} q^j g_j(\mu) \ge R$, where the substitution $s = j+2$, partial fractions, and $\sum_{j\ge1} q^j/j = -\log(1-q) = -\log p = 1$ give

$$\begin{aligned} R &:= \sum_{s=3}^{\infty} q^{s-2}\frac{2(s+p)}{s^2-1} \\ &= \frac{1+p}{q}(1-q) + \frac{1}{q^2}\left(1 - q - \frac{q^2}{2} - \frac{q^3}{3}\right) \\ &= \frac{p(1+p)}{q} + \frac{p}{q^2} - \frac{1}{2} - \frac{q}{3}. \end{aligned}$$

Here the partial fractions are $2(s+p)/(s^2-1) = (1+p)/(s-1) + (1-p)/(s+1)$. For clarity, the two shifted sums are

$$\sum_{s=3}^{\infty} \frac{q^{s-2}}{s-1} = \frac{1-q}{q}, \qquad \sum_{s=3}^{\infty} \frac{q^{s-2}(1-p)}{s+1} = \frac{1-q-q^2/2-q^3/3}{q^2}.$$

An expansion now yields

$$6q^2(R-1) = P(p), \qquad P(x) = 36x - 15x^2 - 4x^3 - 11.$$

On $[0,1/2]$, $P'(x) = 36 - 30x - 12x^2 \ge 18 > 0$, and

$$P(32/87) = \frac{8563}{658503} > 0.$$

By (21), $R > 1$. The bracket in (25) is at least $(R-1)/\mu > 0$, proving the claim. $\square$

5.3 A uniform conditional bound for two targets

Lemma 5.4 (Two-target conditional bound)**.** *For every integer $s \geq 4$,*

$$v_{1,s} \geq \frac{18}{5s}.$$

The proof is given in Appendix A. It uses a lower bound on $H_{s,k}$, the arithmetic–geometric mean inequality, and convexity. The four smallest values $s = 4, 5, 6, 7$ are handled by explicit rational expressions and their adjacent differences. No class patterns are enumerated and no executable certificate is used.

The two remaining conditional values have simple closed forms.

Lemma 5.5 (The first two conditional values)**.**

$$v_{1,2} = \frac{3}{2} + \frac{1 - 3p^2}{4q^2}, \qquad v_{1,3} = 1 + \frac{1 - 3p^2 - 4p^3}{6q^3}.$$

Proof. For $s = 2$, the only admissible pair is $a = b = 1$, so $v_{1,2} = 3/2 + H_{2,1}/2$. Since $h_1(t) = -t \log t/(1-t)$,

$$H_{2,1} = \frac{2}{q^2} \int_p^1 -t \log t \, dt = \frac{1 - 3p^2}{2q^2}.$$

Here we used $\log p = -1$ when evaluating the elementary antiderivative.

For $s = 3$, write $X = H_{3,1}$ and $Y = H_{3,2}$. The admissible pairs are $(1,1), (1,2), (2,1)$. Substitution in (15)–(17), followed by the symmetry interchanging V_A and V_B, leaves the four candidate values

$$\frac{7}{6} + \frac{Y}{3}, \quad 1 + \frac{2Y}{3}, \quad \frac{5}{6} + \frac{2X}{3}, \quad \frac{4}{3} + \frac{Y}{3}.$$

We have $Y \leq 1/2$, and subtraction of the positive series gives

$$X - Y \geq 3\left(1 - \frac{1}{2}\right)\left(\frac{1}{3} - \frac{q}{4}\right) = \frac{1 + 3p}{8} > \frac{1}{4}.$$

The first inequality keeps only the $\ell = 0$ summand of the positive series for $X - Y$; the final strict inequality uses $p > 1/3$. These two estimates show explicitly that

$$1 + \frac{2Y}{3} \leq \frac{7}{6} + \frac{Y}{3}, \qquad 1 + \frac{2Y}{3} < \frac{4}{3} + \frac{Y}{3}, \qquad 1 + \frac{2Y}{3} < \frac{5}{6} + \frac{2X}{3}.$$

Thus the second displayed candidate is the minimum. To evaluate Y, the defining integral gives

$$h_2(t) = \frac{t(-\log t - 1 + t)}{(1-t)^2}, \qquad Y = \frac{3}{q^3} \int_p^1 t(-\log t - 1 + t) \, dt = \frac{1 - 3p^2 - 4p^3}{4q^3}.$$

Again using $\log p = -1$, the last equality follows from the antiderivative $-t^2 \log t/2 - t^2/4 + t^3/3$. Substitution into $v_{1,3} = 1 + 2Y/3$ proves the result. $\square$

5.4 Summing the lower bound

Proof of Theorem 5.1. By Lemmas 5.2 and 5.3, it suffices to prove the stated lower bound for $L_{1,\infty}$. By Lemmas 5.4 and 5.5,

$$\begin{aligned}\frac{L_{1,\infty}}{p} &\geq q + q^2 v_{1,2} + q^3 v_{1,3} + \frac{18}{5}\sum_{s=4}^{\infty}\frac{q^s}{s}\\ &= q + \frac{3}{2}q^2 + \frac{1-3p^2}{4} + q^3 + \frac{1-3p^2-4p^3}{6} + \frac{18}{5}\left(p - \frac{q^2}{2} - \frac{q^3}{3}\right)\\ &= \frac{55+228p-129p^2-28p^3}{60}.\end{aligned}$$

Here $\sum_{s\geq 1} q^s/s = -\log(1-q) = -\log p = 1$ and $1-q=p$. The last equality is a direct expansion in powers of p after substituting $q = 1-p$. Set $Q(x) = -65+228x-129x^2-28x^3$. For $0 \leq x \leq 1/2$, $Q'(x) = 228-258x-84x^2 \geq 78 > 0$, while

$$Q(32/87) = \frac{10873}{658503} > 0.$$

The enclosure (21) implies $Q(p) > 0$. Thus $L_{1,\infty}/p \geq 2 + Q(p)/60 > 2$, as required. □

6 Proof of the main theorem

Proof of Theorem 1.1. Lemma 2.3 proves that the algorithm uses only the information allowed in the theorem and always returns an independent set. Fix a realization of the arrival-independent tie-breaking labels, and use the resulting strict priority order throughout the argument. Invoke the arrival-time coupling of Lemma 3.1 and ignore the loop coordinates, which affect neither selection nor the optimum. The remaining arrival times are still independent and uniform; let n be their number. If $n = 0$, every element is a loop, so $w(\mathcal{A}) = \mathrm{OPT}(M, w) = 0$.

Suppose that $n \geq 1$. In this representation, the first selection is the classical record rule on the n nonloops with cutoff p. Lemma 3.3, with $d = n$ and $t = p$, therefore shows that the highest-priority nonloop, element 1, is selected first with probability at least $-p\log p = p$, where $p = e^{-1}$. Every nonempty priority prefix P_j contains element 1. Consequently, whenever $r(P_j) = 1$,

$$\mathbb{E}|\mathcal{A} \cap P_j| \geq \mathbb{P}(1 \in \mathcal{A}) \geq p = p\, r(P_j).$$

This proves all required prefix inequalities when the matroid has rank at most one, and it also handles every rank-one prefix in a rank-two matroid.

It remains to consider a matroid of rank two. Let P_{m+1} be its first rank-two priority prefix, as defined in Section 3. Proposition 4.7 and Theorem 5.1 give

$$\mathbb{E}|\mathcal{A} \cap P_{m+1}| \geq L_{m,n} > 2p.$$

Every rank-two prefix P_j contains P_{m+1}. Hence, outcome by outcome, $|\mathcal{A} \cap P_j| \geq |\mathcal{A} \cap P_{m+1}|$, and so $\mathbb{E}|\mathcal{A} \cap P_j| > 2p = p\, r(P_j)$. Together with the rank-one case, this proves that, for every $1 \leq j \leq n$, $\mathbb{E}|\mathcal{A} \cap P_j| \geq p\, r(P_j)$.

Because the priority order refines nonincreasing weight, each difference $w_j - w_{j+1}$ is nonnegative. The two identities in Lemma 3.4 now yield

$$\mathbb{E}w(\mathcal{A}) \geq \sum_{j=1}^{n}(w_j - w_{j+1})p\, r(P_j) = p\ \mathrm{OPT}(M, w) = \frac{\mathrm{OPT}(M, w)}{e}.$$

This inequality holds conditional on every realization of the arrival-independent tie-breaking labels. Averaging over those labels gives the unconditional assertion of the theorem. □

The prefix-count argument allows performance on the m target elements of C_A to compensate for a smaller inclusion probability of the single target in C_B. The following example shows why the proof cannot be replaced by an assertion that this algorithm selects every member of the optimal basis with probability at least $1/e$.

Example 6.1 (An optimal element with inclusion probability below $1/e$)**.** Let $E = \{a_1, a_2, a_3, b\}$, and let the independent sets be those containing at most one element of $\{a_1, a_2, a_3\}$ and at most one element of $\{b\}$. This is a rank-two partition matroid whose two parallel classes are

$$C_A = \{a_1, a_2, a_3\}, \qquad C_B = \{b\}.$$

Assign weights

$$w(a_1) = 4, \qquad w(a_2) = 3, \qquad w(a_3) = 2, \qquad w(b) = 1.$$

The priority order, from highest to lowest, is therefore a_1, a_2, a_3, b, and the unique maximum-weight independent set is $\{a_1, b\}$.

Use the arrival-time coupling and let J be the priority index of the highest-priority sampled element, with $J = \infty$ when the sample contains no element. On $J = 2$, an event of probability pq, the first selection is a_1 at time T_1. On $J = 3$, an event of probability pq^2, the first selection is the earlier of a_1, a_2, at time T_2. In either case the first selection lies in C_A, while b retains an independent uniform arrival time on $[0, 1]$. Since b is the only element outside C_A, it is selected if and only if it arrives after the first selection. From the density (7),

$$\mathbb{E}[1 - T_s] = \frac{s}{q^s} \int_p^1 (1 - t)^s \, dt = \frac{sq}{s+1}.$$

Thus the conditional probabilities of selecting b on $J = 2$ and $J = 3$ are respectively $q/2$ and $2q/3$.

On $J = 1$, the sampled element a_1 suppresses every later global record, so the algorithm makes no selection. On $J = 4$, the element b itself was sampled and cannot be selected; moreover, it suppresses every possible second selection outside C_A. Finally, on $J = \infty$, an event of probability q^4, the element b is selected with probability one: it is selected first if it is the earliest arrival, and otherwise the earliest arrival lies in C_A and b is selected second. These five cases are disjoint and exhaustive. Therefore

$$\mathbb{P}(b \in \mathcal{A}) = pq\frac{q}{2} + pq^2\frac{2q}{3} + q^4 = \frac{q^2}{2} + \frac{q^3}{6} + \frac{q^4}{3}.$$

The last equality uses $p = 1 - q$. If $f(q) = q^2/2 + q^3/6 + q^4/3$, then $f'(q) = q + q^2/2 + 4q^3/3 > 0$ for $q > 0$. By (21), $p > 4/11 > 1/3$, and hence $q < 2/3$. It follows that

$$\mathbb{P}(b \in \mathcal{A}) = f(q) < f(2/3) = \frac{82}{243} < \frac{4}{11} < p = \frac{1}{e}.$$

Thus a member of the unique optimum can have inclusion probability below $1/e$ under this algorithm. This does not conflict with Theorem 1.1, which concerns the expected total weight of the entire selected set.

A An analytical bound for the two-target class expressions

This appendix proves Lemma 5.4. Throughout, $m = 1$, $s \geq 4$, and all class counts are integers in the domains specified in (18): $a, b \geq 1$, $a + b \leq s$, and, for a V_C candidate, $1 \leq c \leq s - a - b$. We prove that every candidate in that finite set is greater than $18/(5s)$. When $m = 1$, the expressions V_A and V_B are interchanged by swapping a and b, so it suffices to treat V_A and V_C.

A.1 Lower bounds on the second-selection factor

For $k \geq 1$, let B have density $f_B(\beta) = k(1-\beta)^{k-1}$ on $[0, 1]$; this density integrates to one and

$$\mathbb{E}B = k \int_0^1 \beta(1-\beta)^{k-1}\, d\beta = \frac{1}{k+1}.$$

Substituting $u = t + (1-t)\beta$ into (11) gives

$$h_k(t) = \int_0^1 \frac{t(1-\beta)^{k-1}}{t + (1-t)\beta}\, d\beta = \frac{1}{k}\mathbb{E}\frac{t}{t + (1-t)B}.$$

For fixed $t \in [p, 1]$, the function $\beta \mapsto t/(t + (1-t)\beta)$ is convex on $[0, 1]$: its second derivative is $2t(1-t)^2/[t + (1-t)\beta]^3 \geq 0$. Jensen's inequality therefore gives

$$h_k(t) \geq \psi_k(t) := \frac{(k+1)t}{k(1+kt)}.$$

The function ψ_k is concave on $[p, 1]$, since $\psi_k''(t) = -2(k+1)/(1+kt)^3 < 0$. Thus

$$\psi_k(t) \geq \frac{1-t}{q}\psi_k(p) + \frac{t-p}{q}\psi_k(1).$$

Direct integration against (7) gives

$$\mathbb{E}T_s = \frac{s}{q^s}\int_p^1 t(1-t)^{s-1}\, dt = p + \frac{q}{s+1}.$$

Hence the expected coefficients of the two endpoints in the chord are $\mathbb{E}[(1-T_s)/q] = s/(s+1)$ and $\mathbb{E}[(T_s - p)/q] = 1/(s+1)$. Averaging the chord therefore gives

$$H_{s,k} \geq \frac{s}{s+1}\frac{k+1}{k(k+e)} + \frac{1}{s+1}\frac{1}{k} = \frac{1}{k} - \frac{A_s}{k(k+e)}, \qquad A_s = \frac{(e-1)s}{s+1}. \tag{26}$$

Because $0 < A_s < e$, direct subtraction gives

$$\left(\frac{1}{k} - \frac{A_s}{k(k+e)}\right) - \frac{1}{k + A_s} = \frac{A_s(e - A_s)}{k(k+e)(k+A_s)} \geq 0.$$

Thus $H_{s,k} \geq 1/(k + A_s)$. Since $A \mapsto 1/(k+A)$ is decreasing, for any $A \geq A_s$ we have

$$H_{s,k} \geq \frac{1}{k+A}. \tag{27}$$

A.2 Two class-count-free lower bounds

Let $A \geq A_s$ and set

$$X = \frac{s+A}{s-1}, \qquad Y = \frac{s}{s-1}.$$

For V_A, substitute (27) with $k = s - a$ into (15), set $m = 1$, and multiply by s. This gives

$$sV_A \geq 2 + \frac{2s}{s-1} + \left(\frac{a}{s+A-a} - \frac{2a}{s-1}\right) + \left(\frac{b}{s-b} - \frac{2b}{s-1}\right).$$

For the first parenthesis, put $z = s + A - a > 0$. Then

$$\frac{a}{s+A-a} - \frac{2a}{s-1} = \frac{s+A}{z} + \frac{2z}{s-1} - 1 - 2X \geq 2\sqrt{2X} - 1 - 2X,$$

where the last step is the arithmetic–geometric mean inequality applied to the two positive terms $(s+A)/z$ and $2z/(s-1)$. Applying the same calculation to the second parenthesis with $A = 0$ gives $2\sqrt{2Y} - 1 - 2Y$. Adding these estimates to the two leading terms yields

$$sV_A \geq 2\sqrt{2X} + 2\sqrt{2Y} - 2X. \tag{28}$$

This is a lower bound for every admissible pair; it does not assert that the separate minimizers can satisfy the coupling constraint $a + b \leq s$.

For V_C, the same substitution in (17) gives

$$\begin{aligned} sV_C \geq 2 + \frac{2s}{s-1} &+ \left(\frac{a}{s-a} - \frac{2a}{s-1}\right) + \left(\frac{b}{s-b} - \frac{2b}{s-1}\right) \\ &+ \left(\frac{2c}{s+A-c} - \frac{2c}{s-1}\right). \end{aligned}$$

Applying the arithmetic–geometric mean inequality to each parenthesis gives

$$sV_C \geq 4\sqrt{2Y} + 4\sqrt{X} - 2 - 2Y - 2X. \tag{29}$$

Indeed, after substituting $z = s + A - c$, the last parenthesis is

$$\frac{2(s+A)}{z} + \frac{2z}{s-1} - 2 - 2X \geq 4\sqrt{X} - 2 - 2X.$$

The first two parentheses each contribute $2\sqrt{2Y} - 1 - 2Y$; combining all three contributions with the leading terms gives (29).

A.3 Penalized target classes when $s \geq 8$

By (21), we may choose

$$A = \frac{43s}{25(s+1)} \geq A_s.$$

Indeed, $e < 68/25$ implies $e - 1 < 43/25$, and multiplication by $s/(s+1)$ gives the required comparison. Writing $\tau = 1/(s-1) \in (0, 1/7]$, we have

$$Y = 1 + \tau, \qquad X = (1+\tau)\left(1 + \frac{43\tau}{25(1+2\tau)}\right), \qquad 1 \leq Y \leq \frac{8}{7}, \quad 1 \leq X \leq \frac{2144}{1575}.$$

Both factors in the displayed expression for X are increasing in τ, so the last upper bound is its value at $\tau = 1/7$.

For a concave function g on $[1, x_0]$, the graph lies above its endpoint chord. Consequently, if λ is no larger than the chord slope, then $g(x) \geq g(1) + \lambda(x-1)$ throughout that interval. Applying this fact to the following concave functions gives

$$2\sqrt{2X} - 2X \geq 2\sqrt{2} - 2 - \frac{139}{200}(X-1), \tag{30}$$

$$2\sqrt{2Y} \geq 2\sqrt{2} + \frac{273}{200}(Y-1). \tag{31}$$

For completeness, we now verify the two chosen slopes using exact rational comparisons. The inequalities

$$\sqrt{2} > \frac{7071}{5000}, \qquad \sqrt{2144/1575} < \frac{1167}{1000}, \qquad \sqrt{8/7} < \frac{107}{100}$$

follow on squaring from the positive integer differences

$$\begin{aligned} 2 \cdot 5000^2 - 7071^2 &= 959, \\ 1575 \cdot 1167^2 - 2144 \cdot 1000^2 &= 975175, \\ 7 \cdot 107^2 - 8 \cdot 100^2 &= 143. \end{aligned}$$

At the right endpoint, the square-root part of the chord slope in (30) is

$$\frac{2\sqrt{2}}{1 + \sqrt{2144/1575}} > \frac{14142}{10835} > \frac{261}{200};$$

after subtracting the slope 2 of the linear term $2X$, this gives $-139/200$. Similarly, the chord slope in (31) is

$$\frac{2\sqrt{2}}{1 + \sqrt{8/7}} \frac{14142}{10350} > \frac{273}{200}.$$

The two rational comparisons have cross-product differences 465 and 2850, respectively. This proves (30) and (31) on their full claimed intervals.

Combining (28)–(31) gives

$$sV_A \geq 4\sqrt{2} - 2 + F(\tau), \qquad F(\tau) = \frac{67\tau}{100} - \frac{5977}{5000}\frac{\tau(1+\tau)}{1+2\tau}.$$

For $0 \leq \tau \leq 1/7$,

$$\frac{d}{d\tau}\frac{\tau(1+\tau)}{1+2\tau} = \frac{1}{2} + \frac{1}{2(1+2\tau)^2} \geq \frac{65}{81},$$

so

$$F'(\tau) \leq \frac{67}{100} - \frac{5977}{5000}\frac{65}{81} = -\frac{23431}{81000} < 0.$$

Thus $F(\tau) \geq F(1/7)$, and substitution gives

$$sV_A \geq \frac{4571}{1250} + \frac{1}{7}\left(\frac{67}{100} - \frac{5977}{5000}\frac{8}{9}\right) = \frac{567113}{157500} = \frac{18}{5} + \frac{113}{157500} > \frac{18}{5}. \tag{32}$$

Here $4\sqrt{2} - 2 > 4571/1250$ follows from $\sqrt{2} > 7071/5000$, and $\tau(1+\tau)/(1+2\tau) = 8/63$ at $\tau = 1/7$. Symmetry proves the same bound for V_B.

A.4 Penalized other classes for every $s \geq 4$

Use the same A, now with $\tau = 1/(s-1) \leq 1/3$. Then

$$1 \leq Y \leq \frac{4}{3}, \qquad 1 \leq X \leq \frac{224}{125}, \qquad X - 1 \leq \frac{68}{25}(Y-1).$$

Indeed, $X - 1 = \tau + (43/25)\tau(1+\tau)/(1+2\tau)$ and $(1+\tau)/(1+2\tau) \leq 1$. Chords of the two concave functions give

$$4\sqrt{X} - 2X \geq 2 - \frac{29}{100}(X-1), \tag{33}$$

$$4\sqrt{2Y} - 2Y - 2 \geq 4\sqrt{2} - 4 + \frac{5}{8}(Y-1). \tag{34}$$

The same endpoint-chord argument used above applies here. To verify the chosen slopes explicitly, use

$$\sqrt{224/125} < \frac{229}{171}, \qquad \sqrt{4/3} < \frac{5774}{5000}.$$

Their squared cross-product differences are, respectively,

$$125 \cdot 229^2 - 224 \cdot 171^2 = 5141 > 0, \qquad 3 \cdot 5774^2 - 4 \cdot 5000^2 = 17228 > 0.$$

The chord slope in (33) exceeds

$$\frac{4}{1 + 229/171} - 2 = -\frac{29}{100}.$$

The chord slope in (34) exceeds $28284/10774 - 2 > 5/8$, since $8 \cdot 28284 - 21 \cdot 10774 = 18 > 0$.

Combining these bounds with (29) gives

$$\begin{aligned} sV_C &\geq 4\sqrt{2} - 2 + \left(\frac{5}{8} - \frac{29}{100}\frac{68}{25}\right)(Y-1) \\ &\geq \frac{4571}{1250} - \frac{273}{5000} = \frac{18011}{5000} = \frac{18}{5} + \frac{11}{5000} > \frac{18}{5}. \end{aligned} \tag{35}$$

Here the coefficient of $Y-1$ is $-819/5000$, and $Y - 1 \leq 1/3$. The rational lower bound for $4\sqrt{2} - 2$ is the same one used in (32); since the coefficient is negative, replacing $Y - 1$ by its upper bound $1/3$ preserves the displayed lower bound.

A.5 The four small target-class cases

It remains to bound V_A for $s = 4, 5, 6, 7$. From (26) and $e < 11/4$, we have

$$H_{s,k} \geq \widehat{H}_{s,k} := \frac{1}{k} - \frac{7s}{(s+1)k(4k+11)}.$$

To check the direction of this substitution, rewrite the correction term in (26) as

$$\frac{s}{(s+1)k}\frac{e-1}{k+e}.$$

The last factor has derivative $(k+1)/(k+e)^2 > 0$ as a function of e. Replacing e by the larger value $11/4$ therefore increases the quantity being subtracted and preserves a valid lower bound; simplifying gives $\widehat{H}_{s,k}$. Extend the displayed rational expression $\widehat{H}_{s,k}$ to every real $k > 0$, and define, for real $0 < a, b < s$,

$$\phi_s(a) = a\widehat{H}_{s,s-a} - \frac{2a}{s-1}, \qquad \gamma_s(b) = \frac{b}{s-b} - \frac{2b}{s-1}.$$

Then

$$sV_A \geq 2 + \frac{2s}{s-1} + \phi_s(a) + \gamma_s(b). \tag{36}$$

Both functions are strictly convex. For γ_s, $\gamma_s''(b) = 2s/(s-b)^3 > 0$ on $(0, s)$. For ϕ_s, set $E = 11/4$ and $A = (E-1)s/(s+1)$; the identity

$$\widehat{H}_{s,k} = \frac{1 - A/E}{k} + \frac{A/E}{k+E}$$

expresses $a\widehat{H}_{s,s-a}$ as a positive linear combination of the strictly convex functions $a/(s-a)$ and $a/(s+E-a)$. The coefficients are positive because $0 < A/E < 1$, and

$$\frac{d^2}{da^2}\frac{a}{c-a} = \frac{2c}{(c-a)^3} > 0 \qquad (0 < a < c).$$

Subtracting a linear function preserves strict convexity.

For a convex function, the differences between successive integer values are nondecreasing. The following explicit substitutions therefore identify the minimum of ϕ_s over integers $1 \leq a \leq s-1$:

s	a_*	$\phi_s(a_*)$	$\phi_s(a_*) - \phi_s(a_* - 1)$	$\phi_s(a_* + 1) - \phi_s(a_*)$
4	2	$-179/285$	$-280/1311$	$724/1425$
5	2	$-104/207$	$-2957/14904$	$659/15732$
6	3	$-53/115$	$-103/2070$	$501/2185$
7	3	$-121/288$	$-3323/44640$	$431/6624$

Every backward difference is negative and every forward difference is positive. No other integer needs checking because of convexity. Likewise,

$$\gamma_s(b+1) - \gamma_s(b) = \frac{s}{(s-b-1)(s-b)} - \frac{2}{s-1} \qquad (1 \leq b \leq s-2)$$

gives the minima in the next table:

s	b_*	$\gamma_s(b_*)$	Backward difference	Forward difference
4	1	$-1/3$	left endpoint	0
5	2	$-1/3$	$-1/12$	$1/3$
6	2	$-3/10$	$-1/10$	$1/10$
7	2	$-4/15$	$-1/10$	$1/60$

For $s = 4$, both $b = 1$ and $b = 2$ minimize γ_4; the table records $b_* = 1$, and the zero forward difference records the tie. Using these separate minima only weakens the bound, since it drops the constraint $a + b \leq s$: the minimum of the coupled sum is at least the sum of the two separate minima. Substitution into (36) gives

s	Lower bound on sV_A	Excess over $18/5$
4	$352/95$	$2/19$
5	$1517/414$	$133/2070$
6	$837/230$	$9/230$
7	$5251/1440$	$67/1440$

All entries follow by substitution into the displayed rational functions; the positive excesses prove the desired bound. Symmetry covers V_B. Together with (32) and (35), this proves $sV_A, sV_B, sV_C > 18/5$ for every admissible candidate with $s \geq 4$. The candidate set in (18) is finite and nonempty, since $(a, b) = (1, 1)$ supplies admissible V_A and V_B candidates. Taking its minimum and dividing by $s > 0$ gives $v_{1,s} > 18/(5s)$, which proves Lemma 5.4.

Use of AI: We prepared a list of proof strategies and techniques related to Matroid Secretary problem. We fed this to multiple AI tools (Claude, ChatGPT and two open-weight models with open-source harnesses) and prompted these AI tools to improve the existing bounds. When an AI tool was hallucinating, stuck, or going in the wrong direction, we stopped that AI tool, corrected the intermediate `.md` files, added more specific new ideas and switched to a different AI tool with more detailed fresh prompts. Equations written by one AI tool were critiqued, verified and expanded by a completely different AI tool in an iterative manner, till we were satisfied with their correctness and details. The author is solely responsible for the correctness of the proofs.

References

[1] T. S. Ferguson. *Who Solved the Secretary Problem?* Statistical Science 4(3):282–289, 1989. https://doi.org/10.1214/ss/1177012493.

[2] M. Babaioff, N. Immorlica, and R. Kleinberg. *Matroids, Secretary Problems, and Online Mechanisms*. SODA 2007, pp. 434–443. https://www.cs.cornell.edu/~rdk/papers/matsec.pdf.

[3] M. Babaioff, N. Immorlica, D. Kempe, and R. Kleinberg. *Matroid Secretary Problems*. Journal of the ACM 65(6), Article 35, 2018. https://doi.org/10.1145/3212512.

[4] S. Chakraborty and O. Lachish. *Improved Competitive Ratio for the Matroid Secretary Problem*. SODA 2012, pp. 1702–1712. https://doi.org/10.1137/1.9781611973099.135.

[5] O. Lachish. $O(\log \log \text{rank})$ *Competitive-Ratio for the Matroid Secretary Problem*. FOCS 2014. Full version: https://arxiv.org/abs/1403.7343.

[6] M. Feldman, O. Svensson, and R. Zenklusen. *A Simple* $O(\log \log(\text{rank}))$*-Competitive Algorithm for the Matroid Secretary Problem*. Mathematics of Operations Research 43(2):638–650, 2018. https://doi.org/10.1287/moor.2017.0876. https://arxiv.org/abs/1404.4473.

[7] S. Singla. *The Matroid Secretary Conjecture is True*. arXiv:2609.14555v1, 2026. https://arxiv.org/abs/2609.14555v1.

[8] R. Kleinberg. *A Multiple-Choice Secretary Algorithm with Applications to Online Auctions*. SODA 2005, pp. 630–631. https://www.cs.cornell.edu/~rdk/papers/MultSec.pdf.

[9] T. Kesselheim, K. Radke, A. Tönnis, and B. Vöcking. *An Optimal Online Algorithm for Weighted Bipartite Matching and Extensions to Combinatorial Auctions*. ESA 2013, Lecture Notes in Computer Science 8125, pp. 589–600.
https://doi.org/10.1007/978-3-642-40450-4_50.

[10] J. A. Soto, A. Turkieltaub, and V. Verdugo. *Strong Algorithms for the Ordinal Matroid Secretary Problem*. Mathematics of Operations Research 46(2):642–673, 2021.
https://doi.org/10.1287/moor.2020.1083. https://arxiv.org/abs/1802.01997.

[11] S. Im and Y. Wang. *Secretary Problems: Laminar Matroid and Interval Scheduling*. SODA 2011, pp. 1265–1274. https://doi.org/10.1137/1.9781611973082.96.

[12] P. Jaillet, J. A. Soto, and R. Zenklusen. *Advances on Matroid Secretary Problems: Free Order Model and Laminar Case*. IPCO 2013, Lecture Notes in Computer Science 7801, pp. 254–265.
https://doi.org/10.1007/978-3-642-36694-9_22. Full version:
https://arxiv.org/abs/1207.1333v2.

[13] T. Ma, B. Tang, and Y. Wang. *The Simulated Greedy Algorithm for Several Submodular Matroid Secretary Problems*. Theory of Computing Systems 58(4):681–706, 2016.
https://doi.org/10.1007/s00224-015-9642-4. https://arxiv.org/abs/1107.2188v3.

[14] Z. Huang, Z. Parsaeian, and Z. Zhu. *Laminar Matroid Secretary: Greedy Strikes Back*. ESA 2024, Leibniz International Proceedings in Informatics 308, pp. 73:1–73:8.
https://doi.org/10.4230/LIPIcs.ESA.2024.73. https://arxiv.org/abs/2308.09880.

[15] K. Bérczi, V. Livanos, J. A. Soto, and V. Verdugo. *Matroid Secretary via Labeling Schemes*. IPCO 2025, Lecture Notes in Computer Science 15620, pp. 128–141.
https://doi.org/10.1007/978-3-031-93112-3_10. Full version, Section 4.3:
https://arxiv.org/abs/2411.12069v2.