

Knapsack Secretary with Bursty Adversary

Thomas Kesselheim*

Marco Molinaro†

Abstract

The random-order or *secretary* model is one of the most popular beyond-worst case model for online algorithms. While this model avoids the pessimism of the traditional adversarial model, in practice we cannot expect the input to be presented in perfectly random order. This has motivated research on *best of both worlds* (algorithms with good performance on both purely stochastic and purely adversarial inputs), or even better, on inputs that are a *mix* of both stochastic and adversarial parts. Unfortunately the latter seems much harder to achieve and very few results of this type are known.

Towards advancing our understanding of designing such robust algorithms, we propose a random-order model with *bursts* of adversarial time steps. The assumption of burstiness of unexpected patterns is reasonable in many contexts, since changes (e.g. spike in a demand for a good) are often triggered by a common external event. We then consider the *Knapsack Secretary* problem in this model: there is a knapsack of size k (e.g., available quantity of a good), and in each of the n time steps an item comes with its value and size in $[0, 1]$ and the algorithm needs to make an irrevocable decision whether to accept or reject the item.

We design an algorithm that gives an approximation of $1 - \tilde{O}(\Gamma/k)$ when the adversarial time steps can be covered by $\Gamma \geq \sqrt{k}$ intervals of size $\tilde{O}(\frac{n}{k})$. In particular, setting $\Gamma = \sqrt{k}$ gives a $(1 - O(\frac{\ln^2 k}{\sqrt{k}}))$ -approximation that is resistant to up to a $\frac{\ln^2 k}{\sqrt{k}}$ -fraction of the items being adversarial, which is almost optimal even in the absence of adversarial items. Also, setting $\Gamma = \tilde{\Omega}(k)$ gives a constant approximation that is resistant to up to a constant fraction of items being adversarial. While the algorithm is a simple “primal” one, it does not possess the crucial symmetry properties exploited in the traditional analyses. The strategy of our analysis is more robust and significantly different from previous ones, and we hope it can be useful for other beyond-worst-case models.

*Institute of Computer Science, University of Bonn, 53115 Bonn, Germany. Email: thomas.kesselheim@uni-bonn.de

†Computer Science Department, PUC-Rio, Brazil. Email: mmolinaro@inf.puc-rio.br

1 Introduction

In standard competitive analysis of online algorithms, one assumes that an adversary completely defines the input. While this is a useful model for designing algorithms for many problems, for many others this model is too pessimistic and no algorithm can outperform the trivial ones. One classical example is the *Secretary Problem* and its generalizations. In this problem, one is presented a sequence of n items of values v_1, \dots, v_n . Upon each arrival, one has to decide *irrevocably* if one accepts or rejects the item, without knowing the value of future items in the sequence. The goal is to select a single item in order to maximize the value obtained. It is easy to see that in the adversarial model the best guarantee possible is to obtain expected value that is a $\frac{1}{n}$ -fraction of the offline optimum, and this is achieved by the trivial algorithm that chooses one of the n time steps at random and blindly accepts the item in this time step.

In order to avoid the pessimism of this model and allow for the design of non-trivial algorithms with *hopefully* better performance in practice, there has been a push to consider *beyond worst-case* models. One of the most prominent such models is the *random-order model*, where the adversary can choose the set of items in the instance by they are presented in uniformly random order. This model has been studied since at least the 60s and has seen a lot of developments in the past decade, and several problems are now well-understood under this model, such as Knapsack and more generally Packing LPs [1, 2, 3, 15, 17, 21, 29], assignment problems [8, 12, 17], matroid optimization [4, 6, 13, 14, 23], and many more. For example, for the Secretary Problem in the random-order model one can obtain a $\frac{1}{e}$ -fraction of the offline optimal value (as $n \rightarrow \infty$) with the following classical threshold-based algorithm: reject the first $\frac{1}{e}$ -fraction of items but note their maximum value, then select the next element which exceeds this value if such an element appears.

However, in practice we cannot expect the sequence to arrive exactly in random order. This has motivated research on *best of both worlds*, namely algorithms with good performance on both purely stochastic and purely adversarial inputs [22, 25, 26, 28]. Even more interesting are algorithms that work well on inputs that are a *mix* of both stochastic and adversarial parts. But this seems to be much harder to achieve: in online algorithms we are only aware of the results of [11] on budgeted allocation (see Section 1.3 for a description of their model and assumptions), while in online learning results of this type have only been obtained very recently for multi-armed bandits [16, 24, 31, 32]. We note that all these results are for settings in which non-trivial guarantees can be achieved for pure adversarial inputs.

Towards advancing our understanding of designing such robust algorithms, we introduce a model that mixes random-order and adversarial time steps, assuming that the latter comes in *bursts*. The random-order times represent when the environment is in a “stationary” or “predictable” state, while the adversarial times represent “unexpected” patterns. The assumption of burstiness of unexpected patterns is reasonable in many contexts, since changes are often triggered by an external common event, e.g., the surge in gun sales after news of possible changes in gun control regulations. See [7, 9, 19, 20] for examples of the different ways in which burstiness can be modeled and areas of applications.

1.1 The *Bursty Adversary plus Random Order (BARO)* model

We describe more formally the general version of the proposed model BARO. Consider an online problem where decisions are made sequentially and irrevocably at times $1, 2, \dots, n$. In our model, the adversary first chooses some the time steps $Adv \subseteq [n]$ to be “adversarial” and leaves the others $RO = [n] \setminus Adv$ as “random-order” times. In order to capture the burstiness of the adversarial time steps in a clean way, let \mathbb{W} be the partition of $[n]$ into disjoint intervals of length ℓ . We then assume that the adversarial times Adv are covered by at most Γ intervals in \mathbb{W} . Notice that this allows various patterns in the adversarial part of the input, including individual (non-bursty) adversarial times as well as bursts of size much larger than ℓ ,

for a total of up to $\Gamma\ell$ adversarial times. As in the standard random-order model, the items/inputs on the random-order times RO are arbitrary but presented in uniform random order. The *sequence* items/inputs on the adversarial times Adv is fully adversarial that can be adaptively generated based on an algorithm’s behavior and may even depend on the *order* of the items in RO .

It is important to highlight that the algorithm does not know which time steps are adversarial or random-order, and that in each time step only one item arrives (i.e., the adversarial items do not come in batches).

Note that in many problems this adversary can make an instance completely adversarial by sending “dummy” random-order items. For example, in the Secretary Problem the adversary can set the value of all random-order items to be 0; so again no non-trivial guarantees is possible in this case. In order to obtain meaningful guarantees, we compare the algorithm’s performance only to the optimum over the random-order times RO , which we denote by OPT_{RO} . Thus, in a maximization problem we say that an algorithm is α -*competitive* in the BARO model if the expected value of the algorithm is at least αOPT_{RO} .

1.2 Our Results

In this paper we use the BARO model to obtain a more robust algorithm for the *Knapsack Secretary* problem, a well-studied generalization of the Secretary Problem. The offline version of the problem is the standard Knapsack Problem: there are n items, each with a value $v_i \geq 0$ and size $w_i \in [0, 1]$, and we have a knapsack of size k ; the goal is to select a subset of items with total size at most k , and with total value as large as possible.

Our main result is an algorithm for the Knapsack Problem in the BARO model that is resistant to a fraction of items being adversarial.

Theorem 1.1. *There is a $(1 - O(\frac{\Gamma\ell}{n} \ln \frac{n}{\Gamma\ell}))$ -competitive algorithm for the Knapsack Problem in the BARO model where the adversarial times can be covered by $\Gamma \geq \sqrt{k}$ windows of size $\ell = \frac{n \ln k}{k}$.*

Notice that the term $\frac{\Gamma\ell}{n}$ in the guarantee is precisely the fraction of adversarial items that the algorithm can cope with. For example, setting $\Gamma = \sqrt{k}$, our algorithm obtains a $(1 - O(\frac{\ln^2 k}{\sqrt{k}}))$ -approximation in the presence of up to a $O(\frac{\ln^2 k}{\sqrt{k}})$ -fraction of items being adversarial. This approximation is almost optimal: even in the absence of adversarial items (and even when all items are unit-sized) the best approximation possible is $1 - \Omega(\frac{1}{\sqrt{k}})$ [21] (and this is achieved for example by [1, 15, 17, 29]). Note that these competitive ratios go to 1 as the budget $k \rightarrow \infty$ (recall the normalization of sizes being at most 1). Moreover, with $\Gamma = \Omega(\frac{n}{\ell})$ the algorithm achieves a constant approximation in the presence of a constant fraction of adversarial items.

Primal Algorithm with Time-Based Constraints. Our starting point is the *primal* strategy for the random-order model, whose high-level idea is the following: At time t , one solves a knapsack LP with the items seen so far but with budget proportionally scaled to be $\lceil \frac{t}{n} k \rceil$, and pick (a fraction of) the item at time t exactly as prescribed by the optimal LP solution, if there is space available in the full budget of k .

While this strategy obtains the optimal guarantee in the random-order model [17], it fails in the presence of adversarial items. One way in which it fails is by picking “too many items”: Suppose that the first k items are adversarial, have size 1, and they all have infinitesimal values but sorted in increasing order, and the random-order items have all value and size equal to 1; it is easy to see that the primal algorithm will pick all the adversarial items, filling up the budget with items of infinitesimal value. (Similar examples exist where the adversarial items are not in the beginning of the sequence.) To counter this, in our algorithm we add additional restrictions, outside of the LP, that the algorithm can only pick a constant number of items in each window of size $\ell \approx \frac{n}{k}$, which is roughly the behavior of the optimal solution if the n items were in random order.

However, the algorithm may now fail by picking “too few” items: consider the same example as before but now all the adversarial items have value $1 + \varepsilon$, thus slightly more valuable than the random-order items. The algorithm will then only pick 1 of these adversarial items (by the new restriction added) and will not pick any of the random-order items, since the LP will always fill up its budget with the better adversarial items; so the algorithm obtains value $1 + \varepsilon$, while the $\text{OPT}_{RO} = k$. To avoid this, we also add additional constraints *to the LP* that its solution can select at most a constant number of items in each window of size $\ell \approx \frac{n}{k}$ (note there are $\frac{t}{n}k$ disjoint such windows in $[t]$ and the LP selects total size $\approx \frac{t}{n}k$, again on average 1 per window).

The main difficulty is analyzing the algorithm in the presence of the additional restrictions/constraints. Previous analyses of primal-style algorithms crucially relied on the fact the LP (and its optimal solution) was invariant to the permutation of items/coordinates. This brings about some crucial independence properties: Decisions at time t are independent of the *order* of the arrivals at times $1, \dots, t - 1$ and therefore of the respective decisions. This property allows for the direct use of known concentration inequalities to control the total occupation incurred by the algorithm.

Since our new restrictions/constraints are not permutation invariant, we need to use a different type of analysis. The main handle is what we call the *weighted rank* of an item: the sum of the weights of items with higher value density $\frac{v_i}{w_i}$ than this item, divided by the knapsack capacity. That is, it is by how much one would have to scale the knapsack capacity before the offline optimum would start picking this item. The very high-level idea of the analysis is intuitive: The higher the weighted rank of an item, the smaller its probability of being picked by the LP, even with the new constraints. In addition, while there are complicated dependencies between the events “the algorithm picks the item at time t ”, the weighted ranks of the items in the random-order times are almost independent: they are just sampled without replacement. We leverage this to obtain custom concentration inequalities that control the algorithm’s occupation of the different restrictions/constraints.

1.3 Related Work

As already pointed out above, many algorithms have been proposed for online optimization problems with random arrival order. However, these algorithms usually break when moving to the BARO model. For concreteness, let us illustrate the effect on Kleinberg’s algorithm [21] for the multiple-choice secretary problem, a special case of our problem. The algorithm is allowed up to k selections. Throughout the sequence, it never picks items which are not among the best k so far. Therefore, we can construct the following counterexample. Consider a sequence starting with an adversarial burst of k items of very high value, followed by a random-order sequence with items of smaller values. On this sequence, the algorithm will not pick any random-order items at all. If $n \gg k$, then with high probability (over the randomness of the algorithm) none of the adversarial items are picked either (the threshold-based algorithm for the secretary problem is applied to the first $\approx \frac{n}{k}$ items w.h.p., in which case the first $\approx \frac{1}{e} \frac{n}{k} \gg k$ items are rejected). This argument transfers immediately to other algorithms, such as [2, 17]. Other algorithms such as the one by Agrawal and Devanur [1] or by Babaioff et al. [4] use the beginning of the sequence to estimate the optimal value, which also fails in this sequence.

There is only surprisingly little work when it comes to non-uniform random order model. Recently, Kesselheim et al. [18] introduced models where the order of the items is “much less random” than the uniform random order. Among other results, they show that it is possible to obtain constant-competitive algorithms for the Multiple-choice Secretary Problem under these weaker assumptions, and quantify the minimum entropy of the distribution over orders that admits constant-competitive algorithms for the Secretary Problem. We

remark that these models do not explicitly contain adversarial items.

Closer in spirit to our model, Esfandiari et al. [11] consider online budgeted allocation in an online model that mixes both stochastic and adversarial inputs. They provide algorithms that are optimal when the input is totally adversarial, and whose performance improves when the instance becomes “more stochastic”. There are two crucial differences between our proposed model and Esfandiari et al.’s model: in the latter, while the adversarial items may appear at any point in the sequence (i.e., no burstiness assumption), it is assumed that the algorithm *knows* the distribution of the items in the non-adversarial times, unlike in our model. Also, unlike the Knapsack Problem studied here, the budgeted allocation problem has constant-competitive algorithms even in the adversarial model. Thus, while to some extent an algorithm does not need to worry about “losing everything” if it is fooled by the adversarial part of the instance, its design and analysis have to be delicate enough to obtain fine control over the constants in the competitive-ratio in order to yield interesting results.

In a very recent paper, Bradac et al. [5] present several results for robust secretary problems in a mixed model very similar to ours, which was inspired by a discussion about a preliminary version of this present paper. In contrast to our model, there is no assumption on the number or burstiness of adversarial rounds, making the results incomparable. Our focus is understand situations in which we are close to the optimal guarantee without adversarial rounds. Since their adversary is more powerful, the guarantees are worse in two ways: (i) Their benchmark is weakened by leaving out the best item. (ii) The guarantees depend on the overall number of rounds n , whereas ours only depend on k . The techniques are also quite different.

2 BARO Knapsack: model and algorithm

Model. We consider an online knapsack problem. The algorithm knows upfront the knapsack size k and the number of items n , and the items are presented online, one-by-one. In the t -th time step, the current item’s value V_t and size W_t are revealed, and the algorithm needs to irrevocably decide what fraction $X_t^{alg} \in [0, 1]$ of this item to select. Our algorithm’s selection is always integral, i.e., $X_t^{alg} \in \{0, 1\}$, but our point of comparison is the best fractional solution. The selections made by the algorithm need to fit the knapsack, namely $\sum_{t \in [n]} W_t X_t^{alg} \leq k$ with probability 1, and it tries to maximize the total value of its selections: $\sum_{t \in [n]} V_t X_t^{alg}$. Importantly, the choice in the t -th step has to be made only knowing V_1, \dots, V_t and W_1, \dots, W_t (as well as k and n).

The sequences $V_1, \dots, V_n \geq 0$ and $W_1, \dots, W_n \in [0, 1]$ are generated by the following *Bursty Adversary plus Random Order* (BARO) model. Let us fix a window size ℓ , and let \mathbb{W} denote the collection of disjoint windows of size ℓ that partitions the time steps $[n]$, that is, $\mathbb{W} = \{\{1, 2, \dots, \ell\}, \{\ell + 1, \dots, 2\ell\}, \dots\}$. For concreteness we will use window size $\ell := \frac{n \ln k}{k}$. The adversary first partitions the n times steps into sets Adv (adversarial) and RO (random-order) with the property that Adv can be covered by Γ windows in \mathbb{W} ; we use $\mathbb{W}^{adv} \subseteq \mathbb{W}$ to denote one such cover, fixed throughout. The adversary also fixes the items for the random-order times, namely the value/size pairs $(v_1, w_1), (v_2, w_2), \dots, (v_{|RO|}, w_{|RO|})$, with $w_i \in [0, 1]$ for all i . Moreover, for each random-order time $t \in RO$, nature samples *without replacement* an index I_t from $\{1, 2, \dots, |RO|\}$, i.e., randomly chooses which random-order item will appear at that time. Then, for each time step t the adversary outputs an item with value V_t and size $W_t \in [0, 1]$ as follows:

- (Adversarial) If $t \in Adv$, the adversary outputs an item with arbitrary value V_t and size $W_t \in [0, 1]$; this may depend on an algorithm’s behavior and on the I_t ’s.

- (Random-order) If $t \in RO$, the adversary outputs the item indexed by I_t , namely that with value $V_t := v_{I_t}$ and size $W_t := w_{I_t}$.

Note that there is a subtle difference between capital and small letters here. By V_t and W_t , we refer to the value and weight of the item *arriving in the t -th step*. By v_i and w_i we refer to the i -th random-order item specified by the adversary *before the random permutation is applied*. Consequently, V_t and W_t are random variables whereas v_i and w_i are not. Furthermore, since the I_t 's are sampled without replacement, the items $((V_t, W_t))_{t \in RO}$ in the random-order times are precisely the items $(v_1, w_1), (v_2, w_2), \dots, (v_{|RO|}, w_{|RO|})$ randomly permuted.

Again we highlight that the algorithm does not know which time steps are adversarial and which are random-order, and that the adversarial items do not come in batches. As mentioned before, the benchmark for comparison is the offline optimum for the problem on the random-order items alone, namely $\text{OPT}_{RO} := \max\{\sum_i v_i x_i : \sum_i w_i x_i \leq k, x \in [0, 1]^{|RO|}\}$.

Algorithm. The algorithm we propose is a modification of the primal method of [17] and can be described as follows. Let \mathbb{W}_t be the collection of windows \mathbb{W} truncated to the prefix $[t]$, namely $\{1, \dots, \ell\}, \{\ell + 1, \dots, 2\ell\}, \dots, \{\lfloor \frac{t}{\ell} \rfloor \ell + 1, \dots, t\}$. At time t , in order to compute its selection $X_t^{\text{alg}} \in \{0, 1\}$ of the current item, the algorithm first finds an optimal solution X^t to the following (random) linear program LP_t :

$$\begin{aligned} & \max \sum_{t' \leq t} V_{t'} X_{t'} \\ & \text{s.t.} \sum_{t' \leq t} W_{t'} X_{t'} \leq c_t \frac{t}{n} k && \text{(main inner budget)} \\ & \sum_{t' \in B} W_{t'} X_{t'} \leq a_1 \frac{\ell}{n} k, \quad \forall B \in \mathbb{W}_t && \text{(inner constraints)} \\ & X \in [0, 1]^t, \end{aligned}$$

where we introduce the slight budget scaling $c_t := (1 - \frac{4\Gamma\ell}{t})$, and set the constant $a_1 := 601$. If $X_t^t > 0$, we say that the algorithm *tentatively picks* the item at time t . The algorithm checks if it can permanently pick this item by verifying whether its past selections $X_1^{\text{alg}}, \dots, X_{t-1}^{\text{alg}}$ satisfy the following constraints:

$$\begin{aligned} & \sum_{t' < t} W_{t'} X_{t'} \leq k - 1 && \text{(main budget)} \\ & \sum_{t' \in B_{\text{last}}} W_{t'} X_{t'} \leq a_4 \frac{\ell}{n} k - 1, && \text{(outer constraint)} \end{aligned}$$

where B_{last} denotes the last window in \mathbb{W}_{t-1} , and a_4 is a sufficiently large constant (set in Lemma 4.6). If so, the algorithm fully picks the item, namely it sets $X_t^{\text{alg}} = 1$; otherwise we say that it is *blocked* and it does not pick the item at all, setting $X_t^{\text{alg}} = 0$.

To get some intuition why the algorithm is reasonable, let us observe how the ‘‘offline optimum’’ OPT_{RO} builds up over time. We can define random variables X_t^* indicating what fraction of the item arriving at time t is packed in OPT_{RO} . Because the permutation is uniformly random, these random variables are identically distributed for all $t \in RO$. More specifically, we have $\mathbb{E}[V_t X_t^*] = \text{OPT}_{RO}/|RO| \approx \text{OPT}_{RO}/n$ and $\mathbb{E}[W_t X_t^*] \leq k/|RO| \approx k/n$. So, *in expectation*, slightly scaled versions of the random variables fulfill all constraints stated above. Our algorithm, of course, does not know X_t^* but tries to mimic this process.

Particularly, the goal of ([inner constraints](#)) and ([outer constraint](#)) is to spread out the choices made by the algorithm over time so that the consequences of adversarial bursts are mitigated.

Notice that by construction the solution X^{alg} returned by the algorithm is always feasible, namely $\sum_{t \leq n} W_t X_t^{alg} \leq k$. Thus, we only need to argue that it obtains enough value.

Theorem 2.1 (Total value). *The expected value of the solution X^{alg} returned by the algorithm satisfies*

$$\mathbb{E} \left[\sum_{t \in RO} V_t X_t^{alg} \right] \geq \left(1 - O \left(\frac{\Gamma \ell}{n} \ln \frac{n}{\Gamma \ell} \right) \right) \text{OPT}_{RO}.$$

Roadmap of the analysis. In [Section 3](#) we upper bound for each random-order time t the probability that the algorithm tentatively selects that item. Next, we boost this per-time upper bound into concentration inequalities for the volume of the selections made up to a given point, and use it to upper bound the probability that the algorithm is blocked by constraint ([main budget](#)) or ([outer constraint](#)), in which case it would not be able to make permanent its tentative selection ([Section 4](#)). Using this, we lower bound the value obtained by the algorithm in each (free) random-order time step ([Section 5](#)), and add over all such time steps to show that the algorithm obtains the desired value ([Section 6](#)).

Without loss of generality we assume that the random-order times are sorted in decreasing order of value density, namely $\frac{v_1}{w_1} \geq \frac{v_2}{w_2} \geq \dots \geq \frac{v_{|RO|}}{w_{|RO|}}$. Also, we say that an item is **better** than another if it has higher value density. For simplicity, we also assume that no item has value or weight equal to 0 (else automatically exclude/include in the solution), and that the sum of all item sizes is at least the knapsack size k . We also assume that there are no ties in the value densities $\frac{v_i}{w_i}$; this can be accomplished by infinitesimal perturbations to the values, for example. We also assume $\frac{n}{2} \geq k \geq 80$ and that $\frac{\Gamma \ell}{n} \leq \frac{1}{2}$, so at most half of the windows can have adversarial items. With overload of notation, we use I_t to denote the actual item (pair (V_t, W_t)) at time t , even when t is an adversarial time.

3 Controlling tentative selections via weighted rank

We use $T_t := \mathbf{1}(X_t^t > 0)$ to denote the indicator of *tentative* selection by the algorithm at time t . Our goal in this section is to argue that the algorithm does not tentatively select too many items. As mentioned before, the main handle for making this formal is the notion of *weighted rank*. The weighted rank of the random-order item i is a $\frac{1}{k}$ scaling of the sum of the weights of random-order items better than it (recall these items are sorted in decreasing order of value density $\frac{v_i}{w_i}$).

Definition 3.1 (Weighted rank). *The weighted rank of the random-order item i is $r_i := \frac{1}{k} \sum_{i' < i} w_{i'}$ (we also define $r_{|RO|+1} = \frac{1}{k} \sum_i w_i$ for convenience). For a random-order time t , we use $R_t := r_{I_t}$ to denote the total weighted rank of the item I_t at this time.*

As before, one interpretation of the weighted rank r_i is the following: considering the offline problem with only random-order items, r_i is by how much we need to scale the knapsack of size k before the optimal fractional solution wants to pick a strictly positive fraction of item i . Thus, the higher the rank the worse the item is.

The main result of this section says that the worse the item at time t is, the less likely the algorithm is to tentatively pick it. (The extra conditioning on items $(I_{t'})_{t' \in S}$ will be technically useful later and may be ignored throughout at a first read.)

Theorem 3.2 (UB tentative selection). *Consider a random-order time $t \geq 8\ell(\Gamma + 1)$, and a set S of random-order times with $|S| \leq \frac{\ln k}{4}$. Then*

$$\Pr\left(T_t = 1 \mid R_t, (I_{t'})_{t' \in S}\right) \leq \psi(R_t), \quad \text{where } \psi(\gamma) = \begin{cases} 1, & \text{if } \gamma < 1 \\ \frac{2}{k}, & \text{if } \gamma \in [1, 50] \\ 4ke^{-\frac{\gamma}{20} \ln k}, & \text{if } \gamma > 50. \end{cases}$$

For the rest of the section we prove this result. At its heart is the following deterministic monotonicity property of the LP: Fix a scenario (so the LP is deterministic); if there *is* a solution for the LP with only items better than I_t that saturates the main budget, then I_t is not included at all in the *optimal* LP solution. This is clear if we did not have the inner constraints: The optimal LP solution is obtained by the greedy procedure, and if we can saturate the budget with only better items the greedy will stop before reaching I_t . While this does not hold necessarily hold in the presence of general side constraints, we show it still does under the simple inner constraints. To streamline the presentation, the proof is presented in Appendix B.

Lemma 3.3. *Consider a time $t \in [n]$, and fix a scenario I_1, I_2, \dots, I_n . Suppose that there is a feasible solution \bar{X} of LP_t with $\sum_{t' \leq t} W_{t'} \bar{X}_{t'} = c_t \frac{tk}{n}$ and whose support only includes times with items strictly better than I_t (i.e., $\bar{X}_{t'} > 0$ implies that $I_{t'}$ is strictly better than I_t , for all $t' \in [t]$). Then in any optimal solution X^* of LP_t we have $X_t^* = 0$. (Thus, I_t is not tentatively selected by our algorithm.)*

Our next lemma will leverage this result to show that if there are many items in random-order-only windows better than I_t , then the probability of tentatively selecting the latter is small. Before that, we need to introduce the definition of *free time*, the ones we will focus on for most of the analyses.

Definition 3.4 (FREE_t and RO_t). *A time is free if it does not belong to one of the adversarial windows \mathbb{W}^{adv} . We use FREE_t to denote the collection of free times in $[t]$. Furthermore, $\mathbb{W}_t^{\text{free}}$ denotes the windows from \mathbb{W}_t that only contain free times.*

We also use $\text{RO}_t := \text{RO} \cap [t]$ to denote all the random-order times (free or otherwise) in $[t]$. With slight abuse in notation, we also use RO_t to denote the cardinality of RO_t .

The following estimates follow directly from the assumption that there are at most Γ adversarial windows, each of size ℓ .

Observation 3.5. *The following holds: (a) If $t \geq 2\Gamma\ell$ then $|\text{FREE}_t| \geq \frac{t}{2}$; (b) $\frac{1}{\text{RO}_n} \leq \frac{1}{|\text{FREE}_n|} \leq \frac{1}{n} (1 + \frac{2\Gamma\ell}{n})$.*

We can finally state the promised lemma.

Lemma 3.6. *Consider a random-order time $t \geq 2(\Gamma\ell + 1)$. For a value $\gamma \geq 0$, let G_γ be the event that the sum of the sizes of the items in the times FREE_t that are better than I_t equals $\gamma c_t \frac{tk}{n}$ (i.e., $\sum_{t' \in \text{FREE}_t: I_{t'} < I_t} W_{t'} = \gamma c_t \frac{tk}{n}$). Then for any set of random-order times $S \subseteq \text{RO}$ with $|S| \leq \frac{\ln k}{4}$, we have*

$$\Pr(T_t = 1 \mid G_\gamma, I_t, (I_{t'})_{t' \in S}) \leq \frac{1}{2} \psi(\gamma) = \begin{cases} \frac{1}{k}, & \text{if } \gamma \in [1, 50] \\ 2ke^{-\frac{\gamma}{20} \ln k}, & \text{if } \gamma > 50. \end{cases}$$

Proof. Condition on $I_t, (I_{t'})_{t' \in S}$, and on the set of items $\{I_{t'}\}_{t' \in \text{FREE}_{t-1}}$ in the free times in a way that the event G_γ holds; let ω denote this conditioning. It suffices to show the upper bound $\Pr(T_t = 1 \mid \omega) \leq \frac{1}{2} \psi(\gamma)$, and the lemma follows by taking expectation with respect to multiple of these ω 's. Also notice that this conditioning does not fix the *relative order* of the items in $\text{FREE}_{t-1} \setminus S$, thus

$$\text{The items at times } \text{FREE}_{t-1} \setminus S \text{ are in random order even when conditioning on } \omega. \quad (3.1)$$

Let E be the event that there is a feasible solution X for LP_t whose support only has items better than I_t and that saturates the main budget, i.e., $\sum_{t' \leq t} W_{t'} X_{t'} = c_t \frac{tk}{n}$. From Lemma 3.3, whenever E holds I_t is *not* tentatively selected, so it suffices to *lower bound* the probability $\Pr(E \mid \omega)$.

Case 1: $\gamma \in [1, 50]$. If for each of the free windows \mathbb{W}_{t-1}^{free} the total size of items better than I_t in the window is at most $a_1 \ell \frac{k}{n}$ (not “too many good items” in any free window), then any (fractional) selection of these items of total size $c_t \frac{tk}{n}$ gives a feasible solution for LP_t saturating the main budget, so E holds; notice that it is possible to select this much size because we are in the case $\gamma \geq 1$. The intuition is that since the total size of these good items is $\gamma c_t \frac{tk}{n} \leq 50 \frac{tk}{n}$, each window should have about $\frac{\ell}{t} \cdot 50 \frac{tk}{n} = 50 \ell \frac{k}{n}$ of their size in it, so with high probability no window has more than $a_1 \ell \frac{k}{n}$ of their size (recall $a_1 \gg 50$).

More formally, consider a free window $B \in \mathbb{W}_{t-1}^{free}$. Let $Z_{B \setminus S} = \sum_{t' \in B \setminus S} \mathbf{1}(I_{t'} < I_t) \cdot W_{t'}$ be the sum of sizes of items in $B \setminus S$ better than I_t , and let $Z = \sum_{t' \in \text{FREE}_{t-1} \setminus S} \mathbf{1}(I_{t'} < I_t) \cdot W_{t'}$. Notice that under the conditioning ω , Z is a fixed number satisfying $Z \leq \gamma c_t \frac{tk}{n} \leq \gamma \frac{tk}{n}$, and that $Z_{B \setminus S}$ is a sum of terms sampled without replacement from the terms in Z (because of Observation (3.1)). Thus, we have

$$\mathbb{E}[Z_{B \setminus S} \mid \omega] = \frac{|B \setminus S|}{|\text{FREE}_{t-1} \setminus S|} \mathbb{E}[Z \mid \omega] \leq \frac{\ell}{|\text{FREE}_{t-1}| - \ln k} \cdot \frac{\gamma tk}{n} = \frac{t}{|\text{FREE}_{t-1}| - \ln k} \cdot \gamma \ell \frac{k}{n} \leq 3\gamma \ell \frac{k}{n},$$

where the last inequality uses the fact that $t \geq 2\Gamma\ell$, Observation 3.5, and the assumptions $\Gamma \geq \sqrt{k}$ and $k \geq 80$. Moreover, we can apply the concentration inequality for sampling without replacement (Lemma A.1) conditionally to the sum $Z_{B \setminus S}$ (with $\tau = (600 - 3\gamma)\ell \frac{k}{n}$) to obtain

$$\Pr\left[Z_{B \setminus S} \geq 600 \ell \frac{k}{n} \mid \omega\right] \leq 2 \exp\left(-\frac{9}{7} 3\gamma \ell \frac{k}{n}\right) \leq 2 \frac{1}{k^3} \leq \frac{1}{k^2},$$

where in the first inequality we also used that $\tau \geq 3 \cdot 3\gamma \ell \frac{k}{n}$ because $\gamma \leq 50$, and in the last inequality that $k \geq 80$. Since $|S| \leq \frac{\ln k}{4}$ and each item has size at most 1, the items in $B \cap S$ have total size less than $\ell \frac{k}{n}$. Thus, the conditional probability is at most $\frac{1}{k^2}$ that the total size of items in B better than I_t is at least $a_1 \ell \frac{k}{n}$ (“too many good items”). Since there are fewer than k windows, by taking a union bound over all free windows $B \in \mathbb{W}_t^{free}$ we see that with probability at least $1 - \frac{1}{k}$ none of these windows has too many good items. Thus, $\Pr(E \mid \omega) \geq 1 - \frac{1}{k}$.

Case 2: $\gamma > 50$. The number of windows in \mathbb{W}_{t-1}^{free} of size ℓ (i.e., possibly excluding the last window) is at least $num := \frac{t}{\ell} - \Gamma - 1$. If in each such window the total size of items better than I_t is at least $2\ell \frac{k}{n}$ (“good items everywhere”), then one can (fractionally) select up to $2\ell \frac{k}{n}$ -mass of them in each window and get a feasible solution for LP_t that saturates the main budget; this saturation is possible because this can give a total of size $(2\ell \frac{k}{n}) \cdot num \geq c_t \frac{tk}{n}$ of these better items, where the last inequality uses $t \geq 2\ell(\Gamma + 1)$. Since in this case event E holds, it suffices to lower bound the probability of having good items everywhere. The intuition again is that by assumption there is total mass $\gamma c_t \frac{tk}{n} \geq 12 \frac{tk}{n}$ of these better items, so each window should have about $\frac{\ell}{t} \cdot 12 \frac{tk}{n} = 12 \ell \frac{k}{n}$ size in it, and with high probability all of them should have at least $2\ell \frac{k}{n}$ size in it.

Again, consider any fixed window $B \in \mathbb{W}_{t-1}^{free}$ of size ℓ , and define the sums $Z_{B \setminus S}$ and Z as in the previous case. Now conditioned on ω we have $Z \geq \gamma c_t \frac{tk}{n} - |S| \geq \frac{\gamma tk}{2n} - \frac{\ln k}{4}$ (using the definition of c_t and $t \geq 8\sqrt{k}\ell$), and hence

$$\mathbb{E}[Z_{B \setminus S} \mid \omega] = \frac{|B \setminus S|}{|\text{FREE}_t \setminus S|} \mathbb{E}[Z \mid \omega] \geq \frac{\ell - \frac{\ln k}{4}}{t} \cdot \left(\frac{\gamma tk}{2n} - \frac{\ln k}{4}\right) \geq \frac{\gamma \ell k}{3n},$$

where the last inequality uses $t \gg \frac{n \ln k}{k}$. Again employing the concentration inequality for sampling without replacement (Lemma A.1) conditionally to the sum $Z_{B \setminus S}$ (with $\tau = (\frac{\gamma}{3} - 2)\ell \frac{k}{n}$) we get

$$\Pr \left[Z_{B \setminus S} \leq 2\ell \frac{k}{n} \mid \omega \right] \leq 2 \exp \left(-\frac{(\frac{\gamma}{3} - 2)^2}{\frac{5\gamma}{3} - 2} \ln k \right) \leq 2 \exp \left(-\frac{\gamma}{20} \ln k \right),$$

where the last inequality uses $\gamma \geq 50$. Taking a union bound over the at most k such windows, the probability that we have enough good items in each window in $\mathbb{W}_{t-1}^{\text{free}}$ of size ℓ is at least $1 - 2ke^{-\frac{\gamma}{20} \ln k}$. This concludes the proof. \square

In order to remove the conditioning on G_γ from the previous lemma, we show that this event holds with high probability whenever the weighted rank of I_t is high (i.e., there are many items better than it); again this is just a consequence of concentration of measure. Actually we work with the event $G_{\geq \alpha} := \bigvee_{\gamma \geq \alpha} G_\gamma$, namely that the total size of items better than I_t in FREE_t is at least $\alpha c_t \ell \frac{k}{n}$.

Lemma 3.7. *Consider $t \geq 8\Gamma\ell$ and a set S of random-order times with $|S| \leq \frac{\ln k}{4}$. Then for $\alpha \geq 1$,*

$$\Pr [G_{\geq \alpha} \mid R_t = \alpha, (I_{t'})_{t' \in S}] \geq 1 - e^{-\alpha \ln k}.$$

Proof. Again let $Z = \sum_{t' \in \text{FREE}_{t-1} \setminus S} \mathbf{1}(I_{t'} < I_t) \cdot W_{t'}$ be the total size of items better than I_t in $\text{FREE}_t \setminus S$, and let ω denote the conditioning on $R_t = \alpha$ and $(I_{t'})_{t' \in S}$; it suffices to show $\Pr [Z \geq c_t \alpha \ell \frac{k}{n} \mid \omega] \geq 1 - e^{-\alpha \ln k}$.

Conditioned on ω , the total weight of items better than I_t in $RO_n \setminus S$ is at least $k\alpha - |S|$, and since these items are in random order (even conditioning on ω) we have

$$\mathbb{E}[Z \mid \omega] \geq (k\alpha - |S|) \frac{|\text{FREE}_{t-1} \setminus S|}{|RO_n \setminus S|} \geq \frac{k\alpha}{n} (t - 3\Gamma\ell)$$

where in the last inequality we used $\alpha \geq 1$ and the fact

$$|\text{FREE}_{t-1} \setminus S| \geq (t-1) - \Gamma\ell - \ln k \geq t - 2\Gamma\ell.$$

Again we can apply the concentration inequality for sampling without replacement (Lemma A.1) conditionally to the sum Z (with $\tau = \frac{\Gamma\ell}{t} \mathbb{E}[Z \mid \omega]$) to obtain

$$\begin{aligned} \Pr \left(Z < c_t \alpha t \frac{k}{n} \mid \omega \right) &\leq \Pr \left(Z < \underbrace{\mathbb{E}[Z \mid \omega]}_{=: \beta} \left(\frac{t - 4\Gamma\ell}{t - 3\Gamma\ell} \right) \mid \omega \right) \\ &\leq 2 \exp \left(-\frac{\beta^2 \mathbb{E}[Z \mid \omega]}{4 + \beta} \right) \\ &\leq 2 \exp \left(-\frac{\frac{\Gamma\ell}{t} (t - 4\Gamma\ell) \frac{k\alpha}{n}}{5} \right) \\ &\leq 2 \exp \left(-\frac{\frac{1}{2} \Gamma \alpha \ln k}{5} \right) \\ &\leq e^{-\alpha \ln k}, \end{aligned}$$

where in the third and fourth inequalities we used $t \geq 8\Gamma\ell$ to obtain $\beta \geq \frac{\Gamma\ell}{t}$, and the last inequality uses $\alpha \geq 1$, $\Gamma \geq \sqrt{k}$, and $k \geq 80$. This concludes the proof. \square

Putting the previous two lemmas together we finally obtain the proof of Theorem 3.2.

Proof of Theorem 3.2. We lower bound the probability that $T_t = 0$. First, notice that since the bound of Lemma 3.6 is non-increasing in γ , it still holds if we replace the conditioning on G_γ for a conditioning on $G_{\geq\gamma}$ (i.e., we condition on having possible more items better than I_t). Also, by assumption, all items have different weighted rank, so conditioning on $R_t = \alpha$ is equivalent to conditioning on I_t being the item with weighted rank α . Using these observation, we can apply Lemma 3.6 to obtain

$$\begin{aligned} \Pr[T_t = 0 \mid R_t = \alpha, (I_{t'})_{t' \in S}] &\geq \Pr[T_t = 0 \wedge G_{\geq\alpha} \mid R_t = \alpha, (I_{t'})_{t' \in S}] \\ &= \Pr[T_t = 0 \mid G_{\geq\alpha}, R_t = \alpha, (I_{t'})_{t' \in S}] \Pr[G_{\geq\alpha} \mid R_t = \alpha, (I_{t'})_{t' \in S}] \\ &\stackrel{\text{L. 3.6 and 3.7}}{\geq} \left(1 - \frac{1}{2}\psi(\alpha)\right) (1 - e^{-\alpha \ln k}) \geq 1 - \frac{1}{2}\psi(\alpha) - e^{-\alpha \ln k} \geq 1 - \psi(\alpha). \end{aligned}$$

This concludes the proof. \square

4 Controlling the probability of being blocked

In this section we show that with good probability, when the algorithm tentatively selects an item, it also permanently selects it, i.e., it is not blocked by the constraints (**main budget**) and (**outer constraint**). More precisely, let $O_t := W_t X_t^{\text{alg}}$ be the actual occupation incurred by the the algorithm at time t . We use F_t to denote the indicator of the event that the algorithm is *not* blocked at time t , i.e., $F_t = 1$ if

$$\sum_{t' \in B_{\text{last}}} O_{t'} \leq a_4 \frac{\ell}{n} k - 1 \quad \text{and} \quad \sum_{t' < t} O_{t'} \leq k - 1, \quad (4.2)$$

where again B_{last} is the last window in \mathbb{W}_{t-1} . Otherwise $F_t = 0$. The following is the main result of this section.

Theorem 4.1 (Probability of being blocked). *For all free times $t \geq 8\ell(\Gamma + 2)$, the probability of being blocked is upper bounded as $\Pr(F_t = 0 \mid I_t) \leq \frac{O(1)}{k(1 - \frac{t}{n} - a_5 \frac{\Gamma \ln k}{k})^2}$, for some constant a_5 .*

To prove this result, we will upper bound the probability that either of the two parts of (4.2) is violated. This is done respectively in Lemmas 4.6 and 4.10; Theorem 4.1 then follows by a union bound. While the first part of (4.2) only concerns the occupation from free time steps, the second part also includes non-free ones. To control this second part, we will nonetheless focus on the occupation over the free windows; for non-free windows B the outer constraints guarantee $\sum_{t' \in B} O_{t'} \leq O(\frac{\ell k}{n}) = O(\ln k)$, and so all the Γ of these windows combined can consume only $O(\Gamma \ln k)$ of the budget (so, for example, in the important case $\Gamma = \sqrt{k}$ this is negligible). For the free time steps, it suffices to upper-bound bound the (permanent) occupation O_t by the *tentative occupation* $O'_t := W_t T_t$: For the algorithm to select the item at time t , it is necessary but not sufficient that $T_t = 1$. Therefore, we have $O_t \leq O'_t$ and we focus on controlling the O'_t 's from now on.

As a start, we use Theorem 3.2 to show that in each free time step the expected tentative occupation $\mathbb{E}[O'_t]$ is at most $\approx \frac{k}{n}$; thus, essentially both (4.2) hold in expectation. While what we actually need is a generalization of this result, we present it to illustrate the techniques in a clearer way.

Lemma 4.2 (UB tentative occupation). *For all free times $t \geq 8\ell(\Gamma+1)$, we have $\mathbb{E}[O'_t] \leq \frac{k}{RO_n} (1 + O(\frac{1}{k}))$.*

Proof. Since fixing I_t fixes W_t , using Theorem 3.2 we have

$$\begin{aligned} \mathbb{E} O'_t &= \mathbb{E} W_t T_t = \mathbb{E}_{I_t} [W_t \cdot \mathbb{E}[T_t \mid I_t]] = \mathbb{E}_{I_t} [W_t \cdot \Pr(T_t = 1 \mid I_t)] \\ &\stackrel{T.3.2}{\leq} \mathbb{E}_{I_t} [W_t \cdot \psi(R_t)] = \frac{1}{RO_n} \sum_i w_i \psi(r_i). \end{aligned}$$

Since by definition of rank $r_j = \frac{1}{k} \sum_{j' < j} w_{j'}$, we have $r_{i+1} - r_i = \frac{w_i}{k}$, and thus $w_i = k \cdot \int_{r_i}^{r_{i+1}} 1 \, dx$. Applying this to the last displayed inequality we get

$$\mathbb{E} O'_t \leq \frac{k}{RO_n} \sum_i \int_{r_i}^{r_{i+1}} \psi(r_i) \, dx. \quad (4.3)$$

Since the item sizes are at most 1, we have $r_{i+1} \leq r_i + \frac{1}{k}$ and so $x - \frac{1}{k} \leq r_i$ for all $x \in [r_i, r_{i+1}]$. Thus, as the function ψ is nonincreasing, the right-hand side of (4.3) is at most

$$\frac{k}{RO_n} \sum_i \int_{r_i}^{r_{i+1}} \psi(x - 1/k) \, dx \leq \frac{k}{RO_n} \int_0^\infty \psi(x - 1/k) \, dx.$$

Finally, inspecting $\psi(x)$ we see that it takes value 1 for $x < 1$, takes value $\frac{2}{k}$ for $x \in [1, 50]$, and has exponential decay $\leq \frac{e^{-x}}{k}$ after that. Thus, it is easy to see that the integral on the right-hand side is at most $1 + O(\frac{1}{k})$ (see Lemma C.2). This concludes the proof. \square

However, what we actually need is to show that (4.2) (with O'_t 's) holds with good probability; for that we need concentration inequalities for the sums of the tentative occupations O'_t 's. The biggest problem is that the tentative selections induced by the LP are correlated in a non-trivial way. In particular, it is not clear whether they are negatively associated: for example, if the items up to time $t - 1$ are all “very good” the algorithm will not tentatively select at times $t, t + 1$, etc., indicating possibility of positive correlations on these times. Thus, the O'_t 's are also correlated and it is not clear how to apply standard concentration inequalities.

4.1 Concentration I: controlling the outer constraint

However, as the example above illustrates, we still have hopes of obtaining good *upper bounds* on the probability of multiple tentative selections. In fact, the probability of multiple selection of items I_{t_1}, \dots, I_{t_m} is at most the probability that the “worst” of these items is selected; more precisely:

Lemma 4.3. *Consider $m \leq \frac{\ln k}{4}$ random-order times $t_1, \dots, t_m \geq 8\ell(\Gamma + 1)$. Then*

$$\Pr \left(T_{t_1} = \dots = T_{t_m} = 1 \mid R_{t_1}, \dots, R_{t_m} \right) \leq \psi \left(\max_i R_{t_i} \right).$$

Proof. The inequality follows from the fact $\Pr(X_1 = \dots = X_m = 1 \mid E) \leq \min_i \Pr(X_i = 1 \mid E)$, Theorem 3.2, and $\min_i \psi(R_{t_i}) = \psi(\max_i R_{t_i})$ (by the monotonicity of ψ). \square

The main advantage of this bound is that the ranks R_{t_i} are “almost” independent (they would be independent if the input sequence was generated by sampling items *with* replacement). Moreover, this lemma allows us to upper bound products of tentative occupation $\prod_i O'_i$: for this product to be strictly positive, all these items have to be tentatively selected. In fact, one can prove such upper bound using a similar strategy

as in Lemma 4.2, with a main new element: a simple but general comparison for the expectation of a *non-negative* function under sampling with and without replacement (Lemma C.1), that allow us to work with a decoupled (independent) version $\bar{R}_{t_1}, \dots, \bar{R}_{t_m}$ of the ranks.

Lemma 4.4 (Control of products). *Fix a random-order time t . Consider a set of $m \leq \frac{\ln k}{4}$ distinct RO times t_1, \dots, t_m , all of which are at least $8\ell(\Gamma + 1)$ and less than t . Then there are constants $a_2, a_3 > 1$ such that $\mathbb{E} \left[\prod_{i \in [m]} O'_{t_i} \mid I_t \right] \leq \left(1 + \frac{a_2^m}{k}\right) \left(1 + \frac{4m^2}{RO_n}\right) \left(\frac{k}{RO_n}\right)^m \leq \left(a_3 \frac{k}{n}\right)^m$. In particular, choosing $a_2 = 500$ and $a_3 = 8a_2$ is sufficient.*

Proof. (We use the notation $r(I)$ instead of r_I to denote the rank of item I .) First notice that the product of tentative occupations has value $\prod_i W_{t_i}$ if all items at times t_1, \dots, t_m are tentatively selected, and 0 otherwise. Since conditioning on the items at these times fixes their weight, we have

$$\mathbb{E} \left[\prod_i O'_{t_i} \mid I_{t_1}, \dots, I_{t_m}, I_t \right] = \left(\prod_i w_{I_{t_i}} \right) \cdot \Pr(T_{t_1} = \dots = T_{t_m} = 1 \mid I_{t_1}, \dots, I_{t_m}, I_t).$$

From Lemma 4.3 the last term is at most $\psi(\max_i r(I_{t_i}))$. Employing this bound and taking expectation with respect to the items at times t_1, \dots, t_m we obtain

$$\mathbb{E} \left[\prod_i O'_{t_i} \mid I_t \right] \leq \mathbb{E} \left[\left(\prod_i w_{I_{t_i}} \right) \cdot \psi \left(\max_i r(I_{t_i}) \right) \mid I_t \right].$$

In order to continue upper bounding the right-hand side, it will be convenient to pass to the decoupled version of $(I_{t_i})_i$. That is, let $(\bar{I}_{t_i})_i$ be a sequence of independent random variables, each uniformly distributed in RO_n . Then using the comparison bound between sampling with and without replacement from Lemma C.1 (note that if t is not a random-order time we can just ignore the conditioning on I_t) the right-hand side is at most

$$\left(1 + \frac{4m^2}{RO_n}\right) \cdot \mathbb{E} \left[\left(\prod_i w_{\bar{I}_{t_i}} \right) \cdot \psi \left(\max_i r(\bar{I}_{t_i}) \right) \right] = \left(1 + \frac{4m^2}{RO_n}\right) \cdot \frac{1}{RO_n^m} \sum_{j_1, \dots, j_m} \underbrace{\left(\prod_i w_{j_i} \right) \psi(\max_i r_{j_i})}_A, \quad (4.4)$$

where in the last sum each index j_i ranges over RO_n . We need to better understand the term A .

Again from the definition of weighted rank, we have $w_j = k \cdot \int_{r_j}^{r_{j+1}} 1 \, dx$. So defining the m -dimensional box $B(j_1, \dots, j_m) := \prod_i [r_{j_i}, r_{j_{i+1}}]$, we have

$$\prod_i w_{j_i} = k^m \cdot \int_{B(j_1, \dots, j_m)} 1 \, dx.$$

Then the term A in (4.4) equals

$$A = k^m \cdot \int_{B(j_1, \dots, j_m)} \psi(\max_i r_{j_i}) \, (dx)^m.$$

Now notice the adjacent weighted ranks r_j, r_{j+1} differ by at most $\frac{1}{k}$ (recall that their definition has a factor $\frac{1}{k}$). So the sides of the box $B(j_1, \dots, j_m)$ are at most $\frac{1}{k}$, which implies that every point in $B(j_1, \dots, j_m) - \frac{1}{k}$

is pointwise at most the beginning of the box, namely $(r_{j_1}, \dots, r_{j_m})$. Since ψ is non-increasing, for all $x \in B(j_1, \dots, j_m) - \frac{1}{k}$ we then have $\psi(\max_i x_i) \geq \psi(\max_i r_{j_i})$. Thus, we can upper bound A as

$$A \leq k^m \cdot \int_{B(j_1, \dots, j_m)} \psi(\max_i x_i - 1/k) dx.$$

Adding this over all the boxes (which tile a subset of $[0, n]^m$, since the largest weighted rank r_{RO_n} is at most $\frac{RO_n}{k} \leq n$), we obtain

$$\text{LHS of (4.4)} \leq \left(1 + \frac{4m^2}{RO_n}\right) \cdot \left(\frac{k}{RO_n}\right)^m \cdot \int_{[0, n]^m} \psi(\max_i x_i - 1/k) dx.$$

Finally, as in the end of the proof of Lemma 4.2, using the fast decay of ψ it can be shown that the integral on the right-hand side is at most $1 + \frac{a_2^m}{k}$ for some constant $a_2 \geq 1$ (Lemma C.2). So putting the above bounds together gives

$$\mathbb{E} \left[\prod_i \mathcal{O}'_{t_i} \middle| I_t \right] \leq \left(1 + \frac{4m^2}{RO_n}\right) \cdot \left(\frac{k}{RO_n}\right)^m \cdot \left(1 + \frac{a_2^m}{k}\right),$$

proving the first inequality of the lemma.

To prove the second inequality, we use the following estimates:

- $1 + \frac{a_2^m}{k} \leq 2a_2^m$
- $1 + \frac{4m^2}{RO_n} \leq 2$, using Observation 3.5 and $m \leq \ln k$ and $k \ll n$
- $\left(\frac{k}{RO_n}\right)^m \leq \left(\frac{2k}{n}\right)^m$, using Observation 3.5.

This gives us

$$\left(1 + \frac{a_2^m}{k}\right) \left(1 + \frac{4m^2}{RO_n}\right) \left(\frac{k}{RO_n}\right)^m \leq \left(a_3 \frac{k}{n}\right)^m \leq 2a_2^m \cdot 2 \cdot \left(\frac{2k}{n}\right)^m \leq \left(8a_2 \frac{k}{n}\right)^m.$$

This concludes the proof of Lemma 4.4. □

Finally, such product estimates can be converted into raw moments/tail inequalities using reasonably standard estimates (e.g., Section 3.4 of [10]). In particular, these ideas together with a sharp Rosenthal-type inequality gives following multiplicative Chernoff bound for dependent random variables (proved in Appendix D).

Lemma 4.5. *Consider arbitrary random variables $X_1, \dots, X_n \in [0, 1]$, and an integer $m \geq 2$. Let $p \in [0, 1]$ be such that for all sets $A \subseteq [n]$ of size at most m we have $\mathbb{E} \prod_{i \in A} X_i \leq p^{|A|}$. If $m \leq np$, then $\mathbb{E} \left(\sum_{i \in [n]} X_i \right)^m \leq (2e^2 np)^m$. Thus, applying Markov's inequality to $(\sum_i X_i)^m$ we have $\Pr(\sum_i X_i \geq \alpha np) \leq \left(\frac{2e^2}{\alpha}\right)^m$ for all $\alpha > 0$.*

With this we can finally obtain the desired control of the outer constraint's occupation.

Lemma 4.6 (Control of outer constraints). *Consider a free time $t \geq 8\ell(\Gamma + 1)$, and let B be the last window in \mathbb{W}_{t-1} . Then $\Pr\left(\sum_{t' \in B} O'_{t'} > a_4 \ell \frac{k}{n} \mid I_t\right) \leq \frac{1}{k}$, where $a_4 \geq 2e^6 a_3$, and a_3 is the constant from Lemma 4.4.*

Proof. Since B may have size less than ℓ , let \bar{B} be the full window in \mathbb{W} that contains B , of size exactly ℓ ; it suffices to upper bound the probability that $\sum_{t' \in \bar{B}} O'_{t'} > a_4 \ell \frac{k}{n}$. Let $p = a_3 \frac{k}{n}$ and $m = \frac{\ln k}{4}$. Applying Lemma 4.4 and Lemma 4.5 to $(\sum_{t' \in \bar{B}} O'_{t'})|_{I_t}$, we have

$$\Pr\left(\sum_{t' \in \bar{B}} O'_{t'} > a_4 \ell \frac{k}{n} \mid I_t\right) \leq \left(\frac{2e^2}{a_4/a_3}\right)^m \leq \left(\frac{1}{e^4}\right)^{\frac{\ln k}{4}} = \frac{1}{k}.$$

□

4.2 Concentration II: control of main budget

In order to obtain Theorem 4.1 we need to show that the second part of (4.2) holds with reasonable probability *even when $t \approx n$* ; but since $\mathbb{E}O'_t \approx \frac{k}{n}$, the expected cumulative occupation by the end of the game $\mathbb{E}[\sum_{t'=1}^t O'_{t'}]$ is $\approx k$ for $t \approx n$, so we do not have much room. So unlike the previous section, we are interested in “medium deviations”, where the variance is the right quantity to look at. While Lemma 4.4 directly gives that the cumulative variance until time t is $\lesssim (\frac{tk}{n})^2$, we actually need an upper bound of $O(\frac{tk}{n})$, which is what one would expect from independent Bernoulli’s with success probability $\frac{k}{n}$. Since

$$\text{Var}(Z) = \mathbb{E}Z^2 - (\mathbb{E}Z)^2, \quad (4.5)$$

to obtain variance upper bounds we will obtain an upper bound on the second raw moment and a *lower bound on the expectation*.

In order to simplify obtaining the sharp lower bound on the expectation required, we instead work with $\bar{O}'_t = W_t \bar{T}_t$, where $\bar{T}_t := \max\{T_t, \mathbf{1}[R_t \leq 1]\}$, that is \bar{T}_t equals 1 if either $T_t = 1$ or the weighted rank R_t is at most 1. Notice that in every scenario \bar{O}'_t upper bounds O'_t , and thus it suffices to show that the second part of (4.2) holds for the \bar{O}'_t ’s. An important observation is that Lemma 4.4 still holds for the \bar{O}'_t ’s: this is because the properties of the O'_t ’s are only used through Lemma 4.3, which is a direct consequence of Theorem 3.2, and the latter holds for the \bar{O}'_t ’s because the upper bound “gives up” anyway when $R_t \leq 1$ (i.e., $\psi(R_t) = 1$ when $R_t \leq 1$). This then implies the following first step for using (4.5) to control the variance of the \bar{O}'_t ’s.

Lemma 4.7. *Fix a time t and, based on Lemma 4.4, let $S = RO_{t-1} \setminus [8\ell(\Gamma+1)]$. Then $\mathbb{E}\left[\left(\sum_{t' \in S} \bar{O}'_{t'}\right)^2 \mid I_t\right] \leq \left[\left(\frac{|S| \cdot k}{RO_n}\right)^2 + \frac{|S| \cdot k}{RO_n}\right] (1 + O(\frac{1}{k}))$.*

Proof. Using the assumptions that $RO_n \geq \frac{n}{2} \geq k$, notice that when $m = 2$ the first bound from Lemma 4.4 is $(\frac{k}{RO_n})^2 (1 + O(\frac{1}{k}))$, and similarly for $m = 1$. The result follows by applying this lemma to the right-hand side of

$$\mathbb{E}\left[\left(\sum_{t' \in S} \bar{O}'_{t'}\right)^2 \mid I_t\right] \leq \sum_{t', t'' \in S, t' \neq t''} \mathbb{E}[\bar{O}'_{t'} \bar{O}'_{t''} \mid I_t] + \sum_{t' \in S} \mathbb{E}[\bar{O}'_{t'}^2 \mid I_t]. \quad \square$$

But the good thing is that by passing from O'_t to \bar{O}'_t we easily get a strong lower bound on the expectation.

Lemma 4.8. *Fix a time t and let $S = RO_{t-1} \setminus [8\ell(\Gamma + 1)]$. Then: $\mathbb{E} \left[\sum_{t' \in S} \bar{O}'_{t'} \mid I_t \right] \geq \frac{|S| \cdot k}{RO_n} \left(1 - \frac{2}{k}\right)$.*

Proof. It suffices to show $\mathbb{E} \bar{O}'_{t'} \geq \frac{k-2}{RO_n}$ for all random order times t' . Let i^* be the largest item index with weighted rank r_{i^*} at most 1. Since item sizes are at most 1, this implies that $\sum_{i < i^*} w_i \geq k - 1$, and further $\sum_{i < i^*, i \neq I_t} \geq k - 2$. Then since $\mathbb{E}[\bar{T}_{t'} \mid I_{t'} = i, I_t] = 1$ for all $i \leq i^*$, $\mathbb{E}[\bar{O}'_{t'} \mid I_t] = \mathbb{E}[W_{t'} \bar{T}_{t'} \mid I_t] \geq \sum_{i < i^*, i \neq I_t} \mathbb{E}[W_{t'} \bar{T}_{t'} \mid I_{t'} = i, I_t] \Pr(I_t = i \mid I_t) = \frac{1}{RO_n - 1} \sum_{i < i^*, i \neq I_t} w_i \geq \frac{k-2}{RO_n}$. \square

Putting these bounds together in (4.5) (and using Observation 3.5) we can control the variance of the \bar{O}'_t 's.

Lemma 4.9. *Fix a time t and let $S = RO_{t-1} \setminus [8\ell(\Gamma + 1)]$. Then: $\text{Var} \left[\sum_{i \in S} \bar{O}'_i \mid I_t \right] \leq O\left(\frac{tk}{n}\right)$.*

This variance control is enough to upper bound the probability that tentative solution violates the main budget at any point in time.

Lemma 4.10 (Control of main budget). *For every random-order time t , the probability we are blocked by the main budget can be upper bounded as $\Pr \left[\sum_{t' < t} O_{t'} > k - 1 \mid I_t \right] \leq \frac{O(1)}{k \left(1 - \frac{t}{n} - O\left(\frac{\Gamma \ln k}{k}\right)\right)^2}$.*

Proof. Let $S = RO_{t-1} \setminus [8\ell(\Gamma + 1)]$. Then

$$\Pr \left[\sum_{t' < t} O_{t'} > k - 1 \mid I_t \right] \leq \Pr \left[\sum_{t' \in [t-1] \setminus S} O_{t'} + \sum_{t' \in S} O_{t'} > k - 1 \mid I_t \right]. \quad (4.6)$$

Since $[t-1] \setminus S$ is composed of some adversarial times plus the interval $[8\ell(\Gamma + 1)]$, it can be covered with $\Gamma + 8(\Gamma + 1) \leq 10\Gamma$ intervals in \mathbb{W}_{t-1} . Since the outer constraints control the possible occupation over each of these windows, in every scenario the first sum in (4.6) is at most

$$\sum_{t' \in [t-1] \setminus S} O_{t'} \leq 10\Gamma \cdot a_4 \ell \frac{k}{n} = 10a_4 \Gamma \ln k.$$

Then we can upper bound (4.6) as

$$\Pr \left[\sum_{t' < t} O_{t'} > k - 1 \mid I_t \right] \leq \Pr \left[\sum_{t' \in S} O_{t'} > k - 10a_4 \Gamma \ln k \mid I_t \right] \leq \Pr \left[\sum_{t' \in S} \bar{O}'_{t'} > k - 10a_4 \Gamma \ln k \mid I_t \right].$$

Now we apply Chebyshev's Inequality. Using Lemma 4.4 (and Observation 3.5) we can bound the expected value $\mu := \sum_{t' \in S} \mathbb{E} \bar{O}'_{t'} \leq \frac{tk}{n} \left(1 + O\left(\frac{\Gamma \ell}{n}\right)\right)$, and so the gap we have is

$$gap := k - 10a_4 \Gamma \ln k - \frac{tk}{n} \left(1 + O\left(\frac{\Gamma \ell}{n}\right)\right) \geq k \left(1 - \frac{t}{n} - O\left(\frac{\Gamma \ln k}{k}\right)\right).$$

Thus, using Chebyshev's inequality and Lemma 4.9 to control the variance, we have

$$\Pr \left[\sum_{t' \in S} \bar{O}'_{t'} > k - 10a_4 \Gamma \ln k \mid I_t \right] = \Pr \left[\sum_{t' \in S} \bar{O}'_{t'} > \mu + gap \mid I_t \right] \leq \frac{O\left(\frac{tk}{n}\right)}{gap^2} \leq \frac{O(1)}{k \left(1 - \frac{t}{n} - O\left(\frac{\Gamma \ln k}{k}\right)\right)^2}.$$

This concludes the proof. \square

Taking a union bound over Lemma 4.6 and Lemma 4.10 proves Theorem 4.1.

5 Lower bounding the value obtained

Recall that $X_t^{alg} = T_t F_t$, i.e., the item is permanently selected exactly when it is tentatively selected and it fits the budgets, and that V_t is the value of the item at time t . The following is then our main lower bound on the value obtained by the algorithm.

Theorem 5.1 (Value lower bound). *Consider a free time $t \geq 1, 212\Gamma\ell$. Then*

$$\mathbb{E}[V_t T_t F_t] \geq \left(c_t - \varepsilon_t - p_t - \frac{2}{k} \right) \frac{\text{OPT}_{RO}}{RO_n},$$

where p_t is the bound from Theorem 4.1 and $\varepsilon_t = (a_1 + 3) \frac{\Gamma\ell}{t} + \sqrt{10 \ln k} \sqrt{\frac{2n}{tk}}$.

The next lemma says that if up to time t there are not many items better than the item I_t , then this item is fully tentatively picked.

Lemma 5.2. *Consider a free time t , and a fixed scenario where the following hold:*

- *Up to time t , the total size of items in free times strictly better than I_t is strictly less than $c_t \frac{tk}{n} - \frac{a_1 \Gamma \ell k}{n}$*
- *In the last window $B_{last} \in \mathbb{W}_t$ (which only has free times), the total size of items strictly better than I_t is strictly less than $\frac{a_1 \ell k}{n}$.*

Then any optimal solution X^ of LP_t sets $X_t = 1$, i.e., it fully tentatively picks item I_t .*

Proof. Let X^* be an optimal solution of LP_t , and suppose by contradiction that $X_t^* < 1$. Again we use the notation $WX^*(S) := \sum_{t' \in S} W_{t'} X_{t'}^*$.

Case 1: $WX^*(B_{last}) < \frac{a_1 \ell k}{n}$, i.e., the inner constraint for the last window is not tight. If $WX^*([t]) < c_t \frac{tk}{n}$, inner main budget is also not tight, we could just increase X_t^* to obtain a strictly better solution, reaching a contradiction. So assume the main budget is tight, $WX^*([t]) = c_t \frac{tk}{n}$.

Since the non-free times are covered by Γ windows and we have the inner constraints in the LP, the solution X^* picks up mass at most $A = \frac{a_1 \Gamma \ell k}{n}$ of items in non-free times; thus to fill up the main budget, the solution picks up at mass least $c_t \frac{tk}{n} - A$ in free times. By assumption this implies that it fractionally picks an item in a free time that is worse than I_t , i.e., there is $t' \in \text{FREE}_{t-1}$ such that $\frac{V_{t'}}{W_{t'}} < \frac{V_t}{W_t}$ and $X_{t'}^* > 0$. But then we can increase X_t^* by $\frac{\varepsilon}{W_t}$ and decrease $X_{t'}^*$ by $\frac{\varepsilon}{W_{t'}}$ to obtain a feasible solution (using the fact we are in Case 1) with strictly better value, reaching a contradiction.

Case 2: $WX^*(B_{last}) = \frac{a_1 \ell k}{n}$. By assumption X^* fractionally picks a (free-time) item $I_{t'}$ with $t' \in B_{last}$ that is worse than I_t . Then we can swap a bit of these items exactly as in the previous case to obtain a feasible solution with strictly better value, a contradiction (notice this swap preserves the occupation of the inner main budget and of the inner constraint for B_{last} , everywhere else nothing changes). This concludes the proof. \square

Moreover, from concentration, with high probability the conditions in the above lemma hold whenever I_t has low rank; the proof is deferred to Appendix E.

Lemma 5.3. *For any free time $t \geq 1, 212\Gamma\ell$ we have that the probability of fully tentatively selecting item I_t satisfies: $\Pr \left(X_t^t = 1 \mid R_t \leq c_t - \varepsilon_t \right) \geq 1 - \frac{1}{k}$.*

With this lower bound on the probability of selection by the algorithm conditioned on the item being “good” (low rank), we can proceed with the proof of Theorem 5.1.

Proof of Theorem 5.1. Introducing the conditioning on the item at time t and then using the non-negativity of V_t, T_t , and F_t , we have

$$\begin{aligned}\mathbb{E}[V_t T_t F_t] &= \mathbb{E}_{I_t} [V_t \mathbb{E}[T_t F_t | I_t]] \\ &= \mathbb{E}_{I_t} [V_t \mathbb{E}[T_t F_t | I_t] | R_t \leq c_t - \varepsilon_t] \Pr(R_t \leq c_t - \varepsilon_t) \\ &\quad + \mathbb{E}_{I_t} [V_t \mathbb{E}[T_t F_t | I_t] | R_t > c_t - \varepsilon_t] \Pr(R_t > c_t - \varepsilon_t) \\ &\geq \mathbb{E}_{I_t} [V_t \mathbb{E}[T_t F_t | I_t] | R_t \leq c_t - \varepsilon_t] \Pr(R_t \leq c_t - \varepsilon_t).\end{aligned}\tag{5.7}$$

Notice that

$$\mathbb{E}[T_t F_t | I_t] \geq \Pr(T_t = 1 \text{ and } F_t = 1 | I_t) \geq 1 - \Pr(T_t \neq 1 | I_t) - \Pr(F_t = 0 | I_t),$$

the last inequality following from a union bound. Whenever I_t is such that its rank satisfies $R_t \leq c_t - \varepsilon_t$, we can apply Lemma 5.3 and the definition of p_t to lower bound the right-hand side by $1 - \frac{1}{k} - p_t$. Plugging this in (5.7) we get

$$\mathbb{E}[V_t T_t F_t] \geq \left(1 - p_t - \frac{1}{k}\right) \mathbb{E}_{I_t} [V_t | R_t \leq c_t - \varepsilon_t] \Pr(R_t \leq c_t - \varepsilon_t).\tag{5.8}$$

Now let $S = \{i : r_i \leq c_t - \varepsilon_t\}$. The last two terms of (5.8) are just adding over the value of items in S multiplied by $\Pr(I_t = i) = \frac{1}{RO_n}$, namely

$$\mathbb{E}[V_t T_t F_t] \geq \left(1 - p_t - \frac{1}{k}\right) \frac{1}{RO_n} \sum_{i \in S} v_i.\tag{5.9}$$

Equivalently S can be constructed by picking the largest prefix of best random-order items that has total size at most $k(c_t - \varepsilon_t) \approx k$. For this reason and due to the fact that they take up almost the whole knapsack, this should be close to the optimal solution to our knapsack problem over the random-order items; more precisely, we claim that

$$\sum_{i \in S} v_i \geq (c_t - \varepsilon_t - 1/k) \text{OPT}_{RO}.\tag{5.10}$$

To see that, let x^* be the optimal offline solution to our knapsack problem over the random-order items only, where x_i^* indicates the fraction of the random-order item i picked. Let x be the indicator of the set S , i.e., $x_i = 1$ iff $i \in S$. Since x^* is given by the greedy procedure that scans items in order of value density $\frac{v_i}{w_i}$, as in the construction of S , we have $x_i^* = 1$ whenever $x_i = 1$. Thus, by introducing additional terms, we have for $i \in S$

$$v_i = v_i x_i = \left(\frac{v_i}{w_i} x_i^*\right) x_i (w_i x_i^*).$$

Adding over all $i \in S$ and applying Chebyshev’s Sum Inequality (Lemma A.2) with $a_i = \left(\frac{v_i}{w_i} x_i^*\right)$, $b_i = x_i$, and $p_i = w_i x_i^*$, we get

$$\sum_{i \in S} v_i = \sum_i v_i x_i \geq \left(\sum_i \frac{v_i}{w_i} x_i^* w_i\right) \left(\sum_i x_i w_i\right) / \left(\sum_i w_i x_i^*\right).$$

The first term in the RHS is the value of x^* , which by definition is OPT_{RO} . The second term is the total size of x , which by the maximality in its definition is $k(c_t - \varepsilon_t) - 1$ (recall that all items have size at most 1). Finally, the last term is the total size of x^* , which by optimality equals k . This proves (5.10).

Employing this bound to inequality (5.9) and using $(1-a)(1-b) \geq (1-a-b)$, valid for all non-negative a, b , concludes the proof of the theorem. \square

6 Wrapping up: finishing the proof of Theorem 2.1

To finish the proof of the guarantee of the algorithm, we just need to add the lower bound on the value obtained in each time step given by Theorem 5.1 over all free times except the ones very early or very late in the sequence. More precisely, let $t_0 = 1, 212\Gamma\ell$ and $\gamma = 1 - (a_5 + 1)\frac{\Gamma\ell}{n}$, and define $T = \{t \in \text{FREE}_n : t_0 \leq t \leq \gamma n\}$. For $t \notin T$, we use the trivial bound $V_t \geq 0$. For the other time steps we use Theorem 5.1. Together with the fact $RO_n \leq n$, we get

$$\mathbb{E} \left[\sum_{t \in RO} V_t X_t^{\text{alg}} \right] \geq \sum_{t \in T} \mathbb{E}[V_t T_t F_t] \geq \left(\sum_{t \in T} c_t - \sum_{t \in T} \varepsilon_t - \sum_{t \in T} p_t - \frac{2n}{k} \right) \frac{\text{OPT}_{RO}}{n}. \quad (6.11)$$

Just using some arithmetic we bound each of the remaining sums:

- $\sum_{t \in T} c_t = |T| - \sum_{t \in T} \frac{4\Gamma\ell}{t} \geq (n - O(\Gamma\ell)) - \int_{t_0-1}^n \frac{4\Gamma\ell}{t} dt = n - O(\Gamma\ell \ln \frac{n}{\Gamma\ell})$.
- $\sum_{t \in T} p_t \leq O\left(\frac{1}{k}\right) \cdot \int_0^{\gamma n} \frac{1}{\left(1 - \frac{t}{n} - a_5 \frac{\Gamma \ln k}{k}\right)^2} dt = O\left(\frac{n}{k}\right) \int_0^\gamma \frac{1}{(a-x)^2} dx$, where in the last step we set $a = 1 - a_5 \frac{\Gamma \ln k}{k}$ and use change of variables $x = \frac{t}{n}$. The remaining integral equals $\frac{1}{a-x} \Big|_0^\gamma \leq \frac{1}{a-\gamma}$. By our setting of γ we have $a - \gamma = \frac{\Gamma\ell}{n}$, so we obtain $\sum_{t \in T} p_t \leq O\left(\frac{n}{\Gamma\ell}\right)$.
- $\sum_{t \in T} \varepsilon_t \leq \int_{t_0-1}^n (a_1 + 3) \frac{\Gamma\ell}{t} dt + \int_{t_0-1}^n \sqrt{\frac{20n \ln k}{kt}} dt \leq O(\Gamma\ell \ln \frac{n}{\Gamma\ell}) + O\left(\frac{n\sqrt{\ln k}}{\sqrt{k}}\right)$.

Using these bounds on (6.11) and using the assumption that $\Gamma \geq \sqrt{k}$ concludes the proof of the theorem.

7 Conclusions

In this paper, we give a natural algorithm for the knapsack secretary problem, which we show to be robust against bursts of adversarial items. Our analysis is quite robust and possibly applies to other models mixing aspects of stochastic and adversarial arrivals.

A natural follow-up question is how our results could generalize to other settings. In particular, it would be interesting to extend our algorithm and analysis to packing LPs. The difficulty in using our technique is that there is no natural notion similar to the weighted rank for this setting.

It would also be interesting to better understand the limitations and trade-offs in this and similar models. For example, what regimes of parameter allow constant-competitive or $(1 - \varepsilon)$ -competitive algorithms?

References

- [1] Shipra Agrawal and Nikhil R. Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*,

- San Diego, CA, USA, January 4-6, 2015, pages 1405–1424, 2015. doi: 10.1137/1.9781611973730.93. URL <http://dx.doi.org/10.1137/1.9781611973730.93>.
- [2] Shipra Agrawal, Zizhuo Wang, and Yinyu Ye. A dynamic near-optimal algorithm for online linear programming. *Operations Research*, 62(4):876–890, 2014. doi: 10.1287/opre.2014.1289.
- [3] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *APPROX-RANDOM*, 2007.
- [4] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '07*, pages 434–443, 2007. ISBN 978-0-898716-24-5.
- [5] Domagoj Bradac, Anupam Gupta, Sahil Singla, and Goran Zuzic. Robust algorithms for the secretary problem. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPICs*, pages 32:1–32:26. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020. doi: 10.4230/LIPICs.ITCS.2020.32. URL <https://doi.org/10.4230/LIPICs.ITCS.2020.32>.
- [6] Sourav Chakraborty and Oded Lachish. *Improved Competitive Ratio for the Matroid Secretary Problem*, pages 1702–1712. doi: 10.1137/1.9781611973099.135. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611973099.135>.
- [7] G.R. Dattatreya. *Performance Analysis of Queuing and Computer Networks (Chapman & Hall/Crc Computer & Information Science Series)*. Chapman & Hall/CRC, 2008. ISBN 1584889861, 9781584889861.
- [8] Nikhil R. Devanur and Thomas P. Hayes. The adwords problem: online keyword matching with budgeted bidders under random permutations. In *EC*, 2009.
- [9] Qiming Diao, Jing Jiang, Feida Zhu, and Ee-Peng Lim. Finding bursty topics from microblogs. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Long Papers - Volume 1*, ACL '12, pages 536–544, 2012.
- [10] Devdatt Dubhashi and Alessandro Panconesi. *Concentration of Measure for the Analysis of Randomized Algorithms*. Cambridge University Press, New York, NY, USA, 1st edition, 2009. ISBN 0521884276, 9780521884273.
- [11] Hossein Esfandiari, Nitish Korula, and Vahab Mirrokni. Online allocation with traffic spikes: Mixing adversarial and stochastic models. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation, EC '15*, pages 169–186, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3410-5. doi: 10.1145/2764468.2764536. URL <http://doi.acm.org/10.1145/2764468.2764536>.
- [12] Jon Feldman, Monika Henzinger, Nitish Korula, Vahab S. Mirrokni, and Clifford Stein. Online stochastic packing applied to display ad allocation. In *ESA*, 2010.
- [13] Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $o(\log \log(\text{rank}))$ -competitive algorithm for the matroid secretary problem. In *Proceedings of the Twenty-sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '15*, pages 1189–1201, 2015.

- [14] Moran Feldman, Ola Svensson, and Rico Zenklusen. A framework for the secretary problem on the intersection of matroids. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '18, pages 735–752, 2018. ISBN 978-1-6119-7503-1.
- [15] Anupam Gupta and Marco Molinaro. How the experts algorithm can help solve lps online. *Mathematics of Operations Research*, 41(4):1404–1431, 2016. doi: 10.1287/moor.2016.0782. URL <https://doi.org/10.1287/moor.2016.0782>.
- [16] Anupam Gupta, Tomer Koren, and Kunal Talwar. Better algorithms for stochastic bandits with adversarial corruptions. In *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, pages 1562–1578, 2019. URL <http://proceedings.mlr.press/v99/gupta19a.html>.
- [17] Thomas Kesselheim, Andreas Tönnis, Klaus Radke, and Berthold Vöcking. Primal beats dual on online packing lps in the random-order model. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, STOC '14, pages 303–312, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2710-7. doi: 10.1145/2591796.2591810. URL <http://doi.acm.org/10.1145/2591796.2591810>.
- [18] Thomas Kesselheim, Robert D. Kleinberg, and Rad Niazadeh. Secretary problems with non-uniform arrival order. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, pages 879–888, 2015. doi: 10.1145/2746539.2746602. URL <http://doi.acm.org/10.1145/2746539.2746602>.
- [19] Jon Kleinberg. Bursty and hierarchical structure in streams. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, pages 91–101, New York, NY, USA, 2002. ACM. ISBN 1-58113-567-X. doi: 10.1145/775047.775061. URL <http://doi.acm.org/10.1145/775047.775061>.
- [20] Jon Kleinberg, Yuval Rabani, and Éva Tardos. Allocating bandwidth for bursty connections. In *Proceedings of the Twenty-ninth Annual ACM Symposium on Theory of Computing*, STOC '97, pages 664–673, New York, NY, USA, 1997. ACM. ISBN 0-89791-888-6. doi: 10.1145/258533.258661. URL <http://doi.acm.org/10.1145/258533.258661>.
- [21] Robert Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *SODA*, 2005. ISBN 0-89871-585-7.
- [22] Nitish Korula, Vahab Mirrokni, and Morteza Zadimoghaddam. Online submodular welfare maximization: Greedy beats 1/2 in random order. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing*, STOC '15, pages 889–898, 2015. ISBN 978-1-4503-3536-2. doi: 10.1145/2746539.2746626.
- [23] O. Lachish. $O(\log \log \text{rank})$ competitive ratio for the matroid secretary problem. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 326–335, Oct 2014. doi: 10.1109/FOCS.2014.42.
- [24] Thodoris Lykouris, Vahab Mirrokni, and Renato Paes Leme. Stochastic bandits robust to adversarial corruptions. In *STOC 2018*, 2018.

- [25] A. Meyerson. Online facility location. In *Proceedings of the 42Nd IEEE Symposium on Foundations of Computer Science, FOCS '01*, pages 426–, Washington, DC, USA, 2001. IEEE Computer Society. ISBN 0-7695-1390-5. URL <http://dl.acm.org/citation.cfm?id=874063.875567>.
- [26] Vahab S. Mirrokni, Shayan Oveis Gharan, and Morteza Zadimoghaddam. Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '12*, pages 1690–1701, 2012.
- [27] D.S. Mitrinovic, J. Pecaric, and A.M. Fink. *Classical and New Inequalities in Analysis. Mathematics and its Applications*. Springer Netherlands, 1992. ISBN 9780792320647. URL <https://books.google.com.br/books?id=VkfIHKzP5ZEC>.
- [28] Marco Molinaro. Online and random-order load balancing simultaneously. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '17*, pages 1638–1650, 2017.
- [29] Marco Molinaro and R. Ravi. Geometry of online packing linear programs. In *ICALP*. 2012.
- [30] Victor de la Peña and Evarist Giné. *Decoupling: From Dependence to Independence*. Springer-Verlag, New York, NY, USA, 1999. ISBN 978-0-387-98616-6.
- [31] Yevgeny Seldin and Aleksandrs Slivkins. One practical algorithm for both stochastic and adversarial bandits. In *Proceedings of the 31th International Conference on Machine Learning, ICML 2014, Beijing, China, 21-26 June 2014*, pages 1287–1295, 2014. URL <http://proceedings.mlr.press/v32/seldinb14.html>.
- [32] Julian Zimmert and Yevgeny Seldin. An optimal algorithm for stochastic and adversarial bandits. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, pages 467–475, 2019. URL <http://proceedings.mlr.press/v89/zimmert19a.html>.

A Required inequalities

We will need standard concentration inequalities for sampling without replacement of Bernstein-type; the following can be found, for example on Corollary 2.3 of [15].

Lemma A.1. *Let $U = \{u_1, u_2, \dots, u_n\}$ be a set of real numbers in the interval $[0, 1]$. Let Y_1, Y_2, \dots, Y_s be a sequence of draws from U without replacement, and let $\mu = \sum_i \mathbb{E}Y_i$. Then for every $\tau > 0$,*

$$\Pr\left(\left|\sum_i Y_i - \mu\right| \geq \tau\right) \leq 2 \exp\left(-\frac{\tau^2}{4\mu + \tau}\right).$$

We also need a discrete version of the classical Chebyshev's Sum Inequality, which can be found, for example, in Chapter 9 of [27].

Lemma A.2 (Chebyshev's Sum Inequality). *Let a_1, a_2, \dots, a_n and b_1, b_2, \dots, b_n be non-increasing sequences. Then for any non-negative sequence p_1, p_2, \dots, p_n*

$$\sum_i a_i b_i p_i \geq \left(\sum_i a_i p_i\right) \left(\sum_i b_i p_i\right) / \left(\sum_i p_i\right).$$

B Proof of Lemma 3.3

Consider an optimal solution X^* to LP_t and assume by contradiction that $X_t^* > 0$. Let B_{last} denote the last interval in \mathbb{W}_t , namely the one containing t . If there is a time $t' \in B_{\text{last}}$ with $\bar{X}_{t'} > 0$ and $X_{t'}^* < 1$, then we can change the solution X^* by reducing its t -th coordinate by $\frac{\varepsilon}{\bar{W}_t}$ and increasing its t' -th coordinate by $\frac{\varepsilon}{\bar{W}_{t'}}$ to obtain a feasible solution (just need to check main budget and inner constraint for B_{last}) with better value, contradicting the optimality of X^* .

So suppose that for all $t' \in B_{\text{last}}$ with $\bar{X}_{t'} > 0$ we have $X_{t'}^* = 1$; this implies that $X_{t'}^* \geq \bar{X}_{t'}$ for all $t' \in B_{\text{last}}$. Also, by assumption, we have the strict inequality $X_t^* > \bar{X}_t = 0$. Thus $X^*(B_{\text{last}}) > \bar{X}(B_{\text{last}})$ and hence $WX^*(B_{\text{last}}) > W\bar{X}(B_{\text{last}})$, where for any set of times $S \subseteq [n]$ we define $X^*(S) := \sum_{t \in S} X_t$ and $WX^*(S) := \sum_{t \in S} W_t X_t^*$, and similarly for \bar{X} . Also, under our running assumptions that the sum of all item sizes is at least k and that there are no items of value 0, the optimal solution saturates the main budget: $WX^*([t]) = c_t \frac{tk}{n}$. Since by assumption the same holds for \bar{X} , we have $WX^*([t]) = W\bar{X}([t])$. Thus, as the intervals in \mathbb{W}_t partition $[t]$, we have $\sum_{B \in \mathbb{W}_t \setminus \{B_{\text{last}}\}} W\bar{X}(B) > \sum_{B \in \mathbb{W}_t \setminus \{B_{\text{last}}\}} WX^*(B)$; so there is an interval $B \in \mathbb{W}_t$ with $W\bar{X}(B) > WX^*(B)$. One consequence of this is that there is $t' \in B$ with $1 \geq \bar{X}_{t'} > X_{t'}^* \geq 0$; so t' is a strictly better item than t and $X_{t'}^*$ is not at its upper bound. Another consequence is that, since \bar{X} satisfies the inner constraints, X^* is strictly feasible for the inner constraint relative to B . Thus, we can again increase $X_{t'}^*$ and decrease X_t^* to obtain a feasible solution to LP_t with higher and contradict the optimality of X^* . This concludes the proof.

C Lemmas for the proof of Lemma 4.4

We start with the following comparison of integrating a non-negative function over a series sampled with and without replacement.

Lemma C.1. Consider any set S of size n . Let X_1, \dots, X_m, X be sampled without replacement S , and let X'_1, \dots, X'_m be sampled with replacement from S . Then for any non-negative function $f : S^m \rightarrow \mathbb{R}_+$

$$\mathbb{E}[f(X_1, X_2, \dots, X_m) \mid X] \leq \left(1 + \frac{m}{n-m}\right)^m \mathbb{E}f(X'_1, X'_2, \dots, X'_m).$$

Moreover, if $m \ll n$ (having $\frac{m^2}{n-m} \leq 1$ and $m \leq n/2$ suffices), the multiplicative factor in the right-hand side is at most $(1 + \frac{4m^2}{n})$.

Proof. For any $x \in S$ we expand the conditional expectation:

$$\begin{aligned} \mathbb{E}[f(X_1, X_2, \dots, X_m) \mid X = x] &= \frac{1}{(n-1)(n-2)\dots(n-1-(m-1))} \sum_{i_1, i_2, \dots, i_m \in S \setminus x, \text{ distinct}} f(i_1, \dots, i_m) \\ &\stackrel{f \geq 0}{\leq} \frac{n^m}{(n-1)(n-2)\dots(n-1-(m-1))} \cdot \frac{1}{n^m} \sum_{i_1, i_2, \dots, i_m \in S} f(i_1, \dots, i_m) \\ &= \frac{n^m}{(n-1)(n-2)\dots(n-1-(m-1))} \mathbb{E}f(X'_1, \dots, X'_m). \end{aligned}$$

The first factor in the right-hand side is at most $(\frac{n}{n-1-(m-1)})^m = (1 + \frac{m}{n-m})^m$. This gives the first part of the result.

For the second part, use $1 + x \leq e^x$ (which holds for all x) we obtain that this factor is at most $e^{\frac{m^2}{n-m}}$, and then using $e^x \leq 1 + 2x$ (which holds for $x \in [0, 1]$) and we assumption $\frac{m^2}{n-m} \leq 1$ we further upper bound it by $1 + \frac{2m^2}{n-m}$; finally using the assumption $m \leq n/2$ we reach the final upper bound of $1 + \frac{4m^2}{n}$, concluding the proof. \square

Lemma C.2. If $m \leq \frac{\ln k}{4}$, there is a constant $a_2 > 0$ such that

$$\int_{[0, n]^m} \psi(\max\{x_i : i \in [m]\} - 1/k) dx \leq 1 + \frac{a_2^m}{k}.$$

Proof. Let U_1, \dots, U_m be independent random variables uniformly distributed in $[0, n]$, and let $Z = \max_i U_i$. The integral we want to upper bound equals

$$n^m \cdot \mathbb{E}[\psi(Z - 1/k)] = n^m \cdot \int_0^n \psi(z - 1/k) pdf_Z(z) dz. \quad (\text{C.12})$$

Moreover, we know precisely the distribution of Z (obtained by differentiating $\Pr(Z \leq z) = \prod_i \Pr(U_i \leq z) = (\frac{z}{n})^m$):

$$pdf_Z(z) = \frac{m}{n} \left(\frac{z}{n}\right)^{m-1} = \frac{m z^{m-1}}{n^m}.$$

So breaking up into the different cases in the definition of ψ , we can upper bound the integral in (C.12) as

$$\int_0^n \psi(z - 1/k) pdf_Z(z) dz \leq \underbrace{\int_0^{1+1/k} pdf_Z(z) dz}_{int_1} + \underbrace{\int_{1+1/k}^C \frac{2}{k} pdf_Z(z) dz}_{int_2} + \underbrace{\int_C^\infty 4k e^{-\frac{z}{20} \ln k} pdf_Z(z) dz}_{int_3},$$

where $C := 120$ is a sufficiently large constant. The first integral on the right-hand side can be upper bounded

$$\text{int}_1 \leq \frac{1}{n^m} \left(1 + \frac{1}{k}\right)^m \leq \frac{1}{n^m} e^{\frac{m}{k}} \leq \frac{1}{n^m} \left(1 + \frac{2m}{k}\right),$$

where the first inequality uses $1 + x \leq e^x$ (valid for all x), and the second uses $e^x \leq 1 + 2x$ (valid for $x \in [0, 1]$) and the fact $m \leq \ln k$. For the second integral we have directly $\text{int}_2 \leq \frac{2 \cdot C^m}{n^m k}$.

For the last integral, since $m \leq \frac{\ln k}{4}$ and $C = 120$ is a sufficiently large constant, for all $z \geq C$ we have $z^{m-1} \leq e^{\frac{z}{40} \ln k}$, and so

$$\text{int}_3 \leq \frac{4km}{n^m} \int_C^\infty e^{-\frac{z}{40} \ln k} dz \leq \frac{km \ln k}{10n^m} e^{-\frac{C \ln k}{40}} \leq \frac{k \ln^2 k}{40n^m} k^{-\frac{C}{40}} \leq \frac{1}{n^m k},$$

where the last inequality uses again that C is a sufficiently large constant.

Putting these bounds together we obtain

$$\int_0^n \psi(z - 1/k) \text{pdf}_Z(z) dz \leq \frac{1}{n^m} \left(1 + \frac{2m}{k} + \frac{2 \cdot C^m}{k} + \frac{1}{k}\right).$$

Plugging this bound on (C.12) concludes the proof. \square

D Proof of Lemma 4.5

To prove the first inequality: Since the expression $\mathbb{E} \left(\sum_{i \in [n]} X_i \right)^m$ is a positive combination of the expectation of monomials of the form $\mathbb{E} \prod_{i \in A} X_i$ with $|A| \leq m$ (and the same holds for the X'_i 's), it suffices to have $\mathbb{E} \prod_{i \in A} X_i \leq \mathbb{E} \prod_{i \in A} X'_i$ for all A of size at most m ; this holds by our assumption on the X_i 's and the fact the right-hand side equals $\Pr(\bigwedge_{i \in A} (X'_i = 1)) = p^{|A|}$.

For the second inequality, we employ a sharp Rosenthal-type inequality; more precisely, we can use Theorem 1.5.2 and Lemma 1.5.8 (with $p = m$ and $c = \frac{1}{m}$) of [30] to obtain

$$\left(\mathbb{E} \left(\sum_i X'_i \right)^m \right)^{1/m} \leq 2e \cdot \max \left\{ e \cdot np, (1+m) \left(\frac{np}{m} \right)^{1/m} \right\}.$$

Now we claim that when $2 \leq m \leq np$, the maximum is achieved on the first term: In this range, the second term is at most $e \cdot m(np/m)^{1/m} \leq e \cdot m(np/m) = e \cdot np$. This concludes the proof.

E Proof of Lemma 5.3

Let ω denote the conditioning on an I_t satisfying $R_t \leq c_t - \varepsilon_t$. Let B_{last} denote the last window in \mathbb{W}_{t-1} . Let E be the event that there is total size at least $c_t \frac{tk}{n} - \frac{\Gamma a_1 \ell k}{n}$ of items better than I_t in the free times FREE_{t-1} , and let F be the event that there is total size at least $\frac{a_1 \ell k}{n}$ of items better than I_t in B_{last} . Since X^t is an optimal solution for LP_t , from Lemma 5.2 it suffices to show that $\Pr(E \text{ or } F \mid \omega) \leq \frac{1}{k}$.

We start with event F . By definition, under ω there is total size $kR_t \leq k(c_t - \varepsilon_t) \leq k$ of random-order items better than I_t , so the expected (conditioned on ω) mass of such items in the last window B_{last} is $\frac{|B_{\text{last}}|}{RO_{n-1}} kR_t \leq \frac{3\ell k}{n}$ (see Observation 3.5). Since again this quantity can be expressed as the random-order sum

$\sum_{t' \in B_{\text{last}}} \mathbf{1}(I_{t'} < I_t) \cdot W_{t'}$, we can apply the Bernstein's-type inequality of Lemma A.1 (with $\tau = 598 \frac{\ell k}{n}$) to obtain the following upper bound on $\Pr(F \mid \omega)$:

$$\Pr \left(\text{mass of RO items in } B_{\text{last}} \text{ better than } I_t \geq \frac{3\ell k}{n} + \frac{598\ell k}{n} \mid \omega \right) \leq 2 \exp \left(-\frac{(4 \cdot \frac{3\ell k}{n})^2}{8 \cdot \frac{3\ell k}{n}} \right) \leq \frac{1}{2k},$$

where the first inequality uses $\tau \geq 4 \cdot \frac{3\ell k}{n}$.

Now for the event E . Again, the expected (conditioned on ω) mass of items better than I_t in FREE_{t-1} is at most $\frac{|\text{FREE}_{t-1}|}{RO_{n-1}} kR_t \leq \frac{t}{RO_n} kR_t \leq (c_t - \varepsilon_t)(1 + \frac{2\Gamma\ell}{n}) \frac{tk}{n} =: \mu$ (see Observation 3.5). By definition of ε_t , the difference between $c_t \frac{tk}{n} - \frac{\Gamma a_1 \ell k}{n}$ and this expected mass is at least $\tau := \sqrt{10 \ln k} \sqrt{\frac{2tk}{n}}$. So applying the Bernstein's-type inequality of Lemma A.1 we obtain

$$\begin{aligned} \Pr \left(\text{mass of RO items in } \text{FREE}_{t-1} \text{ better than } I_t \geq \mu + \tau \mid I_t \right) &\leq 2 \exp \left(-\frac{\tau^2}{4\mu + \tau} \right) \leq 2 \exp \left(-\frac{\tau^2}{5 \left(\frac{2tk}{n} \right)} \right) \\ &\leq \frac{1}{2k}, \end{aligned}$$

where in the second inequality we used that $t \geq 1$, $212\Gamma\ell$ and $\Gamma \geq \sqrt{k}$.

Taking a union bound over E and F then concludes the proof.