Approximating robust bin-packing with budgeted uncertainty

Marin Bougeret, Noam Goldberg, Michael Poss

► To cite this version:

Marin Bougeret, Noam Goldberg, Michael Poss. Approximating robust bin-packing with budgeted uncertainty. 2020. hal-02119351v2

HAL Id: hal-02119351

<https://hal.science/hal-02119351v2>

Preprint submitted on 11 Feb 2020 (v2), last revised 30 Oct 2020 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Constant-Ratio Approximation for Robust Bin Packing with Budgeted Uncertainty[★]

Aniket Basu Roy¹, Marin Bougeret¹, Noam Goldberg², and Michael Poss¹

¹ LIRMM, University of Montpellier, CNRS, Montpellier, France
{aniket.basu-roy,marin.bougeret,michael.poss}@lirmm.fr

² Department of Management, Bar-Ilan University, Ramat Gan 5290002, Israel
noam.goldberg@biu.ac.il

Abstract. We consider robust variants of the bin packing problem with uncertain item sizes. Specifically we consider two uncertainty sets previously studied in the literature: budgeted uncertainty (the U^F model) in which at most F items deviate, each reaching its peak value, while other items assume their nominal values. The second uncertainty set, the U^Ω model, bounds the total amount of deviation in each scenario. We show that a variant of the next-fit algorithm is a 2 approximation for the U^Ω model, and another variant of this algorithm is a $2F$ approximation for the U^F model. This first result motivates the question of the existence of a constant approximation factor algorithm for the U^F model. Our main result affirms this question by developing a dynamic-programming based algorithm for which a $3 + \rho \leq 4.5$ approximation factor is proven for the U^F model, where ρ is an approximation factor guarantee of a classical bin packing approximation algorithm.

Keywords: Bin-packing · robust optimization · approximation algorithms · Next-fit · dynamic programming

1 Introduction

Bin packing is the problem of assigning a given set of n items, each item of a specified size, to the smallest number of unit capacity bins. The problem has been the subject of study in an extensive body of research initiated by several publications in the 1970s including the work of Johnson et al. [20]. The problem is $\mathcal{NP}$ -hard and in fact a straightforward reduction from the partition decision problem implies that it is $\mathcal{NP}$ -hard to determine whether a bin packing instance has a solution using only two bins. This also shows that the problem cannot be approximated within a factor less than $3/2$ unless $\mathcal{P} = \mathcal{NP}$. An approximation factor guarantee of $3/2$ has been proven for the first-fit decreasing algorithm by Simchi-Levi [26]. In the standard deterministic setting, much of the research has

[★] A preliminary version of this work has been published in [7].

This research has benefited from the support of the ANR project ROBUST [ANR-16-CE40-0018].

concentrated on the asymptotic case where n tends to infinity, and on the online setting where the instance is not given in advance but each item is revealed and packed one at a time. A fully polynomial-time approximation scheme for the offline asymptotic problem is due to Karmarkar and Karp [21]. The best asymptotic and absolute online competitive ratios of 1.578 and $5/3$, respectively, are due to Balogh et al. in [6] and [5], respectively. Approximation results have also been developed for different extensions and generalizations of the classical bin packing problem; for example see [15, 18] and references therein.

In many applications, the sizes of the items to be packed are not fully known at the time that the packing is carried out. In cargo shipping, for example, the actual weight of a container may deviate from its declared weight or its measurements may be inaccurate. Bin packing has been used to model the assignment of elective surgeries to operating room in hospitals [16]. Here a bin is a shift of a properly equipped and staffed operating room for performing a certain type of elective surgeries. The room scheduler has to fit in the bins as many cases (patients) as possible. In this setting clearly the length of time of performing each surgery is subject to uncertainty for example in the event of complications. Bin packing variants has also been used to model the scheduling of internet advertising; see for example [1]. When multiple commercials are packed into advertising breaks of online video services and each commercial may be skipped by the viewer then the length of each ad is subject to uncertainty. One way to model the uncertainty that falls into the framework of robust optimization is to assume that the sizes are uncertain parameters taking any value in a given set $U \subset \mathbb{R}^n$, where each $a \in U$ represents a possible scenario. This leads to the following problem (where the description of U is sometimes not explicit to avoid exponential length in n).

Robust bin packing (RBP)

Input: $U \subset \mathbb{R}^n$

Output: A solution that is a partition of $[n]$ into k bins $b_1, \dots, b_k$ such that $\max_{a \in U} \sum_{i \in b_j} a_i \leq 1$ for each $j \in [k]$

Minimize: k

Classically, robust combinatorial optimization has dealt with uncertain objective, meaning that the cost vector c can take any value in set U , unlike RBP where the uncertainty affects the feasibility of the solutions. In that context, it is well-known that arbitrary uncertainty sets U lead to robust counterparts that that can be more challenging to approximate. For instance, the robust counterpart of the knapsack problem cannot be approximated to within any factor that is a function of the input size (number of items n) unless $\mathcal{P} = \mathcal{NP}$; see [2] and for additional complexity results with specific uncertainty sets see [3, 14]. Moreover, problems whose nominal (deterministic) version can be solved in polynomial time like the shortest path problem, the minimum spanning tree problem, the minimum cut problem and assignment problem can-

not be approximated to within constant factors; see [22, 23]. Furthermore, describing U by an explicit list of scenarios runs the risk of over-fitting so the optimal solutions may become infeasible for small variations outside U . These two drawbacks are usually tackled by using more specific uncertainty sets, defined by simple budget constraints. One of these widely used uncertainty sets, U^Γ , supposes that the size of each item is either its given nominal size $\bar{a}_i$, or its peak value $\bar{a}_i + \hat{a}_i$. Furthermore, in any scenario, at most $\Gamma \in \mathbb{N}$ of the items may assume their peak value simultaneously. Formally, U^Γ can be defined as $U^\Gamma = \{a | \forall i \in [n], a_i \in \{\bar{a}_i, \bar{a}_i + \hat{a}_i\} \text{ and } \sum_{i \in [n]} (a_i - \bar{a}_i) / \hat{a}_i \leq \Gamma\}$.³ Uncertainty set U^Γ has been widely used in robust combinatorial optimization with a constant number of constraints because the set essentially preserves the complexity properties of the nominal problem, and to some extent, its approximability properties as well. Results for this model, including both algorithms and computational complexity for combinatorial optimization problems, were first developed for min-max problems and uncertain objective functions in [10]. In the general robust optimization literature this uncertainty model is viewed as a polyhedral uncertainty set, specifically the intersection of L_1 and L_∞ balls [8, 9], but that literature had mostly focused on results for ellipsoidal uncertainty sets in continuous settings. The results of [10, 11] were later extended to uncertain constraints independently in [4, 17]. We also consider a second uncertainty set (used in [19, 24, 28], among others), characterized again by $\bar{a}$ and $\hat{a}$, as well as the number $\Omega \in [0, 1]$ stating how much deviation can be spread among all sizes, formally $U^\Omega = \{a \in \times_{i \in [n]} [\bar{a}_i, \bar{a}_i + \hat{a}_i] \mid \sum_{i \in [n]} (a_i - \bar{a}_i) \leq \Omega\}$. From the approximability viewpoint, set U^Ω benefits from similar positive results as U^Γ ; see [25].

The above positive complexity results (e.g., [17, 25]) imply, for instance, that under mild assumptions there exists a fully-polynomial time approximation scheme (FPTAS) for the robust knapsack problem with uncertain profits and uncertain weights belonging to U^Ω and/or U^Γ . Interestingly, these positive results do not extend to most scheduling problems (because they involve nonlinearities) and to the bin packing problem (because it involves a non-constant number of robust constraints). While in previous papers [13, 12] (with authors in common) we provided approximability results on robust scheduling, no such results have yet been proposed for the bin packing problem, the only previous work focusing on numerical algorithms [27]. The purpose of this paper is to fill these gaps, as we present constant-factor approximation algorithms the bin packing problem, both for U^Ω and U^Γ .

Notation, problems definitions, and next-fit algorithm In this paper we consider two special cases of RBP. In the first one, Γ RBP, the input is $\mathcal{I} = (\bar{a}, \hat{a}, \Gamma) \in [0, 1]^n \times [0, 1]^n \times_+ \mathbb{N}$ where $n \in \mathbb{N}$, and $U = U^\Gamma$. In the second one,

³ U^Γ is often defined alternatively in the literature, as the polytope $\{a \in \times_{i \in [n]} [\bar{a}_i, \bar{a}_i + \hat{a}_i] \mid \sum_{i \in [n]} (a_i - \bar{a}_i) / \hat{a}_i \leq \Gamma\}$. For the bin packing problem, one readily verifies using classical arguments that the two definitions lead to the same optimization problem.

Ω RBP, the input is $\mathcal{I} = (oa, \hat{a}, \Omega) \in [0, 1]^n \times [0, 1]^n \times [0, 1]$ where $n \in \mathbb{N}$ and $U = U^\Omega$.

Let us now define some notation that may be required for formally stating the Γ RBP and Ω RBP problems. Given $n \in \mathbb{N}$, sets $\{0, 1, \dots, n\}$ and $\{1, \dots, n\}$ are respectively denoted $[n]_0$ and $[n]$. Set $\{i, \dots, j\}$ is denoted by $\llbracket i, j \rrbracket$. Given a vector $v \in [0, 1]^n$ and a subset $X \subseteq [n]$, we define $v(X) = \sum_{i \in X} v_i$. Given two vectors $\bar{a} \in [0, 1]^n$, $\hat{a} \in [0, 1]^n$ and a subset of items $X \subseteq [n]$, we define $\hat{a}_\Omega(X) = \min\{\hat{a}(X), \Omega\}$, $\Gamma(X)$ as the set of Γ items in X with largest $\hat{a}$ values (ties broken by taking smallest indices), or $\Gamma(X) = X$ if $|X| < \Gamma$, and $\hat{a}_\Gamma(X) = \hat{a}(\Gamma(X))$. Accordingly, we define the fill of a bin $b \subseteq [n]$ as $f_\Gamma(b) = \bar{a}(b) + \hat{a}_\Gamma(b)$ for set U^Γ , and $f_\Omega(b) = \bar{a}(b) + \hat{a}_\Omega(b)$ for set U^Ω . The fill of a bin for a general uncertainty set U is denoted as $f_U(b) = \max_{a \in U} a(b)$.

Consider the following example. We are given an ordered set of pairs $(\bar{a}_i, \hat{a}_i)$, $X = \{(0.3, 0.2), (0.4, 0.2), (0.3, 0.1), (0.2, 0.5)\}$ with $\Gamma = 2$ and $\Omega = 0.3$. Thus, $\Gamma(X) = \{(0.3, 0.2), (0.2, 0.5)\}$, $\bar{a}(X) = 1.2$, $\hat{a}_\Gamma(X) = 0.7$, and $f_\Gamma(X) = 1.9$. Similarly, $\hat{a}_\Omega(X) = 0.3$ and $f_\Omega(X) = 1.5$.

Now, observe that $\max_{a \in U} \sum_{i \in b_j} a_i \leq 1$ (the constraint required in RBP) is equivalent to $f_U(b) \leq 1$, and thus to $f_\Gamma(b_j) \leq 1$ for Γ RBP and $f_\Omega(b_j) \leq 1$ for Ω RBP. For example in Γ RBP, $f_\Gamma(b_j) \leq 1$ simply means that the total nominal ($\bar{a}$) size of the items plus the deviating size ($\hat{a}$) of the Γ largest (in $\hat{a}$ values) items must not exceed one. Then, the two optimization problems studied in this paper can be equivalently formulated as follows.

Γ -robust bin packing (Γ RBP)

Input: $\mathcal{I} = (\bar{a}, \hat{a}, \Gamma) \in [0, 1]^n \times [0, 1]^n \times \mathbb{N}$.

Output: A solution that is a partition of $[n]$ into k bins $b_1, \dots, b_k$ such that $f_\Gamma(b_j) \leq 1$ for each $j \in [k]$

Minimize: k

Ω -robust bin packing (Ω RBP)

Input: $\mathcal{I} = (n, \bar{a}, \hat{a}, \Omega)$ where $n \in \mathbb{N}$, $\bar{a} \in [0, 1]^n$, $\hat{a} \in [0, 1]^n$, and $\Omega \in [0, 1]$.

Output: A solution that is a partition of $[n]$ into k bins $b_1, \dots, b_k$ such that $f_\Omega(b_j) \leq 1$ for each $j \in [k]$

Minimize: k

The optimal solution value or cost of either problem is denoted by $\text{OPT}(\mathcal{I}) = k^*$ ($\mathcal{I}$ may be omitted when the instance is clear from the context) and a corresponding optimal solution is denoted by $s^* = \{b_1^*, b_2^*, \dots, b_{k^*}^*\}$. We introduce in Algorithm 1 a variant of the standard next fit algorithm.

Structure of the paper In Sections 2 and 3, we analyze approximation-factor guarantees for NEXT-FIT in the case of Ω RBP and Γ RBP, respectively. For

initialization: $j = 1$

- 1 Pack items (with smaller index first) in b_j until $f_U(b_j) > 1$ or $n \in b_j$. If $n \notin b_j$ then $j \leftarrow j + 1$ and repeat Step 1. Otherwise, $k' \leftarrow j$ proceed to Step 2.
- 2 Pack the last item of each bin in a new bin: for any j , let $i = \max(b_j)$,
 $b_j^1 = b_j \setminus \{i\}$, and $b_j^2 = \{i\}$

return : $\bigcup_{j=1}^{k'} \{b_j^1, b_j^2\}$

Algorithm 1: NEXT-FIT($\mathcal{I}$)

Ω RBP, using ordering (1) (non-increasing ordering on $\frac{\hat{a}_i}{\bar{a}_i}$) the ratio is equal to 2. For Γ RBP, using ordering (2) (non-increasing ordering on $\hat{a}_i$), the ratio is bounded by 2Γ . As Theorem 3 shows that neither ordering (1) or (2) leads to a constant ratio using NEXT-FIT, this raises the question of the existence of a constant approximation for Γ RBP. We answer the question in Section 4 by providing a dynamic programming algorithm (DP) giving a ratio of 4.5 for Γ RBP and any $\Gamma \in \mathbb{N}$, which is our main result.

2 Next-fit for Ω RBP

Unlike the classical bin packing problem, executing NEXT-FIT on arbitrarily ordered items can lead to arbitrarily bad solutions. For example, given ϵ with $0 < \epsilon \leq \frac{1}{2n}$, consider an instance with $\Omega = 1 - \epsilon$, and items $((2\epsilon, 0), (0, 1 - \epsilon), \dots, (2\epsilon, 0), (0, 1 - \epsilon))$, where item $i \in [n]$ is denoted by the pair $(\bar{a}_i, \hat{a}_i)$. Using this ordering, NEXT-FIT will create $n/2$ bins b_j with $f_\Omega(b_j) > 1$ for any $j \in [n]$ (which will be turned into n bins $\{b_j^1, b_j^2\}$), whereas the optimal solution uses 2 bins. This example also illustrates that, unlike in the standard bin packing, the total size argument no longer apply to the robust counterpart as having $f_\Omega(b_j) > 1$ for any $j \in [n]$ does not imply a large (depending on n) lower bound on the optimum.

Next, we consider an ordering of the items such that

$$\hat{a}_1/\bar{a}_1 \geq \dots \geq \hat{a}_n/\bar{a}_n. \quad (1)$$

Lemma 1. *Suppose that the items are ordered according to (1). Then $k' \leq k^*$.*

Proof. Consider an optimal solution $b_1^*, \dots, b_{k^*}^*$ and the subset of optimal bins given by $G^* = \{j \in [k^*] \mid \hat{a}(b_j^*) > \Omega\}$. Let

$$A = \sum_{i \in [n]} (\bar{a}_i + \hat{a}_i) = \sum_{j=1}^{k'} (\bar{a}(b_j) + \hat{a}(b_j)) = \sum_{j=1}^{k^*} (\bar{a}(b_j^*) + \hat{a}(b_j^*)).$$

Let G denote the first $|G^*|$ bins opened in Step 1 of NEXT-FIT. If $k' \in G$ then clearly $k' \leq k^*$. Otherwise, it can be observed that for each $l \in G$, $\bar{a}(b_l) > 1 - \Omega$ (as $\bar{a}(b_l) + \hat{a}_\Omega(b_l) > 1$ and $\hat{a}_\Omega(b_l) \leq \Omega$) and $1 - \Omega \geq \max_{j \in G^*} \bar{a}(b_j^*)$ (as $f_\Omega(b_j^*) \leq 1$). Thus, $\sum_{j \in G} \bar{a}(b_j) > \sum_{j \in G^*} \bar{a}(b_j^*)$ and so by the assumed ordering (1) of

the items, following a standard knapsack argument, $\sum_{j \in G} \hat{a}(b_j) > \sum_{j \in G^*} \hat{a}(b_j^*)$. Letting $\bar{G} = [k'] \setminus G$ and $\bar{G}^* = [k^*] \setminus G^*$, it follows that

$$\begin{aligned} \sum_{j \in \bar{G}} (\bar{a}(b_j) + \hat{a}(b_j)) &= A - \sum_{j \in G} (\bar{a}(b_j) + \hat{a}(b_j)) \leq \\ &A - \sum_{j \in G^*} (\bar{a}(b_j^*) + \hat{a}(b_j^*)) = \sum_{j \in \bar{G}^*} (\bar{a}(b_j^*) + \hat{a}(b_j^*)) \end{aligned}$$

(equality may hold throughout if $G^* = \emptyset$). Further, for each $j \in \bar{G} \setminus \{k'\}$, $\bar{a}(b_j) + \hat{a}(b_j) \geq f_\Omega(b_j) > 1$ and for each $j \in \bar{G}^*$, $\bar{a}(b_j^*) + \hat{a}(b_j^*) \leq 1$. Therefore, $|\bar{G}| \leq \left\lceil \sum_{j \in \bar{G}} (\bar{a}(b_j) + \hat{a}(b_j)) \right\rceil \leq \left\lceil \sum_{j \in \bar{G}^*} (\bar{a}(b_j^*) + \hat{a}(b_j^*)) \right\rceil \leq |\bar{G}^*|$ and $k' \leq k^*$ as claimed. $\square$

The lemma combined with Step 2 of NEXT-FIT immediately imply the following theorem.

Theorem 1. *If the items are ordered according to (1) then NEXT-FIT is a 2-approximation algorithm for Ω RBP.*

3 Next-fit for Γ RBP

From now on, we focus on problem Γ RBP. Remark first that using an arbitrary ordering leads to arbitrarily bad solutions, considering $\Gamma = 1$ and the same items $((2\epsilon, 0), (0, 1 - \epsilon), \dots, (2\epsilon, 0), (0, 1 - \epsilon))$ as in the previous section. Thus, we consider here an ordering of the items such that

$$\hat{a}_1 \geq \dots \geq \hat{a}_n. \quad (2)$$

The main result of this section is the following. We also note that this result has been improved compared with the result that appears in the preliminary extended abstract version of this paper [7].

Theorem 2. *Suppose that the items are ordered according to (2). Then NEXT-FIT is a 2Γ -approximation algorithm for Γ RBP.*

Recall that k' is the number of bins used in Step 1 and let $s' = (b_1, \dots, b_{k'})$ be the bins output at the end of Step 1. Let $s^* = \{b_1^*, \dots, b_{k^*}^*\}$ be an optimal solution. Define $i_j^* = \max(\Gamma(b_j^*))$, and also let $i_j = \max(\Gamma(b_j))$ for each $j \in [k']$.

The key element in proving Theorem 2 is the following counterpart of Lemma 1, a result which immediately implies an approximation-factor guarantee of 2Γ .

Lemma 2. *Suppose that the items are ordered according to (2). Then $k' \leq \Gamma k^*$.*

Let $\mathcal{M}' \subseteq s'$ be the set of bins that contain only items that either do not deviate or are the smallest deviations, i_j^* for some $j \in [k^*]$. So for each $b \in \mathcal{M}'$, $b \subseteq [n] \setminus \bigcup_{j=1}^{k^*} \Gamma(b_j^*) \cup \{i_1^*, \dots, i_{k^*}^*\}$. In what follows, we bound $|\mathcal{M}'|$ and $|s' \setminus \mathcal{M}'|$ by multiples of k^* .

Lemma 3. $|s' \setminus \mathcal{M}'| \leq (\Gamma - 1)k^*$

Proof. As each bin $b \in s' \setminus \mathcal{M}'$ contains at least one item from $\bigcup_{j=1}^{k^*} \Gamma(b_j^*) \setminus \{i_1^*, \dots, i_{k^*}^*\}$, and as $|\bigcup_{j=1}^{k^*} \Gamma(b_j^*) \setminus \{i_1^*, \dots, i_{k^*}^*\}| \leq (\Gamma - 1)k^*$, the result is immediate. $\square$

The case of $\mathcal{M}'$, which is slightly more involved, is addressed in the following lemma.

Lemma 4. $|\mathcal{M}'| \leq k^*$.

Proof. For convenience in the following without loss of generality let $\mathcal{M}' = \{b_1, \dots, b_{k''}\}$ where $k'' = |\mathcal{M}'| \leq k'$, and let $\{b_1^*, \dots, b_{k^{**}}^*\} \subseteq s^*$ be the set of optimal solution bins each containing Γ deviations, with bins ordered in non-increasing smallest deviation size, so $\hat{a}_{i_k^*} \geq \hat{a}_{i_{k+1}^*}$ for each $k \in [k^{**} - 1]$, and where $k^{**} \leq k^*$. Assume for the sake of deriving a contradiction that $k^{**} < k''$. Then, since $\bigcup_{j=1}^{k''} b_j \subseteq \bigcup_{j=1}^{k^{**}} b_j^*$,

$$\sum_{j=1}^{k''} \bar{a}(b_j) \leq \sum_{j=1}^{k^{**}} \bar{a}(b_j^*) \leq k^{**} - \sum_{j=1}^{k^{**}} \hat{a}_\Gamma(b_j^*).$$

We now show that $\sum_{j=1}^{k^{**}} \hat{a}_\Gamma(b_j) \leq \sum_{j=1}^{k^{**}} \hat{a}_\Gamma(b_j^*)$. To do so we show by induction on $k = 1, \dots, k^{**} - 1$ for a fixed instance and corresponding algorithm bins and optimal solution bins, $b_1, \dots, b_k''$ and $b_1^*, \dots, b_{k^{**}}^*$, respectively. For $k = 1$, and all $i \in b_1$ by the fact that these items either do not deviate in s^* or are in the set of smallest deviations, $\{i_1^*, \dots, i_{k^{**}}^*\}$, it follows that $\hat{a}_i \leq \hat{a}_{i_1^*}$, so $\hat{a}_\Gamma(b_1) \leq \hat{a}_\Gamma(b_1^*)$. Now assume $\sum_{j=1}^k \hat{a}_\Gamma(b_j) \leq \sum_{j=1}^k \hat{a}_\Gamma(b_j^*)$ in order to prove that $\sum_{j=1}^{k+1} \hat{a}_\Gamma(b_j) \leq \sum_{j=1}^{k+1} \hat{a}_\Gamma(b_j^*)$. By the induction hypothesis

$$\begin{aligned} \sum_{j=1}^k \bar{a}(b_j) &> 1 - \sum_{j=1}^k \hat{a}_\Gamma(b_j) \\ &\geq 1 - \sum_{j=1}^k \hat{a}_\Gamma(b_j^*) \\ &\geq \sum_{j=1}^k \bar{a}(b_j^*) \\ &\geq \sum_{j=1}^k \bar{a}(b_j^* \setminus \Gamma(b_j^*) \cup \{i_j^*\}). \end{aligned}$$

The above inequality and the ordering of the items imply that

$$\hat{a}_{\max} \equiv \max_{i \in \bigcup_{j=k+1}^{k^{**}} (b_j^* \setminus \Gamma(b_j^*) \cup \{i_j^*\})} \hat{a}_i \geq \max_{i \in b_{k+1}} \hat{a}_i.$$

Since $\hat{a}_{i_{k+1}^*} = \hat{a}_{\max}$ it follows that $\hat{a}_\Gamma(b_{k+1}^*) \geq \hat{a}_\Gamma(b_{k+1})$. Together with the induction hypothesis it implies that $\sum_{j=1}^{k+1} \hat{a}_\Gamma(b_j) \leq \sum_{j=1}^{k+1} \hat{a}_\Gamma(b_j^*)$.

Note that $\hat{a}_\Gamma(b_j) \leq 1$ for $j = 1, \dots, k''$; otherwise if $\hat{a}_\Gamma(b_j) > 1$ for some $j \in [k'']$, then there is some $i \in b_j$ that has $\hat{a}_i > 1/\Gamma$. But then $i \in b_l^*$ for some $l \in [k^{**}]$ and since i is either a nondeviating item and $\hat{a}_\Gamma(b_l^*) > 1$, or $i = i_l^*$ and it implies that b_l^* may contain only $\Gamma - 1$ deviating items, thereby establishing a contradiction. Now, by definition of the algorithm,

$$\sum_{j=1}^{k''} \bar{a}(b_j) > \sum_{j=1}^{k''-1} \bar{a}(b_j) > k'' - 1 - \sum_{j=1}^{k''-1} \hat{a}_\Gamma(b_j) \geq k^{**} - \sum_{j=1}^{k^{**}} \hat{a}_\Gamma(b_j^*).$$

The last inequality followed from $k'' - 1 \geq k^{**}$ and $\hat{a}_\Gamma(b_j) \leq 1$ for $j \in [k'']$, implying that $\sum_{j=k^{**}+1}^{k''-1} \hat{a}_\Gamma(b_j) \leq k'' - k^{**}$.

Proof of Lemma 2. Lemmas 3 and 4 immediately imply that $k' \leq \Gamma k^*$, thus proving the claim of Lemma 2. $\square$

Proof of Theorem 2. After step 2 of the algorithm the total number of bins is at most $2\Gamma k^*$, concluding the proof of Theorem 2. $\square$

To complete the analysis, we establish the following lower bound on the approximation ratio of NEXT-FIT.

Theorem 3. *If the items are ordered according to (2) or (1), then the approximation ratio of NEXT-FIT for Γ RBP is at least $\frac{2\Gamma}{3}$.*

Proof. Let us define an instance where the ordering (2) can lead to Step 1 of NEXT-FIT using $k' = \Gamma$ bins while $\text{OPT} = 3$. Every row of the $\Gamma \times \Gamma$ matrix below corresponds to the set of items in a bin (after the Step 1) of NEXT-FIT algorithm

$$\begin{pmatrix} (\epsilon, 1/\Gamma - \delta_1) & (0, 1/\Gamma - \delta_1) & \dots & (0, 1/\Gamma - \delta_1) \\ \vdots & \vdots & \ddots & \vdots \\ (\epsilon, 1/\Gamma - \delta_\Gamma) & (0, 1/\Gamma - \delta_\Gamma) & \dots & (0, 1/\Gamma - \delta_\Gamma) \end{pmatrix} \quad (3)$$

where $\epsilon \leq 1/\Gamma$ and $\delta_1 \leq \dots \leq \delta_\Gamma < \epsilon/\Gamma$. On the one hand, $\epsilon + \Gamma \cdot (1/\Gamma - \delta_l) > 1$ for each $l \in [\Gamma]$, so step 1 of NEXT-FIT outputs Γ bins. On the other hand, an optimal solution can pack all the items above except the ones in the first column into a single bin because $\Gamma \cdot 1/\Gamma - \delta_1 \leq 1$. Further, the total weight of the first $\Gamma/2$ items of the first column sums up to $\Gamma/2 \cdot (1/\Gamma + \epsilon) - \sum_{l=1}^{\Gamma/2} \delta_l \leq 1 - \sum_{l=1}^{\Gamma/2} \delta_l \leq 1$, and similarly for the last $\Gamma/2$ items, so an optimal solution may pack the first column using two bins.

This instance can be adapted to establish a lower bound for the approximation ratio of NEXT-FIT when items are ordered according to (1). We consider an example that yields a lower bound on the approximation ratio of NEXT-FIT in solving Γ RBP when the items are ordered according to (1). For some $c \geq \Gamma^2$,

$\epsilon' = \frac{2}{\Gamma(1+c)}$, $\epsilon = \frac{1}{\Gamma(\Gamma^2+\Gamma-1)}$ consider the instance given by the following $\Gamma \times \Gamma$ matrix:

$$\begin{pmatrix} (\epsilon', c\epsilon') & (\epsilon, c\epsilon) & \dots & (\epsilon, c\epsilon) \\ \vdots & \vdots & \ddots & \vdots \\ (\epsilon', c\epsilon') & (\epsilon, c\epsilon) & \dots & (\epsilon, c\epsilon) \end{pmatrix}$$

It can be verified that NEXT-FIT opens a bin for each row, since $(1+c)\epsilon' + (\Gamma-1)(1+c)\epsilon > 1$. The optimal solution opens 3 bins, 2 bins to store the items of the first column and another bin to store the rightmost $\Gamma-1$ columns. Although in this example all items $i \in [n]$ are set to have ratios $\hat{a}_i/\bar{a}_i = c$, the example can be extended in a straightforward manner with slight perturbations of the item sizes so that the ratios will be strictly decreasing for the items ordered from left-to-right and top-to-bottom in this matrix. $\square$

We conclude the section by emphasizing our results for the special cases of $\Gamma = 1$ and $\Gamma = 2$, where the straight-forward NEXT-FIT algorithm with ordering (2) obtains the best approximation guarantees using the analysis of the current paper.

Corollary 1. *If the items are ordered according to (2) and $\Gamma = 1$ or $\Gamma = 2$, then NEXT-FIT is a 2-factor or 4-factor approximation algorithm, respectively, for Γ RBP.*

In the section that follows we consider constant-factor approximation factors that do not depend on the parameter Γ .

4 A constant-factor approximation algorithm for Γ RBP

The algorithm presented in this section relies on three main ideas. First, we show in Section 4.1 that we can restrict ourselves to instances of Γ RBP with small items; that is, $\bar{a}_i \leq \frac{1}{\Gamma}$ and $\hat{a}_i \leq \frac{1}{\Gamma}$ for each $i \in [n]$. Specifically, we show how to convert any ρ -approximation algorithm for the latter special case of the problem with small items into a general $(\rho + \rho_{\text{bp}})$ -approximation for Γ RBP, where ρ_{bp} is an approximation factor guarantee of some algorithm for classical bin packing. First note that having only small items, any set of $\lfloor \Gamma/2 \rfloor$ items can always be packed together into a bin. This fact is stated as the following observation.

Observation 1 *Given an instance $\mathcal{I}$ to the Γ RBP satisfying $\hat{a}_i \leq 1/\Gamma$ and $\bar{a}_i \leq 1/\Gamma$ for each $i \in [n]$, any subset $X \subseteq [n]$ can be packed in $\lceil \frac{2|X|}{\Gamma} \rceil$ bins.*

Next, we introduce in Section 4.2 variants of Γ RBP where items are packed into bins as in standard bin packing but some of the items are placed in two designated special bins. The deviations will be ignored so that they are irrelevant in these two bins, while in each of the other regular bins only a single item can deviate. One of the special bins that will be referred to as the trash bin, cannot contain more than $k(\Gamma-1)$ items, where k is the number of regular bins used in the solution. The trash contains and “mimics” the $\Gamma-1$ deviating items of

each bin in an optimal solution. Following Observation 1, items in the trash can be packed into at most $2k$ additional bins. The problem with trash remains hard because of the capacity of the regular bins, so we focus on almost feasible solutions, which are allowed to exceed each regular bin by one item. We show in that section how finding almost feasible solutions no worse than the optimal solution for the problem with trash leads to an approximation algorithm for Γ RBP with small items.

Finally, we present in Section 4.3 a dynamic programming (DP) algorithm that outputs an almost feasible solution using a number of bins that is at most that of the optimal solution of the original Γ RBP problem. Essentially, for each bin the DP algorithm guesses a particular item that deviates. Then it greedily packs the remaining items with largest nominal values into the trash.

We maintain throughout the section the assumption that the items are ordered according to (2).

4.1 Robust bin packing with small items

We define Γ RBP with small values as the Γ RBP problem restricted to inputs where for any $i \in [n]$, $\bar{a}_i \leq \frac{1}{\Gamma}$ and $\hat{a}_i \leq \frac{1}{\Gamma}$. The following proposition motivates our focus on Γ RBP with small values.

Proposition 1. *Any polynomial ρ -approximation algorithm for Γ RBP with input satisfying $\hat{a}_i \leq 1/\Gamma$ and $\bar{a}_i \leq 1/\Gamma$ for each $i \in [n]$, can be turned into a polynomial-time $(\rho + \rho_{bp})$ -approximation algorithm for Γ RBP.*⁴

Proof. Given an instance $\mathcal{I}$ of Γ RBP, we define the small items $\mathcal{S} = \{i \in [n] : \bar{a}_i \leq 1/\Gamma \text{ and } \hat{a}_i \leq 1/\Gamma\}$ and the large item as $\mathcal{B} = [n] \setminus \mathcal{S}$. We use the given ρ -approximation algorithm to pack $\mathcal{S}$ into $k_{\mathcal{S}}$ bins, so that $k_{\mathcal{S}} \leq \rho \text{OPT}(\mathcal{I})$. Then, we observe that in any packing of $\mathcal{B}$, each bin contains at most Γ items, so that all items deviate in these bins. Hence, the least number of bins needed to pack items in $\mathcal{B}$ is given by a solution that is optimal to the standard bin packing problem for this same set of items $\mathcal{B}$ where the size of each item $i \in \mathcal{B}$ is $a'_i = \bar{a}_i + \hat{a}_i$. Let $\mathcal{I}'$ denote this instance of standard bin packing. So, $\text{OPT}_{bp}(\mathcal{I}') \leq \text{OPT}(\mathcal{I})$ (where OPT_{bp} denotes the optimal objective value of standard bin packing), and observe that any optimal solution of standard bin packing for items $\mathcal{B}$ is a solution that is feasible for the Γ RBP instance $\mathcal{I}$ using the same number of bins. Using a ρ_{bp} -approximation algorithm for standard bin packing to pack $\mathcal{B}$ in $k_{\mathcal{B}}$ bins, it follows that $k_{\mathcal{B}} \leq \rho_{bp} \text{OPT}_{bp}(\mathcal{I}') = \rho_{bp} \text{OPT}(\mathcal{I}) \leq \rho_{bp} \text{OPT}(\mathcal{I})$. Thus, the packing of $\mathcal{B}$ and $\mathcal{S}$ is a packing of $\mathcal{I}$ satisfying $k_{\mathcal{S}} + k_{\mathcal{B}} \leq (\rho + \rho_{bp}) \text{OPT}(\mathcal{I})$. $\square$

Notice that instances with small items are not easier to approximate by NEXT-FIT as illustrated by the instance defined by (3), in Section 3, whose items satisfy $a_i, \hat{a}_i \leq 1/\Gamma$ for all $i \in [n]$.

⁴ In general, if we have a polynomial time additive approximation algorithm using $\text{OPT} + f(\text{OPT})$ bins and polynomial time ρ -approximation algorithm for Γ RBP with small values then our algorithm uses $\text{OPT}(\rho + 1) + f(\text{OPT})$ bins for Γ RBP in polynomial time.[NG??]

4.2 Bin-packing with trash

For any $X \subseteq [n]$, we define $\tilde{a}_\Gamma(X) = \Gamma \hat{a}_1(X)$ ($\tilde{a}_\Gamma(X)$ is Γ times the largest deviating value of an item in X) and $\tilde{f}(X) = \bar{a}(X) + \tilde{a}_\Gamma(X)$. We introduce next a decision problem $\Gamma\text{RBP-T}$ that is related to ΓRBP .

$\Gamma\text{RBP-T}$ (Robust bin packing with trash)	
Input:	$(\mathcal{I}, k, t)$ where $\mathcal{I}$ is an instance of ΓRBP and $k, t \in \mathbb{N}$.
Output:	'Yes' if a solution exists, which is a partition of the set of items into $k + 1$ sets $b_1, \dots, b_k$ and T (called the trash) such that: <ul style="list-style-type: none"> – $\tilde{f}(b_j) \leq 1$ for each $j = 1, \dots, k$ – $T \leq t$ and 'No' otherwise.

Notice that although the input of $\Gamma\text{RBP-T}$ is assumed to include only small items, it is possible to have an item $i \in [n]$ such that $\tilde{f}(\{i\}) > 1$, implying that $i \in T$. The following two lemmas suggest how the decision problem $\Gamma\text{RBP-T}$ may be used to determine an approximate solution of ΓRBP .

Lemma 5. *For any input $\mathcal{I}$ of ΓRBP where $k^* = \text{OPT}(\mathcal{I})$, $(\mathcal{I}, k^*, (\Gamma - 1)k^*)$ is a yes instance of $\Gamma\text{RBP-T}$.*

Proof. Given an optimal solution with objective value k^* of ΓRBP we create a solution to $\Gamma\text{RBP-T}$ problem as follows. For some (arbitrary) $j \in [k^*]$, let b_j^* be a bin of the considered optimum. Let $N_j = b_j^* \setminus \Gamma(b_j^*)$ (the non-deviating items of b_j^*). For $j = 1, \dots, k^*$, let

$$X_j = \begin{cases} \max(\Gamma(b_j^*)) & |\Gamma(b_j^*)| = \Gamma \\ \emptyset & \text{otherwise.} \end{cases}$$

We define $b'_j = N_j \cup X_j$, and adjoin items of $Y_j = b_j^* \setminus b'_j$ to the trash. Note that Y_j is either the set of $\Gamma - 1$ largest deviating items of b_j^* , or otherwise $b_j^* = \Gamma(b_j^*)$ and $|\Gamma(b_j^*)| < \Gamma$. So, $(b'_1, \dots, b'_{k^*}, T)$ is a yes instance for the $\Gamma\text{RBP-T}$ problem since evidently $\tilde{f}(b'_j) = \bar{a}(b'_j) + \tilde{a}_\Gamma(b'_j) \leq \bar{a}(b_j^*) + \hat{a}_\Gamma(b_j^*) \leq 1$ and also $|T| \leq (\Gamma - 1)k^*$. $\square$

The next lemma establishes that yes-instances of $\Gamma\text{RBP-T}$ can be used to construct solutions of ΓRBP using at most a number bins that is a constant factor of the number of bins used by $\Gamma\text{RBP-T}$.

Lemma 6. *For any instance $\mathcal{I}$ of ΓRBP satisfying $\hat{a}_i \leq 1/\Gamma$ and $\bar{a}_i \leq 1/\Gamma$ for each $i \in [n]$, and integer k , given a yes-instance of $(\mathcal{I}, k, \Gamma k)$ of $\Gamma\text{RBP-T}$, we can compute in polynomial time a solution of $3k$ bins for $\mathcal{I}$.*

Proof. Given a solution $b_1, \dots, b_k, T$ for $(\mathcal{I}, k, \Gamma k)$ of $\Gamma\text{RBP-T}$ the bins remain feasible in ΓRBP as for each $j \in [k]$, $f_\Gamma(b_j) = \bar{a}(b_j) + \hat{a}_\Gamma(b_j) \leq \bar{a}(b_j) + \tilde{a}(b_j) = \tilde{f}(b_j)$. Then, Observation 1 implies that the trash T can be packed into $\lceil k\Gamma/(\Gamma/2) \rceil \leq 2k$ additional bins. $\square$

Notice that the trash T in the instance constructed in this lemma has $|T| = \Gamma k$, which exceeds the $|T| = (\Gamma - 1)k$ in Lemma 5. The additional k “slots” are necessary for deciding Γ RBP-T as will be illustrated in the analysis that follows.

In order to develop an algorithm and in particular a dynamic program (DP) for deciding Γ RBP-T it is convenient to consider an optimization variant of Γ RBP-T, to be called G- Γ RBP-T for “generalized robust bin packing with trash”. This variant is defined in the following for a fixed instance $\mathcal{I}$ of Γ RBP and a given integer k (so that they are not considered a part of the input).

Generalized robust bin packing with trash (G- Γ RBP-T)

Input: $\mathcal{I}' = (q, t, \ell)$, where $q \in [n]$, $t \in [(\Gamma - 1)k]_0$, and $\ell \in [k + 1]$.

Output: A feasible solution s is a partition of $\llbracket q, n \rrbracket$ into $k - \ell + 3$ sets, given as a triple $(L, \{b_j : j \in \llbracket \ell, k \rrbracket\}, T)$, such that

- for any $j \in \llbracket \ell, k \rrbracket$, $\tilde{f}(b_j) \leq 1$ (the $k - \ell + 1$ regular bins must satisfy the fill constraints of Γ RBP-T)
- $|T| \leq t$ (we only allow t items in the trash)
- $\min(b_\ell) = q$ (meaning that the deviating item of b_ℓ is q)

Minimize: $c(s) = \bar{a}(L)$

The objective of G- Γ RBP-T is to pack a part (defined by $\llbracket q, k \rrbracket$) of the instance $\mathcal{I}$ for Γ RBP-T given a fixed budget of resources (the number of bins and the size of the trash) while minimizing the sum only of nominal sizes of items in the leftover itemset L . The last constraint (the deviating item of b_ℓ is q), which may appear somewhat artificial, will allow determining optimal solutions of Γ RBP-T from optimal solutions of G- Γ RBP-T, by carefully enumerating possibilities of largest deviating item to be packed each bin in an intelligent way; considering only those possibilities corresponding to solutions that are close to being feasible (a notion to be defined more precisely in the following) and whose objective value is at most that of an optimal solution. This enumeration scheme will be shown to be efficiently solvable by a DP. For convenience, the objective value of infeasible solutions s , including for example $s = (\emptyset, \emptyset, \emptyset)$, is defined as $c(s) = \infty$.

The capacity constraints of the bins make G- Γ RBP-T hard to solve in general. Hence, following the spirit of NEXT-FIT introduced previously, we introduce below almost feasible solutions, which can exceed the capacity of each bin by one item.

Definition 1 (almost feasible solution). *We say that a bin b exceeds by at most one item iff $\tilde{f}(b) > 1$ and $\tilde{f}(b \setminus \{i\}) \leq 1$ where $i = \max(b)$. Given an input $\mathcal{I}' = (q, t, \ell)$ of G- Γ RBP-T, we say that a solution is almost feasible iff all the G- Γ RBP-T constraints are satisfied, except that for any $j \in \llbracket \ell, k \rrbracket$, we allow that b_j exceeds by at most one item instead of $\tilde{f}(b_j) \leq 1$.*

Definition 2 (an optimal almost-feasible solution). *Given an input $\mathcal{I}' = (q, t, \ell)$ of G- Γ RBP-T, we say that a solution $s = (L, \{b_1, \dots, b_k\}, T)$ is an optimal almost-feasible solution iff s is almost feasible with $c(s) \leq \text{OPT}(\mathcal{I}')$.*

The relation between G- Γ RBP-T and Γ RBP-T is characterized in the following two lemmas. Let $\text{OPT}(\mathcal{I}')$ be the optimal solution cost of G- Γ RBP-T instance $\mathcal{I}'$ (notice that an optimal solution must also be feasible for the given problem).

Lemma 7. *For any input $\mathcal{I}$ of Γ RBP and k such that $(\mathcal{I}, k, (\Gamma - 1)k)$ is a yes input of Γ RBP-T, there exist a positive integer $q \in [(\Gamma - 1)k]$ such that $\text{OPT}(q, (\Gamma - 1)k - (q - 1), 1) = 0$.*

Proof. Consider a Γ RBP-T yes-instance $(\mathcal{I}, k, (\Gamma - 1)k)$ and corresponding partition solution $\{b_1, \dots, b_k\}, T$. Let $q = \min \left(\bigcup_{j=1}^k b_j \right)$ (the item with smallest index that is packed in a bin). Let $t = (\Gamma - 1)k - (q - 1)$. By definition of q and (2), items in $[q - 1]$ must be in T , and thus it means that the triple $(\emptyset, \{b_1, \dots, b_k\}, T \setminus [q - 1])$ is a feasible solution of G- Γ RBP-T $(q, t, 1)$ with cost 0. $\square$

Lemma 8. *Let us fix $\mathcal{I}$ an input of Γ RBP and k an integer. For any $q \in [(\Gamma - 1)k]$, $t = (\Gamma - 1)k - (q - 1)$, given an almost feasible 0 cost solution for G- Γ RBP-T instance $\mathcal{I}' = (q, t, 1)$, a solution for Γ RBP-T instance $(\mathcal{I}, k, \Gamma k)$ can be determined in polynomial time.*

Proof. Let $(\emptyset, \{b_1, \dots, b_k\}, T)$ be an optimal solution of G- Γ RBP-T instance $\mathcal{I}' = (q, t, 1)$ with cost 0. For $j \in [k]$, let $b'_j = b_j \setminus \max(b_j)$. Let $T' = T \cup [q - 1] \cup \bigcup_{j=1}^k \max(b_j)$. We now have $\tilde{f}(b'_j) \leq 1$ for any $j \in [k]$ (as b_j exceeds by at most one item), and $|T'| \leq \Gamma k$, so $\{b'_1, \dots, b'_k, T'\}$ is a yes-instance for Γ RBP-T $(\mathcal{I}, k, \Gamma k)$ thus concluding the proof. $\square$

The above lemmas imply that any algorithm outputting an (approximately) optimal almost-feasible solution for G- Γ RBP-T can be used to devise approximation algorithm for Γ RBP with small values, which is illustrated by Algorithm 2.

```

1  $s^* \leftarrow ([n], \emptyset, \emptyset), lb \leftarrow 1, ub \leftarrow n$ 
2 while  $lb \neq ub$  do
3    $k \leftarrow \lceil \frac{lb+ub}{2} \rceil$ 
4    $\mathcal{I}' \leftarrow (1, \Gamma(k-1), 1)$ 
5   Let  $s = (L, \{b_1, \dots, b_k\}, T)$ , be an optimal almost-feasible solution
     returned by the G- $\Gamma$ RBP-T algorithm given input  $\mathcal{I}'$ 
6   if  $c(s) = 0$  then
7      $ub \leftarrow k$ 
8      $s^* \leftarrow s$  (accordingly  $T^* \leftarrow T$ )
9   else
10     $lb \leftarrow k + 1$ 
11 Pack items in  $T^*$  into bins  $b_{k+1}^*, \dots, b_k^*$  according to Observation 1.
    return:  $b_1^*, \dots, b_k^*$ 

```

Algorithm 2: Algorithm for Γ RBP with small values.

Proposition 2. *Algorithm 2 is a 3-approximation for Γ RBP with small values.*

Proof. Let A be a G - Γ RBP-T algorithm that is guaranteed to output an optimal almost-feasible solution given instance (q, t, ℓ) , and whose output will be denoted by $c(A(q, t, \ell))$. Let $k^* = \text{OPT}(\mathcal{I})$ be the optimal value of Γ RBP. Lemma 5 implies that $(\mathcal{I}, k^*, (\Gamma - 1)k^*)$ is a *yes*-instance of Γ RBP-T, so Lemma 7 implies that there exists $q \in [(\Gamma - 1)k^*]$ and $t = (\Gamma - 1)k^* - (q - 1)$ for which $\text{OPT}(q, t, 1) = 0$. Thus, the assumption implies that A outputs an almost feasible solution s that satisfies $c(s) \leq 0$, so that $c(s^*) \leq k^*$. Applying Lemma 8, we can compute a solution of $(\mathcal{I}, k^*, \Gamma k^*)$ for Γ RBP-T, which corresponds to a solution using $3k^*$ bins for Γ RBP following Lemma 6. $\square$

The next section describes a dynamic programming algorithm that outputs an optimal almost-feasible solution for G - Γ RBP-T.

4.3 A DP algorithm for G - Γ RBP-T

We now define a DP scheme that given instance $\mathcal{I}' = (q, t, \ell)$ provides an *almost feasible* solution s with $c(s) \leq \text{OPT}(\mathcal{I}')$. We start by providing a gentle description before formally defining the algorithm and proving its correctness. Let s^* be an optimal solution for $\mathcal{I}'$, with bins ordered in non-increasing deviation size. The DP algorithm starts by enumerating $g = (q', t') \in \llbracket q + 1, n \rrbracket \times [t]_0$. One of the enumerated g must necessarily correspond to the optimal solution s^* in the following sense:

- $q' = \min(b_{\ell+1}^*)$
- t' is the number of items trashed from X' in s^* , where $X' = \llbracket q, q' - 1 \rrbracket$

Because of the ordering of the bins in the optimal solution s^* , the items of X' must be packed in bins b_ℓ^* , L^* or T^* . As the DP algorithm mimics the optimal solution, it packs X' in b_ℓ , L and T . Specifically, the DP algorithm:

- Packs q to b_ℓ (as required by the corresponding constraint of G - Γ RBP-T).
- Packs the remaining t' largest nominal value items of X' in the trash.
- Packs the remaining items of X' into b_ℓ until $\tilde{f}(b_\ell) > 1$ or $X' = \emptyset$.
- Packs the remaining items of X' into L until $X' = \emptyset$.

We discuss next where the others items (of $\llbracket q', n \rrbracket$) are packed. Notice that in s^* , bin b_ℓ^* may contain items of $\llbracket q', n \rrbracket$, and thus the DP algorithm may also have to pack items of $\llbracket q', n \rrbracket$ into b_ℓ . To allow for that possibility, we postpone the decision of which items of $\llbracket q', n \rrbracket$ to pack into b_ℓ . Specifically, let Δ_ℓ be the size of the empty space in b_ℓ after packing X' as described above, and let $L^{X'} = L \cap X'$. After the previous steps, the DP algorithm makes a recursive call to get a solution $\tilde{s}$ that packs $\llbracket q', n \rrbracket$ into regular bins, a trash, and a leftover itemset $\tilde{b}_0$. So far solution $\tilde{s}$ has not benefited from the empty space Δ_ℓ . However, we can unpack items from $\tilde{b}_0$ to b_ℓ while ensuring that these items do not deviate in b_ℓ (as all these items have index greater than q).

Our DP scheme is defined by Algorithm 3. An iteration of this algorithm is further illustrated in Figure 1.

```

1  $s \leftarrow (\emptyset, \emptyset, \emptyset)$  // Where  $c((\emptyset, \emptyset, \emptyset)) = \infty$ 
2 if  $\ell = k$  then
3    $X' \leftarrow \llbracket q, n \rrbracket$ 
4    $T \leftarrow \{\min(t, n - q) \text{ largest items in terms of nominal value in } X'\}$ 
5   Pack  $X' \setminus T$  in  $b_k$  until  $\tilde{f}(b_k) > 1$  or  $X' \setminus T = \emptyset$ 
6   Pack the remaining items in  $L$ 
7   return:  $s = (L, \{b_k\}, T)$ 
7 for  $g = (q', t') \in \llbracket q + 1, n \rrbracket \times [\min(t, q' - q)]_0$  do
8   Pack  $q$  in  $b_\ell$ 
9    $X' \leftarrow \llbracket q + 1, q' - 1 \rrbracket$ 
10   $T' \leftarrow \{t' \text{ largest items in terms of nominal value in } X'\}$ 
11  Pack  $X' \setminus T'$  in  $b_\ell$  until  $\tilde{f}(b_\ell) > 1$  or  $X' \setminus T' = \emptyset$ 
12  Pack the remaining items in  $X' \setminus T'$  in  $L$ 
13   $\tilde{s} = (\tilde{b}_0, \{\tilde{b}_{\ell+1}, \dots, \tilde{b}_k\}, \tilde{T}) \leftarrow \text{DP}(q', t - t', \ell + 1)$ 
14  Unpack  $\tilde{b}_0$  into  $b_\ell$  until  $\tilde{f}(b_\ell) > 1$  or all items of  $\tilde{b}_0$  are unpacked,
    then unpack the potentially remaining items of  $\tilde{b}_0$  into  $L$ .
15   $s_g \leftarrow (L, \{b_\ell, \tilde{b}_{\ell+1}, \dots, \tilde{b}_k\}, \tilde{T} \cup T')$ 
16  if  $c(s_g) < c(s)$  then  $s \leftarrow s_g$ 
return:  $s$ 

```

Algorithm 3: DP(q, t, ℓ)

Fig. 1. DP algorithm handling guess (q', t') , starting from item q .

Let us introduce notations to describe the packing of the DP. Let $b_\ell^{X'} = b_\ell \cap X'$ be the items of X' packed in b_ℓ by the DP, and let $\Delta_\ell = 1 - \tilde{f}(b_\ell^{X'})$ (Δ_ℓ could be negative). We define similarly $L^{X'} = L \cap X'$. Let $c_0 = \bar{a}(L^{X'})$. Notice that c_0

corresponds to the total non deviating size of items packed in L after Step 12 of Algorithm 3. The following lemma establishes the correctness of the algorithm.

Lemma 9. *For any $\mathcal{I}' = (q, t, \ell)$ input of G- Γ RBP-T, $\text{DP}(\mathcal{I}')$ determines an optimal almost feasible solution.*

Proof. The proof is by induction on ℓ .

$\ell = k$. Consider a solution $s = (L, \{b_k\}, T)$ output by the algorithm and optimal solution $s^* = (b_0^*, \{b_k^*\}, T^*)$. In step 5, either $b_\ell = \{q\} \cup X' \setminus T$, in which case $L = \emptyset$ and $\bar{a}(L) \leq \bar{a}(b_0^*)$, or $\tilde{f}(b_k) > 1 \geq \tilde{f}(b_k^*)$. Then, as $q = \min(b_k) = \min(b_k^*)$ (since both solutions must satisfy this constraint) it implies that $\Gamma \hat{a}_1(b_k^*) = \Gamma \hat{a}_1(b_k)$, and so $\bar{a}(b_k) > \bar{a}(b_k^*)$. Then, by the definition of T (in step 4 of the algorithm), also $\bar{a}(T) \geq \bar{a}(T^*)$, so it must be that $\bar{a}(L) < \bar{a}(b_0^*)$.

Inductive step. Suppose now that, for some $\hat{\ell} \in [k-1]$, the claim is true for $\ell = \hat{\ell} + 1$ and that $s^* = (b_0^*, \{b_{\hat{\ell}}^*, \dots, b_k^*\}, T^*)$ is optimal for G- Γ RBP-T $(q, t, \hat{\ell})$. Let $q^* = \min(b_{\hat{\ell}+1}^*)$, let $X^* = \llbracket q, q^* - 1 \rrbracket$ and observe that $X^* \neq \emptyset$. Next, observe that $(b_0^* \setminus X^*, \{b_{\hat{\ell}+1}^*, \dots, b_k^*\}, T^* \setminus X^*)$ is optimal for G- Γ RBP-T $(q^*, \hat{t}, \hat{\ell} + 1)$ where $\hat{t} = |T^* \setminus X^*| \leq t$. Otherwise, there would exist some feasible $(\tilde{L}, \{\tilde{b}_{\hat{\ell}+1}, \dots, \tilde{b}_k\}, \tilde{T})$ that has $\bar{a}(\tilde{L}) < \bar{a}(b_0^* \setminus X^*)$ and then $(\tilde{L} \cup (b_0^* \cap X^*), \{\{q\} \cup (b_{\hat{\ell}} \cap X^*), \tilde{b}_{\hat{\ell}+1}, \dots, \tilde{b}_k\}, \tilde{T} \cup (T^* \cap X^*))$ would be feasible for $(q, t, \hat{\ell})$ with $\bar{a}(\tilde{L} \cup (b_0^* \cap X^*)) < \bar{a}(b_0^*)$, thereby contradicting the optimality of s^* .

By the inductive hypothesis $\bar{a}(\tilde{L}) \leq \bar{a}(b_0^* \setminus X^*)$ for some almost feasible $(\tilde{L}, \{\tilde{b}_{\hat{\ell}+1}, \dots, \tilde{b}_k\}, \tilde{T})$ that is output by the algorithm for $(q^*, \hat{t}, \hat{\ell} + 1)$, where $q^* > q$ and $\hat{t} = |T^* \setminus X^*| \leq t$. Therefore,

$$\bar{a}(b_0^*) = \bar{a}(b_0^* \setminus X^*) + \bar{a}(b_0^* \cap X^*) \geq \bar{a}(\tilde{L}) + \bar{a}(b_0^* \cap X^*). \quad (4)$$

Let us now consider the iteration $g = (q^*, t - \hat{t})$ of the main loop (starting in step 7) and L and $b_{\hat{\ell}}$ the corresponding bins of s_g . By steps 8, 11 and 14 of the algorithm, $L \setminus \tilde{L} = (X^* \setminus T') \setminus b_{\hat{\ell}}$. In line 11 of the algorithm if $\tilde{f}(b_{\hat{\ell}}) \leq 1$, then $L \subseteq \tilde{L}$ and $L \setminus \tilde{L} = \emptyset$. Otherwise $\tilde{f}(b_{\hat{\ell}}) > 1 \geq \tilde{f}(b_{\hat{\ell}}^*)$. Recalling that both solutions must satisfy the given constraint, $q = \min(b_{\hat{\ell}}) = \min(b_{\hat{\ell}}^*)$, it implies that $\Gamma \hat{a}_1(b_{\hat{\ell}}) = \Gamma \hat{a}_1(b_{\hat{\ell}}^*)$. Thus, $\bar{a}(b_{\hat{\ell}}) > \bar{a}(b_{\hat{\ell}}^*)$. Also, by definition of T' , $\bar{a}(T') \geq \bar{a}(T^* \cap X^*)$. Thus, $\bar{a}(L \setminus \tilde{L}) \leq \bar{a}(b_0^* \cap X^*)$. Combined with (4) it finally implies that $\bar{a}(b_0^*) \geq \bar{a}(\tilde{L}) + \bar{a}(b_0^* \cap X^*) \geq \bar{a}(L \setminus \tilde{L}) + \bar{a}(\tilde{L}) = \bar{a}(L)$. $\square$

Lemma 10. *Suppose that $a_i, \hat{a}_i \leq 1/\Gamma$ for all $i \in [n]$. Then, Algorithm 2, with Algorithm 3 as a subroutine for determining almost-feasible optimal solutions for G- Γ RBP-T, is a 3-factor approximation algorithm for Γ RBP. Further, the running time complexity bound of this algorithm is in $\mathcal{O}(n^6 \log(n))$.*

Proof. Proposition 2 and Lemma 9 applied to Algorithm 3 imply that there exists a 3-approximation for Γ RBP with small values. The different number of

inputs of the DP is $\mathcal{O}(n^3)$, and the running time for a fixed input is $\mathcal{O}(n^3)$. Then, since the binary search of Algorithm 2 takes $\mathcal{O}(\log n)$ iterations it follows that the total running time is in $\mathcal{O}(n^6 \log n)$. $\square$

Now, Lemma 10 together with Proposition 1 imply the following theorem using the $\rho_{bp} = \frac{3}{2}$ -approximation first-fit decreasing (FFD) algorithm, whose running time complexity bound is clearly dominated by that of the algorithm for Γ RBP with small values, for standard bin packing; see for example [26].

Theorem 4. *There exists a 4.5-approximation algorithm for Γ RBP with an $\mathcal{O}(n^6 \log(n))$ runtime complexity bound.*

In particular, combining FFD to to pack large items $i \in [n]$, for which with $\bar{a}_i > 1/\Gamma$ or $\hat{a}_i > 1/\Gamma$ and Algorithm 2 using Algorithm 3 as a subroutine to pack the other (small) items, is an algorithm that satisfies the claim of Theorem 4. Also note that while not a focus of the current paper, in the asymptotic setting (as n tends to be large) using instead an asymptotic FPTAS for packing the large items, for $\epsilon > 0$, an asymptotic approximation of $4 + \epsilon$ could be guaranteed in running time that is polynomial in n and $1/\epsilon$. Finally, although this result establishes a constant factor approximation for our problem when Γ is a part of the input, it can be observed that in the special case that $\Gamma \leq 2$, the next-fit approximation established by Theorem 1 may be preferred as simple $\mathcal{O}(n \log n)$ approximation with a $2\Gamma \leq 4$ approximation guarantee.

References

1. Adler, M., Gibbons, P., Matias, Y.: Scheduling space-sharing for internet advertising. *Journal of Scheduling* **5**, 103–119 (2002)
2. Aissi, H., Bazgan, C., Vanderpooten, D.: Approximation of minmax and minmax regret versions of some combinatorial optimization problems. *Discrete Optimization* **179**, 281–290 (2007). <https://doi.org/10.1016/j.ejor.2006.03.023>
3. Aissi, H., Bazgan, C., Vanderpooten, D.: Min-max and min-max regret versions of combinatorial optimization problems: A survey. *European Journal of Operational Research* **197**(2), 427–438 (2009)
4. Álvarez-Miranda, E., Ljubic, I., Toth, P.: A note on the Bertsimas & Sim algorithm for robust combinatorial optimization problems. *4OR* **11**(4), 349–360 (2013). <https://doi.org/10.1007/s10288-013-0231-6>
5. Balogh, J., Békési, J., Dósa, G., Sgall, J., van Stee, R.: The optimal absolute ratio for online bin packing. In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms* (2015)
6. Balogh, J., Békési, J., Dósa, G., Epstein, L., Levin, A.: A New and Improved Algorithm for Online Bin Packing. In: Azar, Y., Bast, H., Herman, G. (eds.) *ESA. LIPIcs*, vol. 112, pp. 5:1–5:14. Dagstuhl, Germany (2018)
7. Basu Roy, A., Bougeret, M., Goldberg, N., Poss, M.: Approximating robust bin-packing with budgeted uncertainty (2019), *Algorithms and Data Structures Symposium (WADS)*
8. Ben-Tal, A., Nemirovski, A.: Robust convex optimization. *Mathematics of Operations Research* **23**(4), 769–805 (1998). <https://doi.org/10.1287/moor.23.4.769>

9. Ben-Tal, A., Ghaoui, L.E., Nemirovski, A.: Robust Optimization. Princeton University Press (2009). <https://doi.org/10.1515/9781400831050>
10. Bertsimas, D., Sim, M.: Robust discrete optimization and network flows. *Math. Program.* **98**(1-3), 49–71 (2003)
11. Bertsimas, D., Sim, M.: The price of robustness. *Operations Research* **52**(1), 35–53 (2004)
12. Bougeret, M., Jansen, K., Poss, M., Rohwedder, L.: Approximation results for makespan minimization with budgeted uncertainty. CoRR **abs/1905.08592** (2019), <http://arxiv.org/abs/1905.08592>, accepted in WAOA
13. Bougeret, M., Pessoa, A.A., Poss, M.: Robust scheduling with budgeted uncertainty. *Discrete Applied Mathematics* **261**, 93–107 (2019)
14. Buchheim, C., Kurtz, J.: Robust combinatorial optimization under convex and discrete cost uncertainty. *EURO J. Computational Optimization* **6**(3), 211–238 (2018). <https://doi.org/10.1007/s13675-018-0103-0>, <https://doi.org/10.1007/s13675-018-0103-0>
15. Christensen, H.I., Khan, A., Pokutta, S., Tetali, P.: Approximation and online algorithms for multidimensional bin packing: A survey. *Computer Science Review* **24**, 63–79 (2017)
16. Dexter, F., Macario, A., Traub, R.D.: Which algorithm for scheduling add-on elective cases maximizes operating room utilization? use of bin packing algorithms and fuzzy constraints in operating room management. *Anesthesiology* **91**, 1491–1500 (Nov 1999)
17. Goetzmann, K., Stiller, S., Telha, C.: Optimization over integers with robustness in cost and few constraints. In: WAOA. pp. 89–101 (2011)
18. Goldberg, N., Karhi, S.: Online packing of arbitrary sized items into designated and multipurpose bins. *European Journal of Operational Research* **279**(1), 54–67 (2019)
19. Gounaris, C.E., Wieseemann, W., Floudas, C.A.: The robust capacitated vehicle routing problem under demand uncertainty. *Operations Research* **61**(3), 677–693 (2013)
20. Johnson, D.S., Demers, A., Ullman, J.D., Garey, M.R., Graham, R.L.: Worst-case performance bounds for simple one-dimensional packing algorithms. *SIAM Journal on Computing* **3**(4), 299–325 (1974)
21. Karmarkar, N., Karp, R.M.: An efficient approximation scheme for the one-dimensional bin-packing problem. In: Proc. 23rd Annual Symp. Foundations of Computer Science (sfcs 1982). pp. 312–320 (Nov 1982)
22. Kasperski, A., Zieliński, P.: On the approximability of minmax (regret) network optimization problems. *Information Processing Letters* **109**(5), 262–266 (2009)
23. Kasperski, A., Zielinski, P.: On the approximability of robust spanning tree problems. *Theor. Comput. Sci.* **412**(4-5), 365–374 (2011). <https://doi.org/10.1016/j.tcs.2010.10.006>
24. Pessoa, A.A., Poss, M., Sadykov, R., Vanderbeck, F.: Branch-and-cut-and-price for the robust capacitated vehicle routing problem with knapsack uncertainty (2020), <https://hal.inria.fr/hal-01958184/>
25. Poss, M.: Robust combinatorial optimization with knapsack uncertainty. *Discrete Optimization* **27**, 88–102 (2018). <https://doi.org/10.1016/j.disopt.2017.09.004>
26. Simchi-Levi, D.: New worst-case results for the bin-packing problem **41**, 579–585 (1994)
27. Song, G., Kowalczyk, D., Leus, R.: The robust machine availability problem–bin packing under uncertainty. *IIE Trans.* pp. 1–35 (2018)

28. Tadayon, B., Smith, J.C.: Algorithms and complexity analysis for robust single-machine scheduling problems. *J. Scheduling* **18**(6), 575–592 (2015)