

Algorithms and Complexity Results for the 0–1 Knapsack Problem with Group Fairness

Enrico Malaguti¹, Paolo Paronuzzi¹, and Alberto Santini^{2,3}

¹Department of Electrical, Electronic and Information Engineering, Alma Mater Studiorum University of Bologna, Viale del Risorgimento, 2, 40136 Bologna, Italy

²Department of Economics and Business, Universitat Pompeu Fabra, Carrer Trias-Fargas, 25–27, 08005 Barcelona, Spain

³Data Science Centre, Barcelona School of Economics, Carrer Trias-Fargas, 25–27, 08005 Barcelona, Spain

Abstract

This paper presents an exact algorithm for a generalisation of the classical 0–1 Knapsack Problem, called the Knapsack Problem with Group Fairness. In this problem, items are partitioned into classes, and fairness constraints affect the number of items that can or must be chosen from each class. The problem was introduced by Patel et al. (2021), where approximation algorithms are discussed. This paper describes the first exact solution approaches for the Knapsack Problem with Group Fairness, based on integer linear programming formulations and a dynamic programming algorithm. The latter allows us to establish the complexity of the problem. Finally, a set of computational experiments on benchmark instances derived from the knapsack literature compare the effectiveness of the alternative solution approaches.

Keywords: packing; knapsack problems; fairness; integer linear programming; dynamic programming.

1 Introduction

This paper presents an exact algorithm for a generalisation of the classical 0–1 Knapsack Problem (KP) in which items are partitioned into classes, and fairness constraints affect the number of items that can or must be chosen from each class. In the KP (Martello and Toth 1990), one is given n items, each one having a profit $p_j \in \mathbb{N}$ and weight $w_j \in \mathbb{N}$, to be packed into a knapsack of capacity $c \in \mathbb{N}$ while maximising the packed profit. Patel et al. (2021) introduced a family of problems, collectively called KP with Group Fairness (KPGF), in which the items are partitioned into ℓ classes J_1, \dots, J_ℓ . The main idea behind the KPGF is that the items packed in the knapsack should represent each class fairly. To this end, the authors consider three different problems in which each class has associated upper and lower bounds on the total weight, total profit or number of items chosen from the class. In this paper, we generalise this idea further by associating a resource consumption value $h_j \in \mathbb{N}$ to each item and imposing lower and upper bounds on the total resource consumption of packed items of each class. The three problems introduced by Patel et al. (2021) then correspond to the special cases in which, for each item j , $h_j = w_j$, $h_j = p_j$, and $h_j = 1$, respectively.

The KPGF can be used to determine a fair allocation of economic resources. Patel et al. (2021) and the references therein list some of its most important applications. Here, we highlight one use case which is particularly important due to its timeliness and social impact: participatory budgeting (PB). A participatory budget is a decision-making process in which ordinary citizens decide how to allocate part of a local government’s budget. Contemporary PB was pioneered

in the '90s by the Brazilian city of Porto Alegre (De Sousa Santos 1998; Novy and Leubolt 2005) and is now used in thousand of major and small cities and communities around the world, such as New York City (Su 2017), Barcelona (Barandiaran et al. 2024), Bologna, Stuttgart and Zaragoza (Manes-Rossi et al. 2023). In a typical PB process at the municipal level, citizens propose interventions they would like to see implemented in their neighbourhood, such as a new bike lane or renovating a public space. Municipality technicians then filter the technically feasible interventions and assign an estimated budget to each one. Then, a public voting phase starts, during which citizens can vote on one or multiple projects they would like implemented. At the end of this phase, the municipality selects which interventions to perform, limited by the available budget and considering citizens' preferences. The problem of selecting the subset of interventions to implement can be modelled as a KP in which the items are the interventions, the capacity is the available budget, the weights are the intervention costs, and the profits are the number of votes gathered. Still, such a system could result in one or a few neighbourhoods receiving most interventions while others do not. Therefore, the municipality is usually interested in establishing rules that ensure that the implemented projects affect all parts of the city. In this case, the problem becomes a KPGF in which the item classes are the neighbourhoods, and many different fairness principles can be used. For example, fairness could be quantified by the number of projects (resulting in a KPGF with all $h_j = 1$), the budget (resulting in a KPGF with $h_j = w_j$), or some other criterion, such as the number of positively affected citizens (resulting in a KPGF in which the values h_j do not coincide with other problem data).

In this paper, we generalise the KPGF introduced by Patel et al. (2021) and present models and an exact solution algorithm for this problem. Moreover, we settle the open question on the complexity of the KPGF by proving that it is weakly \mathcal{NP} -hard. After reviewing the related literature in Section 2, we introduce a compact and an extended formulation based on Dantzig-Wolfe reformulation in Section 3. The extended formulation has a number of variables which is exponential in n . However, in Theorem 3.1 we show that only a pseudopolynomial number of them are required to solve the KPGF to optimality. We exploit this property to devise a solution algorithm outperforming the direct application of a state-of-the-art mixed-integer programming solver to the compact formulation, as shown by the results presented in Section 4. In Section 6 we summarise the main findings and propose further research directions. Finally, we publish the first instance set for the KPGF and make the source code of our solver available (Santini and Paronuzzi 2025).

2 Literature review

The KPGF was introduced by Patel et al. (2021). The authors consider six problems stemming from the combination of two base problems and three notions of fairness. The two base problems are the classical KP and the min-Knapsack problem (Babat 1975; Csirik et al. 1991). As mentioned in the introduction, the three notions of fairness impose bounds on the number of items chosen in each class, their total weight, or their total profit. They prove that finding a feasible KPGF solution is weakly \mathcal{NP} -hard when the resource consumption coincides with the weight or the profit. They also provide one approximation algorithm for each of the six problems considered. Each algorithm can produce a sub-optimal solution, violate a resource consumption bound, violate the knapsack capacity, or present a combination of these three properties. The authors do not implement the proposed algorithms, which are mostly of theoretical interest, and no computational experiments are carried out. Compared with (Patel et al. 2021), our paper presents the first exact solution approaches for the KPGF. Furthermore, our approaches solve a generalisation of the problem with an arbitrary resource and establishes complexity results for this generalised version.

Our work continues a recent research stream on generalisations of the KP; see, e.g., Malaguti

et al. (2019), Clautiaux et al. (2021), Schäfer et al. (2021), Al-douri et al. (2021), Monaci et al. (2022), Boeckmann et al. (2023), and Santini and Malaguti (2024); see also the recent surveys by Cacchiani et al. (2022a) and Cacchiani et al. (2022b).

In the rest of this section, we place our work in the current literature along three axes: knapsack variants with similar features (Section 2.1), other optimisation problems with fairness (Section 2.2), and other optimisation problems with similar applications (Section 2.3).

2.1 Knapsack variants with similar features

Cappanera and Trubian (2005) introduced the Multidemand Multidimensional Knapsack Problem (MMKP). In this problem, profits can be either positive or negative, and each item is associated with multiple resources with both upper and lower bounds. The KPGF can be seen as a special case of the MMKP. In particular, we can transform any KPGF instance with ℓ classes into an MMKP instance with $\ell + 1$ resources in the following way. The KPGF objective function is valid for the MMKP. The KPGF’s capacity constraint corresponds to an MMKP resource constraint in which the resource is the weight, the lower bound is zero, and the upper bound is the capacity. Finally, we create an MMKP resource for each KPGF class. Given a class k , the resource consumption of an item j in the MMKP is h_j if $j \in J_k$ in the KPGF or zero otherwise. The resource upper and lower bounds are \underline{h}_k and \bar{h}_k , respectively.

Cappanera and Trubian (2005) propose a heuristic two-stage Tabu Search algorithm for the MMKP. In the first stage, the algorithm explores infeasible solutions until a feasible one is found. In the second stage, the feasible solution is improved upon. The authors also remark that general-purpose mixed-integer solvers struggle to find feasible solutions for instances with a few hundred of items and tens of inequalities. We are not aware of any specialised exact algorithm for the MMKP.

Xu (2013) studies the Knapsack Problem with a Minimum Filling Constraint (KPMFC). The KPMFC extends the classical KP by imposing a minimum weight to pack. It can be seen as a special case of the MMKP when there is only one resource, the weight, with both a lower and upper bound. The author proves that when the lower bound equals the upper bound, the problem of finding a feasible solution for the KPMFC is weakly \mathcal{NP} -hard; the proof is by reduction from the partition problem (Garey and Johnson 1979). Xu (2013) then considers the special case where the bound is strictly smaller than the capacity and develops an approximation algorithm. The paper, however, does not include computational results; therefore, an empirical comparison with the method of Cappanera and Trubian (2005) is impossible. The KPMFC can also appear in algorithms that solve other optimisation problems. For example, the Dynamic Programming (DP) algorithm presented by Furini et al. (2017) for the Minimum-Cost Maximal Knapsack Problem requires solving multiple instances of the KPMFC.

Bettinelli, Ceselli, et al. (2008) consider a related problem in the bin packing context, where one has to minimise the total cost of the bins required to pack all items. Bins can be of different types; the type chosen determines the cost, capacity, and minimum weight that must be packed into each used bin. The authors propose a branch-and-price algorithm whose pricing subproblem is a restricted version of the MMKP with a single resource and both lower and upper bounds. The authors perform heuristic pricing and then resort to a black-box mixed-integer problem when the heuristic fails to produce a negative reduced cost column. The authors also mention that they compared the black-box solver with a DP algorithm, ultimately favouring the former, but they do not provide any details about the DP algorithm.

Finally, we mention the Multidimensional Multiple-Choice Knapsack Problem. In this problem, items are partitioned into classes, and each item is associated with a profit and several resources. The objective is to select one item in each class to maximise the collected profit and ensure

that the lower and upper bounds on each resource are respected. This problem was introduced, with a different name, by van de Velde and Worm (1994). The current state-of-the-art exact algorithm is the YACA procedure by Mansini and Zanotti (2020).

2.2 Optimisation problems with fairness constraints

Fairness is a timely topic that has been subject to investigation in many areas of knowledge, including combinatorial optimisation problems. It usually arises in resource allocation, where a third party must assign desirable resources to a set of agents. In general, there are several ways of defining and quantifying fairness (see, e.g., Luss 1999). For example, one can minimise the difference between the agents who receive the most and the least resources. In assignment problems, a popular criterion is envy-freeness (Foley 1966; Varian 1974), i.e., devising an assignment such that each agent believes their assignment is no worse than any other agent’s. In the KPGF, agents are classes, resources are items, and assigning a resource to an agent means packing an item of the corresponding class in the knapsack. The notion of fairness in the KPGF states that no agent can receive too many or too few resource units.

We remark here that the term “Group Fairness” in the KPGF as introduced by Patel et al. (2021) slightly differs from the common usage in the literature. In its most common meaning (Conitzer et al. 2019; Benabbou, Chakraborty, Elkind, et al. 2019), group fairness refers to the concept of fair allocation among groups of agents rather than among individual agents. The specific way in which group fairness is modelled can vary greatly. For example, this notion can apply *on top* of individual-agent fairness or *instead of* it; in particular applications (Bera et al. 2019), group fairness can also apply to overlapping groups of agents (that is, an agent can belong to multiple groups). On the other hand, the KPGF does not include fairness for groups of agents, i.e., for groups of classes taken together.

Finally, we mention that another stream of research investigates the “price of fairness”, i.e., the loss in resource allocation efficiency when including fairness constraints. This concept was introduced by Bertsimas et al. (2011) and has been widely employed since then (see, e.g., Lodi, Malaguti, et al. 2010; Bei et al. 2021; Lodi, Sankaranarayanan, et al. 2023).

2.3 Optimisation problems with similar applications

Resource allocation problems often arise in the public sector. For example, the government of Singapore has faced the problem of assigning public housing to families while ensuring that each neighbourhood houses a balanced mix of ethnic groups (Benabbou, Chakraborty, Ho, et al. 2018).

Relevant to our key application in PB, Serramia et al. (2019) study the problem of optimally selecting projects in Barcelona’s participatory budget. Similar to our problem, the authors use a knapsack-like model in which the proposals are the items, the costs are the weights, and the votes received are the profits. The authors do not consider fairness; they focus on avoiding the simultaneous selection of conflicting projects. For example, if a project proposes to create a children’s playground in a public square and another proposes to build a skate park in the same space, only one can be implemented. The resulting optimisation problem is a 0–1 Knapsack Problem with Conflict Graph (KPCG; see, e.g., Bettinelli, Cacchiani, et al. 2017). The authors solve this problem with a black-box solver and compare the results with the current selection method used in Barcelona. This method ranks the projects in descending order of votes and selects them up to the available budget, skipping a project if another incompatible one has already been selected. In optimisation terms, this corresponds to a greedy KPCG heuristic sorting the items by profit. The authors show that an approach based on optimal KPCG solutions can increase the number of accepted proposals and their total support, measured as the total number of votes gathered by implemented projects.

3 Mathematical formulations

A compact integer programming formulation for the KPGF uses binary variables $x_j \in \{0, 1\}$ for each $j \in \{1, \dots, n\}$ taking value one if and only if item j is packed. The formulation reads as follows.

$$\max \quad \sum_{j=1}^n p_j x_j \quad (1a)$$

$$\text{subject to} \quad \sum_{j=1}^n w_j x_j \leq c \quad (1b)$$

$$\sum_{j \in J_k} h_j x_j \geq h_k \quad \forall k \in \{1, \dots, \ell\} \quad (1c)$$

$$\sum_{j \in J_k} h_j x_j \leq \bar{h}_k \quad \forall k \in \{1, \dots, \ell\} \quad (1d)$$

$$x_j \in \{0, 1\} \quad \forall j \in \{1, \dots, n\}. \quad (1e)$$

The objective function maximises the collected profit, while constraint (1b) ensures that the knapsack capacity is respected. Constraints (1c) and (1d) ensure that each class's lower and upper bound on the resource consumption are satisfied.

An alternative formulation for the problem can be obtained by applying a Dantzig-Wolfe reformulation (see, e.g., Vanderbeck and Wolsey 2010) to model (1a)–(1e). The resulting model optimises over the convex hull of vectors satisfying constraints (1b)–(1d) which are encoded, for each class k , by the collection \mathcal{R}_k of valid packings of items of class k . A packing in \mathcal{R}_k is a subset of items $R \subseteq J_k$ such that

$$h_k \leq \sum_{j \in R} h_j \leq \bar{h}_k \quad \text{and} \quad \sum_{j \in R} w_j \leq c. \quad (2)$$

Remark that, if $h_k = 0$, then $\emptyset \in \mathcal{R}_k$, i.e., it is possible that no item of class k is packed. Let p_R , w_R and h_R be the profit, weight and resource consumption associated with a generic packing R , i.e.,

$$p_R = \sum_{j \in R} p_j, \quad w_R = \sum_{j \in R} w_j, \quad \text{and} \quad h_R = \sum_{j \in R} h_j.$$

For any class k and packing $R \in \mathcal{R}_k$, let $\eta_R \in \{0, 1\}$ be a variable taking the value one if the items in R are the ones packed for class k . Then, the Dantzig-Wolfe reformulation for the KPGF reads

$$\max \quad \sum_{k=1}^{\ell} \sum_{R \in \mathcal{R}_k} p_R \eta_R \quad (3a)$$

$$\text{subject to} \quad \sum_{k=1}^{\ell} \sum_{R \in \mathcal{R}_k} w_R \eta_R \leq c \quad (3b)$$

$$\sum_{R \in \mathcal{R}_k} \eta_R = 1 \quad \forall k \in \{1, \dots, \ell\} \quad (3c)$$

$$\eta_R \in \{0, 1\} \quad \forall k \in \{1, \dots, \ell\}, \forall R \in \mathcal{R}_k. \quad (3d)$$

The objective function maximises the packed profit, constraint (3b) ensures that the knapsack's capacity is not exceeded, and constraint (3c) imposes that exactly one packing is chosen for each class.

This formulation defines a Multiple-Choice Knapsack Problem (MCKP), a knapsack problem where items are partitioned into classes, and exactly one item per class must be selected. The problem is known to be weakly \mathcal{NP} -hard. In particular, an MCKP with \bar{n} items and capacity c can be solved in $O(\bar{n}c)$ operations by using a DP algorithm (see Martello and Toth 1990; Kellerer et al. 2004).

For each class k , each packing $R \in \mathcal{R}_k$ of the KPGF corresponds to a MCKP item. Because the number of packings is exponential in n , (3a)–(3d) defines a MCKP with a number of items (say, \bar{n}) exponential in the number of KPGF items (n). However, in the following proposition, we show that, given the KPGF’s special structure, considering a subset of variables of pseudo-polynomial size is sufficient to find an optimal solution. In the rest of this section, we show that this subset can be generated using a DP algorithm with pseudo-polynomial time complexity.

Proposition 3.1. *For each class k , consider the sub-collection $\bar{\mathcal{R}}_k \subset \mathcal{R}_k$ including at most one packing for each integer value $W \in \{1, \dots, c\}$, namely, a packing R maximising p_R and having $\underline{h}_k \leq h_R \leq \bar{h}_k$ and $w_R = W$, if one exists. The optimal solution of formulation (3a)–(3d) restricted to the variables associated with these sub-collections is optimal for the KPGF.*

Proof. For each class k , the sub-collection $\bar{\mathcal{R}}_k$ contains one optimal packing for each amount of the capacity which can be allocated to the class. For each class k , the sub-collection $\bar{\mathcal{R}}_k$ can be enumerated by the DP algorithm presented in the next section. \square

3.1 Dynamic programming algorithm for packing enumeration

Let $n_k = |J_k|$ be the number of items of class k , and $J_k = \{j_{k1}, \dots, j_{kn_k}\}$. Let $\zeta_k(H, W, m)$ be the largest profit that can be achieved with the first m items in J_k (i.e., $\{j_{k1}, \dots, j_{km}\}$) such that the selected items have total weight no larger than W and total resource consumption equal to H , and let $\rho_k(H, W, m) \subseteq J_k$ be an associated packing. Furthermore, denote with $H_k(W)$ the resource consumption of a packing whose total weight is W , $m = n_k$, and the profit is maximum for the given weight, i.e.,

$$H_k(W) = \arg \max_{H \in \{\underline{h}_k, \dots, \bar{h}_k\}} \zeta_k(H, W, n_k).$$

The sub-collection $\bar{\mathcal{R}}_k$ is defined as:

$$\bar{\mathcal{R}}_k = \left\{ \rho_k(H_k(W), W, n_k) : W \in \{1, \dots, c\} \right\}. \quad (4)$$

The values of $\zeta_k(H, W, m)$ and the associated packings $\rho_k(H, W, m)$ can be computed via DP by using the following recursion:

$$\zeta_k(H, W, m) = \max \begin{cases} \zeta_k(H, W, m-1), \\ p_{j_{km}} + \zeta_k(H - h_{j_{km}}, W - w_{j_{km}}, m-1). \end{cases} \quad (5)$$

In (5), we build $\bar{\mathcal{R}}_k$ by choosing, for each feasible weight W , the packing giving the maximum profit and respecting the resource consumption bounds (if any).

The first case corresponds to not packing item j_{km} . In contrast, j_{km} is packed in the second case, yielding its profit and decreasing the available capacity and the target resource consumption level. Remarking that, in the recursion, at least one and possibly all three arguments of ζ_k decrease, one can provide the following base cases that are valid for small enough values of H , W , and m :

$$\zeta_k(H, W, m) = -\infty \quad \text{if } H < \min_{i=1, \dots, m} h_{j_{ki}} \text{ or } H > \sum_{i=1}^m h_{j_{ki}}, \quad (6)$$

$$\zeta_k(H, W, m) = -\infty \quad \text{if } W < \min_{i=1, \dots, m} w_{j_{ki}}, \quad (7)$$

$$\zeta_k(H, W, 1) = \begin{cases} p_1 & \text{if } H = h_{j_{k1}} \text{ and } w_{j_{k1}} \leq W \\ -\infty & \text{otherwise.} \end{cases} \quad (8)$$

Condition (6) refers to two trivial cases in which it is not possible to meet the exact resource level H with items $\{j_{k1}, \dots, j_{km}\}$ either because H is too small or too large compared with the resources associated with the items. We remark that it might be impossible to meet level H in other cases not included in (6), namely when no combination of items can yield exactly H resource units. Still, condition $H < \min_{i=1, \dots, m} h_{j_{ki}}$ is sufficient to obtain a valid base case for the DP algorithm because it becomes active when the value of H is small enough. Condition (7) refers to the case when it is impossible to respect the capacity bound because W is too small. Condition (8) considers the first item j_{k1} only and checks whether taking it as the only item yields a feasible solution. Finally, we use the above DP algorithm to settle the question of the KPGF's complexity.

Proposition 3.2. *The KPGF is weakly NP-hard.*

Proof. Let $h = \max_k \{\bar{h}_k - \underline{h}_k\}$. For each class k , we can compute $\bar{\mathcal{R}}_k$ in $O(n_k ch)$, i.e., $O(nch)$ for all classes.

The input for the MCKP has size $|\bigcup_{k=1}^{\ell} \bar{\mathcal{R}}_k| = O(\ell c)$. Therefore, the DP algorithm for the MCKP needs $O(\ell c^2)$ operations.

The overall complexity of our approach is, thus, $O(nch) + O(\ell c^2)$. \square

3.2 Class-specific capacities

For a given class k' , let $\hat{w}_{k'}$ be the minimum weight of a subset of items of this class consuming an amount of resource between $\underline{h}_{k'}$ and $\bar{h}_{k'}$. The maximum capacity that the items of a class k can occupy in the knapsack of the KPGF is therefore reduced by the values of $\hat{w}_{k'}$ of all other classes. By computing the values $\hat{w}_{k'}$ for all classes, we define per-class capacities $\hat{c}_k := c - \sum_{k' \in \{1, \dots, \ell\} \setminus k} \hat{w}_{k'}$ which can be used in (2) and (5), replacing capacity c .

The minimum weight of subsets of items of a class k can be found by solving the following integer program, in which variable $y_j \in \{0, 1\}$ (for $j \in J_k$) indicates that item j is part of the set.

$$\min \quad \sum_{j \in J_k} w_j y_j \quad (9a)$$

$$\text{subject to} \quad \sum_{j \in J_k} h_j y_j \geq \underline{h}_k \quad (9b)$$

$$\sum_{j \in J_k} h_j y_j \leq \bar{h}_k \quad (9c)$$

$$y_j \in \{0, 1\} \quad \forall j \in J_k. \quad (9d)$$

Model (9a)–(9d) describes a min-Knapsack problem, with the addition of the classical KP constraint (9c). If problem (9a)–(9d) is infeasible for some class or its optimal solution value is larger than c , then the entire KPGF is infeasible. Indeed, in this case it is impossible to select items from such a class while respecting the resource consumption bounds and the knapsack's capacity. Moreover, the KPGF is infeasible if $\hat{c}_k < 0$ for any class k because, in this case, the knapsack does not have enough capacity to accommodate items of all classes. We remark that any lower bound on $\hat{w}_{k'}$, for a given class k' , can be used to set the value of \hat{c}_k for $k \neq k'$. Using a lower bound in the computation of $\hat{w}_{k'}$ would result in a larger (and, therefore, looser) value for \hat{c}_k . In our case, a black-box solver optimally solves problems (9a)–(9d) almost instantaneously. Therefore, our implementation uses values \hat{c}_k instead of c .

4 Computational results

In this section, we report on computational experiments comparing our approach based on the extended formulation with the direct application of a state-of-the-art solver to the compact formulation. Regarding the extended formulation, we first generate the pseudopolynomial number of variables required to solve the model to optimality by using the DP algorithm presented in Section 3.1. Recall that the resulting problem is an MCKP. We solve this problem by means of the algorithm described in Pisinger (1995a) and available at (Pisinger 1995b). We solve the compact formulation by using Gurobi version 11.0.0. Both approaches are tested on a machine with four Intel i7 processors running at 2.60GHz and 16GB RAM, running in single-thread mode. All algorithms have a 1-hour time limit.

In the remainder of this section, we describe how we generated the test instances and provide a computational comparison between the two approaches described above. The instances, their generator, the solver’s code, and the scripts used to generate the figures of this section are available at (Santini and Paronuzzi 2025).

4.1 Instance generation

We generate our KPGF instances from base KP instances. We use the thirteen generators introduced by Pisinger (2005) to generate the KP instances. The generators are: *almost strongly correlated*, *circle*, *inverse strongly correlated*, *MSTR*, *profit ceiling*, *span uncorrelated*, *span weakly correlated*, *span strongly correlated*, *strongly correlated*, *subset sum*, *uncorrelated*, *uncorrelated similar*, and *weakly correlated*. All the generators except *uncorrelated similar* require an additional parameter (denoted R) corresponding to the order of magnitude of the values used for the profits and the weights. Following Pisinger (2005), we use 1000 and 10,000 for this value, while the *uncorrelated similar* instances use a hard-coded value of 100,000.

Our generation procedure starts from a KP instance with n items, profits p_j , weights w_j , and capacity c and transforms it into a KPGF instance with ℓ classes ($\ell \leq n$). We first assign to each item i the class $1 + (i|\ell)$, where $i|\ell$ denotes the remainder of the integer division of i by ℓ . This way, we ensure that all classes are non-empty and their sizes are roughly equal. Second, we generate the resource bounds \underline{h}_k and \bar{h}_k by drawing them independently from the uniform integer distributions $\mathcal{U}(50, 150)$ and $\mathcal{U}(150, 250)$, respectively. Third, we generate the resource consumption values h_j . Denote with $r = (\sum_{j=1}^n w_j)/c$ the ratio between the total item weights and the capacity. The resource consumption of an item j belonging to class k is set to $h_j = \min\{\lceil ru/|J_k| \rceil, \bar{h}_k\}$, where notation $\lceil \cdot \rceil$ denotes integer rounding and u is drawn from the uniform integer distribution $\mathcal{U}(125, 175)$. If the total resource consumption of items in a class is strictly smaller than the class’s resource lower bound, we discard the instance because it is trivially infeasible.

The above procedure aims to maintain tight capacity and resource bounds, making the instances more challenging while avoiding generating too many infeasible instances. First, remark that r^{-1} is a rough approximation of the fraction of items that can be packed in the knapsack due to the capacity constraint. For example, if the total weight of the items is twice as large as the capacity, we can expect that only half of the items can be packed. Thus, we expect $r^{-1}n$ items to be packed in total and that $r^{-1}|J_k|$ items of class k will be packed. Remark that $\mathbb{E}[h_k] = 100$ and $\mathbb{E}[\bar{h}_k] = 200$. Using the procedure described above, the expected resource consumption of a subset of size $r^{-1}|J_k|$ of class- k items is 150, thus increasing the likelihood that a solution respecting both the capacity and each class’s resource bounds can be found.

There are $12 \times 2 + 1 = 25$ combinations of generator and parameter R : two per each generator using this parameter, plus one for the *uncorrelated similar* generator, which does not use it. For each of the 25 combinations, we consider values of $n \in \{50, 100, 200, 500, 1000, 2000\}$ and




Figure 1: Performance profile comparing the compact and extended formulations.

$\ell \in \{20, 100, 500\}$ and consider the twelve (n, ℓ) combinations with $n > \ell$. Finally, we generate ten instances for each of the $25 \times 12 = 300$ combinations of generator, R , n and ℓ , for a total of 3000 instances. Out of these, 1787 are feasible, and the remaining 1213 are provably infeasible. The results presented in the rest of this section refer to the feasible instances.

4.2 Computational results

Figure 1 shows a performance profile comparing the extended and the compact formulations. The x axis reports the time required for solving an instance, normalised by the time required by the fastest of the two algorithms. The y axis reports the percentage of instances solved to optimality within the given runtime. The dark blue line refers to the compact formulation, and the light orange line refers to the extended formulation. Interestingly, the compact formulation outperforms the extended one when instances are easy. Indeed, an analysis of the detailed results shows that the compact formulation solves approximately 69% of the instances within one second. In contrast, the extended formulation only solves about 14% of the instances in the same time. However, this 69% of the instances constitutes the majority of the instances the compact formulation can solve, no matter how much time it is given (up to a maximum of one hour). Indeed, the compact formulation can only solve 1400 out of the 1787 instances, i.e., 78.34%. By contrast, the extended formulation solves more instances when given more time: in total, it solves 1763 out of the 1787 instances, i.e., 98.66%.

To summarise, the compact formulation solves fewer instances, but when it does, it takes less time. By contrast, the extended formulation solves more instances and tends to use a larger fraction of the available time. Figures 2 and 3 provide more details and characterise the hard instances for the compact and the extended formulation, respectively. Each figure contains three charts; each chart displays the percentage of instances solved to optimality when aggregating them according to three criteria. The first criterion (left chart) is the KP generator used to build the base knapsack instance. The second criterion (top right chart) is the number of items n , and the third one (bottom right chart) is the data range R , which determines the order of magnitude of weights, profits and capacity. Figure 2 also shows the average percentage gaps reported by Gurobi when timing out while solving the compact formulation. The gaps are computed over the open instances and are displayed above each bar for which there are open instances. It is not possible to compute the gaps for the extended formulation because the timeouts always occur during the column generation phase (Section 3.1), whereas, once this phase is completed, the MCKP solver only takes a few seconds to produce the optimal solution. When the timeout




Figure 2: Results summary for the compact formulation.




Figure 3: Results summary for the extended formulation.

occurs while generating packings, we do not have a valid MCKP instance to solve, and therefore, no valid bounds. Indeed, if we removed the 1-hour time limit and let the packing generation procedure complete, we would be able to solve all instances in at most 6972 seconds using the extended formulation.

Figure 2 shows that the compact formulation is sensitive to the base knapsack instance type, solving all instances of type *weakly correlated*, *uncorrelated similar*, *uncorrelated* and *subset sum*, but less than two thirds of the *inverse strongly correlated*, *strongly correlated*, and *span weakly correlated* ones, and less than half of the *span strongly correlated* instances. As expected, increasing the number of items also increases the instance difficulty: the compact formulation solves all instances with 50 items but less than two-thirds of those with 2000 items. Finally, increasing the data range parameter from 1000 to 10,000 makes the instances harder. The compact formulation solves all instances with a data range of 100,000 to optimality. However, recall that following Pisinger (2005), only instances of type *uncorrelated similar* use this data range, and these instances are particularly easy to solve with the compact formulation despite their larger data range.

The extended formulation generally outperforms the compact one, solving more instances. However, we notice some complementarity among the two approaches: the *uncorrelated similar* instances, which were among the easiest to solve with the compact formulation, are the hardest for the extended one. The reason is that, as remarked in Theorem 3.2, the algorithm we use to solve the extended formulation has complexity $O(nch) + O(\ell c^2)$. The *uncorrelated similar* instances have a capacity c one or two orders of magnitude larger than the other instances (100,000 vs. 1000 or 10,000), thus affecting the algorithm’s runtime.

At the same time, we remark that the compact formulation’s gaps on the open instances are

generally small, due to the large profits appearing in the (Pisinger 2005) instances. When classifying the instances by generator, the *inverse strongly correlated* instances result in the largest average gaps (0.015%). Interestingly, using a capacity of 1000 or 10 000 does not impact the average gaps significantly. Regarding the number of items, while larger instances are harder to solve for the compact formulation, they do not result in larger average gaps. As mentioned above, this is due to the fact that instances with more items have larger profits, resulting in larger absolute values for the upper and lower bounds.

With respect to the extended formulation, as mentioned above, if we increased the time limit from one to two hours, this formulation would close all instances. This further confirms the previous observation prompted by the analysis of the performance profile (Figure 1): the extended formulation seems to benefit from longer computation times more than the compact formulation.

We conclude this section by reporting that new hard knapsack instances have been recently proposed by Jookan et al. (2022). These instances are inherently difficult and have even larger values for the item profits and the knapsack capacities compared to the instances of Pisinger (2005). In particular, they have capacities of one million, 100 million, and 10 billion. We generated a set of KPGF instances starting from the (Jookan et al. 2022) instances with a procedure analogous to that described in Section 4.1. We used the 01-KP instances with the capacity of the order of magnitude relevant for our application, i.e., one million (see also Section 5). Interestingly, the extended formulation could solve all 1685 feasible instances in at most 376 seconds, whereas the compact one only solved 81.5% of the instances within the 1-hour time limit. The generator, the generated instances, the results, and their analysis are available at (Santini and Paronuzzi 2025).

5 Case study

To illustrate the practical relevance of the KPGF, we analyse the participatory budgeting allocation process that took place in Barcelona, Spain, in 2020 and whose projects were implemented during 2020–2023. This case study shows that considering fairness in knapsack problems can result in alternative and improved solutions supporting real-world decision-making. The city of Barcelona allocates specific budgets to each district and implements a selection process in which projects are approved based on their vote counts until the district’s budget is exhausted.

We examine how alternative optimisation approaches may lead to improved outcomes while respecting the same total budget. Specifically, we compare the solution implemented by Barcelona’s municipality (As-is) with two alternative methods. The first method (KP) optimally solves each district’s separate knapsack problem, with the district’s budget serving as the knapsack capacity. The second method (KPGF) aims to achieve a fair allocation across districts using the KPGF. In the KPGF approach, the total budget, i.e., the sum of the districts’ budgets, is used as the capacity of the knapsack, while we removed the resource upper bound constraints. The resource lower bound for each class is set equal to the budget allocated to the corresponding district in the As-is implementation. This ensures every district receives at least as much funding as in the currently implemented solution, while allowing optimisation to redistribute additional resources. Other policy choices could determine lower bounds differently.

Table 1 compares the three solution approaches. Column “#Proj” reports the number of financed projects, “Votes” indicates the number of votes received by selected projects, and “Cost” reports the total cost of the chosen projects in thousands of euros. Each row displays the results for a single district, while the last row contains aggregate sums across all districts. We did not report computing times, as they are negligible and out of the scope of this analysis.

Our analysis reveals that using an optimisation-based approach leads to substantial improvements in overall satisfaction (measured as the number of votes received by the selected projects).

District	Budget	As-is			KP			KPGF		
		# Proj	Votes	Cost	# Proj	Votes	Cost	# Proj	Votes	Cost
Ciutat Vella	3 400	9	12 656	3 080	11	14 254	3 380	11	13 498	3 085
Eixample	3 000	8	18 674	2 995	12	22 512	2 995	13	23 683	3 245
Gràcia	2 400	8	14 576	2 225	8	14 576	2 225	8	14 576	2 225
Horta Guinardó	3 200	7	9 069	3 187	10	11 133	3 157	10	11 006	3 197
Les Corts	2 000	8	6 870	1 970	11	8 241	1 948	11	8 205	2 023
Nou Barris	3 600	9	11 043	3 573	13	13 826	3 460	13	14 272	3 654
Sant Andreu	3 000	9	11 142	2 973	12	13 762	2 853	13	14 308	3 113
Sant Martí	3 600	7	13 625	3 480	11	18 434	3 459	12	18 940	3 669
Sants Montjuïc	3 600	5	12 960	3 570	11	21 441	3 514	12	22 144	3 639
Sarrià	2 200	6	4 309	2 093	8	4 992	2 183	8	4 860	2 143
Total	30 000	76	114 924	29 145	107	143 171	29 173	111	145 492	29 992

Table 1: Solution comparison for a real participatory budgeting process (Barcelona 2020–2023).

Specifically, the KPGF solution demonstrates superior performance across all metrics, achieving the largest total votes (145 492) and selecting more projects (111) than both the As-is (76 projects, 114 924 votes) and KP (107 projects, 143 171 votes) approaches. The larger total cost (29 992 thousands of euros) also indicates efficient budget utilization.

6 Conclusions

This work has generalised the Knapsack Problem with Group Fairness (KPGF) introduced by Patel et al. (2021). We proved that the problem is weakly \mathcal{NP} -hard and devised two mathematical formulations for it. The first one is a compact formulation using n variables, where n is the number of items in the instance. The second is an extended formulation using a number of variables which is exponential in n . However, we proved that a subset of variables of pseudopolynomial size is sufficient to solve the problem optimally. We provided a constructive algorithm based on DP to build this subset. Computational experiments on instances with 50–2000 items demonstrated the superiority of our approach based on the DP algorithm and the extended formulation.

One of the main applications of the KPGF is in participatory budgeting. In Section 5, we present a case study showing how approaches based on the KPGF can improve the currently implemented solutions. Future research could incorporate additional relevant constraints for this application. For example, Serramia et al. (2019) suggest that a decision-support tool for participatory budgeting should include conflict constraints that prevent mutually incompatible projects from being selected simultaneously.

Declarations

Funding. Alberto Santini was funded by the Spanish Ministry of Science, Innovation and Universities and by the European Social Fund Plus (ESF+) through grant RYC2022-035269-I via MCIN/AEI/10.13039/501100011033 of the “Ramón y Cajal” programme and through grant PID2024-161521OA-I00 of the programme “Proyectos de Generación de Conocimiento 2024” of the Spanish Ministry of Science, Innovation and Universities. Enrico Malaguti and Paolo Paronuzzi were funded by the Air Force Office of Scientific Research under award number FA8655-25-1-7013.

Data availability. The instances used in this paper, their generator, the solver’s code, and the scripts used to generate the figures of this section are available at (Santini and Paronuzzi 2025).

References

- Babat, L. (1975). “Linear functionals on the n -dimensional unit cube”. In: *Reports of the Academy of Sciences of the Soviet Union* 221 (4), pp. 761–762.
- Barandiaran, X., A. Calleja-López, A. Monderde, and C. Romero (2024). *Decidim, a Technopolitical Network for Participatory Democracy. Philosophy, Practice and Autonomy of a Collective Platform in the Age of Digital Intelligence*. SpringerBriefs in Political Science. Springer Cham.
- Bei, X., X. Lu, P. Manurangsi, and W. Suksompong (2021). “The Price of Fairness for Indivisible Goods”. In: *Theory of Computing Systems* 65, pp. 1069–1093.
- Benabbou, N., M. Chakraborty, E. Elkind, and Y. Zick (2019). “Fairness Towards Groups of Agents in the Allocation of Indivisible Items”. In: *Proceedings of the 28th International Joint Conference on Artificial Intelligence* (Macao, China, Aug. 10–16, 2019). Ed. by S. Kraus. AAAI Press, pp. 95–101.
- Benabbou, N., M. Chakraborty, X.-V. Ho, J. Sliwinski, and Y. Zick (2018). “Diversity Constraints in Public Housing Allocation”. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems* (Stockholm, Sweden, July 10–15, 2018). Ed. by M. Dastani, G. Sukthankar, E. André, and S. Koenig. International Foundation for Autonomous Agents and Multiagent Systems, pp. 973–981. ISBN: 978-1-4503-5649-7.
- Bera, S., D. Chakraborty, N. Flores, and M. Negahbani (2019). “Fair Algorithms for Clustering”. In: *Proceedings of the 32nd Conference on Advances in Neural Information Processing Systems (NeurIPS)* (Vancouver, Canada, Dec. 8–14, 2019). Ed. by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. B. Fox, and R. Garnett. Curran Associates, pp. 4954–4965. ISBN: 9781713807933.
- Bertsimas, D., V. Farias, and N. Trichakis (2011). “The Price of Fairness”. In: *Operations Research* 59.1, pp. 17–31.
- Bettinelli, A., V. Cacchiani, and E. Malaguti (2017). “A Branch-and-Bound Algorithm for the Knapsack Problem with Conflict Graph”. In: *INFORMS Journal on Computing* 29.3, pp. 457–473.
- Bettinelli, A., A. Ceselli, and G. Righini (2008). “A branch-and-price algorithm for the variable size bin packing problem with minimum filling constraint”. In: *Annals of Operations Research* 179, pp. 221–241.
- Boeckmann, J., C. Thielen, and U. Pferschy (2023). “Approximating single- and multi-objective nonlinear sum and product knapsack problems”. In: *Discrete Optimization* 48 Part 1. Article ID 100771.
- Cacchiani, V., M. Iori, A. Locatelli, and S. Martello (2022a). “Knapsack problems — An overview of recent advances. Part I: Single knapsack problems”. In: *Computers & Operations Research* 143. Article ID 105692.
- Cacchiani, V., M. Iori, A. Locatelli, and S. Martello (2022b). “Knapsack problems — An overview of recent advances. Part II: Multiple, multidimensional, and quadratic knapsack problems”. In: *Computers & Operations Research* 143. Article ID 105693.
- Cappanera, P. and M. Trubian (2005). “A Local-Search-Based Heuristic for the Demand-Constrained Multidimensional Knapsack Problem”. In: *INFORMS Journal on Computing* 17 (1), pp. 82–98.
- Clautiaux, F., B. Detienne, and G. Guillot (2021). “An iterative dynamic programming approach for the temporal knapsack problem”. In: *European Journal of Operational Research* 293 (2), pp. 442–456.
- Conitzer, V., R. Freeman, N. Shah, and J. Wortman Vaughan (2019). “Group Fairness for the Allocation of Indivisible Goods”. In: *Proceedings of the 33th AAAI Conference on Artificial Intelligence* (Honolulu, United States, Jan. 27–Feb. 1, 2019). AAAI Press, pp. 1853–1860.
- Csirik, J., H. Frenk, M. Labbé, and S. Zhang (1991). “Heuristics for the 0–1 min-Knapsack Problem”. In: *Acta Cybernetica* 10.1–2, pp. 15–20.

- De Sousa Santos, B. (1998). “Participatory Budgeting in Porto Alegre: Toward a Redistributive Democracy”. In: *Politics & Society* 26.4, pp. 461–510.
- Al-douri, T., M. Hifi, and V. Zissimpoulos (2021). “An iterative algorithm for the Max-Min knapsack problem with multiple scenarios”. In: *Operational Research* 21, pp. 1355–1392.
- Foley, D. K. (1966). “Resource allocation and the public sector”. PhD Thesis. Yale University.
- Furini, F., I. Ljubić, and M. Sinnl (2017). “An effective dynamic programming algorithm for the minimum-cost maximal knapsack packing problem”. In: *European Journal of Operational Research* 262.2, pp. 438–448.
- Garey, M. and D. Johnson (1979). *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company. ISBN: 0-7167-1045-5.
- Jooker, J., P. Leyman, and P. De Causmaecker (2022). “A new class of hard problem instances for the 0–1 knapsack problem”. In: *European Journal of Operational Research* 301.3, pp. 841–854.
- Kellerer, H., U. Pferschy, and D. Pisinger (2004). “The Multiple-Choice Knapsack Problem”. In: *Knapsack Problems*. Ed. by H. Kellerer, U. Pferschy, and D. Pisinger. Springer, pp. 317–347. ISBN: 978-3-642-07311-3.
- Lodi, A., E. Malaguti, and N. E. Stier-Moses (2010). “Efficient and fair routing for mesh networks”. In: *Mathematical Programming* 124, pp. 285–316.
- Lodi, A., S. Sankaranarayanan, and G. Wang (2023). “A framework for fair decision-making over time with time-invariant utilities”. In: *European Journal of Operational Research* 319 (2), pp. 456–467.
- Luss, H. (1999). “On Equitable Resource Allocation Problems: A Lexicographic Minimax Approach”. In: *Operations Research* 47.3, pp. 361–378.
- Malaguti, E., M. Monaci, P. Paronuzzi, and U. Pferschy (2019). “Integer optimization with penalized fractional values: The Knapsack case”. In: *European Journal of Operational Research* 273 (3), pp. 874–888.
- Manes-Rossi, F., I. Brusca, R. Levy Orelli, P. C. Lorson, and E. Haustein (2023). “Features and drivers of citizen participation: Insights from participatory budgeting in three European cities”. In: *Public Management Review* 25.2, pp. 201–223.
- Mansini, R. and R. Zanotti (2020). “A Core-Based Exact Algorithm for the Multidimensional Multiple Choice Knapsack Problem”. In: *INFORMS Journal on Computing* 32.4, pp. 1061–1079.
- Martello, S. and P. Toth (1990). *Knapsack Problems. Algorithms and computer implementations*. Wiley. ISBN: 978-0471924203.
- Monaci, M., C. Pike-Burke, and A. Santini (2022). “Exact algorithms for the 0–1 time-bomb knapsack problem”. In: *Computers & Operations Research* 145. Article ID 105848.
- Novy, A. and B. Leubolt (2005). “Participatory Budgeting in Porto Alegre: Social Innovation and the Dialectical Relationship of State and Civil Society”. In: *Urban Studies* 42.11, pp. 2023–2036.
- Patel, D., A. Khan, and A. Louis (2021). “Group Fairness for Knapsack Problems”. In: *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (Online, May 3–7, 2021). Ed. by F. Dignum, A. Lomuscio, U. Endriss, and A. Nowé. International Foundation for Autonomous Agents and Multiagent Systems, pp. 1001–1009. ISBN: 978-1-4503-8307-3.
- Pisinger, D. (1995a). “A minimal algorithm for the multiple-choice knapsack problem”. In: *European Journal of Operational Research* 83 (2), pp. 394–410.
- Pisinger, D. (1995b). *David Pisinger’s Optimization Codes*. URL: <https://web.archive.org/web/20250219094749/http://hjemmesider.diku.dk/~pisinger/codes.html> (visited on 02/19/2025).
- Pisinger, D. (2005). “Where are the hard knapsack problems?” In: *Computers & Operations Research* 32 (9), pp. 2271–2284.

- Santini, A. and E. Malaguti (2024). “The min-Knapsack Problem with Compactness Constraints and Applications in Statistics”. In: *European Journal of Operational Research* 312 (1), pp. 385–397.
- Santini, A. and P. Paronuzzi (2025). *Source Code and Instances for the 0–1 Knapsack Problem with Group Fairness*. URL: <https://github.com/alberto-santini/kpgf>.
- Schäfer, L., T. Dietz, M. Barbati, J. R. Figueira, S. Greco, and S. Ruzika (2021). “The binary knapsack problem with qualitative levels”. In: *European Journal of Operational Research* 289 (2), pp. 508–514.
- Serramia, M., M. Lopez-Sanchez, J. Rodríguez-Aguilar, and P. Escobar (2019). “Optimising Participatory Budget Allocation: The Decidim Use Case”. In: *Artificial Intelligence Research and Development*. Ed. by J. Sabater-Mir, V. Torra, I. Aguiló, and M. González-Hidalgo. Vol. 319. Frontiers in Artificial Intelligence and Applications. IOS Press, pp. 193–202.
- Su, C. (2017). “From Porto Alegre to New York City: Participatory Budgeting and Democracy”. In: *New Political Science* 39 (1), pp. 67–75.
- van de Velde, S. and J. Worm (1994). *Multi-period planning of road maintenance: a multiple-choice multiple-knapsack problem*. Tech. rep. 94-6. Laboratory of Production and Operations Management, University of Twente.
- Vanderbeck, F. and L. A. Wolsey (2010). “Reformulation and Decomposition of Integer Programs”. In: *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-Art*. Ed. by M. Jünger, T. M. Liebling, D. Naddef, G. L. Nemhauser, W. R. Pulleyblank, G. Reinelt, G. Rinaldi, and L. A. Wolsey. Springer Berlin Heidelberg, pp. 431–502.
- Varian, H. (1974). “Equity, envy, and efficiency”. In: *Journal of Economic Theory* 9 (1), pp. 63–91.
- Xu, Z. (2013). “The knapsack problem with a minimum filling constraint”. In: *Naval Research Logistics* 60 (1), pp. 56–63.