The (Exact) Price of Cardinality for Indivisible Goods: A Parametric Perspective*

Alexander Lam[1], Bo Li[2], Ankang Sun[2]

[1]Department of Computer Science, City University of Hong Kong
[2]Department of Computing, The Hong Kong Polytechnic University
alexlam@cityu.edu.hk, comp-bo.li@polyu.edu.hk, ankang.sun@polyu.edu.hk

Abstract

We adopt a parametric approach to analyze the worst-case degradation in social welfare when the allocation of indivisible goods is constrained to be *fair*. Specifically, we are concerned with *cardinality*-constrained allocations, which require that each agent has at most k items in their allocated bundle. We propose the notion of the *price of cardinality*, which captures the worst-case multiplicative loss of *utilitarian* or *egalitarian* social welfare resulting from imposing the cardinality constraint. We then characterize tight or almost-tight bounds on the price of cardinality as exact functions of the instance parameters, demonstrating how the social welfare improves as k is increased. In particular, one of our main results refines and generalizes the existing asymptotic bound of $\Theta(\sqrt{n})$ on the price of balancedness. We also further extend our analysis to the problem where the items are partitioned into disjoint categories, and each category has its own cardinality constraint. Through a parametric study of the price of cardinality, we provide a framework which aids decision makers in choosing an ideal level of cardinality-based fairness, using their knowledge of the potential loss of utilitarian and egalitarian social welfare.

1 Introduction

The allocation of indivisible goods to a set of agents is a ubiquitous problem in our society, capturing a number of real-world scenarios. For example, an inheritance may involve indivisible goods such as jewelry, cars, and estates, and food banks are constantly faced with the task of allocating donations to people in need. In a corporate setting, equipment and human resources such as developers and designers need to be assigned to various projects and departments. These real-world scenarios often involve constraints which are imposed on the allocation, which may make the allocation difficult to compute, or prevent it from being socially optimal.

The most commonly studied constraint in resource allocation is the requirement that the allocation should be *fair* to the agents (Amanatidis et al. 2023). Fairness constraints can be expressed in several different ways: a *proportional* allocation guarantees that each of the n agents receives at least $\frac{1}{n}$th of their utility for the entire set of items, and in a *balanced* allocation, the goods are spread out among the agents as evenly as possible, so that the number of items received by each agent differs by at most one. The latter notion of *balancedness* is a natural constraint which may be imposed by the central decision maker due to its simplicity and ease of implementation, without requiring knowledge of the agents' utilities. However, this constraint may severely degrade the social welfare of the allocation. This is particularly true in instances where agents have low utility for a large number of items, rather than highly valuing a small subset of items, motivating the need for a weaker, more variable notion of 'cardinal fairness'.

The main focus of our paper is on the constraint of *cardinality*, a generalization of balancedness which imposes an upper limit of k on the number of items an agent may receive from the allocation.[1] The cardinality constraint is commonly applied in practice. For example, when a university provides funding to support PhD students, it is typical to impose a limit on the number of students each professor can supervise, due to fairness concerns. Intuitively, the potential loss of welfare decreases as k increases. The central decision maker may vary the parameter of k to achieve their desired tradeoff between the level of balancedness and the social welfare of the cardinality-constrained allocation. More generally, the items may also be partitioned into disjoint categories, where each category j has its own cardinality constraint of k_j. However, the exact effect of the value of k on the social welfare objectives remains unclear, leading to the following research question which we address in this paper:

> *What is the worst-case (multiplicative) loss of social welfare when there is a limit on the number of items each agent can receive in an allocation, and how does the loss change as the limit is varied?*

In particular, we aim to quantify this loss in an *exact* sense, as opposed to the asymptotic bounds which are common in the literature. This helps with making a more informed decision on the cardinality constraint values, particularly in scenarios where the number of agents and/or items is small.

*The author names are ordered alphabetically.

[1]Note that the cardinality constraint is equivalent to balancedness when $m \in \{kn-k+1, kn-k+2, \dots, kn\}$, where m is the total number of items.

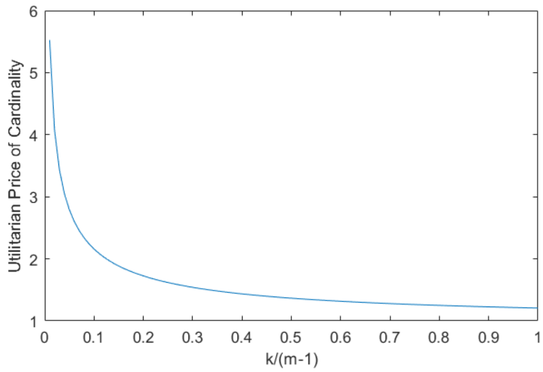

Figure 1: Plot of the utilitarian price of cardinality in the single-category setting as a function of $\frac{k}{m-1}$.

1.1 Our Contribution

In this work, we initiate the study of *price of cardinality* from the parametric perspective. We define the price of cardinality as the worst-case ratio between the welfare of the optimal allocation, and the optimal welfare among all cardinality-constrained allocations. Our work concerns both *utilitarian* and *egalitarian* social welfare, defined as the sum of agents' utilities and worst-off agent's utility respectively. For both objectives, we establish tight or almost-tight bounds for the price of cardinality in both the single-category and multi-category cases, expressing the prices as exact functions of the cardinality parameters, as opposed to the common asymptotic bounds in the literature. A benefit of our parametrized approach is that it enables decision makers to choose their desired level of fairness based on the potential loss of social welfare. We summarize our main results as follows.

Single category. We start with the single category case, where each agent can receive at most k items. We show that for any instance with n agents and m items such that $k \geq \frac{m}{n}$, the utilitarian price of cardinality is $\frac{1}{2}\left(1+\sqrt{1+\frac{m-1}{k}}\right)$. This can be visualized in Figure 1 as a function of the ratio $\frac{k}{m-1}$. We note that this bound is precisely tight for instances where m and k satisfy a divisibility constraint, and tight up to an additive constant that is smaller than 1 for every instance. Furthermore, when $m = kn$, this result coincides with (and refines) the asymptotic bound of $\Theta(\sqrt{n})$ for the *price of balancedness* (Bei et al. 2021).

For the objective of egalitarian social welfare, we present an exact bound of $\max\{\frac{m-n+1}{k}, 1\}$ which is tight for all instances. This result shows that when k is small compared to m, the cardinality constraint may adversely affect the egalitarian fairness of the allocation, particularly when the allocation is constrained to be balanced (i.e., $m = kn$).

Multiple categories. For utilitarian social welfare, we first focus on the case of two agents, giving an exact and tight bound of $\frac{2m_1m_2}{m_2k_1+m_1k_2}$ for the utilitarian price of cardinality.[2] For the case of general n agents, we establish a utilitarian price of cardinality of $\frac{m_1}{k_1}$, which is tight for instances where there are n categories and $\frac{k_1}{m_1} = \cdots = \frac{k_n}{m_n}$.

Finally, in the multi-category case, we establish an exact bound for egalitarian price of cardinality as a function of n, the cardinality constraints k_j, and the number of items in each category m_j, and this bound is tight for all instances.

For all of our main results, we also describe how a cardinal allocation with guaranteed welfare corresponding to the respective price of cardinality can be found.

1.2 Related Work

The problem of allocating a set of indivisible goods to self-interested agents has been extensively studied in computer science and economics in recent years, with a primary focus on finding allocations which are constrained to be 'fair' in some axiomatic sense (see, e.g., the seminal papers by Gamow and Stern (1958); Varian (1974); Steinhaus (1948) and the surveys by Walsh (2020); Amanatidis et al. (2023)). Common fairness notions have been shown to be compatible with cardinality constraints: for example, Biswas and Barman (2018) and Hummel and Hetland (2022) considered how to compute EF1 and approximate MMS allocations under cardinality constraints.

While most related work is concerned with computing fair allocations, there is a growing body of research on quantifying the degradation of welfare when fairness constraints are imposed. This was first proposed by Caragiannis et al. (2012), and results on the price of fairness have since been extended by Barman, Bhaskar, and Shah (2020), Li, Li, and Wu (2022), and Li et al. (2024). These papers generally examine the effect of fairness constraints from an asymptotic perspective, e.g., the prices of EF1 and $\frac{1}{2}$-MMS are $\Theta(\sqrt{n})$, and since computing the exact bound can be challenging for every value of n, the literature often restricts to the case where $n = 2$. In particular, the paper by Bei et al. (2021) gives an asymptotic bound of $\Theta(\sqrt{n})$ (and a precise bound of $4/3$ for $n = 2$) for the price of *balancedness* for utilitarian social welfare. As asymptotic bounds may obscure the precise effects of constraints (particularly in scenarios with a small number of agents), this paper aims to generalize the existing asymptotic bound on the price of balancedness to the exact bound on the price of cardinality, and extend the results to multiple categories.

Beyond cardinality constraints, other types of constraints such as connectivity (Bouveret et al. 2017; Bei et al. 2024), geometric (Segal-Halevi et al. 2017), and separation (Elkind, Segal-Halevi, and Suksompong 2022) have also been studied in the context of fair division. In particular, the cardinality constraint in the single-category case is equivalent to budget-feasibility (as studied by Wu, Li, and Gan (2021)) when agents have identical budgets and items have identical costs. For a recent overview on constraints in fair division, we refer the readers to the survey by Suksompong (2021).

[2]We order the categories such that $\frac{k_1}{m_1} \leq \frac{k_2}{m_2} \leq \cdots \leq \frac{k_h}{m_h}$, where k_j and m_j denote the cardinality constraint and number of items in category j, respectively.

2 Preliminaries

For any $k \in \mathbb{N}^+$, denote $[k] \coloneqq \{1, \dots, k\}$. We have a set N of n agents who are to receive a set $M = \{g_1, \dots, g_m\}$ of m indivisible goods. Each agent $i \in N$ has a utility function $u_i : 2^M \to \mathbb{R}_{\geq 0}$ over each possible subset of goods. Let $\mathcal{U} = (u_1, u_2, \dots, u_n)$ denote the agents' utility profiles. Throughout the paper, we assume the utilities are *additive* (i.e., $u_i(S) = \sum_{g \in S} u_i(\{g\})$) and *normalized* (i.e., $u_i(\emptyset) = 0$ and $u_i(M) = 1$ for each $i \in N$). We slightly abuse notation and write $u_i(g)$ instead of $u_i(\{g\})$.

An allocation $\mathcal{A}$ is a partition of the items into n disjoint bundles $\mathcal{A} = (A_1, \dots, A_n)$, with agent i receiving A_i. The set of goods is partitioned into h different non-overlapping categories $C = \{C_1, \dots, C_h\}$, with respective cardinality constraints $k_1, \dots, k_h$. The cardinality constraint k_j specifies the maximum number of items from category C_j that each agent can have under a *cardinal* allocation. In other words, an allocation $\mathcal{A}$ is *cardinal* if for each bundle A_i and category C_j, $|A_i \cap C_j| \leq k_j$ holds. Our prices of cardinality will be expressed in terms of the number of agents n, the cardinality constraints k_j, and the number of items in each category, so for simplicity, we denote, for each $j \in [h]$, $m_j \coloneqq |C_j|$ and $m = \sum_j m_j$. We also assume that for each $j \in [h]$, $k_j \geq \frac{m_j}{n}$ so that every item can be allocated. In the single-category scenario (i.e., $h = 1$), we slightly abuse notation and use m and k instead of m_1 and k_1 to refer to the number of items and the cardinality constraint, respectively.

We refer to a problem setting with agents N, goods M, utility profile $\mathcal{U}$ and partition of goods into categories C as an *instance*, denoted as $I = \langle N, M, \mathcal{U}, C \rangle$.

Our results are concerned with the following objectives of *utilitarian social welfare* and *egalitarian social welfare*.

Definition 2.1. Given an instance I and an allocation $\mathcal{A} = (A_1, \dots, A_n)$ of the instance,

- the *utilitarian social welfare*, or the sum of the agents' utilities, is denoted by $\text{USW}(\mathcal{A}) \coloneqq \sum_{i \in N} u_i(A_i)$.
- the *egalitarian social welfare*, or the utility of the worst-off agent, is denoted by $\text{ESW}(\mathcal{A}) \coloneqq \min_{i \in N} u_i(A_i)$.

We denote the optimal utilitarian (resp. egalitarian) social welfare over all possible allocations of an instance I as OPT-USW(I) (resp. OPT-ESW(I)).

Given an instance I and cardinality constraints $\kappa = (k_1, \dots, k_h)$, let the set of all cardinal allocations be denoted as $\mathcal{C}_\kappa(I)$. We now define the main concept of our paper: the *price of cardinality*, for our objective functions of utilitarian and egalitarian social welfare.

Definition 2.2. The *utilitarian price of cardinality* is defined as

$$\sup_{I = \langle N, M, \mathcal{U}, C \rangle, \kappa = (k_1, \dots, k_h)} \frac{\text{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})}.$$

Definition 2.3. The *egalitarian price of cardinality* is defined as

$$\sup_{I = \langle N, M, \mathcal{U}, C \rangle, \kappa = (k_1, \dots, k_h)} \frac{\text{OPT-ESW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{ESW}(\mathcal{A})}.$$

For an instance I, if OPT-ESW$(I) = 0$, we define that $\frac{0}{0} = 1$ and the egalitarian price of cardinality to be 1.

3 Single Category

We first give results for the case where there is only one category, and therefore only one cardinality constraint k. Note that if $m \leq k$, then the cardinality constraint will have no effect, and the price of cardinality will be equal to 1. We therefore assume that $m > k$ throughout this section.

3.1 Utilitarian Social Welfare

We begin with the utilitarian price of connectivity in the single category setting. Interestingly, the price of cardinality for utilitarian welfare is only directly dependent on the number of items and the cardinality constraint, and not directly dependent on the number of agents (it is not entirely independent of n as we require $n \geq \frac{m}{k}$).

Theorem 3.1. *In the single-category case, the utilitarian price of cardinality is*

$$\frac{1}{2}\left(1 + \sqrt{1 + \frac{m-1}{k}}\right).$$

Proof. The lower bound will be proven in Lemma 3.2, and the upper bound will be proven in Lemmas 3.4 and 3.6. □

It is worth noting that the stated price of cardinality expression is tight in a very strong sense:

- For any instance with cardinality constraint k and m items, the utilitarian price of cardinality is *at most* $\frac{1}{2}\left(1 + \sqrt{1 + \frac{m-1}{k}}\right)$.
- As we will shortly show in Lemma 3.2, for any instance with cardinality constraint k and $m = k(c^2-1)+1$ items, where $c \in \mathbb{N}^+ \setminus \{1\}$, the utilitarian price of cardinality is exactly $\frac{1}{2}\left(1 + \sqrt{1 + \frac{m-1}{k}}\right)$.
- In Lemma 3.3, we further show that for any other instance with cardinality constraint k, the utilitarian price of cardinality is at least $\frac{1}{2}\left(-1 + \sqrt{1 + \frac{m-1}{k}}\right)$.

This is due to the lower bound construction requiring m and k to meet a divisibility constraint. Specifically, there is a set of agents who each value the same item g_m at one utility, and the remaining agents' utilities are such that they each receive the same number of items under the utilitarian-optimal allocation.

Lemma 3.2. *In the single-category case, if* $m = k(c^2 - 1) + 1$ *for some* $c \in \mathbb{N}^+ \setminus \{1\}$*, then the utilitarian price of cardinality is at least* $\frac{1}{2}\left(1 + \sqrt{1 + \frac{m-1}{k}}\right)$.

Proof. Let I be an instance with cardinality constraint k and $m = k(c^2 - 1) + 1$ items for some $c \in \mathbb{N}^+ \setminus \{1\}$, and let $s \coloneqq -1 + \sqrt{1 + \frac{m-1}{k}}$. Note that

- $s = c - 1 \in \mathbb{N}^+$,
- and $\frac{m-1}{s} = \frac{k(c^2-1)}{c-1} \in \mathbb{N}^+$.

The agents' utilities are as follows:

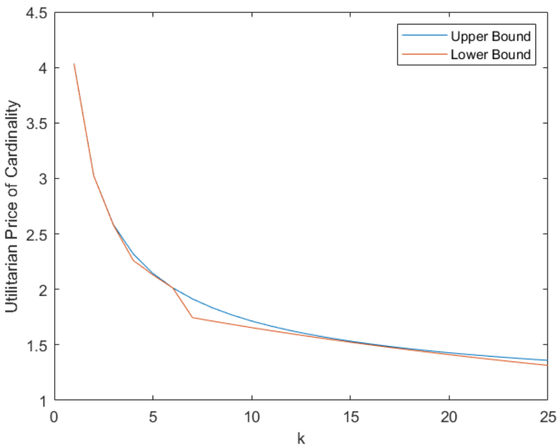

Figure 2: Plot for $m = 50$ showing the gap between the lower bound as described in the main body and proof of Lemma 3.3 for any m and k, and the upper bound from Theorem 3.1.

- for $i = 1, \ldots, s$, if $j = (i-1)\frac{m-1}{s} + 1, \ldots, i\frac{m-1}{s}$, then $u_i(g_j) = \frac{s}{m-1}$; otherwise, $u_i(g) = 0$;
- for $i \geq s+1$, $u_i(g_m) = 1$ and $u_i(g) = 0$ for each $g \in M \setminus \{g_m\}$.

We have $\text{OPT-USW}(I) = 1 + s$ as in the utilitarian-optimal allocation, each agent $i \in [s]$ receives utility 1, and agents $s+1, \ldots, n$ have a total utility of 1. We also have $\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A}) = 1 + \frac{ks^2}{m-1}$ by letting every agent $i \in [s]$ keep their k most valued items; note $\frac{m-1}{s} > k$. Dividing these terms and substituting $s = -1 + \sqrt{1 + \frac{m-1}{k}}$, we get our price of cardinality lower bound of $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} = \frac{1}{2}\left(1 + \sqrt{1 + \frac{m-1}{k}}\right)$. ☐

Note that if the divisibility constraint is not met by m and k (i.e., $m \neq k(c^2 - 1) + 1$ for all $c \in \mathbb{N}^+ \setminus \{1\}$), we can still construct a similar lower bound which will be slightly lower than our general upper bound (as exemplified in Figure 2). We take $s = \lfloor -1 + \sqrt{1 + \frac{m-1}{k}} \rfloor$, and let agents $1, \ldots, s$ each equally value a disjoint subset of the items $g_1, \ldots, g_{m-1}$, such that the number of goods they positively value differs by at most 1. The gap between the upper and lower bound here is due to the rounding of s and the partitioning of $m-1$ items among s agents as evenly as possible. Furthermore, the following lemma implies that the price of cardinality lower bound corresponding to this instance construction differs from our stated upper bound by an additive constant of at most 1.

Lemma 3.3. *In the single-category case, if $m \neq k(c^2 - 1) + 1$ for all $c \in \mathbb{N}^+ \setminus \{1\}$, then the utilitarian price of cardinality is at least $\frac{1}{2}\left(-1 + \sqrt{1 + \frac{m-1}{k}}\right)$.*

We now prove the upper bound of Theorem 3.1, which holds for any m and k. First, we fix an arbitrary instance I with m items and cardinality constraint k, and denote the utilitarian-optimal allocation[3] of I by $\mathcal{A}^* = (A_1^*, \ldots, A_n^*)$. We then define the following subsets of agents. Under $\mathcal{A}^*$,

- The set of agents receiving less than k items is $R = \{i \in [n] : |A_i^*| < k\}$.
- The set of agents receiving exactly k items is $T = \{i \in [n] : |A_i^*| = k\}$.
- The set of agents receiving more than k items is $S = \{i \in [n] : |A_i^*| > k\}$.

We can also assume that $\mathcal{A}^*$ does not satisfy the cardinality constraint, and hence $S \neq \emptyset$, implying $R \neq \emptyset$. We next divide the proof of the upper bound into two cases depending on whether $\sum_{i \in R \cup T} u_i(A_i^*) \geq 1$ or $\sum_{i \in R \cup T} u_i(A_i^*) < 1$. We begin with the former case.

Lemma 3.4. *Let I be a (possibly non-normalized) instance satisfying $0 < u_i(M) \leq 1$ for each $i \in S$ and $\sum_{i \in N} \setminus S u_i(A_i^*) \geq 1$, where S is the set of agents who receive more than k items under utilitarian-optimal allocation $\mathcal{A}^* = (A_1^*, \ldots, A_n^*)$. Then,*

$$\frac{\textit{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \textit{USW}(\mathcal{A})} \leq \frac{1 + \sum_{i \in S} u_i(A_i^*)}{1 + k \sum_{i \in S} \frac{u_i(A_i^*)}{|A_i^*|}} \leq \frac{1+s}{1 + \frac{ks^2}{m-1}},$$

where $s = -1 + \sqrt{1 + \frac{m-1}{k}}$.

Proof Sketch. We present a proof sketch for the inequality $\frac{1 + \sum_{i \in S} u_i(A_i^*)}{1 + k \sum_{i \in S} \frac{u_i(A_i^*)}{|A_i^*|}} \leq \frac{1+s}{1 + \frac{ks^2}{m-1}}$. By taking derivatives, we find that the LHS is maximized when every $u_i(A_i^*)$ is either 1 or 0. This gives us

$$\frac{1 + \sum_{i \in S} u_i(A_i^*)}{1 + k \sum_{i \in S} \frac{u_i(A_i^*)}{|A_i^*|}} \leq \frac{1 + \sum_{i \in S'} 1}{1 + k \sum_{i \in S'} \frac{1}{A_i^*}}$$
$$\leq \frac{1 + |S'|}{1 + \frac{k|S'|^2}{m-1}} \leq \frac{1+s}{1 + \frac{ks^2}{m-1}},$$

where $s = -1 + \sqrt{1 + \frac{m-1}{k}}$. Here, the second inequality follows from the arithmetic mean-harmonic mean (AM-HM) inequality, and the final inequality follows from taking derivatives with respect to $|S'|$. ☐

We now address the remaining case where $\sum_{i \in R \cup T} u_i(A_i^*) < 1$. In this case, we assume that instance I is *preprocessed* such that for each $j \in R$, $\sum_{i \in S} u_j(A_i^*) := 1 - \sum_{i \in R \cup T} u_i(A_i^*)$. As $\mathcal{A}^*$ is the utilitarian-optimal allocation, we have $\sum_{i \in R \cup T} u_i(A_i^*) \geq \sum_{i \in R \cup T} u_j(A_i^*)$ for all $j \in R$, and therefore before processing, we have $\sum_{i \in S} u_j(A_i^*) \geq 1 - \sum_{i \in R \cup T} u_i(A_i^*)$. Accordingly, this preprocessing can be achieved by reducing the utility that each agent $j \in R$ has for items $\{A_i^*\}_{i \in S}$ until we reach

[3]In the case of ties, the tiebreak procedure will be described in the upcoming proofs when necessary.

$\sum_{i \in S} u_j(A_i^*) = 1 - \sum_{i \in R \cup T} u_i(A_i^*)$. The new/preprocessed instance is not necessarily normalized, meaning there may exist $i \in R$ such that $0 \leq u_i(M) \leq 1$. Note that the optimal utilitarian welfare does not change before and after the preprocessing, but the optimal utilitarian welfare among cardinal allocations weakly decreases after preprocessing, meaning that $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})}$ weakly increases. We also remark that the preprocessing does not affect S, R, or T.

Before proving the price of cardinality upper bound for the case where $\sum_{i \in R \cup T} u_i(A_i^*) < 1$, we find a lower bound on the utilitarian social welfare of a cardinal allocation, for any arbitrary instance I.[4]

Lemma 3.5. *For an arbitrary instance I, there exists a cardinal allocation $\mathcal{A}$ such that USW($\mathcal{A}$) $\geq 1 + \sum_{i \in S} \frac{k}{|A_i^*|}(u_i(A_i^*) - u_{i^\dagger}(A_i^*))$ for some agent $i^\dagger \in R$.*

Proof Sketch. In the full proof, we show that there exists a cardinal allocation $\mathcal{A}$ such that $\text{USW}(\mathcal{A}) \geq \sum_{i \in R \cup T} u_i(A_i^*) + \sum_{i \in S} \frac{k}{|A_i^*|} u_i(A_i^*) + \sum_{i \in S} \frac{|A_i^*| - k}{|A_i^*|} u_{i^\dagger}(A_i^*)$, which suffices because we have $\sum_{i \in R \cup T} u_i(A_i^*) = 1 - \sum_{i \in S} u_{i^\dagger}(A_i^*)$. (Recall that I is preprocessed). Specifically, this cardinal allocation can be achieved by the following greedy procedure.

The procedure starts from $\mathcal{A}^*$, and at each step, reassigns the item with the least utility loss from some *unsatisfied* agent's bundle to some *active* agent; an agent is *unsatisfied* if she receives more than k items, and is *active* if she receives less than k items.

- Step 1: Set $\mathcal{B} \leftarrow \mathcal{A}^*$ as the initial allocation, $P \leftarrow S$ as the initial set of unsatisfied agents, and $Q \leftarrow R$ as the initial set of active agents;
- Step 2: If there are no unsatisfied agents, then terminate and output the underlying allocation $\mathcal{B}$ (this will be $\mathcal{A}^k$). Otherwise, find the item $e^* \in \bigcup_{i \in P} B_i$ and an active agent $i^* \in Q$ such that reassigning e^* to agent i^* causes the minimum utilitarian social welfare loss among all single-item reassignments from items of unsatisfied agents to active agents. Reassign e^* to agent i^*, and update $\mathcal{B}$ accordingly.
- Step 3: Update P and Q, and return to Step 2.

As $m \leq kn$, the procedure can terminate and the returned allocation $\mathcal{A}^k$ is cardinal. Moreover, during the reassignment process, an active agent can never become unsatisfied and any unsatisfied agent can never become active. □

We now prove the upper bound for Theorem 3.1 for the case where $\sum_{i \in R \cup T} u_i(A_i^*) < 1$.

Lemma 3.6. *If $\sum_{i \in R \cup T} u_i(A_i^*) < 1$, then*

$$\frac{\textit{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \textit{USW}(\mathcal{A})} \leq \frac{1+s}{1 + \frac{ks^2}{m-1}},$$

where $s = -1 + \sqrt{1 + \frac{m-1}{k}}$.

[4]We remark that the lemma holds regardless of whether $\sum_{i \in R \cup T} u_i(A_i^*) < 1$ or $\sum_{i \in R \cup T} u_i(A_i^*) \geq 1$, and whether or not I is preprocessed.

Partial Proof. We first describe the tiebreak procedure for the utilitarian-optimal allocation $\mathcal{A}^*$. If multiple agents are tied for having the highest utility for an item, we pick the allocation $\mathcal{A}^*$ based on the following criteria;

- if it is possible to allocate each item to the agent that values it most, such that all m items are owned by agents with strictly more than k items, then $\mathcal{A}^*$ is defined as this allocation,
- otherwise, the tie is broken in favour of the agent with less than k items.

We divide the remainder of the proof into two cases, depending on whether or not all of the goods are allocated to agents in S under $\mathcal{A}^*$.

Case 1: $\sum_{i \in S} |A_i^*| < m$. We present the full proof for the case where not all items are allocated to agents in S under $\mathcal{A}^*$. Recall that by Lemma 3.5, there exists an agent $i^\dagger \in R$ and a cardinal allocation $\mathcal{A}^k$ such that $\text{USW}(\mathcal{A}^k) \geq 1 + \sum_{i \in S} \frac{k}{|A_i^*|}(u_i(A_i^*) - u_{i^\dagger}(A_i^*))$.

Consider the agent $i^\dagger \in R$ and a specific item $g^\dagger \in \bigcup_{i \in R \cup T} A_i^*$; the existence of $g^\dagger$ is guaranteed due to $\sum_{i \in S} |A_i^*| < m$. We construct another (possibly non-normalized) instance I' which differs from I only by the agents' utilities. Below, we describe the utility function u' that each agent has in I',

- for $i \in R \cup T$, $u_i'(g^\dagger) = 1$ and $u_i'(g) = 0$ for all $g \neq g^\dagger$;
- for $i \in S$, $u_i'(g) = u_i(g) - u_{i^\dagger}(g)$ if $g \in A_i^*$ and $u'(g) = 0$ otherwise.

Denote by $\mathcal{A}'$ the utilitarian-optimal allocation of I' and by S' the set of agents receiving more than k items in $\mathcal{A}'$; note that $S' = S$. Due to $\sum_{i \in S} |A_i^*| < m$, when picking $\mathcal{A}^*$, if multiple agents are tied for having the highest utility for an item, then the tie is broken in favour of the agent with less than k items. As a consequence, for any $i \in S$ and $g \in A_i^*$, $u_i'(g) > 0$ due to $i^\dagger \in R$.

Now we show that in I', $0 < u_i'(A_i') \leq 1$ for each $i \in S'$. From the construction of $u_i'()$, we immediately have $u_i'(A_i') \leq 1$. To prove $0 < u_i'(A_i')$, since agent i is the only one with positive utility on the items in bundle A_i^*, we have $A_i^* \subsetneq A_i'$ and hence $0 < u_i'(A_i')$; note that $u_i'(g) > 0$ for every $g \in A_i^*$.

We now present the upper bound of the ratio regarding I for Case 1 as follows,

$$\begin{aligned} &\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} \\ &\quad \leq \frac{\sum_{i \in R \cup T} u_i(A_i^*) + \sum_{i \in S} u_i(A_i^*)}{\text{USW}(\mathcal{A}^k)} \\ &\quad = \frac{\text{OPT-USW}(I')}{\max_{\mathcal{A} \in \mathcal{C}_\kappa(I')} \text{USW}(\mathcal{A})} \\ &\quad \leq \frac{1}{2}\left(\sqrt{1 + \frac{m-1}{k}} + 1\right), \end{aligned}$$

where the first equality results from the property of the preprocessed instance and the fact that $i^\dagger \in R$; the last inequality transition follows from Lemma 3.4. □

Finally, we conclude the section with the following result on computing utilitarian-optimal cardinal allocations.

Proposition 3.7. *Given a single-category instance I and cardinality constraint k, the utilitarian-optimal cardinal allocation can be found in polynomial time, and has a utilitarian social welfare of at least $\frac{2}{1+\sqrt{1+\frac{m-1}{k}}} \cdot$ OPT-USW(I).*

Proof. Consider a complete bipartite graph $G = (U, V, E)$, where U represents k copies of each agent, and V represents the m goods, with zero-valued dummy items added such that $|U| = |V|$. Also, an edge between agent $i \in U$ and item $g \in V$ has weight equal to $u_i(g)$. Our desired allocation can be found by computing a maximum weight bipartite matching, such as by using the Hungarian algorithm (Kuhn 1955). The utilitarian social welfare guarantee follows immediately from Theorem 3.1. □

3.2 Egalitarian Social Welfare

We now move to the objective of egalitarian social welfare, where the worst-case degradation of worst-case fairness objective is quantified by our exact and tight bounds on the egalitarian price of cardinality. Note that in addition to the assumption that $m > k$, we also assume in this subsection that $m \geq n$, because if $m < n$, then OPT-ESW$(I) = 0$ and consequently, the egalitarian price of cardinality will be 1.

Theorem 3.8. *In the single-category case, the egalitarian price of cardinality is $\max\left\{\frac{m-n+1}{k}, 1\right\}$.*

Note that this bound is tight for all feasible values of m, n, and k. This result shows that when m is large compared to n and k, there may be a significant reduction in egalitarian fairness when cardinality constraints are naively imposed in pursuit of a fair allocation. We also remark that although computing an egalitarian-optimal (possibly non-cardinal) allocation is well-known to be NP-hard (Karp 1972), if we are provided with such an allocation, we can find, in linear time, a cardinal allocation with an egalitarian social welfare guarantee corresponding to the egalitarian price of cardinality. This is simply achieved by letting each agent keep their k most valued items from the starting egalitarian-optimal allocation.

4 Multiple Categories

We now extend our analysis to the setting where the items are partitioned into multiple categories. Recall that there are h categories, where category $j \in [h]$ has m_j items to be allocated and a cardinality constraint of k_j. We also ensure that all items can be assigned, $\frac{m_j}{k_j} \leq n$ holds for each $j \in [h]$. Without loss of generality, we order categories such that $\frac{k_1}{m_1} \leq \frac{k_2}{m_2} \leq \cdots \leq \frac{k_h}{m_h}$ and break ties in favour of the category with a smaller number of items (i.e., if $\frac{k_i}{m_1} = \frac{k_j}{m_j}$ and $m_i < m_j$, then set $i < j$).

4.1 Utilitarian Social Welfare

For utilitarian social welfare, we first consider the case of two agents. Before stating the main result, we establish a key reduction which restricts the space of instances to those with weakly higher $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})}$ ratio.

Lemma 4.1. *Given an instance I with two agents and cardinality constraints κ, there exists another instance I' which only differs from I in the utility functions, where:*

- *under the utilitarian-optimal allocation $\mathcal{A}'^*$,*
  - *both agents exceed the cardinality constraint in exactly one category each,*
  - *neither agent receives any utility from any category where they do not exceed the cardinality constraint,*
- *and $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} \leq \frac{\text{OPT-USW}(I')}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I')} \text{USW}(\mathcal{A})}$ holds.*

Following this reduction, we are now ready to present the utilitarian price of cardinality for two agents, which is exact and tight for all $\kappa = (k_1, \ldots, k_h)$ and $m_1, \ldots, m_h$.

Theorem 4.2. *For two agents and $h \geq 2$, the utilitarian price of cardinality is $\frac{2}{\frac{k_1}{m_1}+\frac{k_2}{m_2}}$.*

Proof. By Lemma 4.1, it suffices to focus on the case where in $\mathcal{A}^*$, agent 1 (resp. agent 2) exceeds the cardinality constraint of category j_1 (resp. j_2). Note that due to the ordering of our categories, it is 'weakly better' to consider categories 1 and 2. Moreover in $\mathcal{A}^*$, each agent i only receives non-zero utility from C_{j_i}. Note that we must have $j_1 \neq j_2$ due to $\frac{m_j}{k_j} \leq 2$ for all j.

We first prove the upper bound. Consider another (possibly non-normalized) instance I' that only differs from I in utility functions. In I', agent i has utility function $u'_i(g) = u_i(g)$ if $g \in A^*_{ij_i}$ and $u'_i(g) = 0$ otherwise. One can verify that the welfare of utilitarian-optimal allocation of I is equal to that of I', while the maximum welfare of cardinal allocations is weakly decreased in I'. Accordingly, we have $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} \leq \frac{\text{OPT-USW}(I')}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I')} \text{USW}(\mathcal{A})}$. We then convert I' into a normalized instance I'' by increasing agent i's utility for $A^*_{ij_i}$ to 1 in a way such that the 'price of cardinality' ratio weakly increases; this can be done by increasing the utility of items valued the most by both agents. Note that I'' is a normalized instance where in $\mathcal{A}^*$, each agent i receives utility 1 from obtaining all of the items which they positively value.

Finally, we have

$$\begin{aligned}
\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} &\leq \frac{\text{OPT-USW}(I'')}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I'')} \text{USW}(\mathcal{A})} \\
&\leq \frac{2}{\frac{k_{j_1}}{|A^*_{1j_1}|} + \frac{k_{j_2}}{|A^*_{2j_2}|}} \\
&\leq \frac{2}{\frac{k_{j_1}}{m_{j_1}} + \frac{k_{j_2}}{m_{j_2}}} \leq \frac{2}{\frac{k_1}{m_1} + \frac{k_2}{m_2}},
\end{aligned}$$

concluding the proof of the upper bound.

For the lower bound, consider the instance I where agent 1 values each item in category 1 at $\frac{1}{m_1}$ utility and agent 2 values each item in category 2 each at $\frac{1}{m_2}$ utility. Clearly, $\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} = \frac{2}{\frac{k_1}{m_1}+\frac{k_2}{m_2}}$ for this instance. □

We give the utilitarian price of cardinality for general n.

Theorem 4.3. *For general n, the utilitarian price of cardinality is $\frac{m_1}{k_1}$.*

Proof. We first prove the upper bound. Given an instance I, let $\mathcal{A}^*$ be its utilitarian-optimal allocation. Then we have

$$\frac{\text{OPT-USW}(I)}{\max_{\mathcal{A}\in\mathcal{C}_\kappa(I)} \text{USW}(\mathcal{A})} \leq \frac{\sum_{j\in[h]}\left(\sum_{i\in S_j} u_i(A^*_{ij}) + \sum_{i\in N\setminus S_j} u_i(A^*_{ij})\right)}{\sum_{j\in[h]}\left(\sum_{i\in S_j} \frac{k_j}{|A^*_{ij}|} u_i(A^*_{ij}) + \sum_{i\in N\setminus S_j} u_i(A^*_{ij})\right)} \leq \frac{\sum_{j\in[h]}\sum_{i\in S_j} u_i(A^*_{ij})}{\sum_{j\in[h]}\sum_{i\in S_j} \frac{k_j}{|A^*_{ij}|} u_i(A^*_{ij})} \leq \frac{m_1}{k_1},$$

where the last inequality transition is because for every $j \in [h]$, $\frac{k_j}{|A^*_{ij}|} \leq \frac{k_j}{m_j} \leq \frac{k_1}{m_1}$.

For the lower bound, consider an instance I with $h = n$ categories, and where each category $j \in [h]$ has the same cardinality constraint of $k_j = k$ and same number of items $m_j = q$ items, where q is divisible by k. Suppose that in this instance, for each agent i, $u_i(g) = \frac{1}{q}$ if $g \in C_i$, and $u_i(g) = 0$ otherwise. Clearly, we have OPT-USW$(I) = n$. In the utilitarian-optimal cardinal allocation $\mathcal{A}^k$, each agent i receives a utility of $\frac{k}{q}$. Thus, we have USW$(\mathcal{A}^k) = \frac{nk}{q}$, and therefore, the utilitarian price of cardinality is at least $\frac{\text{OPT-USW}(I)}{\text{USW}(\mathcal{A}^k)} = \frac{q}{k} = \frac{m_1}{k_1}$, completing the proof. □

This result is roughly tight in the sense that the utilitarian price of cardinality is at most $\frac{m_1}{k_1}$ for any instance with cardinality constraints $\kappa = (k_1, \ldots, k_j)$, and is precisely $\frac{m_1}{k_1}$ for any instance where there are at least n categories, and $\frac{k_1}{m_1} = \frac{k_2}{m_2} = \cdots = \frac{k_n}{m_n}$.

Finally, we mention a result on computing utilitarian-optimal cardinal allocations, similar to Proposition 3.7.

Proposition 4.4. *Given a multiple-category instance I and cardinality constraints κ, the utilitarian-optimal cardinal allocation can be found in polynomial time, and has a utilitarian social welfare of at least $\frac{k_1}{m_1} \cdot$ OPT-USW(I).*

Proof. The proof is almost identical to the proof of Proposition 3.7, but we instead construct a separate complete bipartite graph for each category $j \in [h]$, and compute a maximum weight bipartite matching for each of these graphs. The runtime remains in polynomial time, and the utilitarian social welfare guarantee follows immediately from Theorem 3.1. □

4.2 Egalitarian Social Welfare

Finally, for egalitarian social welfare, we present bounds for the price of cardinality which are exact and tight for any n, $\kappa = (k_1, \ldots, k_h)$, and $m_1, \ldots, m_h$.

Theorem 4.5. *If $n \leq \sum_{j=2}^{h} m_j + 1$, then the egalitarian price of cardinality is $\frac{m_1}{k_1}$. If $n > \sum_{j=2}^{h} m_j + 1$, then the egalitarian price of cardinality is*

$$\max_{j\in[h]} \left\{ \frac{m_j - \max\{n - 1 - \sum_{t\neq j} m_t, 0\}}{k_j} \right\}.$$

Similar to the single category scenario, if we are given an egalitarian-optimal allocation, we can construct a cardinal allocation with an egalitarian welfare guarantee corresponding to the price of cardinality by letting each agent keep their k_j most valued items in each category $j \in [h]$.

5 Discussion

In this work, we introduced the utilitarian and egalitarian prices of cardinality, which quantify the worst-case multiplicative loss of social welfare when cardinality constraints are imposed on the allocation. For both the single- and multi-category cases, we present tight bounds on the prices of cardinality, expressed as an exact (rather than asymptotic) function of the instance and cardinality parameters. Our results enable decision makers to make a clear, well-informed choice of cardinality constraint with respect to the level of balancedness and the potential loss of social welfare.

Our parametrized approach to the price of cardinality can be applied to other parametrized notions of fairness such as envy-freeness up to k items (EF-k) or α-maximin share (α-MMS), providing similar insights to decision makers regarding the tradeoff between the level of fairness and the potential loss of social welfare.

An immediate open question is to find a more precise utilitarian price of cardinality in the case of multiple categories and $n \geq 3$ agents. Ideally, we would like the price to be tight for all possible values of n, $\{m_j\}_{j\in[h]}$, and $(k_j)_{j\in[h]}$, like our egalitarian prices of cardinality. To pursue such a precise price, one approach is to characterize the worst case scenario. Unfortunately, our reduction (in Lemma 4.1) for the case of two agents can not be immediately extended to the case where $n \geq 3$. Furthermore, part of the complete proof of Lemma 3.6 relies on knowing the exact form of a parameter s, which we cannot find in the multi-category case due to the multivariate property of the problem. However, we believe that the greedy procedure in the proof of Lemma 3.5 could be helpful for identifying the worst case structure.

Another possible direction is finding the price of cardinality for the "dual" problem where each agent must receive *at least* k items. However, this type of constraint violates the hereditary property of matroids, and is generally not studied in other related work.

Acknowledgments

The authors would like to thank Siddarth Barman, Xiaohui Bei, and the anonymous AAAI reviewers for their helpful comments. This work is funded by HKSAR RGC under Grant No. PolyU 15224823 and the Guangdong Basic and Applied Basic Research Foundation under Grant No. 2024A1515011524.

References

Amanatidis, G.; Aziz, H.; Birmpas, G.; Filos-Ratsikas, A.; Li, B.; Moulin, H.; Voudouris, A. A.; and Wu, X. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322: 103965.

Barman, S.; Bhaskar, U.; and Shah, N. 2020. Optimal Bounds on the Price of Fairness for Indivisible Goods. In *WINE*, volume 12495 of *Lecture Notes in Computer Science*, 356–369. Springer.

Bei, X.; Lam, A.; Lu, X.; and Suksompong, W. 2024. Welfare Loss in Connected Resource Allocation. In *IJCAI*, 2660–2668. ijcai.org.

Bei, X.; Lu, X.; Manurangsi, P.; and Suksompong, W. 2021. The Price of Fairness for Indivisible Goods. *Theory Comput. Syst.*, 65(7): 1069–1093.

Biswas, A.; and Barman, S. 2018. Fair Division Under Cardinality Constraints. In *IJCAI*, 91–97. ijcai.org.

Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair Division of a Graph. In *IJCAI*, 135–141. ijcai.org.

Caragiannis, I.; Kaklamanis, C.; Kanellopoulos, P.; and Kyropoulou, M. 2012. The Efficiency of Fair Division. *Theory Comput. Syst.*, 50(4): 589–610.

Elkind, E.; Segal-Halevi, E.; and Suksompong, W. 2022. Mind the gap: Cake cutting with separation. *Artificial Intelligence*, 313: 103783.

Gamow, G.; and Stern, M. 1958. *Puzzle-Math*. Viking press.

Hummel, H.; and Hetland, M. L. 2022. Guaranteeing Half-Maximin Shares Under Cardinality Constraints. In *AAMAS*, 1633–1635. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS).

Karp, R. M. 1972. Reducibility Among Combinatorial Problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, 85–103. Plenum Press, New York.

Kuhn, H. W. 1955. The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2): 83–97.

Li, B.; Li, Y.; and Wu, X. 2022. Almost (Weighted) Proportional Allocations for Indivisible Chores. In *WWW '22: The ACM Web Conference 2022, Virtual Event, Lyon, France, April 25 - 29, 2022*, 122–131. ACM.

Li, Z.; Liu, S.; Lu, X.; Tao, B.; and Tao, Y. 2024. A Complete Landscape for the Price of Envy-Freeness. In *AAMAS*, 1183–1191. International Foundation for Autonomous Agents and Multiagent Systems / ACM.

Segal-Halevi, E.; Nitzan, S.; Hassidim, A.; and Aumann, Y. 2017. Fair and square: Cake-cutting in two dimensions. *Journal of Mathematical Economics*, 70: 1–28.

Steinhaus, H. 1948. The problem of fair division. *Econometrica*, 16: 101–104.

Suksompong, W. 2021. Constraints in fair division. *SIGecom Exch.*, 19(2): 46–61.

Varian, H. R. 1974. Equity, Envy and Efficiency. *Journal of Economic Theory*, 9: 63–91.

Walsh, T. 2020. Fair Division: The Computer Scientist's Perspective. In *IJCAI*, 4966–4972. ijcai.org.

Wu, X.; Li, B.; and Gan, J. 2021. Budget-feasible Maximum Nash Social Welfare is Almost Envy-free. In *IJCAI*, 465–471. ijcai.org.