The Price of Equity with Binary Valuations and Few Agent Types

Umang Bhaskar[1], Neeldhara Misra[2], Aditi Sethia[2(✉)], and Rohit Vaish[3]

[1] Tata Institute of Fundamental Research, Mumbai, India
umang@tifr.res.in

[2] Indian Institute of Technology Gandhinagar, Gandhinagar, India
{neeldhara.m,aditi.sethia}@iitgn.ac.in

[3] Indian Institute of Technology Delhi, New Delhi, India
rvaish@iitd.ac.in

Abstract. In fair division problems, the notion of price of fairness measures the loss in welfare due to a fairness constraint. Prior work on the price of fairness has focused primarily on envy-freeness up to one good (EF1) as the fairness constraint, and on the utilitarian and egalitarian welfare measures. Our work instead focuses on the price of equitability up to one good (EQ1) (which we term *price of equity*) and considers the broad class of *generalized p-mean* welfare measures (which includes utilitarian, egalitarian, and Nash welfare as special cases). We derive fine-grained bounds on the price of equity in terms of the *number of agent types* (i.e., the maximum number of agents with distinct valuations), which allows us to identify scenarios where the existing bounds in terms of the number of agents are overly pessimistic.

Our work focuses on the setting with binary additive valuations, and obtains upper and lower bounds on the price of equity for p-mean welfare for all $p \leqslant 1$. For any fixed p, our bounds are tight up to constant factors. A useful insight of our work is to identify the *structure* of allocations that underlie the upper (respectively, the lower) bounds *simultaneously* for all p-mean welfare measures, thus providing a unified structural understanding of price of fairness in this setting. This structural understanding, in fact, extends to the more general class of binary submodular (or matroid rank) valuations. We also show that, unlike binary additive valuations, for binary submodular valuations the number of agent types does not provide bounds on the price of equity.

Keywords: Price of Fairness · Equitability · Optimal Welfare

1 Introduction

Tradeoffs are inevitable when we pursue multiple optimization objectives that are typically not simultaneously achievable. Quantifying such tradeoffs is a fundamental problem in computation, game theory, and economics. Our focus in this work is on the "price of fairness" in the context of fair division problems, which is a notion that captures tradeoffs between *fairness* and *welfare*.

A. Deligkas and A. Filos-Ratsikas (Eds.): SAGT 2023, LNCS 14238, pp. 271–289, 2023.
https://doi.org/10.1007/978-3-031-43254-5_16

Recall that a fair division instance in the discrete setting involves a set of n *agents* $N = \{1, 2, \ldots, n\}$, m indivisible *goods* $M = \{g_1, \ldots, g_m\}$, and $\mathcal{V} \coloneqq \{v_1, v_2, \ldots, v_n\}$, a *valuation profile* consisting of each agent's valuation of the goods. For any agent $i \in N$, its valuation function $v_i : 2^M \rightarrow \mathbb{N} \cup \{0\}$ specifies its numerical value (or *utility*) for every subset of goods in M. We will assume that the valuations are normalised, that is, for all $i \in N$, $v_i(M) = W$, where W is the normalisation constant. Our goal is to devise an *allocation* of goods to agents; defined as an ordered partition[1] of the m goods into n "bundles", where the bundles are (possibly empty) subsets of M, and the convention is that the i^{th} bundle in the partition is the set of goods assigned to the agent i.

The *welfare* of an allocation is a measure of the utility that the agents derive from the allocation. For additive valuations, the individual utility that an agent i derives from their bundle A_i is simply the sum of the values that they have for the goods in A_i. The overall welfare of an allocation A is typically defined by aggregating individual utilities in various ways. Not surprisingly, there are several notions of welfare corresponding to different approaches to consolidating the individual utilities. For instance, the *utilitarian social welfare* is the sum of utilities of agents under A; the *egalitarian social welfare* is the lowest utility achieved by any agent with respect to A; and the *Nash social welfare* is the geometric mean of utilities of agents under A. One may view all of these welfare notions as special cases of the p-mean welfare (where $p \in (-\infty, 0) \cup (0, 1]$), which is defined as the generalized p-mean of utilities of agents under A, i.e., $W_p(A) \coloneqq \left(\frac{1}{n} \sum_{i \in N} \big(v_i(A_i)\big)^p\right)^{1/p}$. Note that for $p > 1$, the p-mean optimal allocation tends to concentrate the distribution among fewer agents (consider the simple case of two identical agents with additive valuations who value each of two goods at 1), which is contrary to the spirit of fairness. Hence we focus on $p \leqslant 1$.

A natural goal for a fair division problem is to obtain an allocation that maximizes the overall welfare. However, observe that optimizing exclusively for welfare can lead to undesirable allocations. To see this, consider an instance with additive valuations where all the valuation functions are the same, i.e., the utility of any good g is the same for all agents in N. In this case, *every* allocation has the same utilitarian welfare. So, when we only optimize for—in this example, utilitarian—welfare, we have no way of distinguishing between, say, the allocation that allocates all goods to one agent and one that distributes the goods more evenly among the agents. To remedy this, one is typically interested in allocations that not only maximize welfare, but are also "fair".

There are several notions of fairness studied in the literature. Consider an allocation $A = (A_1, \ldots, A_n)$. We say that A is *envy-free* (EF) if for any pair of agents i and k, we have that i values A_i at least as much as they value A_k, i.e., $v_i(A_i) \geqslant v_i(A_k)$; and *equitable* (EQ) if every pair of agents i and k value their respective bundles equally, i.e., $v_i(A_i) = v_k(A_k)$. While both these fairness goals are natural, they may not be achievable, such as in a trivial instance with one good valued positively by two agents. This has motivated several approximations,

[1] Unless otherwise specified, we implicitly assume that allocations are *complete*, i.e., every good is assigned to some agent.

and in particular, the notions of *envy-freeness up to one good* and *equitability up to one good* have been widely studied. The allocation A is *envy-free up to one good* (EF1) if for any pair of agents $i, k \in N$ such that $A_k \neq \emptyset$, there is a good $g \in A_k$ such that $v_i(A_i) \geqslant v_i(A_k \setminus \{g\})$. Analogously, A is *equitable up to one good* (EQ1) if for any pair of agents $i, k \in N$ such that $A_k \neq \emptyset$, there is a good $g \in A_k$ such that $v_i(A_i) \geqslant v_k(A_k \setminus \{g\})$. For instances with additive valuations (and somewhat beyond), EF1 (and, similarly, EQ1) allocations are guaranteed to exist.

The price of fairness is informally the cost of achieving a specific fairness notion, where the cost is viewed through the lens of a particular welfare concept. For a fairness notion $\mathcal{F}$ (such as EQ1 or EF1) and a welfare notion $\mathcal{W}$ (such as egalitarian or utilitarian welfare), the price of fairness is the "worst-case ratio" of the maximum welfare (measured by $\mathcal{W}$) that can be obtained by *any* allocation, to the maximum welfare that can be obtained among allocations that are fair according to $\mathcal{F}$. For example, it is known from the work of Caragiannis et al. [15] that under additive valuations, any allocation that maximizes the Nash social welfare satisfies EF1. Thus, the price of fairness of EF1 with respect to Nash social welfare is 1. Further,Barman et al. [5] show that the price of EF1 with respect to utilitarian welfare is $O(\sqrt{n})$ for normalised subadditive valuations.

In this contribution, we focus on bounds for the price of fairness in the context of EQ1, a notion that we will henceforth refer to as the *price of equity* (PoE) when there is no ambiguity. Much of the existing literature on price of fairness analysis focuses on *specific* welfare measures (e.g., utilitarian, egalitarian, and Nash social welfare). Our work deviates from this trend by analyzing the *entire* family of generalized p-mean welfare measures (i.e., for *all* $p \leqslant 1$); recall that this captures the notions of egalitarian, utilitarian, and Nash welfare as special cases. Our results therefore address the behavior of the PoE for a wide spectrum of welfare notions.

Further, we obtain bounds in terms of the *number of agent types*—which we denote by r—rather than the total number of agents. The number of agent types of a fair division instance is the largest number of agents whose valuations are mutually distinct: in other words, it is the number of distinct valuation functions in the instance. Note that the number of agent types is potentially *much* smaller than the total number of agents. The notion of agent types has been popular in the fair division literature for the reason that it is a natural quantification of the "simplicity" of the structure of the instance as given by the valuations. Note that the well-studied special case of identical valuations is equivalent to the class of instances for which $r = 1$, and therefore one might interpret parameterizing by r as a smooth generalization of the case of identical valuations. For a representative selection of studies that focus on instances with a bounded number of agent types, we refer the reader to [10–12, 22].

We restrict ourselves to the setting of *binary submodular* (also known as matroid rank) valuations. A valuation function v_i is submodular if for any subsets of goods $S, S' \subseteq M$ such that $S \subseteq S'$, and for any good $g \notin S'$, $v_i(S \cup g) - v_i(S) \geqslant v_i(S' \cup g) - v_i(S')$. That is, the marginal value of adding g to S is at least that of

Table 1. Summary of results for the price of equity (PoE). Each cell indicates either the lower or the upper bound (columns) on PoE for a specific welfare measure (rows) as a function of the number of *agent types* r. Our contributions are highlighted by shaded boxes. The lower bounds are from Theorem 2, while the upper bounds are shown in Theorem 3 and Theorem 4. The full version of the paper [9] presents the upper and lower bounds graphically as a function of r, for $p = 1$, $p = 0$, $p = -1$, and $p = -10$.

PoE	**Agent types (r)**	
	Lower bound	Upper bound
Utilitarian welfare ($p = 1$)	$r - 1$	r
Nash welfare ($p = 0$)	$\frac{(r-1)/e}{\ln(r-1)}$	$\frac{(r-1)}{\ln(r-1)/e}$
Egalitarian welfare ($p \to -\infty$)	1	1 [34]
$p \in (0, 1)$	$p(r-1)/e$	$2r - 1$
$p \in (-1, 0)$	$2^{1/p}(r-1)^{1/(1-p)}$	$2^{-1/p}(-p)^{1/p(1-p)}(r-1)^{1/(1-p)}$
$p \leqslant -1$	$2^{1/p}(r-1)^{1/(1-p)}$	$2(r-1)^{1/(1-p)}$

adding g to a superset of S. Valuation v_i is binary submodular if for any subset of goods $S \subseteq M$ and any good g, the marginal value $v_i(S \cup g) - v_i(S) \in \{0, 1\}$. Binary submodular valuations are frequently studied in fair division and are considered to be a useful special case such as in allocating items under a budget, or with exogenous quotas [7,36]. It also provides algorithmic leverage: many computational questions of interest that are hard in general turn out to be tractable once we restrict our attention to binary submodular valuations. As an example, while it is NP-hard to compute a Nash social welfare maximizing allocation even for identical additive valuations [30], such an allocation can be computed in polynomial time under binary submodular valuations in conjunction with other desirable properties such as strategyproofness, envy-freeness up to any good, and ex-ante envy-freeness [4].

A strict subset of binary submodular valuations is the class of *binary additive* valuations—this is a subclass of additive valuations wherein each value $v_i(g)$ is either 0 or 1. Binary additive valuations provide a simple way for agents to express their preferences as they naturally align with the idea of agents "approving" or "rejecting" goods. These are also widely studied in the literature on fair division, for example, see [1,3,4,25,29]. In the case of voting too, binary additive valuations play a role. Darmann and Schauer [16] consider the complexity of maximizing Nash social welfare when scores inherent in classical voting procedures are used to associate utilities with the agents' preferences, and find that the case of approval ballots—which happen to lead to binary additive valuations—are a tractable subclass.

Our Contributions and Techniques

We now turn to a discussion of our findings (see Table 1 for a summary of our results for binary additive valuations). Given an instance of fair division with binary submodular valuations, let $A^\star$ be an allocation that maximizes the Nash social welfare. It is implicit from the results of Benabbou et al. [7] that $A^\star$ also has maximum p-mean welfare for all $p \leqslant 1$ (for details, refer to Sect. 3.1). We show an analogous result for EQ1 allocations, by demonstrating that there exists an EQ1 allocation (which we call B, or the *truncated allocation*) that maximizes the p-mean welfare for all p. To this end, in allocation $A^\star$, let i be an agent with the minimum value, and let $\ell = v_i(A_i^\star)$. If the allocation is not already EQ1, then we reallocate "excess" goods from the bundles of agents who value their bundles at more than $\ell + 1$ and give them to agent i. Notice that agent i must have marginal value 0 for all these excess goods, otherwise this reallocation would improve the Nash welfare. It turns out that this allocation B is EQ1 and also has—among EQ1 allocations—the highest p-mean welfare.

Theorem 1 ($\star$). *For any $p \in \mathbb{R} \cup \{-\infty\}$ and binary submodular valuations, the p-mean welfare of the truncated allocation B is at least that of any other EQ1 allocation.*

Notice that together with the result of Benabbou et al. [7], Theorem 1 allows us to focus only on the maximum Nash social welfare allocation $A^\star$ and the truncated allocation B to obtain upper bounds on the PoE for all $p \leqslant 1$ simultaneously.

We now describe our bounds on the PoE for binary additive valuations. Our lower bounds are based on varying parameters in a single basic instance. The parameters are r, the number of agent types, and W, the normalisation constant for the agents. Given r and W, the instance has $m = rW$ goods, divided into r groups of W goods each. The groups are $M_1, M_2, \ldots, M_r$. There are $W + 1$ agents who value all the goods in M_1 at 1 each and everything else at 0. Further, for each $2 \leqslant i \leqslant r$, we have exactly one agent who values the goods in M_i and nothing else. To summarize, we have $W + 1$ agents of the first type, who have a common interest in W goods. Any allocation must leave one of these agents with zero value. Beyond these coveted goods, each of the remaining goods is valued by exactly one agent. A welfare maximizing allocation will allocate each good in $M_2 \cup \cdots \cup M_r$ to the unique agent who values it; however, an EQ1 allocation is constrained by the fact that an agent of the first type must get value 0.[2] It turns out that using this family of instances, we can obtain the following bounds.

Theorem 2 (PoE lower bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at least:*

1. *s, for $p = 1$,*
2. *$\frac{p}{e}s$, for $p \in (0, 1)$,*
3. *$\frac{s}{e \ln s}$, for $p = 0$,*
4. *$2^{1/p} s^{1/(1-p)}$, for $p < 0$.*

[2] For $p \leqslant 0$, we use the standard convention that allocation A is a p-mean optimal allocation if (a) it maximizes number of agents with positive value, and (b) among all allocations that satisfy (a), maximizes the p-mean welfare when restricted to agents with positive value.

We now turn to the upper bounds for binary additive valuations. It turns out that the PoE for utilitarian welfare is bounded by the *rank* of the instance, where the rank is simply the rank of the $n \times m$ matrix $\{v_i(g_j)\}_{1 \leqslant i \leqslant n; 1 \leqslant j \leqslant m}$. Observe that the rank is a lower bound for the number of agent types, so this result also implies an upper bound of r on the PoE. In fact, the rank could be logarithmic in the number of agent types, and hence this is a significantly tighter bound than the number of agent types.

To obtain this upper bound, in allocation B (which, as shown in Theorem 1, maximizes the utilitarian welfare among EQ1 allocations) we show that the number of wasted goods is at most $m(1 - \frac{1}{k})$, where k is the rank of the instance. This implies the theorem.

Theorem 3 (($\star$) Utilitarian PoE upper bound). *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most the rank of the instance.*

Theorem 4 (($\star$) PoE upper bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at most*

1. $1 + s$ *for* $p = 1$
2. $1 + 2s$ *for* $p \in (0, 1)$
3. $\frac{s}{\ln(s/e)}$ *for* $p = 0$ *(i.e., the Nash social welfare)*
4. $s^{1/(1-p)} 2^{-1/p} (-1/p)^{1/p(p-1)}$ *for* $p \in (-1, 0)$
5. $2s^{1/(1-p)}$ *for* $p \leqslant -1$

We note that for any fixed p, the lower bounds (Theorem 2) and upper bounds (Theorem 4) are within a constant factor of each other.

Conceptually, for the proof of the upper bounds, we show that the worst case for the PoE is in fact the family of instances used for showing our lower bounds in Theorem 2. In particular, any instance can be transformed into one belonging to the lower bound family, without improving the PoE. Note that for the PoE, we can focus on the allocations $A^\star$ and B irrespective of the p-mean welfare measure, since these maximize the p-mean welfare for all $p \leqslant 1$ simultaneously. For a given instance, let l be the minimum value of any agent in $A^\star$. We divide the agent types into two groups: types for which every agent has value at most $l + 1$ in $A^\star$, and types for which an agent has value $> l + 1$. Note that for a type in the first group, each agent of this type retains its value in B, while for a type in the second group, the value of each agent of this type is truncated to $l + 1$. Our proof shows that agents in the first group must have total value at least W, as in the lower bound example. We also use W as an upper bound for the total value of each agent type in the second group. Then letting α be the fraction of agents in the first group, and optimizing over α, gives us the required upper bounds.

We then consider the PoE for binary additive valuations with the additional structure that both the rows and the columns are normalised. That is, each agent values exactly W goods, and each good is valued by exactly W_c agents. For such *doubly normalised* instances, we show the PoE is 1.

Theorem 5. *For doubly normalised instances under binary additive valuations, the PoE for the p-mean welfare is 1 for all $p \leqslant 1$.*

Finally, we obtain bounds on the PoE for binary submodular valuations. For identical valuations, it follows from similar results for EF1 that the PoE is 1.

Proposition 1 ($\star$). *When all agents have identical binary submodular valuations, the PoE is 1 for p-mean welfare measure for all $p \leqslant 1$.*

However, this is the best that can be obtained, in the sense that even with just *two* agent types, the PoE for utilitarian welfare is at least $n/6$, where n is the number of agents. Hence we cannot obtain bounds on the PoE that depend on the number of agent types, as we did for binary additive valuations.

Theorem 6 ($\star$). *The PoE for utilitarian welfare when agents have binary submodular valuations is at least $n/6$ (where n is the number of agents), even when there are just two types of agents.*

Nonetheless, we do obtain an upper bound of $2n$ on the PoE for binary submodular valuations.

Theorem 7 ($\star$). *For binary submodular valuations and any $p \leqslant 1$, the PoE for p-mean welfare is at most $2n$.*

Related Work

The notion of *price of fairness* was proposed in the works of Bertsimas et al. [8] and Caragiannis et al. [14]. These formulations were inspired from similar notions in game theory—specifically, *price of stability* and *price of anarchy*—that capture the loss in social welfare due to strategic behavior.[3] Caragiannis et al. [14] studied the price of fairness for divisible and indivisible resources under three fairness notions: *proportionality* [33], *envy-freeness* [19,21], and *equitability* [17]. For indivisible resources, they defined price of fairness only with respect to those instances that admit some allocation satisfying the fairness criterion.

Recently, Bei et al. [6] studied price of fairness for indivisible goods for fairness notions whose existence is guaranteed; in particular, they studied *envy-freeness up to one good* (EF1) and *maximum Nash welfare* allocations.[4] In a similar vein, Sun et al. [35] studied price of fairness for allocating indivisible *chores* for different relaxations of envy-freeness and maximin share. Perhaps closest to

[3] Price of anarchy was defined by Koutsoupias and Papadimitriou [24] and subsequently studied in the notable work of Roughgarden and Tardos [31], while price of stability was defined by Anshelevich et al. [2].

[4] The EF1 notion was formulated by Budish [13] although subsequently it was observed that an algorithm of Lipton et al. [26] achieves this guarantee for monotone valuations. The Nash social welfare function was originally proposed in the context of the bargaining problem [28] and subsequently studied for resource allocation problems by Eisenberg and Gale [18].

our work is a recent paper by Sun et al. [34]. This work studies PoE and price of equitability for any item (EQX) for indivisible goods as well as indivisible chores under utilitarian and egalitarian welfare. The valuations are assumed to be additive but not necessarily binary. For indivisible goods, the PoE is shown to be between $n-1$ and $3n$, where n is the number of agents, while for egalitarian welfare, a tight bound of 1 is provided.

2 Preliminaries

Problem Instance. An instance of the fair division problem is specified by a tuple $\langle N, M, \mathcal{V}\rangle$, where $N = \{1, 2, \ldots, n\}$ is a set of $n \in \mathbb{N}$ *agents*, $M = \{g_1, \ldots, g_m\}$ is a set of m indivisible *goods*, and $\mathcal{V} := \{v_1, v_2, \ldots, v_n\}$ is the *valuation profile* consisting of each agent's valuation function. For any agent $i \in N$, its valuation function $v_i : 2^M \to \mathbb{N} \cup \{0\}$ specifies its numerical value (or *utility*) for every subset of goods in M. For simplicity, for a valuation function v, we will denote $v(\{g\})$ as $v(g)$.

Agents i and j are said to be of the same *type* if their valuation functions are identical, i.e., if for every subset of goods $S \subseteq M$, $v_i(s) = v_j(S)$. We will use r to denote the number of distinct agent types in an instance. Further, an instance is *normalised* if for some constant W, $v_i(M) = W$ for all agents i. Our work focuses on instances with normalised valuations, since there are trivial instances where the PoE for any p-mean welfare for $p \in \mathbb{R}$ is large without this assumption (e.g., the simple instance with 2 agents and k goods, where agent 1 has value 1 for the first good and zero for the others, and agent 2 has value 1 for all goods, has PoE $k/3$ for the utilitarian welfare).

Classes of Valuation Functions. A valuation function v is:

- *monotone* if for any two subsets of goods S and T such that $S \subseteq T$, we have $v(S) \leqslant v(T)$,
- *monotone submodular* (or simply submodular) if it is monotone and for any two subsets of goods S and T such that $S \subseteq T$ and any good $g \in M \setminus T$, we have $v(S \cup \{g\}) - v(S) \geqslant v(T \cup \{g\}) - v(T)$,
- *additive* if for any subset of goods $S \subseteq M$, we have $v(S) = \sum_{g \in S} v(g)$,
- *binary submodular* (or *matroid rank*) if it is submodular and for any subset $S \subseteq M$ and any good $g \in M \setminus S$, we have $v(S \cup \{g\}) - v(S) \in \{0, 1\}$, and
- *binary additive* if it is additive and for any good $g \in M$, $v(\{g\}) \in \{0, 1\}$.

The containment relations between these classes are as follows:

$$\text{Binary additive} \subseteq \text{Additive} \subseteq \text{Submodular} \subseteq \text{Monotone}$$

and

$$\text{Binary additive} \subseteq \text{Binary submodular} \subseteq \text{Submodular} \subseteq \text{Monotone}.$$

The domains of additive and binary submodular valuations are incomparable in the sense that an instance belonging to one class may not belong to the other.

We will primarily focus on binary submodular valuations in Sect. 3 and 7, and on binary additive valuations in Sect. 4, 5, and 6.

Allocation. A *bundle* refers to any (possibly empty) subset of goods. An allocation $A \coloneqq (A_1, \ldots, A_n)$ is a partition of the set of goods M into n bundles; here, A_i denotes the bundle assigned to agent i. Given an allocation A, we say that agent i values good g if $v_i(A_i \cup \{g\}) > v_i(A_i \setminus \{g\})$. Thus if $g \in A_i$, then removing g decreases i's value. Else, assigning g to A_i increases i's value. For additive valuations, specifying an allocation is unnecessary, and we say i values g if $v_i(\{g\}) = 1$. Further, for an allocation A, we say a good $g \in A_i$ is *wasted* if $v_i(A_i \setminus \{g\}) = v_i(A_i)$, i.e., if removing it does not change the value of agent i. For additive valuations, this implies that $v_i(g) = 0$. We say an allocation (possibly partial) is *wasteful* if some good is wasted (and is *non-wasteful* or *clean* otherwise). If A is a clean allocation, then for binary submodular valuations, for each agent i, $v_i(A_i) = |A_i|$.

Fairness Notions. An allocation $A = (A_1, \ldots, A_n)$ is said to be:

- *envy-free* (EF) if for any pair of agents $i, k \in N$, we have $v_i(A_i) \geqslant v_i(A_k)$ [19, 21],
- *envy-free up to one good* (EF1) if for any pair of agents $i, k \in N$ such that $A_k \neq \emptyset$, there is a good $g \in A_k$ such that $v_i(A_i) \geqslant v_i(A_k \setminus \{g\})$ [13, 26],
- *equitable* (EQ) if for any pair of agents $i, k \in N$, we have $v_i(A_i) = v_k(A_k)$ [17], and
- *equitable up to one good* (EQ1) if for any pair of agents $i, k \in N$ such that $A_k \neq \emptyset$, there is a good $g \in A_k$ such that $v_i(A_i) \geqslant v_k(A_k \setminus \{g\})$ [20, 23].

Pareto Optimality. An allocation $A = (A_1, \ldots, A_n)$ is said to be Pareto dominated by another allocation $B = (B_1, \ldots, B_n)$ if for every agent $i \in N$, $v_i(B_i) \geqslant v_i(A_i)$ and for some agent $k \in N$, $v_k(B_k) > v_k(A_k)$. A *Pareto optimal* allocation is one that is not Pareto dominated by any other allocation.

Welfare Measures. We will now discuss various welfare measures associated with an allocation A.

- *Utilitarian social welfare* is the sum of utilities of agents under A, i.e., $\mathcal{W}^{\texttt{util}}(A) \coloneqq \sum_{i \in N} v_i(A_i)$.
- *Egalitarian social welfare* is the utility of the least-happy agent under A, i.e., $\mathcal{W}^{\texttt{egal}}(A) \coloneqq \min_{i \in N} v_i(A_i)$.
- *Nash social welfare* is the geometric mean of utilities of agents under A, i.e., $\mathcal{W}^{\texttt{Nash}}(A) \coloneqq (\Pi_{i \in N} v_i(A_i))^{1/n}$, and
- for any $p \in \mathbb{R}$, the *p-mean welfare* is the generalized p-mean of utilities of agents under A, i.e., $\mathcal{W}_p(A) \coloneqq \left(\frac{1}{n} \sum_{i \in N} v_i^p(A_i)\right)^{1/p}$.

For any $p \in \mathbb{R}$, the p-mean welfare is a strictly increasing, symmetric function of the agent values. It can be observed that utilitarian, egalitarian, and Nash welfare are all special cases of p-mean welfare for $p = 1$, $p \to -\infty$, and $p \to 0$, respectively [27].

Given an instance, it may be that in every allocation, some agent gets zero value. In this case, we need to redefine the p-mean welfare. We fix a largest subset of agents S that can simultaneously get positive value in an allocation, and then define $\mathcal{W}_p(A) = \left(\frac{1}{|S|}\sum_{i \in S} v_i^p(A_i)\right)^{1/p}$. This follows prior work on the Nash welfare, e.g., [7,15].

A *leximin* allocation is one which maximizes the minimum utility, then subject to that, it maximizes the second minimum, and so on. Thus a leximin allocation also maximizes the egalitarian welfare.

Price of Fairness. Given a fairness notion $\mathcal{F}$ (e.g., EF1, EQ1) and a p-mean welfare measure, the price of fairness of $\mathcal{F}$ with respect to a welfare measure $\mathcal{W}_p$ is the supremum over all fair division instances with n agents and m goods of the ratio of the maximum welfare (according to $\mathcal{W}_p$) of any allocation and the maximum welfare of any allocation that satisfies $\mathcal{F}$.

Formally, let $\mathcal{I}_{n,m}$ denote the set of all fair division instances with n agents and m items. Let $\mathcal{A}(I)$ denote the set of all allocations in the instance I, and further let $\mathcal{A}_{\mathcal{F}}(I)$ denote the set of all allocations in the instance I that satisfy the fairness notion $\mathcal{F}$. Then, the price of fairness (PoF) of the fairness notion $\mathcal{F}$ with respect to the welfare measure $\mathcal{W}_p$ is defined as $\mathrm{PoF}(\mathcal{F}, \mathcal{W}_p) \coloneqq \sup_{I \in \mathcal{I}_{n,m}} \frac{\max_{A^* \in \mathcal{A}(I)} \mathcal{W}_p(A^*)}{\max_{B \in \mathcal{A}_{\mathcal{F}}(I)} \mathcal{W}_p(B)}$.

As indicated earlier, throughout this paper we will focus on equitability up to one good (EQ1) as the fairness notion of choice (i.e., $\mathcal{F}$ is EQ1). For notational simplicity, we will just write PoF instead of $\mathrm{PoF}(\mathcal{F}, \mathcal{W})$ whenever the welfare measure $\mathcal{W}$ is clear from context, and we will refer to this ratio as the price of equity (PoE) whenever the fairness notion in question is EQ1.

2.1 Some Properties of p-Mean Welfare

We state here some basic properties of the p-mean welfare that will be useful. For the proof of the results marked $\star$, we refer the reader to the full version of the paper [9].

Claim 1 ($\star$). For all $p < 1$, the p-mean welfare is a concave function of the agent valuations.

Corollary 1. *Given a vector of values for n agents $x \in \mathbb{R}_+^n$ and a subset $S \subseteq N$ of agents, let x' be the vector where $x'_i = x_i$ if $i \notin S$, and $x'_i = \sum_{j \in S} x_j / |S|$ if $i \in S$. Then for all $p \leqslant 1$, $\left(\frac{1}{n}\sum_{i=1}^n (x_i)^p\right)^{1/p} \leqslant \left(\frac{1}{n}\sum_{i=1}^n (x'_i)^p\right)^{1/p}$ i.e., averaging out the value for a subset of agents weakly increases the p-mean welfare.*

Claim 2 ($\star$). Given $l \in \mathbb{N}$, and a vector $(x_1, \ldots, x_l) \in \mathbb{R}_+^l$, for $p \in [0, 1]$, $\frac{1}{l}\sum_{i=1}^l x_i^{1-p} \leqslant \left(\frac{1}{l}\sum_{i=1}^l x_i\right)^{1-p}$ while for $p < 0$, the opposite inequality holds.

3 Optimal Allocations for Binary Submodular Valuations

We first show that for obtaining bounds on the PoE for the class of binary submodular valuations (and hence, for binary additive valuations), we can focus on two allocations: the first is the Nash welfare optimal allocation $A^\star$, which obtains the optimal p-mean welfare for all $p \leqslant 1$, and the second is the truncated allocation B, which obtains the optimal p-mean welfare among all EQ1 allocations for all $p \in \mathbb{R} \cup \{-\infty\}$.

3.1 An Optimal p-Mean Welfare Allocation

Benabbou et al. [7] show the following results.

Proposition 2 ([7], Theorem 3.14). *Let $\Phi : \mathbb{Z}^n \to \mathbb{R}$ be a symmetric strictly convex function, and let $\Psi : \mathbb{Z}^n \to \mathbb{R}$ be a symmetric strictly concave function. Let A be some allocation. For binary submodular valuations, the following statements are equivalent:*

1. *A is a minimizer of Φ over all the utilitarian optimal allocations,*
2. *A is a maximizer of Ψ over all the utilitarian optimal allocations,*
3. *A is a leximin allocation, and*
4. *A maximizes Nash social welfare.*

Proposition 3 ([7], Theorem 3.11). *For binary submodular valuations, any Pareto optimal allocation is utilitarian optimal.*

For $p \leqslant 1$, if the p-mean welfare function was strictly concave, then it would follow immediately that the Nash welfare optimal allocation $A^\star$ in fact simultaneously maximizes the p-mean welfare for all $p \leqslant 1$. However, in general the p-mean welfare is concave (Claim 1), but not strictly concave. E.g., for any $p \leqslant 1$ and any vector of values $v = (v_1, \dots, v_n)$ with $v_i > 0$ for all agents i, let us overload notation slightly and define $\mathcal{W}_p(v) = \left(\frac{1}{n}\sum_{i=1}^n v_i^p\right)^{1/p}$. Then $\mathcal{W}_p(2v) = (\mathcal{W}_p(v) + \mathcal{W}_p(3v))/2$, violating strict concavity. However, we can slightly modify the proof of Theorem 3.14 from [7], to obtain the following result. The modified proof is in the full version [9].

Proposition 4 ($\star$). *For binary submodular valuations, any Nash welfare maximizing allocation (and hence, leximin allocation) simultaneously maximizes the p-mean welfare for all $p \leqslant 1$.*

The following property of leximin allocations is useful in the proof of Proposition 4.

Proposition 5 ($\star$). *For agents with binary submodular valuations, let A be a utilitarian optimal allocation so that* $\max_i v_i(A_i) \leqslant \min_i v_i(A_i) + 1$. *Then A is a leximin allocation.*

3.2 An Optimal p-Mean Welfare EQ1 allocation

We now show that similarly, there exists an EQ1 allocation B that maximizes the p-mean welfare for all p. Given $A^\star$, allocation B is obtained as follows, which we call the *truncated allocation.* Let $l = \min_i v_i(A_i^\star)$ be the smallest value that any agent obtains in $A^\star$, and let i_l be an agent that has this minimum value. Note that for any agent i, if $v_i(A_i^\star) \geqslant l + 2$, then all goods allocated to i must have marginal value 0 for the agent i_l, i.e., for all $g \in A_i^\star$, $v_{i_l}(A_{i_l}^\star \cup \{g\}) = v_{i_l}(A_{i_l}^\star)$ (else we can increase the Nash social welfare by re-allocating any good that violates this to agent i_l).

For the EQ1 allocation that we would like to construct, for any agent i with $v_i(A_i^\star) \geqslant l+2$, we remove goods from $A_i^\star$ until i's value for the remaining bundle is $l+1$. We allocate the removed goods to agent i_l (that has marginal value 0 for these goods). Let B be the resulting allocation. Then clearly, if $v_i(A_i^\star) \in \{l, l+1\}$, then $v_i(B_i) = v_i(A_i^\star)$, else $v_i(B_i) = l+1$. Thus, allocation B, our truncated NSW allocation, is EQ1.

Theorem 1 ($\star$). *For any $p \in \mathbb{R} \cup \{-\infty\}$ and binary submodular valuations, the p-mean welfare of the truncated allocation B is at least that of any other* EQ1 *allocation.*

4 Lower Bounds on the PoE for Binary Additive Valuations

Theorem 2 (PoE lower bounds). *Let $s := r - 1$. The price of equity for binary additive valuations is at least:*

1. *s, for $p = 1$,*
2. *$\frac{p}{e}s$, for $p \in (0, 1)$,*
3. $\frac{s}{e \ln s}$, for $p = 0$,
4. *$2^{1/p} s^{1/(1-p)}$, for $p < 0$.*

Note that as $p \to -\infty$, $2^{1/p} s^{1/(1-p)} \to 1$.

Proof. All our lower bounds are based on varying parameters in a single instance. The parameters are r, the number of agent types, and W, the normalisation constant for the agents. Given r, W, the instance has $m = rW$ goods, divided into r groups of W goods each. The groups are M_1, M_2, ..., M_r. There are $W + 1$ agents of the first agent type, and 1 agent each of the remaining $r - 1$ types (thus, $n = W + r$). Agents of type t have value 1 for the goods in group M_t, and value 0 for all other goods. The instance is thus *disjoint*; no good has positive value for agents of two different types.

We use Λ_p to denote the PoE for this instance. Then Λ_p is exactly

$$\Lambda_p = \left(\frac{\frac{1}{W+r-1} \left(W \times 1^p + (r-1) \times W^p\right)}{\frac{1}{W+r-1} \left(W \times 1^p + (r-1) \times 1^p\right)} \right)^{1/p} = \left(\frac{W + s \times W^p}{W + s} \right)^{1/p}.$$

For each of the cases in the theorem, we will now show how to choose W to obtain the bound claimed. For $p = 1$, choose $W = s^2$. For $p \in (0, 1)$, choose $W = ps$. For $p = 0$, choose $W = s/\ln s$. For $p < 0$, choose W so that $W = sW^p$, or $W = s^{1/(1-p)}$.

Due to lack of space, we defer the details of the calculations to the full version [9]. □

5 Upper Bounds on the PoE for Binary Additive Valuations

5.1 Upper Bounds on the PoE for $p = 1$

Theorem 3 (($\star$) Utilitarian PoE upper bound). *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most the rank of the instance.*

It follows immediately from the theorem that the PoE is also bounded by the number of agent types.

Corollary 2. *Under binary additive valuations and utilitarian welfare as the objective, the price of equity is at most r, the number of agent types.*

5.2 Upper Bounds on the PoE for $p < 1$

From Proposition 2 and Theorem 1, to bound the PoE for any $p < 1$, it suffices to obtain an upper bound on the ratio of the p-mean welfare for the two allocations $A^\star$ (which maximizes the Nash welfare) and B (the truncated allocation). We will use various properties of the allocations $A^\star$ and B in the following proofs. To state these, define T_k as the set of agents of type k, and let S_k be the set of goods allocated to agents in T_k by $A^\star$. That is, $S_k := \cup_{i \in T_k} A_i^\star$. Let $m_k := |S_k|$, and $n_k := |T_k|$. Then note that for each agent $i \in T_k$, $v_i(A_i^\star) = |A_i^\star| \in \left\{ \left\lfloor \dfrac{m_k}{n_k} \right\rfloor, \left\lceil \dfrac{m_k}{n_k} \right\rceil \right\}$.

We reindex the types in increasing order of the averaged number of goods assigned by $A^\star$, so that $m_i/n_i \leqslant m_{i+1}/n_{i+1}$. Now define

$$\lambda := \begin{cases} \lceil \frac{m_1}{n_1} \rceil & \text{if } m_1/n_1 \text{ is fractional} \\ 1 + \frac{m_1}{n_1} & \text{if } m_1/n_1 \text{ is integral}. \end{cases}$$

Thus λ is integral, $\lambda > m_1/n_1$, and $\lambda \geqslant 2$ (since the p-mean welfare is only taken over agents with positive valuation, $m_1 \geqslant n_1$). Note that in $A^\star$, the smallest value of any agent is $\lfloor m_1/n_1 \rfloor$, and $\lambda \leqslant 1 + \lfloor m_1/n_1 \rfloor$. Hence agents with value at most λ in $A^\star$ will retain their value in allocation B, while other agents will have their values truncated to λ. Now let ρ be the highest index so that $\lambda \geqslant m_\rho/n_\rho$. Thus,

$$\lambda \geqslant \frac{\sum_{i=1}^{\rho} m_i}{\sum_{i=1}^{\rho} n_i}. \tag{1}$$

As stated above, any agent of type $k \leqslant \rho$ will retain their value, i.e., $v_i(B_i) = v_i(A_i^\star)$ for an agent i of type $k \leqslant \rho$. We claim that agents of the first ρ types must have got at least W goods under $A^\star$.

Claim 3 ($\star$). $\sum_{i=1}^{\rho} m_i \geqslant W$.

Then from (1) and Claim 3, we obtain

$$\lambda \geqslant W / \sum_{i=1}^{\rho} n_i \,. \tag{2}$$

We now obtain a general expression for bounding the PoE for all $p \leqslant 1$. We will then optimize this expression for different ranges of p, to obtain upper bounds on the PoE.

Lemma 1 ($\star$). *The price of equity for p-mean welfare for instances with r types is at most*

1. $\sup_{\alpha \in [0,1]} \left(\alpha + \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p}$ *for* $p < 0$
2. $\sup_{\alpha \in [0,1]} \left(\frac{s\alpha}{1-\alpha}\right)^{(1-\alpha)}$ *for* $p = 0$,
3. $\sup_{\alpha \in [0,1]} \left(\alpha + 2^p \alpha^p s^p (1-\alpha)^{1-p}\right)^{1/p}$ *for* $p \in (0,1)$.

where as before, $s = r - 1$.

Theorem 4 (($\star$) PoE upper bounds). *Let* $s := r - 1$. *The price of equity for binary additive valuations is at most*

1. $1 + s$ *for* $p = 1$
2. $1 + 2s$ *for* $p \in (0,1)$
3. $\frac{s}{\ln(s/e)}$ *for* $p = 0$ *(i.e., the Nash social welfare)*
4. $s^{1/(1-p)} 2^{-1/p} (-1/p)^{1/p(p-1)}$ *for* $p \in (-1,0)$
5. $2s^{1/(1-p)}$ *for* $p \leqslant -1$

6 PoE Bounds for Doubly Normalised Instances

So far, we have considered instances with binary additive normalised valuations, where each agent values the same number W of goods. In this case, for the utilitarian welfare, we have seen that the PoE can be as bad as r, the number of types of agents. In this section, we consider instances with further structure. In *doubly normalised* instances, each good g is valued by the same number W_c of agents. Thus, $v_i(M) = W$ for all $i \in N$, and $\sum_{i \in N} v_i(g) = W_c$ for every good $g \in M$. The valuation matrix V is thus both row and column normalised. Such instances are intuitively "balanced," and we ask if this balance is reflected in the PoE for such instances. This indeed turns out to be the case.

Theorem 5. *For doubly normalised instances under binary additive valuations, the PoE for the p-mean welfare is 1 for all $p \leqslant 1$.*

For an undirected graph, the edge-incident matrix X has entry $X_{i,e} = 1$ if edge e is incident on vertex i, and $X_{i,e} = 0$ otherwise. We will use the following well-known property of edge-incidence matrices for bipartite graphs.

Proposition 6 (e.g., [32]). *If G is a bipartite graph, then the edge-incidence matrix of G is totally unimodular.*

Proof of Theorem 5. Let V be the valuation matrix for a doubly normalised instance, where each row sums to W and each column sums to W_c. Divide each entry by W_c. Let V^f be the resulting matrix. Then V^f satisfies: (i) each entry is either 0 or $1/W_c$, (ii) each column sums to 1, and (iii) each row sums to W/W_c. We will show that the matrix V^f can be represented as the convex combination of nonnegative integer matrices $X^1, \ldots, X^t$ so that for any matrix X^k in this decomposition, each column sums to 1 and each row sums to either $\lceil W/W_c \rceil$ or $\lfloor W/W_c \rfloor$. Assuming such a decomposition, fix any such matrix X^k in this decomposition. Clearly, due to (ii) and nonnegativity, each entry of X^k is either 1 or 0. Further if the entry $X^k_{i,g} = 1$, then $V^f_{i,g} = 1/W_c$ since V^f is a convex combination of the M-matrices, and hence $V_{i,g} = 1$. Consider then the allocation A that assigns good g to agent i if $X^k_{i,g} = 1$. In this allocation, following the properties of X^k, each good is assigned to an agent that has value 1 for it, and each agent is assigned either $\lceil W/W_c \rceil$ or $\lfloor W/W_c \rfloor$ goods. The allocation is thus EQ1 and maximizes the utilitarian welfare. Further by Proposition 5 this is also a leximin allocation, and hence by Proposition 4 and Proposition 2 this maximizes the p-mean welfare for all $p \leqslant 1$, proving the theorem. It remains to show that V can be decomposed as stated. We defer that to the full version [9]. $\square$

In the full version [9], we offer an alternate proof of Theorem 5, based on a so-called "eating argument" and an extension of Hall's theorem.

7 PoE Bounds for Binary Submodular Valuations

We now consider the more general case of binary submodular valuations. Here we focus on the utilitarian welfare, and show that our results for binary additive valuations that bound the PoE by the number of types of agents do not extend to binary submodular valuations. We first show that from prior work (see Proposition 7 below), it follows that if the agents have identical valuations, then PoE is 1 for the p-mean welfare objective for all $p \leqslant 1$.

Proposition 1 ($\star$). *When all agents have identical binary submodular valuations, the PoE is 1 for p-mean welfare measure for all $p \leqslant 1$.*

As earlier, an allocation $A = (A_1, \ldots, A_n)$ is *clean* if for all agents i, $v_i(A_i) = |A_i|$, that is, no good is wastefully allocated. We note that, given any allocation A, we can obtain a clean (possibly partial) allocation $\hat{A}$ so that $v_i(A_i) = v_i(\hat{A}_i)$ for all agents i by repeatedly removing wasted items from the allocation A. We will use the following result due to Benabbou at al. [7].

Proposition 7 ([7], Corollary 3.8). *For binary submodular valuations, any clean, utilitarian optimal (partial) allocation that minimizes $\Phi(A) := \sum_i v_i(A_i)^2$ among all utilitarian optimal allocations is EF1.*

The bound in Proposition 1 is, in a certain sense, the best that can be obtained. We will now show that with more than one type of agent under binary submodular valuations, the PoE is at least $n/6$ for utilitarian welfare. Hence we cannot obtain bounds on the PoE that depend on the number of agent types for all $p \leqslant 1$, as we did for binary additive valuations.

Theorem 6 **($\star$).** *The PoE for utilitarian welfare when agents have binary submodular valuations is at least $n/6$ (where n is the number of agents), even when there are just two types of agents.*

Theorem 7 **($\star$).** *For binary submodular valuations and any $p \leqslant 1$, the PoE for p-mean welfare is at most $2n$.*

8 Some Concluding Remarks on Chores

Our focus in the paper has been on goods, where agents have non-negative marginal utility for all items. We briefly remark on the case of bads or chores, where all marginal utilities are non-positive. Consider any instance with binary additive valuations, i.e., the value of each item is either 0 or -1. It is not hard to see that in these instances, there is always a utilitarian optimal EQ1 allocation: if chore c has value 0 for an agent i, assign c to i. The remaining chores have value -1 for all agents, and can be assigned using the round robin procedure. This allocation is clearly EQ1 and also achieves the best possible utilitarian welfare.

For more general additive instances with chores, we now show that the PoE is unbounded, even in very simple cases.[5] To this end, consider the following example with $2n$ items and $n+1$ agents. The first n agents mildly dislike the first n chores, at $-\epsilon$ and severely dislike the last n at -1, while it is the opposite for the $(n+1)$th agent, who strongly dislikes the first n items at -1 and mildly dislikes the last n at $-\epsilon$. In this example, the maximum utility is $-2n\epsilon$: assign the first n chores to the first agent and the last n chores to the last agent. On the other hand, in any EQ1 allocation, the last agent can get at most 2 chores, and hence some agent gets a chore that they value at -1. The PoE is thus at least $1/(2n\epsilon)$, which can be made arbitrarily large by choosing ϵ appropriately. Note that this instance has two item types, two agent types, and only two distinct entries in the valuation matrix. Relaxing any of these conditions implies identical valuations, where the PoE is 1; so, in some sense, this is a "minimally complex" example that already exhibits unbounded PoE. There is thus a sharp change in the PoE between instances where the values are in $\{0, -1\}$ and those where the values are in $\{-\epsilon, -1\}$. While the PoE is unbounded as ϵ approaches 0, it "snaps back" to 1 at $\epsilon = 0$.

To conclude, we obtain nearly tight bounds on the PoE in terms of agent types for the p-mean welfare spectrum. This captures, as special cases, the notions of

[5] For chores, we adopt the natural definition of PoE: the ratio of the utilitarian welfare of the best EQ1 allocation, to the maximum utilitarian welfare obtainable in any allocation. Note that if the denominator is 0, then so is the numerator (and this can be identified in polynomial time).

utilitarian, egalitarian, and Nash welfare. Our bounds are in terms of agent types (r) rather than the number of agents. Overall, our results provide a fine-grained perspective on the behavior of the PoE parameterized by p and r. In future work, it would be interesting to extend the insights that we obtain in this work beyond the domain of binary valuations. We also propose obtaining bounds on the PoE parameterized by other structural parameters, such as the number of item types. We note that for additive valuations, the rank of the valuation matrix is a lower bound on the number of item types, and hence Theorem 3 bounds the PoE in this case by the number of item types as well.

Acknowledgments. We thank the anonymous reviewers from SAGT 2023 and COMSOC 2023 for helpful comments. We are especially grateful to an anonymous reviewer from COMSOC 2023 who pointed out an error in the proof of Lemma 1 in an earlier version of the draft and provided several other helpful comments. RV acknowledges support from DST INSPIRE grant no. DST/INSPIRE/04/2020/000107 and SERB grant no. CRG/2022/002621.

References

1. Amanatidis, G., Birmpas, G., Filos-Ratsikas, A., Hollender, A., Voudouris, A.A.: Maximum nash welfare and other stories about EFX. Theoret. Comput. Sci. **863**, 69–85 (2021)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. SIAM J. Comput. **38**(4), 1602–1623 (2008)
3. Aziz, H., Rey, S.: Almost group envy-free allocation of indivisible goods and chores. In: Proceedings of the 29th International Conference on International Joint Conferences on Artificial Intelligence, pp. 39–45 (2021)
4. Babaioff, M., Ezra, T., Feige, U.: Fair and truthful mechanisms for dichotomous valuations. In: Proceedings of the 35th AAAI Conference on Artificial Intelligence, vol. 35, pp. 5119–5126 (2021)
5. Barman, S., Bhaskar, U., Shah, N.: Optimal bounds on the price of fairness for indivisible goods. In: Proceedings of the 16th International Conference on Web and Internet Economics, pp. 356–369 (2020)
6. Bei, X., Lu, X., Manurangsi, P., Suksompong, W.: The price of fairness for indivisible goods. Theory Comput. Syst. **65**(7), 1069–1093 (2021). https://doi.org/10.1007/s00224-021-10039-8
7. Benabbou, N., Chakraborty, M., Igarashi, A., Zick, Y.: Finding fair and efficient allocations for matroid rank valuations. ACM Trans. Econ. Comput. **9**(4), 1–41 (2021)
8. Bertsimas, D., Farias, V.F., Trichakis, N.: The price of fairness. Oper. Res. **59**(1), 17–31 (2011)
9. Bhaskar, U., Misra, N., Sethia, A., Vaish, R.: The Price of Equity with Binary Valuations and Few Agent Types. arXiv preprint arXiv:2307.06726 (2023)
10. Bliem, B., Bredereck, R., Niedermeier, R.: Complexity of efficient and envy-free resource allocation: few agents, resources, or utility levels. In: Proceedings of the 25th International Joint Conference on Artificial Intelligence, pp. 102–108 (2016)

11. Bouveret, S., Cechlárová, K., Elkind, E., Igarashi, A., Peters, D.: Fair division of a graph. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence, pp. 135–141 (2017)
12. Brânzei, S., Lv, Y., Mehta, R.: To give or not to give: fair division for single minded valuations. In: Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, pp. 123–129 (2016)
13. Budish, E.: The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. J. Polit. Econ. **119**(6), 1061–1103 (2011)
14. Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., Kyropoulou, M.: The efficiency of fair division. Theory Comput. Syst. **50**(4), 589–610 (2012)
15. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J.: The unreasonable fairness of maximum nash welfare. ACM Trans. Econ. Comput. **7**(3), 1–32 (2019)
16. Darmann, A., Schauer, J.: Maximizing nash product social welfare in allocating indivisible goods. Eur. J. Oper. Res. **247**(2), 548–559 (2015)
17. Dubins, L.E., Spanier, E.H.: How to cut a cake fairly. Am. Math. Mon. **68**(1P1), 1–17 (1961)
18. Eisenberg, E., Gale, D.: Consensus of subjective probabilities: the Pari-Mutuel method. Ann. Math. Stat. **30**(1), 165–168 (1959)
19. Foley, D.: Resource Allocation and the Public Sector. Yale Economic Essays, pp. 45–98 (1967)
20. Freeman, R., Sikdar, S., Vaish, R., Xia, L.: Equitable allocations of indivisible goods. In: Proceedings of the 28th International Joint Conference on Artificial Intelligence, pp. 280–286 (2019)
21. Gamow, G., Stern, M.: Puzzle-Math (1958)
22. Garg, J., Kulkarni, P., Murhekar, A.: On fair and efficient allocations of indivisible public goods. In: Proceedings of the 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science. LIPIcs, vol. 213, pp. 22:1–22:19 (2021)
23. Gourvès, L., Monnot, J., Tlilane, L.: Near fairness in matroids. In: Proceedings of the 21st European Conference on Artificial Intelligence, pp. 393–398 (2014)
24. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. Comput. Sci. Rev. **3**(2), 65–69 (2009)
25. Kyropoulou, M., Suksompong, W., Voudouris, A.A.: Almost envy-freeness in group resource allocation. Theoret. Comput. Sci. **841**, 110–123 (2020)
26. Lipton, R.J., Markakis, E., Mossel, E., Saberi, A.: On approximately fair allocations of indivisible goods. In: Proceedings of the 5th ACM Conference on Electronic Commerce, pp. 125–131 (2004)
27. Moulin, H.: Fair Division and Collective Welfare. MIT Press, Cambridge (2004)
28. Nash Jr., J.F.: The bargaining problem. Econometrica 155–162 (1950)
29. Ortega, J.: Multi-unit assignment under dichotomous preferences. Math. Soc. Sci. **103**, 15–24 (2020)
30. Roos, M., Rothe, J.: Complexity of social welfare optimization in multiagent resource allocation. In: Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems, pp. 641–648 (2010)
31. Roughgarden, T., Tardos, É.: How bad is selfish routing? J. ACM **49**(2), 236–259 (2002)
32. Schrijver, A.: Theory of Linear and Integer Programming. Wiley, Hoboken (1998)
33. Steinhaus, H.: The problem of fair division. Econometrica **16**(1), 101–104 (1948)
34. Sun, A., Chen, B., Doan, X.V.: Equitability and welfare maximization for allocating indivisible items. Auton. Agents Multi-Agent Syst. **37**(8) (2023)

35. Sun, A., Chen, B., Vinh Doan, X.: Connections between fairness criteria and efficiency for allocating indivisible chores. In: Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, pp. 1281–1289 (2021)
36. Viswanathan, V., Zick, Y.: A general framework for fair allocation with matroid rank valuations. In: Proceedings of the 24th ACM Conference on Economics and Computation, pp. 1129–1152 (2023)