The Price of Fairness for Indivisible Goods

Xiaohui Bei[1] , **Xinhang Lu**[1] , **Pasin Manurangsi**[2] and **Warut Suksompong**[3]

[1] School of Physical and Mathematical Sciences, Nanyang Technological University
[2]Department of Electrical Engineering and Computer Sciences, UC Berkeley
[3]Department of Computer Science, University of Oxford
xhbei@ntu.edu.sg, xinhang001@e.ntu.edu.sg, pasin@berkeley.edu, warut.suksompong@cs.ox.ac.uk

Abstract

We investigate the efficiency of fair allocations of indivisible goods using the well-studied *price of fairness* concept. Previous work has focused on classical fairness notions such as envy-freeness, proportionality, and equitability. However, these notions cannot always be satisfied for indivisible goods, leading to certain instances being ignored in the analysis. In this paper, we focus instead on notions with guaranteed existence, including envy-freeness up to one good (EF1), balancedness, maximum Nash welfare (MNW), and leximin. We mostly provide tight or asymptotically tight bounds on the worst-case efficiency loss for allocations satisfying these notions.

1 Introduction

The allocation of scarce resources among interested agents is a problem that arises frequently and plays a major role in our society. We often want to ensure that the allocation that we select is *fair* to the agents—the literature of *fair division*, which dates back to the design of cake-cutting algorithms over half a century ago [Steinhaus, 1948; Dubins and Spanier, 1961], provides several ways of defining what fair means. An issue orthogonal to fairness is *efficiency*, or *social welfare*, which refers to the total happiness of the agents. A fundamental question is therefore how much efficiency we might lose if we want our allocation to be fair.

This question was first addressed by Caragiannis *et al.* [2012], who introduced the *price of fairness* concept to capture the efficiency loss due to fairness constraints. For any fairness notion and any given resource allocation instance with additive valuations, they defined the price of fairness of the instance to be the ratio between the maximum social welfare over all allocations and the maximum social welfare over allocations that are fair according to the notion. The overall price of fairness for this notion is then defined as the largest price of fairness across all instances. Caragiannis et al. considered the classical fairness notions of *envy-freeness*, *proportionality* and *equitability*, and presented a series of results on the price of fairness with respect to these notions. As an example, they showed that for the allocation of indivisible goods among n agents, the price of proportionality is $n - 1 + 1/n$, meaning that the efficiency of the best proportional allocation can be a linear factor away from that of the best allocation overall.

Caragiannis et al.'s work sheds light on the trade-off between efficiency and fairness in the allocation of both divisible and indivisible resources. However, a significant limitation of their study is that while an allocation satisfying each of the three fairness notions always exists when goods are divisible, this is not the case for indivisible goods. Indeed, none of the notions can be satisfied in the simple instance with at least two agents and a single good to be allocated. Caragiannis et al. circumvented this issue by ignoring instances in which the fairness notion in question cannot be satisfied. As a result, their price of fairness analysis, which is meant to capture the worst-case efficiency loss, fails to cover certain scenarios that may arise in practice.[1] In addition, the fact that certain instances are not taken into account in the price of fairness have seemingly contradictory consequences. For example, since envy-free allocations are always proportional when valuations are additive, it may appear at first glance that the price of envy-freeness must be at least as high as the price of proportionality. This is not necessarily the case, however, because there are instances that admit proportional but no envy-free allocations.[2]

To address these limitations, in this paper we study the price of fairness for indivisible goods with respect to fairness notions that can be satisfied in every instance. Among other notions, we consider envy-freeness up to one good (EF1), balancedness, maximum Nash welfare (MNW), and leximin.[3] In addition to deriving bounds on the price of fairness for these notions, we also introduce the concept of *strong price of fairness*, which captures the efficiency loss in the worst fair allocation as opposed to that in the best fair allocation. The relationship between the price of fairness and the strong price

[1]From the above example, one may think that such scenarios are rare exceptions. However, for envy-freeness, these scenarios are in fact common if the number of goods is not too large compared to the number of agents [Dickerson *et al.*, 2014; Manurangsi and Suksompong, 2019].

[2]Indeed, the instance that Caragiannis et al. used to show that the price of proportionality is at least $n - 1 + 1/n$ admits no envy-free allocation. Thus, it is still possible that the price of envy-freeness is lower than the price of proportionality.

[3]See Section 2 for the definitions of these notions.

Property P	**Price of** P		**Strong price of** P	
	General n	$n = 2$	General n	$n = 2$
Envy-freeness up to one good (EF1)	LB: $\Omega(\sqrt{n})$ UB: $O(n)$	LB: $8/7$ UB: $2/\sqrt{3}$	∞	∞
Envy-freeness up to any good (EFX)	–	$3/2$	–	∞
Round-robin (RR)	n	2	n^2	4
Balancedness (BAL)	$\Theta(\sqrt{n})$	$4/3$	∞	∞
Maximum Nash welfare (MNW)	$\Theta(n)$	LB: $27/23$ UB: $5/4$	$\Theta(n)$	LB: $27/23$ UB: $5/4$
Maximum egalitarian welfare (MEW)	$\Theta(n)$	$3/2$	∞ for $n \geq 3$	$3/2$
Leximin (LEX)	$\Theta(n)$	$3/2$	$\Theta(n)$	$3/2$
Pareto optimality (PO)	1	1	$\Theta(n^2)$	3

Table 1: Summary of our results. LB denotes lower bound and UB denotes upper bound. We do not consider the (strong) price of EFX for $n > 2$ because it is not known whether an EFX allocation always exists. If we allow dependence on the number of goods m, we have an upper bound of $O(\sqrt{n}\log(mn))$ on the price of EF1.

of fairness is akin to that between the price of stability and the price of anarchy for equilibria. While the strong price of fairness is too demanding to yield any nontrivial guarantee for some fairness notions, as we will see, it does provide meaningful guarantees for other notions.

1.1 Our Results

The majority of our results can be found in Table 1; we highlight a subset of these next. For the price of EF1, we provide a lower bound of $\Omega(\sqrt{n})$ and an upper bound of $O(n)$. We then show that two common ways to obtain an EF1 allocation—the round-robin algorithm and MNW—have a price of fairness of linear order (for round-robin the price is exactly n), implying that these methods cannot be used to improve the upper bound for EF1. We also show that improving this upper bound would yield a corresponding improvement on the price of envy-freeness gap for divisible goods left open by Caragiannis *et al.* [2012]. On the other hand, if we allow dependence on the number of goods m, the price of EF1 is $O(\sqrt{n}\log(mn))$—this means that the $\Omega(\sqrt{n})$ lower bound is almost tight unless the number of goods is huge compared to the number of agents. For MNW, maximum egalitarian welfare (MEW), and leximin, we prove an asymptotically tight bound of $\Theta(n)$ on the price of fairness. Moreover, with the exception of EF1 and MNW, we establish exactly tight bounds in the case of two agents for all fairness notions.

On the strong price of fairness front, we show via a simple instance that the strong price of EF1 and balancedness are infinite, meaning that there are arbitrarily bad EF1 and balanced allocations. Nevertheless, a round-robin allocation, which satisfies these two properties, always has welfare within a factor n^2 of the optimal allocation, and this factor is exactly tight. For MNW and leximin, the strong price of fairness, like the price of fairness, is of linear order. However, while the price of MEW is also $\Theta(n)$, the strong price of MEW is infinite for $n \geq 3$ (and $3/2$ for $n = 2$). Finally, we consider Pareto optimality, for which the price of fairness is trivially 1. We show that the strong price of Pareto optimality is $\Theta(n^2)$.

1.2 Related Work

The price of fairness was introduced independently by Bertsimas *et al.* [2011] and Caragiannis *et al.* [2012]. Bertsimas et al. studied the concept for divisible goods with respect to fairness notions such as proportional fairness and max-min fairness. Caragiannis et al. presented a number of bounds for both goods and *chores* (i.e., items that yield negative utility), both when these items are divisible and indivisible. The price of fairness has subsequently been examined in several other settings, including for *contiguous* allocations of divisible goods [Aumann and Dombb, 2015], indivisible goods [Suksompong, 2019], and divisible chores [Heydrich and van Stee, 2015], as well as in the context of machine scheduling [Bilò *et al.*, 2016].

Typically, the price of fairness study focuses on quantifying the efficiency loss solely in terms of the number of agents. A notable exception to this is the work of Kurz [2014], who remarked that certain constructions used to establish worst-case bounds for indivisible goods require a large number of goods. As a result, Kurz investigated the dependence of the price of fairness on both the number of agents and the number of goods, and found that the price indeed improves significantly if we limit the number of goods.

2 Preliminaries

Denote by $N = \{1, 2, \ldots, n\}$ the set of agents and $M = \{1, 2, \ldots, m\}$ the set of goods. Each agent i has a nonnegative utility $u_i(j)$ for each good j. The agents' utilities are additive, meaning that $u_i(M') = \sum_{j \in M'} u_i(j)$ for every agent i and subset of goods $M' \subseteq M$. Following Caragiannis *et al.* [2012], we normalize the utilities across agents by assuming that $u_i(M) = 1$ for all i. We refer to a setting with agents, goods, and utility functions as an *instance*. An *allocation* is a partition of M into bundles $(M_1, \ldots, M_n)$ such that agent i receives bundle M_i. The *(utilitarian) social welfare* of an allocation $\mathcal{M}$ is defined as $\mathsf{SW}(\mathcal{M}) := \sum_{i=1}^{n} u_i(M_i)$. The optimal social welfare for an instance I, denoted by $\mathsf{OPT}(I)$, is the maximum social welfare over all allocations for this instance.

A *property* P is a function that maps every instance I to a (possibly empty) set of allocations $P(I)$. Every allocation in $P(I)$ is said to satisfy property P.

We are now ready to define the price of fairness concepts.

Definition 2.1. For any given property P of allocations and any instance, we define the *price of P* for that instance to be the ratio between the optimal social welfare and the maximum social welfare over allocations satisfying P:

$$\text{Price of } P \text{ for instance } I = \frac{\mathsf{OPT}(I)}{\max_{\mathcal{M}\in P(I)} \mathsf{SW}(\mathcal{M})}.$$

The overall *price of* P is then defined as the supremum price of fairness across all instances.

Similarly, the *strong price of* P for a given instance is the ratio between the optimal social welfare and the *minimum* social welfare over allocations satisfying P:

$$\text{Strong price of } P \text{ for instance } I = \frac{\mathsf{OPT}(I)}{\min_{\mathcal{M}\in P(I)} \mathsf{SW}(\mathcal{M})}.$$

The overall *strong price of* P is then defined as the supremum price of fairness across all instances.

We will only consider properties P such that $P(I)$ is nonempty for every instance I, so the (strong) price of fairness is always well-defined. With the exception of Theorem 3.7, we will be interested in the price of fairness as a function of n, and assume that m can be arbitrary.

Next, we define the fairness properties that we consider. The first two properties are relaxations of the classical envy-freeness notion.

Definition 2.2 (EF1)**.** An allocation is said to satisfy *envy-freeness up to one good (EF1)* if for every pair of agents i, i', there exists a set $A_{i'} \subseteq M_{i'}$ with $|A_{i'}| \leq 1$ such that $u_i(M_i) \geq u_i(M_{i'} \backslash A_{i'})$.

Definition 2.3 (EFX)**.** An allocation is said to satisfy *envy-freeness up to any good (EFX)* if for every pair of agents i, i' and every good $g \in M_{i'}$, we have $u_i(M_i) \geq u_i(M_{i'} \backslash \{g\})$.

It is clear that EFX imposes a stronger requirement than EF1. An EF1 allocation always exists [Lipton *et al.*, 2004], while for EFX the existence question is still unresolved [Caragiannis *et al.*, 2016]. As such, we will only consider EFX in the case of two agents, for which existence is guaranteed [Plaut and Roughgarden, 2018].

The round-robin algorithm, which we describe below, always computes an EF1 allocation (see, e.g., [Caragiannis *et al.*, 2016]).

Definition 2.4 (RR)**.** The *round-robin algorithm* works by arranging the agents in some arbitrary order, and letting the next agent in the order choose her favorite good from the remaining goods.[4] An allocation is said to satisfy *round-robin (RR)* if it is the result of applying the algorithm with some ordering of the agents.

[4]In case there are ties between goods, we may assume worst-case tie breaking, since it is possible to obtain an instance with infinitesimal difference in welfare and any desired tie-breaking between goods by slightly perturbing the utilities.

Our next property is balancedness, which means that the goods are as spread out among the agents as possible. Balancedness and similar cardinality constraints have been considered in recent work [Biswas and Barman, 2018]. In addition to satisfying EF1, an allocation produced by the round-robin algorithm is also balanced.

Definition 2.5 (BAL)**.** An allocation is said to be *balanced (BAL)* if $|M_i - M_j| \leq 1$ for any i, j.

Next, we define a number of welfare maximizers.

Definition 2.6 (MNW)**.** The *Nash welfare* of an allocation is defined as $\prod_{i\in N} u_i(M_i)$. An allocation is said to be a *maximum Nash welfare (MNW)* allocation if it has the maximum Nash welfare among all allocations.[5]

Definition 2.7 (MEW)**.** The *egalitarian welfare* of an allocation is defined as $\min_{i\in N} u_i(M_i)$. An allocation is said to be a *maximum egalitarian welfare (MEW)* allocation if it has the maximum egalitarian welfare among all allocations.

Definition 2.8 (LEX)**.** An allocation is said to be *leximin (LEX)* if it maximizes the lowest utility (i.e., the egalitarian welfare), and, among all such allocations, maximizes the second lowest utility, and so on.

Finally, we define Pareto optimality. While this is an efficiency notion rather than a fairness notion, we also consider it as it is a fundamental property in the context of resource allocation.

Definition 2.9 (PO)**.** Given an allocation $(M_1, \ldots, M_n)$, another allocation $(M'_1, \ldots, M'_n)$ is said to be a *Pareto improvement* if $u_i(M'_i) \geq u_i(M_i)$ for all i with at least one strict inequality. An allocation is *Pareto optimal (PO)* if it does not admit a Pareto improvement.

Caragiannis *et al.* [2016] showed that a MNW allocation always satisfies EF1 and Pareto optimality. It is clear from the definition that any leximin allocation is Pareto optimal and maximizes egalitarian welfare. The problem of computing a MEW allocation has been studied by Bezáková and Dani [2005] and Bansal and Sviridenko [2006]. Leximin allocations were studied by Bogomolnaia and Moulin [2004] and shown to be applicable in practice by Kurokawa *et al.* [2015].

All omitted proofs can be found in the full version of this paper [Bei *et al.*, 2019].

3 Envy-Freeness

In this section, we consider envy-freeness relaxations and the round-robin algorithm, which always produces an EF1 allocation. We begin with a lower bound on the price of EF1.

Theorem 3.1. *The price of EF1 is* $\Omega(\sqrt{n})$.

Proof. Let $m = n, r = \lfloor\sqrt{n}\rfloor$, and assume that the utilities are as follows:

- For $i = 1, \ldots, r-1$: $u_i((i-1)r + j) = \frac{1}{r}$ for $j = 1, \ldots, r$, and $u_i(j) = 0$ otherwise.

[5]In the case where the maximum Nash welfare is 0, an allocation is a MNW allocation if it gives positive utility to a set of agents of maximal size and moreover maximizes the product of utilities of the agents in that set.

- $u_r(j) = \frac{1}{n-r(r-1)}$ for $j = r(r-1)+1, \ldots, n$, and $u_r(j) = 0$ otherwise.
- For $i = r+1, \ldots, n$: $u_i(j) = \frac{1}{n}$ for all j.

Consider the allocation that assigns goods $ir-r+1, \ldots, ir$ to agent i for $i = 1, \ldots, r-1$ and the remaining goods to agent r. The social welfare of this allocation is r. On the other hand, in any EF1 allocation, each of the agents $i = r+1, \ldots, n$ must receive at least one good—otherwise some agent would receive at least two goods and agent i would envy her. This means that the social welfare is at most $r \cdot \frac{1}{r} + (n-r) \cdot \frac{1}{n} < 2$. Hence the price of EF1 is at least $\frac{r}{2} = \frac{\lfloor\sqrt{n}\rfloor}{2}$. □

For two agents, we establish an almost tight bound on the price of EF1 and a tight bound on the price of EFX.

Theorem 3.2. *For $n = 2$, the price of EF1 is at least $\frac{8}{7} \approx 1.143$ and at most $\frac{2}{\sqrt{3}} \approx 1.155$.*

Theorem 3.3. *For $n = 2$, the price of EFX is $3/2$.*

Next, we give a simple instance showing that EF1 and EFX allocations can have arbitrarily bad welfare.

Theorem 3.4. *The strong price of EF1 is ∞. For $n = 2$, the strong price of EFX is ∞.*

Proof. Let $m = n$, and assume that $u_i(i) = 1$ for all i and $u_i(j) = 0$ otherwise. The allocation that assigns good i to agent i for every i has social welfare n. On the other hand, the allocation that assigns good $i-1$ to agent i for $i = 2, \ldots, n$ and good n to agent 1 is EF1 and EFX, but has social welfare 0. The conclusion follows. □

We now turn our attention to the round-robin algorithm. We show that it is always possible to order the agents to obtain a welfare of 1.

Lemma 3.5. *For any instance, there exists an ordering of the agents such that the round-robin algorithm implemented with this ordering produces an allocation with social welfare at least 1, and this bound is tight.*

Proof. We claim that if we choose the ordering of the agents uniformly at random, the expected social welfare is at least 1. The desired bound immediately follows from this claim.

To prove the claim, consider an arbitrary agent i, and assume without loss of generality that $u_i(1) \geq u_i(2) \geq \ldots \geq u_i(m)$. Note that if the agent is ranked jth in the ordering, her utility is at least $u_i(j) + u_i(n+j) + u_i(2n+j) + \cdots + u_i(kn+j)$, where $k = \lfloor (m-j)/n \rfloor$. Hence, the agent's expected utility is at least

$$\frac{1}{n} \cdot \sum_{j=1}^{n} \sum_{r=0}^{\lfloor (m-j)/n \rfloor} u_i(rn+j) = \frac{1}{n} \cdot \sum_{j=1}^{m} u_i(j) = \frac{1}{n}.$$

It follows from linearity of expectation that the expected social welfare is at least $n \cdot \frac{1}{n} = 1$, as claimed.

The tightness of the bound follows from the instance where every agent has utility 1 for the same good. □

Lemma 3.5 yields a linear price of fairness for round-robin.

Theorem 3.6. *The price of round-robin is n. Consequently, the price of EF1 is at most n.*

Proof. For the upper bound, consider an arbitrary instance. Since every agent receives utility at most 1, the optimal social welfare is at most n. On the other hand, by Lemma 3.5, there exists an ordering of the agents such that the round-robin algorithm yields welfare at least 1. Hence the price of round-robin is at most n.

We now turn to the lower bound. Let $m = x^n$ for some large x that is divisible by n, and assume that the utilities are such that for each agent i, $u_i(j) = 1/x^i$ for $j = 1, \ldots, x^i$ and $u_i(j) = 0$ otherwise.

Consider the allocation that assigns goods $1, \ldots, x$ to agent 1, and $x^{i-1}+1, \ldots, x^i$ to agent i for every $i \geq 2$. In this allocation, agent 1 gets utility 1, while each remaining agent gets utility $(x^i - x^{i-1})/x^i = 1 - 1/x$. The social welfare is therefore $n - (n-1)/x$. This converges to n for large x.

On the other hand, consider the round-robin algorithm with an arbitrary ordering of the agents, and assume without loss of generality that agents always break ties in favor of goods with lower numbers. Hence, regardless of the ordering, the goods get chosen in the order $1, 2, \ldots, m$. As a result, every agent gets exactly $1/n$ of their valued goods, so her utility is $1/n$, and the social welfare is 1. Hence the price of round-robin is n. □

The argument for the lower bound in Theorem 3.6 works even if we can choose a new ordering of the agents in every round. This means that the fixed order is not a barrier to obtaining a better price of fairness, but rather the "each agent picks exactly once in every round" aspect of the algorithm.

One may notice that the lower bound construction uses an exponential number of goods. This is in fact necessary to obtain an instance with a high price of round-robin. As we show next, the $\Omega(\sqrt{n})$ lower bound on the price of EF1 is almost tight as long as m is not too large compared to n.

Theorem 3.7. *The price of round-robin is $O(\sqrt{n}\log(mn))$. Consequently, the price of EF1 is $O(\sqrt{n}\log(mn))$.*

Proof. Consider any instance I. We claim that there exists an ordering for which the round-robin algorithm produces an allocation with social welfare at least $\frac{\mathsf{OPT}(I)}{65\sqrt{n}\log_2(mn)}$. First, observe that if $\mathsf{OPT}(I) \leq 65\sqrt{n}\log_2(mn)$, then Lemma 3.5 immediately yields the desired claim. Henceforth, we will only focus on the case where $\mathsf{OPT}(I) > 65\sqrt{n}\log_2(mn)$.

Fix an optimal allocation $\mathcal{M} = (M_1, \ldots, M_n)$, and let $r := \lceil \log_2(m\sqrt{n}) \rceil$. For each $i \in N$, let us partition M_i into $M_i^0 \cup M_i^1 \cup \cdots \cup M_i^r$, where M_i^ℓ is defined by

$$M_i^\ell = \begin{cases} \{j \in M_i \mid u_i(j) \in (2^{-\ell-1}, 2^{-\ell}]\} & \text{if } \ell \neq r; \\ \{j \in M_i \mid u_i(j) \in [0, 2^{-\ell}]\} & \text{if } \ell = r. \end{cases}$$

Define $M^\ell := \cup_{i=1}^n M_i^\ell$ and $\mathsf{SW}_\ell(\mathcal{M}) := \sum_{i=1}^n u_i(M_i^\ell)$.

Let $\ell^* := \arg\max_{\ell \in \{0,\ldots,r-1\}} \mathsf{SW}_\ell(\mathcal{M})$. We have

$$\mathsf{SW}_{\ell^*}(\mathcal{M}) \geq \frac{1}{r}\left(\sum_{\ell=0}^{r-1} \mathsf{SW}_\ell(\mathcal{M})\right) = \frac{\mathsf{OPT}(I) - \mathsf{SW}_r(\mathcal{M})}{r}.$$

However, since agent i values each item in M_i^r at most $2^{-r} \leq \frac{1}{m\sqrt{n}}$, we have $u_i(M_i^r) \leq 1/\sqrt{n}$. This implies that $\mathsf{SW}_r(\mathcal{M}) \leq \sqrt{n}$, which is no more than $\mathsf{OPT}(I)/65$. Hence,

$$\mathsf{SW}_{\ell^*}(\mathcal{M}) \geq \frac{64}{65r} \cdot \mathsf{OPT}(I) \geq \frac{32 \cdot \mathsf{OPT}(I)}{65 \log_2(mn)}. \qquad (1)$$

Thus, it suffices to show the existence of an ordering such that round-robin produces an allocation with social welfare at least $\mathsf{SW}_{\ell^*}(\mathcal{M})/\sqrt{n}$.

Observe that (1) implies that $\mathsf{SW}_{\ell^*}(\mathcal{M}) > 32\sqrt{n}$. We only consider the case $T := |M^{\ell^*}| > 2n$ here and leave the case $T \leq 2n$ to the full version of this paper [Bei *et al.*, 2019]. Since $u_i(M_i^{\ell^*}) \leq 2^{-\ell^*}|M_i^{\ell^*}|$ for each i, we have $\mathsf{SW}_{\ell^*}(\mathcal{M}) \leq 2^{-\ell^*}T$.

Assume that $T > 2n$. We will show that the round-robin algorithm with arbitrary ordering yields an allocation with social welfare at least $\mathsf{SW}_{\ell^*}(\mathcal{M})/\sqrt{n}$.

To see this, let us consider the round-robin procedure with arbitrary ordering, and consider the set of goods that are picked in the first $t := \lfloor T/(2n) \rfloor$ rounds; let $S_t \subseteq M$ denote this set. Now, observe that

$$\sum_{i=1}^{n} |M_i^{\ell^*} \setminus S_t| \geq T - |S_t| = T - n \cdot t \geq \frac{T}{2}.$$

This implies that

$$\sum_{i=1}^{n} u_i(M_i^{\ell^*} \setminus S_t) \geq \frac{T}{2} \cdot 2^{-\ell^*-1} \geq \frac{\mathsf{SW}_{\ell^*}(\mathcal{M})}{4} > 8\sqrt{n}.$$

Since $u_i(M_i^{\ell^*} \setminus S_t) \leq 1$, there must be more than $8\sqrt{n}$ agents such that $M_i^{\ell^*} \not\subseteq S_t$. Let N^* denote the set of such agents.

We claim that, in each of the first t rounds, every agent $i \in N^*$ must receive an item she values at least $2^{-\ell^*-1}$. The reason is that agent i picks her favorite good, which she must value at least as much as the good(s) left unpicked in $M_i^{\ell^*} \setminus S_t$. Moreover, she values the latter at least $2^{-\ell^*-1}$, so this must also be a lower bound of her utility for the former.

From the claim in the previous paragraph, we can conclude that the social welfare of the allocation produced is at least

$$|N^*| \cdot t \cdot 2^{-\ell^*-1} > 8\sqrt{n} \cdot \frac{T}{4n} \cdot 2^{-\ell^*-1} \geq \frac{\mathsf{SW}_{\ell^*}(\mathcal{M})}{\sqrt{n}}$$

as desired. Note that we use the assumption $T > 2n$ to conclude that $t \geq T/(4n)$ in the first inequality above. □

While Theorem 3.7 shows that the price of EF1 is close to $\Theta(\sqrt{n})$ unless the number of goods is huge, if we are only interested in the dependence on the number of agents, the gap still remains between $\Omega(\sqrt{n})$ and $O(n)$. In fact, Caragiannis *et al.* [2012] left exactly the same gap on the price of envy-freeness for *divisible* goods. In the full version of this paper [Bei *et al.*, 2019], we exhibit an interesting connection between the indivisible and divisible goods settings by showing that the price of EF1 for indivisible goods is always at least the price of envy-freeness for divisible goods. This implies that improving the $O(n)$ upper bound on the price of EF1 would also yield a corresponding improvement on the price of envy-freeness.

We end this section by establishing an exact bound on the strong price of round-robin.

Theorem 3.8. *The strong price of round-robin is* n^2.

4 Balancedness

In this section, we consider balancedness. We begin by establishing an asymptotically tight bound on the price of balancedness.

Theorem 4.1. *The price of balancedness is* $\Theta(\sqrt{n})$.

Proof. For the lower bound, consider the instance in Theorem 3.1. The social welfare can be as high as $r = \lfloor \sqrt{n} \rfloor$, while a similar argument shows that the social welfare of any balanced allocation is at most 2. The conclusion follows.

We now turn to the upper bound. We claim that for any instance I, the maximum social welfare of a balanced allocation is always within a factor $4\sqrt{n}$ of the optimal social welfare; this claim implies the desired upper bound. If $\mathsf{OPT}(I) \leq 4\sqrt{n}$, the claim follows immediately from Lemma 3.5. We therefore assume that $\mathsf{OPT}(I) > 4\sqrt{n}$. We will show that there is a balanced allocation $\mathcal{M}$ such that $\mathsf{SW}(\mathcal{M}) \geq \frac{\mathsf{OPT}(I)-\sqrt{n}}{2\sqrt{n}}$; this suffices for our claim because $\frac{\mathsf{OPT}(I)-\sqrt{n}}{2\sqrt{n}} \geq \frac{\mathsf{OPT}(I)}{4\sqrt{n}}$. We only consider the case $m \geq n$ here and defer the case $m < n$ to the full version of this paper [Bei *et al.*, 2019].

Assume that $m \geq n$. Fix an optimal allocation, and let A be the set of agents who receive at least $\frac{m}{\sqrt{n}}$ goods in the optimal allocation, and B the complement set of agents. Since there are at most $\sqrt{n}$ agents in A, they contribute at most $\sqrt{n}$ to $\mathsf{OPT}(I)$, so the agents in B contribute at least $\mathsf{OPT}(I) - \sqrt{n}$. We let each agent in B keep her $\lceil \frac{m}{2n} \rceil$ most valuable goods (or all of her goods, if she has fewer than this number of goods). This yields a total utility of at least $\frac{\mathsf{OPT}(I)-\sqrt{n}}{2\sqrt{n}}$. Since $\lceil \frac{m}{2n} \rceil \leq \lfloor \frac{m}{n} \rfloor$ due to the assumption $m \geq n$, the remaining goods can be reallocated to obtain a balanced allocation, which has social welfare at least $\frac{\mathsf{OPT}(I)-\sqrt{n}}{2\sqrt{n}}$, as desired. □

For two agents, we give an exact bound on the welfare that can be lost due to imposing balancedness.

Theorem 4.2. *For* $n = 2$*, the price of balancedness is* $4/3$.

Finally, the same construction as in Theorem 3.4 shows that balanced allocations can have arbitrarily bad welfare.

Theorem 4.3. *The strong price of balancedness is* ∞.

5 Welfare Maximizers

In this section, we consider allocations that maximize different measures of welfare. To start with, we show that every MNW and leximin allocation yields a decent welfare.

Lemma 5.1. *For any instance, every MNW allocation and every leximin allocation has social welfare at least* 1*, and both bounds are tight.*

Proof. We first establish the bound for MNW. Consider any MNW allocation where agent i receives bundle M_i, and assume for contradiction that $\sum_{k=1}^{n} u_k(M_k) < 1$. Fix any agent i. Since $\sum_{k=1}^{n} u_i(M_k) = 1$, there exists $j \neq i$ such that $u_i(M_j) > u_j(M_j)$. Construct a directed graph with vertices $1, 2, \ldots, n$, and add an edge from i to j if $u_i(M_j) > u_j(M_j)$. Since every vertex has at least one outgoing edge, the graph consists of a directed cycle. For every edge $i \to j$ in the cycle, we give M_j to agent i instead of agent j. If we consider the change in the multiset of the n utilities between the old and new allocations, at least one number increases while others remain the same. This means that either we have decreased the number of agents who get zero utility, or keep this number fixed and increase the product of utilities of the agents who get nonzero utility. Either case contradicts the definition of an MNW allocation.

To show the bound for leximin, we apply the same argument. An improvement in the multiset of utilities as described in the last step contradicts the definition of leximin.

Finally, the tightness of the bounds follows from the instance where every agent has utility 1 for the same good. □

Lemma 5.1 allows us to show that the price of MNW and the strong price of MNW are both of linear order. Similar techniques can be used for the price of MEW and both prices of leximin, as we establish in the two subsequent theorems.

Theorem 5.2. *The price of MNW and the strong price of MNW are* $\Theta(n)$.

Proof. It suffices to show that the price of MNW is $\Omega(n)$ and the strong price of MNW is $O(n)$.

For the lower bound, let $m = n$ and $0 < \epsilon < 1$, and assume that the utilities are as follows:

- $u_1(1) = 1$ and $u_1(j) = 0$ otherwise.
- For $i = 2, \ldots, n$: $u_i(i-1) = 1 - \epsilon$, $u_i(i) = \epsilon$, and $u_i(j) = 0$ otherwise.

Consider the allocation that assigns good $i-1$ to agent i for $i = 2, \ldots, n$, and good n to agent 1. The social welfare of this allocation is $(n-1)(1-\epsilon)$. On the other hand, the unique MNW allocation assigns good i to agent i for every i. The social welfare of this allocation is $1 + (n-1)\epsilon$. Taking $\epsilon \to 0$, we find that the price of MNW is $\Omega(n)$.

To show a matching upper bound, consider an arbitrary instance. Since every agent receives utility at most 1, the optimal social welfare is at most n. On the other hand, by Lemma 5.1, the social welfare of any MNW allocation is at least 1. The conclusion follows. □

Theorem 5.3. *The price of MEW is* $\Theta(n)$.

Theorem 5.4. *The price of leximin and the strong price of leximin are* $\Theta(n)$.

Surprisingly, MEW allocations can be arbitrarily bad when there are at least three agents.

Theorem 5.5. *For* $n > 2$, *the strong price of MEW is infinite.*

We now turn to the case of two agents. For MNW, we establish almost tight bounds on both prices of fairness.

Theorem 5.6. *For* $n = 2$, *the price of MNW and the strong price of MNW are at least* $27/23 \approx 1.174$ *and at most* $5/4 = 1.25$.

Finally, we derive the exact bound for MEW and leximin with two agents. Note that since all leximin allocations are MEW, Theorem 5.7 immediately implies Theorem 5.8.

Theorem 5.7. *For* $n = 2$, *the price of MEW and the strong price of MEW are* $3/2$.

Theorem 5.8. *For* $n = 2$, *the price of leximin and the strong price of leximin are* $3/2$.

6 Pareto Optimality

In this section, we consider Pareto optimality. Since any allocation that maximizes social welfare is necessarily Pareto optimal, the price of Pareto optimality is trivially 1. By establishing a tight lower bound on the welfare of a Pareto optimal allocation, we show that the strong price of Pareto optimality is quadratic. Our result indicates that while Pareto optimality is sometimes referred to as 'efficiency', it does not necessarily fare well if efficiency is measured in terms of social welfare.

Lemma 6.1. *For any instance, every Pareto optimal allocation has social welfare at least* $1/n$, *and this bound is tight.*

Theorem 6.2. *The strong price of Pareto optimality is* $\Theta(n^2)$.

We also show an exact bound for the case of two agents.

Theorem 6.3. *For* $n = 2$, *the strong price of Pareto optimality is* 3.

7 Discussion

In this paper, we study the price of fairness for indivisible goods using several fairness notions that can always be satisfied. For most cases, we exhibit tight or asymptotically tight bounds on the worst-case efficiency loss that can occur due to fairness constraints. Interestingly, both the round-robin and MNW allocations, which are EF1, can have social welfare a linear factor away from the optimum, but not worse. In future research, it would be useful to close the gaps that remain after this work, the most intriguing of which is perhaps the EF1 gap between $\Omega(\sqrt{n})$ and $O(n)$. As we mentioned, settling this question would also have consequences on the price of envy-freeness gap in the divisible goods setting left open by Caragiannis *et al.* [2012].

Another direction for future work is to study the price of fairness for the *chore division* problem, where chores refer to items that yield negative utility for the agents. Indeed, almost all of the notions that we consider in the goods setting have direct analogs in the chore setting, and it would be interesting to see whether the corresponding bounds in the two settings turn out to be similar as well.

Acknowledgments

This work is partially supported by NSF Award CCF-1815434, and by the European Research Council (ERC) under grant number 639945 (ACCORD). We thank the anonymous reviewers for helpful comments.

References

[Aumann and Dombb, 2015] Yonatan Aumann and Yair Dombb. The efficiency of fair division with connected pieces. *ACM Transactions on Economics and Computation*, 3(4):23, 2015.

[Bansal and Sviridenko, 2006] Nikhil Bansal and Maxim Sviridenko. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 31–40, 2006.

[Bei *et al.*, 2019] Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, and Warut Suksompong. The price of fairness for indivisible goods. *CoRR*, abs/1905.04910, 2019.

[Bertsimas *et al.*, 2011] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.

[Bezáková and Dani, 2005] Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005.

[Bilò *et al.*, 2016] Vittorio Bilò, Angelo Fanelli, Michele Flammini, Gianpiero Monaco, and Luca Moscardelli. The price of envy-freeness in machine scheduling. *Theoretical Computer Science*, 613:65–78, 2016.

[Biswas and Barman, 2018] Arpita Biswas and Siddharth Barman. Fair division under cardinality constraints. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 91–97, 2018.

[Bogomolnaia and Moulin, 2004] Anna Bogomolnaia and Hervé Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

[Caragiannis *et al.*, 2012] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.

[Caragiannis *et al.*, 2016] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

[Dickerson *et al.*, 2014] John P. Dickerson, Jonathan Goldman, Jeremy Karp, Ariel D. Procaccia, and Tuomas Sandholm. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1405–1411, 2014.

[Dubins and Spanier, 1961] Lester E. Dubins and Edwin H. Spanier. How to cut a cake fairly. *The American Mathematical Monthly*, 68(1):1–17, 1961.

[Heydrich and van Stee, 2015] Sandy Heydrich and Rob van Stee. Dividing connected chores fairly. *Theoretical Computer Science*, 593:51–61, 2015.

[Kurokawa *et al.*, 2015] David Kurokawa, Ariel D. Procaccia, and Nisarg Shah. Leximin allocations in the real world. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 345–362, 2015.

[Kurz, 2014] Sascha Kurz. The price of fairness for a small number of indivisible items. In *Operations Research Proceedings*, pages 335–340, 2014.

[Lipton *et al.*, 2004] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC)*, pages 125–131, 2004.

[Manurangsi and Suksompong, 2019] Pasin Manurangsi and Warut Suksompong. When do envy-free allocations exist? In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 2019. Forthcoming.

[Plaut and Roughgarden, 2018] Benjamin Plaut and Tim Roughgarden. Almost envy-freeness with general valuations. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2584–2603, 2018.

[Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1):101–104, 1948.

[Suksompong, 2019] Warut Suksompong. Fairly allocating contiguous blocks of indivisible items. *Discrete Applied Mathematics*, 260:227–236, 2019.