ECAI 2023
K. Gal et al. (Eds.)

doi:10.3233/FAIA230522

On the Price of Fairness in the Connected Discrete Cake Cutting Problem

Ankang Sun[a] and Bo Li[a]

[a]Department of Computing, The Hong Kong Polytechnic University, Hong Kong, China
ankang.sun@polyu.edu.hk, comp-bo.li@polyu.edu.hk

Abstract. Discrete cake cutting is a fundamental model in fair resource allocation where the indivisible resources are located on a path. It is well motivated that, in reality, each agent is interested in receiving a contiguous block of items. An important question therein is to understand the economic efficiency loss by restricting the allocations to be fair, which is quantified as price of fairness (PoF). Informally, PoF is the worst-case ratio between the unconstrained optimal welfare and the optimal welfare achieved by fair allocations. Suksompong [Discret. Appl. Math., 2019] has studied this problem, where fairness is measured by the ideal criteria such as proportionality (PROP). A PROP allocation, however, may not exist in discrete cake cutting settings. Therefore, in this work, we revisit this problem and focus on the relaxed notions whose existence is guaranteed. We study both utilitarian and egalitarian welfare, and our results show significant differences between the PoF of guaranteed fairness notions and that of the ideal notions.

1 Introduction

The study of cake-cutting problem originated in the fields of economics and mathematics [16, 23, 24, 35], and has received increasing more attention from the communities of multi-agent systems and artificial intelligence over the past two decades [7, 20]. The cake is often used metaphorically to refer to heterogeneous resources, and its practical applications often impose constraints on the number of pieces each agent can receive. For example, when the resource has a temporal or spatial structure, every agent desires a connected block of the resource, such as allocating conference sessions to organizers, time slots in a fitness room to members, offices to research groups, or roads to districts. Moreover, real-life applications of cake-cutting are often related to indivisible resources, each of which can only be allocated to one agent, as shown by the aforementioned examples. The discrete cake-cutting problem with connectivity constraints can be alternatively regarded as a model of indivisible items where items are aligned on a path and the allocation of each agent should form a contiguous block. This framework has been considered in various recent papers [13, 28, 30, 31, 36].

Apart from fairness, the social welfare, a measurement of allocation efficiency, is also a major consideration in resources division, with the objective of maximizing the utilization of resources. Fair solutions can be very inefficient, and conversely, socially-optimal allocations may disregard fairness considerations. These two extremes naturally raise the questions of what is the trade-off (if any) between these two social concepts? How much efficiency needs to be sacrificed in order to ensure fairness? To investigate the fairness and efficiency trade-off, Bertsimas et al. [12] and Caragiannis et al. [17] respectively proposed *price of fairness* (PoF) to quantify the efficiency loss under fairness constraints. Thereafter, a line of research established the bounds of PoF for the indivisible items allocation without connectivity constraint [9, 11, 38] and for the divisible cake-cutting with connectivity constraint [5, 27].

In the setting of discrete cake cutting with connectivity constraint, Suksompong [36] considers the three ideal fairness criteria: *envy-freeness*, *proportionality* and *equitability*, which shed light on the fairness and efficiency trade-off by establishing a sequence of results regarding the PoF ratio. However, one limitation of this study is that the ideal fairness notions are not always satisfiable when the items are not divisible[1]. To circumvent the non-existence issue, the instances that do not admit these fair allocations are excluded from consideration. Back to the setting without connectivity constraints, Bei et al. [11] and Barman et al. [9] observed that neglecting the instances where ideally fair allocations do not exist can make the established PoF ratios fail to capture the true picture of the fairness and efficiency trade-off, and thus pointed out the importance of studying guaranteed fairness. Following the above line of research, in this work, we revisit the problem of discrete cake cutting with connectivity constraint and focus on the relaxations of the ideal fairness notions.

To highlight the differences between the ideal fairness and its relaxations and to avoid cumbersome presentation, we restrict our attention to proportional fairness and leave the study for the envy-freeness and equitability for future research. An allocation is proportional (PROP) if every agent's value is at least $\frac{1}{n}$, where n is the number of agents and agents' valuations are normalized to 1. Two widely accepted relaxations of proportionality are *proportionality up to one item* (PROP1) [8, 21] and *maximin share fairness* (MMS) [4, 29], which are known to be satisfiable in the discrete cake-cutting problem with connectivity constraints [6] [14]. Informally, PROP1 requires the condition of PROP to be satisfied after virtually adding an item from/to the agent's bundle. MMS fairness ensures that each agent's value is at least her max-min utility if she is to allocate the items. As for system efficiency, we care about both utilitarian and egalitarian welfare, where the former is the sum of all agents' utilities, and the latter of an allocation is the utility of the worst-off agent. Our main results are summarized in the following subsection.

[1] None of them can be satisfied in the example of assigning one valuable item to two agents.

1.1 Main results

The main contribution of this work is to bound the price of fairness for the discrete cake-cutting problem with connectivity constraint, where fairness is measured by either MMS or PROP1 and welfare is either utilitarian or egalitarian. For the utilitarian welfare, the PoF with respect to MMS is at least $\frac{\sqrt{n}}{2}$ and at most $4n$, and for PROP1, the PoF ratio is at least $\frac{\sqrt{n}}{2}$ and at most $2n$. For the egalitarian welfare, the PoF regarding MMS is at least $\frac{n}{2}$ and at most $2n$, and the PoF regarding PROP1 is infinite. We remark that all the upper bounds are proven by designing polynomial time algorithms that return fair allocations with the desired efficiency guarantees. The main results are summarized in Table 1, in which for ease of comparison, known PoF results regarding PROP are also provided.

As we can see from Table 1, for egalitarian welfare, when the fairness notions are changed from PROP to MMS and PROP1, the PoF changes significantly. This result seems to be contradictory at first glance since any PROP allocation must be MMS/PROP1. As we have discussed, this counter-intuitive result is due to the fact that the instance in which no PROP allocation exists is neglected when studying PROP, but the welfare maximizing solution of this instance can be unfair. As shown by Suksompong [36], achieving proportionality (if such an allocation exists) does not cause any loss of egalitarian welfare since every agent has utility at least $\frac{1}{n}$ in a PROP allocation and an allocation with higher egalitarian welfare must also be PROP. However, when the instance does not admit a PROP allocation and the fairness notions are changed to MMS or PROP1, a significant amount of egalitarian welfare has to be sacrificed. This result shows that the ignored instances are actually critical for quantifying welfare loss under fairness constraints and thus exhibits the importance of studying the fairness notions with guaranteed existence.

Although the bounds of PoF regarding utilitarian welfare do not show significant difference, we have extra challenges. Actually, given any PROP allocation (if any), it is straightforward that if the valuations are normalized, every agent has a value at least $\frac{n}{2}$, which is sufficient to achieve the tight bound of $\Theta(n)$ since the optimal utilitarian welfare is at most n. However, this argument does not carry over to the notions of MMS and PROP1. To find the MMS and PROP1 allocations with the best possible efficiency guarantee, we propose two parametric subroutines to balance the individual agent's value and the economic welfare. The first one is a matching procedure, where some items are very valuable for some agents so it suffices for us to allocate each such agent one of her valuable items. We need to be careful at this step since unallocated items on the left and right of the matched item cannot be allocated together to a single agent later due to the connectivity constraint. The second subroutine is motivated by the moving knife algorithm, where we stand at the very left and collect items until the first time that there exist agents who are satisfied regarding the designed parameters. By carefully integrating the two subroutines into algorithm design and properly choosing the parameters, we obtain the upper bound results as shown in Table 1. The lower bounds are proven by identifying hard instances where the welfare is inevitably sacrificed by enforcing the allocations to be fair.

Besides studying the general case with arbitrary number of agents, we are also interested in the two-agent model – a simple but important special case that has been widely investigated in the literature [1, 2]. We provide tight analyses for all settings, and the results coincide with those regarding PROP proven in [36]. Specifically, we show that the price of fairness is $\frac{3}{2}$ for utilitarian welfare and 1 for egalitarian welfare, no matter what fairness notions we have. Interestingly, we show that there exists a single allocation that is simultaneously MMS and PROP1 and achieves the corresponding optimal ratios.

1.2 Other Related Works

The traditional research on fair division is centered around allocating a divisible resource, denoted by the real interval $[0, 1]$, among a set of heterogeneous agents, i.e., cake cutting problem. It is well known that an envy-free (and thus proportional) allocation always exists [16] and can be found in finite steps [7]. Su [24] considered the case where it is required that every agent receives a contiguous piece, and Aumann and Dombb [5] analyzed the resulting price of fairness. However, when the items become indivisible, it is a different story. On one hand, people want to understand how to achieve approximate envy-freeness and proportionality for both unconstrained settings and the settings when there are restrictions on the allocations [32, 33, 34]. On the other hand, a line of research focuses on investigating how to achieve maximal efficiency while ensuring fairness [11, 17]. Besides the price of fairness, the compatibility between Pareto optimality and fairness is also widely studied [3, 10, 18]. In a recent couple of years, motivated by real-world applications, constraints have been investigated in the company with fair division; we refer the readers to the survey by Suksompong [37]. One of the most frequently studied constraints is *connectivity*, where the items are assumed to be distributed on a graph [13, 14, 19] and each agent should receive a connected subgraph. Line structure as we considered in this work is an important special case which admits some interesting positive results. Finally, besides the study of goods, Höhne and van Stee [28] considered the same problem for chores, i.e., undesirable cake.

2 Preliminaries

Denote by $[k] = \{1, \ldots, k\}$ for any positive integer k. A fair division instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$ is composed of n agents $N = \{1, \ldots, n\}$ and m indivisible items $E = \{e_1, \ldots, e_m\}$. The items are placed on a path in the order $e_1, \ldots, e_m$ from the left to the right. For simplicity, denote by $L(k) = L(e_k) = \{e_1, \ldots, e_k\}$ the items on the left of e_k including e_k and by $R(k) = R(e_k) = \{e_k, \ldots, e_m\}$ the items on the right of e_k including e_k. A feasible allocation assigns each agent a contiguous bundle of items. Let $\mathcal{S}$ be the set of all contiguous bundles. Each agent i is associated with a valuation function $v_i : \mathcal{S} \to \mathbb{R}^+ \cup \{0\}$ and $\mathcal{V} = \{v_i\}_{i\in N}$. For simplicity, we use $v_i(e_j)$ to represent $v_i(\{e_j\})$. The valuations are additive, i.e., $v_i(S) = \sum_{e\in S} v_i(e)$ for any $S \in \mathcal{S}$. It is assumed that the valuations are normalized, i.e., for all $i \in N$, $v_i(\emptyset) = 0$, and $v_i(E) = 1$.

A feasible allocation $\mathbf{A} = (A_1, \ldots, A_n)$ is an n-partition of E where every bundle is contiguous, i.e., $A_i \cap A_j = \emptyset$ for any $i \neq j$, $\cup_{i\in N} A_i = E$ and $A_i \in \mathcal{S}$ for each $i \in N$. If not explicitly stated otherwise, all the allocations in this paper are assumed to be contiguous. For any bundle S and any positive integer k, denote by $\Pi_k(S)$ the set of all k-contiguous partitions of S, and by $|S|$ the number of items in S. Given an allocation $\mathbf{A}$, the *utilitarian welfare* (UW) of $\mathbf{A}$ is $\mathsf{UW}(\mathbf{A}) = \sum_{i\in N} v_i(A_i)$, and the *egalitarian welfare* (EW) of $\mathbf{A}$ is $\mathsf{EW}(\mathbf{A}) = \min_{i\in N} v_i(A_i)$.

2.1 Fairness Notions

We next introduce the solution concepts. Note that the original definitions do not have any constraints, but to adapt to our setting, all allocations are required to be contiguous.

An allocation $\mathbf{A} = (A_1, \ldots, A_n)$ of a given instance $\mathcal{I}$ is *proportional* (PROP) if for any $i \in N$, $v_i(A_i) \geq \frac{1}{n}$.

General n	PROP	MMS	PROP1	
PoF	$\Theta(n)$	LB: $\Omega(\sqrt{n})$ UB: $O(n)$ (Theorems 4 and 14)		Utilitarian
	1	$\Theta(n)$ (Theorem 10)	∞ (Theorem 17)	Egalitarian

Table 1. The price of fairness regarding PROP, MMS, PROP1. The ratios for PROP are proved in [36].

Algorithm 1 Matching$(\mathcal{I}, \boldsymbol{\alpha})$

Input: Instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$ and vector $\boldsymbol{\alpha} = (\alpha_i)_{i\in N}$.
Output: A partial allocation $\mathbf{A}'$ where each agent is allocated at most one item.

1: Construct a weighted bipartite graph $G = (N \cup E, N \times E)$ where agents are vertices on one side and items are vertices on the other side. For each $i \in N$ and $j \in E$, there is an edge (i, j) with weight $v_i(e_j)$ if $v_i(e_j) \geq \alpha_i$.
2: Compute a maximum *weighted* matching μ of G. For every agent $i \in N$, denote by $\mu(i)$ the item matched to i and set $A'_i \leftarrow \{\mu(i)\}$; If i is unmatched, $A'_i \leftarrow \emptyset$.
3: **return** $\mathbf{A}' = (A'_1, \ldots, A'_n)$.

Definition 1 (PROP1) *A contiguous allocation* $\mathbf{A} = (A_1, \ldots, A_n)$ *is* proportional up to one item *(PROP1) if for any* $i \in N$*, there exists* $e \in E \setminus A_i$ *such that* $A_i \cup \{e\} \in \mathcal{S}$ *and* $v_i(A_i \cup \{e\}) \geq \frac{1}{n}$.

An alternative relaxation of PROP is maximin share (MMS) fairness. Given an instance $\mathcal{I}$, the *maximin share* (MMS) of agent $i \in N$ is the maximum value she can guarantee if she partitions E into n contiguous bundles but receives the smallest one. Formally,

$$\mathsf{MMS}_i(E, n) = \max_{\mathbf{X}\in\Pi_n(E)} \min_{j\in N} v_i(X_j).$$

If the instance $\mathcal{I}$ is clear from the context, we write $\mathsf{MMS}_i(\mathcal{I})$ or MMS_i for simplicity. Moreover, it is not hard to verify that for any instance $\mathcal{I}$, $\mathsf{MMS}_i(\mathcal{I}) \leq \frac{1}{n}$. For any agent i and a contiguous n-partition $\mathbf{A} = (A_1, \ldots, A_n)$, if $v_i(A_j) \geq \mathsf{MMS}_i$ for all j, $\mathbf{A}$ is called an MMS_i-defining partition. Note that although the computation of MMS values without connectivity constraints is NP-hard [39], in our setting, it can be computed efficiently [14].

Definition 2 (MMS) *An allocation* $\mathbf{A} = (A_1, \ldots, A_n)$ *is* maximin share *(MMS) fair if* $v_i(A_i) \geq \mathsf{MMS}_i(\mathcal{I})$ *for any* $i \in N$.

2.2 Price of Fairness

The *price of fairness* (PoF) is the supremum ratio over all instances between the maximum welfare of all allocations and maximum welfare of all fair allocations. Formally, given an instance $\mathcal{I}$ and a welfare function $W \in \{\mathsf{EW}, \mathsf{UW}\}$, denote by $\mathsf{OPT}_W(\mathcal{I})$ the maximum welfare with respect to W among all allocations of $\mathcal{I}$. For simplicity, if the instance is clear from the context, we write OPT_E and OPT_U to refer to $\mathsf{OPT}_{EW}(\mathcal{I})$ and $\mathsf{OPT}_{UW}(\mathcal{I})$, respectively. Letting $F \in \{\text{PROP1}, \text{MMS}\}$ be a fairness criterion, denote by $F(\mathcal{I})$ the set of all allocations satisfying F. Since we have different welfare and fairness notions, we sometimes use $W \mid F$ to denote the specific setting, where W is either egalitarian or utilitarian welfare, and F is either MMS or PROP1.

Definition 3 (PoF) *The price of fairness with respect to fairness criterion F and welfare function W is*

$$\mathsf{PoF}(W \mid F) = \sup_{\mathcal{I}} \min_{\mathbf{A}\in F(\mathcal{I})} \frac{\mathsf{OPT}_W(\mathcal{I})}{W(\mathbf{A})}.$$

If the setting $W \mid F$ is clear from the context, we simply write PoF.

We remark that if no fair allocations can achieve non-zero welfare, the PoF is infinity. The PoF with respect to fairness criterion F is also called *price of F*, i.e., price of MMS or price of PROP1.

3 Price of MMS for Indivisible Goods

We start with MMS fairness in this section. We first introduce two subroutines that will be used to design fair allocation algorithms with high welfare in Section 3.1, and then prove our main results in Sections 3.2 and 3.3.

3.1 Useful Subroutines in the Algorithms

We first design the subroutines, Matching and MovingKnife, as shown in Algorithms 1 and 2. Intuitively, given an instance $\mathcal{I}$, Matching$(\mathcal{I}, \boldsymbol{\alpha})$ uses a threshold vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n)$ to identify a set of valuable items for each agent, and then uses a maximum matching to assign large items to the agents so that each agent gets at most one item. Particularly, if $\alpha_i \geq \mathsf{MMS}_i$, then the agents who get allocated by Matching are happy with MMS fairness. MovingKnife$(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$ is motivated by the moving-knife algorithm [22], where we stand at the left-end of all unallocated items and find the closest item such that there exists an agent who is happy (regarding the parameters α_i's) with the contiguous block between left-end and this item. As we will see, MovingKnife can ensure MMS fairness by setting $\alpha_i = \mathsf{MMS}_i(\mathcal{I})$, but may produce low utilitarian and egalitarian welfare. We can adjust α_i's to increase the welfare guarantee, however, we need to be careful since if α_i's are too large, we may not find a contiguous allocation to satisfy all agents. In the following sections, we show how to properly choose the parameters and combine Matching and MovingKnife so that the returned contiguous allocations meet the fairness criteria and satisfy the desired welfare guarantee, simultaneously.

Algorithm 2 MovingKnife$(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$

Input: Instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$, partial allocation $\mathbf{A}'$ of $Q \subseteq E$ on $P \subseteq N$ with $|P| = |Q| < n$, and vector $\boldsymbol{\alpha} = (\alpha_i)_{i\in N\setminus P}$.
Output: A complete allocation $\mathbf{A}$.

1: Initialize $\bar{N} \leftarrow N \setminus P$, $\bar{E} \leftarrow E \setminus Q$, $p_0 \leftarrow 0$;
2: **while** $|\bar{N}| \geq 1$ & $\bar{E} \neq \emptyset$ **do**
3: Find the smallest index $p > p_0$ such that there is an agent $i \in \bar{N}$ and a contiguous bundle $S \subseteq L(p) \cap \bar{E}$ with $v_i(S) \geq \alpha_i$. If there are multiple agents, choose the one with the largest value on S. If no such p is found, break and return "**Error**".
4: Set $A_i \leftarrow S$ and update $\bar{N} \leftarrow \bar{N} \setminus \{i\}$, $\bar{E} \leftarrow \bar{E} \setminus A_i$, $p_0 \leftarrow p$.
5: **end while**
6: **return** $\mathbf{A} = (A_1, \ldots, A_n)$.

3.2 Price of MMS with Utilitarian Welfare

We first design an algorithm (as shown in Algorithm 3) to compute an MMS fair allocation that ensures constant utilitarian welfare, which implies the $O(n)$ upper bound since the valuations are normalized so that the utilitarian welfare of any allocation is no more than n. Our first failed attempt is to run $\mathsf{MovingKnife}$ by setting $\alpha_i = \max\{\mathsf{MMS}_i, \frac{1}{4n}\}$ for all agents i and $\mathbf{A}' = \{\emptyset\}$, so that every agent is satisfied regarding MMS fairness and has value at least $\frac{1}{4n}$ (and hence $O(n)$ approximation). However, by setting $\alpha_i = \max\{\mathsf{MMS}_i, \frac{1}{4n}\}$, $\mathsf{MovingKnife}$ itself may fail to return a feasible allocation. For example, consider an instance with $N = \{1, 2\}$, $E = \{e_1, e_2\}$, and both agents have value $\epsilon \ll \frac{1}{8}$ for e_1 and $1-\epsilon$ for e_2. Since $\mathsf{MMS}_i = \epsilon$ and $\alpha_i = \frac{1}{8}$, $\mathsf{MovingKnife}$ allocates both e_1 and e_2 to one of the agents and the other agent gets nothing. The failure of this attempt is because there is one item, i.e., e_2 in the example, having very value for the agents so that if such an item together with some items on its left are assigned to one agent, the remaining items may not be sufficient to ensure MMS values for the rest agents.

Algorithm 3 Utilitarian | MMS

Input: Instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$.
Output: A contiguous allocation $\mathbf{A}$.

1: Run $\mathsf{Matching}(\mathcal{I}, \boldsymbol{\alpha})$ with $\alpha_i = \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{4n}\}$ and obtain a partial allocation $\mathbf{A}'$ on agents $N_0 \subseteq N$ and items $E_0 \subseteq E$. Let $N' = N \setminus N_0$, $E' = E \setminus E_0$.
2: **if** $\mathsf{UW}(\mathbf{A}') \geq \frac{1}{4}$ **then**
3: Run $\mathsf{MovingKnife}(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$ with $\alpha_i = \mathsf{MMS}_i(\mathcal{I})$ for $i \in N'$ and obtain allocation $\mathbf{A}$.
4: **else**
5: Letting $\alpha_i = \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{4n}\}$ for $i \in N'$, run $\mathsf{MovingKnife}(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$ and obtain allocation $\mathbf{A}$.
6: **end if**
7: If there are still items remaining, arbitrarily assign them to the agent whose bundle forms a contiguous block with them. (The remaining items must be a set of contiguous blocks.)
8: **return** $\mathbf{A}$

Therefore, in Algorithm 3, we first use $\mathsf{Matching}$ to allocate large items by setting each $\alpha_i = \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{4n}\}$. Note that if the partial utilitarian welfare of the agents who get allocated by $\mathsf{Matching}$ is already at least $\frac{1}{4}$, then it suffices to ensure solely MMS values for the remaining agents and $\mathsf{MovingKnife}$ with $\alpha_i = \mathsf{MMS}_i$ is able to find such an allocation. If the welfare from $\mathsf{Matching}$ is smaller than $\frac{1}{4}$, to ensure constant total utilitarian welfare, we turn to feed $\alpha_i = \max\{\mathsf{MMS}_i, \frac{1}{4n}\}$ to $\mathsf{MovingKnife}$ so that every agent's value is at least $\frac{1}{4n}$, which implies constant welfare as well. Fortunately, given the welfare from $\mathsf{MovingKnife}$ being smaller than $\frac{1}{4}$, we manage to prove that the previous failure will not happen and are able to allocate each remaining agent a contiguous block with value no smaller than $\max\{\mathsf{MMS}_i, \frac{1}{4n}\}$.

In summary, we have the following theorem.

Theorem 4 *For Utilitarian welfare and MMS fairness, the price of fairness is at least $\Omega(\sqrt{n})$ and at most $O(n)$.*

The upper bound in Theorem 4 relies on the following Lemmas 5 and 7 which show that Algorithm 3 can always return a feasible allocation with utilitarian welfare at least $\frac{1}{4}$.

Lemma 5 *For any instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$, if the allocation $\mathbf{A}$ returned by Algorithm 3 is constructed in Step 3, then $\mathbf{A}$ is contiguous and $v_i(A_i) \geq \mathsf{MMS}_i$ for agent $i \in N$.*

Proof. By the design of the algorithm, A_i is contiguous for all i. For the MMS fairness, we first present Claim 6 which states that if we remove an (arbitrary) agent and a contiguous block of items whose value is small for every remaining agent, in the reduced instance the MMS value of each remaining agent does not decrease.

Claim 6 *Let $p_i \leq m$ be the smallest index such that $v_i(L(p_i)) \geq \mathsf{MMS}_i(\mathcal{I})$, then $\mathsf{MMS}_i(E \setminus X, n-1) \geq \mathsf{MMS}_i(\mathcal{I})$ for $X \subseteq L(p_i)$.*

Proof of Claim 6. It suffices to show $\mathsf{MMS}_i(E \setminus L(p_i), n-1) \geq \mathsf{MMS}_i(\mathcal{I})$. Consider an arbitrary $\mathsf{MMS}_i(\mathcal{I})$-partition $(X_1, \ldots, X_n)$ for agent i, where the bundles are ordered from left to right. Then $v_i(X_j) \geq \mathsf{MMS}_i(\mathcal{I})$ holds for any bundle X_j. The condition of p_i being the smallest index such that $v_i(L(p_i)) \geq \mathsf{MMS}_i(\mathcal{I})$ implies that $L(p_i) \subseteq X_1$ and hence $E \setminus L(p_i) \supseteq \bigcup_{j=2}^n X_j$. Note that $(X_2, \ldots, X_n)$ is one $(n-1)$-partition of $\bigcup_{j=2}^n X_j$, and thus

$$\mathsf{MMS}_i(E \setminus L(p), n-1) \geq \mathsf{MMS}_i(\bigcup_{j=2}^n X_j, n-1) \geq \mathsf{MMS}_i(\mathcal{I}),$$

which completes the proof of the claim. ■

Since the assignment in Step 7 does not reduce agents' values, we can without loss of generality assume that all items are assigned before Step 7. In Step 1, for every matched agent $i \in N_0$, we have $v_i(A_i) \geq \mathsf{MMS}_i$. We then focus on $i \in N'$. Suppose that $A_1, \ldots, A_n$ are ordered from left to right, and let agent k be the one receiving the right-most block assigned by Step 3. Then, it suffices to show $v_k(A_k) \geq \mathsf{MMS}_k(\mathcal{I})$. For every $j \in [k-1]$, denote by $\mathcal{I}^j$ the reduced instance right after removing agent j and bundle A_j, and let $\mathcal{I}^0 = \mathcal{I}$. Then, for $j \in [k-1]$, we have either A_j acts as a bundle of $\mathsf{MMS}_k(\mathcal{I}^{j-1})$-defining partition (when A_j is found by Step 1) or $v_k(A_j) \leq \mathsf{MMS}_k(\mathcal{I}^{j-1})$ (when A_j is found by Step 3). Then according to Claim 6, $\mathsf{MMS}_k(\mathcal{I}^j) \geq \mathsf{MMS}_k(\mathcal{I}^{j-1})$ holds. By induction, we have $\mathsf{MMS}_k(\mathcal{I}^{k-1}) \geq \mathsf{MMS}_k(\mathcal{I}^0) = \mathsf{MMS}_k(\mathcal{I})$. In the reduced instance $\mathcal{I}^{k-1}$, there exist $n-k+1$ agents and a set of items composed by bundle A_k and a number $n-k$ of individual items (each of them is at the right of A_k). Note that in $\mathcal{I}^{k-1}$ if an item e^* has value no less than $\mathsf{MMS}_k(\mathcal{I}^{k-1})$, then in the $\mathsf{MMS}_k(\mathcal{I}^{k-1})$-defining partition, e^* solely forms a contiguous bundle. As a consequence, it must hold that $v_k(A_k) \geq \mathsf{MMS}_k(\mathcal{I}^{k-1})$, and therefore, $v_k(A_k) \geq \mathsf{MMS}_k(\mathcal{I})$. ■

Lemma 7 *For any instance $\mathcal{I} = \langle N, E, \mathcal{V}\rangle$, if the allocation $\mathbf{A}$ returned by Algorithm 3 is constructed in Step 5, then $\mathbf{A}$ is contiguous and $v_i(A_i) \geq \max\{\mathsf{MMS}_i, \frac{1}{4n}\}$ for all $i \in N$.*

Proof. According to $\mathsf{Matching}(\mathcal{I}, \boldsymbol{\alpha})$ in Step 1 of Algorithm 3, it holds that $v_i(A_i) \geq \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{4n}\}$ for each $i \in N_0$. Then, we show that it is also the case for every $i \in N \setminus N_0$, which is equivalent to prove that each agent $i \in N \setminus N_0$ is able to receive a bundle from $\mathsf{MovingKnife}(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$ in Step 5 of Algorithm 3.

We now fix $i \in N \setminus N_0$. If $\mathsf{MMS}_i \geq \frac{1}{4n}$, then by arguments similar to that in the proof of Lemma 5, one can verify that agent i must receive a bundle from Step 5 of Algorithm 3 with value $v_i(A_i) \geq \max\{\mathsf{MMS}_i, \frac{1}{4n}\}$. We can further focus on the case of $\mathsf{MMS}_i < \frac{1}{4n}$. For a contradiction, assume that agent i does not receive a bundle in Step 5. Since agent i is not matched in $\mathsf{Matching}(\mathcal{I}, \boldsymbol{\alpha})$, we can claim $v_i(e) < \frac{1}{4n}$ for every $e \in E'$. Denote by $\widetilde{N}$ the set of agents receiving bundles in Step 5, and we have $|\widetilde{N}| \le n - |N_0| - 1$; otherwise, the statement already holds. For each $j \in \widetilde{N}$, let B_j be the bundle assigned to agent j in Step 5 and $e_{j_R} \in B_j$ be the right-most item of B_j. Due to $v_i(e_{j_R}) < \frac{1}{4n}$ and Subroutine $\mathsf{MovingKnife}$ in Step 5, we have $v_i(B_j) = v_i(B_j \setminus \{e_{j_R}\}) + v_i(e_{j_R}) < \frac{1}{2n}$. We then present an upper bound of agent i's value on assigned items. As in Step 1, $\mathsf{Matching}(\mathcal{I}, \boldsymbol{\alpha})$ computes the maximum weighted matching, we have $v_i(E_0) \le \mathsf{UW}(\mathbf{A}') < \frac{1}{4}$. Then, agent i values the assigned items

$$v_i(E_0) + v_i(\cup_{j \in \widetilde{N}} B_j) < \frac{1}{4} + \frac{n - |N_0| - 1}{2n}.$$

Accordingly, agent i has value at least $\frac{n+2|N_0|+2}{4n}$ on the unassigned items due to normalized valuations. Since there are at most $n-1$ agents receiving a bundle, unassigned items are then composed by at most n contiguous blocks. According to the pigeonhole principle, agent i values on one of the unassigned contiguous blocks at least $\frac{n+2|N_0|+2}{4n^2} > \frac{1}{4n}$, based on which agent i should receive a bundle in Step 5, a contradiction. ■

By Lemmas 5 and 7, the returned allocation $\mathbf{A}$ has a utilitarian welfare at least $\mathsf{UW}(\mathbf{A}) \ge \frac{1}{4}$. Since the valuations are normalized, the utilitarian welfare of any allocation is at most n, and thus, the price of MMS with respect to utilitarian welfare is at most $4n$. The lower bound in Theorem 4 comes from the following lemma.

Lemma 8 *There is an instance in which no contiguous MMS allocation has better than $\Theta(\frac{1}{\sqrt{n}})$ of the optimal utilitarian welfare.*

Proof. Consider an instance with n agents and a set $E = \{e_1, \dots, e_{2n}\}$ of goods. For $i = 1, \dots, \sqrt{n}$, the valuation function $v_i(\cdot)$ is: $v_i(e_j) = \frac{1}{2\sqrt{n}}$ for $2(i-1)\sqrt{n} + 1 \le j \le 2i\sqrt{n}$ and $v_i(e_j) = 0$ for other j. For $i \ge \sqrt{n} + 1$, the valuation functions is: $v_i(e_j) = \frac{1}{2n}$ for any $j \in [2n]$. One can compute $\mathsf{MMS}_i = 0$ for $i \le \sqrt{n}$ and $\mathsf{MMS}_i = \frac{1}{n}$ for $i \ge \sqrt{n}+1$. In a utilitarian welfare maximizing allocation $\mathbf{O}$, each agent $i \le \sqrt{n}$ receives all goods on which she has positive value and agent $i > \sqrt{n}$ receives none. We can compute $\mathsf{OPT}_U = \mathsf{UW}(\mathbf{O}) = \sqrt{n}$. But for agent $i \ge \sqrt{n} + 1$, since $v_i(O_i) = 0 < \mathsf{MMS}_i$, she violates MMS under allocation $\mathbf{O}$. To make such an agent satisfy MMS fairness, two adjacent goods must be assigned to her. Thus, in total $2(n - \sqrt{n})$ goods need to be assigned to the latter $n - \sqrt{n}$ agents, which makes that only $2\sqrt{n}$ goods can be assigned to the first $\sqrt{n}$ agents in an MMS allocation. Consequently, for an arbitrary MMS allocation $\mathbf{A}$, we have $\mathsf{UW}(\mathbf{A}) \le 2 - \frac{1}{\sqrt{n}}$ and therefore,

$$\mathsf{PoF} \ge \frac{\sqrt{n}}{2 - \frac{1}{\sqrt{n}}} = \Omega(\sqrt{n}),$$

which completes the proof of the lemma. ■

In addition, the proof of Lemma 5 indeed results in a general reduction property regarding contiguous MMS allocations, which may be of independent interest. Informally, if we assign k arbitrary items to k arbitrary agents, we can still satisfy the remaining $n-k$ agents with the remaining items.

Corollary 9 *For any instance $\mathcal{I}$, any $P \subseteq N$ and $Q \subseteq E$ with $|P| = |Q| = k \le n$, there exists $(X_i)_{i \in N \setminus P} \in \Pi_{n-k}(E \setminus Q)$ such that $v_i(X_i) \ge \mathsf{MMS}_i(\mathcal{I})$ and $X_i \in \mathcal{S}$ for all $i \in N \setminus P$.*

Note that in the setting without connectivity constraints, both Lemma 1 in Bouveret et al. [15] and Lemma 3.4 in Amanatidis et al. [4] state that one can remove a single good to an agent without decreasing the maximin share of other agents on the reduced instance. Corollary 9 generalizes this monotonicity property to the setting where the connectivity constraint is required.

3.3 Price of MMS with Egalitarian Welfare

Next, we discuss egalitarian welfare. For an instance $\mathcal{I}$, if $\mathsf{OPT}_E \ge \frac{1}{n}$, it means that the egalitarian welfare maximizing allocation guarantees MMS fairness for every agent since $\mathsf{MMS}_i(\mathcal{I}) \le \frac{1}{n}$ for all i. Then it suffices to consider the case when $\mathsf{OPT}_E < \frac{1}{n}$. We observe that Subroutine $\mathsf{MovingKnife}$ alone may result in a poor approximation. The reason is that the MMS value of an agent can be negligible in some instances, and when such an agent receives a bundle with value MMS, $\mathsf{MovingKnife}$ stops assigning more items to her. Therefore, in Algorithm 4 we balance fairness and efficiency by first allocating items of large values to agents whose MMS values are small via a matching procedure. Particularly, we set $\alpha_i = +\infty$ if $\mathsf{MMS}_i(\mathcal{I}) \ge \frac{1}{2n} \cdot \mathsf{OPT}_E$ and $\alpha_i = \frac{1}{2n} \cdot \mathsf{OPT}_E$ otherwise. We remark that, to improve the value of the worst-off agent, subroutine $\mathsf{Matching}'$ is implemented in Step 1 of Algorithm 4. Subroutine $\mathsf{Matching}'$ is similar to $\mathsf{Matching}$, with the only difference being that, instead of maximum weighted matching, $\mathsf{Matching}'$ applies maximum *cardinality* matching. Thereafter, Algorithm 4 uses $\mathsf{MovingKnife}$ with proper parameters to allocate the remaining items. Seemingly we are quite conservative on the selection of the parameter $\frac{1}{2n} \cdot \mathsf{OPT}_E$, but it turns out to induce the (asymptotically) best possible PoF ratio.

Algorithm 4 Egalitarian | MMS

Input: An instance $\mathcal{I} = \langle N, E, \mathcal{V} \rangle$ with $\mathsf{OPT}_E < \frac{1}{n}$.
Output: A contiguous allocation $\mathbf{A}$.

1: Run $\mathsf{Matching}'(\mathcal{I}, \boldsymbol{\alpha})$ with $\alpha_i = +\infty$ if $\mathsf{MMS}_i(\mathcal{I}) \ge \frac{1}{2n} \cdot \mathsf{OPT}_E$ and $\alpha_i = \frac{1}{2n} \cdot \mathsf{OPT}_E$ otherwise, and obtain a partial allocation $\mathbf{A}'$ on agents $N_0 \subseteq N$ and items $E_0 \subseteq E$. Let $N' = N \setminus N_0$, $E' = E \setminus E_0$.
2: Letting $\alpha_i = \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{2n}\mathsf{OPT}_E\}$ for $i \in N'$, run $\mathsf{MovingKnife}(\mathcal{I}, \mathbf{A}', \boldsymbol{\alpha})$ and obtain allocation $\mathbf{A}$.
3: If there are still items remaining, arbitrarily assign them to the agent whose bundle forms a contiguous block with them.
4: **return** $\mathbf{A}$

Theorem 10 *For Egalitarian welfare and MMS fairness, the price of fairness is $\Theta(n)$.*

To prove Theorem 10, we prove Lemmas 11 and 13.

Lemma 11 *For any instance $\mathcal{I} = \langle N, E, \mathcal{V} \rangle$, Algorithm 4 returns an allocation $\mathbf{A}$ with $v_i(A_i) \ge \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{2n} \cdot \mathsf{OPT}_E\}$ for all agents $i \in N$.*

Proof. We first prove the following claim.

Claim 12 *In Subroutine $\mathsf{Matching}'(\mathcal{I}, \boldsymbol{\alpha})$, if agent i's vertex has degree at least one in G, then i is matched by μ.*

Proof. Denote by $\bar{N}$ the set of agents with a degree at least one in G. Clearly, $\mathsf{MMS}_i < \frac{1}{2n}\mathsf{OPT}_E$ holds for every $i \in \bar{N}$. For the sake of contradiction, assume the matching μ is not $\bar{N}$-perfect. Then according to Hall's theorem, there exists a subset $N^* \subseteq \bar{N}$ satisfying $|N^*| > |D_G(N^*)|$ where $D_G(N^*)$ is the neighbourhood of N^* in

G. We then focus on the set $E \setminus D_G(N^*)$ and claim that no subset $P \subseteq E \setminus D_{\bar{G}}(N^*)$ is able to bring value $v_t(P) \geq \mathsf{OPT}_E$ for some agent $t \in N^*$. Suppose not, and assume that agent $j \in N^*$ has value $v_j(P^*) \geq \mathsf{OPT}_E$ where $P^* \subseteq E \setminus D_G(N^*)$ is contiguous. Due to the choice of $\boldsymbol{\alpha}$, we have $v_j(e) < \frac{1}{2n}\mathsf{OPT}_E$ for each $e \in P^*$. Thus, set P^* can be partitioned into n subsets $\{P_l^*\}_{l=1}^n$ such that $v_j(P_l^*) \geq \frac{1}{2n}\mathsf{OPT}_E$ for all l, which then leads to $\mathsf{MMS}_j \geq \frac{1}{2n}\mathsf{OPT}_E$, a contradiction. Thus, given an agent $i \in N^*$ and a bundle $S \in \mathcal{S}$, $v_i(S) \geq \mathsf{OPT}_E$ holds if and only if $S \cap D_G(N^*) \neq \emptyset$. Note that $|N^*| > |D_G(N^*)|$, then it is impossible to make every agent in N^* receive value at least OPT_E, a contradiction. Therefore, matching μ must be $\bar{N}$-perfect. ■

Claim 12 implies that if agent i with $\mathsf{MMS}_i < \frac{1}{2n}\mathsf{OPT}_E$ values a single item $v_i(e) \geq \frac{1}{2n}\mathsf{OPT}_E$, then she will be matched to a single item in Step 1 of Algorithm 4 and has value

$$v_i(A_i) \geq v_i(\mu(i)) \geq \frac{1}{2n}\mathsf{OPT}_E \geq \mathsf{MMS}_i,$$

where μ is the maximum cardinality matching computed in Subroutine $\mathsf{Matching}'(\mathcal{I}, \boldsymbol{\alpha})$. Thus, if $N_0 = N$, the statement is proven, and we can further assume $N_0 \subsetneq N$. Fix agent $i \in N \setminus N_0$ and split the proof into two cases.

Case 1: $\mathsf{MMS}_i \geq \frac{1}{2n}\mathsf{OPT}_E$. The proof of this case is similar to the proof of Lemma 5 and we omit it.

Case 2: $\mathsf{MMS}_i < \frac{1}{2n}\mathsf{OPT}_E$. It suffices to show that agent i receives a bundle in Step 2 of Algorithm 4. Based on Claim 12, agent i has value $v_i(e) < \frac{1}{2n}\mathsf{OPT}_E$ for each $e \in E$, and thus, $v_i(\mu(j)) < \frac{1}{2n}\mathsf{OPT}_E$ holds for $j \in N_0$ where μ is the matching computed in Subroutine $\mathsf{Matching}'$. For each $j \in N \setminus N_0$ and $j \neq i$, denote by B_j (if exists) the bundle received by agent j in Step 2 of Algorithm 4, and accordingly, $v_i(B_j) < \frac{1}{n}\mathsf{OPT}_E$ holds; otherwise, agent i will receive a bundle at the time when agent j is picked in Subroutine $\mathsf{MovingKnife}$ in Step 2 of Algorithm 4. Therefore, even after all other $n-1$ agents receiving bundles in either Step 1 or Step 2, agent i still has value no less than $1 - \mathsf{OPT}_E > 1 - \frac{1}{n}$ on the unassigned items that are composed by at most n contiguous subsets. By the pigeonhole principle, there exists a subset with value at least $\frac{n-1}{n^2} > \mathsf{MMS}_i$ for agent i. Hence, agent i receives a bundle B_i in Subroutine $\mathsf{MovingKnife}$ in Step 2 of Algorithm 4 and has value

$$v_i(A_i) \geq v_i(B_i) \geq \max\{\frac{1}{2n}\mathsf{OPT}_E, \mathsf{MMS}_i\}.$$

Therefore, Algorithm 4 outputs an allocation $\mathbf{A}$ such that $v_i(A_i) \geq \max\{\frac{1}{2n}\mathsf{OPT}_E, \mathsf{MMS}_i\}$ for each $i \in N$. ■

Note that Algorithm 4 uses the value of OPT_E which is NP-hard to compute [14].[2] To make it run in polynomial time, we adopt the following technique by guessing the value of OPT_E, denoted by o. If $o \leq \mathsf{OPT}_E$, by Lemma 11, we can find a feasible allocation where $v_i(A_i) \geq \max\{\mathsf{MMS}_i(\mathcal{I}), \frac{1}{2n} \cdot o\}$. Thus, we can start with a trivial upper bound of OPT_E by setting $o = 1$, and run Algorithm 4. If we do not obtain a feasible allocation, we decrease the value of o by a small constant $\epsilon > 0$ and repeat Algorithm 4. We stop until we obtain a feasible allocation regarding o and it is guaranteed that $o \geq \mathsf{OPT}_E - \epsilon$. To ensure $\mathsf{OPT}_E - \epsilon = \Theta(\frac{1}{n}) \cdot \mathsf{OPT}_E$ for all agents, the value of ϵ cannot be too large. One possible way is to set $\epsilon = \frac{1}{2} \cdot \min_{i \in N} \mathsf{MMS}_i$, since $\mathsf{OPT}_E \geq \min_{i \in N} \mathsf{MMS}_i$ and thus $\mathsf{OPT}_E - \epsilon \geq \frac{1}{2} \cdot \mathsf{OPT}_E$.

[2] It is proven in Bouveret et al. [14] that deciding whether an instance admits a PROP allocation (i.e., whether it admits an allocation with egalitarian welfare no smaller than $\frac{1}{n}$) is NP-hard, which is a special case of our problem.

Lemma 13 *There is an instance in which no contiguous MMS allocation achieves better than $\Theta(\frac{1}{n})$ of the optimal egalitarian welfare.*

4 Price of PROP1 for Indivisible Goods

4.1 Price of PROP1 with Utilitarian Welfare

Theorem 14 *For Utilitarian welfare and PROP1 fairness, the price of fairness is at least $\Omega(\sqrt{n})$ and at most $O(n)$.*

We first prove that there exists a PROP1 allocation with a utilitarian welfare at least $\frac{1}{2}$. Technically, such a desired PROP1 allocation is achieved by rounding a contiguous proportional allocation of the related divisible cake-cutting instance.

Lemma 15 *A* PROP1 *allocation with utilitarian welfare at least $\frac{1}{2}$ can be computed efficiently.*

Proof. Note that if an agent i has valuation $v_i(e_j) \geq 1/n$ on good e_j, then she satisfies PROP1 when she receives empty bundle. Accordingly, we can further focus on the instance $\mathcal{I}$ in which $v_i(e_j) < 1/n$ holds for any i, j. We then construct a corresponding cake-cutting instance $\mathcal{I}'$ with n agents and the cake being the interval $[0, m]$ where $m = |E|$. In $\mathcal{I}'$, the value of interval $[a, b]$ for each agent i is equal to $\int_a^b f_i(x)dx$ where f_i is agent i's density function. Each agent i has a piecewise constant density function $f_i(x) = v_i(e_j)$ on interval $[j-1, j]$, and thus, agent i has value $v_i(e_j)$ on piece $[j-1, j]$. According to Dubins and Spanier [22], instance $\mathcal{I}'$ admits a proportional allocation π, in which w.l.o.g, agents $1, \ldots, n$ receive the 1st, 2nd, ..., n-th piece of cake from left to right, and each agent i receives interval $[\pi_{i-1}, \pi_i]$. We then transfer π into allocations of $\mathcal{I}$ that satisfy PROP1 and have an absolute welfare guarantee.

Since each agent i has value at least $1/n$ on her piece $[\pi_{i-1}, \pi_i]$, we claim that $[\pi_{i-1}, \pi_i] \notin [j-1, j]$ for any pair of i, j due to the property of $\mathcal{I}$. As a result, for each $j \in [m]$, interval $[j-1, j]$ is either covered by an interval $[\pi_{p-1}, \pi_p]$ or not covered by a single interval but intersects with two contiguous pieces received by agents in π. We then construct two allocations of $\mathcal{I}$. In allocation $\mathbf{A}^L$ (resp. $\mathbf{A}^R$), each good e_j is assigned to agent i if $[j-1, j] \in [\pi_{i-1}, \pi_i]$, and e_j is assigned to agent p (resp. q) if $[j-1, j]$ intersects[3] with two contiguous pieces $[\pi_{p-1}, \pi_p]$, $[\pi_{q-1}, \pi_q]$ with $p < q$. In the following, we first show both $\mathbf{A}^L$ and $\mathbf{A}^R$ are PROP1 allocations and then prove that one of them has utilitarian welfare at least $1/2$.

The connectivity of $\mathbf{A}^L$ and $\mathbf{A}^R$ comes from the connectivity of π. We then prove the PROP1 of allocation $\mathbf{A}^L$ and fix an agent i who receives the piece of cake $[\pi_{i-1}, \pi_i]$. If $\pi_{i-1} \in \mathbb{N}^+$, then the bundle received by agent i is $A_i^L = \{e_j \mid \pi_{i-1} \leq j \leq \lceil \pi_i \rceil\}$. Accordingly, we have the following

$$v_i(A_i^L) = \int_{\pi_{i-1}}^{\lceil \pi_i \rceil} f_i(x)dx \geq \int_{\pi_{i-1}}^{\pi_i} f_i(x)dx \geq \frac{1}{n},$$

where the last inequality is due to the proportionality of π. If $\pi_{i-1} \notin \mathbb{N}^+$, agent i receives $A_i^L = \{e_j \mid \lceil \pi_{i-1} \rceil \leq j \leq \lceil \pi_i \rceil\}$. Note that item $A_i^L \cup \{e_{\lfloor \pi_{i-1} \rfloor}\} \in \mathcal{S}$ and we have

$$v_i(A_i^L \cup \{e_{\lfloor \pi_{i-1} \rfloor}\}) = \int_{\lfloor \pi_{i-1} \rfloor}^{\lceil \pi_i \rceil} f_i(x)dx \geq \int_{\pi_{i-1}}^{\pi_i} f_i(x)dx \geq \frac{1}{n},$$

[3] If the intersection of two intervals is a single point, then we regard their intersection as an empty set.

which then implies that agent i also satisfies PROP1. Therefore, we can conclude that $\mathbf{A}^L$ satisfies PROP1. By a similar argument, one can prove that $\mathbf{A}^R$ is also a PROP1 allocation.

As each item e_j of $\mathcal{I}$ corresponds to an interval $[j-1, j]$ of $\mathcal{I}'$, then any contiguous subsets $S \subseteq E$ also corresponds to a subinterval of $[0, m]$. For each agent i, denote by $[x^L_{i-1}, x^L_i]$ and $[x^R_{i-1}, x^R_i]$ the corresponding intervals of A^L_i and A^R_i, respectively. By the construction of $\mathbf{A}^L$ and $\mathbf{A}^R$, for each $i \in [n]$, we have $\min\{x^R_{i-1}, x^L_{i-1}\} \le \pi_{i-1}$ and $\max\{x^R_i, x^L_i\} \ge \pi_i$. Then, for each $i \in [n]$, we have the following inequality

$$\int_{x^L_{i-1}}^{x^L_i} f_i(x)dx + \int_{x^R_{i-1}}^{x^R_i} f_i(x)dx \ge \int_{\pi_{i-1}}^{\pi_i} f_i(x)dx,$$

which then implies

$$\sum_{t=1}^{n} \int_{x^L_{t-1}}^{x^L_t} f_t(x)dx + \sum_{t=1}^{n} \int_{x^R_{t-1}}^{x^R_t} f_t(x)dx \ge \sum_{t=1}^{n} \int_{\pi_{t-1}}^{\pi_t} f_t(x)dx.$$

The right hand side of the last inequality is the utilitarian welfare of allocation π and should be at least one because π is a proportional allocation. Furthermore, the left hand side is actually equals to $\mathsf{UW}(\mathbf{A}^L) + \mathsf{UW}(\mathbf{A}^R)$ due to the construction of the density function. Consequently, we can conclude one of $\mathbf{A}^L$ and $\mathbf{A}^R$ has utilitarian welfare at least $\frac{1}{2}$.

The above-mentioned contiguous proportional cake-cutting solution π can be found efficiently. Then, this construction proof can be easily transferred to an efficient algorithm on computing the PROP1 allocation with utilitarian welfare at least $\frac{1}{2}$. ■

Lemma 16 *There is an instance in which no contiguous PROP1 allocation has better than $\Theta(\frac{1}{\sqrt{n}})$ of the optimal utilitarian welfare.*

4.2 Price of PROP1 with Egalitarian Welfare

When concerning egalitarian welfare, we have the following infinite PoF ratio. Recall that the PoF ratio regarding PROP and egalitarian welfare is proven as 1 in [36], which indicates that some PROP allocation achieves OPT_E, and hence so does PROP1. However, as the result below turns out, PROP1 cannot provide any guarantee regarding egalitarian welfare. This fact confirms that if a certain set of instances is excluded, the computed PoF ratios fail to capture the fairness and efficiency trade-off.

Theorem 17 *For egalitarian welfare and PROP1 fairness, the price of fairness is unbounded.*

Proof. Consider an instance with $n \ge 3$ agents and $n+1$ items $E = \{e_1, \dots, e_{n+1}\}$. Agents have identical valuation functions, where $v_i(e_j) = \frac{1}{n^2}$ for $j = 1, \dots, n$ and $v_i(e_{n+1}) = 1 - \frac{1}{n}$. It is straightforward that the optimal egalitarian welfare is at least $\frac{1}{n^2}$. There are three ways to satisfy PROP1 for each agent i: (1) obtaining e_{n+1} (possibly with some items on its left) so that PROP is satisfied, (2) obtaining e_n (possibly with some items on its left) so that $v_i(\{e_n, e_{n+1}\}) \ge \frac{1}{n}$, (3) obtaining nothing so that $v_i(\{e_{n+1}\}) \ge \frac{1}{n}$. Thus, in any PROP1 allocation, at least $n-2$ agents receive nothing and achieve PROP1 via (3), which means the egalitarian welfare is 0 and the price of PROP1 is infinite. ■

Note that the hard instance designed in the proof of Theorem 17 shows an interesting fact and a significant difference between MMS and PROP1. If an agent cannot receive items e_n or e_{n+1}, with MMS, she wants to have one item from $\{e_1, \dots, e_{n-1}\}$; however, with PROP1, she actually prefers to receive nothing due to the connectivity requirement in the definition of PROP1.

5 Two Agents

In previous sections, we have shown that, for general n, when the fairness notions are changed from PROP to MMS and PROP1, PoF ratios change dramatically. Contrarily, in the case of $n = 2$, an interesting special case that has been widely studied in the literature, we find that tight PoF ratios on MMS and PROP1 coincide with the these on PROP proven in [36]. In particular, we first show that there exists an allocation that simultaneously achieves fairness (MMS and PROP1) and significant welfare guarantee, and then derives the tight PoF ratio based on that allocation.

Lemma 18 *For any instance with two agents, there is an allocation that is MMS, PROP1, maximizes egalitarian welfare and achieves utilitarian welfare at least 1.*

Proof sketch. We will consider the allocation $\mathbf{O} = (O_1, \dots, O_n)$ constructed as follows:

- $\mathbf{O}$ first maximizes the egalitarian welfare among all allocations; If there is a tie, $\mathbf{O}$ maximizes the number of items allocated to the agent with smaller value.

We can without loss of generality assume $v_1(O_1) \le v_2(O_2)$. If $v_1(O_1) \ge 1/2$, then $\mathbf{O}$ satisfies all properties mentioned in the statement. We can further assume $v_1(O_1) \le 1/2$. Note that either $v_2(O_2) < 1/2$ or $v_2(O_1) > v_1(O_1)$ violates the construction of $\mathbf{O}$, then we can focus on the case of $v_2(O_1) \le v_1(O_1) < \frac{1}{2} \le v_2(O_2)$.

The fact that $\mathbf{O}$ achieves maximum egalitarian welfare directly follows from the construction. Moreover, facts that $v_2(O_1) \le v_1(O_1)$ and valuations are normalized to one can imply $\mathsf{UW}(\mathbf{O}) \ge 1$. For PROP1, we consider another allocation $\mathbf{O}' = (O'_1, O'_2)$ with $O'_1 = O_1 \cup \{e^*\}$ and $O'_2 = O_2 \setminus \{e^*\}$ where $e^* \in O_2$ is the item such that $O_1 \cup \{e^*\} \in \mathcal{S}$. We prove $v_2(O'_2) < \frac{1}{2}$ and $v_2(O'_1) > \frac{1}{2}$, which then implies $v_1(O'_1) > 1/2$ as $\mathbf{O}$ maximizes egalitarian welfare. Lastly, to prove MMS fairness, we consider $\mathbf{T} = \{T_1, T_2\}$ the MMS_1-defining partition and analyze the bundle-wise inclusion relationship to $\mathbf{O}$. ■

By Lemma 18, we have the following result on the price of MMS and of PROP1.

Theorem 19 *When $n = 2$, $\mathsf{PoF} = \frac{3}{2}$ for utilitarian welfare, no matter the fairness notion is MMS or PROP1; and $\mathsf{PoF} = 1$ for egalitarian welfare, no matter the fairness notion is MMS or PROP1.*

6 Conclusion

In this work, we bound the ratios of PoF regarding both utilitarian and egalitarian welfare, and MMS and PROP1 fairness, for the discrete cake cutting problem when the agents need to receive contiguous blocks of items. For both fairness notions, we provide tight ratios regarding the egalitarian welfare, but regarding the utilitarian welfare, there is still a gap. An immediate open problem is to explore the tight ratio for this setting. There are several future directions. One can consider alternative fairness notions (and their relaxations) such as equitability [25, 26] and envy-freeness [30]. We can also extend the line structure to other graphs, and study the allocation of chores.

Acknowledgements

This work is funded by GDSTC (No. 2023A1515010592), NSFC (No. 62102333) and HKSAR RGC (No. PolyU 25211321).

References

[1] Alessandro Agnetis, Bo Chen, Gaia Nicosia, and Andrea Pacifici, 'Price of fairness in two-agent single-machine scheduling problems', *Eur. J. Oper. Res.*, **276**(1), 79–87, (July 2019).
[2] Allesandro Agnetis, Pitu B. Mirchandani, Dario Pacciarelli, and Andrea Pacifici, 'Scheduling Problems with Two Competing Agents', *Oper. Res.*, **52**(2), 229–242, (April 2004).
[3] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A. Voudouris, 'Maximum nash welfare and other stories about EFX', *Theor. Comput. Sci.*, **863**, 69–85, (2021).
[4] Georgios Amanatidis, Evangelos Markakis, Afshin Nikzad, and Amin Saberi, 'Approximation algorithms for computing maximin share allocations', *ACM Trans. Algorithms*, **13**(4), (dec 2017).
[5] Yonatan Aumann and Yair Dombb, 'The efficiency of fair division with connected pieces', *ACM Trans. Economics and Comput.*, **3**(4), 23:1–23:16, (2015).
[6] Haris Aziz, Ioannis Caragiannis, Ayumi Igarashi, and Toby Walsh, 'Fair allocation of indivisible goods and chores', *Auton. Agents Multi Agent Syst.*, **36**(1), 3, (2022).
[7] Haris Aziz and Simon Mackenzie, 'A discrete and bounded envy-free cake cutting protocol for any number of agents', in *FOCS*, pp. 416–427. IEEE Computer Society, (2016).
[8] Haris Aziz, Hervé Moulin, and Fedor Sandomirskiy, 'A polynomial-time algorithm for computing a pareto optimal and almost proportional allocation', *Oper. Res. Lett.*, **48**(5), 573–578, (2020).
[9] Siddharth Barman, Umang Bhaskar, and Nisarg Shah, 'Optimal bounds on the price of fairness for indivisible goods', in *WINE*, volume 12495 of *Lecture Notes in Computer Science*, pp. 356–369. Springer, (2020).
[10] Siddharth Barman, Sanath Kumar Krishnamurthy, and Rohit Vaish, 'Finding fair and efficient allocations', in *EC*, pp. 557–574. ACM, (2018).
[11] Xiaohui Bei, Xinhang Lu, Pasin Manurangsi, and Warut Suksompong, 'The price of fairness for indivisible goods', *Theory Comput. Syst.*, **65**(7), 1069–1093, (2021).
[12] Dimitris Bertsimas, Vivek F. Farias, and Nikolaos Trichakis, 'The Price of Fairness', *Oper. Res.*, **59**(1), 17–31, (February 2011).
[13] Vittorio Bilò, Ioannis Caragiannis, Michele Flammini, Ayumi Igarashi, Gianpiero Monaco, Dominik Peters, Cosimo Vinci, and William S. Zwicker, 'Almost envy-free allocations with connected bundles', *Games Econ. Behav.*, **131**, 197–221, (2022).
[14] Sylvain Bouveret, Katarína Cechlárová, Edith Elkind, Ayumi Igarashi, and Dominik Peters, 'Fair division of a graph', in *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI'17*, pp. 135–141. ijcai.org, (2017).
[15] Sylvain Bouveret and Michel Lemaître, 'Characterizing conflicts in fair division of indivisible goods using a scale of criteria', *Auton. Agents Multi Agent Syst.*, **30**(2), 259–290, (2016).
[16] Steven J Brams and Alan D Taylor, 'An envy-free cake division protocol', *The American Mathematical Monthly*, **102**(1), 9–18, (1995).
[17] Ioannis Caragiannis, Christos Kaklamanis, Panagiotis Kanellopoulos, and Maria Kyropoulou, 'The efficiency of fair division', *Theory Comput. Syst.*, **50**(4), 589–610, (2012).
[18] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang, 'The unreasonable fairness of maximum nash welfare', *ACM Trans. Economics and Comput.*, **7**(3), 12:1–12:32, (2019).
[19] Ioannis Caragiannis, Evi Micha, and Nisarg Shah, 'A little charity guarantees fair connected graph partitioning', (2022).
[20] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia, 'Truth, justice, and cake cutting', *Games Econ. Behav.*, **77**(1), 284–297, (January 2013).
[21] Vincent Conitzer, Rupert Freeman, and Nisarg Shah, 'Fair public decision making', in *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17*, pp. 629–646. ACM, (2017).
[22] L. E. Dubins and E. H. Spanier, 'How to Cut A Cake Fairly', *The American Mathematical Monthly*, **68**(1), 1–17, (1961).
[23] Lester E Dubins and Edwin H Spanier, 'How to cut a cake fairly', *The American Mathematical Monthly*, **68**(11), 1–17, (1961).
[24] Francis Edward Su, 'Rental harmony: Sperner's lemma in fair division', *The American mathematical monthly*, **106**(10), 930–942, (1999).
[25] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia, 'Equitable allocations of indivisible goods', in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI'19 2019*, pp. 280–286. ijcai.org, (2019).
[26] Rupert Freeman, Sujoy Sikdar, Rohit Vaish, and Lirong Xia, 'Equitable allocations of indivisible chores', in *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '20*, pp. 384–392. International Foundation for Autonomous Agents and Multiagent Systems, (2020).
[27] Sandy Heydrich and Rob van Stee, 'Dividing connected chores fairly', *Theor. Comput. Sci.*, **593**, 51–61, (August 2015).
[28] Felix Höhne and Rob van Stee, 'Allocating contiguous blocks of indivisible chores fairly', *Inf. Comput.*, 104739, (March 2021).
[29] Xin Huang and Pinyan Lu, 'An algorithmic framework for approximating maximin share allocation of chores', in *EC '21: The 22nd ACM Conference on Economics and Computation*, pp. 630–631. ACM, (2021).
[30] Ayumi Igarashi, 'How to cut a discrete cake fairly', *CoRR*, **abs/2209.01348**, (2022).
[31] Ayumi Igarashi and Dominik Peters, 'Pareto-optimal allocation of indivisible goods with connectivity constraints', in *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI'19*, pp. 2045–2052. AAAI Press, (2019).
[32] David Kurokawa, Ariel D. Procaccia, and Junxing Wang, 'Fair enough: Guaranteeing approximate maximin shares', *J. ACM*, **65**(2), 8:1–8:27, (2018).
[33] Richard J. Lipton, Evangelos Markakis, Elchanan Mossel, and Amin Saberi, 'On approximately fair allocations of indivisible goods', in *EC*, pp. 125–131. ACM, (2004).
[34] Benjamin Plaut and Tim Roughgarden, 'Almost envy-freeness with general valuations', *SIAM J. Discret. Math.*, **34**(2), 1039–1068, (2020).
[35] Hugo Steinhaus, 'The problem of fair division', *Econometrica*, **16**, 101–104, (1948).
[36] Warut Suksompong, 'Fairly allocating contiguous blocks of indivisible items', *Discrete Applied Mathematics*, **260**, 227–236, (May 2019).
[37] Warut Suksompong, 'Constraints in fair division', *SIGecom Exch.*, **19**(2), 46–61, (2021).
[38] Ankang Sun, Bo Chen, and Xuan Vinh Doan, 'Connections between fairness criteria and efficiency for allocating indivisible chores', in *AAMAS '21: 20th International Conference on Autonomous Agents and Multiagent Systems*, pp. 1281–1289. ACM, (2021).
[39] Gerhard J. Woeginger, 'A polynomial-time approximation scheme for maximizing the minimum machine completion time', *Oper. Res. Lett.*, **20**(4), 149–154, (1997).