Fair Division of Time: Multi-layered Cake Cutting

Hadi Hosseini
Rochester Institute of Technology
hhvcs@rit.edu

Ayumi Igarashi
University of Tokyo
igarashi@mist.i.u-tokyo.ac.jp

Andrew Searns
Rochester Institute of Technology
abs2157@rit.edu

ABSTRACT

We initiate the study of multi-layered cake cutting with the goal of fairly allocating multiple divisible resources (layers of a cake) among a set of agents. The key requirement is that each agent can only utilize a single resource at each time interval. Several real-life applications exhibit such restrictions on overlapping pieces, for example, assigning time intervals over multiple facilities and resources or assigning shifts to medical professionals. We investigate the existence and computation of envy-free and proportional allocations. We show that envy-free allocations are guaranteed to exist for up to three agents with two types of preferences, when the number of layers is two. We further devise an algorithm for computing proportional allocations for any number of agents when the number of layers is factorable to three and/or some power of two.

1 INTRODUCTION

Consider a group of students who wish to use multiple college facilities such as a conference room and an exercise room over different periods of time. Each student has a preference over what facility to use at different time of the day: Alice prefers to set her meetings in the morning and exercise in the afternoon, whereas Bob prefers to start the day with exercising for a couple of hours and meet with his teammates in the conference room for the rest of the day.

The fair division literature has extensively studied the problem of dividing a heterogeneous divisible resource (aka a *cake*) among several agents who may have different preference over the various pieces of the cake [5, 17, 18]. These studies have resulted in a plethora of axiomatic and existence results [3, 12] as well as computational solutions [2, 16] under a variety of assumptions, and were successfully implemented in practice (see [4, 15] for an overview). In the case of Alice and Bob, each facility represents a layer of the cake in a *multi-layered cake cutting* problem, and the question is how to allocate the time intervals (usage right) of the facilities according to their preferences in a fair manner.

One naive approach is to treat each cake independently and solve the problem through well-established cake-cutting techniques by performing a fair division on each layer separately. However, this approach has major drawbacks: First, the final outcome, although fair on each layer, may not necessarily be fair overall. Second, the allocation may not be feasible, i.e., it may assign two overlapping pieces (time intervals) to a single agent. In our example, Alice cannot simultaneously utilize the exercise room and the conference room at the same time if she receives overlapping intervals. Several other application domains exhibit similar structures over resources: assigning nurses to various wards and shifts, doctors to operation rooms, and research equipment to groups, to name a few.

In multi-layared cake cutting, each layer represents a divisible resource. Each agent has additive preferences over every disjoint (non-overlapping) intervals. A division of a multi-layered cake is *feasible* if no agent's share contain overlapping intervals, and is contiguous if each allocated piece of a layer is contiguous. There has been some recent work on dividing multiple cakes among agents [7, 11]. Yet, none of the previous work considered the division of multiple resources under feasibility and contiguity constraints. Therefore, in this setting we ask the following research question:

> *What fairness guarantees can be achieved under feasibility and contiguity constraints for various number of agents and layers?*

1.1 Our Results

We initiate the study of the multi-layered cake cutting problem for allocating divisible resources, under contiguity and feasibility requirements. Our focus is on two fairness notions, *envy-freeness* and *proportionality*. Envy-freeness (EF) requires that each agent believes no other agent's share is better than its share of the cake. Proportionality (Prop) among n agents requires that each agent receives a share that is valued at least $\frac{1}{n}$ of the value of the entire cake. For efficiency, we consider *complete* divisions with no leftover pieces.

Focusing on envy-free divisions, we show the existence of envy-free and complete allocations for two-layered cakes and up to three agents with at most two types of preferences. These cases are particularly appealing since many applications often deal with dividing a small number of resources among few agents (e.g. assigning meeting rooms). We then show that proportional complete allocations exist for three agents and three layers and can be computed efficiently. Subsequently, we show that although this result cannot be immediately extended to any number of agents and layers, a proportional complete allocation exists when the number of layers is factorable to a product of three and powers of two. We defer some proofs to the full version of the paper, due to space constraints.

1.2 Related Work

In recent years, cake cutting has received significant attention in artificial intelligence and economics as a metaphor for algorithmic approaches in achieving fairness in allocation of resources [1, 6, 10, 16]. Recent studies have focused on the fair division of resources when agents have requirements over multiple resources that must be simultaneously allocated in order to carry out certain tasks (e.g. CPU and RAM) [8, 9, 14]. The most relevant work to

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 2020, Auckland, New Zealand

© 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

https://doi.org/doi

Agents (n)	Layers (m)	EF	Prop
2	2	✓(Thm. 4.1)	✓
3	2	✓(Thm. 4.3†)	✓
$n \geq m$	2^a, $a \in \mathbb{Z}_+$	?	✓(Thm. 5.8)
$n \geq m$	$2^a 3^b$, $a \in \mathbb{Z}_+, b \in \{0,1\}$	?	✓(Thm. 5.7◇)

Table 1: The overview of our results. † assumes two types of agents' preferences. ◇ indicates that existence holds without contiguity requirement. Note that when $m > n$, no complete and feasible (non-overlapping) solution exists.

ours is the envy-free multi-cake fair division that considers dividing multiple cakes among agents with linked preferences over the cakes. Here, agents can simultaneously benefit from all allocated pieces with no constraints. They show that envy-free divisions with only few cuts exist for two agents and many cakes, as well as three agents and two cakes [7, 11, 13]. In contrast, a multi-layered cake cutting requires non-overlapping pieces. Thus, [7]'s generalized envy-freeness notion on multiple cakes does not immediately imply envy-freeness in our setting and no longer induces a feasible division.

2 OUR MODEL

Our setting includes a set of *agents* denoted by $N = [n]$, a set of *layers* denoted by $L = [m]$, where for a natural number $s \in \mathbb{N}$, $[s] = \{1, 2, \ldots, s\}$. Given two real numbers $x, y \in \mathbb{R}$, we write $[x, y] = \{ z \in \mathbb{R} \mid x \leq z \leq y \}$ to denote an interval. We denote by $\mathbb{R}_+$ (respectively $\mathbb{Z}_+$) the set of non-negative reals (respectively, integers) including 0. A *piece* of cake is a finite set of disjoint subintervals of $[0, 1]$. We say that a subinterval of $[0, 1]$ is a *contiguous piece* of cake. An *m-layered cake* is denoted by $C = (C_j)_{j \in L}$ where $C_j \subseteq [0, 1]$ is a contiguous piece for $j \in L$. We refer to each $j \in L$ as j-th *layer* and C_j as j-th *layered cake*.

Each agent i is endowed with a non-negative *integrable density function* $v_{ij} : C_j \to \mathbb{R}_+$. For a given piece of cake X of j-th layer, $V_{ij}(X)$ denotes the value assigned to it by agent i, i.e., $V_{ij}(X) = \sum_{I \in X} \int_{x \in I} v_{ij}(x)dx$. These functions are assumed to be normalized over layers: $\sum_{j \in L} V_{ij}(C_j) = 1$ for each $i \in N$. A *layered piece* is a sequence $\mathcal{X} = (X_j)_{j \in L}$ of pieces of each layer $j \in L$; a layered piece is said to be *contiguous* if each X_j is a contiguous piece of each layer. We assume valuation functions are *additive* on layers and write $V_i(\mathcal{X}) = \sum_{j \in L} V_{ij}(X_j)$.

A layered contiguous piece is said to be *non-overlapping* if no two pieces from different layers overlap, i.e, for any pair of distinct layers $j, j' \in L$ and for any $I \in X_j$ and $I' \in X_{j'}$, $I \cap I' = \emptyset$. For two layered pieces $\mathcal{X}$ and $\mathcal{X}'$, we say that agent i *weakly prefers* $\mathcal{X}$ to $\mathcal{X}'$ if $V_i(\mathcal{X}) \geq V_i(\mathcal{X}')$.

A *multi-allocation* $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n)$ is a partition of the m-layered cake C where each $\mathcal{A}_i = (A_{ij})_{j \in L}$ is a layered piece of the cake allocated to agent i; we refer to each $\mathcal{A}_i$ as a *bundle* of i. For a multi-allocation $\mathcal{A}$ and $i \in N$, we write $V_i(\mathcal{A}_i) = \sum_{j \in L} V_{ij}(A_{ij})$ to denote the value of agent i for $\mathcal{A}_i$. A *multi-allocation* $\mathcal{A}$ is said to be

- *contiguous* if each $\mathcal{A}_i$ for $i \in N$ is contiguous;
- *feasible* if each $\mathcal{A}_i$ for $i \in N$ is non-overlapping.

We focus on *complete* multi-allocations where the entire cake must be allocated. Notice that some layers may be disjoint (see Figure 1), and the number of agents must exceed the number of layers, i.e. $n \geq m$; otherwise the multi-allocation will contain overlapping pieces.

Fairness. A multi-allocation is said to be *envy-free* if no agent *envies* the others, i.e., $V_i(\mathcal{A}_i) \geq V_i(\mathcal{A}_{i'})$ for any pair of agents $i, i' \in N$. A multi-allocation is said to be *proportional* if each agent gets his *proportional fair share*, i.e., $V_i(\mathcal{A}_i) \geq \frac{1}{n}$ for any $i \in N$. The following implication, which is well-known for the standard setting, holds in our setting as well.

Lemma 2.1. *An envy-free complete multi-allocation satisfies proportionality.*

Proof. Consider an envy-free complete multi-allocation $\mathcal{A}_i = (A_{ij})_{j \in L}$ and an agent $i \in N$. By envy-freeness, we have that $V_i(\mathcal{A}_i) \geq V_i(\mathcal{A}_j)$ for any $j \in N$. Summing over $j \in N$, we get $V_i(\mathcal{A}_i) \geq \frac{1}{n} \sum_{j \in N} V_i(\mathcal{A}_j) = \frac{1}{n}$ by additivity. □

Example 2.2 (Resource sharing). Suppose that there are three meeting rooms r_1, r_2, and r_3 with different capacities, and three researchers Alice, Bob, and Charlie. The first room is available all day, the second and the third rooms are only available in the morning and late afternoon, respectively (see Fig. 1). Each researcher has a preference over the access time to the shared rooms. For example, Alice wants to have a group meeting in the larger room in the morning and then have an individual meeting in the smaller one in the afternoon.

The m-layered cuts. In order to cut the layered cake while satisfying the non-overlapping constraint, we define a particular approach for partitioning the entire cake into diagonal pieces. Consider the m-layered cake C where m is an even number. For each point x of the interval $[0, 1]$, we define

- $LR(x, C) = (\bigcup_{j=1}^{\frac{m}{2}} C_j \cap [0, x]) \cup (\bigcup_{j=\frac{m}{2}+1}^{m} C_j \cap [x, 1])$;
- $RL(x, C) = (\bigcup_{j=1}^{\frac{m}{2}} C_j \cap [x, 1]) \cup (\bigcup_{j=\frac{m}{2}+1}^{m} C_j \cap [0, x])$.

$LR(x, C)$ consists of the top-half subintervals of points left of x and the lower-half subintervals of points right of x; similarly, $RL(x, C)$ consists of the top-half subintervals of points right of x and the lower-half subintervals of points left of x (Fig. 2). We abuse the notation and write $LR(x) = LR(x, C)$ and $RL(x) = RL(x, C)$ if C is clear from the context.

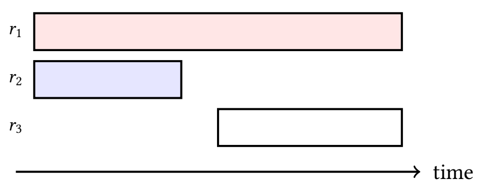

Figure 1: Example of a multi-layered cake. There are three meeting rooms r_1, r_2, and r_3 with different capacities, shared among several research groups.

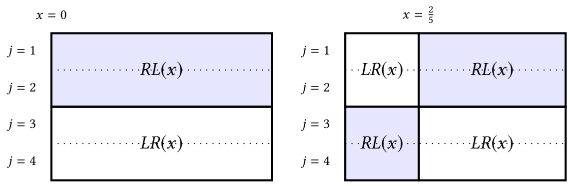

Figure 2: Examples of the partitions induced by $x = 0$ and $x = \frac{2}{5}$ for a four-layered cake.

Computational model. Following the standard *Robertson-Webb Model* [17], we introduce two types of queries: those for a cake on each layer (called a *short knife*) and those for the entire cake (called a *long knife*).

Short knife. Short eval query: given an interval $[x, y]$ of the j-th layered cake C_j, $eval_j(i, x, y)$ asks agent i for its value $[x, y]$, i.e., $V_{ij}([x, y])$. Short cut query: given a point x and $r \in [0, 1]$, $cut_j(i, x, r)$ asks agent i for the minimum point y such that $V_{ij}([x, y]) = r$.

Long knife. Long eval query: given a point x, $eval(i, x)$ asks agent i for its value $LR(x)$, i.e., $V_i(LR(x))$. Long cut query: given $r \in [0, 1]$, $cut(i, r)$ asks agent i for the minimum point x such that $V_i(LR(x)) = r$ if such point x exists.

3 EXISTENCE OF A SWITCHING POINT

We start by showing the existence of a point x that equally divides the entire cake into two pairs of diagonal pieces, both for the individuals and for the majority; these will serve as a fundamental property in our problem. We say that $x \in [0, 1]$ is a *switching point* for agent i over a layered cake C if $V_i(LR(x)) = V_i(RL(x))$.

Lemma 3.1. *Suppose that the number m of layers is even. Let $r \in \mathbb{R}$ be such that $V_i(LR(0)) \geq r$ and $V_i(RL(0)) \leq r$ for some agent $i \in N$. Then, there exists a point $x \in [0, 1]$ such that i values $LR(x)$ exactly at r, i.e.$V_i(LR(x)) = r$. In particular, a switching point for i always exists.*

Proof. Suppose that $V_i(LR(0)) \geq r$ and $V_i(RL(0)) \leq r$. Consider the function $f(x) = V_i(LR(x))$ for $x \in [0, 1]$. Recall that $f(x)$ is a continuous function written as the sum of continuous functions: $f(x) = \sum_{j=1}^{\frac{m}{2}} V_{ij}(C_j \cap [0, x]) + \sum_{j=\frac{m}{2}+1}^{m} V_{ij}(C_j \cap [x, 1])$. Since $f(0) \geq r$ and $f(1) \leq r$, there is a point $x \in [0, 1]$ with $f(x) = r$ by the intermediate value theorem, which proves the claim. Further, by taking $r = \frac{1}{2}$, the point x where $V_i(LR(x)) = \frac{1}{2}$ is a switching point for agent i. □

We will generalize the notion of a switching point from the individual level to the majority. For layered contiguous pieces $\mathcal{I}$ and $\mathcal{I}'$, we say that the majority weakly prefer $\mathcal{I}$ to $\mathcal{I}'$ (denoted by $\mathcal{I} \overset{m}{\succeq} \mathcal{I}'$) if there exists $S \subseteq N$ such that $|S| \geq \lceil \frac{n}{2} \rceil$ and each $i \in S$ weakly prefers $\mathcal{I}$ to $\mathcal{I}'$. We say that $x \in [0, 1]$ is a *majority switching point* over C if $LR(x) \overset{m}{\succeq} RL(x)$ and $RL(x) \overset{m}{\succeq} LR(x)$. The following lemma guarantees the existence of a majority switching point, for any even number of layers and any number of agents.

Lemma 3.2. *Suppose that the number m of layers is even. Then, there exists a majority switching point for any number $n \geq m$ of agents.*

Proof. Suppose without loss of generality that the majority of agents weakly prefer $LR(0)$ to $RL(0)$. Since $LR(0) = RL(1)$ and $RL(0) = LR(1)$, this means that by the time when the long knife reaches the right-most point, i.e., $x = 1$, the majority preference switches.

Formally, consider the following set of points $x \in [0, 1]$ where the majority weakly prefer $LR(x)$ to $RL(x)$:

$$M := \{\, x \in [0, 1] \mid LR(x) \overset{m}{\succeq} RL(x) \,\}.$$

We will first show that M is a compact set. Clearly, M is bounded. To show that M is closed, consider an infinite sequence as follows $X = \{x_k\}_{k=1,2,\dots} \subseteq M$ that converges to x^*. For each $k = 1, 2, \dots$, we denote by S_k the set of agents who weakly prefer $LR(x_k)$ to $RL(x_k)$; by definition, $|S_k| \geq \lceil \frac{n}{2} \rceil$. Since there are finitely many subsets of agents, there is one subset $S_k \subseteq N$ that appears infinitely often; let S^* be such subset and $\{x^*_k\}_{k=1,2,\dots}$ be an infinite subsequence of X such that for each k, each agent in S^* weakly prefers $LR(x^*_k)$ to $RL(x^*_k)$. Since the valuations V_i for $i \in S^*$ are continuous, each agent $i \in S^*$ weakly prefers $LR(x^*)$ to $RL(x^*)$ at the limit x^*, which implies that $x^* \in M$ and hence M is closed. Now since M is a compact set, the supremum $t^* = \sup M$ belongs to M. By the maximality of t^*, at least $\lceil \frac{n}{2} \rceil$ agents weakly prefer $RL(t^*)$ to $LR(t^*)$. Since $t^* \in M$, at least $\lceil \frac{n}{2} \rceil$ agents weakly prefer $LR(t^*)$ to $RL(t^*)$ as well. Thus, t^* corresponds to a majority switching point. □

4 ENVY-FREE MULTI-LAYERED CAKE CUTTING

Now we will look into the problem of obtaining complete envy-free multi-allocations, while satisfying non-overlapping constraints. When there is only one layer, it is known that an envy-free contiguous allocation exists for any number of agents under mild assumptions on agents' preferences [19, 20]. Given the contiguity and feasibility constraints, the question is whether it is possible to guarantee an envy-free division in the multi-layered cake-cutting model.

4.1 Two agents and two layers

We answer the above question positively for a simple, yet important, case of two agents and two layers. The standard protocol that achieves envy-freeness for two agents is known as the *cut-and-choose* protocol: Alice divides the entire cake into two pieces of equal value. Bob selects his preferred piece over the two pieces, leaving the remainder for Alice.

We extend this protocol to the multi-layered cake cutting using the notion of a switching point. Alice first divides the layered cake into two *diagonal pieces*: one that includes the top left and lower right parts and another that includes the top left and lower left parts of the cake. Our version of the cut-and-choose protocol is specified as follows:

Cut-and-choose protocol for $n = 2$ agents over a two-layered cake C:
Step 1. Alice selects her switching point x over C.
Step 2. Bob chooses a weakly preferred layered contiguous piece among $LR(x)$ and $RL(x)$.
Step 3. Alice receives the remaining piece.

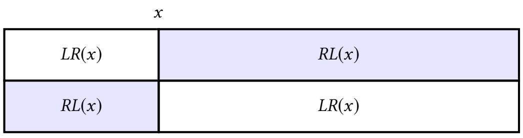

Figure 3: Cut-and-Choose for two-layered cake

Theorem 4.1. *The cut-and-choose protocol yields a complete envy-free multi-allocation that is feasible and contiguous for two agents and a two-layered cake using $O(1)$ number of long eval and cut queries.*

Proof. It is immediate to see that the protocol returns a complete multi-allocation where each agent is assigned to a non-overlapping layered contiguous piece. The resulting allocation satisfies envy-freeness: Bob does not envy Alice since he chooses a preferred piece among $LR(x)$ and $RL(x)$. Alice does not envy Bob by the definition of a switching point. □

As we noted in Section 2, the existence result for two agents does not extend beyond two layers: if there are at least three layers, there is no feasible multi-allocation that completely allocates the cake to two agents.

4.2 Three agents and two layers

We now move on to the case of three agents and two layers. We will design a variant of Stromquist's protocol that achieves envy-freeness for one-layered cake [19]: The referee moves two knives: a short knife and a long knife. The short knife moves from left to right over the top layer and gradually increases the left-most top piece (denoted by Y), while the long knife keeps pointing to the point x, which can partition the remaining cake into two diagonal pieces $LR(x)$ and $RL(x)$ in an envy-free manner. Each agent shouts when the left-most top piece Y becomes at least as highly valuable as the preferred piece among $LR(x)$ and $RL(x)$. Some agent, say s, shouts eventually (before the left-most top piece becomes the top layer), assuming that there is at least one agent who weakly prefers the top layer to the bottom layer. We note that x may be positioned left to y; see Figure 4 for some possibilities of the long knife's locations.

We will show that the above protocol works, for a special case when there are at most two types of preferences: In such cases, the majority switching points coincide with the switching points of an agent with the majority preference.

Lemma 4.2. *Suppose that $m = 2$, $n = 3$, and there are two different agents $i, j \in N$ with the same valuation V. Then, x is a majority switching point over C if and only if x is a switching point for i.*

An obvious implication of the above lemma is that when performing Stromquist's protocol, one can point out to a switching point of an individual, instead of a majority one. This allows the referee to move a long knife continuously. For a given two-layered cake C, we write $C^{-y} = (C_1^{-y}, C_2)$ as a two-layered cake obtained from C where the first segment $[0, y]$ of the top layer is removed, i.e., $C_1^{-y} = C_1 \setminus [0, y]$. For each majority switching point x over C^{-y}, we select three different agents $\ell(x)$, $m(x)$, and $r(x)$ as follows:

- $\ell(x)$ is an agent who weakly prefers $LR(x, C^{-y})$ to $RL(x, C^{-y})$;
- $m(x)$ is an agent who is indifferent between $LR(x, C^{-y})$ and $RL(x, C^{-y})$; and
- $r(x)$ and agent who weakly prefers $RL(x, C^{-y})$ to $LR(x, C^{-y})$.

Moving-knife protocol for $n = 3$ agents over a two-layered cake C: w.l.o.g. assume that at least one agent weakly prefers the top layer ($j = 1$) over the bottom layer ($j = 2$)
Step 1. The referee continuously moves a short knife from the left-most point ($y = 0$) to the right-most point ($y = 1$) over the top layer, while continuously moving a long knife pointing to a switching point over C^{-y} for i. Let y be the position of the short knife and Y be the top layer piece to its left. Let x be the position of the long knife.
Step 2. The referee stops moving the short knife when some agent *s shouts*, i.e., Y becomes at least as highly valuable as the preferred piece among $LR(x, C^{-y})$ and $RL(x, C^{-y})$.
Step 3. We allocate the shouter s to the left-most top piece Y and partitions the rest into $LR(x, C^{-y})$ and $RL(x, C^{-y})$.

- If $s = \ell(x)$, then we allocate $LR(x, C^{-y})$ to $m(x)$ and $RL(x, C^{-y})$ to $r(x)$.
- If $s = m(x)$, then we allocate $LR(x, C^{-y})$ to $\ell(x)$ and $RL(x, C^{-y})$ to $r(x)$.
- If $s = r(x)$, then we allocate $LR(x, C^{-y})$ to $\ell(x)$ and $RL(x, C^{-y})$ to $m(x)$.

Theorem 4.3. *Suppose that $m = 2$, $n = 3$, and for each $i \in N$ and $j \in L$, v_{ij} is continuous. If there are two different agents with the same valuation, an envy-free complete multi-allocation that is feasible and contiguous exists.*

Proof. Assume w.l.o.g. that at least one agent prefers the top layer over the bottom layer. This means that such agent weakly prefers the top layer to any of the pieces $LR(z, C^{-y})$ and $RL(z, C^{-y})$ when $y = 1$. Suppose that $i \in N$ is one of the two different agents with the same valuations. We design the following protocol for three agents over a two-layered cake:

By our assumption, some agent eventually shouts and thus the protocol returns an allocation $\mathcal{A}$. Clearly, $\mathcal{A}$ is feasible, contiguous, and complete. Also, it is easy to see that the shouter s who receives a bundle Y does not envy the other two agents. The agents $i \neq s$ do not envy s because the referee continuously moves both a short and a long knife. Finally, the agents $i \neq s$ do not envy each other by the definition of a majority switching point and by Lemma 4.2. □

In the general case, achieving envy-free multi-allocations deems to be challenging due to the non-monotonicity of valuations over diagonal pieces.[1] Therefore, we focus our attention on the less demanding fairness notion of proportionality.

[1] See Section 6 for an extensive discussion.

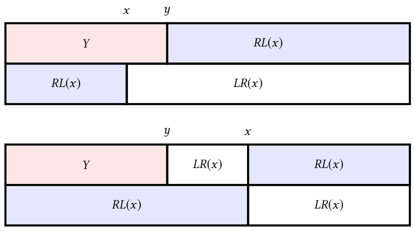

Figure 4: Moving knife protocol for three agents over a two-layered cake. Note that the position of x may appear before y.

5 PROPORTIONAL MULTI-LAYERED CAKE CUTTING

Focusing on a less demanding fairness notion, it turns out that a complete proportional multi-allocation that is feasible exists for a wider class of instances, i.e., when the number m of layers is a product of three and some power of two, and the number n of agents is at least m. Notably, we show that the problem can be decomposed into smaller instances when the number of agents is greater than the number of layers, or the number of layers is a power of two. Building up on the *base cases* of two and three layers, our algorithm recursively calls the same algorithm to decide on how to allocate the cake of the sub-problems. We first proceed by showing how to solve the base case: For three layers and three agents, a feasible proportional multi-allocation exists and can be efficiently computed.

Theorem 5.1. *A proportional complete multi-allocation that is feasible exists for three layers and three agents and can be computed using $O(1)$ number of short eval queries and long eval and cut queries. Further, each bundle of the resulting multi-allocation includes at most two contiguous pieces within each layer.*

We start by showing two auxiliary lemmas. We define a *merge* of two disjoint contiguous pieces I_j and $I_{j'}$ of layers j and j' as replacing the j-th layered cake with the union $I_j \cup I_{j'}$ and removing j'-th layered cake. The *merge* of a finite sequence of mutually disjoint contiguous pieces $(I_1, \dots, I_k)$ can be defined inductively: merge $(I_1, \dots, I_{k-1})$ and then apply the merge operation to the resulting outcome and I_k. Now we observe that if there are two disjoint layers, one can safely merge these layers and reduce the problem size.

Lemma 5.2. *Suppose that C_j and $C_{j'}$ are two disjoint layers of a layered cake C, and C' is obtained from C by merging C_j and $C_{j'}$. Then, each non-overlapping contiguous layered piece of C' is a non-overlapping contiguous layered piece of the original cake C.*

The above lemma can be generalized further: Let C be a $2m$-layered cake and $x \in [0, 1]$. We define a *merge* of $LR(x) = (S_j)_{j \in L}$ by merging the pair (S_j, S_{j+m}) for each $j \in [m]$. A *merge* of $RL(x)$ can be defined analogously. Such operation still preserves both feasibility and contiguity.

Corollary 5.3. *Let C be a $2m$-layered cake and $x \in [0, 1]$. Suppose that C' is a m-layered cake obtained by merging $LR(x)$ or $RL(x)$. Then, each non-overlapping contiguous layered piece of C' is a non-overlapping contiguous layered piece of the original cake C.*

Below, we show that each agent can divide the entire cake into n equally valued layered pieces. A multi-allocation $\mathcal{A}$ is *equitable* if for each agent $i \in N$, $V_i(\mathcal{A}_i) = \frac{1}{n}$. We design a recursive algorithm that iteratively finds two layers for which one has value at most $\frac{1}{m}$ and at least $\frac{1}{m}$ and removes a pair of diagonal pieces of value exactly $\frac{1}{m}$ from the two layers.

Lemma 5.4. *For any number m of layers and any number $n = m$ of agents with the identical valuations, an equitable complete multi-allocation that is feasible and contiguous exists and can be found using $O(m^2)$ number of short eval queries and $O(m)$ number of long cut queries.*

Proof. We denote by $V = V_i$ the valuation function for each agent $i \in N$. Consider the following recursive algorithm $\mathcal{D}$ that takes a subset N' of agents with $|N'| \geq 1$, a $|L'|$-layered cake C', and a valuation profile $(V_i)_{i \in N'}$, and returns an equitable complete multi-allocation of the layered cake to the agents. When $|L'| = |N'| = 1$, then the algorithm allocates the entire cake to the single agent. Suppose that $|L'| = |N'| \geq 2$. The algorithm first finds a layer j whose entire value is at most $\frac{1}{m}$ and another layer j' whose entire value is at least $\frac{1}{m}$. The algorithm $\mathcal{D}$ then finds a point $x \in [0, 1]$ where $V(S_j \cup S_{j'}) = \frac{1}{m}$ for $S_j = C_j \cap [0, x]$ and $S_{j'} = C_j \cap [x, 1]$; such point exists due to Lemma 3.1. We allocate $S_j \cup S_{j'}$ to one agent and apply $\mathcal{D}$ to the remaining cake C'' with $|N'| - 1$ agents where C'' is obtained from merging the remaining j-th layered cake $C_j \setminus S_j$ and the j'-th layered cake $C_{j'} \setminus S_{j'}$. The correctness of the algorithm as well as the bound on the query complexity are immediate. □

Equipped with Lemma 5.4, we will prove Theorem 5.1.

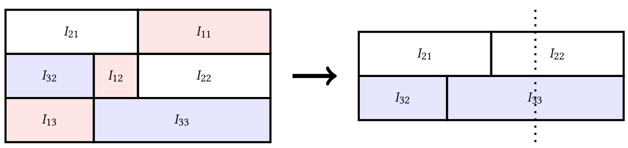

Figure 5: Protocol for proportionality for three agents and three layers. Alice divides the entire cake into three equally valued layered pieces $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$ (the left picture). Here, $\mathcal{I}_i = (I_{ij})_{j=1,2,3}$ for each $i = 1, 2, 3$ where $I_{23} = I_{31} = \emptyset$. After allocating $\mathcal{I}_1$ to either Alice or Charlie, the algorithm merges $\mathcal{I}_2$, and $\mathcal{I}_3$ and applies the cut-and-choose among the remaining agents (the right picture).

Proof of Thm. 5.1. Suppose there are three agents: Alice, Bob, and Charlie. Our procedure for achieving proportionality on a three-layered cake works as follows. See Figure 5 for an illustration.

A protocol for proportionality for $n = 3$ agents over a three-layered cake C:
Step 1. Alice partitions the cake into three non-overlapping layered contiguous pieces $\mathcal{I}_1$, $\mathcal{I}_2$, and $\mathcal{I}_3$ which she considers of equal value, using the algorithm in the proof of Lemma 5.4.
Step 2. Assume w.l.o.g. that $\mathcal{I}_1$ is the piece for which Bob has value at most $\frac{1}{3}$.
Step 3. Decide on the agent who is assigned to $\mathcal{I}_1$ depending on the preference of Charlie.
Step 3-1. If Charlie values $\mathcal{I}_1$ at least his proportional fair share $\frac{1}{3}$, then allocate it to him.
Step 3-2. Otherwise, allocate $\mathcal{I}_1$ to Alice.
Step 4. Merge all the disjoint contiguous pieces in $\mathcal{I}_2$ and $\mathcal{I}_3$, respectively, and create a two-layered cake C' consisting of each merge of $\mathcal{I}_2$ and $\mathcal{I}_3$. Apply the cut-and-choose protocol to C' among the remaining agents.

We will first show that the resulting multi-allocation $\mathcal{A}$ is proportional. Clearly, the agent who gets $\mathcal{I}_1$ has proportional fair share for his bundle. Further, each i of the remaining agents have value at least $\frac{2}{3}$ for the remaining instance; thus, $V_i(\mathcal{A}_i)$ is at least $\frac{V_i(\mathcal{I}_2)+V_i(\mathcal{I}_3)}{2} \geq \frac{1}{3}$. Further, each bundle contains at most two contiguous pieces from each layer of C: If $\mathcal{A}_i = \mathcal{I}_1$, then $\mathcal{A}_i$ is contiguous. If $\mathcal{A}_i = LR(x, C')$ for some $x \in R$, then each layer of C' contains at most one contiguous piece of the original layers in C; since $\mathcal{A}_i$ is contiguous with respect to C', it contains at most two contiguous pieces of every distinct layer of C. A similar argument applies to the case when $\mathcal{A}_i = RL(x, C')$ for some $x \in R$. Finally, it can be easily verified that each bundle $\mathcal{A}_i$ is non-overlapping. This completes the proof. □

We are now ready to prove that a proportional complete multi-allocation exists for any $n = m$ when m is a product of some power of 2 and 3. In essence, the existence of a majority switching point, as proved in Lemma 3.2, allows us to divide the problem into two instances. We will repeat this procedure until the number of layers of the subproblem becomes either 2 or 3, for which we know the existence of a proportional, feasible multi-allocation by Theorem 4.1 and Theorem 5.1.

Theorem 5.5. *A proportional complete multi-allocation that is feasible exists for any number m of layers and any number $n = m$ of agents where $m = 2^a 3^b$ for some $a \in \mathbb{Z}_+$ and $b \in \{0, 1\}$.*

Proof. We design the following recursive algorithm $\mathcal{D}$ that takes a subset N' of agents with $|N'| \geq 2$, a $|L'|$-layered cake C', and a valuation profile $(V_i)_{i \in N'}$, and returns a proportional complete multi-allocation of the cake to the agents which is feasible. Suppose that $m = n$. If $m = n = 1$, then we allocate the entire cake to the single agent. If $m = n = 2$, we run the cut-and-choose algorithm as described in the proof of Theorem 4.1. If $m = n = 3$, we run the procedure as described in the proof of Theorem 5.1. Now consider the case when $m = n = 2^a 3^b$ for some integers $a \geq 1$ and $b \in \{0, 1\}$. Then the algorithm finds a majority switching point x over C'. We let $\mathcal{I}_1 = LR(x)$ and $\mathcal{I}_2 = RL(x)$. By definition of a majority switching point and the fact that n is even, we can partition the set of agents N' into N_1 and N_2 where N_1 is the set of agents who weakly prefer $\mathcal{I}_1$ to $\mathcal{I}_2$, N_2 be the set of agents who weakly prefer $\mathcal{I}_2$ to $\mathcal{I}_1$, and $|N_k| = \frac{|N'|}{2}$ for each $k = 1, 2$. We apply $\mathcal{D}$ to the merge of $\mathcal{I}_k$ with the agent set N_k for each $k = 1, 2$, respectively.

We will prove by induction on m that the complete multi-allocation $\mathcal{A}$ returned by $\mathcal{D}$ satisfies proportionality as well as feasibility. This is clearly true when $m = n = 2$ due to Lemma 2.1 and Theorem 4.1. The claim also holds for $m = n = 3$, due to Theorem 5.1. Suppose that the claim holds for $m = n = 2^a 3^b$ with $1 \leq a \leq k - 1$; we will prove it for $a = k$. Suppose that the algorithm divides the input cake C' via a majority switching point x into $\mathcal{I}_1 = LR(x)$ and $\mathcal{I}_2 = RL(x)$. Suppose that (N_1, N_2) is a partition of the agents where N_1 is the set of agents who weakly prefer $\mathcal{I}_1$ to $\mathcal{I}_2$, N_2 is the set of agents who weakly prefer $\mathcal{I}_2$ to $\mathcal{I}_1$, and $|N_k| = \frac{|N'|}{2}$ for each $k = 1, 2$. Observe that each agent $i \in N_1$ weakly prefers $\mathcal{I}_1$ to $\mathcal{I}_2$ and thus $V_i(\mathcal{I}_1) \geq \frac{1}{2} V_i(C')$. Similarly, $V_i(\mathcal{I}_2) \geq \frac{1}{2} V_i(C')$ for each $i \in N_2$. Thus, by the induction hypothesis, each agent i has value at least $\frac{1}{|N'|} V_i(C')$ for its allocated piece $\mathcal{A}_i$. By Corollary 5.3, the feasibility of $\mathcal{A}$ is immediate. □

It remains open whether a proportional contiguous multi-allocation exists when the number of layers is three. A part of the reason is that our algorithm for finding an equitable multi-allocation (Lemma 5.4) may not return a 'balanced' partition: The number of pieces contained in each layered piece may not be the same when the number of layers is odd. For example, one layered piece may contain pieces from three different layers while the other two parts may contain pieces from two different layers, as depicted in Figure 5. However, we are able to avoid this problem when the number of layers is a power of two: Indeed, in such cases, proportionality, contiguity, and feasibility are compatible with each other.

Theorem 5.6. *A proportional complete multi-allocation that is feasible and contiguous exists for any number m of layers and any number $n = m$ of agents where $m = 2^a$ for some $a \in \mathbb{Z}_+$.*

We will generalize the above theorems to the case when the number of agents is strictly greater than the number of layers. Intuitively, when $n > m$, then there is at least one layer whose sub-piece can be 'safely' allocated to some agent without violating the non-overlapping constraint.

Theorem 5.7. *A proportional complete multi-allocation that is feasible exists for any number m of layers and any number $n \geq m$ of agents where $m = 2^a 3^b$ for some $a \in \mathbb{Z}_+$ and $b \in \{0, 1\}$.*

Proof. We design the following recursive algorithm $\mathcal{D}$ that takes a subset N' of agents with $|N'| \geq 2$, a $|L'|$-layered cake C', and a valuation profile $(V_i)_{i \in N'}$, and returns a proportional complete multi-allocation of the layered cake to the agents which is feasible. For $n = m$, we apply the algorithm described in the proof of Theorem 5.5. Suppose that $n > m$. The algorithm first identifies a layer C_j whose entire valuation is at least $\frac{1}{n}$ for some agent; assume w.l.o.g. that $j = 1$. We move a knife from left to right over the top cake C_1 until some agent i *shouts*, i.e., agent i finds the left contiguous piece Y at least as highly valued as his proportional fair share $\frac{1}{n}$. The algorithm $\mathcal{D}$ then gives the piece to the shouter. To decide on the allocation of the remaining items, we apply $\mathcal{D}$ to the reduced instance $(N' \setminus \{i\}, (C'_j)_{j \in L}, (V_{i'})_{i' \in N' \setminus \{i\}})$ where $C'_j = C_j \setminus Y$ for $j = 1$ and $C'_j = C_j$ for $j \neq 1$. One can prove by induction on $|N'|$ that the multi-allocation $\mathcal{A}$ returned by $\mathcal{D}$ satisfies proportionality as well as feasibility.

We will prove by induction on $|N'|$ that the complete multi-allocation $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \ldots, \mathcal{A}_n)$ returned by $\mathcal{D}$ satisfies proportionality as well as feasibility. This is clearly true when $m = |N'|$, due to Theorem 5.5. Suppose that the claim holds for $|N'|$ with $m \leq |N'| \leq k-1$; we will prove it for $|N'| = k$. Suppose agent i is the shouter who gets the left contiguous piece Y. Clearly, agent i receives her proportional share under $\mathcal{A}$. Observe that all remaining agents have the value at least $\frac{|N'|-1}{|N'|}V_i(C')$ for the remaining cake. Thus, by the induction hypothesis, each agent $i' \neq i$ has value at least $\frac{1}{|N'|}V_i(C')$ for its allocated piece $\mathcal{A}_{i'}$. This completes the proof. □

Theorem 5.8. *A proportional complete multi-allocation that is feasible and contiguous exists for any number m of layers and any number $n \geq m$ of agents where $m = 2^a$ for some $a \in \mathbb{Z}_+$.*

6 DISCUSSION

We provided protocols that find an envy-free multi-allocation of a two-layered cake for two or three agents with at most two types of preferences. An obvious question is whether such allocation also exists for any number n of agents over a m layered cake when $n \geq m$. One might expect that the Simmon-Su's technique [20] using Sperner's Lemma can be adopted to our setting by considering all possible diagonal pieces. However, this approach may not work because multi-layered cake-cutting necessarily exhibits non-monotonicity in that the value of a pair of diagonal pieces may decrease when the knife moves from left to right.

For the case when the contiguity constraint is relaxed, one can show the existence of an envy-free feasible multi-allocation, when each density function v_{ij} is continuous and $m \leq n$: We can reduce the problem to finding a "perfect" allocation of a one-layered cake. We will defer the formal proof to the full version of the paper.

With respect to proportionality, one intriguing future direction is extending our positive algorithmic results to any m, which requires careful consideration of contiguity and feasibility, which are often at odds with completeness. Lastly, the compatibility of the fairness notions with a more demanding efficiency requirement of Pareto optimality, and studying its query complexity is open in the multi-layered cake-cutting problem.

ACKNOWLEDGEMENT

Hadi Hosseini was supported by the National Science Foundation (grant IIS-1850076). Ayumi Igarashi was supported by the KAKENHI Grant-in-Aid for JSPS Fellows no. 18J00997 and JST, ACT-X. The authors would like to thank Yuki Tamura for introducing the problem to us and for the fruitful discussions. The authors also acknowledge the helpful comments by the GAIW and IJCAI reviewers. We are grateful to an anonymous GAIW reviewer for the proof of the existence of an envy-free multi-allocation in the general case.

REFERENCES

[1] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 416–427. IEEE, 2016.

[2] Haris Aziz and Simon Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, pages 454–464. ACM, 2016.

[3] Julius B Barbanel. *The geometry of efficient fair division*. Cambridge University Press, 2005.

[4] Steven J. Brams and Alan D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.

[5] Steven J. Brams, Michal A. Jones, and Christian Klamler. Better ways to cut a cake. *Notices of the American Mathematical Society*, 53(11):1314–1321, 2006.

[6] Simina Brânzei and Noam Nisan. Communication complexity of cake cutting. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 525–525. ACM, 2019.

[7] John Cloutier, Kathryn L Nyman, and Francis Edward Su. Two-player envy-free multi-cake division. *Mathematical Social Sciences*, 59(1):26–37, 2010.

[8] Ali Ghodsi, Matei Zaharia, Benjamin Hindman, Andy Konwinski, Scott Shenker, and Ion Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, NSDI'11, pages 323–336, Berkeley, CA, USA, 2011. USENIX Association.

[9] Avital Gutman and Noam Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pages 719–728. International Foundation for Autonomous Agents and Multiagent Systems, 2012.

[10] David Kurokawa, John K Lai, and Ariel D Procaccia. How to cut a cake before the party ends. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, pages 555–561. AAAI Press, 2013.

[11] Nicolas Lebert, Frédéric Meunier, and Quentin Carbonneaux. Envy-free two-player m-cake and three-player two-cake divisions. *Operations Research Letters*, 41(6):607–610, 2013.

[12] Hervé Moulin. *Fair division and collective welfare*. MIT press, 2004.

[13] Kathryn Nyman, Francis Edward Su, and Shira Zerbib. Fair division with multiple pieces. *Discrete Applied Mathematics*, 2020.

[14] David C Parkes, Ariel D Procaccia, and Nisarg Shah. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)*, 3(1):3, 2015.

[15] Ariel D. Procaccia and Hervé Moulin. *Cake Cutting Algorithms*, pages 311 —- 330. Cambridge University Press, 2016.

[16] Ariel D Procaccia. Cake cutting: not just child's play. *Communications of the ACM*, 56(7):78–87, 2013.

[17] Jack Robertson and William Webb. *Cake-cutting Algorithms: Be Fair if You Can*. A. K. P.s, Ltd, 1998.

[18] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16(1), 1948.

[19] Walter Stromquist. How to cut a cake fairly. *The American Mathematical Monthly*, 87(8):640–644, 1980.

[20] Francis E. Su. Rental harmony: Sperner's lemma in fair division. *The American Mathematical Monthly*, 106(10):930–942, 1999.