

Mind the Gap: Cake Cutting With Separation

Edith Elkind,¹ Erel Segal-Halevi,² Warut Suksompong³

¹ Department of Computer Science, University of Oxford

² Department of Computer Science, Ariel University

³ School of Computing, National University of Singapore

Abstract

We study the problem of fairly allocating a divisible resource, also known as cake cutting, with an additional requirement that the shares that different agents receive should be sufficiently separated from one another. This captures, for example, constraints arising from social distancing guidelines. While it is sometimes impossible to allocate a proportional share to every agent under the separation requirement, we show that the well-known criterion of maximin share fairness can always be attained. We then establish several computational properties of maximin share fairness—for instance, the maximin share of an agent cannot be computed exactly by any finite algorithm, but can be approximated with an arbitrarily small error. In addition, we consider the division of a pie (i.e., a circular cake) and show that an ordinal relaxation of maximin share fairness can be achieved.

1 Introduction

The end of the year is fast approaching, and members of a city council are busy planning the traditional New Year’s fair on their city’s main street. As usual, a major part of their work is to divide the space on the street among interested vendors. Each vendor naturally has a preference over potential locations, possibly depending on the proximity to certain attractions or the estimated number of customers visiting that space. Additionally, this year is different from previous years due to the social distancing guidelines issued by the government—vendors are required to be placed at least two meters apart. How should the city council allot the space so that all vendors feel fairly treated and at the same time everyone stays safe and sound under the new guidelines?

The problem of fairly allocating a heterogeneous divisible good among a set of agents in a fair manner has a long history and is commonly known as *cake cutting* (Brams and Taylor 1996; Robertson and Webb 1998; Procaccia 2016). A typical fairness criterion in cake cutting is *proportionality*, which means that each agent should receive her proportionally fair share—specifically, this amounts to $1/n$ of the agent’s value for the whole cake, where n denotes the total number of agents. For any set of agents with arbitrary valuations, a proportional allocation in which each agent receives a single connected piece is guaranteed to exist. Better still,

Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

such an allocation can be found by a simple and efficient algorithm (Dubins and Spanier 1961).

In this paper, we initiate the study of cake cutting with separation requirements. Besides the social distancing example that we mentioned, our setting captures the task of allocating machine processing time, where we need time to erase data from the previous process before the next process can be started, as well as land division, where we want space between different plots in order to avoid cross-fertilization. When separation is imposed, it is no longer the case that proportionality can always be satisfied—an extreme example is when all agents place their entire value on a small piece of length less than the minimum gap required. A similar failure of proportionality has notably been observed in the allocation of *indivisible* items (without separation), and a solution that has been proposed and widely studied in that context is *maximin share fairness* (Budish 2011; Kurokawa, Procaccia, and Wang 2018). Maximin share fairness requires each agent to receive her “maximin share”, which is the best share that the agent can secure by dividing the items into n bundles and getting the worst bundle. In this work, we demonstrate that maximin share fairness is an appropriate substitute for proportionality in cake cutting with separation, and analyze it from a computational perspective. To the best of our knowledge, this is the first use of maximin share fairness in connected cake cutting.

1.1 Our Results

As is commonly done in cake cutting, we assume that the cake is represented by an interval and each agent is to be allocated a single subinterval of the cake. We further require the pieces of any two agents to be separated by distance at least s , where $s > 0$ is a given separation parameter.

In Section 3, we begin by proving that an allocation that gives every agent at least her maximin share always exists, meaning that maximin share fairness can be guaranteed. Such an allocation can be found by a simple algorithm provided that the algorithm knows the maximin share of each agent. Unfortunately, we show that *no* finite algorithm can compute the maximin share of an agent exactly in the standard Robertson-Webb model—this impossibility holds even when $n = 2$ and the agents have piecewise constant valuations. Nevertheless, we design an algorithm that approximates the maximin share up to an arbitrarily small error,

Task	Cake Cutting	Pie Cutting
Decide whether $\text{MMS}_i = r$	Yes (Cor. 3.7)	No (Thm. 4.2)
Decide whether $\text{MMS}_i > r$	Yes (Thm. 3.6)	No (Thm. 4.2)
Decide whether $\text{MMS}_i \geq r$	Yes (Thm. 3.4)	No* (Thm. 4.5)
Compute the maximin share of an agent	No (Thm. 3.2)	No (Cor. 4.6)
Approximate the maximin share up to ε	Yes (Cor. 3.5)	Yes (Thm. 4.7)
Compute a maximin partition of an agent	No (Cor. 3.3)	No (Cor. 4.6)
Approximate a maximin partition up to ε	Yes (Cor. 3.5)	Yes (Thm. 4.7)

Table 1: Summary of the tasks that can and cannot be accomplished by finite algorithms in the Robertson-Webb model for cake cutting and pie cutting. All negative results hold even if the valuations of the agents are piecewise constant (but not given explicitly). The result with an asterisk holds when we do not allow the number of queries that the algorithm makes to depend on the separation parameter s .

which also allows us to compute an allocation wherein each agent obtains an arbitrarily close approximation of her maximin share. In addition, we present algorithms that decide whether the maximin share of an agent is greater than, less than, or equal to a given value, and show that if the agents have piecewise constant valuations that are given *explicitly* as part of the input, then we can exactly compute their maximin shares, and therefore a maximin allocation, in polynomial time.

In Section 4, we consider the allocation of a “pie”, which is a one-dimensional circular cake and serves to model, for example, the shoreline of an island or daily time slots for using a facility. In contrast to cake cutting, maximin share fairness cannot necessarily be guaranteed in pie cutting, and even the commonly studied cardinal multiplicative approximation cannot be obtained. Therefore, we focus instead on an *ordinal* relaxation of the maximin share, which allows each agent to partition the pie into k pieces for some parameter $k > n$. We show that when $k = n + 1$, the resulting notion—called the *1-out-of-(n+1) maximin share*—can be satisfied. We then investigate computational properties of maximin share fairness in pie cutting and demonstrate similarities and differences with cake cutting. For instance, while we can still approximate the maximin share of an agent, deciding whether the maximin share is equal to a given value is no longer possible for any finite algorithm.

A summary of our results can be found in Table 1.

1.2 Related Work

Cake cutting has long been studied by mathematicians and economists, and more recently attracted substantial interest from computer scientists, as it suggests a plethora of computational challenges. In particular, a long line of work in the artificial intelligence community in recent years has focused on cake cutting and its variants (Brânzei et al. 2016; Alijani et al. 2017; Bei et al. 2017; Menon and Larson 2017; Bei, Huzhang, and Suksompong 2018; Segal-Halevi 2018; Arunachaleswaran, Barman, and Rathi 2019; Goldberg, Hollender, and Suksompong 2020; Hosseini, Igarashi, and Searns 2020; Bei and Suksompong 2021).

In order to prevent agents from receiving a collection of

tiny pieces, it is often assumed that each agent must receive a connected piece of the cake (Dubins and Spanier 1961; Stromquist 1980; Su 1999; Bei et al. 2012; Cechlárová and Pillárová 2012; Cechlárová, Doboš, and Pillárová 2013; Aumann and Dombb 2015; Arunachaleswaran et al. 2019; Goldberg, Hollender, and Suksompong 2020). Indeed, when we divide resources such as time or space, non-connected pieces (e.g., disconnected time intervals or land plots) may be hard to use, or even totally useless. Note that we impose the connectivity constraint not only on the allocation but also in the definition of the maximin share benchmark. Similar conventions have been used in the context of indivisible items, where the items are vertices of an undirected graph and every agent must be allocated a connected subgraph (Bouveret et al. 2017; Lonc and Truszczynski 2020).¹

Most previous works on cake cutting did not explicitly consider the maximin share. This is because a proportional allocation is also a maximin allocation, since each agent’s maximin share is always at most $1/n$ of the agent’s value for the entire cake. In particular, without separation constraints, classic algorithms for proportional cake cutting (Steinhaus 1948; Dubins and Spanier 1961; Even and Paz 1984) attain maximin share fairness. The maximin share becomes interesting when a proportional allocation does not exist, for example, when the cake is a collection of disconnected intervals and each agent should receive a bounded number of connected pieces. Segal-Halevi (2021, Appendix B) briefly considered maximin share fairness in this setting.

2 Preliminaries

In cake cutting, the cake is represented by the interval $[0, 1]$. The set of agents is denoted by $N = [n]$, where $[k] := \{1, 2, \dots, k\}$ for any positive integer k . The preference of each agent i is represented by an integrable *density function* $f_i : [0, 1] \rightarrow \mathbb{R}_{\geq 0}$, which captures how the agent values different parts of the cake. A *piece of cake* is a finite union of disjoint intervals of the cake; it is said to be *connected* if it consists of a single interval. Agent i ’s value for

¹Bei et al. (2021) explored the relations between the constrained and unconstrained versions of the maximin share in that context.

a piece of cake X is given by $v_i(X) := \int_{x \in X} f_i(x) dx$. For $0 \leq x \leq y \leq 1$, we simplify notation and write $v_i(x, y) = v_i([x, y])$. As is standard in cake cutting, we assume that the density functions are normalized so that $v_i(0, 1) = 1$ for all $i \in N$. A valuation function is said to be *piecewise constant* if it is represented by a piecewise constant density function. An *allocation* of the cake is denoted by a vector $\mathbf{A} = (A_1, \dots, A_n)$, where each A_i is a piece of cake, and A_i and A_j are disjoint for all $i \neq j$. The piece A_i is allocated to agent i .

Let $s \in (0, \frac{1}{n-1})$ be the separation parameter. We will be interested in allocations that are *connected*—each A_i is a connected piece—and moreover have the property that any two pieces are separated by length at least s ; we call such allocations *s-separated*. We define *partitions* and *s-separated partitions* in a similar manner, with the difference being that for partitions, we have a set $\mathbf{P} = \{P_1, \dots, P_n\}$ instead of a vector $\mathbf{A} = (A_1, \dots, A_n)$. The *min-value* of partition P for agent i is defined as $\min_{j=1}^n v_i(P_j)$. Assume without loss of generality that the pieces P_1, \dots, P_n are in increasing order from left to right, and denote by $\Gamma_{n,s}$ the set that consists of all s -separated partitions of the cake. Note that an s -separated allocation or partition is incomplete since some of the cake necessarily remains unallocated. An *instance* consists of the agents, cake, density functions, and separation parameter.

A standard method for a cake-cutting algorithm to access agents' valuations is through queries in the model of Robertson and Webb (1998), which supports two types of queries:

- $\text{EVAL}_i(x, y)$: Asks agent i to evaluate the interval $[x, y]$ and return the value $v_i(x, y)$.
- $\text{CUT}_i(x, \alpha)$: Asks agent i to return the leftmost point y such that $v_i(x, y) = \alpha$, or state that no such point exists.

We now define the main fairness criterion of our paper.

Definition 2.1. The *maximin share* of agent i is defined as $\text{MMS}_i^{n,s} := \sup_{\mathbf{P} \in \Gamma_{n,s}} \min_{j \in [n]} v_i(P_j)$. When n and s are clear from context, we omit them from the notation and write MMS_i instead of $\text{MMS}_i^{n,s}$.

Let $\Gamma'_{n,s} \subseteq \Gamma_{n,s}$ be the set of s -separated partitions of $[x, y]$ such that every pair of consecutive pieces is separated by length exactly s . We claim that the definition of the maximin share can be simplified by replacing $\Gamma_{n,s}$ with $\Gamma'_{n,s}$; intuitively, for every partition in which the distance between some adjacent pieces is larger than s , there is a partition with at least the same min-value in which the distance between all adjacent pieces is exactly s . We also claim that the supremum in the definition can be replaced with a maximum, i.e., a maximizing partition always exists.

Proposition 2.2. *For every agent i , the following holds: $\text{MMS}_i^{n,s} = \max_{\mathbf{P} \in \Gamma'_{n,s}} \min_{j \in [n]} v_i(P_j)$.*

The proof of Proposition 2.2, as well as all other omitted proofs, can be found in the full version of our paper (Elkind, Segal-Halevi, and Suksompong 2020).

From now on, we will work with this new definition of the maximin share. An s -separated partition is said to be a *maximin partition* for agent i if every piece in the partition

yields value at least $\text{MMS}_i^{n,s}$. Proposition 2.2 implies that every agent has at least one maximin partition. Similarly, an s -separated allocation is said to be a *maximin allocation* if every agent i receives value at least $\text{MMS}_i^{n,s}$ from the allocation.

3 Cake Cutting

In this section, we consider cake cutting with separation, both in the Robertson-Webb query model and in a model where the agents' valuations are given explicitly.

3.1 Robertson-Webb Query Model

We begin by showing that the maximin share is an appropriate fairness criterion in our setting: it is always possible to fulfill this criterion for every agent using a quadratic number of queries in the Robertson-Webb model. Our algorithm is similar to the famous Dubins-Spanier protocol for finding proportional allocations when separation is not required (Dubins and Spanier 1961): we process the cake from left to right and, at each stage, allocate a subsequent piece of cake to an agent who demands the smallest piece.

Theorem 3.1. *For any instance, there exists a maximin allocation. Moreover, given the maximin share of each agent, such an allocation can be computed using $O(n^2)$ queries in the Robertson-Webb model.*

Proof. We ask each agent i to mark the leftmost point x_i such that $v(0, x_i) = \text{MMS}_i$. The agent who marks the leftmost x_i is allocated the piece $[0, x_i]$ (with ties broken arbitrarily); we then remove this agent along with the piece $[x_i, x_i + s]$, and recurse on the remaining agents and cake. If there is only one agent left, that agent receives all of the remaining cake. Since we make $n - j$ cut queries when there are $n - j$ agents left (and no eval queries), our algorithm uses $\sum_{j=0}^{n-1} (n - j) = O(n^2)$ queries.

We now prove the correctness of the algorithm. Consider any agent i and her maximin partition. If agent i receives the first piece allocated by the algorithm, she receives value MMS_i . Else, the allocated piece is no larger than the first piece of her maximin partition. Since the algorithm inserts a separator of length exactly s , the right endpoint of the first separator is either the same or to the left of the corresponding point in agent i 's maximin partition. Applying a similar argument repeatedly, we find that if agent i is not allocated any of the first $n - 1$ pieces, then the remaining cake contains the n th piece of the agent's partition. Hence, agent i receives a value of at least MMS_i in this case too. \square

The algorithm in Theorem 3.1 crucially relies on knowing the maximin share of each agent. Unfortunately, we show next that this knowledge is impossible to achieve in finite time, even if the valuations are piecewise constant but are not given explicitly as part of the input. Our result is similar in spirit to the non-finiteness results for connected envy-free cake cutting (Stromquist 2008), equitability (Cechlárová and Píllárová 2012; Procaccia and Wang 2017) and average-proportionality (Segal-Halevi and Nitzan 2019). However, all previous impossibility results relied on the fact that there

are two or more agents with possibly different valuations. In contrast, our impossibility result is attained even for a single agent who wants to cut the cake into two s -separated pieces.

Theorem 3.2. *There is no algorithm that can always compute the maximin share of an agent by asking the agent a finite number of Robertson-Webb queries. This holds even when $n = 2$ and the agent's valuation is piecewise constant and strictly positive (but is not given explicitly).*

Proof. Let v be the agent's valuation function and let $g(x) := v(0, x)$. By assumption the density function is strictly positive, so $g(x)$ is strictly increasing in the range $x \in [0, 1 - s]$. We will refer to this property as *monotonicity*. Similarly, $h(x) := v(x + s, 1) = 1 - g(x + s)$ is strictly decreasing in the same range. Since both functions are continuous and $g(0) = h(1 - s) = 0$, there exists a unique point $x_0 \in [0, 1 - s]$ such that $g(x_0)$ and $h(x_0)$ are equal; the value $g(x_0) = h(x_0)$ is the agent's maximin share. Equivalently, the maximin share is the unique value $g(x_0)$ such that $g(x_0) + g(x_0 + s) = 1$.

For simplicity we assume that the algorithm only asks queries of the form $\text{EVAL}(0, x)$ —which should return the value of $g(x)$ —and $\text{CUT}(0, \alpha)$ —which should return the value of $g^{-1}(\alpha)$. Every standard query can be implemented using two such simplified queries, so this assumption changes the total number of queries by a factor of at most 2.

During the run, there is always a finite set of points $x \in [0, 1]$ for which the algorithm knows the value of $g(x)$; we say that such points are *recorded*. Initially only points 0 and 1 are recorded. Given a point $x \in [0, 1]$, we denote by x_- the largest recorded point that is at most x , and by x_+ the smallest recorded point that is at least x (if x itself is recorded, then $x_- = x_+ = x$).

Assume for contradiction that there exists an algorithm for finding the maximin share using finitely many queries. We will show how an adversarial agent can answer the queries made by the algorithm so that after any finite number of queries, for any value that the algorithm may answer as the agent's maximin share, there exists a piecewise constant valuation function consistent with the answers for which the maximin share is different from the algorithm's answer. This is sufficient to obtain the desired contradiction.

When asked $\text{EVAL}(0, x)$, if x is recorded then the adversary replies $g(x)$; else, the adversary chooses a value for $g(x)$ satisfying the following properties:

- (i) Monotonicity is preserved, i.e., $g(x_-) < g(x) < g(x_+)$.
- (ii) If the point $x + s$ is recorded, then $g(x) \neq 1 - g(x + s)$.
- (iii) If the point $x - s$ is recorded, then $g(x) \neq 1 - g(x - s)$.

Since condition (i) allows infinitely many values to choose from, and each of the conditions (ii) and (iii) rules out at most one value, the adversary can make a choice satisfying these conditions.

When asked $\text{CUT}(0, \alpha)$, if there is a recorded point y such that $g(y) = \alpha$ then the adversary replies y ; else, the adversary chooses a point y satisfying the following properties:

- (i) Monotonicity is preserved, i.e., $g(y_-) < \alpha < g(y_+)$.

- (ii) None of the points $y - s$, $y + s$, and y is recorded.

Again, the former condition allows infinitely many points, and the latter forbids only a finite number of them.

Since the algorithm is finite, it must eventually return some number $r \in [0, 1]$ as the agent's maximin share. Now, in order to falsify the algorithm's answer, the adversary voluntarily answers two more queries by the same rules as above: a $\text{CUT}(0, r)$ query, returning a point y_r such that $g(y_r) = r$, and an $\text{EVAL}(0, y_r + s)$ query. Now both y_r and $y_r + s$ are recorded. The answering rules guarantee that $g(y_r) + g(y_r + s) \neq 1$, and therefore $g(y_r) = r$ cannot be the correct MMS value.²

Finally, to complete the valuation function, the adversary simply makes the density function between any two consecutive recorded points uniform. Hence the valuation is piecewise constant, but the corresponding maximin share is different from r , completing the proof. \square

Given a maximin partition of an agent, we can compute the agent's maximin share by simply taking the minimum among the agent's values for the pieces in the partition. Moreover, for two agents with identical valuations, an allocation in which each agent receives at least their (common) maximin share corresponds to a maximin partition for the common valuation. Theorem 3.2 therefore yields the following corollary, which also implies that an allocation whose existence is guaranteed by Theorem 3.1 cannot be computed without the knowledge of the agents' maximin shares.

Corollary 3.3. *There is no finite algorithm in the Robertson-Webb model that can always (a) compute a maximin partition of a single agent, or (b) compute a maximin allocation for n agents. This holds even when $n = 2$ and the agents' valuations are piecewise constant (but not given explicitly).*

Despite these negative results, we show next that it is possible to get an arbitrarily good approximation of the maximin share, partition, and allocation.

Theorem 3.4. *Given an agent i and a number $r > 0$, it is possible to decide whether $\text{MMS}_i \geq r$ (and, if so, compute a partition with value at least r for agent i) using at most n queries in the Robertson-Webb model.*

Proof. The idea is similar to that in Theorem 3.1. Ask the agent to mark the leftmost point x such that $v(0, x) = r$, make $[0, x]$ one of the pieces in a potential partition, add a separator $[x, x + s]$, and repeat starting from $x + s$. If there is still value at least r left after $n - 1$ iterations, answer yes; else, answer no. It is clear that at most n queries are used.

If the algorithm answers yes, then it finds a partition with value at least r , so $\text{MMS}_i \geq r$. Conversely, suppose that $\text{MMS}_i \geq r$, and consider a maximin partition. The right endpoint of the first piece in this partition is either the same or to the right of our first marked point x . In addition, since our algorithm inserts a separator of length exactly s , the right endpoint of our first separator is no further to the right than

²If $y_r + s > 1$ then the adversary cannot ask $\text{EVAL}(0, y_r + s)$, but then $g(y_r)$ cannot be the correct MMS value anyway, since making a cut at y_r would make the rightmost piece empty.

the corresponding point in the maximin partition. Applying a similar argument $n - 1$ times, we find that the right endpoint of our $(n - 1)$ st separator is no further to the right than the corresponding point in the maximin partition. Since the final piece of the partition has value at least MMS_i , our remaining piece also has value at least $\text{MMS}_i \geq r$. Hence the algorithm answers yes, as claimed. \square

Combining the algorithm in Theorem 3.4 with binary search allows us to approximate the maximin share.

Corollary 3.5. *Given an agent i and a number $\varepsilon > 0$, it is possible to find a number r for which $\text{MMS}_i - \varepsilon \leq r \leq \text{MMS}_i$ (together with a partition with min-value at least r for agent i) using $O(n \log(1/\varepsilon))$ Robertson-Webb queries.*

If instead of the exact MMS_i , we are given a number $r_i \leq \text{MMS}_i$ for each agent i , the algorithm for computing a maximin allocation in Theorem 3.1 still computes an allocation in which agent i receives value at least r_i ; the proof is essentially the same as before. It therefore follows from Corollary 3.5 that for any $\varepsilon > 0$, we can compute an allocation in which agent i receives value at least $\text{MMS}_i - \varepsilon$ using at most $O(n^2 + n \log(1/\varepsilon))$ queries.

Next, we consider the question of deciding whether the maximin share of an agent is *strictly greater than* a given number r . At first glance, it may seem that to this end, we can simply run the algorithm from Theorem 3.4 and answer yes exactly when the value left after $n - 1$ iterations is strictly greater than r . While this modification indeed works if the density function is positive throughout the cake, it may fail when intervals with zero value are present. Concretely, suppose that $v_i(0, 1/3) = 0.4$, $v_i(1/3, 2/3) = 0$, $v_i(2/3, 1) = 0.6$, and the value is distributed uniformly within each interval. Moreover, $s = 1/3$ and we want to decide whether $\text{MMS}_i^{2,s} > 0.4$. Even though there is leftover value at the right end of the cake when we run the algorithm, which may tempt us to believe that the maximin share can go above 0.4, the zero-valued middle part in fact renders this belief false.

This example suggests a simple modification to the algorithm from Theorem 3.4: instead of marking the leftmost point such that the value of the resulting piece is r , we should mark the *rightmost* point with this property. Indeed, we have $\text{MMS}_i > r$ if and only if after executing this modified algorithm we are left with a positive-value piece. In particular, in the first iteration of the algorithm, we want to mark the rightmost point x such that $v(0, x) = r$. Unfortunately, and perhaps surprisingly, this task *cannot* be done using the queries available via the Robertson-Webb model.³ Nevertheless, we can get around this issue by instead going over the cake from right to left. Then in the first iteration we want the leftmost point x such that $v(x, 1) = r$. This is equivalent to finding the leftmost point x such that $v(0, x) = 1 - r$, which can be done with a standard cut query.

Theorem 3.6. *Given an agent i and a number $r > 0$, it is possible to decide whether $\text{MMS}_i > r$ using at most n queries in the Robertson-Webb model.*

³See the full version of our paper for details (Elkind, Segal-Halevi, and Suksompong 2020).

Theorems 3.4 and 3.6 immediately imply the following:

Corollary 3.7. *Given an agent i and a number $r > 0$, it is possible to decide whether $\text{MMS}_i = r$ using at most $2n$ queries in the Robertson-Webb model.*

3.2 Explicit Piecewise Constant Valuations

Now, suppose that all agents have piecewise constant valuations and, moreover, these valuations are given *explicitly*. That is, for an agent $i \in N$ we are given a list of breakpoints (p_0, p_1, \dots, p_s) with $p_0 = 0$, $p_s = 1$, and a list of densities $(\gamma_1, \dots, \gamma_s)$, so that for each $j \in [s]$ and each $x \in [p_{j-1}, p_j]$ the valuation function v_i of agent i satisfies $v_i(0, x) = \sum_{\ell=1}^{j-1} \gamma_\ell (p_\ell - p_{\ell-1}) + \gamma_j (x - p_{j-1})$; moreover, all breakpoints and densities are rational numbers, represented as fractions whose numerators and denominators are given in binary. Note that it is straightforward to implement both types of Robertson-Webb queries in this model, so every problem that can be solved in polynomial time in the Robertson-Webb model can also be solved in polynomial time in this model. It turns out that using the explicit representation offers additional benefits: we can compute the agents' maximin shares exactly rather than approximately. This is in contrast to Theorem 3.2, which shows that this task is impossible if the piecewise constant valuations are not given explicitly and have to be queried through the Robertson-Webb model.

Theorem 3.8. *Given an agent i with a piecewise constant valuation function given explicitly, we can compute $\text{MMS}_i^{n,s}$ in time polynomial in the size of the input.*

At a high level, the proof of Theorem 3.8 proceeds by formulating a linear program whose solution corresponds to MMS_i . The challenge is that in order to have a linear program that returns a correct answer, we need to find out the intervals to which each endpoint of a maximin partition belongs. To accomplish this, we proceed from left to right, determining the interval for one endpoint at a time.

Combined with Theorem 3.1, Theorem 3.8 implies that when agents have piecewise constant valuations given explicitly, we can compute a maximin allocation efficiently (cf. Corollary 3.3).

Corollary 3.9. *For agents with piecewise constant valuations given explicitly, a maximin allocation can be computed in time polynomial in the size of the input.*

4 Pie Cutting

In the canonical model of cake cutting the cake is assumed to be linear. By contrast, in this section we assume that it is circular. In other words, our resource is represented by the interval $[0, 1]$ with the two endpoints identified with each other. The respective division problem is known in the literature as *pie cutting*; its applications include dividing the shoreline of an island among its inhabitants and splitting a daily cycle for using a facility (Thomson 2007; Brams, Jones, and Klamler 2008; Barbanel, Brams, and Stromquist 2009).

The definitions of s -separated partitions and allocations can be readily adjusted to pie cutting—the only difference

is that, due to the circular structure, there are n separators in pie cutting rather than $n - 1$ (so we can assume that $s < 1/n$). Note that, since the pie is one-dimensional, distances are measured along the circumference of the pie. We denote by $\Pi_{n,s}$ the set of s -separated partitions with respect to the pie, and by $\Pi'_{n,s} \subseteq \Pi_{n,s}$ the subset of partitions for which every pair of consecutive pieces is separated by length exactly s . The maximin share can then be defined similarly to how it is defined in cake cutting (Definition 2.1); just as in Proposition 2.2, we can show that $\text{MMS}_i^{n,s} = \max_{\mathbf{P} \in \Pi'_{n,s}} \min_{j \in N} v_i(P_j)$.

However, in pie cutting, unlike in cake cutting, a maximin allocation does not necessarily exist. This is evident in the example in Figure 1, where $s > 1/4$ and $0 < \varepsilon < \min\{s - 1/4, 1/2 - s\}$, and Alice values pieces of length ε centered at the top and bottom of the pie while Bob values similar pieces on the left and right. Since $s < 1/2 - \varepsilon$, both agents have a maximin share of $1/2$. However, since the distance between any point in Alice's piece and any point in Bob's piece is at most $1/4 + \varepsilon < s$, no s -separated allocation gives both agents a positive value. Hence, not only is there no maximin allocation in this instance, but we cannot even ensure any positive (multiplicative) approximation to the maximin share.




Figure 1: Example of a pie cutting instance with no maximin allocation. Each of the two agents uniformly values the bold part of the pie, $s > 1/4$, and $0 < \varepsilon < \min\{s - 1/4, 1/2 - s\}$.

Given this negative result, it may seem unclear whether any meaningful fairness guarantee can be achieved in pie cutting with separation. However, it turns out that we can obtain positive results if, instead of using *cardinal* approximations of the maximin share, we relax the criterion in an *ordinal* manner. Specifically, when each agent computes her maximin share, we allow partitioning into k pieces, where k is a parameter greater than n —we refer to the resulting notion as the *1-out-of- k* maximin share and write $\text{MMS}_i^{k,s}$ or simply MMS_i^k for the share of agent i . Ordinal approximations were introduced by Budish (2011), who considered the case $k = n + 1$.⁴ It turns out that this relaxation is precisely what we need for pie cutting.

⁴One way to think about this relaxation is that we pretend that there are $k > n$ agents when computing the maximin share. 1-out-of- k maximin share is a special case of the ℓ -out-of- k maximin share notion introduced by Babaioff, Nisan, and Talgam-Cohen (2019) and further studied by Segal-Halevi (2020), which takes the ℓ pieces of minimum value in a partition into k pieces.

Theorem 4.1. *For any pie cutting instance with n agents, there exists an allocation in which every agent i receives a piece of value at least MMS_i^{n+1} . Moreover, given the 1-out-of- $(n + 1)$ maximin share of each agent, such an allocation can be computed using $O(n^2)$ queries in the Robertson-Webb model.*

The idea behind our algorithm is similar to that of the analogous result for cake cutting (Theorem 3.1). The difference is that, because of the circular structure, when we start proceeding over the pie from a certain point (which we take to be the point 0), we may destroy one of the pieces in each agent's partition. This is why we need $n + 1$ pieces in the partition rather than n .

Proof. We ask each agent i to mark the leftmost point x_i such that $v(0, x_i) = \text{MMS}_i^{n+1}$. The agent who marks the leftmost x_i is allocated the piece $[0, x_i]$ (with ties broken arbitrarily); we then remove this agent along with the piece $[x_i, x_i + s]$, and recurse on the remaining agents and pie. If there is only one agent left, we still allocate to that agent a piece worth MMS_i^{n+1} (as opposed to the entire remaining pie). Since we make $n - j$ cut queries when there are $n - j$ agents left (and no eval queries), our algorithm uses $\sum_{j=0}^{n-1} (n - j) = O(n^2)$ queries.

We now prove the correctness of the algorithm. Consider any agent i and her 1-out-of- $(n + 1)$ maximin partition. When we turn the pie into a cake by cutting at the point 0 (equivalently, the point 1), we may break one of the pieces in the partition. Nevertheless, at least n pieces remain intact. If agent i receives the first piece allocated by the algorithm, she receives value MMS_i^{n+1} . Else, the allocated piece is no larger than the first intact piece of her maximin partition. Since the algorithm inserts a separator of length exactly s , the right endpoint of the first separator is no further to the right than the left endpoint of the agent's second intact piece. Applying a similar argument repeatedly, we find that if agent i is not allocated any of the first $n - 1$ pieces, then after removing the $(n - 1)$ st piece and the following separator, the remaining cake contains the n th intact piece of agent i 's partition as well as a positive amount of the $(n + 1)$ st piece (which may be intact or not). In particular, after allocating a piece of value MMS_i^{n+1} to the agent, the separator between the n th and $(n + 1)$ st pieces still remains. This means that there is a gap of length at least s between the first and last pieces of our allocation, and therefore the allocation is s -separated. \square

Theorem 4.1 can be generalized to guarantee each agent her k -out-of- $(kn + 1)$ maximin share, for any integer $k \geq 1$. As an example where this can be useful, suppose $s = 1/6$ and $n = 2$, and consider an agent who values the regions $[0, 1/30]$, $[6/30, 7/30]$, $[12/30, 13/30]$, $[18/30, 19/30]$, $[24/30, 25/30]$ uniformly at $1/5$ each (and has no value for the remaining pie). Then her 2-out-of-5 MMS is $2/5$, which is higher than her 1-out-of-3 MMS. On the other hand, if she values the regions $[0, 1/6]$, $[2/6, 3/6]$, $[4/6, 5/6]$ uniformly at $1/3$ each, then her 1-out-of-3 MMS is $1/3$, which is higher than her 2-out-of-5 MMS. Interestingly, the following algorithm allows each agent i to get the

k_i -out-of- $(k_i n + 1)$ MMS for her optimal k_i : Reduce the pie into a cake by breaking it in an arbitrary point (say, the point 0). This destroys, for each agent i , at most a single part in her $(k_i n + 1)$ -maximin partition. Thus, at least $k_i n$ parts (with their adjacent separators) remain intact. A procedure similar to the one described in Theorem 4.1 then guarantees each agent at least $k_i n / n = k_i$ of these parts.

For cake cutting, there exists an algorithm that, given an agent i and a number r , decides whether $\text{MMS}_i > r$ and whether $\text{MMS}_i = r$ (Theorem 3.6 and Corollary 3.7). In contrast, for pie cutting this is not the case.

Theorem 4.2. *Fix any $k \geq 2$. For pie cutting, there is no finite algorithm in the Robertson-Webb model that can decide, for any agent i and real number r , whether $\text{MMS}_i^k > r$ or whether $\text{MMS}_i^k = r$, even when the valuation of this agent is piecewise constant (but not given explicitly).*

Theorem 4.2 leaves open the question of whether it is possible to decide whether $\text{MMS}_i^k \geq r$ for a given r .

Open question 4.3. *For pie cutting, does there exist a finite algorithm in the Robertson-Webb model that, given an integer $k \geq 2$ and a real number $r > 0$, decides whether $\text{MMS}_i^k \geq r$?*

So far we have found two partial answers. First, we have a positive result for the special case $r = 1/k$. Note that we always have $\text{MMS}_i^k \leq 1/k$, and, moreover, $\text{MMS}_i^k = 1/k$ only if there is a partition where each separator has value 0. These observations turn out to be very useful for the analysis of the case $r = 1/k$.

Theorem 4.4. *For pie cutting, there exists an algorithm that, given an agent i and any $k \geq 2$, decides whether $\text{MMS}_i^k \geq 1/k$ (and if so, computes a maximin partition) using $O(k/s)$ queries in the Robertson-Webb model.*

The number of queries made by the algorithm in Theorem 4.4 scales linearly with $1/s$. This is in contrast to Theorem 3.4 for cake cutting, where the number of queries is independent of s . We show that for pie cutting the number of queries must depend on s ; this result holds even for $k = 2$ and $r = 1/k$.

Theorem 4.5. *Let c be any constant not depending on s . For pie cutting, there is no algorithm using at most c Robertson-Webb queries that, given an agent i , can always decide whether $\text{MMS}_i^2 \geq 1/2$, even when the agent's valuation is piecewise constant (but not given explicitly).*

We now turn to the problem of computing the maximin share and a maximin partition of a pie. Theorem 4.2 obviously rules out the possibility of exact computation:

Corollary 4.6. *Fix any $k \geq 2$. For pie cutting, there is no finite algorithm in the Robertson-Webb model that, given an agent i , can either (a) compute MMS_i^k , or (b) compute a maximin partition into k pieces for i . This holds even when the valuation of this agent is piecewise constant (but not given explicitly).*

Despite these negative results, we show that it is possible to approximate the maximin share of an agent up to an arbitrary error. The idea is to mark points on the pie so that any

piece between two adjacent marks has value at most $\varepsilon/2$, and, for each piece between two (not necessarily adjacent) marks, try to construct an s -separated partition with min-value equal to the value of this piece, by means of a greedy algorithm.

Theorem 4.7. *Fix any $k \geq 2$. For pie cutting, given an agent i and a number $\varepsilon > 0$, it is possible to find a number r such that $\text{MMS}_i^k - \varepsilon \leq r \leq \text{MMS}_i^k$, along with an s -separated partition with min-value r , using $O(1/\varepsilon)$ queries in the Robertson-Webb model.*

An interesting question is whether it is possible to make the dependence on $1/\varepsilon$ logarithmic instead of linear, as is possible for cake cutting (Corollary 3.5).

Open question 4.8. *In pie cutting, is it possible to compute an ε -approximation of MMS_i^k using $O(\text{poly}(k, \log(1/\varepsilon)))$ Robertson-Webb queries?*

5 Conclusion and Future Work

In this paper, we have initiated the study of cake cutting under separation requirements, and established several existence and computational results on maximin share fairness in this setting. Even though the cake in cake cutting is typically represented by an interval, certain applications of divisible resource allocation may require different representations. Indeed, this is the motivation behind the model of pie cutting that we have addressed in Section 4. Other recent works have studied models capturing the division of land (Segal-Halevi et al. 2017; Segal-Halevi, Hassidim, and Aumann 2020) as well as road networks (Bei and Suksompong 2021); one could investigate the effects of separation in these settings.

While the canonical maximin share is a reasonable fairness requirement when agents have equal entitlements to the resource, in certain situations the agents may be endowed with different entitlements (Chakraborty et al. 2020; Cseh and Fleiner 2020). Various extensions of the maximin share have been proposed (Aziz, Chan, and Li 2019; Babaioff, Nisan, and Talgam-Cohen 2019; Farhadi et al. 2019), and it may be interesting to study them in the context of cake cutting with separation. For different entitlements, a connected cake allocation may not exist even without separation (Segal-Halevi 2019; Crew, Narayanan, and Spirkl 2020).

Another avenue for future work is to consider other fairness criteria in light of separation constraints. Besides proportionality and maximin share fairness, two important criteria are *envy-freeness*—every agent believes that her piece is at least as valuable as any other allocated piece—and *equitability*—each agent receives the same value for his own piece. While these notions can be trivially satisfied by not allocating any of the cake, it would be interesting to determine whether they can be fulfilled subject to natural desiderata such as allocating the maximum possible amount of cake.

At a higher level, separation requirements represent one type of constraints that arise in a number of applications of cake cutting. Examining the interplay between such constraints and fairness considerations is an important direction that will likely lead to fruitful research.

Acknowledgments

This work was partially supported by the European Research Council (ERC) under grant number 639945 (ACCORD), by the Israel Science Foundation under grant number 712/20, and by an NUS Start-up Grant. We would like to thank the anonymous reviewers for their valuable comments.

References

- Alijani, R.; Farhadi, M.; Ghodsi, M.; Seddighin, M.; and Tajik, A. S. 2017. Envy-free mechanisms with minimum number of cuts. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 312–318.
- Arunachaleswaran, E. R.; Barman, S.; Kumar, R.; and Rathi, N. 2019. Fair and efficient cake division with connected pieces. In *Proceedings of the 15th Conference on Web and Internet Economics (WINE)*, 57–70.
- Arunachaleswaran, E. R.; Barman, S.; and Rathi, N. 2019. Fair division with a secretive agent. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, 1732–1739.
- Aumann, Y.; and Dombb, Y. 2015. The efficiency of fair division with connected pieces. *ACM Transactions on Economics and Computation* 3(4): 23:1–23:16.
- Aziz, H.; Chan, H.; and Li, B. 2019. Weighted maxmin fair share allocation of indivisible chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 46–52.
- Babaioff, M.; Nisan, N.; and Talgam-Cohen, I. 2019. Fair allocation through competitive equilibrium from generic incomes. In *Proceedings of the 2nd ACM Conference on Fairness, Accountability, and Transparency (ACM FAT*)*, 180.
- Barbanel, J. B.; Brams, S. J.; and Stromquist, W. 2009. Cutting a pie is not a piece of cake. *American Mathematical Monthly* 116(6): 496–514.
- Bei, X.; Chen, N.; Hua, X.; Tao, B.; and Yang, E. 2012. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI)*, 1263–1269.
- Bei, X.; Chen, N.; Huzhang, G.; Tao, B.; and Wu, J. 2017. Cake cutting: Envy and truth. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 3625–3631.
- Bei, X.; Huzhang, G.; and Suksompong, W. 2018. Truthful fair division without free disposal. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 63–69.
- Bei, X.; Igarashi, A.; Lu, X.; and Suksompong, W. 2021. The price of connectivity in fair division. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- Bei, X.; and Suksompong, W. 2021. Dividing a graphical cake. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*. Forthcoming.
- Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 135–141.
- Brams, S. J.; Jones, M. A.; and Klamler, C. 2008. Proportional pie-cutting. *International Journal of Game Theory* 36(3–4): 353–367.
- Brams, S. J.; and Taylor, A. D. 1996. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press.
- Brânzei, S.; Caragiannis, I.; Kurokawa, D.; and Procaccia, A. D. 2016. An algorithmic framework for strategic fair division. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 418–424.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy* 119(6): 1061–1103.
- Cechlárová, K.; Doboš, J.; and Pillárová, E. 2013. On the existence of equitable cake divisions. *Information Sciences* 228: 239–245.
- Cechlárová, K.; and Pillárová, E. 2012. On the computability of equitable divisions. *Discrete Optimization* 9(4): 249–257.
- Chakraborty, M.; Igarashi, A.; Suksompong, W.; and Zick, Y. 2020. Weighted envy-freeness in indivisible item allocation. In *Proceedings of the 19th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 231–239.
- Crew, L.; Narayanan, B.; and Spirkl, S. 2020. Disproportionate division. *Bulletin of the London Mathematical Society* 52(5): 885–890.
- Cseh, A.; and Fleiner, T. 2020. The complexity of cake cutting with unequal shares. *ACM Transactions on Algorithms* 16(3): 29:1–29:21.
- Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *American Mathematical Monthly* 68(1): 1–17.
- Elkind, E.; Segal-Halevi, E.; and Suksompong, W. 2020. Mind the gap: Cake cutting with separation. *arXiv preprint arXiv:2012.06682*.
- Even, S.; and Paz, A. 1984. A note on cake cutting. *Discrete Applied Mathematics* 7(3): 285–296.
- Farhadi, A.; Ghodsi, M.; Hajiaghayi, M. T.; Lahaie, S.; Pennock, D.; Seddighin, M.; Seddighin, S.; and Yami, H. 2019. Fair allocation of indivisible goods to asymmetric agents. *Journal of Artificial Intelligence Research* 64: 1–20.
- Goldberg, P. W.; Hollender, A.; and Suksompong, W. 2020. Contiguous cake cutting: Hardness results and approximation algorithms. *Journal of Artificial Intelligence Research* 69: 109–141.
- Hosseini, H.; Igarashi, A.; and Searns, A. 2020. Fair division of time: Multi-layered cake cutting. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 182–188.

- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2018. Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM* 64(2): 8:1–8:27.
- Lonc, Z.; and Truszczynski, M. 2020. Maximin share allocations on cycles. *Journal of Artificial Intelligence Research* 69: 613–655.
- Menon, V.; and Larson, K. 2017. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 352–358.
- Procaccia, A. D. 2016. Cake cutting algorithms. In Brandt, F.; Conitzer, V.; Endriss, U.; Lang, J.; and Procaccia, A. D., eds., *Handbook of Computational Social Choice*, chapter 13, 311–329. Cambridge University Press.
- Procaccia, A. D.; and Wang, J. 2017. A lower bound for equitable cake cutting. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 479–495.
- Robertson, J.; and Webb, W. 1998. *Cake-Cutting Algorithms: Be Fair if You Can*. Peters/CRC Press.
- Segal-Halevi, E. 2018. Redividing the cake. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 498–504.
- Segal-Halevi, E. 2019. Cake-cutting with different entitlements: How many cuts are needed? *Journal of Mathematical Analysis and Applications* 480(1): 123382.
- Segal-Halevi, E. 2020. Competitive equilibrium for almost all incomes: Existence and fairness. *Autonomous Agents and Multi-Agent Systems* 34(1): 26:1–26:50.
- Segal-Halevi, E. 2021. Fair multi-cake cutting. *Discrete Applied Mathematics* 291: 15–35.
- Segal-Halevi, E.; Hassidim, A.; and Aumann, Y. 2020. Envy-free division of land. *Mathematics of Operations Research*. 45(3): 896–922.
- Segal-Halevi, E.; and Nitzan, S. 2019. Fair cake-cutting among families. *Social Choice and Welfare* 53(4): 709–740.
- Segal-Halevi, E.; Nitzan, S.; Hassidim, A.; and Aumann, Y. 2017. Fair and square: Cake-cutting in two dimensions. *Journal of Mathematical Economics* 70(8): 1–28.
- Steinhaus, H. 1948. The problem of fair division. *Econometrica* 16(1): 101–104.
- Stromquist, W. 1980. How to cut a cake fairly. *American Mathematical Monthly* 87(8): 640–644.
- Stromquist, W. 2008. Envy-free cake divisions cannot be found by finite protocols. *Electronic Journal of Combinatorics* 15: #R11.
- Su, F. E. 1999. Rental harmony: Sperner’s lemma in fair division. *American Mathematical Monthly* 106(10): 930–942.
- Thomson, W. 2007. Children crying at birthday parties. Why? *Economic Theory* 31(3): 501–521.