# The Complexity of Cake Cutting with Unequal Shares 

ÁGNES CSEH, Institute of Economics, Centre for Economic and Regional Studies, Hungary and Hasso Plattner Institute, University of Potsdam, Potsdam, Germany<br>TAMÁS FLEINER, Department of Computer Science and Information Theory, Budapest University of Technology and Economics, Hungary and Institute of Economics, Centre for Economic and Regional Studies, Hungary


#### Abstract

An unceasing problem of our prevailing society is the fair division of goods. The problem of proportional cake cutting focuses on dividing a heterogeneous and divisible resource, the cake, among $n$ players who value pieces according to their own measure function. The goal is to assign each player a not necessarily connected part of the cake that the player evaluates at least as much as her proportional share.

In this article, we investigate the problem of proportional division with unequal shares, where each player is entitled to receive a predetermined portion of the cake. Our main contribution is threefold. First we present a protocol for integer demands, which delivers a proportional solution in fewer queries than all known protocols. By giving a matching lower bound, we then show that our protocol is asymptotically the fastest possible. Finally, we turn to irrational demands and solve the proportional cake cutting problem by reducing it to the same problem with integer demands only. All results remain valid in a highly general cake cutting model, which can be of independent interest.


CCS Concepts: • Mathematics of computing $\rightarrow$ Combinatorial algorithms;
Additional Key Words and Phrases: Cake cutting, fair division, unequal shares, proportional division

## ACM Reference format:

Ágnes Cseh and Tamás Fleiner. 2020. The Complexity of Cake Cutting with Unequal Shares. ACM Trans. Algorithms 16, 3, Article 29 (May 2020), 21 pages.
https://doi.org/10.1145/3380742

## 1 INTRODUCTION

In cake cutting problems, the cake symbolizes a heterogeneous and divisible resource that shall be distributed among $n$ players. Each player has her own measure function, which determines the value of any part of the cake for her. The aim of proportional cake cutting is to allocate each player a piece that is worth at least as much as her proportional share, evaluated with her measure function [35]. The measure functions are not known to the protocol.

[^0]The efficiency of a fair division protocol can be measured by the number of queries. In the standard Robertson-Webb model [30], two kinds of queries are allowed. The first one is the cut query, in which a player is asked to mark the cake at a distance from a given starting point so that the piece between these two is worth a given value to her. The second one is the eval query, in which a player is asked to evaluate a given piece according to her measure function.

If shares are meant to be equal for all players, then the proportional share is defined as $\frac{1}{n}$ of the whole cake. In the unequal shares version of the problem (also called cake cutting with entitlements), proportional share is defined as a player-specific demand, summing up to the value of the cake over all players. The aim of this article is to determine the query complexity of proportional cake cutting in the case of unequal shares. Robertson and Webb [30] write in their seminal book "Nothing approaching general theory of optimal number of cuts for unequal shares division has been given to date. This problem may prove to be very difficult." Instead of the number of physical cuts, we now settle the issue for the number of queries, which is the standard measure of efficiency for cake cutting protocols.

### 1.1 Related Work

Equal shares. Possibly the most famous cake cutting protocol belongs to the class of Divide and Conquer algorithms. Cut and Choose is a two-player equal-shares protocol that guarantees proportional shares. It already appeared in the Old Testament, where Abraham divided Canaan into two equally valuable parts and his nephew Lot chose for himself the one he valued more. The first $n$-player variant of this protocol is attributed to Banach and Knaster [35] and it requires $O\left(n^{2}\right)$ cut and eval queries. Other methods include the continuous (but discretizable) Dubins-Spanier protocol [15] and the Even-Paz protocol [18]. The latter authors show that their method requires $O(n \log n)$ queries at most. The complexity of proportional cake cutting has been studied in a setting where a circle instead of an interval represents the cake [3, 6], and also for higher dimensional cakes, where cuts are tailored to fit the shape of the cake [4, 22, 23, 33].

Unequal shares. The problem of proportional cake cutting with unequal shares is first mentioned by Steinhaus [35]. Motivated by dividing a leftover cake, Robertson and Webb [30] define the problem formally and offer a range of solutions for two players. More precisely, they list cloning players, using Ramsey partitions [27] and most importantly, the Cut Near-Halves protocol [30]. The last method computes a fair solution for two players with integer demands $d_{1}$ and $d_{2}$ in $2\left\lceil\log _{2}\left(d_{1}+d_{2}\right)\right\rceil$ queries. Robertson and Webb also show how any two-player protocol can be generalized to $n$ players in a recursive manner. The number of physical cuts Cut Near-Halves makes for two players can be beaten for certain demands, as Robertson and Webb [30] also note. For some demands, Carney [10] and Lohr [25] design such protocols utilizing number-theoretic approaches. Proportional allocation with unequal shares is also discussed in the context of indivisible items instead of a cake [1, 19].
Irrational demands. The case of irrational demands in the unequal shares case is interesting from the theoretical point of view, but beyond this, solving it might be necessary, because other protocols might generate instances with irrational demands. For example, in the maximumefficient envy-free allocation problem with two players and piecewise linear measure functions, any optimal solution must be specified using irrational numbers, as Cohler et al. [12] show. Barbanel [2] studies the case of cutting the cake in an irrational ratio between $n$ players and presents a protocol that constructs a proportional division. Shishido and Zeng [34] solve the same problem with the objective of minimizing the number of resulting pieces. Their protocol is simpler than that of Barbanel [2].

Lower bounds. The drive toward establishing lower bounds on the complexity of cake cutting protocols is coeval with the cake cutting literature itself [35]. For proportional cake cutting with
equal shares, Even and Paz [18] conjectured that their protocol is the best possible, while Robertson and Webb explicitly write that "they would place their money against finding a substantial improvement on the $n \log _{2} n$ bound." After approximately 20 years of no breakthrough in the topic, Magdon-Ismail et al. [26] showed that any protocol must make $\Omega(n \log n)$ comparisons-but this was no bound on the number of queries. Essentially simultaneously, Woeginger and Sgall [36] came up with the lower bound $\Omega(n \log n)$ on the number of queries for the case where contiguous pieces are allocated to each player. Not much later, this condition was dropped by Edmonds and Pruhs [16] who completed the query complexity analysis of proportional cake cutting with equal shares by presenting a lower bound of $\Omega(n \log n)$. Brams et al. [7] study the minimum number of physical cuts in the case of unequal shares and proved that $n-1$ cuts might not suffice-in other words, they show that in some instances, no proportional allocation exists with contiguous pieces. Crew et al. [13] and Segal-Halevi [31] improve this lower bound and show that at least $2 n-2$ cuts may be necessary, and $3 n-4$ cuts are always sufficient. However, no lower bound on the number of queries has been known in the case of unequal shares.

Generalizations in higher dimensions. There are two sets of multiple-dimensional generalizations of the proportional cake cutting problem. The first group focuses on the existence of a proportional division, without any constructive proof. The existence can be shown easily using Lyapunov's theorem, as stated by Dubins and Spanier [15]. Berliant et al. [5] investigate the existence of envy-free divisions. Dall'Aglio [14] considers the case of equal shares and defines a dual optimization problem that allows one to compute a proportional solution by minimizing convex functions over a finite dimensional simplex. Complexity issues are not discussed in these papers; in fact, queries are not even mentioned in them.

The second group of multiple-dimensional generalizations considers problems where certain geometric parameters are imposed on the cake or the pieces (see Barbanel et al. [3], Beck [4], Brams et al. [6], Hill [22], Iyer and Huhns [23], Segal-Halevi et al. [33]). Also, some of these have special extra requirements on the output, such as contiguousness or envy-freeness. These works demonstrate the interest in various problems in multi-dimensional cake cutting, for which we define a very general framework.

### 1.2 Our Contribution

We provide formal definitions in Section 2 and present the query analysis of the fastest known protocol for the $n$-player proportional cake cutting problem with total demand $D \geq n$ in Section 3 . Readers familiar with proportional cake cutting might consider skipping these two sections. In Section 4, we focus on our protocol for the problem, which is our main contribution in this article. The idea is that we recursively render the players in two batches so that these batches can simulate two players who aim to cut the cake into two approximately equal halves. Our protocol requires only $2(n-1) \cdot\left\lceil\log _{2} D\right\rceil$ queries. Other known protocols reach $D \cdot\left\lceil\log _{2} D\right\rceil$ and $n(n-1) \cdot\left\lceil\log _{2} D\right\rceil$ queries, thus ours is the fastest procedure that derives a proportional division for the $n$-player cake cutting problem with unequal shares. Moreover, our protocol also works on a highly general cake (introduced in Section 5), extending the traditional notion of the cake to any finite dimension.

We complement our positive result by showing a lower bound of $\Omega(n \cdot \log D)$ on the query complexity of the problem in Section 6. Our proof generalizes, but does not rely on, the lower bound proof given by Edmonds and Pruhs [16] for the problem of proportional division with equal shares. Moreover, our lower bound remains valid in the generalized cake cutting and query model, allowing a considerably more powerful notion of a query even on the usual, $[0,1]$ interval cake.

In Section 7, we turn to irrational demands and solve the proportional cake cutting problem by reducing it to the same problem with integer demands only. By doing so, we provide a novel and simple approach to the problem. Moreover, our method works in the generalized query model as well.

## 2 PRELIMINARIES

We begin with formally defining our input. Our setting includes a set of players of cardinality $n$, denoted by $\left\{P_{1}, P_{2}, \ldots, P_{n}\right\}$, and a heterogeneous and divisible good, which we refer to as the cake and project to the unit interval $[0,1]$. Each player $P_{i}$ has a non-negative, absolutely continuous measure function $\mu_{i}$ that is defined on Lebesgue-measurable sets. We remark that absolute continuity implies that every zero-measure set has value 0 according to $\mu_{i}$ as well. In particular, $\mu_{i}((a, b))=\mu_{i}([a, b])$ for any interval $[a, b] \subseteq[0,1]$. Besides measure functions, each player $P_{i}$ has a demand $d_{i} \in \mathbb{Z}^{+}$, representing that $P_{i}$ is entitled to receive $\left.d_{i} / \sum_{j=1}^{n} d_{j} \in\right] 0,1[$ part of the whole cake. The value of the whole cake is identical for all players; in particular, it is the sum of all demands:

$$
\forall 1 \leq i \leq n \quad \mu_{i}([0,1])=D=\sum_{j=1}^{n} d_{j} .
$$

We remark that an equivalent formulation is also used sometimes, where the demands are rational numbers that sum up to 1 , the value of the full cake. Such an input can be transformed into the above form simply by multiplying all demands by the least common denominator of all demands. As opposed to this, if demands are allowed to be irrational numbers, then no ratio-preserving transformation might be able to transform them to integers. That is why the case of irrational demands is treated separately.

The cake $[0,1]$ will be partitioned into subintervals in the form $[x, y), 0 \leq x \leq y \leq 1$. A finite union of such subintervals forms a piece $X_{i}$ allocated to player $P_{i}$. We would like to stress that a piece is not necessarily connected.

Definition 2.1. A set $\left\{X_{i}\right\}_{1 \leq i \leq n}$ of pieces is a division of the cake $[0,1]$ if $\bigcup_{1 \leq i \leq n} X_{i}=[0,1]$ and $X_{i} \cap X_{j}=\emptyset$ for all $i \neq j$. We call division $\left\{X_{i}\right\}_{1 \leq i \leq n}$ proportional if $\mu_{i}\left(X_{i}\right) \geq d_{i}$ for all $1 \leq i \leq n$.

In words, proportionality means that each player receives a piece with which her demand is satisfied. We do not consider Pareto optimality or alternative fairness notions such as envy-freeness in this article.

We now turn to defining the measure of efficiency of a cake cutting protocol. We assume that $1 \leq i \leq n, x, y \in[0,1]$, and $0 \leq \alpha \leq 1$. Oddly enough, the Robertson-Webb query model was not formalized explicitly by Robertson and Webb first, but by Woeginger and Sgall [36], who attribute it to the earlier two. In their query model, a protocol can ask agents the following two types of queries.

- Cut query $\left(P_{i}, \alpha\right)$ returns the leftmost point $x$ so that $\mu_{i}([0, x])=\alpha$. In this operation, $x$ becomes a so-called cut point.
-Eval query $\left(P_{i}, x\right)$ returns $\mu_{i}([0, x])$. Here, $x$ must be a cut point.
Notice that this definition implies that choosing sides, sorting marks, or calculating any other parameter than the value of a piece are not counted as queries and thus they do not influence the efficiency of a protocol. Our protocols do not abuse the model by performing a large number of such operations. We also remark that the "leftmost point" criterion in the cut query can be omitted if the measure $\mu_{i}$ is not only absolutely continuous with respect to the Lebesgue measure, but it is equivalent to it-meaning that the Lebesgue measure is also absolutely continuous with respect to $\mu_{i}$.

Definition 2.2. The number of queries in a protocol is the number of eval and cut queries until termination. We denote the number of queries for an $n$-player protocol with total demand $D$ by $T(n, D)$.

The query definition of Woeginger and Sgall is the strictest of the type Robertson-Webb. We now outline three options to extend the notion of a query, all of which have been used in earlier papers $[16,18,30,36]$ and are also referred to as Robertson-Webb queries.
(1) The query definition of Edmonds and Pruhs. There is a slightly different and stronger formalization of the core idea, given by Edmonds and Pruhs [16] and also used by Procaccia [28, 29]. The crucial difference is that they allow both cut and eval queries to start from an arbitrary point in the cake.

- Cut query $\left(P_{i}, x, \alpha\right)$ returns the leftmost point $y$ so that $\mu_{i}([x, y])=\alpha$ or an error message if no such $y$ exists.
-Eval query $\left(P_{i}, x, y\right)$ returns $\mu_{i}([x, y])$.
These queries can be simulated as trivial concatenations of the queries defined by Woeginger and Sgall. To pin down the starting point $x$ of a cut query ( $P_{i}, x, \alpha$ ), we introduce the cut point $x$ with the help of a dummy player's Lebesgue-measure, ask $P_{i}$ to evaluate the piece $[0, x]$, and then we cut query with value $\alpha^{\prime}=\alpha+\mu_{i}([0, x])$. Similarly, to generate an eval query ( $P_{i}, x, y$ ), one only needs to artificially generate the two cut points $x$ and $y$ and then ask two eval queries of the Woeginger-Sgall model, $\left(P_{i}, x\right)$ and $\left(P_{i}, y\right)$. We remark that such a concatenation of Woeginger-Sgall queries reveals more information than the single query in the model of Edmonds and Pruhs.
(2) Proportional cut query. The term proportional cut query stands for $n$-player Cut queries of the sort " $P_{i}$ cuts the piece $[x, y]$ in ratio $a: b$," where $a, b$ are integers. As Woeginger and Sgall also note it, two eval queries and one cut query with ratio $\alpha=\frac{a}{a+b} \cdot \mu_{i}([x, y])$ are sufficient to execute such an operation if $x, y$ are cut points, otherwise five queries suffice. Notice that the eval queries are only used by $P_{i}$ when she calculates $\alpha$, and their output does not need to be revealed to any other player or even to the protocol.
(3) Reindexing. When working with recursive protocols, it is especially useful to be able to reindex a piece $[x, y]$ so that it represents the interval $[0,1]$ for $P_{i}$. Any further cut and eval query on $[x, y]$ can also be substituted by at most five queries on the whole cake. Similarly as above, there is no need to reveal the result of the necessary eval queries addressed to a player.

These workarounds ensure that protocols require asymptotically the same number of queries in both model formulations, even if reindexing and proportional queries are allowed. We opted for utilizing all three extensions of the Woeginger-Sgall query model in our upper bound proofs, because the least restrictive model allows the clearest proofs. Regarding our lower bound proof, it holds even if we allow a highly general query model including all of the above extensions, which we define in Section 5.2.

## 3 KNOWN PROTOCOLS

To provide a base for comparison, we sketch the known protocols for proportional cake cutting with unequal shares and bound their query complexity.

The most naive approach to the case of unequal shares is the cloning technique, where each player $P_{i}$ with demand $d_{i}$ is substituted by $d_{i}$ players with unit demands. In this way, a $D$-player equal shares cake cutting problem is generated, which can be solved in $O(D \log D)$ queries [18].

As Robertson and Webb [30] point out, any two-player protocol can be generalized to an $n$ player protocol. They list two, two-player protocols, Cut Near-Halves and the Ramsey Partition Algorithm [27], and also remark that for two players, Cut Near-Halves is always at least as efficient as the Ramsey Partition Algorithm. Therefore, we restrict ourselves to analyzing the complexity of the $n$-player Cut Near-Halves protocol.

Cut Near-Halves is a simple procedure, in which the cake of value $D$ is repeatedly cut in approximately half by players $P_{1}$ and $P_{2}$ with demands $d_{1} \leq d_{2}$ as follows. $P_{1}$ cuts the cake into two near-halves, more precisely, in ratio $\left\lfloor\frac{D}{2}\right\rfloor:\left\lceil\frac{D}{2}\right\rceil$. Then, $P_{2}$ picks a piece that she values at least as much as $P_{1}$. This piece is awarded to $P_{2}$ and her claim is reduced accordingly, by the respective near-half value of the cake. In the next round, the same is repeated on the remaining part of the cake, and so on, until $d_{1}$ or $d_{2}$ is reduced to zero. Notice that the cutter is always the player with the smaller current demand, and thus this role might be swapped from round to round.

The recursive $n$-player protocol of Robertson and Webb runs as follows. We assume that $k-1<$ $n$ players, $P_{1}, P_{2}, \ldots, P_{k-1}$, have already divided the whole cake of value $D=d_{1}+d_{2}+\ldots+d_{n}$ in ratio $d_{1}: d_{2}: \ldots: d_{k-1}$. The next player $P_{k}$ then challenges each of the first $k-1$ players separately to redistribute the piece already assigned to them. In these rounds, $P_{k}$ claims $\frac{d_{k}}{d_{1}+d_{2}+\ldots+d_{k-1}}$ part of each piece. This results in a division of the whole cake of value $D$ in ratio $d_{1}: d_{2}: \ldots: d_{k}$. This generates $k-1$ rounds of the Cut Near-Halves protocol, each with two players. Notice that this protocol tends to assign a highly fractured piece of cake to every player.

The following theorem summarizes the results known about the complexity of the two-player and $n$-player versions of the Cut Near-Halves protocol.

Theorem 3.1 (Robertson and Webb [30]). The two-player Cut Near-Halves protocol with demands $d_{1}, d_{2}$ requires $T(2)=2\left\lceil\log _{2}\left(d_{1}+d_{2}\right)\right\rceil$ queries at most. The recursive $n$-player version is finite.

Here we give an estimate for the number of queries of the recursive protocol.
Theorem 3.2. The number of queries in the recursive n-player Cut Near-Halves protocol is at most

$$
T(n, D)=\sum_{i=1}^{n-1}\left[2 i \cdot\left\lceil\log _{2}\left(\sum_{j=1}^{i+1} d_{j}\right)\right]\right] \leq n(n-1) \cdot\lceil\log D\rceil .
$$

Proof. The first round consists of players $P_{1}$ and $P_{2}$ sharing the cake using $2\left\lceil\log _{2}\left(d_{1}+d_{2}\right)\right\rceil$ queries. The second round then has two, two-player runs, each of them requiring $2\left\lceil\log _{2}\left(d_{1}+d_{2}+d_{3}\right)\right\rceil$ queries. In general, the $i$ th round terminates after $i \cdot 2\left\lceil\log _{2}\left(\sum_{j=1}^{i+1} d_{j}\right)\right\rceil$ queries at most. The number of rounds is $n-1$. Now we add up the total number of queries.

$$
\sum_{i=1}^{n-1}\left[2 i \cdot\left\lceil\log _{2}\left(\sum_{j=1}^{i+1} d_{j}\right)\right]\right] \leq \sum_{i=1}^{n-1}\left[2 i \cdot\left\lceil\log _{2}\left(\sum_{j=1}^{n} d_{j}\right)\right]\right]=\sum_{i=1}^{n-1}\left[2 i \cdot\left\lceil\log _{2} D\right\rceil\right]=n(n-1) \cdot\lceil\log D\rceil
$$

The following example proves that the calculated bound can indeed be reached asymptotically in instances with an arbitrary number of players.

Example 3.3. The estimation for the query number in Theorem 3.2 is asymptotically sharp if $\left\lceil\log _{2}\left(\sum_{j=1}^{i+1} d_{j}\right)\right]=\left\lceil\log _{2} D\right\rceil$ holds for at least a fixed portion of all $1 \leq i \leq n-1$, say, for the third of them. This is easy to reach if $n$ is a sufficiently large power of 2 and all but one player have demand 1, while there is another player with demand 2 . Notice that this holds for every order for the agents. If one sticks to a decreasing order of demand when indexing the players, then not only asymptotic, but also strict equality can be achieved by setting $d_{1}$ much larger than all other demands.

## 4 OUR PROTOCOL

In this section, we present a simple and elegant protocol that beats all three above mentioned protocols (cloning, Ramsey partitions, Cut Near-Halves) in query number. Our main idea is that we recursively render the players in two batches so that these batches can simulate two players


Fig. 1. The steps performed by our protocol on Example 4.1. The colored thick intervals are the pieces already allocated to the player in their label. The ratio players cut the current cake can be seen in the framed labels.
who aim to cut the cake into two approximately equal halves. For now, we work with the standard cake and query model defined in Section 2. Later, in Section 5.3, we will show how our protocol can be extended to a more general cake. We remind the reader that cutting near-halves means to cut in ratio $\left\lfloor\frac{D}{2}\right\rfloor:\left\lceil\frac{D}{2}\right\rceil$.

To ease the notation, we assume that the players are indexed so that when they mark the nearhalf of the cake, the marks appear in an increasing order from 1 to $n$. In the subsequent rounds, we reindex the players to keep this property intact. Based on these marks, we choose "the middle player" $P_{j}$, this being the player whose demand reaches the near-half of the cake when summing up the demands in the order of marks from left to right. This player cuts the cake and each player is ordered to the piece her own mark falls to. The middle player $P_{j}$ is cloned if necessary so that she can play on both pieces. The protocol is then repeated on both generated subinstances, with adjusted demands. In the subproblem, the players' demands are according to the ratios listed in the pseudocode. The base case of the recursion is a subproblem with one player only, in which case she is allocated the piece.

## Proportional division with unequal shares:

Each player marks the near-half of the cake $X$.
Sort the players according to their marks.
Calculate the smallest index $j$ such that $\left\lfloor\frac{D}{2}\right\rfloor \leq \sum_{i=1}^{j} d_{i}=: m$.
Cut the cake in two along $P_{j}$ 's mark.
Define two instances of the same problem and solve them recursively.
(1) Players $P_{1}, P_{2}, \ldots, P_{j}$ share piece $X_{1}$ on the left. Demands are set to $d_{1}, d_{2} \ldots, d_{j-1}, d_{j}-m+\left\lfloor\frac{D}{2}\right\rfloor$, while measure functions are set to $\mu_{i}$. $\left\lfloor\frac{D}{2}\right\rfloor / \mu_{i}\left(X_{1}\right)$, for all $1 \leq i \leq j$.
(2) Players $P_{j}, P_{j+1}, \ldots, P_{n}$ share piece $X_{2}=X \backslash X_{1}$ on the right. Demands are set to $m-\left\lfloor\frac{D}{2}\right\rfloor, d_{j+1}, d_{j+2}, \ldots, d_{n}$, while measure functions are set to $\mu_{i} \cdot\left\lceil\frac{D}{2}\right\rceil / \mu_{i}\left(X_{2}\right)$, for all $j \leq i \leq n$.

Example 4.1. We present our protocol on an example with $n=3$. Every step of the protocol is depicted in Figure 1 . Let $d_{1}=1, d_{2}=3, d_{3}=1$.

Row 1: Since $D=5$ is odd, all players mark the near-half of the cake in ratio 2:3. The cake is then cut at $P_{2}$ 's mark, since $d_{1}<\left\lfloor\frac{D}{2}\right\rfloor$, but $d_{1}+d_{2} \geq\left\lfloor\frac{D}{2}\right\rfloor$.
Row 2: The first subinstance will consist of players $P_{1}$ and $P_{2}$, both with demand 1 , whereas the second subinstance will have the second copy of player $P_{2}$ alongside $P_{3}$ with demands 2 and 1 , respectively. In the first instance, both players mark half of the cake and the one who marked it closer to 0 will receive the leftmost piece, while the other player is allocated the remaining piece. The players in the second instance mark the cake in ratio $1: 2$. Suppose that the player demanding more marks it closer to 0 . The leftmost piece is then allocated to her.
Row 3: The two players in the second subinstance play further: they share the remaining piece in ratio $1: 1$. The player with the mark on the left will be allocated the piece on the left, while the other players takes the remainder of the piece.
Row 4: The whole cake is divided proportionally.
These rounds require $3+2+2+2=9$ proportional cut queries and no eval query.
Theorem 4.2. Our "Protocol for proportional division with unequal shares" terminates with a proportional division.

Proof. We need to show that the base case with one player only in a generated subinstance is indeed always reached in a finite number of rounds. Observe that each subinstance has strictly less demand for its players together than its parent instance had. Since all demands are positive integers, after sufficiently many instance divisions, the subinstances must reduce to the base case instance.

We provide detailed calculations for the first subinstance only, because analogous calculations can easily be obtained for the second subinstance. First we observe that $d_{j}-m+\lfloor D / 2\rfloor=$ $\lfloor D / 2\rfloor-\sum_{1}^{j-1} d_{i}$ is positive by the definition of $j$. Now we have to ensure that the subinstance is generated in such a manner that all players evaluate the full cake $X_{1}$ of the first subinstance equally and to the sum of all their demands. In the case of the first subinstance, the sum of demands is $\sum_{i=1}^{j} d_{i}-m+\left\lfloor\frac{D}{2}\right\rfloor=\left\lfloor\frac{D}{2}\right\rfloor$. This needs to be the measure of the cake $X_{1}$ for all players. To achieve this, $\mu_{1}, \mu_{2}, \ldots, \mu_{j}$ need to be adjusted. Each $\mu_{i}$ will become in this subinstance

$$
\mu_{i 1}=\mu_{i} \cdot \frac{\left\lfloor\frac{D}{2}\right\rfloor}{\mu_{i}\left(X_{1}\right)} .
$$

If $P_{i}, 1 \leq i<j$, receives a piece of worth $d_{i}$ according to $\mu_{i 1}$, then in the original instance, it is of worth

$$
d_{i} \cdot \frac{\mu_{i}\left(X_{1}\right)}{\left\lfloor\frac{D}{2}\right\rfloor} \geq d_{i},
$$

because $\mu_{i}\left(X_{1}\right) \geq \mu_{j}\left(X_{1}\right)=\left\lfloor\frac{D}{2}\right\rfloor$, due to the cutting rule in our protocol. With this we have shown that every player appearing only in the first subinstance is guaranteed to gain her proportional share. An analogous proof works for players $P_{j+1}, P_{j+2}, \ldots, P_{n}$. The last step is to show that $P_{j}$ collects her proportional share from the two subinstances.

The only player whose measure function certainly need not be adjusted is $P_{j}$. It is because $\mu_{j}\left(X_{1}\right)=\left\lfloor\frac{D}{2}\right\rfloor$, thus $\mu_{j 1}=\mu_{j} \cdot \frac{\left\lfloor\frac{D}{2}\right\rfloor}{\mu_{j}\left(X_{1}\right)}=\mu_{j}$. Therefore, if $P_{j}$ receives her proportional share $d_{j}-m+$ $\left\lfloor\frac{D}{2}\right\rfloor$ and $m-\left\lfloor\frac{D}{2}\right\rfloor$ in the two subinstances, then in the original instance her piece is worth $d_{j}$ at least.

Having shown its correctness, we now present our estimation for the number of queries our protocol needs.

Theorem 4.3. For any $2 \leq n$ and $n<D$, the number of queries in our $n$-player protocol on a cake of total value $D$ is $T(n, D) \leq 2(n-1) \cdot\left\lceil\log _{2} D\right\rceil$.

Proof. If $n=2$, then our protocol simulates the Cut Near-Halves protocol-except that it uses cut queries exclusively-and according to Theorem 3.1 it requires $2\left\lceil\log _{2} D\right\rceil$ queries at most. This matches the formula stated in Theorem 4.3. From this we prove by induction. For $n>2$, the following recursion formula corresponds to our rules.

$$
T(n, D)=n+\max _{1 \leq i \leq n}\left\{T\left(i,\left\lfloor\frac{D}{2}\right\rfloor\right)+T\left(n-i+1,\left\lceil\frac{D}{2}\right\rceil\right)\right\}
$$

We now substitute our formula into the right side of this expression.

$$
\begin{aligned}
n+\max _{1 \leq i \leq n}\left\{T\left(i,\left\lfloor\frac{D}{2}\right\rfloor\right)+T\left(n-i+1,\left\lceil\frac{D}{2}\right\rceil\right)\right\} & = \\
n+\max _{1 \leq i \leq n}\left\{2(i-1)\left\lceil\log _{2}\left\lfloor\frac{D}{2}\right\rfloor\right\rceil+2(n-i)\left\lceil\log _{2}\left\lceil\frac{D}{2}\right\rceil\right\rceil\right\} & \leq(*) \\
n+\max _{1 \leq i \leq n}\left\{2(i-1)\left(\left\lceil\log _{2} D\right\rceil-1\right)+2(n-i)\left(\left\lceil\log _{2} D\right\rceil-1\right)\right\} & = \\
n+2(n-1)\left(\left\lceil\log _{2} D\right\rceil-1\right) & = \\
-n+2+2(n-1)\left\lceil\log _{2} D\right\rceil & \leq \\
2(n-1) \cdot\left\lceil\log _{2} D\right\rceil & =T(n, D) .
\end{aligned}
$$

The inequality marked by $(*)$ is trivially correct if $D$ is even. For odd $D$, we rely on the fact that $\log _{2} D$ cannot be an integer.

$$
\left\lceil\log _{2}\left\lfloor\frac{D}{2}\right\rfloor\right\rceil \leq\left\lceil\log _{2}\left\lceil\frac{D}{2}\right\rceil\right\rceil=\left\lceil\log _{2} \frac{D+1}{2}\right\rceil=\left\lceil\log _{2}(D+1)-\log _{2} 2\right\rceil=\left\lceil\log _{2}(D+1)\right\rceil-1=\left\lceil\log _{2} D\right\rceil-1
$$

With a query number of $O(n \log D)$, our protocol is more efficient than all known protocols. On the other hand, our protocol is not truthful if players maximize their utility instead of just caring about receiving their proportional share. Take the two-player equal shares case with nonzero measure functions on any nonzero measure interval. If the player whose mark is at the left knows the measure function of the other player, she can easily manipulate the outcome by marking the half of the cake just before the mark of the other player. As a result, her piece will be larger than what she receives if she reports the truth, unless their measure functions are special. This is not a shortcoming of our protocol, but rather of the problem itself: none of the classic protocols are strategyproof [8,11, 24]; moreover, it is also known that no protocol producing a proportional division can have this property [9].

We will now point out an essential difference in fairness when comparing to the fastest known protocol before our result, the $n$-player Cut Near-Halves. Our protocol treats players equally, while the $n$-player Cut Near-Halves does not. Equal treatment of players is a clear advantage if one considers the perception of fairness from the point of view of a player.

Remark 1. In the "Protocol for proportional division with unequal shares"

- each player answers the exact same queries as the other players in the same round and same subinstance; and
- no player is asked to disclose the outcome of an eval query.

Proof. In any subinstance, our protocol asks each player to answer the same proportional cut query, namely cutting the current cake to near-halves. Eval queries in these proportional queries are only utilized as technical workarounds to determine the value of the piece that plays the cake in the current subinstance. Their result is never revealed to any other player or even the protocol
itself. The only outcome of the proportional cut query is a mark at the near-half of the current cake. Moreover, there is no difference in the role of the players when queries are asked, and no player is doomed to receive her exact share, like the cutter in Cut-and-Choose. If we consider Cut-NearHalves, being the cutter in the first round is the most undesired role, followed by being a cutter in the second round, and so on. Our protocol forgoes this differentiation between the players, since it addresses the same queries to each player in a round, and the cake will be cut at the mark of the player whose demand happens to reach $\left\lfloor\frac{D}{2}\right\rfloor$ when the demands are summed up in order of the marks on the cake.

The $n$-player Cut Near-Halves protocol fails to satisfy both of the above points. It addresses both eval and cut queries to players and treats players differently based on which type of query they got. In the two-player version of Cut Near-Halves, only one player marks the cake and the other player uses an eval query to choose a side. This enables the second player to have a chance for a piece strictly better than half of the cake, while the first player is only entitled for her exact proportional share and has no chance to receive more than that. Besides this, the player who is asked to evaluate a piece might easily speculate that she was offered the piece because the other player cut it off the cake-and thus gain information about the measure function of the other player.

However, the remark is true for the Even-Paz protocol for proportional division with equal shares, which can be utilized in our problem through the cloning technique. As mentioned in Section 3, it needs $O(D \log D)$ proportional cut queries. The more efficient $n$-player Cut Near-Halves protocol only needs $O\left(n^{2} \log D\right)$ queries, but it treats players differently. Our protocol adheres to the equal treatment of players principle and beats both protocols in efficiency.

## 5 GENERALIZATIONS

In this section, we introduce a far generalization of cake cutting, where the cake is a measurable set in arbitrary finite dimension and cuts are defined by a monotone function. At the end of the section, we prove that even in the generalized setting, $O(n \log D)$ queries suffice to construct a proportional division.

### 5.1 A General Cake Definition

Our players remain $\left\{P_{1}, P_{2}, \ldots, P_{n}\right\}$ with demands $d_{i} \in \mathbb{Z}^{+}$, but the cake is now a Lebesguemeasurable subset $X$ of $\mathbb{R}^{k}$ such that $0<\lambda(X)<\infty$. Each player $P_{i}$ has a non-negative, absolutely continuous measure function $\mu_{i}$ defined on the Lebesgue-measurable subsets of $X$. An important consequence of this property is that for every $Z \subseteq X, \mu_{i}(Z)=0$ if $\lambda(Z)=0$. The value of the whole cake is identical for all players; in particular, it is the sum of all demands:

$$
\forall 1 \leq i \leq n \quad \mu_{i}(X)=D=\sum_{j=1}^{n} d_{j}
$$

A measurable subset $Y$ of the cake $X$ is called a piece. The volume of a piece $Y$ is the value $\lambda(Y)$ taken by the Lebesgue-measure on $Y$. The cake $X$ will be partitioned into pieces $X_{1}, \ldots, X_{n}$.

Definition 5.1. A set $\left\{X_{i}\right\}_{1 \leq i \leq n}$ of pieces is a division of $X$ if $\bigcup_{1 \leq i \leq n} X_{i}=X$ and $X_{i} \cap X_{j}=\emptyset$ holds for all $i \neq j$. We call division $\left\{X_{i}\right\}_{1 \leq i \leq n}$ proportional if $\mu_{i}\left(X_{i}\right) \geq d_{i}$ holds for all $1 \leq i \leq n$.

We will show in Section 5.3 that a proportional division always exists.

### 5.2 A Stronger Query Definition

The more general cake clearly requires a more powerful query notion. Cut and eval queries are defined on an arbitrary piece (i.e., measurable subset) $I \subseteq X$. Beyond this, each cut query specifies
a value $\alpha \in \mathbb{R}^{+}$and a mapping $f:[0, \lambda(I)] \rightarrow 2^{I}$ (representing a moving knife) such that $f(x) \subseteq$ $f(y)$ and $\lambda(f(x))=x$ holds for every $0 \leq x \leq y \leq \lambda(I)$. This function $f(x)$ is essentially identical to the "knife function" from [32], except that we do not assume $S$-continuity, a notion defined in [32]. Instead, we show that it is always possible to make a cut at $\alpha$ for any chosen $0 \leq \alpha \leq \mu(I)$.

Lemma 5.2. If $\mu: 2^{I} \rightarrow \mathbb{R}_{\geq 0}$ is an absolutely continuous measure defined on the Lebesguemeasurable set $I$, and $f:[0, \lambda(I)] \rightarrow 2^{I}$ is a mapping with $f(x) \subseteq f(y)$ and $\lambda(f(x))=x$ for every $0 \leq x \leq y \leq \lambda(I)$, then for any $\alpha$ with $0 \leq \alpha \leq \mu(I)$ there exists a $t \in[0,1]$ such that $\mu(f(t))=\alpha$.

Proof. Define $X:=\{x: \mu(f(x)) \leq \alpha\}, Y:=\{y: \mu(f(y)) \geq \alpha\}$ and let $L:=\bigcup\{f(x): x \in X\}$, $U:=\bigcap\{f(y): y \in Y\}$. Note that the union and intersection defining $L$ and $U$, respectively, can be taken as countable (e.g., by monotone sequences converging to $\sup X$ and $\inf Y$, respectively), as $f$ is monotone in $x$. Thus, $L$ and $U$ are measurable. Finally, $A:=U \backslash L$. From $X \cup Y=[0,1]$ follows that $\sup X=\inf Y$. Let $t$ denote this common value. Due to the monotonicity of $f$,

$$
\mu(L)=\mu\left(\bigcup_{n=\left\lceil\frac{1}{t}\right\rceil}^{\infty} f\left(t-\frac{1}{n}\right)\right)=\lim _{n \rightarrow \infty} \mu\left(f\left(t-\frac{1}{n}\right)\right) \leq \alpha
$$

and $\mu(U) \geq \alpha$ follows similarly. As $A \subseteq f(t+\varepsilon) \backslash f(t-\varepsilon)$, we know that $\lambda(A) \leq \lambda(f(t+\varepsilon))-$ $\lambda(f(t-\varepsilon))=2 \varepsilon$ holds for any $\varepsilon>0$. Hence, $\lambda(A)=0$ and by the absolute continuity of $\mu$, $0=\mu(A)=\mu(U)-\mu(L)$ follows. Consequently, $\alpha \geq \mu(L)=\mu(U) \geq \alpha$, hence, $\mu(f(t))=\alpha$ as we claimed.

Note that as $\mu(f(x))$ is a monotone function of $x$, Lemma 5.2 also implies that $\mu(f(x))$ is continuous in $x$. We now turn to defining the two queries in our generalized model.

- Eval query $\left(P_{i}, I\right)$ returns $\mu_{i}(I)$.
- Cut query $\left(P_{i}, I, f, \alpha\right)$ returns the smallest $x \leq \lambda(I)$ with $\mu_{i}(f(x))=\alpha$ or an error message if such an $x$ does not exist.

As queries involve an arbitrary measurable subset $I$ of $X$, our generalized queries automatically cover the generalization of the previously discussed Edmonds-Pruhs queries, proportional queries, and reindexing. If we restrict our attention to the usual unit interval cake $[0,1]$, generalized queries open up a number of new possibilities for a query, as Example 5.3 shows.

Example 5.3. On the unit interval cake the following rules qualify as generalized queries.
-Evaluate an arbitrary measurable set.

- Cut a piece of value $\alpha$ surrounding a point $x$ so that $x$ is the midpoint of the cut piece.
-For disjoint finite sets $A$ and $B$, cut a piece $Z$ of value $\alpha$ such that $Z$ contains the $\varepsilon$ neighborhood of $A$ and avoids the $\varepsilon$-neighborhood of $B$ for a maximum $\varepsilon$.
-Determine $x$ such that the union of intervals $[0, x],\left[\frac{1}{n}, \frac{1}{n}+x\right], \ldots,\left[\frac{n-1}{n}, \frac{n-1}{n}+x\right]$ is of value $\alpha$.

The new notions also allow us to define cuts on a cake in higher dimensions.
Example 5.4. Defined on the generalized cake $X \subseteq \mathbb{R}^{k}$, the following rules qualify as generalized queries.
-Evaluate an arbitrary measurable set.

- Cut a piece of value $\alpha$ of piece $I$ so that the cut is parallel to a given hyperplane.
-Multiple cut queries on the same piece $I \subset \mathbb{R}^{2}$ : one player always cuts $I$ along a horizontal line, the other player cuts the same piece along a vertical line. The two players have different
$f(x)$ functions in this example; the first one has a horizontal "knife" that is moved from top to bottom (or the other way around), while the second player has a vertical "knife" that is moved from left to right (or the other way around).


### 5.3 The Existence of a Proportional Division

Our protocol "Proportional division with unequal shares" in Section 4 extends to the above described general setting and hence proves that a proportional division always exists.

Theorem 5.5. For any $2 \leq n<D$ and for any function $f:[0, \lambda(I)] \rightarrow 2^{I}$ with $f(x) \subseteq f(y)$ and $\lambda(f(x))=x$ for every $0 \leq x \leq y \leq \lambda(I)$, the number of generalized queries in our $n$-player protocol on the generalized cake of total value $D$ is $T(n, D) \leq 2(n-1) \cdot\left\lceil\log _{2} D\right\rceil$.

Proof. The proof of Theorem 3.1 carries over without essential changes, thus we only discuss the differences here. First, we observe that proportional queries in ratio $a: b$ can still be substituted by a constant number of eval and cut queries. In the generalized model, proportional query $\left(P_{i}, I, f, a, b\right)$ returns $x \leq \lambda(I)$ such that $b \cdot \mu_{i}(f(x))=a \cdot \mu_{i}(I \backslash f(x))$. Similarly as before, $P_{i}$ first measures $I$ by a single eval query and then uses the cut query $\left(P_{i}, I, f, \alpha\right)$ with $\alpha=\frac{a}{a+b} \cdot \mu_{i}(I)$. In the first round of our generalized protocol, all players are asked to cut the cake $X$ in near-halves using the same $f$ function. Then $P_{j}$ is calculated, just as in the simpler version, and we cut $X$ into the two near-halves according to $P_{j}$ 's $f$-cut and clone $P_{j}$ if necessary. Due to the monotonicity of $f$, this sorts each player to a piece she values at least as much as the full demand on all players sorted to that piece. Subsequent rounds are played in the same manner.

The query number for $n=2$ follows from the fact that each of the two players are asked a proportional cut query in every round until recursively halving $\left\lceil\frac{D}{2}\right\rceil$ reaches 1 , which means $\left\lceil\log _{2} D\right\rceil$ queries in total. The recursion formula remains intact in the generalized model, and thus the query number $T(n, D)=2(n-1)\left\lceil\log _{2} D\right\rceil$ too.

We remark that the existence of a proportional division heavily relies on the fact that $\mu_{i}$ is a measure, and that the monotone function $f$ can be chosen uniformly for all agents. In other papers for multidimensional cake cutting, these are not always assumed. For example, Segal-Halevi et al. [33] study the problem of fair two-dimensional division wherein the allotted pieces must be of some restricted two-dimensional geometric shapes, particularly squares and fat rectangles. Their impossibility results do not contradict our Theorem 5.5, because their $\mu_{i}$ functions are not even additive.

## 6 THE LOWER BOUND

In this section, we prove our lower bound on the number of queries any deterministic protocol needs to make when solving the proportional cake cutting problem with unequal shares. This result is valid in two relevant settings: (1) on the $[0,1]$ cake with Robertson-Webb or with generalized queries and (2) on the general cake and queries introduced in Section 5.

The lower bound proof is presented in two steps. In Section 6.1, we define a single-player cake cutting problem where the goal is to identify a piece of small volume and positive value for the sole player. For this problem, we design an adversary strategy and specify the minimum volume of the identified piece as a function of the number of queries asked. In Section 6.2, we turn to the problem of proportional cake cutting with unequal shares. We show that in order to allocate each player a piece of positive value, at least $\Omega(n \log D)$ queries must be addressed to the players-otherwise the allocated pieces overlap.


Fig. 2. The crumble partition after two queries in Example 6.1. The first and third pictures illustrate the two queries, while the second and fourth pictures illustrate the generated crumbles.

### 6.1 The Single-Player Problem

We define our single-player problem on a generalized cake of value $D$, a player $P$, and her unknown measure function $\mu$. The aim is to identify a piece of positive value according to $\mu$ by asking queries from $P$. The answers to these queries come from an adversary strategy we design. We would like to point out that the single-player thin-rich game of Edmonds and Pruhs [16] defined on the unit interval cake has a different goal. There, the player needs to receive a piece that has value not less than 1 and width at most 2 . Moreover, their proof for the $n$-player problem is restricted to instances with $n=2 \cdot 3^{\ell}, \ell \in \mathbb{Z}^{+}$, whereas ours is valid for any $n \in \mathbb{Z}^{+}$.

In our single-player problem, a set of queries reveals information on the value of some pieces of the cake. Each generalized eval query $(P, I)$ partitions the cake into two pieces; $I$ and $X \backslash I$. An executed cut query ( $P, I, f, \alpha$ ) with output $x$ partitions the cake into three; $f(x), I \backslash f(x)$ and $X \backslash I$. To each step of a protocol, we define the currently smallest building blocks of the cake, which we call crumbles. Two points of $X$ belong to the same crumble if and only if they are in the same partition in all queries asked so far. At start, the only crumble is the cake itself and every new query can break an existing crumble into more crumbles. More precisely, $q$ queries can generate $3^{q}$ crumbles at most. Crumbles at any stage of the protocol partition the entire cake. The exact value of a crumble is not necessarily known to the protocol and no real subset of a crumble can have a revealed value. As a matter of fact, the exact same information is known about the value of any subset of a crumble.

Example 6.1. In Figure 2, we illustrate an example for crumbles on the unit interval cake after two queries.

Row 1: The upper picture depicts a cut query defined on the thick green interval $I$. It generates a piece of value $\alpha$ so that it contains the $\varepsilon$-neighborhood of points $A_{1}, A_{2}, A_{3}$ for maximum $\varepsilon$. This piece is marked by the two red intervals within the thick green region in the figure.
Row 2: The cut query generates three crumbles: the two red intervals form $C_{1}$, the remainder of the thick green $I$ is $C_{2}$, while $X \backslash I$, the black segments outside of the thick green interval, form $C_{3}$. Each crumble is colored and also labeled in the second picture.
Row 3: The second query evaluates the thick blue piece $I$ in the third picture.
Row 4: The second query cuts the existing crumbles into six crumbles in total, as depicted in the bottom picture. These six crumbles are generated from the intersecting and
non-intersecting parts of the previous three crumbles with the thick blue $I$ set of the second query. The first index in a crumble label indicates the previous crumble it belonged to.

We now proceed to construct an adversary strategy that bounds the volume of any crumble $C$ with $\mu(C)>0$. Our adversary can actually reveal more information than asked; we allow her to disclose the value of each crumble in the cake. When a query is asked, the answer is determined based on the parameters of the query and the current set of crumbles, which we denote by $\mathfrak{C}$. Together with the answer to the query, the adversary also specifies the new set of crumbles $\mathfrak{C}^{n e w}$ together with $\mu\left(C^{\text {new }}\right)$ for each $C^{\text {new }} \in \mathfrak{C}^{\text {new }}$. In the next query, this $\mathfrak{C}^{n e w}$ will serve as the current set of crumbles $\mathfrak{C}$. The adversary answers the queries in accordance with the following rules, which are also stated in a pseudocode "Adversary strategy" below.

- eval query (C, $I$ )

This query changes the structure of the crumble set $\mathfrak{C}$ in such a way that each crumble $C \in \mathbb{C}$ is split into exactly two new crumbles $C \cap I$ and $C \backslash I$, either of which might be empty (lines 1 and 2). If the part inside the crumble is at least as large as the other part, then the adversary assigns the full value of $C$ to $C \cap I$ (lines 3-5). Otherwise, the outer part $C \backslash I$ will get the entire value (lines $6-8$ ). The answer to the eval query is the total value of new crumbles that lie in $I$ (line 11).

- cut query ( $\mathfrak{C}, I, f, \alpha$ )

Each cut query is executed in two rounds. In the first round (lines $12-22$ ), we define new crumbles $C \backslash I$ (line 13) and intermediate crumbles $C \cap I$ (line 14) for all crumbles $C \in \mathbb{C}$. If $\lambda(C \cap I) \geq 2 / 3 \cdot \lambda(C)$, then $C \cap I$ inherits the entire value of $C$ (lines 15-17), otherwise $C \backslash I$ carries all the value of $C$ (lines 18-20). The new crumbles $C \backslash I$ are set aside until the next query arrives, while the intermediate crumbles $C \cap I$ will be the crumbles of the second round (lines 23-45).
If the total value of these intermediate crumbles is less than $\alpha$, then an error message is returned indicating that $I$ is not large enough to be cut off a piece of value $\alpha$ (lines 23 and 24). Otherwise, for each intermediate crumble $C_{i}^{i n t}$, we define the value $x_{i}$ for which $f\left(x_{i}\right)$ halves $C_{i}^{\text {int }}$ in volume (lines 26-28). Such an $x_{i}$ value must exist, due to Lemma 5.2. We then reorder the indices of intermediate crumbles according to these $x_{i}$ values (line 29). Now we find the index $k$ for which $\sum_{i=1}^{k-1} \mu\left(C_{i}^{i n t}\right)<\alpha \leq \sum_{i=1}^{k} \mu\left(C_{i}^{i n t}\right)$.
The set of new crumbles will now be completed by adding sets $\left\{C_{i}^{\text {int }} \cap f\left(x_{k}\right)\right\}$ and $\left\{C_{i}^{i n t} \backslash f\left(x_{k}\right)\right\}$ to it (lines 31 and 32). The value of these new crumbles is specified depending on the index $i$ of $C_{i}^{i n t}$. If $i<k$, then the crumble in $f\left(x_{k}\right)$ inherits the full value of the intermediate crumble (lines 34-36). If $i>k$, then the crumble outside of $f\left(x_{k}\right)$ inherits the value of the intermediate crumble (lines 37-39). Finally, for $i=k$, the crumble inside $f\left(x_{k}\right)$ receives all of $\alpha$ that has not been assigned to new crumbles inside $f\left(x_{k}\right)$ with a smaller index (line 41). After this, the crumble outside of $f\left(x_{k}\right)$ gets the remainder of $\mu\left(C_{k}^{\text {int }}\right)$ (line 42). At last, the protocol returns $x=x_{k}$ (line 44).

Once all queries have been answered according to the above rules, the player is allocated a piece $Z \subseteq X$. The adversary specifies $\mu(Z)$ as the total value of those crumbles that are subsets of $Z$.

Having described and demonstrated our adversary strategy, we now turn to proving our key lemma on the volume of pieces that carry a positive value.

Lemma 6.2. After q queries in the single-player problem, the volume of any piece with positive value is at least $\frac{D}{3 q}$.
ACM Transactions on Algorithms, Vol. 16, No. 3, Article 29. Publication date: May 2020.

```
ALGORITHM 1: Adversary strategy
    Eval query ( \(\mathfrak{C}, I\) )
    Initialize \(\mathfrak{C}^{\text {new }}=\mathfrak{C}\)
    for \(\forall C \in \mathbb{C}\) do
        \(\mathfrak{C}^{\text {new }} \leftarrow \mathfrak{C}^{\text {new }} \cup\{C \cap I\} \cup\)
        \(\{C \backslash I\}\)
        if \(\lambda(C \cap I) \geq \frac{1}{2} \lambda(C)\) then
            \(\mu(C \cap I) \leftarrow \mu(C)\)
            \(\mu(C \backslash I) \leftarrow 0\)
        else
            \(\mu(C \cap I) \leftarrow 0\)
                \(\mu(C \backslash I) \leftarrow \mu(C)\)
        end
    end
    return \(\sum_{C^{n e w} \in \mathbb{C}^{\text {new }}, C^{\text {new }} \subseteq I} \mu\left(C^{\text {new }}\right)\)
```

```
Cut query (C, \(I, f, \alpha\) )
Initialize \(\mathfrak{C}^{\text {new }}=\mathfrak{C}, \mathfrak{C}^{\text {int }}=\emptyset\)
for \(\forall C \in \mathbb{C}\) do
    \(\mathfrak{c}^{\text {new }} \leftarrow \mathfrak{C}^{\text {new }} \cup\{C \backslash I\}\)
    \(\mathfrak{C}^{\text {int }} \leftarrow \mathfrak{C}^{\text {int }} \cup\{C \cap I\}\)
    if \(\lambda(C \cap I) \geq \frac{2}{3} \lambda(C)\) then
        \(\mu(C \cap I) \leftarrow \mu(C)\)
        \(\mu(C \backslash I) \leftarrow 0\)
    else
        \(\mu(C \cap I) \leftarrow 0\)
        \(\mu(C \backslash I) \leftarrow \mu(C)\)
    end
end
if \(\sum_{C^{\text {int }} \in \mathbb{C} \text { int }} \mu\left(C^{\text {int }}\right)<\alpha\) then
    return error
else
    for \(\forall C_{i}^{\text {int }} \in \mathbb{C}^{\text {int }}\) do
            find \(x_{i} \in \mathbb{R}\) so that \(\lambda\left(C^{\text {int }} \cap f\left(x_{i}\right)\right)=\frac{1}{2} \lambda\left(C_{i}^{\text {int }}\right)\)
    end
    reorder [i] in \(C_{i}^{\text {int }}\) so that \(x_{1} \leq x_{2} \leq \ldots\)
    find \(k \in \mathbb{Z}\) so that \(\sum_{i=1}^{k-1} \mu\left(C_{i}^{i n t}\right)<\alpha \leq \sum_{i=1}^{k} \mu\left(C_{i}^{i n t}\right)\)
    for \(\forall C_{i}^{i n t} \in \mathbb{C}^{i n t}\) do
        \(\mathfrak{C}^{\text {new }} \leftarrow \mathbb{C}^{\text {new }} \cup\left\{C_{i}^{\text {int }} \cap f\left(x_{k}\right)\right\} \cup\left\{C_{i}^{\text {int }} \backslash f\left(x_{k}\right)\right\}\)
    end
    if \(i<k\) then
        \(\mu\left(C_{i}^{\text {int }} \cap f\left(x_{k}\right)\right) \leftarrow \mu\left(C_{i}^{\text {int }}\right)\)
        \(\mu\left(C_{i}^{i n t} \backslash f\left(x_{k}\right)\right) \leftarrow 0\)
    else if \(i>k\) then
            \(\mu\left(C_{i}^{\text {int }} \cap f\left(x_{k}\right)\right) \leftarrow 0\)
            \(\mu\left(C_{i}^{i n t} \backslash f\left(x_{k}\right)\right) \leftarrow \mu\left(C_{i}^{i n t}\right)\)
        else
            \(\mu\left(C_{k}^{\text {int }} \cap f\left(x_{k}\right)\right) \leftarrow \alpha-\sum_{i=1}^{k-1} \mu\left(C_{i}^{\text {int }}\right)\)
            \(\mu\left(C_{k}^{i n t} \backslash f\left(x_{k}\right)\right) \leftarrow \mu\left(C_{k}^{\text {int }}\right)+\sum_{i=1}^{k-1} \mu\left(C_{i}^{i n t}\right)-\alpha\)
        end
        return \(x_{k}\)
end
```

Proof. Due to the last rule on $\mu(Z)$ of the adversary strategy, the volume of any piece with positive value is bounded from below by the volume of any crumble with positive value. We will now argue that eval and cut queries assign positive value to crumbles whose volume is at least a third of the volume of the previous crumble.

At the very beginning of the protocol, for $q=0$, the only crumble is $X$ itself, with volume $\frac{D}{3^{0}}$. Eval queries assign positive value to crumbles that are at least as large as half of the previous crumble they belonged to prior to the query. If a new crumble with positive value was created in the first round of a cut query, then its volume was at least one-third of a previous crumble (lines 13 and 20). Otherwise, the new crumble with positive value was an intermediate crumble $C_{i}^{\text {int }}$ in the
second round. The first round of our protocol assigns positive value to an intermediate crumble only if it was at least two-thirds of the old crumble in the input of the cut query (lines 14 and 15). This round will now cut $C_{i}^{i n t}$ into two new crumbles (line 32). If $i \neq k$, then the larger of these will inherit the value of the intermediate crumble (lines 35 and 39). Otherwise, if $i=k$, then $C_{i}^{\text {int }}$ is cut into exact halves (lines 41 and 42). All in all, new crumbles that are assigned a positive value in the second round are of volume at least half of two-thirds of the volume of the original crumble in the input of the query.

### 6.2 The $n$-Player Problem

We now place our single-player problem into the framework of the original problem. The instance we construct has $c_{1} n$ players whose demand sums up to $c_{2} n$, where $c_{1} \leq c_{2}$ are arbitrary constants between 0 and 1 . We call these players humble, because their total demand is modest compared to the number of them. The remaining $\left(1-c_{1}\right) n$ players share a piece of worth $D-c_{2} n$. These players are greedy, because their total demand is large. The simplest such instance is where $n-1$ humble players have demand 1 , and the only greedy player has demand $D-(n-1)$. We fix the measure function of every greedy player to be the Lebesgue-measure, and so the value of the cake equals the volume of it. This enforces humble players to share a piece of volume $c_{2} n$ among themselves. Lemma 6.2 guarantees that after $q_{i}$ queries addressed to the humble player $P_{i}$, the volume of any piece carrying positive value for $P_{i}$ is at least $\frac{D}{3 q_{i}}$. We now sum up the volume of the pieces allocated to humble players in any proportional division.

$$
\sum_{i=1}^{c_{1} n} \frac{D}{3 q_{i}} \leq c_{2} n
$$

We divide both sides by $c_{1} n D$.

$$
\frac{1}{c_{1} n} \sum_{i=1}^{c_{1} n} 3^{-q_{i}} \leq \frac{c_{2}}{c_{1} D}
$$

For the left side of this inequality, we use the well known inequality for the arithmetic and geometric means of non-negative numbers.

$$
\sqrt[c_{1} n]{3^{-\sum_{i_{1}}^{c_{1} n} q_{i}}} \leq \frac{c_{2}}{c_{1} D}
$$

Taking the logarithm of both sides leads to the following:

$$
\frac{1}{c_{1} n}\left(-\sum_{i_{1}}^{c_{1} n} q_{i}\right) \leq \log _{3} \frac{c_{2}}{c_{1}}-\log _{3} D
$$

With this, we have arrived at a lower bound on the number of queries.

$$
\sum_{i_{1}}^{c_{1} n} q_{i} \geq c_{1} n\left(\log _{3} D+\log _{3} \frac{c_{1}}{c_{2}}\right) \sim \Omega\left(n \log _{3} D\right)
$$

This proves that one needs $\Omega(n \log D)$ queries to derive a proportional division for the humble players in the instance. Moreover, if $c_{1}$ and $c_{2}$ are known, a more accurate bound can be determined using our formula $c_{1} n\left(\log _{3} D-\log _{3} c_{2}+\log _{3} c_{1}\right)$. This suggests that the problem becomes harder to solve if the $c_{1} n$ humble players vastly outnumber the greedy players. In the $c_{1} n=n-1$ case we mentioned earlier, the query number is at least $(n-1) \log _{3} D$.

We can now conclude our theorem on the lower bound.

Theorem 6.3. To construct a proportional division in an n-player unequal shares cake cutting problem with demands summing up to $D$ one needs $\Omega(n \log D)$ queries.

This bound asymptotically matches the complexity of our "Protocol for proportional division with unequal shares." As we described in Section 2, there are various notions widely accepted as Robertson-Webb queries, and thus, an analysis on the exact number of queries exclusively depends on the allowances one makes on the query model.

## 7 IRRATIONAL DEMANDS

In this section, we consider the case when some demands are irrational numbers. Apart from this, our setting is exactly the same as before. The termination of our protocol in Section 4 is guaranteed by the fact that each generated subinstance has integer demands that sum up to the initial $D$ for all subinstances of the same round. After sufficiently many instance divisions, the subinstances must reduce to base case instances, for no demand can go below 1, the smallest integer. Such a lower bound would not exist if the protocol is started with irrational demands.

Even though two direct protocols have been presented for the problem of proportional cake cutting with irrational demands $[2,34]$, we feel that our protocol sheds new light on the topic.

The complexity of all known protocols for irrational shares, including ours, falls into the same category: finite but unbounded. Another common feature of all three protocols is that they utilize disagreement between agents; more precisely, a piece of the cake that is evaluated differently by two players. Barbanel [2] assumes that the measures are not identical for all players, and that we are given a witness of the disagreement of some two measures in the input. After solving the two-player case, the recursive method for $n$ players is referred to, in which the $n$th player shares each of the $n-1$ pieces the other players have already obtained. Shishido and Zeng [34] present a protocol that is claimed to be simpler than the one of Barbanel [2], and it forgoes the assumption on a disagreement given in the input-instead, they produce one. First they present a two-player protocol, in which $P_{2}$ marks the cake at her proportional share, and if $P_{1}$ agrees, then the cake is cut, otherwise a disagreement is found. The authors then refer to the usual recursive method to the case of $n$ players. The two-player case with disagreement-given or calculated-is solved similarly in both papers. Let $\mu_{2}(X)-\mu_{1}(X)>0$ on some piece $X$ for which $\mu_{2}(X)=d_{1}$. It is shown that $P_{2}$ is able to mark a large number of overlapping intervals on the cake that are all worth $d_{1}$ for her, so that at least one out of these is evaluated by $P_{1}$ to $d_{1}$ at least.

These procedures are cumbersome compared to our protocol, which in each round either reduces the problem to one with rational demands or decreases the number of players. Moreover, our method works on our generalized cake and query model.

Our way of utilizing disagreement is as follows. Let us choose an arbitrary piece $A \subseteq X$ such that $\mu_{i}(A)>0$ for all players $P_{i}$. If the players share $A$ and $X \backslash A$ in two separate instances, both in their original ratio $d_{1}: d_{2}: \ldots: d_{n}$, then the two proportional divisions will give a proportional $d_{1}: d_{2}: \ldots: d_{n}$ division of $X$ itself. Assume now that $\mu_{i}(A)<\mu_{j}(A)$ for some players $P_{i}$ and $P_{j}$, and some piece $A \subseteq X$. When generating the two subinstances on $A$ and $X \backslash A$, we reduce $d_{i}$ on $A$ to 0 and increase it in return on $X \backslash A$ and swap the roles for $d_{j}$, increasing it on $A$ and decreasing it on $X \backslash A$. The first generated instance thus has $n-1$ players with irrational demands, while the second instance has $n$ players with irrational demands. We will show in Lemma 7.2 that if we set the right new demands in these instances, the two proportional divisions deliver a proportional division of $X$. The key point we prove in Lemma 7.3, which states that the demands in the second subinstance sum up to slightly below all players' evaluation of $X \backslash A$. Redistributing the slack as extra demand among players gives us the chance to round the demands up to rational numbers in the second subinstance and keep proportionality in the original instance. Iteratively breaking up
the instances into an instance with fewer players and an instance with rational demands leads to a set of instances with rational demands only.

We now describe our protocol in detail. Without loss of generality, we can assume that $d_{1} \leq$ $d_{2} \leq \ldots \leq d_{n}$. As a first step, $P_{1}$ answers the cut query with $x=d_{1}$ and $I=X$. We denote the piece in $f\left(d_{1}\right)$ by $A$ and ask all players to evaluate $A$. Let $P_{j}$ be one of the players whose evaluation is the highest. Notice that $\mu_{j}(A) \geq d_{1}$, because $\mu_{1}(A)=d_{1}$. We distinguish two cases from here, which are also summarized in the framed description below.
(1) If $\mu_{j}(A)=d_{1}$, then $\mu_{i}(A) \leq d_{1}$ for all players. We allocate $A$ to $P_{1}$ and continue with an instance $\mathcal{I}_{1}$ with $n-1$ players having the same demands as before. The measure functions need to be normalized to $\frac{D-d_{1}}{D-\mu_{i}(A)} \cdot \mu_{i}$ for all $i \neq 1$ so that all players of $\mathcal{I}_{1}$ evaluate $X \backslash A$ to $D-d_{1}$.
(2) Otherwise, $\mu_{j}(A)=d_{1}+\varepsilon$, where $\varepsilon>0$. We generate instances $I_{2 a}$ and $I_{2 b}$.
(a) In the first instance $\mathcal{I}_{2 a}$, the cake is $A, P_{1}$ 's demand is $0, P_{j}$ 's demand is $d_{j}+d_{1}$, while all other players keep their original $d_{i}$ demand. In order to make all players evaluate the full cake to the sum of their demands $D$, measure functions are modified to $\frac{D}{\mu_{i}(A)} \cdot \mu_{i}$.
(b) In the second instance $\mathcal{I}_{2 b}$, the cake is $X \backslash A, P_{1}$ 's demand is $d_{1}+\frac{d_{1}^{2}}{D-d_{1}}, P_{j}$ 's demand is $d_{j}-\frac{d_{1}\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}$, while the original $d_{i}$ demands are kept for all other players. In order to make all players evaluate the full cake to $D$, we set $\frac{D}{D-\mu_{i}(A)} \cdot \mu_{i}$.

## Proportional division with irrational demands:

$P_{1}$ marks $d_{1} \rightarrow A$. All players evaluate $A$. $P_{j}$ has the highest evaluation.
If $\mu_{j}(A)=d_{1}$, then allocate $A$ to $P_{1}$ and continue with $n-1$ players on $\mathcal{I}_{1}$.

Otherwise, $\mu_{j}(A)=d_{1}+\varepsilon$. Define two instances $\mathcal{I}_{2 a}$ and $\mathcal{I}_{2 b}$. While $\mathcal{I}_{2 a}$ has $n-1$ players, demands in $I_{2 b}$ sum up to below $D$ and thus can be rationalized.

|  | $\mathcal{I}_{1}$ | $\mathcal{I}_{2 a}$ | $\mathcal{I}_{2 b}$ |
| :--- | :---: | :---: | :---: |
| cake | $X \backslash A$ | $A$ | $X \backslash A$ |
| $d_{1}$ | 0 | 0 | $d_{1}+\frac{d_{1}^{2}}{D-d_{1}}$ |
| $d_{j}$ | $d_{j}$ | $d_{j}+d_{1}$ | $d_{j}-\frac{d_{1}\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}$ |
| $d_{i}$ | $d_{i}$ | $d_{i}$ | $d_{i}$ |
| $\mu_{i}$ | $\frac{D-d_{1}}{D-\mu_{i}(A)} \mu_{i}$ | $\frac{D}{\mu_{i}(A)} \mu_{i}$ | $\frac{D}{D-\mu_{i}(A)} \mu_{i}$ |

The upcoming three lemmas justify our instance transformations to $\mathcal{I}_{1}, \mathcal{I}_{2 a}$, and $I_{2 b}$.
Lemma 7.1. A proportional division in $I_{1}$ extends to a proportional division in the original problem once $P_{1}$ 's allocated piece $A$ is added to it.

Proof. Clearly $P_{1}$ is satisfied with $A$, since $\mu_{1}(A)=d_{1}$. In any proportional division in $\mathcal{I}_{1}$, every player $P_{i}, i \neq 1$ is guaranteed to receive a piece that is worth at least $\frac{D-\mu_{i}(A)}{D-d_{1}} \cdot d_{i} \geq d_{i}$ for her in the original instance.

Lemma 7.2. If each player receives her demanded share in $\mathcal{I}_{2 a}$ and $\mathcal{I}_{2 b}$, then the union of these pieces gives a proportional division in the original problem.

Proof. We calculate the share of each player for the case when each player receives a piece satisfying her demand in $\mathcal{I}_{2 a}$ and $I_{2 b}$.

- $d_{1}: 0+\left(d_{1}+\frac{d_{1}^{2}}{D-d_{1}}\right) \cdot \frac{D-d_{1}}{D}=d_{1}$
- $d_{j}:\left(d_{j}+d_{1}\right) \cdot \frac{d_{1}+\varepsilon}{D}+\left(d_{j}-\frac{d_{1} \cdot\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}\right) \cdot \frac{D-\left(d_{1}+\varepsilon\right)}{D}=d_{j}$
- $d_{i}, i \notin\{1, j\}: d_{i} \cdot \frac{d_{1}+\varepsilon}{D}+d_{i} \cdot \frac{D-\left(d_{1}+\varepsilon\right)}{D}=d_{i}$

Lemma 7.3. By slightly increasing all demands, $I_{2 b}$ can be transformed into an instance of proportional cake cutting with rational demands.

Proof. The key observation here is that there is a slack in the demands, meaning that demands in $I_{2 b}$ sum up to strictly below $D$, which is the evaluation of all players of the full cake $X \backslash A$. The sum of the demands is the following.

$$
d_{1}+\frac{d_{1}^{2}}{D-d_{1}}+d_{j}-\frac{d_{1} \cdot\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}+\sum_{i \notin\{1, j\}} d_{i}=\sum_{i=1}^{n} d_{i}+\frac{d_{1}^{2}}{D-d_{1}}-\frac{d_{1} \cdot\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}<D
$$

The inequality above follows from the fact that $\frac{d_{1}^{2}}{D-d_{1}}<\frac{d_{1}\left(d_{1}+\varepsilon\right)}{D-\left(d_{1}+\varepsilon\right)}$ for all $\varepsilon>0$.
The slack can be distributed as extra demand among all players so that all demands are rational. An implementation of this could be that we round up the irrational demands at a sufficiently insignificant digit.

Theorem 7.4. Any instance of the proportional cake cutting problem with $n$ players and irrational demands can be transformed into at most $n-1$ proportional cake cutting problems with rational demands and thus can be solved using a finite number of queries.

Proof. We prove this theorem by induction. For $n=2$, we need to show that the problem with irrational demands can be reduced to at most one proportional cake cutting problem with rational demands. For two players, our protocol proceeds as follows. First, $P_{1}$ marks a piece $X$ that is worth $d_{1}$ for her. Now we ask $P_{2}$ to evaluate $X$. If $\mu_{2}(X) \leq d_{1}$, then $P_{1}$ is allocated $X$ and $P_{2}$ is satisfied with $I \backslash X$, because $\mu_{2}(I \backslash X) \geq d_{2}$. Otherwise, if $\mu_{2}(X)=d_{1}+\varepsilon$ for some $\varepsilon>0$, then $P_{2}$ is allocated $X$ and the two players share $I \backslash X$ with demands $d_{1}+\frac{d_{1}^{2}}{d_{2}}$ and $d_{2}-\frac{d_{1}\left(d_{1}+\varepsilon\right)}{d_{2}-\varepsilon}$. These demands ensure a proportional share to both players, as we show in Lemma 7.2. Moreover, applying Lemma 7.3 to two players proves that they sum up to strictly below $d_{1}+d_{2}$ and thus can be rounded up to rational numbers, which gives us the single two-player problem that must be solved in order to derive a proportional division for the original problem.

Assume now that an $n-1$-player proportional cake cutting problem with irrational demands can be solved by transforming it into $n-2$ proportional cake cutting problems with rational demands. If we are given an $n$-player proportional cake cutting problem with irrational demands, our "Proportional division with irrational demands" protocol transforms it into either $I_{1}$ or into a problem with two instances, $I_{2 a}$ and $I_{2 b}$. Solving either of those problems will lead to a proportional division in the original $n$-player problem, as Lemmas 7.1 and 7.2 show. The first instance is an $n-1$ player proportional cake cutting problem with irrational demands, which can be solved via $n-2$ proportional cake cutting problems with rational demands by our assumption. The same is true for $I_{2 a}$. As of $I_{2 b}$, Lemma 7.3 proves that its demands can be rounded up to rational numbers. Even in the worst case, when our protocol generates $I_{2 a}$ and $I_{2 b}$ in every recursive step, it ends up constructing $n-1$ instances with rational demands.

Just as our earlier protocol, this protocol is also not strategyproof. Players in this protocol play distinct roles; for example, it is crucial to start with $P_{1}$, who has the smallest demand. We would like to emphasize that even though we have transformed any proportional cake cutting problem with irrational demands into a set of problems with rational demands, we did not show any upper bound on its query complexity. When the problems with rational demands are created, $D$ might grow arbitrarily large, which hugely affects the query number.

Open questions. Deciding whether there is a bounded protocol for irrational demands might be the most intriguing question left open by this article. To construct a bounded protocol, one would need to bound the number of queries in $n$ and at least one other measure corresponding to $D$ in
our setting. Another possible research direction leads to the topic of chore division, where players share "bads" instead of goods [17,21]. A division is proportional if no player $P_{i}$ receives a piece that is worth more than $d_{i}$. Our protocol from Section 4 carries over to this problem variant with minimal modifications, namely that we need to allocate each player to the subinstance where her mark is not present. We do not see a similarly simple argument to save the validity of our lower bound from Section 6. For the proportional chore division problem with equal shares, Farhadi and Hajiaghayi [20] prove the same bound as the corresponding cake cutting problem has. They design a dual measure function for each player, based on which function the adversary strategy is defined. We conjecture that combining their approach with ours can lead to an $\Omega(n \log D)$ lower bound for the proportional chore division problem with unequal shares.

## ACKNOWLEDGMENT

We thank Simina Brânzei, Péter Koltai, Erel Segal-Halevi, and the reviewers of an earlier version of the article for their generous and insightful advice.

## REFERENCES

[1] Moshe Babaioff, Noam Nisan, and Inbal Talgam-Cohen. 2017. Competitive equilibrium with indivisible goods and generic budgets. arXiv preprint arXiv:1703.08150.
[2] Julius B. Barbanel. 1996. Game-theoretic algorithms for fair and strongly fair cake division with entitlements. In Colloquium Mathematicae, Vol. 69:1. 59-73.
[3] Julius B. Barbanel, Steven J. Brams, and Walter Stromquist. 2009. Cutting a pie is not a piece of cake. The American Mathematical Monthly 116, 6 (2009), 496-514.
[4] Anatole Beck. 1987. Constructing a fair border. The American Mathematical Monthly 94, 2 (1987), 157-162.
[5] Marcus Berliant, William Thomson, and Karl Dunz. 1992. On the fair division of a heterogeneous commodity. Fournal of Mathematical Economics 21, 3 (1992), 201-216.
[6] Steven J. Brams, Michael A. Jones, and Christian Klamler. 2008. Proportional pie-cutting. International fournal of Game Theory 36, 3 (2008), 353-367.
[7] Steven J. Brams, Michael A. Jones, and Christian Klamler. 2011. Divide-and-conquer: A proportional, minimal-envy cake-cutting algorithm. SIAM Review 53, 2 (2011), 291-307.
[8] Simina Brânzei and Peter Bro Miltersen. 2013. Equilibrium analysis in cake cutting. In Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems. International Foundation for Autonomous Agents and Multiagent Systems, 327-334.
[9] Simina Brânzei and Peter Bro Miltersen. 2015. A dictatorship theorem for cake cutting. In Proceedings of the 29th AAAI Conference on Artificial Intelligence. AAAI Press, 481-488.
[10] Edward Carney. 2012. A new algorithm for the cake-cutting problem of unequal shares for rational ratios: The divisor reduction method. Scientific Terrapin 3, 2 (2012), 15-22.
[11] Yiling Chen, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2013. Truth, justice, and cake cutting. Games and Economic Behavior 77, 1 (2013), 284-297.
[12] Yuga J. Cohler, John K. Lai, David C. Parkes, and Ariel D. Procaccia. 2011. Optimal envy-free cake cutting. In Proceedings of the T25th AAAI Conference on Artificial Intelligence. AAAI Press, 626-631.
[13] Logan Crew, Bhargav Narayanan, and Sophie Spirkl. 2019. Disproportionate division. arXiv preprint arXiv:1909.07141.
[14] Marco Dall'Aglio. 2001. The Dubins-Spanier optimization problem in fair division theory. Fournal of Computational and Applied Mathematics 130, 1 (2001), 17-40.
[15] Lester E. Dubins and Edwin H. Spanier. 1961. How to cut a cake fairly. The American Mathematical Monthly 68, 1 (1961), 1-17.
[16] Jeff Edmonds and Kirk Pruhs. 2011. Cake cutting really is not a piece of cake. ACM Transactions on Algorithms (TALG) 7, 4 (2011), 51.
[17] Francis Edward Su. 1999. Rental harmony: Sperner's lemma in fair division. The American Mathematical Monthly 106, 10 (1999), 930-942.
[18] Shimon Even and Azaria Paz. 1984. A note on cake cutting. Discrete Applied Mathematics 7, 3 (1984), 285-296.
[19] Alireza Farhadi, Mohammad Ghodsi, Mohammad Taghi Hajiaghayi, Sebastien Lahaie, David Pennock, Masoud Seddighin, Saeed Seddighin, and Hadi Yami. 2019. Fair allocation of indivisible goods to asymmetric agents. Journal of Artificial Intelligence Research 64 (2019), 1-20.
[20] Alireza Farhadi and Mohammad Taghi Hajiaghayi. 2017. On the complexity of chore division. arXiv preprint arXiv:1710.00271.
[21] Martin Gardner. 1978. Aha! Aha. Insight. Scientific American, New York.
[22] Theodore P. Hill. 1983. Determining a fair border. The American Mathematical Monthly 90, 7 (1983), 438-442.
[23] Karthik Iyer and Michael N. Huhns. 2009. A procedure for the allocation of two-dimensional resources in a multiagent system. International fournal of Cooperative Information Systems 18, $03 n 04$ (2009), 381-422.
[24] David Kurokawa, John K. Lai, and Ariel D. Procaccia. 2013. How to cut a cake before the party ends. In Proceedings of the 27th AAAI Conference on Artificial Intelligence. AAAI Press, 555-561.
[25] Andrew Lohr. 2012. Tight lower bounds for unequal division. arXiv preprint arXiv:1206.1553.
[26] Malik Magdon-Ismail, Costas Busch, and Mukkai S. Krishnamoorthy. 2003. Cake-cutting is not a piece of cake. In 20th Annual Symposium on Theoretical Aspects of Computer Science. Springer, Berlin,, 596-607.
[27] Kevin McAvaney, Jack Robertson, and William Webb. 1992. Ramsey partitions of integers and pair divisions. Combinatorica 12, 2 (1992), 193-201.
[28] Ariel D. Procaccia. 2013. Cake cutting: Not just child's play. Communications of the ACM 56, 7 (2013), 78-87.
[29] Ariel D. Procaccia. 2015. Cake cutting algorithms. In Handbook of Computational Social Choice, chapter 13. Cambridge University Press.
[30] Jack Robertson and William Webb. 1998. Cake-cutting Algorithms: Be Fair If You Can. AK Peters, Natick.
[31] Erel Segal-Halevi. 2019. Cake-cutting with different entitlements: How many cuts are needed? Journal of Mathematical Analysis and Applications 480, 1 (2019), 123382. DOI : https://doi.org/10.1016/j.jmaa.2019.123382
[32] Erel Segal-Halevi, Avinatan Hassidim, and Yonatan Aumann. 2015. Envy-free cake-cutting in two dimensions. In 29th AAAI Conference on Artificial Intelligence.
[33] Erel Segal-Halevi, Shmuel Nitzan, Avinatan Hassidim, and Yonatan Aumann. 2017. Fair and square: Cake-cutting in two dimensions. Journal of Mathematical Economics 70 (2017), 1-28.
[34] Harunor Shishido and Dao-Zhi Zeng. 1999. Mark-choose-cut algorithms for fair and strongly fair division. Group Decision and Negotiation 8, 2 (1999), 125-137.
[35] Hugo Steinhaus. 1948. The problem of fair division. Econometrica 16 (1948), 101-104.
[36] Gerhard J. Woeginger and Jiří Sgall. 2007. On the complexity of cake cutting. Discrete Optimization 4, 2 (2007), 213220.

Received January 2019; revised January 2020; accepted January 2020


[^0]:    Supported by the Cooperation of Excellences Grant (KEP-6/2019), the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2019), its János Bolyai Research Fellowship, and OTKA grant K128611.
    Authors' addresses: Á. Cseh, Centre for Economic and Regional Studies, Institute of Economics, Tóth Kálmán utca 4, Budapest 1097, Hungary; email: cseh.agnes@krtk.mta.hu; T. Fleiner, Department of Computer Science and Information Theory, Budapest University of Technology and Economics, Magyar tudósok kőrútja 2, Budapest 1117, Hungary; email: fleiner@cs.bme.hu.
    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
    © 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.
    1549-6325/2020/05-ART29 \$15.00
    https://doi.org/10.1145/3380742

