A Closer Look at the Cake-Cutting Foundations through the Lens of Measure Theory

Peter Kern, Daniel Neugebauer, Jörg Rothe,
René L. Schilling, Dietrich Stoyan, and Robin Weishaupt

Abstract

Cake-cutting is a playful name for the fair division of a heterogeneous, divisible good among agents, a well-studied problem at the intersection of mathematics, economics, and artificial intelligence. The cake-cutting literature is rich and edifying. However, different model assumptions are made in its many papers, in particular regarding the set of allowed pieces of cake that are to be distributed among the agents and regarding the agents' valuation functions by which they measure these pieces. We survey the commonly used definitions in the cake-cutting literature, highlight their strengths and weaknesses, and make some recommendations on what definitions could be most reasonably used when looking through the lens of measure theory.

1 Introduction

Across the research field of cake-cutting (see, e.g., the books by Brams and Taylor [16] and Robertson and Webb [39] and the chapters by Procaccia [36] and Lindner and Rothe [32]) there exist several different assumptions about the underlying model used. Our goal is to review thoroughly and comprehensively all the different models currently applied in the literature. Furthermore, we study the relationships between these models and formulate some results about them. It turns out that some of these models are problematic and should not be used as they are formulated. We highlight these models' problems and provide specific examples showing why they are problematic. Our overall goal is to determine a model as general as possible, yet which fixes these problems and is compatible with as many as possible of the currently used models.

To motivate our study, consider first the following fun example. If you enter a pastry shop and ask the confectioner whether he is able to cut a continuous piece of cake for you out of a chocolate cake, the confectioner probably agrees. However, if you ask the confectioner whether he might cut a Cantor-like piece of cake for you, you most likely must explain yourself. To obtain a Cantor-like piece of cake, one proceeds as follows, depicted in Figure 1. We start with the complete cake as a piece, i.e., $A_0 = X = [0, 1]$. Now, cut out the middle third of A_0 to obtain the intermediate piece $A_1 = A_0 \setminus (1/3, 2/3)$. Next, cut out the middle third of both remaining pieces in A_1 to obtain $A_2 = [0, 1/9] \cup [2/9, 1/3] \cup [2/3, 7/9] \cup [8/9, 1]$. Repeat this procedure of cutting out the middle third of each piece infinitely often. The intersection of all these pieces, i.e., $C = \bigcap_{i=1}^{\infty} A_i$, then is a Cantor-like piece of cake (in technical terms, it is the *Cantor set*, a.k.a. *Cantor dust*). Cutting out the middle third of a cake in a first step is no problem. But, continuing this procedure with the remaining pieces infinitely often will let the confectioner fail (perhaps not immediately, but as time goes on). Now, let us consider another example a bit more formally.

Example 1.1 (Vitali [49]; see also Appendix G, Theorem G.4 in the monograph [42] and § 7.22 in [43])**.** *As usual, we denote the cake by* $X = [0, 1]$*. Commonly, in cake-cutting theory a valuation function* $v\colon \mathcal{P}(X) \to [0, 1]$ *is represented as shown in Figure 2: The cake* X *is split horizontally into multiple pieces and the number of vertically stacked boxes per piece*

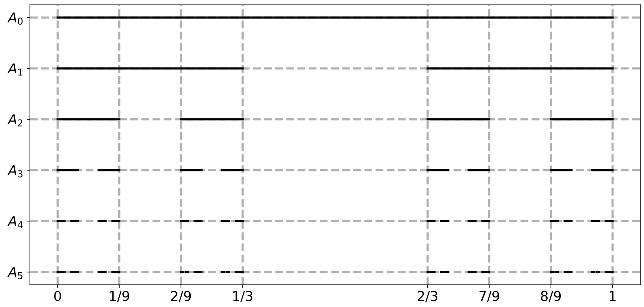

Figure 1: Step-wise pieces to be cut for a Cantor-like piece of cake.

describe the piece's valuation from some agent's perspective. For example, the valuation function v in Figure 2 evaluates the piece $X' = [0, 2/6]$ with $v(X') = 3/17$.

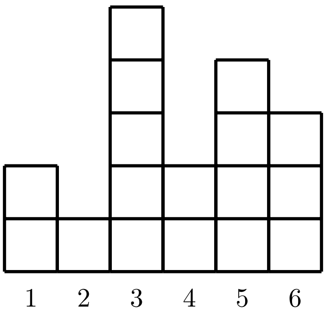

Figure 2: Common representation for a valuation function in cake-cutting.

As previously defined and common in cake-cutting, the valuation function's domain is defined as $\mathcal{P}(X) = \{X' \mid X' \subseteq X\}$, i.e., the power set of X. Consequently, v must be able to evaluate every *piece of cake in $\mathcal{P}(X)$.*

Let us define the relation $*$ as follows. We say that two real numbers $x, y \in \mathbb{R}$ satisfy the relation $*$ if and only if $x - y \in \mathbb{Q}$. One can show that $*$ is an equivalence relation and hence, we can define the equivalence classes $[x] = \{y \in \mathbb{R} \mid x * y\} \subseteq \mathbb{R}$, partitioning $\mathbb{R}$ disjointly. Using the axiom of choice, we can construct the set $V \subset X$, which contains exactly one representative in $X = [0, 1]$ of every equivalence class $[x]$. A set like V is called a Vitali set. *Since $V \in \mathcal{P}(X)$ holds, it follows from the definition of our valuation function that we must be able to calculate $v(V)$. Obviously, this is not possible and hence, v is not a valid valuation function in a strictly formal sense.*

To sum up, both the Cantor-like pieces of cake and the Vitali sets show that there are many possible pieces of cake contained in $\mathcal{P}(X)$ which we cannot evaluate with our common valuation functions: either the valuation function's domain $\mathcal{P}(X)$ is too large or our valuation functions are too simplistic. Sure enough, authors proposing cake-cutting protocols often abstain from making formal assumptions or from formalizing their model in detail. For example, Brams et al. [17, p. 553] write:

> "Many feel that the informality adds to the subject's simplicity and charm, and we would concur. But charm and simplicity are not the only factors determining the direction in which mathematics moves or should move. Our analysis in this paper raises several issues that may only admit a resolution via some negative results. While such results may not require complete formalization of what is

permissible, they do appear to require partial versions. We will refer to such partial limitations as theses."

It would thus be desirable to have some common consensus on which models are useful for which purpose and which are not. Therefore, we will have a closer look at the cake-cutting foundations through the lens of measure theory.

2 Preliminaries

Let $X = [0, 1]$ denote a cake and define $\mathcal{X} \subseteq \mathcal{P}(X)$ to be the set of all possible pieces of X that (a) can be allocated to some players via a cake-cutting protocol, and (b) can be evaluated by the players using their valuation functions. We start by formulating requirements for $\mathcal{X}$ regarding the possible pieces of cake. Obviously, we want to be able to allocate the complete cake X as well as an empty piece $\emptyset$ to a player and therefore, $X \in \mathcal{X}$ and $\emptyset \in \mathcal{X}$ must hold. If $A \subseteq X$ is already allocated to some player, i.e., $A \in \mathcal{X}$, then we want to be able to give the remainder of the cake to another player; so for all $A \in \mathcal{X}$, we demand $\overline{A} = X \setminus A \in \mathcal{X}$. Furthermore, we want to be able to cut and combine pieces of cake; so for all $A, B \in \mathcal{X}$, we require $A \cap B \in \mathcal{X}$. Note that $A \cup B = \overline{\overline{A} \cap \overline{B}}$ and $A \setminus B = A \cap \overline{B}$, so our previously formulated requirements also allow us to join a finite number of pieces of cake and to calculate the difference of two pieces of cake.

Definition 2.1. *Let $\mathbb{S}$ be a set. A set $\mathcal{A} \subseteq \mathcal{P}(\mathbb{S})$ is called an **algebra** over $\mathbb{S}$ if $\emptyset \in \mathcal{A}$, $\mathbb{S} \in \mathcal{A}$, and for all $A, B \in \mathcal{A}$, it holds that $A \setminus B \in \mathcal{A}$ and $A \cup B \in \mathcal{A}$.*

It is worth noting that only by the formulation of intuitive requirements with respect to the set of all possible pieces of cake, we ended up with a well-studied, structured set from math: an *algebra*. It is easy to check that, for instance, $\{\emptyset, X\} \subseteq \mathcal{P}(X)$ is an algebra. However, it is also obvious that this algebra is useless for our purpose, as then only two possible pieces can be allocated, the complete cake and an empty piece. At the other extreme, we might also take $\mathcal{P}(X)$ as our set for the possible pieces of X, which is an algebra, too. However, when choosing $\mathcal{X}$, we must also ensure that meaningful valuation functions can exist for this set. Therefore, we now list requirements with respect to the players' valuation functions.

A valuation function v shall evaluate an arbitrary piece of cake $Z \in \mathcal{X}$ with a nonnegative real number, i.e., $v\colon \mathcal{X} \to [0, \infty]$. However, in order to normalize the players' valuations and keep them comparable, we demand that $v(\emptyset) = 0$ and $v(X) = 1$ hold. Hence, we can further limit the valuation function's image to $[0, 1]$, i.e., we can even write $v\colon \mathcal{X} \to [0, 1]$. The next definition lists further requirements for a valuation function.

Definition 2.2. *Let X be a cake. We denote a player's* valuation function *by $v\colon \mathcal{X} \to [0, 1]$ and demand it to satisfy the following axioms:*

(N) Normalization*: $v(\emptyset) = 0$ and $v(X) = 1$ must hold.*

(A) Additivity*: For all $A, B \in \mathcal{X}$, it must hold that $v(A \cup B) = v(A) + v(B) - v(A \cap B)$.*

(D) Divisibility*: For every $A \in \mathcal{X}$ and for every λ, $0 \leq \lambda \leq 1$, there must exist some $A' \in \mathcal{X}$ with $A' \subseteq A$ such that $v(A') = \lambda v(A)$.*

Since we want $\mathcal{X}$ to be an algebra, (N) and (A) imply that a valuation function v shall be a finite content on $\mathcal{X}$.

Definition 2.3. *Let $\mathbb{S}$ be a set and $\mathcal{A} \subseteq \mathcal{P}(\mathbb{S})$ an algebra over $\mathbb{S}$. A function $\mu\colon \mathcal{A} \to [0,\infty]$ is called a **content** if $\mu(\emptyset) = 0$ and*

$$\mu\left(\bigcup_{i=1}^{n} A_i\right) = \sum_{i=1}^{n} \mu(A_i)$$

holds for any choice of finitely many pairwise disjoint $A_1, \dots, A_n \in \mathcal{A}$.

Requirement (D) demands even more from v than just to be a content on $\mathcal{X}$. Specifically, (D) requires every valuation function to be *atom-free*, i.e., for every $x \in X$ it must hold that $v(\{x\}) = 0$ is true.

Lemma 2.4. *Every valuation function $v\colon \mathcal{X} \to [0,1]$ which satisfies* (D) *is also atom-free.*

Proof. Since there is no proper nonempty subset of a singleton $\{x\}$, it is clear that the requirement of assumption (D), $v(B) = \lambda v(\{x\})$, can only hold for $\lambda = 0$ (and $v(B) = 0$) or $\lambda = 1$ and $B = \{x\}$. $\square$

Besides, we require valuation functions to satisfy *continuity*, a property that is crucial for so-called moving-knife cake-cutting protocols to work.

Definition 2.5. *Let $v\colon \mathcal{X} \to [0,1]$ be a valuation function. We say that v is **continuous** if for all a and b with $0 \le a < b \le 1$ satisfying $v([0,a]) = \alpha$ and $v([0,b]) = \beta$ and for every $\gamma \in [\alpha,\beta]$, there exists some $c \in [a,b]$ such that $v([0,c]) = \gamma$.*

In other words, a valuation function v is continuous if its distribution function $F : [0,1] \to [0,1]$, given by $F(t) = v([0,t])$, is a continuous function. Note that a continuous valuation function on an algebra cannot have atoms as $v(\{x\}) = v([0,x] \setminus [0,x)) = F(x) - F(x-) = 0$, with the left-hand limit $F(x-) = \lim_{y \uparrow x} F(y)$.

2.1 Banach Limits

We will need a nonconstructive way to extend linear maps. The key result is the standard Hahn–Banach theorem, which is well-known from functional analysis (see, e.g., Rudin [40, Theorem 3.2]).

Theorem 2.6. *Assume that $(X, \|\cdot\|)$ is a normed vector space and $L : M \to \mathbb{R}$ a linear functional, which is defined on a linear subspace $M \subset X$ satisfying $|Lx| \le \kappa\|x\|$ for all $x \in M$ with a universal constant $\kappa = \kappa_L \in (0,\infty)$. Then there is an extension $\hat{L} : X \to \mathbb{R}$ such that $\hat{L}$ is again linear and satisfies $|\hat{L}x| \le \kappa\|x\|$ for all $x \in X$ with the same constant $\kappa = \kappa_L$ as before.*

With a little more effort, but essentially the same proof, we can replace the norm $\|x\|$, resp. $\kappa\|x\|$, by a general sublinear map $p : X \to \mathbb{R}$. Sublinear means that $p(\lambda x) = \lambda p(x)$ and $p(x+y) \le p(x) + p(y)$ for all $x, y \in X$ and $\lambda \ge 0$. In this case, the extension of $Lx \le p(x)$ satisfies $-p(-x) \le \hat{L}x \le p(x)$. Note that p is only positively homogeneous, i.e., $-p(-x) \ne p(x)$.

The proof is nonconstructive and, at least for nonseparable spaces X, relies on the axiom of choice.

We will use the Hahn–Banach theorem for the space of bounded sequences $\ell^\infty([0,\infty)) = \{x = (x_n)_{n\in\mathbb{N}} \subset [0,\infty) \mid \|x\|_\infty < \infty\}$, where $\|x\|_\infty = \sup_{n\in\mathbb{N}} x_n$ is the uniform norm.

A prime example of a bounded linear functional is the limit: Consider those $x = (x_n)_{n\in\mathbb{N}} \in \ell^\infty([0,\infty))$ where $L(x) := \lim_{n\to\infty} x_n = x$ exists in the usual sense. It is common to write $c([0,\infty)) = \{x \in \ell^\infty([0,\infty)) \mid \lim_{n\to\infty} x_n \text{ exists}\}$. Clearly, $\lim_{n\to\infty} x_n =$

$\limsup_{n\to\infty} x_n \le \sup_{n\in\mathbb{N}} x_n$, so that L is a bounded linear functional on $M = c([0,\infty)) \subset X = \ell^\infty([0,\infty))$, and we can extend it to all of X as the **Banach limit**, i.e.,

$$\underset{n\to\infty}{\mathrm{LIM}}\, x_n := \begin{cases} \lim_{n\to\infty} x_n & \text{if } x \in c([0,\infty)), \\ \hat{L}(x) & \text{if } x \in \ell^\infty([0,\infty)) \setminus c([0,\infty)). \end{cases}$$

Using the addition to the Hahn–Banach theorem with $p(x) := \limsup_{n\to\infty} x_n$ and the observation that $\lim_{n\to\infty} x_n$ exists if, and only if, $\liminf_{n\to\infty} x_n = \limsup_{n\to\infty} x_n \in [0,\infty)$, we can choose the extension $\hat{L}$ in such a way that

$$\liminf_{n\to\infty} x_n \le \underset{n\to\infty}{\mathrm{LIM}}\, x_n \le \limsup_{n\to\infty} x_n.$$

2.2 Measure Theory

We now introduce some basics from measure theory, which we need in the subsequent discussion of the cake-cutting literature. For more background of measure theory, in much more detail and greater depth, we refer to the monographs [42] and [43].

First of all, let $\mathbb{S} \subseteq \mathbb{R}$ be an interval in the real numbers. Then we denote by

- $\mathcal{C}(\mathbb{S}) = \{[a,b] \mid a,b \in \mathbb{S}, a \le b\}$ the *set of all closed intervals over* $\mathbb{S}$,
- $\mathcal{O}(\mathbb{S}) = \{(a,b) \mid a,b \in \mathbb{S}, a \le b\}$ the *set of all open intervals over* $\mathbb{S}$,
- $\mathcal{LO}(\mathbb{S}) = \{(a,b] \mid a,b \in \mathbb{S}, a < b\}$ the *set of all left-open intervals over* $\mathbb{S}$,
- $\mathcal{RO}(\mathbb{S}) = \{[a,b) \mid a,b \in \mathbb{S}, a < b\}$ the *set of all right-open intervals over* $\mathbb{S}$, and
- $\mathcal{I}(\mathbb{S}) = \mathcal{C}(\mathbb{S}) \cup \mathcal{O}(\mathbb{S}) \cup \mathcal{LO}(\mathbb{S}) \cup \mathcal{RO}(\mathbb{S})$ the *set of all intervals over* $\mathbb{S}$.

In Definition 2.1 we defined the notion of algebra. However, in measure theory one is typically working with the notion of a σ*-algebra* instead of an algebra.

Definition 2.7. *Let $\mathbb{S}$ be a set. A subset $\mathcal{A} \subseteq \mathcal{P}(\mathbb{S})$ is called a* ***σ-algebra*** *over $\mathbb{S}$ if $\mathcal{A}$ is an algebra over $\mathbb{S}$ and for all sequences $(A_n)_{n\in\mathbb{N}}$ with $A_n \in \mathcal{A}$, $\bigcup_{n\in\mathbb{N}} A_n$ is in $\mathcal{A}$, too.*

We now explain what we mean by a σ*-operator on some set.* Given $\mathbb{S} \subseteq \mathbb{R}$, we write $\sigma(\mathbb{S})$ to denote the *smallest* σ-algebra containing $\mathbb{S}$; clearly, $\mathbb{S} \subset \sigma(\mathbb{S})$. Note that $\sigma(\mathbb{S})$ is well-defined since $\mathcal{P}(X)$ is a σ-algebra and the intersection of an arbitrary number of σ-algebras is again a σ-algebra.

The next lemma is a standard result from set theory (see, e.g., [42]).

Lemma 2.8. *For some set $\mathbb{S} \subseteq \mathbb{R}$, it holds that*

$$\sigma(\mathcal{O}(\mathbb{S})) = \sigma(\mathcal{C}(\mathbb{S})) = \sigma(\mathcal{LO}(\mathbb{S})) = \sigma(\mathcal{RO}(\mathbb{S})) = \sigma(\mathcal{I}(\mathbb{S})).$$

This set, $\sigma(\mathcal{O}(\mathbb{S}))$, is the smallest σ-algebra containing all open intervals from $\mathbb{S}$ and thus has some special name.

Definition 2.9. *For a set $\mathbb{S} \subseteq \mathbb{R}$, we denote by $\mathcal{B}(\mathbb{S}) = \sigma(\mathcal{O}(\mathbb{S}))$ the smallest σ-algebra on $\mathbb{S}$ containing all open intervals from $\mathbb{S}$ and call it the* ***Borel σ-algebra*** *over $\mathbb{S}$.*

The following definition is taken (and slightly adjusted) from [42, Definition 4.1].

Definition 2.10. *Let $\mathbb{S}$ be some set and $\mathcal{A}$ a σ-algebra on $\mathbb{S}$. A (positive)* ***measure*** *μ on $\mathbb{S}$ is a map $\mu\colon \mathcal{A} \to [0,\infty]$ satisfying* *(i)* $\mu(\emptyset) = 0$ *and*

(ii) for every pairwise disjoint sequence of elements $(A_n)_{n\in\mathbb{N}}$ *in* $\mathcal{A}$*, it holds that*

$$\mu\left(\bigcup_{n\in\mathbb{N}} A_n\right) = \sum_{n\in\mathbb{N}} \mu(A_n).$$

Property (ii) is called $\boldsymbol{\sigma}$***-additivity****.*

The idea behind Lebesgue measure is to have a set-function $A \mapsto \lambda(A)$ in $\mathbb{R}^n$ with all properties of the familiar "volume" from geometry; in particular, we want a volume that is invariant under shifts and rotations. This motivates and explains the additivity requirement (A). Having in mind triangulations of more complicated shapes, it also makes sense to require countable σ-additivity, i.e., the property that for countably many disjoint sets $(A_n)_{n\in\mathbb{N}}$ with $A = \bigcup_{n\in\mathbb{N}} A_n$, we have $\lambda(A) = \sum_{n\in\mathbb{N}} \lambda(A_n)$. The restriction to countable unions is natural, as we exhaust a given shape by nontrivial sets A_n, having nonempty interior: Each of them contains a rational point $q \in \mathbb{Q}^n$, hence there are at most countably many non-overlapping A_n. In order to construct Lebesgue measure, we prescribe an elementary volume on "easy" shapes, say n-dimensional rectangles, by

$$\lambda(R) = \prod_{i=1}^{n}(b_i - a_i) \quad \text{if} \quad R = (a_1, b_1] \times \cdots \times (a_n, b_n]$$

in the hope that we can exhaust all possible sets $B \subset \mathbb{R}^n$ by (necessarily at most countable many) non-overlapping rectangles. At this point we encounter a problem: General sets $A \subset \mathbb{R}^n$ are way too complicated to get a well-defined and unique extension of λ from the rectangles to $\mathcal{P}(\mathbb{R}^n)$. The way out is the notion of measurable sets and Carathéodory's extension theorem (see further down).

This works as follows: In view of the σ-additivity property of λ, it makes sense to consider a family $\mathcal{A} \subset \mathcal{P}(\mathbb{R}^n)$ which contains the rectangles and is closed under the formation of countable unions of its members and the formation of complements (this allows us to consider differences and intersections of sets). Thus, we naturally arrive at the notion of the (Borel) σ-algebra as the natural domain of Lebesgue measure. Unfortunately, there are so many Borel sets that we cannot build them constructively from rectangles – we would need transfinite induction for this – and this is one of the reasons why cutting a cake is not always a piece of cake.

The question of whether *every* set $A \subset \mathbb{R}^n$ has a unique "geometric" volume (in the above sense) is dimension-dependent. If $n = 1$ or $n = 2$, we can extend the notion of length and area to all sets, but not in a unique way. In dimension 3 and higher, we'll end up with contradictory statements ("Banach–Tarski paradox") if we try to have an additive content for all sets. This conundrum can be resolved by looking at the Borel sets or the Lebesgue sets – these are the Borel sets enriched by all subsets of Borel sets with Lebesgue measure zero.

A good reference and introduction to measure theory are the monographs [42] and [43].

3 Five Possible Definitions from the Literature

In the cake-cutting literature, a wide variety of different definitions have been used for the set $\mathcal{X}$ of possible pieces of cake. We first collect the most commonly used definitions used for $\mathcal{X}$, along with the corresponding references and discuss them in detail. Then we show several relations among these definitions and discuss what this implies for a most reasonable choice of $\mathcal{X}$.

The set $\mathcal{X}$ containing all possible pieces over a cake X can be defined as

1. all finite unions of intervals from X, i.e., $\mathcal{I}(X)^n = \{\bigcup_{i=1}^n I_i \mid I_i \in \mathcal{I}(X), n \in \mathbb{N}\}$,
2. all countable unions of intervals from X, i.e., $\mathcal{I}(X)^{\mathbb{N}} = \{\bigcup_{i \in \mathbb{N}} I_i \mid I_i \in \mathcal{I}(X)\}$,
3. the Borel σ-algebra over X, i.e., $\mathcal{B}(X)$,
4. the set of all Lebesgue-measurable sets over X, i.e., $\mathcal{L}(X)$, or
5. the power set $\mathcal{P}(X)$ of X.

Assuming $\mathcal{X} = \mathcal{I}(X)^n$ is common among papers that consider only finite cake-cutting protocols. Such protocols can only make a finite number of cuts, thus producing a finite set of contiguous pieces, i.e., intervals, to be evaluated by the players. Authors that make this assumption and use $\mathcal{X} = \mathcal{I}(X)^n$ include Woeginger and Sgall [52], Stromquist [47], Lindner and Rothe [31], Procaccia [35], Walsh [50], Cohler et al. [27], Bei et al. [9], Cechlárová and Pillárová [24], Brams et al. [10], Cechlárová et al. [22], Chen et al. [25], Brânzei and Miltersen [19], Aziz and Mackenzie [6], Edmonds and Pruhs [30], and Aziz and Mackenzie [5].

As a special case, valuation functions may even be restricted to single intervals, which is done by Cechlárová and Pillárová [23] and Aumann and Dombb [2]. Even though the restriction to finite unions of intervals is sensible from a practical perspective, it may artificially constrain results that could hold also in a more general setting.

Brânzei et al. [20] extend $\mathcal{X}$ to contain countably infinite unions of intervals, i.e., $\mathcal{I}(X)^{\mathbb{N}}$.

Authors assuming $\mathcal{X} = \mathcal{B}(X)$ include Stromquist and Woodall [48], Deng et al. [28], and Segal-Halevi et al. [45].

Works using $\mathcal{X} = \mathcal{L}(X)$ include Reijnierse and Potters [37], Arzi et al. [1], and Robertson and Webb [38]. Additionally, several authors do not explicitly make the assumption $\mathcal{X} = \mathcal{L}(X)$, but they define valuation functions based on (Lebesgue-)measurable sets only, most prominently, a valuation function is often defined as the integral of a given probability density function on X. This or a similar assumption is made by Brams et al. [11, 14, 12, 13], Robertson and Webb [39], Webb [51], Aumann et al. [3], Brânzei et al. [18], and Caragiannis et al. [21].

Papers that assume $\mathcal{X} = \mathcal{P}(X)$ include those by Maccheroni and Marinacci [33], Sgall and Woeginger [46], Saberi and Wang [41], Manabe and Okamoto [34], and Aumann et al. [4].

Finally, several works, including Dubins and Spanier [29], Barbanel [8], Zeng [53], Brams and Taylor [15], and Barbanel [7], define the set of pieces of cake to be any (σ-)algebra (not necessarily Borel) over X.

Note that each of the sets $\mathcal{I}(X)^n$, $\mathcal{B}(X)$, $\mathcal{L}(X)$, and $\mathcal{P}(X)$ is an algebra over X, and all, except $\mathcal{I}(X)^n$, are also σ-algebras over X. That $\mathcal{I}(X)^n$ is an algebra is shown in Lemma 4.4 below. However, $\mathcal{I}(X)^{\mathbb{N}}$ is not an algebra, as the proof of the following theorem shows.

Having introduced all the different approaches currently used in the literature, we will now prove the strict inclusions among these sets stated in the following theorem.

Theorem 3.1. $\mathcal{I}(X)^n \overset{(a)}{\subsetneq} \mathcal{I}(X)^{\mathbb{N}} \overset{(b)}{\subsetneq} \mathcal{B}(X) \overset{(c)}{\subsetneq} \mathcal{L}(X) \overset{(d)}{\subsetneq} \mathcal{P}(X)$.

Proof. We start with proving (a): $\mathcal{I}(X)^n \subsetneq \mathcal{I}(X)^{\mathbb{N}}$. Obviously, $\mathcal{I}(X)^n \subseteq \mathcal{I}(X)^{\mathbb{N}}$ is true, as every finite union of intervals is a countable union of intervals. To see that the two sets are not equal, look at $I = \bigcup_{i=0}^{\infty}[3/2^{i+2}, 1/2^i]$. It is clear that $I \in \mathcal{I}(X)^{\mathbb{N}}$ is true, as I is a countable union of intervals. However, it holds that $I = [3/4, 1] \cup [3/8, 1/2] \cup \cdots$, i.e., I

cannot be written as a finite union of intervals, as all these subintervals are pairwise distinct. Hence, $I \notin \mathcal{I}(X)^n$, so $\mathcal{I}(X)^n \subsetneq \mathcal{I}(X)^{\mathbb{N}}$.

In Section 2 we defined $\mathcal{B}(X) = \sigma(\mathcal{O}(X))$ and stated as Lemma 2.8 the result that $\sigma(\mathcal{O}(X)) = \sigma(\mathcal{C}(X)) = \sigma(\mathcal{LO}(X)) = \sigma(\mathcal{RO}(X)) = \sigma(\mathcal{I}(X))$ holds. Consequently, as $\mathcal{E} \subseteq \sigma(\mathcal{E})$ holds by definition for any set $\mathcal{E}$, we immediately obtain $\mathcal{I}(X) \subseteq \mathcal{B}(X)$. Furthermore, as $\mathcal{B}(X)$ is a σ-algebra, it is closed under countable unions and we thus obtain $\mathcal{I}(X)^{\mathbb{N}} \subseteq \mathcal{B}(X)$. However, since $\mathbb{Q} \cap X \in \mathcal{B}(X)$ is true, as $\mathbb{Q} \cap X$ can be written as a countable union of intervals, each containing one element, it must hold that $\overline{\mathbb{Q} \cap X} \in \mathcal{B}(X)$ by the definition of a σ-algebra. However, the irrational numbers $\overline{\mathbb{Q} \cap X}$ in X cannot be written as a countable union of intervals, since every interval containing more than one element immediately contains a rational number. Therefore, $\mathcal{I}(X)^{\mathbb{N}}$ is not an algebra and $\mathcal{I}(X)^{\mathbb{N}} \neq \mathcal{B}(X)$ holds, proving (b).

The inclusion $\mathcal{B}(X) \subseteq \mathcal{L}(X)$ holds by definition, as all Borel sets are Lebesgue-measurable. However, there are Lebesgue-measurable sets that are not Borel-measurable. This follows from the observation that the cardinality of $\mathcal{L}(X)$ is the cardinality of $\mathcal{P}(X)$ (which is $2^{\mathfrak{c}} > \mathfrak{c}$), whereas there are only continuum-many (i.e., $\mathfrak{c}$, the cardinality of X) Borel sets, see [42, Appendix G, Corollary G.7]. This proves (c).

Finally, the power set $\mathcal{P}(X)$ trivially contains all other families of sets considered earlier. Nevertheless, there are sets in $\mathcal{P}(X)$ that are not Lebesgue-measurable, for example the Vitali set that we introduced in Section 1, so $\mathcal{L}(X) \neq \mathcal{P}(X)$ and we have (d). □

4 Discussion

Taking $\mathcal{X} = \mathcal{I}(X)^{\mathbb{N}}$ is problematic, as $\mathcal{I}(X)^{\mathbb{N}}$ is not even an algebra and thus does not even satisfy the minimum requirements for $\mathcal{X}$ as described in the first paragraph of Section 2. By contrast, taking $\mathcal{X} = \mathcal{I}(X)^n$ is, of course, fine; this choice provides a sufficient range of sets to be used with common finite cake-cutting protocols.

From a theoretical point of view, however, this choice may be unnecessarily restrictive, especially in light of the fact that we also want to use infinite cake-cutting protocols. Therefore, a larger set $\mathcal{X}$ may be desirable, perhaps even larger than $\mathcal{I}(X)^{\mathbb{N}}$, which (as we have seen) has disqualified itself.

We start our discussion by explicating why $\mathcal{X} = \mathcal{P}(X)$ is a bad choice and we then provide arguments for a better option, namely $\mathcal{X} = \mathcal{B}(X)$.

4.1 Taming $\mathcal{P}(X)$ with Exotic Contents via Banach Limits

If one boldly desires to define valuation functions on the set $\mathcal{P}(X)$ of all subsets of the cake, it remains to be shown that this indeed is possible. We have seen that the commonly used valuation functions represented via blocks, as depicted in Figure 2, are not capable of evaluating every piece of cake in $\mathcal{P}(X)$. Hence, in this subsection we aim to define a valuation function capable of evaluating *every* possible piece of cake in $\mathcal{P}(X)$.

In other words, we formally define a finitely additive content μ on $\mathcal{P}(X)$ satisfying all requirements from Definition 2.2. To do so, in a first step, we must choose an arbitrary sequence $(x_i)_{i\in\mathbb{N}}$ of pairwise distinct elements from X. For every $A \subseteq X$, we define a mapping $f_A \colon \mathbb{N} \to [0, 1]$ with

$$n \mapsto \frac{|A \cap \{x_1, \ldots, x_n\}|}{n},$$

($|B|$ denotes the cardinality of the set B), i.e., $f_A(n)$ describes the relative frequency of the first n elements of $(x_i)_{i\in\mathbb{N}}$ being in A. For some sets A the limit $\lim_{n\to\infty} f_A(n)$ does exist.

We can, however, use the Banach limits introduced in Section 2.1 to define

$$\mu\colon \mathcal{P}(X) \to [0,1], \quad A \mapsto \operatorname*{LIM}_{n\to\infty} f_A(n).$$

It is clear that $\mu(A)$ is additive since both the limit and the Banach limit are additive, so (A) is satisfied. In the following lemma we show why μ satisfies (D).

Lemma 4.1. *For every $A \in \mathcal{P}(X)$ with $\mu(A) > 0$ and every $\lambda \in [0,1]$, there exists a subset $B \subseteq A$ in $\mathcal{P}(X)$ such that $\mu(B) = \lambda\mu(A)$.*

Proof. If $\mu(A) > 0$ then A must contain an infinite number of points of the underlying sequence, say $A \cap \{x_1, x_2, \dots\} = \{x_{i(1)}, x_{i(2)}, \dots\}$ for some increasing sequence $(i(k))_{k\in\mathbb{N}}$ of integers. In order to keep notation simple, we write $x_{i(k)} = x'_k$. By assumption, $\mu(A) = \lim_{n\to\infty} |\{x'_1, \dots, x'_n\}|/n$. We have to construct a set $B \in \mathcal{P}(A)$ such that $\lim_{n\to\infty} f_B(n) = \lambda\mu(A)$ for fixed $\lambda \in [0,1]$.

The key observation in this proof is the fact that for any nonnegative rational number k/n with $k < n$, we have

$$\frac{k}{n+1} < \frac{k}{n} < \frac{k+1}{n+1},$$

i.e., the quantity $f_B(n) = |B \cap \{x'_1, \dots, x'_n\}|/n$ decreases if we jack up $n \to n+1$ and $x'_{n+1} \notin B$, and increases if we jack up $n \to n+1$ and $x'_{n+1} \in B$.

Fix $\lambda \in [0,1]$ and observe that we can assume $0 < \lambda < 1$: If $\lambda = 0$, we take $B = \emptyset$, and for $\lambda = 1$, we use $B = \{x'_n \mid n \in \mathbb{N}\}$. For $\lambda \in (0,1)$, we use a recursive approach.

Define $B_1 = \{x'_1\}$ with $f_{B_1}(1) = 1$ and assume that we have already found a set B_n such that $f_{B_n}(n) \geq \lambda\mu(A)$. Because of the observation at the beginning of the proof, the numbers

$$\begin{aligned}
\ell_{n+1} &:= \min\{k > n \mid f_{B_n}(k) \leq \lambda\mu(A)\} \text{ and} \\
u_{n+1} &:= \min\left\{k > \ell_{n+1} \mid f_{A_k}(k) \geq \lambda\mu(A) \text{ for } A_k = B_n \cup \{x'_{\ell_{n+1}}, x'_{\ell_{n+1}+1}, \dots, x'_k\}\right\}
\end{aligned}$$

are well-defined and satisfy $\ell_n < u_n < \ell_{n+1} < u_{n+1}$ and $\ell_{n+1} \to \infty$. Setting

$$B_{n+1} = A \cup \{x'_{\ell_{n+1}}, x'_{\ell_{n+1}+1}, \dots, x'_{u_{n+1}}\}$$

finishes the recursion, and we can define $B = \bigcup_{n\in\mathbb{N}} B_n$.

By construction, $|f_B(n) - \lambda\mu(A)| \leq \ell_{n+1}^{-1}$ holds, completing this proof. □

We now provide a counterexample that shows that μ is not σ-additive. We do this to emphasize that μ, as earlier defined, indeed is a finite content on $\mathcal{P}(X)$, which only satisfies all requirements defined earlier but not more. To do so, we define $A_0 = X \setminus \{x_i \mid i \in \mathbb{N}\}$ and $A_i = \{x_i\}$ for $i \in \mathbb{N}$. Obviously, for all $j \in \mathbb{N} \cup \{0\}$, it holds that $\mu(A_j) = 0$, while at the same time we have

$$\mu\left(\bigcup_{j\in\mathbb{N}\cup\{0\}} A_j\right) = \mu(X) = 1,$$

which means that μ is not σ-additive.

Having this content, it seems to be an attractive valuation function for cake-cutting. We can interpret the sequence $(x_i)_{i\in\mathbb{N}}$ as countably many points, which are used to evaluate arbitrary pieces of the cake. However, multiple drawbacks exist. First of all, the existence of a Banach limit is only guaranteed if one is willing to accept the validity of the axiom of choice, as already mentioned in Section 2.1. Furthermore, until now no explicit nontrivial example of a Banach limit is known. Hence, we cannot calculate $\mu(A)$ for $A \in \mathcal{P}(X)$ if the ordinary limit of $f_A(n)$ does not exist, as we do not know what the Banach limit looks like.

Thus, although μ is theoretically capable of evaluating all pieces of cake in $\mathcal{X} = \mathcal{P}(X)$, it is actually not useful for us. Besides the previously listed mathematical problems, there are also practical problems related to cake-cutting itself. If we would use μ as a valid valuation function in cake-cutting, all players would be obliged to precisely define a countable sequence $(x_i)_{i\in\mathbb{N}}$ of pairwise distinct elements in X and some Banach limit they are using for their valuation functions. When we think of common approaches and results in the cake-cutting literature, this approach seems simply impractical and not feasible to use.

Hence, finding practically usable valuation functions defined on $\mathcal{P}(X)$ seems to remain an open problem. Nevertheless, this section showed that defining more complex valuation functions (compared to the valuation functions represented via boxes) does not solve our initial problem. Therefore, in the next section we discuss an alternative solution, namely, reducing $\mathcal{X}$ in size from $\mathcal{P}(X)$ to a smaller subfamily contained in $\mathcal{P}(X)$.

4.2 Borel σ-Algebra

We suggest to use $\mathcal{X} = \mathcal{B}(X)$ as set of all possible pieces of cake. As shown in Theorem 3.1, the Borel σ-algebra $\mathcal{B}(X)$ (strictly) contains $\mathcal{I}(X)^n$ as well as $\mathcal{I}(X)^{\mathbb{N}}$ but is strictly smaller than $\mathcal{P}(X)$.

In general, the Borel σ-algebra can become quite large and complicated if the base set is not countable, as is the case for $X = [0,1] \subset \mathbb{R}$. This means that in general we cannot explicitly assign a value $\mu(A)$ to every element A of a σ-algebra for a measure μ, as the σ-algebra is simply too large. Instead, one can use Carathéodory's extension theorem as a solution.

Recall that a set function $\mu\colon \Sigma \to [0,1]$ is called *σ-additive relative to* Σ if for any sequence of pairwise disjoint elements $(A_n)_{n\in\mathbb{N}} \subset \Sigma$ such that $\bigcup_{n\in\mathbb{N}} A_n \in \Sigma$, we have

$$\mu\left(\bigcup_{n\in\mathbb{N}} A_n\right) = \sum_{n\in\mathbb{N}} \mu(A_n).$$

For a σ-algebra, the condition "$\bigcup_{n\in\mathbb{N}} A_n \in \Sigma$" is trivially satisfied.

Theorem 4.2 (Carathéodory's extension theorem, cf. [42, Remark following Theorem 5.7, Theorem 6.1])**.** *Let Σ be an algebra over some set $\mathbb{S}$ and consider a finitely additive content $\mu\colon \Sigma \to [0,1]$. If μ is σ-additive relative on Σ, then μ can be extended to a measure $\hat{\mu}$ defined on $\sigma(\Sigma)$ such that for all $A \in \Sigma$, we have $\mu(A) = \hat{\mu}(A)$. If μ is finite or σ-finite, then the extension to $\hat{\mu}$ is unique.*

Let us show here that the block-based valuation functions are σ-additive contents on $\mathcal{I}(X)^n$ and that $\mathcal{I}(X)^n$ is an algebra. In this case, we can use the previous theorem to extend the valuation functions to measures on $\sigma(\mathcal{I}(X)^n) = \mathcal{B}(X)$. Since the valuation functions are σ-finite, it follows that this extension is unique. Hence, by providing a block-based valuation function, we obtain a unique measure on $\mathcal{B}(X)$. Thus $\mathcal{X} = \mathcal{B}(X)$ is a good solution to our problem.

Let us formalize the box-based valuation functions. A box-based valuation function μ partitions the complete cake X into a finite number of pairwise disjoint subintervals, where each subinterval is allocated a finite number of boxes of equal height. We denote the set of all subintervals which μ uses by

$$\mathcal{I}_\mu = \{I_1 = [a_1, b_1), \ldots, I_{n-1} = [a_{n-1}, b_{n-1}), I_n = [a_n, b_n]\},$$

where $\bigcup_{i=1}^{n} I_i = X$ and we have $I_i \cap I_j = \emptyset$ for all i and j, $1 \leq i < j \leq n$. Furthermore, denote by $\psi_i \in \mathbb{N}$, for $1 \leq i \leq n$, the number of boxes allocated to an interval I_i by μ and denote by $\psi_\mu = \sum_{i=1}^{n} \psi_i$ the total number of boxes.

As previously shown, μ cannot be defined on $\mathcal{P}(X)$. Hence, we define μ only on $\mathcal{I}(X)^n$, i.e., on finite unions of intervals from X, that is, we can write $\mu \colon \mathcal{I}(X)^n \to [0,1]$. Now, for a piece of cake $X' \in \mathcal{I}(X)^n$, which is a finite union of pairwise disjoint intervals in X, we can write

$$X' \cap I_i = \bigcup_{j=1}^{n_{X'}^i} [c_j^i, d_j^i]$$

for $1 \leq i \leq n$. Consequently, we can formulate

$$\mu(X') = \frac{1}{\psi_\mu} \sum_{i=1}^{n} \left[\frac{\psi_i}{b_i - a_i} \sum_{j=1}^{n_{X'}^i} (d_j^i - c_j^i) \right]$$

and thus, compute the valuation of a piece X' by μ.

Example 4.3. *Referring back to the box-based valuation function ν from Figure 2, we obtain*

$$\mathcal{I}_\nu = \{I_1 = [0, 1/6), I_2 = [1/6, 2/6), \ldots, I_6 = [5/6, 1]\}.$$

Also, we have $\psi_1 = 2$, $\psi_2 = 1$, $\psi_3 = 5$, $\psi_4 = 2$, $\psi_5 = 4$, $\psi_6 = 3$, and $\psi_\nu = 17$. For $X' = [0, 2/6]$, we obtain

$$\begin{aligned}
\nu(X') &= \frac{1}{\psi_\nu} \sum_{i=1}^{6} \left[\frac{\psi_i}{b_i - a_i} \sum_{j=1}^{n_{X'}^i} (d_j^i - c_j^i) \right] \\
&= \frac{1}{17} \left(\frac{2}{1/6} \cdot (1/6 - 0) + \frac{1}{1/6} \cdot (2/6 - 1/6) + \frac{5}{1/6} \cdot 0 + \frac{2}{1/6} \cdot 0 + \frac{4}{1/6} \cdot 0 + \frac{3}{1/6} \cdot 0 \right) \\
&= \frac{3}{17}.
\end{aligned}$$

Now, let us prove that $\mathcal{I}(X)^n$ is an algebra over X.

Lemma 4.4. *$\mathcal{I}(X)^n$ is an algebra over X.*

Proof. Since $\mathcal{I}(X)^n = \{\bigcup_{i=1}^{n} I_i \mid I_i \in \mathcal{I}(X), n \in \mathbb{N}\}$, obviously, it holds that $\emptyset, X \in \mathcal{I}(X)^n$. Now, let $A, B \in \mathcal{I}(X)^n$. Then we can write $A = \bigcup_{i=1}^{n} I_i$ and $B = \bigcup_{j=1}^{m} I_j$. Furthermore, we have

$$A \cup B = \left(\bigcup_{i=1}^{n} I_i \right) \cup \left(\bigcup_{j=1}^{m} I_j \right) = \bigcup_{i=1}^{k} J_i \in \mathcal{I}(X)^n$$

for appropriate $J_i \in \mathcal{I}(X)$, $1 \leq i \leq k$, as the union of finitely many intervals can again be written as finite union of pairwise distinct intervals. Lastly, $A \setminus B \in \mathcal{I}(X)^n$ holds by the same argument as before. Therefore, $\mathcal{I}(X)^n$ is an algebra over X. $\square$

What is left is to show that a box-based valuation function μ is σ-finite and σ-additive. If that is the case, it follows by the automatic extension theorem of Carathéodory that μ can be uniquely extended to a measure on $\mathcal{B}(X)$.

By the normalization axiom (N), $\mu(X) = 1 < \infty$ always holds for a valuation function μ. Thus it immediately follows that μ is σ-finite. Furthermore, σ-additivity follows directly from the construction. In fact, the normalized box-based valuation function can be interpreted as a piecewise constant probability density with respect to the Lebesgue measure, which directly gives Carathéodory's extension.

To conclude, $\mathcal{B}(X)$ is recommended as a very good choice for $\mathcal{X}$, since this choice enables us to use box-based valuation functions and their extensions as measures.

5 Conclusion and Some Further Technical Remarks

Among the questions we have tried to answer are:

1. Which subsets of $[0, 1]$ should be considered as pieces of cake? Only finite unions of intervals or more general sets?

2. If valuation functions are considered as set-functions as studied in measure theory, should they be *σ-additive measures* or *finitely additive contents*?

A related interesting question that may be tackled in future work is:

3. Which continuity property should be used? Either *divisibility* (D) or *absolute continuity*[1] with respect to Lebesgue measure? (There should be a relation between both, i.e., if we decide to go with Borel sets then absolute continuity implies divisibility.) First steps to answer this question were taken by Schilling and Stoyan [44].

While we can define the Dirac and counting measures for all sets in $\mathcal{P}(X)$, there is no way to define a geometrically sensible notion of "volume" for all sets – if we accept the validity of the axiom of choice. (One can even show that the axiom of choice is equivalent to the existence of nonmeasurable sets, cf. Ciesielski [26, p. 55].)

For a measure we have that divisibility (D) is equivalent to atom-freeness. However, for a content this does not necessarily hold, i.e., there are contents that satisfy atom-freeness but not divisibility. And there are divisible contents that are not continuous in the above sense.

To conclude, we have surveyed the existing rich literature on cake-cutting algorithms and have identified the most commonly used choices of sets consisting of what is allowed as pieces of cake. After showing that these five most commonly used sets are distinct from each other, we have discussed them in comparison. In particular, we have argued that $\mathcal{P}(X)$ is too general to define a (practically or theoretically) useful measure on it. And finally, we have reasoned why we recommend the Borel σ-algebra $\mathcal{B}(X)$ as a very good choice.

Acknowledgments

This work was supported in part by DFG grants RO-1202/14-2, RO-1202/21-1 and SCHI-419/11-1 (DFG-NCN Beethoven Classic 3). The second author was a member of the PhD-programme "Online Participation," supported by the North Rhine-Westphalian funding scheme "Forschungskollegs."

[1] A finite measure μ is called *absolutely continuous with respect to Lebesgue measure* if every Lebesgue null set is also a μ null set; this is equivalent to the existence of an integrable density function $f \geq 0$ such that $\mu(B) = \int_B f(x)dx$ for all $B \in \mathcal{A}$. μ is said to be *continuous* if its distribution function $F_\mu(x) = \mu(-\infty, x]$ is continuous. Note that the notions of *continuity* and *atom-freeness* coincide in the one-dimensional case.

References

[1] O. Arzi, Y. Aumann, and Y. Dombb. Throw one's cake—and eat it too. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, pages 69–80. Springer, 2011.

[2] Y. Aumann and Y. Dombb. The efficiency of fair division with connected pieces. In *Proceedings of the 6th International Workshop on Internet & Network Economics*, pages 26–37. Springer-Verlag *Lecture Notes in Computer Science #6484*, December 2010.

[3] Y. Aumann, Y. Dombb, and A. Hassidim. Computing socially-efficient cake divisions. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 343–350. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[4] Y. Aumann, Y. Dombb, and A. Hassidim. Auctioning a cake: Truthful auctions of heterogeneous divisible goods. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*, pages 1045–1052. International Foundation for Autonomous Agents and Multiagent Systems, 2014.

[5] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science*, pages 416–427. IEEE, 2016.

[6] H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for four agents. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 454–464, 2016.

[7] J. Barbanel. Game-theoretic algorithms for fair and strongly fair cake division with entitlements. In *Colloquium Mathematicae*, volume 69, pages 59–73, 1996.

[8] J. Barbanel. Super envy-free cake division and independence of measures. *Journal of Mathematical Analysis and Applications*, 197(1):54–60, 1996.

[9] X. Bei, N. Chen, X. Hua, B. Tao, and E. Yang. Optimal proportional cake cutting with connected pieces. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1263–1269. AAAI Press, 2012.

[10] S. Brams, M. Feldman, J. Lai, J. Morgenstern, and A. Procaccia. On maxsum fair cake divisions. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence*, pages 1285–1291. AAAI Press, 2012.

[11] S. Brams, M. Jones, and C. Klamler. Perfect cake-cutting procedures with money. Technical report, mimeo, 2003.

[12] S. Brams, M. Jones, and C. Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006.

[13] S. Brams, M. Jones, and C. Klamler. Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4):353–367, 2008.

[14] S. Brams, M. Jones, and C. Klamler. N-person cake-cutting: There may be no perfect division. *The American Mathematical Monthly*, 120(1):35–47, 2013.

[15] S. Brams and A. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995.

[16] S. Brams and A. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution.* Cambridge University Press, 1996.

[17] S. Brams, A. Taylor, and W. Zwicker. A moving-knife solution to the four-person envy-free cake-division problem. *Proceedings of the American Mathematical Society*, 125(2):547–554, 1997.

[18] S. Brânzei, I. Caragiannis, D. Kurokawa, and A. Procaccia. An algorithmic framework for strategic fair division. Technical Report arXiv:1307.2225v2 [cs.GT], ACM Computing Research Repository (CoRR), July 2016.

[19] S. Brânzei and P. Miltersen. Equilibrium analysis in cake cutting. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 327–334. International Foundation for Autonomous Agents and Multiagent Systems, 2013.

[20] S. Brânzei, A. Procaccia, and J. Zhang. Externalities in cake cutting. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 55–61. AAAI Press, 2013.

[21] I. Caragiannis, J. Lai, and A. Procaccia. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, volume 22, page 127, 2011.

[22] K. Cechlárová, J. Doboš, and E. Pillárová. On the existence of equitable cake divisions. *Information Sciences*, 228:239–245, 2013.

[23] K. Cechlárová and E. Pillárová. A near equitable 2-person cake cutting algorithm. *Optimization*, 61(11):1321–1330, 2012.

[24] K. Cechlárová and E. Pillárová. On the computability of equitable divisions. *Discrete Optimization*, 9(4):249–257, 2012.

[25] Y. Chen, J. Lai, D. Parkes, and A. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.

[26] K. Ciesielski. How good is Lebesgue measure? *The Mathematical Intelligencer*, 11(2):54–58, 1989.

[27] Y. Cohler, J. Lai, D. Parkes, and A. Procaccia. Optimal envy-free cake cutting. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pages 626–631. AAAI Press, 2011.

[28] X. Deng, Q. Qi, and A. Saberi. On the complexity of envy-free cake cutting. Technical Report arXiv:0907.1334v1 [cs.GT], ACM Computing Research Repository (CoRR), July 2009.

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

[30] J. Edmonds and K. Pruhs. Cake cutting really is not a piece of cake. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 271–278, 2006.

[31] C. Lindner and J. Rothe. Degrees of guaranteed envy-freeness in finite bounded cake-cutting protocols. In *Proceedings of the 5th International Workshop on Internet and Network Economics*, pages 149–159. Springer, 2009.

[32] C. Lindner and J. Rothe. Cake-cutting: Fair division of divisible goods. In J. Rothe, editor, *Economics and Computation. An Introduction to Algorithmic Game Theory, Computational Social Choice, and Fair Division*, Springer Texts in Business and Economics, chapter 7, pages 395–491. Springer-Verlag, 2015.

[33] F. Maccheroni and M. Marinacci. How to cut a pizza fairly: Fair division with decreasing marginal evaluations. *Social Choice and Welfare*, 20(3):457–465, 2003.

[34] Y. Manabe and T. Okamoto. Meta-envy-free cake-cutting protocols. In *International Symposium on Mathematical Foundations of Computer Science*, pages 501–512. Springer, 2010.

[35] A. Procaccia. Thou shalt covet thy neighbor's cake. In *IJCAI*, pages 239–244, 2009.

[36] A. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endriss, J. Lang, and A. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13, pages 311–329. Cambridge University Press, 2016.

[37] J. Reijnierse and J. Potters. On finding an envy-free Pareto-optimal division. *Mathematical Programming*, 83(1-3):291–311, 1998.

[38] J. Robertson and W. Webb. Near exact and envy free cake division. *Ars Combinatoria*, 45:97–108, 1997.

[39] J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be Fair If You Can.* CRC Press, 1998.

[40] W. Rudin. *Functional Analysis.* McGraw Hill, 2 edition, 1991.

[41] A. Saberi and Y. Wang. Cutting a cake for five people. In *Proceedings of the 5th International Conference on Algorithmic Applications in Management*, pages 292–300. Springer, 2009.

[42] R.L. Schilling. *Measures, Integrals and Martingales.* Cambridge University Press, 2 edition, 2017.

[43] R.L. Schilling and F. Kühn. *Counterexamples in Measure and Integration.* Cambridge University Press, 2021.

[44] R.L. Schilling and D. Stoyan. Continuity assumptions in cake-cutting. Technical Report arXiv:1611.04988v1 [cs.DS], ACM Computing Research Repository (CoRR), November 2016.

[45] E. Segal-Halevi, S. Nitzan, A. Hassidim, and Y. Aumann. Fair and square: Cake-cutting in two dimensions. *Journal of Mathematical Economics*, 70:1–28, 2017.

[46] J. Sgall and G. Woeginger. An approximation scheme for cake division with a linear number of cuts. *Combinatorica*, 27(2):205–211, 2007.

[47] W. Stromquist. Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15(1):R11, 2008.

[48] W. Stromquist and D. Woodall. Sets on which several measures agree. *Journal of Mathematical Analysis and Applications*, 108(1):241–248, 1985.

[49] G. Vitali. *Sul problema della misura dei Gruppi di punti di una retta: Nota.* Tip. Gamberini e Parmeggiani, 1905. Reprinted in: L. Pepe, *Giuseppe Vitali e l'analisi reale*, Rendiconti del Seminario Matematico e Fisico di Milano 54: 187–201, 1985.

[50] T. Walsh. Online cake cutting. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, pages 292–305. Springer, 2011.

[51] W. Webb. How to cut a cake fairly using a minimal number of cuts. *Discrete Applied Mathematics*, 74(2):183–190, 1997.

[52] G. Woeginger and J. Sgall. On the complexity of cake cutting. *Discrete Optimization*, 4(2):213–220, 2007.

[53] D. Zeng. Approximate envy-free procedures. In *Game Practice: Contributions from Applied Game Theory*, pages 259–271. Springer, 2000.

Peter Kern
Mathematisches Institut
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: `kern@hhu.de`

Daniel Neugebauer
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: `daniel.neugebauer@hhu.de`

Jörg Rothe
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: `rothe@hhu.de`

René L. Schilling
Institut für Mathematische Stochastik
TU Dresden
Dresden, Germany
Email: `rene.schilling@tu-dresden.de`

Dietrich Stoyan
Institut für Stochastik
TU Bergakademie Freiberg
Freiberg, Germany
Email: `stoyan@math.tu-freiberg.de`

Robin Weishaupt
Institut für Informatik
Heinrich-Heine-Universität Düsseldorf
Düsseldorf, Germany
Email: `robin.weishaupt@hhu.de`