

CUTTING A CAKE IS NOT ALWAYS A ‘PIECE OF CAKE’: A CLOSER LOOK AT THE FOUNDATIONS OF CAKE-CUTTING 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

Since the groundbreaking work of Steinhaus [1948], cake-cutting is a metaphor for the so-called *fair division problem for a divisible, heterogeneous good*, which addresses the problem to split a contested quantity (a ‘cake’) in a fair way among several parties A, B, C, \dots ; each party may have its own idea about the value of the different parts of the cake. While mainly mathematicians and economists were concerned with the study of cake-cutting early on, “*in recent years, cake cutting has emerged as a major research topic in artificial intelligence,*” as Balkanski et al. [2014, p. 567] note. They substantiate their claim by listing ten papers on cake-cutting five of which appeared in AAI (e.g., Cohler et al. [2011]), three in IJCAI (e.g., Procaccia [2009]), and the remaining two in AAMAS proceedings (e.g., Aumann et al. [2013]). For more than a decade now, AAI and IJCAI (the two top AI conferences) and AAMAS (the leading venue for research on multiagent systems) have published numerous research papers on fair division and, in particular, on cake-cutting. Balkanski et al. [2014, p. 567] go on to write, “*The growing interest in cake cutting, and fair division more broadly, is partly motivated by potential applications in AI, such as industrial procurement, manufacturing and scheduling, and airport traffic management [Chevalyere et al., 2006]. For example, concrete applications to the allocation of multiple computational resources in shared computing systems have recently received significant attention [Gutman and Nisan, 2012, Kash et al., 2013].*” The main purpose of our paper, however, are neither applications nor novel protocols for cake-cutting; instead, we will have a closer look at the mathematical foundations of cake-cutting, establishing the connection to measure theory. Under very modest assumptions (e.g., the possibility to make continuous cuts) this will empower researchers in the field with new tools which are potent enough to deal with situations that are currently seen as “exotic” or “theoretical.”

A traditional way of fair division between two parties A and B would be to let A divide the cake into two pieces (depending on their own valuation) while B has the right to choose one of the pieces, the so-called *cut & choose* protocol. There are other possibilities for two parties as well as extensions to more than two parties (see, e.g., Procaccia, 2016, Lindner and Rothe, 2015, for an overview). Yet, while the basic rules of the game are pretty clear, the assumptions on the actual cutting process are often treated in a gentlemanlike manner. If the whole cake is represented by an interval, say $[0, 1]$, many authors think of the pieces as ‘intervals,’ without specifying whether the intervals are open $(a, b) \subset [0, 1]$, half-open $(a, b], [a, b) \subset [0, 1]$, or closed

Key words and phrases. cake-cutting protocol, admissible piece of cake, finitely additive measure, continuity properties of measures.

$[a, b] \subseteq [0, 1]$, and how to treat the – possibly twice counted – end points, i.e., $[0, 1/2) \cup [1/2, 1]$ vs. $[0, 1/2] \cup [1/2, 1]$; this is, of course, not an issue if a one-point set like $\{1/2\}$ has zero value for all parties. However, this simple example shows that a formal mathematical approach to cake-cutting needs to address questions like:

- Are (open, closed, half-open) intervals the only possible pieces of cake?
- Do we allow for finitely many or infinitely many cuts? A ‘cut’ means the split of any subset of $[0, 1]$ at a single point; it depends on the particular protocol to which interval the point will belong.
- Which properties should a valuation function (by which an agent individually evaluates the pieces of cake) have, and how does it interact with the family of admissible pieces of cake?

For some cases, there is an obvious answer: If we use only finitely many cuts, finite unions of intervals of the form $\langle a, b \rangle$ – where the angular braces indicate either open or closed ends – is all we can get; and if, in addition, any single point $a \in [0, 1]$ has zero value, we do not have to care about the open or closed ends anymore. We will see in Section 2.1 below that this rather implicit assumption brings us in a much more potent framework that can effectively deal with a countably infinite number of cuts.

Let us briefly discuss situations where an infinite number of cuts may actually be inevitable. While most research in cake-cutting has focused on finite protocols and on minimizing the required number of cuts, there are also some impossibility results that show that no finite cake-cutting protocol (even if unbounded) can guarantee all players their fair share of the cake (for various notions of fairness such as proportionality, exactness, or envy-freeness). For example, Stromquist [2008] shows that no finite cake-cutting protocol can guarantee an envy-free division of a cake among three or more players who each are to receive a single connected piece. As another example, in contrast to the moving-knife procedure due to Austin [1982] that guarantees two players an exactly proportional share, Robertson and Webb [1998] show that no finite cake-cutting protocol, bounded or unbounded, can guarantee an exactly proportional division of the cake for two players; see also the related nearly exact, envy-free, finite unbounded protocol by Robertson and Webb [1997]. Such impossibility results indicate that infinite cake-cutting protocols are unavoidable when one has to deal with certain valuations of the cake (which, admittedly, are usually constructed specifically for the purpose of proving the desired impossibility result).

As soon as we allow for countably infinitely many cuts, things change dramatically, as the following example shows.

Example 1 (Cantor dust; Cantor’s ternary set). Start with the complete cake as a single piece, i.e., $A_0 = [0, 1]$. Now, cut out the middle third of A_0 to obtain the intermediate piece $A_1 = A_0 \setminus (1/3, 2/3) = [0, 1/3] \cup [2/3, 1]$ comprising two closed intervals. Next, cut out the middle third of both remaining pieces in A_1 to obtain a union of four closed intervals $A_2 = [0, 1/9] \cup [2/9, 1/3] \cup [2/3, 7/9] \cup [8/9, 1]$, see Figure 1. If this procedure is repeated on and on, we will remove countably many open intervals, and the remainder set is $C_{1/3} = \bigcap_{i=1}^{\infty} A_i$. The set $C_{1/3}$ is the **Cantor (ternary) set** (see, e.g., Schilling and Kühn, 2021, § 2.5), and one can show that this is a closed set, which has more than countably many points, does not contain any interval, and is dense in itself, i.e., each of its points is a limit point of a sequence inside $C_{1/3}$. In the usual measuring scale, the original cake had length 1, and the recursively removed pieces have total length

$$\frac{1}{3} + \left(\frac{1}{9} + \frac{1}{9} \right) + \left(\frac{1}{27} + \frac{1}{27} + \frac{1}{27} + \frac{1}{27} \right) + \cdots = \sum_{i \in \mathbb{N}} \frac{2^{i-1}}{3^i} = 1,$$

so that $C_{1/3}$ has zero ‘length,’ but it still contains more than countably many points.

The same construction principle, removing at each stage 2^{i-1} identical open middle intervals, each having length p^i for some p , $0 < p \leq 1/3$, leads to the Cantor set C_p , which is, again, closed, uncountable, and does not contain any interval. If, say, $p = 1/4$, the removed intervals have total length $1/2$ and the remaining Cantor dust has ‘length’ $1 - 1/2 = 1/2$. This is not quite expected.




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

While it is intuitive that the removed intervals should have a certain length, it feels unnatural to speak of the ‘length’ of a dust-like set as C_p . In fact, we are dealing here with (one-dimensional) Lebesgue measure, which is the mathematically formal extension of the familiar notion of ‘length.’

An alternative, slightly more formal way of illustrating the Cantor dust is given in the appendix as Example 18.

This example shows that, as soon as we allow for countably many cuts, there can appear sets which may not be written as a countable union of intervals; moreover, although these sets consist of limit points only, they may have strictly positive length.

An important feature of this example is the fact that we extend the family of intervals to a family of subsets which contains (i) finite unions, (ii) countable intersections, and (iii) complements of its members, leading to fairly complicated subsets as, e.g., C_p . Moreover, when calculating the length of all removed intervals, we tacitly assumed

- the (finite) additivity of length: The length of two disjoint sets is the sum of their lengths;
- the countable or σ -additivity which plays the role of a continuity property: The length of a countable union is the limit of the length of the union of the first N sets as $N \rightarrow \infty$.

As it will turn out, these are two far-reaching assumptions on the interplay of the valuation function (here: length) with its domain; we will see how this relates to the desirable property that we can cut off pieces of arbitrary length ℓ , $0 \leq \ell \leq 1$, from the cake $[0, 1]$ (allowing for any valuation values of the cut-off pieces).

Commonly, in cake-cutting theory (see, e.g., Brams and Taylor, 1996, Procaccia, 2016, Lindner and Rothe, 2015) a (piece-wise constant) valuation function $v: \mathcal{P} \rightarrow [0, 1]$, where \mathcal{P} is some family of subsets of the cake, is represented as shown in Figure 2: The cake $[0, 1]$ is split horizontally into multiple pieces and the number of vertically stacked boxes per piece describes 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 $v(X') = 3/17$. Having this example in mind, one is not aware of any limitations and might assume that all possible sets are indeed admissible pieces, i.e., $\mathcal{P} = \mathfrak{P}([0, 1]) = \{A \mid A \subseteq [0, 1]\}$.

The following classical example from measure theory shows that there cannot exist a valuation function that assigns to intervals $\langle a, b \rangle \subseteq [0, 1]$ their natural length $b - a$, and which is additive, σ -additive (in the sense explained above), and able to assign a value to every set $A \subseteq [0, 1]$. Things are different if we do not require σ -additivity (see the discussion in Schilling and Kühn, 2021, § 7.31).




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

Example 2 (Vitali, 1905; see also, e.g., Schilling and Kühn, 2021). Let $[0, 1]$ be the standard cake, and assume that the valuation function v is σ -additive (see Definition 2 on page 6), assigning to any interval its natural length. This means, in particular, that v is invariant under translations and evaluates the complete cake with $v([0, 1]) = 1$. 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, $x - y \in \mathbb{Q}$, i.e., their difference is rational. The relation $*$ is an equivalence relation and the corresponding equivalence classes $[x] = \{y \in \mathbb{R} \mid x * y\} \subseteq \mathbb{R}$ lead to a disjoint partitioning of \mathbb{R} . By the axiom of choice, there is a set $V \subset [0, 1]$ which contains exactly one representative of every equivalence class $[x]$. A set like V is called a **Vitali set**. Clearly, $V \in \mathfrak{P}([0, 1])$ and the sets $q + V = \{q + x \bmod 1 \mid x \in V\}$, $q \in \mathbb{Q}$, are a disjoint partition of $[0, 1]$; thus $\bigcup_{q \in \mathbb{Q}} (q + V) = [0, 1]$. By assumption, v is σ -additive and assigns to each $q + V$ the same value (translation invariance). Hence, we end up with the contradiction

$$1 = v([0, 1]) = v\left(\bigcup_{q \in \mathbb{Q}} (q + V)\right) = \sum_{q \in \mathbb{Q}} v(q + V) = \begin{cases} 0 & \text{if } v(V) = 0, \\ \infty & \text{if } v(V) > 0. \end{cases}$$

Thus v cannot have the power set of the cake $[0, 1]$ as its domain *if we assume that v is σ -additive*. We will see below that certain commonly used divisibility assumptions are equivalent to the σ -additivity of the valuation.

Example 3 (Cantor function). Let us return to Example 1 and interpret the points in the set $C_{1/3}$ as valuable assets which need to be priced. We may assume that the total value of the cake $C_{1/3}$ is 1. We want to construct a ‘cumulative valuation function V ’ which has the property that for $0 \leq a \leq b \leq 1$ the difference $V(b) - V(a)$ is the value of the points contained in $C_{1/3} \cap (a, b]$. Clearly, $x \mapsto V(x)$ is a (not necessarily strictly) increasing function with $V(0) = 0$ and $V(1) = 1$.

If we agree that the assets should be ‘homogeneously’ priced, then we are automatically led to the following scheme: As the total value of $C_{1/3}$ is one, the value of $C_{1/3} \cap [0, 1/2]$ and $C_{1/3} \cap [1/2, 1]$ should be the same, i.e., $1/2$. Since $C_{1/3} \cap (1/3, 2/3) = \emptyset$, we see that both the left third and the right third of $C_{1/3}$ has the value $1/2$. This means that $V(x) = 1/2$ on the whole middle third $(1/3, 2/3)$.

Now we can repeat this argument in the two remaining sets $C_{1/3} \cap [0, 1/3]$ and $C_{1/3} \cap [2/3, 1]$. Since these pieces are scaled-down versions of the original set $C_{1/3}$, we can repeat our argument to the three thirds of the scaled sets and so we see that the cumulative valuation function $V(x)$ takes the values $1/4$ and $3/4$ on the intervals $(1/9, 2/9)$ and $(7/9, 8/9)$, respectively.

Iterating this procedure *ad infinitum*, the remaining values of V at interfaces of the intervals are uniquely determined by monotonicity and we end up with the so-called *Cantor function* or *devil’s staircase*, which is monotone, increasing, continuous, and it is flat (i.e., constant) on all middle-thirds removed in the construction process of $C_{1/3}$ in Example 1; a precise mathematical description can be achieved, e.g., using the alternative representation of the Cantor set in Example 18 of Appendix 6, but at this point the pictures in Figure 3 tell it all: The Cantor function is the typical function where the fundamental theorem of integral and differential calculus fails: $V(x)$ is constant on all intervals contained in $[0, 1] \setminus C_{1/3}$. Since V is increasing, we can set




FIGURE 3. Step-by-step construction of the Cantor function: In each step, we subdivide the top right and bottom left squares into nine smaller squares. We keep only the three squares along the diagonal, and discard (gray out) the off-diagonal squares. The middle square is halved by a horizontal line (here V is constant). Repeat.

$V'(x) := \limsup_{h \rightarrow 0} \frac{1}{h}(V(x+h) - V(x))$, and this is the usual derivative whenever it exists. In particular, $V'(x) = 0$ on $[0, 1] \setminus C_{1/3}$. On the other hand, we have

$$\int_0^x V'(t) dt = \int_{C_{1/3} \cap [0, x]} V'(t) dt = 0 \neq V(x) - V(0) \quad \text{for any } x \in (0, 1].$$

This happens because the ‘length’ of $C_{1/3}$ is zero, i.e., we integrate over ‘too small a set’ (no matter how big the integrand may be, it could even take the value $+\infty!$) so as to pick up any strictly positive value.

The above examples highlight some of the problems when evaluating sets. A Cantor-like piece can only be evaluated if the valuation function is not too simplistic. On the other hand, a Vitali set cannot be evaluated at all if we request too many properties of a valuation function, i.e., the domain $\mathfrak{P}([0, 1])$ consisting of all possible pieces of cake is, in general, too large.

Across the research field of cake-cutting (see, e.g., the textbooks by Brams and Taylor, 1996, Robertson and Webb, 1998, and the book chapters by Procaccia, 2016, Lindner and Rothe, 2015), there exist several different assumptions on the underlying model. Our goal is to review thoroughly and comprehensively all the different models that are currently applied in the literature. Furthermore, we study the relationships between these models and formulate some related results. 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, which is as simple as possible, yet powerful enough to cope with these problems and still compatible with many of the currently used models.

Frequently, authors proposing cake-cutting protocols abstain from making formal assumptions or from formalizing their model in detail. For example, Brams et al. [1997, 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 any given purpose, and which are not. If we allow only a fixed number of cuts, splitting the cake $[0, 1]$ into a finite number of pieces of the type $\langle a, b \rangle \subseteq [0, 1]$, a naive approach is always possible: The valuation should be additive and its domain contains unions of finitely many intervals. If, on the other hand, there are potentially infinitely many cuts – e.g., if the players play a game resulting in an *a priori* not fixed number of rounds (such as the finite unbounded envy-free cake-cutting

protocol of Brams and Taylor [1995b]) – the limiting case cannot any longer be treated by a finitely additive valuation and a domain containing only finite unions, see Example 1.

We propose to use ideas from measure theory, which provides the right toolbox to tackle the issues described above. We will see that, at least for the cake $[0, 1]$, even the naive approach plus the requirement that we can split every piece $\langle a, b \rangle$ by a single cut into any proportion (in fact, a slightly weaker requirement will do, cf. Definition 2 (D)), automatically leads to the measure-theoretic point of view. That is to say that in many natural situations the naive standpoint is ‘practically safe’ since its obvious shortcomings are automatically ‘fixed by (measure) theory,’ if one uses the correct formulation.

2. THE RULES OF THE GAME

Throughout this paper, $[0, 1]$ denotes a standard cake, and the power set $\mathfrak{P}([0, 1]) = \{S \mid S \subseteq [0, 1]\}$ are all **possible** pieces of cake from a set-theoretic point of view. We define $\mathcal{P} \subseteq \mathfrak{P}([0, 1])$ as the set of all **admissible** pieces of $[0, 1]$, i.e., those pieces which (a) can be allocated to some players via a cake-cutting protocol, and (b) can be evaluated by the players using their valuation functions. Sometimes it is necessary to consider an ‘abstract’ cake X , with its possible and admissible pieces $\mathfrak{P}(X)$ and $\mathcal{P} \subseteq \mathfrak{P}(X)$. Some results for the standard cake $[0, 1]$ remain true for abstract cakes. For example, an abstract cake X might be contained in the n -dimensional unit cube: $X \subseteq [0, 1]^n$.

2.1. Dividing a Cake with Finitely Many Cuts. We start by formulating requirements for \mathcal{P} regarding the admissible pieces of cake. The discussion in this section applies both to the standard cake $[0, 1]$ and the abstract cake X . Obviously, we want to be able to allocate the complete cake X as well as an empty piece \emptyset to a player and therefore, $X \in \mathcal{P}$ and $\emptyset \in \mathcal{P}$ must hold. If $A \subseteq X$ is already allocated to some player, i.e., $A \in \mathcal{P}$, then we want to be able to give the remainder of the cake to another player; so for all $A \in \mathcal{P}$, we demand that the complement of A , denoted by $\overline{A} = X \setminus A$, is in \mathcal{P} . Furthermore, we want to be able to cut and combine pieces of cake; so for all $A, B \in \mathcal{P}$, we require $A \cup B \in \mathcal{P}$. Note that $A \cap B = \overline{\overline{A} \cup \overline{B}}$ and $A \setminus B = A \cap \overline{B}$, so our previously formulated requirements also allow us to allocate the intersection of a finite number of pieces of cake and to evaluate the difference of two pieces of cake.

Definition 1. Let X be a(n abstract) cake. A family $\mathcal{A} \subseteq \mathfrak{P}(X)$ is called an **algebra** over X if $\emptyset \in \mathcal{A}$ and for all $A, B \in \mathcal{A}$ it holds that \overline{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 admissible pieces of cake, we ended up with a well-studied, structured concept from measure theory: an algebra.

Example 4. If $X = [0, 1]$, then $\mathfrak{P}(X)$ and $\{\emptyset, X\}$ are algebras – in fact these are the largest possible and the smallest possible algebras over $[0, 1]$. Another useful algebra is the family $\mathcal{I}([0, 1])$ of all unions of finitely many intervals in $[0, 1]$ – and it is easy to check that $\mathcal{I}([0, 1])$ is the smallest algebra containing all closed (or all open or all half-open) intervals from $[0, 1]$. While it is obvious that $\{\emptyset, [0, 1]\}$ is useless for our purpose, as then only two possible pieces can be allocated, the complete cake and an empty piece, we might – at the other extreme – also take $\mathfrak{P}([0, 1])$ as the set for the admissible pieces of $[0, 1]$. However, when choosing \mathcal{P} , we must also ensure that meaningful valuation functions can exist for this set, and Example 2 shows that for a rather natural valuation function – geometric length – $\mathfrak{P}([0, 1])$ is too big.

Let us list the common requirements for the players’ valuation functions. A **valuation function** v shall assign to any admissible piece of cake $A \in \mathcal{P}$ some positive real number. In order to normalize the players’ valuations and keep them comparable, we map the positive real numbers onto $[0, 1]$ continuously, bijectively, and preserving the natural order. Hence, we can further limit the valuation function’s range to $[0, 1]$, i.e., we have $v: \mathcal{P} \rightarrow [0, 1]$. The next definition lists desirable properties for a valuation function.

Definition 2. Let X be a(n abstract) cake and \mathcal{A} the algebra of admissible pieces. A **valuation function** is a function $v: \mathcal{A} \rightarrow [0, 1]$, which is normalized, i.e., $v(\emptyset) = 0$ and $v(X) = 1$. Moreover, v is called

- (M) **monotone** if for $A, B \in \mathcal{A}$ with $A \subseteq B$, one has $v(A) \leq v(B)$;
- (A) **additive** or **finitely additive** if for all $A, B \in \mathcal{A}$ such that $A \cap B = \emptyset$, one has $v(A \cup B) = v(A) + v(B)$;
- (Σ) **σ -additive** or **countably additive** if for any sequence $(A_n)_{n \in \mathbb{N}}$ of pieces in \mathcal{A} such that $A_i \cap A_j = \emptyset$ ($i \neq j$) and $\bigcup_{i \in \mathbb{N}} A_i \in \mathcal{A}$, one has $v(\bigcup_{i \in \mathbb{N}} A_i) = \sum_{i \in \mathbb{N}} v(A_i)$;
- (D) **divisible** if for every $A \in \mathcal{A}$ and for every real number α , $0 \leq \alpha \leq 1$, there exists some $A_\alpha \in \mathcal{A}$ with $A_\alpha \subseteq A$ such that $v(A_\alpha) = \alpha v(A)$.

Clearly, (Σ) implies (A) – take $A_1 = A$, $A_2 = B$, and $A_i = \emptyset$ for $i \geq 3$ – and (A) is equivalent to the so-called **strong additivity**, defined as $v(A \cup B) = v(A) + v(B) - v(A \cap B)$: Just observe that $A \cup B = [A \setminus (A \cap B)] \cup [B \setminus (A \cap B)] \cup [A \cap B]$, i.e., $A \cap B \neq \emptyset$ counts towards both $v(A)$ and $v(B)$ but only once in $v(A \cup B)$, hence the correction $-v(A \cap B)$. Finally, (strong) additivity implies monotonicity.

The assumption that \mathcal{A} is an algebra makes sure that we can indeed perform all of the above manipulations with sets without ever leaving \mathcal{A} . Note, however, that (Σ) and (D) require a certain richness assumption on \mathcal{A} , which need not be satisfied for an algebra; for example, a union of countably many member sets need not be in the algebra. In other words: The properties (Σ) and (D) affect both v and \mathcal{A} .

Remark 5. Let X be a(n abstract) cake and $\mathcal{A} \subseteq \mathfrak{P}(X)$ an algebra over X . Any additive valuation is a **finitely additive measure** with total mass $v(X) = 1$ (see, e.g., Schilling, 2017, Chapter 4).

Requirement (D) not only demands more from \mathcal{A} but also from v . Specifically, (D) entails that any $N \in \mathcal{A}$ which does not contain a nonempty and strictly smaller piece of cake – this is an **atom**, i.e., an indivisible piece of cake – must have zero valuation.

Definition 3. Let \mathcal{A} be an algebra over a(n abstract) cake X and v be a finitely additive valuation. A set $A \in \mathcal{A}$ is an **atom** if $v(A) > 0$ and every $B \subseteq A$, $B \in \mathcal{A}$, satisfies $v(B) = \alpha v(A)$ with $\alpha = 0$ or $\alpha = 1$.

Clearly, a valuation v which enjoys property (D) cannot have atoms.

2.2. Dividing the Standard Cake. Let us briefly discuss the consequences of the notions introduced in the previous section if X is the standard cake $[0, 1]$. If, in addition, \mathcal{A} contains all intervals of type $\langle a, b \rangle$, then all singletons $\{a\} = [a, b] \setminus \langle a, b \rangle$ are in \mathcal{A} , and they are the only possible atoms. In this case, (D) entails that v does not charge single points: $v(\{a\}) = 0$ for all $a \in [0, 1]$. This is the proof of the following lemma.

Lemma 6. *Let $[0, 1]$ be the standard cake and \mathcal{A} an algebra of admissible sets. Every additive valuation function $v: \mathcal{A} \rightarrow [0, 1]$ that satisfies (D) is atom-free. In particular, if $\mathcal{A} \supset \mathcal{I}([0, 1])$ contains all intervals, then $v(\{a\}) = 0$ for all $a \in [0, 1]$.*

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

Definition 4. Let $v: \mathcal{A} \rightarrow [0, 1]$ be a finitely additive valuation function on the algebra $\mathcal{A} = \mathcal{I}([0, 1])$ of finite unions of intervals from $[0, 1]$.

- (1) The function $x \mapsto F_v(x) := v([0, x])$, $x \in [0, 1]$, is the **distribution function** of the valuation v .
- (2) The valuation v is said to be **continuous** if $x \mapsto F_v(x)$ is continuous.

Since v is additive, $F_v: [0, 1] \rightarrow [0, 1]$ is positive, monotonically increasing, and bounded by $F_v(1) = 1$. Note that a continuous valuation function on $\mathcal{I}([0, 1])$ cannot have atoms, as

$$v(\{x\}) = v([0, x] \setminus [0, x)) = F_v(x) - F_v(x-) = 0, \text{ where } F_v(x-) = \lim_{y \uparrow x} F_v(y).$$

The continuity of v can also be cast in the following way: For all a and b with $0 \leq a < b \leq 1$ 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 $v([0, c]) = \gamma$. This explains the close connection between continuity and divisibility of v . In fact, assuming divisibility (D) of v , it can be shown that the distribution function is necessarily continuous. The following proof of this statement is inspired by Schilling and Stoyan [2016, Example 3.4].

Lemma 7. *Let v be an additive valuation for the standard cake $[0, 1]$, where $\mathcal{I}([0, 1])$ denotes the family of admissible pieces. If v is divisible, then the distribution function $F = F_v$ is a continuous function with $F(0) = 0$.*

Proof. We have seen in Lemma 6 that a divisible additive valuation v has no atoms, so $F(0) = v(\{0\}) = 0$. Since F is monotone and bounded, the one-sided limits $F(t-) := \lim_{s \uparrow t} F(s)$ and $F(t+) := \lim_{s \downarrow t} F(s)$ exist for all $t \in (0, 1]$ and $u \in [0, 1]$.

Assume that F is not continuous. Then there exists some $t_0 \in [0, 1]$ such that $F(t_0-) < F(t_0)$ or $F(t_0+) > F(t_0)$. If $F(t_0) - F(t_0-) = \varepsilon > 0$, then there exists some $t_1 < t_0$ such that $F(t_0) - F(t_1) \leq \frac{3}{2}\varepsilon$. Set $I := (t_1, t_0]$ and observe that $v(I) = F(t_0) - F(t_1) \in [\varepsilon, \frac{3}{2}\varepsilon]$. Pick an arbitrary $J \in \mathcal{I}([0, 1])$ which is contained in I . Since J is a finite union of intervals, J differs from its closure \bar{J} by at most finitely many points; as $v(\{x\}) = 0$ for any $x \in [0, 1]$, we have $v(J) = v(\bar{J})$.

We distinguish two cases: If $t_0 \in \bar{J}$ is not an isolated point, then $v(J) = v(\bar{J}) \geq \varepsilon$. If $t_0 \notin \bar{J}$ or if $t_0 \in \bar{J}$ is an isolated point, then we have due to $v(\{t_0\}) = 0$ that

$$\begin{aligned} v(J) &= v(\bar{J}) \leq F(t_0-) - F(t_1) = (F(t_0) - F(t_1)) - (F(t_0) - F(t_0-)) \\ &= v(I) - \varepsilon \leq \frac{1}{2}\varepsilon. \end{aligned}$$

Hence, it is not possible to select a piece of cake $J \in \mathcal{I}([0, 1])$ with $J \subseteq I$ and $v(J) = \frac{3}{4}\varepsilon \in [\frac{1}{2} \cdot v(I), \frac{3}{4} \cdot v(I)]$, which contradicts divisibility.

If $F(t_0+) - F(t_0) = \varepsilon > 0$, a similar argument applies. \square

Conversely, if the distribution function F_v of a finitely additive valuation v defined on $\mathcal{I}(X)$ is continuous with $F_v(0) = 0$, then it is easy to see that v is divisible. Hence we get:

Corollary 8. *A finitely additive valuation v on $\mathcal{I}([0, 1])$ is divisible if, and only if, its distribution function F_v is continuous with $F_v(0) = 0$. This is also equivalent to v being atom-free.*

Corollary 8 establishes a one-to-one correspondence between divisible valuations and monotonically increasing, continuous functions on $[0, 1]$ which are 0 at the origin and 1 at $x = 1$. This shows that the identity $x \mapsto x$ gives rise to a valuation (it assigns every interval $\langle a, b \rangle$ its natural length $b - a$) but also the Cantor function $V(x)$ from Example 3 can be viewed as a valuation function.

We will see in the next section that every finitely additive, divisible valuation can be extended to become and identified with a unique σ -additive measure that is defined on the Borel σ -algebra $\mathcal{B}(X)$; this is the smallest family of sets that contains all intervals and that is stable under complements and countable unions of its members. This enables us to evaluate sets in $\mathcal{B}(X)$ that are not finite unions of intervals, such as the Cantor set in Example 1.

2.3. Measure Theory: The Art of Dividing a Cake by Countably Many Cuts. In Sections 2.1 and 2.2, we have focused on finitely many cuts when dividing the cake. But we may easily come into the situation where the number of cuts is not limited; not all protocols in the cake-cutting literature are finite.¹ Thus we are led to consider unions of countably many pieces

¹For example, prior to the celebrated finite bounded envy-free cake-cutting protocol due to Aziz and Mackenzie [2016b, 2020], the cake-cutting protocol of Brams and Taylor [1995b] was the best protocol known to guarantee envy-freeness for any number of players. They argue that the allocation must become envy-free at some (unknown) finite stage, which is why their protocol is considered to be a *finite unbounded* envy-free procedure only. And yet, being open-ended, it is in some sense even an *infinite* procedure that describes an infinite process. Similarly, it is reasonable to conjecture that some moving-knife procedures can be converted to discrete procedures that require infinitely many cuts.

and the valuation of such countable unions, see also property (Σ) in Definition 2. To deal with such situations, measure theory provides the right tools.

We will now introduce some basics from measure theory, which we need in the subsequent discussion of the cake-cutting literature. Our standard references for measure theory are the monographs by Schilling [2017] and Schilling and Kühn [2021], where also further background information can be found.

Definition 5. Let $[0, 1]$ be a cake. A subset $\mathcal{A} \subseteq \mathfrak{P}([0, 1])$ is called a **σ -algebra** over $[0, 1]$ if \mathcal{A} is an algebra over $[0, 1]$ and, for all sequences $(A_n)_{n \in \mathbb{N}}$ with $A_n \in \mathcal{A}$, the countable union $\bigcup_{n \in \mathbb{N}} A_n$ is in \mathcal{A} , too.

Every algebra in $[0, 1]$ containing finitely many sets is automatically a σ -algebra. On the other hand, $\mathfrak{P}([0, 1])$ is both an algebra and a σ -algebra, whereas the family $\mathcal{I}([0, 1])$ is an algebra, but not a σ -algebra: For instance, the Cantor dust C_p (cf. Example 1) is not in $\mathcal{I}([0, 1])$. Recall that we defined $\mathcal{I}([0, 1])$ to be the smallest algebra containing all (finite unions of) intervals in $[0, 1]$; thus it is natural to consider the smallest σ -algebra containing all (finite unions of) intervals in $[0, 1]$.

To see that this is well-defined, we need a bit more notation. Recall that $\langle a, b \rangle$ stands for any (open, closed, or half-open) interval of $[0, 1]$. We denote by

$$\mathcal{Q}([0, 1]) = \{\langle a, b \rangle \mid a, b \in [0, 1]\}.$$

the family of all intervals within $[0, 1]$.

Moreover, if $\mathcal{P} \subseteq \mathfrak{P}([0, 1])$ is any family, then $\sigma(\mathcal{P})$ denotes the smallest σ -algebra containing \mathcal{P} . This can be a fairly complicated object and its existence is not really obvious. To get an idea as to why $\sigma(\mathcal{P})$ makes sense, we note that $\mathcal{P} \subseteq \mathfrak{P}([0, 1])$, that $\mathfrak{P}([0, 1])$ is a σ -algebra, and that the intersection of any number of σ -algebras is still a σ -algebra.

The next lemma is a standard result from measure theory.

Lemma 9. Let \mathcal{P} denote any of the four families of open intervals, closed intervals, left-open intervals, or right-open intervals within $[0, 1]$. It holds that

$$\sigma(\mathcal{P}) = \sigma(\mathcal{Q}([0, 1])).$$

The fact that $\sigma(\mathcal{Q}([0, 1]))$ coincides with the σ -algebra generated by all closed intervals in $[0, 1]$ can be used to generalize Lemma 9 to abstract cakes, which carry a topology, hence a family of open and of closed sets. The thus generated ‘topological’ σ -algebra plays a special role and has a special name.

Definition 6. We denote by $\mathcal{B}([0, 1])$ the smallest σ -algebra on $[0, 1]$ containing all closed intervals from $[0, 1]$ and call it the **Borel or topological σ -algebra** over $[0, 1]$.

The following definition is also well-known. We state it only for the standard cake, but it is clear how to extend it to abstract cakes.

Definition 7. Let $[0, 1]$ be a cake and \mathcal{A} a σ -algebra on $[0, 1]$. A (positive) **measure** μ on $[0, 1]$ is a map $\mu: \mathcal{A} \rightarrow [0, \infty]$ satisfying that $\mu(\emptyset) = 0$ and μ is σ -additive.

It is useful to see a measure μ as a function defined on the sets. If the set-function is additive, then σ -additivity is, in fact, a continuity requirement on μ , as it allows to interchange the limiting process in the infinite union $\bigcup_{i \in \mathbb{N}} A_i$ of pairwise disjoint sets with a limiting process in the sum. To wit:

$$\mu\left(\bigcup_{n \in \mathbb{N}} A_n\right) = \lim_{N \rightarrow \infty} \mu\left(\bigcup_{n=1}^N A_n\right) = \lim_{N \rightarrow \infty} \sum_{n=1}^N \mu(A_n) = \sum_{n \in \mathbb{N}} \mu(A_n);$$

since all terms are positive, the value of the sum is well-defined, i.e., it is either convergent in $[0, \infty)$ or improperly convergent yielding $+\infty$. Equivalently, we can state σ -additivity as $B_1 \subset B_2 \subset B_3 \subset \dots \uparrow B = \bigcup_{n \in \mathbb{N}} B_n$, then $\mu(B_n) \uparrow \mu(B)$ (for any measure μ) or as $C_1 \supset C_2 \supset C_3 \supset \dots \downarrow C = \bigcap_{n \in \mathbb{N}} C_n$, then $\mu(C_n) \downarrow \mu(C)$ (for finite measures μ).

Sometimes (and a bit provocatively) it is claimed that there are essentially only two measures on $[0, 1]$ (or on \mathbb{R} or \mathbb{R}^n): **Lebesgue measure** $A \mapsto \lambda(A)$ and **Dirac measure** $A \mapsto \delta_x(A)$, where $x \in [0, 1]$ is a fixed point. Let us briefly discuss these two extremes and explain as to why the claim is incorrect but still sensible.

Dirac Measure. Let $a \in [0, 1]$ be a fixed point and set $A \mapsto \delta_a(A) = 1$ or $= 0$ according to $a \in A$ or $a \notin A$, respectively. This definition works for any $A \subseteq [0, 1]$, and it is easy to see that this set-function is indeed a measure (in the sense of Definition 7 on the σ -algebra $\mathcal{A} = \mathfrak{P}([0, 1])$ – or any smaller σ -algebra over $[0, 1]$).

We call $\{a\}$ the **support** of δ_a since, by definition, δ_a charges only sets such that $\{a\} \subseteq A$. If we compare Dirac measure with Lebesgue’s measure, the problem is that the support of δ_a is a degenerate interval $\{a\} = [a, a]$ of length zero, see below.

Lebesgue Measure. The idea behind Lebesgue measure is to have a set-function $A \mapsto \lambda(A)$ in $[0, 1]$ (or in \mathbb{R} or \mathbb{R}^n) with all properties of the familiar volume from geometry; in particular, we want a volume that is additive and invariant under shifts and rotations. Thus it is natural to define for a simple set Q like an interval $Q = (a, b) \subset [0, 1]$ (or an n -dimensional ‘cube’ $Q = \prod_{i=1}^n (a_i, b_i]$)

$$\lambda(Q) = b - a$$

$$\left(\text{respectively, } \lambda(Q) = \prod_{i=1}^n (b_i - a_i) = \text{length} \times \text{width} \times \text{height} \times \cdots \right).$$

Invariance under shifts together with the σ -additivity (Σ) allow us to exhaust (‘triangulate’) more complicated shapes like a circle with countably many disjoint sets $(Q_n)_{n \in \mathbb{N}}$ such that with $A = \bigcup_{n \in \mathbb{N}} Q_n$, we have $\lambda(A) = \sum_{n \in \mathbb{N}} \lambda(Q_n)$. The restriction to countable unions is natural, as we exhaust a given shape by nontrivial sets Q_n , having nonempty interior: Each of them contains a rational point $q \in \mathbb{Q}^n$; hence, there are at most countably many nonoverlapping Q_n .

There are immediate questions with this approach: Which types of sets can be ‘measured’? Is the procedure unique? Is the process of measuring more complicated sets constructive? At this point we encounter a problem: General sets $A \subseteq \mathbb{R}^n$ are way too complicated to get a well-defined and unique extension of λ from the rectangles to $\mathfrak{P}(\mathbb{R}^n)$. In dimension $n = 1$ and for the standard cake $[0, 1]$, the Cantor sets C_p from Example 1 were already challenging, but the Vitali set from Example 2 shows that the cocktail of shift invariance and σ -additivity becomes toxic.

The way out is the notion of measurable sets and Carathéodory’s extension theorem (stated as Theorem 10 further down). This works as follows: In view of the σ -additivity property of λ , it makes sense to consider the σ -algebra $\mathcal{A} \subseteq \mathfrak{P}(\mathbb{R}^n)$ which contains the intervals (respectively, cubes). Thus we naturally arrive at the notion of the Borel σ -algebra as the canonical 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 $A \subseteq \mathbb{R}^n$ 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 (such as the *Banach–Tarski paradox*; see, e.g., Wagon, 1985) if we try to have a finitely additive geometric volume 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.

General Measures. Let us return to the assertion that λ and δ_a are ‘essentially the only measures’ on \mathbb{R}^n . To keep things simple, we discuss here only the standard cake $[0, 1]$.

Lebesgue’s decomposition theorem shows that all σ -additive measures μ on $[0, 1]$ with the Borel σ -algebra $\mathcal{B}([0, 1])$ are of the form $\mu = \mu^{\text{ac}} + \mu^{\text{sc}} + \mu^{\text{d}}$ where ‘ac,’ ‘sc,’ and ‘d’ stand for absolutely continuous, singular continuous, and discontinuous. This is best explained by looking at the distribution function $F(x) = F_\mu(x)$. Since $x \mapsto F(x)$ is increasing, it is either

continuous or discontinuous (with at most countably many discontinuities), accounting for the parts $(\mu^{\text{ac}}, \mu^{\text{sc}})$ and μ^{d} , respectively. At the points where F is continuous, we have again two possibilities: F is either differentiable ($F'(x) = f(x)$) or it isn't, yielding 'ac' vs. 'sc.' From Lebesgue's differentiation theorem it is known that the points with 'sc' or 'd' must have Lebesgue measure zero. Thus, we finally arrive at the decomposition

$$(1) \quad \mu(dx) = f(x) dx + \mu^{\text{sc}}(dx) + \sum_i (F(x_i) - F(x_i-)) \delta_{x_i}(dx),$$

where x_1, x_2, \dots are the at most countably many discontinuities (jump points) of F and $f(x) = \frac{d}{dx} F(x)$.

Here are four typical **examples** for valuations corresponding to these cases.

- Purely **ac**: $F^{\text{ac}}(x) := x$ is absolutely continuous since $\frac{d}{dx} F^{\text{ac}}(x) = 1$ exists and $F^{\text{ac}}(x) = \int_0^x 1 dt$. This F^{ac} corresponds to Lebesgue measure. In general, an absolutely continuous $F^{\text{ac}}(x)$ is always of the form $F^{\text{ac}}(x) - F^{\text{ac}}(0) = \int_0^x f(t) dt$ and $f(t) = \frac{d}{dt} F^{\text{ac}}(t)$.
- Purely **sc**: The Cantor function $F^{\text{sc}}(x) := V(x)$ from Example 3 is continuous, but it is not absolutely continuous. $V'(x)$ exists (in a classical sense) only in the points $[0, 1] \setminus C_{1/3}$ and $V(x) \neq \int_0^x V'(t) dt$. The corresponding valuation is nevertheless of the form $v((a, b]) = V(b) - V(a)$, but it cannot be represented in the form $\int_a^b f(t) dt$ for any function f .
- Purely **d**: Any increasing step-function with jumps of size $\Delta_i > 0$, $i = 1, 2, \dots$, at the points $x_i \in [0, 1]$ corresponds to the discontinuous case: We have atoms exactly at the points x_i where $F^{\text{d}}(x)$ is discontinuous (i.e., jumps). The general form of such functions is $F^{\text{d}}(x) = \sum_{i=1}^{\infty} \Delta_i \mathbf{1}_{[x_i, 1]}(x)$ where $\mathbf{1}_{[x_i, 1]}(x)$ is the indicator function (taking the values 1 and 0 according to $x \in [x_i, 1]$ or $x \notin [x_i, 1]$, respectively) and $\sum_{i=1}^{\infty} \Delta_i = 1$.
- Mixed **ac+sc+d**: Let $p_{\text{ac}}, p_{\text{sc}}, p_{\text{d}} \in [0, 1]$ be such that $p_{\text{ac}} + p_{\text{sc}} + p_{\text{d}} = 1$ and let $F^{\text{ac}}, F^{\text{sc}}, F^{\text{d}}$ be as in the previous examples. Then the convex combination $F(x) = p_{\text{ac}} F^{\text{ac}}(x) + p_{\text{sc}} F^{\text{sc}}(x) + p_{\text{d}} F^{\text{d}}(x)$ corresponds to a valuation which combines all three types of (dis-)continuity properties.

Let us close this section with the central result on the extension of valuations defined on an algebra \mathcal{A} to measures on the σ -algebra $\sigma(\mathcal{A})$ generated by \mathcal{A} . We state it only for the standard cake; the formulation for more abstract cakes is obvious.

Theorem 10 (Carathéodory's extension theorem). *Let \mathcal{A} be the algebra of admissible pieces of the cake $[0, 1]$ and $v : \mathcal{A} \rightarrow [0, 1]$ be a valuation such that $v(\emptyset) = 0$. If v is additive and σ -additive relative to \mathcal{A} , i.e., v satisfies (Σ) , then there is a unique extension of v , defined on $\sigma(\mathcal{A})$, which is a σ -additive measure on $\sigma(\mathcal{A})$.*

2.4. Abstract Cakes. Let us briefly discuss more general cakes X than $[0, 1]$. In this section, $X \neq \emptyset$ will be a general set, \mathcal{A} an algebra of admissible pieces. The notion of σ -algebra is, *mutatis mutandis*, the same as in the case of the standard cake (Definition 5) and we denote by $\sigma(\mathcal{A})$ the smallest σ -algebra that contains the algebra \mathcal{A} . The definition and the properties of a valuation $v : \mathcal{A} \rightarrow [0, 1]$ (cf. Definition 2) still work in this general setting, but since X is abstract, there may not be (an equivalent of) a distribution function; this means that the connection between divisibility and σ -additivity, cf. Lemma 7 and Corollary 8, might fail in an abstract setting.

We begin with a new definition of (D) for finitely additive valuations on abstract cakes.

Definition 8. A finitely additive valuation v on an abstract cake X and an algebra of admissible pieces \mathcal{A} has the property (DD) if for every $A \in \mathcal{A}$ and $\alpha \in (0, 1)$, there is an increasing sequence of sets $B_\alpha^1 \subset B_\alpha^2 \subset B_\alpha^3 \subset \dots$, $B_\alpha^n \in \mathcal{A}$, such that $B_\alpha^n \subset A$ and $\sup_{n \in \mathbb{N}} v(B_\alpha^n) = \alpha v(A)$.

Property (DD) essentially says that for every value $\alpha v(A) \in [0, 1]$ we can find an admissible piece of cake $B_\alpha^n \subset A$ whose valuation $v(B_\alpha^n)$ is close to $\alpha v(A)$. The limiting piece $\bigcup_n B_\alpha^n$, which should produce the value $\alpha v(A)$ exactly, may not be admissible if we are restricted to finitely many cuts.

If v is a σ -additive valuation and \mathcal{A} a σ -algebra, then $B_\alpha := \bigcup_{n \in \mathbb{N}} B_\alpha^n$ is again in \mathcal{A} , and, because of σ -additivity, we see that $v(B_\alpha) = \sup_{n \in \mathbb{N}} v(B_\alpha^n)$. Thus the properties (D) and (DD) are indeed equivalent for σ -additive valuations (or, in view of Corollary 8, for finitely additive valuations on the standard cake $[0, 1]$ and $\mathcal{A} \supset \mathcal{I}([0, 1])$).

We will also need the opposite of the property (DD); to this end, recall Definition 3 of an atom. If A and B are atoms, then we have either $v(A \cap B) = 0$ or $v(A \cap B) = v(A) = v(B) > 0$; in the latter case, if $v(A \cap B) > 0$, we call the atoms *equivalent*. If A and B are nonequivalent, then A and $B \setminus A$ are still nonequivalent and disjoint. Iterating this procedure, we can always assume that countably many nonequivalent atoms $(A_n)_{n \in \mathbb{N}}$ are disjoint: Just replace the atoms by $A_1, A_2 \setminus A_1, \dots, A_{n+1} \setminus \bigcup_{i=1}^n A_i, \dots$.

Since $v(X) = 1$, a finitely additive valuation v can have at most n nonequivalent atoms such that $v(A) \geq \frac{1}{n}$, and so there are at most countably many atoms. Comparing Definition 8 which defines property (DD) with Definition 3 of an atom, it is clear that (DD) implies that v has no atoms. We will see in Theorem 11 that the converse implication holds as well.

Definition 9. Let v be a finitely additive valuation on the algebra \mathcal{A} over a(n abstract) cake X . The valuation v is **sliceable** if for any $\varepsilon > 0$, there are finitely many disjoint sets $B_i \in \mathcal{A}$, $i = 1, \dots, n$, $n = n(\varepsilon)$, such that $0 < v(B_i) \leq \varepsilon$ and $X = B_1 \cup \dots \cup B_n$.

A set $B \in \mathcal{A}$ is **v -sliceable** if the set-function $A \mapsto v(A \cap B)$ is sliceable.

We will now see that a sliceable finitely additive valuation enjoys property (DD), and *vice versa*, i.e., sliceability, atom-freeness, and property (DD) are pairwise equivalent for finitely additive valuations.

Theorem 11. Let v be a finitely additive valuation on an algebra \mathcal{A} over a(n abstract) cake X . The conditions (DD), ‘ v is sliceable,’ and ‘ v has no atoms’ are pairwise equivalent.

Proof. We start by showing that atom-freeness implies sliceability. Fix $\varepsilon > 0$.

Step 1: Let $Y \subseteq X$ be *any* subset, and assume that there is some $B \subseteq Y$, $B \in \mathcal{A}$, such that $v(B) > 0$. Define

$$\mathcal{F}^Y := \mathcal{F}_\varepsilon^Y := \{F \in \mathcal{A} \mid F \subseteq Y, 0 < v(F) \leq \varepsilon\}.$$

We claim that for the special choice $Y = B \in \mathcal{A}$ the family \mathcal{F}^B is not empty.

Since B is not an atom, there is some $F \subseteq B$, $F \in \mathcal{A}$, with $0 < v(F) < v(B)$.

If $v(F) \leq \varepsilon$, then $F \in \mathcal{F}^B$, and we are done.

If $v(F) > \varepsilon$, we assume, to the contrary that there is no subset $F' \subseteq F$, $F' \in \mathcal{A}$, with $0 < v(F') \leq \varepsilon$. Since F cannot be an atom, there is a subset $F' \subseteq F$ with $\varepsilon < v(F') < v(F)$ and $v(F \setminus F') > \varepsilon$. Iterating this with $F \rightsquigarrow F \setminus F'$ furnishes a sequence of disjoint sets $F_1 = F', F_2, F_3, \dots$ with $v(F_i) > \varepsilon$ for all $i \in \mathbb{N}$. This is impossible since $v(F) < \infty$. So we can find some $F' \subseteq F \subseteq B$ with $0 < v(F') \leq \varepsilon$, i.e., \mathcal{F}^B is not empty.

Step 2: Define a(n obviously monotone) set-function $c(Y) := \sup_{C \in \mathcal{F}^Y} v(C)$ for any $Y \subseteq X$; as usual, $\sup \emptyset = 0$. Since \mathcal{F}^X is not empty, we can pick some $B_1 \in \mathcal{F}^X$ such that $\frac{1}{2}c(X) < v(B_1) \leq \varepsilon$.

If $v(X \setminus B_1) \leq \varepsilon$, we set $B_2 := X \setminus B_1$; otherwise, we can pick some $B_2 \in \mathcal{F}^{X \setminus B_1}$ such that $\frac{1}{2}c(X \setminus B_1) < v(B_2) \leq \varepsilon$.

In general, if $v(X \setminus (B_1 \cup \dots \cup B_n)) \leq \varepsilon$, we set $B_{n+1} = X \setminus (B_1 \cup \dots \cup B_n)$; otherwise, we pick

$$(2) \quad B_{n+1} \in \mathcal{F}^{X \setminus (B_1 \cup \dots \cup B_n)} \quad \text{such that} \quad \frac{1}{2}c(X \setminus (B_1 \cup \dots \cup B_n)) < v(B_{n+1}) \leq \varepsilon.$$

We are done if this procedure stops after finitely many steps; otherwise, we get a sequence of disjoint sets B_1, B_2, \dots satisfying (2). Define $B_\infty := X \setminus \bigcup_n B_n$. This set need not be in \mathcal{A} , but we still have, because of (2),

$$c(B_\infty) \leq c(X \setminus (B_1 \cup \dots \cup B_m)) \leq 2v(B_{m+1}) \xrightarrow{m \rightarrow \infty} 0$$

since the series

$$\sum_{n \in \mathbb{N}} v(B_n) = \sup_N \sum_{n=1}^N v(B_n) = \sup_N v\left(\bigcup_{n=1}^N B_n\right) \leq v(X)$$

converges. In particular, $\lim_{n \rightarrow \infty} v(X \setminus \bigcup_{i=1}^n B_i) = 0$.

Using again the convergence of the series $\sum_n v(B_n)$, we find some $N = N(\varepsilon)$ such that $\sum_{n > N} v(B_n) \leq \varepsilon$, hence B_1, B_2, \dots, B_N and $X \setminus \bigcup_{n=1}^N B_n$ are the desired small pieces of X . This completes the proof that v is sliceable.

We now show that sliceability implies condition (DD). Let $B \in \mathcal{A}$ with $v(B) > 0$. Since the ‘relative’ finitely additive valuation $v_B(A) := v(A \cap B)/v(B)$ inherits the nonatomic property from v , it is clearly enough to show that for every $\alpha \in (0, 1)$, there is an increasing sequence

$$B_\alpha^1 \subset B_\alpha^2 \subset B_\alpha^3 \subset \dots, \quad B_\alpha^n \in \mathcal{A} : \sup_{n \in \mathbb{N}} v(B_\alpha^n) = \alpha,$$

which is the property (DD) relative to the full cake X only.

Since v is sliceable, there are mutually disjoint sets $C_1^n, \dots, C_N^n \in \mathcal{A}$, where $N = N(n)$, $X = \bigcup_{i=1}^N C_i^n$, and $v(C_i^n) < \frac{1}{n}$.

Let $k = \lfloor 1/\alpha \rfloor + 1$. Set $B_k := C_1^k \cup \dots \cup C_{M(k)}^k$, where $M(k) \in \{1, \dots, N(k)\}$ is the unique number such that

$$\sum_{i=1}^{M(k)} v(C_i^k) \leq \alpha < \sum_{i=1}^{M(k)+1} v(C_i^k) \leq \sum_{i=1}^{M(k)} v(C_i^k) + \frac{1}{k}.$$

By construction, $\alpha \geq v(B_k) = \sum_{i=1}^{M(k)} v(C_i^k) > \alpha - \frac{1}{k}$. Thus, we can iterate this procedure, considering $X \setminus B_k$ and constructing a set $D_{k+1} \subseteq X \setminus B_k$ that satisfies

$$(\alpha - v(B_k)) \geq v(D_{k+1}) > (\alpha - v(B_k)) - \frac{1}{k+1}.$$

For $B_{k+1} := B_k \cup D_{k+1}$, we get $\alpha \geq v(B_{k+1}) > \alpha - \frac{1}{k+1}$.

The sequence B_{k+i} , $i \in \mathbb{N}$, satisfies $v(B_{k+i}) \uparrow \alpha$, i.e., $B_\alpha^n = B_{k+n}$ is the sequence of sets we need to have property (DD).

As mentioned earlier, (DD) implies atom-freeness, which completes this proof. \square

Since for a σ -additive valuation on a σ -algebra \mathcal{A} , properties (D) and (DD) are equivalent, we immediately get:

Corollary 12. *Let v be a σ -additive valuation on a σ -algebra \mathcal{A} over an abstract cake X . The conditions (D), (DD), ‘ v is sliceable,’ and ‘ v has no atoms’ are pairwise equivalent.*

If A_1, A_2, \dots is an enumeration of the nonequivalent atoms of the σ -additive valuation v , then $A_\infty := X \setminus \bigcup_{n \in \mathbb{N}} A_n \in \mathcal{A}$, and we can restate Corollary 12 in the form of a decomposition theorem.

Corollary 13. *Let v be a σ -additive valuation on a σ -algebra \mathcal{A} over a(n abstract) cake X . Then X can be written as a disjoint union of a v -sliceable set A_∞ and at most countably many atoms A_1, A_2, \dots .*

3. WHICH PIECES SHOULD BE ADMISSIBLE?

In the cake-cutting literature, a great variety of different definitions have been used for the set \mathcal{P} of admissible pieces of cake. We first collect the most commonly used definitions for \mathcal{P} , 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{P} .

Typical choices for the set \mathcal{P} containing all admissible pieces of a standard cake $[0, 1]$ are

- (1) all finite unions of intervals from $[0, 1]$, i.e., the family $\mathcal{I}([0, 1])$ defined earlier on page 6;
- (2) all countable unions of intervals from $[0, 1]$, i.e., $\mathcal{I}([0, 1])^{\mathbb{N}} = \{\bigcup_{i \in \mathbb{N}} I_i \mid I_i \in \mathcal{I}([0, 1])\}$;
- (3) the Borel σ -algebra over $[0, 1]$, i.e., $\mathcal{B}([0, 1])$;

- (4) the set of all Lebesgue sets over $[0, 1]$, i.e., $\mathcal{L}([0, 1])$;² or
 (5) the power set $\mathfrak{P}([0, 1])$ of $[0, 1]$.

Assuming $\mathcal{P} = \mathcal{I}([0, 1])$ is common among papers that consider only finite cake-cutting protocols. Such protocols can make only 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 $\mathcal{P} = \mathcal{I}([0, 1])$ include Woeginger and Sgall [2007], Stromquist [2008], Lindner and Rothe [2009], Procaccia [2009], Walsh [2011], Cohler et al. [2011], Bei et al. [2012], Cechlárová and Pillárová [2012a], Brams et al. [2012], Cechlárová et al. [2013], Chen et al. [2013], Brânzei and Miltersen [2013], Aziz and Mackenzie [2016a,b, 2020], Edmonds and Pruhs [2006], and Aziz and Mackenzie [2016b].

As a special case, valuation functions may even be restricted to single intervals, which is done by Cechlárová and Pillárová [2012b] and Aumann and Dombb [2010]. 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. [2013] extend \mathcal{P} to contain countably infinite unions of intervals, i.e., $\mathcal{I}([0, 1])^{\mathbb{N}}$.

Authors assuming $\mathcal{P} = \mathcal{B}([0, 1])$ include Stromquist and Woodall [1985], Deng et al. [2009], and Segal-Halevi et al. [2017].

Works using $\mathcal{P} = \mathcal{L}([0, 1])$ include those by Reijnierse and Potters [1998], Arzi et al. [2011], and Robertson and Webb [1997]. Additionally, several authors do not explicitly make the assumption $\mathcal{P} = \mathcal{L}([0, 1])$, 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 $[0, 1]$. This or a similar assumption is made by Brams et al. [2003, 2006, 2008, 2013], Robertson and Webb [1998], Webb [1997], Aumann et al. [2013], Brânzei et al. [2016], and Caragiannis et al. [2011].

Papers that assume $\mathcal{P} = \mathfrak{P}([0, 1])$ include those by Maccheroni and Marinacci [2003], Sgall and Woeginger [2007], Saberi and Wang [2009], Manabe and Okamoto [2010], and Aumann et al. [2014].

Finally, several works, including those by Dubins and Spanier [1961], Barbanel [1996a,b], Zeng [2000], and Brams and Taylor [1995a], define the set of admissible pieces of cake to be some (σ -)algebra (not necessarily Borel) over $[0, 1]$.

Note that each of the sets $\mathcal{I}([0, 1])$, $\mathcal{B}([0, 1])$, $\mathcal{L}([0, 1])$, and $\mathfrak{P}([0, 1])$ is an algebra over $[0, 1]$, and all, except $\mathcal{I}([0, 1])$, are also σ -algebras over $[0, 1]$. However, $\mathcal{I}([0, 1])^{\mathbb{N}}$ 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 14. $\mathcal{I}([0, 1]) \stackrel{(a)}{\subsetneq} \mathcal{I}([0, 1])^{\mathbb{N}} \stackrel{(b)}{\subsetneq} \mathcal{B}([0, 1]) \stackrel{(c)}{\subsetneq} \mathcal{L}([0, 1]) \stackrel{(d)}{\subsetneq} \mathfrak{P}([0, 1])$.

Proof. We start with proving (a): $\mathcal{I}([0, 1]) \subsetneq \mathcal{I}([0, 1])^{\mathbb{N}}$. Obviously, $\mathcal{I}([0, 1]) \subseteq \mathcal{I}([0, 1])^{\mathbb{N}}$ 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 $I = \bigcup_{i \in \mathbb{N} \cup \{0\}} [3 \cdot 2^{-i-2}, 2^{-i}]$. It is clear that $I \in \mathcal{I}([0, 1])^{\mathbb{N}}$ is true, as I is a countable union of intervals. However, it holds that $I = [3/4, 1] \cup [3/8, 1/2] \cup \dots$, i.e., I cannot be written as a finite union of intervals, as all these subintervals are pairwise disjoint. Hence, $I \notin \mathcal{I}([0, 1])$, so $\mathcal{I}([0, 1]) \subsetneq \mathcal{I}([0, 1])^{\mathbb{N}}$, and we have shown (a).

In Lemma 9 and Definition 6, we have seen that $\mathcal{B}([0, 1]) = \sigma(\mathcal{Q}([0, 1]))$ where $\mathcal{Q}([0, 1])$ is the family of all intervals within $[0, 1]$. Since a σ -algebra is stable under (finite and countable) unions, we get $\mathcal{I}([0, 1]) \subseteq \sigma(\mathcal{Q}([0, 1])) = \mathcal{B}([0, 1])$. Using again the stability of a σ -algebra under countable unions, we arrive at $\mathcal{I}([0, 1])^{\mathbb{N}} \subseteq \mathcal{B}([0, 1])$.

Since, however, $\mathbb{Q} \cap [0, 1] \in \mathcal{B}([0, 1])$ is true, as $\mathbb{Q} \cap [0, 1]$ can be written as a countable union of intervals that each contain one element, it must hold that $\overline{\mathbb{Q} \cap [0, 1]} \in \mathcal{B}([0, 1])$ by the definition of a σ -algebra. However, the irrational numbers $\overline{\mathbb{Q} \cap [0, 1]}$ in $[0, 1]$ cannot be written as a countable

²Recall that a set \tilde{B} is a Lebesgue set if, and only if, there is a Borel set B such that the symmetric difference $\tilde{B} \Delta B := (B \setminus \tilde{B}) \cup (\tilde{B} \setminus B) \subseteq N$ is contained in a Borel set N with Lebesgue measure $\lambda(N) = 0$. We will see in Theorem 14 that there are indeed Lebesgue sets that are not Borel sets.

union of intervals, since every interval containing more than one element immediately contains a rational number. Therefore, $\mathcal{I}([0, 1])^{\mathbb{N}}$ is not an algebra and $\mathcal{I}([0, 1])^{\mathbb{N}} \neq \mathcal{B}([0, 1])$ holds, proving (b).

The inclusion $\mathcal{B}([0, 1]) \subseteq \mathcal{L}([0, 1])$ holds by definition, as all Borel sets are also Lebesgue sets. However, there are Lebesgue sets that are not Borel sets: Observe that the cardinality of $\mathcal{L}([0, 1])$ is the cardinality of $\mathfrak{P}([0, 1])$ (which is $2^{\mathfrak{c}} > \mathfrak{c}$), whereas there are only continuum-many (i.e., \mathfrak{c} , the cardinality of $[0, 1]$) Borel sets (see Schilling, 2017, Appendix G, Corollary G.7). This proves (c). An alternative direct construction can be based on the Cantor function, also known as the *devil's staircase* (see Schilling and Kühn, 2021, p. 153, Example 7.20).

Finally, the power set $\mathfrak{P}([0, 1])$ trivially contains all other families of sets considered earlier. Nevertheless, there are sets in $\mathfrak{P}([0, 1])$ that are not Lebesgue sets, for example the Vitali set that we introduced in Example 2, so $\mathcal{L}([0, 1]) \neq \mathfrak{P}([0, 1])$, and we have (d). \square

4. DISCUSSION

Taking $\mathcal{P} = \mathcal{I}([0, 1])$ as domain for a valuation v and a protocol involving a finite number of cuts is always possible; this remains true for open-ended protocols that stop after a finite but *a priori* unknown number of steps. If the protocol is infinite, the naive choice $\mathcal{P} = \mathcal{I}([0, 1])^{\mathbb{N}}$ is problematic, as $\mathcal{I}([0, 1])^{\mathbb{N}}$ is not an algebra and thus does not even satisfy the minimum requirements for \mathcal{P} as described in the first paragraph of Section 2.1.

From a theoretical point of view, however, the choice $\mathcal{P} = \mathcal{I}([0, 1])$ 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{P} may be desirable, perhaps even larger than $\mathcal{I}([0, 1])^{\mathbb{N}}$, which (as we have seen) has disqualified itself.

We start our discussion by explicating why $\mathcal{P} = \mathfrak{P}([0, 1])$ is a bad choice and we then provide arguments for a better option, namely the Borel σ -algebra $\mathcal{P} = \mathcal{B}([0, 1])$.

4.1. Taming $\mathfrak{P}([0, 1])$ with Exotic Valuations via Banach Limits. If one boldly desires to define valuation functions on the set $\mathfrak{P}([0, 1])$ 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 boxes, as depicted in Figure 2, are not capable of evaluating every piece of cake in $\mathfrak{P}([0, 1])$. Hence, in this section we aim to define a valuation function capable of evaluating *every* possible piece of cake in $\mathfrak{P}([0, 1])$.

Let us begin with a negative result.

4.1.1. A Negative Result. Using axiomatic set theory one can show that there cannot be a valuation v of the standard cake $[0, 1]$ which

- a) is defined on all of $\mathfrak{P}([0, 1])$,
- b) is σ -additive, and
- c) is divisible, hence satisfies $v(\{x\}) = 0$ for any $x \in [0, 1]$.

The requirements a)–c) are a consequence of a result by Ulam, and it requires that the continuum hypothesis holds true, see the books by Oxtoby [1980, p. 26, Proposition 5.7] or Schilling and Kühn [2021, pp. 132–3, Example 6.15].

We may relax on b), i.e., the σ -additivity, and look for **finitely** additive valuations if we want to admit all pieces of cake. Let us formally define a valuation function μ on $\mathfrak{P}([0, 1])$ satisfying the requirements (M), (A), and (D) from Definition 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 $[0, 1]$. For every $A \subseteq [0, 1]$, we define a mapping $f_A: \mathbb{N} \rightarrow [0, 1]$ with

$$(3) \quad n \mapsto f_A(n) = \frac{|A \cap \{x_1, \dots, x_n\}|}{n},$$

where $|B|$ denotes the cardinality of any set B . That is, $f_A(n)$ 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 $\lim_{n \rightarrow \infty} f_A(n)$ does exist, but it may not exist for other sets A . We can, however, use the Banach limits, that we will now introduce.

4.1.2. *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, 1991, Theorem 3.2), so we need to go on a quick excursion into functional analysis.

Theorem 15. *Assume that $(Y, \|\cdot\|)$ is a normed vector space and $L : M \rightarrow \mathbb{R}$ a linear functional, which is defined on a linear subspace $M \subseteq Y$ satisfying $|Lx| \leq \kappa\|x\|$ for all $x \in M$ with a universal constant $\kappa = \kappa_L \in (0, \infty)$. Then there is an extension $\hat{L} : Y \rightarrow \mathbb{R}$ such that \hat{L} is again linear and satisfies $|\hat{L}x| \leq \kappa\|x\|$ for all $x \in Y$ with the same constant $\kappa = \kappa_L$ as before.*

With a little more effort, but essentially the same proof, we can replace the norm $\|x\|$ (respectively, $\kappa\|x\|$) by a general sublinear map $p : Y \rightarrow \mathbb{R}$. Sublinear means that $p(\alpha x) = \alpha p(x)$ and $p(x + y) \leq p(x) + p(y)$ for all $x, y \in Y$ and $\alpha \geq 0$. In this case, the extension of $Lx \leq p(x)$ satisfies $-p(-x) \leq \hat{L}x \leq p(x)$. Note that p is only positively homogeneous, i.e., it may happen that $-p(-x) \neq p(x)$.

The proof is nonconstructive and, at least for nonseparable spaces Y , 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. Note that $(\ell^\infty([0, \infty)), \|\cdot\|_\infty)$ is a nonseparable space.

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 $L(x) := \lim_{n \rightarrow \infty} x_n = x$ exists in the usual sense. It is common to write $c([0, \infty)) = \{x \in \ell^\infty([0, \infty)) \mid \lim_{n \rightarrow \infty} x_n \text{ exists}\}$. Clearly, $\lim_{n \rightarrow \infty} x_n = \limsup_{n \rightarrow \infty} x_n \leq \sup_{n \in \mathbb{N}} x_n$, so that L is a bounded linear functional on $M = c([0, \infty)) \subset Y = \ell^\infty([0, \infty))$, and we can extend it to all of Y as the **Banach limit**, i.e.,

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

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

$$\liminf_{n \rightarrow \infty} x_n \leq \text{LIM}_{n \rightarrow \infty} x_n \leq \limsup_{n \rightarrow \infty} x_n.$$

The construction of Banach limits is a typical application of the Hahn–Banach extension theorem, hence the axiom of choice. The appearance of these two concepts in this context is not an accident. The seminal paper of Banach [1923] (see also Banach, 1932, Chapter II.§1) proves what we now call the ‘Hahn–Banach extension theorem for linear functionals’ in order to solve the *problème de la mesure* by Lebesgue [1904, Chapter VII.ii] which asks for the existence of an additive, or σ -additive, translation invariant measure on $\mathfrak{P}(\mathbb{R}^n)$. The answer depends on the dimension: In dimension $n \geq 3$, it is always negative (because of the Banach–Tarski paradox), whereas in dimensions 1 and 2 it is negative if the measure is to be σ -additive (because of Vitali-type constructions, cf. Example 2). More on this can be found in the books by Wagon [1985, Chapter 10] and Schilling and Kühn [2021, Example 7.31].

There is a deep connection between the underlying group structure of the space \mathbb{R}^n and Lebesgue’s measure problem (this was discovered by von Neumann, 1929). Following M. M. Day, a group \mathbb{G} which allows for finitely additive, (left-)translation invariant measures on all of $\mathfrak{P}(\mathbb{G})$ is nowadays called **amenable** – a pun combining the actual meaning of the word (‘nice, comfortable’) with its pronunciation which reminds of ‘mean value’ or measure. The axiom of choice, which is needed for Hahn–Banach, can also be used to construct extensions of measures defined on a sub-algebra \mathcal{A}_0 of an algebra \mathcal{A} . It is known that this extendability, essentially, is equivalent to the Hahn–Banach theorem (cf. Wagon, 1985, Theorem 10.11 and Corollary 13.6) describing its axiomatic strength.

4.1.3. *From Banach Limits to Valuation Functions.* Having defined and discussed Banach limits, we will now use them to construct, based on the function f_A defined in (3), a valuation function

$$\mu: \mathfrak{P}([0, 1]) \rightarrow [0, 1], \quad A \mapsto \text{LIM}_{n \rightarrow \infty} f_A(n).$$

It is clear that $\mu(A)$ is additive since $A \mapsto f_A(n)$ is additive (for every fixed n) and both the limit and the Banach limit are additive, so property (A) from Definition 2 is satisfied. In the following lemma we show that μ satisfies property (D). At first glance, this seems to contradict Corollary 8. But divisibility (D) involves the domain of the valuation, and the proof of the lemma shows that we have almost no control on the set $A_\alpha \subseteq A$ which achieves divisibility. That means, the following phenomenon is symptomatic for having a ‘too big domain.’

Lemma 16. *For every $A \in \mathfrak{P}([0, 1])$ with $\mu(A) > 0$ and every real number $\alpha \in [0, 1]$, there exists a subset $A_\alpha \subseteq A$ in $\mathfrak{P}([0, 1])$ such that $\mu(A_\alpha) = \alpha\mu(A)$.*

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. By assumption, $A \cap \{x_1, x_2, \dots, x_n\} = \{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n\}$, and so $\mu(A) = \text{LIM}_{n \rightarrow \infty} \frac{1}{n} |\{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n\}|$. We have to construct a set $B \in \mathfrak{P}(A)$ such that $\text{LIM}_{n \rightarrow \infty} f_B(n) = \alpha\mu(A)$ for fixed $\alpha \in [0, 1]$.

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

$$\text{on the one hand } \frac{k}{n+1} < \frac{k}{n} \quad \text{and on the other hand } \frac{k}{n} < \frac{k+1}{n+1},$$

i.e., the quantity $f_B(n) = \frac{1}{n} |B \cap \{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n\}|$ decreases if we jack up $n \rightarrow n+1$ and the numerator does not increase, i.e., if $i(m+1) > n+1$ or if $x_{i(m+1)} = x_{n+1} \notin B$, and it increases if we jack up $n \rightarrow n+1$ and $x_{i(m+1)} = x_{n+1} \in B$.

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

Since $\mu(A) = \text{LIM}_{n \rightarrow \infty} \frac{1}{n} |\{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n\}|$ and $0 < \alpha < 1$ we must have $\frac{1}{n(1)} |\{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n(1)\}| \geq \alpha\mu(A)$ for some $n(1) \in \mathbb{N}$. Define $B_{n(1)} = \{x_{i(1)}, \dots, x_{i(m)} \mid i(m) \leq n(1)\}$, and assume that we have already found a set B_n such that $f_{B_n}(n) \geq \alpha\mu(A)$. Because of the observation at the beginning of the proof, the numbers

$$\ell_{n+1} := \min \{k > n \mid f_{B_n}(k) \leq \alpha\mu(A)\} \quad \text{and}$$

$$u_{n+1} := \min \{k > \ell_{n+1} \mid f_{A_k}(k) \geq \alpha\mu(A) \text{ for } A_k = B_n \cup \{x_{i(\ell_{n+1}+1)}, \dots, x_{i(k)}\}\}$$

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

$$B_{n+1} = B_n \cup \{x_{i(\ell_{n+1}+1)}, \dots, x_{i(u_{n+1})}\}$$

finishes the recursion, and we can define $B = \bigcup_{n \geq n(1)} B_n$.

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

We now provide a counterexample that shows that μ is not σ -additive. To do so, we define $A_0 = [0, 1] \setminus \{x_i \mid i \in \mathbb{N}\}$ 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

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

which means that μ is not σ -additive.

The valuation function μ defined above may seem to be attractive 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, there are multiple drawbacks. 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 4.1.2. Furthermore, until now no explicit nontrivial example of a Banach

limit is known. Hence, we cannot calculate $\mu(A)$ for $A \in \mathfrak{P}([0, 1])$ 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 $\mathcal{P} = \mathfrak{P}([0, 1])$, it is actually not useful for our purposes. 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 $[0, 1]$ 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 impractical and not feasible to use.

Hence, finding practically usable valuation functions defined on $\mathfrak{P}([0, 1])$ seems to remain an open problem. Nonetheless, this section showed that defining more complex valuation functions (compared to the valuation functions represented via boxes) does not solve our initial problem on $\mathfrak{P}([0, 1])$. Therefore, in the next section we discuss an alternative solution, namely, reducing \mathcal{P} in size from $\mathfrak{P}([0, 1])$ to a smaller subfamily contained in $\mathfrak{P}([0, 1])$.

4.2. Borel σ -Algebra. We recommend to use $\mathcal{P} = \mathcal{B}([0, 1])$ as the most useful family of all admissible pieces of cake. As shown in Theorem 14, the Borel σ -algebra $\mathcal{B}([0, 1])$ (strictly) contains $\mathcal{I}([0, 1])$ as well as $\mathcal{I}([0, 1])^{\mathbb{N}}$, but is strictly smaller than $\mathcal{L}([0, 1])$ and $\mathfrak{P}([0, 1])$.

In general, the Borel σ -algebra can become quite large and complicated if the base set is not countable, as is the case for $[0, 1] \subset \mathbb{R}$. In particular, one needs transfinite induction to ‘construct’ all Borel sets. This means that, in general, we cannot construct a valuation μ on $\mathcal{B}([0, 1])$ by explicitly assigning a value $\mu(A)$ to every element $A \in \mathcal{B}([0, 1])$ nor give a recursive algorithm to construct $\mu(A)$, as the σ -algebra is simply too large. Instead, one can describe the valuation on a suitable generator of the σ -algebra and use Carathéodory’s extension theorem, stated previously as Theorem 10.

Let us show here that the box-based valuation functions are σ -additive valuations on $\mathcal{I}([0, 1])$ and that $\mathcal{I}([0, 1])$ is an algebra. In this case, we can use Carathéodory’s extension theorem to extend the valuation functions to measures on $\sigma(\mathcal{I}([0, 1])) = \mathcal{B}([0, 1])$. Since the valuation functions are finite, it follows that this extension is unique. Hence, by providing a box-based valuation function, we obtain a unique measure on $\mathcal{B}([0, 1])$. Thus $\mathcal{P} = \mathcal{B}([0, 1])$ 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 $[0, 1]$ 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

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

where $\bigcup_{i=1}^n I_i = [0, 1]$ 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. This gives the following weight function

$$p(x) := \frac{1}{\psi_\mu} \sum_{i=1}^n \frac{\psi_i}{\lambda(I_i)} \mathbf{1}_{I_i}(x) = \frac{1}{\psi_\mu} \sum_{i=1}^n \frac{\psi_i}{b_i - a_i} \mathbf{1}_{I_i}(x).$$

Note that $p(x)$ is an integrable function, which is a probability density, i.e., $\int_0^1 p(x) dx = 1$. Since μ is a Lebesgue measure with a weight, we cannot define it on all of $\mathfrak{P}([0, 1])$, but we may extend it easily onto $\mathcal{B}([0, 1])$ using integration: Define $\mu: \mathcal{B}([0, 1]) \rightarrow [0, 1]$ as

$$B \mapsto \mu(B) := \int_B p(x) dx.$$

In particular, if $B \in \mathcal{I}([0, 1])$ is a finite union of intervals in $[0, 1]$, we see that

$$B \cap I_i = \bigcup_{j=1}^{n(B,i)} \langle c_j^i, d_j^i \rangle, \quad i = 1, 2, \dots, n$$

for suitable $n(B, i) \in \mathbb{N}$, and

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

Example 17. Referring back to the box-based valuation function ν from Figure 2 on page 3, we obtain

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

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 $B = [1/10, 1/4]$, we obtain

$$\begin{aligned} \nu(B) &= \frac{1}{\psi_\nu} \sum_{i=1}^6 \left[\frac{\psi_i}{b_i - a_i} \sum_{j=1}^{n(B,i)} (d_j^i - c_j^i) \right] \\ &= \frac{1}{17} \left(\frac{2}{1/6} \cdot (1/6 - 1/10) + \frac{1}{1/6} \cdot (1/5 - 1/6) + \frac{5}{1/6} \cdot 0 + \frac{2}{1/6} \cdot 0 + \frac{4}{1/6} \cdot 0 + \frac{3}{1/6} \cdot 0 \right) \\ &= \frac{1}{17}. \end{aligned}$$

Summing up, $\mathcal{B}([0, 1])$ is recommended as a very good choice for \mathcal{P} , since this choice enables us to use box-based valuation functions and their extensions as measures. It also enables us to use any probability density – not only piecewise continuous densities – to define a divisible valuation function on $\mathcal{B}([0, 1])$ which is an absolutely continuous probability measure with respect to Lebesgue measure. Since Lebesgue null sets are subsets of Borel null sets, and since an absolutely continuous valuation attaches value zero to any Borel null set, a further extension to $\mathcal{L}([0, 1])$ is also possible, but the enrichment by subsets of Borel null sets (which are evaluated zero) has no additional benefit.

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 or only finitely additive?

A related interesting question is:

- (3) Which continuity property should be used for a valuation?

For the standard cake $[0, 1]$, the natural choices are either *divisibility* (D) or *absolute continuity* with respect to Lebesgue measure, see p. 10. Obviously, absolute continuity implies continuity. There is a partial converse to this assertion: The notions of continuity and divisibility coincide (cf. Corollary 8) and the distribution function $F_\nu(x)$ of a continuous valuation can be represented as a sum of the form $F_\nu(x) = \int_0^x f(t) dt + v^{\text{sc}}([0, x])$; this means that it has an absolutely continuous part and a continuous-singular part, see the discussion in the paragraph on ‘General Measures’ following Definition 7. For an abstract cake, one should replace divisibility (D) by the notion of sliceability, which is equivalent to condition (DD) by Theorem 11, see Section 2.4 and Schilling and Stoyan [2016].

While one can define the Dirac and counting measures for all sets in $\mathfrak{P}(X)$, there is no way to define a geometrically sensible (and σ -additive [in dimensions one and two] or finitely additive [in all higher dimensions]) 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 non-measurable sets (cf. Ciesielski, 1989, p. 55).

Our findings result in concrete recommendations for cake-cutters. For a finitely additive valuation ν on the standard cake $[0, 1]$ (or indeed any one-dimensional cake) equipped with the

algebra $\mathcal{I}([0, 1])$ generated by the intervals, divisibility (D) is equivalent to atom-freeness or the continuity of the distribution function $F_v(x) = v([0, x])$, cf. Corollary 8. For an abstract cake and a finitely additive valuation v , divisibility (D) should be replaced by sliceability (DD), which is equivalent to v being atom-free; if v is even σ -additive, conditions (D) and (DD) coincide, see Theorem 11 and Corollary 12.

All of this breaks down, however, if we consider finitely additive valuations on too big domains, say $\mathcal{P} = \mathfrak{P}(X)$: Even for the standard cake there are divisible, finitely additive but not σ -additive valuations, see Lemma 16.

We have also discussed in detail the measure-theoretic notions and results that are relevant for the foundations of cake-cutting, for both the standard cake and abstract cakes, including the notions of σ -additivity, the Borel σ -algebra, and Carathéodory's extension theorem (Theorem 10). We emphasized the importance of the Hahn–Banach theorem and the underlying axiom of choice if one needs to evaluate arbitrary pieces of cake which are not Borel or Lebesgue sets.

Banach, who can be seen as one of the founding fathers of the field of cake-cutting,³ might perhaps have appreciated the close connection between his work in measure theory and in cake-cutting. For future work, we suggest to study which implications our findings may have on existing or on yet-to-be-designed cake-cutting algorithms.

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 $\mathfrak{P}(X)$ is too general to define a (practically or theoretically) useful valuation function on it. And finally, we have reasoned why we recommend the Borel σ -algebra $\mathcal{B}(X)$ as a very good choice and how to construct, using Carathéodory's extension theorem, a measure on $\mathcal{B}(X)$ that cake-cutters can use to handle their box-based and even more general valuation functions.

For a pragmatic approach to cake-cutting on the standard cake $[0, 1]$, the following five points are important:

1. If one is interested in a fixed number of players and a fixed number of cuts, any additive valuation v defined on the algebra of intervals $\mathcal{I}([0, 1])$ will do.
2. If the players take rounds and if the protocol is open-ended (i.e., finite unbounded) or even infinite (recall the examples mentioned in Section 1 and in Footnote 1 of Section 2.3), the finite additivity of the valuation v needs to be strengthened to σ -additivity, and the domain of the valuation should contain the Borel σ -algebra $\mathcal{B}([0, 1])$ – this is the smallest σ -algebra containing $\mathcal{I}([0, 1])$.
3. If the valuation v on $\mathcal{I}([0, 1])$ is divisible, measure theory guarantees that one is automatically in the situation described in item 2, i.e., the proper domain of (the extension of) v is the Borel σ -algebra $\mathcal{B}([0, 1])$.
4. If one wants to extend the domain of the valuation v beyond $\mathcal{B}([0, 1])$, things become difficult: On the one hand, it is quite tricky to ‘construct’ sensible valuations – unless we are happy with ‘rather simple’ valuations like countable sums of point masses $v = \sum_{i \in \mathbb{N}} p_i \delta_{x_i}$, $\sum_{i \in \mathbb{N}} p_i = 1$, $(x_i)_{i \in \mathbb{N}} \subseteq [0, 1]$, but these are obviously not divisible – and, on the other hand, they are not well-behaved, touching the very basis of axiomatic set theory.
5. The tools provided by measure theory are powerful enough to handle even abstract cakes.

Acknowledgement. We thank William S. Zwicker for helpful discussions. We are grateful for the constructive refereeing process which helped to improve the presentation of our contribution.

Funding. This work was supported in part by grants of the Deutsche Forschungsgemeinschaft (DFG) RO-1202/14-2, RO-1202/21-1 and the DFG and Narodowe Centrum Nauki (NCN) joint initiative ‘Beethoven 3 classic’ SCHI-419/11-1, NCN 2018/31/G/ST1/02252. The second author

³Indeed, Steinhaus [1948] presents the so-called last-diminisher procedure that is due to his students Banach and Knaster and guarantees a proportional division of the cake among any number of players.

was a member of the PhD-programme ‘Online Participation,’ supported by the North Rhine-Westphalian funding scheme ‘Forschungskollegs.’

6. APÜENDIX: AN ALTERNATIVE EXAMPLE ILLUSTRATING THE CANTOR DUST

Example 18 (Cantor dust; Cantor’s ternary set). Write the elements $x \in [0, 1]$ of the cake $[0, 1]$ as ternary numbers, i.e., in the form

$$x = \sum_{n \in \mathbb{N}} \frac{x_n}{3^n} \simeq 0.x_1x_2x_3\dots, \quad \text{where } x_n \in \{0, 1, 2\},$$

and consider the set $C_{1/3}$ comprising all x whose ternary expansion contains the digits ‘0’ or ‘2’ only. To enforce uniqueness, identify expressions of the form $0.***1000\dots$ with $0.***0222\dots$. The set $C_{1/3}$ is not countable since there is a bijection between $C_{1/3}$ and $[0, 1]$: Take any $x = 0.x_1x_2x_3\dots \in C_{1/3}$ and read $\hat{x} := 0.\frac{x_1}{2}\frac{x_2}{2}\frac{x_3}{2}\dots$ as **dyadic** expansion of an arbitrary element $\hat{x} \in [0, 1]$.

The set $C_{1/3}$ is the so-called **Cantor set** from Example 1. Think of its elements as ‘cream’ pieces within the cake $[0, 1]$, and imagine two players, taking turns in picking pieces of cake; for some reason (that their cardiologist elaborated on in detail) they have to avoid the cream altogether. For this, they are allowed to make two cuts, taking out an interval from the cake.⁴

The optimal strategy is to take, in each round, the largest (necessarily open) interval between two cream pieces. From the triadic expansion, we see that, at each stage of the game, the maximum distance between two cream pieces is $0.\underbrace{***}_n2000\dots - 0.\underbrace{***}_n0222\dots = 0.\underbrace{0001}_n000\dots \simeq 3^{-n}$, and this situation appears exactly 2^{n-1} times, since we have 2^{n-1} choices for the leading $n-1$ digits denoted by the wildcard ‘***’ – to wit, the pieces taken out are always the middle thirds of the largest remaining interval of cake:

$$\begin{aligned} A_0 = [0, 1] &\xrightarrow{(*)} A_1 = A_0 \setminus (1/3, 2/3) = [0, 1/3] \cup [2/3, 1] \\ &\xrightarrow{(**)} A_2 = [0, 1/9] \cup [2/9, 1/3] \cup [2/3, 7/9] \cup [8/9, 1]. \end{aligned}$$

At the step marked (*) Player 1 takes the first middle third, at the (double) step marked (**) Player 2 and then Player 1 take the middle thirds of the remaining intervals, etc.

If this procedure is repeated on and on, we remove countably many intervals from $[0, 1]$ and end up with the **Cantor (ternary) set** $C_{1/3} = \bigcap_{n \in \mathbb{N}} A_n$ from Example 1, see Figure 1.

We can use the triadic expansion also to assign a unique code to the removed piece: $I_{t_1t_2\dots t_n2}$ denotes the newly removed piece of cake at stage $n+1$ and the $t_1, \dots, t_n \in \{0, 2\}$ mark the right-end point of the interval using the triadic expansion: $\sup I_{t_1t_2\dots t_n2} = \sum_{i=1}^n t_i 3^{-i} + 2 \cdot 3^{-n-1}$. This allows us to come up with a formula for the Cantor function from Example 3:

$$(4) \quad \text{on all of } I_{t_1t_2\dots t_n2} \text{ the function } V \text{ has the value } \sum_{i=1}^n \frac{t_i}{2} 2^{-i} + 3^{-n-1}.$$

We refer to [Schilling and Kühn, 2021, Sec. 2.5, 2.6] for a full discussion of this.

REFERENCES

- 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.
- 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.
- 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.

⁴This is, of course, a non-standard cake-cutting protocol.

- 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.
- A. Austin. Sharing a cake. *Mathematical Gazette*, 66(437):212–215, 1982.
- 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. ACM Press, June 2016a.
- H. Aziz and S. Mackenzie. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *Proceedings of the 57th IEEE Symposium on Foundations of Computer Science*, pages 416–427. IEEE Computer Society Press, October 2016b.
- H. Aziz and S. Mackenzie. A bounded and envy-free cake cutting algorithm. *Communications of the ACM*, 63(4):119–126, 2020.
- E. Balkanski, D. Kurokawa, S. Brânzei, and A. Procaccia. Simultaneous cake cutting. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*, pages 566–572. AAAI Press, July 2014.
- S. Banach. Sur le problème de la mesure. *Fundamenta Mathematica*, 4:7–33, 1923.
- S. Banach. *Théorie des Opérations Linéaires*. Wydawnictwo Naukowe PWN, Warsaw, Poland, 1932.
- J. Barbanel. Game-theoretic algorithms for fair and strongly fair cake division with entitlements. In *Colloquium Mathematicae*, volume 69, pages 59–73, 1996a.
- J. Barbanel. Super envy-free cake division and independence of measures. *Journal of Mathematical Analysis and Applications*, 197(1):54–60, 1996b.
- 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.
- S. Brams and A. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995a.
- S. Brams and A. Taylor. An envy-free cake division protocol. *The American Mathematical Monthly*, 102(1):9–18, 1995b.
- S. Brams and A. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- 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.
- S. Brams, M. Jones, and C. Klamler. Perfect cake-cutting procedures with money. Technical report, mimeo, 2003.
- S. Brams, M. Jones, and C. Klamler. Better ways to cut a cake. *Notices of the AMS*, 53(11):1314–1321, 2006.
- S. Brams, M. Jones, and C. Klamler. Proportional pie-cutting. *International Journal of Game Theory*, 36(3-4):353–367, 2008.
- 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.
- 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.
- 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.
- 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.
- 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.

- I. Caragiannis, J. Lai, and A. Procaccia. Towards more expressive cake cutting. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, pages 127–132. AAAI Press/IJCAI, 2011.
- K. Cechlárová and E. Pillárová. On the computability of equitable divisions. *Discrete Optimization*, 9(4):249–257, 2012a.
- K. Cechlárová and E. Pillárová. A near equitable 2-person cake cutting algorithm. *Optimization*, 61(11):1321–1330, 2012b.
- K. Cechlárová, J. Doboš, and E. Pillárová. On the existence of equitable cake divisions. *Information Sciences*, 228:239–245, 2013.
- Y. Chen, J. Lai, D. Parkes, and A. Procaccia. Truth, justice, and cake cutting. *Games and Economic Behavior*, 77(1):284–297, 2013.
- Y. Chevaleyre, P. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.
- K. Ciesielski. How good is Lebesgue measure? *The Mathematical Intelligencer*, 11(2):54–58, 1989.
- 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.
- 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.
- L. Dubins and E. Spanier. How to cut a cake fairly. *American Mathematical Monthly*, 68(1):1–17, 1961.
- 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.
- A. Gutman and N. Nisan. Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 719–728. IFAAMAS, June 2012.
- I. Kash, A. Procaccia, and N. Shah. No agent left behind: Dynamic fair division of multiple resources. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 351–358. IFAAMAS, May 2013.
- H. Lebesgue. *Leçons sur l'Intégration*. Gauthier-Villars, Paris, France, 1904.
- 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.
- 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.
- 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.
- 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.
- J. Oxtoby. *Measure and Category*. Graduate Texts in Mathematics. Springer, New York, USA, 2nd edition, 1980.
- A. Procaccia. Thou shalt covet thy neighbor's cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 239–244, 2009.
- 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.
- J. Reijnierse and J. Potters. On finding an envy-free Pareto-optimal division. *Mathematical Programming*, 83(1-3):291–311, 1998.
- J. Robertson and W. Webb. Near exact and envy free cake division. *Ars Combinatoria*, 45:97–108, 1997.
- J. Robertson and W. Webb. *Cake-Cutting Algorithms: Be Fair If You Can*. CRC Press, 1998.

- W. Rudin. *Functional Analysis*. McGraw Hill, 2nd edition, 1991.
- 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.
- R. Schilling. *Measures, Integrals and Martingales*. Cambridge University Press, Cambridge, UK, 2nd edition, 2017.
- R. Schilling and F. Kühn. *Counterexamples in Measure and Integration*. Cambridge University Press, Cambridge, UK, 2021.
- 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.
- 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.
- J. Sgall and G. Woeginger. An approximation scheme for cake division with a linear number of cuts. *Combinatorica*, 27(2):205–211, 2007.
- H. Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- W. Stromquist. Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15(1):R11, 2008.
- W. Stromquist and D. Woodall. Sets on which several measures agree. *Journal of Mathematical Analysis and Applications*, 108(1):241–248, 1985.
- 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.
- J. von Neumann. Über die analytischen Eigenschaften von Gruppen linearer Transformationen und ihrer Darstellungen. *Mathematische Zeitschrift*, 30:3–42, 1929.
- S. Wagon. *The Banach–Tarski Paradox*. Cambridge University Press, Cambridge, UK, 1985.
- T. Walsh. Online cake cutting. In *Proceedings of the 2nd International Conference on Algorithmic Decision Theory*, pages 292–305. Springer, 2011.
- W. Webb. How to cut a cake fairly using a minimal number of cuts. *Discrete Applied Mathematics*, 74(2):183–190, 1997.
- G. Woeginger and J. Sgall. On the complexity of cake cutting. *Discrete Optimization*, 4(2): 213–220, 2007.
- D. Zeng. Approximate envy-free procedures. In *Game Practice: Contributions from Applied Game Theory*, pages 259–271. Springer, 2000.

(P. Kern) MATHEMATISCHES INSTITUT, HEINRICH-HEINE-UNIVERSITÄT DÜSSELDORF, GERMANY.
E-MAIL: kern@hhu.de

(D. Neugebauer) INSTITUTE FOR COMPUTER SCIENCE, HEINRICH-HEINE-UNIVERSITÄT DÜSSELDORF, GERMANY.
E-MAIL: daniel.neugebauer@hhu.de

(J. Rothe) INSTITUTE FOR COMPUTER SCIENCE, HEINRICH-HEINE-UNIVERSITÄT DÜSSELDORF, GERMANY.
E-MAIL: rothe@hhu.de

(R.L. Schilling) INSTITUT FÜR MATHEMATISCHE STOCHASTIK, TU DRESDEN, GERMANY.
E-MAIL: rene.schilling@tu-dresden.de

(D. Stoyan) INSTITUT FÜR STOCHASTIK, TU BERGAKADEMIE FREIBERG, GERMANY.
E-MAIL: stoyan@math.tu-freiberg.de

(R. Weishaupt) INSTITUTE FOR COMPUTER SCIENCE, HEINRICH-HEINE-UNIVERSITÄT DÜSSELDORF, GERMANY.
E-MAIL: robin.weishaupt@hhu.de