# THE BANACH-TARSKI PARADOX 

ULRICH MEYER


#### Abstract

Émile Borel regards the Banach-Tarski Paradox as a reductio ad absurdum of the Axiom of Choice. Peter Forrest instead blames the assumption that physical space has a similar structure as the real numbers. This paper argues that Banach and Tarski's result is not paradoxical and that it merely illustrates a surprising feature of the continuum: dividing a spatial region into disjoint pieces need not preserve volume.


Keywords: Banach-Tarski Paradox, Axiom of Choice, Vitali Sets, Measure

## 1. Introduction

Stefan Banach and Alfred Tarski (1924) prove that one can cut a solid sphere into finitely many pieces, shift some of them to a different location, and then rotate all the pieces, each by a different angle, to obtain two solid spheres of the same radius. ${ }^{1}$ Unlike transformations that stretch or compress, the shifts and rotations used by the Banach-Tarski Paradox (BTP) do not change lengths or angles, and are provably volume-preserving. Banach and Tarski appear to have shown that one can double the volume of a region by means of volume-preserving transformations.

To obtain this result, the sphere needs to be cut into pieces with rather peculiar shapes. At this juncture, Banach and Tarski make essential use of the Axiom of Choice, which claims that any set of non-empty sets has a choice set that contains one element of each of the sets. Choice is a non-constructive principle that postulates the existence of choice sets without identifying a method for finding them. Likewise, BTP establishes the existence of paradoxical divisions of the sphere without providing cutting instructions. But it does not matter whether there is a physical realization of BTP; the problem is that Banach and Tarski appear to have shown that there is a volume-preserving bijection from one spatial region to another spatial region with twice the volume.

If that is right, then we need to give up something to avoid this contradictory result. Émile Borel $(1946,210)$ suggests that BTP is a reductio ad absurdum of the Axiom of Choice. Rejecting Choice would prevent the paradox from arising but the
${ }^{1}$ See also Stromberg (1979), Tomkowicz \& Wagon (2016), and Jech (1973, sec. 1.3). Wilson (2005) shows that the pieces can be moved continuously while remaining disjoint.
theoretical costs of doing so would be prohibitive. It is almost impossible to have an interesting theory in any area of mathematics without appealing to Choice at some point. In their book-length survey of the role of the axiom, Paul Howard and Jean Rubin (1998) spend over four hundred pages listing merely the more important consequences of the Axiom of Choice.

Peter Forrest (2004) accepts the Axiom of Choice, and instead blames what I want to call the $\mathbb{R}$-thesis. This is the widely accepted view that the set $\mathbb{R}^{3}$ of triples of real numbers provides a good model of physical space. We know from the general theory of relativity that space does not look exactly like $\mathbb{R}^{3}$ because the presence of masses causes space to curve. But just as curved lines can be locally approximated by the straight lines that serve as their tangents, the curved spaces of general relativity can be locally approximated by the flat tangent space $\mathbb{R}^{3}$. This is all that BTP needs, and this is what Forrest denies. Without the $\mathbb{R}$-thesis, Forrest argues, we merely get a curious result about sets of triples of real numbers, but no genuine paradox about regions of physical space.

Forrest concludes that space must either be discrete ("gritty") or not be made up of points at all ("gunky"). ${ }^{2}$ Either option would prevent paradoxical decompositions of the sphere, but Forrest's proposal is as radical as Borel's. The dominant view since Isaac Newton has been that lines in space have the same structure as the real numbers, and it was for this very purpose that mathematicians developed the theory of $\mathbb{R}$ in the first place. If BTP is a reductio of the $\mathbb{R}$-thesis then the last three hundred years of physics were based on a colossal mistake: it would be logically impossible for space to look the way everybody thought it did.

I want to advocate a more conservative assessment that retains both the Axiom of Choice and the $\mathbb{R}$-thesis. My solution focuses on what happens before and after the shifts and rotations. Banach and Tarski first cut the sphere into pieces, shift and rotate the pieces individually, and then merge the transformed pieces to form two solid spheres. What is at issue is not a simple combination of shifts and rotations, but a transformation of the form cut-shift \& rotate-merge. Given that its output has twice the volume as its input, this combined transformation does not preserve volume. Since shifts and rotations are provably volume-preserving, it follows that the cutting and merging changes the volume. That is my thesis.

## 2. The Inflated Sphere

The problem with the Banach-Tarski transformation is not just that its output has twice the volume as the input. Many other transformations do the same. Consider a solid sphere of radius one that is located in the origin of our coordinate system, and that consists of all points with position vectors $\langle x, y, z\rangle$ such that $x^{2}+y^{2}+z^{2} \leq 1$. If we stretch each of these position vectors by a factor $\lambda>1$, by mapping $\langle x, y, z\rangle \mapsto \lambda \cdot\langle x, y, z\rangle$, then we get a larger sphere with $\lambda^{3}$ times the volume. Nobody thinks that this result is paradoxical because nobody expects such
${ }^{2}$ Arntzenius (2012, sec. 4.4) makes a similar point but does not cite the earlier article by Forrest.
stretching transformation to preserve volume. On the contrary, the inflated sphere is a proof that stretchings are not volume-preserving.

Suppose we first cut the unit sphere into point-sized regions and then individually shift each point by a different amount, by adding $(\lambda-1) \cdot\langle x, y, z\rangle$ to $\langle x, y, z\rangle$. Once all the separate points are at their new locations, we merge them again to generate a new region. If all cut-shift-merge transformations preserve volume then this new region would have the same volume as the unit sphere, and it does not. Since $\langle x, y, z\rangle+(\lambda-1) \cdot\langle x, y, z\rangle=\lambda \cdot\langle x, y, z\rangle$, the net result of our cut-shift-merge transformation is the same as stretching by a factor $\lambda$.

My thesis is that BTP is another illustration of this phenomenon. Banach and Tarski first cut the unit sphere into pieces, individually shift and rotate the pieces, each by different amounts, and then merge the transposed pieces to form two separate spheres of radius one. The shifts and rotations in the middle might be volume-preserving, but the combined cut-shift \& rotate-merge transformation is like a volume-changing stretching. Once we turn our attention from the shifts and rotations to the cuts and mergers, the case becomes less obviously paradoxical. Nobody is denying that shifts and rotations preserve volume; the question is whether cuts and mergers do the same. I claim that BTP shows that they do not.

Suppose we divide a region into pairwise disjoint and jointly exhaustive subregions. Then it might seem obvious that the volume of the entire region must equal the sum of the volume of the subregions. If volume were additive in this way then any cut—shift \& rotate-merge transformation would have to be volume-preserving. By additivity, the volume of the original region equals the sum of the volume of its subregions. Since shifts and rotations are provably volume-preserving, this sum equals the sum of the volumes of the transposed subregions, which equals the volume of the merged region, again by additivity. The inflated sphere and BTP show that this is false: volume can fail to be additive.

The inflated sphere does not use the Axiom of Choice but requires a division into uncountably many pieces. BTP only needs finitely many pieces but makes essential use of Choice. I want to argue that these differences do not matter, and that failures of additivity are a characteristic feature of the continuum. Additivity fails across the board. For any cardinal $1<\alpha \leq \beth_{1}$, there is a way of dividing a region in the continuum into $\alpha$ disjoint subregions such that the sum of the volumes of the $\alpha$ many subregions differs from the volume of the entire region.

## 3. Measure Theory

Geometry tells us about the volume of spatial regions with regular shapes, such as cubes, spheres, or cylinders. Measure theory aims to assess the volume of arbitrarily shaped regions, by approximating them with regions of known volume.

Henri Lebesgue (1902) uses rectangular boxes, whose volume is given by the product of the lengths of their sides. Given a region $A$, a box cover is a collection of boxes whose union contains all of $A$. The corresponding cover volume is the sum of the volume of all the boxes. As an assessment of the volume of $A$, many cover
volumes are far off the mark. A box cover may extend far outside the region $A$, or it may contain overlapping boxes, whose intersections would be double-counted. But we can always arrange for a tighter fit by using smaller boxes with less overlap. The Lebesgue measure $\mu(A)$ is defined as the infimum (greatest lower bound) of all the cover volumes of $A$. We get a three-dimensional measure by approximating volume with cuboids, a two-dimensional measure by approximating areas with rectangles, and a one-dimensional measure by approximating lengths with line segments.

Infinitely extended regions can lack a Lebesgue measure if all their cover volumes diverge. But all finitely extended regions possess a Lebesgue measure, no matter how oddly shaped they might be. Any finitely extended region can be covered with a single large box and thus has at least one finite cover volume. The set of all its cover volumes is therefore non-empty and is bounded from below by zero. By the completeness of the real numbers, it follows that there is a greatest lower bound of the cover volumes of the region, which is its Lebesgue measure.

If we shift or rotate a box, we get another box with the same volume. Let $C$ be a box cover of $A$. If we shift or rotate $A$ to $A^{\prime}$ then we get a box cover $C^{\prime}$ of $A^{\prime}$ by shifting or rotating all the boxes in $C$ by the same amount. Since $C$ and $C^{\prime}$ have the same cover volume, the infimum of the cover volumes of $A$ and $A^{\prime}$ are the same. Lebesgue measure is thus invariant under shifts and rotations, $\mu(A)=\mu\left(A^{\prime}\right)$. Since any box cover of a region is a box cover of its sub-regions, the Lebesgue measure is also monotone in the sense that $\mu(A) \leq \mu(B)$ whenever $A \subseteq B$.

Some people call $\mu(A)$ the outer Lebesgue measure because it approximates the volume of $A$ from above. We could have taken the opposite approach, and defined an inner measure that approximates the volume of $A$ from below. Instead of covering $A$ with boxes, consider (possibly empty) collections of non-overlapping open boxes that are covered by $A$. If we call the sum of the volumes of these boxes the inclusion volume then we can define the inner measure $\mu_{*}(A)$ as the supremum (least upper bound) of the inclusion volumes. ${ }^{3}$ I say more about the inner measure in Sec. 7. Until then, I will always use 'measure' to refer to the outer measure $\mu$.

## 4. Null Sets

Questions about the additivity of volume are as old as philosophy itself. Zeno of Elea suggested twenty-five hundred years ago that extended regions cannot be made up of extensionless points (Skyrms 1983). If we put one extensionless point next to another, we again get something without extension: zero plus zero equals zero. Finite sums of zeros are also zero, as are (countably) infinite sums of zeros.

The same issue arises here. If we divide a box into pointlike subregions then all these individual subregions have zero measure. Take an infinite sequence of nested boxes that converges on a single point $x$. Each box in the sequence has non-zero measure, but the greatest lower bound of all their volumes is zero, which means
${ }^{3}$ Lebesgue $(1902,7)$ and most modern textbooks adopt a different approach. They define the inner measure in terms of the outer measure. Suppose that $A$ is contained in a region $X$ with finite outer measure. Then the inner measure of $A$ is defined as $\mu_{*}(A)=\mu(X)-\mu(X-A)$.
that $\mu(\{x\})=0$. Yet the volume of the entire box, which is made up of individual points like $x$, is somehow supposed to be non-zero.

These worries ultimately got resolved by modern measure theory. The key insight is that any theory of measure admits null sets, which are non-empty regions with measure zero. We already noted that any individual point is null. It is easy to show that the same holds for finite collection of points. What is a little bit more surprising is that any countably infinite set of points is also null. ${ }^{4}$ This does not contradict the claim that boxes have non-zero measure because we know from Cantor's Theorem that boxes contain uncountably many points. Instead of a paradox, we get a profound insight into the nature of the continuum. To get a region with non-zero measure, we need to combine at least $\beth_{1}$ many points. Sets of $\beth_{0}$ or fewer points are of negligible size compared to extended regions.

Null sets lead to failures of additivity. Let $f$ be a function from the real numbers to the positive real numbers. Then the definite Lebesgue integral $\int_{a}^{b} f(x) \mathrm{d} x$ over the interval $[a, b]$ is defined as the two-dimensional Lebesgue measure of the area between the $x$-axis, the graph of $f$, and vertical lines at $x=a$ and $x=b$. We can think of this integral as the weighted sum of the $\beth_{1}$ many points between $a$ and $b$, with the weight of point $x$ given by the value of $f(x)$. If we assign all points the same weight, by choosing the constant function $f(x)=1$, then we recover the one-dimensional Lebesgue measure of the underlying interval, $\mu([a, b])=\int_{a}^{b} \mathrm{~d} x=b-a$, which is non-zero as long as $b \neq a$.

If the end-points coincide then the interval consists of a point, and its measure is zero because single points are null, $\mu(\{x\})=\mu([x, x])=\int_{x}^{x} \mathrm{~d} y=x-x=0$. If we add up the individual measures of the $\beth_{1}$ points in the interval $[a, b]$, by using integration again, then we still get zero, $\int_{a}^{b} \mu(\{x\}) \mathrm{d} x=\int_{a}^{b} 0 \mathrm{~d} x=0$. This shows that the Lebesgue measure is not uncountably additive:

$$
\mu([a, b]) \neq \int_{a}^{b} \mu(\{x\}) \mathrm{d} x
$$

The measure of an extended interval does not equal the sum (i.e., integral) of the individual measures of the uncountably many points that make up the interval.

This might be an astonishing result but it is a discovery about the continuum, not a paradox. Null sets are not pathologies that ought to be avoided; they are the very thing that allows us to make sense of the view that extended regions are composed of extensionless points. If there are null sets then failures of additivity are unavoidable, and we can easily generate results like the inflated sphere. We cut a region into null sets, rearrange the pieces with shifts and rotations, and then merge the transposed pieces again to get something of a different volume.
${ }^{4}$ Here is the standard proof. Suppose that $x_{1}, x_{2}, x_{3}, \ldots$ is a countably infinite collection of points. Given an arbitrarily small real number $\varepsilon>0$, we cover the $i$-th point $x_{i}$ with a box of volume $\varepsilon / 2^{i}$. This box cover has volume $\varepsilon$ because $1=\frac{1}{2}+\frac{1}{4}+\frac{1}{8}+\ldots$. Since $\varepsilon$ is arbitrarily small, the infimum of the volume of all box covers of the set is therefore zero. Hence any set of fewer than $\beth_{0}$ points is null. The converse is not true. Cantor's ternary set is an example of a null set with cardinality $\beth_{1}$.

## 5. The Vitali Sets

To prove that measure can also fail to be countably and finitely additive, we need the Axiom of Choice. The basic insight is due Giuseppe Vitali (1905), rather than the later paper by Banach and Tarski. Given Choice, which he uses but does not mention, Vitali shows that there is a (non-unique) subset $V$ of the unit interval $[0,1]$ in $\mathbb{R}$ that has the following two properties:
(a) Up to addition of a rational number, $V$ contains all the real numbers. For any $r \in \mathbb{R}$, there is a $v \in V$ and a $q \in \mathbb{Q}$ such that $r=v+q$.
(b) $V$ is a minimal set with property (a) and the difference between any two elements of $V$ is always an irrational number.

The unit interval itself satisfies (a); the hard part is to find a set that satisfies (b). This is where Vitali appeals to Choice. He divides $\mathbb{R}$ into equivalence classes in which the difference between any two numbers is always rational, and he then intersects each of these classes with the unit interval. Any choice set of the resulting set of sets of real numbers then satisfies conditions (a) and (b).

Given a Vitali set $V$, let $q_{1}, q_{2}, \ldots$ be an enumeration of the $\beth_{0}$ many rational numbers between -1 and 1 . Let $V_{i}$ be the result of shifting the elements of $V$ by adding the rational number $q_{i}$. Since shifts preserve measure, all these rationally shifted Vitali sets have the same measure, $\mu(V)=\mu\left(V_{1}\right)=\mu\left(V_{2}\right)=\ldots$ Thanks to (b), the shifted Vitali sets $V_{1}, V_{2}, \ldots$ are pairwise disjoint, and (a) ensures that their union covers the unit interval [0, 1]. By elementary arithmetic, their disjoint union is contained in the larger interval $[-1,2]$. This gives us two inclusions:

$$
[0,1] \subseteq \bigcup_{i=1}^{\infty} V_{i} \subseteq[-1,2]
$$

Since the Lebesgue measure is monotone, these inclusions deliver upper and lower bounds on the measure of the union of the shifted Vitali sets:

$$
1 \leq \mu\left(\bigcup_{i=1}^{\infty} V_{i}\right) \leq 3
$$

Now suppose, for sake of reductio, that the Lebesgue measure is countably additive. The measure of the disjoint union of the shifted Vitali sets would then equal the infinite sum of their individual measures, $\mu\left(\bigcup_{i=1}^{\infty} V_{i}\right)=\sum_{i=1}^{\infty} \mu\left(V_{i}\right)$. Since all the shifted Vitali sets have the same measure, $\mu\left(V_{i}\right)=\mu(V)$, we thus get:

$$
1 \leq \sum_{i=1}^{\infty} \mu(V) \leq 3
$$

It is impossible to satisfy both inequalities. If $\mu(V)=0$ then the infinite sum is zero, violating the left inequality. If $\mu(V)>0$ then the infinite sum diverges, violating the right inequality. Hence the Lebesgue measure is not countably additive.

Vitali interprets his result differently. He assumes that measure is countably additive and concludes that $V$ has no measure. This strikes me as misguided. If we define the measure of a region as the infimum of its cover volumes then the existence of $\mu(V)$ follows from the completeness of the real numbers. Since $V$ is contained in the unit interval, it is finitely extended, and we noted earlier that all finitely extended region have a Lebesgue measure. Moreover, we already know that measure is not uncountably additive, so we cannot just assume that it is countably additive; we should ask for a proof. If we learned anything in the last few centuries, it is that "intuition" is a poor guide to truth about the infinite.

We can use the Vitali sets to define a cut—shift-merge transformation that doubles measure. Let $A$ be the union of $V_{1}, V_{2}, V_{3}, \ldots$ Suppose we cut $A$ into these disjoint pieces, and then shift the odd pieces $V_{1}, V_{3}, V_{5}, \ldots$ to the right, by adding the number 4 to all their elements. Next, we take the even pieces $V_{2}, V_{4}$, $V_{6}, \ldots$, which are still in their initial position, and individually shift each of them to close the gaps left by the odd pieces. By shifting $V_{2 n}$ to the original position of $V_{n}$, we recover a full sequence $V_{1}, V_{2}, V_{3}, \ldots$ on the left, which we then merge to recover $A$. We do the same with the odd pieces on the right and recover another full sequence of shifted Vitali sets that we then merge to generate a second copy of $A$.

## 6. Finite Additivity

It remains to be shown that measure is not finitely additive, either. If we accept the Axiom of Choice then we can prove that measure is at least countably sub-additive. The measure of the union of countably many sets $A_{1}, A_{2}, \ldots$ is never greater than the infinite sum of their individual measures, $\mu\left(\bigcup_{i=1}^{\infty} A_{i}\right) \leq \sum_{i=1}^{\infty} \mu\left(A_{i}\right)$.

Since we generated $V$ by using the non-constructive Axiom of Choice, we cannot say much about the precise value of $\mu(V)$. But we do know that the shifted Vitali sets jointly cover the unit interval $[0,1]$. Hence $1=\mu([0,1]) \leq \mu\left(\bigcup_{i=1}^{\infty} V_{i}\right)$ by monotonicity and thus $1 \leq \sum_{i=1}^{\infty} \mu\left(V_{i}\right)=\sum_{i=1}^{\infty} \mu(V)$ by countable sub-additivity and the fact that the shifted Vitali sets all have the same measure. This tells us something we did not know before: the Vitali set is not null, $\mu(V)>0$.

This also means that we can divide by $\mu(V)$. Let $n$ be a natural number greater than $4 / \mu(V)$. If measure were finitely additive on disjoint sets then the first $n$ shifted Vitali sets would satisfy $\mu\left(V_{1} \cup \ldots \cup V_{n}\right)=\mu\left(V_{1}\right)+\ldots+\mu\left(V_{n}\right)=n \cdot \mu(V)$. Given how we picked $n$, this entails $\mu\left(V_{1} \cup \ldots \cup V_{n}\right)>4$, which cannot be true. The shifted Vitali sets are contained in $[-1,2]$, which has length 3 . Hence $\mu\left(V_{1} \cup \ldots \cup V_{n}\right) \leq 3$ by monotonicity. So $\mu$ is not finitely additive.

Since any division into $n$ pieces can be achieved by $n-1$ successive divisions into two pieces, this means that not even a division into two disjoint pieces is guaranteed to preserve measure. Given the Axiom of Choice, additivity thus fails across the board. It can fail for uncountable, countably infinite, and finite cuts. For any $1<\alpha \leq \beth_{1}$, there is a way of dividing a region into $\alpha$ disjoint subregions so that the sum of the volume of the subregions differs from the volume of the region.

These results easily generalize to higher dimensions. To demonstrate the failure
of countable and finite additivity in $n+1$ dimensions, we take our Vitali construction in one dimension and multiply everything by $n$ copies of the unit interval. For example, to prove that the three-dimensional Lebesgue measure fails to be additive, we would use the unit cube $[0,1] \times[0,1] \times[0,1]$ instead of the unit interval, and the set of "Vitali squares" $\{\{v\} \times[0,1] \times[0,1]: v \in V\}$ instead of $V$.

Viewed from this perspective, Banach and Tarski merely provide a another illustration of how the failure of finite additivity can lead to volume-changing cutshift \& rotate—merge transformation in three dimensions. Their example might be especially compelling, but it does not add to the fundamental insight that, given the Axiom of Choice, volume in the continuum fails to be finitely additive.

## 7. The Inner Measure

All the non-additivity results discussed so far concerned the outer measure $\mu$. An anonymous referee for this journal noted that $\mu$ nicely captures our pre-theoretic habits of approximating volume from the outside. For example, if you buy a bushel of grain, you may complain if the merchant does not fill the bushel to the brim, but you must accept the tiny spaces between the individual grains. This might speak in favor of adopting $\mu$ rather than $\mu_{*}$ as the regimentation of the folk notion of volume, but not much depends on this. While the inner Lebesgue measure does have slightly different properties than the outer measure, we do not get any more additivity of volume by using $\mu_{*}$ instead of $\mu$.

Given Choice, one can prove that the inner measure is countably super-additive on disjoint sets, $\mu_{*}\left(\bigcup_{i=1}^{\infty} A_{i}\right) \geq \sum_{i=1}^{\infty} \mu_{*}\left(A_{i}\right)$. We previously used the countable sub-additivity of the outer measure to show that the Vitali set has non-zero outer measure, $\mu(V)>0$. In a similar way, we can use the countable super-additivity of $\mu_{*}$ to prove that $V$ is null relative to the inner measure, $\mu_{*}(V)=0$.

The inner measure is translation and rotation invariant for the same reasons as the outer measure, which means that all the shifted Vitali sets have the same inner measure, $\mu_{*}(V)=\mu_{*}\left(V_{1}\right)=\mu_{*}\left(V_{2}\right)=\ldots$ We also know that the union of the shifted Vitali sets is contained in the interval $[-1,2]$ with inner measure 3. Suppose we write $[-1,2]$ as the disjoint union of $\bigcup_{i=1}^{\infty} V_{i}$ and the "extra piece" $[-1,2]-\bigcup_{i=1}^{\infty} V_{i}$. By finite super-additivity, and by ignoring the measure of the extra piece, we get $3 \geq \mu_{*}\left(\bigcup_{i=1}^{\infty} V_{i}\right)$. The countable super-additivity of the inner measure then entails $3 \geq \sum_{i=1}^{\infty} \mu_{*}\left(V_{i}\right)=\sum_{i=1}^{\infty} \mu_{*}(V)$, which can only be true if $\mu_{*}(V)=0$.

Like the outer measure, the inner measure is thus not countably additive. All the shifted Vitali sets have inner measure zero, so the infinite sum of their inner measures is also zero, $\sum_{i=1}^{\infty} \mu_{*}\left(V_{i}\right)=0$. But since their union $\bigcup_{i=1}^{\infty} V_{i}$ covers the unit interval, its inner measure is greater than one. Hence $\mu_{*}\left(\bigcup_{i=1}^{\infty} V_{i}\right) \neq \sum_{i=1}^{\infty} \mu_{*}\left(V_{i}\right)$.

The countable sub-additivity of the outer measure entails that we cannot divide a non-null region of space into merely countably many null sets. Relative to the outer measure, the countable union of null sets is therefore always another null set. This is not true for the inner measure. Even though all the shifted Vitali sets are null relative to the inner measure, their countable union is not null.

Since $\mu_{*}(V)=0$, we cannot recycle our Vitali-inspired proof from the previous section to show that $\mu_{*}$ fails to be finitely additive. But we can easily derive this from Banach and Tarski's result. Since the solid spheres and their parts are all contained in a finite region of space, the completeness of the real numbers ensures that they all possess inner Lebesgue measures. It is also easy to show that the two spheres have twice the inner measure of a single sphere. Given that the inner measure is invariant under shifts and rotation, finite additivity would thus entail that Banach and Tarski's cut-shift \& rotate-merge transformation preserves inner measure, and it does not. Hence the inner measure is also not finitely additive.

## 8. Measurable Sets

The question is what to do with all these failures of additivity. Measure theory deals with the issue by ignoring the sets that cause the problems. It gives up on uncountable additivity altogether and then tries to identify a class of well-behaved sets for which measure can be shown to be countably additive. These well-behaved sets are called measurable sets, and they alone form the subject matter of the theory.

There are different ways of characterizing measurability. The Carathéodory Condition calls a set $A$ measurable if and only if $\mu(X)=\mu(X \cap A)+\mu(X-A)$ for any other set $X$, where $X$ need not be measurable. If $A$ satisfies this condition then it divides any other set $X$ into two well-behaved pieces $X \cap A$ and $X-A$. Another option is to call a set measurable just in case its inner and outer measure coincide, which is provably equivalent to the Carathéodory Condition.

Either way, the class of measurable sets is chosen with the sole purpose of maximizing countable additivity. The aim is to prove that whenever we divide a measurable region into at most countably many disjoint measurable parts, the volume of the region equals the sum of the volume of its parts. Since that is all we want from a definition of 'measurable', it might have been more appropriate to call the excluded sets 'non-additive' rather than 'non-measurable'.

The Vitali sets are measurable in the straightforward sense of possessing both inner and outer Lebesgue measures. These measures just happen to be different: their inner measures are zero and their outer measures are not. While this difference might illustrate the peculiarities of the Vitali sets, deciding to ignore them does not change the fact that these sets exist, and that we can define volume-changing cut-shift-merge transformations in terms of them.

## 9. Double-Edged Choice

We already noted the prohibitive cost of rejecting the Axiom of Choice. But it is not just other branches of mathematics that need the axiom; it also plays a crucial role in measure theory itself. On the one hand, we can only prove the existence of the Vitali sets, or establish Banach and Tarski's result, if we assume the Axiom of Choice. This leads to a failure of countable and finite additivity that we would not get otherwise. Robert Solovay (1970) proves that there are models of set theory
without Choice in which every set of real numbers is measurable.
In other respects, though, we get less additivity without Choice, not more. To prove countable additivity on measurable sets, we need to prove countable subadditivity, which requires Choice. Solomon Feferman and Azriel Lévy (1963) show that, without Choice, it is possible that the real numbers are the countable union of countable sets. Countable sets are always null, and null sets are easily shown to be measurable. Since the set of real numbers is not null, measure would thus fail to be countably sub-additive even on measurable sets.

While Choice does give rise to some failures of finite and countable additivity, it is also needed to turn measure into the useful notion that it is. If we are interested in maximizing additivity then we must accept Choice and limit ourselves to countable additivity on measurable sets. But that does not undermine the fundamental insight that measure in the continuum fails to be additive across the board.

There is room for an intermediate position. To prove countable additivity on measurable sets, we only need an Axiom of Countable Choice that asserts the existence of a choice set for any countable set of sets. Solovay shows that Countable Choice is too weak to entail the existence of non-measurable sets. We could thus avoid the Vital sets and BTP, by rejecting unrestricted Choice, and still recover standard measure theory, by retaining Countable Choice.

Yet this neither gives us the full Axiom of Choice, which we still need for many other applications, nor does it restore the full additivity of volume, which still fails for uncountable cuts. Instead of making a principled stand against failures of additivity and Choice, the intermediate position accepts a bit of both. Given that everybody agrees that measure is not uncountably additive, it is not clear why we should be especially concerned about failures of countable and finite additivity, and it is even less clear why we should give up any part of Choice.

One could try to motivate the intermediate position in a different way. Advocates of set-theoretic parsimony might insist that we ought not postulate more sets than are needed to formulate our best scientific theories. If it turns out that we can make do with Countable Choice then that would be reason to abandon the unrestricted axiom. We would no longer get the Vitali sets and BTP, but that would be a sideeffect of set-theoretic parsimony. We would give up Uncountable Choice because it yields too many sets, not because it leads to failures of additivity.

## 10. Conclusion

Instead of a genuine paradox that needs to be resolved by a radical revision of our mathematics or our physics, I propose that we accept the failure of additivity as a feature of the continuum. It is a discovery about the infinite, on a par with the discovery that infinite sets have the same cardinality as some of their proper subsets, which also used to be mis-characterized as being paradoxical.

In one sense, the failure of additivity is pervasive: it fails for finite, countable, and uncountable cuts. In another sense, though, it is highly localized. The success of measure theory shows that failures of additivity can be confined to an isolated
class of pathological regions, and that it does not spread to anything that we care about. We can retain the additivity of volume for divisions into at most countably many measurable parts, but not for arbitrary divisions. Some divisions cut too finely (into null sets), some cut too weirdly (into non-measurable sets).

The existence of such pathological divisions might be a peculiar feature of the continuum, but they are a feature, not a problem. If cuts need not preserve volume any more than stretchings then Banach and Tarski's result is not paradoxical. The shifts and rotations they use in the middle of their transformation are guaranteed to preserve volume, but not the cuts and mergers at the beginning and end.

## Acknowledgements

Earlier versions of this paper were presented at the University of Rochester and the Université de Neuchâtel. I would like to thank my audiences, and especially Paul Audi and Claudio Calosi, for very helpful comments. My work was supported by the Marion Hoeflich Memorial Endowment for Advances in Philosophy.

## References

Arntzenius, F. (2012), Space, Time, \& Stuff, Oxford University Press, New York.
Banach, S. \& Tarski, A. (1924), 'Sur la décomposition des ensembles de points en parties respectivement congruentes', Fundamenta Mathematica 6, 244-77.

Borel, É. (1946), Les paradoxes de l'infini, Gallimard, Paris.
Feferman, S. \& Lévy, A. (1963), 'Independence results in set theory by Cohen's method II', Notices of the American Mathematical Society 10, 593.

Forrest, P. (2004), 'Grit or gunk: Implications of the Banach-Tarski paradox', The Monist 87, 351-370.

Howard, P. \& Rubin, J. (1998), Consequences of the Axiom of Choice, Vol. 59 of Mathematical Surveys and Monographs, American Mathematical Society, Providence, R.I.

Jech, T. (1973), The Axiom of Choice, North-Holland, Amsterdam.
Lebesgue, H. (1902), Intégrale, longueur, aire, Bernandon de C. Rebeschini, Milan.
Skyrms, B. (1983), 'Zeno’s paradox of measure', in Cohen, R. \& Laudan, L., eds., Physics, Philosophy, and Psychoanalysis: Essays in Honor of Adolf Grünbaum, Vol. 76 of Boston Studies in the Philosophy of Science, Reidel, Dordrecht and Boston, 223-254.

Solovay, R. (1970), 'A model of set-theory in which every set of reals is Lebesgue measurable', Annals of Mathematics 92, 1-56.

Stromberg, K. (1979), 'The Banach-Tarski paradox', American Mathematical Monthly 86, 151-61.

Tomkowicz, G. \& Wagon, S. (2016), The Banach-Tarski Paradox, second edn, Cambridge University Press, Cambridge.

Vitali, G. (1905), Sul problema della misura dei gruppi di punti di una retta, Gamberini e Parmeggiani, Bologna.

Wilson, T. (2005), 'A continuous movement version of the Banach-Tarski paradox: A solution to de Groot's problem', Journal of Symbolic Logic 70, 946-52.

Colgate University
Hamilton, NY 13346, USA
umeyer@colgate.edu www.ulrichmeyer.org

