The Banach–Tarski Paradox

Contents

1 Introduction

The Banach–Tarski Paradox, which Jan Mycielski called "the most surprising result of theoretical mathematics" [1, p. xi] is a consequence of two aspects of mathematics from the last two hundred years—Cantorian infinity and the axiom of choice—which are now universally accepted but were once highly divisive among mathematicians. In layman's terms, it states that a three-dimensional ball can be decomposed into finitely many "pieces", and then, without any form of stretching, recomposed into two copies of the original ball.

Cantor's relevance is mainly limited to the notion of countable and uncountable infinite sets—that is, that infinite sets can have more than one cardinality—and though he had many contemporary detractors, his work was generally accepted by the time of Banach–Tarski. The axiom of choice, however, was far more fresh a development, and a major component of a sea change in mathematics. After Russell's paradox among others destroyed naive set theory, Ernst Zermelo (1871-1953)—who had independently discovered the paradox before Russell published it—began formulating a new axiomatic set theory. This, with the suggested extra axioms of Abraham

Fraenkel (1891-1965) and Thoralf Skolem (1887-1963), became Zermelo–Fraenkel set theory, the current standard in mathematics.

Generally, the Zermelo–Fraenkel axioms are fairly simple statements, for example the existences of unions and power sets and that the image of a set under a function is itself a set. The crux of Banach–Tarski, and the controversy, is the following:

Axiom of Choice (AC). *If X is a set of non-empty sets, there exists a function f, known as a* choice function *on X such that, for any $A \in X$, $f(A) \in A$. That is, the image of X under f consists of exactly one member of each of the sets in X.*

This is trivially true for finite X, and can be derived from the other axioms, but becomes problematic when considering infinite X. AC was later proven to be logically independent from the rest of Zermelo–Fraenkel set theory, and therefore its validity rests entirely on being accepted for a specific piece of mathematics. The key issue is that the choice function does not have to be constructed, merely assumed to exist, and this existence is taken as an axiom rather than shown by proof. To a mathematical constructivist, this is absolute anathema; to a Platonist, however, its existence is obvious, and that is all that matters.

An obvious question in response to such a counter-intuitive result stemming from such a controversial axiom is why it doesn't invalidate ZFC (Zermelo–Fraenkel + Choice) in the same way that Russell's paradox invalidated naive set theory. This is not merely a layman's question. In the first contemporary reaction to the Hausdorff Paradox, a crucial building block of Banach–Tarski, Émile Borel (1871-1956), explicitly stated that [2, p. 188]:

> "The contradiction has its origin in the application [...] of Zermelo's Axiom of Choice. [...] The paradox results from the fact that A is not defined, in the logical and precise sense of the word defined. If one scorns precision and logic, one arrives at contradictions."

The distinction between paradoxes in the sense of Banach–Tarski and paradoxes in the sense of Russell's, is that Banach–Tarski is entirely logically consistent with everything that must be assumed to prove it, even if not with our intuitions of what *should* be true. Russell's paradox by contrast is not only inconsistent with any formulation of naive set theory, it is inconsistent with itself. This is what forces a piece of mathematics to be totally discarded. Banach–Tarski certainly forces one to reconsider ZFC, but it can be—and indeed has been—accepted as one of its quirks in a way that Russell's paradox cannot.

Now we have set the stage for this bewildering result, to reach it properly we need to formalise it:

Banach–Tarski Paradox (AC). $\mathbb{R}^3$ *is* G_3*-paradoxical,* G_3 *being the set of isometries in* $\mathbb{R}^3$*. Equivalently, any ball in* $\mathbb{R}^3$ *is equidecomposable with two copies of itself.*

Then we begin by defining several key terms.

2 Groundwork

2.1 Starting Definitions

We begin with a fundamental definition:

Definition 2.1. Let G be a group acting on a set X and $E \subseteq X$ be a non-empty subset. Then E is G*-paradoxical* if for some pairwise disjoint proper subsets $A_1, \ldots, A_m, B_1, \ldots, B_n$ of E and $g_1, \ldots, g_m, h_1, \ldots, h_n \in G$, $E = \bigcup g_i(A_i) = \bigcup h_j(B_j)$.

With this definition, we can now see that Banach–Tarski says that anything in $\mathbb{R}^3$, for example a ball, can be taken apart into subsets of points, then using only isometries some of the subsets can be transformed into the entire ball, with other subsets left over, which can also be transformed into the entire ball. For an example of a paradoxical set, which ends up working similarly to the final paradox on the sphere and ball, we can consider a free non-Abelian group with two generating elements.

2.1.1 Free Non-Abelian Groups

Definition 2.2. A *free group* F with a generating set M is the group of words with letters in M. That is, if $M = \{\sigma, \tau\}$, an example of an element in F would be $\sigma\tau^{-1}\sigma\tau^2\sigma$.

A word generated by a free group can always be expressed in infinitely many ways by adding pairs of elements and their inverses anywhere in the word, for example

$$\sigma\tau = \sigma\tau\sigma\sigma^{-1} = \sigma\tau\sigma^{-1}\sigma = \sigma\sigma\sigma^{-1}\tau$$

so in order to avoid discussing equivalence classes, we only include words in F in their *reduced* forms, with all adjacent pairs of elements and inverses removed. If the generators have finite order, the reduced form also replaces every instance of σ^n with $\sigma^{n \pmod{|\sigma|}}$—later we will do this even for inverses.

The key aspect of this group which makes it free is the fact that the only reduced element equivalent to the identity is the identity itself. This will become relevant and non-trivial later when we make a free group out of rotations in $\mathbb{R}^3$, and it will no longer be obvious that, for example, $\sigma\tau\sigma^{-1}\tau^{-1} \neq e$.

Theorem 2.1. *A free group F with two generators is F-paradoxical (F acts on itself by left multiplication).*

Proof. Let the two generating elements be σ and τ, and consider each element of F only in terms of σ, σ^{-1}, τ and τ^{-1}, for example express $\tau^2\sigma^{-2}$ as $\tau\tau\sigma^{-1}\sigma^{-1}$. For $\lambda \in \{\sigma, \sigma^{-1}, \tau, \tau^{-1}\}$, let $W(\lambda)$ be the set of words with λ as the leftmost element. We immediately notice that $F = e \cup W(\sigma) \cup W(\sigma^{-1}) \cup W(\tau) \cup W(\tau^{-1})$ and $w(\lambda) \cap W(\mu) = \emptyset$ for $\lambda \neq \mu$. We can then see that if $h \in F \setminus W(\sigma)$, $\sigma^{-1}h \in W(\sigma^{-1})$ and so $\sigma\sigma^{-1}h = h \in \sigma W(\sigma^{-1})$. The identity e is included in $\sigma W(\sigma^{-1})$, and therefore $F = W(\sigma) \cup \sigma W(\sigma^{-1}) = W(\tau) \cup \tau W(\tau^{-1})$. □

We now need another key definition to understand the second sentence of the formal statement of Banach–Tarski and why it is equivalent.

2.1.2 Equidecomposability

Definition 2.3. Let G be a group acting on a set X and $A, B \subseteq X$. A and B are *G-equidecomposable*, denoted by $A \sim_G B$, if they can each be partitioned into pairwise disjoint subsets $A_1, \ldots, A_n, B_1, \ldots, B_n$ such that $A = \bigcup A_i$, $B = \bigcup B_i$ and for all $1 \leq i \leq n$, there exists a $g_i \in G$ such that $g_i(A_i) = B_i$.

Having defined equidecomposability, it is now in our interests to establish that shapes are equidecomposable with the same shapes minus some points. This both foreshadows Banach–Tarski and ends up playing a part in its proof.

Theorem 2.2. *Let x be a point on the unit circle S^1. $S^1 \setminus \{x\}$ is SO_2-equidecomposable with S^1 (recall that SO_2 is the group of rotations on $\mathbb{R}^2$).*

Proof. Without loss of generality we can let $x = (1, 0)$; if it is not, simply rotate $S^1 \setminus \{x\}$ to make this so. From now on, consider every point in S^1 both in terms of points in $\{(x, y) \in \mathbb{R}^2 : x^2+y^2 = 1\}$ and in terms of complex numbers $\{z \in \mathbb{C} : |z| = 1\}$, with $z = x+iy$ and $(x, y) = (Re(z), Im(z))$. Let θ be any real number which is not a rational multiple of π, then consider the countably infinite set $D = \{e^{in\theta} \ : n \in \mathbb{N}^+\}$ and its $\mathbb{R}^2$ equivalent. It is clear that within the natural numbers, $e^{in\theta} \neq e^{im\theta}$ for all $n \neq m$, since that would imply that $\theta(n - m) \mid 2\pi$. Therefore if we define the function $\rho(z) = ze^{-i\theta}$ and its equivalent on $\mathbb{R}^2$, $\rho(D) = \{1\} \cup D$, or in terms of $\mathbb{R}^2$, $\rho(D) = \{x\} \cup D$. Finally, we have:

$$S^1 \setminus \{x\} = D \cup S^1 \setminus (\{x\} \cup D) \sim_{SO_2} \rho(D) \cup S^1 \setminus (\{x\} \cup D) = S^1$$

□

This is easily shown to be equivalent to a solid disc with a radial line segment missing. Simply replace each point in S^1, expressed as before as a complex number $e^{i\theta}$, with a line segment $\{\alpha e^{i\theta} : 0 < \alpha \leq 1\}$. Then add a centre point, which is not disturbed for the entire proof, and the rest follows. We can now use this notion of "filling the gap" to show something perhaps more surprising, which ends up being a crucial step for Banach–Tarski.

Theorem 2.3. *Let D be a countably infinite subset of the unit sphere S^2. Then $S^2 \setminus D \sim_{SO_3} S^2$ (recall that SO_3 is the group of rotations on $\mathbb{R}^3$).*

Proof. Let l be a line through the origin which does not intersect any points in D; there are uncountably many such lines. Now when rotating S^2 around the axis l, each point in D will trace a circle. Let A be the set of angles θ such that for some $n \in \mathbb{N}^+$ and $x \in D$, $r_{n\theta}(x) \in D$, where $r_{n\theta}$ is a rotation through $n\theta$ radians about l. A is a countable set, so we can choose a $\theta \notin A$ and define $\rho = r_\theta$. By definition, $\rho^n(D) \cap D = \emptyset$ for all $n \in \mathbb{N}^+$, and therefore $\rho^n(D) \cap \rho^m(D) = \emptyset$ for $n \neq m$. Now, similarly to in Theorem 2.2, we consider the countable set $R = \bigcup_{n \in \mathbb{N}^+} \rho^n(D)$. Then $\rho^{-1}(R) = R \cup D$, and we conclude similarly to Theorem 2.2:

$$S^2 \setminus D = R \cup S^2 \setminus (R \cup D) \sim_{SO_3} \rho^{-1}(D) \cup S^2 \setminus (R \cup D) = S^2$$

□

Here, and with the paradoxical free group, we already have two examples which seem similar to Banach–Tarski, taking proper subsets and translating them into the whole, and neither of which even use the Axiom of Choice.[1]. The distinction here, however, is that all of the sets we have managed to "translate into existence" have been countable. The final leap comes in using this approach on an uncountably infinite set with no overlaps, a leap we are now in the position to begin taking and a leap which will require the Axiom of Choice.

3 The Road to Banach–Tarski

3.1 Finding Rotations

We can now begin making inroads towards the final result via the Hausdorff Paradox. From here on out, we will be considering the unit sphere S^2 centred at the origin, and instead of considering the whole isometry group G_3, we will only consider $SO_3 \subset G_3$. Adjusting the proof for a non-unit sphere is trivial, extending the proof from S^2 to $\mathbb{B}^3$ is similar to how we extended

[1]This should not be surprising—Cantor in fact defined infinite sets precisely as sets which have bijections with proper subsets of themselves. Consider the naturals compared to the even naturals for the simplest possible example.

Theorem 2.2 from a circle to a disc with only a single extra step needed, and the other isometries can be ignored by translating the ball to the origin before beginning.

Our first aim is to find two rotations σ and τ in SO_3 which generate a free group G. The words in G will all represent rotations, and again we only consider reduced words. As mentioned before, to establish that G is a free group, we need to establish the uniqueness property: if a word in σ and τ does not reduce to the identity e, it does not represent a rotation equivalent to the identity. This will then imply that each word in G represents a unique rotation in SO_3 since for $g_1, g_2 \in G$, $g_1 = g_2 \iff g_1 g_2^{-1} = e$.

There are many such pairs of rotations we can choose. When Hausdorff originally published his paradox, he showed that if we take two rotations through π and $\frac{2\pi}{3}$ about two axes with an angle θ between them, and $\cos 2\theta$ is transcendental, then these two rotations satisfy the uniqueness property. However, in 1978, Barbara Osofsky simplified this [3] by showing that the property is also satisfied by the much simpler $\theta = \frac{\pi}{4}$. These rotations are very easy to visualise, so we will let σ be the rotation through π about the z axis and τ be the rotation through $\frac{2\pi}{3}$ about the line $z = x$.

Now that our free group has generators of finite order, we can think differently about how we reduce words. From now on, we will reduce all words in G to words made up of $\{\sigma, \tau, \tau^2\}$ only. Since $\sigma^2 = \tau^3 = e$, we can replace all inverses in any word with equivalent non-inverses: $\sigma^{-1} = \sigma$, $\tau^{-1} = \tau^2$ and $\tau^{-2} = \tau$.

Theorem 3.1 (Uniqueness Property of σ and τ)**.** *No nontrivial word in $\{\sigma, \tau, \tau^2\}$ corresponds to the identity rotation. (Adapted from [3] and [4])*

Proof. We begin by noting that every reduced word in G other than e, σ, τ and τ^2 can be expressed in exactly one of the four forms:

$$\begin{aligned} \alpha &= \tau^{p_1}\sigma\tau^{p_2}\sigma \dots \tau^{p_n}\sigma \\ \beta &= \sigma\tau^{p_1}\sigma\tau^{p_2} \dots \sigma\tau^{p_n} \\ \gamma &= \tau^{p_1}\sigma\tau^{p_2}\sigma \dots \sigma\tau^{p_n} \\ \delta &= \sigma\tau^{p_1}\sigma\tau^{p_2} \dots \sigma\tau^{p_n}\sigma \end{aligned} \tag{1}$$

where $n \geq 1$ (except for γ where $n > 1$) and p_i is 1 or 2.

First we show that $\alpha \neq e$. By considering the matrix representations of σ and τ

$$\sigma = \begin{bmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{bmatrix}, \tau = \begin{bmatrix} -\frac{1}{2} & -\frac{\sqrt{3}}{2} & 0 \\ \frac{\sqrt{3}}{2} & -\frac{1}{2} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

we find that, for $p \in \{1, 2\}$

$$\tau^p\sigma = \frac{1}{2}\begin{bmatrix} 0 & (-1)^{p-1}\sqrt{3} & -1 \\ 0 & 1 & (-1)^{p-1}\sqrt{3} \\ 2 & 0 & 0 \end{bmatrix}$$

It follows by induction that

$$\alpha = \frac{1}{2^n} \begin{bmatrix} m_{11} & m_{12}\sqrt{3} & m_{13} \\ m_{21}\sqrt{3} & m_{22} & m_{23}\sqrt{3} \\ m_{31} & m_{32}\sqrt{3} & m_{33} \end{bmatrix}$$

where m_{11}, m_{21}, m_{31}, m_{32} and m_{33} are even integers and m_{12}, m_{13}, m_{22} and m_{23} are odd integers. $m_{12} \neq 0$, therefore $\alpha \neq I_3$, and returning to viewing α as a word in a free group, $\alpha \neq e$.

Now consider β as in (1). For any given word of the form β, $\sigma\beta\sigma$ is a corresponding word of the form α. Therefore if there exists a $\beta = e$, then its corresponding $\alpha = \sigma\beta\sigma = \sigma^2 = e$, which we have shown cannot be the case, so $\beta \neq e$.

Next, assume that γ, again as in (1), is the word in its form that is equal to e (recall that $n > 1$ for γ). First we consider the case $p_1 = p_n$. For any γ with $p_1 = p_n$, $\tau\gamma\tau^2$ is a corresponding word of the form α (if $p_1 = p_n = 1$) or β (if $p_1 = p_n = 2$). Therefore if there exists a $\gamma = e$, then its corresponding α or $\beta = \tau\gamma\tau^2 = \tau^3 = e$, neither of which are possible.

If instead $p_1 \neq p_n$, assume not only that $\gamma = e$ but also that it has minimal n. We immediately notice that if $p_1 \neq p_n$, $\tau^{p_1+p_n} = e$. Therefore if $n > 3$, for $\gamma = \tau^{p_1}\sigma\tau^{p_2}\sigma \ldots \sigma\tau^{p_{n-1}}\sigma\tau^{p_n}$

$$\sigma\tau^{p_n}\gamma\tau^{p_1}\sigma = \sigma\tau^{p_n}\tau^{p_1}\sigma\tau^{p_2}\sigma \ldots \sigma\tau^{p_{n-1}}\sigma\tau^{p_n}\tau^{p_1}\sigma = \tau^{p_2}\sigma \ldots \sigma\tau^{p_{n-1}}$$

is another word of the form γ with a lesser n, contradicting the assumption that γ has minimal n. If $n = 3$, then for $\gamma = e = \tau^{p_1}\sigma\tau^{p_2}\sigma\tau^{p_3}$

$$e = \sigma\tau^3\sigma = \sigma\tau^{p_3}\gamma\tau^{p_1}\sigma = \sigma\tau^{p_3}\tau^{p_1}\sigma\tau^{p_2}\sigma\tau^{p_3}\tau^{p_1}\sigma = \tau^{p_2}$$

and if $n = 2$, then for $\gamma = e = \tau^{p_1}\sigma\tau^{p_2}$

$$e = \tau^{p_2}\tau^{p_1} = \tau^{p_2}\gamma\tau^{p_1} = \tau^{p_2}\tau^{p_1}\sigma\tau^{p_2}\tau^{p_1} = \sigma$$

neither of which make sense, so $\gamma \neq e$.

Finally, δ is similar to β. Once again consider δ as in (1). For any given word of the form δ, $\sigma\delta\sigma$ is a corresponding word of the form γ. Therefore if there exists a $\delta = e$, its corresponding $\gamma = \sigma\delta\sigma = \sigma^2 = e$ which we have shown cannot be the case, so $\delta \neq e$. We have proven all the cases, so we are done. □

3.2 The Hausdorff Paradox

With Theorem 3.1 we have established that G, generated with $\{\sigma, \tau, \tau^2\}$, represents a countably infinite set of unique rotations. Our next step to get to Banach–Tarski is to break G down into three partitions, which we will call G_1, G_2 and G_3. For these subsets to be partitions, we must have

$G_n \cap G_m = \emptyset$ for $n \neq m$ and $\bigcup G_n = G$. With a countably infinite set, the two ways that it makes sense to decide which subset an element should be assigned to are as follows: to take a property which is invariant on all but one part of an element, or to work recursively. We will do the second.

We begin by assigning the identity e to G_1, σ and τ to G_2 and τ^2 to G_3. Each element $g \in G$ can be expressed as $\alpha_1\alpha_2 \ldots \alpha_n$, with $\alpha_i \in \{\sigma, \tau, \tau^2\}$; we will call n the *length* of the element and denote it by $l(g)$. We have already assigned all elements of length 0 (the identity) and 1. Then for each g with $l(g) = 1$ which has already been assigned, we will assign σg, τg and/or $\tau^2 g$ to a subset, depending on the leftmost letter in g. Then we repeat this for every element of length 2, then length 3, etc. The rules are as follows:

1. If $g \in G_1$:
   (a) If its leftmost letter is σ, then we assign τg to G_2 and $\tau^2 g$ to G_3
   (b) If its leftmost letter is τ or τ^2, then we assign σg to G_2
2. If $g \in G_2$:
   (a) If its leftmost letter is σ, then we assign τg to G_3 and $\tau^2 g$ to G_1
   (b) If its leftmost letter is τ or τ^2, then we assign σg to G_1
3. If $g \in G_3$:
   (a) If its leftmost letter is σ, then we assign τg to G_1 and $\tau^2 g$ to G_2
   (b) If its leftmost letter is τ or τ^2, then we assign σg to G_1

Because nothing of length n is assigned until everything of length $n-1$ is assigned, this will assign the entire countably infinite G to exactly one of the subsets. We now have

$$\begin{aligned} G_1 &= \{e, \sigma\tau, \sigma\tau^2, \tau^2\sigma, \sigma\tau\sigma, \ldots\} \\ G_2 &= \{\sigma, \tau, \tau\sigma\tau, \sigma\tau^2\sigma, \tau\sigma\tau^2, \ldots\} \\ G_3 &= \{\tau^2, \tau\sigma, \tau^2\sigma\tau, \tau^2\sigma\tau^2, \ldots\} \end{aligned} \tag{2}$$

with relations which will lead to the Hausdorff Paradox.

Lemma 3.2. *For the G_1, G_2, G_3 we have defined, some of the elements of which can be seen in (2), we have the following relations: $\tau G_1 = G_2$, $\tau^2 G_1 = G_3$ and $\sigma G_1 = G_2 \cup G_3$. [4]*

Proof. We want these relations to be bijections, so first note that we can rewrite them as $\tau^2 G_2 = G_1$, $\tau G_3 = G_1$ and $\sigma(G_2 \cup G_3) = G_1$.

Now we begin an inductive proof by verifying this relation for the elements of length 0 and 1. There is only one of length 0: the identity, which is in G_1. We easily verify that $\tau e \in G_2$, $\tau^2 e \in G_3$ and $\sigma e \in G_2 \subset G_2 \cup G_3$.

Now we verify the relations for σ, $\tau \in G_2$ and $\tau^2 \in G_3$. We can begin by reversing the previous relations to show that $\sigma\sigma = \tau^2\tau = \tau\tau^2 = e \in G_1$. Then we can individually verify that $\tau^2\sigma$, $\sigma\tau$ and $\sigma\tau^2$ are all in G_1.

Assume that for $n > 1$, the relations hold for all elements g with $l(g) < n$. Choose a new arbitrary g of length n.

Case 1 (g has leftmost letter σ). (1a) in the assignment process implies that $g \in G_1 \iff \tau g \in G_2, \tau^2 g \in G_3$, (2a) implies that $g \in G_2 \iff \tau^2 g \in G_1$ and (3a) implies that $g \in G_3 \iff \tau g \in G_1$. $l(\sigma g) = n-1$ so our assumption implies:

$$g \notin G_1 \iff \sigma(\sigma g) = g \in G_2 \cup G_3 \iff \sigma g \in G_1 \iff \sigma g \notin G_2 \cup G_3$$

Case 2 (g has leftmost letter τ). (1b), (2b) and (3b) all imply that $g \in G_1 \iff \sigma g \in G_2 \cup G_3$. $\tau g = \tau^2 h$, where $l(h) = n-1$ and h has leftmost letter σ, so our assumption and (3a) imply:

$$\tau g = \tau^2 h \in G_2 \iff \tau^2 g = h \in G_3 \iff g = \tau h \in G_1$$

Case 3 (g has leftmost letter τ^2). Again, (1b), (2b) and (3b) all imply that $g \in G_1 \iff \sigma g \in G_2 \cup G_3$. $\tau g = h$, where $l(h) = n-1$ and h has leftmost letter σ. Similarly, our assumption and (2a) imply:

$$\tau g = h \in G_2 \iff g = \tau^2 h \in G_1 \iff \tau^2 g = \tau h \in G_3$$

Thus the relations hold in all cases for an arbitrary element of length n, and by induction hold for the entire sets. □

These relations immediately bring the paradoxical sets from earlier to mind. Their remarkable nature should already be apparent, but now we can finally apply it to the unit sphere and bring it right into conflict with our geometric intuition. Here is also where we finally use the Axiom of Choice.

Theorem 3.3 (Hausdorff Paradox)**.** *The sphere S^2 can be partitioned into four sets, which we will call D, K_1, K_2 and K_3, such that D is countably infinite, and $K_1 \cong K_2 \cong K_3 \cong K_2 \cup K_3$, with $\cong$ signifying congruence.*[2]

Proof. Each rotation in G has two poles which are not moved when it is applied to the sphere. We remove these from consideration in order to keep our points well defined—otherwise any point could be moved to a pole, be rotated but not moved, then moved to the same point it would have gone to without that middle rotation, leading to two different rotations taking one point to the same endpoint. We let $D = \{x \in S^2 : \exists \rho \neq e \in G \; s.t. \; \rho(x) = x\}$ be the set of poles, and since it has at most two points for each rotation in

[2]As an aside beyond the scope of this essay, this implies that on S^2 there is no finitely additive measure defined on all subsets preserved by congruence, since that would imply that the measure of a given K_n is simultaneously $\frac{1}{3}$, $\frac{1}{2}$ and $\frac{2}{3}$.

the countably infinite G, it is also countable. We now only consider $S^2 \setminus D$, which is clearly uncountable.

Any given $x \in S^2 \setminus D$ is taken to a unique point for each $g \in G$, which only results in a countably infinite number of points. To solve this, we consider G's orbits; that is, sets $G \cdot x = \{gx : g \in G\}$ for every $x \in S^2 \setminus D$. Any two orbits are either identical or disjoint, since if $t = a(x) = b(y) \in G \cdot x \cap G \cdot y$ and $z = c(x) \in G \cdot x$, then $z = c(x) = ca^{-1}(t) = ca^{-1}b(y) \in G \cdot y$. We can therefore partition $S^2 \setminus D$ into orbits of G.

Now we use the Axiom of Choice to create a choice set $C \subset S^2 \setminus D$ containing exactly one point from each orbit. The important properties of C are that no point in C can be rotated to a different point in C using rotations in G, but *any* point in $S^2 \setminus D$ can be reached by taking *one* point in C and rotating it by some rotation in G.

We define $K_n = \{g(c) : g \in G_n, c \in C\}$ for $n = 1, 2, 3$. Since G_1, G_2 and G_3 partition G, K_1, K_2 and K_3 partition $S^2 \setminus D$. It is clear that the G_n relations in Lemma 3.2 have equivalents with K_n: applying the rotation τ to K_1 yields K_2 and so $K_1 \cong K_2$, etc. □

The difficult part of Banach–Tarski is well behind us now; we are almost finished. It may now be clearer why the Axiom of Choice was controversial: we simply defined the choice set into existence, seemingly without any rigour whatsoever. This is not hand-waving, this is an application of the axiom as intended.

3.3 Final Steps

Now we have successfully partitioned the uncountably infinite S^2 into a countably infinite subset D, plus three uncountably infinite subsets K_1, K_2 and K_3. We can use the relations we have found to rotate each K_n into $K_2 \cup K_3$; the specific rotations are σK_1, $\sigma\tau^2 K_2$ and $\sigma\tau K_3$. We do this one at a time, starting with K_1, then splitting our new $K_2 \cup K_3$ into K_2 and K_3 by taking $\sigma K_1 \cap K_2 = K_2$ and $\sigma K_1 \cap K_3 = K_3$, and repeating this for the original K_2 and K_3. Rotate one of our K_2s and one of our K_3s into $K_1 = \tau^2 K_2 = \tau K_3$. Now we have decomposed the entire sphere into four subsets, then translated three of those subsets into two copies of each of them, giving us the subsets necessary to reconstruct the original sphere, plus an extra sphere missing a countably infinite number of points.

However, we already know from Theorem 2.3 that a sphere missing countably many points is equidecomposable with the entire sphere. We apply this to "translate the poles into existence" for the second sphere, and we have completed the proof for a sphere. Similarly to how we used a proof on a circle and constructed an equivalent for a disc, we can augment our proof on the sphere to one on a ball.

We begin by converting each point into a radial line. Our entire proof on the sphere works exactly the same on the closed ball minus the centre if we substitute $x \to \{tx : 0 < t \leq 1\}$ (the second inequality is strict if we want to prove the theorem for an open ball). If we think of taking a ball, then removing the centre, we can "clone" what remains. The final step is to put the centre we removed back, then translate a new centre into existence.

Now that we are considering a ball missing only $\mathbf{0}$, there are infinitely many circles inside the ball with only $\mathbf{0}$ missing. We apply Theorem 2.2 to rotate one of these circles missing a point into its respective entire circle. We have now filled in the centre, and we are done.

4 Implications

The first mathematical implication of the Banach–Tarski Paradox is that previous notions of volume-preserving measurability have to be treated far more carefully if not totally discarded. To focus on this, however, would be myopic.

Banach–Tarski shows that we must choose between the seemingly innocuous Axiom of Choice, which virtually all mathematicians accept outside of specific contexts in set theory, and the very idea that three dimensional space in mathematics perfectly corresponds to reality. How can it when we need to jump through hoops to define something so simple as measurement without leading to what seems like total nonsense? But what then do we make of the symbiotic relationship between mathematics and physics, which uses experiments as well as theory? Differential calculus is built on understanding movement, and consistently had physics as one of its key motivations, from simple distance-time gradients to curl and divergence. How do we reconcile this incongruity between theoretical work and the experimentally verifiable reality which motivates it?

Of course, the orbits that we consider are so dense and precise that no human could construct them from an actual physical ball, but that misses the forest for the trees, and moreover feels like a way to avoid the questions that Banach–Tarski forces to the table. Do we accept that a hypothetical infinite-fidelity machine could indeed duplicate items? If we do, then it is worth mentioning that the version of Banach–Tarski proven here is the weak form. Before we state and prove the strong form, we will state (but not prove) another theorem:

Theorem 4.1 (Banach–Schröder–Bernstein Theorem)**.** *Suppose G acts on X and A, $B \subseteq X$. If A is equidecomposable with a subset of B, ($A \preceq B$) and vice versa, then $A \sim_G B$.*

The strong form of Banach–Tarski is as follows:

Banach–Tarski Paradox (AC) (Strong Form)**.** *For any two bounded A, $B \subset \mathbb{R}^3$ with non-empty interior, $A \sim B$. [1, p. 31]*

Proof. We will prove that the strong form follows from the weak form and Banach–Schröder–Bernstein. Choose two solid balls K, L such that $A \subseteq K$, $B \subseteq L$. Let $n \in \mathbb{N}$ be large enough that K is covered by n overlapping copies of L. Now for the set S of n disjoint copies of L, by Banach–Tarski and translating the obtained copies we find that $L \succeq S$. Therefore $A \subseteq K \preceq S \preceq L \subseteq B$, so $A \preceq B$. The same argument works to show that $B \preceq A$, and so $A \sim_{G_3} B$. □

This is possibly even more absurd to imagine in the physical world. Something must be the fundamental issue. Generally critics point to the Axiom of Choice, and indeed the Axiom of Choice is logically equivalent to Banach–Tarski. It is also necessary for Vitali sets, which like Banach–Tarski cause a particular notion of measure to fail. It is not, however, necessary for the Sierpiński–Mazurkiewicz Paradox—which shows that $\mathbb{R}^2$ has a paradoxical subset using the fact that e is transcendental—or the Mycielski–Wagon Paradox [1, §4.3], an equivalent to the Hausdorff Paradox in the hyperbolic $\mathbb{H}^2$. Banach–Tarski happens to act on the domain which we expect to represent our experience of the world, but it is not the only counter-intuitive result in mathematics by any stretch.

The Axiom of Choice is not the only notion implicated in Banach–Tarski. As has been alluded to, the Banach–Tarski and Hausdorff Paradoxes both wreak havoc on any attempt to define measures in ways which preserve the properties we expect them to have (e.g. area) under the operations we would expect to preserve them (isometries). If the subsets we take have no sensible notion of volume, does it make sense to say that the new ball has "the same volume" as the original? If not, why does this new ball with every point in the original ball not make sense when attempting to measure it? If there were answers to all of these questions, then this would not be called a paradox. It may not have an explicit logical contradiction baked into it like Russell's paradox, and one could argue based on a strict personal definition that "Banach–Tarski Paradox" is a misnomer, but it still forces us to go over our foundational assumptions in much the same way.

In the end, disputes over Banach–Tarski may be as simple as the physicist seeing a contradiction in the answer, but the mathematician seeing a contradiction in the question.

References

[1] Grzegorz Tomkowicz and Stan Wagon. *The Banach-Tarski Paradox*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2nd edition, 2016.

[2] Gregory H. Moore. *Zermelo's Axiom of Choice: Its Origins, Development, and Influence*. Studies in the History of Mathematics and Physical Sciences 8. Springer-Verlag New York, 1st edition, 1982. Translation of excerpt from Émile Borel's *Leçons sur la theorie des fonctions* (1914).

[3] Barbara Osofsky and Scot Adams. Problem 6102 and solution. *The American Mathematical Monthly*, 85:504–505, 1978.

[4] Karl Stromberg. The Banach-Tarski Paradox. *The American Mathematical Monthly*, 86:151–161, 1979.