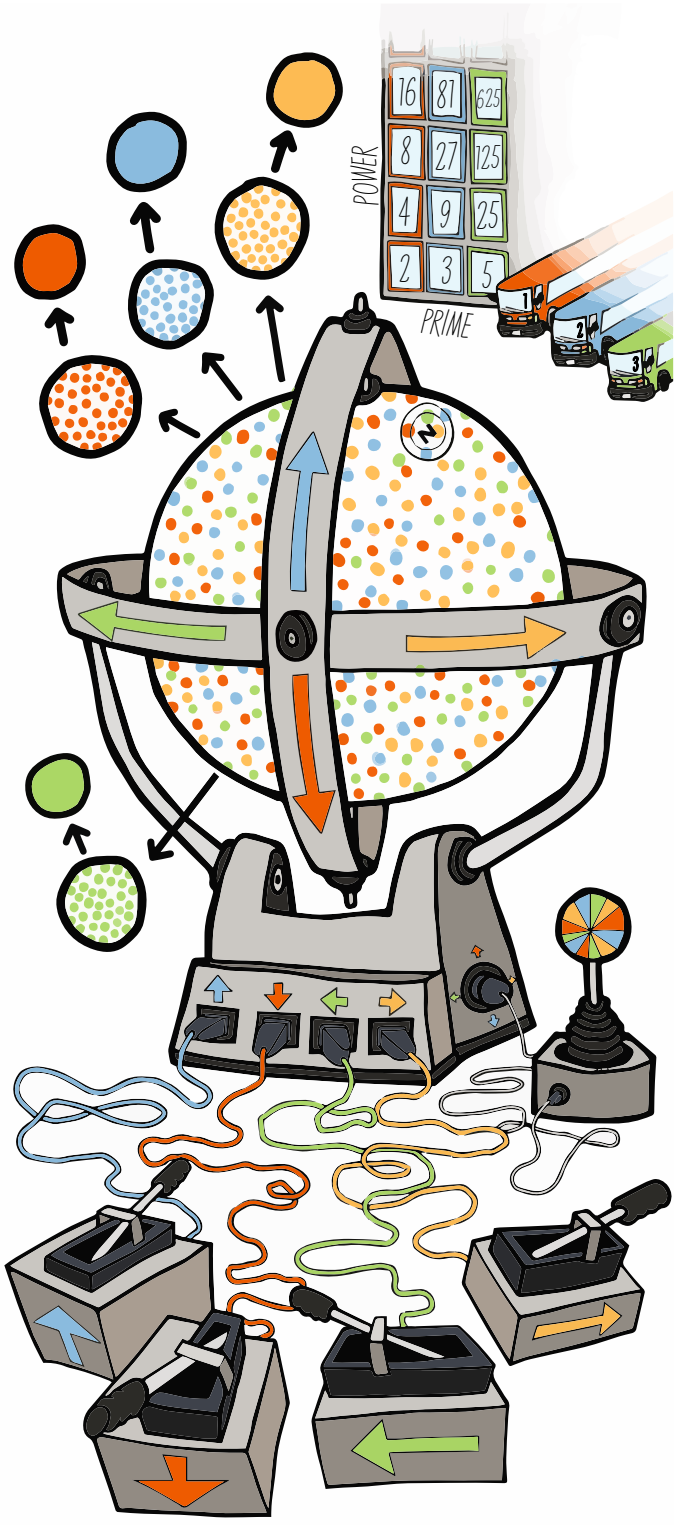
An Axiom of Your Choice

Hilbert's Hotel with Infinite Buses: An infinite number of buses, each carrying infinitely many passengers, arrives at a hotel with countably infinite rooms. To accommodate everyone, each passenger n on bus m is assigned to room $(p_n)^m$, where p_n is the nth prime number. No two passengers ever collide, since distinct prime powers are never equal: for $p \neq q$, there are no integers k, ℓ such that $p^k = q^\ell$.

Banach-Tarski as a Card Game: Imagine a deck with cards A, B, A^{-1}, B^{-1}, where you can make any pile of cards, under one rule: A and A^{-1} cancel, B and B^{-1} cancel (the cards disappear if they are next to one another). Each card represents a rigid rotation: A, A^{-1} rotate around one axis; B, B^{-1} around another. In the diagram, blue and red are inverses, as are green and yellow.

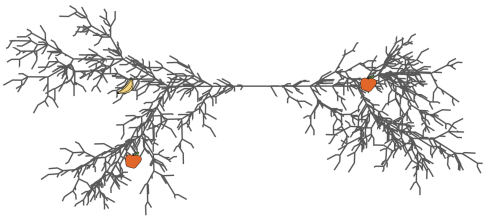
Four players each take all reduced sequences (no canceling cards neighbor each other) beginning with a specific card, forming four disjoint point sets: $S_A, S_{A^{-1}}, S_B, S_{B^{-1}}$.

Now prepend an A to every sequence in $S_{A^{-1}}$, practically rotating the whole set, forming $A \cdot S_{A^{-1}}$. This transforms sequences that began with A^{-1} into sequences that no longer do. For instance, $A^{-1}BA \mapsto BA$, and $A^{-1} \mapsto$ identity. The union $S_A \cup A \cdot S_{A^{-1}}$ now covers all sequences — forming one full copy of the sphere. Similarly, $S_B \cup B \cdot S_{B^{-1}}$ forms a second.



An Axiom of Your Choice

The Banach–Tarski paradox shows that a solid sphere can be partitioned into finitely many disjoint pieces and, using only rigid motions, reassembled into two spheres identical to the original. This construction depends on the Axiom of Choice and the existence of non-measurable sets, whose behavior diverges from intuitions about volume. While not physically realizable, the result illustrates how certain set-theoretic assumptions allow decompositions that defy standard notions of size and conservation.



Axiom of Choice ∘ Non-Measurable Sets ∘ Sphere Paradox ∘ Free Group F_2 ∘ Orbit Decomposition ∘ Rigid Motion Invariance ∘ Countable Additivity ∘ von Neumann Naturals ∘ Hilbert's Hotel ∘ Volume Non-Conservation ∘ Mathematical Constructivism

„Das ist nicht Mathematik, das ist Theologie!“
(“This is not mathematics, this is theology!”)
— Paul Gordan (1888), protesting Hilbert's use of non-constructive existence proofs

„... auch die Theologie hat ihre Verdienste.“
(“... even theology has its merits.”)
— David Hilbert's reply, defending non-constructive reasoning

An Axiom of Your Choice

In the second half of the nineteenth century, questions about the foundations of analysis led mathematicians to examine the basic assumptions underlying number, function, and space. Dedekind formalized the real numbers via cuts, while Weierstrass removed appeals to geometric intuition from calculus. At the same time, mathematicians encountered pathologies — such as functions continuous everywhere but differentiable nowhere, or nowhere-dense sets of positive measure — that challenged classical notions of size and shape. These developments revealed that intuitive notions of length, area, and convergence required formal clarification, especially in the context of infinite processes.

In the late nineteenth century, Georg Cantor changed mathematics with his work on infinite sets, laying the groundwork for new perspectives on measure and cardinalities. In 1905, Giuseppe Vitali introduced the first example of a non-measurable set, suggesting that some subsets of $\mathbb{R}^n$ cannot be assigned an intuitive notion of size. Building on this foundation, Felix Hausdorff presented a paradox in 1914, showing that a sphere could be decomposed in a way that hinted at even more surprising outcomes.

A decade later, in 1924, Stefan Banach and Alfred Tarski formulated it further in a concrete result, known as the Banach–Tarski paradox. They demonstrated that a solid ball in three-dimensional space could be split into a finite number of pieces and then reassembled, through rigid motions, into two full copies of the original in a process that requires the Axiom of Choice, introduced two decades prior (Zermelo, 1904) and still a subject of research. Though it does not apply to physical objects, the Banach–Tarski paradox remains a powerful example of how set-theoretic assumptions can lead to unexpected and extraordinary results in geometry.

A mathematical system begins with a specification of its elements: which objects exist, which operations are defined on them, and which relations must hold. These specifications are encoded in the system's axioms. An axiom is a formal assumption that serves as the foundation for the system. Within that system, no statement can be derived unless it is implied by the axioms in conjunction with the rules of logical inference.

Once a set of axioms is fixed, all further reasoning, including definitions, proofs, and theorems, must proceed within the structure they determine. The consistency and character of the system depend entirely on these initial choices.

Different axiomatic systems describe different mathematical worlds. In one system, every set may have a well-ordering (every non-empty subset has a least element with respect to the order). In another, it may not. In one geometry, parallel lines exist; in another, they do not.

Even familiar objects depend on axiomatic choices. One way to see this is through the von Neumann construction of the natural numbers (von Neumann, 1923) inside set theory. Think of sets as "bags" that can hold distinct objects: duplicates vanish, so $\{\text{dog}, \text{dog}, \text{cat}\} = \{\text{dog}, \text{cat}\}$. The union of two bags (denoted by $A \cup B$) merges their contents, ignoring duplicates. With this picture, the naturals are nested bags, because all you have in this

©2026 David H. Silver, CC BY-NC 4.0

universe are empty bags. Mark an empty bag as $\{\}$.

$$0 := \{\}, \quad 1 := \{0\} = \{\{\}\}, \quad 2 := \{0, 1\} = \{\{\}, \{\{\}\}\}, \quad 3 := \{0, 1, 2\}, \ \ldots .$$

Notice the pattern. Zero is an empty bag. One is a bag containing the empty bag. Two is a bag containing both zero and one. Three contains zero, one, and two. Each number n is the bag containing all smaller numbers. This means $m \in n$ (the set m is an element of the set n) exactly when $m < n$. The natural numbers are the smallest set that satisfies the axioms.

Addition is defined recursively: $m + 0 = m$ and $m + S(n) = S(m + n)$, where $S(n) = n \cup \{n\}$ is the successor. Since n is the set of all smaller numbers, $S(n)$ contains all those smaller numbers plus the set n itself as a new element. This makes even the simplest arithmetic fact a theorem rather than a definition. For instance, one proves that $1 + 1 = 2$. Since $1 + 1 = 1 + S(0) = S(1 + 0) = S(1) = 2$, the familiar statement is established within the axioms.

This illustrates how set theory provides foundations for all mathematics. Relations such as $<$ become sets of ordered pairs, and axioms govern the construction of increasingly abstract objects. Some axioms describe intuitive operations such as forming power sets, while others assert the existence of entities that cannot be explicitly constructed, such as inaccessible cardinals or non-measurable sets.

The Axiom of Choice is one such axiom (Zermelo, 1904). It asserts that for any collection of non-empty sets, there exists a function that selects exactly one element from each set. In finite cases, such selections can be written down explicitly or proved to exist using elementary methods. In infinite settings, this is not always possible. For example, consider an infinite collection of drawers, each containing a left and a right shoe. A rule such as "choose the right shoe" provides a well-defined selection and does not require the Axiom of Choice. But if each drawer contains a pair of identical socks with no distinguishing features, then no explicit rule can be formulated. The existence of a function that selects one sock from each drawer in this case depends on accepting the Axiom of Choice.

Let us now move down the ladder of abstraction, from axioms to the notion of size. Understanding how we measure things — length, area, volume — will be important for seeing why the Banach-Tarski paradox is so surprising.

How do we assign "size" to things? In everyday life, it's intuitive — a room's area is its length times width, and if you combine two non-overlapping rooms, their total area is simply the sum of their individual areas.

Measure theory formalizes this intuition. A *measure* μ is a function that assigns a non-negative real number to certain subsets of a space. The central requirement is **additivity**: if two disjoint measurable sets A and B are combined, then their measure is the sum of the measures of the two sets: $\mu(A \cup B) = \mu(A) + \mu(B)$ (provided that $A \cap B = \emptyset$, i.e., they share no elements). This principle extends to infinite collections: if a set is decomposed into countably infinite disjoint measurable subsets $\{A_i\}_{i=1}^{\infty}$, then $\mu\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mu(A_i)$. This is called **countable additivity**.

Not every subset can be assigned a measure. Some sets resist consistent size assignment — they are *non-measurable*. The existence of non-measurable sets depends on accepting axioms such as the Axiom of Choice.

In familiar settings, the standard measure corresponds to area in the 2D space $\mathbb{R}^2$ or volume in the 3D space $\mathbb{R}^3$. But even in these cases, not all subsets are measurable.

When measuring spatial objects, **invariance under rigid motions** is required. Translating or rotating a measurable set leaves its measure unchanged. Rigid motions preserve distances and angles; they don't stretch, tear, or compress. This encodes the expectation that volume is an intrinsic property — not dependent on where the object is or how it's oriented.

Non-measurable sets lead to paradoxical results, including the Banach–Tarski paradox (Banach & Tarski, 1924). The paradox states that a 3-dimensional ball can be partitioned into five disjoint subsets, which can then be recombined — using only rigid motions — into two balls congruent to the original. This result demonstrates the failure of volume to be preserved when applied to non-measurable sets. In physical systems such as a stone or a fluid body, each component contributes additively to the whole. The Banach–Tarski construction defines a setting where this principle fails and volume is not additive.

Consider Hilbert's Hotel — an infinite hotel (Hilbert, 1926) with rooms numbered $1, 2, 3, \ldots$, all occupied. To accommodate one new guest, shift everyone from room n to room $n + 1$, freeing room 1. To accommodate infinitely many new guests, move everyone from room n to room $2n$, freeing all odd-numbered rooms.

The Banach–Tarski construction can be presented as a combinatorial game. Consider a deck with four types of cards: A, B, A^{-1}, and B^{-1}. These cards can be arranged in sequences, with one rule: if A and A^{-1} appear next to each other, they annihilate. The same applies to B and B^{-1}. A sequence like $ABA^{-1}B$ is stable, but a sequence like $B^{-1}BA$ immediately reduces to just A. The set of all irreducible sequences forms the free group F_2 — all sequences of the cards A and B, with the rule that A and A^{-1} annihilate, and B and B^{-1} annihilate.

Now here's where the cards become geometric transformations. Each card corresponds to a rotation of a sphere: A rotates by an irrational angle (e.g., $\sqrt{2}$ degrees) around the Z-axis, A^{-1} rotates back by the same angle. Similarly, B rotates by an irrational angle around the X-axis, B^{-1} rotates back.

Why irrational angles? Because they ensure that no finite sequence of these rotations will ever bring a point back to exactly where it started (unless all the cards cancel out).

When we fix two rotations A and B, we can build arbitrary sequences of them and their inverses to move points around the sphere. Pick a point p on the sphere. Apply every possible sequence of A, A^{-1}, B, and B^{-1} to p. The resulting collection of points is called the orbit of p under this group of rotations.

Two points lie in the same orbit if one can be turned into the other by some sequence of these rotations. Orbits are disjoint. A point belongs to exactly one orbit, because sequences of rotations either connect two points or they don't. The union of all orbits is the whole sphere, so the orbits form a partition.

Because A and B are chosen carefully (rotations about different axes by irrational angles), the group they generate is free and each orbit is infinite, spreading densely across the sphere.

At this point the Axiom of Choice enters. From each orbit, pick a single representative point. Call the set of these chosen representatives R. Every other point on the sphere can be written uniquely as $g \cdot r$, where g is some sequence of rotations and $r \in R$ is the representative of its orbit. In other words, the entire sphere is recovered by "shuffling" the representatives through all possible rotation sequences.

Now comes the partitioning that makes the paradox work. Four players will divide the entire sphere among themselves by each taking the representatives R and applying only certain rotation sequences:

- *Player 1*: Gets all points reached by sequences starting with card A (avoiding immediate A^{-1} cancellation). This creates the point set S_A.
- *Player 2*: Gets all points from sequences starting with A^{-1} (avoiding A), creating $S_{A^{-1}}$.
- *Player 3*: Gets all points from sequences starting with B (avoiding B^{-1}), creating S_B.
- *Player 4*: Gets all points from sequences starting with B^{-1} (avoiding B), creating $S_{B^{-1}}$.

Every possible rotation sequence must start with one of these four cards (A, A^{-1}, B, or B^{-1}). This means the four sets are disjoint — no point belongs to two different players — and together they cover the entire sphere (except for fixed points on rotation axes, which are handled separately).

First Sphere Reconstruction:

Player 2 gives their entire collection $S_{A^{-1}}$ to Player 1. Player 1 then rotates every point in $S_{A^{-1}}$ by applying rotation A before each sequence, creating the new set $A \cdot S_{A^{-1}}$.

When you rotate a point reached by sequence $A^{-1}BA$ using rotation A, you get $A \cdot (A^{-1}BA) = (AA^{-1})BA = BA$ — the A and A^{-1} cancel. Similarly, rotating the point from sequence A^{-1} gives $A \cdot A^{-1} =$ identity (the original representative point).

Player 1 now has S_A (all points reached by sequences starting with A) and $A \cdot S_{A^{-1}}$ (all points reached by sequences *not* starting with A). Together, these form a complete sphere!

Meanwhile, Players 3 and 4 perform the same trick. Player 4 gives $S_{B^{-1}}$ to Player 3, who rotates it by B to get $B \cdot S_{B^{-1}}$. Combining S_B and $B \cdot S_{B^{-1}}$ gives another complete sphere.

We started with four disjoint pieces that made up one sphere. By rotating two of those pieces, we reassembled them into two identical spheres. No points were added or removed — only rearranged. That is the heart of the paradox.

This paradoxical outcome depends on selecting representative points from each orbit, which requires the Axiom of Choice. The resulting sets are non-measurable — they cannot be assigned consistent volumes that preserve both countable additivity and invariance under rotation. The non-measurability arises from the paradoxical structure of the free group's action on the sphere.

This construction illustrates the core ideas of the Banach–Tarski paradox. The full theorem decomposes a solid three-dimensional ball into five pieces, which can then be rotated into two balls identical to the original. The essential ingredients — free groups, orbit decompositions, the Axiom of Choice, and non-measurable sets — remain the same.

The Axiom of Choice occupies a unique position in modern mathematics. Gödel proved that if ZF (the Zermelo–Fraenkel axioms) (Gödel, 1938) is consistent, then so is ZFC (ZF plus Choice). Cohen (1963) later proved that ZF with the negation of Choice is also consistent, establishing that Choice is independent of ZF — neither provable nor disprovable from the other axioms. This independence was revolutionary. Classical theorems across mathematics depend on it — "every vector space has a basis" and "the product of compact spaces is compact" are each equivalent to accepting Choice.

Before dismissing Choice to avoid the paradox, consider that its absence permits equally frightening results. When you partition a set into disjoint subsets, selecting one representative from each part typically creates an injection from the index set into the original — the number of parts cannot exceed the number of elements. Without Choice, this intuition fails. In ZF alone, it is consistent that a partition $X = \bigsqcup_{i \in I} B_i$ exists with no injection $I \hookrightarrow X$, so the number of parts $|I|$ can exceed $|X|$. A set can be split into more parts than it has elements. Rejecting Choice trades one counterintuitive result for another.

Mathematics versus Reality

> This chapter forces a distinction between mathematical and physical reasoning. The Banach–Tarski construction is not a paradox in the sense of contradiction or physical impossibility, but it is a good example of the consequences of adopting the Axiom of Choice, showing that intuitive notions such as volume are not preserved across all set decompositions. The result is clean and formally sound, yet incompatible with empirical modeling. That gap — between internally consistent mathematics and physically grounded expectation — illustrates the conceptual boundaries explored throughout this book. As with other chapters that emphasize when simplifications fail (relativity in gold, curvature in gravity, topology in voting), this example shows that what appears absurd may instead be a well-posed feature of a chosen formal system.



For the following trick we will require the axiom of choice

Mathematical Foundations of the Banach-Tarski Paradox

Overview. The Banach-Tarski paradox states that a solid ball in $\mathbb{R}^3$ can be partitioned into finitely many disjoint subsets, and, using only rigid motions, reassembled into two identical copies of the original. The result depends on the Axiom of Choice, the existence of a free subgroup of $\mathrm{SO}(3)$, and the absence of a finitely additive rotation-invariant measure on all subsets of the sphere.

Construction Outline

- **Group-Theoretic Basis.** The free group $F_2 = \langle a, b \rangle$ is equidecomposable with two disjoint copies of itself under left multiplication. This violates the property of amenability, which forbids such duplications.
- **Geometric Embedding.** Rotations S and T in $\mathrm{SO}(3)$, each by angle $\theta = \arccos(1/3)$ around orthogonal axes, generate a subgroup isomorphic to F_2. No nontrivial reduced word in these rotations fixes any point off the axes.
- **Removing Fixed Points.** The set D of all points fixed by nontrivial group elements is countable. One can remove D from the sphere by applying a suitable rotation to each of its images.
- **Orbit Decomposition.** The action of F_2 on $S^2 \setminus D$ partitions it into orbits. The Axiom of Choice selects one representative per orbit to form a set M. Then $S^2 \setminus D = \bigsqcup_{g \in F_2} gM$.
- **Duplication via Equidecomposition.** Since $F_2 = A \sqcup B$ with $A \cong F_2$ and $B \cong F_2$, define bijections $\phi_i : F_i \to F_2$, and extend them to $S^2 \setminus D$ via $x = gm \mapsto \phi_i(g)m$. This yields two disjoint subsets each equidecomposable with the whole.
- **Lifting to the Ball.** The unit ball is viewed as concentric spherical shells. The paradoxical decomposition is applied to each shell simultaneously. The center is treated separately.

Amenability and Non-Amenability. A group is amenable if there exists a finitely additive, invariant probability measure on all its subsets. Free groups on two or more generators are non-amenable: they admit no such measure. This failure is what allows equidecomposition of a set with two disjoint isometric copies of itself. The Banach-Tarski paradox is a geometric manifestation of non-amenability.

Cayley Graph Interpretation. The Cayley graph of F_2 is a 4-regular tree. Each branch corresponds to a left coset, and the entire group acts by translation. It is somewhat easy to see that a careful merger of two opposite branches of the tree will result in a tree with the same structure as the original full tree, in a fractal way.

Rotations Generating the Free Group. Define $R = \frac{1}{3}\begin{bmatrix} 1 & -2\sqrt{2} \\ 2\sqrt{2} & 1 \end{bmatrix}$. Then:

$$S = \begin{bmatrix} R & 0 \\ 0 & 1 \end{bmatrix}, \quad T = \begin{bmatrix} 1 & 0 \\ 0 & R \end{bmatrix}.$$

Any reduced word W in S, T, and their inverses maps $(1, 0, 0)$ to a vector of the form $\left(\frac{a}{3^n}, \frac{\sqrt{b}}{3^n}, \frac{c}{3^n}\right)$, with $b > 0$, confirming that the group acts freely on the orbit of $(1, 0, 0)$.

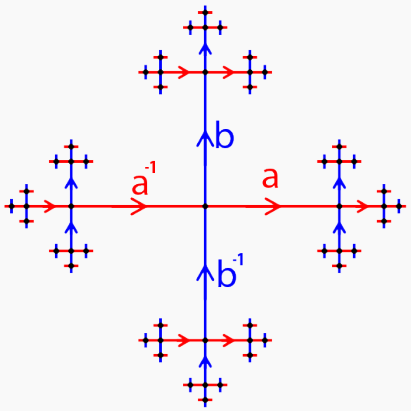


Figure: Cayley graph of F_2, visualized as a tree with no cycles and uniform branching.

References:

Banach, S. and Tarski, A. (1924). *Fund. Math.* **6**
Wu, A. (2008). *The Banach-Tarski Paradox.*