The Banach-Tarski paradox is a result of research in to set theory.
Stefan Banach and Alfred Tarski, after which the paradox is named, were Polish mathematicians.

It states that it is possible to take a solid sphere (a "ball"), cut it up into a finite number of pieces, rearrange them using only rotations and translations, and re-assemble them into two identical copies of the original sphere, effectively doubling the volume of the sphere.

It's important to be clear what we're talking about. The most important point is that when we mention a sphere, we're talking about a mathematical sphere. Physical spheres and mathematical spheres are not entirely congruent, the most important distinction being that a mathematical sphere is infinitely divisible. As a physical sphere contains a finite (albeit large) number of atoms, it is not so.

For simplicity, we shall assume that our mathematical sphere, S, has radius 1. As such, it can be defined as follows:

$$S = \{(x,y,z) \mid x^2+y^2+z^2 \leq 1 \}$$

That is, S is defined as the set of points that lie within a 3-dimensional spherical area in $\square^3$, where $\square$ is the set of all real numbers.

The Definition:

$A_i \cap A_j = \square$ for each i and j between 1 and n such that $i \neq j$ (no two pieces overlap each other)

$A_1 \square A_2 \square ... \square A_n = S$ (assembling all the pieces yields the original sphere S)

There exist $T_1, T_2, ... T_n$, where each $T_i$ represents some finite sequence of rotations and translations, such that if we apply each Ti to each $A_i$ (let's call the result $A_i'$), then:

$A_1' \square A_2' \square ... \square A_m' = S$ (a subset of the original pieces forms S)
$A_{m+1}' \square A_{m+2}' \square ... A_n = S'$ (the remaining pieces forms a copy of S)

where m is some number between 1 and n, and S' is S translated by some finite amount so that S and S' are disjoint.


So effectively, the jargon on the board states that, that it is possible to take S, our previously defined sphere, cut it up into n disjoint pieces, which we've called $A_1$, $A_2$, ..... $A_n$ whose union is S itself, where n is *finite*. Upon performing a finite number of translations and rotations to each piece, we will end up with two copies of S, though one is translated so as not to overlap with the original

The previous statement(s) allow all sorts of unintuitive things to happen. You can cut up a sphere in to n pieces, remove some of the pieces and then reassemble those remaining in to your original sphere, without any change in volume. Obviously, this is impossible with a physical sphere, but such are the joys of mathematics.

This may seem to be utter trite, but when working with mathematically infinite objects, results are certainly not always intuitive. For example, we can look at some slightly more familiar ground. If we take **N = { 0, 1, 2, 3 ... }** , the (infinitely large) set containing all integers, we can split the set into two subsets E and F, containing all the odd and all the even numbers respectively. E = {0, 2, 4, 6 ...}, F = { 1, 3, 5 .. }. Now one would assume that E and F are each smaller than N, but what

happens if we rename each member of E so that a number, x, is renamed to x divided by two? $f(x) = (x/2)$ We find that E is equal to N. Similarly, rename all members of F such that a number, y, is renamed to (y-1)/2 $f(y)=(y-1)/2$ and once again, we have the set of all integers, so F=N. Much like in the Banach-Tarski paradox, we have just duplicated the set of integers using nothing more than the original integers, and some simple operations.

Coming back to spheres, it is helpful to keep in mind that each of these pieces $A_i$ are potentially infinitely complex so that they do not have any well-defined volume. Now, this whole paradox may seem remotely possible if we had, say, required 1,000,000 pieces to achieve our feat; there's intuitively more room for a sleight of hand if we had a million pieces to play with. However, the kicker about this whole paradox is that we don't need more than **five** pieces to achieve this feat. And unlike our odd/even number example, we do not need to play tricks with renaming individual points of each piece; we can perform the miracle by merely using well-behaved operations like rotations and translations. Furthermore, one of these pieces only needs to contain the single point at the centre of the sphere.

In other words, it is mathematically possible to cut S into a mere *four* pieces (if we disregard the one centre point), and to reassemble two of these pieces into the original S, and reassemble the other two into a copy of S.

The catch, of course, is that each of these four pieces are so complex that they do not have any "measure" (i.e., their respective volumes are not well-defined), and that we do not know how to mathematically describe them other than the fact that they exist and exhibit the strange re-assembly property. In fact, it is quite possible that each of those pieces consists of isolated points spread out throughout the entire volume of the original sphere S.

Now, accepting this is really pretty tough, but it's possible to attempt to make an "intuitive" rationalization of it. I'll do it by way of an analogy with a physical sphere.

Let's forget for the moment the mathematical sphere S, which has infinite density. Let's consider a real, physical sphere B (for "ball"), also of radius 1. B is identical to S except that it consists of a finite number of atoms. These atoms are lain out in a crystalline structure. Usually, the crystalline structure of an object is a simple geometric relationship between neighbouring atoms.

Notice that although the geometric relationship between atoms define its crystalline structure, the *precise distance* between atoms in different molecules may vary. This leads to materials of different densities.

Now, we perform the equivalent of a Banach-Tarski decomposition on our physical sphere B: we "atomize" B into four spherical clouds of atoms, let's call them $C_1$, $C_2$, $C_3$, and $C_4$. (We'll ignore the central atom in B, just as in the mathematical version of this decomposition.) Let's assume that each of these clouds are sparse enough that they are gaseous, no longer solid by themselves (to try and imitate the immeasuarability of the four pieces of S ). Furthermore, let's say that the atoms in each of these clouds are laid out in a regular pattern, so that if we rotate $C_1$ by some angle G, and put it together with $C_2$ in the same spherical region, the atoms in both clouds line up into the same crystalline structure as B, except that now the distance between atoms is greater (to account for the missing atoms now in clouds $C_3$ and $C_4$). Similarly, assume we can do the same with $C_3$ and $C_4$: we just translate them away from the original spherical region of B so that they don't interfere with $C_1$ and $C_2$, and reassemble them into another sphere.

Now, we have successfully built two (physical!) spheres with the same radius as B, using only material from B itself. Each of the two spheres have the same crystalline structure as B. The only difference between these spheres and B is that they each have only half the density of B.

To bring this analogy back to the mathematical sphere S: we can think of the infinitely complex pieces $A_1$, ... $A_4$ as the equivalent of "atom clouds", which are non-solid (immeasurable). For those that know something of topology, which excludes myself, you can think of the crystalline structure of S as the topological structure of points in $\Box^3$. These "clouds" lack this "crystalline structure" (i.e., they are unmeasurable); but by suitable rearrangement of them, we can form them into two identical spheres, with half the density of the original, so that they *do* have the same "crystalline structure" (i.e., the resultant two spheres are well-behaved, measurable sets). These two spheres are identical to S, except for having only half the density of S. However, S is infinitely dense, and so are its pieces $A_1$, ... $A_4$. This means the two resultant spheres are *still* infinitely dense. That is to say, they are *identical* to S.

Et voila! There is no paradox here after all. We are merely seeing the logical consequence of mathematical sets like S being infinitely dense. In fact, if you think about it, this is not any stranger than how we managed to duplicate the set of all integers, by splitting it up into two halves, and renaming the members in each half so they each become identical to the original set again. It is only logical that we can continually extract more volume out of an infinitely dense, mathematical sphere S.


**Epilogue**

Now, having convinced you that the Banach-Tarski Paradox isn't really *that* strange after all, I'd like to mention that the derivation of this paradox depends on the Axiom of Choice, and although most mathematicians accept the Axiom of Choice, not all agree with it. There has been much debate over the merit of adopting this axiom, as well as research into the consequences of choosing either way: it does simplify a lot of mathematical proofs, but it also introduces strange results like the Banach-Tarski paradox which we just discussed.

If you're unfamiliar with the Axiom of Choice, it basically goes like this: if you have a collection of sets C (which may potentially contain an uncountably large number of sets), then there exists a set H, called the *choice set*, which contains precisely one element from each (non-empty) set in C. H is called the "choice set" because you are essentially going through each set in C and *choosing* one element from it. One feature of the Axiom of Choice is that H is simply assumed to *exist*; there is no algorithm given which might tell you how to construct an example of H.

In the case of the Banach-Tarski paradox, each of the infinitely complex "pieces" of the sphere S is built from these choice sets. Since we do not know of any algorithm to actually construct these sets, we can only indirectly infer some of the properties of the "pieces", such as their not having a (Lebesgue-) measure (i.e., they have intractible geometric complexity). Some of the debate surrounding the Axiom of Choice revolves around whether these non-constructible sets are mathematically admissible. The reader is encouraged to make good use to Google for more information about this debate; it is too vast a topic to explore in this article. It suffices to say that most mathematicians adopt the Axiom of Choice, simply because of the usefulness of results that can be derived.

One might wonder, then, about what would happen if we *didn't* assume the Axiom of Choice. We do know that we would likely be unable to derive the Banach-Tarski paradox; however, we *also* know that paradoxical sets *do* exist even without the Axiom of Choice. These paradoxical sets exhibit the same "weird" behaviour of the Banach-Tarski spheres, in that you can decompose these sets into a finite number of parts, and reassemble them into multiple copies of the original.