The Banach-Tarsky Paradox for S^2

Derek Sorensen

December 22, 2017

This document contains a (pretty standard) proof of the Banach-Tarski Paradox. It is meant to be a fun read for any person who has the concepts of groups (including free groups) and group actions under her belt. These, along with a firm understanding of countable vs uncountable infinities, are the only prerequisites to understanding the result we present here—even though the content we present here feels geometric in nature, or even topological. A happy surprise to students who heard about this paradox in one of their introductory math courses but have not yet seen the proof. Since it doesn't come up in classes very often, I thought I would include it here for fun.

This "paradox" is really nothing to be scared of. Nor is it very paradoxical. You've already learned that there are the same number of even integers as there are (even and odd) integers. Also, there are as many real numbers in the entire number line as there are between 0 and 1. Considering these initially unintuitive observations about infinity, the Banach-Tarski Paradox is less of a paradox and more of just a fun theorem that you can wow your non-math friends with.

We're going to present this proof in two parts: We will first prove the theorem for S^2 minus a certain countable set D, which we will define. That's our stepping-stone, from there, we can prove the result for all of S^2. There is an analagous result for B^3, and a generalized version of this theorem. We hope to add these to this document in the future.

I suppose we should start by stating the theorem.

Theorem 1 (Banach-Tarski for S^2)**.** *There exists a partition of S^2 into finitely many pieces which can be rotated to form two copies of S^2.*

1 The Hausdorff Paradox

To begin, we'll first make a rather interesting observation about free groups. Take a free group on two generators, $F := \mathcal{F}(x, y)$. Recall that any element in F is a *reduced word* with the alphabet x, y, x^{-1}, y^{-1}. As a set, we can partition F by what element comes first in each word (or element) of F. That is, we can form F_x to be the set of words in F which begin with x, F_y to be the set of words in F which begin with y, and so on for $F_{x^{-1}}$ and $F_{y^{-1}}$. This gives us a partition:

$$F = F_x \sqcup F_y \sqcup F_{x^{-1}} \sqcup F_{y^{-1}} \sqcup e \tag{1}$$

(the square union brackets denote disjoint unions). Let's take a closer look at this partition. Notice that if I take $F_{x^{-1}}$ and multiply each of its elements on the left by x, then I get all of F, minus those elements that begin with x (nothing in $F_{x^{-1}}$ can begin with $x^{-1}x$ because the words have to be reduced). So I also have a partition of F given by

$$F = F_x \sqcup xF_{x^{-1}},$$

and likewise for F_y and $yF_{y^{-1}}$.

This means that I can take F, break it up into five pieces like in (1). Then if I multiply $F_{x^{-1}}$ on the left by x and $F_{y^{-1}}$ by y, I can rearrange the pieces of F into two copies of F, plus an extra point $\{e\}$. The idea of breaking up F strategically, multiplying some of the pieces by the correct element, and then rearranging the pieces into two copies of F is exactly the idea of the Banach-Tarski paradox, but applied to the elements of S^2. The way these two things connect is through rotations of S^2, i.e. the natural group action of $SO(3)$ on S^2. For convenience, we will write the group action on the right.

1.1 The Action of a Free Subgroup of $SO(3)$ on S^2

Take a copy of S^2 and consider the natural group action of $SO(3)$ on S^2 by rotations. We'd like to partition S^2 using orbits of some kind via the action of $SO(3)$ on S^2. If we use the whole of $SO(3)$, this won't yield anything interesting, because for any two points x and y on S^2, there exists some rotation of S^2, i.e. an element of $SO(3)$, that sends x to y. Thus we want to restrict ourselves to some special rotations of S^2, or some subgroup of $SO(3)$. Recall that we're trying to partition S^2 by the group action, so we're going to try to find a *free* subgroup of $SO(3)$ and then restrict our attention to that group.

Recall from before that

$$F = F_\varphi \sqcup F_{\varphi^{-1}} \sqcup F_\psi \sqcup F_{\psi^{-1}} \sqcup e$$

and that

$$F = F_\varphi \sqcup \varphi F_{\varphi^{-1}} = F_\psi \sqcup \psi F_{\psi^{-1}},$$

and so for our purpose, we would ideally like to consider a free group of rotations, on two generators. Fortunately for us, one such group exists!

It is generated by φ, a counterclockwise rotation about the x-axis by $\cos^{-1}(1/3)$, and ψ, a counterclockwise rotation about the y-axis by $\sin^{-1}(1/3)$. In terms of matrices, this translates to

$$\varphi = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{-2\sqrt{2}}{3} \\ 0 & \frac{2\sqrt{2}}{3} & \frac{1}{3} \end{pmatrix} \quad \psi = \begin{pmatrix} \frac{1}{3} & \frac{-2\sqrt{2}}{3} & 0 \\ \frac{2\sqrt{2}}{3} & \frac{1}{3} & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

As it turns out, φ and ψ generate a free subgroup of $SO(3)$. I will include this proof in the appendix (if it's not there now, it will be in the future). In any case, the specifics of φ and ψ are not important; we could choose any free subgroup of $SO(3)$ on two generators that we like. The free group on two generators is all we need here.

Let $F = \mathcal{F}(\varphi, \psi)$, the free group generated by φ and ψ. Now that we have F, let's consider F-orbits of S^2, via the natural action. Note that a finitely-generated free group is countable, and so any

F-orbit of S^2 can only be countable. Therefore, there are uncountably many F-orbits on S^2 (the countable union of countable sets is countable). Let's see if we can partition this set into finitely many pieces. There are considerably more than finitely many orbits, so we'll have to get creative here.

Using the Axiom of Choice, choose one representative of each orbit, and put them all into a set which we will call B'_0 (we'll define an improved set, B_0, later). We can think of this of having taken F, chosen a starting point, hit it with F, and then chosen a new starting point which was not hit in the orbit of our first, then hitting the new one with F, and so on, "uncountably many times"—whatever that means. After finishing that, we've taken all the starting points and put them into B'_0. (This explanation is for intuition, not rigor!)

Note that, after choosing B'_0, each point in $x \in S^2$ can be reached by hitting an element of B'_0 with an element, say ρ, of F. We know that any element of F is a word with the letters φ, φ^{-1}, ψ, ψ^{-1} and so we can define the following set:

$$B'_\varphi := \{x \in S^2 \,|\, x = b\rho, \text{ where } b \in B'_0 \text{ and } \varphi \text{ is the first letter of } \rho.\}$$

Intuitively, this is the set of elements of S^2 that can be reached from B'_0 by a rotation which starts with φ. We can do the same things for $B'_{\varphi^{-1}}$, B'_ψ and $B'_{\psi^{-1}}$. And because of our decomposition of F, we can write down the following decomposition:

$$S^2 = B'_0 \cup B^{-1}_\varphi \cup B'_{\varphi^{-1}} \cup B'_\psi \cup B'_{\psi^{-1}}.$$

And we can use our trick from before to rotate two of the sets and rearrange them to get

$$S^2 = B'_\varphi \cup \varphi B'_{\varphi^{-1}} = B'_\psi \cup \psi B'_{\psi^{-1}}$$

where

$$\varphi B'_{\varphi^{-1}} = \{b\varphi\rho \,|\, b\rho \in B'_{\varphi^{-1}}\}, \text{ and } \psi B'_{\psi^{-1}} = \{b\psi\rho \,|\, b\rho \in B'_{\psi^{-1}}\}.$$

This is just what we wanted! Wait...

This is exciting and all, but we run into a slight problem here. If you noticed, for this decomposition, I used $\cup$ instead of $\sqcup$ to denote the unions. What was meant by this is that these unions need not be disjoint. There is a set of points in S^2 that gives us problems, and that we have to account for. These points are those fixed by rotations in F; each rotation of S^2 has two antipodal points which are fixed by the rotation—the axes of rotation.

To show why this is an issue, suppose that x is on the axis of rotation of $\rho \in F$, and without loss of generality, suppose that the first rotation in ρ is φ. If we had the misfortune of choosing x to be in B'_0, then x is in B_0 (the set of "starting points") but x is also in B'_φ (the set of elements in S^2 which can be reached from B'_0 by a rotation which starts with φ). We might try to choose B'_0 carefully, so that nothing in B'_0 lies on a axis of rotation, but this is actually impossible.

To show this, suppose that we can choose B'_0 in this way, and suppose that x is on the axis of rotation of some non-identity element $\rho \in F$ (thus $x \notin B'_0$). In other words, $x\rho = x$, where $\rho \neq e$. Then there is some $b \in B'_0$ and some non-identity $\sigma \in F$ such that $b\sigma = x$, and $(b\sigma)\rho = b\sigma$. Note that $\sigma\rho\sigma^{-1}$ is not the identity. This gives us that b is fixed by $(\sigma\rho\sigma^{-1})$, since:

$$b(\sigma\rho\sigma^{-1}) = (b\sigma)\rho\sigma^{-1} = (b\sigma)\sigma^{-1} = b.$$

Thus b was on the axis of rotation of an element of F, and this gives us a contradiction. So we definitely have a problem with the unions from before not being disjoint.

Notice that what we just proved here was that if I have an F-orbit which contains one point of S^2 lying on an axis of rotation of an element of F, then all of them lie on some axis of rotation of some (different) element of F. This means that if I take out all the points of S^2 that lie on some axis of rotation, then I can still consider F-orbits on the remaining points without any problems. Cool!

This means that, if we write D to be the set of all points in S^2 fixed by some element of F (for those familiar with group actions, D is the union of stabilizers of elements of F), then we can rewrite our decomposition of S^2 before, but with disjoint unions. To that end, define the sets B_0, B_φ, $B_{\varphi^{-1}}$, B_ψ and $B_{\psi^{-1}}$ to be the intersection of their prime counterparts with $S^2 \setminus D$. That is, $B_0 := B_0' \cap S^2 \setminus D$, $B_\varphi := B_\varphi' \cap S^2 \setminus D$, etc. The reason we can do this without having to rechoose B_0 is that when we removed D from S^2, we took out whole F-orbits at a time, and so the structure that we wanted from these sets is still there.

This gives us:

$$S^2 \setminus D = B_0 \sqcup B_\varphi^{-1} \sqcup B_{\varphi^{-1}} \sqcup B_\psi \sqcup B_{\psi^{-1}}.$$

And we can use our trick from before to rotate two of the sets and rearrange them to get

$$S^2 = B_\varphi \sqcup \varphi B_{\varphi^{-1}} = B_\psi \sqcup \psi B_{\psi^{-1}}$$

where

$$\varphi B_{\varphi^{-1}} = \{b\varphi\rho \,|\, b\rho \in B_{\varphi^{-1}}\}, \text{ and } \psi B_{\psi^{-1}} = \{b\psi\rho \,|\, b\rho \in B_{\psi^{-1}}\}.$$

What this means is that we took $S^2 \setminus D$, carefully divided it up into five parts, rotated two of them, and reassembled them to get two identical copies of $S^2 \setminus D$, and an extra B_0 to spare. This is much closer to what we wanted, because the unions are finally disjoint. Now all we have to do is work out something with D, which is where it really gets interesting!

1.2 Resolving D

This is actually a quick, relatively painless trick. To illustrate what we're about to do, consider the unit circle S^1 in $\mathbb{C}$. Start with the point $e^0 = 1$ and choose points by rotating counter-clockwise by any rational angle, say by 1/10. This gives us a sequence of points, $e^0, e^{i/10}, e^{2i/10}, e^{3i/10}$, and so on. Note that since 2π is irrational, and 1/10 is rational, this sequence will never repeat itself. If we call the sequence D', we can do a funny trick: Take D' in S^1 and multiply everything in D' by $e^{i/10}$. In other words, rotate each element in D' by the angle 1/10. All that we've done is move around some points of S^1 (we haven't deleted any of them), but what we end up with is S^1, minus the point at 1. If we continue rotating, we can take out as many points as we like—indeed, by simply rotatiting D' we can remove any finite number of points from S^1 by rearranging the points, and not actually deleting them. This can be done for any infinite sequence D', and only requires that we move elements of D' along in the sequence. Furthermore, if we start with any finite number of points missing from S^1, we can fill them by using a sequence. If the sequence follows the pattern of D' (i.e. its terms are made from rational rotations), you can fill the empty points by rotating a subset of S^1.

With this in mind, let's get back to our case with S^2 and D. We're going to do a trick similar to the one we just did for S^1, but for S^2. Take a rotation wich does not fix any points of D. This is done by just taking a rotation of S^2 in $SO(3) \setminus F$, as its axis of rotation, by definition D, does not contain any points of D. Let that asix be ℓ. (One can also argue that D is countable, and there are uncountably many antipodal points in S^2, so such an axis must exist.) Using ℓ, we'd like to find an angle, just as we did in the case of S^1, that allows us to rotate the points of D about ℓ without ever hitting the same points twice.

Formally, let ℓ_θ be the rotation of S^2 about ℓ, counterclockwise, by the angle $\theta \in [0, 2\pi)$. What we want is to find θ such that, for all $n \neq m$, $n, m \geq 0$,

$$D\ell_\theta^n \cap D\ell_\theta^m = \emptyset. \tag{2}$$

Let's first prove a slightly weaker result, which is that we can find a θ such that, for all $n > 0$,

$$D \cap D\ell_\theta^n = \emptyset. \tag{3}$$

Let's prove this by constructing the following set:

$$A = \{\alpha \in [0, 2\pi) \mid \text{ there is some } d \in D \text{ and } n > 0 \text{ such that } d\ell_\alpha^n \in D.\}$$

Now, what is the cardinality of A? Since any $d \in D$ can be rotated onto another point $d' \in D$ at most a countable number of ways, and D is countable, for each d, there are only countably many $\alpha \in A$ which rotate d onto another point in D. Again, since D is countable and the countable union of countable sets is countable, A must be countable. Therefore, $[0, 2\pi) \setminus A$ is nonempty, meaning that we can choose some θ that satisfies Equation (3). Great!

Now why does (2) follow? If there are $m, n \geq 0$ such that $n \neq m$ and $D\ell_\theta^n \cap D\ell_\theta^m$ is nonempty, then there are elements $d, d' \in D$ such that $d\ell_\theta^n = d'\ell_\theta^m$. Without loss of generality, suppose that $n < m$. Then hitting both sides by $\ell_\theta^{-n} = (\ell_\theta^n)^{-1}$, we get that $d = d'\ell_\theta^{m-n}$, which contradics Equation (3).

Now we're ready to apply the trick. Let's take $\mathcal{D}$ to be the union of D and all its powers under $\sigma = \ell_\theta$. That is,

$$\mathcal{D} = \bigcup_{n=0}^{\infty} D\sigma^n.$$

Note that $\mathcal{D}\sigma = \mathcal{D} \setminus D$, and thus that $(\mathcal{D} \setminus D)\sigma^{-1} = \mathcal{D}$.

This means that if we have $S^2 \setminus D$, we can fill in the points of D by simply rotating the set $\mathcal{D}$ within S^2 via σ^{-1}. Conversely, if we have S^2 we can erase the points of D by simply rotating $\mathcal{D}$ by σ. That is, anything we can do to $S^2 \setminus D$ we can also do to S^2, and visa versa.

So, how does the decomposition go in the Banach-Tarski paradox? The idea is to take S^2, rotate $\mathcal{D}$ to get $S^2 \setminus D$, apply the decomposition from before to get two copies of $S^2 \setminus D$, and then on each, rotate $\mathcal{D}$ by σ^{-1} and we get two copies of S^2.

On the next page is a diagram to illustrate:

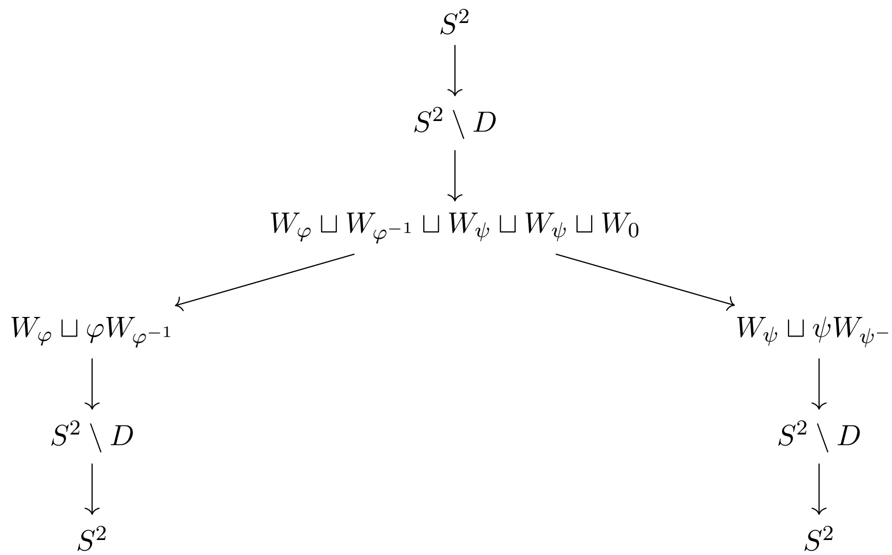

The Banach-Tarski Paradox is given in a very specific form, namely that we need to divide up S^2, rotate some of the pieces, and then put them back together to get two copies of S^2. What we've just states is good, but it doesn't give us explicit rotations that we can use to satisfy the statement of the paradox. What this gives us is a partition which can be put back together to get two copies of S^2. It's close, but not quite there. To complete this result for S^2, we just have to do some calculations and break up the S^2 into a few more pieces than what we've shown. This will give explicit rotations which compose to give the operation that we see in the diagram here. We note here that it is possible to do this with only five pieces, but our proof will have to give a few more.

1.3 Completing the proof for S^2

The problem with trying to pass of the above as suitable for the proof is that we need to account for the portion of $\mathcal{D}$ which is not in D. Essentially, we're using the same decomposition as before, but we're going to be careful about the points inside of $\mathcal{D}$ and those in $S^2 \setminus \mathcal{D}$.

So first of all, let's write S^2 in terms of $\mathcal{D}$ instead of D. Note first that:

$$S^2 = [(S^2 \setminus D) \cap (S^2 \setminus \mathcal{D})] \sqcup \sigma^{-1}[(S^2 - D) \cap \mathcal{D}]$$

because $S^2 = (S^2 \setminus \mathcal{D}) \sqcup \mathcal{D}$ and $\mathcal{D} = \sigma^{-1}[(S^2 \setminus D) \cap \mathcal{D}]$ and $S^2 \setminus \mathcal{D} = [(S^2 \setminus D) \cap (S^2 \setminus \mathcal{D})]$ (this because $D \subset \mathcal{D}$). Let's use our decompositions of $S^2 \setminus D$ from before to rewrite this. Remember that

$$\begin{aligned} S^2 \setminus D &= B_0 \sqcup B_\varphi^{-1} \sqcup B_{\varphi^{-1}} \sqcup B_\psi \sqcup B_{\psi^{-1}} \\ &= B_\varphi \sqcup \varphi B_{\varphi^{-1}} \\ &= B_\psi \sqcup \psi B_{\psi^{-1}}. \end{aligned}$$

By substituting each of our three decompositions of $S^2 \setminus D$, we get:

$$\begin{aligned} S^2 &= [(S^2 \setminus D) \cap (S^2 \setminus \mathcal{D})] \sqcup \sigma^{-1}[(S^2 - D) \cap \mathcal{D}] \\ &= [(B_0 \sqcup B_\varphi^{-1} \sqcup B_{\varphi^{-1}} \sqcup B_\psi \sqcup B_{\psi^{-1}}) \cap (S^2 \setminus \mathcal{D})] \sqcup \sigma^{-1}[(B_0 \sqcup B_\varphi^{-1} \sqcup B_{\varphi^{-1}} \sqcup B_\psi \sqcup B_{\psi^{-1}}) \cap \mathcal{D}] \\ &= [(B_\varphi \sqcup \varphi B_{\varphi^{-1}}) \cap (S^2 \setminus \mathcal{D})] \sqcup \sigma^{-1}[(B_\varphi \sqcup \varphi B_{\varphi^{-1}}) \cap \mathcal{D}] \\ &= [(B_\psi \sqcup \psi B_{\psi^{-1}}) \cap (S^2 \setminus \mathcal{D})] \sqcup \sigma^{-1}[(B_\psi \sqcup \psi B_{\psi^{-1}}) \cap \mathcal{D}] \end{aligned}$$

In order to account for $\mathcal{D}$ in all of the decompositions, we need to decompose each of $B_0, B_\varphi, B_{\varphi^{-1}}, B_\psi$, and $B_{\psi^{-1}}$ into the points inside of $\mathcal{D}$ and the points out of $\mathcal{D}$. This way, we treat $\mathcal{D}$ on its own. Furthermore, when we twist and put back together to get two copies of S^2, we need to do the same thing over again on $B_{\varphi^{-1}}$ and $B_{\psi^{-1}}$, but with $\mathcal{D}\varphi$ and $\mathcal{D}\psi$, repsectively.

Thus, we decompose B_0, B_φ and B_ψ into:

$$\begin{aligned} B_0 :& B_0^1 = B_0 \cap (S^2 \setminus \mathcal{D}) \\ & B_0^2 = B_0 \cap \mathcal{D} \\ B_\varphi :& B_\varphi^1 = B_\varphi \cap (S^2 \setminus \mathcal{D}) \\ & B_\varphi^2 = B_\varphi \cap \mathcal{D} \\ B_\psi :& B_\psi^1 = B_\psi \cap (S^2 \setminus \mathcal{D}) \\ & B_\psi^2 = B_\psi \cap \mathcal{D}, \end{aligned}$$

so $B_0 = B_0^1 \sqcup B_0^2$, and so on with B_φ and B_ψ.

In order to take into account the rotations by φ and ψ respectively, we're going to have to split $M_{\varphi^{-1}}$ and $M_{\psi^{-1}}$ into four pieces each, instaead of just two, as follows: We first split with regards to $\mathcal{D}$, and then split with regards to, respectively, $\mathcal{D}\varphi$ and $\mathcal{D}\psi$.

$$\begin{aligned} B_{\varphi^{-1}} :& B_{\varphi^{-1}}^{11} = B_{\varphi^{-1}} \cap (S^2 \setminus \mathcal{D}) \cap (S^2 \setminus \mathcal{D})\varphi^{-1} \\ & B_{\varphi^{-1}}^{12} = B_{\varphi^{-1}} \cap (S^2 \setminus \mathcal{D}) \cap \mathcal{D}\varphi^{-1} \\ & B_{\varphi^{-1}}^{21} = B_{\varphi^{-1}} \cap \mathcal{D} \cap (S^2 \setminus \mathcal{D})\varphi^{-1} \\ & B_{\varphi^{-1}}^{22} = B_{\varphi^{-1}} \cap \mathcal{D} \cap \mathcal{D}\varphi^{-1} \end{aligned}$$

$$\begin{aligned} B_{\psi^{-1}} :& B_{\psi^{-1}}^{11} = B_{\psi^{-1}} \cap (S^2 \setminus \mathcal{D}) \cap (S^2 \setminus \mathcal{D})\psi^{-1} \\ & B_{\psi^{-1}}^{12} = B_{\psi^{-1}} \cap (S^2 \setminus \mathcal{D}) \cap \mathcal{D}\psi^{-1} \\ & B_{\psi^{-1}}^{21} = B_{\psi^{-1}} \cap \mathcal{D} \cap (S^2 \setminus \mathcal{D})\psi^{-1} \\ & B_{\psi^{-1}}^{22} = B_{\psi^{-1}} \cap \mathcal{D} \cap \mathcal{D}\psi^{-1} \end{aligned}$$

Thus $B_{\varphi^{-1}} = B_{\varphi^{-1}}^{11} \sqcup B_{\varphi^{-1}}^{12} \sqcup B_{\varphi^{-1}}^{21} \sqcup B_{\varphi^{-1}}^{22}$, and likewise for $B_{\psi^{-1}}$.

Now, by just plugging in our decompositions and simplifying, we get the following equation:

$$\begin{aligned} S^2 &= B_0^1 \sqcup \sigma^{-1} B_0^2 \sqcup B_\varphi^1 \sqcup \sigma^{-1} B_\varphi^2 \sqcup B_\psi^1 \sqcup \sigma^{-1} B_\psi^2 \sqcup B_{\varphi^{-1}}^{11} \sqcup B_{\varphi^{-1}}^{12} \sqcup \\ &\quad \sigma^{-1} B_{\varphi^{-1}}^{21} \sqcup \sigma^{-1} B_{\varphi^{-1}}^{22} \sqcup B_{\psi^{-1}}^{11} \sqcup B_{\psi^{-1}}^{12} \sqcup \sigma^{-1} B_{\psi^{-1}}^{21} \sqcup \sigma^{-1} B_{\psi^{-1}}^{22} \\ \\ &= \varphi B_{\varphi^{-1}}^{11} \sqcup (\varphi\sigma^{-1})\sigma B_{\varphi^{-1}}^{21} \sqcup \sigma^{-1}\varphi B_{\varphi^{-1}}^{12} \sqcup (\sigma^{-1}\varphi\sigma^{-1})\sigma B_{\varphi^{-1}}^{22} \sqcup B_\varphi^0 \sqcup \sigma^{-1} B_\varphi^1 \\ \\ &= \psi B_{\psi^{-1}}^{11} \sqcup (\psi\sigma^{-1})\sigma B_{\psi^{-1}}^{21} \sqcup \sigma^{-1}\psi B_{\psi^{-1}}^{12} \sqcup (\sigma^{-1}\psi\sigma^{-1})\sigma B_{\psi^{-1}}^{22} \sqcup B_\psi^0 \sqcup \sigma^{-1} B_\psi^1 \end{aligned}$$

If you look carefully at these equations, you'll see that they are giving the explicit rotations which compose to form the transformation from S^2 into two copies of S^2 that we gave in the big diagram before. This completes the proof of the Banach-Tarski Paradox for S^2.