The Banach-Tarski 'paradox': a basic magic of group theory

JACOB L HOOVER, AFTER NIR AVNI, 2 May 2013

The following is a description of the famous Banach-Tarski paradox, which shows that you can conjure things out of thin air *...with the axiom of choice (you also need an infinitely divisible object, but only on short loan, because after you cut it up and reässemble it, you'll have two!). It assumes only basic knowledge of group theory, and is based on my notes from a presentation by Nir Avni, in Mathematics 122 (Algebra I), 10 Nov 2011, Harvard College.*

Very brief summary

Take the unit ball in three dimensions. We will cut up the ball into pieces, and use a group called F_2 to shuffle the pieces, using only rotations about the origin, and find them reassembled into two identical copies of the original ball.

Yes, it works! In fact, nothing about this is contradictory. It is just contracommonsensical.

Background

This is the part you can skim. I will lay out a theorem to be proven, some definitions, and a couple propositions to get started.

Theorem 1. *There exists an injective homomorphism $\phi: F_2 \hookrightarrow SO_3(\mathbb{R})$, where F_2 is the free group on two letters, and $SO_3(\mathbb{R})$ is the 3-dimensional special orthogonal group over the reals.*[1]

While this fact is very important, the proof is not particularly relevant, so I left it for the Appendix.

Definition 1 (disjoint partition)**.** *A* disjoint partition *of set S is a sequence of subsets, $S_1,\dots,S_n \subseteq S$, such that $S_1 \cup \dots \cup S_n = S$, and $S_i \cap S_j = \varnothing, \forall i,j$. Let such a partition be written as $S = S_1 \amalg \dots \amalg S_n$.*

Definition 2 (paradoxical)**.** *Suppose group G acts on*[2] *set X. This action is* paradoxical *if there are disjoint subsets (not necessarily a partition) $X_1,\dots,X_n \subseteq X$, some (not necessarily distinct) elements $g_1,\dots,g_n \in G$, and integer $1 \le m \le n$ such that*

$$g_1X_1 \amalg \dots \amalg g_mX_m = g_{m+1}X_{m+1} \amalg \dots \amalg g_nX_n = X.$$

This is called a paradoxical decomposition *of X.*

Here is an easy, useful, demonstation of a paradoxical action:

Proposition 1. *The action of F_2 on itself by left multiplication is paradoxical.*

Proof. Let the two generating elements of F_2 be called a and b (i.e. $F_2 = \langle\{a,b\}\rangle$).

Now let $A_+ = \{\text{words in } F_2 \text{ starting with } a\}$, and $A_- = \{\text{words in } F_2 \text{ starting with } a^{-1}\}$, and define sets B_+, B_- likewise. These four subsets are disjoint, and the only element of F_2 not contained in one of the four is the identity, 1, so $F_2 = \{1\} \amalg A_+ \amalg A_- \amalg B_+ \amalg B_-$.

Now let a^{-1} act on A_+. The set $a^{-1}A_+$ is precisely the set of all words in F_2 not beginning with a^{-1}. That is, $a^{-1}A_+ = \{1\} \amalg A_+ \amalg B_+ \amalg B_-$. Likewise, $b^{-1}B_+ = \{1\} \amalg A_+ \amalg A_- \amalg B_+$.

Thus,

$$a^{-1}A_+ \amalg A_- = b^{-1}B_+ \amalg B_- = F_2$$

is a paradoxical decomposition, and this action of $F_2 \curvearrowright F_2$ is paradoxical. □

Now, a similar, but more general result, again about the action of F_2. This one will be directly useful to the Banach-Tarski proof.

Proposition 2. *Suppose that F_2 acts on a set X such that for all nonidentity $g \in F_2$, g has no fixed points in X. This action is paradoxical.*

Proof. Invoke the axiom of choice to choose a single point of X from each orbit[3] of under this action, and collect them to form a set, Y. Let $Y_{A_+} = \{g \cdot y \mid g \in A_+, y \in Y\}$, and define $Y_{A_-}, Y_{B_+}, Y_{B_-}$ likewise, for the subsets $A_+, A_-, B_+, B_- \subset F_2$ as defined in Proposition 1, above.

Now, the proof relies on the following two propositions:

1. $Y_{A_+}, Y_{A_-}, Y_{B_+}, Y_{B_-} \subset X$ are all disjoint.
2. $a^{-1}Y_{A_+} \amalg Y_{A_-} = b^{-1}Y_{B_+} \amalg Y_{B_-} = X$.

Proof of 1. Assume that they are not all disjoint. Without loss of generality, assume that $Y_{A_+} \cap Y_{A_-} \neq \varnothing$. Then, $gy = g'y'$ for some $g \in A_+, g' \in A_-, y, y' \in Y$, so $(g')^{-1}gy = y'$. The element $(g')^{-1}g \in F_2$, thus $y, y' \in Y$ are in the same orbit, so $y = y'$, and is a fixed point of $(g')^{-1}g$.

By assumption, if $(g')^{-1}g$ has a fixed point, it is the idenity, so $g' = g$. But, then they cannot be one in A_+ and the other in A_-. Contradiction. □

Proof of 2. Showing $X = a^{-1}Y_{A_+} \amalg Y_{A_-}$ will be sufficient. And, for this, it will be clearly be enough just to show that $X \subseteq a^{-1}Y_{A_+} \cup Y_{A_-}$.

Given some $x \in X$, the orbit of x intersects Y at a point y, by assumption. Then there is some $g \in F_2$ such that $gx = y$, so $x = g^{-1}y$. And, $g^{-1} \in F_2 = a^{-1}A_+ \amalg A_-$, so there are only two possibilities: either $g^{-1} \in a^{-1}A_+$, so $g^{-1} = a^{-1}h$, for $h \in A_+$ (in which case $x = a^{-1}hy$ so $x \in a^{-1}Y_{A_+}$), or $g^{-1} \in A_-$ (in which case $x \in Y_{A_-}$). □

So, this action of F_2 on X is paradoxical. □

The Banach-Tarski Magic

The point is to show that action of SO_3 on the unit ball is paradoxical.

Since we already know something about making paradoxical actions of F_2, it would be a nice start to have F_2 act on the unit ball in $\mathbb{R}^3$. By Theorem 1, there is an injection $\phi: F_2 \hookrightarrow SO_3$, so, if SO_3 acts on set X, an action of F_2 on X can be formed just by applying ϕ first. Well, here is a helpful fact, then:

Proposition 3. *The group $SO_3(\mathbb{R})$ is isomorphic to the group of rotations in $\mathbb{R}^3$.*

[1] If you are unfamiliar with these entities, see Appendix.

[2] An *action* of group G on set X (written $G \curvearrowright X$) is a function from $G \times X \to X$, satisfying associativity and preserving identities. It should be a rather intuitive notion. Basically, it is just the formal way to allow a group to manipulate a set; it specifies the operation by which they interact.

[3] For any element $x \in X$, the *orbit* is defined as $F_2x = \{g \cdot x \mid g \in F_2\}$. Note that, of course, distinct elements in X will be in the same orbit, and every orbit will have more than one element in it, by assumption. There may be an uncountable infinity of orbits.

[If you are used to thinking of SO_3 as the rotation group by definition, then no proof is necessary. If you are used to it being defined abstractly as in Definition 4 below, then see the proof in the Appendix.]

So SO_3 acts simply as rotations on the unit ball. Nice. Next, we make F_2 act through rotations on the ball, and show that it is a paradoxical action—one that can do the magic of Banach-Tarski.

In order to do that, let's first limit ourselves to the surface of the ball, the 2-sphere, S^2. We have an action $F_2 \curvearrowright S^2$. Now, all we need to do is show that it is paradoxical, observe that it can be done with just rotations in 3 dimensions, and then extend inward from the sphere on the surface to the entire, 3D, ball.

Proposition 4. *Action by F_2 on (most of) S^2 is paradoxical.*

Proof. For each nonidentity $g \in F_2$, $\phi(g)$ is a rotation about some axis ℓ_g. This line intersects the sphere in two points, $\ell_g \cap S^2 = \{\pm p_g\}$. Let $D = \{\pm p_g \mid g \in F_2 \setminus \{1\}\}$.

Take $g,h \in F_2 \setminus \{1\}$. The axis of g is the line $\ell_g = \mathbb{R}p_g$. And the axis of the conjugate of g by h (the element hgh^{-1}) is $\mathbb{R}\phi(h)p_g$. But, this axis is the line $\ell_{hgh^{-1}} = \mathbb{R}p_{hgh^{-1}}$, so, $p_{hgh^{-1}} = \phi(h)p_g$ (with a possible sign difference). Therefore, $\phi(h)D = D$; that is to say, poles are mapped only to poles.

The set $S^2 \setminus D$ is invariant under the action by F_2, which is without fixed points. So, by Proposition 2, this action of $F_2 \curvearrowright S^2 \setminus D$ is paradoxical. □

And, since $\phi(F_2) \subset SO_3$, we have that $SO_3 \curvearrowright S^2 \setminus D$ is paradoxical. Now, to the main theorem!

Theorem 2. *Action by $SO_3(\mathbb{R})$ on S^2 is paradoxical.*

Proof. The set D is countable. S^2 is not. There are uncountably many lines through the origin that are not the axis of any rotation in $\phi(F_2)$.

Since the interval $[0,2\pi]$ is (continuous, and the continuum is) uncountable, there exists a rotation $r \in SO_3$, such that $D \cap r(D) \cap r^2(D) \cap \cdots = \varnothing$.

Take

$$S^2 \setminus D = S^2 \setminus (D \cap r(D) \cap r^2(D) \cap \ldots) \amalg (r(D) \cap r^2(D) \cap \ldots)$$

So, letting the identity act on the first part of the above disjoint union, and letting r^{-1} act on the second, and taking the union of the results, we get

$$\big(S^2 \setminus (D \cap r(D) \cap \ldots)\big) \amalg (D \cap r(D) \cap \ldots) = S^2.$$

We have just created S^2 by cutting up and rotating $S^2 \setminus D$.

So, since $SO_3 \curvearrowright S^2 \setminus D$ is paradoxical, just add this procedure to the end of this paradoxical action, and we have a paradoxical action of $SO_3 \curvearrowright S^2$. □

So there you have it. You can take two subsets of points of S^2, rotate them, and turn them each into the entire sphere!

Now, to be complete, let's extend this business inward so we can do it to a three dimensional ball. This would be trivial (just extend inwards), but the center point, $\{0\}$, creates a little difficulty, so I'll be explicit:

Theorem 3. *Action by rotation on the three dimensional unit ball, B, is paradoxical.*

Proof. Define a small 2-sphere $S \subset B$ with $\{0\} \in S$. Taking $B = (B \setminus S) \amalg S$, and acting on $B \setminus S$ with the identity, and on S in a familiar cut-up-and-rotate manner[4] to get $S \setminus \{0\}$, we get $(B \setminus S) \amalg (S \setminus \{0\}) = B \setminus \{0\}$.

Now, create the obvious map $S^2 \to B \setminus \{0\}$ made by extending inwards from each point on the surface. Given the paradoxical decomposition $S^2 = X_1 \amalg \cdots \amalg X_n$, define an analogous one $B \setminus \{0\} = \tilde{X}_1 \amalg \cdots \amalg \tilde{X}_n$ where each $\tilde{X}_i = \{v \in B \setminus \{0\} \mid v/\|v\| \in X_i\}$, for $i = 1,\ldots,n$.

We have now succeded in duplicating the ball minus the center point. Cut-up-and-rotate one last time to get from each copy of $B \setminus \{0\}$ to B, and the magic is complete. □

Appendix

Definition 3 (Free group). *Take a set $S = \{a,b,c\ldots\}$, and call the elements "letters". Now create $S' \supset S$, by adding "inverses" $a^{-1},\ldots$, and an "identity", 1. The* free group *F, spanned by S, is the set of all "words" (strings) of nonidentity elements of S', wherein no letter is adjacent to its inverse, plus the identity; the group's operation is concatenation. The only rule is this: when concatenating, if a letter appears adjacent to its inverse, the pair is deleted. Explicitly, the group presentation is $F = \langle S \mid \varnothing \rangle$ (*i.e. *there are no relations—that is, the the identity cannot be reached except trivially).*

For $S = \{a,b\}$, the elements of the free group $F_2 = \langle S \rangle$ are all the reduced words that can be made from these letters: $\{1,a,b,a^{-1},\ldots,ba^{-1}b^{-1}a\ldots\}$.

Definition 4 (Special Orthogonal Group). *The Orthogonal Group, $O_n(F)$ is defined as the set of all $n \times n$ norm-preserving invertible matrices with entries in field F, with the operation of matrix multiplication. The* Special Orthogonal Group *$SO_n(F)$ is the subgroup of $O_n(F)$ comprised of all elements whose determinant is* 1.

For $F = \mathbb{R}, n = 3$, $SO_3(\mathbb{R})$ is the group of rotations in 3 dimensions, by Proposition 3.

Proof of Proposition 3 ($SO_3(\mathbb{R}) \cong$ rotations in $\mathbb{R}^3$)

Proof. Suppose $g \in SO_3$. The characteristic polynomial of g has degree 3. Odd polynomials always have a root (by the intermediate value theorem), so g has an eigenvalue that is real. Let v be the eigenvector of this eigenvalue, giving $gv = \lambda v$, and thus $\|gv\| = \|\lambda\|$ since all elements of SO_3 have norm one, so $\lambda = \pm 1$.

If $\lambda = +1$, then g is a rotation around the line $\mathbb{R}v$. There are two cases:

(i) Only one eigenvalue is real: the three eigenvalues are $\lambda, z, \bar{z}$. The determinant $\det(g) = \lambda z\bar{z} = 1$, and $z\bar{z}$ is positive, so λ is positive.

(ii) There are three real eigenvalues, all ± 1: their product, $\det(g) = 1$, so at least one is $+1$.

□

Proof of Theorem 1 ($\exists\ \phi : F_2 \hookrightarrow SO_3(\mathbb{R})$)

Proof. Choose two orthogonal axes of $\mathbb{R}^3$, say the x and the y. Let $R_x, R_y \in SO_3$ be the matrices of rotation about the x- and y- axes by angle $\theta = \arccos(1/3)$. The subgroup of SO_3 freely generated by these two rotations has the same structure as F_2. So, define $\phi : F_2 \hookrightarrow SO_3$ by $a \mapsto R_x, b \mapsto R_y, 1 \mapsto \mathrm{Id}$. □

(This is not the only map that would work; there are many other possibilities, but the point is just that a map exists. It doesn't matter to us what it is, in particular.)

[4] in the proof of Theorem 3, we cut up $S^2 \setminus D$ and rotated bits to get S^2. Just do this in reverse, with $S, \{0\}$ instead of S^2, D.