The Banach-Tarski Paradox

Ziv Scully

MIT Educational Studies Program

Contents

1 Group Theory **2**
1.1 Introduction . . . 2
1.2 Group Actions and Orbits . . . 4
1.3 Generators . . . 6

2 Explanation of the Paradox **8**
2.1 The Free Group on Two Generators . . . 8
2.2 Equidecomposition . . . 10
2.3 Poking Holes in Spheres . . . 12
2.4 The Paradox Itself . . . 14

3 Linear Algebra **17**
3.1 Kernels and Images . . . 18
3.2 A Free Group of Rotations . . . 20

Hi, everyone!

As promised, here is my Banach-Tarski Paradox writeup, as those of you who took a strategic look at the title of this document probably figured out. It covers the entire lecture with some additional background. Section 1 has important, easier material that I didn't go over in the lecture. Section 2 has the the lecture. Section 3 has less important, less easier material that I didn't go over in the lecture which you may find interesting, though you'll want to know a bit about matrices. In particular, we show that we can pick two rotations of a sphere that generate a free group, which is perhaps the most important detail I leave out of the lecture.

1 Group Theory

A wise person once said, "What are these strange words doing in my math?" That wise person, if you read the table of contents, was probably you. For the sake of stuffing the entire paradox into an hour-long class, I skipped a little. (This may not have been a good idea. The most common feedback you gave me was to make the class an hour longer.) The purpose of this section is to unskip that until-recently-skipped material, which does indeed involve a bunch of strange words in your math.

1.1 Introduction

Speaking of strange words, let's define one!

Definition 1.1. A *group* is a set G with a binary operation $\circ\colon G^2 \to G$ (i.e., it takes two things in G and gives back something in G) satisfying the following conditions:

1. For all $a, b, c \in G$, $(a \circ b) \circ c = a \circ (b \circ c)$. That is, the operation is *associative*.
2. There exists $e \in G$ such that $a \circ e = e \circ a = a$ for all $a \in G$. We call e the *identity*.
3. For all $a \in G$, there exists $a^{-1} \in G$ such that $a \circ a^{-1} = a^{-1} \circ a = e$. We call a^{-1} the *inverse* of a.

Most of the time we write ab instead of $a \circ b$ and leave out parentheses, which don't matter anyway because of associativity.

One way of thinking about groups is as a generalization of number systems. For example, the nonzero rational numbers are a group under multiplication: the operation is associative, the identity is 1, and the inverse of $\frac{a}{b}$ is $\frac{b}{a}$. The integers under addition are also a group: the operation is associative, the identity is 0, and the inverse of a is $-a$. However, groups can look very different from any ordinary number system. For instance, these two groups have *commutative* operations, in which $ab = ba$, but this is not in general true of groups. The ubiquitous example of such a group is the symmetries of an equilateral triangle, which we'll take a look at now.

What are symmetries, and how do they make a group? In this case, we can think of each symmetry as a motion the triangle can go through that doesn't change its image. There are six symmetries of the triangle: rotations about the center in each direction by $\frac{1}{3}$ of a turn, which we'll call ρ_1 and ρ_2; reflections along each line going through a vertex and the center, which we'll call r_X, r_Y, and r_Z; and the identity symmetry that leaves the triangle alone, which we'll call e. These are our six group elements.

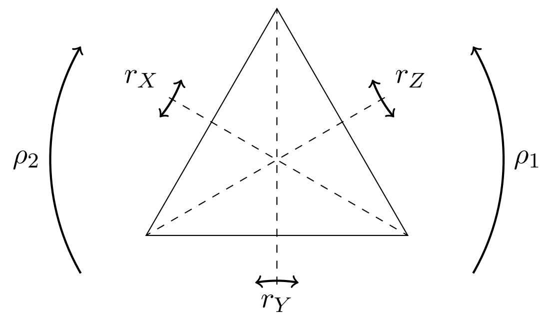

There's a natural way to turn two movements into a third: first do one, then do the other. This is often called "composition". For example, if we first do ρ_1 and then do r_X, we get the same result as if we had done just r_Y. By convention, we write composition backwards to what you might expect[1]—that is, ab means "do b, then do a"—so we have $r_X\rho_1 = r_Y$. However, if we reflect first and then rotate, we find that $\rho_1 r_X = r_Z$, so the group operation is not commutative.

We do this for every pair to make a multiplication table for the group, which we can use to check the operation is associative[2] and has inverses.

	e	ρ_1	ρ_2	r_X	r_Y	r_Z
e	e	ρ_1	ρ_2	r_X	r_Y	r_Z
ρ_1	ρ_1	ρ_2	e	r_Z	r_X	r_Y
ρ_2	ρ_2	e	ρ_1	r_Y	r_Z	r_X
r_X	r_X	r_Y	r_Z	e	ρ_1	ρ_2
r_Y	r_Y	r_Z	r_X	ρ_2	e	ρ_1
r_Z	r_Z	r_X	r_Y	ρ_1	ρ_2	e

There are several patterns to find in the above multiplication table, most of them important in some aspect of group theory. One that particularly jumps out is that the upper-left 3×3 block is self-contained. Put another way, if we multiply rotations together (counting the identity as a rotation

[1] This is because each movement is actually a function that brings each point in the the triangle to its destination, and doing one movement then another is composing functions. We write function composition in reverse because we want $(fg)(x) = f(g(x))$, even though fg is really "first do g, then do f". Some people write gf instead, but they typically also write something like xf instead of $f(x)$ for function application.

[2] Even better—"better", of course, meaning "far less work"—our group operation is, as mentioned above, function composition, so we get that the operation is associative for free. Function composition is associative because $(h(gf))(x) = h(g(f(x))) = ((hg)f)(x)$.

by zero), we get another rotation. Upon closer inspection, we see that the block is a valid multiplication table for a group!

If you think that this might be important, then thousands of mathematicians throughout history have agreed with you. They've even given the concept a proper name.

Definition 1.2. Let G be a group. A *subgroup* of G is a subset $H \subset G$ such that $a^{-1} \in H$ and $ab \in H$ for all $a, b \in H$.

A subgroup is a group in its own right. Notice that every group has the following two subgroups: itself and the trivial group, which has only an identity element. We'll talk more about a specific type of subgroup in a bit.

1.2 Group Actions and Orbits

One way to think about groups is as a way to generalize number systems. However, there isn't an easy way to think about the symmetries of a triangle as a number system[3]. It turns out that it's generally easier to think about groups as the set of symmetries of something. For example, groups can describe the symmetries of a polygon, polyhedron, or tessellation. The symmetries of a circle—that is, all rotations about a given axis—make a group, though unlike the group we just looked at, it has uncountably many elements. This detail may or may not be important later[4].

As with our triangle example, when we talk about symmetries, we're really talking about movements that leave a given object invariant in some way. We could also try applying those same movements to a different object, leaving it different under some movements and the same other others. For example, we could take just vertex X our equilateral triangle. It stays put when we do e or r_X but moves around if we do one of the other movements. The following definition is a way to formalize this concept.

Definition 1.3. A *group action*[5] of group G is a set S and an operation $*\colon G \times S \to S$ (i.e., it takes a group element and something in the set and gives us something in the set) satisfying the following conditions:

1. For all $a, b \in G$ and $x \in S$, $a * (b * x) = ab * x$.
2. For all $x \in S$, $e * x = x$.

[3]On the other hand, the subgroup with just the rotations is isomorphic (i.e., identical in all but name) to the integers mod 3 under addition. If you don't know what that means, Google "modular arithmetic".

[4]Cue scary music!

[5]Specifically, a left group action. Right group actions have an operation $S \times G \to S$ instead of $G \times S \to S$ but are basically the same thing.

We often say that G "acts on" S. One thing this definition makes clear is that every group acts on itself in a pretty obvious way: just let $a * x = ax$ for all $a, x \in G$. (Because G is acting on itself, $S = G$ in this example.)

As before, we will sometimes be lazy and write ax instead of $a * x$, even if S isn't G. We also have a few notational shortcuts for operating over many possible combinations of group elements and set elements.

Definition 1.4. Let G act on S. Given $a \in G$, $A \subset G$, $x \in S$, and $X \subset S$,

$$\begin{aligned} Ax &= \{ax \mid a \in A\} \\ aX &= \{ax \mid x \in X\} \\ AX &= \{ax \mid a \in A, x \in X\}. \end{aligned}$$

We also use this notation for group multiplication, which we can think of as a case where $S = G$. Given $x \in S$, there's one set like this that will be of particular importance to us.

Definition 1.5. Let G act on S. The *orbit* of $x \in S$ is Gx.

In other words, the orbit of a set element is all the other set elements you can get to from it. For example, the group of rotations about a single axis can act on a sphere, and the orbit of a point on the sphere will be the circle of points on the same latitude. Notice that each point is in exactly one orbit. This is because if we do some rotation to get from y to x, meaning x is in the orbit of y, we can do the reverse rotation to get from x to y, meaning y is in the orbit of x. Furthermore, if z is in the orbit of y, the doing both the rotation to get from x to y and from y to z shows us that z is in the orbit of x. This reasoning, as we'll soon show, works for any group and gives us a useful property of orbits[6].

Theorem 1.6. *Let G act on S. The orbits of the action partition S. That is, $\bigcup_{x \in S} Gx = S$, and for all $x, y \in G$, either $Gx = Gy$ or $Gx \cap Gy = \emptyset$.*

Proof. Each element of S is in its some orbit—namely, its own—so the union of all the orbits will contain all of S. That takes care of the first condition.

[6] You may have noticed that we just showed that "being in the orbit of" is an equivalence relation, but our proposition talks about partitions. In general, any equivalence relation yields a partition and vice versa. To see this, consider equivalence classes of a relation $\sim$ on S, sets of the form $\{y \in S \mid x \sim y\}$. The equivalence classes partition S. In the other direction, if we have a partition, defining $x \sim y$ if and only if x and y are in the same part of the partition gives an equivalence relation.

For the second condition, let $x, y \in S$. Suppose that $z \in Gx \cap Gy$. Then there exist $g, h \in G$ such that $gx = z$ and $hy = z$, so $h^{-1}gx = h^{-1}hy = y$. This means $Gy = G(h^{-1}gx)$. Remember that

$$Gh^{-1}g = \{ah^{-1}g \mid a \in G\}.$$

Given any $b \in G$, if we let $a = bg^{-1}h$, then

$$ah^{-1}g = bg^{-1}hh^{-1}g = bg^{-1}g = b.$$

This means $Gh^{-1}g$ contains all of G, so $Gh^{-1}g = G$. This gives us,

$$Gy = G(h^{-1}gx) = (Gh^{-1}g)x = Gx.$$

We didn't show that we can move the parentheses around like that, but it follows pretty simply from Definition 1.4 and associativity. Finally, remember that to get to this, we assumed that $Gx \cap Gy$ had at least one element z, but it could instead be the empty set. Therefore, either $Gx = Gy$ or $Gx \cap Gy = \emptyset$. $\square$

There are lots of other cool things to discover about orbits[7], but this is the result we care most about.

1.3 Generators

Let's consider the group $\mathbb{Z}^2$ of ordered pairs of integers under vector addition (i.e., $(p, q) + (r, s) = (p + r, q + s)$). How can we find a subgroup H of $\mathbb{Z}^2$?

Let's suppose we want $(p, q) \in H$. We know that a subgroup has to have inverses, so $(-p, -q) \in H$. Furthermore, the subgroup has to be closed under the group operation, so $(p, q) + (p, q) = (2p, 2q) \in H$. Similarly, $(3p, 3q)$, $(4p, 4q)$, $(-2p, -2q)$, and so on must all be in H. That is, $(np, nq) \in H$ for all integers n. Notice that if we add two vectors of this form, we get another of this form, so nothing else has to be in H. In other words, $\{(np, nq) \mid n \in \mathbb{Z}\}$ is the smallest possible subgroup that contains (p, q). If we wanted to include another pair (r, s) as well, we'd find that the smallest subgroup doing so is

$$\{(np + mr, nq + ms) \mid n, m \in \mathbb{Z}\}.$$

In general, we can show that the smallest subgroup containing some set of pairs is all possible sums of (possibly negative) multiples of those pairs. This way of constructing subgroups generalizes to other groups, too.

[7] If you think this might be your cup of tea (or different and probably inferior caffeinated beverage), check out orbit-stabilizer theorem. I couldn't find a link that doesn't use terms not defined here, but if you look up what words mean as you come across them, you'll be able to understand it. It uses the same sort of intuition as Theorem 1.6.

Definition 1.7. Let G be a group and $S \subset G$. The subgroup *generated* by S, written $\langle S \rangle$, is the subgroup of of G consisting of all finite products of elements in S and their inverses. (The empty product counts; it's the identity of G.)

Of course, we have to check that this definition actually defines a subgroup, but that isn't too difficult. Let all $s_i \in S$. If $a = s_1 \cdots s_n$ and $b = s_{n+1} \cdots s_m$, then $ab = s_1 \cdots s_m$, and if $a = s_1 \cdots s_n$, then we can check that $a^{-1} = s_n^{-1} \cdots s_1^{-1}$.

Continuing our laziness with notation, we often drop curly braces around a set. For example, we'd write the subgroups of $\mathbb{Z}^2$ we talked about earlier as $\langle (p, q) \rangle$ and $\langle (p, q), (r, s) \rangle$. Note that these generated groups can be a lot more complicated when the group operation is not necessarily commutative.

We can also use generators to define groups from scratch. When we do this, we give a set of generators and a set of rules describing how the generators interact. This is called a group *presentation*, which we write as $\langle S \mid R \rangle$, where S is the set of generators and R the set of rules. For example,

$$\langle \rho_1, r_X \mid \rho_1^3 = e, \rho_1 r_X = r_X \rho_1^2 \rangle$$

gives the group of the symmetries of an equilateral triangle discussed earlier. If our presentation had no rules, then every product of the generators that didn't include an element next to its inverse would be distinct. There's a special name for this special type of group.

Definition 1.8. The *free group* on S is $F_S = \langle S \mid \emptyset \rangle$.

One way to think about the free group is to consider S as alphabet. With this intuition, F_S is the group of all finite words using that alphabet, and the group operation is concatenation (i.e., sticking end-to-end) of words. There are a few ways to think about inverses in this way. We could consider them as different letters of the alphabet and add the cancelation of inverses as part of concatenation. The way I like to think of it is that, just as we could have a appear n in a row, we could also have it appear $-n$ times in a row. Sure, writing a letter a negative number of times is hard to do in real life, but if you still think that everything in math corresponds exactly to things in real life, then you may have missed the part of class where we took a ball, broke it into pieces, spun them around a bit and put them back together to get two copies of the original ball[8]. But if you did miss that bit, that's okay, because we're about to explain it!

[8]Or maybe you think this is possible in real life, in which case every physicist ever wants to give you a stern talking-to about this thing called "conservation of energy".

2 Explanation of the Paradox

In this section, we actually demonstrate the Banach-Tarski Paradox. Isn't that exciting? Of course, if you read the previous section, you may be wondering how all that stuff about groups and actions could possibly relate to anything we did in class, in which case you may also be wondering why you bothered to read all those pages that stuck all those strange words in your math. It turns out, as you may very well notice, that we use concepts from group theory all over the place, though I glossed over the group-theoretic details in class. I'll point them out as we go along.

2.1 The Free Group on Two Generators

We concern ourselves first with F_2, the free group on a set of two elements, which we'll call a and b. We're going to show that, in some sense, F_2 contains two copies of itself.

Let[9]

$$S(x) = \{\text{fully simplified words in } F_2 \text{ that start with } x\},$$

where $x \in \{a, a^{-1}, b, b^{-1}\}$. If $s \in S(x)$, then we can write s as xt with $t \notin S(x^{-1})$, because the second letter of a word starting with x cannot be x^{-1}. This means

$$x^{-1}s = x^{-1}xt = t.$$

In other words, for all $s \in S(x)$, $x^{-1}s \notin S(x^{-1})$, which is the same as saying $x^{-1}s \in F_2 - S(x^{-1})$. Furthermore, for any $t \in F_2 - S(x^{-1})$, there exists $s \in S(x)$ such that $x^{-1}s = t$—namely, $s = xt$. Therefore,

$$x^{-1}S(x) = F_2 - S(x^{-1}).$$

That is, the set of words that start with x when multiplied by x^{-1} on the left gives us the set of words that don't start with x^{-1}. (There's a good picture of this coming up.) With this tool, we're ready to prove the following crucial fact about F_2, which is at the heart of the Banach-Tarski Paradox.

One piece of notation: if sets A and B satisfy $A \cup B = C$ and $A \cap B = \emptyset$, we say that C is the *disjoint union* of A and B and write $A \sqcup B = C$. When we say something is the disjoint union of more than two sets, we require the intersection of every pair of sets we're taking the union of to be empty.

[9] I honestly can't think of a better definition for $S(x)$. That's right, mathematicians sometimes need to know how to read normal English writing! I'm not going to rigorously define exactly what a word is, but it's easy enough to do if you really want to—and likely tedious and unnecessary enough to not be worth your time.

Step 1. *There exist $W, X, Y, Z \subset F_2$ such that*

$$W \sqcup X \sqcup Y \sqcup Z = a^{-1}W \sqcup X = b^{-1}Y \sqcup Z = F_2.$$

Proof. It helps to draw a picture of F_2 as follows: we put e, the empty word, in the middle. From any word, we go right to multiply by a, left to multiply by a^{-1}, up to multiply by b, or down to multiply by b^{-1}, with all of these multiplications happening on the right, at the end of the word.

We know that $x^{-1}S(x) = F_2 - S(x^{-1})$ for all $x \in \{a, a^{-1}, b, b^{-1}\}$, so we can let $W = S(a)$ and $X = S(a^{-1})$. This gives us $a^{-1}W = F_2 - X$. That is, $a^{-1}W$ contains exactly the elements of F_2 not in X, so $a^{-1}W \sqcup X = F_2$.

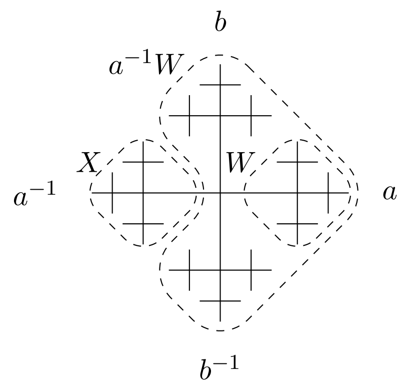

We'd like to define a Y and Z similarly. However, we can't just let $Y = S(b)$ and $Z = S(b^{-1})$, because then we'd leave out e from all of W, X, Y, and Z. Suppose we want $e \in Y$. This means that $b^{-1} \in b^{-1}Y$, so we must have $b^{-1} \notin Z$. But we can't leave out b^{-1}, so $b^{-1} \in Y$. We keep going to show that if $e \in Y$, we need $b^{-n} \in Y$ for all $n \in \mathbb{N}$. This suggests the following definition:

$$\begin{aligned} Y &= S(b) \cup \{b^{-n} \mid n \in \mathbb{N}\} \\ Z &= S(b^{-1}) - \{b^{-(n+1)} \mid n \in \mathbb{N}\}. \end{aligned}$$

(Whether 0 is a natural number is a matter of convention that differs by author and by field of math. Here we say that $0 \in \mathbb{N}$.) This gives us

$$\begin{aligned} b^{-1}Y &= b^{-1}S(b) \cup b^{-1}\{b^{-n} \mid n \in \mathbb{N}\} \\ &= (F_2 - S(b^{-1})) \cup \{b^{-(n+1)} \mid n \in \mathbb{N}\} \\ &= F_2 - (S(b^{-1}) - \{b^{-(n+1)} \mid n \in \mathbb{N}\}) \\ &= F_2 - Z. \end{aligned}$$

We leave confirming $W \sqcup X \sqcup Y \sqcup Z = F_2$ as a mostly harmless exercise for the reader. $\square$

You can probably see in this lemma some resemblance to the paradox itself: we can split F_2 into four pieces, "move" two of them (by multiplying by a^{-1} or b^{-1}), and put them all back together to get two copies of F_2. Our strategy from here is to try to find a way that F_2 acts on a ball, with each element of F_2 corresponding to a rotation. A key word here is "try"[10].

2.2 Equidecomposition

We're trying to prove that we can break something apart, move around the pieces, and put them back together again to make something else. To do this, we need to be slightly more precise than "break something apart, move around the pieces, and put them back together again to make something else". We define what "move around" means first.

Definition 2.1. A *rigid motion* in a space S—which, for our purposes, is $\mathbb{R}^2$ or $\mathbb{R}^3$, i.e. a plane or three-dimensional space—is a finite composition of translations and rotations, which are functions $S \to S$.

Surprise application of group theory number one: rigid motions in S form a group that acts on S. We're not rigorously defining what rigid motions are, so we won't prove this, but it's easy to see intuitively why, for example, inverses exist: to invert a sequence of translations and rotations, do each one backwards in the reverse order.

Definition 2.2. Let $A, B \subset S$, where S is $\mathbb{R}^2$ or $\mathbb{R}^3$. We say that A and B are *equidecomposable* in S if and only if there exist "pieces" $X_1, \dots, X_n \subset S$ and rigid motions $\alpha_1, \dots, \alpha_n$ such that

$$\begin{aligned} A &= X_1 \sqcup \cdots \sqcup X_n \\ B &= \alpha_1 X_1 \sqcup \cdots \sqcup \alpha_n X_n. \end{aligned}$$

We write $A \sim B$.

We use the fact that rigid motions are a group to prove some useful things about equidecomposition.

Proposition 2.3. *Equidecomposability in S is an equivalence relation. That is, for all $A, B, C \subset S$:*
1. *$A \sim A$.*
2. *If $A \sim B$, then $B \sim A$.*
3. *If $A \sim B$ and $B \sim C$, then $A \sim C$.*

[10]More dramatic music!

Proof. Letting $X_1 = A$ and $\alpha_1 = e$, the rigid motion that does nothing, gives $A \sim A$. If $A \sim B$ with pieces $X_1, \dots, X_n$ and rigid motions $\alpha_1, \dots, \alpha_n$, then $B \sim A$ with pieces $\alpha_1 X_1, \dots, \alpha_1 X_n$ and rigid motions $\alpha_1^{-1}, \dots, \alpha_n^{-1}$.

The last condition is slightly tricker. Suppose $A \sim B$ with pieces $X_1, \dots, X_n$ and rigid motions $\alpha_1, \dots, \alpha_n$, and suppose $B \sim C$ with pieces $Y_1, \dots, Y_m$ and rigid motions $\beta_1, \dots, \beta_m$. Intuitively, partitioning B into pieces is like making several cuts. To give a way to decompose A into C, we're going to make the cuts in B implied by the X_is and the cuts implied by the Y_js at the same time. Let

$$\begin{aligned} Z_{i,j} &= X_i \cap \alpha_i^{-1} Y_j \\ \gamma_{i,j} &= \beta_j \alpha_i. \end{aligned}$$

We use the notation $\bigsqcup_i$ to mean the disjoint union over all i from 1 to n, $\bigsqcup_j$ for the same over all j from 1 to m, and $\bigsqcup_{i,j}$ for the same over all i and j in those ranges. Because $\alpha_i X_i \subset B$, we have

$$X_i = \alpha_i^{-1} \alpha_i X_i = \alpha_i^{-1} (\alpha_i X_i \cap B) = \alpha_i^{-1} \bigsqcup_j (\alpha_i X_i \cap Y_j) = \bigsqcup_j Z_{i,j}.$$

By similar reasoning, $Y_j = \bigsqcup_j \alpha_i Z_{i,j}$. (We skip the proof of some simple facts, such as group actions applied to a set distributing over disjoint union and intersection.) This means

$$A = \bigsqcup_i X_i = \bigsqcup_{i,j} Z_{i,j}.$$

What we did to get the $Z_{i,j}$s was, roughly, to transfer the cuts the Y_js make in B back to A by applying α_i^{-1}s. All we showed was that, even though we cut up A into more pieces, those pieces still make A when stuck together. The reason we want all these pieces is to do this:

$$C = \bigsqcup_j \beta_j Y_j = \bigsqcup_j \beta_j \bigsqcup_i \alpha_i Z_{i,j} = \bigsqcup_{i,j} \gamma_{i,j} Z_{i,j},$$

which means $A \sim C$. $\square$

Somewhat hidden in this proof are the required properties for group operations and group actions. We used the existence of an identity to prove $A \sim A$, inverses to prove $A \sim B$ implies $B \sim A$, and associativity to prove $A \sim B$ and $B \sim C$ implies $A \sim C$.

We can now state our goal precisely. Let $p \in \mathbb{R}_3$, and let

$$B_p = \{x \in \mathbb{R}^3 \mid \|x - p\| \leq 1\}$$

be the ball of radius 1 centered around p. (If you haven't seen this sort of notation with vectors, all you need to know is that $\|x - p\|$ is the distance from x to p.) We want to show $B_p \sim B_p \sqcup B_q$, where $q \in \mathbb{R}^3$ and $\|p - q\| > 2$.

2.3 Poking Holes in Spheres

We examine the following two spheres:

$$S^1 = \{(x, y) \in \mathbb{R}^2 \mid x^2 + y^2 = 1\}$$
$$S^2 = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 = 1\}.$$

You may know S^1 by the more familiar name "circle", but it's the analogue of a sphere on a plane[11].

We're going to nonchalantly prove some results about S^1 and S^2, pretending that they're just warmups[12].

Step 2. *$S^1 \sim S^1 - \{x\}$, where $x \in S^1$.*

Proof. Let ρ be a rotation about the origin. Because S^1 is centered at the origin, the rotation brings points in S^1 to other points in S^1. Put another way, $\rho S^1 = S^1$. Let

$$P = \{\rho^n x \mid n \in \mathbb{N}\},$$

and notice that

$$\rho P = \{\rho^{n+1} x \mid n \in \mathbb{N}\}.$$

Let $Q = S^1 - P$, which means $S^1 = P \sqcup Q$. If $x \notin \rho P$, then we're done, because $S^1 - \{x\} = \rho P \sqcup Q$. Therefore, to complete the proof, it suffices to show that there exists ρ such that $x \neq \rho^{n+1} x$ for all $n \in \mathbb{N}$. Call such a rotation "good", and call a rotation that isn't good "bad".

If $x = \rho^{n+1} x$, then ρ must be a rotation by $\frac{2\pi k}{n+1}$ radians, where $k \in \mathbb{N}$. That is, every bad rotation corresponds to a pair $(k, n) \in \mathbb{N}^2$. Because $\mathbb{N}^2$ is countable, there are at most countably many bad rotations. However, there are uncountably many rotations—any real number can be the angle of rotation—so there must exist a good rotation. $\square$

[11]The superscripts in S^1 and S^2 don't have anything to do with exponentiation; they're just parts of the names.

[12]I realize now that I've foreshadowed so obtusely at upcoming surprises that the biggest surprise would be to have no surprise. But I could still do that, so stay on your toes.

We could outline the proof of the next warmup with two words: "same thing". However, we can afford a little more detail than that.

Step 3. *$S^1 \sim S^1 - X$, where $X \subset S^1$ is countable.*

Proof. Let ρ be a rotation about the origin, and let

$$P = \{\rho^n x \mid n \in \mathbb{N}, x \in X\}.$$

It suffices to show that there exists ρ such that $x \neq \rho^{n+1}y$ for all $x, y \in X$, because if so, then $\rho P = P - X$ and we complete the proof in the same fashion as Step 2 above. Again, call such a rotation "good", and call a rotation that isn't good "bad".

Let $x, y \in X$, and let θ be the angle of the arc between them. If $x = \rho^{n+1}y$, then ρ rotates by $\frac{\theta+2\pi k}{n+1}$ radians for some $k \in \mathbb{N}$. Therefore, each bad rotation corresponds to some tuple $(x, y, n, k) \in X^2 \times \mathbb{N}^2$. The product of countable sets is countable, so there are only countably many bad rotations, which means there must exist a good one. $\square$

The temptation to write "same thing" in the next warmup is even stronger than it was for the previous. We almost completely succumb.

Step 4. *$S^2 \sim S^2 - X$, where $X \subset S^2$ is countable.*

Proof. Because there are only countably many points in X, we can pick an axis of rotation through the origin that doesn't go through a point in X. Consider the longitude L_x of each point $x \in S^2$, which is the semicircle containing x that intersects the axis of rotation twice. Note that each point is not on the axis of rotation, so for all $x, y \in X$, either $L_x = L_y$ or $L_x \cap L_y = \emptyset$.

As we showed while proving Step 3, there exists a "good" rotation ρ about the chosen axis of rotation such that for all $x, y \in X$ and all $n \in \mathbb{N}$,

$$L_x \cap \rho^{n+1} L_y = \emptyset,$$

This is because the position of each longitude can be represented by the point the longitude intersects a circle of latitude, such as the equator. Looking at only those intersection points gives us the exact same situation as in Step 3. Finally, note that $\rho^{n+1}L_y = L_{\rho^{n+1}y}$—it doesn't matter if we find the longitude or rotate first—so $x \neq \rho^{n+1}y$. We define P as in Step 3, observe that $\rho P = P - X$, and complete the proof in the usual way[13]. $\square$

All right. I think we're ready.

[13]This, for those of you who might complain that I didn't use the two magic words, is a slightly classier way of saying "same thing".

2.4 The Paradox Itself

At the end of Subsection 2.2, we phrased exactly what we were trying to prove as $B_p \sim B_p \sqcup B_q$, where B_p is the ball of radius 1 centered at p and p and q are far enough apart that the balls around them. We're not going to worry about the translation; it will be easy to add it into the proofs where it should be, if you'd like to do so. As such, we'll write what we're going to prove as $B \sim 2 \times B$, where B is the ball of radius 1 centered at the origin.

We don't want to deal with all the guts inside the ball just yet, so we'll start by talking about S^2. This is probably the step with the most meat.

Step 5. $S^2 \sim 2 \times S^2$.

Proof. By Theorem 3.5 (who sends greetings from the future), there exist two rotations a, b about axes going through the origin such that $\langle a, b \rangle$ is[14] F_2. Recall that for $x \in S^2$, F_2x is the orbit of x—that is, all points x can reach by applying rotations in F_2—and that, by Theorem 1.6,

$$\mathcal{O} = \{F_2x \mid x \in S^2\},$$

the set of all orbits, partitions S^2. Because F_2 is countable[15], each orbit is countable. However, S^2 is uncountable, so $\mathcal{O}$ must be uncountable[16].

Suppose temporarily that we just had to deal with one orbit. Pick some point t in the orbit to be an "anchor". Suppose that for all $f, g \in F_2$, $ft \neq gt$ unless $f = g$. This means that for each $x \in F_2t$, there is a unique $f \in F_2$ such that $ft = x$. By Step 1, there exist $W, X, Y, Z \subseteq F_2$ such that

$$W \sqcup X \sqcup Y \sqcup Z = a^{-1}W \sqcup X = b^{-1}Y \sqcup Z = F_2.$$

Because each $x \in F_2t$ corresponds with exactly one $f \in F_2$, we get

$$F_2t = Wt \sqcup Xt \sqcup Yt \sqcup Zt$$
$$2 \times F_2t = a^{-1}Wt \sqcup Xt \sqcup b^{-1}Yt \sqcup Zt,$$

so $F_2t \sim 2 \times F_2t$, a promising halfway result.

We're going to try to do this for all the orbits at the same time. However, there's a problem: there might some orbits where, after picking an anchor

[14]Technically, we should say "is isomorphic to", which more or less means that things could have different names but are still pretty much the same.

[15]It's the union of words of length n for all $n \in \mathbb{N}$, a countable union of countable sets.

[16]If $\mathcal{O}$ were countable, then expressing S^2 as the union of all the orbits would show it to be a countable union of countable sets, which would be countable.

t, there exists $f \in F_2$ with $ft = t$ but $f \neq e$, the identity[17]. If this is the case, then for any $x \in F_2 t$, which we can write as $x = gt$ for some $g \in F_2$, we have

$$gfg^{-1}x = g^{-1}fg^{-1}gt = gft = gt = x.$$

That is, if $ft = t$, then for every $x \in F_2 t$, there exists $\rho \in F_2$ such that $\rho x = x$. Call a point $x \in S^2$ "bad" if it satisfies $\rho x = x$ for some $\rho \in \mathbb{F}_2$. Remember that elements of F_2 act on S^2 as rotations. There are only two bad points associated with each rotation: those that lie where the axis of rotation intersects the sphere. We conclude two things from this reasoning:

1. If a point is bad, then all points in its orbit are bad.
2. There are only countably many bad points in S^2.

Let $S' \subset S^2$ be the points in S^2 that are "good" (i.e., not bad), and let

$$\mathcal{O}' = \{F_2 x \mid x \in S'\}.$$

By the first item above, each orbit in $\mathcal{O}'$ is a subset of S', because each orbit is either all good or all bad. By the second item, $S^2 \sim S'$, because we can remove the countable subset of bad points with Step 4. This means that all we have to show is $S' \sim 2 \times S'$, because then

$$S^2 \sim S' \sim 2 \times S' \sim 2 \times S.$$

Remembering how we showed that $F_2 t \sim 2 \times F_2 t$ for all $t \in S'$ (because the condition we asked t to satisfy is, in fact, the same as being good), we might want to find $T \subset S'$ such that $F_2 T = S'$ and do something similar. The key before was that for each $x \in F_2 t$, there was a unique $f \in F_2$ such that $ft = x$. The analogous condition here is that for every $x \in S'$, there is a unique $f \in F_2$ such that $x \in fT$. This can't be satisfied if T contains two or zero points in $F_2 x$. Therefore, we define

$$T = \{t_\Omega \mid \Omega \in \mathcal{O}'\},$$

where $t_\Omega \in \Omega$ for all $\Omega \in \mathcal{O}'$. We can think of T as a set containing exactly one anchor for each orbit. For all $x \in S'$, there is a unique $f \in F_2$ such that $x \in fT$ because there is a unique $f \in F_2$ such that $x = ft_{F_2 x}$. That is, each point's orbit is represented exactly once in T, and because only good points are left, there's only one rotation bringing the anchor of the point's orbit to

[17] Earlier, we required that $ft \neq gt$ for all distinct $f, g \in F_2$, but we only really need to require that $ft \neq t$ for all $f \in F_2 - \{e\}$. This is because if $ft = gt$, then $g^{-1}ft = t$.

the point itself. This means

$$F_2T = WT \sqcup XT \sqcup YT \sqcup ZT$$
$$2 \times F_2T = a^{-1}WT \sqcup XT \sqcup b^{-1}YT \sqcup ZT$$

and $S' = F_2T$, so $S' \sim 2 \times S'$. □

Notice what we just did to define T. For each orbit, we chose a point in it. The orbits have a nice structure, but they look the same everywhere. That is, given an orbit, we could pick any point to be the anchor, and there isn't a natural way to choose one. We might try something along the lines of picking the point in each orbit closest to a chosen point on the sphere, but it might be that there's a sequence of points in the orbit that gets arbitrarily close to that point without reaching it, meaning there is no closest point[18]. In other words, to construct T, we have to make uncountably many arbitrary choices. It isn't immediately obvious that we're even allowed to do this. There's an axiom of set theory called, appropriately, the Axiom of Choice, that says we can do this sort of thing. However, despite being intuitively an obvious thing to be able to do, repeated choices like this lead to some unsettling results—like this one! Most mathematicians accept the Axiom of Choice, but, like a common allergen, it's the sort of thing you'd put on a mathematical ingredients label in bold letters.

Step 6. *$B - \{0\} \sim 2 \times (B - \{0\})$. (Recall that $0 \in \mathbb{R}^3$ is the origin and $B = \{(x, y, z) \in \mathbb{R}^3 \mid x^2 + y^2 + z^2 \leq 1\}$ is the ball of radius 1 centered at the origin.)*

Proof. We can repeat the proof of Step 5, but wherever we have a point $x \in S^2$, we replace it with the line segment $\{\lambda x \in S^2 \mid \lambda \in (0, 1]\}$. ($\lambda x$ is the result of multiplying each component of x by λ.) These line segments are "half-open", including one boundary, x, but not the other, 0, so we only cover $B - \{0\}$ with this. □

Remember how Step 2 was "just" a warmup?

Step 7 (The Banach-Tarski Paradox)**.** *$B \sim 2 \times B$.*

Proof. Because $\sim$ is an equivalence relation, it suffices to show that $B \sim B - \{0\}$, because then, by Step 6,

$$B \sim B - \{0\} \sim 2 \times (B - \{0\}) \sim 2 \times B.$$

To show this, draw a circle inside B containing 0, then use Step 2. □

[18] This is like how the sequence $1, \frac{1}{2}, \frac{1}{3}, \ldots$ has no minimum.

Oh, dude! We did it! Yay[19]! I drew a picture of what just happened.

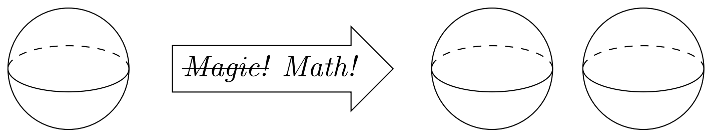

That covers everything we went over in class. If you're interested how we find an embedding of F_2 in the group of rotations of a sphere, read on!

3 Linear Algebra

In this section, we show that there exist rotations a, b such that the subgroup $\langle a, b \rangle$ of the group of all rotations of a sphere is free, which is crucial to Step 5. This requires some background with matrices[20]. For instance, you should know how to add and multiply matrices. You should know that the adding is really easy but that the multiplication is a bit more complicated. You should know that this multiplication is associative but not commutative. You should know that for each n there's an $n \times n$ identity matrix I such that $AI = IA = A$ for all $n \times n$ matrices A. You should know that some but not all $n \times n$ matrices A have an inverse A^{-1} such that $AA^{-1} = A^{-1}A = I$. You should know that this means that the invertible $n \times n$ matrices form a group. You should know that matrix multiplication distributes over addition. You should know that the transpose of a matrix A, which we write $A^\top$, is the matrix flipped over the diagonal. For example,

$$\begin{bmatrix} 1 & 2 & 3 \\ 4 & 5 & 6 \\ 7 & 8 & 9 \end{bmatrix}^\top = \begin{bmatrix} 1 & 4 & 7 \\ 2 & 5 & 8 \\ 3 & 6 & 9 \end{bmatrix}.$$

You should know that $(AB)^{-1} = B^{-1}A^{-1}$ and $(AB)^\top = B^\top A^\top$. You should know that we can take the dot product of two vectors with the same number of entries, and that we can write the dot product of x and y as $x^\top y$. (This gives us a 1×1 matrix whose only entry is $x \cdot y$.) You should know that vectors x and y are orthogonal (also called perpendicular) if and only if $x^\top y = 0$. You should know that we can multiply matrices and vectors by scalars.

[19] Here's a video of my reaction to proving this for the first time.

[20] A willingness to believe whatever I say about them might also suffice. In this case, when I say "you should know that X", instead of reading it as "you should already know that X", think of it as "you should take my word that X".

Finally, you should know (not from prior knowledge, but because I'm telling you) that throughout this section, we write matrices with uppercase letters (and, in one place, lowercase greek letters), vectors with lowercase letters to the right of matrices, and scalars with lowercase letters to the left of matrices. When we write 0, we mean a scalar, vector, or matrix, all of whose entries are 0, with appropriate dimensions for the context.

You will also need to know what "mod 5" means, but this is definitely something you can learn quickly[21].

3.1 Kernels and Images

An $m \times n$ matrix (m rows and n columns[22]) with real entries represents a linear transformation $\mathbb{R}^n \to \mathbb{R}^m$, and a $n \times 1$ matrix represents a vector in $\mathbb{R}^n$. To apply the transformation represented by A to the vector represented by x, we just multiply Ax. Following our tradition of lazy notation, we blur the distinction between matrices and transformations and that between one-column matrices and vectors.

A very common theme throughout algebra[23] is to study the set of elements of a space that a function sends to an identity element. In all of these contexts, this set is called the "kernel" of the function. For our purposes, we define it just for linear transformations, but be aware that there are many related definitions with the same name.

Definition 3.1. The *kernel* of $m \times n$ matrix A is $\{x \in \mathbb{R}^n \mid Ax = 0\}$, which we write as $\ker A$.

Another common notion is to study the "image", or range, of a function. An example of this from earlier is our notation aX for applying one element of a group action to many set elements at once[24].

Definition 3.2. The *image* of $m \times n$ matrix A is $\{Ax \mid x \in \mathbb{R}^n\}$, which we write as $\operatorname{im} A$.

For an $m \times n$ matrix A, which represents a transformation $\mathbb{R}^n \to \mathbb{R}^m$, note that $\ker A \subset \mathbb{R}^n$ and $\operatorname{im} A \subset \mathbb{R}^m$. That is, the kernel contains vectors

[21] I suggested in an earlier footnote that, if you don't know what thas means, you Google "modular arithmetic". This suggestion still stands. The basics are not very hard, but even simple things with it can be really powerful—and useful on math contests.

[22] I have to look up which dimension is which every time.

[23] I mean a significantly more general study of algebra—one that discusses groups, vector spaces, rings, fields, and so on—than classes called "Algebra n" in most high schools.

[24] We defined a group action with a function $G \times S \to S$. We could have also used a definition with a function $G \to \{\text{functions } S \to S\}$. This is an example of currying.

that we apply A to, while the image contains vectors that A has already been applied to.

We haven't defined what a *vector space* is. The quick version is that it is a group with a commutative operation—vector addition—and an extra operation called scalar multiplication which distributes over addition of both vectors and scalars. A *subspace* of a vector space is analogous to a subgroup. It must be closed under vector addition and scalar multiplication. (Inverses are taken care of by multiplying a vector by the scalar -1.) For example, lines and planes through the origin are subspaces of $\mathbb{R}^3$. The kernel and image of a linear transformation[25] are both subspaces.

Definition 3.3. Let S be a subspace of vector space V. The *orthogonal complement* of S is $\{x \in V \mid s^\top x = 0 \text{ for all } s \in S\}$, the set of vectors orthogonal to every vector in s. We write this as $S^\perp$.

For example, a line and a plane that intersect at a right angle are orthogonal complements. One can show that if S is a subspace, then $S^\perp$ is also a subspace. We care about orthogonal complements because of a specific theorem.

Theorem 3.4. *Let A be an $m \times n$ matrix. Then* $\ker A = (\operatorname{im} A^\top)^\perp$.

Proof. Let $x \in \ker A$. By definition, $Ax = 0$, which means that $y^\top Ax = 0$ for all $y \in \mathbb{R}^m$. (Remember that $x \in \mathbb{R}^n$ and $Ax \in \mathbb{R}^m$.) But

$$(A^\top y)^\top x = y^\top Ax = 0,$$

and $A^\top y$ can be anything in $\operatorname{im} A^\top$. Therefore, everything in $\operatorname{im} A^\top$ is orthogonal to everything in $\ker A$. $\square$

Though the proof is quick and easy, this result is a little tricky to understand intuitively. The example I always think of is that of an orthogonal projection. One can show that projecting onto a plane going through the origin in $\mathbb{R}^3$ is a linear transformation represented by a symmetric matrix A, which means $A = A^\top$, in which case the theorem reduces to $\ker A = (\operatorname{im} A)^\perp$. The image of A is the plane it projects onto. The kernel of A is the set of vectors that get projected to the origin, and because A is an orthogonal projection, these vectors lie on a line going through the origin that's orthogonal to the plane A projects onto. The plane and the line are orthogonal complements.

[25]We never defined what a linear transformation actually is. For our purposes, it's a function that can be represented by matrix multiplication.

This sort of intuition can be helpful, but it's sometimes at least as easy to think about things algebraically, write out the formulas, and watch as stuff magically becomes zero[26]. This is especially likely when working in vector spaces with more than three dimensions and other hard-to-visualize things. We take this approach from here on.

3.2 A Free Group of Rotations

We gave a sketch of what a vector space was in the previous subsection, saying it was a group with an extra operation called scalar multiplication. That scalar multiplication must involve scalars from some number system. Specifically, our scalars when working with a particular vector space all come from the same *field*, which is a set that is simultaneously two groups: one under addition, and one under multiplication (excluding the additive identity). We require both operations to be commutative and that multiplication distribute over addition. For example, the rational numbers, real numbers, and complex numbers are all fields, as are the integers mod p for prime p. In the previous subsection, we only considered the field $\mathbb{R}$, but a second look at our definitions and proof of Theorem 3.4 finds that nowhere did we use the fact that the field we were working with was $\mathbb{R}$. We work in the next proof with the field $\mathbb{F}_5$, the integers mod 5, and Theorem 3.4 will be crucial.

We're not going to prove this here, but the following matrices are rotations in $\mathbb{R}^3$, each about one of the coordinate axes:

$$A = \frac{1}{5}\begin{bmatrix} 5 & 0 & 0 \\ 0 & 4 & -3 \\ 0 & 3 & 4 \end{bmatrix} \qquad A^{-1} = A^\top = \frac{1}{5}\begin{bmatrix} 5 & 0 & 0 \\ 0 & 4 & 3 \\ 0 & -3 & 4 \end{bmatrix}$$

$$B = \frac{1}{5}\begin{bmatrix} 4 & -3 & 0 \\ 3 & 4 & 0 \\ 0 & 0 & 5 \end{bmatrix} \qquad B^{-1} = B^\top = \frac{1}{5}\begin{bmatrix} 4 & 3 & 0 \\ -3 & 4 & 0 \\ 0 & 0 & 5 \end{bmatrix}.$$

In fact, all matrices with inverse equal to their transpose represent transformations that are either rotations or reflections. We can use the *determinant*[27] of the matrix to figure out which one it is. Specifically, determinant

[26] Sometimes, like in the following section, what you want to show requires stuff to magically become nonzero. This is also a frequent occurrence.

[27] This is a quantity that, among other things, is used to calculate inverses of matrices. One way to think about the determinant of a transformation is as the ratio of the volumes of a shape after and before a transformation. For example, the determinant of a transformation that doubles one coordinate would be 2.

1 indicates a rotation, and determinant -1 indicates a reflection[28].

Theorem 3.5. *$\langle A, B\rangle$ is*[29] *F_2, the free group on two generators.*

Proof. We want to show that there is no product of As, Bs and their inverses that equals the identity matrix I but doesn't contain a matrix next to its inverse. We'll do this by showing that if such a product exists, then we conclude something false, which means that no such product exists.

To be precise, suppose that there exist $M_1, \dots, M_n \in \{A, A^\top, B, B^\top\}$ with $n > 0$ satisfying:

1. $M_1 \cdots M_n = I$.
2. $M_i^\top \neq M_{i+1}$ for all $i \in \{1, \dots, n-1\}$.

If we replace each M_i with $5M_i$, we get that $(5M_1)\cdots(5M_n) = 5^n I$. Note that each of these matrices has integer entries because $5A$ and $5B$ do.

Let μ_i be $5M_i$, but interpret the entries as elements of $\mathbb{F}_5$, the integers mod 5. Define α and β similarly from $5A$ and $5B$, respectively. Interpreting the entries as elements of $\mathbb{F}_5$, $5^n I$ becomes 0, because each entry of it is a multiple of 5. Therefore,

$$\mu_1 \cdots \mu_n x = 0.$$

for all vectors $x \in \mathbb{F}_5^3$. There must be a minimum $n > 0$ for which this can happen, so we work with that n. This means there exists $x \in \mathbb{F}_5^3$ such that

$$\mu_2 \cdots \mu_n x \neq 0.$$

Let $y = \mu_2 \cdots \mu_n x$. Notice that $y \in \operatorname{im} \mu_2$ and $y \in \ker \mu_1$. By Theorem 3.4, $\ker \mu_1 = (\operatorname{im} \mu_1^\top)^\perp$. Remember that $\mu_1^\top \neq \mu_2$.

It seems like we need to know the images of α, β, and their transposes. We do this by multiplying each matrix with an arbitrary vector. For α, we calculate

$$\begin{bmatrix} 0 & 0 & 0 \\ 0 & 4 & -3 \\ 0 & 3 & 4 \end{bmatrix} \begin{bmatrix} a \\ b \\ c \end{bmatrix} = \begin{bmatrix} 0 \\ 4b - 3c \\ 3b + 4c \end{bmatrix} = \begin{bmatrix} 0 \\ 4b + 2c \\ -2b - c \end{bmatrix} = (2b + c) \begin{bmatrix} 0 \\ 2 \\ -1 \end{bmatrix}.$$

That is, $\operatorname{im} \alpha$ consists entirely of scalar multiples of $(0, 2, -1)^\top$. (This is just a way to compactly write one-column matrices.) We could make $2b + c$ be

[28] This is because reflections flip orientation, which, in one way of thinking, makes the volume of the shape negative. Note that not all matrices with determinant 1 or -1 are necessarily rotations or reflections; we also need to have the inverse equal the transpose.

[29] As mentioned previously, we really should say "is isomorphic to".

any scalar in $\mathbb{F}_5$. This and similar calculations for the other matrices yield

$$\begin{aligned}
\operatorname{im}\alpha &= \{kv_1 \mid k \in \mathbb{F}_5\} &\quad& \text{where } v_1 = (0,2,-1)^\top \\
\operatorname{im}\alpha^\top &= \{kv_2 \mid k \in \mathbb{F}_5\} && \text{where } v_2 = (0,2,1)^\top \\
\operatorname{im}\beta &= \{kv_3 \mid k \in \mathbb{F}_5\} && \text{where } v_3 = (2,-1,0)^\top \\
\operatorname{im}\beta^\top &= \{kv_4 \mid k \in \mathbb{F}_5\} && \text{where } v_4 = (2,1,0)^\top.
\end{aligned}$$

Consider y once more. We know that $y \in \operatorname{im}\mu_2$ and $y \in (\operatorname{im}\mu_1^\top)^\perp$, where $\mu_1^\top \neq \mu_2$, so there must exist v_i and v_j with $i \neq j$ such that $v_i^\top v_j = 0$ (where v_i spans $\operatorname{im}\mu_1^\top$ and v_j spans $\operatorname{im}\mu_2$). But

$$v_1^\top v_2 = 3 \quad v_1^\top v_3 = 3 \quad v_1^\top v_4 = 2 \quad v_2^\top v_3 = 3 \quad v_2^\top v_4 = 2 \quad v_3^\top v_4 = 3.$$

Therefore, we can't have $y \in \operatorname{im}\mu_2$ and $y \in (\operatorname{im}\mu_1^\top)^\perp$. This is a contradiction, so our initial assumption that we can have $M_1 \cdots M_n = I$ without adjacent inverse matrices must be wrong, which means $\langle A, B\rangle$ is free. $\square$

As a wise person—actually, in the earliest Looney Tunes, it was a pig, and Porky is admittedly not "wise" by all definitions of the word—once said, "That's all, folks!"