

The Language of Nature

Mathematics and the Laws of Physics

Written By: Michael Enciso

Contents

Preface	1
0.1 Note to the Reader	3
PART 1	3
1 Sets, Functions, Infinities	5
1.1 Introduction	5
1.2 Sets and Subsets	5
1.3 Set Constructions	8
1.4 Functions	9
1.5 Counting and Infinities	11
2 Groups, Subgroups, and Homomorphisms	17
2.1 Introduction	17
2.2 Groups	17
2.3 Subgroups	23
2.4 Abelian Groups	24
2.5 Group Constructions	26
2.6 Homomorphisms	28
2.7 *Equivalence Relations	31
2.8 *Conjugation	33
3 Real Vector Spaces, Linear Maps, Matrices	39
3.1 Introduction	39
3.2 Real Vector Spaces	39
3.3 Linear Combinations of Vectors	44
3.4 Vector Subspaces	46
3.5 Bases and Dimensionality	46
3.6 *Linear Maps	52
3.7 *Matrices	58
4 Complex Numbers and Complex Vector Spaces	69
4.1 Introduction	69
4.2 Complex Numbers	70
4.3 Geometry of Complex Numbers	72
4.4 Complex Conjugates	73
4.5 Complex Vector Spaces	75
4.6 Outro	77

PART 2	78
5 The Classical World	79
5.1 Introduction	79
5.2 Classical Space-time	80
5.3 The (Very) Basic Geometry of Classical Space-time	82
5.4 *Group Properties of Classical Space	87
5.5 *Movement Through Classical Space-Time	88
6 The Relativistic World	93
6.1 Introduction	93
6.2 The Assumptions of Special Relativity	93
6.3 The Consequences of Those Assumptions	95
6.4 Time Dilation, Length Contraction	96
6.5 *The Group Structure of Relativistic Space-time	103
7 The Quantum World	105
7.1 Introduction	105
7.2 The Importance of Measurement	106
7.3 The Loss of Determinacy	108
7.4 The 2-State System	111
7.5 The 2^N -State System	116
7.6 Entanglement	120
7.7 Concluding Remarks	122
Further Reading	124

Preface

This work is meant to achieve two closely related goals. The first is to help high school level mathematics students transition from learning and doing the math that they're likely used to from high school math courses, to learning and doing the math that they'll find in any standard undergraduate math major's curriculum. This transition is a notoriously difficult one, since the modes of reasoning and that type of work one does in the latter is profoundly and often times shockingly different from those of the former. The problem of finding the proper way of transitioning students at this stage of their mathematical careers is also a notoriously difficult problem to solve. It is the goal of this work to provide yet another possible solution.

The second goal of this work—one that is closely related to the first—is to introduce its readers to a wide variety of mathematical structures and physical concepts that likely haven't been discussed in high school math or physics courses. On the purely mathematical side, this goal is closely related to the first, as I'll discuss in the next paragraph. On the physics side, the hopeful achievement of this goal is meant to both prepare its readers for a possible future in physics (since the author of this work is primarily concerned with math's role in physics) as well as to see how mathematics is applied to areas outside of its pure, abstract existence. Undoubtedly, the physics in this work takes a backseat to the math, and when we explore physics here we'll do so primarily from a mathematical point of view. Thus, this work is largely a mathematical one.

As any expert in these fields will immediately notice, this work does not dive too deeply into any of the topics that it covers, and instead gives a sort of ankle-deep tour through lots of different concepts. This is not accidental. And although we don't dive too deeply into our topics, we most certainly do not sacrifice any rigor in our development (at least in the first part of the text). In fact, rigor is precisely what we do put emphasis on, as this is usually the most difficult aspect of transitioning from high school math to college math. The emphasis on proof and rigor at the college level is almost totally absent from a high school curriculum, and the level of abstraction in a college-level math course only adds to the difficulty a high school student faces in this transition.

As it stands, there are two primary solutions to this problem. One is to have students take part in a sort of "Introduction to Proofs" class, which is meant to teach fresh college mathematicians how mathematical proofs work. While some (many?) of these courses are very useful, some also focus only on teaching students a bag of tricks—like proof by induction and proof by contradiction. I believe this method is precisely the opposite of what a fresh college mathematician needs, because the most important thing to learn when making this transition is that mathematics is **not** about a bag of tricks, but rather a mode of thinking. When "Intro to Proofs" classes do not focus on the modes of thinking, then they are only adding to the problem.

The second solution to the "transition problem", as I'll call it, is to simply throw students into the deep end and assume they'll learn how to swim. Namely, just enroll them into a college math course and have them spend countless sleepless nights trying to navigate this foreign abstract land of theorems and proofs

on their own. After all, a college math major is likely a smart and hard working person, so all will be fine. While I'm not trying to say that sleepless nights are not in the future of any mathematician (because they most certainly are), I **am** trying to say that this second option is still not the best. Even when a student is sufficiently well motivated and doesn't decide to drop the subject and take up something else, there still exists the problem that the student may find him/herself feeling inadequate, especially if other students are seemingly making the transition more seamlessly. After all, a genuine college math course will likely have lots of students for whom the course is **not** their first college math course, and will naturally be better prepared for it. Thus, when a bright student who happens to have never taken a math course sees himself working ten times harder than his peers, it is not uncommon for this student to become caught in a cycle of self-doubt.

Therefore, it is important to make clear to students that there **is** a transition in methodology and modes of reasoning that is taking place and that struggling with this transition is completely natural, and arguably inevitable. Moreover, we need to do so in a way that doesn't simply just replace one bag of tricks (from high school) with another (for college). This is precisely what this text aims to do, and it aims to do so by finding a sort of happy medium between the above two "solutions". In particular, we will explore **real** college-level math, like set theory and group theory and abstract vector spaces, and we'll prove **real** theorems about them with complete rigor. This way the student will be able to see first-hand how this sort of math works and will be able to see for herself how the modes of reasoning need to pan out. The only difference between this text and a college-level textbook is that we simply don't do **too much** of any one thing. By covering lots of different topics with complete rigor (and minimal depth), we'll be exposed to many different modes of proof and the thought processes behind them. Moreover, by actually covering the topics that we do with the same level of rigor that is found in college math classes, the methods of proof that are presented do not appear as a "bag of tricks", but rather the necessary methods used for building mathematics from the ground up.

It is also hoped that by moving from topics as fundamental as set theory to the less fundamental topics of complex vector spaces, it will become clear to the reader that mathematics is a **creative** field. The "rules" and "tricks" that one learns in mathematics were not written in stone somewhere—we had to **find** them. Mathematics is about both discovery and creation. It is an art form, where the artist needs to decide what is beautiful and what is interesting, and the artist's medium is abstraction and logic. This is a side of math that is almost never seen in high school, and which will hopefully start to show itself here.

The final thing we'll address in this introductory remark is why there is any physics included at all, if the primary aim of this goal is to solve the **math** transition problem. To answer this I must first address the more basic question of why I should write this text at all. The motivation for writing this text came from an opportunity that I had to teach a math and/or physics course to third and fourth year high school students, where the subject matter that we were to cover was almost entirely for me to decide. The decision to create a class that I thought would help transition the students into their futures was obvious to me, but the methods for doing this were less so. Namely, I asked myself the same question of whether or not physics should be included, and the answer to this question soon became clear.

Being a math and physics student myself, I've seen how cool physics is and I've experienced how my love of each subject influences my love for the other. Some students enjoy pure math for its own sake, and others find varying degrees of pleasure by seeing it applied to other fields. Physics is one such field that pure math can be applied to, and it is the field that I happen to work in. By including it here, I hope to encourage the reader to not only further pursue education in pure math, but also to continue to be open to the possibility (and likelihood) that pure math will be able to affect the world in other ways (whether it be in physics, computer science, economics, engineering, or some new field that has yet to be discov-

ered). Therefore, we include the physics here not only because the subjects are themselves interesting—and mathematical—but also to remind the reader that there is a world outside of pure math that might be just as interesting. To a pure mathematician this may be blasphemous, and for that I apologize, but I also don't apologize **too** much.

0.1 Note to the Reader

Admittedly, my explanations of certain concepts may not be entirely clear upon the first reading. When this occurs, there are two actions that I would encourage. The first is to continue reading even with a somewhat unclear understanding. Often times I'll make definitions and or say certain things and then follow them up with something like "In other words," or "Note in particular that", or "Namely, " or some other potentially clarifying phrase. It is my goal that the sentences following such phrases will make clearer the topics that came before. If upon this extended reading the topics are still not clear, then I would suggest rereading the confusing passage (not too dissimilar to how the reader likely already deals with such situations). It is my goal to both minimize the number of places that are extremely confusing, and also to alert the reader when any of the following might occur: 1) a possible confusion with other topics, 2) a particularly obscure idea emerges, and/or 3) a topic is purposefully not fully developed or explained. If, however, there are passages in which I have not fully met my goal, then I would encourage the reader to alert me of this via email at truebeautyofmath@gmail.com.

The above email address is associated with my website—www.truebeautyofmath.com—which covers a small portion of the topics presented here (and some that aren't) in **much** greater detail, and may be a decent resource for the interested in reader.

Part 1

This text is designed to serve as an introduction both to higher, abstract mathematics as well as to various ideas in modern physics. In the first part, we will focus on presenting some key ideas in pure mathematics as well as developing the very elusive quality of "mathematical sophistication". Depending on the reader's background with abstract, pure mathematics, this type of reasoning may be extremely new. We will therefore lean on the side of "over-explanation" in order to try to make this transition as smooth as possible. After exploring these concepts, we'll see how they're used in various parts of modern physics. The vast majority of what is now considered "modern physics" lies far outside the realm of what is "intuitive", or what we can visualize. Therefore, we must rely on the abstraction and rigor of mathematics to lead us to the "right answer" as we continue to find better and more accurate models of Nature.

As mentioned, we'll begin by introducing the mathematical side of things (and this will take up the majority of these notes). The concepts that we'll introduce are likely foreign, and very different from what one is exposed to in a typical high school curriculum. They are not particularly difficult or complex ideas. Instead, we'll be building mathematics from the ground up, focusing on complete rigor as we do so—mathematics is nothing without its rigor. As we'll see, math is built in such a way that one cannot help but think that it "couldn't be any other way". We hope to get a feel for what this means as we go, and in the process begin to develop the mathematical sophistication mentioned above.

Some statements or concepts may seem extremely obvious or trivial (and some may not). The reason we're sometimes forced to consider and study even the most seemingly obvious ideas is that in order to be completely rigorous, we must not take anything for granted. This means that no truth can go unscrutinized. We must ask ourselves what the most basic object(s) of our understanding is, and attempt to build as much as possible with as few assumptions as possible. In building mathematics, we can't assume that anything is true due to its "obviousness". We must make seemingly obvious or straightforward concepts as mathematically precise as possible. Sometimes ideas that seem obvious are indeed that way, but sometimes there is hidden subtlety and intricacy. It is for this reason that we must firmly and rigorously establish every truth that we uncover.

The general layout of advanced (undergrad-level, grad-level, and beyond) mathematics textbooks and papers is to introduce new concepts in the form of definitions, to introduce new facts about these new concepts in the form of theorems, propositions, claims, and lemmas, and to follow up each theorem, proposition, claim, or lemma with a proof of its truth. We will therefore adopt such a layout, in order to prepare the reader for more advanced texts. However, we will attempt to infuse this layout with words (sometimes lots of them) that don't fall into any of the three categories of definition, theorem, or proof. This will often be to elaborate on the intuitive concepts behind the formal definitions and mathematical ideas, and/or to motivate the study of the objects that we're currently considering, and/or to study a particular example of a given abstract construction. This may be a new and/or foreign layout for a math text to the reader, and we hope that the additional verbiage will help ease the transition process. So without further ado, let us begin.

Chapter 1

Sets, Functions, Infinities

1.1 Introduction

Our starting point in these lectures will be the assumption that "an element" exists. This is by no means a fully rigorous notion, nor is it the final answer to the question of what the fundamentals of mathematics are built on. These questions very quickly get wildly obscure, and they will not help us arrive at where we want to go. Our chosen starting point is, as it seems to me, a healthy mixture of complete rigor and intuitive understanding. It will hopefully highlight the power of abstraction, all while not falling too deep into the rabbit-hole of philosophy and epistemology. So, without further ado, let us begin with mathematics.

1.2 Sets and Subsets

What is the most basic thing that we can think of as existing? What is the most blatantly obvious starting point from which we can build up mathematics? Well, we've already answered that question: anything. Anything at all. We certainly must assume that "things" exist, and that will be our starting point.

Definition 1.1. An **element** is a thing.

That's it. Pretty much as basic as we can get.

Example 1.2. A number is an element, a person is an element, an idea is an element. A color is an element, a computer is an element, a donkey is an element. Hopefully the picture is clear.

Definition 1.3. A **set** is a collection of distinct elements.

Example 1.4. The collection "my dog, this cup of coffee, Kobe Bryant" is a set, and it has three elements. The collection "all dogs in the world" is a set, and it has lots of elements (though not infinitely many). The collection "all positive whole numbers" is a set, and it has infinitely many elements.

Example 1.5. There is one particularly important example of a set, and that is the empty set, which has no elements at all. After all, the collection of "no elements" is a perfectly well-defined collection of elements. We often denote the empty set by \emptyset .

Let us now set up some notation and terminology that will prove to be very useful for us in the future. This will greatly minimize the amount of writing we need to do, and will maintain the essential features of sets and elements.

If a is an element in the set A , we say that " a is in A " and we denote this by $a \in A$. Note that a

can be any element at all, and not literally the letter " a ", and that A can be any set at all. In other words, the letter a just stands for whatever we want it to stand for. So I could say "let a be my cup of coffee, and let A be the set of all things on the table in front of me". Then " $a \in A$ " would be a true statement, and would be read as " a is in A " because "my cup of coffee is in the set of things on the table in front of me". This may seem overly pedantic right now, and is most certainly **not** how we should think about everyday things like cups of coffee, but this notation and terminology will eventually be very useful for us.

We now introduce a better notation for how we specify the elements in our sets. We usually specify a set by putting some condition on the elements that are in it, and grammatically we can use the model sentence (if we let the set be A) " A is the set of elements **such that** some condition holds on those elements". We can write this out more symbolically as

$$A = \{\text{the set of elements} \mid \text{some condition holds}\}$$

where the vertical line is read as "such that".

Example 1.6. We often denote the set of integers (positive and negative whole numbers, as well as zero) as \mathbb{Z} , so that $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$. We can then specify the set of positive integers (often denoted by \mathbb{N}) as $\mathbb{N} = \{a \in \mathbb{Z} \mid a > 0\}$. This reads as " \mathbb{N} is the set of all elements a that are in \mathbb{Z} such that a is also greater than 0". Similarly, we could leave what goes to the left of the vertical line more general, and put all of the conditions on the right side of the vertical line, as follows: $\mathbb{N} = \{a \mid a \in \mathbb{Z} \text{ and } a > 0\}$. This then reads as " \mathbb{N} is the set of all elements a such that a is an integer and such that a is greater than 0".

Exercise 1.7. Using this notation, specify the set of all positive, even numbers. Specify the set of all negative, odd numbers.

Example 1.8. Here's a weirder example. If I let A be the set of all donkeys in the world, then I'd have $A = \{a \mid a \text{ is a donkey}\}$. Note that I could let this set (or any set) be represented by any symbol that I'd like, such as B or C or Ξ , and that I just keep choosing A arbitrarily.

Let us now make another definition, before introducing more notation.

Definition 1.9. Let A be a set. A **subset** of A is a set B such that every element in B is also in A .

Note that, when we adhere strictly to the definition of a subset, we find that any set is a subset of itself, and that the empty set is a subset of every set. We denote the statement " B is a subset of A " by $B \subseteq A$. Thus, the preceding sentence is summarized by $A \subseteq A$ and $\emptyset \subseteq A$ for any set A . If it is known that B is a subset of A , and that A is **strictly larger than** B (i.e., there is at least one element in A that is not in B), then we can denote this by $B \subset A$. Note that the former symbol works just as well in this case, and that we have to be more careful when using the symbol " \subset ". Accordingly, we'll usually just use " \subseteq " even if it is known that we have strict inclusion. Note the similarity between these symbols and the symbols used to represent inequalities between numbers: " \leq " and " $<$ ". As with the inequality symbols, we can also reverse the direction of the set inclusion symbols (to write things like $A \supseteq B$, just as we could write $5 \geq 3$), but this is usually in bad taste, or at least somehow less common (though not wrong).

We now have enough machinery to talk about what it means for two sets to be equal to each other. Clearly, we'd like to consider two sets to be "equal" when they are "exactly the same", i.e., when they have exactly the same elements. We can now make this formal.

Definition 1.10. We say that two sets A and B are **equal**, and write $A = B$, when the following two conditions hold: i) $A \subseteq B$ and ii) $B \subseteq A$.

The first condition says that everything that is in A is also in B , and the second condition says that everything that is in B is also in A , and therefore it should be clear that the only possible way **both** of these conditions can hold is if A and B are the exact same sets. Now, it might seem like this idea of equality is hardly worth its own definition, but recall that we must be as precise and rigorous as possible. We have been able to derive a notion of "equality" using our machinery of subsets, and we can't assume that this notion of equality is already understood, so we therefore must give it its own definition.

Before talking about set constructions in the next subsection, let us describe an important kind of set. Some of the most commonly studied sets in mathematics are sets whose elements are themselves sets. After all, a set is "a thing" in its own right, and we can therefore consider it as a single element, and therefore build sets up out of these elements. So, for example, we could consider the following sets of numbers: $\{1, 2, 3\}$, $\{4, 5, 2\}$, and $\{1, 5, 6\}$. Then we can consider each of these sets as a single element, and form the following set (call it A): $A = \{\{1, 2, 3\}, \{4, 5, 2\}, \{1, 5, 6\}\}$, whose elements are themselves sets. Thus, A actually only has 3 elements, and not 9!

There are some subtleties that come with talking about sets of sets, as the example following the next exercise shows. For now, though, just note that such sets can, and should, exist.

Exercise 1.11. Which of the following sets are equal? $A_1 = \{1, 2, 3, 4, 5\}$, $A_2 = \{\{1, 2\}, 3, 4, 5\}$, $A_3 = \{\{2, 1\}, 4, 5, 3\}$, $A_4 = \{2, 5, 3, 4, 1\}$

Example 1.12. One thing we need to be careful about when defining sets of sets is the notion of "self-referentiality". What this amounts to is that there are some restrictions on which sets we can define. For suppose we considered something like "the set of all sets that contain themselves". We could also consider "the set of all sets that **don't** contain themselves". Most sets don't contain themselves. For example, the set $\{1, 2, 3\}$ doesn't contain itself because it only contains the three elements 1, 2, and 3, and obviously none of these elements equal the set $\{1, 2, 3\}$. But if we consider "the set of all sets", then this set **does** contain itself, because it is a set, and therefore it is an element of itself. Thus, it appears that "the set of all sets that contain themselves" is actually an interesting set (i.e., it is not the empty set).

But now we can ask the following question: Does the set of all sets that don't contain themselves contain itself? I.e., is the set of all sets that don't contain themselves a set that does indeed contain itself as an element, or not? From what we've said so far, there needs to be an answer one way or the other. Namely, we should be able to take any element in the world and any set in the world and say definitively whether or not that set contains that element. Sometimes it will, sometimes it won't, but never neither, and never both. So if we view our set as "the set of all sets that don't contain themselves", and if we view our element as being "the set of all sets that don't contain themselves" (where now we view this set as an individual element), then we should be able to answer definitively whether or not this element is in this set (even though they're the same!).

Let's suppose the set of all sets that don't contain themselves does contain itself as an element. Then it lies in the set of all sets that do contain themselves and not in the set of all sets that don't contain themselves. But since this **is** the set of all sets that don't contain themselves, this means that this set doesn't contain itself, which contradicts the supposition that started this paragraph! Thus it must be the case that this set doesn't contain itself. But if this were the case, then this set would be an element in the set of all sets that don't contain themselves, which is itself! Therefore it does contain itself! Therefore this set is neither in, nor not in, a particular set, which is impossible.

This is known as Russell's paradox and when it was first discovered it led to a huge explosion in the study of set theory, and progress in this field is still made today. The goal is to find axioms for how we define sets that ensure that no such paradoxes can occur, but such that the math that we know and love can still be constructed. If this is an interesting endeavor to the reader, then a career in the foundations of math and/or mathematical logic might be a fruitful one. For now, we'll put these subtleties behind us and

take on faith that these problem won't plague our future development (which is true).

1.3 Set Constructions

We now briefly talk about how we can build up new sets using the data of old sets. The three main constructions we'll discuss are unions, intersections, and Cartesian products. The general theme is to take two sets and form a new, third set from the information that we have from the first two sets. Some of these constructions carry over to the case where we want to build up new sets from many old ones (and not just two), but some of the constructions don't carry over in an obvious way. We'll discuss these in turn.

Definition 1.13. Let A and B be sets. The **union** of A and B , often denoted by $A \cup B$, is the set of elements that are either in A or B or both, so that $A \cup B = \{a | a \in A \text{ or } a \in B\}$.

Example 1.14. Let $A = \{1, 2, 3, 4\}$ and $B = \{5, 6, 7, 8\}$, then $A \cup B = \{1, 2, 3, 4, 5, 6, 7, 8\}$.

Example 1.15. Let $A = \{a, b, \text{DONKEY}\}$ and $B = \{a, 1, 2, 3\}$, then $A \cup B = \{1, 2, 3, a, b, \text{DONKEY}\}$.

Now we move on to our second construction—the intersection.

Definition 1.16. Let A and B be sets. The **intersection** of A and B , often denoted by $A \cap B$, is the set of elements in both A and B , so that $A \cap B = \{a | a \in A \text{ and } a \in B\}$.

Example 1.17. Let $A = \{1, 2, 3, 4, 5\}$ and $B = \{4, 5, 6, 7\}$. Then $A \cap B = \{4, 5\}$.

Example 1.18. Let $A = \{4, 5, 6\}$ and $B = \{1, 2, 3\}$, then $A \cap B = \emptyset$.

Exercise 1.19. Show that for any sets A and B , it is always the case that $A \subseteq A \cup B$ and that $A \cap B \subseteq A$.

The above two constructions were relatively straightforward, but now we're going to introduce a third construction that is noticeably different from these first two. In particular, unions and intersections build new sets out of old ones while keeping the individual elements the same. I.e., the elements that make up unions and intersections are exactly the same as those that make up the original sets, they're just "re-arranged" in some way. In the final construction that we'll introduce, however, we're going to actually be forming new **elements** from the sets that we're initially given. In particular, we're going to take the elements of the two "old sets" and form new elements, and these new elements will be "pairs". One element in each pair will come from one of the original sets, and the other element in each pair will come from the other original set. Thus, we can and should think about the Cartesian product (which is what this construction is called) as the set whose **individual elements** are pairs of elements. Perhaps more intuitively, we can think of each element in the Cartesian product as "a way of picking one element from one set, and one element from the other set". Let us now make the following definition.

Definition 1.20. Let A and B be sets. The **Cartesian product** of A and B , often denoted by $A \times B$, is the set of pairs of elements from A and B , so that $A \times B = \{(a, b) | a \in A, b \in B\}$.

Note that we could write these elements in many different ways, not necessarily in the form " (a, b) ". I.e., we could have written $[a, b]$ or $[(\{a, b\})]$, or whatever, but it would clearly be the case that all of these sets would somehow "be the same", or at least "contain the same information". Thus, we simply choose to use the " (\cdot, \cdot) " notation and notice that the truly important information is that each element contains exactly one element from one set (A in this case) and one element from the other (B in this case).

Exercise 1.21. Let $A = \{1, 2, 3\}$ and $B = \{a, b, c\}$. Write out all of the elements in $A \times B$.

One important difference between our set constructions lies in whether or not we can combine **several** sets. In particular, if we take two sets and construct a third using one of the above methods (union, intersection, or Cartesian product), we're left with a perfectly good set. Therefore, we can take this new set and form the union, intersection, or Cartesian product of this set with another set. The question is, then, whether or not we're left with the same set if we apply these constructions in different orders. Namely, if A, B , and C are three sets, is it the case that $(A \cup B) \cup C = A \cup (B \cup C)$? Is it the case that $(A \cap B) \cap C = A \cap (B \cap C)$? Is it the case that $(A \times B) \times C = A \times (B \times C)$? If the answer to these questions is yes, then we can form several-fold unions, intersections, and/or products, since the order in which we construct these many-fold constructions is irrelevant. The next exercise will answer two of these questions.

Exercise 1.22. Let A, B , and C be sets. Show that $(A \cup B) \cup C = A \cup (B \cup C)$ and that $(A \cap B) \cap C = A \cap (B \cap C)$.

However, it is **not** the case that $(A \times B) \times C = A \times (B \times C)$. This is a somewhat trivial issue, but it's important to note. The reason these two sets aren't equal is that we have to **change** the elements when forming the Cartesian product of two sets. Thus, on the left hand side we have elements that look like $((a, b), c)$, whereas on the right hand side we have elements that look like $(a, (b, c))$. This may look rather harmless, but strictly speaking these are different kinds of elements. Now, it happens to be the case that these two sets are **effectively** equivalent, as may be clear, but they're not **strictly** equivalent because their elements are different. It turns out that we can still form many-fold products due to the "effective equivalence" that we just mentioned, but to see how to make this rigorous, we need to learn what a bijective function is, and for this we need to press on to the next section.

This will end our brief discussion of set constructions. We'll get much more use out of these constructions as we move on, and the ideas are relatively straightforward, so we'll save their more detailed study for the future.

1.4 Functions

Once we have some sets lying around, it's nice to be able to "relate" them to each other, and that's exactly what we turn our attention to now. Our weapon of choice is a "function", and it is quite possibly one of if not the most important notions in all of mathematics.

Definition 1.23. A **function** from a set A to a set B is an assignment to every element in A some element in B .

We can think of functions as being "dynamical", in that they "send" elements in one set to elements in another set. It need not be the case that every element in B is "hit" by something in A , but it must be the case that each element in A is "sent" somewhere, and that each element in A is "sent" only to one element in B . This analogy should only be taken so far, though, because clearly there's nothing really "moving" out there in the world of abstract mathematics. Instead, we're just setting up an association between the elements of two sets, in a particular way. The dynamical nature of a function is a helpful analogy, however, and is the basic analogy which lies at the heart of how mathematics models the real world. This is because the real world is most certainly dynamical, and we often view dynamical phenomena as "functions of time". We'll see this more clearly when we turn to the physics side of things later on. For now, just know that the analogy is there.

We often denote the statement " f is a function from A to B " by $f : A \rightarrow B$. Thus, when we write "let $f : A \rightarrow B$ ", we should read this as "let f be a function from A to B ", and we must always be sure that the symbolic expressions that we write down are always grammatically correct. Believe it or not, grammar

is actually quite important in mathematics! Using our above analogy in which functions are "dynamical", we denote the idea that " f sends $a \in A$ to $b \in B$ " by $f(a) = b$.

The next thing we might like to know is what to do when we have multiple functions around. Namely, how do we make new functions out of two old ones? We can think of this as making "function constructions", and the next definition partially answers this question.

Definition 1.24. Let $f : A \rightarrow B$ and let $g : B \rightarrow C$. Then the **composition** of f and g is the function $g \circ f : A \rightarrow C$ that maps $a \in A$ to $g(f(a)) \in C$.

I.e., the **composition** of two functions is the function that "does the first function first, and the second function second". Note that when the sets A , B , and C are all different, there is only one order in which the functions can be applied. Namely, the "second" function always needs to "start from" wherever the "first" function ended. For example, in the above definition, the expression $g(f(a))$ is well-defined because it takes an element from A to B and then from B to C , but there is no expression along the lines $f(g(a))$, because if we apply g first then we're starting at B and ending at C , whereas f starts at A .

Example 1.25. Let $A = \{a, b, c\}$, $B = \{1, 2, 3, 4\}$, and $C = \{this, that\}$. Also, let $f : A \rightarrow B$ be the function that has $f(a) = 1$, $f(b) = 2$, and $f(c) = 1$. Finally, let $g : B \rightarrow C$ be the function that has $g(1) = this$, $g(2) = that$, $g(3) = this$, and $g(4) = this$. Then the function $g \circ f : A \rightarrow C$ is such that $g \circ f(a) = g(f(a)) = this$, $g \circ f(b) = g(f(b)) = that$, and $g \circ f(c) = g(f(c)) = this$.

Let us now make some more definitions. These won't be particularly deep, but they'll give us some useful vocabulary for dealing with these concepts so that we don't have to keep using words like "start" and "finish" in reference to where functions go "from" and "to".

Definition 1.26. Let $f : A \rightarrow B$ (note that this should be read in the only grammatically correct way possible). Then A is called the **domain** of f , and B is called the **codomain** of f .

That's it, just words. Now, however, we can consider something slightly more subtle. Namely, we see that "the set of points 'hit' by f " is not necessarily the same as the set B , simply because it's not required in the definition of a function that the function "hits" everything in B . Thus, we might want some terminology to distinguish the two ideas. The following definition provides that for us. First, we need more notation. The symbol " \exists " should be read as "there exists" and the symbol " \forall " should be read as "for all". Remember, don't be scared, it's just notation!

Definition 1.27. Let $f : A \rightarrow B$. The set $\{b \in B \mid \exists a \in A \text{ such that } f(a) = b\}$ (this is read "the set of elements in B such that there exists an element a in A such that a is sent to b by f ") is called the **image** of f , and is often denoted by $Im(f)$.

Now, we can naturally ask when it is the case that $Im(f) = B$, i.e., when the function "hits" everything in its codomain. We thus make the following definition:

Definition 1.28. Let $f : A \rightarrow B$. If $Im(f) = B$, then f is called **surjective**.

Thus, a surjective function is one in which every element in the codomain is hit by something in the domain.

Example 1.29. Let $A = \{1, 2, 3, 4, 5\}$ and $B = \{a, b, c\}$. Then the function $f : A \rightarrow B$ that has $f(1) = a$, $f(2) = a$, $f(3) = a$, $f(4) = a$, and $f(5) = b$ is not surjective, because c is never "hit" by the function f , but the function $g : A \rightarrow B$ such that $g(1) = a$, $g(2) = c$, $g(3) = a$, $g(4) = b$, and $g(5) = c$ is surjective. Note that it is impossible to find a surjective function $f : B \rightarrow A$. Can you see why?

Another natural question to ask about a function is whether or not it is the case that each $a \in A$ is sent to a different $b \in B$. A general function is not required to have this property, but it will prove useful at times to consider those functions that do. We'll therefore give these functions a name, but first we need some more notation.

If P and Q are statements, we write " $P \Rightarrow Q$ " to mean " P implies Q ". Thus, for example, an expression like " $A = B \Rightarrow C = D$ " means that " $A = B$ implies $B = C$ ", so that if we know that $A = B$, then we automatically know that $C = D$ also.

Definition 1.30. A function $f : A \rightarrow B$ is said to be **injective** if it is the case that $f(a) = f(b) \Rightarrow a = b$.

Note that this is really just an abstract way of writing "each element in A is sent to a different element in B ", because the definition states that if two elements (a and b) in A are sent to the **same** element in B ($f(a)$ and $f(b)$), then they were actually the same element to begin with!

Example 1.31. Let $A = \{1, 2, 3, 4\}$ and $B = \{a, b, c, d, e\}$, and let $f : A \rightarrow B$ be the function that has $f(1) = b, f(2) = c, f(3) = a$, and $f(4) = e$. Then f is injective. However, if $g : A \rightarrow B$ is the function that has $g(1) = a, g(2) = a, g(3) = c$, and $g(4) = e$, then g is not injective. This is because 1 and 2 both get sent to the same place in B . I.e., it is not true that for all elements $x, y \in A$, $g(x) = g(y) \Rightarrow x = y$, because it is obviously the case that even though $g(1) = g(2)$, $1 \neq 2$. Note that it is impossible to find an injective function $B \rightarrow A$ in this case. Can you see why?

We now have the following extremely important definition, as a natural extension of the previous two definitions:

Definition 1.32. A function is said to be **bijective** if it is both injective and surjective.

We sometimes call a bijective function a "one-to-one correspondence" and/or a bijection.

Example 1.33. Let $A = \{1, 2, 3, 4\}$ and $B = \{a, b, c, d\}$. Then the function $f : A \rightarrow B$ that has $f(1) = c, f(2) = b, f(3) = d$, and $f(4) = a$ is bijective.

Exercise 1.34. Find at least three other bijective functions $A \rightarrow B$, with A and B as they are in the above example.

Exercise 1.35. Show that if there is a bijective function $f : A \rightarrow B$, then there exists a bijective function $g : B \rightarrow A$. Show also that this is not the case for functions that are either **only** injective, or **only** surjective.

Exercise 1.36. We now put to rest the issue of $A \times (B \times C) \neq (A \times B) \times C$ that we saw above. Namely, even though these sets aren't strictly equal, they're always in a bijective correspondence with each other, and can therefore be viewed as equivalent for all practical purposes. So, let A, B , and C be three sets. Show that there exists a bijective function $f : A \times (B \times C) \rightarrow (A \times B) \times C$.

The idea of a bijection also gives us an abstract way to count, and a precise way to deal with (the various kinds of infinities), as we'll see in the next section.

1.5 Counting and Infinities

Now that we know about bijective functions, we can define "counting" in an abstract and rigorous way. I.e., we can make precise the notion of "how many elements a set has".

Namely, if the enemy hands you a set like $A = \{a, b, c, \text{apple}\}$, which has finitely many elements, we

want an abstract way of defining "how many" elements the set has. Clearly, A has 4 elements. But how do we KNOW that? Well, we just count them! But what are we "really" doing when we count? How do we make this logically and mathematically rigorous?

Let's first define the following sets: $S_i = \{1, 2, \dots, i\}$ where i is some non-negative whole number. Thus, $S_1 = \{1\}$, $S_2 = \{1, 2\}$, $S_5 = \{1, 2, 3, 4, 5\}$, and $S_0 = \emptyset$. Now, we can easily see that the set A above can be put into a one-to-one correspondence with S_4 . I.e., there exists a bijective function $f : A \rightarrow S_4$. In fact, there exist many bijective functions between these two sets, but that doesn't matter—all that matters is that there **is** such a function. Note that since A can be bijectively mapped to S_4 , it cannot be bijectively mapped to any S_i such that $i \neq 4$. A little bit of playing around with this and trying to define a bijective function to other S_i 's should be enough to convince the reader of this fact. Thus, the following definition of **cardinality** is well-defined.

Definition 1.37. Let A be a set with finitely many elements. We say that A has **cardinality** N if there exists a bijective function $f : A \rightarrow S_N$.

Thus, cardinality is nothing but a way of saying how many elements a set has. Note that we make no assumptions about how many elements A has, simply because the notion of "how many" is not well-defined until now. I.e., we are **defining** the notion of "how many" to be "which S_i can it be put into bijection with?". But note that we have to specify that A has finitely elements in the definition of cardinality, because clearly a set with infinitely many elements cannot be put into bijection with any of the S_i . But, if we're not supposed to know what "how many" even means before this definition, how can we know that there are only finitely many elements in A ? Well, we can simply **define** a set to have finitely many elements if there exists some N such that such a bijection exists. Moreover, we can **define** a set to have infinitely many elements if there is no N such that a bijection can exist.

And that's all well and good, but as things stand these ideas seem unnecessarily complicated and abstract. Why can't we just trust our ability to count normally, even if we don't have an abstract, rigorous formulation of it? The power of this abstraction comes from how it allows us to deal with infinities. Let us see how this works.

Let's denote by S_{\aleph_0} the set $\{1, 2, 3, 4, \dots\}$ where the " \dots " simply means "this goes on forever". Thus, S_{\aleph_0} has infinitely many elements, and moreover these elements are all of the "counting" numbers. We then make the following definition.

Definition 1.38. A set A is said to have cardinality \aleph_0 (pronounced "aleph-naught") if there exists a bijective function $f : A \rightarrow S_{\aleph_0}$. We then say that A is "countably infinite".

A couple things. First, we've already seen that if there exists a bijective function $f : A \rightarrow B$, then there also exists a bijective function $g : B \rightarrow A$. Thus, I'll sometimes switch which way our bijective functions go, but that's okay, because a bijective function one way **defines for us** a bijective function the other way.

Note also that the above definition gives us a way of "counting infinity", in the sense that we can now assign a meaningful cardinality to infinity. But why do I give it this particular symbol, and why don't I just say that all sets with infinitely many elements have cardinality "infinity"? Well, it turns out that there are sets that are in fact **more infinite** than the infinity of S_{\aleph_0} . And by "more infinite" I mean the only thing that I can possibly mean—that there are sets that are infinite that **cannot** be put into bijection with S_{\aleph_0} .

Now, for any infinite set A , there is **always** a surjective function from A to S_{\aleph_0} . To see why this is the case, just consider taking any element in A and assigning it to $1 \in S_{\aleph_0}$, then taking any other element in A and assigning it to $2 \in S_{\aleph_0}$, and then taking any other element in A and assigning it to $3 \in S_{\aleph_0}$, and so

on. We know that we'll never run out of elements in A because it's infinite, and therefore we'll eventually hit all of the elements in S_{\aleph_0} . However, it's not always the case that there is a surjective function from S_{\aleph_0} to some infinite set A . Thus, S_{\aleph_0} is somehow "the smallest infinity". In fact, we'll soon see that there is an infinitely high tower of infinities above this "smallest" infinity!

First note that there are some subtleties here. Let me use this opportunity to introduce more notation. First, we write \mathbb{Z} to denote the set of integers: $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$. We write \mathbb{N} for the set $\{1, 2, 3, \dots\}$ (i.e., $\mathbb{N} = S_{\aleph_0}$). We denote by \mathbb{Q} the set of all fractions, and we denote by \mathbb{R} the set of all real numbers, of which we'll have more to say shortly.

Now, it's relatively easy to show that the cardinality of \mathbb{Z} is the same as the cardinality of $\mathbb{N} = S_{\aleph_0}$. Just define a bijective function between the two.

Exercise 1.39. Define a bijective function $f : \mathbb{N} \rightarrow \mathbb{Z}$.

Exercise 1.40. Show that if $f : A \rightarrow B$ and $g : B \rightarrow C$ are both bijective, then so is the composition $g \circ f : A \rightarrow C$.

Exercise 1.41. (this is tricky) Show that there is a bijective function $\mathbb{Q} \rightarrow \mathbb{N}$. Note that the above exercise says that it's sufficient to find a bijective function $\mathbb{Q} \rightarrow \mathbb{Z}$.

From the above exercises, we've found that the cardinalities of \mathbb{N} , \mathbb{Z} , and \mathbb{Q} are all "the same" in the sense defined above using bijective functions to "count infinity". But now we're going to find a new infinity. I.e., we'll find an infinite set such that it is absolutely impossible to find a bijective function from \mathbb{N} to it. To do so, however, we need to introduce some new elements.

Consider a number of the form

$$A.a_1a_2a_3a_4a_5a_6\dots \quad (1.1)$$

where $A \in \mathbb{Z}$, each $a_i \in \{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$, and the "..." as usual mean that these numbers can go on forever. Each such number is called a real number, and the set of all numbers of the form (1.1) is known as the set of real numbers, and denoted by \mathbb{R} . Note that these numbers include \mathbb{Z} , simply by considering all of those elements that have the digits after the decimal point set to zero, and that \mathbb{R} also includes \mathbb{Q} because any fraction can be expressed either as a finite decimal, or as a repeating decimal. It is known, however, that there are elements in \mathbb{R} that are not in \mathbb{Q} , as for example $\sqrt{2}$ (i.e., that there is no fraction equal to $\sqrt{2}$). We'll prove this below. For now, we take that fact on faith, and note that we have the following chain of **strict** inclusions:

$$\mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R}. \quad (1.2)$$

We've shown that the first three sets all have the same cardinality, and so perhaps it's the case that indeed, there is only one type of infinity. Let us show now that \mathbb{R} is indeed, in a very real sense, **bigger** than \mathbb{N} (and therefore bigger than \mathbb{Z} and \mathbb{Q} as well).

We proceed using proof by contradiction. We suppose that there is a bijection from \mathbb{N} to \mathbb{R} and arrive at a logical contradiction, thus making the initial assumption impossible to make, and thus proving that no such bijection can exist. So what we happen if there were such a bijective function? It would give us an

infinitely long list of real numbers, "labeled" by elements of \mathbb{N} , as follows:

$$\begin{aligned} A_1.a_{11}a_{12}a_{13}a_{14}a_{15}\dots \\ A_2.a_{21}a_{22}a_{23}a_{24}a_{25}\dots \\ A_3.a_{31}a_{32}a_{33}a_{34}a_{35}\dots \\ A_4.a_{41}a_{42}a_{43}a_{44}a_{45}\dots \\ \vdots \end{aligned} \tag{1.3}$$

where each individual real number goes off infinitely far to the right, and the list goes infinitely far down. Moreover, we've expressed the "labeling" by labeling each "integer part" of the real number (left of the decimal) by the element in \mathbb{N} that maps to it, and we've labeled each digit to the right of the decimal point by two numbers: the label on left corresponds to the same label as the integral part, namely, the natural number that maps to it, and the label on the right labels the distance of that digit from the decimal point.

We keep the labeling general like this, despite the fact that it's a bit confusing, to ensure that this argument will hold for **any** bijection $\mathbb{N} \rightarrow \mathbb{R}$. Thus, we keep the real numbers in the image of the bijection completely general.

Now recall that since this map is supposed to be bijective, then it must also be surjective, which means that **every** real number should be in this list somewhere. We'll get our contradiction by finding a real number that is not in the list.

Let us construct the following real number, written in the form $B.b_1b_2b_3b_4b_5\dots$ just as in (1.1). We take our list in (1.3) and use it to construct our real number by first letting B be any integer (it won't end up mattering), and then letting $b_1 = a_{11} + 1$, $b_2 = a_{22} + 1$, $b_3 = a_{33} + 1$, $b_4 = a_{44} + 1$. If it happens to be the case that some $a_{ii} = 9$, then we set the corresponding $b_i = 0$. Thus, all we're doing is going down the hypothetical list (1.3) along the diagonal, changing that entry by 1, and using it as the corresponding entry in the decimal expansion of the real number that we're constructing.

Now we can ask where in our list $B.b_1b_2b_3b_4b_5\dots$ appears, because after all, it's a real number and therefore should be somewhere. However, we know that it differs from the first entry in our list in the first decimal place (because $b_1 \neq a_{11}$), and that it differs from the second number in our list in the second decimal place (because $b_2 \neq a_{22}$), and that it differs from the third number in our list in the third decimal place (because $b_3 \neq a_{33}$). This goes on forever, and the number we've constructed differs from the N^{th} number in our list in precisely the N^{th} decimal place. Thus, our number is nowhere in the list, which is a contradiction to the fact that our list was supposed to hit every real number! Thus, no such list can exist. I.e., no such bijection is possible, since we've now seen that for **any** hypothetical list the enemy hands us, we can find some real number not in the list. We've therefore found a **new** infinity, i.e., one with a **larger** cardinality than the infinite cardinality that we defined above.

Now it turns out that there's actually an infinite tower of larger and larger infinities. In order to build this tower, however, we need one more tool. That tool is the power set.

Definition 1.42. Let A be a set. The **power set** of A , denoted by $P(A)$, is the set of subsets of A .

Exercise 1.43. Let A be a set with finitely many elements, and denote that number by N . Show that $P(A)$ has 2^N elements.

Just as the above exercise explored the power sets of set with finitely many elements, we can ask about the power sets of sets with infinitely many elements. Clearly, if the set A has infinitely many elements, then $P(A)$ also does (what is the easiest proof of this that you can think?). But now that we have the machinery to explore whether or not two infinities are "the same" (i.e., if two infinite sets can be bijectively related to each other), we might as well ask how the infinity of A relates to that of $P(A)$.

Claim 1.44. If A is a set with infinitely many elements, then there is no bijective function $A \rightarrow P(A)$.

Proof Let A be a set with infinitely many elements, and let $P(A)$ denote its power set. We prove this claim via contradiction. Suppose that there is a bijective function $f : A \rightarrow P(A)$. Then $\forall a \in A$, $f(a)$ is an element of $P(A)$, and is therefore a subset of A . Thus, for any element $a \in A$, we can ask whether or not $a \in f(a)$. We can then define the following subset S of A (i.e. $S \subseteq A$). We define S to be the subset that **does** contain $a \in A$ if $a \notin f(a)$, and which **does not** contain $a \in A$ if $a \in f(a)$. This is a perfectly good definition of a subset of A . Since f is supposed to be a bijective function, it is therefore surjective, and so there must be some element $b \in A$ such that $f(b) = S$, since S is just another subset of A . Let us then ask whether or not b is in S . Well, if b is in S , then b is in $f(b)$, because $S = f(b)$. But S was defined to be the subset such that if $a \in f(a)$, then $a \notin S$, and so it would be a contradiction if b were in S . Thus b must not be in S . But then if it's not in S , then $b \notin f(b)$, and so by the definition of S it **is** in S ! And since $S = f(b)$, we have yet another contradiction. So b can't be in S . But then b is both in and not in S , which is impossible! Thus our initial assumption—that such a bijective function exists—must be wrong, i.e., impossible. \square (box means QED, which means "proof is done!", or something, in Latin)

We've therefore established that no infinite set can be put into bijection with its power set. And since it's obvious that there are surjective functions from $P(A)$ to A (can you see why?), we have established that $P(A)$ has a fundamentally larger (albeit infinite) cardinality than A . But now the interesting part comes when we realize that since $P(A)$ is itself infinite (if A is), we can take " P " of it again. I.e., we can consider the power set of the power set of A , denoted by $P(P(A))$. This will have a larger cardinality than $P(A)$ and of A . We can then consider $P(P(P(A)))$, and $P(P(P(P(A))))$, and so on. Each of these has a fundamentally larger cardinality than those that came before it.

We've already seen that S_{\aleph_0} has is the "smallest infinity", and we've denoted its cardinality by \aleph_0 . We then denote the cardinality of $P(S_{\aleph_0})$ by \aleph_1 , and the cardinality of $P(P(S_{\aleph_0}))$ by \aleph_2 , and so on. We can genuinely think of this as "adding 1" to infinity, in the sense that taking the power set of an infinite set jumps us up to "the next" infinity. We could then even consider doing this infinitely many times:

$$\dots P(P(P\dots P(P(A))\dots))\dots \quad (1.4)$$

and denote that cardinality by, say, ω_0 . But then we can consider the operation of taking infinitely many P 's as a single operation in itself, and let that be the next operation that takes us up to the next class of infinities, say ω_1 . There are, however, infinitely many infinities "in between" ω_0 and ω_1 , obtained by hitting ω_0 by a single " P ". This is analogous to how there are infinitely many numbers between 1 and \aleph_0 . This wonderful field of math is called ordinal theory, or ordinal analysis. We'll be leaving this subject for now, because it quickly gets **very** abstract, but hopefully it has piqued the reader's curiosity.

Before ending this chapter, let's first show that real numbers are, in fact, real. By this I mean that we truly **need** them, and that \mathbb{Q} is not expansive enough. I'll do this by showing that the square root of 2, $\sqrt{2}$, is not rational. We most certainly want a number system that can describe this number (i.e., the number that, when multiplied to itself, gives 2), for the following reason. Suppose I draw a right triangle, where the two legs are both of length 1. Pythagorean's theorem tells us that the hypotenuse is then of length $\sqrt{1^2 + 1^2} = \sqrt{2}$. Since we certainly want to be able to talk about such triangles, we also certainly want to be able to talk about such numbers. It turns out that **most** real numbers are not rational

numbers—i.e., they're not $\in \mathbb{Q}$ and therefore not expressible as fractions—but simply showing that there is **one** such number will be enough to warrant our use of real numbers. We therefore have the following claim.

Claim 1.45. $\sqrt{2} \notin \mathbb{Q}$.

Proof: Suppose it were. Then $\sqrt{2} = \frac{a}{b}$ for some $a, b \in \mathbb{Z}$. Without loss of generality, we can assume that a and b are "completely reduced", meaning that there is no common factor between the two. In other words, we know for example that $\frac{2}{4} = \frac{1}{2}$, but the left hand side is not completely reduced in the way that we want, whereas the right hand side is. We can always turn a fraction into a "completely reduced" fraction by simply dividing out all common factors, as is hopefully clear.

Now, if $\sqrt{2} = \frac{a}{b}$, then after squaring both sides we have that $2 = \frac{a^2}{b^2}$. Multiplying both sides by b^2 , we have that $a^2 = 2b^2$, and therefore a^2 is an even number (since it's "2 times some number"). This means that a is also an even number, because an odd number times an odd number is an odd number, and therefore a must be even. Thus $a = 2k$ for some $k \in \mathbb{Z}$, which then implies that $(2k)^2 = 2b^2$. This then implies that $4k^2 = 2b^2$, which implies that $b^2 = 2k^2$. This implies that b^2 is even, which similarly implies that b must be even too (by the above logic for a). But this means that both a and b are even, which means that our fraction is reducible, which contradicts our assumption that we had already reduced our fraction! Thus $\sqrt{2}$ cannot be expressed this way, and is therefore $\notin \mathbb{Q}$. \square

Just like that we've proved that there are very "real" and necessary numbers that are not fractions, and therefore that the first new infinity that we found in this chapter—the infinity of \mathbb{R} —is a very "real" and necessary new infinity!

Chapter 2

Groups, Subgroups, and Homomorphisms

2.1 Introduction

We now want to add structure to our sets, i.e., to give them life. We're motivated to do this for several reasons, but the two that will concern us here are as follows. First, it appears that the world around us is described—at least with remarkable accuracy—by mathematical structures which are themselves nothing but "sets with structure". Indeed, an overwhelming amount of mathematics in general is nothing but the study of sets with various sorts of structure added to them. Which takes us to our second point. With added structure to our sets, we'll be able to prove many more things about them, as we have a lot more to work with. We'll see what all of this means as we go.

The first structure that we'll look at is that of a group. A group is nothing but a set that comes equipped with a particular way of "combining" elements in the set to form a new element in the set, as well as a special element (called an identity) that leaves all other elements "fixed" under this process of combination, and such that each element has a partner (called its inverse) element that, when combined with it, takes you back to the identity. Our motivating example will be the integers. For this set, the obvious form of "combination" is addition, which takes two integers and gives back an integer. The special element—i.e., the identity—is zero, because zero plus anything is again zero. The inverse of any number is simply its negative, since anything plus its negative is zero, the identity.

As we learn about groups, it is important to remember that the integers are merely a motivating example. Groups are one of the most versatile and pervasive structures that arise in mathematics and physics, and this power might be hidden if one only thinks about groups as an abstract way to talk about addition. Indeed, we'll see lots of examples of groups that have nothing to do with adding integers. Let us therefore just dive right in.

2.2 Groups

Definition 2.1. A **group** is a set G together with a function $g : G \times G \rightarrow G$ that is associative, so that $\forall a, b, c \in G, g(g(a, b), c) = g(a, g(b, c))$. Moreover, there is a special element $e \in G$ such that $\forall a \in G, g(a, e) = a$ and $g(e, a) = a$. Lastly, it must be the case that $\forall a \in G$ there must exist an element, which we denote by a^{-1} and call "the inverse of a ", such that $g(a, a^{-1}) = g(a^{-1}, a) = e$.

Now this might all seem a bit abstract, but when we compare it with the motivating example of the integers and identify all of the abstract parts of the definition with their corresponding concepts the \mathbb{Z} , we see that it's actually not that bad. To see this, we note that the function in the above definition is nothing but addition—namely, it takes two integers (or one element from $\mathbb{Z} \times \mathbb{Z}$) and gives another integer, and it

is clearly associative. The identity element of this special function is then 0, and the inverse of any $a \in \mathbb{Z}$ is nothing but $-a \in \mathbb{Z}$. This is clearly a **much** more abstract way to talk about adding integers, but as we'll see, this abstract definition is actually **extremely** powerful and beautiful. For now, let me go on to establish some notation that we'll end up using, and which is quite helpful.

When we add numbers together, we don't necessarily think in terms of this "special function" g and write $g(5, 3)$ when we want to express $5 + 3$. In fact, we can clearly see that the $+$ sign is simply acting as a different notation for this special function, and that it's a particularly useful one. We also know that the expression $(a + b) + c$ for "add a and b first, then add c ", and the expression $a + (b + c)$ for "add b and c first, then add a " are equal. This is simply a reflection of the fact that addition of integers is an associative operation, as are all group operations (by definition). It also happens to be the case that $a + b = b + a$, which in the notation of definition 2.1 corresponds to $g(a, b) = g(b, a)$, but we note that this quality (that of commutativity) is **not** in the definition of a group. I.e., we often have groups with elements a and b such that $g(a, b) \neq g(b, a)$. We'll see some of these soon.

We've seen in the example of the integers that the notation involving $+$ is quite useful, and that it strips away the unnecessarily clunky notation of $g(\cdot, \cdot)$. Thus, let us employ a similar notation **for any** group. What we do is let the symbol \cdot denote the "combining operation", so that $g(a, b)$ ends up being written $a \cdot b$. We could even be more lazy if we'd like, and simply write ab . Then, if our group is \mathbb{Z} with addition, it would be the case that $a \cdot b = a + b$, or $ab = a + b$. However, we can use this notation for other groups as well. Moreover, the seemingly scary expression for associativity simply becomes $(a \cdot b) \cdot c = a \cdot (b \cdot c)$ instead of $g(g(a, b), c) = g(a, g(b, c))$. Thus, we can write $a \cdot b \cdot c$ without any ambiguity, since it doesn't matter in which order we combine these elements. Note, though, that we cannot freely move these elements past each other (unless the group is commutative, which we are not assuming). And if we want to be really lazy, we can unambiguously write abc for the previous expression.

We also sometimes write (G, \cdot) to denote a group, along with its abstract form of composition (which we call its "group multiplication", even if it has nothing to do with actual multiplication, as in the case of the addition of integers (we'll see why this terminology is a good one later)).

To begin to get a feel for groups and how general of a definition this is, let us look at some examples.

Example 2.2. Consider the set $\{0, 1, 2, 3, 4\}$, and let us define an operation on them that we'll call "addition modulo 5". What this will mean is that when we add two elements together, we take their normal sum, then divide by 5, and take the remainder of this division as our result. So, for example, $4 + 4 = 8$, and the remainder of 8 divided by 5 is 3. Thus, $4 + 4 \bmod 5 = 3$ (where we tack on the *mod* 5 to denote "modulo 5"). It is important to note that we can define this type of addition on any pair of integers, even those not in the above set. For example, $20 + 4 = 24$, and the remainder of 24 when divided by 5 is 4, so we therefore have that $20 + 4 \bmod 5 = 4$. We can even define the "modulo 5" function on the integers themselves, where each integer is simply mapped to its remainder when divided by 5. Thus, $4 \bmod 5 = 4$, $20 \bmod 5 = 0$, and $73 \bmod 5 = 3$. In order to see how this works on negative integers, it is important to note that when we divide by 5 we're taking the **positive** remainder. In other words, we always "stop" at the multiple of 5 that is **less than** the number under consideration, and then subtract this number from the number we started from. Thus $73 \bmod 5 = 73 - 70 = 3$, and so on. When we ask about $-3 \bmod 5$, we must find the first multiple of 5 that is **less than** -3 and then subtract. The desired multiple of 5 in this case is of course -5 , and then we have that $-3 \bmod 5 = -3 - (-5) = 2$. This way we have a more obvious relationship between integers and integers modulo 5:

$$\begin{array}{ll} \text{integers :} & \dots, \quad -6, \quad -5, \quad -4, \quad -3, \quad -2, \quad -1, \quad 0, \quad 1, \quad 2, \quad 3, \quad 4, \quad 5, \quad 6, \quad \dots \quad (2.1) \\ \text{modulo 5 :} & \dots, \quad 4, \quad 0, \quad 1, \quad 2, \quad 3, \quad 4, \quad 0, \quad 1, \quad 2, \quad 3, \quad 4, \quad 0, \quad 1, \quad \dots \end{array}$$

Now let us get back to our set $\{0, 1, 2, 3, 4\}$ with "addition modulo 5" defined on it. This is a group. It is easily checked that this operation is associative and that 0 is our identity element (namely, it's still the case that "0 plus anything (modulo 5) is still the same anything". We then have that the inverse of 0 is 0, the inverse of 1 is 4, the inverse of 2 is 3, the inverse of 3 is 2, and the inverse of 4 is 1.

It should also be relatively clear how we could generalize this to addition modulo any positive integer N . In particular, what we can do is consider the set $\{0, 1, 2, 3, \dots, N-1\}$ with "addition modulo N " defined on it. This will still be an associative operation, 0 is still our identity element, and we have that the inverse of 0 is 0, the inverse of 1 is $N-1$, the inverse of 2 is $N-2$, the inverse of 3 is $N-3$, and so on. This is called **modular arithmetic** with addition modulo N , and it is a very important concept to understand fully. We note here that we often denote this set with this abstract group multiplication \mathbb{Z}_N , or $\mathbb{Z}/N\mathbb{Z}$, for reasons that we'll see later.

Example 2.3. Let's now examine another extremely important group, known as either the **Symmetric group (on n elements)** or the **Permutation group (on n elements)**. This turn out to be a group that has elements a and b such that $a \cdot b \neq b \cdot a$, which is a significant difference between the permutation group and, say, the group of integers with addition.

To define this group let us consider the following scenario. Suppose we have a set of n elements (these elements can be anything at all), and suppose we have a set of precisely n boxes labelled 1 to n . Finally, suppose we've distributed our n elements amongst these n boxes in the most democratic way possible, namely where each box has one and only one element in it (i.e., no empty boxes and no boxes with 2 or more elements in it). Now suppose we ask the following question: How can we "shuffle" these elements around in such a way that when we're done shuffling, each box still has one and only one element in it? Answering this question motivates the definition of the permutation group. A more abstract (and completely equivalent) way of thinking about this question is to realize that we're asking about what kinds of bijective functions can be defined from this set of n elements to itself.

We call each possible shuffle of the elements in these boxes (or equivalently each bijective function from the set to itself) a **permutation** of the set. The most trivial permutation of this set is to just leave all the elements alone in the box that they started in. Let's call this permutation "do nothing" and let's denote it by e . The next simplest permutation is to take two boxes and swap their contents. For example, we could take the element in box 1 and put it into box 2, and take the element in box 2 and put it in box 1. We can call this permutation "swapping boxes 1 and 2", and let's denote this permutation by $(1, 2)$. There are two things to be said about this choice of notation. The first is that whenever a number **does not** appear in the expression, it should be understood that the element in the box corresponding to that number is left alone. Thus, $(1, 2)$ means "swap 1 and 2's elements, and leave the rest alone". Secondly, we should read each individual parenthetical expression as a cyclic expression, meaning that it cycles through the end and comes back to the beginning. We often start on the left of the parentheses and move to the right, taking the elements of one box and putting it in the box whose number appears on the right. Thus, $(1, 2)$ is interpreted as follows: take the element in box 1 and put it in box 2 (because 2 is directly to the right of 1), then take the element that was already in box 2 and put it in box 1 (because 1 is directly to the right of 2, when we cycle through the end of the parentheses). This notation will become more clear with some more examples.

Along with swaps, we can also consider more complicated permutations. For example, we could consider the permutation "take the element in box 3 and put it in box 5, then take what was already in box 5 and put it in box 2, then take what was already in box 2 and put it in box 17, then take what was in box 17 and put it in 3". With the notation that we've developed above, this permutation will be represented by $(3, 5, 2, 17)$ (note that this also includes the part of the permutation that takes the element out of 17 and puts it in 3 because we cycle back through the end of the parentheses). Note also that the cyclic nature of our notation ensures that the functions that they define are truly bijective, in that our elements never "pile-up" anywhere. In particular, since our parenthetical expressions are interpreted cyclicly, it is always the case that each box has precisely one element put in it and one element removed from it. Another thing

to note is that the cyclicity of these expressions means that we can "cycle the numbers through" without changing the permutation that they represent. For example, $(1, 2) = (2, 1)$, $(1, 2, 3) = (3, 1, 2) = (2, 3, 1)$, and so on (but note that, for example, $(1, 3, 2) \neq (1, 2, 3)$).

We can also put these parenthetical expressions next to each other when we need to. For example, suppose our set has at least 4 elements so that we can meaningfully ask about the permutation "swap 2 and 4" and "swap 1 and 3". Note that this should be viewed as a **single permutation**, since it defines a perfectly good bijective function from the set to itself. We want our notation to be able to handle cases like these, and so we extend our notation to allow for expressions like $(1, 3)(2, 4)$, where we only cycle through each individual parenthetical expression. There is some subtlety here, though, and we'll return to it shortly. First, let us ask why we're even doing all of this.

The goal here is to be able to talk about the "group of permutations" of these n elements, where we view the **permutations themselves** as the elements of our group. Thus, the set that we're permuting plays only a background role in this story, since all we need it for is to give us a collection of things to shuffle around. What makes up the actual group is the more abstract set of **actions** (permutations) on this set. In order to make the set of permutations into a group, we need a way of combining two permutations into one permutation. The obvious choice for this is to say that the abstract group product of two permutations is the composition of them. That is, if we have two permutations a and b , then we can compose them into the single permutation of "first do a , then do b ". We can denote this by $b \cdot a$. Note that our convention here is to put the permutation that "happens first" on the right, so that $b \cdot a$ means "do a then do b ", whereas $a \cdot b$ means "first do b , then do a ".

Let's see how this works with some of the permutations we've already considered. Suppose we want to compose the two permutations "swap 1 and 2" and "swap 3 and 1", so that we're composing $(1, 2)$ with $(1, 3)$. Moreover, suppose we want to do $(1, 2)$ first, then do $(1, 3)$. We then have the product permutation $(1, 3) \cdot (1, 2)$. Now let's write this as a single expression, without the "." and so that each number only appears once (after all, this should be expressible as a **single** permutation). Let's see what this single permutation is. We're first taking the element in 1 and putting it in 2 (this comes from the first permutation), and then the second permutation doesn't touch the box 2, so we know that the net effect of this product is to take the element in box 1 and put it in box 2. Thus, we'll have something like $(1, 2, \dots)$, where we don't know what is in the "..." yet because we don't know where the element that was originally in box 2 goes. Let us now find out. We first take the element originally in box 2 and put it in box 1 (from the first permutation), but then the second permutation tells us to put the element in box 1 into box 3. Thus, the net effect of this product on the element in box 2 is to send it to 3. Therefore we have $(1, 2, 3, \dots)$. All that remains is to see where the element in box 3 goes. Our guess is that it must cycle back through to 1, since that's the only possibility for it, but let's just make sure our notation is consistent with this. Indeed it is, because the first permutation doesn't touch box 3, and the second permutation takes the element in box 3 and puts it in box 1, so that our permutation indeed does cycle back appropriately and we have that $(1, 3) \cdot (1, 2) = (1, 2, 3)$.

Let us now suppose that we want to take the product of $(1, 2)$ and $(1, 3)$ in the reverse order, so that we're first doing $(1, 3)$ and then doing $(1, 2)$. Thus, we're asking about $(1, 2) \cdot (1, 3)$. Well, the first permutation takes the element in box 1 to box 3, and the second permutation doesn't touch box 3, so we have something like $(1, 3, \dots)$. The first permutation also takes the element in box 3 to box 1, and then the second permutation takes the element in box 1 to box 2, so that the net effect is that 3 goes to 2, so we have $(1, 3, 2, \dots)$. Finally, the first permutation doesn't touch box 2 and the second permutation takes the element in box 2 to box 1, so that the net effect is 2 goes to 1, so we have $(1, 2) \cdot (1, 3) = (1, 3, 2)$. Note, however, that $(1, 2, 3) \neq (1, 3, 2)$, so that $(1, 3) \cdot (1, 2) \neq (1, 2) \cdot (1, 3)$. We've thus found two permutations $a = (1, 2)$ and $b = (1, 3)$ such that $a \cdot b \neq b \cdot a$! This is not a peculiar phenomenon, but rather precisely what makes group theory so interesting!

Before going on to look for an identity element and inverses, we need to make sure that this abstract group multiplication is indeed associative. But this is actually trivially easy, for the following reason. Suppose we have three permutations a, b , and c and suppose we wanted to form the two products $c \cdot (b \cdot a)$

and $(c \cdot b) \cdot a$. In other words, the first product is the permutation that first does $b \cdot a$ and then does c , and the second product is the product that first does a and then does $(c \cdot b)$. We need these net products to be equal, for all a, b, c , in order for this product rule to be associative. But these most certainly are equal, because both of them are simply "first do a , then do b , then do c ". To see this, note that the product $c \cdot (b \cdot a)$ says to first do $b \cdot a$, which is in itself short for "first do a then do b ", and then we compose c onto this so that the net effect is " a then b then c ". Similarly, $(c \cdot b) \cdot a$ is "first do a " and then do the product $c \cdot b$, which is itself "first do b then do c ", so that the net effect is also "first a then b then c ". The only difference in these products is our choice of which pair of permutations to view as a single permutation. Namely, we have three permutations that we do in order, and we can choose any pair of consecutive permutations and view their product as a single permutation. Clearly, our choice of what to count as a single permutation doesn't affect the permutation itself. The next paragraph gives a more humble example of this idea.

Suppose a woman named Sally is extremely well organized and suppose that every day she makes a to-do list for herself. Suppose also that she follows her to-do list exactly, including the order in which she does her actions. It's a Saturday so she's not very busy, and her to do list is to 1) wash the dog, 2) go for a run, and 3) call her father. If she desired, she could reduce her list to a list of only two if she clumped them together as 1) wash the dog then go for a run, and 2) call her father. Equivalently, she could have written 1) wash the dog, and 2) go for a run then call her father. Clearly, no matter what choice she makes, her actions and the order in which she does them are the exact same. Thus, her day is unaffected by this arbitrary choice. This is exactly what our composition is doing for us—it just clumps together two actions and views them as a single one, all while maintaining their order. Thus, no matter how we clump consecutive actions, our total product will be the same. Thus, our product is associative.

Now that we have an associative product on our set, we need to know if there is an identity element. If we let e = "do nothing", as we have above, then this will be a perfect identity element. This is because e leaves every element in our set alone under our product rule. If a is a permutation, then $a \cdot e = e \cdot a = a$, since "do nothing then permute" is the same as "permute then do nothing", which is also the same as just permuting. So there we have it—our identity element.

All that remains now is to show that for any permutation a , there is a permutation b such that $a \cdot b = b \cdot a = e$. If we can always find such an element, we'll call it the inverse of a , and our set of permutations will indeed be a group under composition. Luckily, it's relatively easy to find our inverses. Let's first consider the permutation $(1, 2, 3)$ and see if we can find its inverse (if we can't find it for this simple case, then we're in trouble if we want to find it for the general case). This permutation takes the element in box 1 to 2, that in box 2 to 3, and that in box 3 to 1. Thus, if we want to arrive back at the permutation that "does nothing", we need to reverse this permutation one step at a time. Namely, we need to put the element in box 1 back into box 3, the element in box 3 back into box 2, and the element in box 2 back into box 1. This permutation is precisely $(3, 2, 1)$, and we have that first doing $(1, 2, 3)$ then doing $(3, 2, 1)$ is the same as doing nothing, so that $(3, 2, 1) \cdot (1, 2, 3) = e$. Similarly, it is easy to check that $(1, 2, 3) \cdot (3, 2, 1) = e$ as well, so that $(3, 2, 1)$ is indeed the inverse of $(1, 2, 3)$.

In fact, it is easy to check that for any parenthetical expression like $(i_1, i_2, i_3, \dots, i_N)$, the inverse of it will be the expression in reverse order, namely $(i_N, i_{N-1}, \dots, i_3, i_2, i_1)$. We've therefore found the inverse of any permutation that can be expressed as a single parenthetical expression, but what about those that involve several parenthetical expressions, like $(1, 3, 4)(2, 5, 6)(7, 8)$? The subtlety here is that we still view this total expression as beginning with parenthetical expression furthest to the right. Since we want to end up reversing this permutation to give us the "do nothing" permutation, we want to reverse the order not only of the numbers within a given parenthetical expression, but also the order of the parenthetical expressions themselves. Thus, we're really reversing this permutation one step at a time. Therefore, the inverse of $(1, 3, 4)(2, 5, 6)(7, 8)$ is $(8, 7)(6, 5, 2)(4, 3, 1)$, because then their product is

$$(8, 7)(6, 5, 2)(4, 3, 1) \cdot (1, 3, 4)(2, 5, 6)(7, 8).$$

We can then see that we're first doing the permutation $(1, 3, 4)(2, 5, 6)(7, 8)$ and then simply reversing each

move one step at a time. This process will clearly work for any permutation made up of any number of parenthetical expressions. We simply reverse the order of the parenthetical expressions, and then reverse the order of the numbers within each given parenthetical expression! It is therefore the case that any permutation has an inverse permutation.

To recap, we've seen in this extremely long example that the set of all permutations of N objects, for any positive integer N , is a group where the abstract group multiplication is composition. This group is commonly denoted S_N . The identity element is the "do nothing" permutation. Moreover, we saw that there are elements $a, b \in S_N$ such that $a \cdot b \neq b \cdot a$, thus showing the true generality of groups, as this group is radically different from anything involving the addition or multiplication of numbers.

Exercise 2.4. What is the cardinality of S_n , in terms of n ?

One thing that needs to be discussed, and which the keen reader may have picked up on, is whether or not we are correct in referring to "the" identity element, and "the" inverse of an element. Namely, how do we know that there's only one identity element, or that for any given element there is only one inverse? Well, we don't know this yet, but I claim it's true. Let us see how this works.

Claim 2.5. Let G be a group. Then the identity element e is unique.

Proof: We use a standard method of proof for showing that an object with certain qualities is unique, and that is to first consider two objects with these qualities and showing that they're actually equal. This clearly means that there must only be one such object. Accordingly, let us suppose f and e are both identity elements for G . This means that $e \cdot f = e$, since f is an identity element and therefore fixes anything that it is multiplied by. But e is an identity element also, so it fixes everything it multiplies as well. In particular, we know that $e \cdot f = f$. We therefore have that $e = e \cdot f = f$, and therefore that $e = f$. This is exactly what we wanted to show. \square

We've therefore shown that we are warranted in referring to "the" identity element, because we've just shown that there is only one of them. Thus, once we find one identity element we can stop looking for others. This may seem obvious, since the identity element in the integers under addition is 0, and it would seem foolish to keep looking around for other integers that keep everything fixed under addition (because we know that any non-zero integer changes whatever it's added to). However, what's important here is that our definition doesn't **assume** that identities are unique. Instead, we had to **prove** it from the axioms of a group. Moreover, we now know that **any** group that the enemy hands us will have only one identity element, which is a much more general statement than just saying that 0 is the only integer that doesn't change things when it's added to them.

Let us now also show that we are indeed warranted to talk "the" inverse of any given element.

Claim 2.6. Let G be a group and let $a \in G$ be arbitrary. Then the inverse of a is unique.

Proof: Again, we consider two inverses of a and show that they're equal. Let b and b' both be inverses of a . Then we know that $b \cdot a = e$ and that $b' \cdot a = e$ (note that we can set both of these expressions equal to e because we just proved that there is only one identity e in any given group). Therefore $b \cdot a = b' \cdot a$. If we multiply both sides of this equation by b on the right, we have $(b \cdot a) \cdot b = (b' \cdot a) \cdot b$, and then associativity lets us move these parentheses around so that we have $b \cdot (a \cdot b) = b' \cdot (a \cdot b)$. But since b is an inverse of a , we have that $a \cdot b = e$, so that we get $b \cdot e = b' \cdot e$, and this implies, by definition of the identity, that $b = b'$, which is exactly what we wanted to prove. \square

This again may seem obvious in the case of integers with addition, because we somehow already just **know** that the only additive inverse of 4 is -4 . However, we've now proved it from first principles, and we've proved it in much greater generality. This is the power of abstract algebra (of which group theory is a part), and this is the power of abstract mathematics as a whole.

2.3 Subgroups

Just as subsets provided a notion of "substructure" to a set, there exists a proper notion of "substructure" for groups as well. But what do I mean by "proper"? Obviously we can ask about subsets of groups, because groups are indeed sets themselves (just sets with extra structure). But clearly there is something different about the subset $\{\dots, -6, -4, -2, 0, 2, 4, 6, \dots\}$ (where again the dots mean the pattern is continued in both directions) of \mathbb{Z} , and the subset $\{3, -17, 100\}$. One of these subsets (the former) seems to have more structure than the other (the latter)—it seems to maintain some sense of "interesting-ness", whereas the latter doesn't seem like it will be very interesting to study. This is because the former subset itself forms a group. You can check that this subset satisfies all of the requirements for a group. In fact, its special function (addition), identity element (0), and inverses ($-a$) are all the same as those for \mathbb{Z} itself. We therefore say that this subset is a **subgroup** of \mathbb{Z} , and this motivates the following definition.

Definition 2.7. Let G be a group, and let H be a subset of G . We call H a **subgroup** of G if it is itself a group with the same group multiplication, identity element, and inverses as G .

One important consequence of this definition is the fact that subgroups are what's called **closed** under the multiplication that it inherits from G . Namely, it must be the case that if I take any two elements within the subset H and combine them (using the multiplication law from G , which is well defined because any element in H is also an element in G), it must be the case that their product is also in H , for otherwise H wouldn't be a group on its own. Note that it might be the case that I can multiply something in H by something out of H (but also in G) and get something that is also out of H . Note also that if $a \in H \subseteq G$, and if H is a subgroup of G , then $a^{-1} \in H$ also. This is because for H to be a group, the inverses of all of its elements must also be in H (and we know that inverses are unique, and so we never need to ask "which inverse" is in H because there's only one).

Example 2.8. Let's consider the group $\{0, 1, 2, 3, 4, 5, 6, 7, 8, 9\}$ with addition modulo 10. One subgroup of this group is the subset of elements $\{0, 5\}$. This is because $0 + 0 = 0$, $0 + 5 = 5 + 0 = 0$, and $5 + 5 = 0$ modulo 10. Thus this set is closed under the group multiplication (which is addition modulo 10 here), and it includes the identity as well as the inverses of all of its elements. Similarly, the subset $\{0, 2, 4, 6, 8\}$ is a subgroup.

We can actually do this more generally. Suppose we consider the set $\mathbb{Z}_N = \{0, 1, 2, 3, \dots, N - 1\}$ with addition modulo N . Now suppose that a is some positive integer that divides N evenly. Then the subset S_a of \mathbb{Z}_N consisting of all multiples of a is a subgroup. This is because S_a is closed under addition, since the sum of any two multiples of a is again a multiple of a . Moreover, we have that $0 \in S_a$ because 0 is a multiple of a (namely, $0 = 0 \times a$), and for every $x \in S_a$, the inverse of x is also in S_a . To see this, we consider some $x \in S_a$. Since x is a multiple of a , we know that $x = ka$ for some positive integer k . We also know that a divides N evenly, so that there is some positive integer m such that $N = ma$. Let's let $y = (m - k)a$. Then $x + y = ka + (m - k)a = ma = N = 0$, where the last equality comes from the fact that we're doing addition modulo N . Thus y is the additive inverse of x , and y is a multiple of a , so $y \in S_a$. Thus S_a is a subgroup. We note that the first example in the paragraph prior to this one had $a = 5$, and the second example had $a = 2$.

Exercise 2.9. Write down five subgroups of \mathbb{Z}_{36} .

Example 2.10. Let us now consider some subgroups of the permutation groups S_n (sorry for the confusing notation of using " S " both here and in the last example—please note that these sets have no relation to each other). One thing we can notice right away is that for any "swapping" of elements, namely for any (i, j) where $i \neq j$, we can form the subgroup $\{e, (i, j)\}$. This is because $(i, j) \cdot (i, j) = e$, so that this set is closed under composition, has the identity, and has inverses.

There are in fact tons of subgroups of permutation groups. For example, if we consider S_4 , the set of permutations of 4 elements, then the subset $\{e, (1, 2), (3, 4), (1, 2)(3, 4)\}$ is a subgroup. It is easy to check that this is closed under composition (this is why we needed to include the permutation $(1, 2)(3, 4)$, because otherwise $(1, 2) \cdot (3, 4)$ would not be in this subset and it therefore wouldn't be closed under composition), it clearly contains the identity, and every element is its own inverse, so this subset contains inverses too.

In fact, we can take any permutation and form a subgroup using it, by simply **forcing** our subset to be closed under composition. For example, take the element $(1, 3, 2, 4) \in S_4$ and suppose we want to construct a subgroup H containing this element. All we do is put this element as well as all powers of it (namely, this element composed with itself over and over again until we get back to the identity) in H . Thus, we start by putting $(1, 3, 2, 4)$ in H . Then we put $(1, 3, 2, 4) \cdot (1, 3, 2, 4) = (1, 2)(3, 4)$ in H as well. Then we put $(1, 3, 2, 4) \cdot (1, 3, 2, 4) \cdot (1, 3, 2, 4) = (1, 3, 2, 4) \cdot (1, 2)(3, 4) = (1, 4, 2, 3)$ in H as well. Then we put $(1, 3, 2, 4) \cdot (1, 3, 2, 4) \cdot (1, 3, 2, 4) \cdot (1, 3, 2, 4) = (1, 3, 2, 4) \cdot (1, 4, 2, 3) = e$ in H , and that completes the subgroup (namely, because multiplying by $(1, 3, 2, 4)$ again will just get us back to $(1, 3, 2, 4)$). Thus $H = \{e, (1, 3, 2, 4), (1, 2)(3, 4), (1, 4, 2, 3)\}$ is a subgroup containing $(1, 3, 2, 4)$. There are of course other subgroups containing this element, but this is the smallest, since it has the bare minimum number of elements required to be closed under multiplication.

Exercise 2.11. Find a subgroup of S_5 containing the element $(1, 2)(3, 4, 5)$.

2.4 Abelian Groups

We've always been very careful to point out that in a general group it is not necessarily the case that $a \cdot b = b \cdot a$ (sometimes we'll be even more lazy and drop the "." and simply write ab for $a \cdot b$). In fact, lots of very interesting groups have elements such that $ab \neq ba$ (like the permutation groups that we've already seen). Another example of such a group is known as the **dihedral group** of some regular polygon (i.e., a polygon such that the angle between all adjacent sides are the same, and all side lengths are the same). Let us take a brief detour and explore these dihedral groups.

Consider a regular polygon, and for simplicity let us take a square (i.e., a regular polygon with four sides). Draw a square on a piece of paper (seriously, please draw it). Pick some corner of it and label that vertex with a "1". Now move in a **counterclockwise** direction from that vertex and label the next vertex (i.e., the next corner) with a 2, and the next corner with a 3, and the last corner with a 4, so that at this point the next corner should be the "1" corner. We can (and should) think of this as a square napkin lying on a table, where we're writing the numbers down **on the table**, so that they provide a way of defining how to rotate/flip the napkin by saying "move the corner at location i to location j " (where i and j are numbers between 1 and 4). This is analogous to how we number the boxes in the case of the permutation groups as a means of specifying which permutation we were talking about.

Now consider the group of (say, counterclockwise) rotations of this square. This group has four elements: the identity ("do nothing"), the rotation counterclockwise by 90 degrees, the rotation counterclockwise by 180, and the rotation counterclockwise by 270. This is a group such that $ab = ba$ for any a and b in the group, since it doesn't matter whether we do rotation 1 then rotation 2, or rotation 2 and then rotation 1. This is because we're simply adding the number of degrees that we're rotating by in the exact same way as we add numbers, and the addition of numbers is commutative. Think of this as just adding minutes on a clock, where clearly it doesn't matter if we first add 10 minutes and then 15, or first 15 and then 10.

But we can include more complicated symmetries of the square, and in so doing we can consider a group that is not commutative. Namely, we now not only consider all of the square's rotations, but also its reflections. I.e., we can reflect the square through various axes of symmetry and still obtain an object identical

to the original square. In order to describe this, we need to understand where the vertices go under these transformations. In the case of a 90 degree counterclockwise rotation, we have that vertex 1 goes to vertex 2, 2 goes to 3, 3 to 4, and 4 to 1. For a 180 degree rotation, we have $1 \rightarrow 3, 2 \rightarrow 4, 3 \rightarrow 1, 4 \rightarrow 2$.

Exercise 2.12. Write down where all of the vertices go for a 270 degree counterclockwise rotation.

Now consider the axis of symmetry that connects the midpoint of the edge 1-2, and the edge 3-4 (by "1-2" I mean the side of the square whose ends are vertices 1 and 2). The line connecting these two midpoints is such that if we flip the square **about that axis** (i.e., through the third dimension), then we end up with a square again. In this case, we have that $1 \rightarrow 2, 2 \rightarrow 1, 3 \rightarrow 4, 4 \rightarrow 3$. Let us call this motion of the square, i.e., this reflection, r_1 . Now consider the line connecting the midpoints of the sides 1 – 4 and 2 – 3. This is another axis of symmetry, and we can also flip the square about this axis (again through the third dimension) so that we have $1 \rightarrow 4, 2 \rightarrow 3, 3 \rightarrow 2, 4 \rightarrow 1$. Let us call this reflection r_2 .

Let me make a note of why we're calling these flips "reflections". We do so because we notice that the act of flipping the napkin in this way is precisely the same (in terms of where the corners of the napkin are sent to) as "reflecting" the napkin through itself along the axis of rotation that we're considering. Namely, instead of flipping the napkin around this axis, let's lay a mirror down along this axis and send the points of the napkin to their corresponding reflection points. We would then get the same exact assignment of corners of the napkin. Even though we call these "reflections", I think it's more natural to think of them as "flips" (mathematical terminology apparently doesn't agree with me, however).

We have two more reflections to consider. These are the reflections about the two diagonals of the square. I.e., we get one more reflection by flipping the square about the line connecting the corners 1 and 3, and we get another one by rotating the square about the line connecting the corners 2 and 4. The first reflection, about the diagonal connecting 1 and 3, leaves 1 and 3 fixed and switches 2 and 4, so that we have $1 \rightarrow 1, 2 \rightarrow 4, 3 \rightarrow 3, 4 \rightarrow 2$. Let us call this r_3 . Similarly, the reflection about the diagonal connecting 2 and 4 gives $1 \rightarrow 3, 2 \rightarrow 2, 3 \rightarrow 1, 4 \rightarrow 2$, and let us call this r_4 (note that by writing these expressions for where the vertices go, we're not really relabeling any vertices, but rather denoting the operation on the square that configures the square in such a way that the vertex on the left of the arrow ends up wherever the vertex on the right of the arrow was before the operation was made. Thus $1 \rightarrow 4$ **must** be accompanied by $2 \rightarrow 3$, since it's impossible to "move 1 to 4" without also moving 2 to 3).

Now we consider the group of all of these motions of the square, i.e., the flips r_1, r_2, r_3 , and r_4 and the rotations that we had before. Let's write the identity as e and the 90 degree rotation as ρ_1 , the 180 degree rotation as ρ_2 , and the 270 degree rotation as ρ_3 . Then $\rho_1\rho_3 = e$, $\rho_2\rho_2 = e$, $\rho_1\rho_2 = \rho_3$, and so on (where we view the element on the right as "happening first", so that $\rho_1\rho_2$ is short for "first rotate by 180, then by 90"). We noted above that any combination of ρ 's is commutative, meaning that $\rho_i\rho_j = \rho_j\rho_i$ for any i, j . Now let us see that this group as a whole is **not** commutative. In order to do this, we clearly need to consider the reflections.

Let's ask what happens when we do the action $\rho_1 \cdot r_2$. Let's first see where the first vertex goes. Under r_2 , we see from above that $1 \rightarrow 4$. Then, after rotating by 90 degrees (applying ρ_1) we see that $4 \rightarrow 1$, so that the composite action of $\rho_1 r_2$ actually leaves 1 fixed: $1 \rightarrow 1$.

Exercise 2.13. Work out where the other three vertices go.

However, if we consider the action of $r_2\rho_1$, i.e., the element obtained by these two transformations in the reverse order, we'll find a different total transformation. Let us again see where vertex 1 goes. Now, it is first rotated counterclockwise by 90 degrees (ρ_1), so we have $1 \rightarrow 2$, and then applying r_2 , we see that $2 \rightarrow 3$, so that in total this transformation gives $1 \rightarrow 3$. This is certainly different from what we got in the previous case!

Exercise 2.14. Work out where the other three vertices go.

We've thus seen another example of a group that is non-commutative, and these often turn out to be some of the most interesting groups. Similar groups can be defined and explored by using regular polygons of more or less sides (equilateral triangles, regular pentagons, regular hexagons, and so on), and considering their various rotations and reflections about axes of symmetry. Let us end this section by giving a name to the groups that are **not** interesting in this way (although they'll be interesting in other ways). This all comes from the fact that one of the first guys to study these kinds of groups was a dude named Abel, so he gets his name on it.

Definition 2.15. A group G is said to be **Abelian** if $ab = ba \forall a, b \in G$. If this is not the case, then the group is **non-Abelian**.

Exercise 2.16. What are the respective inverses of r_1, r_2, r_3 , and r_4 ? (Hint 1: once you've found the inverse of one of them, the rest will follow similarly. Hint 2: Don't over think it. If I do r_2 , how can I get back to the configuration of having done nothing?)

Exercise 2.17. What are some subgroups of this group? Which of those subgroups are Abelian?

Finally, we note that \mathbb{Z}, \mathbb{Q} , and \mathbb{R} are all Abelian under addition, and $\mathbb{Q} \setminus \{0\}$ and $\mathbb{R} \setminus \{0\}$ (which means \mathbb{Q} without the element 0, etc.) are Abelian under addition.

2.5 Group Constructions

Now that we've defined groups, seen how we can find groups "within" groups, and explored some interesting examples of groups, we now seek to find ways that we can build new groups from old ones. This is very similar to what we did in the set theoretic discussion of lecture 1, where we defined sets and subsets, and then proceeded to define unions, intersections, and Cartesian products. For groups, we'll only explore one such construction here, as most of the other constructions require a bit more machinery. For one such construction, however, we will indeed go on to develop the necessary machinery in due time.

The most immediate definition that we can make is that of the **direct product group**. This is obtained by taking the Cartesian product of two groups and extending their respective multiplications in the only way possible. Let's make this precise with the following definition.

Definition 2.18. Let (G, \cdot_G) and (H, \cdot_H) be groups (note that G and H are completely arbitrary, and one is not necessarily a subgroup of the other. Note also that the symbols for their respective abstract multiplication rules include a subscript of the group itself, and that is because these are different groups and so their multiplications may be completely different from each other, and using different symbols for them will remind us of this). Then the **direct product group** of G and H , denoted by $(G \times H, \cdot)$, is the Cartesian product of G and H with the multiplication rule $(a, b) \cdot (c, d) = (a \cdot_G c, b \cdot_H d)$.

There are several remarks that need to be made about this definition, so if any of it is confusing don't worry, as the following might clear some things up. Let's first note that regardless of whether or not G and H are groups, we can take their Cartesian product. I.e., the Cartesian product itself is a construction on sets and knows nothing about the group structure that G and H have. Thus, what is important about this definition is that it gives a way of giving a group structure to the Cartesian product of G and H **when G and H are groups**.

But what exactly is this group structure? Remember that it needs to be a way of combining elements

in $G \times H$, and so we need a way of combining elements of the form (a, b) where $a \in G$ and $b \in H$, with elements of the form (c, d) with $c \in G$ and $d \in H$. The most natural thing to do would be to just let the respective multiplication rules that we know we have from G and H simply "extend" to the respective "slots" in the Cartesian product—that is all we're doing in the last line of the definition with the expression $(a, b) \cdot (c, d) = (a \cdot_G c, b \cdot_H d)$. The left hand side of this equality involves a \cdot without any subscript because this is the abstract multiplication on $G \times H$ itself and so we adopt the usual notation for it. However, what it instructs us to do is take two elements $(a, b), (c, d) \in G \times H$ and multiply them "component-wise" by multiplying the first slot using the multiplication from G and the second slot using the multiplication from H . Remember, we're turning the **set** $G \times H$ into a **group** using the information that we have from the "building block" groups.

It still remains to show that this definition is indeed a definition—i.e., we need to show that it is **well-defined**. The notion of "well-defined" in mathematics is completely pervasive, and fully understanding what it means in every situation is sometimes difficult and usually just comes with experience. In this case, it simply means that we need to check that Definition 2.3 actually defines a group. I.e., we can't define something to be a group without checking that it actually is a group. For example, saying the phrase "let \mathbb{Q} be a group under multiplication" is not well-defined because we know that \mathbb{Q} is not a group under multiplication—we either have to consider it under addition or remove 0 from it to consider it under multiplication. Thus, we need to make sure that the direct product group of any two groups is indeed a group.

This can easily be done by showing that the group multiplication as defined is indeed associative, then noting that (e_G, e_H) is the identity (where e_G is the identity of G and e_H is the identity of H), and lastly noting that for any $(a, b) \in G \times H$, (a^{-1}, b^{-1}) is its inverse (again, where $a, a^{-1} \in G$ and $b, b^{-1} \in H$).

Exercise 2.19. Show that for any groups G and H , $G \times H$ is a group (i.e., take the previous sentence and fill in the details by explicitly proving all of the things we said to "note").

Example 2.20. Consider the two groups \mathbb{Z}_3 and \mathbb{Z}_4 . Since these are both groups, we know we can form the product group $\mathbb{Z}_3 \times \mathbb{Z}_4$. As a set, this is just the Cartesian product of $\{0, 1, 2\}$ and $\{0, 1, 2, 3\}$, so there are twelve elements. $\{0, 0\}$ is the identity for this group, since $\forall (a, b) \in \mathbb{Z}_3 \times \mathbb{Z}_4$, we have that $(a, b) + (0, 0) = (a + 0, b + 0) = (a, b)$ and $(0, 0) + (a, b) = (0 + a, 0 + b) = (a, b)$. Just for fun, let's see what happens when we add the element $(1, 1)$ to itself over and over. We get $(1, 1) + (1, 1) = (1 + 1, 1 + 1) = (2, 2)$, then (adding $(1, 1)$ to this result) $(2, 2) + (1, 1) = (0, 3)$ (since in \mathbb{Z}_3 , $2 + 1 = 0$), then $(0, 3) + (1, 1) = (1, 0)$ (since in \mathbb{Z}_4 , $3 + 1 = 0$), then $(1, 0) + (1, 1) = (2, 1)$, then $(2, 1) + (1, 1) = (0, 2)$, then $(0, 2) + (1, 1) = (1, 3)$, then $(1, 3) + (1, 1) = (2, 0)$, then $(2, 0) + (1, 1) = (0, 1)$, then $(0, 1) + (1, 1) = (1, 2)$, then $(1, 2) + (1, 1) = (2, 3)$, then $(2, 3) + (1, 1) = (0, 0)$, and then adding $(1, 1)$ again gets us back to where we started. Wow, this little thing that we did just for fun ended up generating all twelve elements in our group! Namely, by simply adding $(1, 1)$ to itself over and over, we were able to "hit" every element in the group $\mathbb{Z}_3 \times \mathbb{Z}_4$. This means that every element $(a, b) \in \mathbb{Z}_3 \times \mathbb{Z}_4$ is of the form $(a, b) = (1, 1) + (1, 1) + \dots + (1, 1)$ for some finite number of terms in the sum. We therefore call the element $(1, 1)$ a **generator** of $\mathbb{Z}_3 \times \mathbb{Z}_4$, and we equivalently say that $\mathbb{Z}_3 \times \mathbb{Z}_4$ is **generated by** $(1, 1)$.

Exercise 2.21. Show that the product group $\mathbb{Z}_4 \times \mathbb{Z}_6$ is **not** generated by the element $(2, 3)$.

There are other group constructions that we can make, such as the **quotient group**, but we'll need some more machinery before introducing these. Let us start to develop that machinery.

2.6 Homomorphisms

Now that we've defined groups, subgroups, and a way to build new groups from old ones, let us now ask how to relate two groups to each other (following the same line of development that we did in the set theoretic case). Recall that when defining subgroups we wanted to retain some kind of "group quality" about subgroups, so that we could make a distinction between a subset of \mathbb{Z} such as $\{\dots, -9, -6, -3, 0, 3, 6, 9, \dots\}$ and $\{12, 25, -1\}$. We'd like to do the same with functions between groups. Let us consider the two groups \mathbb{Z} and \mathbb{Q} , both under addition. We could map \mathbb{Z} to the "copy" of \mathbb{Z} that "lies in" \mathbb{Q} by sending $a \in \mathbb{Z}$ to $a \in \mathbb{Q}$, or we could send $a \in \mathbb{Z}$ to $2 \times a \in \mathbb{Q}$, so that $0 \mapsto 0$, $1 \mapsto 2$, $2 \mapsto 4$, and so on. This seems to retain a certain structure. For we could also consider a map $\mathbb{Z} \rightarrow \mathbb{Q}$ that sends 3 to $\frac{17}{6}$, 0 to 65, -4 to $\frac{109}{57}$, and a whole bunch of other completely random associations. And clearly this second map doesn't have any kind of "structure" the way the first one does.

It turns out that an extremely useful kind of map between groups is what's called in the industry a **homomorphism**. This is a map which "preserves the group structure" of the original group. What this means is that if we first combine two elements in the domain of the map, and then map the product over to the other group, we always end up at the same element as we would have if we just sent the initial two elements over to the other group and combined them over there. I.e., if we combine the two elements and then map them, it's the same as mapping them and then combining them. This will become more clear after seeing the precise definition and a few examples. Let's start with the definition.

Definition 2.22. Let G and H be groups and let f be a function from G to H , so that $f : G \rightarrow H$. f is said to be a **homomorphism** if $\forall a, b \in G$, it is the case that $f(a \cdot b) = f(a) \cdot f(b)$.

Note that in the last expression of the definition, we use "." on both sides of the equal sign even though technically this is an abuse of notation. The reason for this is that on the LHS of the equal sign we're combining the elements **in** G first, and then mapping them by f , and therefore we're using G 's multiplication. On the RHS, however, we're first mapping a and b using f , and then we're combining them **in** H , so that we're actually using H 's multiplication rule. We use this notation out of laziness, as these sorts of concepts should eventually be gotten used to, with the implied meaning understood. Note that we could be even more lazy and write this as $f(ab) = f(a)f(b)$ with it "understood" that placing the elements next to each other means "combining them", and that we're combining them in the only groups for which the expression makes sense.

Example 2.23. Let's consider the groups \mathbb{Z} and \mathbb{Z}_8 , and some homomorphisms $f : \mathbb{Z} \rightarrow \mathbb{Z}_8$.

The first and perhaps most obvious homomorphism that we'll study is the one that sends each element in \mathbb{Z} to itself modulo 8. In other words, $f(a) = a \bmod 8$, where the "mod 8" operation is the same as that which we defined above, sending each integer to some integer greater than or equal to 0, and less than or equal to 7. So $f(6) = 6$ because $6 \bmod 8$ is still 6, but $f(25) = 1$ and $f(-3) = 5$. We need to show that this is indeed a homomorphism. I.e., we need to show that adding two integers and then taking that sum modulo 8 is the same as first taking the two integers modulo 8, and then adding them (modulo 8). Let's first try a couple examples, since if we can find some integers for which this isn't true, then we will know that this is not a homomorphism and we can save ourselves the trouble of trying to prove something that's not true.

Let's consider the integers 117 and -34 . $117 + (-34) = 83$ and $83 \bmod 8 = 3$, so we have $f(117 + (-34)) = f(83) = 3 \in \mathbb{Z}_8$. Now let's first take the "modulo 8" of both of these integers (i.e., let's first use f to send them into \mathbb{Z}_8). We have $f(117) = 5$ and $f(-34) = 6$, so $f(117) + f(-34) = 5 + 6$ where now our "+" sign is addition modulo 8. Since $5 + 6 = 11 = 3 \pmod{8}$, we indeed have that $f(117 + (-34)) = f(117) + f(-34)$ for this special case. However, we chose these integers sufficiently randomly that hopefully this is a general phenomenon and not dependent on our special choice of integers. Indeed it is, and we'll now see why.

Let's consider $f(a+b)$ for any $a, b \in \mathbb{Z}$. There is a unique pair of integers m and r such that $a+b = 8m+r$,

where $0 \leq r < 8$ (this is a fact that we have been assuming all along in our definition of modular arithmetic, and although it can be proved rigorously, we'll just take it on faith because it's rather obvious). By the definition of addition modulo 8, we have that $f(a + b) = r$. Now we also know that $a = 8k + r_1$ and $b = 8l + r_2$ for some unique set of integers k, l, r_1 , and r_2 , where we have $0 \leq r_1, r_2 < 8$. We then know that $f(a) = r_1$ and $f(b) = r_2$, so $f(a) + f(b) = r_1 + r_2 \pmod{8}$. We also know that we can write $r_1 + r_2$ as $r_1 + r_2 = 8M + R$, where M, R are a unique pair of integers such that $0 \leq R < 8$ (for example, if $r_1 + r_2 < 8$ then $M = 0$ and $R = r_1 + r_2$ (because we know that both r_1 and r_2 are greater than or equal to 0)). Thus we have $f(a) + f(b) = R$. But using our new expressions for a and b , we also know that $a + b = 8(k + l) + r_1 + r_2 = 8(k + l + M) + R$, which is of the exact same form as $a + b = 8m + r$ that we had above. Moreover, we know that this expression is unique, so it must be true that $k + l + M = m$ and $r = R$. Thus we have that $f(a + b) = R$, and therefore that $f(a + b) = f(a) + f(b)$. Since a and b were completely arbitrary integers, we've shown that this f is indeed a homomorphism. In fact, it can easily be seen that this homomorphism is always a surjective one (and never an injective one).

In fact, the above can easily be generalized to \mathbb{Z}_N for any N by simply replacing 8 with N everywhere. In particular, by making this replacement and following through the proof in the exact same way, we can show that the map $f : \mathbb{Z} \rightarrow \mathbb{Z}_N$ defined by sending each integer to itself modulo N is always a homomorphism.

Let's now consider one more homomorphism $g : \mathbb{Z} \rightarrow \mathbb{Z}_8$. Let's define g to take every even integer to $0 \in \mathbb{Z}_8$ and every odd integer to $4 \in \mathbb{Z}_8$. Thus, $g(a) = 0$ if a is even and $g(a) = 4$ if a is odd. This is clearly not surjective nor injective, but let's see that it is indeed a homomorphism. We need to show that $g(a + b) = g(a) + g(b)$ for any $a, b \in \mathbb{Z}$, and there are clearly three cases. The first case is when both a and b are even, the second case is when exactly one of them are even, and the third case is when they're both odd.

Case 1) a and b both even means that $a + b$ is even, so $g(a + b) = 0$ by definition of g . However, $g(a) = g(b) = 0$, and $0 + 0 = 0$, so we have that $g(a + b) = g(a) + g(b)$ in this case.

Case 2) We can without loss of generality suppose that a is even and b is odd. Then $a + b$ is odd since the sum of an even number and an odd number is odd. Thus $g(a + b) = 4$. Additionally, we know that $g(a) = 0$ and that $g(b) = 4$, so that $g(a) + g(b) = 0 + 4 = 4$, so in this case we also have that $g(a + b) = g(a) + g(b)$.

Case 3) If a and b are both odd, then their sum is even, so that $g(a + b) = 0$. Now, we know that $g(a) = g(b) = 4$, so that $g(a) + g(b) = 4 + 4 = 8$, but this is addition modulo 8, so we have that $8 = 0$. Thus $g(a + b) = g(a) + g(b)$ in this case as well. Since this covers all three cases, we've completed the proof that this is a homomorphism.

Exercise 2.24. Define a homomorphism from \mathbb{Z}_8 to \mathbb{Z}_2 and show that it's a homomorphism. Then define a homomorphism from \mathbb{Z}_3 to \mathbb{Z}_9 and show that it's a homomorphism.

Claim 2.25. Let G and H be arbitrary groups, let e_G and e_H denote the their respective identity elements, and let $f : G \rightarrow H$ be a homomorphism. Then $f(e_G) = e_H$.

Proof Since e_G is the identity for G , we know that $e_G \cdot e_G = e_G$. Thus, $f(e_G) = f(e_G \cdot e_G)$. Since we know that f is a homomorphism, we know that $f(e_G \cdot e_G) = f(e_G) \cdot f(e_G)$. Thus we have $f(e_G) = f(e_G \cdot e_G) = f(e_G) \cdot f(e_G)$. Now we simply multiply both sides by $f(e_G)^{-1}$, the inverse of $f(e_G)$, which we know exists because $f(e_G) \in H$ and H is a group. We then have

$$f(e_G)^{-1} \cdot f(e_G) = f(e_G)^{-1} \cdot f(e_G) \cdot f(e_G). \quad (2.2)$$

But $f(e_G)^{-1} \cdot f(e_G)$ is, by definition, e_H , and so we have

$$e_H = e_H \cdot f(e_G) = f(e_G). \quad \square \quad (2.3)$$

We've thus shown that a homomorphism from one group to another always maps the identity element of the domain to the identity element of the codomain. This allows us to note that whenever we're handed two

groups, we're always able to construct at least one homomorphism from one to the other, where we simply send every element of one group to the identity of the other. This certainly satisfies all of the requirements of a homomorphism, since it doesn't matter whether we compose and then map, or map and then compose, because we'll always just end up with the identity element of the codomain! We call a homomorphism that sends its entire domain to the identity element of its codomain the **trivial homomorphism**, for obvious reasons.

We can use the result of the previous claim to establish some other important results.

Claim 2.26. Let G, H, e_G, e_H , and f be as in Claim 2.6. Then $\forall a \in G, f(a^{-1}) = f(a)^{-1}$ (i.e., for any element in G , the image of its inverse is the same as the inverse of its image).

Proof We know from the previous claim that $e_H = f(e_G)$. Using the fact that for any element $a \in G$, $e_G = a \cdot a^{-1}$, we have the following chain of equalities, using the homomorphism property as well:

$$e_H = f(e_G) = f(a \cdot a^{-1}) = f(a) \cdot f(a^{-1}). \quad (2.4)$$

Now, multiplying both side on the left by $f(a)^{-1}$, we have

$$f(a)^{-1} \cdot e_H = f(a)^{-1} \cdot f(a) \cdot f(a^{-1}). \quad (2.5)$$

which gives, after noting that identity elements keep things fixed, $f(a)^{-1} = f(a^{-1})$, as desired. \square

Thus we've now shown that homomorphisms not only map identities to identities, but they also map inverses to inverses (i.e., they preserve the relationship between an element and its inverse). Note that all of this was derived only from the definition of a homomorphism, with no extra assumptions necessary.

Exercise 2.27. Let G and H be groups, and let $f : G \rightarrow H$ be a homomorphism. Show that $Im(f)$ is a subgroup of H . (Hint: what are the things you need to show for this? You need to show that it is closed under multiplication, that it contains the identity, and that if $a \in Im(f)$, then so is a^{-1} . Note that associativity follows from the fact that H is a group and so its multiplication is automatically associative).

Definition 2.28. Let G and H be groups and let $f : G \rightarrow H$ be a homomorphism. If f is bijective, then we call f an **isomorphism**.

Note that when two groups are isomorphic (i.e., when there exists an isomorphism between them), then they are essentially "the same" group. That is, they have "the same" structure, since whatever I can do with one group I can do with the other simply by mapping one to the other with the isomorphism. This is because I have a one-to-one "set-theoretic" correspondence that **also** preserves all of the structure of the group.

Exercise 2.29. Show that \mathbb{Z} and $\{\dots, -6, -4, -2, 0, 2, 4, 6, \dots\}$ are isomorphic (where both groups have their "abstract multiplication" being addition).

Exercise 2.30. Show that the rotations of a square (i.e., without the flips) and the group $\{0, 1, 2, 3\}$ (with modular arithmetic) are isomorphic.

Closely related to the image of any homomorphism is what's known as the "kernel" of a homomorphism. Given some homomorphism $f : G \rightarrow H$ between two groups, the kernel of f , often denoted as $Ker(f)$, is the set of elements in G that get sent to H 's identity element. Let us give this a precise definition:

Definition 2.31. Let G and H be groups and $f : G \rightarrow H$ a homomorphism. Then the **kernel** of f , denoted by $Ker(f)$, is the following set: $Ker(f) = \{a \in G | f(a) = e_H\}$, where $e_H \in H$ is the identity element.

Exercise 2.32. Show that for any homomorphism $f : G \rightarrow H$, $\text{Ker}(f)$ is a subgroup of f .

Example 2.33. In the previous example we looked at two explicit examples of homomorphisms. The first homomorphism $f : \mathbb{Z} \rightarrow \mathbb{Z}_8$ sent every integer to itself modulo 8, so that $f(-5) = 3, f(-4) = 4, f(-3) = 5, \dots, f(4) = 4, f(5) = 5, \dots, f(8) = 0, f(9) = 1, \dots$, and the second homomorphism $g : \mathbb{Z} \rightarrow \mathbb{Z}_8$ sent even integers to 0 and odd integers to 4. Let's now see what their kernels look like. For g , the kernel is easy—it's simply the set of all even numbers. This is because these are precisely the integers that get mapped to the identity element in \mathbb{Z}_8 , which is 0. The kernel of f is the set of all multiples of 8, since it is only these numbers that get sent to 0 under the map "itself modulo 8". Thus, $\text{Ker}(f) = \{\dots, -24, -16, -8, 0, 8, 16, 24, \dots\}$. It will prove useful to denote this set by $8\mathbb{Z}$ (the notation is suggesting "8 times \mathbb{Z} ", which is good because every element in $8\mathbb{Z}$ is "8 times some number in \mathbb{Z} ". We then have that $\text{Ker}(f) = 8\mathbb{Z}$. More generally, we can define the set $N\mathbb{Z}$ to be the set of all integer multiples of N , so that $4\mathbb{Z} = \{\dots, -12, -8, 4, 0, 4, 8, 12, \dots\}$ and $45\mathbb{Z} = \{\dots, -135, -90, -45, 0, 45, 90, 135, \dots\}$, and so on. It can then be easily checked that for the more general homomorphism $f : \mathbb{Z} \rightarrow \mathbb{Z}_N$ defined by sending every integer to itself modulo N (as we discussed at the end of the previous example), we'll have $\text{Ker}(f) = N\mathbb{Z}$.

Exercise 2.34. Define a homomorphism $f : \mathbb{Z}_{25} \rightarrow \mathbb{Z}_5$, show that it's a homomorphism, and find $\text{Ker}(f)$. Now find a homomorphism $g : \mathbb{Z}_5 \rightarrow \mathbb{Z}_{25}$, show that it's a homomorphism, and find $\text{Ker}(g)$.

There is one more group construction that we will look at, and we'll use homomorphisms and kernels for it, but before doing so we'll need to develop some more machinery. Let us therefore go on to discuss the incredibly important notion of an equivalence relation, which will be vital for our future discussion.

2.7 *Equivalence Relations

A very important notion in all of mathematics, and especially in group theory, is that of an equivalence relation. This is simply an extra structure that can be put on a given set, but it often "interacts nicely" with other structures that are placed on that set. It is a very useful tool for constructing new objects from old ones, as we'll be seeing shortly. Let us go ahead and get on with the definition and then explore some of its properties afterwards.

Definition 2.35. Let S be a set. An **equivalence relation** " \sim " on S is a relation between the elements of the set that satisfy the following three requirements:

- 1) reflexivity: $\forall x \in S, x \sim x$
- 2) symmetry: $\forall x, y \in S, x \sim y \Rightarrow y \sim x$
- 3) transitivity: $\forall x, y, z \in S, x \sim y \text{ and } y \sim z \Rightarrow x \sim z$.

For example, one can define the equivalence relation on the set \mathbb{Z} by fixing some element $a \in \mathbb{Z}$ and then saying that two elements in \mathbb{Z} are related by \sim if they're both less than or equal to a , or if they're both greater than a . I.e., if $x \leq a$ and $y \leq a$, then $x \sim y$, or if $x > a$ and $y > a$, then $x \sim y$. All three requirements of an equivalence relation are trivially satisfied here. In this case, \sim is reflexive because if x is less than or equal to a , then so is x , because $x = x$! (The same obviously holds if x is greater than a). The fact that \sim is symmetric is equally trivial, because if x and y are both less than or equal to a , then so are y and x . Transitivity also holds for the following reason. Suppose that $x \sim y$ and $y \sim z$. Then either both x and y are less than or equal to a , or they're both greater than a . Suppose they're both less than or equal to a . Then the fact that $y \sim z$ implies that z is related to a the same way y is, and since y is less than or equal to a , we know that z is less than or equal to a as well. But then both x and z are less than or equal to a , so we have that $x \sim z$. the same exact reasoning holds if we started by assuming that both x and y are greater than a .

Let us now consider a much less trivial example of an equivalence relation on \mathbb{Z} . Let us define \sim by

saying that $x \sim y$ if $x \bmod 5 = y \bmod 5$. I.e., two elements are equivalent if they're remainders after being divided by 5 are the same.

Exercise 2.36. Prove that this \sim is actually an equivalence relation.

We then have that $-5 \sim 0 \sim 5 \sim 10 \sim 15$, and so on, as well as $-4 \sim 1 \sim 6 \sim 11$, and $-3 \sim 2 \sim 7$, etc.

Now, one of the most important facts about equivalence relations is that they always partition the set on which they're defined into completely disjoint subsets in a very canonical way (the word "canonical" is widely used in mathematics, and it really just means "obvious" and/or "natural"). To be precise, let's make a clear definition of what we mean by a partition. First, though, we need to introduce the notion of a "labeling set". This is a straightforward concept, and one should try to avoid getting caught up in the abstraction. What we do is take one set (any set), which for some reason we usually label by Λ (pronounced "lambda"), and use it to label some other collection of objects. For example, we could say that $\Lambda = \{A, 4, PEANUT\}$, and then label three sets as S_A, S_4 , and S_{PEANUT} . Then, an expression like $\cup_{\alpha \in \Lambda} S_\alpha$ actually makes sense. All it means is to form the union of all the sets labeled by Λ , so that $\cup_{\alpha \in \Lambda} S_\alpha = S_A \cup S_4 \cup S_{PEANUT}$. Now, this is admittedly a pretty dumb choice of labeling set, and usually labeling sets are sets of numbers like \mathbb{N} or \mathbb{Z} . The point is, though, that these labeling sets can be **any** set (even \mathbb{R} , which is uncountably infinite), and all these labeling sets do for us is give us a completely abstract and general way to label other objects (usually collections of sets, which is how we'll use it in the next definition).

Definition 2.37. Let S be a set, and let $\{S_\alpha\}_{\alpha \in \Lambda}$ be a collection of subsets of S (where Λ is just some labeling set). We say that $\{S_\alpha\}_{\alpha \in \Lambda}$ is a **partition** of S if $S = \cup_{\alpha \in \Lambda} S_\alpha$ and if for any $\alpha, \beta \in \Lambda$ such that $S_\alpha \neq S_\beta$, we have $S_\alpha \cap S_\beta = \emptyset$.

What this definition says is that a partition of a set S is a collection of subsets of S that "cover" S and don't overlap without equaling each other. I.e., there are no "partial overlaps", only "all or nothing". Thus, $\{\{1, 2, 4, 5\}, \{3\}\}$ is a partition of $\{1, 2, 3, 4, 5\}$, as is $\{\{1, 2\}, \{3, 5\}, \{4\}\}$, whereas $\{\{1, 2, 4, 5\}, \{3, 4\}\}$ is not, nor is $\{\{1, 2, 4\}, \{3\}\}$.

We now make one more definition, and then prove something nice about all of these definitions.

Definition 2.38. Let (S, \sim) be a set with an equivalence relation \sim . For each $x \in S$, define the subset $S_x \subseteq S$ as $S_x = \{a \in S \mid x \sim a\}$. We then call S_x the **equivalence class** of x (under " \sim ").

For example, with the set \mathbb{Z} and the equivalence relation "remainder modulo 5", we have that $S_1 = \{\dots, -9, -4, 1, 6, 11, 16, \dots\}$ and $S_{13} = \{\dots, -7, -2, 3, 8, 13, 18, \dots\}$. It is important to note that $S_1 = S_6 = S_{11}$, and that in general two equivalence classes S_x and S_y will equal each other if and only if $x \sim y$. Note that this does not necessarily mean that $x = y$, and it is this fact that makes equivalence classes so interesting. We'll now show that the equivalence classes of a set under some equivalence relation always partition the set.

Let (S, \sim) be a set S together with an equivalence relation on it (any set, and any equivalence relation). Now for each $x \in S$, define the subset $S_x := \{a \in S \mid x \sim a\}$. Note that every element in S will be in **some** subset of this form, namely because by definition of an equivalence relation we always know that $x \sim x$, and so $x \in S_x$ for all $x \in S$. Now I want to show that every element $x \in S$ is in **precisely** one such subset of S , so that these equivalence classes really do partition our set S .

Let $x \in S$ and suppose that $x \in S_y$ for some $y \in S$. We want to show that $S_y = S_x$. We do this by first showing that $S_y \subseteq S_x$, and then that $S_x \subseteq S_y$, thus implying that $S_y = S_x$. To show that $S_y \subseteq S_x$,

we need to show that every element in S_y is also in S_x . So choose an element $a \in S_y$ arbitrarily. Then, by definition of S_y , we know that $a \sim y$. But, since $x \in S_y$, we also know that $x \sim y$. Thus, by the definition of an equivalence relation (namely, by using transitivity), we know that $a \sim x$. Thus $a \in S_x$ also. Since a was an arbitrary element of S_y , we now know that anything in S_y is also in S_x , so that $S_y \subseteq S_x$.

Conversely, pick any element $b \in S_x$. Then $b \sim x$. But we also have that $x \sim y$, so that we know $b \sim y$, and thus $b \in S_y$. Thus $S_x \subseteq S_y$, and thus $S_x = S_y$.

We've therefore established the following:

Proposition 2.39. If (S, \sim) is a set with an equivalence relation defined on it, then the equivalence classes of S form a partition of S .

2.8 *Conjugation

Let us now discuss a very important concept in group theory, known as conjugation. One very interesting class of objects to study are homomorphisms from a group to itself, namely, the various kinds of homomorphisms $f : G \rightarrow G$. For each element in G , we have a canonical type of homomorphism, known as conjugation. Namely, fix some element $g \in G$ and consider the map $f : G \rightarrow G$ defined by taking some element $h \in G$ and sending it to the element ghg^{-1} (note the lazy notation) in G . I.e., $f(h) = ghg^{-1}$ for all $h \in G$. For a more suggestive notation, we could write this map as $C_g : G \rightarrow G$ instead of $f : G \rightarrow G$, to remind us that it's "C"onjugation, and that we're doing the conjugation by the element $g \in G$.

As it stands, we've only defined C_g to be a function, and have not shown that it's a homomorphism. Let us do this now.

Proposition 2.40. Let G be a group and $C_g : G \rightarrow G$ be the map "conjugation by g ", so that $C_g(h) = ghg^{-1}$. Then C_g is a homomorphism.

Proof Let $a, b \in G$. Then $C_g(ab) = gabg^{-1}$. Now consider $C_g(a)C_g(b)$. We have $C_g(a)C_g(b) = (gag^{-1})(gbg^{-1}) = ga(g^{-1}g)bg^{-1} = gabg^{-1}$, so that $C_g(a)C_g(b) = C_g(ab)$, and thus C_g is a homomorphism. (Note that in the chain of equalities, we relied heavily on the associativity of multiplication in G). \square

We may now make some interesting definitions. If we fix some element $a \in G$, we can ask which other elements we can "get to" by conjugating a by various elements in G . Namely, we can define the set S_a of elements that can be written as gag^{-1} for some $g \in G$. For example, if we think of rotations and reflections, we can fix some transformation, call it "transformation a" (which might be some combination of rotations and reflections) and then ask which other transformations can be obtained by making the composite transformation "the inverse of transformation b, then transformation a, then transformation b", for some "transformation b". In symbols, we have $S_a = \{b \in G \mid \exists g \in G \text{ such that } b = gag^{-1}\}$.

It is then clear that for any $a \in G$, we know that $a \in S_a$, since we can conjugate a by the identity element e and return back to a , seeing as $ea e^{-1} = a$ (note/recall that the inverse of the identity is just the identity).

Example 2.41. If our group G is Abelian, then for any $a \in G$, we have that $S_a = \{a\}$. This is because for any $g \in G$, $gag^{-1} = gg^{-1}a = a$, where the first equality comes from the fact that G is Abelian and so we can commute g^{-1} past a . For example, we know that \mathbb{Z} is Abelian, and if we fix any integer a , then gag^{-1} corresponds to simply adding and subtracting the integer g , which clearly doesn't change a .

Example 2.42. Consider the set S_3 of permutations of 3 objects. This set has six elements:

$$\{e, (1, 2), (1, 3), (2, 3), (1, 2, 3), (1, 3, 2)\}.$$

Let's fix $a = (1, 2)$ and ask about the set of conjugates of $(1, 2)$, i.e., $S_{(1,2)}$. In other words, let's conjugate $(1, 2)$ by every element in S_3 :

$$\begin{aligned} e \cdot (1, 2) \cdot e^{-1} &= (1, 2) \\ (1, 2) \cdot (1, 2) \cdot (1, 2)^{-1} &= (1, 2) \cdot (1, 2) \cdot (1, 2) = (1, 2) \\ (1, 3) \cdot (1, 2) \cdot (1, 3)^{-1} &= (1, 3) \cdot (1, 2) \cdot (1, 3) = (1, 3) \cdot (1, 3, 2) = (2, 3) \\ (2, 3) \cdot (1, 2) \cdot (2, 3)^{-1} &= (2, 3) \cdot (1, 2) \cdot (2, 3) = (2, 3) \cdot (3, 1, 2) = (1, 3) \\ (1, 2, 3) \cdot (1, 2) \cdot (1, 2, 3)^{-1} &= (1, 2, 3) \cdot (1, 2) \cdot (3, 2, 1) = (1, 2, 3) \cdot (3, 1) = (2, 3) \\ (1, 3, 2) \cdot (1, 2) \cdot (1, 3, 2)^{-1} &= (1, 3, 2) \cdot (1, 2) \cdot (2, 3, 1) = (1, 3, 2) \cdot (2, 3) = (1, 3) \end{aligned}$$

Therefore the only elements that we can get by conjugating $(1, 2)$ are $(1, 2)$ itself, $(2, 3)$, and $(1, 3)$. Thus, $S_{(1,2)} = \{(1, 2), (2, 3), (1, 3)\}$, and there are elements in S_3 that we can't get to by conjugating $(1, 2)$. Moreover, it's important to note that $S_{(1,2)}$ isn't even a subgroup of S_3 , since it doesn't have the identity element in it and is therefore not a group.

Now, if two elements $a, b \in G$ are related by $a = gb g^{-1}$ for some $g \in G$, we can say that a and b are **conjugates** to each other. Note that we can say that they're mutually conjugate (as opposed to having to specify " a is conjugate to b " or " b is conjugate to a ") because if $a = gb g^{-1}$ for some $g \in G$, then it's also the case that $b = hah^{-1}$ for some $h \in G$. In fact, if we just take $h = g^{-1}$, then it is clear from $a = gb g^{-1}$ that $b = hah^{-1}$.

Exercise 2.43. Show this explicitly.

Exercise 2.44. Let G be a group. Show that S_e , the set of conjugates of the identity e , is nothing but $\{e\}$. I.e., show that $S_e = \{e\}$.

Now this might all be starting to smell like something vaguely familiar, because we now have a relation that is "symmetric" in the same way that an equivalence relation is. Namely, we have some relation such that if $a \sim b$, then $b \sim a$.

Proposition 2.45. The relation of conjugacy is an equivalence relation on any group G .

Proof It is clear that $x \sim x$ for any $x \in G$, because x is conjugate to itself via conjugation by the identity element. Now, we've already shown that if $x \sim y$ (i.e., if $x = gy g^{-1}$ for some $g \in G$), then $y \sim x$. Now all that remains to be shown is transitivity. Suppose $x \sim y$, and $y \sim z$. Then $x = gy g^{-1}$ for some $g \in G$, and $y = hzh^{-1}$ for some $h \in G$. But then, plugging in for y , we have that $x = g(hzh^{-1})g^{-1}$. Now, we notice that the inverse of gh is $h^{-1}g^{-1}$, since $gh(h^{-1}g^{-1}) = h^{-1}g^{-1}(gh) = e$, so that $(gh)^{-1} = h^{-1}g^{-1}$. We therefore have $x = (gh)z(gh)^{-1}$, which is of the required form $x = kzk^{-1}$, where now $k = gh \in G$. Thus $x \sim z$, and so transitivity has been established. Thus, the relation of conjugacy is an equivalence relation. \square

We can now take this one step further and make a similar construction using an entire **subset** of G . Namely, if we let $A \subseteq G$ be some subset of G , we can ask which elements in G can be "gotten to" by conjugating some element in A by some element in G . In other words, we can define the following subset:

$$T_A = \{h \in G | \exists g \in G, a \in A, \text{ such that } gag^{-1} = h\}. \quad (2.6)$$

(The subscript A is just there to help us remember that this set is defined relative to A , i.e., that information about A is contained in this definition).

Now, if we can make this construction for a subset of G , then we can make the same construction for a **subgroup** of G , simply because a subgroup is a special kind of subset. Thus, if we let H be a subgroup of G , then we can define the set (in analogy with T_A above) $T_H = \{h \in G \mid \exists g \in G, a \in H, \text{ such that } gag^{-1} = h\}$.

Now, in order to make our next definition, we need to use conjugation in essentially the only way that we haven't used it yet. Namely, suppose we take some subgroup H of G , and some particular element $g \in G$. We can then define the set of all elements of the form ghg^{-1} , with $h \in H$. Let us denote this set by gHg^{-1} . Note that this notation does not imply that there is any kind of multiplication going, and note that such an interpretation would be impossible because we haven't defined what it means for an element of G to multiply an entire subset (or subgroup) of G . This is, therefore, just notation for some subset of G . Namely, we have $gHg^{-1} = \{k \in G \mid k = ghg^{-1} \text{ for some } h \in H\}$. Namely, we're holding $g \in G$ fixed and letting $h \in H$ vary over all possibilities (in H), and taking all of the elements defined in this way and forming the set gHg^{-1} with them.

Now there are very special kinds of subgroups, called **normal subgroups** (oddly named, because they're not necessarily "normal"), that have the property that $gHg^{-1} \subseteq H$ for all $g \in G$. Namely, a subgroup H is called **normal** if it contains all of its conjugates. We now make this a definition.

Definition 2.46. A subgroup H of a group G is called **normal** if $gHg^{-1} \subseteq H$ for all $g \in G$, where we have adopted the notation gHg^{-1} used above.

Note that $gHg^{-1} \subseteq H$ most certainly does not mean that $ghg^{-1} = h$ for all $h \in H$. Rather, it means that for any $h \in H$ and for any $g \in G$, we have that $ghg^{-1} = h'$ where h' is just some other element in H .

Let us backtrack a bit and see where the notation gHg^{-1} really comes from.

Definition 2.47. Let H be an arbitrary subset (i.e., not necessarily a normal subgroup) of a group G . For every $g \in G$, we define a **left coset** gH of H to be $gH = \{k \in G \mid k = gh \text{ for some } h \in H\}$, and we define a **right coset** Hg of H to be $Hg = \{k \in G \mid k = hg \text{ for some } h \in H\}$.

Note that cosets are themselves subsets of G , and so we can take cosets of cosets, in the sense that if we have some (left) coset gH , then we can form the coset $k(gH)$ as follows. Since gH is the set of all elements of the form gh for some $h \in H$, the coset $k(gH)$ is simply the set of all elements of the form kgh for some $h \in H$. Again, we're holding k and g fixed, and letting h be any element in H . It should then be clear that $k(gH) = (kg)H$, which is an equality of **sets**. In other words, this says that if we first form the coset gH and then the coset $k(gH)$, then we get the same set as we would if we first multiplied k and g to get kg and then formed the coset $(kg)H$ with this element. This should be clear, since both are simply the set of all elements of the form kgh for some $h \in H$. We're thus allowed to unambiguously write either $k(gH)$, $(kg)H$, or kgH , and they'll all mean the same thing.

Now we can show that a normal subgroup N is simply a special kind of subgroup, such that for any $g \in G$, the left and right cosets gN and Ng are equal **as sets**. Let's see how this works. As usual, we want to show that for any $g \in G$, $gH \subseteq Hg$ and that $Hg \subseteq gH$.

Proposition 2.48. Let H be a normal subgroup of a group G , and let $g \in G$ be some arbitrary element. Then $gH = Hg$ (note that this is an equality of **sets**).

Proof Let $k \in gH$ be an arbitrary element. We want to show that $k \in Hg$, so that $gH \subseteq Hg$. $k \in gH \Rightarrow k = gh$ for some $h \in H$. Multiplying the right hand side of this equality by g^{-1} , we have that

$kg^{-1} = ghg^{-1}$. But since H is a normal subgroup, we know that $ghg^{-1} \in H$ (because $gHg^{-1} \subseteq H$ when H is normal). Thus, we know that $ghg^{-1} = h'$ for some $h' \in H$. Accordingly, we have that $kg^{-1} = h'$, and so multiplying the right side of this equality by g , we find that $k = h'g$. But this, by definition, means that $k \in Hg$. Thus, any element in gH is also in Hg , and so $gH \subseteq Hg$.

Now we want to show that $Hg \subseteq gH$, and the logic is the exact same as above. Namely, take $k \in Hg$ arbitrary. Then $k = hg$ for some $h \in H$. Then $g^{-1}k = g^{-1}hg$. Now, we know that $g^{-1}hg \in H$, since H is normal (where we're using the fact that if $gHg^{-1} \subseteq H$ for all $g \in G$, then it is also the case that $g^{-1}Hg \subseteq H$ for all $g \in G$, since g^{-1} is another perfectly good element of G). Thus $g^{-1}hg = h'$ for some $h' \in H$, and so we have that $g^{-1}k = h'$. Multiplying the left of this equality by g , we get that $k = gh'$, which is precisely the statement that $k \in gH$. Thus anything in Hg is also in gH , so $Hg \subseteq gH$. Since we have that both $gH \subseteq Hg$ and $Hg \subseteq gH$, we have shown that $gH = Hg$. \square

We've thus shown that if N is a normal subgroup of a group G (we use " N " instead of " H " now), then for any element $g \in G$, the sets gN and Ng are equal as sets. Note that this does not mean that for every $n \in N$ it is the case that $gn = ng$, but rather that for any $n \in N$, there is some $n' \in N$ such that $gn = n'g$, and vice versa. Now, it might be the case that $n = n'$, so that $gn = ng$ (as, say, would happen if $g = e$, the identity, for then $en = ne$), but this is not generally the case.

What this allows us to do, however, is form a new set using the normal subgroup N . Let us employ the notation $N \trianglelefteq G$ for the expression " N is a normal subgroup of G ". Namely, we can take each $g \in G$ and view the set gN as a single element of some new set. In doing so, we must first ask which elements in G give rise to the same set. I.e., when is it the case that two different element $g, g' \in G$ give rise to the same sets, so that $gN = g'N$?

Exercise 2.49. Let N be a subgroup of G . Show that $gN = g'N \Leftrightarrow g = g'n$ for some $n \in N$. Similarly, show that $Ng = Ng' \Leftrightarrow g = ng'$ for some $n \in N$. (This one's a little tricky). (Note that for normal subgroups, these two facts are equivalent).

We now define the set whose elements are those subsets of G that are of the form gN , where $g \in G$ and N is a subgroup of G (the above exercise shows that there are several different ways of writing down the set gN , since there are different elements of G that give the same coset. Thus, we have to remember that we're only dealing with a particular representation of the set, and we need to make sure that any constructions we make with these subsets is independent of which representative we use). Let us denote this set by G/N . I.e., we have $G/N = \{gN | g \in G\}$. Thus, the **set** gN is a single element of G/N .

Now, the important part of all of this is that, when N is a normal subgroup of G , the set G/N is actually a group. If N weren't a normal subgroup, we could still define the **set** G/N , but it wouldn't be a group. Let us now see how we can give group structure to the set G/N .

The first thing we need to do is define a multiplication on the elements in G/N . So let gN and $g'N$ be two elements $\in G/N$. **Define** $gN \cdot g'N := (gg')N$. I.e., the product of the coset gH with $g'H$ is the coset $(gg')H$, where (gg') denotes the normal multiplication within G . We now need to check that this is well-defined, i.e., that it doesn't matter which representatives for gN we use. So let's pick some other elements $k, k' \in G$ such that $kN = gN$ and $k'N = g'N$. Then we know that $k = gn$ for some $n \in N$, and that $k = g'n'$ for some $n' \in N$. Then we have $kN \cdot k'N = (kk')N = (gng'n')N$. But we know that for any $n \in N$ and $g \in G$, $gnN = gN$, thus we have $(gng'n')N = (gng')N$. Then using the fact that N is normal, we know that $g'N = Ng'$, so we have $(gng')N = gnNg' = gNg'$, and then using $Ng' = g'N$ again, we have $gNg' = gg'N$, so that following all of these equalities along we indeed have $kk'N = gg'N$, meaning that the result of multiplying cosets is independent of which representatives we use. (Yikes, this may have been

a confusing paragraph, and if so then take a deep breath and re-read it slowly and try to understand **why** each step was necessary—if it wasn't confusing then all the better).

We're now finally in a position to show how G/N is a group when N is a normal subgroup. We've defined the multiplication on this set, and it is easy to show that this multiplication is associative.

Exercise 2.50. Show that the multiplication on G/N defined above is associative.

Thus, all we need to do is show that there is an identity element, and that there are inverses for every element in G/N . I claim that the identity element in G/N is simply N , i.e., the coset formed by "multiplying" N by the identity e or by any element in N (since $nN = N$ for any $n \in N$). To see this, let gN be any coset $\in G/N$, and let us multiply it on the left by N . We then have $N \cdot gN = eN \cdot gN = (eg)N = gN$, and the same logic works if we multiply it on the right by N . Thus we've found our identity element. Now we need to find an inverse for every coset gN . But this is easy, since it just follows from the group structure itself. Namely, the inverse of gN is simply $g^{-1}N$, i.e., the coset formed by multiplying N by the inverse of g . To see this, we calculate: $gN \cdot g^{-1}N = (gg^{-1})N = eN = N$ and $g^{-1}N \cdot gN = g^{-1}gN = eN = N$. Thus it is the case that every coset has an inverse, and therefore that G/N is a group.

(Note that we **require** N to be a normal subgroup because otherwise the multiplication on the group is not well-defined, in that it depends on the particular representation of the cosets. This sort of representation-dependence is not allowed for a well-defined structure.)

We therefore have the following definition.

Definition 2.51. Let G be a group and $N \trianglelefteq G$. Then G/N , the set of cosets, is a group with the group structure as described above. We call G/N the **quotient group**, or the **quotient** of G by N .

Note that in the definition above we didn't have to specify whether this was the set of left or right cosets, because when N is normal (as it is in this definition), the left cosets are precisely the same as the right cosets.

Exercise 2.52. Let G be an Abelian group and show that every subgroup of G is a normal subgroup.

Example 2.53. One of the most common quotient groups is actually a group that we've already seen before, namely \mathbb{Z}_N , the set $\{0, 1, 2, 3, \dots, N-1\}$ with addition modulo N . I claim that this group can actually be viewed as a quotient group of the integers, in a very natural way. Let's see how this works for the particular group \mathbb{Z}_6 .

Take the integers \mathbb{Z} with their usual addition. This is an Abelian group, and so by the above exercise we know that every subgroup of it is normal, and we therefore know that we can form quotient groups from subgroups of \mathbb{Z} . Consider the subgroup of \mathbb{Z} formed by the set of all integers that are multiples (positive and negative) of 6. This is easily seen to be a subgroup (it has 0, as well as inverses). Let's denote this subgroup, as we did before, by $6\mathbb{Z}$. Then we have $6\mathbb{Z} = \{\dots, -12, -6, 0, 6, 12, 18, \dots\}$, and we know this is a normal subgroup of \mathbb{Z} . We therefore can ask about the quotient of \mathbb{Z} by $6\mathbb{Z}$, which, using the above notation for quotient groups, is denoted by $\mathbb{Z}/6\mathbb{Z}$.

Recall that the **elements** of $\mathbb{Z}/6\mathbb{Z}$ are themselves sets, and in particular they are the cosets of $6\mathbb{Z}$. Let us therefore see what these cosets look like. Let us consider the element $1 \in \mathbb{Z}$, and let us denote the coset formed by 1 as $1 + 6\mathbb{Z}$. Note that $1 + 6\mathbb{Z}$ is an entire set of elements in \mathbb{Z} , and these elements are precisely those that are of the form $1 + x$ where $x \in 6\mathbb{Z}$. Thus, $1 + 6\mathbb{Z}$ is the exact analogue of the more abstract expression gN for some $g \in G$, where now our normal subgroup N is $6\mathbb{Z}$. The reason we use this notation now is that the "abstract group multiplication" of \mathbb{Z} is really just addition, and this notation reminds us of that.

So what is the coset $1 + 6\mathbb{Z}$? Well, as mentioned above, it's the set of elements that are of the form $1 + x$ where x is some multiple of 6. Thus, $1 + 6\mathbb{Z} = \{\dots, -17, -11, -5, 1, 7, 13, 19, \dots\}$. Note that this is **not** a subgroup. Similarly, we can form the coset $2 + 6\mathbb{Z}$, which would look like $2 + 6\mathbb{Z} = \{\dots, -16, -10, -4, 2, 8, 14, 20, \dots\}$, i.e., the set of all elements of the form $2 + x$ where x is some multiple of 6. Similarly, we can form $3 + 6\mathbb{Z}$, $4 + 6\mathbb{Z}$, and $5 + \mathbb{Z}$. Note, though, that if we form $6 + 6\mathbb{Z}$, we're left with the set of elements of the form $6 + x$ with x a multiple of 6, which is just the set of multiples of 6 again! Thus, $6 + 6\mathbb{Z} = 0 + 6\mathbb{Z} = 6\mathbb{Z}$. It is a simple matter to check that $7 + 6\mathbb{Z} = 1 + 6\mathbb{Z} = 13 + 6\mathbb{Z} = \dots$. I.e., if we form some coset $a + 6\mathbb{Z}$, then this coset is the same as the coset we form by adding to a any multiple of 6 (i.e., $a + 6\mathbb{Z}$ is the same coset as $(a + 6b) + 6\mathbb{Z}$ for any $b \in \mathbb{Z}$, where the addition in the parentheses $a + 6b$ is just normal addition).

Thus there are precisely 6 elements in our quotient group: $\mathbb{Z}/6\mathbb{Z} = \{6\mathbb{Z}, 1 + 6\mathbb{Z}, 2 + 6\mathbb{Z}, 3 + 6\mathbb{Z}, 4 + 6\mathbb{Z}, 5 + 6\mathbb{Z}\}$. (We could have chosen other representative for these cosets, like $7 + 6\mathbb{Z}$ instead of $1 + 6\mathbb{Z}$, but we chose the simplest and most obvious representatives here.) Now we know that since $6\mathbb{Z}$ is a normal subgroup, there is a well-defined "abstract multiplication" on these cosets. Namely, it is that $(a + 6\mathbb{Z}) + (b + 6\mathbb{Z}) = (a + b) + 6\mathbb{Z}$. So for example we have $(1 + 6\mathbb{Z}) + (3 + 6\mathbb{Z}) = 4 + 6\mathbb{Z}$. (Note that I am slightly abusing notation because I'm using the same "+" sign for two different groups, namely, the integers and the quotient group. This should hopefully not cause confusion, however, as long as we remember that there really are two different group operations involved here).

But now what happens when we add, say, $5 + 6\mathbb{Z}$ to $4 + 6\mathbb{Z}$? Well, we get $(5 + 6\mathbb{Z}) + (4 + 6\mathbb{Z}) = 9 + 6\mathbb{Z}$. But we know that $9 + 6\mathbb{Z}$ is the same coset as $3 + 6\mathbb{Z}$, so we can just as well write $(5 + 6\mathbb{Z}) + (4 + 6\mathbb{Z}) = 3 + 6\mathbb{Z}$. Look familiar? Perhaps it will look more familiar if we adjust our notation. Let us denote the cosets $6\mathbb{Z}, 1 + 6\mathbb{Z}, 2 + 6\mathbb{Z}, 3 + 6\mathbb{Z}, 4 + 6\mathbb{Z}$, and $5 + 6\mathbb{Z}$ simply as $0, 1, 2, 3, 4, 5$, respectively. Then our quotient group is simply $\{0, 1, 2, 3, 4, 5\}$ with the same exact "abstract multiplication" as our clock arithmetic group, which in this case is simply addition modulo 6. We've therefore given precise and rigorous meaning to the phrase "clock arithmetic", which we now see to be nothing but the operation on the quotient group of \mathbb{Z} by some normal subgroup of the form $N\mathbb{Z}$ for some integer N (in the case we just considered, $N = 6$). Accordingly, we can identify "addition modulo n " as simply the addition on the quotient group $\mathbb{Z}/n\mathbb{Z}$.

Exercise 2.54. Repeat the above construction for $\mathbb{Z}/10\mathbb{Z}$. Write out what the cosets look like. How many cosets are there? Show/convince yourself that addition of elements in this quotient group is the exact same as addition modulo 10.

This ends our preliminary discussion on groups. We've given their definition, seen how to map groups to each other (homomorphisms), found substructure in groups (subgroups), and seen how to construct new groups from old ones (product groups and quotient groups). There are tons more that can be said about groups, and I've listed some references for further study in group theory at the end of this text. We'll see more examples of groups in what is to come, but in order to do so we need to introduce some more mathematical structures.

Chapter 3

Real Vector Spaces, Linear Maps, Matrices

3.1 Introduction

Two of the most fundamental and important mathematical structures that arise in Nature are groups, and vector spaces. Additionally, the notion of "linearity" is absolutely essential in our mathematical description of Nature. Seeing as we've already addressed groups, let us now turn to the later two topics (which are closely related).

3.2 Real Vector Spaces

In order to talk about general vector spaces, we would need to introduce a lot more machinery than we currently have available to us. Therefore, let us restrict ourselves to talking about **real** vector spaces (the meaning of this will become clear shortly).

We are likely already familiar with the notion of "a line", or "the plane," or of "three-dimensional space", but perhaps the abstract and rigorous mathematical definition of these ideas are not so familiar. At the heart of these ideas are the real numbers \mathbb{R} . These numbers form what we refer to as "the continuum", because they are distributed "continuously". Again, the rigorous definition of what this means will have to be skipped, as it requires more machinery. Intuitively, though, this simply means that one can zoom in infinitely far on the real numbers and never see a "gap" in the numbers. This can be contrasted with the integers, for there most certainly is a "gap" between, say, 1 and 2. Even the fractions are not "continuous" in the sense that the real numbers are, because the fractions can be "counted" (in the sense of lecture 1) whereas the real numbers are uncountable. Let me leave these technicalities behind and simply trust that the intuition for real numbers is in place—namely, that the real numbers form "the number line".

Using the real numbers, we can model "the line", "the plane", "three-dimensional space", and even higher dimensional spaces, as follows. A single copy of \mathbb{R} models the line, where we can view "walking up and down the line" as simply increasing or decreasing the value at which we're sitting in \mathbb{R} . Similarly, we can model the plane as simply the Cartesian product of \mathbb{R} with itself, i.e., $\mathbb{R} \times \mathbb{R}$. Each element $(a, b) \in \mathbb{R} \times \mathbb{R}$ is then a description of one's location, where we can perhaps define the first element to denote our "left-right" location, and then second element to denote our "up-down" location. Similarly, three-dimensional space can be modeled by $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$, where each element is written (a, b, c) and we can perhaps define a to denote our "left-right" position, b to denote our "up-down" position, and c our "forward-backward" position.

But surely whatever we use to model space like this, we would also want to model motion in that space. This is precisely (one example of) what we use vector spaces for (which I'll define shortly). First, let's

develop some intuition. Suppose I'm sitting in the place at location (a, b) . We must first appreciate the fact that this already implies that we know where the point $(0, 0)$ is, for in order for you to know where I am when I give you my coordinates (a, b) , we must both know where $(0, 0)$ is to begin with. Then, when I say that I'm sitting at location (a, b) , I mean that I am sitting a units horizontal to $(0, 0)$, and b units vertical to $(0, 0)$. We could first define our coordinates so that if a is positive, it is known that I'm sitting to the right of $(0, 0)$, and if a is negative, then I am to the left of $(0, 0)$. Similarly, we could say that if b is positive then I'm "above" the point $(0, 0)$ and if it's negative then I'm below $(0, 0)$.

Now suppose I move 1 unit to the right, and 1 unit up, from my initial location (a, b) . I'll clearly then be at the location $(a + 1, b + 1)$. If I had moved 1 unit to the right and 1 unit down, then I'd be at the location $(a + 1, b - 1)$, etc. We then notice that my motion throughout the plane is very similar to the usual addition defined on the product group $\mathbb{R} \times \mathbb{R}$, where two elements (a, b) and (c, d) are added "component-wise", so that $(a, b) + (c, d) = (a + c, b + d)$.

Notice also that my coordinates are dependent on how I'm measuring things, i.e., what units I'm using (notice that I used the word "unit" a lot in the above discussion—this will become more clear when we do some physics later). For example, suppose I said that my coordinates were $(5, 3)$. You'd not only have to know where $(0, 0)$ is for that to make any sense, but you'd also have to know what units I'm using. Am I 5 feet to the right of $(0, 0)$, or am I 5 miles to the right? Clearly this makes a difference! But suppose I tell you that I'm at $(5, 3)$ and that I'm using yards (there are 3 feet to each yard) to describe my location. Then suppose that you only know how to deal with feet, and you therefore want to convert my location into the more familiar units of feet. What would you do? You would have to multiply **both** of my coordinates by 3, since there are 3 feet in each yard. Thus, in your units, I'd be sitting at $(15, 9)$.

We therefore see that in order to meaningfully talk about space (lines, planes, etc.), we need a way of adding things together, as well as a way of multiplying these things by numbers. This might all seem extremely obvious, and hardly requiring any kind of abstract formalism. However, as we've seen, abstraction is a powerful tool that lets us talk about things that aren't necessarily "obvious" or "intuitive". In fact, as we'll see when we talk a bit about quantum mechanics, we truly and genuinely need abstract vector spaces, which in no way resemble the vector spaces of "physical space" (like lines, planes, or 3-D space). Additionally, one of the immediate benefits of defining vector spaces abstractly is that we can immediately generalize them to higher and higher dimensions, which makes vector spaces extremely fundamental for physics where we often want to describe many-dimensional concepts (that have nothing to do with physical space).

Before I write the entire definition of a real vector space (the word "real" should not be viewed as "the opposite of fake", but rather as a manifestation of the fact that we'll be dealing with vector spaces that are intimately related to real numbers (when we introduce complex numbers, we'll also introduce complex vector spaces)), let me break it up into smaller, more digestible pieces. First and foremost, a vector space V is a set which is also an Abelian group, and we call the "abstract group multiplication" on this Abelian group "addition". We want our vector spaces to be an Abelian group under this addition to reflect how our elements in $\mathbb{R} \times \mathbb{R}$ added. Namely, we had $(a, b) + (c, d) = (a + c, b + d)$, which is clearly an Abelian operation (since addition of real numbers is Abelian). Since the set is also a group, we know that this operation of addition is associative, so that $u + (v + w) = (u + v) + w$ for all $u, v, w \in V$, and we can therefore unambiguously write $u + v + w$ without any parentheses.

Let us now turn to defining how we can multiply the elements of V by real numbers. We saw in the above discussion that we want such an operation, and now let us endeavor to make this an abstract and rigorous construction. In particular, what we want is a way of taking a real number a and a vector (a

"vector" is simply a fancier word for an element in the vector space, i.e., an element of the set V) and creating a new vector that we can identify as "the vector v times the number a ". After all, we saw that in a very clear sense the vector $(15, 9)$ was simply the vector $(5, 3)$ multiplied by the number 3. More abstractly, we could have thought of "3" as a function in its own right, mapping every vector to "3 times that vector", so that not only does it send $(5, 3)$ to $(15, 9)$, but it also sends $(1, 0)$ to $(3, 0)$, $(0, 4)$ to $(0, 12)$, $(5.2, -2.4)$ to $(15.6, -7.2)$, and so on. Moreover, we get such a "function" for each number in \mathbb{R} , when we view a real number as a function from $\mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \times \mathbb{R}$. In this case, and real number a maps any vector (x, y) to (ax, ay) .

When we put it like this, though, it becomes a bit more clear how we should define our abstract multiplication by real numbers, which we'll call "scalar multiplication". It turns out that what we want of our set V , when it comes to scalar multiplication, is a function from $\mathbb{R} \times V \rightarrow V$, which takes a real number and a vector, and gives us back a vector. Let us denote this function simply by a dot " \cdot ", so that if $a \in \mathbb{R}$ and $v \in V$, we have that $a \cdot v \in V$. If we really want this function to represent what we mean by "a vector times a number", then we need to add some extra requirements to it, since there are lots and lots of different functions $\mathbb{R} \times V \rightarrow V$, and many of them will not reflect what we want them to. For example, if $V = \mathbb{R} \times \mathbb{R}$ as in our above example of the plane, we could define the function $\mathbb{R} \times (\mathbb{R} \times \mathbb{R}) \rightarrow \mathbb{R} \times \mathbb{R}$ which takes $(a, (x, y))$ to $(ax + 5, y - 4)$, but this obviously does not reflect what we want when we say "the vector (x, y) multiplied by a ".

One of the things we most certainly want when describing "a vector multiplied by a number" is that if we add two vectors together and then multiply them, that should be the same as first multiplying them individually and then adding them. I.e., if I move 3 yards to the right and 1 yard up, and then you convert to feet, that should be the same as if you first converted to feet, and then I moved the same distance. In other words, the order of vector addition (the Abelian group addition) and scalar multiplication doesn't matter. This is also called "distributivity", and the corresponding law is the "distributive law". Thus, we require of our scalar multiplication that for any $a \in \mathbb{R}$ and for any $u, v \in V$, we have that $a \cdot (u + v) = a \cdot u + a \cdot v$. (Note that sometimes we'll be lazy (just like in the case of groups) and denote scalar multiplication by just putting the scalar and the vector next to each other when it's clear which is the scalar and which is the vector, so that $a \cdot v$ becomes av when $a \in \mathbb{R}$ and $v \in V$).

Another very obvious requirement that we would ask of our scalar multiplication is that if we multiply a vector by a number, and then multiply it again by another number, the final vector that we're left with should be the same as that which we'd get if we first multiplied our two numbers together, and then multiplied our vector by that number. Namely, if we took some vector $v \in V$ and multiplied it by $a \in \mathbb{R}$ to get the vector $a \cdot v$ (recall that this vector can also be denoted by av), and then we took this vector and multiplied it again by another element $b \in \mathbb{R}$ to get the vector $b \cdot (av)$ (or equivalently $b(av)$), then we would like this final vector to be the same vector that we'd get if we first multiplied a and b together as usual, and then multiplied the vector v by the product ba . I.e., we want $(ba) \cdot v = b \cdot (av)$, and we want this to hold $\forall v \in V$ and $\forall a, b \in \mathbb{R}$.

Lastly, we want the addition of vectors to be suitably compatible with the addition of real numbers. Namely, it is natural to require that if we take a vector $v \in V$ and multiply it by the sum of two real numbers $a + b$ to get the vector $(a + b) \cdot v$, then we should get the same vector that we'd get if we first multiplied v by a to get av and added it to bv . I.e., we want $(a + b)v = av + bv$, and we want this to hold $\forall v \in V$ and $\forall a, b \in \mathbb{R}$. Note that there is a slight abuse of notation here. Namely, on the left hand side of the equation the "+" sign in " $(a + b)$ " is referring to the usual addition of real numbers, whereas on the right hand side the "+" sign in " $av + bv$ " is referring to the addition of our abstract vectors (the addition of these vectors is that given to us by the fact that V is an Abelian group).

We now have everything that we need to define a real vector space rigorously. This is a somewhat long definition, with seemingly a lot of axioms, but when we recall the above discussion we notice that all of these axioms are extremely natural and obvious—they just may seem obscure when phrased in this abstract language (but that's what we're trying to get used to here!).

Definition 3.1. A **real vector space** V is an Abelian group $(V, +)$ equipped with a special function $\cdot : \mathbb{R} \times V \rightarrow V$ called "scalar multiplication" which, $\forall a, b \in \mathbb{R}$ and $\forall v, w \in V$, satisfies the following requirements:

- 1) $a \cdot (v + w) = a \cdot v + a \cdot w$ (distributivity of scalar multiplication over vector addition)
- 2) $(a + b) \cdot v = a \cdot v + b \cdot v$ (distributivity of scalar addition over scalar multiplication)
- 3) $(ab) \cdot v = a \cdot (bv)$ (compatibility of scalar multiplication with regular multiplication).
- 4) $0 \cdot v = 0$ (compatibility of $0 \in \mathbb{R}$ with $0 \in V$)
- 5) $1 \cdot v = v$ (just an extra condition that makes vector spaces much nicer)

There are a couple more things to note. First, the notation " $(V, +)$ " is used here to denote what the group operation on V is, and we choose to use the symbol "+" because we like calling this operation "vector addition". Secondly, note that since $(V, +)$ is an Abelian group, we know that there is an identity element and that there are inverses for all elements. Namely, there is a unique element that, when added to any vector in V , gives us that element right back. We denote this element by "0" and call it the **zero vector**, but it is important to note that this element might not be the number 0 itself. Rather, it is the abstract identity element in this abstract group, and we choose to denote this element by the symbol "0" simply because we call this operation "vector addition" and this coincides with the identity element of usual addition (namely, the actual number 0). We also denote the inverse of any vector $v \in V$ by $-v$, so that $v + (-v) = 0$. We can use the more standard shorthand and write this as $v - v = 0$ which, when written this way, may seem extremely obvious. It is for this reason that we must simply remember that the expression $v - v = 0$ is really the statement "any vector plus its inverse gives the identity vector". Finally, it should be noted that conditions 4) and 5) were not previously discussed, and are primarily there for simplicity's sake. In almost all examples of vector spaces that we'll see, those conditions are extremely obvious and natural. Requiring that those hold make it so that vector spaces have some very nice properties, some of which we'll derive now.

Exercise 3.2. Let V be a real vector space. Using only the axioms in the definition of a real vector space, show that $a \cdot v = 0 \Leftrightarrow a = 0$ or $v = 0$ (or both). Make sure you're clear on the distinction between the zero "scalar" (i.e., the $0 \in \mathbb{R}$) and the zero vector (i.e., the $0 \in V$).

Exercise 3.3. Let $a \in \mathbb{R}$ and $v \in V$. Show that the additive inverse (i.e., the inverse using the group structure of V) of av is $(-a)v$.

It is important to note that although these properties may seem obvious, they're in no way "guaranteed" by the definition of a vector space, and we need to derive them to be sure of their truth. As always, we want our definitions to have as few axioms as possible, since it's always better to prove things from axioms than it is to simply insert more axioms. We note again that we can only prove properties about any mathematical structure using only the information that is either in the defining axioms themselves, or in the set of properties that have already been rigorously proven from those axioms.

We'll end this subsection with a very important example, and this example should be held in the back of one's mind as we continue to explore vector spaces. In fact, this example is in many ways what motivated the definition of a vector space in the first place, and it turns out that a very large class of vector spaces are all, in a very general way, identical to this example.

Example 3.4. Consider the set $\mathbb{R} \times \mathbb{R}$, and call this set V (for consistency with the above abstract formalism). It turns out that this set is a vector space, but in order to see this we need to first define what it means to add two elements in this set together. After all, this set does not "come with" a rule book for how to add things together—that's something that we need to come up with.

We know that we can represent any element in V as (a, b) where $a, b \in \mathbb{R}$. Thus, to represent two such elements, we can write $(a, b), (c, d)$, with $a, b, c, d \in \mathbb{R}$. Now, there's an obvious way to define an addition function on this set, and that is to just add the elements together "component-wise", or "slot-wise". That is, we **define** addition on this set to be that function which takes the two vectors (a, b) and (c, d) and sends them to the vector $(a + b, c + d)$. Lastly, we need to define a scalar multiplication as well. Again, though, an obvious choice suggests itself. Namely, we define our scalar multiplication in a "slot-wise" manner as well, so that it takes a real number $k \in \mathbb{R}$ and a vector $(a, b) \in V$ and sends them to the vector (ka, kb) . I.e., it's just slot-wise multiplication. We leave it as a trivial (but important) exercise to check that all of the axioms of a vector space are satisfied by these definitions.

Finally, we note that this definition can immediately be generalized to the Cartesian product of N copies of \mathbb{R} in a very straightforward way. Namely, if we let N be some (finite) integer greater than zero, then we can form the set $V = \mathbb{R} \times \dots \times \mathbb{R}$, where there are N copies of \mathbb{R} in this product. Then any vector in V can be represented as (a_1, a_2, \dots, a_N) , where each $a_i \in \mathbb{R}$. Then if we have two such vectors, namely $(a_1, a_2, \dots, a_N), (b_1, b_2, \dots, b_N)$, we can add them together to form the vector $(a_1 + b_1, a_2 + b_2, \dots, a_N + b_N)$. Also, scalar multiplication goes through in exactly the same way, with $k \cdot (a_1, a_2, \dots, a_N) = (ka_1, ka_2, \dots, ka_N)$. Again, it is very easy (but important) to check that these more general definitions still satisfy the axioms of a vector space.

It is important to note that the above example shows us how we can (and should) view vectors geometrically as arrows with some length and pointing in some direction, all with their tails glued to a fixed point (the zero vector). Since we can draw pictures of the sets \mathbb{R} , $\mathbb{R} \times \mathbb{R}$, and $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$ (i.e. the line, the plane, and three-dimensional space), we can thus get a good visualization of these vector spaces and how vector addition and scalar multiplication act. A useful way to get such a visualization is as follows. Suppose we take $\mathbb{R} \times \mathbb{R}$, so that we can draw our pictures in the plane. For any vector $(a, b) \in \mathbb{R} \times \mathbb{R}$, we choose to draw an arrow from the origin $(0, 0)$ to the point (a, b) , where the tip of the arrow is hitting (a, b) . With this visualization, a vector **is** an arrow with a length and a direction. Note that this visualization can be carried over to $\mathbb{R} \times \mathbb{R} \times \mathbb{R}$ as well. Then, when we add two vectors, we're simply moving the "back" of one of the two arrows (i.e., the side of the vector touching the origin) to the tip of the other arrow, while keeping both arrows "pointing in the same direction". We then draw an arrow from the origin to the point at the tip of this new two-arrow system, and this new arrow corresponds to precisely the arrow of the vector that is the sum of the original two.

For example, by adding the arrow corresponding to $(1, 1)$ (which points "to the right and up" (both by a unit of 1)) to the arrow corresponding to $(-1, 1)$ (which points "to the left and up" (both by a unit of 1)), we get an arrow that points straight up for two units, and this precisely corresponds to the vector $(0, 2)$, which is indeed the sum of the above two vectors. Note that we've chosen our axes so that the "first slot" corresponds to the horizontal axis, and the second slot to the vertical axis. By playing around with this visualization, it should be clear that it doesn't matter which arrow is added to which, and this is a reflection of the fact that vector addition is Abelian. We call this visualization scheme "head-to-tail" addition.

This method of visualization motivates a lot of the terminology that will be used in describing vectors, for we now have an intuitive notion of what it means for vectors to "lie on the same line" or to "point in different directions". In fact, "lying on the same line" happens precisely when "pointing in different directions" **doesn't** happen. I strongly encourage the reader to play around with this visualization scheme in 2 and 3 dimensions (it works in 1 dimension too, but then it's kind of dull). In particular, I encourage the reader to see what scalar multiplication does to this visualization of vectors. Namely, how does multiplying

a vector by 2 affect the drawing? For now, let's go back to the abstract stuff.

Exercise 3.5. Let $V = \mathbb{R} \times \mathbb{R}$. What is the zero vector in V ? Now let $V = \mathbb{R} \times \dots \times \mathbb{R}$ be an N -fold product of \mathbb{R} 's. What is the zero vector in V now?

Exercise 3.6. Let us go back to 2 copies of \mathbb{R} for a second, and let $V = \mathbb{R} \times \mathbb{R}$. Suppose we defined vector addition to be $(a, b) + (c, d) = (ac + b, b + d)$. Show why this does not satisfy the axioms of a vector space. Now suppose we picked the right vector addition but we defined scalar multiplication to be $k \cdot (a, b) = (ka, b)$. Show why this does not satisfy the axioms of a vector space. Lastly, again suppose that we've chosen the right vector addition, but now suppose that we've defined scalar multiplication to be $k \cdot (a, b) = (k^2a, k^2b)$. Show why this does not satisfy the axioms of a vector space.

3.3 Linear Combinations of Vectors

We now move on to discuss the extremely important notions of linear combinations, linear dependence, and linear independence. The first thing we need to do is define what we mean by a "linear combination" of vectors.

(In what follows, every time the phrase "vector space" is mentioned, it should be understood that we mean "real vector space". Only when we introduce complex numbers in the next chapter will we see vector spaces that aren't "real vector spaces").

Suppose that we have three vectors $u, v, w \in V$, and suppose moreover that there are real numbers $a, b \in \mathbb{R}$ such that $w = au + bv$. In this case, we say that w is a **linear combination** of the vectors u and v . We call a and b the coefficients of u and v , respectively. For example, if $w = u + v$ then w is a linear combination of u and v with 1 as a coefficient of both. Similarly, if $w = 5u - 4v$ then w is a linear combination of u and v with respective coefficients 5 and -4 .

Exercise 3.7. Let $V = \mathbb{R} \times \mathbb{R}$. Is it possible for the vector $(1, 1)$ to be written as a linear combination of $(-1, -2)$ and $(-4, -8)$? What about as a linear combination of $(-1, -2)$ and $(2, 1)$?

More generally, we can consider linear combinations of arbitrarily many vectors. In particular, if $\{v_1, v_2, \dots, v_N\}$ is some finite set of vectors in a vector space V , then any vector of the form $a_1v_1 + a_2v_2 + \dots + a_Nv_N$ with each $a_i \in \mathbb{R}$ is called a **linear combination** of the vectors v_1, v_2, \dots , and v_N . It is important to note that a linear combination of vectors must be a sum of only a finite number of terms.

Now that we have the machinery of linear combinations at our fingertips, we can (and should) discuss the notion of linear dependence and independence. We want to eventually be able to talk about whether or not a whole set of vectors are linearly dependent, but to do so we must first introduce what we mean by a pair of vectors being linearly dependent. We therefore make the following definition.

Definition 3.8. Let V be a vector space, and let $v, w \in V$ be non-zero vectors. We say that v and w are **linearly dependent** if there is an $a \in \mathbb{R}$ such that $v = aw$. If v and w are not linearly dependent (i.e., if there is no such $a \in \mathbb{R}$), then we say that v and w are **linearly independent**.

As usual, there are a couple of things to note here. First, we note that $v = aw$ means that v and w lie on the same line, at least in the examples that we've considered above. This motivates the terminology. Second, and perhaps most importantly, we don't need to worry about the asymmetrical nature of this definition—namely, why did we choose to make the definition " $v = aw$ " as opposed to " $w = av$ ". After all, before the " $v = aw$ " part of the definition our vectors v and w were on equal footing, and neither was more "special" than the other (all we knew about them was that they're both non-zero). But don't fear,

because if we know that $v = aw$ for some $a \in \mathbb{R}$, i.e. if we know that v and w are linearly dependent, then we also know that $w = bv$ for some $b \in \mathbb{R}$, and so the definition that we made above really accounts for both possibilities. This is because we know that v and w are both not zero, and therefore we know that a isn't zero either (because we've already shown that $0 \cdot v = 0$ for any $v \in V$). Thus, we can multiply both sides by $\frac{1}{a}$ to get that $w = \frac{1}{a}v$, which is what we wanted.

Now, let us notice that this definition can be phrased in a more general way—a way that will more naturally generalize to more vectors. For this, we notice that we can just as well have first defined linear **independence** and then said that a pair of vectors which are not linearly independent are then linearly dependent (it's just sort of a double negative this way, but that's okay). What we want to notice is that we could just as well have defined linear independence of v and w as follows: v and w are linearly independent if $av + bw = 0 \Rightarrow a = 0$ and $b = 0$. In other words, v and w are linearly dependent if the only linear combination of them that adds to the zero vector is the trivial linear combination $0 \cdot v + 0 \cdot w$.

To see that this is equivalent to the above definition, we first suppose that v and w are not linearly independent in the sense just described. That means that there are $a, b \in \mathbb{R}$ such that $av + bw = 0$, where $a \neq 0$ and $b \neq 0$ (the reason we know that both a and b are non-zero follows from the fact that we know that both v and w are non-zero). It then follows that $v = \frac{-b}{a}w$, and that $w = \frac{-a}{b}v$, so that v and w are linearly dependent in the original sense. Thus the two definitions are equivalent, which means that two vectors are linearly (in)dependent in the first sense if and only if they are linearly (in)dependent in the second sense, and so we can take either definition as "the" definition and use it to prove the other definition as a theorem. The reason we made the first definition is that it is slightly less abstract than the second one, but now we'll use the second one because it generalizes to more vectors in a very nice way.

Definition 3.9. Let V be a vector space, and let $\{v_1, v_2, \dots, v_N\} \subseteq V$ be a finite set of non-zero vectors. This set is a **linearly independent** set of vectors if $a_1v_1 + a_2v_2 + \dots + a_Nv_N = 0 \Rightarrow a_1 = a_2 = \dots = a_N = 0$. If the set of vectors is not linearly independent, then we say that the set is **linearly dependent**. Equivalently, this set is a linearly dependent set of vectors if there exists a set of real numbers $\{a_1, a_2, \dots, a_N\} \subseteq \mathbb{R}$ such that $a_1v_1 + a_2v_2 + \dots + a_Nv_N = 0$ and some $a_i \neq 0$.

Notice that what this definition is really saying is that a set of non-zero vectors is linearly independent if no vector in the set can be written as a linear combination of other vectors in the set. The following exercises make this point clear.

Exercise 3.10. Let V be a vector space and $\{v_1, v_2, \dots, v_N\}$ be a linearly dependent set of vectors in V . Then we know that there exists a set of real numbers $\{a_1, a_2, \dots, a_N\} \subseteq \mathbb{R}$ such that $a_1v_1 + a_2v_2 + \dots + a_Nv_N = 0$ and some $a_i \neq 0$. Show that there is also at least one other $a_j \neq 0$, where $i \neq j$.

Exercise 3.11. Let V be a vector space and $\{v_1, v_2, \dots, v_N\}$ be a linearly dependent set of vectors in V . Show that at least one vector in this set can be written as a linear combination of the other vectors. Hint: use the previous exercise.

Exercise 3.12. Let $V = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$. Say whether or not the following sets of vectors are linearly independent. If they are linearly dependent, find at least one non-trivial linear combination of them that adds up to the zero vector.

- 1) $\{(1,0,0), (0,1,0), (0,0,1)\}$
- 2) $\{(1,0,0), (0,1,0), (0,0,1), (2,4,-1)\}$
- 3) $\{(2,-2,0), (3,-7,0), (0,1,0)\}$

3.4 Vector Subspaces

As usual, once we define a mathematical structure we want to see what kind of meaningful substructure it might have. For sets, this substructure was simply the subset, and neither sets nor subsets have very much "added structure" associated with them. We can think of sets as being very "flabby" mathematical objects due to this lack of extra structure. For groups, we found that the meaningful substructure we were after was the notion of a subgroup, and we made sure that these subgroups maintained the same "group-ness" that the "parent group" had. In other words, once the mathematical structure itself (the group) has a particular "added structure" associated with it, we want to make sure that its substructures also have this added structure. We can think of subgroups as "inheriting" their group structure from the parent group, since the group operation, identity, and inverses all came from the original parent group.

In exactly the same way we want our vector space substructures to have their own identity as a vector space, and we want this identity to be inherited from the parent vector space. Namely, we want whatever we call a "vector subspace" to be a vector space in its own right (just as a subgroup of a group was a group in its own right), and we want the "vector space structure" of this subspace (namely, the vector addition and scalar multiplication rules) to be inherited from the parent vector space. It turns out that this is actually a very easy definition to make, and so we simply make it.

Definition 3.13. Let V be a vector space and $W \subseteq V$ be a non-empty subset of V . We say that W is a **vector subspace** of V if W is closed under vector addition and scalar multiplication, where the vector addition and scalar multiplication is inherited from V . This means that if $u, v \in W$, then $u + v \in W$, and that $\forall k \in \mathbb{R}$ and $\forall v \in W$, $k \cdot v \in W$.

Exercise 3.14. Show that if V is a vector space and W is a vector subspace of V , then $0 \in W$, where 0 is the zero vector in V .

It is trivial to check that the subset $W \subseteq V$ is a vector space in its own right, i.e., that it satisfies all the axioms of a vector space. The reason this is so trivial to check is that we already know that V satisfies those axioms, since we know that V is a vector space. Since W inherits its addition and scalar multiplication rules from V , and since it is closed under these operations, it is immediate that W is a vector space in its own right.

Exercise 3.15. Show that for any vector space V , the following subsets are both vector subspaces of V :

- 1) V itself
- 2) $\{0\}$, the subset consisting of only the 0 vector.

Exercise 3.16. Let $V = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$. Show that the following subsets of V are both vector subspaces of V :

- 1) $\{(a, b, 0) \mid a, b \in \mathbb{R}\}$
- 2) $\{(a, 0, a) \mid a \in \mathbb{R}\}$

We end this subsection with a very quick and easy definition.

Definition 3.17. Let V be a vector space and let W be a vector subspace of V . If $W \neq \{0\}$ and $W \neq V$, then W is a **non-trivial** vector subspace of V .

This definition should be clear enough, as there is something very "trivial" about the vector subspaces V and $\{0\}$ of a vector space V .

3.5 Bases and Dimensionality

We now introduce one of the most important properties of a vector space: its dimension. After all, we know that there is something "two-dimensional" about $\mathbb{R} \times \mathbb{R}$, and something "three-dimensional" about

$\mathbb{R} \times \mathbb{R} \times \mathbb{R}$, but how can we make this mathematically rigorous? In order to gain access to this concept, however, we first need to define the concept of a basis. Before making the formal definition, let us make a few preparatory remarks.

Our definition of a basis for a vector space is perhaps best motivated by considering the vector space $V = \mathbb{R} \times \mathbb{R}$ and noticing that every vector in this space can be written as a linear combination of the vectors $(1, 0)$ and $(0, 1)$. In particular, since any vector in V can be written as (a, b) (with $a, b \in \mathbb{R}$), we have $(a, b) = a \cdot (1, 0) + b \cdot (0, 1)$. Thus there is a set of vectors in V such that any other vector in V can be written as a linear combination of these vectors. We note that this is a very non-unique construction, for we very well could have chosen the vectors $(5, 0)$ and $(0, 3)$, in which case any vector (a, b) would be written as $(a, b) = \frac{a}{5} \cdot (5, 0) + \frac{b}{3} \cdot (0, 3)$. It doesn't matter which vectors we choose, all that matters is that it is **possible** to choose a set of vectors, all linear combinations of which can "hit" every vector in the vector space.

We notice that this works for other vector spaces as well. Consider $V = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$. We can again find vectors that can "generate" every vector in the space via their linear combinations. The canonical choice of vectors would be $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 0, 1)$. However we could just as well have chosen $(16, 0, 0)$, $(0, 1, -1)$, and $(0, 1, 1)$. This clearly works for any \mathbb{R}^n , which is short-hand notation for the n -fold product of \mathbb{R} 's.

The final thing that we'll remark on before making some precise definitions is that once we find a set of vectors that "works" in this way (i.e., generates all the vectors in the space via linear combinations), we can always add in more vectors to this set and not lose our ability to generate all of our vectors. For example, with $V = \mathbb{R} \times \mathbb{R} \times \mathbb{R}$, we can use the vectors $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$, and $(1, 1, 0)$, since we know that the first three already generate all of our vectors for us. It is clear, then, that adding in the vector $(1, 1, 0)$ to our "generating set" was unnecessary and redundant. One way to see this is that we can write this extra vector as a linear combination of the others: $(1, 1, 0) = 1 \cdot (1, 0, 0) + 1 \cdot (0, 1, 0)$. Thus, this generating set of 4 vectors is not linearly independent. It is this lack of linear independence of our generating set that alerts us that we're using a redundant set of generating vectors. This is a key observation for how we define a basis, and what the dimension of a space is. Let us turn to this now.

Definition 3.18. Let V be a vector space. A subset $B \subseteq V$ of linearly independent vectors is called a **basis** for V if any vector in V can be expressed as a linear combination of vectors in B .

It turns out that **every** vector space has a basis, but we won't prove this. We'll primarily be dealing with vector spaces that have a very natural choice of basis, and so proving the existence of bases for any vector space is not strictly necessary (though it is interesting, and may be discussed in a later edition). Thus, we'll simply accept on faith (at least for now) that every vector space has a basis. Note also that the definition of a basis implies that no vector in the basis can be the zero vector, because that would contradict the requirement that all the vectors in B need to be linearly independent (and if we check the definition of linear independence, we notice that we require the vectors to be non-zero).

There is one other way to define a basis, and although this other way is completely equivalent to what we've done above, it will allow us to introduce some new and important terminology. Therefore, let us take a moment to make this alternative definition. In fact, it will soon be clear that this "other definition" is nothing but a rewording of the above definition, lumping some of the terms in that definition into their own definition. The new terminology that we want to introduce is that of a "span". Basically, if the enemy hands us a set of vectors, we can define the "span" of these vectors to be the set of all linear combinations of these vectors. Thus, we've really already been discussing the concept of a "span", but we just haven't been calling it that yet. Before making this into a precise definition, I'll use this opportunity to introduce

some new notation. This notation may at first sight seem kind of scary, but in reality it's quite simple.

The notation I want to introduce is what's known in the business as "sigma sum notation", where we use the following scary looking symbol: \sum (which is pronounced "sigma", hence the name). We use this symbol to minimize the amount of writing that we need to do. For example, I've written the expression $a_1v_1 + \dots + a_Nv_N$ many times already, and that quickly becomes annoying. Moreover, it quickly becomes clear that I don't **really** need to write this all out to get the same information across to the reader. Namely, all the reader needs to know is that I'm summing up the terms that look like a_iv_i , and that I'm starting from $i = 1$ and ending at $i = N$ (where it clearly must already be understood what N , a_i , and v_i are, at least in the abstract). We therefore introduce the sigma notation to streamline this idea. Namely, the expression

$$\sum_{i=1}^N a_iv_i$$

will mean precisely the same thing as $a_1v_1 + \dots + a_Nv_N$. The subscript to the symbol \sum denotes both the index over which we're summing (in this case, it's i) as well as the starting value of this index in the sum. The superscript (in this case N) for \sum is the final value that our summation index i will take, and the expression that follows \sum is precisely the thing that we're summing. We note that the i index is completely arbitrary, and we could have written

$$\sum_{DONKEY=1}^N a_{DONKEY}v_{DONKEY}$$

where it is then understood that the index $DONKEY$ takes integral values between 1 and N . This would be a dumb choice of index because it increases the amount of writing we need to do (the opposite of which was the whole motivation for introducing this notation), but it is indeed mathematically correct. Due to this arbitrariness in the index i , we call this index the "dummy index" or the "dummy variable", because we're free to change it to anything we want. We're then free to manipulate these objects in any way that is well-defined. For example,

$$\sum_{i=1}^m a_iv_i + \sum_{i=m+1}^N a_iv_i = \sum_{i=1}^N a_iv_i,$$

and if $c \in \mathbb{R}$, then

$$\sum_{i=1}^N (ca_iv_i) = c \sum_{i=1}^N a_iv_i,$$

which is nothing but the result of distributivity of addition over multiplication. We'll see much more complicated uses of this notation in due time, and admittedly some of these uses may look pretty terrifying. But one should always remember that this is really nothing but addition, and so examining any complicated-looking expression carefully will uncover its true simplicity. One very common use of this notation, which I'll expose the reader to now, involves having indices on both sides of the equal sign. We'll see this in the following context. Let $\{v_1, v_2, \dots, v_M\}$ be a set of vectors in some vector space V , and let $\{w_1, w_2, \dots, w_N\}$ be another set of vectors in V . Suppose that each v_i can be written as a linear combination of the w_j . This means that we have a set of scalars $c_{ij} \in \mathbb{R}$ such that

$$v_i = \sum_{j=1}^N c_{ij}w_j.$$

What this says is that the i^{th} vector in $\{v_1, v_2, \dots, v_M\}$ is written as a linear combination of the w_j , with coefficients c_{ij} . We note that we're summing over the j index, so that the indices on both sides "match up", which means that the only "free index" on both sides of the equation is the i index. This is a very important property that always needs to hold, simply because comparing objects with different index structures is ill-defined.

We have enough at hand now to make the definition of the **span** of a subset of vectors rigorous.

Definition 3.19. Let V be a vector space and let S be a subset of V . The **span** of S , denoted by $Span(S)$, is the subset of V of all linear combinations of vectors in S . Thus, $Span(S) = \{v \in V \mid v = \sum_{i=1}^N a_i w_i, N < \infty, w_i \in S \forall 1 \leq i \leq N\}$ (the statement " $N < \infty$ " just means that N is some finite number).

It now becomes clear that the definition of a basis is equivalent to the following definition.

Definition 3.20. Let V be a vector space and B a subset of V . B is a **basis** of V if all the vectors in B are linearly independent and if $Span(B) = V$.

We now make a definition that sort of jumps the gun in terms of the terminology that is used, but we'll soon see why the words we use make sense.

Definition 3.21. Let V be a vector space with a basis B . If B has finitely many elements, then we say that V is **finite dimensional**.

This definition will become more clear once we prove the following very important theorem (the proofs of the following theorem and lemmas are a bit intricate, and so don't worry if they're not completely clear upon the first reading).

Theorem 3.22. *If V is a finite dimensional vector space, then any basis for V has the same (finite) cardinality.*

Note that this may be clear in the case of $\mathbb{R}^N = \mathbb{R} \times \dots \times \mathbb{R}$ (where there are N factors of \mathbb{R}), but now we can actually **prove** this in general. In order to prove this, however, we'll need to first prove two more very important results, both of which will be shown as lemmas.

Lemma 3.23. Let V be a vector space and B a basis for V . Then for each $v \in V$ there is a **unique** expansion (up to reordering) in terms of basis elements: $v = \sum_{i=1}^N k_i b_i$, where each $k_i \in \mathbb{R}$ and each $b_i \in B$. (Note that we're not assuming that B has finite cardinality, but rather we're recalling that a linear combination must have only finitely many terms in its sum, and it is for that reason that we know our sum only involves N terms (for some finite integer N)).

Proof of Lemma: Let V and B be as in the statement of the lemma. We know that the zero vector 0 has a unique expansion in terms of basis elements, and that is the expansion with 0 as the coefficient of every term (this follows from the fact that we know that all of the vectors in B are linearly independent). Thus, we only need to show that the Lemma holds for all non-zero vectors in V . Choose $v \in V$ as an arbitrary non-zero vector. Since B is a basis for V , we know that there is at least one expansion of v in terms of elements in B , so let us denote this expansion as $v = \sum_{i=1}^N k_i b_i$ where N is some finite integer, each $k_i \in \mathbb{R}$, and each $b_i \in B$. Since v is non-zero, we know that at least one k_i is also non-zero. Without loss of generality, we can suppose that each k_i is non-zero, for otherwise we're just adding the zero vector for no reason.

Now we suppose that there is another expansion of v , which we write as $v = \sum_{i=1}^M l_i c_i$, where M is again some finite integer (not necessarily the same as N), each $l_i \in \mathbb{R}$, and each $c_i \in B$. Similarly to our assumption above, we can assume without losing any generality that each $l_i \neq 0$. Then we have the following chain of equalities:

$$0 = v - v = \sum_{i=1}^N k_i b_i - \sum_{j=1}^M l_j c_j. \quad (3.1)$$

But we recall that all the vectors in B are linearly independent, so that if a linear combination of them equals zero it must be the case that each individual coefficient in that linear combination is zero. Thus, if there is a c_j that does not equal any of the b_i 's in the sum, its coefficient must be zero. But this contradicts our assumption that all of the l_j are non-zero. Thus there can be no c_j that doesn't appear as a b_i . Similarly, there can be no b_i that doesn't appear as a c_j somewhere, because we assumed that each k_i was non-zero. To recap, what we've shown so far is that each c_j equals **some** b_i , and vice versa. Thus, possibly after reordering the c_j 's, we can write the above sum as

$$0 = v - v = \sum_{i=1}^N (k_i - l_i) b_i, \quad (3.2)$$

and since each coefficient of this sum must individually vanish (by definition of linear independence), it must be the case that $k_i = l_i$ for each i (again, this is after a possible reordering). Thus, after reordering, these two expansions of v in terms of basis elements are equivalent. \square

Lemma 3.24. Let V be a vector space, A be a subset of m linearly independent vectors in V , and B a set of n vectors in V such that $\text{Span}(B) = V$. Then $m \leq n$.

Proof of Lemma: Without loss of generality, we can write $A = \{a_1, \dots, a_m\}$, and $B = \{b_1, \dots, b_n\}$. Now, since $\text{Span}(\{b_1, \dots, b_n\}) = V$, we know that $a_1 = \sum_{i=1}^n l_i b_i$, where each $l_i \in \mathbb{R}$ and where at least one $l_t \neq 0$ (because we know $a_1 \neq 0$ since it is an element of a linearly independent set of vectors). Then we have that $b_t = \frac{1}{l_t}(a_1 - \sum_{i=1, i \neq t}^n l_i b_i)$. Additionally, we know that $\text{Span}(\{a_1, b_1, \dots, b_n\}) = V$, since adding vectors to a subset never diminishes that subset's span. However, the previous line has shown that a_1 and b_t are linearly dependent, so we can remove b_t from $\{a_1, b_1, \dots, b_n\}$ without diminishing its span. If we let $B_1 = \{a_1, b_1, \dots, b_{t-1}, b_{t+1}, \dots, b_n\}$, then we have that $\text{Span}(B_1) = V$. But now we can repeat this process using the sets $A_1 = \{a_2, \dots, a_m\}$ and the B_1 just defined. Moreover, since the a_i 's are all linearly independent, we know that the vector that we'll remove from B_1 is not a_1 . More generally, as we continue to repeat this process, we'll never have to delete a_i from B_i . Thus, after m repetitions of this process, we'll have a set that looks like $B_m = \{a_1, \dots, a_m, b_{l_1}, \dots, b_{l_k}\}$ where the b_{l_i} 's are the elements of B that never got deleted in any of the repetitions. It may be the case that there are no such left over elements, but we'll always be able to repeat this process at least m times. Since it is the case that we'll either have 0 elements left over, or some positive number of elements left over, it is true that $m \leq n$. \square

(We note that this proof works only because the sets of vectors in question had finitely many elements. This was an algorithmic proof, and in order to know that such a proof actually works, we need to know that the algorithm will eventually end, and we only know this if we know that we'll only need to do a finite number of steps in the algorithm)

We can now use this lemma to prove the theorem.

Proof of theorem: Since V is a finite dimensional vector space, we know that there exists at least one basis B for it. We know that this basis is a set of finitely many vectors, so let us denote it as follows: $B = \{b_1, b_2, \dots, b_N\}$ where N is some non-negative integer. Moreover, we know that any vector in V can be written as a linear combination of these b_i 's, and that the b_i 's are all linearly independent. Now let C be some other basis for V . We want to show that C has N elements as well.

Let us first show that C doesn't have infinitely many elements. We first suppose that C does have

infinitely many elements. Next, we note that since C is a basis, any vector in B can be expressed (uniquely, by one of the above lemmas) as a linear combination of elements in C as follows:

$$\begin{aligned} b_1 &= \sum_{i=1}^{M_1} k_{1i} c_{1i}, \\ b_2 &= \sum_{i=1}^{M_2} k_{2i} c_{2i}, \\ &\dots, \\ b_N &= \sum_{i=1}^{M_N} k_{Ni} c_{Ni} \end{aligned} \tag{3.3}$$

(here we're just labeling our coefficients as $k_{ij} \in \mathbb{R}$, where the i subscript refers to the corresponding b_i , and the j subscript refers to the corresponding c_j , and we need to label each M with a different subscript because each sum might have a different number of terms).

What's important to note is that we are only using a finite number of elements in C to express the elements in B as linear combinations of these basis elements, since we know that there are only finitely many elements in B and only finitely many elements in each linear combination. Thus, there are elements of C that are not used in any of these linear combinations (in fact, there are infinitely many such elements!). Let us therefore choose one of the elements in C that isn't used in any of these linear combinations, and let us call this element c' . Then, since B is a basis, we know that c' can be written as a linear combination of elements in B as $c' = \sum_{i=1}^N l_i b_i$ (now we use N as the upper limit of the sum because we know how many elements B has—namely N). But now we can substitute in the above expressions for the b_i 's as follows:

$$c' = \sum_{i=1}^N l_i \left(\sum_{j=1}^{M_j} k_{ji} c_{ji} \right) = \left(\sum_{i=1}^N l_i \sum_{j=1}^{M_j} k_{ji} \right) c_{ji} \tag{3.4}$$

where in the second equality we have simply moved the parentheses around (and we're allowed to do this because all of our sums only involve finitely many terms—namely, the property of distributivity that numbers have ensures that for any sums with finitely many terms, products and sums commute (we do need to be careful when there are infinitely many terms in a sum, though, but we won't see such infinite sums in this text)). We can then view the parenthetical expression in the right-most equality of (3.4) as the coefficients in a linear combination of basis elements in C . Moreover, we know that $c' \neq c_{ji}$ for any i, j , because that is precisely how we defined c' . Finally, we also know that at least some of the coefficients in this linear combination are non-zero. Thus, we've expressed one basis element in C as a linear combination of other basis elements in C , which is a contradiction to the fact that the vectors in C must be linearly independent (by definition of a basis). Thus our supposition that C has infinitely many elements is impossible, and so we know that C must be a finite collection of vectors.

To recap, we've now shown that if a vector space V has a basis B with finitely many elements, then every other basis for V will also have finitely many elements. What we want to show now is that every other basis for V has precisely the same number of elements that B has. Luckily for us, the second of the above two lemmas finishes this theorem for us. For now we have two finite sets of vectors B and C , both of which are linearly independent and both of which span V . We know that B has N elements, and we can suppose that C has M elements. Then we can apply the above lemma with B playing the role of A , and C playing the role of B to get that $N \leq M$. But then we can also apply the same lemma with the roles of B and C reversed to get that $M \leq N$. Thus it must be the case that $N = M$, and we're done! \square

What we've done is show that if a vector space V has a basis B with N elements, then every basis for V will have exactly N elements. Thus, we've found a number, N , that is somehow associated to V in a very intimate way. We call this number the **dimension** of V , and it turns out that the dimension of a vector space is the most important number associated to that space. Let us make this definition precise.

Definition 3.25. Let V be a vector space. If a basis B of V has finitely many elements, then the cardinality of B is called the **dimension** of V . (Note that we need the above theorem to make this a well-defined notion).

Although there are many important infinite dimensional vector spaces "out there", the ones that will interest us the most here are the finite dimensional ones. Therefore it should be assumed from here on out (unless otherwise stated) that all of our vector spaces are finite dimensional.

Exercise 3.26. What is the dimension of the following vector spaces:

- 1) $\mathbb{R} \times \mathbb{R}$,
- 2) \mathbb{R}^n ,
- 3) \mathbb{R} ?

One final remark on dimensionality is that we simply **define** the trivial vector space consisting of one vector, $\{0\}$, to have dimension 0. We need to make this a **definition** (as opposed to being able to derive it) because the definition of dimension would imply that a 0-dimensional space should have a basis with 0 elements in it. But this basis would clearly be the empty set, and the span of the empty set is also empty. Thus, the vector space $\{0\}$ most certainly **does not** equal the span of a 0-element basis, and therefore it needs its own definition.

This ends our discussion of the dimensionality of vector spaces, and now we move on to discuss the important functions between vector spaces...

3.6 *Linear Maps

Let us recap what we've done with vector spaces so far. We first motivated their definition and then proceeded to define them. Once we defined them, we uncovered a few special properties about them, and then we defined their substructure—namely, vector subspaces. This parallels what we did with sets, where we first defined sets and then subsets. Similarly, when discussing groups we first defined what they were and then defined subgroups. Subgroups were nothing but subsets with certain special group properties in relation to their parent group, which is to be expected since groups are nothing but sets with special properties themselves. Finally, vector subspaces are just special subsets of vector spaces, carrying over all of the vector space information from its parent space.

In the case of both sets and groups, we then proceeded to ask how we can relate different sets or groups to each other. In the case of sets, the answer was a generic function, and in the case of a group the answer was a homomorphism. Again, a homomorphism is just a function with certain extra requirements, reflecting the fact that a group is just a set with extra requirements. Namely, we wanted our "special functions" (i.e., our homomorphisms) to be compatible with the group structures of the two groups in question, so that "composing then sending" was the same as "sending then composing", where composition happened **within** the respective groups, and the "sending" part was done by the function **between** groups.

Recall finally that in the case of a group, we could "forget" about the group structure and only view the object as a set. Then we could consider functions between groups simply as regular set functions, and we'd then have many more options at our disposal. In particular, requiring that a function between two

groups be a homomorphism is a much more stringent requirement than asking for a function between the two groups when viewed simply as sets. We could of course talk about regular set functions between groups, but then since we've lost all of what makes groups interesting (i.e., their "added structure"), we've then lost most of the interesting things to talk about. This is why homomorphisms are so important in group theory.

Since vector spaces have so much more structure than a group (Abelian, scalar multiplication, distributivity, etc.), we'd like the functions that relate two vector spaces to each other to preserve all of this structure. We could of course just talk about regular set functions, but then we'd be losing **even more** structure than we would in the case of groups. Thus, we want the analogue of a "homomorphism" for a vector space, only there is now more structure to preserve. Let us therefore start to see what kinds of things we would naturally ask of these special functions.

First let us set the stage. Suppose we have two vector spaces $(V, +)$ and $(W, +)$. Note that these are two different sets with two different rules of composition (vector addition), but that we're using the same symbol "+" for denoting both. This is simply because both rules are called "vector addition" regardless of whichever vector space we're doing the addition in, and also because in any given expression it will be clear which "+" sign we're dealing with (i.e., in which vector space we're adding).

Now let $f : V \rightarrow W$ be the function that we're going to start putting some restrictions on. The first thing we'd like is for f to be a homomorphism in the strict sense, seeing as both V and W have a structure as an Abelian group. Thus, we want it to be the case for all $v_1, v_2 \in V$, $f(v_1 + v_2) = f(v_1) + f(v_2)$. This is nothing but the property of being a homomorphism. Note that on the left of this equality we are adding v_1 and v_2 in V first, and then sending the sum over to W , whereas on the right we are first sending v_1 over to W , as well as sending v_2 over to W , and then we're adding these images up in W . Note also that even though the "+" signs are the same on both sides of the equation, they refer to addition in **different** vector spaces, but that deciphering which is which is an easy task (namely, the "+" signs are assigned in the only way that makes sense).

There is another key feature of vector spaces that we need to consider, and that is scalar multiplication. Since this is so important in the definition of (i.e., construction of) vector spaces, we'd like our special "vector space to vector space" functions to be compatible with it. And what do we mean by "compatible"? Well, just as our homomorphisms are such that "compose then send" is the same as "send then compose", we want our vector space functions to be such that "multiply then send" is the same as "send then multiply". I.e., in symbols, we want $f(a \cdot v) = a \cdot f(v)$ for all $a \in \mathbb{R}$ and for all $v \in V$. Note again that the "." on the left and right of this equality are scalar multiplication rules for **different** vector spaces, where on the left we're multiplying a to v (in V , since $v \in V$) and then sending the product to W , whereas on the right we're first sending v to W via f , and then multiplying the **vector** $f(v) \in W$ by a , using W 's scalar multiplication.

It turns out that this is all we need in our definition, as other qualities like distributivity follow simply from ensuring the above two requirements. Thus, we have enough at hand to make a formal definition. In the following definition, we'll attach the subscripts " V " and " W " to the symbols "+" and "." to remind us in which space the vector addition and scalar multiplication is being carried out, but in all practical uses these subscripts will be omitted since we can determine this information solely from context. We do this only so that there is not an abuse of notation in a formal definition (something that should be avoided at all reasonable times).

Definition 3.27. Let $(V, +_V, \cdot_V)$ and $(W, +_W, \cdot_W)$ be vector spaces with respective vector addition rules $+_V$ and $+_W$, and respective scalar multiplication rules \cdot_V and \cdot_W . Let $f : V \rightarrow W$. f is said to be a **linear map** if the following conditions hold $\forall v_1, v_2 \in V$ and $\forall a \in \mathbb{R}$:

- 1) $f(v_1 +_V v_2) = f(v_1) +_W f(v_2)$
- 2) $f(a \cdot_V v_1) = a \cdot_W f(v_1)$.

By seeing how confusing (or at least cluttered) the notation in this definition is, we can appreciate the utility of simply using "+" and "." for both vector spaces' vector addition rules and scalar multiplication rules, and letting context tell us "which is which" when we see more than one "+" and/or "." around.

We note that "linear maps" are also sometimes called "linear functions" or "linear transformations".

Now let's find some properties that all linear maps have in common with each other. I.e., we can prove the following results using only the axioms of a linear map, and thus any linear map will have the following traits in common.

Exercise 3.28. Let V and W be two vector spaces, let $f : V \rightarrow W$ be a linear map, and let 0_V and 0_W denote the two spaces' respective zero vectors. Show that $f(0_V) = 0_W$.

Exercise 3.29. Let V and W be two vector spaces, let $f : V \rightarrow W$ be a linear map, let $\{e_1, \dots, e_n\}$ be a set of n vectors in V , and let $\{a_1, \dots, a_n\}$ be a set of n real numbers. Finally, let $v \in V$ be the linear combination $v = a_1 e_1 + \dots + a_n e_n$. Show that $f(a_1 e_1 + \dots + a_n e_n) = a_1 f(e_1) + \dots + a_n f(e_n)$.

Exercise 3.30. Let V and W be two vector spaces, let $f : V \rightarrow W$ be a linear map, and let B be a basis for V . Show that f 's behaviour is completely determined by its behaviour on elements in B . Namely, show that if we know $f(b)$ for all $b \in B$, then we know $f(v)$ for all $v \in V$.

Now that we have our analogue of a homomorphism for vector spaces, i.e. linear maps, we can start to ask even more questions about them. In the case of homomorphisms, we were interested in finding out if certain of them were surjective, injective, or bijective. We then gave a special name to the bijective homomorphisms—an isomorphism. If there exists an isomorphism between two groups then we say that the two groups are isomorphic, and we found that isomorphic groups could, in a very concrete way, be viewed as being "equivalent" to each other—the technical term would be "equivalent up to isomorphism". The reason we can view isomorphic groups as being "equivalent" is that anything that we can do with one group—or anything that is true of one group—will be true of the other, simply because every element in one group has a corresponding element in the other (a "partner", so to speak), that behaves exactly the same way. We can find this partner using the bijectivity of the function, and we know it plays the same role in the other group by using the fact that the map is a homomorphism (which preserves group behavior).

We can make a completely analogous definition for vector spaces by considering bijective linear maps. We can, of course, also consider linear maps that are only surjective, or only injective, and although there are lots of interesting things to said about these cases, we'll be most interested for now in the case when the linear map is bijective. We therefore make the following definition (it's really two definitions in one).

Definition 3.31. Let V and W be vector spaces. If $f : V \rightarrow W$ is a bijective linear map, then we say that f is a **linear isomorphism**. Moreover, if there exists a linear isomorphism between V and W , then we say that V and W are linearly isomorphic.

(Note that sometimes we drop the "linear" and only say that the map is an isomorphism, or that two spaces are isomorphic, if it's understood that we're talking about vector spaces).

Exercise 3.32. Let V and W be vector spaces and let $f : V \rightarrow W$ be a linear isomorphism. Show that $f^{-1} : W \rightarrow V$ is also a linear isomorphism (here, f^{-1} is the map that takes each element in $w \in W$ to the element in V that maps to w under f , which we showed was well-defined in the "Sets" lecture). (Note: this one is a bit tricky)

The above exercise shows that it is indeed well defined to say that two spaces are isomorphic even though any given isomorphism goes from one space to another. In other words, the phrase "these two spaces are isomorphic" implies a sort of symmetry between the two spaces, in that they're both isomorphic to each other. However, any particular linear isomorphism breaks that symmetry because we know that a function must go **from** one space **to** another, thus putting one space on a different footing than the other. However, the above exercise shows that if there's a linear isomorphism "one way", then there's a linear isomorphism "the other way", so that the symmetric notion of two spaces being isomorphic "to each other" is well-defined.

Now it just so happens, rather remarkably, that any two finite-dimensional vector spaces with same dimension are equivalent to each other, up to isomorphism. The precise statement of this fact is that any two finite dimensional vector spaces of the same dimension are linearly isomorphic to each other. This fact is so important, and so key to the study of vector spaces, that we make it a theorem.

Theorem 3.33. *If V and W are finite dimensional vector spaces, and if they are of equal dimension, then there exists a linear isomorphism $f : V \rightarrow W$.*

Proof: Since V and W are both finite dimensional vector spaces with the same dimension, we know that they both have bases with the same finite number of elements. Let us call this number N . Thus, there is a set B_V of N linearly independent vectors in V that any other vector in V can be written as a linear combination of, and a set B_W of N linearly independent vectors in W that do the same in W . Let us denote these vectors as follows: $B_V = \{x_1, x_2, \dots, x_N\}$ and $B_W = \{y_1, y_2, \dots, y_N\}$.

We now want to define a linear isomorphism from V to W (which we recall is equivalent to finding a linear isomorphism from W to V). With our bases in hand, though, this will be easy. We recall from an above exercise that defining our linear map $f : V \rightarrow W$ on the basis elements of V is sufficient for defining f on all of V . Thus, we'll define f only on the x_i and let it be extended to the rest of V via linearity. Namely, let us define f by the following: $f(x_i) = y_i$. I.e., f sends the i^{th} basis element in V to the i^{th} basis element in W . We're **defining** this map to be a linear map, so that for any linear combination of x_i 's (which covers all the vectors in V), we have $f(a_1x_1 + \dots + a_Nx_N) = a_1f(x_1) + \dots + a_Nf(x_N) = a_1y_1 + \dots + a_Ny_N$ (where each $a_i \in \mathbb{R}$).

Thus, the only thing we need to do now is show is that this linear map is bijective. We begin by showing that it is surjective. To do this, recall that we need to take an arbitrary vector $w \in W$ and show that there is some $v \in V$ such that $f(v) = w$. Accordingly, we first note that w can be written as a linear combination of the y_i (since $\{y_i\}$ is a basis), so that for some set of N real numbers $\{b_1, \dots, b_N\} \subset \mathbb{R}$, we have $w = b_1y_1 + \dots + b_Ny_N$. Now we simply notice the following chain of equalities:

$$w = b_1y_1 + \dots + b_Ny_N = b_1f(x_1) + \dots + b_Nf(x_N) = f(b_1x_1) + \dots + f(b_Nx_N) = f(b_1x_1 + \dots + b_Nx_N) \quad (3.5)$$

where the first equality comes from the preceding sentence, the second equality comes from the fact that $f(x_i) = y_i$, the third equality comes from the fact that linear maps are compatible with scalar multiplication, and the last equality comes from the fact that linear maps have the property that $f(v) + f(w) = f(v + w)$. Thus, we've found a vector in V that maps to $w \in W$, and that vector is the linear combination $b_1x_1 + \dots + b_Nx_N$. Therefore f is surjective.

The last thing we need to do is prove that f is injective. To do this, we need to show that for all $v, w \in V$, $f(v) = f(w) \Rightarrow v = w$. So let's choose vectors $v, w \in V$ arbitrarily, and let's suppose that $f(v) = f(w)$. We first note that both v and w can be expanded as linear combinations of elements B_V , so let us write $v = a_1x_1 + \dots + a_Nx_N$ with all of the $a_i \in \mathbb{R}$, and $w = b_1x_1 + \dots + b_Nx_N$ with all of the $b_i \in \mathbb{R}$. Then we have the following:

$$f(v) = f(w) \Rightarrow f(a_1x_1 + \dots + a_Nx_N) = f(b_1x_1 + \dots + b_Nx_N) \Rightarrow a_1f(x_1) + \dots + a_Nf(x_N) = b_1f(x_1) + \dots + b_Nf(x_N). \quad (3.6)$$

Now we take the last equality and subtract $b_1f(x_1) + \dots + b_Nf(x_N)$ from both sides, so that we have

$$a_1f(x_1) - b_1f(x_1) + \dots + a_Nf(x_N) - b_Nf(x_N) = 0, \quad (3.7)$$

and this implies that

$$(a_1 - b_1)f(x_1) + \dots + (a_N - b_N)f(x_N) = 0. \quad (3.8)$$

We then recall that $f(x_i) = y_i$, so that the above equality implies that

$$(a_1 - b_1)y_1 + \dots + (a_N - b_N)y_N = 0. \quad (3.9)$$

However, we know that the vectors $\{y_1, \dots, y_N\}$ are all linearly independent, so that the only linear combination of them that equals zero is that which has all of its coefficients equal to zero. I.e., $k_1y_1 + \dots + k_Ny_N = 0 \Rightarrow k_1 = k_2 = \dots = k_N = 0$. Thus we have that $a_1 - b_1 = 0, a_2 - b_2 = 0, \dots, a_N - b_N = 0$, or more simply that for all $1 \leq i \leq N, a_i - b_i = 0$. Thus, for all $1 \leq i \leq N$, we have that $a_i = b_i$. But then this means that $a_1x_1 + \dots + a_Nx_N = b_1x_1 + \dots + b_Nx_N$, and this means that $v = w$. Thus it is the case that $f(v) = f(w) \Rightarrow f = w$, and so f is injective.

Therefore f is a bijective linear map, and so is a linear isomorphism, and so V and W are linearly isomorphic. \square

We now know that if the enemy hands us two vector spaces, both of the same finite dimension, then they'll be linearly isomorphic to each other and so we can effectively view them as "the same" vector spaces. Now, we already know of a bunch of finite dimensional vector spaces, namely \mathbb{R}^n for any finite positive integer n (recall that \mathbb{R}^n is the set of n -tuples (a_1, \dots, a_n) with each $a_i \in \mathbb{R}$ with the vector space structure (vector addition and scalar multiplication) as discuss above). Moreover, it is clear that for each n , \mathbb{R}^n is of dimension n . This is because we can form the basis $\{(1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)\}$ which is easily seen to span \mathbb{R}^n and which is easily seen to be linearly independent. Since this basis has n vectors in it, the dimension of \mathbb{R}^n is n . Therefore, by the above theorem, we know that any vector space of dimension n will be linearly isomorphic to \mathbb{R}^n . We've therefore classified all finite dimensional real vector spaces up to linear isomorphisms. That is, we know what all finite dimensional real vector spaces "look like", and we know that they all look like \mathbb{R}^n for some n .

Accordingly, our study of finite dimensional real vector spaces is really a study of \mathbb{R}^n . By focusing on \mathbb{R}^n , we are naturally led to a concept that is pervasive in mathematics and physics, and that is the concept of a matrix. Now, these matrices will have nothing to do with the action-packed sci-fi movie series starring Keanu Reeves, but they will be pretty interesting objects to study in their own right. This is what we turn our attention to after the next example and the subsequent afterword, and this is the subject that will finish (for now) our discussion of real vector spaces.

Example 3.34. So far we've only discussed the vector space \mathbb{R}^n , and although this is indeed an important vector space, this subject would be kind of dull if it were the only example of a vector space. Namely, there exists lots of abstract vector spaces that aren't as "concrete" as \mathbb{R}^n , so we'll take a look at one now. Of course, we know that if any such abstract vector space is finite dimensional, then it will be **isomorphic to** \mathbb{R}^n for some n , but often times we're handed structures that are secretly vector spaces, and without seeing that they have a vector space structure we may have never known that these structures behaved like \mathbb{R}^n .

Let $X = \{x_1, x_2, x_3\}$ be some generic set with three elements, and consider the set V of all functions $f : X \rightarrow \mathbb{R}$. I've purposefully called this set V because it turns out this set has the structure of a vector space. Moreover, it's a vector space of dimension 3. Let's see how this goes.

Let's define our vector addition to be the operation that takes two functions $f, g \in V$ and forms the new function $h = f + g$ defined by mapping each x_i to $f(x_i) + g(x_i)$. Thus $h(x_i) = f(x_i) + g(x_i) \in \mathbb{R}$ is a perfectly good function from X to \mathbb{R} , and is therefore a perfectly good element (vector) in V . Scalar

multiplication by an element $k \in \mathbb{R}$ will then be defined as taking each function $f \in V$ to the function $kf \in V$, which is defined as $(kf)(x_i) = k \times f(x_i)$. Note that on the left of this equality we're applying the **function** denoted by kf to the element x_i , whereas on the right we're multiplying the **number** k to the **number** $f(x_i)$. With this definition of scalar multiplication and vector addition, we're able to use the properties of adding and multiplying numbers in \mathbb{R} to prove that V does indeed have a vector space structure—i.e., that it satisfies all the requirements of a vector space.

To see that V is 3-dimensional, we notice that we can define a basis of functions with three elements. Namely, we can define our basis functions to be e_1, e_2 , and $e_3 \in V$ as usual, where $e_i(x_j) = 1$ if $i = j$ and $e_i(x_j) = 0$ if $i \neq j$. Thus each basis element maps one element in X to 1 and every other element to 0. To see that this actually forms a basis, we need to first show that any function $f \in V$ can be written as a linear combination of these e_i , and then we need to show that the $\{e_i\}$ are all linearly independent. Choose $f \in V$ arbitrary. Then let $f(x_1) = a_1 \in \mathbb{R}$, $f(x_2) = a_2 \in \mathbb{R}$, and $f(x_3) = a_3 \in \mathbb{R}$. We then have that $f = a_1e_1 + a_2e_2 + a_3e_3$, because then, for example, $f(x_1) = a_1e_1(x_1) + a_2e_2(x_1) + a_3e_3(x_1) = a_1 \times 1 + a_2 \times 0 + a_3 \times 0 = a_1$. It is simple to check that this works for x_2 and x_3 as well, so that we've found the right linear combination of the e_i 's for f . To see that these e_i 's are linearly independent, we suppose that there are some real numbers c_1, c_2 , and $c_3 \in \mathbb{R}$ such that $c_1e_1 + c_2e_2 + c_3e_3 = 0$, where the zero vector in V is the function that sends all three elements in X to $0 \in \mathbb{R}$. We then simply apply the function $c_1e_1 + c_2e_2 + c_3e_3$ to the three elements in X . Namely, we have that $c_1e_1(x_1) + c_2e_2(x_1) + c_3e_3(x_1) = c_1$, which therefore means that $c_1 = 0$ since this function is supposed to send everything to 0. We can similarly show that c_2 and c_3 both are zero, therefore satisfying the requirement for linear independence of the vectors e_1, e_2 , and e_3 (namely, if any linear combination of them is zero, then it must be the case that each individual coefficient of the linear combination is zero). We've thus shown that this V is 3-dimensional.

As usual, we can generalize this quite a bit. Namely, if we let $X = \{x_1, x_2, \dots, x_N\}$ be a set with finitely many elements, then the above constructions carry over in the same exact way and we notice that we have an N -dimensional vector space. What's more, we can even let X be any set at all (finite or infinite) and do the same, only now our vector space V of functions will be infinite dimensional. We can see this by again defining a basis of functions with each basis function sending one element in X to $1 \in \mathbb{R}$, and all other elements in X to 0. It is clear then that the cardinality of the basis will be the same as the cardinality of X , and so if X is infinite then so is the dimension of V . But that's okay, this just means we've seen our first infinite dimensional vector space!

(Afterword: this process of "classifying" all finite dimensional vector spaces is a very common process in mathematics. What we've done here is take an abstract definition of an abstract structure (that of a vector space), ascribe to that structure a notion of "equality" (that of a bijective linear map, or linear isomorphism) where we view two different cases of the abstract structure as being equivalent in a well-defined way, and then seek to find a set of concrete realizations (namely, \mathbb{R}^n) that completely embody all of the possible instances of that abstract structure. Now, often times this is hard to do with the most general type of abstract structure, and therefore certain restrictions will have to be put on the structure. In our case here, we had to require that our abstract vector spaces are finite dimensional. Once we put that extra requirement on our abstract structure, we were able to completely classify this subclass of structures by showing that they're all equivalent to the more concrete set \mathbb{R}^n . The reason this is so fruitful is that often times in mathematics (and physics, and elsewhere) we'll have motivations to make an abstract structure, and often times these abstract structures are hard to manipulate when they're in their abstract form. By classifying these structures in terms of more concrete structures, like \mathbb{R}^n , we then gain the ability to manipulate and understand the abstract ones. For instance, we know (and will continue to learn) how to manipulate and calculate things in the set \mathbb{R}^n , whereas a general abstract n -dimensional vector space is more difficult.

Lots of mathematicians spend their time working to classify various abstract structures. Most of these lines of study are much more difficult than that which we underwent here. For example, it is sometimes difficult to know what the right notion of "equality" that we want to define is, and it is often very difficult

to know what the right extra conditions that we want to impose on the structure are. We often want to impose as few conditions as possible, for then our classification scheme will be more general and more powerful, but this needs to be balanced by the fact that the more conditions we put on a structure the easier it is to prove things about them.

The important thing to take away from this afterword is that the classification of finite dimensional vector spaces that we just witnessed is only one of many such lines of reasoning in mathematics, and many of them are much more involved than what we've seen here.)

3.7 *Matrices

(A word of warning: the algebra in this section gets nasty very quickly, but try not to let this scare you. If the reader attacks it calmly and slowly, she will find that this algebra is nothing but addition and multiplication as usual!)

We now restrict ourselves to dealing with only finite dimensional vector spaces, and therefore—by the discussion ending the previous section—to spaces like \mathbb{R}^n for some n . What we'll be most interested in for now is linear maps between finite dimensional spaces, and therefore linear maps from \mathbb{R}^n for some n to \mathbb{R}^m for some m . We've seen before that a linear map is completely specified once it is specified on the basis elements of some basis for the domain of the map (assuming that we "force" linearity on the map, meaning that we ensure that it satisfies all the axioms of a linear map). In our case, we have a linear map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and this map will be completely specified once we define it on some basis for \mathbb{R}^n .

However, we've seen before that any vector space has lots of different bases. For example, one basis for \mathbb{R}^2 is $\{(1, 0), (0, 1)\}$, but so are $\{(2, 0), (0, 2)\}$, $\{(-1, 0), (0, -1)\}$, and $\{(1, -1), (1, -2)\}$. In fact, there are infinitely many different bases, as can be easily seen. What we'll do here is first choose a particular basis, and then later see what happens when we decide to change our basis from the one we started with to another one.

Let us, for the sake of being concrete, focus on linear maps of the form $g : \mathbb{R}^3 \rightarrow \mathbb{R}^2$, i.e., from a 3-dimensional vector space to a 2-dimensional vector space. Let us also choose the natural basis for \mathbb{R}^3 , which is $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. Let us give these vectors names: $e_1 = (1, 0, 0)$, $e_2 = (0, 1, 0)$, $e_3 = (0, 0, 1)$. Our map is then fully defined by where it sends these three vectors. Let us therefore denote the image of these three vectors in \mathbb{R}^2 by $w_1, w_2, w_3 (\in \mathbb{R}^2)$, so that $w_1 = g(e_1)$, $w_2 = g(e_2)$, $w_3 = g(e_3)$. Let us now choose the natural basis for \mathbb{R}^2 , which is $\{(1, 0), (0, 1)\}$, and let us give these vectors names as well: $f_1 = (1, 0)$, $f_2 = (0, 1)$. Then we know that each w_i can be expressed as a linear combination of these two vectors, and we have the following:

$$g(e_1) = a_{11}f_1 + a_{21}f_2 \tag{3.10}$$

$$g(e_2) = a_{12}f_1 + a_{22}f_2$$

$$g(e_3) = a_{13}f_1 + a_{23}f_2$$

$$\tag{3.11}$$

where each $a_{ij} \in \mathbb{R}$, and where the right subscript on the a_{ij} corresponds to the subscript of the e_j that is being sent to \mathbb{R}^2 , and where the left subscript on the a_{ij} corresponds to the subscript of the f_i it is a coefficient of. It is important that we keep straight "what is what" here. Namely, the a_{ij} 's are real numbers, whereas the e_j 's are vectors in \mathbb{R}^3 and the f_i 's are vectors in \mathbb{R}^2 .

Now that we've defined g on the basis elements of \mathbb{R}^3 , we can know where it sends any element \mathbb{R}^3 . For if $v \in \mathbb{R}^3$, then $v = b_1e_1 + b_2e_2 + b_3e_3 = (b_1, b_2, b_3)$ for some $b_1, b_2, b_3 \in \mathbb{R}$. Then we have (after extending g using linearity)

$$\begin{aligned} g(v) &= g(b_1e_1 + b_2e_2 + b_3e_3) = b_1g(e_1) + b_2g(e_2) + b_3g(e_3) \\ &= b_1(a_{11}f_1 + a_{21}f_2) + b_2(a_{12}f_1 + a_{22}f_2) + b_3(a_{13}f_1 + a_{23}f_2) \end{aligned} \quad (3.12)$$

which, after collecting the coefficients of each f_i , gives

$$g(v) = (b_1a_{11} + b_2a_{12} + b_3a_{13})f_1 + (b_1a_{21} + b_2a_{22} + b_3a_{23})f_2 = ((b_1a_{11} + b_2a_{12} + b_3a_{13}), (b_1a_{21} + b_2a_{22} + b_3a_{23})). \quad (3.13)$$

It is therefore the case that the six real numbers a_{ij} fully determine our linear map g .

It turns out that there's a more concise way to present the data of this linear map, and it is to combine the six real numbers a_{ij} into an array called a **matrix**. What we do is set up a 2×3 grid, where the first digit (in this case, 2) corresponds to the number of rows of the grid, and the second digit corresponds to the number of columns. We then fill in this grid as follows:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix}$$

We then write the vectors in \mathbb{R}^3 vertically, as opposed to horizontally like we have been so far, so that

$$e_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, e_3 = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}$$

where scalar multiplication and vector addition go through component-wise as usual:

$$\begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = \begin{pmatrix} a_1 + b_1 \\ a_2 + b_2 \\ a_3 + b_3 \end{pmatrix}, \quad k \cdot \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = \begin{pmatrix} kb_1 \\ kb_2 \\ kb_3 \end{pmatrix}$$

We now need to define a new kind of multiplication—that between a matrix and a vector—which will encapsulate the behavior of our linear map. Let us start by first placing the matrix next to a general vector $v = (b_1, b_2, b_3)$, since most multiplication rules involve putting the things that are being multiplied next to each other:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

This notation is short hand for "the linear map defined by the six numbers a_{ij} applied to the vector (b_1, b_2, b_3) ", or in other words, the above expression is nothing but another (albeit much longer) way of writing the expression $g(v)$. But we know exactly what the action of g is on v , and so we know that we want the above expression to behave as follows:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{pmatrix} \cdot \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = \begin{pmatrix} b_1a_{11} + b_2a_{12} + b_3a_{13} \\ b_1a_{21} + b_2a_{22} + b_3a_{23} \end{pmatrix}$$

This actually serves to define our multiplication for us, without us having to do any more work. For we can see exactly how this multiplication needs to work: we start in the top row of the matrix and multiply the left-most element with the top-most element of the vector, and we place that in the top row of our resulting vector. We then take one step to the right in matrix (over to element a_{12}) and one step down

in the vector (down to element b_2), multiply them, and add them to the first row as well, then take one more step to the right in the matrix (to element a_{13}) and one more step down in the vector (to element b_3), multiply them, and add them to the first row. We then move down a row and all the way back to the left in the matrix, and move back up to the top in the vector, and do that procedure again only now we're placing the products in the second row of the vector on the right hand side of the equal sign. Thus, we're always placing the products of the matrix elements with the vector elements in the same row of the product vector (on the right hand side of the equal sign) that we got the matrix element from.

In more symbolic notation, this means that the i^{th} row in the product vector is the following sum:

$$\sum_{j=1}^3 b_j a_{ij}$$

and this can (and should) be checked in the above example. Note that in this expression we're only summing over the j index, so that the i index is "left-over", which corresponds to the fact that we haven't described a particular row, but rather that we're talking about the i^{th} row in the abstract. I.e., once we give a concrete value to i (for example, if we specified that we were talking about the 6^{th} row) we then have a concrete expression for the corresponding sum.

It's no coincidence that this linear map from \mathbb{R}^3 to \mathbb{R}^2 was able to be specified by 6 numbers, and the reason for this is that $6 = 3 \times 2$. In particular, we had to specify where each basis element of our basis for \mathbb{R}^3 was sent by the map, and in order to do this we needed two real numbers to specify the linear combination of the two basis elements for \mathbb{R}^2 . Since there are three basis elements for \mathbb{R}^3 , and since each of these requires two real numbers to specify its image, we find that we need 6 real numbers to specify the map. Similarly, if we were to consider maps that went the other way, i.e. from \mathbb{R}^2 to \mathbb{R}^3 , we would also need exactly 6 numbers. This is because we would need 3 numbers to specify where each element in a basis for \mathbb{R}^2 goes, and there are 2 basis elements for which we would do this.

We can start to see now how this process generalizes to any linear map from \mathbb{R}^n to \mathbb{R}^m . We first pick a basis for \mathbb{R}^n and a basis for \mathbb{R}^m . Then any linear map is determined by where it sends the basis elements of \mathbb{R}^n , and for each basis element we need m numbers (the coefficients of the m basis elements for \mathbb{R}^m) to specify where it goes. Since there are n basis elements that we need to do this for, we can specify any linear map from \mathbb{R}^n to \mathbb{R}^m by $n \times m$ real numbers.

We can then arrange these numbers in an array just as we did in the $n = 3, m = 2$ case. This array will be an $m \times n$ array, so it will have m rows and n columns (just as in our example above). Notice that in the above example, the a_{ij} which specified the linear map were labeled by their subscripts according to their placement in the array. Namely, the left-most subscript labeled that element's row, and the right-most subscript labeled the element's column. This is simply a convention that we chose to conform to, and we did so because it's a very useful one. The reason it's useful is that it helps us see immediately how to generalize all of this to the general n and m case that we're now considering. In particular, our $m \times n$ array will now look like the following:

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \cdots & a_{mn} \end{pmatrix}$$

We also know that any vector v in \mathbb{R}^n will be of the form

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix}$$

where each $b_i \in \mathbb{R}$. We can then define the multiplication of the above matrix by the above vector as

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \cdots & a_{mn} \end{pmatrix} \cdot \begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = \begin{pmatrix} b_1 a_{11} + b_2 a_{12} + \cdots + b_n a_{1n} \\ b_1 a_{21} + b_2 a_{22} + \cdots + b_n a_{2n} \\ \vdots \\ b_1 a_{m1} + b_2 a_{m2} + \cdots + b_n a_{mn} \end{pmatrix}$$

where we can check that the general procedure here is exactly the same as in the example above with $n = 3$ and $m = 2$. In particular, we start in the first row of the matrix, at the left-most element of that row, and multiply that by the top-most element of the vector. We then move one step to the right in the matrix and one step down in the vector, multiply those components together, and add them to the first row of the vector on the right. We keep repeating this until we're at the end of the row of the matrix (and equivalently at the bottom of the column of the vector), then move down a row in the matrix (and go back to the top of the vector) and repeat, putting these terms in the second row of the vector on the right. We therefore have again that the i^{th} row of the image vector (i.e., the vector $g(v) \in \mathbb{R}^m$) is

$$\sum_{j=1}^n b_j a_{ij}.$$

We note that this is the exact same expression as that in our above example, only with "3" replaced with the more general n . We've therefore found that there is a set of mn real numbers ("m times n") $\{a_{ij}\}$ such that if $v = (b_1, \dots, b_n) \in \mathbb{R}^n$, then $g(v) = ((b_1 a_{11} + b_2 a_{12} + \cdots + b_n a_{1n}), (b_1 a_{21} + b_2 a_{22} + \cdots + b_n a_{2n}), \dots, (b_1 a_{m1} + b_2 a_{m2} + \cdots + b_n a_{mn})) \in \mathbb{R}^m$. This fact, along with the next exercise, shows that there is a one-to-one correspondence between the set of all linear maps $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$. In particular, the following exercise reverses the logic that was used above. What we did above was consider a general linear map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and found—after choosing bases for the two spaces—that there is an $m \times n$ array of real numbers that specifies the linear map. What we do in the next exercise is first suppose that we're handed an $m \times n$ array of real numbers and then we find a linear map that corresponds to it.

Exercise 3.35. Let

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \cdots & a_{mn} \end{pmatrix}$$

be an $m \times n$ array of real numbers. Show that this array defines a linear map from \mathbb{R}^n to \mathbb{R}^m by sending each $v = (b_1, \dots, b_n) \in \mathbb{R}^n$ to

$$\begin{pmatrix} b_1 a_{11} + b_2 a_{12} + \cdots + b_n a_{1n} \\ b_1 a_{21} + b_2 a_{22} + \cdots + b_n a_{2n} \\ \vdots \\ b_1 a_{m1} + b_2 a_{m2} + \cdots + b_n a_{mn} \end{pmatrix}$$

in \mathbb{R}^m . In other words, show that the above multiplication rule is truly linear. Hint: all you need to do is show that it's compatible with vector addition and scalar multiplication, meaning that adding two vectors and then multiplying their sum by a matrix is the same as multiplying them both individually by the matrix and then adding them (and that multiplying it by a scalar and then a matrix is the same as multiplying it by a matrix and then a scalar).

What we've now shown is that the study of linear maps from an n -dimensional space to an m -dimensional vector space is the same as the study of $m \times n$ arrays of real numbers. One extremely important example of this arises when we ask about compositions of linear maps. Before we explore the relationship between compositions of linear maps and our arrays of numbers, let us prove the following important result.

Proposition 3.36. Let V, W , and X be finite dimensional vector spaces of dimension l, m , and n , respectively. Let $f : V \rightarrow W$ be a linear map, and let $g : W \rightarrow X$ be a linear map. Then the composition of f and g is a linear map. I.e., the map $g \circ f : V \rightarrow X$ which sends $v \in V$ to $g(f(v)) \in X$ is a linear map.

Before proving this result, let us first recall what the composition actually means. All we're doing here is first taking our element in V and sending it to W using f , and then taking the element that it lands on, $f(v) \in W$, and sending it to X using g . We can then view this as a single function from V to X .

Proof: Let us denote the function $g \circ f$ by h , so that $h : V \rightarrow X$ and $h(v) = g(f(v))$. Then all we need to do is prove that for all $v, w \in V$ and for all $a \in \mathbb{R}$, $h(v + w) = h(v) + h(w)$ and that $h(av) = ah(v)$. We'll do the former first, and we (obviously) need to rely heavily on the fact that g and f are both linear maps:

$$h(v + w) = g(f(v + w)) = g(f(v) + f(w)) = g(f(v)) + g(f(w)) = h(v) + h(w),$$

where the first equality comes from the definition of h , the second equality comes from the linearity of f , the third equality comes from the linearity of g , and the last equality again comes from the definition of h . We use the exact same logic to prove the second result:

$$h(a \cdot v) = g(f(a \cdot v)) = g(a \cdot f(v)) = a \cdot g(f(v)) = a \cdot h(v)$$

where again we first used the definition of h , then the linearity of f , then the linearity of g , then the definition of h again. This completes the proof that h is a linear map from V to W . \square

Now that we know that the composition of linear maps is again a linear map, the natural question to ask would be whether or not this fact can be reflected in our new "array-vector multiplication" formalism. The answer is yes.

For suppose we have a linear map $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and another linear map $g : \mathbb{R}^m \rightarrow \mathbb{R}^l$ (note that it is extremely important for the "inner" vector space (namely, the codomain of f and the domain of g) to be the same, for otherwise our composition isn't defined). We then know, from the above discussion, that once we pick a basis for our spaces $\mathbb{R}^n, \mathbb{R}^m$, and \mathbb{R}^l , we'll have an $m \times n$ array of real numbers corresponding to f , and an $l \times m$ array of real numbers corresponding to g . But we also know, from the above proposition, that we have a linear map $g \circ f : \mathbb{R}^n \rightarrow \mathbb{R}^l$, and therefore we should have an $l \times n$ array of real numbers that correspond to the linear map $h = g \circ f$. Moreover, it seems natural that we should be able to find this $l \times n$ array of numbers using only our other two arrays of numbers (the $m \times n$ one and the $l \times m$ one), since the composition of the two functions only contains the information of the two functions, and all of this information is contained in the first two arrays of numbers. We therefore need a way of generating our $l \times n$ array from the other two arrays, and this array needs to send $v \in \mathbb{R}^n$ to the same vector in \mathbb{R}^l that it would go to if we first applied f and then applied g .

To see how to do this, we simply do it! (Note: the following algebra might look pretty scary, but it's actually not too bad! Just take a deep breathe and follow it through one step at a time). First let us suppose that the $m \times n$ array of real numbers

$$\begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \cdots & a_{mn} \end{pmatrix}$$

corresponds to $f : \mathbb{R}^n \rightarrow \mathbb{R}^m$, and let us suppose that the $l \times m$ array of real numbers

$$\begin{pmatrix} b_{11} & b_{12} & b_{13} & \cdots & b_{1m} \\ b_{21} & b_{22} & b_{23} & \cdots & b_{2m} \\ \vdots & \cdots & \cdots & \cdots & \vdots \\ b_{l1} & b_{l2} & b_{l3} & \cdots & b_{lm} \end{pmatrix}$$

corresponds to $g : \mathbb{R}^m \rightarrow \mathbb{R}^l$. Then let us choose an arbitrary vector $v = (c_1, c_2, \dots, c_n) \in \mathbb{R}^n$ and see where it lands in \mathbb{R}^l . After first acting on this vector with f (i.e., by multiplying it by the a_{ij} matrix), we're left with the vector

$$\begin{pmatrix} c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n} \\ c_1 a_{21} + c_2 a_{22} + \dots + c_n a_{2n} \\ \vdots \\ c_1 a_{m1} + c_2 a_{m2} + \dots + c_n a_{mn} \end{pmatrix}$$

in \mathbb{R}^m . But this is a perfectly good vector in \mathbb{R}^m , and therefore we can now hit this with the b_{ij} matrix to send it into \mathbb{R}^l . Now, however, we simply have a more messy expression for the vector. Instead of dealing with a vector of the form (b_1, \dots, b_m) , say, we're dealing with the above vector, whose individual components is some complicated sum of products of numbers. But what's important to notice is that the expression $c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n}$ is a perfectly good real number. Let us "wrap up" these expressions to simpler looking things so that we can manipulate the simpler expressions, and then "unwrap" those expressions at the end. This way it'll be a bit less scary.

In particular, let's let $d_i = c_1 a_{i1} + c_2 a_{i2} + \dots + c_n a_{in}$, so that we have

$$\begin{pmatrix} c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n} \\ c_1 a_{21} + c_2 a_{22} + \dots + c_n a_{2n} \\ \vdots \\ c_1 a_{m1} + c_2 a_{m2} + \dots + c_n a_{mn} \end{pmatrix} = \begin{pmatrix} d_1 \\ d_2 \\ \vdots \\ d_m \end{pmatrix}$$

which is a much nicer expression. Now we can manipulate this easier expression and then "unwrap" the d_i 's at the end. Accordingly, let's hit this vector with our b_{ij} matrix to get the \mathbb{R}^l vector

$$\begin{pmatrix} d_1 b_{11} + d_2 b_{12} + \dots + d_m b_{1m} \\ d_1 b_{21} + d_2 b_{22} + \dots + d_m b_{2m} \\ \vdots \\ d_1 b_{l1} + d_2 b_{l2} + \dots + d_m b_{lm} \end{pmatrix}.$$

But now we can (and should) unwrap the expressions that the d_i are standing for, and when we do so we get the following extremely gross expression:

$$\begin{pmatrix} (c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n}) b_{11} + (c_1 a_{21} + c_2 a_{22} + \dots + c_n a_{2n}) b_{12} + \dots + (c_1 a_{m1} + c_2 a_{m2} + \dots + c_n a_{mn}) b_{1m} \\ (c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n}) b_{21} + (c_1 a_{21} + c_2 a_{22} + \dots + c_n a_{2n}) b_{22} + \dots + (c_1 a_{m1} + c_2 a_{m2} + \dots + c_n a_{mn}) b_{2m} \\ \vdots \\ (c_1 a_{11} + c_2 a_{12} + \dots + c_n a_{1n}) b_{l1} + (c_1 a_{21} + c_2 a_{22} + \dots + c_n a_{2n}) b_{l2} + \dots + (c_1 a_{m1} + c_2 a_{m2} + \dots + c_n a_{mn}) b_{lm} \end{pmatrix}.$$

Wow. That's REALLY gross. But remember what we're after. We're after an $l \times m$ array of numbers that sends the vector $v = (c_1, \dots, c_n) \in \mathbb{R}^n$ to the vector in \mathbb{R}^l that is arrived at after first applying f and then applying g . But the above nasty expressions **is** the vector that (c_1, \dots, c_n) arrives at after applying these two linear maps. Thus, what we want to do is shake the above nasty expression into the same form that we had when we had a single matrix multiplying a vector. Namely, we want to shake the above gross expression into the following form:

$$\begin{pmatrix} c_1(\text{something}) + c_2(\text{something}) + \dots + c_n(\text{something}) \\ c_1(\text{something}) + c_2(\text{something}) + \dots + c_n(\text{something}) \\ \vdots \\ c_1(\text{something}) + c_2(\text{something}) + \dots + c_n(\text{something}) \end{pmatrix}.$$

In order to shake our expression into this form, we just have to collect all the coefficients to each individual c_i . In other words, the above nasty expression can simply be rewritten as follows, where all we do is move around some of the terms:

$$\begin{pmatrix} (b_{11}a_{11} + b_{12}a_{21} + \dots + b_{1m}a_{m1})c_1 + (b_{11}a_{12} + b_{12}a_{22} + \dots + b_{1m}a_{m2})c_2 + \dots + (b_{11}a_{1n} + b_{12}a_{2n} + \dots + b_{1m}a_{mn})c_n \\ (b_{21}a_{11} + b_{22}a_{21} + \dots + b_{2m}a_{m1})c_1 + (b_{21}a_{12} + b_{22}a_{22} + \dots + b_{2m}a_{m2})c_2 + \dots + (b_{21}a_{1n} + b_{22}a_{2n} + \dots + b_{2m}a_{mn})c_n \\ \vdots \\ (b_{l1}a_{11} + b_{l2}a_{21} + \dots + b_{lm}a_{m1})c_1 + (b_{l1}a_{12} + b_{l2}a_{22} + \dots + b_{lm}a_{m2})c_2 + \dots + (b_{l1}a_{1n} + b_{l2}a_{2n} + \dots + b_{lm}a_{mn})c_n \end{pmatrix}$$

so that the i^{th} row of this vector is given by $\sum_{k=1}^l (\sum_{j=1}^m b_{ij}a_{jk})c_k$. But this immediately tells us what our $l \times n$ matrix is. Namely, this sum is the exact form of our previous arrays, where our entries are now entire sums. Namely, the entry in the i^{th} row and k^{th} column of our $l \times n$ matrix is the real number $\sum_{j=1}^m b_{ij}a_{jk}$. Note that the subscripts actually work out, meaning that the i subscript on b_{ij} ranges from 1 to l , and the k subscript on a_{jk} ranges from 1 to n , so that this actually does define for us an $l \times n$ matrix of real numbers. Moreover, it is clear that if we first construct this array and then multiply our vector (c_1, \dots, c_n) by it, we'll end up at the same vector in \mathbb{R}^l as we would have if we first multiplied by the $\{a_{ij}\}$ matrix and then by the $\{b_{ij}\}$ matrix. Thus, we've found our "composition matrix", which is the matrix of the linear function obtained from the composition of the matrices of the two linear functions that we started with.

In symbols, we've found that the following holds:

$$\begin{pmatrix} b_{11} & b_{12} & b_{13} & \cdots & b_{1m} \\ b_{21} & b_{22} & b_{23} & \cdots & b_{2m} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ b_{l1} & b_{l2} & b_{l3} & \cdots & b_{lm} \end{pmatrix} \cdot \begin{pmatrix} a_{11} & a_{12} & a_{13} & \cdots & a_{1n} \\ a_{21} & a_{22} & a_{23} & \cdots & a_{2n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & a_{m3} & \cdots & a_{mn} \end{pmatrix} \cdot \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix} = \quad (3.14)$$

$$\begin{pmatrix} (b_{11}a_{11} + b_{12}a_{21} + \dots + b_{1m}a_{m1}) & (b_{11}a_{12} + b_{12}a_{22} + \dots + b_{1m}a_{m2}) & \cdots & (b_{11}a_{1n} + b_{12}a_{2n} + \dots + b_{1m}a_{mn}) \\ (b_{21}a_{11} + b_{22}a_{21} + \dots + b_{2m}a_{m1}) & (b_{21}a_{12} + b_{22}a_{22} + \dots + b_{2m}a_{m2}) & \cdots & (b_{21}a_{1n} + b_{22}a_{2n} + \dots + b_{2m}a_{mn}) \\ \vdots & \vdots & \ddots & \vdots \\ (b_{l1}a_{11} + b_{l2}a_{21} + \dots + b_{lm}a_{m1}) & (b_{l1}a_{12} + b_{l2}a_{22} + \dots + b_{lm}a_{m2}) & \cdots & (b_{l1}a_{1n} + b_{l2}a_{2n} + \dots + b_{lm}a_{mn}) \end{pmatrix} \cdot \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_n \end{pmatrix}$$

Moreover, this holds for all $(c_1, c_2, \dots, c_n) \in \mathbb{R}^n$. This suggests not only that we can define a "matrix multiplication" which takes the two matrices on the left of the above equation and gives back the matrix on the right as a product, but it also tells us exactly **how** we want to define this multiplication. Namely, if

our two matrices are given by $\{a_{ij}\}$ with $1 \leq i \leq m$ and $1 \leq j \leq n$ and $\{b_{ij}\}$ with $1 \leq i \leq l$ and $1 \leq j \leq m$, then the product matrix is an $l \times n$ matrix with elements

$$c_{ij} = \sum_{k=1}^m b_{ik}a_{kj}.$$

It is important to see how this relates to our rule of multiplying a vector by a matrix. Namely, multiplying two matrices together is very similar to multiplying a matrix to a vector, and in fact can be viewed as **exactly** the same if we view a vector as an $n \times 1$ matrix. General matrix multiplication goes as follows. When we're considering the element in the i^{th} row and j^{th} column of our matrix (as we are in the above expression), we simply go to the i^{th} row of the matrix on the left of the product (in this case the $\{b_{kl}\}$ matrix) and the j^{th} column of the matrix on the right of the product (in this case the $\{a_{kl}\}$ matrix) and multiply them just as we would in the case of a vector and a matrix. Namely, we start on the far left of the row in the left matrix and the very top of the column in the right matrix, multiply those two entries, then take one step to the right and one step down (in the two respective matrices), multiply those two entries, and so on, until we reach the end of the row (which will also mean we've reached the end of the column). We sum up all of these individual products and put this sum into the corresponding slot of the product matrix (namely, the slot in the i^{th} row and the j^{th} column). It should be checked that the above expression is in alignment with this method.

We must note how important it is for the right hand subscript of the b_{ij} 's to run over the same number (namely m) as the left hand subscript of the a_{ij} 's. This is the only way that this definition is well-defined. This is nothing but the statement that the composition of two linear maps is only defined when the codomain of the "first" map is the same as the domain of the "second" map. For recall that our two linear maps were from \mathbb{R}^n to \mathbb{R}^m and from \mathbb{R}^m to \mathbb{R}^l , and this means that the index structure of our two matrices will always work out properly. However, if we want to try to define this matrix multiplication in a purely abstract way, without reference to linear maps, then we're forced to limit ourselves to matrices whose index structures work out properly. Namely, if we have two matrices A and B and we want to multiply them together using the above matrix multiplication to get the matrix $A \cdot B$, then it must be the case that the number of columns of A is the same as the number of rows of B .

This last fact leads us immediately to conclude that the above multiplication rule is not commutative, so that for any two matrices A and B it is not always the case that $A \cdot B = B \cdot A$. In fact, even if $A \cdot B$ exists, it is not guaranteed that $B \cdot A$ does, for it might not be well defined in the sense of having the right index structure. We could limit ourselves to only considering "square" matrices—meaning matrices that have the same number of rows as columns—for this will ensure that the index structure will always work out to give a well-defined product. However even this doesn't ensure that this multiplication rule is commutative, as the next exercise shows.

Exercise 3.37. Show that

$$\begin{pmatrix} 3 & 2 \\ 4 & 1 \end{pmatrix} \cdot \begin{pmatrix} 5 & 3 \\ 2 & 2 \end{pmatrix} \neq \begin{pmatrix} 5 & 3 \\ 2 & 2 \end{pmatrix} \cdot \begin{pmatrix} 3 & 2 \\ 4 & 1 \end{pmatrix}$$

Exercise 3.38. Let A and B both be $n \times n$ matrices defined respectively by $\{a_{ij}\}$ and $\{b_{ij}\}$, with $1 \leq i, j \leq n$. Show that, in general, $A \cdot B \neq B \cdot A$. Give some conditions for which $A \cdot B = B \cdot A$.

It should be noted that we can take this matrix multiplication that we've defined and apply it to square matrices (so that we always know that the index structure "works" and that the product will always be defined) and turn the set of square matrices of a particular dimension (i.e., $n \times n$ matrices for some particular n) into a group. However, we need to leave some matrices out of this group because some matrices don't have an inverse under this multiplication rule, much like 0 doesn't have an inverse under normal

multiplication. In order to figure out whether or not a matrix has an inverse (and therefore to figure out which matrices we need to exclude from the group), we need to study something called the **determinant** of a matrix. We won't go into this here, although it is indeed a very important subject. It will most certainly be covered in any linear algebra or remotely advanced physics course. For now, we simply note that such a construction of matrix groups is possible, and indeed such matrix groups (especially matrices with complex entries, the meaning of which will become clear in the next chapter) are some of the most important objects in all of physics (believe it or not).

The final thing that we'll discuss in regards to vector spaces and matrices (for now) is what happens to the formalism that we've developed when we change the basis that we're considering. After all, many of the ideas and proofs that we've found have relied on first choosing one of an infinite number of bases, and doing our calculations in reference to this choice. For example, we recall that a matrix representing a linear map is full of numbers that correspond to a particular way of writing one vector space's basis elements in terms of another vector space's basis elements. However, if we change which basis we chose, then surely the numbers in the matrix representing this map will change as well. Thankfully it turns out that we can fully determine precisely how our matrix will change when we change our basis by simply using more matrix machinery. Let us see how this is done.

(Note: the remaining discussion in this section is considerably more fast-paced and more algebraically complex than much of the above discussion. If the following material is rather abstruse on the first reading, have no fear, because it is not meant to provide full details of every calculation. Therefore, only worry about the remaining material in this section if you fully understand all of the above material and are up for a challenge and willing to work out some of the details on your own. If not, don't worry, we won't need the material in the remainder of this chapter for what is to come in the future of this text.)

Let $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map, let $E = \{e_i\}$ ($1 \leq i \leq n$) be a basis for \mathbb{R}^n , and let $F = \{f_i\}$ ($1 \leq i \leq m$) be a basis for \mathbb{R}^m . Then let us suppose that A is the $m \times n$ matrix (whose elements are $\{a_{ij}\}$ with $1 \leq i \leq m$, $1 \leq j \leq n$) representing g with respect to the two bases that we've chosen. Now let's suppose that we pick another basis for \mathbb{R}^n , and let's call this basis $E' = \{e'_i\}$. There will then be another matrix representing g with respect to the bases E' and F , and let's call this matrix $A' = \{a'_{ij}\}$. What we seek is some relationship between A and A' , and our motivation for why there must be one is that they are matrices that represent the same linear map, albeit with respect to two different bases.

The key to finding this relationship is realizing that we can express one basis in terms of linear combinations of the other basis. Before seeing how this goes, let us first recall precisely what the preceding paragraph tells us. What we have is two different realizations of the same linear map in terms of two different matrices, A and A' . The matrix A tells us where g takes the basis E with respect to the basis F , and the matrix A' tells us where g takes the basis E' with respect to the basis F . In particular, we have that

$$g(e_i) = \sum_{j=1}^n a_{ji} f_j$$

and that

$$g(e'_i) = \sum_{j=1}^n a'_{ji} f_j.$$

Now let us suppose that the basis E is the standard basis that we've used above, so that a vector $v = (c_1, \dots, c_n) \in \mathbb{R}^n$ can also be written $v = c_1 e_1 + \dots + c_n e_n = \sum_{j=1}^n c_j e_j$. Then we have that

$$g(v) = g\left(\sum_{k=1}^n c_k e_k\right) = \sum_{k=1}^n c_k g(e_k) = \sum_{k=1}^n c_k \left(\sum_{j=1}^n a_{jk} f_j\right) = \sum_{j=1}^n \left(\sum_{k=1}^n c_k a_{jk}\right) f_j,$$

where we have used the linear property of g to write $g(v)$ as a linear combination of the basis elements in F , with coefficients $\sum_{k=1}^n c_k a_{jk}$. We can now do the same thing using the basis E' , but we have to first note that there will now be a new set of coefficients for $v \in \mathbb{R}^n$ when it is expressed as a linear combination of the $\{e'_i\}$. We therefore must write $v = d_1 e'_1 + \dots + d_n e'_n = \sum_{j=1}^n d_j e'_j$. We then follow the exact same logic as above to find that

$$g(v) = \sum_{j=1}^n \left(\sum_{k=1}^n d_k a'_{jk}\right) f_j.$$

We then note that

$$0 = g(v) - g(v) = \sum_{j=1}^n \left(\sum_{k=1}^n c_k a_{jk} - d_k a'_{jk}\right) f_j,$$

where we have simply expanded $g(v)$ in terms of the two different linear combinations that we have, with respect to the two different bases. Due to the linear independence of the $\{f_j\}$, we know that each coefficient must vanish independently, and so we therefore have that

$$\sum_{k=1}^n c_k a_{jk} = \sum_{k=1}^n d_k a'_{jk}.$$

Now we use the fact that each basis can be written as a linear combination of the other. In particular, we have that for some $\{p_{ij}\} \subset \mathbb{R}$, with $1 \leq i, j \leq n$,

$$e'_j = \sum_{i=1}^n p_{ij} e_i.$$

We then have that

$$v = \sum_{k=1}^n d_k e'_k = \sum_{k=1}^n d_k \left(\sum_{l=1}^n p_{lk} e_l\right) = \sum_{l=1}^n \left(\sum_{k=1}^n d_k p_{lk}\right) e_l.$$

But we also know that $v = \sum_{l=1}^n c_l e_l$, so that we can use the same trick of adding and subtracting the same thing, but using two different expressions for it to find the following:

$$0 = v - v = \sum_{l=1}^n \left(c_l - \sum_{k=1}^n d_k p_{lk}\right) e_l.$$

Then, using the linear independence of the $\{e_l\}$, we know that each coefficient must vanish individually, so that we have ($\forall 1 \leq l \leq n$)

$$c_l = \sum_{k=1}^n d_k p_{lk}.$$

We now rename our indices and plug this into

$$\sum_{k=1}^n c_k a_{jk} = \sum_{k=1}^n d_k a'_{jk}$$

to get that

$$\sum_{l=1}^n \left(\sum_{k=1}^n d_k p_{lk} \right) a_{jl} = \sum_{k=1}^n d_k a'_{jk}.$$

Note that between the above two expressions we simultaneously made the aforementioned substitution as well as renamed our indices. We then rearrange this sum to get

$$\sum_{k=1}^n d_k \left(\sum_{l=1}^n p_{lk} a_{jl} \right) = \sum_{k=1}^n d_k a'_{jk},$$

and this implies that

$$\sum_{k=1}^n d_k (a'_{jk} - \sum_{l=1}^n a_{jl} p_{lk}) = 0.$$

Now we use one final clever bit of logic, and that is the following. We recall that the $\{d_k\}$ are nothing but the coefficients of the vector $v \in \mathbb{R}^n$ in terms of the basis E' , and that v was an arbitrary vector. Therefore, the above expression should hold for all possible d_k , since nothing that we've done here has depended on which particular v we chose initially. Thus, the only way for the above sum to be zero is if it is zero for each index k , therefore giving us that

$$a'_{jk} = \sum_{l=1}^n a_{jl} p_{lk}.$$

We now note that if we form an $n \times n$ matrix P out of the entries $\{p_{lk}\}$, then the above expression is nothing but the statement that $A' = A \cdot P$, where " \cdot " is precisely the matrix multiplication that we defined above. In particular, the above expression says that the j, k component of A' is precisely the j, k component of $A \cdot P$. This can also be written $(A')_{jk} = (A \cdot P)_{jk}$. Note that we could have written the E basis in terms of the E' basis, instead of writing the E' basis in terms of the E basis as we did here, and we would find a similar expression for the matrices A and A' . Moreover, we also could have changed our basis F to some other basis F' for the codomain, and we would again get a similar expression relating A and A' (though now the "change of basis matrix" will be on the other side!). We could discuss these issues in much greater detail, but for the sake of eventually learning about complex numbers, we'll rather arbitrarily stop the discussion here and end this section with an extremely vague and large exercise.

Exercise 3.39. Carry through the calculations mentioned in the above paragraph, finding the expressions relating A and A' in terms of a "change of basis matrix" for the various changes of basis suggested here.

Chapter 4

Complex Numbers and Complex Vector Spaces

4.1 Introduction

We're now going to introduce one of the biggest superstars in all of mathematics and physics, and that is the complex number. It will be difficult to see in the extremely short development that we provide here, but it turns out that the very simple and seemingly unphysical definition of a complex number was one of the most important breakthroughs in the history of mathematics. When we define it, it will not seem real. It will seem like a trick at best, and a lie at worst. In fact, an unfortunate bit of mathematical terminology has dubbed half of each complex number as "imaginary" (this will make sense shortly). This is unfortunate because the beauty of complex numbers in mathematics and the importance of complex numbers in describing our physical world at its most fundamental level have both provided strong evidence that complex numbers are hardly just a mathematical trick. In fact, as one studies these wonderful numbers and the role that they play in physics, it very quickly becomes obvious that there is something very deep and mystifying about their existence.

Unfortunately we will not be able to explore much of this beauty here. The reason for this is two-fold. First, it is likely that the mathematical background needed to fully comprehend some of the most beautiful facts about these numbers has not been developed at this stage. With infinite time, we could develop that background here, but unfortunately such quantities of time are not available to us. The second reason we won't see much of the beauty of these numbers is that we won't need all of these beautiful details for what we'll do in the rest of the text. I've therefore made the choice that would be blasphemous to many mathematicians, which is to sacrifice beauty for pragmatism. Accordingly, we will only develop our theory of complex numbers to a point that we can use them in complex vector spaces, as complex vector spaces are of fundamental importance to quantum mechanics (as we shall see in our short description of some quantum mechanical systems in the last chapter of this text).

Therefore, the reader will have to take much of what has been said about the profound beauty of this mathematical structure on faith, and/or wait until this beauty shows itself in other, more advanced courses. Hopefully our development here will be sufficient to show that these numbers are perfectly well-defined and that they can and should be viewed as being very "real". Without further ado, let us start to unravel complex numbers.

4.2 Complex Numbers

Complex numbers include every type of number that we've already studied, but extends them in a very important way. As we've already seen, we have $\mathbb{N} \subset \mathbb{Z} \subset \mathbb{Q} \subset \mathbb{R}$, and now I claim that there is a set of numbers, the complex numbers, denoted by \mathbb{C} , such that $\mathbb{R} \subset \mathbb{C}$.

Let us first recall how we built up the numbers that we already know about. The natural numbers $\mathbb{N} = \{0, 1, 2, 3, 4, 5, \dots\}$ need to, in some sense, be **assumed** to exist (as opposed to being derived from more fundamental principles (although there are people who investigate those sorts of questions, and make great progress towards them as well)). This is obviously a natural and important assumption to make. If we want to make this set into a group under addition, so that 0 is our identity element, then we need to include the additive inverses of each number. This means that we need to include the negatives of each whole number, and in so doing we find ourselves with the set of integers $\mathbb{Z} = \{\dots, -3, -2, -1, 0, 1, 2, 3, \dots\}$. If we then want to make the integers a group under multiplication, so that 1 is the identity element (and where we have to exclude the element 0), then we need to include the fractions $\frac{1}{a}$ where $a \in \mathbb{Z}$. Clearly, for any a other than 1 and -1 , these fractions will not be in \mathbb{Z} , and we therefore need to extend \mathbb{Z} to a new set. This motivates us to include **all** fractions, and not just those that are "1 over an integer", and in so doing we've found ourselves with the set \mathbb{Q} of all fractions.

We then can ask about our ability to solve the equation $x^2 - 2 = 0$. I.e., do we have a system of numbers S such that there is an $x \in S$ whose square is equal to 2? It turns out, as we saw at the end of the first chapter, that if we stick only with \mathbb{Q} then the answer to this question is no. This is because there is no fraction whose square is 2. It turns out that the number whose square is 2 is

$$1.41421356237\dots$$

where the " \dots " means that this decimal continues on forever while never repeating any kind of pattern at all. If we want a system of numbers that includes the solution to $x^2 - 2 = 0$, then we need to extend \mathbb{Q} . This extension is precisely \mathbb{R} , which we've defined to be the set of all decimals of infinite extent, where fractions and whole numbers can be recovered by letting the decimals repeat, so that $1 = 1.0000\dots$ and $\frac{1}{3} = .3333\dots$

(Note: this is admittedly a very crude way of generating our various number systems. There has been lots of work done, and lots of success had, at putting these number systems on much more rigorous foundations. In fact, there is still work left to do in this regard, like deriving the existence of the natural numbers from first principles, and if this sort of thing fascinates the reader then a career studying the foundations of mathematics would be a good one!)

We may now ask whether or not our numbers can be extended again, and the answer is yes. This may be surprising, seeing as the set of numbers that we've called "real numbers" already seems to cover all possible meaningful numbers that we can imagine. In fact, we can already get the sense that our numbers are somehow "not real", in that we can second guess just how "real" an infinite decimal is. Either way, we do get a sense that real numbers are very "real" due to the very geometric motivation for expanding our set of fractions, as for example a right triangle with two sides of length 1 forces us to say that the hypotenuse is of length $\sqrt{2}$. Moreover, when we draw the number line \mathbb{R} and stare at it, we may find it difficult to imagine a possible way of extending it. After all, every "extension" we've found so far has included "filling in" the number line (from natural numbers to integers, from integers to fractions, and from fractions to real numbers), and now the number line seems full!

However, it turns out that there is another meaningful extension of our numbers, and moreover that there really is only **one** more such extension (there are indeed **many** other extensions, but each one of

them forces us to sacrifice certain properties of numbers that we normally want to keep, and unfortunately we won't be able to explore these other cool extensions in this work (for the interested reader, a couple keywords are "quaternions" and "octonions"). It turns out the the million dollar question to ask is whether or not we can solve the following equation in our system of numbers: $x^2 + 1 = 0$. That's it. That's the question that sparks what is one of the most elegant fields in all of mathematics—the study of complex numbers.

We note that the equation $x^2 + 1 = 0$ is equivalent to the equation $x^2 = -1$, and that indeed the number system that is currently our most expansive— \mathbb{R} —does not solve this equation. This easily seen when we note that any real number squared is non-negative. In particular, any positive number times a positive number is a positive number, and any negative number times a negative number is positive (and of course $0^2 = 0$). However, our equation is asking for a number that, when multiplied to itself, gives a negative number (namely, -1). This is clearly impossible in the real numbers, because we know that x can't be zero, positive, or negative and that **every** real number is either zero, positive, or negative.

But generalizing the real numbers is no easy task, for it's pretty hard to think of anything else that we can do to them! Recall that in our first extension we included negative numbers, then fractions, and then infinite decimals. What else can we do to our numbers? What other freedom can we give them?

It turns out that we shouldn't (and can't) **change** the numbers we have, but rather we should **add** directly to them. In other words, we should simply **define** a new number and add it in to our system of numbers. What we'll do is **define** a number to be the square root of -1 , and see where it gets us. This may seem rather arbitrary, and at first glance it is, but when mankind figured this out a few centuries ago it was one of the most significant advancements in all of mathematics.

Accordingly, let us **define** a new number, denoted by i , whose square is negative one, so that $i^2 = -1$. Now let us define a multiplication between these number and real numbers, so that a product like $5 \times i$ can exist. We define this product to be nothing but the number $5i (= i5)$. We can then force the multiplication of numbers of this kind to be both commutative and associative, so that: $5i \times 4i = (5 \times 4) \times (i \times i) = 20 \times -1 = -20$. Now, all of the sudden we not only can solve the equation $x^2 = -1$ by letting $x = i$, but we can also solve the equation $x^2 = -a$ for any $0 < a \in \mathbb{R}$ by simply letting $x = i\sqrt{a}$. Let us call any number that is of the form "a real number times i " an **imaginary number**. In other words, the set of imaginary numbers is precisely $\{ia \mid i^2 = -1, a \in \mathbb{R}\}$.

One of the most fundamental qualities about numbers, though, is that we can add them together to get other numbers. This is so obvious that it hardly needs mentioning, until we realize that we have yet to discuss how to add numbers of this new form. For example, can we add 5 and i ? We can, if we define how to do so. In particular, let us simply **define** the addition of 5 and i to be the number $5 + i$. Well, that's a little anticlimactic, but the important thing to note is that we're now viewing $5 + i$ as a **single number**. But what type of number is this? It's precisely a **complex number**. Let us now just go ahead and make the general definition of a complex number.

Definition 4.1. A **complex number** is a number of the form $a + ib$ with $a, b \in \mathbb{R}$. Given a complex number $a + ib$, we say that the **real part** of the complex number is a , and denote this by $Re(a + ib) = a$. Similarly, we say that the **imaginary part** of a complex number is b , and denote this by $Im(a + ib) = b$.

We note that given any real number $a + ib$, we have that $Re(a + ib) \in \mathbb{R}$ and $Im(a + ib) \in \mathbb{R}$. We also note that we recover our real numbers by considering the following subset of complex numbers: $\{a + ib \mid b = 0\}$. When viewed in this way, it is clear that $\mathbb{R} \subset \mathbb{C}$. In fact, it is clear that there is a one-to-one correspondence between \mathbb{C} and $\mathbb{R} \times \mathbb{R}$, as the next exercise shows.

Exercise 4.2. Define a bijective function $f : \mathbb{C} \rightarrow \mathbb{R} \times \mathbb{R}$.

We now **define** the addition of two complex numbers to be the addition of their real and imaginary parts, respectively, so that $(a + ib) + (c + id) = (a + b) + i(c + d)$. Therefore the real (imaginary) part of the sum of two complex numbers is equal to the sum of the real (imaginary) parts. Thus the addition of complex numbers, as we've so defined them, is rather straightforward. What makes these numbers so gorgeous is the weird way that we multiply them. Let us explore this now.

We know how to multiply real numbers and we know how to multiply imaginary numbers, but how should we define the multiplication of complex numbers that have non-zero real and imaginary parts? Let's look to real numbers for motivation. Recall that if we have a sum of real numbers like $(a + b)$ and another sum like $(c + d)$, then the product of these two sums is distributive:

$$(a + b) \times (c + d) = ac + ad + bc + bd.$$

Let us therefore **force** our multiplication of complex numbers to be distributive in the same way. By doing so, we find a very interesting multiplication rule:

$$(a + ib) \times (c + id) := ac + iad + ibc + i^2bd = (ac - bd) + i(ad + bc).$$

(We now drop the " \times " symbol and simply recall that when two numbers are placed next to each other it is implied that multiplication is involved. We've already done this in places, so let's just do this everywhere). Therefore, $Re((a + ib)(c + id)) = ac - bd$ and $Im((a + ib)(c + id)) = ad + bc$.

Exercise 4.3. Calculate the following:

- 1) $(5 + 3i)(-4 - 2i)$
- 2) $((2 + i) + (-3 + 4i))(7 + 4i)$.

Exercise 4.4. Find two square roots of i .


One thing to note is that we often write complex numbers using a single symbol, such as z . When we say that $z \in \mathbb{C}$, we're implying that there are two real numbers $a, b \in \mathbb{R}$ such that $z = a + bi$. When speaking of the abstract number it is easier to refer to (and write) one symbol— z —but when doing explicit calculations we often need to resort to the "inner workings" of z , which we know is a real and an imaginary part: $a + bi$. In other words, the statement " $z \in \mathbb{C}$ " is equivalent to the statement " $\exists a, b \in \mathbb{R}$ such that $z = a + ib$ ".

4.3 Geometry of Complex Numbers

As we know, whenever we draw the 2-dimensional plane on a piece of paper, we're "really" drawing a pictorial representation of the set $\mathbb{R} \times \mathbb{R} = \mathbb{R}^2$. However, we saw above that the sets \mathbb{R}^2 and \mathbb{C} are in a bijective correspondence with each other. By exploiting this correspondence, we can find a very nice geometric representation of complex numbers. This will be completely analogous to the geometric representation of real numbers that we use by drawing a line (which we often refer to as "the number line") and depicting a real number as a dot or tick on this line.

All we do is simply identify \mathbb{C} with $\mathbb{R} \times \mathbb{R}$ by identifying $Re(z)$ with the first factor of \mathbb{R} and $Im(z)$ with the second factor, for all $z \in \mathbb{C}$. Thus, if we've drawn our 2-dimensional plane with one \mathbb{R} going horizontally and the other going vertically (which is the standard way of drawing our " x - y axes"), then $Re(z)$ is plotted along the horizontal line and $Im(z)$ is plotted along the vertical line. We then draw a single dot at the point $(Re(z), Im(z)) \in \mathbb{R} \times \mathbb{R}$, as shown in figure 4.1, where we plot the complex number $z = a + bi \in \mathbb{C}$, and where $Re(z) = a$ is plotted on the horizontal axis, and $Im(z) = b$ is plotted on the vertical axis.

Figure 4.1:



This single dot is analogous to the single dot that we draw on a number line to represent a real number, and this single dot similarly represents a single complex number. Thus, the analogue of "the real number line" is "the complex number plane". We have a full plane's worth of numbers, and each dot that we draw on this plane is a **single** number.

One quality about complex numbers that we'll see come up several times is that we should truly view them as being single-dimensional. This means that even though we **pictorially represent** complex numbers on a plane—which we're used to thinking of as 2-dimensional—we should really view this as a one dimensional object, namely the complex version of a line. The difference, however, is that this single dimension is a single **complex** dimension. Thus, a plane is only 2-dimensional when we view it as an object that is described by **real numbers**, but it is one-dimensional when we view it as being described by **complex numbers**. This is because we define dimensionality, roughly, as the number of coordinates required to tell us "where we are". On a plane, it takes two real numbers to specify where a point is, and therefore when we're talking about real dimensions, a plane is two-dimensional. But it only takes a single complex number to specify where a point on the plane is, and thus a plane is one-complex-dimensional.

One important bit of complex geometry is the fact that the Pythagorean Theorem plays an important role in the description of complex numbers. In particular, any real number $z \in \mathbb{C}$ is described by the two real coordinates $(\text{Re}(z), \text{Im}(z))$ in \mathbb{R}^2 . We can then ask how far away this complex number is from the origin $(0, 0)$. This means that if we were to draw the point $z \in \mathbb{C}$ in the plane and then draw a straight line from $(0, 0)$ to $(\text{Re}(z), \text{Im}(z))$, we want to know about the length of this line. Well, this is easy! Namely, the Pythagorean theorem tells us that the **square** of this length $\text{Re}(z)^2 + \text{Im}(z)^2$. This idea motivates the following subsection, which explores the complex conjugate of any given complex number.

4.4 Complex Conjugates

Let us take what will look like a detour, even though we'll very quickly wind up back to where we just left off, which is the distance from any complex number to the origin.

Take some arbitrary complex number $z = a + ib$. Let us define the **complex conjugate** of this number to be the same number with i replaced by $-i$. If we denote the complex conjugate of z by \bar{z} , then we have that $\bar{z} = a - ib$. This definition seems harmless enough. We note that in the complex plane, the complex conjugate of any number is the "reflection about the $\text{Re}(z)$ -axis of the original number, as can be seen in figure 4.1 by considering the complex number $a - ib$.

Now let us ask what the product of any complex number times its complex conjugate is. Namely, given any $z = a + ib \in \mathbb{C}$, what is $z\bar{z}$? Let's simply work it out:

$$z\bar{z} = (a + ib)(a - ib) = a^2 - i^2b^2 + i(ab - ab) = a^2 + b^2 = \operatorname{Re}(z)^2 + \operatorname{Im}(z)^2. \quad (4.1)$$

Thus, we've found that the product $z\bar{z}$ is nothing but the squared distance from the point in the plane representing the complex number to the origin. Therefore, the positive square root of this product ($+\sqrt{z\bar{z}}$) is precisely the distance from the complex number's representative point to the origin. This number is so important we give it a special name and notation. We call it the **modulus** of z , and denote it by $|z|$, so that $|z| = \sqrt{z\bar{z}}$ and $|z|^2 = z\bar{z}$ (the latter is often referred to as " z mod squared").

Exercise 4.5. Show that the modulus of z is equal to the modulus of \bar{z} (i.e., $|z| = |\bar{z}|$), the complex conjugate of z , for all $z \in \mathbb{C}$. Also show that for all $z \in \mathbb{R} \subset \mathbb{C}$, $\bar{z} = z$ (where we view $\mathbb{R} \subset \mathbb{C}$ as the set of complex numbers with zero imaginary part).

We're now in a position to talk about the multiplicative inverses of complex numbers. Recall that for any non-zero real number $a \in \mathbb{R}$, the multiplicative inverse of a is $\frac{1}{a}$, because $a \cdot \frac{1}{a} = 1$ and 1 is the multiplicative identity. There is no multiplicative inverse for 0, and therefore if we want to talk about a group structure on \mathbb{R} where usual multiplication is our "abstract group multiplication", then we need to exclude the element 0, so that $\mathbb{R} \setminus \{0\}$ is a group under multiplication.

Well, when we extend to complex numbers we still don't pick up a multiplicative inverse for 0, so any group structure on \mathbb{C} where the operation is multiplication will have the element $\{0\}$ excluded. What remains to be shown, however, is that every other element in \mathbb{C} has a multiplicative inverse. It turns out the notion of a modulus is important for this, as is the following proposition.

Proposition 4.6. Let $z \in \mathbb{C}$ and suppose $z \neq 0$. Then $|z|^2 > 0$.

Proof: $z = a + ib \neq 0$ means that either $a \neq 0$, $b \neq 0$, or both. Thus, there are three cases to consider. The first case is when $a \neq 0$ and $b = 0$. We know that $|z|^2 = a^2 + b^2 = a^2$, where the last equality comes from the fact that $b = 0$. But, since $a \neq 0$, we have that $a^2 \neq 0$. Moreover, any real number squared is non-negative. Thus we've shown that $|z|^2 = a^2$ is neither negative, nor zero, so it must be positive. This proves the proposition for this case.

The next case ($a = 0, b \neq 0$) is the exact same, since then $|z|^2 = b^2 > 0$. This is because in this case $|z|^2 = ib \times (-ib) = i \times (-i) \times b^2 = b^2$ since $i \times (-i) = -1 \times i^2 = -1 \times -1 = 1$. Finally, the third case where both $a \neq 0$ and $b \neq 0$ is even easier, since then $|z|^2 = a^2 + b^2$ where both $a^2 > 0$ and $b^2 > 0$, and we know that the sum of positive numbers is again a positive number. This completes the proof. \square

(Note: if you found the above proof extremely trivial, fine. It is primarily here to demonstrate the technique of breaking up proofs into cases that need to be enumerated).

We now make the following proposition, which shows that every non-zero complex number has a multiplicative inverse.

Proposition 4.7. Let $z \in \mathbb{C}$ and $z \neq 0$. Then the multiplicative inverse of z is $\frac{\bar{z}}{|z|^2}$. Namely, $z \cdot \frac{\bar{z}}{|z|^2} = 1$.

Proof: The proof of this proposition is also trivial. The main point to make is that it is well-defined because we know that $z \neq 0 \Rightarrow |z|^2 \neq 0$, so that we're allowed to put $|z|^2$ in the denominator in the statement of the proposition. For once we know this is well-defined, we simply have that

$$z \cdot \frac{\bar{z}}{|z|^2} = \frac{(z \cdot \bar{z})}{|z|^2} = \frac{|z|^2}{|z|^2} = 1.$$

□

What the above proposition tells us is that the set $\mathbb{C} \setminus 0$ is a group under multiplication, with multiplicative identity 1.

We now have all the tools that we need to talk about one of the most important mathematical structures in all of physics, and that is the complex vector space. We'll give a very brief account of this structure, since every single statement and result follows directly from our discussion of real vector spaces simply by replacing \mathbb{R} by \mathbb{C} everywhere. Let us begin.

4.5 Complex Vector Spaces

The first thing we do is recall the definition of a **real** vector space, which should be familiar to us by now.

Recall: A **real vector space** V is an Abelian group $(V, +)$ equipped with a special function $\cdot : \mathbb{R} \times V \rightarrow V$ called "scalar multiplication" which, $\forall a, b \in \mathbb{R}$ and $\forall v, w \in V$, satisfies the following requirements:

- 1) $a \cdot (v + w) = a \cdot v + a \cdot w$ (distributivity of scalar multiplication over vector addition)
- 2) $(a + b) \cdot v = a \cdot v + b \cdot v$ (distributivity of scalar addition over scalar multiplication)
- 3) $(ab) \cdot v = a \cdot (bv)$ (compatibility of scalar multiplication with regular multiplication).
- 4) $0 \cdot v = 0$ (compatibility of $0 \in \mathbb{R}$ with $0 \in V$)
- 5) $1 \cdot v = v$ (just an extra condition that makes vector spaces much nicer).

Now we ask the simple question: Is there anything stopping us from replacing \mathbb{R} with \mathbb{C} everywhere in that definition? All we needed to know about real numbers in this definition is that we can add them, multiply them, that there's an additive identity, and that there's a multiplicative identity. Everything else was abstract—namely, we **defined** how to multiply a real number by an element in the set V , and then enforced certain conditions on it. Since complex numbers have an additive identity and a multiplicative identity just like the real numbers do, we can simply transcribe this definition word for word, but replacing \mathbb{R} with \mathbb{C} everywhere. Let's do this, and call the resulting object a complex vector space.

Definition 4.8. A **complex vector space** V is an Abelian group $(V, +)$ equipped with a special function $\cdot : \mathbb{C} \times V \rightarrow V$ called "scalar multiplication" which, $\forall a, b \in \mathbb{C}$ and $\forall v, w \in V$, satisfies the following requirements:

- 1) $a \cdot (v + w) = a \cdot v + a \cdot w$ (distributivity of scalar multiplication over vector addition)
- 2) $(a + b) \cdot v = a \cdot v + b \cdot v$ (distributivity of scalar addition over scalar multiplication)
- 3) $(ab) \cdot v = a \cdot (bv)$ (compatibility of scalar multiplication with regular multiplication).
- 4) $0 \cdot v = 0$ (compatibility of $0 \in \mathbb{C}$ with $0 \in V$)
- 5) $1 \cdot v = v$ (just an extra condition that makes vector spaces much nicer).

(In fact, there's a much more general definition that we can make that includes both real and complex vector spaces, and this goes vaguely as follows. We note that all we really need from our scalars (in these cases either \mathbb{R} or \mathbb{C}) is that they have two associative and commutative operations—addition and multiplication—defined on them, that they have additive and multiplicative identities and inverses, and that addition and multiplication are properly distributive. Such an object has a more general, abstract existence, and is called a **field**. These objects can be studied in their own right, though we haven't done so in these notes. If we had studied these, we'd be able to define a vector space over **any** field at all. Thus, the fact that we can freely change our definition of a real vector space to that of a complex vector space is not coincidental, because both \mathbb{R} and \mathbb{C} are fields.)

We will skip over **lots** of the general theory about complex vector spaces that we developed in the case of real vector spaces, simply because literally **every** result that we obtained for real vector spaces holds for complex vector spaces when every \mathbb{R} is replaced by a \mathbb{C} . This can (and should) be checked, but we simply won't do so here because it's tedious, boring, and in no way enlightening. The key to doing this checking is to notice that there is no property of \mathbb{R} that we used that isn't shared by \mathbb{C} . This is, in fact, a very important process in mathematics. We make one definition—perhaps that which seems the most obvious—and then we ask how we can generalize it. In attempting to generalize it, we look at the properties of the mathematical structures in question and see which of them we actually used. In this case, a complex vector space is not a **generalization** of a real vector space in the strict sense, but rather a **different** definition. The important point is that we can make the definition of a complex vector space by simply noting that we used nothing in the definition of a real vector space that doesn't carry over to the complex case.

In particular, the definition of linear combinations, linear (in)dependence, bases, dimensionality, linear maps, linear isomorphisms, and matrices all carry over immediately to the complex case. We therefore also have that any two finite dimensional complex vector spaces of the same dimension are linearly isomorphic to each other. Thus, after doing the next exercise, we'll see that every finite dimensional vector space is linearly isomorphic to $\mathbb{C}^n = \mathbb{C} \times \mathbb{C} \times \dots \times \mathbb{C}$ for some n , where the right hand side of this equality has n products of \mathbb{C} .

Exercise 4.9. Show that for any integer $n > 0$, $\mathbb{C}^n = \mathbb{C} \times \mathbb{C} \times \dots \times \mathbb{C}$ is a complex vector space of dimension n , where scalar multiplication and vector addition are defined in the obvious way.

One of the important distinctions to make between real and complex vector spaces is that the dimension of a complex vector space is its **complex** dimension, which is twice that of its real dimension. Namely, \mathbb{C}^n can be viewed not only as a complex vector space of **complex** dimension n , but it can also be viewed as a real vector space of **real** dimension $2n$. Let us see how this works in an example.

Example 4.10. Consider $\mathbb{C}^2 = \mathbb{C} \times \mathbb{C}$. This is a complex vector space of complex dimension 2, since it has a complex basis $\{(1, 0), (0, 1)\}$. This is because any $v \in \mathbb{C}^2$ can be written (a, b) with $a, b \in \mathbb{C}$, and can therefore be written as the complex linear combination $v = a(1, 0) + b(0, 1)$.

However, we can also view \mathbb{C}^2 as a real vector space of **real** dimension 4. To see this, we recall that we can write any $a \in \mathbb{C}$ as $a = a_1 + ia_2$ with $a_1, a_2 \in \mathbb{R}$. We then have that any $v \in \mathbb{C}^2$ can be written as $v = (a_1 + ia_2, b_1 + ib_2) = a_1(1, 0) + a_2(i, 0) + b_1(0, 1) + b_2(0, i)$, with $a_1, a_2, b_1, b_2 \in \mathbb{R}$. Moreover, the set of vectors $\{(1, 0), (i, 0), (0, 1), (0, i)\}$ are all linearly independent **over the real numbers**, where by this we mean that any linear combination of these four vectors with **real number** coefficients that equals zero must be the combination where all four coefficients are zero. This can (and should) be checked. Thus, this set of four vectors forms a basis for \mathbb{C}^2 as a **real** vector space, and thus \mathbb{C}^2 has **real** dimension 4.

The above example can be easily extended to the case of general \mathbb{C}^n , and not just $n = 2$. Thus, we have that any finite dimensional complex vector space of **complex** dimension n is linearly isomorphic to a real vector space of **real** dimension $2n$. This is because any finite dimensional complex vector space is linearly isomorphic to \mathbb{C}^n for some n , and we now know that \mathbb{C}^n can be viewed as a real vector space of dimension $2n$.

Note that the converse is most definitely not true: real vector spaces are not, in general, linearly isomorphic to complex vector spaces. The obvious counter example is that of odd dimensional real vector spaces, since any finite dimensional complex vector space can only be linearly isomorphic to real vector spaces of even dimension.

The final note that we'll make about complex vector spaces is that just as a linear map from \mathbb{R}^n to \mathbb{R}^m can be completely described by an $m \times n$ matrix of real numbers, so too can a linear map from \mathbb{C}^n to \mathbb{C}^m by an $m \times n$ matrix of complex numbers. In other words, instead of having a matrix with entries $\{a_{ij}\}$

where each $a_{ij} \in \mathbb{R}$, we now have a matrix with entries $\{a_{ij}\}$ where each $a_{ij} \in \mathbb{C}$. Additionally, the product of a matrix and a vector is exactly the same. Namely, if $v = (c_1, \dots, c_n) \in \mathbb{C}^n$ (so that each $c_i \in \mathbb{C}$), and if A is a matrix with entries $\{a_{ij}\}$ where each $a_{ij} \in \mathbb{C}$, then the product $A \cdot v$ is exactly the same as in the real case, where the i^{th} row of the product vector is given by

$$(A \cdot v)_i = \sum_{j=1}^n a_{ij} c_j.$$

The only difference now is that each a_{ij} and each c_i has its own real and imaginary parts, but this makes no difference in the development of these ideas. In the exact same way, and for the exact same reasons, the product of two matrices A and B (corresponding to the composition of two linear maps), where $(A)_{ij} = a_{ij} \in \mathbb{C}$ and $(B)_{ij} = b_{ij} \in \mathbb{C}$, is the exact same (where A is an $n \times m$ matrix and B is an $m \times l$ matrix):

$$(A \cdot B)_{ij} = \sum_{k=1}^m a_{ik} b_{kj}.$$

We will end our discussion of complex vector spaces here. We hope that the reader is convinced that every word after the definition of a complex vector space has been superfluous, seeing as literally every general result and expression that was derived for real vector spaces carries over to complex vector spaces by simply replacing every \mathbb{R} with \mathbb{C} . We only briefly covered how this goes so as not to leave the reader completely to herself to uncover how this translation goes, but it truly is as easy as it sounds. Only later will we see more important ways in which these two types of vector spaces differ.

4.6 Outro

We've now developed enough mathematical machinery and sophistication to start modeling the physical world. The ideas that we've developed so far are very general, and we've only barely begun to scratch the surface of what they have to offer—it would be an understatement to say that the title of every subsection discussed so far could easily warrant a life's worth of study. We've built up every structure that we've discussed so far from the ground up, and as we press on to see how the physical world around us uses these structures, we'll have to endow them with ever-more-detailed structure. For example, we'll end up endowing our real and complex vector spaces with more specific structure, and this structure will have more readily available analogues to physical quantities.

We therefore end our general discussion about these basic mathematical building blocks and shift our focus to seeing how we should specialize them to describe the physical world. As we go, we'll see how important these abstract structures are, and how studying them can give us access to ideas and phenomena that our intuition would otherwise struggle to grasp. This ends the first (and largest) part of these lectures.

PART 2

Here starts the second part of these lecture notes, where we begin to explore how to use abstract mathematics to model the physical universe in which we happen to find ourselves. As with our presentation of the mathematics that took up the first part of this text, we make no attempt to give a fully detailed or completely thorough account of the ideas that we'll explore. In part 1 of these lectures, we focused on abstraction and rigor, as opposed to providing a full exposition of the ideas presented. For the latter, there are tons of great textbooks, and we list a few in the "Further Reading" section.

In this part of the text, we'll similarly choose to focus not on a full and self-contained presentation of the ideas, but rather the thought processes behind them. In particular, we'll focus on the transition from abstract and pure mathematics to things that we can readily associate with real-world physical quantities. We'll see both how our physical intuition can motivate various mathematical definitions as well as how abstract mathematical results can tell us things about our universe that our intuition may have missed.

Our senses deceive us, and our intuition cannot be fully trusted. Thus, in seeking to understand the universe in a completely robust and trustworthy way, physics seeks to craft all of its laws as abstract mathematical frameworks. The difference between physics and pure math, though, is that a physicist assigns physical meaning to various abstract mathematical structures and lets the properties of these structures tell her something about the universe. A pure mathematician sees no such physical significance in her mathematical structure, and only sees it for its abstract existence. Both viewpoints are extremely important, and often times (especially in modern physics) both are necessary. We've spent the first part of these lectures developing the abstract perspective, and now we seek to develop the physical one (still with an eye towards abstraction, however).

This part of the lectures is split up into three chapters, and these three chapters correspond to the three different world views that physics currently offers. The first chapter—and that which came first historically—is the chapter that explores the classical world of physics. This will be the world that is most intuitively obvious, and it will serve as a nice warm-up for seeing how mathematical structures can relate to the world around us. The second chapter will explore the much less intuitive relativistic world. In this world, objects move close to the speed of light, and space and time become combined into one object and need to be viewed on equal footing. The third chapter explores the whacky world of the quantum. This world is so counter-intuitive and unbelievable that to this day, no physicist fully understands even its most basic properties. In this world we need to rely almost fully on the machinery of abstract mathematics to tell us something about the world. Accordingly, we'll focus on the abstract mathematical side of all of these worlds, almost as a warm-up for studying the quantum world. We'll see how we can take the structures that we've found in the first part of these lectures and adapt them to suit our physical needs. Let us begin...

Chapter 5

The Classical World

5.1 Introduction

We begin our journey into physics in the same setting that humankind did thousands of years ago—with the obvious. Namely, when we look around with our naked eye and observe the universe around us, what do we see? How can we describe what we see mathematically?

The important thing to keep in the back of our minds is that there is no reason why this method needs to be "right". In particular, the things that we can see, hear, feel, taste, and touch are merely the aspects of the universe with which we can directly interact, without the need of fancy technology. It's important to note, though, that there is no reason why our universe should care about what we can hear or see. Namely, we are a very particular organism with very particular sense organs, and therefore the physics that we "derive" using nothing but these sense organs (and maybe some basic technology) should not be viewed as the "be all end all" when it comes to the ultimate laws of the universe.

In fact, in the following two chapters we'll see that the physics of the very small and/or the physics of the very fast is **wildly** different from what we experience with our human senses. It may therefore seem tempting to think that these new laws of physics are somehow "wrong". But the "correctness" of the laws of physics are (and should be) determined solely by what our experiments tell us. Thus, if I have a beautiful idea of how the world works but it disagrees with experiment, then it's wrong. PERIOD. Moreover, if I have an idea that I **don't** like, but it agrees with experiment, then it is right. PERIOD. The physics that we'll "derive" in this section will all be found by "common sense" experience, and/or experiments that can be run at human scales (like that of a table top). This is a very particular subset of phenomena that the universe has to offer, and we must retain in the back of our minds the possibility that the universe behaves very differently on different scales. Indeed it does, and so this is a forewarning so as not to be so attached to the world of classical physics that we can't envision it coming crashing down.

We'll begin by exploring the seemingly simple idea of classical space and time, and afterwards we'll see how things move around within space and time. We'll use empirical knowledge and abstract it to a mathematical language reminiscent of that which we've already established. Our presentation will still be as mathematical as possible, because as physics gets harder and harder (i.e., less and less intuitive), we need to rely more and more heavily on mathematics.

5.2 Classical Space-time

Suppose I'm sitting down at the table in the kitchen writing this text on my laptop, and all of the sudden I'm overcome by an overwhelming sensation of thirst. Suppose also that I've recently undergone a knee surgery, and therefore getting up, walking to my room, grabbing the water bottle that I know is half-full, walking back to the table, and sitting back down, is extremely difficult for me. Suppose that I also have an extremely generous flatmate who is willing to help me in my current situation, and that he has offered to go to my room to grab my water bottle and bring it back for me. Suppose also that my flatmate is blind. It is therefore not possible for him to simply walk into my room and see my water bottle, and so he has to rely on very precise directions from me. Luckily (or unluckily?), he has been blind for quite a while and has therefore developed a very sophisticated awareness of space and length. I can therefore give him very detailed information about where in my room my water bottle is, and he'll be able to get it.

But how do I begin? What is the information that my generous flatmate needs to successfully retrieve my water bottle? Suppose my room is a perfect cube, that it's perfectly empty except for my water bottle and a thin rod holding my water bottle up off the ground somewhere in my room (okay, as if this example wasn't crazy enough already...just bear with me). There would then be three questions that I need to answer for my flatmate in order for him to find my water bottle on his first try. The first question is, upon arriving in my room, how many steps forward he should take. The second is how many steps left or right he should take, and the third is how high off the ground the bottle is being held. With those three pieces of data, my bottle is perfectly located. Thus, three numbers $a, b,$ and c , having some physical dimension (let's call them "meters") will perfectly specify my water bottle's location.

This weird example immediately generalizes to the universe as a whole. Namely, if we could take a "snapshot" of the entire universe, it seems as though any object's location can be specified by three numbers. In particular, suppose we took a snapshot of the whole universe and suppose I was able to move around in this "frozen instance" of the universe. Suppose I start in my flat in Oxfordshire. Then any object's location in the universe can be specified by telling me how far I have to walk a) forward or backward, b) left or right, and c) up or down (assuming that I can move around empty space without any difficulties, i.e., gravity is not a problem for me). This seems like a safe enough assumption. It works for my room (the water bottle story), it works on the scale of a city (specifying, say, a building on a particular cross-street (left/right, forward/back) as well as a particular story of that building (up/down)), and it seemingly works for the universe as a whole.

So what does this mean mathematically? Is there a space that we've seen that might do a good job of modeling the property that any point (object) can be specified by three numbers? In fact, there's a perfect space for this, and that is $\mathbb{R} \times \mathbb{R} \times \mathbb{R} = \mathbb{R}^3$. If we suppose that any snapshot of our universe is nothing but a copy of \mathbb{R}^3 , then each object in our universe will sit at some point in this space, and will therefore be specified by some set of three numbers $(a, b, c) \in \mathbb{R}^3$. These numbers must have some kind of physical dimension associated to them, like feet or inches or meters or kilometers or lightyears or millions of lightyears.

We'd also need to specify some arbitrary point as our "starting point" so that we know where to measure these locations from. For example, let's say that I modeled a snapshot of our universe as \mathbb{R}^3 and chose the point $(0, 0, 0)$ to be me sitting at my kitchen table. Suppose additionally that I chose to use feet as my unit of measurement. Finally, suppose I assigned the following meaning to these three numbers: the "first slot" corresponds to the number of feet to the right of me (so that a negative number corresponds to the left of me), the "second slot" corresponds to the number feet to the front of me (so that a negative number corresponds to the back of me), and suppose the "third slot" corresponds to the number of feet above me (so that a negative number corresponds to below me). Then I could theoretically give my flatmate

the three numbers (15, 12, 10) (where there's a 10 in the third slot and therefore my flat must have some stairs in it) and he should know exactly where my water bottle is. Now, there might be some walls and/or doors in the way, but that doesn't change the fact that my water bottle has these coordinates in the universe.

One thing that's important to note is that the three numbers (15, 12, 10) are **highly dependent** on where I put my coordinate system (namely, where I declared (0, 0, 0) to be) and what units I used. The location of my water bottle of course does not depend on my choice of origin or units, so if I chose unit of, say kilometers, I'd have very different numbers. Additionally, if I chose my origin to be the middle of the living room—as opposed to where I'm sitting in the kitchen—then the three numbers that I hand my flatmate will also be different. Finally, we must note that my flatmate must know where I put (0, 0, 0), and he must know which units I'm using. For example, if he thought I was using the units of kilometers, then he'd think that I'm sending him on a much longer journey than I truly am. Additionally, if he thought that I chose my (0, 0, 0) to be somewhere in the kitchen of my parent's house in California, and that I was still using the units of feet, then he'd think I was sending him on a trans-atlantic trip for a water bottle.

Regardless of the arbitrariness of the units that we use and the location of the origin (0, 0, 0), it is (seems?) obvious that the physical world we see around us labels its spatial locations with three numbers $(a, b, c) \in \mathbb{R}^3$. Therefore, we'll simply assume this to be the case (for now).

Now we'd like to ask what happens when we decide to view our universe "in time", i.e., when we loosen the requirement that we're looking at a particular "snapshot" of the universe. A little bit of thought makes us realize that our discussion above—which led us to the conclusion that we should model a spatial snapshot of the universe as \mathbb{R}^3 —was completely independent of **which** particular snapshot we took. In other words, we never specified exactly **when** we'd take our snapshot of the universe, yet we still came to the conclusion that our snapshot should be modeled by \mathbb{R}^3 . This must mean that **every** snapshot of the universe can be modeled the same way, namely by \mathbb{R}^3 . We therefore conclude that each and every snapshot of the universe can be modeled by \mathbb{R}^3 .

We now appeal to one last bit of intuition, and that is that time "flows". Of course, as we go about our day-to-day business we're constantly observing this fact about the universe—namely, that time goes forward. We'll therefore assume that time flows in a succession of "snapshots", one right after another. In fact, we'll further assume that this flowing is continuous. As we've discussed before, a rigorous notion of "continuous" is beyond the scope of the mathematics that we've developed so far (though it's possible). We'll therefore take "continuous" to simply be synonymous with \mathbb{R} in this case. Namely, we'll assume that as time flows, we move along \mathbb{R} (as opposed to, say, \mathbb{Z} , which would model time as moving along in discrete steps).

We therefore have the following picture of classical space-time. At each instant of time, the universe looks like \mathbb{R}^3 , and there is an \mathbb{R} 's worth of "instants of time". We can then recall that we have a very natural description of such a space, and that is \mathbb{R}^4 . This is because we can view our construction as $\mathbb{R} \times (\mathbb{R} \times \mathbb{R} \times \mathbb{R})$, where the left-most \mathbb{R} represents time and the remaining three \mathbb{R} 's represent the copy of physical space associated to each instant of time $t \in \mathbb{R}$.

There are a few points to be made about these assumptions. The first is the "continuous" nature of this space. Suppose we focus on the \mathbb{R}^3 of space at a given time. Modeling space this way means that we can "zoom in" infinitely far and space will always "look" the same. In particular, if we choose our spatial units to be feet, then the nature of space between the points (0, 0, 1) and (0, 0, 2) (which are a distance of 1 foot from each other) is qualitatively no different than the space between the points (0, 0, 0.0000000000000001) and (0, 0, 0.0000000000000002) (which are much closer). Moreover, we can zoom in **literally** forever and

never come across any qualitative difference in the structure of space. We could similarly do this with time, in that between any two moments of time there are infinitely many more moments of time, and in between any two of those moments there are infinitely many more, and so on, forever, and that at any stage of this game our moments of time are all qualitatively exactly the same.

This may seem all well and good, since on the scales that we're used to this is true. Namely, space does indeed look similar whether we're considering a cubic inch of it, a cubic foot of it, or a cubic kilometer of it. Time intervals also "look the same" on our everyday time scales, in that the passing of a second seems quite similar to the passing of an hour, or a week, or a millisecond (the only difference being that some of these are shorter in duration than others). Therefore our every day experience tells us that describing our universe as \mathbb{R}^4 with one \mathbb{R} associated with time and the other three \mathbb{R} 's associated with space is obvious and necessary. However, it's fundamentally wrong. Space and time are not "infinitely zoom-able" as we've described it here, and they're also not so totally separated from each other (i.e., one \mathbb{R} for time and the other three for space). So even though this continuous picture of space and time is a **remarkably** good approximation of reality, it is not the full picture.

We will see reasons for the incompleteness of our description later, but for now we should just note that even though our discussion so far has been **intuitively** obvious, it is not **actually** true on all scales. And we note this just to point out that the job of physics is to find the mathematics that describes the physical universe that our **experiments** uncover, and not the physical universe that our **intuitions** are familiar with, simply because our intuitions are flawed (or at least incomplete). Our intuitions have been crafted to help us do things like hunt and survive and reproduce, and **not necessarily** to uncover the mathematical structure of our universe.

With that said, we can indeed use our intuition a great deal, and often times our physical and mathematical intuitions can help us uncover structure about the universe (even if this structure is counter intuitive!). Therefore, we'll continue to use this intuitive description of our universe, primarily because on the scale of physics that is relevant for our intuitions (namely, of things that don't move too fast, are pretty big (but not too big), and hot (but not too hot)) this description is wildly successful. Therefore, let us press on.

5.3 The (Very) Basic Geometry of Classical Space-time

Let us go on to discuss some very basic properties of space. In doing so, we'll see that we'll have a perfect arena for applying some of the abstract mathematics that we developed in part 1 of these notes. In some cases, we'll need to add some more structure to these mathematical structures, and we'll find that our physical intuition gives us great motivation for doing so.

The first thing we'll study is based off the very obvious realization that any two objects in space have a certain distance between them. Namely, if I choose some units (let's say meters) and some origin (doesn't matter where), and if I have two objects in my universe, where one is located at (a_1, a_2, a_3) and the other is located at (b_1, b_2, b_3) , then I know that there is a number that tells me how far apart these objects are. This distance will be the length of the straight line connecting the two objects. Moreover, we know that this distance will be a number that is greater than or equal to zero, simply because the closest any two objects can get to each other is "being in the exact same location", and then their distance from each other is zero. Additionally, there is no such thing (in this case) as "negative distance".

Since all we know about these two objects are their locations in space, namely (a_1, a_2, a_3) and (b_1, b_2, b_3) , and since we know that the sheer existence of these objects in space defines a number that we call "the

distance between these objects", it would be nice if these two sets of coordinates (the a_i 's and the b_i 's) alone could specify this number that we call distance. It turns out that these coordinates are indeed enough to specify this distance. Let's see how this works.

First let's suppose that my two objects are both on the same flat table, so that if I view my "up-down" coordinate as the third coordinate of the (x_1, x_2, x_3) , then I know that $a_3 = b_3$. Moreover, I could also make a smart choice of origin and make it so that the origin $(0, 0, 0)$ also sits on the table, so that $a_3 = b_3 = 0$. Then the coordinates of our two objects are, respectively $(a_1, a_2, 0)$ and $(b_1, b_2, 0)$, and therefore we can view these objects as simply lying in a plane together. Let's suppose that the distance separating them in the "left-right" direction is x and the distance separating them in the "forward-backward" directions is y (where we know that the distance separating them in the "up-down" direction is zero), then Pythagorean's theorem tells us that the total distance d separating them is $d = +\sqrt{x^2 + y^2}$, or equivalently that $d^2 = x^2 + y^2$.

But what is the distance separating them in "left-right" and "forward-backward" directions? This is nothing but $|a_1 - b_1|$ and $|a_2 - b_2|$, respectively, where the expression $|a_1 - b_1|$ is called "the absolute value" of $a_1 - b_1$, and this simply means "always take the non-negative part of the answer". This should be relatively straightforward to see, since a_1 says how far object 1 is from the origin in the left-right direction and b_1 says how far object 2 is from the origin in the left-right direction. Thus, the **difference** of these two numbers is precisely their distance **from each other** in this direction. Thus, when we need to square the distance in the "left-right" direction in order to get the total distance, the quantity that we need to square is $|a_1 - b_1|$. However, any real number squared is non-negative, even if the real number we started with is negative, and so we can simply plug in the expression $a_1 - b_1$ into Pythagorean's theorem. In other words, we have that $(a_1 - b_1)^2 = (b_1 - a_1)^2$, and this is closely related to the fact that $|a_1 - b_1| = |b_1 - a_1|$.

Thus, when our two objects are sitting on the same flat table, we know that their spatial distance (squared) from each other is

$$d^2 = (a_1 - b_1)^2 + (a_2 - b_2)^2.$$

In fact, it can be easily proved that this formula generalizes to the case when the two objects are **not** on the same table, and this proof follows from applying the Pythagorean twice. We'll skip the proof of this, and if the proof is unfamiliar to the reader then it can very easily be looked up. Assuming that we're familiar with this generalization of Pythagorean's theorem, it then follows that the distance between any two objects in space, with coordinates (a_1, a_2, a_3) and (b_1, b_2, b_3) , is

$$d^2 = (a_1 - b_1)^2 + (a_2 - b_2)^2 + (a_3 - b_3)^2.$$

We can check our sanity by noting how this formula aligns with our intuition in two different ways. First, we note that the distance d is always greater than or equal to zero (since we have a sum of real numbers squared, and a real number squared is always non-negative), and that $d = 0$ if and only if $a_1 = b_1$, $a_2 = b_2$, and $a_3 = b_3$, which then means that the two objects are in the exact same location, as we'd expect.

We're now in a position to ask how or if this corresponds to any of the mathematical structures we've talked about so far. In fact it does, so long as we add some extra structure to them. Let's see what this structure is.

Let's focus only on space at a given instant of time, i.e., at a given snapshot of the universe. We are therefore dealing with \mathbb{R}^3 . This is a three dimensional real vector space, but now we've found that there's even more structure on it than what we've considered before. Namely, it has the special structure that says that any two vectors have a certain "distance" between them (I use scare quotes because we haven't made "distance" mathematically rigorous yet, and we like to put non-mathematically rigorous terms in quotes).

But how do we make this notion of distance mathematically rigorous? Let's find out.

Before discussing the distance "between" two vectors, let us first notice that we can also talk about the length of a single vector. Namely, if I put the origin of my coordinates somewhere, perhaps directly in my lap, and I know an object is sitting at the location (a_1, a_2, a_3) , then I'll immediately know how far away this object is from me. This is because I can use the formula above with $b_1 = b_2 = b_3 = 0$, since I'm assuming the origin $(0, 0, 0)$ is at my location. Plugging this into the above equation, I get that the distance d that this object is away from me is

$$d = +\sqrt{a_1^2 + a_2^2 + a_3^2}.$$

But now I note that this distance only depends on the objects location. In other words, once I know the location of the origin, then the distance of any object from that origin depends only on the object's coordinates, i.e., only depends on one vector. We can call this the "length" of that vector. It might be clear that whatever we use to characterize the "length" of a vector will be closely related to the "length" **between** two vectors, so let us consider this first.

What we've found is that our vector space comes with a special function $d : V \rightarrow \mathbb{R}$ that sends every vector to the real number corresponding to its distance from the origin. Namely, we have a function d that sends each $(a, b, c) \in \mathbb{R}^3$ to the real number $(a^2 + b^2 + c^2)^{1/2} \in \mathbb{R}$. We can let this fact about the physical interpretation of the physical space \mathbb{R}^3 motivate a much more general definition by first noting what the key features of this function are. In other words, let's strip everything away and only keep the absolutely necessary features about this function. To do this, we need to find out what the key properties are that make this function correspond to a "length", and not worry so much about the explicit form of the Pythagorean theorem.

The first thing we note is that the length of any vector is non-negative, and we certainly want to retain this in any generalized notion of "length". Secondly, we note that the only vector that has zero length is the zero vector itself. We again want to maintain this in a more general definition of length, since we want it to be the case that any vector with any real "extent" has a non-zero (and positive) length. Another thing we note is that the length of a vector scales in precisely the way we'd want it to if we multiplied the vector by some real number (using scalar multiplication). Namely, if we multiplied a vector by 2, then its length doubles, and if we multiply a vector by a real number $a \in \mathbb{R}$, then its length is multiplied by $|a|$ (we use absolute values here because we don't want a negative length). Namely, the vector $(2, 2, 2)$ will be twice as long as $(1, 1, 1)$, as it should be, the vector $(-2, -2, -2)$ will also be twice as long as $(1, 1, 1)$ (as it should be), the vector $(-1, -1, -1)$ will be equally as long as $(1, 1, 1)$ (as it should be), and the vector (ab, ac, ad) will be $|a|$ times as long as (b, c, d) . To see this in the Pythagorean case, we simply note that

$$d((ab_1, ab_2, ab_3)) = d(a \cdot (b_1, b_2, b_3)) = ((ab_1)^2 + (ab_2)^2 + (ab_3)^2)^{1/2} = (a^2(b_1^2 + b_2^2 + b_3^2))^{1/2} = |a|(b_1^2 + b_2^2 + b_3^2)^{1/2}.$$

This is again something that we want to maintain in any definition of length. The final thing that we'll want to maintain, and which is slightly less obvious, is what's known as the "triangle equality". This will simply reflect the fact that the shortest distance between two points is a straight line, which is something that we all know and love. The way this is reflected in our function d is that the length of any sum of vectors is less than or equal to sum of the length of vectors. In other words, if we're interesting in adding two vectors together and then asking about the resulting vector's length, we'll always get a larger (or equal) number if we first take the lengths of the vectors individually and then add them (as opposed to adding them and then taking the resulting vector's length). By drawing the right picture, we can see that this means that the sum of the lengths of two sides of a triangle will always be greater than the length of the third side—namely, by using the head-to-tail method of drawing the addition of vectors, as discussed in

the chapter on real vector spaces. The proof of the following claim requires some results that we won't develop here, but it's very easy to convince oneself of when the right picture is drawn. Namely, the following claim simply says that when we add two vectors using the head-to-tail method (remember, this isn't a **new** method, but rather just a way of visualizing **the** method), then the length of the final vector that we draw will always be less than or equal to the length that we'd get by first measuring one of the vectors that we added up, and then adding this length to the length of the second vector that we added up. This of course isn't a proof in the strict sense, but since we won't develop the necessary tools for a strict proof, we'll just have to take the following on (good) faith and know that it **can be** proven rigorously.

Claim 5.1. Let $d : \mathbb{R}^3 \rightarrow \mathbb{R}$ be the map $d(v) = (v_1^2 + v_2^2 + v_3^2)^{1/2}$ where $v = (v_1, v_2, v_3)$. Then for any $v, w \in \mathbb{R}^3$, we have that $d(v + w) \leq d(v) + d(w)$. \square

(Boxes go directly after a claim when we either have already proved the claim before the statement of it, or if we simply won't prove the claim, which is the case here.)

This will be the final thing that we'd like to maintain when we generalize our notion of length. To recap, we now have four things that we want from our notion of length, and which seemingly characterize all the essential qualities that length should have. First, no vector can have a negative length. Second, **only** the zero vector can have zero length. Third, the length of our vectors should scale properly with scalar multiplication. Fourth, the triangle equality should hold, so that the sum of the lengths of two vectors (two sides of a triangle) should always be greater than or equal to the length of the sum of the pair of vectors. In fact, the only time these two things are equal is when the two vectors are "collinear", meaning that there is some scalar $a \in \mathbb{R}$ such that $v = aw$ (where v, w were the pair of vectors that we were considering). It turns out, though, that by only supposing the last three, the first condition can be derived. Namely, our definition will only require the last three requirements as axioms, and the first requirement can be proved from these three (and it will be left as an exercise immediately following the definition).

Definition 5.2. Let V be a vector space. Then V is a **normed vector space** if there is a function $d : V \rightarrow \mathbb{R}$ such that the following conditions hold for all $v, w \in V$ and for all $a \in \mathbb{R}$:

- 1) $d(av) = |a|d(v)$
- 2) $d(0) = 0$
- 3) $d(v + w) \leq d(v) + d(w)$.

The function d is then called the **norm** of V , and for any $v \in V$, the number $d(v)$ is called the **norm** of v .

Exercise 5.3. Let (V, d) be a normed vector space. Using only the axioms for a normed vector space as given above, show that for all $v \in V$ we have $d(v) \geq 0$.

This definition allows us to generalize our discussion to many more vector spaces and many more norms. For example, the function $g : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by $g((v_1, v_2, v_3)) = (v_1^4 + v_2^4 + v_3^4)^{1/4}$ is also a norm on \mathbb{R}^3 . As we've seen, it isn't the norm of physical space, but it is a norm in the mathematical sense, and may have its own interesting existence in the world of mathematics. Additionally, we could also extend our norm to higher dimensional vector spaces. For example, we could consider the vector space \mathbb{R}^N with the norm $d : \mathbb{R}^N \rightarrow \mathbb{R}$ defined by $d((v_1, v_2, \dots, v_N)) = (v_1^2 + v_2^2 + \dots + v_N^2)^{1/2}$. This is nothing but our 3-dimensional norm extended to higher dimensions, so by making this generalization we have gained the ability to describe the distance between objects in universes of higher dimension!

To summarize, we've found that classical space-time is described by $\mathbb{R}^4 = \mathbb{R} \times \mathbb{R}^3$, and that at each instant in time (the first factor of \mathbb{R}) we have a copy of \mathbb{R}^3 , which is a normed vector space with norm $d : \mathbb{R}^3 \rightarrow \mathbb{R}$ defined by $d((a_1, a_2, a_3)) = (a_1^2 + a_2^2 + a_3^2)^{1/2}$. Then, to describe the distance **between** objects, we simply do the following. Suppose object 1 is located at $v = (v_1, v_2, v_3)$ in space, and object 2 is located at $w = (w_1, w_2, w_3)$ in space. Then, the distance between them is simply the length of the vector $v - w$. This is

because $d(v-w) = ((v_1-w_1)^2 + (v_2-w_2)^2 + (v_3-w_3)^2)^{1/2}$, which is precisely the expression for the distance between two objects that we found above! Notice also that $d(v-w) = d(w-v)$, just as we'd expect (namely, it doesn't matter from which object we start measuring at—the distance between the two will be the same).

Now it turns out that there's one last generalization that we can make, and that is to simply drop the vector space structure from the definition of a normed vector space. When we drop this structure, we lose our ability to add elements in V to each other, because now V is just an ordinary set with no "addition" (or scalar multiplication) operation. Still, though, we can simply **force** our construction to give us a notion of "length" between two points in an arbitrary set. To do this, we simply require 1) positivity of our lengths, 2) the distance between an element and itself is zero, and the only time an element is "zero distance" from another element is when these elements are actually the same, 3) the distance from an element a to an element b is the same as the distance from the element b to the element a , and 4) the triangle equality. Note that these requirements are all motivated by what we did above in the case of a normed vector space, only now we can make them slightly more general.

Definition 5.4. Let S be a set. Then S is a **metric space** if it has a function $d : S \times S \rightarrow \mathbb{R}$ such that the following conditions hold for all $x, y, z \in S$:

- 1) $d(x, y) \geq 0$ (positivity)
- 2) $d(x, y) = 0 \Leftrightarrow x = y$ (only 0 when the same)
- 3) $d(x, y) = d(y, x)$ (distance from one to the other the same as from the other to the one)
- 4) $d(x, z) \leq d(x, y) + d(y, z)$ (triangle inequality).

When these conditions are satisfied, d is called the **metric** for S .

Thus, a metric is simply an abstract way of defining "how far away" two elements in a set are, and it may or may not have anything to do with any kind of physical space. The following example shows this.

Example 5.5. This example may seem a little strange and non-physical, primarily because it is. But it will hopefully show just how general the definition of a metric space is. Let $S = \{a, b, c, d\}$ and let D be a metric on S (where we use a capital "D" since we're already using the lower-case "d" in the set S) defined by the following relations:

$$D(a, b) = D(a, c) = D(a, d) = D(b, c) = D(b, d) = D(c, d) = 1.$$

In other words, every element in S is 1 unit of distance from every other element. We note that we don't have to define D for (a, a) , for example, because we're forcing D to be a metric so we know that $D(a, a) = 0$. Moreover, we also don't need to define D on (b, a) because if it's defined for (a, b) , then the symmetry of a metric determines it for (b, a) . Namely, we force $D(a, b) = D(b, a) (= 1)$.

To see that D actually is a metric is relatively simple. We first note that we've satisfied the first requirement, since $D(x, y) = 1 \geq 0$ for all $x, y \in S$. We've satisfied the second requirement automatically, by simply forcing it to be true, and we've also forced the third requirement to be true as well (i.e., by **defining** D to be such that the second and third axioms of the definition of a metric are satisfied). Therefore, the last thing we need to check is that the triangle equality holds, and this is pretty easy as well.

We first take two arbitrary elements $x, z \in S$. There are two cases: either $x = z$ or $x \neq z$. The first case is trivial, because in this case the left hand side of the fourth axiom of the definition of a metric is 0, and we know that both terms on the right hand side are ≥ 0 , and therefore the inequality is satisfied. Now we need to consider the case where $x \neq z$. Then we have $D(x, z) = 1$. Now let $y \in S$ be arbitrary. If $x = y$, then $y \neq z$ (since $y = x \neq z$), and so we know that $D(y, z) = 1$. Thus $D(x, y) + D(y, z) = 1$, and since 1 is indeed greater than or equal to 1 (namely, equal to), we have that $D(x, z) \leq D(x, y) + D(y, z)$. The same logic applies when $y = z$, because then we know that $y \neq x$, and we repeat this argument. Finally, if $y \neq x$ and $y \neq z$, then $d(x, y) + D(y, z) = 2 > 1$, so the triangle inequality is satisfied in this case as well. Therefore the triangle inequality is always satisfied, so D is indeed a metric.

In fact, S can have as many elements as we want, since nothing we've done has been dependent on the number of elements in S . Thus, so long as we define D to assign the value 1 to any pair (a, b) such that $a \neq b$, and so long as we define $D(a, a) = 0$ for all $a \in S$ and $D(a, b) = D(b, a)$, we'll automatically turn S into a metric space. Now, this metric space might not be very interesting (it isn't), but it does show how general this definition is.

The following proposition shows that metric spaces are strictly more general than normed vector spaces, in that any normed vector space can be turned into a metric space, but not vice versa.

Proposition 5.6. Let V be a normed vector space with norm $d : V \rightarrow \mathbb{R}$. Then V has a natural structure as a metric space.

Proof: Define a metric $D : V \times V \rightarrow \mathbb{R}$ on V as follows: $D(v, w) = d(v - w)$ where d is the norm on V . Note that this is well-defined, in that D does indeed take two vectors (one from each factor of $V \times V$) and spits out a real number, namely, the real number $d(v - w)$, since d takes vectors into real numbers. We now need to show that D satisfies all of the requirements of a metric. The positivity (or really the non-negativity) of D follows immediately from the positivity (non-negativity) of d , and the fact that $D(v, w) = 0 \Leftrightarrow v = w$ follows from the fact that $d(x) = 0 \Leftrightarrow x = 0$, so that $D(v, w) = d(v - w) = 0 \Leftrightarrow v - w = 0 \Leftrightarrow v = w$, thus proving the second requirement of a metric space. Now, $D(v, w) = d(v - w) = d((-1) \cdot (w - v)) = |-1|d(w - v) = 1d(w - v) = d(w - v) = D(w, v)$, where we've used the fact that norms are consistent with scalar multiplication in the way we defined. Thus, D also satisfies the third requirement of a metric space. Finally, we have the triangle inequality. After choosing arbitrary $x, y, z \in V$, we use the following chain of (in)equalities:

$$D(x, y) = d(x - y) = d(x - z + z - y) = d((x - z) + (z - y)) \leq d(x - z) + d(z - y) = D(x, z) + D(z, y),$$

where in the second equality we simply added zero in the form of $z - z$, and where the inequality comes from the triangle inequality that we know is satisfied by d . This completes our proof. \square

We've now shown that every normed vector space can equally well be viewed as a metric space, but that there are metric spaces that cannot be viewed as vector spaces. Namely, in the example preceding the previous proposition, we saw a metric space that has no vector space structure.

To recap, we let the physical fact that any two objects in space have a distance between them motivate an abstract mathematical definition. This abstract definition generalized our notion of distance and length far beyond the original functions that motivated them. In this case, the generalizations that we made did not lead to any new physics, but this is not always the case. In fact, there have been (and still are) many cases in physics where a physical idea motivates an abstract mathematical definition, and that reapplying this added abstraction to motivate new physical questions and/or solutions leads to new physical insight. We won't see too many examples of this here, but we will (shortly) see how some physical ideas are **only** accessible via abstract mathematics. For now, let's explore a bit more classical physics.

5.4 *Group Properties of Classical Space

Suppose I'm holding a box—a regular cardboard cube. Suppose I measure the distance between, say, two opposite corners. Then suppose I rotate this box around in space, and I do so very carefully, making sure I don't crush the box or bend it in any way. Maybe I throw the box up in the air (either the box is empty, or I'm extremely strong, or both), I let it rotate around, and I very gently catch it so that the cardboard doesn't bend or get crushed in any way. Now suppose I remeasure the distance between the same two corners. What will my result be?

Surely the answer to this question is so obvious that the reader may be surprised that I'm even asking it. Of course the distance between these two corners will be the same. In fact, the distance between any two points on the box will be the same before and after said flipping/rotating. This is so obvious that there can't possibly be any deep mathematical statement in regards to this, right? Wrong.

Before we can see this, though, let us note the following more general statement. Suppose I "lay down my coordinates" for the universe, meaning I choose an origin (let's say, in my lap) for \mathbb{R}^3 and I choose some units of measurement (only considering one snapshot for now). Now suppose I measure the distance between two points in space. Now suppose I'm sitting in one of those fun rotating chairs. Now suppose I rotate around in my chair, and I "bring my coordinates with me". By this I mean the origin stays where it is (in my lap), but I rotate around the rest of the axes. Clearly the coordinates of the two points that I measured the distance between will change, because I'm changing the coordinates I'm using to describe them, but it's also clear that the distance between them won't change (just as the distances between points on the box don't change).

Now, consider the set of things that I can do to my coordinates that will a) keep my origin sitting in my lap (i.e., not move the origin), and b) not change the distance between any two points in space. Clearly simply rotating my axes the way that I've described (swiveling around in my chair) will satisfy these requirements. But how about the transformation $(a_1, a_2, a_3) \rightarrow (ba_1, ba_2, ba_3)$, where I "rescale" my coordinates by some real number b . Then, since the two points (a_1, a_2, a_3) and (c_1, c_2, c_3) were initially (in my original coordinate system) $((a_1 - c_1)^2 + (a_2 - c_2)^2 + (a_3 - c_3)^2)^{1/2}$ units away from each other, in our new rescaled coordinates they're now described by (ba_1, ba_2, ba_3) and (bc_1, bc_2, bc_3) , and their distance from each other will now be $|b|((a_1 - c_1)^2 + (a_2 - c_2)^2 + (a_3 - c_3)^2)^{1/2}$. Thus, unless $b = \pm 1$, the transformation $(a_1, a_2, a_3) \rightarrow (ba_1, ba_2, ba_3)$ does **not** satisfy the requirements.

Let us therefore focus only on the pure rotations of space. These are the transformations that simply move all of space together while keeping the origin fixed and without "rescaling" at all. This actually forms a group, as is relatively straightforward to see. Namely, the rotation "don't rotate at all" is the identity element, the multiplication of rotation 1 and rotation 2 is simply "first do rotation 1 then do rotation 2", and the inverse of any rotation is simply "rotate the exact same amount but in the exact opposite direction". Thus, the set of rotations of space that keep the distance between any two points fixed is a group. In fact, this group is extremely important in physics, though we will not be able to see how in this text. Moreover, the group theory of rotations of space has a very nice expression in terms of matrices acting on vectors, but we won't explore this yet and instead we'll leave it for a future course for the reader. For now, we simply note this group structure and notice how groups really show up all over the place, and often times have nothing to do with adding or multiplying numbers!

5.5 *Movement Through Classical Space-Time

To end our (incredibly incomplete) discussion of classical physics, we'll talk briefly about how objects move through space in the classical picture. After all, physics is ultimately trying to describe how things move and interact. This is the ultimate, abstract goal of physics, and really all of science—if I know some things about the universe now, how much can I know about it in the future? In classical physics, the answer to this question is "everything". Namely, if I know everything about the universe now, I can theoretically know everything about it at any time in the future. Let us see a bit of how all of this works.

Classical physics (and all of physics, for that matter) works with two main ideas: objects and forces acting on objects. Objects make up the universe, and forces act on objects to move them around in various ways (okay, in this sense, force "makes up the universe" as well, but hopefully you catch my drift). Before

we can describe how forces act on objects, though, we first need to establish how objects might even be able to be acted upon in the first place.

To do this, we need to understand the relationship between position, velocity, and acceleration. Position is something we're already familiar with. At any fixed time $t \in \mathbb{R}$, any object in the universe has some location $(a_1, a_2, a_3) \in \mathbb{R}^3$. So that covers the position topic. Now, the velocity of an object is that quantity that tells us how the position changes with time. Namely, $100\text{km}/\text{hour}$ is a velocity, and it tells us that the position changes by 100km for every hour of time that passes. Let us put this mathematically. Suppose we have a particle that is constrained to move along a 1-dimensional space, so that it can only move left or right, say. Suppose also that at some time $t = t_0$ our particle is at position x_0 (on the number line \mathbb{R}). Finally, suppose that our particle is moving with some velocity v , which can be either positive or negative. By this, I mean that the velocity is negative if the particle is moving to the left (for example) and positive if it's moving to the right. Then, at some later time $t = t_1 > t_0$, our particle is at the position $x_1 = x_0 + v(t_1 - t_0)$. This is because our particle has been moving for $t_1 - t_0$ seconds (if we've chosen seconds as our units of time), and so $v(t_1 - t_0)$ is its total change in position.

Now this can easily be generalized to the non-one-dimensional case, by simply letting all of our physical quantities (except time) be vectors in \mathbb{R}^3 instead of simply numbers in \mathbb{R} . In particular, we already know that position is a vector in \mathbb{R}^3 , and we can similarly let our velocity v be a vector $\in \mathbb{R}^3$ as follows. When we write $v = (v_1, v_2, v_3) \in \mathbb{R}^3$, and when we interpret this as a velocity, we then mean that the first coordinate v_1 is the velocity—or change in position per time—of the first coordinate for the position vector $x = (x_1, x_2, x_3)$. Similarly, v_2 is the velocity of the x_2 component, and v_3 is the velocity of the x_3 component. Thus, if we write the position vector at time $t = t_0$ as $x = (x_1, x_2, x_3)$ and the position of the particle at time $t = t_1$ as $x' = (x'_1, x'_2, x'_3)$, then we have the more general 3-dimensional version of the above expression: $x' = x + (t_1 - t_0)v$. Or, to make explicit the 3-dimensionality of this equation, we have

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} + (t_1 - t_0) \cdot \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} = \begin{pmatrix} x_1 + (t_1 - t_0)v_1 \\ x_2 + (t_1 - t_0)v_2 \\ x_3 + (t_1 - t_0)v_3 \end{pmatrix}$$

We therefore see that this is nothing but the one-dimensional expression repeated in each of the three components.

We're now in a position to start talking about acceleration, which is the final component of kinematics in classical physics. Just as velocity is the change in position with time, acceleration is the change in velocity with time. We can think of this as follows. Suppose you're driving a car down a road that is completely frictionless. This means that the friction of the tires on the road won't slow you down at all, so that if you're cruising at some speed v_0 , you won't slow down or speed up at all. Acceleration occurs when you put your foot on the pedal. This changes your velocity. Now you might wonder why we have to keep our foot on the pedal when driving down the road, even though it seems like we're maintaining a (relatively) constant velocity. This is because we constantly have the friction of the tires on the road and the air of our atmosphere pushing back against us, so we have to keep re-accelerating in order to keep an almost constant velocity. Thus, what might **appear** as a constant velocity is really the constant battle between two opposing accelerations—one acceleration (friction) is constantly changing our velocity for the negative, and the gas pedal is constantly changing it for the positive.

We thus find a completely analogous expression between velocity and acceleration that we did between position and velocity. Namely, if we're currently (at time $t = t_0$) traveling with velocity v_0 and we accelerate with constant acceleration until time $t = t_1$, then our final velocity v_1 will be $v_0 + (t_1 - t_0)a$ (where a can be positive or negative depending on whether or not we're speeding up or slowing down). Our velocity

has therefore changed to "our initial velocity plus our acceleration times the amount of time we accelerated for". We can similarly view all of these as vector equations, since we're free to "speed up" or "slow down" in any of our 3 directions of space, and so we get the completely analogous expressions

$$\begin{pmatrix} v'_1 \\ v'_2 \\ v'_3 \end{pmatrix} = \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} + (t_1 - t_0) \cdot \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix} = \begin{pmatrix} v_1 + (t_1 - t_0)a_1 \\ v_2 + (t_1 - t_0)a_2 \\ v_3 + (t_1 - t_0)a_3 \end{pmatrix}.$$

Now it would be nice if we could have all of this in terms of position, since "the position of our particle" is what we're usually interested in. Namely, if I know where my particle initially is, how fast it's initially moving, and what my constant acceleration is, can I find an expression for the final position of my particle? The answer is yes, and to derive this result we need to use the following crucial result. Suppose I'm driving down the road at $10\text{km}/\text{hour}$ (pretty slow, but I'm cautious), and I accelerate at a constant rate until I reach the daredevil speed of $20\text{km}/\text{hour}$. It is hopefully clear that from the moment in time that I started accelerating to the moment I stopped accelerating, I will have traveled the same distance as I would have if I had started traveling at $15\text{km}/\text{hour}$ and simply maintained this speed for that whole duration. That is, if I constantly accelerate for some amount of time Δt (pronounced "delta t ") from the velocity v_0 to the velocity v_1 , then I will have traveled the same distance as I would have if I traveled at my **average** velocity $\bar{v} = \frac{1}{2}(v_0 + v_1)$ for the same amount of time. Thus, the distance I travel in this amount of time, which I'll call Δx , is

$$\Delta x = \bar{v}t = \frac{1}{2}(v_0 + v_1)\Delta t.$$

But I also know that $v_1 = v_0 + ta$, where a is my constant acceleration. I thus have

$$\Delta x = \bar{v}\Delta t = \frac{1}{2}(v_0 + v_1)\Delta t = \frac{1}{2}(v_0 + v_0 + \Delta ta)\Delta t = v_0\Delta t + \frac{1}{2}a(\Delta t)^2.$$

Now noting that $\Delta x = x_1 - x_0$, since "change in position" is nothing but "final position minus initial position", and that $\Delta t = t_1 - t_0$ for similar reasons, we find the final expression

$$x_1 = x_0 + v_0(t_1 - t_0) + \frac{1}{2}a(t_1 - t_0)^2.$$

This does exactly what we want it to: it takes in the data of our initial position, initial velocity, and acceleration, and gives us our final position. We can then translate this to our vector notation by noting that this equation simply gets copied into each component, since the various directions in space "don't talk to each other". Denoting our initial position by $x = (x_1, x_2, x_3)$, our initial velocity by $v = (v_1, v_2, v_3)$, our acceleration by $a = (a_1, a_2, a_3)$, and our final position by $x' = (x'_1, x'_2, x'_3)$, we have

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} + (t_1 - t_0) \cdot \begin{pmatrix} v_1 \\ v_2 \\ v_3 \end{pmatrix} + \frac{1}{2}(t_1 - t_0)^2 \cdot \begin{pmatrix} a_1 \\ a_2 \\ a_3 \end{pmatrix}.$$

As a final ingredient for actually calculating something physical, we need to use one of Newton's famous laws. While entire volumes can be (and have been) written about these laws, we'll simply and briefly state this one and say what it means for our purposes here. The law we need is that $F = ma$, where F is the force acting on a particle, m is the mass of a particle, and a is the acceleration of the particle. This can be interpreted several ways, and they're all important. First, this means that if a force with strength F is applied to a particle of mass m , then it's acceleration will be $a = \frac{F}{m}$. This can also be interpreted as telling us that if a particle of mass m is accelerating with acceleration a , then there must be a force with strength $F = ma$ acting on it. Finally, this means that if a force with strength F is applied to a particle and it accelerates with acceleration a , then its mass must be $m = \frac{F}{a}$.

We're almost ready to bring this all together to solve some actual physics problems, but first we need to make some notes about units. First, we need to see what **kinds of** units our various quantities, like position, velocity, and force, need to have, and then once that's done we need to agree on which units we'll use. We also need to note that any equation involving units needs to have **the same** units on both side of the equal sign. Namely, we can only compare distances to distances, times to times, and so on. For example, it makes no sense to say "5 grams of water=3 meters". That's just completely absurd. So, once we give units to each of our quantities, checking to make sure our units "match up" will provide a nice sanity check on our work.

So let's begin with distances, and let our units for distance be meters. Thus, if $x \in \mathbb{R}^3$ is a position, we'll say that x is $d(x)$ meters away from the origin (where d is the standard physical metric), and if $x, y \in \mathbb{R}^3$ are both positions, then we can say that they're $d(x - y)$ meters away from each other. Now, let's use seconds for our unit of time. This means that a velocity has units of "length per time", and so in our units this is "meters per second". We note that the equation $\Delta x = v\Delta t$ "works" in that its units work out—on the left we have "length" (in meters), and on the right we have "length/time" \times "time" (meters/second \times seconds), so that we're left with "length" (meters). This then means that acceleration has units of "length per time per time", which in our units is meters per second per second, or "meters/second²". This is because when we multiply acceleration by a time, we're left with a velocity (a la the equation $\Delta v = a\Delta t$). Finally, this means that a force must have the units of "mass \times distance per time per time", because it has the same units as acceleration only with an extra factor of mass multiplied to it. If we choose our units of mass to be kilograms (kg), then force has units of kilograms \times meters per second². We call this unit a Newton (since he was the guy who came up with a lot of this stuff), so that a single Newton is the amount of force that will accelerate one kilogram by one meter/second².

We're finally now in a position to do some physics.

Exercise 5.7. It is a known (observational) fact that gravity accelerates all objects equally, regardless of their mass. Near the surface of the Earth, the acceleration due to gravity is approximately $9.8m/sec^2$ (m =meters, sec =seconds), so let us assume this is the much nicer number $10m/sec^2$. What, then, is the gravitational force being applied to an object of mass $m = 5kg$ near the Earth's surface? What is the gravitational force being applied to an object of mass $m = 10kg$ near the Earth's surface? Now suppose we have a particle of mass $m = 10kg$ initially falling with speed $5m/sec$. How fast is this object falling after being in free fall for 10 more seconds? Suppose, moreover, that this object started out 1200 meters above the surface of the Earth. How far off the ground is this object after these 10 seconds? (Hint/Challenge: There is some superfluous information hidden in this problem, can you identify it?)

Exercise 5.8. Suppose an object of mass $m = 10kg$ is sitting on a massless table at rest. Suppose I apply a force of $5N$ to this object in the horizontal direction, and suppose I do so for 4 seconds. How far has this object moved in this process?

This will end our discussion of classical physics. It is incredibly under-developed and incomplete, as there are many more interesting topics that can be investigated. With finite time, however, we're forced to accept a qualitative understanding of how this stuff works. We've found a way to model classical space-time as \mathbb{R}^4 , where three of these factors of \mathbb{R} have a natural normed vector space structure at each instant of time. This motivated the definition not only of a normed vector space, but also of a metric space, which is much more general (and often not even physical). From there, we briefly mentioned how the set of rotations of space that leave distances unaffected form a group. Finally, we briefly introduced Newton's laws and how they determine how objects move through classical space-time.

Everything that we've done so far has been very intuitive. We've been able to draw pictures of almost all physical ideas that we've discussed, and the mental motion picture of what's going on is almost always clear. In the next chapter, where we introduce special relativity and the mixing of space and time, we'll lose some intuition for what's going on. Nevertheless, pictures can still be drawn and a new intuition can be slowly developed. We'll start to see, though, how mathematics can help when our intuitions are underdeveloped. In the final chapter, however, when we introduce some quantum mechanical principles, we'll have to discard all intuitive and preconceived notions of how the universe behaves. Our ability to draw pictures of what's going on will be almost completely lost, and we'll have to rely solely on the mathematical framework that we develop. Therefore, although classical physics aligns itself nicely with our intuitions, this is not a general phenomenon of the laws of Nature. This is not a problem, though, because (fortunately or unfortunately) it's not up to us to **determine** the laws of Nature, but rather to **understand** them. As we've seen and as we'll continue to see more and more, mathematics seems to be the only language through which we can do so.

Chapter 6

The Relativistic World

6.1 Introduction

Einstein's theory of relativity—both the Special and the General one—are two of the most successful and elegant scientific theories to ever exist. From only a few simple assumptions—either experimental or philosophical—and from the immense genius of Einstein (plus some others), a vast mathematical framework, modeling a wide array of phenomena and predicting a huge collection of phenomena that were previously never even hypothesized to occur, arose. As the names suggest, "special relativity" is a "special" case of "general relativity", so in reality there is only one "true" theory of relativity and that is the "general" one. Although we'll only be exploring part of special relativity, I'll say a few introductory words about both.

Special relativity is the physical theory that tells us how space and time mix to become two inseparable parts of one larger whole—space-time. With the discovery of special relativity came the realization that our previous notions of "space" being one thing and "time" being another were simply and profoundly wrong. General relativity is the physical theory that tells us how this new single object—space-time—is curved by the masses and energies that reside in it, and how objects move through this curved space-time. This physical theory of space-time curvature gives us a remarkably successful and immensely mathematically beautiful theory of gravitation, and predicts all sorts of new gravitational phenomena like black holes and gravitational waves (waves **of** space-time curvature (contrasted with waves (like ocean waves) that exist **within** space-time)).

The language of general relativity is differential geometry, which we are still very far away from seeing. However, although we haven't developed all of the tools necessary to study special relativity in any deep way, we do have the tools available to see how the basic assumptions of special relativity give rise to the **necessity** of viewing time and space on equal footing, and seeing things like time slowing down for certain observers. Let us therefore be content with this somewhat "shallow" exploration into this wondrous world, and simply move on to see some of the cool things it has to offer.

6.2 The Assumptions of Special Relativity

There are two main assumptions of special relativity. That's right, only two. One of them is both observationally motivated as well as philosophically motivated, and the other is purely observationally motivated and **highly** counterintuitive and non-obvious. Therefore, let's start with the first assumption.

For some reason, almost all special relativistic analogies take place on trains, or pairs of trains, so let us adopt this convention as well. Suppose you are sitting on a train, and that there are two other trains,

one on either side of you, next to you on neighboring pairs of tracks. Suppose these trains are all extremely long and high, so that all you can see out of your windows are the windows of the neighboring trains. Suppose you boarded early and immediately fell asleep, thus losing track of time. Suppose also that you have no wristwatch or any time-telling device at all, so you can't see if it's past your train's departure time, and finally suppose that you're riding a very high class rail line, and you know that the rails are **perfectly** smooth.

Now suppose you wake up and want to know whether or not your train has already left, i.e., whether or not it's moving. You look out the window to the left and see that you're overtaking the windows of the train on that side. You therefore can deduce that you're moving in a forward direction **relative to that train**. But can you know whether or not it's your train moving forward, or the left train moving backwards? Maybe it's a train heading in the other direction, and you're still stationary with respect to the station at which you boarded. You then look to your right and you see that that train's windows are overtaking you, so you can deduce that you're moving backward **relative to that train**. But can you deduce whether or not you're genuinely going in the opposite direction that you'd be (maybe you boarded the wrong train?) or if it's the **other** train that's simply moving faster than you?

You, being the clever reader that you are, note that if you experienced the train go from a stopped position to a moving position, then you'd **feel** that motion because it would push you back against your seat (or pull you forward slightly, if you're sitting in one of those opposite-facing seats). But you know you fell asleep for some time, so it's possible that this all occurred without you experiencing it, and now you're left to deduce which way you're moving (or if you're moving at all), based on what the other trains around you are doing. But you quickly realize that, unless you feel the sort of **acceleration** that comes with going from a stopped state to a moving one, then you have no hope of figuring out which trains are "really" moving.

The point is that if two observers are moving relative to each other with a **constant velocity**, then there is no objective definition of who is "stationary" and who is "moving". To me, the other trains are moving, but to them, I'm the one moving. And no one is wrong. Thus, if my train is cruising at the constant rate of $100\text{km}/\text{hour}$, then there's no experiment that I can run **in my train** that could determine this to be true. After all, I'm only moving with this speed **relative to the sheep in the field that we're passing**, but relative to a train moving in the same direction at $120\text{km}/\text{hour}$, I'm actually moving **backwards** at 20km per hour. This is a large part of where the name "relativity" comes from—there is no **objective** notion of constant velocity motion, only **relative** motion. It's meaningless to say that my train is **objectively** moving at $100\text{km}/\text{hour}$ —the only meaningful statements I can make is how fast my train is moving **relative** to certain other objects.

This may or may not seem like an obvious statement when put in this way, but it turns out to be a very deep statement. The take-away from this is the following: if two observers are moving with a constant velocity relative to each other, then the physics that they respectively observe is identical. Finally, note the importance of the statement **constant velocity**. We know from the train example that we **can** distinguish **objectively** between accelerated motion and non-accelerated motion. This is because accelerated motion requires a force, (remember Newton: $F = ma$), and we can **feel** a force being applied. Namely, there is an objective definition of whether or not an object is being acted on by a force, whereas there is no such definition for whether or not an object is moving with a constant velocity.

The second, and much less intuitive assumption of special relativity is that the speed of light, which we denote as c , is the same **for all observers**. The speed of light happens to be about $3 \times 10^8 = 300,000,000\text{ m}/\text{sec}$, but that doesn't really matter. What matters is that **everyone**, no matter their state of motion, sees the same speed of light (in vacuum, that is). Therefore, if I'm moving in my train relative to

your train with a constant velocity of $1,000\text{km}/\text{hour}$ (a **very** fast train), and if a sheep in the field happens to step on a flashlight, then we'll still **both** measure that flash of light to zoom by us at the same speed. That's an extremely non-obvious fact. We know it's true because hundreds upon thousands of experiments have confirmed it, without the slightest shred of possible exceptions.

In fact, Einstein was around for one of the first experiments that proved this weird fact about the speed of light, and Einstein took this result very seriously—so seriously, in fact, that he found that we either have to believe that time can move faster or slower for observers moving in different ways, or that the results of the experiment were wrong. It is one of many examples of Einstein's genius that he actually took this result for what it was and showed us how time can in fact speed up or slow down. Now, almost a century later, we've confirmed, re-confirmed, and re-re-confirmed this result with marvelous precision. Every day, in every particle collider in the world, this fact about the speed of light is confirmed without the slightest amount of doubt. We therefore need to accept its counter-intuitive consequences for what they are—the truth. Let us then begin to explore some of these weird consequences.

6.3 The Consequences of Those Assumptions

Before seeing the physical consequences of the above assumptions in mathematical form, let's first describe them qualitatively so that we can prepare ourselves for what is to come. The first assumption (about how physics is the same for observers moving at constant relative velocities) is not all that new or crazy, and although it does come into play in relativity theory, we won't focus on it too much for now. Instead, we'll focus on the second assumption (the fact that every observer views the same speed of light), which as we'll see has some seriously weird consequences.

First, though, we should note that these two assumptions are very closely related. For suppose it **wasn't** the case that observers moving with constant relative velocity recorded the same speed of light. Then it **would** be the case that these two observers could determine who was **objectively** moving, which our first assumption says is impossible. For example, we could define the speed of light to be one number, and then to find out whether or not we're moving at a constant velocity (objectively), we can simply measure the speed of light. We'd then know if we're moving or if we're stationary by seeing what this measured speed of light is and comparing it to "the" speed of light that we've defined.

Note that the speed of light is different than all other speeds. For example, there is no problem with measuring one train to be moving faster than another, so why is there a problem with potentially measuring light to be faster to one observer than to another? The difference lies in the fact that light travels through the **vacuum**, whereas everything else (at least, all other kinds of waves) travel through some kind of medium. Ocean waves need water, sound waves need air, trains need tracks. Thus, when we measure the differences in speeds between these other objects, we can observe the reference objects from which we're measuring them and adjust our physical description of the world accordingly. Therefore this is in line with "all physics is the same for constant relative velocity observers". However, with light, there is no such "reference object" because light doesn't need any medium to move in or on. So if we take the first assumption seriously (which we must), and if we find out that light is truly moving in the vacuum and not any kind of background material (which is the experimental finding that motivated Einstein), then it **must** be the case that the speed of light is the same for anyone who observes it.

So what does this all actually mean in terms of physics? Well, suppose I'm on a train (as most special relativity students are) and I throw a baseball down the length of the train car, from end to end. Suppose I throw the ball with the constant velocity of $10\text{m}/\text{sec}$, and suppose I have a "baseball speed measuring device" in my train. Obviously my measuring device reads $10\text{m}/\text{sec}$. Now suppose that you're

standing on the train platform as all of this happens. Suppose my train is barreling through the platform, not stopping (I'm on the express line), so you just see my train whizz by with constant velocity. Suppose my train is moving at $100m/sec$ and suppose you also have a "baseball speed measuring device" with you on the platform. What will your measuring device measure if I throw the ball right when the train is passing you? Of course, it will measure $110m/sec$, since the train is moving at $100m/sec$ and the ball is moving (in the train) at $10m/sec$.

Now suppose that instead of throwing a baseball from one end of the train car to the other, I shine a flashlight. Suppose I also have a "light speed" measuring device. My measuring device will read $c = 3 \times 10^8 m/s$, because that's the speed of light. Suppose someone else (who doesn't know anything about special relativity), on the platform, also has such a device and measures the light as my train passes him. As my train passes, he expects to see the reading $(3 \times 10^8 + 100)m/sec$, because he knows the speed of light and he knows how fast my train is going. What he doesn't know, however, is that he wasted his money on a "light speed" measuring device (we both did), because we know what this device will measure every single time (in fact, the device would give the correct answer to the speed of light measurement just as often as it would if it were just a cardboard box with the expression " $3 \times 10^8 m/sec$ " written on it). Indeed, his measuring device **also** reads $3 \times 10^8 m/sec$!

Take a moment to really consider how weird that is. I'm in my train seeing the light move at $3 \times 10^8 m/sec$, and you're on the platform as my train comes barreling down at $100m/sec$, yet you see the light—the **same** light that I'm seeing—**also** move at $3 \times 10^8 m/sec$. This seems impossible. Usually, when we see something moving at speed w and something else that is located **on** or **in** that thing moving at speed v , then the total speed of the second object that we observe is $w + v$. Here, however, the light is moving at speed $c (= 3 \times 10^8 m/s)$ on an object that moves at speed v , yet the platform observer **still** sees the light moving with a **total** speed c .

So how is it possible for both of us to see light move with the same speed, even though we're in relative motion with each other? How can we observe an object moving at the same speed in the same direction, even though we're moving relative to each other? It turns out that the only way out of this predicament, as we'll see, is to let one of the observer's time **slow down** relative to the other. In other words, if we insist on the fact that we're both observing the same value for c , then the passage of time itself is different for each of us (if and only if we're in relative motion).

This is not only highly counterintuitive, but also pretty crazy at first sight. How can we be so committed to having the speed of light be the same to all observers that we're willing to make time itself speed up and slow down? Well, for one, it's not up to us to determine what light does or how it behaves or how it looks to its observers. It's an experimental fact that has been tested and retested for over a century, and therefore it's simply a fact of the universe that we have to get used to. But we can only say this now with confidence because of the 100+ years of experimental support for this fact. This was much more difficult to accept at the time, and it was a true sign of Einstein's profound genius to be able to see this idea through and develop it into a full physical theory—special relativity. This theory has also been tested in almost every way, and it has withstood every challenge by predicting the values that we see in labs and particle colliders with extraordinary precision. Let us therefore go on to see some of the mathematics that lies within this great theory.


6.4 Time Dilation, Length Contraction

Rather remarkably, we can derive some of the main results of special relativity using only some basic algebra and Pythagorean's theorem.

As usual, we find ourselves in a situation where I'm on a train and you're on a platform (a platform that is stationary with respect to the Earth). Relative to the platform, where you're standing, let's suppose the train is moving to the right with speed v . Suppose also that I've built a clock of sorts, albeit a very weird clock. In particular, suppose my clock consists of two mirrors, where one mirror is placed on the floor of the train and the other is on the ceiling. Supposing the train is L meters high, we can then be sure that the mirrors are L meters away from each other (suppose they're negligibly thin mirrors).


Now, my "clock" involves simply putting a single photon (an individual "particle" of light (for now, since a photon is actually a rather advanced notion, let's just consider this to be a particle that moves at the speed of light)) into my mirror system and letting it bounce up and down in the train car. This is a clock because we know that the particle will always move at speed c , and so the time it takes to get from the bottom mirror to the top one is $\frac{L}{c}$ seconds. Therefore we know how much time passes between each "tick" of the "clock"—namely, $\frac{L}{c}$ seconds. See figures 6.1 and 6.2 for the physical set up (sorry for my poor drawing skills).

Figure 6.1:



Now let's ask the following seemingly harmless question: how long does it take for the particle to get from the bottom mirror to the top one (just one trip)? From my perspective in the train car, moving along with this clock, the answer is simply what we had before: $t = \frac{L}{c}$. But from your perspective on the platform, this is a very different question. This is because now the train car, the mirrors, and the particle all have a horizontal component to their velocities. In particular, since the whole train (and therefore all of its contents) is moving with speed v to the right, it is also the case that the particle needs to travel with this horizontal speed as well. Thus, in addition to moving the L meters in the vertical direction, it also must travel vt' meters to the right, where t' is the amount of time that **you** measure the process to take. We don't know what t' is yet, but whatever it is, it will tell us how far the particle has moved to the right.

Figure 6.2:



Therefore, the particle has to travel the route as shown in figure 6.3.

A simple application of Pythagorean's theorem tells us that according to you on the platform, the particle has to move a distance of $d = \sqrt{L^2 + (vt')^2}$. But remember that the particle is still traveling with speed c according to you (and everyone/everything in the universe). So since the time t' you measure equals the distance the particle goes divided by the speed with which it moves, we have

$$t' = \frac{\sqrt{L^2 + (vt')^2}}{c}.$$

We now do some algebra. The above expression implies

$$ct' = \sqrt{L^2 + (vt')^2}$$

and squaring both sides gives

$$(ct')^2 = L^2 + (vt')^2$$


which means, after doing some more algebra, that

$$t'^2 = \frac{L^2}{c^2 - v^2} = \frac{L^2}{c^2} \times \frac{1}{1 - \frac{v^2}{c^2}},$$

where we pulled the factor of $\frac{L^2}{c^2}$ out front because we know that this equals t , the time I measured. Thus, taking the square root of both sides and plugging in our above value for t , we find that

$$t' = t \times \sqrt{\frac{1}{1 - \frac{v^2}{c^2}}}.$$

Figure 6.3:



Now, since $1 - \frac{v^2}{c^2}$ is always less than 1, we have that the factor multiplying t in the above expression is always **greater than 1**. What this means is that the time that time passes more slowly for me than it does for you, since you will measure **more time** pass than I will (as long as my train is actually moving). This phenomenon is called **time dilation**, and it's very real.

(Note: the keen reader may be worried that although the expression $1 - \frac{v^2}{c^2}$ is always less than one, it might also seem like this expression can be negative, which would mean we're taking the square root of a negative number and therefore somehow introduce complex numbers into this whole discussion. This reader doesn't need to worry, though, because it can be proven (using only special relativity) that nothing moves faster than light (which might be a familiar fact), and therefore we always have that $v \leq c$, so that $v^2 \leq c^2$, so that $\frac{v^2}{c^2} \leq 1$, so that $1 - \frac{v^2}{c^2}$ is greater than or equal to zero. We therefore have that $0 \leq 1 - \frac{v^2}{c^2} \leq 1$, and so taking the square root of this expression doesn't require using complex numbers. We won't prove why nothing can move faster than light, and instead we'll leave this for a full class in special relativity.)

Using the above expression, we can now understand why it is that we never experience time dilation in our daily lives. I.e., why did it take humans so long to figure out that time is actually different for different observers? The reason is now clear: the speeds of all the things we're used to interacting with are so small in comparison with c that this effect is never even felt. This is why we still measure the baseball moving at the speed $110m/s$ in the above example. In reality, its velocity is indeed being altered by time dilation effects, but this alteration is so fantastically tiny that we couldn't possibly detect it (with our eyes or even our fanciest machinery). For example, a commercial airplane travels at about $250m/sec$, which we can round up to $300m/sec$. Now let's suppose that we consider something moving **ten times faster than this**, which is $3,000m/sec = 3km/sec$. This speed is **certainly** faster than most (all?) things we experience in any given day, or lifetime. Let us then suppose we do the above experiment on this super-fast jet, and

let's see how much more time you experience than me (if I'm the one on the jet). Using $v = 3,000m/sec$ and $c = 3 \times 10^8 m/sec$, we have that

$$\frac{v^2}{c^2} = \frac{9 \times 10^6}{9 \times 10^{16}} = 10^{-10},$$

so that

$$1 - \frac{v^2}{c^2} = 0.9999999999,$$

and that

$$\frac{1}{1 - \frac{v^2}{c^2}} = 1.0000000001$$

so that

$$\sqrt{\frac{1}{1 - \frac{v^2}{c^2}}} \approx 1.00000000005.$$

Therefore if $t = 10^{11} sec$, then $t' = (10^{11} + 5) sec$. Since 10^{11} seconds is roughly equal to 3,000 years, and since t is the time I measure in my jet and t' is the time you measure on the ground, I'd have to stay in my super-fast jet for over 3,000 years before your wristwatch ticks off five more seconds than mine, or equivalently I'd have to fly around in my jet for over 600 years before our clocks differ by a **single full second**. This is clearly not something I would ever notice, and seeing as we don't even spend all that much time on jets that travel 10 times as fast as commercial airplanes, this effect is totally negligible and in no way observable to us at the speeds we usually operate at. This effect only becomes relevant (and then it becomes **very** relevant) at speeds that are very close to the speed of light, which are speeds we have no intuition for. This is why we have so little intuition for special relativity, but this has nothing to do with the rightness of wrongness of the theory.

We should also note that this result is completely independent of how I built my clock. The point is that this is just **one way** of seeing that in order for our laws of physics to be consistent with our two defining assumptions about special relativity, I **must** experience time differently than others who are in relative motion. Our clock example is one way of deriving the result, but the result itself is completely general. My wristwatch, heartbeat, brain signals, and breathing patterns all slow down, since all of my internal organs are obviously (hopefully?) moving along with me. All the physics in the train car slows down with me as well (assuming they're moving with me), so that while I'm in the car I don't notice anything strange at all. I.e., I notice "the same physics" as the guy on the platform, and to me **he's** the one who looks like he's moving. I therefore have no way of knowing who's "really" moving, as should be the case.

Before going on to length contraction, we'll look at one example of how relevant this effect is at high enough velocities.

Exercise 6.1. If I'm in a rocket ship traveling at 99.999999% the speed of light (i.e., $.99999999c$ meters per second), how long, in my time, will I have to travel before your clock (here on Earth) has gone 10 years (approximately 3×10^8 seconds) into the future of mine?

The above example shows that time travel—at least into the future—is entirely allowed by our laws of physics. In fact, if you go outside for a jog, you'll have technically "traveled into the future" of those who aren't jogging with you, since your time will slow down compared to theirs. The only caveat is that the amount of "extra time" that you have is so mind-blowingly small that it is completely negligible, and not felt even in the slightest—sorry.

We now go on to discuss one more very real physical effect of special relativity, known as length contraction. To make my job of typing more easy, and to coincide with standard physics conventions, let's make the following definition:

$$\gamma = \sqrt{\frac{1}{1 - \frac{v^2}{c^2}}}.$$

This way, the time dilation equation is written $t' = \gamma t$, which is much easier to write/type (so long as we remember what γ is).

Now suppose we have some rod of length D (I don't want to use L because that's how tall our train car is) lying on the floor of our train, "lengthwise", in the sense that it's parallel with the train car. Now suppose I want to measure it, and that I don't have a ruler. Using the constancy of the speed of light, I can get around my lack-of-ruler problem by setting up the following (perfectly good) measuring device. Let me put a mirror on both ends of the rod, and then let me bounce a photon from one end to the other, and back. Since I know the speed of light is c , all I have to do is time the photon's roundtrip journey to get the length of the rod. In particular, the photon will make a journey of $2D$ meters (out and back) in time t (according to my own wristwatch), and so we have $D = \frac{ct}{2}$.

Now let's see what our platform observer measures as the length of the rod, where he gets to borrow my measuring apparatus. Let's suppose the length this observer measures is D' (we don't yet know what D' is, i.e., it could be that $D = D'$, but we'll soon see that this **isn't** the case). He starts his clock when the photon leaves the mirror towards the back of the car, so that the photon at first starts heading the same direction that the train is heading. But since the train is also moving, the photon doesn't just need to travel the distance D' , it instead needs to travel the distance $D' + vt'_1$, where t'_1 is the time this observer measures between the photon leaving the back mirror and hitting the first.

Now let's look at the photon's return journey, from the view of the platform. As the photon leaves the front mirror and heads towards the back of the train car, the photon and the train are moving in opposite directions. Thus, the photon now only needs to travel the distance $D' - vt'_2$ (because now the photon and the mirror are moving toward each other), where t'_2 is the time measured by the platform observer between the photon leaving the front mirror and hitting the back mirror. We therefore have that the total distance the photon travels in this roundtrip journey is $2D' + vt'_1 - vt'_2$. The total time it travels for (in the platform observer's frame), is by definition $t'_1 + t'_2$. Using "distance=time times velocity" for each of the times t'_1 and t'_2 , we find

$$D' + vt'_1 = ct'_1 \quad \text{and} \quad D' - vt'_2 = ct'_2$$

so that

$$t'_1 = \frac{D'}{c - v},$$

and similarly that

$$t'_2 = \frac{D'}{c + v}.$$

Thus the total roundtrip time $t'_1 + t'_2$ is

$$t'_1 + t'_2 = \frac{D'}{c - v} + \frac{D'}{c + v}$$

which, after some algebra, gives

$$t'_1 + t'_2 = 2c \frac{D'}{c^2 - v^2} = \frac{2}{c} \frac{D'}{1 - \frac{v^2}{c^2}} = \frac{2}{c} \gamma^2 D'.$$

Now $t'_1 + t'_2$ is the analogous time as t , which was my measured roundtrip time. Thus, using our time dilation equation, we know that $t'_1 + t'_2 = \gamma t$, so we have

$$\gamma t = \frac{2}{c} \gamma^2 D',$$

and now using the fact that $t = \frac{2D}{c}$, we have

$$D = \gamma D',$$

or equivalently that

$$D' = \frac{D}{\gamma}.$$

This means the length measured by the platform observer is **shorter** than my length by a factor of γ (which is always greater than or equal to 1). Again, due to the fact that γ is usually absurdly close to 1 for speeds that we're used to, we never observe moving things shrink along the direction of motion, but at speeds close enough to c , we would. The following exercise shows this.

Exercise 6.2. How fast do I need to be going (in terms of c) relative to some observer in order for an object of length L to contract to the length $\frac{L}{2}$ in his frame?

One thing to note is that this contraction only occurs **along the direction of motion** of the object, so that if the rod had some kind of vertical extent (i.e., towards the ceiling of the train car), then this length would not be affected at all by the rod's horizontal motion. This is simply due to the fact that the various directions in space don't know what's going on in the other directions, and are (and should be) independent of each other. Of course, if the train also started to have a vertical velocity (i.e., if the train started to fly), then we'd also have length contraction occurring in this other direction as well. Therefore, if one of our observers were a skydiver who had already reached terminal velocity and is therefore falling at a constant rate, then it will observe (a negligible amount of) Lorentz contraction in the vertical extent of the rod.

These are just a few of the weird things that occur in the world of things that move fast. It is also our first exposure to physics that is completely beyond our intuitive grasp. Sure, some people learn to develop an "intuition" for these sorts of ideas, but at the end of the day our mental picture of how these things work will never be as clear as, say, a ball rolling down a hill. One can (and should) get used to the mathematics, which gets progressively more complicated, and one can (and should) have an understanding of where these concepts come from and what motivated them, but one will never be "at home" in the relativistic world, simply because our minds are not wired to do so. As we saw, even some of the most advanced propulsion technology that we have (times 10) is nowhere near capable of letting us move fast enough to observe these effects, and so they're totally foreign to us. Nonetheless, we can make fantastic progress by finding the correct math to describe this stuff and by holding on tightly to the concepts that we can understand deeply. Before ending this chapter, we'll take a brief look at the group structure of relativistic space-time—much as we did in the classical space-time case. As we'll see, this group structure is very counterintuitive, but that should be expected since it needs to reflect very counterintuitive ideas.

6.5 *The Group Structure of Relativistic Space-time

Recall that our description of classical space-time involved using \mathbb{R}^4 , where we picked one special factor of \mathbb{R} and said that for each element $t \in \mathbb{R}$ (where we viewed t as a moment in time) there is a copy of \mathbb{R}^3 which is a normed vector space, with norm $d(x_1, x_2, x_3) = \sqrt{a_1^2 + a_2^2 + a_3^2}$. We then talked briefly about the group of rotations of space (at a particular moment in time), and this was the group of linear transformations of \mathbb{R}^3 that kept all (squared) distances between points $(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2$ the same. Well, in the relativistic case, lots of things will change.

First and foremost, we've now seen how motion in space affects one's motion "in time". I.e., spatial motion has an effect on one's passage of time. Therefore, it seems unlikely that we can "split up" our \mathbb{R}^4 into "time" on one hand, and "space" on the other. This is simply because now time and space are intimately related to each other. Now, it's clear that we've already assumed that space-time is still \mathbb{R}^4 , and this is the correct assumption to make since labeling a time and 3 coordinates of space still labels every event in the history and future of the universe. However, we simply need to change the **structure** of this \mathbb{R}^4 to make it consistent with special relativity.

Before seeing this, let's note again **why** we can't split time and space up the way we did before. This is because if we put coordinates down in our universe, including time, and if we've split time and space up like we did in the classical case, then one observer's "time" will be different from another's if they're moving relative to each other. We would then be comparing quantities at different times and therefore different "copies" of \mathbb{R}^3 , and we wouldn't have a natural way of comparing quantities in one "copy" of \mathbb{R}^3 to those in a different copy of \mathbb{R}^3 . This is why, when we realize that spatial motion affects time passage (as we've realized in special relativity), we must view all of space-time "as a whole", i.e., we have to view it all "at the same time" (pun intended).

It turns out, for reasons that we won't go in to here, that instead of considering the real number $(x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2 \in \mathbb{R}$ as the distance squared between two points $(x_1, y_1, z_1), (x_2, y_2, z_2) \in \mathbb{R}^3$, we should **actually** be considering the number

$$-(t_1 - t_2)^2 + (x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2 \in \mathbb{R}$$

as the "distance" between two space-time points $(t_1, x_1, y_1, z_1), (t_2, x_2, y_2, z_2) \in \mathbb{R}^4$. There are a couple things to note here. First, there's a negative sign in front of the $(t_1 - t_2)^2$ term, which means that this "distance" can be negative! This is why I used quotes here, since we usually think of distances as being non-negative. This is what's counterintuitive about it all, but that's okay, counterintuitive-ness should be expected. Second, just as spatial rotations keep the distances $(x_1 - y_1)^2 + (x_2 - y_2)^2 + (x_3 - y_3)^2$ fixed, it turns out that the space-time "rotations" that we care about in special relativity are precisely those that keep the "distances" $-(t_1 - t_2)^2 + (x_1 - x_2)^2 + (y_1 - y_2)^2 + (z_1 - z_2)^2$ fixed between any two space-time points. In other words, just as spatial rotations would change the coordinates of our points but not change this special number which we called distance, it now is the case that space-time rotations will change the coordinates of our space-time points (including mixing up the space and time coordinates), but they'll do so in a way to keep this new "distance" unchanged. Finally, we note that if we let $t_1 = t_2$ so that we **are** comparing things "at the same time", then we recover the usual distance on the spatial coordinates.

It turns out the set of all linear maps that preserve this new "distance" on space-time is a group, and it's a very important group. We can, in some ways, view this group as a product of the group of spatial rotations that we already know and love (which only mix up the spatial coordinates) and the group of rotations that mix space and time. In other words, just as we can rotate the y-axis into the x-axis in physical space, we can now rotate the x-axis into the t-axis, for example. So time is still "special" in the

sense that it has the minus sign in front of its "distance" term, but other than that it is viewed on equal footing as the other three dimensions of space-time.

This group is much harder to visualize or understand when given as short of an introduction as has been given here, but it is worth mentioning because it's a beautiful subject when it's learned properly. Hopefully this gives enough of a flavor of this subject to motivate the reader to continue on and learn this in detail. What it amounts to is studying a couple special matrix groups, and this is just one other way in which matrix groups (which we mentioned in the chapter on vector spaces, and which we won't be studying in this text) play a fundamental role in physics.

We've left a lot of interesting and accessible special relativity still undiscussed. For the sake of space (since these lecture notes are already quite lengthy), I'll only mention these topics now so that the interested reader can a) further pursue these ideas on her own and/or b) have something to look forward to in future classes. These are i) Space-time diagrams and ii) the relativity of simultaneity. Any Google search of either of those two phrases will unleash lots of explanations (some better and some worse) of some more topics in special relativity. This is a profoundly rich subject and is truly at the heart of **all** modern physics, so any future physicist (and even a lot of future mathematicians) would do well to learn anything and everything they can about this weird and gorgeous physical theory.

Our next and final stop on this wild ride is into the whacky world of quantum physics. Quantum physics is just as important and fundamental as special relativity, but it is even less intuitive. I'll give a fair warning to the reader at the start of the next chapter, so let's just dive right in!

Chapter 7

The Quantum World

7.1 Introduction

We now take a brief and as usual incomplete tour through the whacky and wonderful world of the quantum. We started our journey in the classical world, where all of our results were rather intuitive and obvious. We found some of the mathematical structures that, with the proper interpretation, modeled this world well. From there we went on to the relativistic world, where objects move close to the speed of light, time and space are fused into one, and our passages through space-time are all unique and dependent on our motion in space-time. In this world we saw phenomena that are very counterintuitive, since we have no sensual experience of time slowing down or lengths changing due to relative motion.

Now, however, our step away from what is intuitive or obvious will be about 10 times as large as our step from the classical world to the relativistic one. Our step into the quantum realm involves not only opening our minds up to new possibilities (analogous to opening our minds up to the possibility (and reality) of time slowing down), but **also** involves **letting go** of some of our most fundamental preconceived ideas about how the universe behaves. In order to do quantum physics, we need to break our minds free from the shackles that our classical intuitions have put on them. Since we are organisms that are much, much, much larger than individual atoms, our intuitions have not been formed with any ability to "visualize" the world at which quantum phenomena take place. However, the universe is—at its most fundamental level—a quantum mechanical world, and we therefore cannot ignore the consequences of quantum phenomena.

We will be studying only the most simple quantum mechanical system, and therefore it should be clear that there is **lots** of interesting material that will not even be mentioned here. But from this one system alone, we will see lots of the weirdness that lies at the heart of quantum mechanics, and how this weirdness manifests itself to us. Many of the axioms of quantum mechanics will be presented completely axiomatically—and by that I mean without motivation. A full description of quantum mechanics and its corresponding mathematical model took lots of extremely brilliant people a long time to work out, and even today there are still some (lots of?) subtleties that are not fully understood. We therefore will not go into what motivated the things that we describe, but will merely state their mathematical structure and physical interpretation. For now, the reader will just have to take on faith that these ideas are **extremely** well motivated by hundreds of experimental findings, and have been **remarkably** successful in describing and predicting physical phenomena. This brief tour will hopefully motivate the reader to pursue further study on her own.

Our last introductory note will be that the "understanding" of quantum mechanics that we'll attain here will be very different from the understanding of, say, classical physics that we had. In fact, there is no amount of study that will inject the reader with the same kind of understanding of the quantum world,

simply because it is highly likely that no deep understanding of this world really exists. In particular, we can understand the mathematics of this theory and use it to make the correct predictions—we can even use this mathematics to invent new technologies (like the computer I'm typing this on right now)—but we'll never be able to "see" the entanglement between two particles or the superposition of one particle in two states (these terms will be "understood" soon) the same way that we can "see" a ball rolling down an inclined plane or an object falling under the influence of gravity. As one of the greatest physicists of all time, Richard Feynman, once said, "If you think you understand quantum mechanics, then you don't really understand quantum mechanics". This is to say that if you think you have the correct "mental picture" of what's going on, then you have the wrong mental picture, since there is simply no mental picture or classical analogy that can be obtained for most quantum phenomena.

Now that we've (hopefully) opened our minds up to some wildly new possibilities and freed them of their classical shackles, let us go forward.

7.2 The Importance of Measurement

One of the most important aspects of quantum mechanics is the newly found emphasis on the act of measuring. It is very difficult to make any progress in this field without appreciating the importance of the act of measurement, but "the act of measurement" is also one of the most subtle and difficult concepts to fully understand. In fact, it is our lack of a deep understanding of this process that motivates quotes like Feynman's above. There are volumes written about the philosophy of measurement and the different world-views one can take regarding what a quantum mechanical measurement "really is". To avoid this rabbit-hole, we'll try to keep our language as loose and intuitive as possible, while still holding on to some kind of mathematical rigor.

In short, we'll define a measurement to be the act of asking a system a question. Now this may seem like a somewhat strange way to phrase things, but we'll see that it actually works rather nicely. Of course, it should be understood that no one is really "asking a question" in the human sense of the phrase, but rather in the more abstract setting that we'll explore now.

Suppose we have some physical system and that it is (at least to a high approximation) isolated from the rest of the universe. For example, if I have a box full of gas particles sitting on my table, and if the box is sufficiently well sealed, then we can view it as being sufficiently well isolated from the universe. This is because what goes on in the box is sufficiently unaffected by the physics going on elsewhere in the room (except for maybe the temperature of the room), and it is **certainly** unaffected by, say, the physics going on at the surface of Mars, or the Sun, or in some galaxy halfway across the universe. So when we say "a physical system", we mean one in which all the parameters that determine its evolution in time are constrained to lie within the system itself, and we don't need to consider the affects of physics going on elsewhere.

If we consider the box of gas on the table again, then it is clear that we can very easily make it dependent on the room's physics. Namely, we can open the box, thus letting the gas that used to be isolated in the box begin to interact with the physics of the room. Now, we must view the entire room-box system as one single physical system. Of course, if I open the door of the room, then I'm now exposing the room to the physics of the hallway, so I now must consider the hallway-room-box system as one single physical system. We can obviously keep going. What we're doing is "coupling one system to another" in all of these cases, and we'll use this terminology from time to time.

Now let's say I have some physical system evolving in time in some way, and suppose I know exactly

how it started out. For example, suppose I have a particle at position $x = (x_1, x_2, x_3)$ at time t_1 , that it's moving with velocity $v = (v_1, v_2, v_3)$, and that it's subject to a force F . These are the so-called "initial conditions" for my system. Now, if I want to know where this particle is at some later time t_2 , I need to make a measurement. In fact, with the definition of measurement that I've given above (of asking the system a question), I'll need to make infinitely many measurements. Namely, at time t_2 I need to go to every position in the universe and ask my system "is the particle here?". For example, if I want to know if the particle has gone to position (y_1, y_2, y_3) , then I need to go to that position and ask the system if the particle is there.

This may seem like a pretty inefficient way of doing things, but in the next section we'll see that it's actually a necessary way to phrase "a measurement". This is because, as we'll see later, a quantum mechanical measurement is very different from a classical one. For now, though, let's say a few more words about quantum measurements.

Before we try to say what an act of measurement actually is, let us first consider the role measurement plays in our physical theories. Perhaps the best way to see this role is to ask ourselves why we haven't talked about measurements in either the classical or the relativistic regime. Namely, if we can understand why we **didn't** need to discuss measurements before, then we might gain some insight into why we **do** need to do so now.

Let us therefore consider a classical system, and to be precise let's consider a particle falling under the influence of a gravitational force. If we know it's position and its velocity (i.e., its initial conditions) at some time t_0 , then we know what the system is doing for all times thereafter. But in any real experiment, we have to pick a time to measure the system, and only then do we actually gain information about where the particle is. Of course, our laws of physics told us before we made the measurement precisely where the particle would be, so making the measurement didn't give us any new information. The point is, however, that to actually "see" where the particle is at, say, time t_1 , we have to look at it at that time. We'll see in the next section how intimately related this all is to the fact that our classical laws of physics are completely "deterministic", which we'll describe later. For now, though, the important thing to take away is the distinction between "knowing where something is" in the abstract, and actually "seeing" something. It is only by measuring the system that we "see" something, and the rest of the time we can only "suppose" that we knew where the particle was.

But "knowing where the particle was" at a time that we weren't actually measuring the particle is a very subtle issue, and it turns out that quantum mechanics exploits this issue fully. For what does it really mean for us to "know where the particle is" or to "know what the particle is doing" if we're not actually measuring the particle doing it? Namely, our laws of physics may be such that it **seems** as though we know what the particle is doing even when we're not observing it, but in order to check that we're right we have to make an observation. But any observation is done at some point in time, so that all we're really doing is checking that our laws of physics were "working" at the times at which we measured them. If we wanted to know what our system was doing at some intervening time, then we should measure our system at **that** time. This is a subtle point, and may need some reflection (perhaps for a lifetime). But to summarize, a measurement is a way of obtaining information about a system at a specific time. To ask what the system was doing "before that time" is a pointless and ill-defined question, because if we wanted to know what the system was doing at this earlier time, then we should have measured it then. In classical and relativistic physics we never had to make this distinction because our laws gave us clear answers to the question of what our system was doing even when we weren't looking, but this is merely a peculiarity of these physical theories and in no way a basic assumption of physical theory. In other words, there is no reason why we **should** be able to know what a system is doing when we're not looking at it. As we'll see, quantum mechanics gives us something new to work with.

Finally, it is important to discuss how measurements are actually done. If we consider our box of gas on the table again, then we realize that if we want to obtain any information from it (i.e., ask it a question), then we need to look inside, or pick it up and shake it, or heat it up and see what happens, or cool it down, or do something to it that alters its current state and see how it responds. This is a very general and important phenomenon when it comes to measurements—that doing a measurement requires altering our system **in some way**. For if we don't alter our system at all (by shining light on it or "opening the box"), then the system will continue to sit there, isolated from the universe, forever, and we'll never know anything about it. This may seem obvious or trivial, but in quantum mechanics this act of altering a system is crucial. We also note that each of these examples of a possible measurement (opening the box, shining light on a system, etc.) involves "coupling" the system we're trying to measure to some other (usually larger) system. In fact, as it stands, our best description of a measurement is in terms of this sort of coupling, but this is a very subtle issue and certain problems are still not fully resolved, so we'll leave this issue for future individual study for the reader.

Again, let's consider why we **didn't** need to worry about altering our systems in classical physics when we measure them. If a ball is falling through a gravitational field, and we want to know where it is, we might consider shining a flashlight on it. Well, believe it or not, this alters our ball-gravity system, because the flashlight is shooting photons of light at the ball, which proceed to reflect some light from the ball to our eyes, transferring some momentum to the ball and thus altering its trajectory through space. However, for any ball of reasonable size, this alteration is completely and utterly negligible. We might see the ball's trajectory be affected by the wind (as anyone who's played basketball outdoors will know), but the wind is a stronger force than bouncing photons by a factor so utterly gigantic that we usually don't even consider "turning the lights on" as affecting the trajectory of anything through space.

In quantum mechanics, however, where things are extremely tiny (like a single atom, or smaller), these sorts of phenomena play a huge role. We'll see what role this is shortly, but for now we should just focus on seeing how it's possible for measurement to play a completely different role in quantum mechanics than it does in classical mechanics, as this discussion hopefully made somewhat clear.

7.3 The Loss of Determinacy

Perhaps the most startling and troublesome aspect of quantum mechanics is the fact that it is a probabilistic theory. Now, by "probabilistic" I don't mean in the same way that rolling a pair of dice is. This is because we view a pair of dice as being a classical system, and the probabilistic behavior of a given role is based solely on the fact that we would need to calculate the precise forces with which they were thrown—and their precise initial conditions—with such an absurd level accuracy that it's just **easier** to view the system as probabilistic. The point is, though, that "underneath" this approximation is nothing but a set of deterministic laws of physics, all working in concord to bring about some result that, had we spent the time and the energy to calculate it precisely, we would be able to predict.

This warrants a brief definition of "deterministic". For a more drawn out, lengthy, and probably convoluted definition, one might want to ask a philosopher. For a practical definition, a "deterministic" system is one such that given its initial conditions and knowledge of all of the forces acting on it, we'll be able to predict the answer to any question asked about it—with certainty—for any subsequent time. For example, if we know where a particle is initially and how fast it's moving, and we know that the only force acting on it is gravity, then we'll be able to predict **with certainty** its position for all subsequent time (until, at least, it runs into some other object, but then that's just extending the system that we were initially talking about and is not inconsistent with determinism, for had we known about the presence of this other object and

included it into our description of the system, then we'd still be able to make predictions **with certainty**).

The reason I keep putting "with certainty" in bold is that when we do quantum mechanics, we lose our certainty and learn that the fact that classical physics describes things in "certain" terms is merely a peculiarity of that theory. Quantum mechanics is very different. If we know everything that there is to know about a quantum mechanical system at some time, then it is **not** the case that we'll be able to predict **with certainty** what it will be doing at some later time (i.e., some later measurement). What we **will** be able to do, however, is predict the **probabilities** with which it will be doing something, and we'll be able to predict these probabilities with immense accuracy.

For example, suppose we know exactly what state our system is in (i.e., we know everything there is to know about it) at time t_0 . If we then leave the system alone and let it evolve naturally in time, it will evolve into what's called a "superposition" of states. We'll see what this means mathematically in a short while, but physically it means that our state is **simultaneously** in several different states at the same time. This seems impossible, but it's how the world works. Of course, when we "open the box" and measure the system, we're going to see the system in only one state. Quantum mechanics tells us the **probabilities** with which we'll observe the system to be in a particular state.

What happens upon measurement is called the "collapse of the wave function" or the "collapse of the state vector" (again, we'll see why we call it this shortly). The state goes from being "in many states at once" to "being in the state we observe" instantaneously, and it does so because our measurement has, by definition, disturbed the system. I.e., our disturbance of the system by measuring it causes the state to "jump" from a collection of simultaneous states to one state—the one state that we observe after the measurement—**instantaneously**.

This warrants a bit more comment. One might be tempted to think of this "collapse" from "many states at once" to the observed state as follows. One might view this as our system **really being** in only one state, and that we just don't know **which** state it's in until we measure it. Upon measuring it, we then find out which state it **was** in before the measurement. But this is wrong. It is not the case that the system "was in" a particular state before the measurement, and that our measurement simply "uncovered" what that state was. It truly is the case that the system existed as many states simultaneously, and our measurement **changed** this state of affairs to one in which only a single state is observed. There is mathematics and experimental evidence proving that this is the only correct interpretation, and that the interpretation of this as simply "uncovering" what the state was is absolutely wrong, and impossible. We won't go into what this experimental evidence or mathematical framework is, and only take on faith (for now, until a proper quantum mechanics course), that we **must** interpret quantum mechanical systems as **truly being in more than one state at a time, until a measurement disturbs it and "kicks" it into the single state that we observe**.

At this point in time it might seem like quantum mechanics is a pretty weak theory, since it doesn't tell us what we're going to find on any given experiment. But this is not so. Quantum mechanics tells us a lot about a system, it just doesn't do so in a deterministic way. This should be expected, though, because **the universe itself** is not deterministic.

What quantum mechanics does is tell us which collection of states our system is in a superposition of—i.e., it tells us all the possible states that our system can "collapse" to when we measure it. Moreover, for each such state, it tells us the probability with which our system will collapse to it. We can then test quantum mechanics by seeing how accurate these probabilities are.

In order to do this, though, we need to set up several hundred (or more) identical copies of the same system, and then perform the identical measurement on each one of them. Each individual measurement will "collapse" the system to exactly one of the states that it is existing in simultaneously (because we can't actually observe the system doing two different things at the same time), and after several hundred measurements a probability distribution for how many times it collapsed to various states will build up, and we'll be able to see how accurate our predictions were. It's important to note that we have to set up hundreds of **different** copies of the **same** system because measuring the system will disturb it, as we've described above. Therefore, it isn't okay to simply perform the same measurement on the same system over and over again because each measurement changes the system, and therefore our collection of measurements wouldn't be truly "on the same system". We'll see how this plays out in more detail in the following section.

Finally, we should note how this sort of probabilistic behavior is different from our case of rolling dice. In that case, we knew that there were deterministic laws "underneath the surface", we were just too lazy to calculate out in detail what they told us. In the case of quantum mechanics, **there is no such deterministic underbelly**. So if a system has some set of initial conditions and later evolves into a superposition of five states, there is no physical theory underneath the surface that could predict with certainty which outcome we'll get for any individual measurement. This is a proven fact. It has been shown in incredible generality that there is no possible deterministic physical theory that could reproduce the results of certain quantum mechanical experiments that have been run.

Theories that suppose such a deterministic underbelly are called "hidden variable theories", and the way one goes about proving things about these theories is to suppose that there is some physical theory that is governed by some set of deterministic variables (we can even allow for a theory with infinitely many such deterministic variables). They're called "hidden variables" because we don't even need to know anything about them. All we need to do is suppose that there is some deeper physical theory—one that is governed only by deterministic variables (which can be phrased in a completely mathematically rigorous way)—and we can prove things about it regardless of whether or not we know about all of its details. And in remarkable fashion, a guy named John Bell took results from certain quantum mechanical experiments and proved that there is **no possible** hidden variable theory that could ever exist that could predict the results of the experiment. Therefore, since any deeper physical theory must be able to predict the results of experiments that have already been done, John Bell showed that any fundamental theory of Nature will be **inherently** probabilistic.

This means that the universe itself is probabilistic at its core. There is no deeper level of complete determinism, and our constraint to only be able to calculate probabilities is not a sign of our ignorance, but rather a **fundamental and deep** statement about how the universe works. This is something that is very hard to come to grips with and fully internalize. It is said that Einstein took with him to his grave a firm belief that determinism was still hidden under it all somewhere, even though John Bell proved and stated his results during Einstein's lifetime (motivated, actually, by a thought experiment that Einstein and others did in an attempt to "prove" that the probabilistic nature of quantum mechanics was wrong). No one can blame Einstein for his stubbornness, though, because he was there when all these realizations occurred. Now, with a little hindsight and slightly more open minds (not to say that Einstein's mind was particularly "closed") we can simply accept the truth of John Bell's rigorous mathematical statement about highly scrutinized and firmly supported experimental evidence. The universe is **fundamentally** probabilistic—in ways that rolling dice is not—and this will not change no matter how passionately we wish it otherwise, or how counterintuitive it is. We can't determine the laws of Nature, we can just study them.

At the end of this chapter we'll say a few words about why we see a pretty deterministic world around us, and why this probabilistic behavior is confined to scales we don't experience daily. In order to fully

appreciate these words, however, we need to actually **do** some quantum mechanics, and so this is where we'll turn our attention now.

7.4 The 2-State System

As we now plunge fully into some quantum mechanics, we need to set the stage a bit. What we need at our disposal is an abstract mathematical structure known as "state space". In short, this is the space (which is a mathematical space, not a physical one, and which can usually be thought of as a set with some sort of structure, like a vector space or a group) of all possible states that a system can be in. It is important to emphasize that this space is an abstract one, and the individual points of this space are entire states of a physical system. Let's look at some examples from classical physics first.

Consider a classical system of a single particle in empty space. For some reason, we call the "state space" for a classical system "phase space", so we'll stick with that terminology and leave "state space" for quantum mechanical systems. We'll still refer to "states" of a classical system, though. Now, if we consider this "single particle in empty space" system, we find that there are two things that determine its state completely. This particle will have a location in space as well as a velocity. We define the **momentum** of a classical particle to be its mass times its velocity, and since a classical particle's mass is a constant that "comes with" the particle, we have that velocity and momentum are completely equivalent. For various reasons, we tend to use momentum as much as possible, so we have that it's the particle's position and momentum (not velocity) that determine its state completely.

So if we consider the space of all possible states of this simple classical system, we find that it will always have some location in \mathbb{R}^3 , and some momentum, which again is an element of \mathbb{R}^3 (because we know that velocity can be described by an \mathbb{R}^3 and that velocity and momentum are effectively synonymous (modulo a factor of mass)). Since the particle can be at any one of these locations, and with any possible momentum, we find that the state space S for a one particle classical system is $S = \mathbb{R}^6$. This is a 6-dimensional space, and so it is important to understand that it's completely abstract—don't try to visualize it (no one can). Now, if we specify a point in S , we're **completely characterizing** the system at some time, because we'll have assigned to the particle some location and some momentum, and that's all there is to know about the particle at any given time. Now, as time progresses, this point in S will move around in S as the state of the particle changes, and this path that the point takes will be completely determined by the forces acting on and in our system.

If we now consider a two-particle (classical) system, then the dimensionality of our phase space doubles. This is because to specify a single state of our system, we need to specify the position of particle 1 (three dimensions) as well as particle 1's momentum (3 more dimensions), and then we need to do the same for particle 2 (6 more dimensions), thus making for a phase space $S = \mathbb{R}^{12}$. It should now be clear that the classical phase space for an N particle system is \mathbb{R}^{6N} , because any point in this $6N$ -dimensional space assigns a position and momentum to every single particle in the system, thus completely characterizing everything there is to know about the system.

Suppose we know what all the forces acting on and in our system are. Then, choosing a set of initial conditions should completely characterize the future dynamics of our system (because this is classical physics where everything is still deterministic). But choosing a set of initial conditions is nothing but assigning to each particle in our system an initial position and an initial velocity (momentum), and so choosing a set of initial conditions is the same as picking **a single point** in our phase space S . Then, as the system evolves in time, this point that we've picked out will move around S , sweeping out a well-defined and continuous path, and measuring the system at a later time just means seeing at which point in S we've moved to. We

call this path "well-defined" because the laws of physics here are deterministic, so there's only one path that the point in S can take, and we call this path "continuous" because the system never "jumps" from one state to another in an infinitely short amount of time, but rather it smoothly transitions from one state to the next.

In classical physics, the "extra structure" of phase space is what's called a symplectic manifold, and we haven't and won't explore these gorgeous structures here. When we assign to this space the forces that are acting in and on it, we're assigning extra structure to this manifold, but again we won't see the details of all of this here (the reader may be motivated to continue to pursue math and physics and find out for herself!). Instead, we're just setting the stage for quantum mechanics, which takes "state space" to be a very fundamental structure as well, but which assigns a very different kind of "extra structure" to it. We're now ready to see how this goes. As usual, we'll leave a lot of details out and do our best to a) get to some interesting stuff while b) maintaining a solid amount of mathematical rigor.

It turns out that the best way to model quantum mechanical behavior is to give the state space S the extra structure of a complex vector space. Or, to put it another way, the state spaces of quantum mechanical systems turn out to all be complex vector spaces. Most of the time these vector spaces will be infinite dimensional, but the system that we'll study is only two dimensional. In the event that the reader has forgotten some of our discussion about complex vector spaces, it may be wise to go back now and reread the bits about real vector spaces and complex numbers (since a complex vector space is the same as a real vector space only with scalar multiplication being done with numbers in \mathbb{C} and not in \mathbb{R}).

We now have a state space that is a vector space, so that each possible state for a given quantum mechanical system is fully determined by a single point in S , and is therefore a "state" as well as a "vector". Actually, that is a lie. It turns out that there are some redundancies in this description, and we therefore need to put some extra structure on this space S . It turns out that the correct thing to do is to put an equivalence relation on S , where we view two vectors (two states) as equivalent if there is some non-zero complex number that, when multiplied to one, takes it to the other. In symbols, we have that if $w, v \in S$ and if there is an $a \in \mathbb{C} \setminus \{0\}$ such that $v = aw$ (remember, these are vectors, so we can scalar multiply), then we say that $v \sim w$.

Exercise 7.1. Show that this \sim is really an equivalence relation.

Now that we have an equivalence relation on S , we can define the set of equivalence classes of S , and denote this new set by " S/\sim ". Now we view each point of S/\sim as a distinct state of our quantum mechanical system. In practice, it is much easier to deal with actual vectors, and so this is what we'll do and we'll just remember that any two vectors $v, w \in S$ such that there is a non-zero complex number a such that $v = aw$ will describe **the same** physical state.

The reason that vector spaces are a natural way to describe quantum mechanical states is that in a vector space we can construct linear combinations of vectors. With our current interpretation of vectors as representing (equivalence classes of) states, this means that we can express one state as a linear combination of other states. Why do we want to do this? Well, because in quantum mechanics it is possible for (and often the case that) a single particle or system is **simultaneously in** more than one state (as we discussed above). This is very difficult to wrap one's head around, but it is really the cause of how one system can evolve into a superposition of states. This is in sharp contrast to the classical case, where one state always evolves into one state, and does so deterministically.

At this point, the only way to really move forward is to just state the mathematical model of quantum mechanical systems, and then see some examples of how this plays out. We won't motivate our mathemat-

ics any more than we already have, but we will look at some examples so that we can "get used to" the mathematics.

Let S be the state space of a quantum mechanical system, and let $E \subset S$ be a basis for S (S is possibly infinite dimensional). This means that any $e \in E$ is a state in it's own right, and since all the vectors in E are linearly independent we know that if e and e' are both in E and represent the same equivalence class in S/\sim , then $e = e'$ (or equivalently, any two $e, e' \in E$ such that $e \neq e'$ represent **different** equivalence classes in S/\sim , and thus they represent different states).

Exercise 7.2. Prove the previous statement.

Therefore what we have is a "basis of states", and any other state in S can be expressed as a linear combination of these "basis states". Therefore, if $v \in S$, then there exists some set $\{a_1, \dots, a_N\}$ of complex numbers such that $v = a_1 e_1 + \dots + a_N e_N$ where each $e_i \in E$. We then say that the state v is in a **linear superposition** of the states $\{e_i\}$ (thus, "linear superposition" is the physicists' term for "linear combination", and this is the reason we've used the terminology of "superposition" above). Now, the axioms of quantum mechanics (which took mankind a long time to figure out) tell us that if we perform a measurement on the state v , we'll collapse our system into one of the states e_i , and we'll collapse the system into the state e_i with probability $|a_i|^2$ (where we recall that for any complex number z , $|z|^2 = z\bar{z}$ where \bar{z} is the complex conjugate of z). Remember, these are just axioms that we must accept for now. It is obviously true that whenever we measure a system, we'll always find it to be in only one state, but we simply can't determine which state we'll find it in with certainty in any given measurement.

There are a couple extremely important things to note. First, we note a simple fact about probabilities, and that is that probabilities always need to add up to 1. This is because we're guaranteed that we'll find the system in **some** state, and therefore the sum over all states of the probabilities of finding it in each state should add up to 1 (this is analogous to say, basketball, where the sum of "missed shot" percentages and "made shot" percentages should always be 100%, since **something** happens (make or miss) on every single shot). Now, in the previous paragraph we didn't make sure that this would happen, but this is where we take advantage of our freedom to multiply our vectors by non-zero complex numbers and still retain the same physical state. For if we know that v is not the zero vector (which is always the case for physical states), then we know that at least one of the a_i 's is non-zero in the linear superposition. Therefore we know that the sum $|a_1|^2 + |a_2|^2 + \dots + |a_N|^2$ is strictly greater than zero, and therefore we can multiply v by the non-zero complex number

$$\frac{1}{(|a_1|^2 + |a_2|^2 + \dots + |a_N|^2)^{1/2}}$$

and still be talking about the same physical state. If we do this and define

$$v' = \frac{1}{(|a_1|^2 + |a_2|^2 + \dots + |a_N|^2)^{1/2}} \cdot v = \frac{a_1}{(|a_1|^2 + |a_2|^2 + \dots + |a_N|^2)^{1/2}} e_1 + \dots + \frac{a_N}{(|a_1|^2 + |a_2|^2 + \dots + |a_N|^2)^{1/2}} e_N,$$

then we see that if we take the coefficients of the linear superposition of v' , take the "mod square" of each of them (i.e., the probabilities of finding the state in e_i upon measurement) and adding them all up, we get precisely 1.

Exercise 7.3. Show this.

We call any vector with this property (that the sum of the mod squares of its coefficients is 1) a **normalized** vector. What we've just shown is that we can take any non-zero vector $v \in S$ and find a normalized vector v' that represents the same physical state. We can then speak meaningfully about probabilities, since

any individual probability of finding our system in a particular state e_i will be less than or equal to 1 (as it should be), and the sum of all such probabilities is equal to 1 (because we know that we'll find **something**).

Now, quantum mechanics is the business of calculating v in terms of basis vectors e_i . Quantum mechanics gives us a set of tools to do this with **complete precision**. I.e., we can in principle calculate v **exactly**. Therefore, this is not where the probabilistic nature of quantum mechanics comes in. The probabilistic nature of quantum mechanics comes in when we realize that even when we know v precisely, we still don't know what an **individual** measurement will give us. We can only calculate the **probabilities** for what an individual measurement will give us, and there's nothing we can do about this.

Let us now see how to put this into practice and actually calculate something. The system that we'll consider is a single electron. Now, it turns out that an electron has an extra property in the quantum world known as "spin", and that there is no real analogue for "spin" in the classical world. We will not discuss the details of spin, and for now we'll take the grossly inaccurate picture of an object actually spinning around its axis (the way the Earth spins on its axis to give us night and day). It can't be stressed enough, though, that this is very very wrong. Do not get too attached to this picture of an electron, because it's the wrong picture. It's a useful picture at times (like right now), but it's not what's "really" going on and there are several reasons for this that the reader will learn about in any thorough quantum mechanics course. For now, we'll take this picture seriously and know that for what we'll do, we can uncover all the same truths by using this picture, and that spin is a very real thing that we're just choosing to not go into detail about.

It turns out that an electron has only two possible states of spin. For our purposes, we can think of the electron as either spinning "to the left" or "to the right". This may seem obvious because, for example, these are the only two ways that the Earth can rotate on its axis, but we must note that a) we're using the wrong picture of an electron here and b) there are particles with more than two states of spin, and so "left" and "right" aren't the only two options (this is one of many reasons why our picture is wrong). For conventional reasons, let's say that an electron spinning to the right is in the "up" state and that an electron spinning to the left is in the "down" state. Suppose we don't care about where the electron is or how fast it's moving. Then, there are only two different possible states it can be in: up or down. Therefore, there are only two linearly independent states in S , and so S is a two-dimensional complex vector space.

We know that any finite dimensional complex vector space is effectively "the same" as \mathbb{C}^n for some n , so we can view S as being simply \mathbb{C}^2 . Let's take the basis

$$e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

and say that e_1 corresponds to the state "up" and that e_2 corresponds to the state "down". Then, if we have an electron in a box, we know that its state vector will be in some linear combination of these two states. In other words, the electron will be in some combination of "spinning to the left" and "spinning to the right" **at the same time**. When we measure the state, we'll collapse it into one or the other, since we can only observe it **either** spinning left **or** spinning right ("down" or "up").

So suppose we have a single electron, and suppose this electron is described by the state vector v , and the expansion of v in terms of the basis vectors we just defined is

$$v = \frac{1}{\sqrt{2}}e_1 + \frac{1}{\sqrt{2}}e_2 = \begin{pmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{pmatrix}.$$

Then if we perform a measurement, we'll find our electron in the spin up state with probability $(\frac{1}{\sqrt{2}})^2 =$

$\frac{1}{2}$, (note that since $\frac{1}{\sqrt{2}}$ is a real number, the "mod square" of it is just the number squared, since any real number is equal to its own complex conjugate), and similarly we'll measure it in the spin down state with probability $\frac{1}{2}$. Note that these two probabilities—which are the only two possibilities—add up to 1, as they should. Thus, if we set up a 100 different but identical copies of this same system, where the electron is in precisely the same state in each system, then approximately 50 times we'll measure the electron to be in the "up" state, and the rest of the times we'll measure it to be in the "down" state.

Exercise 7.4. Suppose now the electron is in the state $v = \frac{\sqrt{3}}{2}e_1 + \frac{1}{2}e_2$. Calculate the probability for a measurement to give spin up, and for a measurement to give spin down.

One thing to note about the above exercise and the example that preceded it is that the electron states that we considered were both originally normalized. I.e., the corresponding probabilities already added up to 1. If we consider the following electron state:

$$v = 2e_1 + \frac{i}{4}e_2,$$

then one can easily check that this state is **not** normalized.

Exercise 7.5. Check this.

We then need to normalize it, using the above prescription. Namely, we multiply v

$$\frac{1}{\sqrt{2^2 + |\frac{i}{4}|^2}} = \frac{1}{\sqrt{4 + \frac{1}{16}}} = \sqrt{\frac{16}{65}} = \frac{4}{\sqrt{65}},$$

so that

$$v' = \frac{4}{\sqrt{65}}v = \frac{8}{\sqrt{65}}e_1 + i\frac{1}{\sqrt{65}}e_2$$

which can readily be checked to be normalized.

Exercise 7.6. Check this.

That's it. This is quantum mechanics. Granted, this is by far the simplest of all quantum mechanical systems, and we've skipped over **lots** of interesting points regarding even this simplest of system. In fact, we can just as well not even view this system as a "real" system, and rather just view it as a toy-model of quantum mechanics. This wouldn't be entirely fair, though, because there are lots of things that we've described here that are fully true. Namely, that electrons are found in one of two spin states and that their physical description relies on the mathematics of complex superpositions in a two-dimensional complex vector space.

What we haven't described here is what we do when we change bases, since we know from our study of vector spaces that any vector space has lots of different bases available. A change of basis in this case does reflect **very** deep quantum mechanical phenomena, but a thorough description of this is out of place at the moment. Instead, let us continue to use what we have to move forward. In particular, we have enough available to us now to examine systems that have more than 2 states available to them, namely by considering systems of more than one electron. In generalizing to this case, we find a natural use of certain mathematical structures that we have yet to define, and so we'll define them in the next section. For now, the reader should be content with evolving in time in a superposition of both a) having learned a lot about quantum mechanics in a relatively rigorous way and b) knowing that although this discussion has been rigorous, it has been seriously incomplete and for a more complete discussion one should refer to a genuine quantum mechanics course and/or quantum mechanics textbook.

7.5 The 2^N -State System

We now move on to consider the case where we have several electrons lying around. Suppose, for example, that we have two electrons in a box, and that we don't care about where they are or what their momentums are—the only thing we care about is their spin states. With only two electrons around, it is relatively clear that there are four possible spin states for this two-electron system. One spin state is the state in which both electrons are in the "up" state, another spin state is when electron one is in the state "up" and electron two is in the state "down", another spin state is when electron one is in the state "down" and electron two is in the state "up", and the final (fourth) spin state is when both electrons are in the "down" spin state. Any general state of this system will be in a linear superposition of these four basis states. Our state space S is therefore four-dimensional, and so we can say it's \mathbb{C}^4 and assign to each one of these four basis states the basis vectors $(1, 0, 0, 0)$, $(0, 1, 0, 0)$, $(0, 0, 1, 0)$, and $(0, 0, 0, 1)$.

Perhaps one more example will make it clear how to generalize this to arbitrary numbers of electrons. Suppose we have three electrons in a box. There are then 8 possible spin states: (up, up, up), (up, up, down), (up, down, up), (down, up, up), (up, down, down), (down, up, down), (down, down, up), (down, down, down). Any general state of this 3-electron system will then be a linear superposition of these 8 basis states, and we have that our state space $S = \mathbb{C}^8$. We can then assign each of these basis states a basis vector in \mathbb{C}^8 and carry on with our analysis as usual.

It should be clear now that an N -electron system will have a state space of 2^N dimensions and will be linearly isomorphic to \mathbb{C}^{2^N} . It should also be clear that as N gets large, writing these vectors out explicitly will be a huge pain. For example, if we only have 5 electrons around, then our state space is $\mathbb{C}^{2^5} = \mathbb{C}^{32}$, which means each vector that we write out will have 32 components. Any equation written for this system will be an absolute mess, and each line of the equation would take about a full page just to write down.

It turns out that there's a nice way to solve this problem, and it involves introducing a more user friendly notation. Before doing so, however, we'll introduce some new abstract mathematical structures to help make clear what our new notation "really means". In particular, we'll see a couple different ways to take two vector spaces and combine them to make a new vector space. This is completely analogous to how we formed the product group of two groups. Now, we'll see that there are two obvious possibilities for constructing new vector spaces from old ones, and that one of these two possibilities is perfectly suited for what we want to do with our electron systems. It should be clear that the state space of a 5-electron system, or an N -electron system, is intimately related to the state space of a 1-electron system. Namely, we might have the feeling that our N -electron state space is in some way "built up from" our 1-electron state space. This would be the correct feeling to have, and we'll now show how to make this rigorous.

For now, we'll assume that all of our vector spaces are finite dimensional, and **either real or complex** (we don't need to specify which, because we know that the abstract theory of both are equivalent). The infinite dimensional cases are similar, but sometimes have some extra subtleties that won't help us here, and therefore we'll leave them for a proper course on those sorts of things.

The question we now seek to answer is the following: given two vector spaces V and W , can we make a new vector space using the data of these two "old" vector spaces? The answer is yes, and that there are a couple different possibilities (just as we could, in the case of sets, form the union, intersection, or Cartesian product of two "old sets" to form a "new set"). The first case we'll explore is the **direct product** of two vector spaces.

Definition 7.7. Let V and W be finite dimensional vector spaces (both real or both complex) of dimension n and m , respectively. The **direct sum** of V and W , denoted by $V \oplus W$, is the set of ordered pairs (v, w) ,

with $v \in V$ and $w \in W$. On $V \oplus W$ we define vector addition as $(v_1, w_1) + (v_2, w_2) = (v_1 + v_2, w_1 + w_2)$, and scalar multiplication as $k \cdot (v, w) = (k \cdot v, k \cdot w)$, thus turning $V \oplus W$ into a vector space.

Before we actually prove that $V \oplus W$ is a vector space, we must first note that we used the usual abuse of notation in the definition of our vector addition and scalar multiplication. Namely, on the left of those equations we're defining the "+" and "." in $V \oplus W$ in terms of the corresponding operations that are assumed to already be defined in V and W (which are the symbols used on the right hand side of those equations).

Exercise 7.8. Show that $V \oplus W$ is indeed a vector space.

The next proposition shows that the direct sum of vector spaces is **not** what we want for our multi-electron systems.

Proposition 7.9. Let V and W be finite dimensional vector spaces of respective dimension n and m . Then the dimension of $V \oplus W$ is $n + m$.

Proof: Let $\{e_1, e_2, \dots, e_n\}$ be a basis for V and $\{f_1, f_2, \dots, f_m\}$ be a basis for W . We then have that $E = \{(e_1, 0), (e_2, 0), \dots, (e_n, 0), (0, f_1), (0, f_2), \dots, (0, f_m)\}$ is a basis for $V \oplus W$. To see this, we note first that these vectors are all clearly linearly independent. Now we just need to show that any vector in $V \oplus W$ can be expressed as a linear combination of these vectors. Let $(v, w) \in V \oplus W$, then $v \in V$ and $w \in W$, so $v = a_1 e_1 + \dots + a_n e_n$ for some scalars $\{a_i\}$ and $w = b_1 f_1 + \dots + b_m f_m$ for some scalars $\{f_i\}$. We then have that

$$(v, w) = (a_1 e_1 + \dots + a_n e_n, b_1 f_1 + \dots + b_m f_m) = a_1(e_1, 0) + \dots + a_n(e_n, 0) + b_1(0, f_1) + \dots + b_m(0, f_m),$$

which is our desired linear combination. We thus have that $\text{Span}(E) = V \oplus W$ and that all of the vectors in E are linearly independent. E is therefore a basis for $V \oplus W$, and since there are $n + m$ vectors in E , our proposition is proved. \square

Exercise 7.10. Let V , W , and U be three finite dimensional vector spaces. Show that

$$(V \oplus W) \oplus U \simeq V \oplus (W \oplus U),$$

where " \simeq " means "is linearly isomorphic to". I.e., find a bijective linear map between these spaces.

We now know that the direct sum is **not** what we want to consider, because it's now clear that a 5-electron system is not the direct sum of 5 1-electron systems. This is because the direct sum of 5 1-electron systems will be $5 \times 2 = 10$ dimensional, since each 1-electron system contributes two dimensions and a direct sum simply adds up these dimensions (moreover, the above exercise shows that we can take many-fold direct sums and that they'll be equivalent up to linear isomorphism). We know that what we want is a $2^5 = 32$ -dimensional system. For this, we need to study what's called the **tensor product** of two vector spaces. There is an abstract definition of this vector space in terms of its two constituent vector spaces, but this abstract definition is very opaque and not very instructive. Therefore, we'll define this space in terms of bases for its two constituent vector spaces, as this is the best way to get an idea of what the tensor product of two vector spaces really is.

Definition 7.11. Let V and W be finite dimensional vector spaces of dimension n and m , respectively, and let $E_V = \{e_1, e_2, \dots, e_n\}$ be a basis for V and $E_W = \{f_1, f_2, \dots, f_m\}$ be a basis for W . The **tensor product** of V and W , often denoted by $V \otimes W$, is the vector space defined by the span of basis elements of the form $e_i \otimes f_j$ with $1 \leq i \leq n$ and $1 \leq j \leq m$.

In order for this definition to click, it is best to see an example.

Example 7.12. Let V be a 2-dimensional vector space with basis $\{e_1, e_2\}$ and W a 3-dimensional vector space with basis $\{f_1, f_2, f_3\}$. Then there are 6 basis elements for $V \otimes W$, and they are

$$e_1 \otimes f_1, e_1 \otimes f_2, e_1 \otimes f_3, e_2 \otimes f_1, e_2 \otimes f_2, e_2 \otimes f_3,$$

and any vector $x \in V \otimes W$ is a linear combination of the form

$$x = a_1(e_1 \otimes f_1) + a_2(e_1 \otimes f_2) + a_3(e_1 \otimes f_3) + a_4(e_2 \otimes f_1) + a_5(e_2 \otimes f_2) + a_6(e_2 \otimes f_3),$$

where each a_i is a scalar (either real or complex depending on if V and W are both real or complex (note that we can't do any of this unless both vector spaces are over the same set of scalars)). A better, more succinct notation is

$$x = \sum_{\substack{1 \leq i \leq 2 \\ 1 \leq j \leq 3}} a_{ij} e_i \otimes f_j,$$

where now each a_{ij} is a scalar (instead of the single-subscript scalars that we used above).

It is important to note that in a tensor product of vector spaces, we're viewing each $e_i \otimes f_j$ as a single basis vector. Moreover, these basis vectors are defined to be linearly independent, so that $e_i \otimes f_j$ is linearly independent from $e_k \otimes f_l$ unless both $i = k$ and $j = l$ (i.e., unless they're the same pair of vectors).

Now, to make our above definition precise, we really should have come up with a symbol like " \otimes_{E_V, E_W} " to reflect the fact that this tensor product is defined **with respect to** the bases E_V and E_W . However, as the next exercise shows, we can construct our tensor product from any pair of bases and the resulting vector spaces will be linearly isomorphic to each other, and we can therefore view them as being equivalent.

Exercise 7.13. Let V be a finite dimensional vector space and let E_V and F_V be two bases for V . Let W be a finite dimensional vector space and let E_W and F_W be two bases for W . Show that

$$V \otimes_{E_V, E_W} W \simeq V \otimes_{F_V, F_W} W.$$

Due to the above exercise, we can simply view the tensor product of two vector spaces V and W with respect to any basis as "the" tensor product, and not worry about which bases we chose to construct it.

We now have a way of constructing a new vector space from two old ones in such a way that the new vector space will be of dimension nm when the two old vector spaces are of dimension n and m , respectively. The following exercise will now be the final step in showing that the tensor product is precisely the vector space construction that we want for our multi-electron systems.

Exercise 7.14. Let V , W , and U be three finite dimensional vector spaces. Show that

$$(V \otimes W) \otimes U \simeq V \otimes (W \otimes U).$$

The above exercise shows that, up to linear isomorphism, we can consider n -fold tensor products of vector spaces. It is then clear that this is precisely what we want for constructing our state space of a multi-electron system. Namely, for every electron in our system, we simply add another factor of the individual electron state space $S = \mathbb{C}^2$ into our tensor product. Thus, a 2-electron system's states space will be $S \otimes S = \mathbb{C}^2 \otimes \mathbb{C}^2$, which is a 2×2 dimensional vector space. Similarly, a 3-electron system's state space will be $S \otimes S \otimes S = \mathbb{C}^2 \otimes \mathbb{C}^2 \otimes \mathbb{C}^2$, which is a $2 \times 2 \times 2 = 8$ dimensional vector space, and is therefore linearly isomorphic to \mathbb{C}^8 , just as we'd expect. We can then see that an n -electrons system's state space

will be the n -fold tensor product of $S = \mathbb{C}^2$, which will be of dimension $2 \times 2 \times \dots \times 2$ (where there are n total factors of 2), and is therefore linearly isomorphic to \mathbb{C}^{2^n} , just as we'd hoped for.

Now, in order to do calculations within these many-fold tensor product vector spaces, we need to introduce some new notation. Let's reconsider our single-electron system with the two dimensional state space $S = \mathbb{C}^2$. Instead of writing

$$e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

let's instead write

$$\uparrow = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \downarrow = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

This is nothing but a change in notation, but it should be clear what the relationship is between these basis vectors and the states "up" and "down". Now let's consider the two-electron system with state space $S \otimes S \simeq \mathbb{C}^4$, which would have the four basis vectors

$$e_1 \otimes f_1, e_1 \otimes f_2, e_2 \otimes f_1, e_2 \otimes f_2,$$

where we let

$$f_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, f_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

be the basis vectors for the "second copy" of S . We could then assign these basis vectors to basis vectors of \mathbb{C}^4 as follows:

$$e_1 \otimes f_1 = \begin{pmatrix} 1 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \quad e_1 \otimes f_2 = \begin{pmatrix} 0 \\ 1 \\ 0 \\ 0 \end{pmatrix}, \quad e_2 \otimes f_1 = \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}, \quad e_2 \otimes f_2 = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}.$$

But this is really tedious. After all, what we really need to keep track of is the **coefficients** of the various basis vectors (and not the "inner workings" of the basis vectors themselves) since these are what tell us the probabilities of collapsing the state into the corresponding basis state upon measurement. We must note that this is completely independent of how we choose to write down our basis vectors, so if we can find a better way to do so we might save ourselves a lot of trouble. It turns out that one way of making such an improvement is to simply extend our " \uparrow " and " \downarrow " notation to tensor products. We can do this in a completely straightforward way—all we do is assign " \uparrow " to each basis vector with a subscript 1, and " \downarrow " to each basis vector with a subscript 2. We thus have

$$e_1 \otimes f_1 = \uparrow\uparrow, \quad e_1 \otimes f_2 = \uparrow\downarrow, \quad e_2 \otimes f_1 = \downarrow\uparrow, \quad e_2 \otimes f_2 = \downarrow\downarrow,$$

and we get the added benefit that this notation tells us which electrons are up and which are down. So once we become comfortable viewing an expression like " $\uparrow\downarrow$ " as a **vector**, so that $a \uparrow\uparrow + b \uparrow\downarrow = (a + b) \uparrow\uparrow$, etc., then we can simply manipulate these much more transparent expressions. We should also quickly note the importance of order here, since $\uparrow\downarrow \neq \downarrow\uparrow$, which is simply a reflection of the fact that $e_1 \otimes f_2 \neq e_2 \otimes f_1$.

We can then easily generalize this to n -fold tensor products as follows. Suppose we have n -electrons. Then the 2^n basis state vectors will be represented as all possible expressions of the form " $\uparrow\uparrow\downarrow \dots \downarrow\uparrow$ " (where order still matters here, in that an expression with the same number of up and down arrows, in a different order, are inequivalent expressions), where there are a total of n arrows and each one is either up or down. It is easy to convince oneself that there will be 2^n such expressions, and therefore there will be 2^n basis state vectors since each such expression is **defined to be** linearly independent from every

other distinct expression. We then scalar multiply such expressions "component-wise", so that for example $a \cdot (\uparrow\uparrow\downarrow\uparrow + \downarrow\uparrow\downarrow\downarrow) = a \uparrow\uparrow\downarrow\uparrow + a \downarrow\uparrow\downarrow\downarrow$, and vector addition is also done component-wise, so that $\uparrow\uparrow + \uparrow\uparrow = 2 \uparrow\uparrow$.

The last thing we need to do is come to an agreement about which arrow corresponds to which electron, so let us say that "electron 1" will correspond to the left-most arrow, and "electron n " will correspond to the right-most arrow, and the intervening electrons are labeled accordingly. With this agreement, the following exercise becomes possible.

Exercise 7.15. Suppose we have a 6-electron system and that our system is in the state

$$v = \frac{1}{\sqrt{3}} \uparrow\uparrow\downarrow\uparrow\uparrow + \frac{1}{\sqrt{3}} \uparrow\uparrow\uparrow\downarrow\downarrow + \frac{1}{\sqrt{3}} \uparrow\uparrow\uparrow\uparrow\uparrow.$$

What is the probability that a measurement of our system will yield a state with the third electron in the "down" state? What about in the "up" state? What is the probability that a measurement will yield a state with the fifth electron in the "up" state? What is the probability that a measurement will yield a state with the first electron in the "down" state?

Now suppose our system is in the state

$$w = \frac{2}{3} \uparrow\uparrow\uparrow\uparrow\uparrow + \uparrow\uparrow\uparrow\downarrow\downarrow + \frac{1}{3} \uparrow\uparrow\uparrow\uparrow\downarrow + \downarrow\downarrow\downarrow\downarrow\downarrow.$$

What is the probability of a measurement yielding a state in which electron 4 is in the "down" state? Hint: first check if this state is normalized, and if not, normalize it.

And that's basically it. This is quantum mechanics. Again, this is some of the simplest quantum mechanics that one can do—and we've swept a lot of details under the rug—so please don't doubt that there is still **lots** to learn about this marvelous yet whacky world. To end this chapter, and this text, we'll look at quite possibly the weirdest consequence of quantum mechanics that Nature has in store for us. Namely, we'll take a look at the phenomenon of entanglement. Entanglement is still, to this day, quite possibly the most mind-boggling and philosophically strange phenomenon in all of physics. Believe it or not, we already have enough mathematical machinery to describe it with a decent amount of rigor. Therefore, it is a surprisingly simple concept, yet a deep understanding of it is still far from us, if not impossible to attain. Let us take a quick look.

7.6 Entanglement

We end this text with a discussion of entanglement, and use this phenomenon as a sort of climax to all the work we've put in so far. We've developed a good deal of mathematical and physical knowledge about our world, and we can now understand one of the weirdest aspects of it with a good deal of mathematical sophistication. Let's dive right in.

Suppose we have 2 single-electron systems. Namely, even though there are two electrons present, we're going to view each electron as a separate system. Then for each electron we have a state space S that is isomorphic to \mathbb{C}^2 . One electron is then in the state

$$v_1 = a \uparrow + b \downarrow \in S_1 = \mathbb{C}^2$$

where S_1 is the first electron's state space (and where we assume v is properly normalized so that $|a|^2 + |b|^2 = 1$), and the other electron is in the state

$$v_2 = c \uparrow + d \downarrow \in S_2 = \mathbb{C}^2$$

where S_2 is the second electron's state space and again v_2 is properly normalized. We then found in the previous section that if we want to view these two electrons as a **single** 2-electron system, then we simply need to take the tensor product of their state spaces: $S_1 \otimes S_2 \simeq \mathbb{C}^4$.

When we view these two electrons as part of a single 2-electron system, we have that their state in $S_1 \otimes S_2$ is nothing but $v_1 \otimes v_2$, which has the expansion in terms of basis vectors as follows:

$$v_1 \otimes v_2 = (a \uparrow + b \downarrow) \otimes (c \uparrow + d \downarrow) = ac \uparrow\uparrow + ad \uparrow\downarrow + bc \downarrow\uparrow + bd \downarrow\downarrow$$

where we used the fact that "the tensor product of a sum of vectors is the same as the sum of the tensor product of vectors" (we haven't proved this fact, but some thought will either prove that this is true or prove that this is the only reasonable thing to do, or both (either way, it **is** the right thing to do)). We then have a set of probabilities for finding the four different possible configurations of these electrons in any given measurement, and these probabilities are dependent on the state of the **individual** electrons. We call a state $s \in S_1 \otimes S_2$ that can be "split up" in this way as a tensor product of two "individual" states a **separable** state. The fact that we can split a state up into a tensor product of single-electron states implies that we can view each individual electron as exactly that—individual electrons.

What is perhaps surprising, and immensely profound, is that there are (lots of) states in $S_1 \otimes S_2$ that **can't** be split up in this way. In other words, there are states of the combined system that **can't** be split up into tensor products of states "living in" the individual state spaces. We call such states **entangled states**. So a state is either separable or entangled—never both, and never neither. We'll first show that entangled states exist, and then we'll talk about the consequences of their existence.

Proposition 7.16. The state $\frac{1}{\sqrt{2}}(\uparrow\uparrow + \downarrow\downarrow) \in S_1 \otimes S_2$ is an entangled state.

Proof: We first note that the factor of $\frac{1}{\sqrt{2}}$ that is out front is there simply to make the state normalized (as you can check). Now, to prove that this state is entangled, we need to prove that it's **not** separable. To prove that it's not separable, we use our most powerful tool for proving negative statements—proof by contradiction. I.e., we'll suppose the state **is** separable, and find a contradiction.

If the above state is separable, then that means that there are two single-electron states $a \uparrow + b \downarrow$ and $c \uparrow + d \downarrow$ such that

$$\frac{1}{\sqrt{2}}(\uparrow\uparrow + \downarrow\downarrow) = (a \uparrow + b \downarrow) \otimes (c \uparrow + d \downarrow) = ac \uparrow\uparrow + ad \uparrow\downarrow + bc \downarrow\uparrow + bd \downarrow\downarrow.$$

But the state $\frac{1}{\sqrt{2}}(\uparrow\uparrow + \downarrow\downarrow)$ has no " $\uparrow\downarrow$ " term, and no " $\downarrow\uparrow$ " term. This means that $ad = 0$ and $bc = 0$. Now, $ad = 0$ implies that either $a = 0$ or $d = 0$ (since two non-zero complex numbers multiply to a non-zero complex number). But if $a = 0$, then there would be no " $\uparrow\uparrow$ " term, and we clearly need that term. Similarly, if $d = 0$, then there'd be no " $\downarrow\downarrow$ " term, and we clearly also need that term. Therefore we already have enough information to show that such a splitting of the state $\frac{1}{\sqrt{2}}(\uparrow\uparrow + \downarrow\downarrow)$ is impossible, but just to really put this proposition to bed we'll note that we could also get this result by seeing that $bc = 0 \Rightarrow b = 0$ or $c = 0$, and both of these cases also give impossibilities. \square

The simplicity of the above proof may make it likely that the reader has not fully absorbed the fact that we've just proved what is in my mind the most significant result of this entire text. On the surface it seems completely harmless. Proving it really only depends on the fact that the product of two non-zero complex numbers can never equal zero, yet this is one of the most significant physical phenomena—with some of the deepest philosophical consequences—that mankind has yet to discover.

So what's the big deal? The big deal is that we have a two-electron system that **cannot** be viewed as

two one-electron systems. Suppose we put our two electrons in a box this size of our solar system, and suppose we prepare them in the state $\frac{1}{\sqrt{2}}(\uparrow\uparrow + \downarrow\downarrow)$. (To date, there has been absolutely no experimental evidence hinting towards the possibility that this would be impossible, and there has been **tons** of experimental evidence proving that entangled states can be constructed.) If we've set up our two electrons in this state, then it turns out that they will actually behave as a **single** system. In other words, we must view these two particles as **one combined indivisible system**. To see this, we see that if we measure one electron, then we **automatically** know **with certainty** which state the other electron is in. For if we "open the box" and see one electron in the state \uparrow , then we know that the other electron is also in \uparrow simply because there is a zero probability of finding the $\uparrow\downarrow$ or $\downarrow\uparrow$ states. Similarly, if we find one electron in the \downarrow state, then we know **for sure** that the other electron is in the \downarrow state as well.

The point is that we simply cannot view these two particles as independent entities. They are not two particles moving about in the universe however they wish, but rather **one indivisible system** that knows about the other's behavior **at all times**, and the observation of one immediately determines the result of an observation on the other, no matter how far apart they are physically. There is absolutely no classical analogue to the entangled state. As of today, there is no deeper explanation as to **why** entanglement exists in the universe, and there has been **substantial** evidence that entanglement **does** exist. Namely, experimentalists have found effects of entanglement between particles even when they're hundreds of kilometers away from each other! Think about this. This means that two particles—separated by several kilometers (or more)—**must** be viewed as a **single** system.

Philosophers have had a blast exploring the consequences of entanglement, and volumes have been written about this for the interested reader. For a more pragmatic reader, it turns out that entanglement is actually **completely vital** for various quantum algorithms in the theory of quantum computation—an extremely fascinating field that explores how to take advantage of the weird properties of quantum mechanics to find algorithms that are either completely new, or exponentially faster than their classical counterparts.

Entangled states are in many ways what defines quantum mechanics as a completely distinct theory from classical mechanics. Nowhere else in physics do we see objects that can be so distantly separated in space yet still behave as a single system. We tend to think of physics as being a purely "local" thing, meaning that what happens "here" is only determined by what happens in the parts of space that are extremely close to "here". But quantum entanglement shows that this is simply not how the universe behaves at its most fundamental level. Once again, Nature proves to us that she is more subtle, more sneaky, more devious, and more beautiful than we can possibly imagine.

7.7 Concluding Remarks

We have now reached the end (for now) of the mathematical and physical concepts that we will develop. As I've tried to make clear throughout this text, we have moved through this material in such a way as to only introduce the bare minimum necessary to gain an understanding of the underlying concepts. Therefore, in absolutely **every single** topic that we have covered, there is an entire lifetime's worth of further study that could be enjoyed. There are also infinitely many other mathematical and physical concepts that one can study that we have yet to even mention, and that hold gorgeous and profound secrets for all those willing to find them.

What we hope these notes have provided is a brief tour through some of the big ideas in math and physics,

with enough depth to appreciate their rigorous and abstract existence, in order to inspire the reader to continue his or her studies in these amazing fields. Mathematics is the art form that creates beauty from the mediums of abstraction and logic. It may seem perhaps miraculous that Nature has chosen this art form as the language of her behavior. So far, throughout the course of human history, whenever we challenge her with a complete description of her behavior using her own language of mathematics, she surprises us right back with more unforeseen behavior.

It is hoped that these notes have inspired the reader to find some of the joys that come along with learning the language of abstraction, and/or with seeing how that language can be applied towards understanding the universe that we find ourselves in (and possibly even those universes that we **don't** find ourselves in). To the reader who plans on going forward, you still have lots to look forward. Not only will you learn about all of these subjects in much greater depth, but you can (will?) also learn about a) geometries of curved spaces of various dimensions, b) number systems that you couldn't possibly imagine existed (and that make complex numbers look like child's play), c) structures so abstract that the entirety of all mathematics that can ever exist can lie within just one instance of the structure, d) physical theories of extra dimensions of space-time, e) new formulations of the space-time that we already know and love, f) particles that exist for mind-blowingly short amounts of time, but that shape the dynamics of the universe that we experience, g) fluids that practically jump out of their containers, h) black holes evaporating, and so, so much more. This is only the very beginning!

Further Reading

For more on set theory:

"Naive Set Theory" by Paul R. Halmos.

For more on group theory:

"Abstract Algebra" by David Dummit and Richard Foote

"Algebra" by Michael Artin

For more on vector spaces (the theory of which is called "linear algebra"):

"Linear Algebra and Its Applications" by David Lay

For more on complex numbers/analysis:

"Visual Complex Analysis" by Tristan Needham

For more on classical physics:

"Fundamentals of Physics Extended" by David Halliday, Robert Resnick, and Jearly Walker

"Classical Mechanics" by John Taylor

"Introduction to Electrodynamics" by David Griffiths

For more on relativistic physics:

"Introduction to Special Relativity" by Wolfgang Rindler

"A First Course in General Relativity" by Bernard Schutz

For more on quantum physics:

"Introduction to Quantum Mechanics" by David Griffiths

"Quantum Physics of Atoms, Molecules, Solids, Nuclei, and Particles" by Robert Eisberg and Robert Resnick.