

Chapter 2

Mathematical Induction: “And so on . . .”

2.1 Introduction

This chapter marks our first big step toward investigating mathematical proofs more thoroughly and learning to construct our own. It is also an introduction to the first significant **proof technique** we will see. As we describe below, this chapter is meant to be an appetizer, a first taste, of what **mathematical induction** is and how to use it. A couple of chapters from now, we will be able to rigorously define induction and *prove* that this technique is mathematically valid. That’s right, we’ll actually prove how and why it works! For now, though, we’ll continue our investigation of some interesting mathematical puzzles, with these particular problems hand-picked by us for their use of inductive techniques.

2.1.1 Objectives

The following short sections in this introduction will show you how this chapter fits into the scheme of the book. They will describe how our previous work will be helpful, they will motivate why we would care to investigate the topics that appear in this chapter, and they will tell you our goals and what you should keep in mind while reading along to achieve those goals. Right now, we will summarize the main objectives of this chapter for you via a series of statements. These describe the skills and knowledge you should have gained by the conclusion of this chapter. The following sections will reiterate these ideas in more detail, but this will provide you with a brief list for future reference. When you finish working through this chapter, return to this list and see if you understand all of these objectives. Do you see why we outlined them here as being important? Can you define all the terminology we use? Can you apply the techniques we describe?

By the end of this chapter, you should be able to . . .

- Define what an inductive argument is, as well as classify a presented argument as an inductive one or not.
- Decide when to use an inductive argument, depending on the structure of the problem you are solving.
- Heuristically describe mathematical induction via an analogy.
- Identify and describe different kinds of inductive arguments by comparing and contrasting them, as well as identify the underlying structures of the corresponding problems that would yield these similarities and differences.

2.1.2 Segue from previous chapter

As in the previous chapter, we won't assume any familiarity with more advanced mathematics beyond basic algebra and arithmetic, and perhaps some visual, geometric intuition. We will, however, make use of summation and product notation fairly often, so if you feel like your notational skills are, go back and review Section 1.3.5.

2.1.3 Motivation

Look back at the Puzzle in Section 1.4.3, where we proved that the sum of the first n odd natural numbers is exactly n^2 . We first observed this pattern geometrically, by arranging the terms of the sums (odd integers) as successively larger “corner pieces” of a square. The first way we proved the result, though, didn't seem to depend on that observation and merely utilized a previous result (about sums of even *and* odd integers) in an *algebraic* way; that is, we did some tricky manipulation of some equations (multiplying and subtracting and what have you) and then—voilà!—out popped the result we expected. What did you think about that approach? Did it feel satisfying? In a way, it didn't quite match the geometric interpretation we had, at first, so it might be surprising that it worked out so nicely. (Perhaps there is a *different* geometric interpretation of this approach. Can you find one?)

Our second approach was to model that initial geometric observation. We transformed visual pieces into algebraic pieces; specifically, a sum was related to the area of a square, and the terms of the sum were related to particular pieces of that square. We established a *correspondence* between different interpretations of the same problem, finding a way to relate one to the other so that we could work with either interpretation and know that we were proving something about the overall result. The benefit of the visual interpretation is that it allowed us to take advantage of a general proof strategy known as **mathematical induction**, or sometimes just **induction**, for short. (The word *induction* has some non-mathematical meanings, as well, such as in electromagnetism or in philosophical arguments, but within the context of this book, when we say *induction*, we mean

mathematical induction.) What exactly is induction? How does it work? When can we use this strategy? How do we adapt the strategy to a particular puzzle? Are there variations of the strategy that are more useful in certain situations? These are all questions that we hope to answer in this chapter.

The first topic we'd like to address is a question that we didn't just ask in the last paragraph, namely, "*Why* induction? *Why* bother with it?" Based on that puzzle in Section 1.4.3, it would seem that mathematical induction isn't entirely necessary since there might be other ways of proving something, instead of by induction. Depending on the context, this very well may be true, but the point we'd like to make clear from the beginning is that *induction is incredibly useful!* There are many situations where a proof by induction is the most concise and clear approach, and it is a well-known general strategy that can be applied in a variety of such situations. Furthermore, applying induction to a problem requires there to be a certain *structure* to the problem, a dependence of one "part" of the result on a "previous part". (The "parts" and the "dependence" will depend on the context, of course.) Recognizing that induction applies, and actually going through the subsequent proof process, will usually *teach* us something about the inherent structure of the problem. This is true even when induction fails! Perhaps there's a particular part of a problem that "ruins" the induction process, and identifying that particular part can be helpful and insightful.

We hope to motivate these points through some illustrative examples first, after which we will provide a reasonably thorough *definition* of mathematical induction that will show how the method works, in generality. (A completely *rigorous* definition will have to be put off until a little bit later, after we have defined and investigated some relevant concepts, like set theory and logical statements and implications. For now, though, the definition we give will suffice to work on some interesting puzzles and allow us to discuss induction as a general proof strategy.)

2.1.4 Goals and Warnings for the Reader

Do keep in mind that we are still building towards our goal of mathematical rigor, or as much as is possible within the scope and timing of this book and course. Some of the claims we make in this Chapter will be clarified and technically proven later on, once we have properly discussed the natural numbers and some basic mathematical logic. All in due time!

That said, this chapter is still very important, since we are continuing to introduce you to the process of solving mathematical problems, applying our existing knowledge and techniques to discover new facts and explain them to others. In addition, mathematical induction is a fundamental proof technique that will likely appear in every other mathematics course you take! This is because of its usefulness and the prevalence of inductive properties throughout the mathematical world.

2.2 Examples and Discussion


2.2.1 Turning Cubes Into Bigger Cubes

To motivate the overall method of mathematical induction, let's examine a geometric puzzle and solve it together. This example has been chosen carefully to illustrate how **mathematical induction** is relevant when a puzzle has a particular type of structure; specifically, some truth or fact or observation *depends* or *relies* or can be *derived* from a “previous” fact. This dependence on a previous case (or cases) is what makes a process *inductive*, and when we observe this phenomenon, applying *induction* is almost always a good idea.

1-Cube into a 2-Cube


Let's examine cubic numbers and, specifically, let's try to describe a cubic number *in terms of* the previous cubic number. Imagine a $1 \times 1 \times 1$ cube, just one building block. How can we build the “next biggest” cube, of size $2 \times 2 \times 2$, by adding $1 \times 1 \times 1$ building blocks? How many do we need to add? Arithmetically, we know the answer: $2^3 = 8$ and $1^3 = 1$, so we need to add 7 blocks to have the correct volume. Okay, that's a specific answer, but it doesn't quite tell us *how* to arrange those 7 blocks to make a cube, nor does it give us any insight into how to answer this question for *larger* cubes. Ultimately, we would like to say how many blocks are required to build a $100 \times 100 \times 100$ cube into a $101 \times 101 \times 101$ cube without having to perform a lot of tedious arithmetic; that is, we are hoping to eventually find an answer to the question: given an $n \times n \times n$ cube, how many blocks must we add to build it into a $(n+1) \times (n+1) \times (n+1)$ cube? With that in mind, let's think carefully about this initial case and try to answer it with a general argument.

Given that single building block, and knowing we have to add 7 blocks to it, let's try to identify exactly where those 7 blocks should be placed to make a $2 \times 2 \times 2$ cube. (For simplicity, we will refer to a cube of size $n \times n \times n$ as an n -cube, for any value of n . We will only need to use values of n that are *natural numbers*, i.e. non-negative whole numbers, in this example.) Look at the pictures of the 1-cube and 2-cube below and try to come up with an explanation of constructing one from the other.




Here's one reasonable explanation that we want to use because it will guide us in the general explanation of building an $(n+1)$ -cube from an n -cube, and because


it is a mathematically *elegant* and simple explanation. Start with the 1-cube positioned as it is above and “enlarge” the 3 exposed faces by the appropriate amount, in this case by one block. This accounts for 3 of the 7 blocks, thus far: $2^3 = 1^3 + 3 + \underline{\hspace{1cm}}$. Where are there “holes” now?



The blocks we just added have created “gaps” between each pair of them, and each of those “gaps” can be filled with one block. This accounts for 3 more of the 7 total blocks: $2^3 = 1^3 + 3 + 3 + \underline{\hspace{1cm}}$. Now what?



There is just one block left to be filled, and it’s the very top corner. Adding this block completes the 2-cube and tells us how to mathematically describe our construction process with the following picture and equation:




$$2^3 = 1 + \textcolor{blue}{3} + \textcolor{red}{3} + \textcolor{green}{1}$$

2-Cube into a 3-Cube


Okay, we might now have a better idea of how to describe this process in general, but let’s examine another case or two just to make sure we have the full idea.

Let's start with a 2-cube and construct a 3-cube from it. (You can even try this out by hand if you happen to own various sizes of Rubik's Cubes!) We can follow a process similar to the steps we used in the previous case and just change the numbers appropriately. Starting with a similar picture




we see that we need to “enlarge” the three exposed faces of the 2-cube but, in this case, the amount by which we need to enlarge them is *different* than before (with the 1-cube) since we are working with a larger initial cube. Specifically, each face must be enlarged by a 2×2 *square* of blocks (whereas, in the previous case, we added a 1×1 square of blocks). Thus, an equation to account for this addition is

$$3^3 = 2^3 + 3 \cdot 2^2 + \underline{\hspace{1cm}}$$




After we do this, we see that we need to fill in the gaps between those enlarged faces with 2×1 of blocks (whereas, in the previous case, we added 1×1 rows of blocks). An equation to account for the additions thus far is

$$3^3 = 2^3 + 3 \cdot 2^2 + 3 \cdot 2 + \underline{\hspace{1cm}}$$




After we do this, we see that there is only the top corner left to fill in. Accordingly, we can depict our construction process and its corresponding equation:



$$3^3 = 2^3 + 3 \cdot 2^2 + 3 \cdot 2 + 1$$


n -Cube into an $(n + 1)$ -Cube

Do you see now how this process will generalize? What if we started with an n -cube? How could we construct an $(n + 1)$ -cube? Let's follow the same steps we used in the previous two cases. First, we would enlarge the three exposed faces by adding three *squares* of blocks. How big is each square? Well, we want each square to be the same size as the exposed faces, so they will be $n \times n$ squares, accounting for n^2 blocks for each face:



$$(n+1)^3 = n^3 + 3n^2 + \underline{\hspace{1cm}}$$

Next, we would fill in the gaps between these enlarged faces with rows of blocks. How long are those rows? Well, they each lie along the edges of the squares of blocks we just added, so they will each be of size $n \times 1$, accounting for n blocks for each gap:



$$(n+1)^3 = n^3 + 3n^2 + 3n + \underline{\hspace{1cm}}$$

Finally, there will only be the top corner left to fill in! Therefore,

$$(n+1)^3 = n^3 + 3n^2 + 3n + 1$$

“Wait a minute!” you might say, abruptly. “We already knew that, right?” In a way, yes; the equation above is an algebraic identity that we can also easily see by just expanding the product on the left and collecting terms:

$$\begin{aligned} (n+1)^3 &= (n+1) \cdot (n+1)^2 \\ &= (n+1) \cdot (n^2 + 2n + 1) \\ &= (n^3 + 2n^2 + n) + (n^2 + 2n + 1) \\ &= n^3 + 3n^2 + 3n + 1 \end{aligned}$$

So what have we really accomplished? Well, the main point behind deriving this identity in this geometric and visual way is that it exhibits how this identity represents some kind of *inductive* process. We sought to explain how to derive one “fact” (a cubic number, $(n+1)^3$) from a previously known “fact” (the next smallest cubic number, n^3) and properly explained how to do just that. Compare this to one of the methods we used to investigate the fact that the sums of odd integers yield perfect squares, too. That observation also belies an

inductive process and, although we didn't describe it as such at the time, we encourage you to think about that now. Look back at our discussion and try to write out how you could write $(n+1)^2$ in terms of n^2 by looking at squares of blocks. Does it look anything like an "obvious" algebraic identity? (If you're feeling ambitious, think about what happens with writing $(n+1)^4$ in terms of n^4 . Is there any geometric intuition behind this? What about higher powers?)

The benefit of the method we've used is that we now know how to describe cubic numbers in terms of smaller cubic numbers, *all the way down to 1*; that is, any time we see a cubic number in an expression, we know precisely how to write that value in terms of a smaller cubic number and some leftover terms. Furthermore, each of those expressions and leftover terms have an inherent *structure* to them that depends on the cubic number in question. Thus, by iteratively replacing any cubic number, like $(n+1)^3$, with an expression like the one we derived above, and continuing until we can't replace any more, should produce an equation that has some built in *symmetry*. This idea is best illustrated by actually doing it, so let's see what happens. Let's start with the expression we derived, for some arbitrary value of n ,

$$(n+1)^3 = n^3 + 3n^2 + 3n + 1$$

and then recognize that we now know a similar expression

$$n^3 = (n-1)^3 + 3(n-1)^2 + 3(n-1) + 1$$

We proved that this equation holds when we gave a general argument for the expression above for n^3 , since that only relied on the fact that $n \geq 1$. We can follow the same logical steps, throughout replacing n with $n-1$, and end up with the second expression above, for $(n-1)^3$. (Does this keep going, for *any* value of n ? Think about this for a minute. Does our argument make any sense when $n \leq 0$? Would it make physical sense to talk about, say, constructing a $(-2) \times (-2) \times (-2)$ cube from a different cube?)

Therefore, we can replace the n^3 term in the line above

$$\begin{array}{ccccccc} (n+1)^3 = & \cancel{n^3} & + & 3n^2 & + & 3n & + & 1 \\ & + & (n-1)^3 & + & 3(n-1)^2 & + & 3(n-1) & + & 1 \end{array}$$

This is also an algebraic identity, but it's certainly not one that we would easily think to write down just by expanding the product on the left-hand side and grouping terms. Here, we are taking advantage of the *structure* of our result to apply it over and over and obtain new expressions that we wouldn't have otherwise thought to write down. Let's continue with this substitution process and see where it takes us! Next, we replace $(n-1)^3$ with the corresponding expression and find

$$\begin{array}{ccccccc} (n+1)^3 = & & 3n^2 & + & 3n & + & 1 \\ & \cancel{(n-1)^3} & + & 3(n-1)^2 & + & 3(n-1) & + & 1 \\ & + & (n-2)^3 & + & 3(n-2)^2 & + & 3(n-2) & + & 1 \end{array}$$

Perhaps you see where this is going? We can do this substitution process over and over, and the columns that we've arranged above will continue to grow, showing us that there is something deep and mathematically symmetric going on here. But where does this process stop? We want to write down a concise version of this iterative process and be able to explain all of the terms that arise, so we need to know where it ends. Remember the very first step in our investigation of the cubic numbers? We figured out how to write $2^3 = 1^3 + 3 + 3 + 1$. Since this was our *first* step in *building* this inductive process, it should be the *last* step we apply when building backwards, as we are now. Accordingly, we can write

$$\begin{array}{rcccccl}
 (n+1)^3 = & & 3n^2 & + & 3n & + & 1 \\
 & + & 3(n-1)^2 & + & 3(n-1) & + & 1 \\
 & + & 3(n-2)^2 & + & 3(n-2) & + & 1 \\
 & + & 3(n-3)^2 & + & 3(n-3) & + & 1 \\
 & & \vdots & + & \vdots & + & \vdots \\
 & + & 3 \cdot 2^2 & + & 3 \cdot 2 & + & 1 \\
 + 1^3 & + & 3 \cdot 1^2 & + & 3 \cdot 1 & + & 1
 \end{array}$$

This is *definitely* an identity we wouldn't have come up with off the top of our heads! In addition to being relatively pretty-looking on the page like this, it also allows us to apply some of our previous knowledge and simplify the expression. To see how we can do that, let's apply summation notation to the columns above and collect a bunch of terms into some simple expressions:

$$(n+1)^3 = 1^3 + 3 \cdot \sum_{k=1}^n k^2 + 3 \cdot \sum_{k=1}^n k + \sum_{k=1}^n 1$$

In the last chapter, we saw a couple of different proofs that told us

$$\sum_{k=1}^n k = \frac{n(n+1)}{2}$$

Let's use that fact in the line above, and also simplify the term on the far right, to write

$$(n+1)^3 = 1 + 3 \cdot \sum_{k=1}^n k^2 + \frac{3n(n+1)}{2} + n$$

What does this tell us? What have we accomplished after all this algebraic manipulation? Well, we previously proved a result about the sum of the first n natural numbers, so a natural question to ask after that is: What is the sum of the first n natural numbers *squared*? How could we begin to answer that? That's a trick question, because *we already have*! Let's do one or two more algebraic steps with the equation above by isolating the summation term and

then dividing:

$$(n+1)^3 - 1 - n - \frac{3n(n+1)}{2} = 3 \cdot \sum_{k=1}^n k^2$$

$$\frac{1}{3}(n+1)^3 - \frac{1}{3}(n+1) - \frac{n(n+1)}{2} = \sum_{k=1}^n k^2$$

This is what we've accomplished: we've derived a formula for the sum of the first n square natural numbers! Of course, the expression on the left in the line above isn't particularly nice looking and we could perform some further simplification, and we will leave it to you to verify that this yields the expression below:

$$\sum_{k=1}^n k^2 = \frac{1}{6}n(n+1)(2n+1)$$

“And so on” is not rigorous!

There are a couple of “morals” that we'd like to point out, based on all of this work. The first moral is that generalizing an argument is a good method for discovering new and interesting mathematical ideas and results. Did you think about how this puzzle is related to the sums of odd natural numbers? If not, we encourage you strongly to try that now, as well as think about generalizing this even further to four or five dimensional “cubes” and so on. In addition to giving you some other interesting results, it will also be incredibly instructive for learning to think abstractly and apply inductive processes. The second moral is more like an admission: we have *not* technically *proven* the formula above for the sum of the first n square natural numbers. It seems like our derivation is valid and tells us the “correct answer” but there is a glaring issue: ellipses! In expanding the equation for $(n+1)^3$ and obtaining those columns of terms that we collected into particular sums, writing \vdots in the middle of those columns was helpful in guiding our intuition, but *this is not a mathematically rigorous technique*. How do we *know* that all of the terms in the middle are exactly what we'd expect them to be? How can we be so sure that all of our pictures of cubes translate perfectly into the mathematical expressions we wrote down? What do we really mean by “and keep going all the way down to 1”?

As an example, consider this:

$$1, 2, 3, 4, \dots, 100$$

What is that list of numbers? You probably interpreted it as “all the natural numbers between 1 and 100, inclusive”. That seems reasonable. But what if we *actually* meant this list?

$$1, 2, 3, 4, 7, 10, 11, 12, 14, \dots, 100$$

Why, of course, we meant the list of natural numbers from 1 to 100 that don't have an “i” in their English spelling! Wasn't it obvious?

The point is this: when talking with a friend, and *verbalizing* some ideas, it might be okay to write “1, 2, 3, . . . , 100” and ensure that whoever is listening knows *exactly* what you mean. In general, though, we can’t assume that a reader would just naturally *intuit* whatever we were trying to convey; we should be as *explicit* and *rigorous* as possible.

It may seem to you now like we’re nit-picking, but the larger point is that there is a mathematical way of making this argument more *precise*, so that it constitutes a completely valid *proof*. Everything we have done so far is useful in guiding our intuition, but we will have to do a little more work to be sure our arguments are completely convincing. There are a few other concepts required to make this type of argument rigorous, in general, and we will investigate those in the next chapter and return to this subject immediately after that. However, in the meantime, let’s examine one more example to practice this intuitive argument style and recognizing when induction is an applicable technique.


2.2.2 Lines On The Plane

Take a clean sheet of paper and a pen and a ruler. How many regions are on your sheet? Just one, right? Draw a line all the way across the paper. Now there are two regions. Draw another straight line that intersects your first line. How many regions are there? You should count four in total. Draw a third line that intersects both of the first two, but *not* at the point where the first two intersect. (That is, there should be three intersection points, in total.) How many regions are there? Can you predict the answer before counting? What happens when there are 4 lines? Or 5? Or 100? How do we approach this puzzle and, ultimately, solve it? Let’s give a more formal statement to be sure we’re thinking the same way:

Consider n lines on an infinite plane (two-dimensional surface) such that no two lines are *parallel* and no more than two lines *intersect* at one point. How many distinct regions do the lines create?


We can draw a few examples by hand when n is small (up to, say, $n = 5$ is reasonable), and let’s use this to guide our intuition into making a general argument for an *arbitrary* value of n . (Notice that this strategy is very similar to what we did in the previous puzzle: identify a pattern with small cases, identify the relevant components of those cases that can generalize, then abstract to an arbitrary case.) Specifically, we want to attempt to identify how the number of regions in one drawing *depends* on the number of regions in a drawing with fewer lines. What happens when we draw a new line? Can we determine how this changes the already existing regions? Can we somehow count how many regions this creates? Do some investigation of this puzzle on your own before reading on. If you figure out some results, compare your work to the steps we follow below.

Let’s start with a small case, say $n = 2$. We know one line divides a plane into 2 regions; what happens when we add a second line? We know we get 4 regions, because we can just look and count them:




However, we are only looking at *one specific case* of two intersecting lines. How do we know that we will *always* find four regions, no matter how we draw those two lines appropriately? That is, can we describe *how* this happens in a way that somehow incorporates the fact that the number of lines is $n = 2$? Think about it!

Here's our approach. Notice that each of the already existing regions is split into two when we add a second line, and that this is true *no matter how you choose to draw the lines*; as long as we make sure the two lines aren't parallel, they will always behave this way. That is, if we take one line that splits the plane into two regions,




then adding a new line will split each of those existing regions in two. This adds two new regions to the whole plane, giving four regions in total:




What about when $n = 3$? In this case, we want to think about adding a third line to a diagram with two lines and four regions. We want to make an argument that doesn't depend on a *particular* arrangement of the lines, so eventually the only facts we should use are that no lines are *parallel* and any point of *intersection* only lies on two lines (not three or more). For now, though,

it helps to look at a particular arrangement of lines so that we are talking about the same diagram; we can use our observations of this specific diagram to guide our general argument. Let's start with a two-line diagram, on the left below, and add a third line, but let's choose the third line so that all of the intersection points are "nearby" or within the scope of the diagram, so that we don't have to rescale the picture:




We certainly have 7 regions now, but we made the third line a separate color so that we can identify where the "new" regions appear: one region (the lower quadrant, Region 4) remains unchanged, but the three other regions are split in two and each of those "splits" adds 1 to our count (where there was 1 region, now there are 2). What if we had placed the line differently?



The same phenomenon occurs, where one quadrant remains untouched but the other three are split in two. (How do we know there aren't any other regions not depicted within the scale of our diagram? This is not as easy a question to answer as you might think at first blush, and it's worth thinking about.) Experiment with other arrangements of the three lines and try to convince yourself that this always happens; also, think about *why* this is the case and *how* we could explain that this must happen. Before giving a general explanation, though, let's examine another small case.

When $n = 4$, we start with three lines and 7 regions and add a fourth line that is not parallel to any of the existing lines and doesn't pass through any existing intersection points. Again, we will want to make an argument that isn't tied to a specific arrangement of the lines, but looking at the following specific diagram will help guide our intuition into making that argument:




Notice that three of the original regions remain unchanged (Regions 3 and 5 and 7), and the other four are split in two. Do you notice a pattern here? It seems like for every n we've examined, adding the n -th line leaves exactly $n - 1$ regions unchanged while the rest are split in two. Let's try to explain why this happens. Remember that we're trying to identify how many regions appear when we draw n lines, so let's assign that value a "name" so we can refer to it; let's say $R(n)$ represents the number of regions created by drawing n lines on the plane so that no two lines are parallel and no intersection point belongs to more than two lines. In these examples we've considered for small values of n , we've looked at what *changes* when we add a new line; that is, we've figured out what $R(n)$ is by already knowing $R(n - 1)$. Let's try to adapt our observations so that they apply to any *arbitrary* value of n .

Assume that we know $R(n)$ already. (Why can we do this? Do we know any particular value of $R(n)$ for sure, for some specific n ? Which? How?) Say we have an *arbitrary* diagram of n lines on the plane that satisfy the two conditions given in the puzzle statement above. How many regions do these lines create? Yes, exactly $R(n)$. Now, what happens when we add the $(n + 1)$ -th line? What can we say *for sure* about this line and how it alters the diagram? Well, the only information we really have is that (a) this new line is not parallel to any of the existing n lines and (b) this new line does not intersect any of the already existing intersection points. Now, condition (a) tells us that this new line must intersect *all* of the existing n lines; parallel lines don't intersect, and non-parallel lines must intersect somewhere. Thus, we must create n new intersection points on the diagram. Can any of those intersection points coincide with any existing intersection points? No! This is precisely what condition (b) tells us. These two pieces of information together tell us that, no matter how we draw this new line, as long as it satisfies the requirements of the puzzle, we *must* be able to identify n "special" points along that line. Those special points are precisely the points where the new line intersects an existing line.

We'd now like to take these special points and use them to identify new regions in the diagram. Look back to the cases we examined above: identify the new intersection points and see if you can associate them with new regions. Perhaps it would help to label those intersections with a large dot and mark the new regions with an X to make them all stand out. We'll show you one example below, where $n = 4$. What do you notice? Can you use these dots to

help identify how many new regions are created with the addition of that n -th line? Think about this for a minute and then read on.



Exactly! Between any two of the new intersection points, we have a line *segment* that splits a region in two! All that remains is to identify how many new such segments we've created. Since each one corresponds to exactly *one* existing region split in two, this will tell us exactly how many new regions we've created. We've already figured out that this $(n+1)$ -th line creates n new intersection points. Think about how these points are arranged on the line. Any two "consecutive" points create a segment, but the extreme points also create infinite segments (that continue past those extreme points forever). How many are there in total? Exactly $n+1$. (Look at the diagram above, for $n=3$. We see that there are 3 new intersection points and 4 new segments, with two of them being infinite rays.) This means there are $n+1$ line segments that split regions in two, so we have created exactly $n+1$ new regions! This allows us to say that

$$R(n+1) = R(n) + n + 1$$

Phew, what an observation! It took a bit of playing around with examples and making some geometric arguments, but here we are. We've identified some *inductive structure* to this puzzle; we've found how one case depends on another one. That is, we've found how $R(n+1)$ depends on $R(n)$. This hasn't *completely* solved the puzzle, but we are now much closer. All that remains is to replace $R(n)$ with a similar expression, and continually do this until we reach a value we know, $R(1) = 2$. Observe:

$$\begin{aligned}
 R(n+1) &= \\
 &= \\
 &= \cancel{R(n-2)} + \frac{\cancel{R(n-1)}}{(n-1)} + \frac{\cancel{R(n)}}{n} + n + 1 \\
 &\vdots \\
 &= \cancel{R(2)} + 3 + \cdots + n + n + 1 \\
 &= R(1) + 2 + 3 + \cdots + n + (n+1)
 \end{aligned}$$

Since we know $R(1) = 2$, we can say

$$R(n+1) = 2 + (2 + 3 + \cdots + n + (n+1)) = 2 + \left(\sum_{k=1}^{n+1} k \right) - 1 = 1 + \sum_{k=1}^{n+1} k$$

and this is a sum we have investigated before! (Also notice that we had to subtract 1 because of the missing first term of the sum in parentheses.) Recall that $\sum_{k=1}^n k = \frac{n(n+1)}{2}$, and to represent the sum we have in the equation above, we just replace n with $n + 1$. Therefore,

$$R(n+1) = 1 + \frac{(n+1)(n+2)}{2}$$

One final simplification we would like to make is to replace $n+1$ with n throughout the equation, because it makes more sense to have an expression for $R(n)$ (For what values of n is this valid?)

$$R(n) = 1 + \frac{n(n+1)}{2}$$

Finally, we have arrived at an answer to the originally-posed puzzle! In so doing, we employed an *inductive* technique: we explained how one “fact”, namely the value of $R(n+1)$, *depends* on the value of a “previous fact”, namely $R(n)$, and used these iterative dependencies to work backwards until we reached a particular, *known* value, namely $R(1)$.

We want to point out, again, that the derivation we followed and the observations we made in this section have guided our intuition into an answer, but this has not *rigorously proven* anything. The issue is with the “...”, which is not a concrete, “officially” mathematical way of capturing the inductive process underlying our technique. Furthermore, our method with the “lines in the plane” problem had us *starting* with a diagram of $n - 1$ lines and *building* a new diagram with n lines; is this okay? Why does this actually tell us anything about an *arbitrary* diagram of n lines? Do all such diagrams come from a smaller diagram with one fewer line?

We will, in the next two chapters, learn the necessary tools to fully describe a *rigorous* way of doing what we have done so far, and in the chapter after that, we will employ those tools to make **mathematical induction** officially rigorous. For now, though, we want to give a heuristic definition of induction and continue to examine interesting puzzles and observations that rely on inductive techniques. Practicing with these types of puzzles—learning when to recognize an inductive process, how to work with it, how to use that structure to solve a problem, and so on—will be extremely helpful in the future, and we have no need to delve into technical mathematical detail. (At least, not just yet!)

2.2.3 Questions & Exercises

Remind Yourself

Answering the following questions briefly, either out loud or in writing. These are all based on the section you just read, so if you can’t recall a specific definition or concept or example, go back and reread that part. Making sure you can confidently answer these before moving on will help your understanding and memory!

- (1) What properties characterize an *inductive* process?
- (2) How did we prove that $\sum_{k=1}^n k = \frac{n(n+1)}{2}$ is correct? How was our method inductive? (Reread Section 1.4.2 if you forget!)
- (3) Why can we take the sum formula mentioned in the previous question and “replace” n with $n+1$ and know that it still holds true? Can we also replace n with $n-1$?
- (4) Work through the algebraic steps to obtain our final expression for the sum of the first n squares; that is, verify that

$$\frac{1}{3}(n+1)^3 - \frac{1}{3}(n+1) - \frac{n(n+1)}{2} = \frac{1}{6}n(n+1)(2n+1)$$

- (5) Try to recall the argument that adding the $(n+1)$ -th line on the plane created *exactly* $n+1$ new regions. Can you work through the argument for a friend and convince him/her that it is valid?
- (6) To find the sum of the first n squares, why couldn’t we just square the formula for the sum of the first n numbers? Why is that wrong?

Try It

Try answering the following short-answer questions. They require you to actually write something down, or describe something out loud (to a friend/classmate, perhaps). The goal is to get you to practice working with new concepts, definitions, and notation. They are meant to be easy, though; making sure you can work through them will help you!

- (1) Draw 5 lines on the plane (that satisfy the two conditions of the puzzle) and verify that there are 16 regions. Can you also verify that 6 lines yield 22 regions?
- (2) Come up with another description of a sequence that goes $1, 2, 3, 4, \dots, 100$ that is *not* simply all of the numbers from 1 to 100. (Recall the example we gave: all the numbers from 1 to 100 with no “i” in their English spelling.)
- (3) Come up with an algebraic expression that relates $(n+1)^4$ to n^4 , like we did with cubes.

(Challenge: Can you come up with a *geometric* interpretation for the expression you just derived?)

- (4) **Challenge:** Let’s bump the “lines in the plane” puzzle up one dimension! Think about having n planes in three-dimensional space. How many regions are created? Assume that no two planes are parallel, and no three of them intersect in one line. (Think about how these two conditions are directly analogous to the specified conditions for the “lines” puzzle.)

2.3 Defining Induction

To properly motivate the forthcoming definition of **mathematical induction** as a proof technique, we want to emphasize that the above examples used some intuitive notions of the structure of the puzzle to develop a “solution”, where we use quotation marks around *solution* to indicate that we haven’t officially proven it yet. In that sense, we ask the following question: What if we had been *given* the formula that we derived and asked to verify it? What if we had not gone through any intuitive steps to derive the formula and someone just told us that it is correct? How could we check their claim? The reason we ask this is because we are really facing that situation now, except the person telling us the formula is . . . the very same intuitive argument *we* discovered above!

Pretend you have a skeptical friend who says, “Hey, I heard about this formula for the sum of the first n natural numbers squared. Somebody told me that they add up to $\frac{1}{6}n(n+1)(2n+1)$. I checked the first two natural numbers, and it worked, so it’s gotta be right. Pass it on!” Being a logical thinker, but also a good friend, you nod along and say, “I did hear that, but let’s make *sure* it’s correct for *every* number.” How would you proceed? Your friend is right that the first few values “work out” nicely:

$$\begin{aligned}1^2 &= 1 = \frac{1}{6}(1)(2)(3) \\1^2 + 2^2 &= 5 = \frac{1}{6}(2)(3)(5) \\1^2 + 2^2 + 3^2 &= 14 = \frac{1}{6}(3)(4)(7) \\1^2 + 2^2 + 3^2 + 4^2 &= 30 = \frac{1}{6}(4)(5)(9)\end{aligned}$$

and so on. We could even check, by hand, a large value of n , if we wanted to:


$$1^2 + 2^2 + 3^2 + 4^2 + 5^2 + 6^2 + 7^2 + 8^2 + 9^2 + 10^2 = 385 = \frac{1}{6}(10)(11)(21)$$

Remember, though, that this formula is claimed to be valid for *any* value of n . Checking individual results by hand would take forever, because there are an *infinite* number of natural numbers. No matter how many individual values of n we check, there will always be larger values, and how do we *know* that the formula doesn’t break down for some large value? We need a far more *efficient* procedure, mathematically and temporally speaking, to somehow verify the formula for all values of n in just a few steps. We have an idea in mind, of course (it’s the upcoming rigorous version of mathematical induction), and here we will explain how the procedure works, in a broad sense.

2.3.1 The Domino Analogy

Pretend that we have a set of dominoes. This is a special set of dominoes because we have an infinite number of them (!) and we can imagine anything we want

written on them, instead of the standard array of dots. Let's also pretend that they are set up in an infinite line along an infinite tabletop, and we are viewing the dominos from the side and we can see a label under each one so that we know where we are in the line:



For this particular example, to verify the formula

$$\sum_{k=1}^n k^2 = \frac{1}{6}n(n+1)(2n+1)$$

we will imagine a particular “fact” written on each domino. Specifically, we will imagine that the 1st domino has the expression

$$\sum_{k=1}^1 k^2 = \frac{1}{6}(1)(2)(3)$$

written on it, and the 2nd domino has the expression


$$\sum_{k=1}^2 k^2 = \frac{1}{6}(2)(3)(5)$$

written on it. In general, we imagine that the n -th domino in the infinite line has the following “fact” written on it:

$$\sum_{k=1}^n k^2 = \frac{1}{6}n(n+1)(2n+1)$$

Since we're dealing with dominos that are meant to fall into each other and knock each other over, let's pretend that whenever a domino falls, that means the corresponding “fact” written on it *is a true statement*. This is how we will relate our physical interpretation of the dominos to the mathematical interpretation of the validity of the formula we derived.

We did check the sum for $n = 1$ by hand: $1^2 = \frac{1}{6}(1)(2)(3)$. Thus, the fact written on the first domino is a true statement, so we know that the first domino will, indeed, fall over. We also checked the sum for $n = 2$ by hand, so we know that the second domino will fall over:



However, continuing like this brings us back to the same problem as before: we don't want to check *every individual* domino to make sure it falls. We would really like to encapsulate our physical notion of the line of dominoes—namely, that when a domino falls it will topple into the next one and knock that over, and so on—and somehow relate the “facts” that are written on adjacent dominoes.

Let's look at this situation for the first two dominoes. Knowing that Domino 1 falls, can we guarantee that Domino 2 falls without rewriting all of the terms of the sum? How are the statements written on the two dominoes related? Each statement is a sum of squared natural numbers, and the one on the second domino has exactly one more term. Thus, knowing *already* that Domino 1 has fallen, we can use the *true statement* written on Domino 1 to *verify* the truth of the statement written on Domino 2:

$$\sum_{k=1}^2 k^2 = 1^2 + 2^2 = 1 + 2^2 = 5 = \frac{1}{6}(2)(3)(5)$$

Now, this may seem a little silly because the only “work” we have saved is not having to “do the arithmetic” to write $1^2 = 1$. Let's use this procedure on a case with larger numbers so we can more convincingly illustrate the benefit of this method. Let's *assume* that Domino 10 has fallen. (In case you are worried about this assumption, we wrote the full sum a few paragraphs ago and you can verify it there.) This means we *know* that

$$\sum_{k=1}^{10} k^2 = \frac{1}{6}(10)(11)(21) = 385$$

is a *true statement*. Let's use this to verify the statement written on Domino 11, which is

$$\sum_{k=1}^{11} k^2 = \frac{1}{6}(11)(12)(23)$$

The sum written on Domino 11 has 11 terms, and the first 10 are exactly the sum written on Domino 10! Since we know something about that sum, let's just separate that 11th term from the sum and apply our knowledge of the other

terms:

$$\begin{aligned}
 \sum_{k=1}^{11} k^2 &= (1^2 + 2^2 + \cdots + 10^2) + 11^2 \\
 &= \left(\sum_{k=1}^{10} k^2 \right) + 11^2 \\
 &= 385 + 121 \\
 &= 506 \\
 &= \frac{1}{6} 3036 = \frac{1}{6} (11)(12)(23)
 \end{aligned}$$

Look at all of the effort we saved! Why bother reading the first 10 terms of the sum if we know something about them already?

Now, imagine if we could do this procedure for *all* values of n , *simultaneously*! That is, imagine that we could prove that any time Domino n falls, we are *guaranteed* that Domino $(n + 1)$ falls. What would this tell us? Well, think about the infinite line of dominos again. We *know* Domino 1 will fall, because we checked that value by hand. Then, because we verified the “Domino n knocks over Domino $(n + 1)$ ” step for *all* values of n , we know Domino 1 will fall into Domino 2, which in turn falls into Domino 3, which in turns falls into Domino 4, which \dots The entire line of dominos will fall! In essence, we could collapse the whole line of dominos falling down into just *two* steps:

- (a) Make sure the first domino topples;
- (b) Make sure every domino knocks over the one immediately after it in line.

With only these two steps, we can *guarantee* every domino falls and, therefore, *prove* that all of the facts written on them are true. This will prove that the formula we derived is valid for *every* natural number n .

We have already accomplished step (a), so now we have to complete step (b). We have done this for specific cases in the previous paragraphs (Domino 1 topples Domino 2, and Domino 10 topples Domino 11), so let’s try to follow along with the steps of those cases and generalize to an arbitrary value of n . We *assume*, for some *specific but arbitrary* value of n , that Domino n falls, which tells us that the equation

$$\sum_{k=1}^n k^2 = \frac{1}{6} n(n+1)(2n+1)$$

is a *true statement*. Now, we want to relate this to the statement written on Domino $(n + 1)$ and apply the knowledge given in the equation above. Let’s do what we did before and write a sum of $n + 1$ terms as a sum of n terms plus the last term:

$$\sum_{k=1}^{n+1} k^2 = 1^2 + 2^2 + \cdots + n^2 + (n+1)^2 = \left(\sum_{k=1}^n k^2 \right) + (n+1)^2$$

Next, we can apply our assumption that Domino n has fallen (which tells us that the fact written on it is true) and write

$$\sum_{k=1}^{n+1} k^2 = \frac{1}{6}n(n+1)(2n+1) + (n+1)^2$$

Is this the same as the fact written on Domino $(n+1)$? Let's look at what that is, first, and then compare. The “fact” on Domino $(n+1)$ is similar to the fact on Domino n , except everywhere we see “ n ” we replace it with “ $n+1$ ”:

$$\sum_{k=1}^{n+1} k^2 = \frac{1}{6}(n+1)((n+1)+1)(2(n+1)+1) = \frac{1}{6}(n+1)(n+2)(2n+3)$$

It is not clear yet whether the expression we have derived thus far is actually equal to this. We could attempt to simplify the expression we've derived and factor it to make it “look like” this new expression, but it might be easier to just expand both expressions and compare all the terms. (This is motivated by the general idea that expanding a factored polynomial is far easier than recognizing a polynomial can be factored.) For the first expression, we get

$$\begin{aligned} \frac{1}{6}n(n+1)(2n+1) + (n+1)^2 &= \frac{1}{6}n(2n^2 + 3n + 1) + (n^2 + 2n + 1) \\ &= \frac{1}{3}n^3 + \frac{1}{2}n^2 + \frac{1}{6}n + n^2 + 2n + 1 \\ &= \frac{1}{3}n^3 + \frac{3}{2}n^2 + \frac{13}{6}n + 1 \end{aligned}$$

and for the second expression, we get

$$\begin{aligned} \frac{1}{6}(n+1)(n+2)(2n+3) &= \frac{1}{6}(n+1)(2n^2 + 7n + 6) \\ &= \frac{1}{6}[(2n^3 + 7n^2 + 6n) + (2n^2 + 7n + 6)] \\ &= \frac{1}{3}n^3 + \frac{3}{2}n^2 + \frac{13}{6}n + 1 \end{aligned}$$

Look at that; they're identical! Also, notice how much easier this was than trying to rearrange one of the expressions and “morph” it into the other. We proved they were identical by manipulating them both and finding the same expression, ultimately. Now, let's look back and assess what we have accomplished:

1. We likened proving the validity of the formula

$$\sum_{k=1}^n k^2 = \frac{1}{6}n(n+1)(2n+1)$$

for *all* values of n to knocking over an infinite line of dominos.

2. We verified that Domino 1 will fall by checking the formula corresponding to that case by hand.
3. We proved that Domino n will fall into Domino $(n+1)$ and knock it over by *assuming* the fact written on Domino n is true and using that knowledge to show that the fact written on Domino $(n+1)$ must also be true.
4. This guarantees that *all* dominos will fall, so the formula is true for all values of n !

Are you convinced by this technique? Do you think we've *rigorously proven* that the formula is valid for all natural numbers n ? What if there were a value of n for which the formula didn't hold? What would that mean in terms of our domino scheme?

Remember that this domino analogy is a good intuitive guide for how induction works, but it is not built on mathematically rigorous foundations. That will be the goal of the next couple of chapters. For now, let's revisit the other example we've examined in this section: lines in the plane. Again, the use of *ellipses* in our derivation of the formula $R(n)$ is troublesome and we want to avoid it. Let's try to follow along with the domino scheme in the context of this puzzle.

Imagine that we have defined the expression $R(n)$ to represent the number of distinct regions in the plane created by n lines, where no two lines are parallel and no three intersect at one point. Also, imagine that on Domino n we have written the "fact" that " $R(n) = 1 + \frac{n(n+1)}{2}$ ". Can we follow the same steps as above and verify that all the dominos will fall?

First, we need to check that Domino 1 does, indeed fall. This amounts to verifying the statement: " $R(1) = 1 + \frac{1(2)}{2} = 1 + 1 = 2$ ". Is this a true statement? Yes, of course, we saw this before; one line divides the plane into two regions. Second, we need to prove that Domino n will topple into Domino $(n+1)$ for any *arbitrary* value of n . That is, let's *assume* that " $R(n) = 1 + \frac{n(n+1)}{2}$ " is a *true* statement for some value of n and *show* that " $R(n+1) = 1 + \frac{(n+1)(n+2)}{2}$ " must also be a true statement. How can we do this? Well, let's follow along with the argument we used before to relate $R(n+1)$ to $R(n)$. By considering the geometric consequences of adding an extra line to *any* diagram with n lines (that also fit our rules about the lines) we proved that $R(n+1) = R(n) + n + 1$. Using this knowledge *and* our assumption about Domino n falling, we can say that

$$R(n+1) = R(n) + n + 1 = 1 + \frac{n(n+1)}{2} + n + 1$$

Is this the same expression as what is written on Domino $(n+1)$? Again, let's simplify *both* expressions to verify they are the same. We have

$$1 + \frac{n(n+1)}{2} + n + 1 = 2 + n + \frac{n^2 + n}{2} = \frac{1}{2}n^2 + \frac{3}{2}n + 2$$

and

$$1 + \frac{(n+1)(n+2)}{2} = 1 + \frac{n^2 + 3n + 2}{2} = \frac{1}{2}n^2 + \frac{3}{2}n + 2$$

Look at that; they're identical! Thus, we have shown that Domino n is *guaranteed* to fall into Domino $(n + 1)$, for *any* value of n . Accordingly, we can declare that *all* dominos will fall!

Think about the differences between what we have done with this “domino technique” and what we did before to derive the expressions we just proved. Did we use any ellipses in this section? Why is it better to prove a formula this way? Could we have used the domino induction technique to *derive* the formulas themselves?

2.3.2 Other Analogies

The Domino Analogy is quite popular, but it's not the only description of how induction works. Depending on what you read, or who you talk to, you might learn of a different analogy, or some other kind of description altogether. Here, we'll describe a couple that we've heard of before. It will help solidify your understanding of induction (at least as far as we've developed it) to think about how these analogies are all *equivalent*, fundamentally.

Mojo the Magical, Mathematical Monkey

Imagine an infinite ladder, heading straight upwards into the sky. There are infinitely-many rungs on this ladder, numbered in order: 1, 2, 3, and so on going upwards. Our friend Mojo happens to be standing next to this ladder. He is an intelligent monkey, very interested in mathematics but also a little bit magical, because he can actually climb up this infinite ladder!

If Mojo makes it to a certain rung on the ladder, that means the fact corresponding to that number is **True**. How can we make sure he climbs up the entire ladder? It would be inefficient to check each rung individually. Imagine that: we would have to stand on the ground and make sure he got to Rung 1, then we would have to look up a bit and make sure he got to Rung 2, and then Rung 3, and so on . . . Instead, let's just confirm two details with Mojo before he starts climbing. Is he going to start climbing? That is, is he going to make it to Rung 1? If so, great! Also, are the rungs close enough together so that he can *always* reach the next one, no matter where he is? If so, even greater! These are exactly like the conditions established in our Domino Analogy. To ensure that Mojo gets to *every* rung, we just need to know he gets to the first one and that he can always get to the next one.

Doug the Induction Duck

Meet Doug. He's a duck. He also loves bread, and he's going to go searching through everyone's yards to find more bread. These yards are all along Induction Street in Math Town, where the houses are numbered 1, 2, 3, and so on, all in a row.

Doug starts in the yard of house #1, looking for bread. He doesn't find any, so he's still hungry. Where else can he look? The house next door, #2, has a

backyard, too! Doug heads that way, his tummy rumbling. He doesn't find any bread there, either, so he has to keep looking. He already knows house #1 has no bread, so the only place to go is next door to house #3. We think you see where this is going . . .

If we were keeping track of Doug's progress, we might wonder whether he eventually gets to *every* yard. Let's say we also knew ahead of time that *no-body* has any bread. This means that whenever he's in someone's yard, he will definitely go to the next house, still searching for a meal. This means that he will definitely get to every house! That is, no matter which house we live in, no matter how large the number on our front door might be, at some point we will see Doug wandering around our backyard. (Unfortunately, he will go hungry all this time, though! Poor Doug.)

2.3.3 Summary

Let's reconsider what we've accomplished with the two example puzzles we've seen thus far, and the analogies we've given. In our initial consideration of each puzzle, we identified some aspect of the *structure* of the puzzle where a "fact" depended on a "previous fact". In the case of the cubic numbers, we found a way to express $(n + 1)^3$ in *terms* of n^3 ; in the case of the lines in the plane, we described how many regions were added when an extra line was added to a diagram with n lines. From these observations, we applied this encapsulated knowledge over and over until we arrived at a "fact" that we knew, for a "small" value of n (in both cases, here, $n = 1$). This allowed us to derive a formula or equation or expression for a general fact that should hold for *any* value of n .

This work was interesting and essential for deriving these expressions, but it was *not enough* to *prove* the validity of the expressions. In doing the work described above, we identified the presence of an inductive process and utilized its structure to derive the expressions in question. This was beneficial in two ways, really; we actually found the expressions we wanted to prove and, by recognizing the inductive behavior of the puzzle, we realized that proving the expressions by *mathematical induction* would be prudent and efficient.

For the actual "proof by induction", we followed two main steps. First, we identified a "starting value" for which we could check the formula/equation by hand. Second, we *assumed* that one particular value of n made the corresponding formula hold true, and then used this knowledge to show that the corresponding formula for the value $n + 1$ must also hold true. Between those two steps, we could rest assured that "all dominos will fall" and, therefore, the formulas would hold true for all relevant values of n .

One Concern: What's at the "top" of the ladder?

You might be worried about something, and we'll try to anticipate your question here. (We only bring this up because it's a not uncommon observation to make. If you *weren't* thinking about this, try to imagine where the idea would come from.) You might say, "Hey now, I think I see how Mojo is climbing the ladder,

but how can he actually get *all the way to the top*? It's an infinite ladder, right? And he never gets there . . . right?"

In a way, you would be right. Since this magical ladder really does go on *forever*, then there is truly no *end* to it and Mojo will never get "there". However, that isn't the point; we don't care about any "end" of the ladder (and not just because there *isn't* one). We just need to know that Mojo actually gets to *every possible* rung. He doesn't have to surpass all of them and stand at the top of the ladder, looking down at where he came from. That wasn't the goal!

Think of it this way: pretend you have a vested interest in some particular fact that we're proving. Let's say it's Fact #18,458,789,572,311,000,574,003. (Some huge number. It doesn't matter, really.) Its corresponding rung is waaayyyyyy up there on the ladder, and all you care about is whether or not Mojo gets there on his journey. Does he? . . . You bet he does! It might take a long time (how many steps will it take?), but in this magical world of monkeys and ladders, who cares about time anyway! You know that he'll eventually get there, and that makes you happy. Now, just imagine that for each fact, there's somebody out there in that magical world that cares about only that fact. Surely, everyone will be happy with the knowledge that Mojo will get to their rung on his journey. Nobody cares about whether he gets to the top; that isn't their concern. Meanwhile, out here in our regular, non-magical world, we are extremely happy with the fact that everyone in *that* world will eventually be happy. That entire infinite process of ladder-climbing was condensed into just two steps, and with only those two steps, we can rest assured that every rung on that ladder will be touched. Every numbered fact is true.

Think about this in terms of the Domino Analogy, as well. Do we care whether or not there is some "ending point" of the line of dominoes, so that they all fall into a wall somewhere? Of course not; the line goes on forever. Every domino will eventually fall over, and we don't even care how "long" that takes. Likewise, we know Doug will get to everyone's yard; we don't care "when" he gets to any *individual* yard, just that he gets to *all* of them.

2.3.4 Questions & Exercises

Remind Yourself

Answering the following questions briefly, either out loud or in writing. These are all based on the section you just read, so if you can't recall a specific definition or concept or example, go back and reread that part. Making sure you can confidently answer these before moving on will help your understanding and memory!

- (1) How are the Domino, Mojo, and Doug analogies all *equivalent*? Can you come up with some "function" that describes their relationship, that converts one analogy into another?
- (2) Find a friend who hasn't studied mathematical induction before, and try to describe it. Do you find yourself using one of the analogies? Was it helpful?

- (3) Why is it the case that our work with the cubes didn't *prove* the summation formula? Why did we still need to go through all that work?
- (4) Think about the Domino Analogy. Is it a problem that the line of dominoes goes on forever? Does this mean that there are some dominoes that will never fall down? Try to describe what this means in terms of the analogy.

Try It

Try answering the following short-answer questions. They require you to actually write something down, or describe something out loud (to a friend/classmate, perhaps). The goal is to get you to practice working with new concepts, definitions, and notation. They are meant to be easy, though; making sure you can work through them will help you!

- (1) Work through the inductive steps to prove the formula

$$\sum_{k=1}^n k = \frac{n(n+1)}{2}$$

- (2) Work through the inductive steps to prove the formula

$$\sum_{k=1}^n (2k-1) = n^2$$

- (3) Work through the inductive steps to prove the formula

$$\sum_{k=1}^n k^3 = \left(\frac{n(n+1)}{2} \right)^2$$

- (4) Suppose we have a series of facts that are indexed by natural numbers. Let's use the expression " $P(n)$ " to represent the n -th fact.
 - (a) If we want to prove *every* instance is **True**, for every natural number n , how can we do this?
 - (b) What if we want to prove that only every *even* value of n makes a **True** statement? Can we do this? Can you come up with a modification of one of the analogies we gave that would describe your method?
 - (c) What if we want to prove that only every value of n greater than or equal to 4 makes a **True** statement? Can we do this? Can you come up with a modification of one of the analogies we gave that would describe your method?

2.4 Two More (Different) Examples

This short section serves a few purposes. For one, we don't want you to get the idea, right away, that induction is all about proving a *numerical formula* with numbers and polynomials. Induction is so much more useful than that! One of the following examples, in particular, will be about proving some abstract property is true for any “size” of the given situation. You will see how it still falls under the umbrella of “induction”, but you will also notice how it is different from the previous examples. Furthermore, these examples will illustrate that sometimes we need to know “more information” to knock over some dominoes. In the previous examples, we only needed to know that Domino n fell to *guarantee* that Domino $n + 1$ will fall. Here, though, we might have to know about several previous dominoes. After these two examples, we will summarize how this differs from the domino definition given above, and preview a broader definition of the technique of induction, as it applies to these examples.


2.4.1 Dominos and Tilings

This next example is a little more complicated than the first two. We will still end up proving a certain numerical *formula*, but the problem is decidedly more visual than just manipulating algebraic expressions. Furthermore, we'll notice an interesting “kink” in the starting steps, where we have to solve a couple of “small cases” before being able to generalize our approach. This will be our first consideration of how the technique of induction can be generalized and adapted to other situations.


The question we want to answer is nicely stated as follows:

Given a $2 \times n$ array of squares, how many different ways can we tile the array with dominoes? A *tiling* must have every square covered by one—and *only* one—domino.

For example, the following are proper tilings



whereas the following are *not* proper tilings




As before, let's examine the first few cases—where $n = 1, 2, 3$, and so on—and see if we notice any patterns. Try working with the problem yourself, before reading on, even!


When $n = 1$, we have an array that is exactly the shape of one domino, so surely there is only one way to do this. Let's use the notation $T(n)$ to represent the number of tilings on a $2 \times n$ array. Thus, $T(1) = 1$.




When $n = 2$, we have a 2×2 array. Since the orientation of the array matters, we have each of the following distinct tilings. Thus, $T(2) = 2$.



What about when $n = 3$? Again, we can enumerate these tilings by hand and be sure that we aren't missing any. We see that $T(3) = 3$.




Okay, one more case, when $n = 4$. We see that $T(4) = 5$.




Can we start to find a pattern now? Writing out larger arrays will just be tiresome! Let's think about how we could have used the fact that $T(1) = 1$ to deduce something about $T(2)$... Well, wait a minute ... We couldn't, right? There was something fundamentally different about those two cases. Specifically, because dominoes are 2×1 in size, the fact that we only added one row to the array didn't help us.

Alright, let's consider $n = 3$, then. Could we use the fact that $T(2) = 2$ at all? In this case, yes! Knowing there were two tilings of the 2×2 array, we could immediately build two tilings of the 2×3 array without much thought. Specifically, we can just *append a vertical domino* to each of those previous tilings. But we know now that $T(3) = 3$. Where did the third tiling come from? Look at that tiling again and how it compares to the other two. In that third tiling, the dominoes on the right side are horizontal, as opposed to the vertical one in the other two tilings. If we remove those two parallel, horizontal dominoes, we are left with precisely the situation when $n = 1$. Put another way, we can build a tiling of a 2×3 array by *appending a square of two horizontal dominoes* to the right side. In total, then, we have described all of the tilings of a 2×3 board in terms of boards of smaller sizes, namely 2×2 and 2×1 :



Now you might see how the pattern develops! Let's show you what happens when $n = 4$, how we can construct *all* of the tilings that make up $T(4)$ by appending a vertical domino to each of the tilings that make up $T(3)$, or by appending two horizontal dominoes to each of the tilings that make up $T(2)$:



Notice, as well, that no tiling for the 2×4 array was produced *twice* in this way. (Think carefully about why this is true. How can we characterize the two types of tilings in a way that will guarantee they don't coincide at all?) With this information, we can immediately conclude that $T(4) = T(3) + T(2)$.

Furthermore, we can generalize this argument; nothing was special about $n = 4$, right? For any particular n , we can just consider all possible tilings, and look at what happens on the *far right-hand side* of the array: either we have one vertical domino (which means the tiling came from a $2 \times (n - 1)$ array) or two horizontal dominoes (which means the tiling came from a $2 \times (n - 2)$ array). With confidence in this argument, we can conclude that

$$T(n) = T(n-1) + T(n-2)$$

for all of the values of n for which this expression makes sense. What values are those? Remember that we had to identify $T(1)$ and $T(2)$ separately; this argument doesn't apply to those values. Accordingly, we have to add the restriction $n > 3$ for the equation above to hold true.

With this information, we can then easily figure out $T(n)$ for any value of n , given enough time. We could write a computer program fairly easily, even. It was this *inductive* argument, though—the pattern that we noticed and our thorough description of why it occurs—that allowed us to make the conclusion in the first place. In this case, too, it just so happens that the value of every term, $T(n)$, depends on the value of *two* previous terms, $T(n-1)$ and $T(n-2)$. This did *not* happen in our previous examples in this chapter, and it hints at something deeper going on here. Do you see how our previous definition of induction, and the domino analogy, doesn't exactly apply here anymore? How might you try to amend our analogy to explain this kind of situation? Think

about these issues for a bit and then read on. We'll talk about them more in-depth after the next example.

By the way, did you notice something interesting about our solution to this example? Do you know any other sequences of numbers that behave similarly? Think about it ...

2.4.2 Winning Strategies

This example will be our first induction puzzle that *doesn't* prove a numerical formula! It might seem strange to think about that, but it's true, as you'll see. This is actually more common in mathematics than you might think, too: a problem or mathematical object might have some underlying inductive structure without depending on something algebraic or arithmetic.

In fact, we will be discussing a *game*. It's a game in the usual sense—there are rules to be followed by two players and there is a clear winner and loser—but it's also a game in the mathematical sense, where we can formulate the rules and playing situations using mathematical notation and discuss formal *strategies* in an abstract way. We can even *solve* the game. This is very different than say, the game of baseball.


Let's discuss the rules for this game, which we shall call "Takeaway", for now. There are two players, called P_1 and P_2 . The player P_1 goes first every time. The players start with two piles of stones in the middle of a table, each pile containing exactly n stones, where n is some natural number. (To distinguish the different versions of the game, we will say the players are "playing T_n " when there are n stones per pile.) On each player's move, they are allowed to remove *any* number of stones from *either* pile. It is illegal, though, to remove stones from both piles at once. The player who removes the *final* stone from the piles is the *winner*.

Try playing Takeaway with some friends. Use pennies or candies or penny candy as stones. Try it for different values of n . Try switching roles so you are P_1 and then P_2 . Try to come up with a winning *strategy*, a method of playing that maximizes your chances of winning. Try to make a conjecture for what happens for different values of n . Who is "supposed" to win? Can you *prove* your claim? Seriously, play around with this game and attempt to prove something before reading on for our analysis thereof. You might be surprised by what you can accomplish!

As with the other examples, let's use some small values of n to figure out what's really going on, then try to generalize. When $n = 1$, this game is rather silly. P_1 must empty one pile of its only stone, then P_2 gets the only remaining stone in the other pile. Thus, P_2 wins. (Notice that it doesn't matter which of the two piles P_1 picks from, P_2 will always get the other one. We might say that P_1 picks the pile on the left "without loss of generality" because it doesn't matter either way; the situations are equivalent, so we might as well say it's the left pile to have something concrete to say. We will explore this idea of "without loss of generality" later on when we discuss mathematical logic, too.)



When $n = 2$, we now have a few cases that might appear. Think about P_1 's possible moves. Again, they might act on either the left or right pile, but because they're ultimately identical and we can switch the two piles, let's just say (without loss of generality) that P_1 removes some stones from the left pile. How many? It could be one or two stones. Let's examine each case separately.




If P_1 removes both stones, how should P_2 react? Duh, they should take the other pile, so P_1 probably shouldn't have made that move in the first place. However, P_1 might not be thinking straight or something and, besides, we need to consider all possible situations here to fully analyze this game. Thus, in this case (the top line in the above diagram) P_2 wins. Okay, that's the easy situation.

What if P_1 removes just one stone from the left pile (the bottom line above)? How should P_2 react? We now have some options:

- If P_2 removes the other stone from the left pile ... well, P_1 takes the other pile and P_1 wins.
- If P_2 removes both stones from the right pile ... well, P_1 takes the last stone from the left pile and P_1 wins.

- However, if P_2 removes just one stone from the right pile ...




Now we have exactly the same situation presented by T_1 , which we already analyzed! It is, again, P_1 's move first, so we know what will happen: P_2 wins no matter what. If you are player P_2 , this is obviously the best move: you win *no matter how P_1 responds!*

Stepping back for a second, let's think about what this has shown: no matter what P_1 does first (remove one or two stones from either pile), there is *some possible response* that P_2 can make that will *guarantee* a win for P_2 , regardless of P_1 's subsequent response. Wow, P_2 is sitting pretty! Let's see if this happens for other values of n .


When $n = 3$, we will again assume (without loss of generality) that player P_1 acted on the left pile. They could remove one, two, or three stones:

- If P_1 removes all three, P_2 responds by taking the other pile completely and wins.
- If P_2 removes two stones ... well, what should player P_2 do?

Finishing off that left pile is stupid (because P_1 can take the whole right pile and win), and pulling the entire right pile is similarly stupid (because P_1 can take the whole left pile and win), so something in between is required. Now, if P_2 removes just one stone from the right pile, notice that P_1 can respond with the same action; this leaves exactly one stone in both piles, but the roles reversed. With P_2 going first in such a situation, they are now bound to lose, per our previous analysis. Bad move, P_2 !



Let's try again. If P_2 removes two stones from the right pile instead ... look at that! We now have exactly one stone in each pile, with P_1 up first, so we know P_1 is going to lose. P_2 strikes again!



Think about the case where $n = 4$ for a minute, and you'll find the exact same analysis occurring. You'll another possibility to consider: player P_1 can remove one, or two, or three, or four stones from the left pile. Whatever they do, though, you'll find that P_2 can just *mimic that action* on the other pile, reducing the whole game to a previous, *smaller* version of the game, where P_2 was shown to be guaranteed a win! It looks like P_2 is in the driver's seat the whole time, since they can respond to whatever P_1 does, making an identical move on the other pile. No matter what P_1 does, there is always a response for P_2 that means they win, regardless of P_1 's subsequent moves. In this sense, we say " P_2 has a winning strategy". There is a clear and describable method for P_2 to assess the game situation and choose a specific move to *guarantee a win*.

How might we prove this? How does this even fit into this chapter on induction? It might be hard to see, at the moment. What are we really even proving here? What are the dominoes or rungs in our analogy for this problem? In wrapping your brain around this example, you should hopefully realize the following: induction is *not* about algebraic formulas all the time; induction represents some kind of "building-up" structure, where larger situations depend on smaller ones; we have to prove some initial fact, and then argue how an arbitrary, larger fact can be reduced so that it depends on a previous fact. This is really what the dominoes analogy is meant to accomplish. It just so happens that this analogy is particularly illustrative for certain induction problems (but not all) and is visualizable and memorable. It does not perfectly apply to *all* situations, though.

Read back through these four examples from this chapter and think about how they are similar and how they are different. Try to come up with a more precise, mathematical description of mathematical induction using some better terminology, perhaps of your own invention. (By this, we mean something better than our intuitive analogy. You'd be surprised at how well you might be able to describe induction without really knowing what you "ought" to say, and you'll actually learn a lot, in the process!) In due time, we will have a rigorous statement to make, and prove, about mathematical induction and its various forms. In the meantime, we need to take a trip through some other areas of mathematics to build up the necessary language, notation, and knowledge to come back and tackle this problem. Before we go, though, we should mention a few useful applications of mathematical induction.

2.4.3 Questions & Exercises

Remind Yourself

Answering the following questions briefly, either out loud or in writing. These are all based on the section you just read, so if you can't recall a specific definition or concept or example, go back and reread that part. Making sure you can confidently answer these before moving on will help your understanding and memory!

- (1) How are these two examples *inductive*? In what ways are they similar to the previous examples, with the cubes and lines? In what ways are they different?
- (2) With the domino tilings, how *many* previous values did we need to know to compute $T(n)$?
- (3) What is the difference between writing $T(n) = T(n-1) + T(n-2)$ and $T(n+2) = T(n+1) + T(n)$?
- (4) What is the winning strategy in the Takeaway game? Try playing with a friend who doesn't know the game, and use that strategy as player P_2 . How frustrated do they get every time you win? Do they start to catch on?

Try It

Try answering the following short-answer questions. They require you to actually write something down, or describe something out loud (to a friend/classmate, perhaps). The goal is to get you to practice working with new concepts, definitions, and notation. They are meant to be easy, though; making sure you can work through them will help you!

- (1) What is $T(5)$? Can you draw all of those tilings?
- (2) Work through the possibilities for takeaway with two piles of 4 stones. Can you make sure that player P_2 always has a winning move?
- (3) **Challenge:** What happens if you play Takeaway with *three* piles of equal sizes? Can you find a winning strategy for either player? Try playing with a friend and see what happens!
- (4) Look up the *Fibonacci numbers*. How are they related to the sequence of numbers $T(n)$ we found in the domino tiling example?

2.5 Applications

2.5.1 Recursive Programming

The concepts behind mathematical induction are employed heavily in computer science, as well. Think back to how we first derived the formula for $\sum_{k=1}^n k^2$.

Once we had a way to represent a cubic number in terms of a smaller cube and some leftover terms, we repeated this substitution process over and over until we arrived at the “simplest” case, namely, the one that we first observed when starting the problem: $2^3 = 1+3+3+1$. Recursive programming takes advantage of this technique: to solve a “large” problem, identify how the problem depends on “smaller” cases, and reduce the problem until one reaches a simple, known case.

A classical example of this type of technique is seen in writing code to compute the *factorial* function, $n!$, which is defined as the product of the first n natural numbers:

$$n! = 1 \cdot 2 \cdot 3 \cdots (n-1) \cdot n$$

This is a simple definition that we, as humans, intuitively understand, but telling a computer how to perform this product doesn’t work quite the same way. (Try it! How do you say “and just keep going until you reach n ” in computer code?) A more efficient way to program the function, and one that models the mathematically inductive definition, in fact, is to have one program *recursively call itself* until it reaches that “simple” case. With the factorial function, that case is $1! = 1$. For any other value of n , we can simply apply the knowledge that

$$n! = (n-1)! \cdot n$$

over and over to compute $n!$. Consider the following *pseudocode* that represents this idea:

```
factorial(n):
    if n = 1
        return 1
    else
        return n * factorial(n-1)
    end
```

We know that $1! = 1$, so if the program is asked to compute that, the correct value is returned right away. For any larger value of n , the program refers to *itself* and says, “Go back and compute $(n-1)!$ for me, then I’ll add a factor of n at the end, and we’ll know the answer.” To compute $(n-1)!$, the program asks, again, if the input is 1; if not, it calls itself and says, “Go back and compute $(n-2)!$ for me, then I’ll add a factor of $n-1$ at the end.” This process continues until the program returns $1! = 1$. From there, it knows how to find $2! = 1 \times 2$, and then $3! = 2! \times 3$, and so on, until $n! = (n-1)! \times n$.

Another example involving recursive programming arises with the *Fibonacci numbers*. You may have seen this sequence of numbers before in a mathematics course. (In fact, we even mentioned them in the last section, with the domino tilings!) You also might have heard about how they appear in nature in some interesting and strange ways. (The sequence was first “discovered” by the Italian mathematician Leonardo of Pisa while studying the growth of rabbit populations.) The first two numbers in the sequence are specified to be 1, and any

number in the sequence is defined as the sum of the previous two. That is, if we say $F(n)$ represents the n -th Fibonacci number, then

$$F(1) = 1 \quad \text{and} \quad F(2) = 1 \quad \text{and} \quad F(n) = F(n-1) + F(n-2) \text{ for every } n \geq 3$$

Now, what is $F(5)$? Or $F(100)$? Or $F(10000)$? This can be handled quite easily by a recursive program. The idea is the same: if the program refers to either one of the “simple cases”, i.e. $F(1)$ or $F(2)$, then it will know to return the correct value of 1 immediately. Otherwise, it will call itself to compute the previous two numbers and then add those together. Look at the pseudocode below and think about how it works. What would happen if we used this program to compute $F(10)$? How would it figure out the answer?

```

Fibonacci(n):

  if n = 1 or n = 2
    return 1
  else
    return Fibonacci(n-1) + Fibonacci(n-2)
  end

```

This follows the same idea as the **factorial** program above (let the program call itself to compute values for “smaller” cases of the function until we reach a known value) but there’s something a little deeper going on here. If we were to input $n = 10$ into the program, it would recognize that it does not know the output value yet, and it will call itself to compute **Fibonacci(9)** and **Fibonacci(8)**. In each of those calls to the program, it would again recognize the value is as yet unknown. Thus, it would call upon itself again to compute **Fibonacci(8)** and **Fibonacci(7)**, but also **Fibonacci(7)** and **Fibonacci(6)**. That’s right, the program calls itself multiple times with the *same input value*. To compute $F(9)$, we need to know $F(8)$ and $F(7)$, but meanwhile, to compute $F(8)$, we also need to know $F(7)$ and $F(6)$. In this way, we end up calling the program **Fibonacci** many times.


Try to compare the programs **Fibonacci** and **factorial**, especially in regards to the inductive processes we have been investigating in this chapter. Do they use similar ideas? How do they relate to the “domino” analogy of mathematical induction that we outlined? Think of the “fact” written on Domino n as being the computation of the correct value of $n!$ or $F(n)$. How does the analogy work in each case? Will all the dominos fall? Keep these questions in mind as you read on. There is some very powerful mathematics underlying all of these ideas.

2.5.2 The Tower of Hanoi

Let’s take a short break and play a game. Well, it’s not exactly a break because this is, in a sense, an *inductive* game, so it’s completely relevant. But it is a

game, nevertheless! The *Tower of Hanoi* is a very popular puzzle, partly because it involves such simple equipment and rules. Solving it is another matter, though!

Imagine that we have three vertical rods and three disks of three different sizes (colored blue, green, and red) stacked upon each other like so:




The goal is to move all three disks to another rod (either the middle or the right one, it doesn't matter) by following these rules:

1. A single move consists of moving one (and *only* one) disk from the top of the stack on any rod and moving it to the top of the stack on another rod.
2. A disk cannot be placed on top of a smaller disk.

That's it! Two simple rules, but a difficult game to play. Try modeling the game with a few coins or playing cards or whatever you have handy. (You can even buy Tower of Hanoi sets at some games stores.) Can you solve it? How many moves did it take you? Is your solution the "best" one? Why or why not?

We mentioned that this is an *inductive* game, so let's explore that idea now. We want to consider how many *moves* it takes to solve the puzzle (where one *move* accounts for moving one disk from one rod to another) and, more specifically, identify the *smallest possible number of moves* it would take to solve the puzzle. To solve the puzzle with three disks, we could keep moving the smallest disk back and forth between two rods and generate 100 moves, if we wanted to, and then solve it, but that's certainly not the best way to do it, right? Let's say we found a way to solve the puzzle in a certain number of moves; how could we show that the number of moves we used is the *smallest possible* number of moves?

To address this question, we want to break down the method of solving the puzzle *recursively*. In doing so, we are actually going to answer a far more general question: What is the smallest number of moves required to solve the Tower of Hanoi puzzle with n disks on 3 rods? We posed the puzzle above with just 3 disks to give you a concrete version to think about and work with, but we can answer this more general question by thinking carefully. To make sure we are on the same page, we will show you how we solved the version with 3 disks:



Notice that the largest disk is essentially “irrelevant” for most of the solution. Since we are allowed to place any other disk on top of it, all we need to do is “uncover” that disk by moving the other disks onto a different rod, move the largest disk to the only empty rod, then move the other disks on top of the large one. In essence, we perform the same procedure (shifting the two smaller disks from one rod to another) twice and, in between those, we move the large disk from one rod to another. If the largest disk hadn’t been there at all, what we actually did was solve the version of the puzzle with 2 disks, but twice! (Think carefully about this and make sure you see why this is true. Follow along with the moves in the diagrams above and pretend the large, blue disk isn’t there.)

This shows that the way to solve the 3-disk puzzle involves two iterations of solving the 2-disk puzzle, with one extra move in between (moving the largest disk). This indicates a *recursive* procedure to solve the puzzle, in general. To optimally solve the n -disk puzzle, we would simply follow the procedure to optimally solve the $(n - 1)$ -disk puzzle, use one move to shift the largest, n -th disk, then solve the $(n - 1)$ -disk puzzle again.

Now that we have some insight into *how* to optimally solve the puzzle, let’s identify how many *moves* that procedure requires. Recognizing that solving this puzzle uses a *recursive* algorithm, we realize that *proving* anything about the optimal solution will require *induction*. Accordingly, we would need to identify a “starting point” for our line of dominos, and it should correspond to the “smallest” or “simplest” version of the puzzle. For the Tower of Hanoi, this is the 1-disk puzzle. Of course, this is hardly a “puzzle” because we can solve it in one move, by simply shifting the only disk from one rod to any other rod. If we let $M(n)$ represent the number of *moves* required to optimally solve the n -disk puzzle, then we’ve just identified $M(1) = 1$. To identify $M(2)$, we can use our observation from the previous paragraph and say that

$$\underbrace{M(2)}_{\text{solve 2-disk}} = \underbrace{M(1)}_{\text{solve 1-disk}} + \underbrace{1}_{\text{shift largest disk}} + \underbrace{M(1)}_{\text{solve 1-disk}} = 1 + 1 + 1 = 3$$

and then it must be that

$$M(3) = M(2) + 1 + M(2) = 3 + 1 + 3 = 7$$

and

$$M(4) = M(3) + 1 + M(3) = 7 + 1 + 7 = 15$$

and so on. Do you notice a pattern yet? Each of these numbers is one less than a power of 2, and specifically, we notice that $M(n) = 2^n - 1$, for each of the cases we have seen thus far. It’s important to point out that observing this pattern doesn’t *prove* the pattern; just because it works for the first 4 cases does not mean the trend will continue, but that’s exactly what an induction proof would accomplish. Also, recognizing that pattern and “observing” that $M(n) = 2^n - 1$ is a non-trivial matter, itself. We happened to know the answer and had no problem identifying the formula for you. You should probably try, on your own, to “solve” the following relationship

$$M(n) = 2M(n - 1) + 1 \quad \text{and} \quad M(1) = 1$$

and see if you can derive the formula $M(n) = 2^n - 1$. The reason such a formula is *nicer* than the above relationship is that, now, $M(n)$ depends only on n , and not on previous terms (like $M(n-1)$, for example). This relationship and others like it are known as *recurrence relations*, and they can be rather difficult to solve, in general!

We know how to solve this one, though, and it yields $M(n) = 2^n - 1$. We will leave it to you to verify this. You can do so by checking a few values in the equation above, but we all know that isn't a *proof*. Try working through the inductive steps to actually prove it! We have already done most of the work, but it will be up to you to arrange everything carefully and clearly. Remember that you should identify what the "fact" on each domino is, ensure that Domino 1 falls, and then make a general argument about Domino n toppling into Domino $(n+1)$. Try to write that proof. Do the details make sense to you? Try showing your proof to a friend and see if they understand it. Did you need to tell them anything else or guide them through it? Think about the best way to *explain* your method and steps so that the written version suffices and you don't have to add any verbal explanations.

2.5.3 Questions & Exercises

Remind Yourself

Answering the following questions briefly, either out loud or in writing. These are all based on the section you just read, so if you can't recall a specific definition or concept or example, go back and reread that part. Making sure you can confidently answer these before moving on will help your understanding and memory!

- (1) How is a recursive program inductive?
- (2) What is the inductive structure of the Tower of Hanoi? Where did we solve the 2-disk puzzle while solving the 3-disk puzzle?

Try It

Try answering the following short-answer questions. They require you to actually write something down, or describe something out loud (to a friend/classmate, perhaps). The goal is to get you to practice working with new concepts, definitions, and notation. They are meant to be easy, though; making sure you can work through them will help you!

- (1) Follow the steps of the psuedocode `factorial` to compute $5!$.
- (2) Follow the steps of the psuedocode `Fibonacci` to compute $F(5)$.
- (3) Solve the Tower of Hanoi puzzle with 4 disks. Make sure that you can do it in the *optimal* number of moves, $2^4 - 1 = 15$.

2.6 Summary

We have now seen some examples of **inductive arguments**. We realized that some of the puzzles we were solving used similar argument styles, and explored several examples to get a flavor for the different issues that might come up in such arguments. Specifically, we saw how inductive arguments are *not* always about proving a summation formula or an equation: inductive arguments can apply to *any* situation where a fact depends on a “previous instance” of that fact. This led us into developing an analogy for how induction works, mathematically speaking. We are comfortable with thinking of induction in terms of the “Domino Analogy” for now, but one of our main goals in moving forwards is rigorously *stating* and *proving* a principle of induction. For now, let’s get lots of practice working with these kinds of arguments. This is what this chapter’s exercises are meant to achieve. Later on, once we’ve formalized induction, we’ll be better off for it, and we’ll have a thorough understanding of the concept!

2.7 Chapter Exercises

Here are some problems to get you comfortable working with inductive-style arguments. We aren’t looking for fully rigorous proofs here, just a good description of what is going on and a write-up of your steps. We’ll come back to some of these later and rigorously prove them, once we’ve established the Principle of Mathematical Induction (PMI) and a corresponding proof strategy.

Problem 2.7.1. Prove the following summation formula holds for every natural number, and for $n = 0$, as well:

$$\sum_{i=0}^n 2^i = 2^{n+1} - 1$$

Follow-up question: use this result to state how many games are required to determine a winner in a single-elimination bracket tournament with 2^n teams. (For example, the NCAA March Madness Tournament uses this format, with $n = 6$.)

Problem 2.7.2. Prove that $3^n \geq 2^{n+1}$ for every natural number n that is greater than or equal to 2.

Problem 2.7.3. For which natural numbers n do the following inequalities hold true? State a claim and then prove it.

1. $2^n \geq (n+1)^2$
2. $2^n \geq n!$
3. $3^{n+1} > n^4$
4. $n^3 + (n+1)^3 > (n+2)^3$

Problem 2.7.4. The December 31 Game: Two players take turns naming dates from a calendar. On each turn, a player can increase the month or date but not both. The starting position is January 1, and the winner is the person who says December 31. Determine a winning strategy for the first player.

For example, a sequence of moves that yields Player 1 winning is as follows:

(1) January 10, (2) March 10, (1) August 10, (2) August 25, (1) August 28, (2) November 28, (1) November 30, (2) December 30, (1) December 31

By *winning strategy* we mean a method of play that Player 1 follows that *guarantees* a win, no matter what Player 2 does.

Problem 2.7.5. Find and prove a formula for the sum of a *geometric series*, which is a series of the form

$$\sum_{i=0}^{n-1} q^i$$

for some real number q and some natural number n . (Hint: be careful when $q = 1$.)

Problem 2.7.6. Write a sentence that depends on n such that the sentence is true for all values of n from 1 to 99 (inclusive), but such that the sentence is false when $n = 100$.

Problem 2.7.7. What is wrong with the following “spoof” of the claim that $a^n = 1$ for every n ?

“Spoof”: Let a be a nonzero real number. Notice that $a^0 = 1$. Also, notice that we can inductively write

$$a^{n+1} = a^n \cdot a = a^n \cdot \frac{a^n}{a^{n-1}} = 1 \cdot \frac{1}{1} = 1$$

“□”

Problem 2.7.8. In a futuristic society, there are only two different denominations of currency: a coin worth 3 Brendans, and a coin worth 8 Brendans. There is also a nation-wide law that says shopkeepers can only charge prices that can be paid in **exact change** using these two coins.

What are the legal costs that a shopkeeper could charge you for a cup of coffee?

Hint: Try a bunch of small values and see what happens.

Problem 2.7.9. Consider a chessboard of size $2^n \times 2^n$, for some arbitrary natural number n . Remove **any** square from the board. Is it possible to tile the remaining squares with L -shaped triominoes?

If your answer is **Yes**, prove it.

If your answer is **No**, provide a counterexample argument. (That is, find an n such that no *possible* way of tiling the board will work, and show why this is the case.)

Problem 2.7.10. Consider an $n \times n$ grid of squares. How many sub-squares, of any size, exist within this grid? For example, when $n = 2$, the answer is 5: there are 4 1×1 squares and 1 2×2 square. Find a formula for your answer and try to prove it is correct.

Problem 2.7.11. Prove that, in a line of at least 2 people, if the 1st person is a woman and the last person is a man, then somewhere in the line there is a man standing immediately behind a woman.

Problem 2.7.12. Prove that $n^3 - n$ is a multiple of 3, for every natural number n .

Problem 2.7.13. A **binary n -tuple** is an ordered string of 0s and 1s, with n total numbers in the string. Provide an *inductive argument* to explain why there are 2^n possible binary n -tuples.

Problem 2.7.14. Recall that the **Fibonacci Numbers** are defined by setting $f_0 = 0$ and $f_1 = 1$ and then, for every $n \geq 2$, setting $f_n = f_{n-1} + f_{n-2}$. This produces the sequence 0, 1, 1, 2, 3, 5, 8, 13, 21, 34, ...

You might not know that the Fibonacci Numbers also have a *closed form*; that is, there is a specific *formula* that defines them, in addition to the usual recursive definition given above. Here it is:

$$f_n = \frac{1}{\sqrt{5}} \left[\left(\frac{1 + \sqrt{5}}{2} \right)^n - \left(\frac{1 - \sqrt{5}}{2} \right)^n \right]$$

Prove that this formula is correct for all values of $n \geq 0$.

Problem 2.7.15. Again, considering the Fibonacci Numbers, f_n , prove the following:

1. $\sum_{i=0}^n f_i = f_{n+2} - 1$
2. $\sum_{i=0}^n f_i^2 = f_n \cdot f_{n+1}$
3. $f_{n-1} \cdot f_{n+1} - f_n^2 = (-1)^n$
4. $f_{m+n} = f_n \cdot f_{n+1} + f_{m-1} \cdot f_n$
5. $f_n^2 + f_{n+1}^2 = f_{2n+1}$

Problem 2.7.16. Try to provide an inductive argument that explains why every natural number $n \geq 2$ can be written as a product of prime numbers. Can you also show that this product is *unique*? That is, can you also explain why there is *exactly one way* to factor a natural number into primes?

Problem 2.7.17. Prove that

$$\sum_{k=1}^n k \cdot k! = 1 \cdot 1! + 2 \cdot 2! + 3 \cdot 3! + \cdots + n \cdot n! = (n+1)! - 1$$

Problem 2.7.18. What is wrong with the following “spoof” that all pens have the same color.

“Spoof”: Consider a group of pens with size 1. Since there is only 1 pen, it certainly has the same color as itself.

Assume that any group of n pens has only one color represented inside the group. (Note: we explained why this assumption is valid for $n = 1$ already, so we can make this assumption.) Take any group of $n + 1$ pens. Line them up on a table and number them from 1 to $n + 1$, left to right. Look at the first n of them, i.e. look at pens 1, 2, 3, \dots , n . This is a group of n pens so, by assumption, there is only one color represented in the group. (We don’t know what color that is yet.) Then, look at the last n of the pens; i.e. look at pens 2, 3, \dots , $n + 1$. This is also a group of n pens so, by assumption, there is only one color represented in this group, too. Now, pen #2 happens to belong to both of these groups. Thus, whatever color pen #2 is, that is also the color of every pen in *both* groups. Thus, all $n + 1$ pens have the same color.

By induction, this shows that any group of pens, of any size, has only one color represented. Looking at the finite collection of pens in the world, then, we should only find one color. \square

Problem 2.7.19. ★ This problem is *extremely difficult* to analyze, and is taken from the blog of the famous mathematician Terence Tao ([link here](#)).

There is an island upon which a tribe resides. The tribe consists of 1000 people, with various eye colours. Yet, their religion forbids them to know their own eye color, or even to discuss the topic; thus, each resident can (and does) see the eye colors of all other residents, but has no way of discovering his or her own (there are no reflective surfaces). If a tribesperson does discover his or her own eye color, then their religion compels them to commit ritual suicide at noon the following day in the village square for all to witness. All the tribespeople are highly logical and devout, and they all know that each other is also highly logical and devout (and they all know that they all know that each other is highly logical and devout, and so forth).

(For the purposes of this logic puzzle, “highly logical” means that any conclusion that can logically deduced from the information and observations available to an islander, will automatically be known to that islander.)

Of the 1000 islanders, it turns out that 100 of them have blue eyes and 900 of them have brown eyes, although the islanders are not initially aware of these statistics (each of them can of course only see 999 of the 1000 tribespeople).

One day, a blue-eyed foreigner visits to the island and wins the complete trust of the tribe.

One evening, he addresses the entire tribe to thank them for their hospitality.

However, not knowing the customs, the foreigner makes the mistake of mentioning eye color in his address, remarking *how unusual it is to see another blue-eyed person like myself in this region of the world*.

What effect, if anything, does this faux pas have on the tribe?

2.8 Lookahead

In this chapter, we have introduced you to the concept of **mathematical induction**. We looked at a few examples of puzzles where an inductive process guided our solution, and then we described how a *proof by induction* would follow to *rigorously verify* that solution. With the mathematical techniques and concepts we have at hand thus far, we had to rely on a non-technical analogy to describe this process to you. Thinking of an infinite line of dominos with “facts” written on them knocking into each other is a perfectly reasonable interpretation of this process, but it fails to represent the full mathematical extent of induction. In a way, it’s like having a friend describe to you how to swing a golf club, even though you’ve never played golf before. Sure, they can provide you with some mental imagery of what a swing “feels like”, but without getting out there and practicing yourself, how will you truly understand the mechanics of the golf swing? How will you learn how to adapt your swing, or tell the differences between using a driver and a five iron and a sand wedge? By investigating the underlying mechanics and practicing with those concepts, we hope to gain a better understanding of mathematical induction so that, in the future, we can use it appropriately, identify situations where it would be useful, and, eventually, learn how to *adapt* it to other situations. Of course, it will help to have that domino analogy in mind to guide our intuition, but we should also remember that it is not rigorous mathematics. It also doesn’t perfectly describe the other examples we discussed, where a falling domino depended on not only the one immediately behind it, but several others before it.

In the next chapter, we will explore some relevant concepts needed to rigorously state and prove mathematical induction as a proof technique. Specifically, we will study some ideas of *mathematical logic* and investigate how to break down complicated mathematical statements and theorems into their constituent parts, and also how to build interesting and complex statements out of basic building blocks. Along the way, we will introduce some new notation and shorthand that will let us condense some of the wordy statements we make into concise (and precise) mathematical language. With that in hand, we will explore some more fundamental proof strategies, that we will then apply to *everything else we do* in this course, including the induction technique, itself! We will also study some of the ideas of *set theory*, a branch of mathematics that forms the foundation for all other branches. This will be extremely useful for organizing our ideas in the future, but it will also help us define the *natural numbers* in a rigorous manner. With some concepts and knowledge from these two branches of mathematics under our collective belts, we will be able to build mathematical induction on a solid foundation and continue to use it properly.