

INTRODUCTION TO MATHEMATICAL THINKING

Linear algebra can be viewed as the mathematical apparatus needed to solve systems of linear equations, understand their underlying structure, and to apply what is learned in other contexts. Unlike your first brush with the subject in high school, which probably emphasized matrices and Euclidean spaces like the real plane $\mathbb{R}^2$, we will focus on abstract vector spaces and linear transformations. These terms will be defined later, so do not worry if you are not familiar with them. These notes start from the beginning of the subject, assuming no knowledge for linear algebra. The key point is that you are about to immerse yourself in serious mathematics, with an emphasis on your attaining a deep understanding of the definitions, theorems, and proofs.

Upper-level mathematics has much in common with high school mathematics, and students who have been accepted onto a Biomedical Engineering major already have an array of mathematical skills that will serve them well. On the other hand, upper-level mathematics also differs in some important respects. This means that most students need to extend and adapt their existing skills in order to continue to do well. Making such extensions and adaptations can be difficult for those who have never really reflected on the nature of their skills.

One thing you certainly have learned to do is to apply mathematical procedures to calculate answers to standard questions, perhaps by matching key words in the statement of a question with a memorized recipe of calculations from within a set of many recipes. This might have been useful to achieve a good grade in the standardized entry exam to university, which is a controlled and predictable environment. Moreover, some people enjoy this type of work. They like the satisfaction of arriving at a page of correct answers, and they like the security of knowing that if they do everything right then their answers will, indeed, be correct. Sometimes they compare mathematics favorably with other subjects in which things seem to be more a matter of opinion and “there are no right answers.”

Other people dislike this aspect of mathematics. They find it dull to do lots of repetitious exercises, and they get more satisfaction from learning about *why* the various procedures work and how they fit together. However, note that knowing how to apply procedures is extremely important because, without fluency in calculations, it is hard to focus your attention on higher-level concepts.

When you start taking upper-level courses, professors will expect you to be fluent in using the procedures you have already learned. They will expect you to be able to accurately manipulate algebraic expressions, to solve equations, and so on. They will expect you to be able to do these things without having

to stop each time to look up a rule, and they might not be patient with students who are not able to do so. This is not because they are impatient with students in general – most professors will be happy to spend a long time talking with you about new mathematics, or responding to students who say, “I know how to do this, but I’ve never understood why we do it this way.” But they will not expect to re-teach things you have already studied. So you should brush up your knowledge prior to beginning a course, especially if, say, you’ve done no mathematics all summer.

Once you do begin, you’ll find that some upper-level mathematics involves learning new procedures. These procedures, unsurprisingly, will be longer and more complicated than those you met in earlier work. We are not worried about your ability to apply long and complicated procedures, however, because to have got this far you must be able to do that kind of thing. Here, we want to focus on more substantive changes in the ways in which you have to interact with the procedures.

The first substantive difference is that you will have more responsibility for deciding which procedure to apply. Of course, you have learned to do this to some extent already. For instance, you have learned how to multiply out brackets and write things like:

$$(x+2)(x-5) = x^2 - 3x - 10.$$

But hopefully you have also learned that it is *not* sensible to multiply out when trying to simplify a fraction like this:

$$\frac{x^2(x+2)(x-5)}{x^2+2x}.$$

For the fraction, simplification is easier if you keep the factors “visible.” Nonetheless, many students automatically multiply out, probably because multiplying out was one of the first things they learned to do when studying algebra. They would, however, be more successful in mathematics courses like calculus and linear algebra if they learned to stop and think first about what would allow them to make the most progress. If this doesn’t apply to you for this particular type of problem, does it apply in others? Have you ever done a long calculation and then realized that you didn’t need to? Could you have avoided it if you’d stopped to think first? Part of deciding which procedures to apply is giving yourself a moment to think about it before you leap in and do the first thing that comes to mind.

This might not sound like a big deal, but think for a moment about how often you *don’t* have to make a choice about what procedure to apply. Often, questions

in books or on tests tell you exactly what to do. They say things like, "Use Cramer's rule to solve this system of equations." Even when a question doesn't tell you outright, it is sometimes obvious from the context. In high school, if your teacher spent a lesson showing you how to apply the Rouché-Frobenius theorem, then gave you a set of questions to do, it was probably safe for you to assume that these would involve automatically applying the Rouché-Frobenius theorem. This helped you out, but it means that much of the time you didn't have to decide what procedure to apply. In the wider world, and in advanced mathematics, making decisions is more highly valued and more often expected. This means that questions presented to you on problem sheets and exams will usually just say "solve this problem" rather than "solve this problem using this procedure."

Another part of deciding which procedures to apply is being able to distinguish between cases that look similar but are best approached in different ways. For example, consider integration, and more specifically integration by parts. You may know that this is used when we want to integrate a product of two functions, one of which gets simpler when we differentiate it, and the other of which does not get any more complicated when we integrate it. For instance, in $\int xe^x dx$, x gets simpler if we differentiate it, and e^x gets no more complicated if we integrate it. You might also know, however, that sometimes mathematical situations look superficially similar, but are best tackled using different procedures. In the integration case, integration by substitution might sometimes be more appropriate. For instance, in $\int xe^{x^2} dx$, we would probably want to use integration by substitution instead of by parts. Can you see why?

Integration by substitution is a good case for another point we would like to make, this time about making decisions *within* procedures. It might be that you read the end of the last paragraph and thought, "But what substitution should I use?" Perhaps your teachers or books always told you what to use, but we would argue that they shouldn't necessarily have to. After a while, you should notice that certain substitutions are useful in certain cases. If you pay attention to the structures of these cases then, even if you wouldn't be sure that you could pick a good substitution for a new case, you should have an idea of some sensible things to try. If you haven't deliberately thought about this before, we suggest you do so now. Get out some questions on integration by substitution and, without actually doing the problems, look at the suggested substitutions. Can you anticipate why those substitutions will work? Can you then anticipate what would work in similar cases?

So how can a student improve their ability to make decisions about and within procedures? We have two suggestions. The first is to try doing exercises from a source where the procedure to be applied is not obvious. A good place to look is in the books listed in the bibliography section of the teaching guide. The second suggestion is to turn ordinary exercises into opportunities for

reflection. When you finish an exercise, instead of just moving on to the next one, stop and think about these questions:

1. Why did that procedure work?
2. What could be changed in the question so that it would still work?
3. What could be changed in the question so that it would *not* work?
4. Could I modify the procedure so that it would work for some of these cases?

All of these questions should help you build flexibility in applying what you know.

We want to come back to the idea that knowing how to apply such procedures is only one part of understanding mathematics. It is an important part, but most students can recognize the difference between learning to apply a procedure mechanically, and understanding why it works. Learning mechanically has some advantages: it is generally quick and relatively straightforward. But it also has disadvantages: if you learn procedures mechanically, it is easier to forget them, to misapply them, and to mix them up (and it will be harder for you to be as successful in this course as you would like to be). Developing a proper understanding of why things work is generally harder and more time-consuming, but the resulting knowledge is easier to remember and more supportive of flexible and accurate reasoning (and will certainly help you get a good grade in this course).

To summarize, before taking upper-level courses like calculus and algebra, you should brush up your knowledge of standard procedures because your professors will expect you to be able to use these fluently. As you progress, you will be expected to take more responsibility for deciding which procedure to apply; it might be a good idea to practice this by working on exercises from sources that do not tell you exactly what to do. You will also be expected to adapt procedures in sensible ways, and to work out how theorems or definitions can be applied without necessarily having seen many worked examples. You will not succeed in upper-level mathematics if you always try to solve problems by finding something that looks similar and copying it. You have to be more thoughtful than that. Mathematics is not just about procedures. Fluency with procedures is important, but in many cases you should aim for a deeper understanding of why procedures work.

Definitions

We will encounter many *definitions*, *theorems*, and *proofs* throughout these notes.

A definition has nothing to do with something being true. It just tells us what a mathematical word means. A definition might define a type of object, or it might define a property. The definition about defines a property about linear equations.

You already know that a definition tells us what a word means because you've been using dictionaries for a very long time. But there are two very, very important differences between dictionary definitions and mathematical definitions. If you want to understand the material in these notes, it is vital that you understand these differences.

The first difference is that when mathematicians (in particular, your algebra professor) state a definition, *they really mean it*. They don't mean that this is a good description of the majority of cases, but that there might be exceptions out there somewhere. This is not how dictionary definitions work. If you took two dictionaries and looked up an everyday concept, either a concrete one (like "table") or an abstract one (like "justice"), would you expect the definitions to be identical? Probably not. Probably, in fact, you would expect to be able to find things in the world that satisfy one definition but not the other. Or to find something that you would want to call a "table" or "justice" but that didn't really satisfy either definition*. Or to find something for which, even though there is a definition, people would disagree.

Now, compare this with what would happen if you took a mathematics book and looked up a definition of "even number." Would you expect to be able to find an even number that did not satisfy the definition? Or an non-even number that did, nonetheless, satisfy the definition? Absolutely not. Exceptions simply don't exist. It's not like you could find a number so big that it could be even without being divisible by 2. If you took two textbooks, you might not expect the two definitions to be phrased in exactly the same way, but you would expect them to be logically equivalent. That is, you would expect that all and only the things that satisfy the first definition would also satisfy the second. So, that's one difference: mathematical definitions mean exactly, *exactly* what they say. There are no exceptions, and different phrasings might exist but these must be logically equivalent.

The other difference, which is partly a consequence of the first, is that mathematical definitions are precise and operable in a way that dictionary definitions are not. This means that they contain some information that we can actually manipulate in an algebraic or logical argument. For instance, consider the following simple definition:

* For example, is a 10-meter tall "table" at an art exhibition a table? Sort of, but it wouldn't satisfy functional criteria such as being a flat surface you could rest things on – no-one would be able to reach, and anyway touching the exhibit is probably not allowed.

Definition: A number is *even* if it is divisible by 2.

"Well, yes," you'll be thinking, "obviously." In fact, though, this is probably not how your professor will write this definition. You're more likely to see something like this:

Definition: A number n is *even* if and only if there is an integer k such that $n = 2k$.

You could be forgiven for thinking that this over-complicates things. But it has some advantages. The first is precision. We can see this by comparing with the kind of thing that students sometimes write when they try to define *even*. They tend to say things like "Even is when it's divisible by 2." Clearly that captures the key idea, but it isn't very precise. For a start, what is "it"? Clearly "it" is supposed to be a number, but this number isn't introduced properly. In contrast, the better definition tells us explicitly that we are dealing with a number and gives it a name, n . For another thing, the student definition contains the locution "is when." This tends to sound clunky, in mathematics but also in other fields. In mathematics, if you find yourself writing "it" or "is when," you should probably consider rephrasing.

The second advantage is operability. We can operate with the mathematical definition to prove things. The better definition gives us a way of capturing and manipulating even numbers because it states what divisibility means in algebraic terms. We could, for instance, use it to prove that any integer multiple of an even number must also be even, perhaps by writing something like this:

Suppose that n is an even number. Then (by definition) there is an integer k such that $n = 2k$. Now, let z be any integer and $y = zn$. Then $y = z(2k)$. But we can rewrite this as $y = 2(zk)$. Now zk is an integer, so y is even because it can be written in the required form.

We could have the same ideas if we were working with the imprecise student definition. But the better version gives us a leg-up for writing arguments like this by providing some notation. Moreover, we expect the precise definition to *exclude* all the numbers that are not even (for the number 3, there is no appropriate integer k). In this way, the definition of even number captures the notion of evenness in a reasonable way.

A final point here is that we do not "prove" definitions. We can't, because there's nothing to prove: definitions just capture conventions in which we all agree to use a word to mean exactly the same thing. A professor might, at some point, explain to you how a definition captures an intuitive idea, but this isn't the same as proving it.

Theorems

Whereas a definition states what we mean by a word describing a mathematical object or property, a theorem tells us about a relationship between two or more types of objects and properties; it assumes that we already know what those objects and properties are. Of course, if you don't know the meanings of all the words and symbols in the statement of a theorem you will probably feel that you don't understand the theorem. Notice, however, that even if you don't understand a theorem, you should be able to recognize that it has a theorem-like structure of the form:

If this thing is true, then this other thing is true as well.

The bit that goes with the *if* is called the premise or the assumption or the hypothesis. The bit that goes with the *then* is called the conclusion. It's actually not quite this simple in real life, because there are a few ways of phrasing theorems that don't make the "if. . . then. . ." structure so obvious.

Theorems tell us true things about relationships between concepts. Nevertheless, we need to go to the trouble of proving them. Sometimes we do this because theorems are far from obvious. Sometimes, though, theorems are pretty obviously true, and we do it for the more subtle reason that proving them allows us to see how they all fit together to form a coherent theory.

Once definitions are settled, theorems follow by logical necessity. For instance, once we have decided to define even numbers as numbers divisible by 2, we must conclude that any integer multiple of an even number is also even. We don't have any choice about that kind of consequence.

We should say that although we've only used the term *theorem*, there are several words that sometimes get used instead. Some of these are *proposition*, *lemma*, *claim*, *corollary*, and the apparently all-encompassing *result*. *Lemma* is usually used for a small theorem that will then be used to prove a bigger, more important one. *Corollary* is used for a result that follows as a fairly immediate consequence of a big theorem. The other terms are more or less interchangeable (in our opinion).

Thinking about objects can allow us to link new statements to our existing knowledge. However, we can also work towards understanding statements by thinking about their logical structures. To do so, we need to pay attention to mathematical uses of logical language. We will start by looking more carefully at ways in which the word *if* is used in mathematics.

First, consider a statement "If A then B ." This is sometimes written as " $A \Rightarrow B$," which is read out loud as " A implies B ." The use of an arrow makes

it clearer that we can think of this statement as having a direction, which is important because we might have a situation in which $A \Rightarrow B$ is a true statement but $B \Rightarrow A$ is not. Sometimes both implications do hold, as in the case:

$$x \text{ is even} \Rightarrow x^2 \text{ is even (true)}$$

$$x^2 \text{ is even} \Rightarrow x \text{ is even (true)}$$

However, sometimes one is true but the other is not, as in the case:

$$x < 2 \Rightarrow x < 5 \text{ (true)}$$

$$x < 5 \Rightarrow x < 2 \text{ (false)}$$

These statements are actually somewhat imprecise, because we have not specified what type of object x is. You probably assumed that it must be an integer in the squaring case and a real number in the inequalities case, since that would make sense. But, to be clearer, we could write things like:

$$\text{For every } x \in \mathbb{R}, x < 5 \Rightarrow x < 2.$$

Doing so makes it easier to see why this particular statement is false: there are some real numbers that are less than 5 but not less than 2. But mathematicians sometimes omit such phrases when writing in note form or when the intended interpretation is obvious.

When they want to discuss both implications at once, mathematicians use a double-headed arrow meaning "is equivalent to," or write "if and only if" or its abbreviation "iff." So these are different ways of writing the same (true) statement about integers:

$$x \text{ is even} \Leftrightarrow x^2 \text{ is even.}$$

$$x \text{ is even if and only if } x^2 \text{ is even.}$$

$$x \text{ is even iff } x^2 \text{ is even.}$$

We have seen the phrase "if and only if" before, in our definitions. Here it is again:

Definition: A number n is *even* if and only if there is an integer k such that $n = 2k$.

To think about the phrase, it might be illuminating to split up this definition and write each implication separately:

$$\text{A number } n \text{ is even if there exists an integer } k \text{ such that } n = 2k.$$

A number n is even *only if* there exists an integer k such that $n = 2k$.

Can you see how this definition “catches” the numbers that are even, and excludes those that are not?

One final thing to note about a statement of the form “ A if and only if B ” is that, if we want to prove it, we can take one of two approaches. We can either construct a proof in which all the lines are equivalent to each other, or prove the two statements $A \Rightarrow B$ and $B \Rightarrow A$ separately.

The uses of “if” and “ $\Rightarrow$ ” sound straightforward when we are considering simple mathematical statements. But we would like to draw your attention to two potential sources of confusion.

First, some thought is necessary to sort out which of the “if” and the “only if” corresponds to which implication. You should think about this, perhaps by thinking about which one could replace the “implies” arrow in these two versions of the same (true) statement:

$$x < 2 \Rightarrow x < 5 \quad x < 5 \Leftarrow x < 2.$$

Second, it turns out that students do *not* always interpret “if” in a mathematical way in everyday life. In everyday conversation, we tend to speak rather imprecisely, relying on the context to help our listener make the interpretation we intend. For instance, imagine someone says to you,

If you clean the car then you can go out on Friday night.

You could reasonably infer from this that if you don’t clean the car then you can’t go out Friday night. Clearly that is what the speaker intends. And someone else could infer that if you were allowed out on Friday, they you must have cleaned the car. But, in fact, *neither of these is logically equivalent to the original statement*. Perhaps the easiest way to see this is to look at the logic in parallel with our simple statements about inequalities:

clean car $\Rightarrow$ out Friday	$x < 2 \Rightarrow x < 5$
not clean car $\Rightarrow$ not out Friday	$x \geq 2 \Rightarrow x \geq 5$
clean car $\Leftarrow$ out Friday	$x < 2 \Leftarrow x < 5$.

The second and third statements in each case are not logically the same as the first. Alternatively, you could think about the everyday situation, and see that the original statement says nothing at all about what happens if you don’t clean the car, so there would be no contradiction if you didn’t clean it but you were still allowed out. Technically, the person bargaining with you should really say:

You can go out on Friday night *if and only if* you clean the car.

Of course, no one talks like that. Which means that you might have less practice than you think accurately interpreting logical statements. Using logical language in a mathematically correct sense is not too hard, though, because there are cases in which the intended natural language interpretation *is* the same as the mathematical one. Consider the statement:

If Juan is from Sevilla then Juan is from Andalucía.

No one hearing this would dream of inferring either of these:

If Juan is not from Sevilla then Juan is not from Andalucía.

If Juan is from Andalucía then Juan is from Sevilla.

But these inferences are analogous to those we looked at for the statement about cleaning the car. Make sure you can see how.

In the Sevilla case, the natural interpretation of the statement is the same as the mathematical one. We can also use it to illustrate a general point about logical equivalence of different implications. For any statement of the form $A \Rightarrow B$ we can consider three related statements called its *converse*, *inverse*, and *contrapositive*. This is what each one means, using the Sevilla example for illustration.

original	$A \Rightarrow B$	from Sevilla $\Rightarrow$ from Andalucía
converse	$B \Rightarrow A$	from Andalucía $\Rightarrow$ from Sevilla
inverse	not $A \Rightarrow$ not B	not from Sevilla $\Rightarrow$ not from Andalucía
contrapositive	not $B \Rightarrow$ not A	not from Andalucía $\Rightarrow$ not from Sevilla.

This is what each one means using one of our simple mathematical example instead:

original	$A \Rightarrow B$	$x < 2 \Rightarrow x < 5$
converse	$B \Rightarrow A$	$x < 5 \Rightarrow x < 2$
inverse	not $A \Rightarrow$ not B	$x \geq 2 \Rightarrow x \geq 5$
contrapositive	not $B \Rightarrow$ not A	$x \geq 5 \Rightarrow x \geq 2$.

This should help you remember that if $A \Rightarrow B$ is a true statement, then its contrapositive will also be true (in fact, they are logically equivalent), but its inverse and its converse might not be.

Finally, I should point out that your professor will be careful about uses of "if" and "if and only if" in theorems and proofs, but perhaps not so careful in

definitions. In a definition, they might just write “if” instead of “if and only if.” This works for the same reason that everyday communication works: everyone knows that, in a definition, this is what is intended.

Next we want to discuss the phrases “for all” and “there exists.” These are called *quantifiers*, because they tell us how many of something we’re talking about. They are so common in mathematics that we have symbols for them: we use “ $\forall$ ” (the universal quantifier) to mean “for all” and “ $\exists$ ” (the existential quantifier) to mean “there exists.”

In simple statements, quantifiers are easy to think about. Here is a simple quantifier statement:

$$\forall x \in \mathbb{Z}, x^2 \geq 0.$$

In this statement, we write $\forall x \in \mathbb{Z}$ to specify exactly which objects we are talking about. We could just write $\forall x$, and sometimes people do when it is obvious what kind of numbers (or other objects) a statement is about. But doing so could be ambiguous. In this case, the statement could just as well be about real or complex numbers, which raises issues of the truth or otherwise of the statement: “ $\forall x \in \mathbb{Z}, x^2 \geq 0$ ” is true, “ $\forall x \in \mathbb{C}, x^2 \geq 0$ ” is not. So it is good practice to be specific.

More complicated quantified statements can be harder to think about. Here is a definition written in words and in an abbreviated form using the new symbol (which might be more naturally read as “for every” in this case):

Definition: A function $f: \mathbb{R} \rightarrow \mathbb{R}$ is *increasing* if and only if for every $x_1, x_2 \in \mathbb{R}$ such that $x_1 < x_2$, we have $f(x_1) \leq f(x_2)$.

Definition (abbreviated): $f: \mathbb{R} \rightarrow \mathbb{R}$ is *increasing* if and only if $\forall x_1, x_2 \in \mathbb{R}$ such that $x_1 < x_2, f(x_1) \leq f(x_2)$

Here is a simple quantified statement involving the existential quantifier:

$$\exists x \in \mathbb{Z} \text{ such that } x^2 = 25.$$

This is a true statement, because when mathematicians say “there exist,” they mean “there exists at least one.” Here, there are two different integers x that satisfy the statement. In other cases, there might be hundreds. Students sometimes find it strange that we say “there exists” without specifying how many because, if we know exactly how many there are, it seems rude not to say so. However, advanced mathematics is at least partly about general relationships between concepts, rather than about finding “answers” as such. We can see other reasons why it makes sense to use “there exists” without extra

specification if we look at another definition (again shown both in words and in an abbreviated form):

Definition: A number n is *even* if and only if there exists an integer k such that $n = 2k$.

Definition (abbreviated): $n \in \mathbb{Z}$ is *even* if and only if $\exists k \in \mathbb{Z}$ such that $n = 2k$.


This definition gives us an agreed way of deciding whether a number is even or not. To make that decision, we do not care what the particular k is, we just care whether or not there is one. It also allows us to make general arguments about all even numbers. We might start by saying "Suppose n is even, so $\exists k \in \mathbb{Z}$ such that $n = 2k$." In this case, we don't want to specify what the k is because we want the ensuing argument to be general in the sense that it applies to any number that satisfies to the condition.

Some mathematical statements have more than one quantifier. The following definition (again with an abbreviated version) might be described as doubly quantified or as having two nested quantifiers:

Definition: A set $X \subseteq \mathbb{R}$ is *open* if and only if for every $x \in X$ there exists $d > 0$ such that $(x - d, x + d) \subseteq X$.

Definition (abbreviated): $X \subseteq \mathbb{R}$ is *open* if and only if $\forall x \in X \exists d > 0$ such that $(x - d, x + d) \subseteq X$.

When a statement has more than one quantifier, *the order in which they appear is really important*. In this case, the definition says, "for every x , there exists a d ." We should imagine taking a particular x value and finding an appropriate d (maybe depending on x). If we take a different x , we might need a different d (perhaps a smaller one, as in the right-hand diagram below).



If the definition instead said "there exists a d for every x ," mathematicians would read that as meaning that we could select a single d (independent of x) that works for every x . That is not the same thing at all.

To check your understanding of this, consider the following two statements. One is true, and the other is false. Which is which?

$\exists y > 0$ such that $\forall x > 0, y < x$.

$\forall x > 0 \exists y > 0$ such that $y < x$.

It is normal to find this difficult, because everyday life would probably not distinguish between these two statements. Without even realizing it, we would make the interpretation that seems most realistic, disregarding logical correctness. Most students therefore have to concentrate for a while before they get the hang of reading what is literally there and making the mathematical correct interpretation.

Although we have written the theorems in the form “If...then ...,” you might also see different phrasings. Here are some common ones:

Theorem: If f is an even function, then $\frac{df}{dx}$ is an odd function.

Theorem: Suppose that f is an even function. Then $\frac{df}{dx}$ is an odd function.

Theorem: Every even function has an odd derivative.

These would all be interpreted to mean the same thing: the premise in each case is that the function is even, and the conclusion is that its derivative is odd. It might seem strange that we don't just pick one form and stick to it but, sometimes, one version or another sounds more natural, so mathematicians like to have this flexibility.

You will also see theorems of different types. There are, for instance, *existence theorems* like this one:

Theorem: There exists a number x such that $x^3 = x$.

One way to prove a theorem like this is just to produce an object that satisfies it: the number 1 would do, in this case. It's not always that easy, but it's important to recognize that it might be, because sometimes students tie themselves in knots doing complicated things when a simple answer would do.

There are also theorems about non-existence, like this one:

Theorem: There does not exist a largest prime number.

In fact, a bit of thought, non-existence theorems can be restated in our standard form. This one, for instance, could be written with a universal quantifier:

Theorem: For every prime number n there exists a prime number p such that $p > n$.

Then it could be rephrased into our initial form:

Theorem: If n is a prime number then there exists another prime number p such that $p > n$.

These rephrasing possibilities can be very useful when we want to start proving something: sometimes rewriting in a different way can give us different ideas about sensible things to try. However, the fact that we can often rephrase does *not* mean that you can be sloppy about your mathematical writing. Sloppy paraphrasing can easily change the logical meaning of a statement. Once you become fluent in using logical language in a mathematical way, you will find that you can switch forms without doing violence to the meaning. Until that point you should think carefully about logical precision.

One great thing about having precise meanings for logical terms is that it buys us a lot of mechanistic reasoning power. For instance, if we know that $A \Rightarrow B$ and that $B \Rightarrow C$, then we can deduce that $A \Rightarrow C$. We can do this even if A , B , and C are about really complicated objects that we've never met before and we don't understand. Similarly, if we would like to prove a statement of the form $A \Rightarrow B$, but we are not making much progress, we can remember that the contrapositive (not $B \Rightarrow$ no A) is always equivalent to the original, and try proving that instead.

This is what we mean when we say that we can develop valuable understanding by looking at the logical structure of a statement. If we pay attention to constructions involving "if" or quantifier, we can make use of such regularities in our reasoning. This is (at least partly) what people mean when they talk about *formal* work: we can concentrate on the logical form of a sentence and temporarily ignore its meaningful content. We don't *have to* ignore the meaning, of course, and for most of the above discussion you were probably thinking about meanings as well. But attending to logical form is vital for proper understanding. Some students are lax about this; when reading mathematics, they look mostly at the symbols, ignoring or glossing over the words. This can make their understanding faulty and their writing inaccurate, because they mix up important quantifiers or implications. For instance, consider the following theorem:

Rolle's Theorem: Suppose that $f : [a, b] \rightarrow \mathbb{R}$ is continuous on $[a, b]$ and differentiable on (a, b) and that $f(a) = f(b)$. Then $\exists c \in (a, b)$ such that $f'(c) = 0$.

It is quite common to see students make errors, writing things like this:

Rolle's Theorem: Suppose that $f : [a, b] \rightarrow \mathbb{R}$ is continuous on $[a, b]$ and differentiable on (a, b) . Then $f(a) = f(b)$ and $\exists c \in (a, b)$ such that $f'(c) = 0$.

We can see that these are different by looking at their logical forms: one of the premises in the correct version appears instead as part of the conclusion in the second (look carefully to make sure that you can see this). Clearly that must make a very important difference, so the incorrect version cannot possibly be logically equivalent to the correct version. Nonetheless, it might still be a valid theorem. In this case, however, it isn't, which we can see by thinking in terms of examples. In the incorrect version, the premises introduce a function f that is defined on an interval and is continuous and differentiable on this interval. The conclusion claims that the function values are equal at the endpoints of the interval, but this cannot possibly follow in a valid way from the premises, because there are many functions and intervals that satisfy the premises but do not have this property. For example, $f(x) = x^2$ is continuous on $[0,2]$ and differentiable on $(0,2)$, but certainly isn't the case that $f(0) = f(2)$. We can see how a student might write the incorrect version in the first place, but someone who is thinking about the meaning of their writing should recognize such errors when rereading.

This brings us back to the idea that both logical form and example objects can contribute to mathematical understanding, though focusing on each has different advantages and disadvantages. If you look mainly at examples, you might feel that you understand, but you might fail to appreciate the full generality of a statement or find it difficult to see the logical structure of a whole course. If you look mainly at formal arguments, you might be able to see how everything fits together logically, but you might find yourself complaining that it is very abstract and that you don't really understand what is going on. Your experience of these issues will probably vary from course to course, because some professors give lots of examples and draw lots of diagrams, whereas others give a much more formal presentation. If a professor's approach does not match your preferred way of developing understanding, you might find it useful to strengthen your understanding of the links between the example objects and the formal work.

Proofs

You have been constructing mathematical proofs for many years. For instance, to do the calculations needed to prove that the solutions to the equation $x^2 - 20x + 10 = 0$ are $10 + 3\sqrt{10}$ and $10 - 3\sqrt{10}$, you would write something like this:

$$x = \frac{20 \pm \sqrt{400 - 40}}{2} = \frac{20 \pm \sqrt{360}}{2} = \frac{20 \pm 6\sqrt{10}}{2} = 10 \pm 3\sqrt{10}.$$

This calculation uses methods that everyone agrees are valid, so it captures everything we need for a proof that the solutions really are as claimed. To make it look more like an upper-level proof, we could rewrite it like this:

Claim: If $x^2 - 20x + 10 = 0$ then $x = 10 + 3\sqrt{10}$ or $x = 10 - 3\sqrt{10}$.

Proof: Suppose that $x^2 - 20x + 10 = 0$. Then, using the quadratic formula,

$$x = \frac{20 \pm \sqrt{400 - 40}}{2} = 10 \pm 3\sqrt{10}.$$

This version explicitly states the claim, begins with the premise ("Suppose that $x^2 - 20x + 2 = 0$ "), and contains few words to justify those steps that are more sophisticated or less obvious.

My point is that there is nothing inherently mysterious about proofs. It is certainly true that most high school mathematics is not presented in this way, but it's also true that most of it could be. We say this because, in upper-level courses, lots of the mathematics you meet will be presented in this form; these lecture notes will be full of theorems and proofs. This might seem rather an abrupt change, and some students get the idea that proof writing is a mysterious dark art to which only the very privileged have access. It isn't. In cases like the one above, all it really involves is writing in a more mathematically professional way – writing less like a student and more like a textbook, if you like.

This is not to belittle the genuine difficulties that undergraduate students face when handling proofs. Obviously the example above is simple one, and the proofs you are asked to understand and construct in upper-level courses will often (though not always) be much harder. It might take you a while to get used to digesting mathematics presented in this way, and to writing your own mathematics more professionally. But there is no reason to think you won't manage it, and this section discusses some things you could pay attention to in order to get used to it quickly.

One thing you'll often asked to do is to prove that a mathematical object satisfies a definition. The question won't be phrased like that, though. It will just say something like, "Prove that the set $(2,5)$ is open." You will have to interpret this to mean "Prove that the set $(2,5)$ satisfies the definition of open set," and review the definition of open set to make sure you are clear about what this involves. That sounds simple, but We have often seen students who, faced with an instruction like "Prove that the set $(2,5)$ is open," don't know what to do. If you don't know how to start on any proof problem, your first thought should often be, what does the definition say?

We have seen the relevant definition, which says:

Definition: A set $X \subseteq \mathbb{R}$ is *open* if and only if $\forall x \in X, \exists d > 0$ such that $(x - d, x + d) \subseteq X$.

How should we write a proof that $(2,5)$ is an open set? Often the best advice is to follow the structure of the definition itself. We want to prove that $(2,5)$ is open, so we want to show that for every $x \in (2,5)$ there exists $d > 0$ such that $(x-d, x+d) \subseteq (2,5)$. When we want to show something is true for every x in some set, we usually start our proof by introducing one, like this:

Claim: $(2,5)$ is an open set.

Proof: Let $x \in (2,5)$ be arbitrary.

Here, *arbitrary* means that we are taking any old x , not some specific one or one with special properties. It is not necessary to write this – lots of people would just write “Let $x \in (2,5)$ ” – but it does emphasize that the ensuing argument will work for any x in the set.

Now we need to show the existence of an appropriate d . The simplest way to show the existence of something is to produce one. In this case, we need to produce a d that will work for our x . This d will depend on x , and one way to do it is to pick d to be the minimum of the two distances $5-x$ and $x-2$ (think about why). So we might write the rest of our proof like this:

Claim: $(2,5)$ is an open set.

Proof: Let $x \in (2,5)$ be arbitrary. Let d be the minimum of $5-x$ and $x-2$. Then $(x-d, x+d) \subseteq (2,5)$. So $\forall x \in (2,5), \exists d > 0$ such that $(x-d, x+d) \subseteq (2,5)$. So $(2,5)$ is an open set.

Notice that this proof reflects the order and structure of the definition. We are showing something for all x , so we start with an arbitrary one. We are showing that for this x , an appropriate d exists, so we produce one. Because of this, the structure of the proof will be obvious to a mathematician, so you don't really need to write the line “So $\forall x \in (2,5)$...”, but you might find it helpful for your own thinking.

Some people tend to include diagrams along with proofs like this, and some don't. That's because diagrams can be illuminating, but they are not necessary, and they are not a substitute for writing a proof out properly (diagrams are not proofs); mathematicians (in particular you algebra and calculus professors) want to see a written argument that is clearly linked to the appropriate definition.

Another standard proof type is known as *direct proof*. In a direct proof we start by assuming that the premise(s) hold and move, via a sequence of valid manipulations or logical deductions, to the desired conclusion. Here are some theorems for which we've already studied a direct proof:

Theorem: If n is an even number, then any integer multiple of n is even.

Theorem: If $x^2 - 20x + 10 = 0$ then $x = 10 + 3\sqrt{10}$ or $x = 10 - 3\sqrt{10}$.

Theorem: $(2,5)$ is an open set.

Stop and think for a moment here. Can you write out proofs of these statements without looking? If you can't do it immediately, can you remember the gist of how we proceed in each case and reconstruct the rest? If you give yourself a minute for each one, we bet you can remember more than you would initially have thought. Students often have too little faith in their own ability to recall mathematical ideas and reconstitute arguments around them.

One important thing to note is that *direct* describes the eventual proof we write, it does not necessarily describe the process of constructing the proof. You might be able to write down the premises and just follow your nose to a proof, but is more likely that you'll have to try out some of the things like: write everything in terms of definitions, think about some examples, maybe draw a diagram, and so on. You should then, however, work out how to write your final proof in a way that makes its logical structure clear for a reader. It is probably a good idea to treat this writing as a separate task, one that is worthy of your attention over and above simply getting to an answer or a proof.

A second thing to note is that direct proofs can have somewhat more complicated structures within them. The most obvious such structure occurs in a *proof by cases*, which means what it sounds like it means: we divide up the cases we're dealing with into sensible groups and work with each one separately within the main proof. Consider the following theorem:

Theorem: For every $x, y \in \mathbb{R}$, $\max\{x, y\} = \frac{x + y + |x - y|}{2}$.

Of course, we need to define the symbols that appear in the statement.

Definition: For every $x, y \in \mathbb{R}$,

$$\max\{x, y\} = \begin{cases} x, & \text{if } x \geq y, \\ y, & \text{if } y < x. \end{cases}$$

Definition: For every $x \in \mathbb{R}$,

$$|x| = \begin{cases} x, & \text{if } x \geq 0, \\ -x, & \text{if } x < 0. \end{cases}$$

The first definition gives the meaning to $\max\{x, y\}$ as "choose the largest number between x and y ", and the second definition is just the usual absolute value function. The theorem states that $\max$ can be expressed in terms

of the absolute value function. In order to prove the theorem, you need to reinterpret it as follows:

Theorem: For every $x, y \in \mathbb{R}$, the expression $\frac{x + y + |x - y|}{2}$ satisfies the definition of $\max\{x, y\}$.

Notice that the definition of $\max\{x, y\}$ is given in a piece-wise format, where two cases are obvious: $x \geq y$ and $y > x$. The absolute value function is defined in a piece-wise format as well (and probably it is the first time you see it written formally in this way). This should hint you into considering a proof by cases. Before reading the proof below, you should try to come up with one by yourself.

Proof: *Case 1:* Suppose that $x \geq y$. Then $x - y \geq 0$ and, thus, by the definition of the absolute value function, $|x - y| = x - y$. So,

$$\frac{x + y + |x - y|}{2} = \frac{x + y + x - y}{2} = \frac{2x}{2} = x.$$

Case 2: Suppose that $x < y$. Then $x - y < 0$ and, thus, by the definition of the absolute value function, $|x - y| = -(x - y) = y - x$. So,

$$\frac{x + y + |x - y|}{2} = \frac{x + y + y - x}{2} = \frac{2y}{2} = y.$$

Putting both cases together, we have that for all $x, y \in \mathbb{R}$,

$$\frac{x + y + |x - y|}{2} = \begin{cases} x, & \text{if } x \geq y, \\ y, & \text{if } y > x. \end{cases}$$

Therefore, $\frac{x + y + |x - y|}{2}$ satisfies the definition of $\max\{x, y\}$.

So when should you consider a proof by cases? Sometimes you will find that you have to. In this illustration, for instance, we don't have much choice; the $\max\{x, y\}$ values are different depending on whether $x \geq y$ or $x < y$, so we have to handle these cases separately. In other situations, it might not be necessary to use a proof by cases, but it might be convenient anyway because there is some sort of natural split, perhaps between positive and negative numbers (like in the definition of the absolute value function), or between odd and even ones. Finally, it might be worth starting a proof by cases if you think that you can construct an argument for some of the objects

to which the theorem applies but not for others. A start is better than nothing and, once you've got an argument written down for one case, you might find that reflecting on it gives you ideas about how to continue.

One final tip is that, if you have produced a proof by cases, or if you're looking at one produced by someone else, it might be a good idea to ask whether the number of cases could be reduced. You don't have to do this, of course – if your proof is valid as it is, that's fine. But remember that mathematicians also value elegance, and brevity contributes to elegance so it's a worthwhile aim.

The next type of proof we want to talk about is *proof by contradiction*. This is a type of indirect proof, so called because we do not proceed directly from the premises to the conclusion. Instead, we make a temporary assumption that our desired conclusion (or some part of it) is false, and we show that this leads us to a contradiction. From this we can deduce that the temporary assumption must have been wrong, and thus that the desired conclusion is true. This sounds rather convoluted, but you're accustomed to making this kind of argument informally in everyday life. Here's a simple case:

Your friend: Daniel was at home in Vicálvaro all weekend.

You: No he wasn't, I saw him in Alcorcón on Saturday afternoon.

Here you are implicitly using a proof by contradiction. The temporary assumption is that Daniel was in Vicálvaro. Your argument says that if we make that assumption, then we can deduce that he wasn't in Alcorcón (if you like, this uses the "theorem" that people can't be in two places at the same time). But this is contradicted by the fact that you saw him in Alcorcón. So the assumption that he was in Vicálvaro must have been wrong.

Next, we will look at a mathematical example, which involves the definition of *rational number*. The notation $\mathbb{Q}$ denotes the set of all rational numbers and $\mathbb{Z}$ denotes the set of integers. We introduce this definition here:

Definition: $x \in \mathbb{Q}$ if and only if $\exists p, q \in \mathbb{Z}$ (with $q \neq 0$) such that $x = p/q$.

The theorem and proof below use the symbol $\notin$ to mean "is not an element of," and in this case $y \notin \mathbb{Q}$ means that y is an *irrational number*. The theorem and proof both implicitly assumes that all the numbers we are working with are real (this is common in early work with rational and irrational numbers). As with any theorem and proof, you should read everything carefully, checking that you understand what is going on in each step.

Theorem: If $x \in \mathbb{Q}$ and $y \notin \mathbb{Q}$ then $x + y \notin \mathbb{Q}$.

Proof: Let $x \in \mathbb{Q}$, so $\exists p, q \in \mathbb{Z}$ (with $q \neq 0$) such that $x = p/q$. Let $y \notin \mathbb{Q}$. Suppose for contradiction that $x + y \in \mathbb{Q}$. This means that $\exists r, s \in \mathbb{Z}$ (with $s \neq 0$) such that $x + y = \frac{r}{s}$. But then

$$y = \frac{r}{s} - x = \frac{r}{s} - \frac{p}{q} = \frac{rq - ps}{sq}.$$

Now $rq - ps \in \mathbb{Z}$ and $sq \in \mathbb{Z}$ because $p, q, r, s \in \mathbb{Z}$. Also $sq \neq 0$ because $q \neq 0$ and $s \neq 0$. So $y \in \mathbb{Q}$. But this contradicts the theorem premise. So it must be the case that $x + y \notin \mathbb{Q}$.

In this proof, the temporary assumption is this one:

Suppose for contradiction that $x + y \in \mathbb{Q}$.

Making that temporary assumption leads us, by some sensible use of definitions and algebra, to this line:

So $y \in \mathbb{Q}$.

This (as stated) contradicts the theorem premise, so it allows us to conclude that our temporary assumption must have been wrong, like this:

So it must be the case that $x + y \notin \mathbb{Q}$.

It should be clear that in order to properly understand a proof like this, you have to do more than you might have had to do in earlier mathematics. In high school, most of your mathematical reading will have involved checking some algebra. Here, you have to be more sophisticated.

You certainly should check to make sure that you can see how the algebra works and that there are no mistakes. But that's not really where the action is in a proof like this. To fully understand it, you need to understand its global structure. You need to be able to identify what assumptions are made where, identify where the contradiction arises and what exactly is being contradicted, and understand how it all fits together to prove that the theorem is true.

The final standard proof type we want to discuss is *proof by induction*. Depending on your previous experience, you may have met this already. If so, you'll have used it to prove things like this:

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

If not, you might not have seen this notation so here is a quick explanation. Thy symbol $\sum$ is called "sigma" and is a Greek upper case letter Σ , used here to denote a sum from $i = 1$ to $i = n$. Written out, the left-hand side of the above means

$$\sum_{i=1}^n i^2 = 1^2 + 2^2 + 3^2 + \dots + n^2.$$

What we've done here is substitute $i = 1$, then $i = 2$, then $i = 3$, and so on, up to $i = n$, where we stop.

So our original expression actually captures infinitely many propositions:

$$\begin{aligned} P(1) 1^2 &= \frac{1(1+1)(2+1)}{6} \\ P(2) 1^2 + 2^2 &= \frac{2(2+1)(4+1)}{6} \\ P(3) 1^2 + 2^2 + 3^2 &= \frac{3(3+1)(6+1)}{6} \\ P(4) 1^2 + 2^2 + 3^2 + 4^2 &= \frac{4(4+1)(8+1)}{6}, \text{ and so on.} \end{aligned}$$

Having a theorem that captures infinitely many cases isn't unusual—many of the other theorems we've looked at do the same. The difference here is that the form of the statement allows us to put the propositions in an ordered list $P(1), P(2), P(3), P(4), \dots$

Proof by induction works as follows. First we prove $P(1)$. This is often easy. Then we do something clever. We don't try to prove any of the other propositions directly. Instead, we take a general number k and prove that if $P(k)$ is true, then $P(k+1)$ must be true too. This gives us $P(1) \Rightarrow P(2)$ and, since we've already proved $P(1)$, we can conclude that $P(2)$ is true as well. It also gives us $P(2) \Rightarrow P(3)$, so we can conclude that $P(3)$ is true as well. You get the idea. We have made an infinite chain of propositions, which are all true because the first one is true and the implications are all true:

$$P(1) \Rightarrow P(2) \Rightarrow P(3) \Rightarrow P(4) \Rightarrow \dots$$

Proof by induction is one of those ideas that students usually find intuitive straightforward when it is explained in the abstract. However, they often find it difficult to use in any particular case, so we will look closely at the example we started with. In that example, it is easy to prove that $P(1)$ is true:

$$1^2 = 1 \text{ and } 1 = \frac{1(1+1)(2+1)}{6}$$

Proving that the implication $P(k) \Rightarrow P(k+1)$ is a bit harder. We would like to assume that $P(k)$ is true and use this to prove $P(k+1)$ is true. We would start doing some rough work at this point, writing something like this:

Will assume $P(k)$, which means $\sum_{i=1}^k i^2 = \frac{k(k+1)(2k+1)}{6}$.

Want to prove $P(k+1)$, which means

$$\sum_{i=1}^{k+1} i^2 = \frac{(k+1)([k+1]+1)(2[k+1]+1)}{6},$$

which, by rewriting the left-hand side, means we want

$$\left(\sum_{i=1}^k i^2 \right) + (k+1)^2 = \frac{(k+1)([k+1]+1)(2[k+1]+1)}{6},$$

which, by the assumption about $P(k)$, means we want

$$\frac{k(k+1)(2k+1)}{6} + (k+1)^2 = \frac{(k+1)([k+1]+1)(2[k+1]+1)}{6}.$$

Then I've just got some algebra to do to show that the last two things are, in fact, equal (you might like to try it).

This, however, is definitely a situation in which the way you think about constructing a proof is not necessarily the same as the way you should write it out. The thinking above is completely logical, but presenting it like that wouldn't work very well because it doesn't match the structure of what we are trying to prove. When we write out a proof that $P(k) \Rightarrow P(k+1)$, we really want our proof to start with a clear assumption of $P(k)$ and proceed through some nice, tidy deduction to $P(k+1)$.

For proofs by induction, we favor a layout that makes that structure very clear. We would write something like this:

Theorem: $\forall n \in \mathbb{N}, \sum_{i=1}^n i^2 = \frac{n(n+1)(2n+1)}{6}$.

Proof (by induction): Let $P(n)$ be the statement $\sum_{i=1}^n i^2 = \frac{n(n+1)(2n+1)}{6}$.

Note that $1^2 = 1 = \frac{1(1+1)(2+1)}{6}$ so $P(1)$ is true.

Now let $k \in \mathbb{N}$ be arbitrary and assume that $P(k)$ is true, that is, that

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

Then

$$\begin{aligned} \sum_{i=1}^{k+1} i^2 &= \left(\sum_{i=1}^k i^2 \right) + (k+1)^2 \\ &= \frac{k(k+1)(2k+1)}{6} + (k+1)^2 \text{ by the assumption} \\ &= \frac{k(k+1)(2k+1) + 6(k+1)^2}{6} \\ &= \frac{(k+1)(2k^2 + 7k + 6)}{6} \\ &= \frac{(k+1)(k+2)(2k+3)}{6} \\ &= \frac{(k+1)([k+1]+1)(2[k+1]+1)}{6} \end{aligned}$$

So $\forall k \in \mathbb{N}, P(k) \Rightarrow P(k+1)$. Hence, by mathematical induction, $P(n)$ is true $\forall n \in \mathbb{N}$.

There are a couple of things to notice about this. First, the proof contains only a few words, but these help to make the structure clear. Second, all the algebra is in a single chain of equalities that starts with the left-hand side of the statement of $P(k+1)$ and ends with the right-hand side. You might like to think about why it makes sense to do the manipulations in this order, given that we know what we're going for. You might also like to think about why the last expression in the chain is not necessary but might be useful for a reader who wants to link the proof back to the theorem.

Confusion does tend to arise with proof by induction because there are lots of things to think about. We find that students are confused most often by the point which we write, "Assume that $P(k)$ is true." Students often read this and think, "But that's what we want to prove, how come we're allowed to assume it?" In fact, at that stage of the proof, we are not proving that $P(k)$ is true, we are proving that $P(k) \Rightarrow P(k+1)$. Make sure you can see the difference. Also, students sometimes confuse themselves, usually because they have allowed ambiguities to creep into their writing by using the word "it" in phrases like "so it is true for n ." There are many possible candidates for the

meaning of “it” in a typical proof by induction, so you should be more specific. Writing “So $P(n)$ is true” is one way to do that. (Notice that the proof above does not include the word “it” – we are very specific about what we have deduced at each stage.)

So when should you use proof by induction? In some cases this will be obvious, because you are likely to have a section about it in at least one course. Also you will come across cases in which you want to prove that something is true for all $n \in \mathbb{N}$, which is a giveaway that induction is worth a try. Be aware, however, that problems for which induction is useful can vary quite a bit. First, there is no particular reason for a proof by induction to start at $n = 1$. You might be asked to prove that something is true for every $n \in \mathbb{N}$ such that $n > 4$, for instance. In that case, you can just make $P(5)$ your base case and proceed as before, except that at some point, perhaps in the *induction step* (the point in the proof when you apply the assumption that $P(k)$ is true, referred to as the *induction hypothesis*), you will find that you need $n > 4$ to justify some manipulation you want to make. Second, while some of the first problems you meet will involve working with a sum, proof by induction is useful for many other types of problem. All we really need is a situation in which we have infinitely many statements that can be listed in the order of the natural numbers, which can happen in all sorts of ways. For instance, consider these tasks:

Prove that for every natural number $n > 10, 2^n > n^3$

Prove that for every $n \in \mathbb{N}, 5^{3n} + 2^{n+1}$ is divisible by 3.

For the first task, we would write

Let $P(n)$ be the statement that $2^n > n^3$.

Then we would prove directly that $P(11)$ is true, that is, that $2^{11} > 11^3$. Then we would work out how to prove that if $k > 10$ and $2^k > k^3$, then $2^{k+1} > (k+1)^3$.

For the second task, we would write

Let $P(n)$ be the statement that $5^{3n} + 2^{n+1}$ is divisible by 3.

Then we would prove directly that $P(1)$ is true, that is, that $5^3 + 2^2$ is divisible by 3. Then we would work out how to prove that if $5^{3k} + 2^{k+1}$ is divisible by 3, then $5^{3(k+1)} + 2^{(k+1)+1}$ is divisible by 3. Notice, in this case, that the statement $P(k)$ is **not** just “ $5^{3k} + 2^{k+1}$.” In fact, $5^{3k} + 2^{k+1}$ isn’t a statement at all – we couldn’t prove it because it’s just an expression (for any particular k it is a number; you can’t “prove” a number, and it makes no sense to say that one number implies another). The statement is “ $5^{3k} + 2^{k+1}$ is divisible by 3.”

In our experience, starting out with a clear statement of $P(n)$ often makes the difference between success and failure in constructing a proof by induction, especially when dealing with a new type of problem. This isn't that surprising, of course – before you start any problem, you should always make sure that you are clear about what you are trying to do. In any case, you won't always be told what method to use, so you should be on the lookout for less familiar cases like these, and you should train yourself to notice when proof by induction might be useful.

Some advice

Earlier we said that a student should never sit in from of a problem and think “I don't know what to do.” There are always things to try and, to be a good student, you must be willing to try them. In fact, you must be willing to try things that turn out not to work. In our experience, students are sometimes unwilling to do this, for three main reasons.

First, some students dislike the insecurity of not knowing exactly what to do. It makes them nervous. They want to know in advance what is going to work, and sometimes they ask for a teacher's assurance about this (“Is this the right way to do it”). The problem with seeking such assurance all the time is that you never find out what you could have done if you'd had a go, which means that you never get any more confident, which means that you end up in a vicious circle, having to ask for support all the time.

Second, some students do not want to waste time. We understand this – obviously no one wants to spend ages on one thing, especially when there are so many interesting things to do at university. But it is a big mistake to think that trying something that turns out not to work is a waste of time. Time spent learning is never wasted. If you try a method that doesn't work then, provided you are thoughtful about it, you learn *why* it doesn't work, which means you know something new about the applicability of the method. And you might gain some insight about the problem so that you have a better idea about what to try next. Of course, it is a mistake to keep plugging away at a method that clearly isn't working – research shows that good problem solvers stop frequently to re-evaluate whether their current approach seems to be getting them anywhere. But it's an even bigger mistake not to start.

Third, some students do not want to mess up their paper. They want to know that once they begin writing, they will be able to carry on writing and arrive at a nice, neat, correct solution. If this applies to you, then we're afraid you'll have to get over it. Real mathematical thinking is not tidy. It is full of false starts and partial attempts and realizations that what does not seem to be working just here would, in fact, form a useful part of a solution if put together with

something that failed then minutes ago, or yesterday, or last week. It is very important to embrace this if you want to keep improving as a mathematical problem solver. You need to get partial solution attempts on paper for the simple, practical reason that your brain cannot handle many things at once. You have an enormous amount of knowledge stored in what is known as your long-term memory, but your working memory, where you actually do the new thinking, has a seriously limited capacity. It will not be big enough to hold all the information about a complicated mathematical problem while simultaneously working out how to solve it. When you write down definitions or theorems or calculations that might help you solve a problem or construct a proof, you are using the paper to supplement your cognitive powers by, in effect, extending the capacity of your working memory. So don't be worried about writing things that are wrong or that turn out not to be useful. You can always write up a neat version of your solution or proof later.

Your professors will explain proofs that are long and logically complicated. They will also explain proofs that rely on some really clever insight. Sometimes they will explain proofs that are long and logically complicated *and* that rely on some really clever insight. This tends to worry students. They think, "Well, okay, I can see how that works, but *I would never have thought of it.*" This makes them wonder whether they're good enough at mathematics. But you shouldn't worry, because you're not supposed to be able to reinvent the whole of modern mathematics by having all the original ideas yourself. Even a mathematics PhD student wouldn't be expected to have many totally original insights. As an undergraduate, when faced with a proof like this, your job is to appreciate the clever insight, to understand why it works, to think about how modifications of it might work in slightly different circumstances, and to relate it to ideas used elsewhere in the course or in your major. To reassure you further, here is a list of things the you *will* be expected to do.

First, you will be expected to do routine mathematical calculations much like those you have seen in lower-level mathematics. As we said at the beginning of this introduction, you should be prepared for these calculations to be longer and more involved than those you have experienced before, and you should be prepared to have to adapt the calculation procedure if a step in it not valid for a new case.

Second, you will be expected to adapt proofs that you have seen to closely related cases. For instance:

- Having seen the proof that $(2,5)$ is an open set, you might be expected to prove that any interval of the form (a,b) is an open set.
- Having seen the proof that $\max\{x,y\} = \frac{x+y+|x-y|}{2}$, you might be expected to find a formula, using the absolute value, to express the function

$$\phi(x) = \begin{cases} x, & \text{if } x \geq 0, \\ 0, & \text{if } x < 0. \end{cases}$$

In such cases, you will often be able to treat the proof you have seen like a template, and change some numbers appropriately. However, to reiterate one of the main points of this introduction, you shouldn't do this thoughtlessly - you should make sure that each step in the proof really does work for the new cases, and be ready to make minor adjustments if it doesn't. It is important to be careful in cases where some number might be zero, for instance, or when dividing both sides of an inequality by a number that might be negative.

Third, you will be expected to adapt proofs you have seen to cases that are related, but not so closely. For instance:

- Having seen the proof that for all $n \in \mathbb{N}$, $\sum_{i=1}^n i^2 = \frac{n(n+1)(2n+1)}{6}$, you might

be expected to prove that $a^n - b^n = (a-b) \sum_{i=1}^n a^{n-i} b^{i-1}$.

- Having seen the proof that $(2,5)$ is an open set, you might be expected to prove that $A = \{x \in \mathbb{R} : x^2 - 1 < 0\}$ is open.

In cases like these, the proof you have seen will certainly be helpful, but you will not be able to treat like a template. You might be able to construct a proof that is very similar in its basic structure, but you will have to think further to work out exactly what needs changing.

Fourth, you will be expected to show that definitions are satisfied. You will sometimes be expected to do this with definitions you have not seen before, if your professor thinks that they are sufficiently straightforward.

Fifth, you will be expected to construct proofs of theorems for which you haven't seen a closely related model, or to solve problems for which you haven't seen your professor solve in class. As we've said, the biggest mistake you could make here would be to sit around thinking "We haven't been shown how to do this." No one will ask you to do things that are completely beyond you.

Sixth and finally, on an exam you might be asked to solve some challenging problems. Sometimes a question might lead you through a solution in steps, or might offer a fairly big hint to help you with a key idea or useful trick. Sometimes it might just ask outright, which means you'll have to be able to remember the key ideas or useful tricks yourself and come up with the rest. To do so you will need to have effectively read and understood the material in your lecture notes.