

[SQUEAKING]

[RUSTLING]

[CLICKING]

ZACHARY ABEL: Hello. Good morning. Good afternoon. Hope everyone's doing well. Welcome back to 6.1200. Today, we are going to be talking more about proofs. We spent all of last lecture going through 2/3 of our definition of a mathematical proof. And the last term we haven't talked about yet is logical deductions.

So we said that mathematical proofs are built from our verifications of propositions. We talked a lot about what propositions are using a sequence of logical deductions starting from a base set of axioms. And we talked about what axioms are. So finally, we get to talk about logical deductions. This is going to be the heart of how we build and construct and check our proofs. So this is where most of the proofing happens inside of our proofs.

So what is a logical deduction? Let's see. So an inference rule-- let's see-- is a rule for combining true propositions to form other true propositions.

For example, one of the classic ones, and one we've talked about briefly, is called modus ponens. And I'm writing this just to say let's not ever use those words again. The concept is intuitive enough. If we know P and we know P implies Q , if we know these, then we know Q .

Yeah. If we know P , and we know P implies Q , then we know Q . If I think, and thinking implies amming, then I am. Very intuitive. Let's again never mention this by name. That's way too precise for our purposes.

There are some other famous ones. Like if we know P implies Q , and we know not Q -- oh, sorry-- not Q . Then all of this implies, all of that implies not P . This is the opposite direction. If we know P implies Q , but Q is false, then P better be false as well.

Here's another fun one. If we know that not P implies false, then that implies P . Basically, saying if P isn't false, then it's true. Hopefully, intuitive. But we have the tools to not just rely on intuition here. We can verify these statements are actual inference rules, for example, by drawing a truth table. Let's do that last one just because it's smaller.

Let's see. So we have P , and it has two possible values. And then we have our big old formula. So not P implies false. Parens imply P . I want to make sure that this formula on the right is always true no matter what P is.

Well, a useful way to fill out these truth tables is to go one step at a time. And I like to do that by drawing directly under the symbol I'm working with. So like not P -- when P is true, not P is false. And when P is false, not P is true. And I put that right under the P bar, so I know I'm talking about the P bar.

Now let's talk about this implies. Does false imply false? Yeah. Yeah, that's true. Does true imply false? No, that's the case we need to avoid for implication

So finally, this implies is really this whole parenthetical clause there. Does true imply true? Because remember, this is true and false here. Does true imply true? Yeah. So this implies in this case is true.

Does false imply false? Also yes. OK. So the column that corresponds to the whole formula is always true no matter what. And that is how we verify that this statement right here is always true. This is one possible way to verify that.

These other inference rules, and lots of other ones that we're going to use in this class, you can also verify this way if you wish. Please don't do this in the middle of a proof. Let's just use these rules as rules that we're familiar with. You don't have to reprove them every time. I'm just showing you that we have the tools to do it if we wanted to.

In general, when writing proofs in this class, we're going to want each step to be clear and logical. We're going to want to use a valid inference rule. But we just want to use it.

It should be clear from your writing what rule is being used, but you don't have to call it out every time. Like, modus ponens, modus tollens. I didn't even tell you what that second one means. Just go ahead and use it.

We're going to have a familiar set of tools and deductions when writing these proofs. So just-- if you use those or you're doing something else that is clear enough, use that. And if it's less clear, explain it more.

Of course, you shall be explaining your steps. Usually, you're relying on facts that we proved before, or axioms that maybe we haven't explicitly written down yet. And you should be identifying those. When you're relying on something, tell us what you're relying on.

Let's see. Please don't use what's called proof by intimidation. Clearly, this is true. This is obvious. This is intuitive. In fact, I would strongly recommend striking those words from your mathematical vocabulary entirely.

Anything that is clear and obvious to you, well, first of all, it might not be clear and obvious to your readers who might be your graders, who might be your peers. And that can lose you points because you didn't explain yourself fully.

That can be demotivating for whoever's reading it. Like, wait, this isn't obvious to me. Am I not as smart as I thought I was? And, of course, it runs the risk that the thing you said is obvious is wrong. In fact, this is a common source of mistakes because instead of proving it, you just glossed over it.

The proof is there to help you check that things are correct and help you catch mistakes. And if you're skipping the proof, then you're skipping that redundancy. So things are never obvious. Things are never clear. Please explain it. Doo, doo, doo, doo, doo.

I mentioned last time that we're going to take basically all of high school math as our common set of axioms. The one exception to that is if we ask you to prove a theorem, and you say, I know this from high school math. It's an axiom. I'm done. That's maybe not adhering to the spirit of the assignment.

If we ask you to prove a statement that you already know is true, please go and prove it. We're asking you to think more deeply about why it's true. Does that make sense? Cool.

All right. So using these logical deductions and things like them, I want to talk about a couple very common proof outlines-- proof outlines-- oop, that's not a Q, that's a U-- outlines for proving some of the most common shapes that our theorems might take.

For example, if we have a theorem, there exists a natural number N , such that n is at least 10, and n is prime. Here's our theorem. There exists a number that's at least 10 and is prime. There exists a prime bigger than 9. Can someone prove this for me? Yes?

AUDIENCE: n is 17.

ZACHARY ABEL: 17. That's actually the same example in my notes. Well done. Proof-- n equals 17. Is that a full proof? I mean, kind of. Yes?

AUDIENCE: [INAUDIBLE].

ZACHARY ABEL: If we-- one more time.

AUDIENCE: [INAUDIBLE].

ZACHARY ABEL: OK. So the suggestion was we should be going through all of the possible factors of 17 and show that they're not factors to verify that 17 is prime. Absolutely. That is something we haven't done. Also, we haven't checked that explicitly. We probably want to do that.

But let me back up, and let me write down what the general proof usually looks like. Let's see-- proof. Oops. We'll show that n equals fill in the blank works.

This n is in the set we asked for. It's in the natural numbers. In this case, we're going to secretly be using 17. So it's pretty clear. But in general, it might be less clear. So for reasons. And n is prime because reasons.

So this is in general, how will we recommend writing a proof of an exists? If you want to show something exists in a set with a property, you tell us what it is. You show us why it's in the set. You show us why it has the property. Pretty intuitive, right? Hopefully. I know I just used that word that I said we aren't supposed to use.

But this is inextricably tied with what this symbol means in general. There exists a natural number. To prove that there exists a number with this property, you have to show me a number with this property. That's the most straightforward way to do it. And that's what this proof outline is letting us do.

In this particular problem, we're going to say, we'll show n equals 17 works. This number, 17, is a natural number because well, yeah-- I don't know how to convince you of that if you're not sure of that. And n equals 17 is prime because-- and this is where we should, if we want to be super careful, check all the possible factors. It's not divisible by 2, by 3, by 4, et cetera. Eric?

AUDIENCE: Are you missing n greater or equal to 10?

ZACHARY ABEL: I am missing n greater or equal to 10. Yeah, we need to show that n satisfies this whole property, and n is prime and greater or equal to 10 because reasons.

All right. So more abstractly, if our theorem is there exists x in some set such that P of x holds, then our proof can often look like choose X equals some specific value that you have to put in. Then x is in S because big to-do there. These are to-dos. And P of x is true because to-do. All right.

Similarly, what if we're asked to prove a for all theorem? For all numbers, x in the reals-- so for all real numbers x . x squared minus $6x$ is greater than negative 10.

So this theorem is about a for all, not in exists. When we're asking about exist, we need just a single example. But when we're asking about for all, examples aren't enough. Remember, last time when we talked about n squared plus n plus 41 being prime lots of times. Well, it's not prime every time. If we're trying to prove it for all, we need to show that the property is true for every possible choice of x .

And, in general, let's see what this outline might look like. So if our theorem is for all x in some set, P of x is true, our proof usually looks like this.

Proof-- suppose x is a generic element of S . Then P of x is true because reasons. I think it makes sense to put little question marks in there. So we know that we have to figure out what's going on. There we go.

This first line-- suppose x is some element of S . Which element? We're not allowed to know. We're not allowed to choose a particular one. We need to prove that this is true for every possible choice of x , which means in our proof, the only thing we assume about x is that it's in S .

So you start by taking some thing in S and assuming nothing else about it, and then going and concluding the property that we need to show. Again, these might look very similar to you, and they are.

This says everything in S has this thing. This says a generic thing in S has this property. The difference is that suppose x is a generic element of S , or however you want to phrase that, makes it concrete. It gives us a particular x to talk about. So now we have something to use and talk about and reason about inside of our proof.

Whereas, up here, how do we approach every-- that's a little daunting. I only know how to deal with one. So this gives us one. It's called x . We know nothing about it except that it's in S , but at least we have the one, and we can work with it.

So whenever you see a theorem that starts with a for all, your first line is almost always going to be this. In order to bring some thing into scope that we can talk about. Sound good? Awesome.

So how do we prove this theorem? Proof-- let's see. Let's see. Suppose x is an element of the real numbers. Then, let's see. x minus 3 squared is greater or equal to 0 because all reals have non-negative squares.

This can be a fact we're citing from high school. We know how the real numbers multiply to each other. And once we know this-- let's see-- equivalently, this thing says that x squared minus $6x$ is greater or equal to minus 9, which, in turn, is greater or equal to minus 10 as needed.

So we started with pick some element of R without specifying which one. We ended with the property we need is true-- QED. Sound good? QED, by the way, is a common way to mark the end of a proof. It's Latin for quod erat demonstrandum, which was to be shown. So it's a synonym for as needed or as desired. QED-- all the same thing. All right. Question?

AUDIENCE: Could you like show that, like, for all the negative numbers, [INAUDIBLE] all the negative numbers will not [INAUDIBLE]?

ZACHARY ABEL: Sure. So the question was, could we do this a different way? Could we look at all the negative numbers, and they behave one way. And then look at all the numbers between 0 and 3 or 0 and 6 or something, and they behave a different way. And then look at all the rest. Yes, absolutely. we can do something like that. That's called proof by cases. We're going to talk about it in next lecture.

But for now, let's move on to the next common outline, which is proving implications. If we have a theorem that says P implies Q . I mentioned, we're going to be using implies all the time. And often, theorems say, if this, then that. If P , then Q . How do we prove this?

Our proof often looks like assume P , then Q is true because reasons. Again, looks pretty similar to the kind of thing we're talking about. But this first sentence, assume P -- assume P is true-- makes things concrete. It instantiates the thing we're supposed to be assuming. It gives us something concrete to work with.

We can now live in this hypothetical where we know for certain that P is true, and then work with it as if it's true because inside of this proof, for the rest of this proof, it is. Assume P . Use that to show Q . Whenever we're trying to prove an implication, this is a common way to do it. This is called a direct proof.

We're going to see indirect proofs a bit later. But there is another common way-- actually, sorry. Before that, let's do an example of proving something in this fashion.

Theorem-- if n is a multiple of 10, then it is even. If n is a multiple of 10, then it is even. We're proving an implication. n multiple of 10 implies even. P implies Q .

So even without having to think very hard, proof-- assume n is a multiple of 10. We know we're going to start with that. Assume the first thing. We know we're going to end with the last thing. Therefore, n is even.

So question mark, question mark, question mark. Therefore, n is even. And if we can connect the dots from the first line to the last line, we'll have our proof.

Notice, by the way, that was kind of an automatic process. If the theorem is an implication, then this is what our proof is commonly going to look like. So we can just write it down. And now we have a clearer, more concrete sense of the task we're supposed to do.

Starting from the fact that n is even, we're supposed to derive the fact that it's-- sorry. Starting from the fact that it's a multiple of 10, derive the fact that it's even. So let's go ahead and do that.

Assume n is a multiple of 10. What does it mean to be a multiple of 10? So n is 10 times k for some integer k . But then n equals 2 times 5 times k . And so it's 2 times an integer. Therefore, it's even. There's our proof. We're done. Just by identifying where we're supposed to start and where we're supposed to end, we were able to bridge the gap. Sound good? Nice.

So there's another common way of proving implications. This is called proof by contrapositive. Recall last lecture, we said that P -- oops-- P implies Q is equivalent to its contrapositive. Let's see. Not Q implies not P . This right here is the contrapositive-- contrapositive.

And so if we're supposed to be proving P implies Q , sometimes it's more convenient to instead prove not Q implies not P . And so our proof outline would say proof by contrapositive.

So assume Q is false. Assume not Q . Let's see. Assume Q is false. Then P is false because question mark. Same as we did for direct implications, but we're proving this implication instead. So assume not Q . Use it to prove not P . And let's see an example of this.

Theorem-- let's see. So for all integers n -- let's see. If n squared is even, then n is even. For all natural numbers n , if n squared is even, then n is even.

Now automatically, oh, we're proving it for all. I know how to write that proof. Proof-- suppose n is an integer. Now we have an n to work with. We need to prove this implication.

Let's see. So proof by contrapositive. It's always nice to say what proof technique you're using, if it's maybe not as self-evident as it might be. So now we're indicating we're proving this implication by contrapositive. And once we've said that, both you, the writer, and your audience, the readers, know what's supposed to come next.

Proof by contrapositive-- so assume this thing on the right is false. Assume n is odd because that's the opposite of n being even. We want to show n squared is not even. n squared is odd.

This, by the way, WTS-- Want to Show. This, I find, is really useful for me when planning out proofs, thinking about proofs, sometimes even when writing proofs. This is an indicator that this is where I'm trying to go. This is the end of our proof. This is my goal.

And it signals both to you that this is work that's left to do. It signals to your reader that, OK, great. We agree that we're working towards the right thing. So now let's go and do it.

And also, as you're writing scratch work and solving these proofs, and maybe you get distracted and have to come back later, and you see a bunch of statements on your page. Wait, which of these statements did we prove, and which are we trying to work towards?

So as I'm working on these proofs, I really like to be organized and say, we know this. We know that. We want to show this other thing. So I know that we haven't proven it yet, and we want to.

So proof by contrapositive. So assume n is odd. We want to show that n squared is odd because that's what the proof method tells us to do. And now we can go do it.

So assume n is odd. So n is $2k$ plus 1. That's what it means to be odd for some integer k . Then n squared is $4k$ squared plus $4k$ plus 1, which is 2 times an integer plus 1, which is odd. QED. And just to be precise about it, $2k$ -- oops, sorry-- $2k$ squared plus $2k$. But 2 times an integer plus 1 is what it means to be odd. So we're done.

So we have proven implication. We've shown two different proof outlines for how to show P implies Q . One direct-- you assume P . Use it to conclude Q . One flipped around. Assume not Q . Use it to derive not P .

There's a related proof technique called proof by contradiction.

Now, when growing up, you might have been told, don't use double negatives. They're not as clear. It's always much stronger to more directly say what you mean. But sometimes in math, a double negative isn't what we don't want to avoid. I think there were four negatives in there. Got it right. Yeah.

Double negatives can sometimes work to our favor. And proof by contradiction is exactly that technique as a proof outline. So let's talk about what this does-- proof by contradiction.

So theorem P. Whatever P is, the proof by contradiction can be used to prove anything you want sometimes. Sometimes it works. Sometimes it doesn't. Sometimes it's nicer. But it's not assuming a specific form of theorem.

Our proof looks like this. Let's see. For the sake of contradiction, assume P is false. All right. Assume P is false. Then the proof continues. Then some other statement R. Any statement R we want we get to choose. Then R is both true and false, which is bad. That breaks math.

Propositions are supposed to be true or false, not both. That's a sign that our set of axioms is inconsistent. And remember, that's the one thing we don't want them to be. If something's true and false, then everything is true and everything is false, and that's bad. Let's not break math, which is a contradiction.

And then to explain it further. So our assumption is wrong, so P is true. So we started by assuming P is false. We broke math. So the assumption that P is false must be wrong. So if P is false, it isn't true.

Now, I put this last sentence in parentheses because it's clear. It's clearer than if you leave it out. But if you're writing between people that all agree on how proof by contradiction works and understand what's going on, then just including this word contradiction at the beginning and saying what the contradiction that we've reached is. That's usually enough.

You don't have to have that summary statement at the end if you don't want it. That's a matter of personal style. But you should be saying right at the beginning that you're doing a proof by contradiction, and you should say at the end what your contradiction is. Does that make sense? Cool. Yes?

AUDIENCE: [INAUDIBLE]. Are you really just saying that [INAUDIBLE]?

ZACHARY ABEL: Yeah. That's exactly right. This inference rule that we analyzed at the beginning is exactly what's justifying proof by contradiction as a proof outline. If not, P implies that false is true, then P has to be true. That's exactly what this is saying. And that's what the proof technique is doing. Very well-spotted. Thank you. I meant to say that and might have forgotten. So thank you for the reminder.

A couple other things about this. For the sake of contradiction, there are lots of other ways to say that. You can just say proof by contradiction, and that's enough to signal to the reader that you're doing a proof by contradiction.

I've also seen BWOC, By Way of Contradiction, assume P is false. I have seen that. I like it. I don't use it as much because I don't think it's as standard. But it's silly, and I like it.

At the end, when you find your contradiction, and you want to say, this is a contradiction. We've broken math. We're done-- however you want to say that part. I've also seen this symbol-- equals x equals because it looks like two arrows pointing at each other contradicting each other. That's another fun embellishment if you want.

But now that we've seen the proof outline, hey, that looks like a smiley face. Now that we've seen the proof outline, let's use it on a real theorem. And that is that the square root of 2 is not a rational number.

So remember, this blackboard bold $\mathbb{Q}$ is the set of rational numbers of integer divided by nonzero integer-- the set of all such things. And the theorem is that square root of 2 is not rational. So let's prove it.

Often, when your theorem has a not on the outside, it is not true that property. Contradiction is a common way to handle that. Because our proof by contradiction is going to say, well, proof by contradiction, assume--

And, again, I want to emphasize. The moment we say this is a proof by contradiction, what we assume is fixed. There's only one thing we're supposed to assume. The outline tells us what we have to assume, which is that the entire theorem is false. That's the only thing we're supposed to do.

We should still say it because we want you and the reader to agree that the correct procedure is being followed. You're assuming the correct thing. So assume that this is false. In other words, assume that square root of 2 is a rational number. Ah-- is, not isn't. Do what I say, not what I write.

And now let's keep going. What does it mean for square root of 2 to be a rational number? So square root of 2 equals a over b for some integers a and b , where b does not equal 0. And as we might recall from high school math, we can always assume that a fraction is in lowest terms-- in lowest terms. So a and b have no common divisors.

We can always reduce a fraction to lowest terms. So let's take that fraction already written in lowest terms, and let's keep going from there. So a over b equals square root of 2 means a equals square root of 2 times b , which means a squared equals $2b$ squared, which means a squared is even.

And we already proved something about if the square of a number is even, then the number itself is even. It's very convenient that we already have that there. So since a squared is even, this theorem shows us that a is even by the theorem over there on the previous board. Yeah? OK.

What does it mean for a to be even? So a equals 2 times some integer c . For some c -- oops, sorry. For some integer-- I did it again-- c in the integers. That's what it means for a to be even. It's 2 times an integer. So let's keep going. Let's put this back into our formula.

So now, this says that $2c$ squared equals $2b$ squared. So this says that $4c$ squared equals $2b$ squared. We can divide by 2. This says $2c$ squared equals b squared. But now b squared is even.

And as we know by our Lemma over there again, b is now even by the same theorem. Yeah? OK. So a and b are now both even. So they have a common factor of 2.

To be more precise, no common divisor is greater than 1. And if a and b are in lowest terms, then there's no integer bigger than 1 that divides both of them. But we just found that 2 divides both of them. So this is our contradiction. This statement here contradicts this statement here.

We have a single statement that we showed is both true and false. So we've reached a contradiction, and our proof is done. Does that make sense? Awesome. Very nice, very nice.

Let me take one more minute to talk about these proof outline techniques in general. Now, these aren't meant to restrict the way you're allowed to prove things or restrict the way you're allowed to express things. They're there as scaffolds to make sure you're proving the right things and to direct your thinking.

The more you can break down a proof into smaller and smaller to-dos, first of all, the easier it is to make sure that your proof is correct. And secondly, the easier it's going to be to think about those individual tasks. Smaller tasks are easier to handle than larger tasks.

Also, a common mistake when writing proofs in general is when you're presented with a theorem to dive right in. Why is this true? How am I going to figure this out without taking the time to realize that actually, what's going on in your head, what you're trying to prove is not actually what the theorem is asking you to prove. And so that's just a bunch of wasted effort and lost points on your homework.

If you had taken a few minutes, or with practice, a few seconds to do these standard outlining procedures to figure out what the structure of the proof is supposed to be, then now you know what the structure of the proof is supposed to be. And you can spend your precious thinking cycles thinking about the things that you're actually supposed to prove.

And let me demonstrate this, that with proof you can make this automatic. You can make this fast. You can really use it to help inform your proof writing endeavors.

So let's say we have a theorem. Let's see. So an integer n is fooish precisely when n plus 1 is barsome. What do these underlined words mean? Absolutely nothing. But let's imagine they actually mean something. Let's imagine this theorem is meaningful, and we've been asked to prove it.

I claim there's a lot of progress we can make in writing and structuring this proof, even before we go and look up what these two weird terms mean. So let's do that. So this theorem, we can phrase it as for all n in $\mathbb{Z}$, n is fooish if and only if n plus 1 is barsome. Mm. Excuse me.

Now, this-- sorry, I prefer to write it like this. F_n , if and only if B_{n+1} . I don't remember if we talked about if and only if on Tuesday. Great. Let's talk about it now.

So P if and only if Q . also written as P if and only if Q means that P and Q are both true or both false. P and Q both true or both false. It means they're expressing the same condition. They're true in exactly the same times. They're false in exactly the same times. So they're logically equivalent.

One way you can define this-- we could just draw the truth table like we did last time. But also commonly-- oops-- this can be defined as P implies Q and Q implies P . If P is true, then Q is true. And if Q is true, then P is true. And this, if and only if, implies in both directions this concept is what we mean when we say precisely when. These are true or false together. Sound good? OK.

So we have our theorem. Let's prove it. Proof-- OK. The theorem starts with a for all. I know how to handle a for all. Suppose n is an integer. We want to show F_n , if and only if, B_{n+1} . That was one of the first outlines we talked about. If you have a for all, you introduce a generic one, and you want to show the property for that n .

Well, now we want to show an if and only if, which we know is defined like this. So really, instead of all of that, we can be a little more precise with it. We want to show that F_n implies $B_n + 1$, and that $B_n + 1$ implies F_n . And once we've done both of those tasks, we will have proved the if and only if.

Well, now we're trying to prove implications. We know how to prove implications. Let's use a direct method for each of these. So assume F_n is true. Now we want to show $B_n + 1$. So these are separate sections in our proof.

In this part, we are temporarily assuming that F_n is true and using that to prove that $B_n + 1$ is true. And now down here, assume instead that $B_n + 1$ is true. And now we want to show F_n is true.

And now our tasks are very clear to us. We need to show that if F_n , then $B_n + 1$. We need to show-- in addition to that, we need to show that if $B_n + 1$, then F_n .

And now at this point, there's not really much else we can do without going to look up what these terms mean, use their definitions, figure out why this thing is true. But this is where the interesting part of the proof happens, where the unique part of the proof happens, the part that's specific to this problem. Yeah.

And that's a lot of progress that we made even before the theorem even makes sense. So once you get better at this task, the easier it will be to set things up correctly and set yourself up for success. Questions on that? Awesome.

All right. So the other main topic for today, we talked a lot about proof outlines, proof by contradiction. The other main topic is proof by induction. And interesting, I didn't I didn't hear the groan that I usually get.

Proof by induction is one of those topics that some students come in with the impression that it's difficult, that it's hard to grasp. And there are reasons for that. Even before telling you what induction is, I can tell you that it's kind of a shame that we have to teach you induction so early.

Because we're now in the process of telling you that proofs are supposed to be clear and self-evident and easy to follow. And then we come to proof by induction, which has a weird outline that has lots of things you have to say in weird orders for some arcane reason, like we're summoning an ancient demon. Why do we have to say this in that particular way, in this place? I don't know. It's confusing.

And it is confusing. And there's a reason it's confusing that the reason induction works, and the reason that the proof is structured the way it's structured isn't explained in every proof by induction because everyone who's writing a proof by induction assumes that everyone who's reading their proof by induction knows what proof by induction is and why it works.

So every proof by induction you see is completely opaque and makes absolutely no sense until you know what induction is. And I think that's why it can be difficult to approach.

What we're going to talk about right now is-- first of all, before I tell you what proof by induction is, let's solve a problem together that's going to lead us to the why and how and what of induction. It's going to tell us what we need to do and what tool we need to do it. And then induction will be there right at the end to say I'm the tool for the job.

So here is a theorem. Theorem-- for all n and z -- sorry. For all natural numbers n $1 + 2 + 3 + \dots + n$ equals n times $n + 1$ divided by 2. This might be familiar to some of you. There are lots of ways to prove it.

You can prove it by induction. You can prove it by picture. You can prove it by counting. There are lots and lots of ways. We're going to talk through one particular way of making sense of this theorem. And let's just go by intuition. We're not doing induction right now. We're just we're just figuring out why this thing is true.

So, first of all, let's make sure we understand what this thing is even saying. For example, so e.g., when n equals 4, this is claiming that $1 + 2 + 3 + 4$ equals 4×5 over 2, which looks pretty good. This is 10. This is 20 over 2. Yeah, they're equal.

And we're supposed to prove this kind of thing for every natural number n . Yeah? Even all the way down to just n equals 1-- e.g., n equals 1.

What's the sum from 1 up to 1? Just the 1 itself. 1 is the sum of one term. And that's supposed to equal 1×2 over 2, which looks like it does, even down to n equals 0.

What is the sum from 1 up to zero? Oh, that's weird. That is a little weird. Turns out, we can make sense of it in this context. So another way to write this is with what's called sigma notation. k goes from 1 up to n of k .

And you can think of this like a for loop. You start at k equals 1, and you write down the value of k . In that case, 1. Then you go up to k equals 2, and write down the value, k equals 3, write down the value, all the way up to this part. You go up to k equals n . Write down the value.

And the sigma says you add all these together. Sigma is the Greek S for sum. So this is notation for exactly this sum from 1 to n . And by convention, and there are good reasons for this convention.

But by convention, the sum from k equals 1 to 0, the sum of 0 terms of whatever you want in here, but now, we're talking about the sum k , is just zero. The empty sum is zero.

If you're adding nothing, you have a total of zero. And, in that case, the empty sum is zero, which is supposed to equal 0×1 divided by 2. And it does.

Hooray. This theorem is true when n equals 0. And 1 and 4-- we verified three cases, which is a pretty good 0% of all the cases we're trying to show. But it's a good start. So let's see how to continue.

So we are trying to show that 0 equals 0×1 divided by 2. 1 equals 1×2 divided by 2. $1 + 2$ is 2×3 divided by 2. $1 + 2 + 3$ is 3×4 divided by 2. I'll stop in a minute, I promise. $1 + 2 + 3 + 4$ equals 4×5 divided by 2. Let's stop there for now. But we're supposed to prove all of these statements. That's what the theorem is asking from us.

And we can think about this a little bit. If we stare at this, we can notice that on the left side, to get from this to this, we're adding 1. To get from this to this, we're adding 2. It's right there in how this sum is created.

To get from here to here, we're adding 3. Again, we're just adding one more term to the sum. Then we're adding 4. Then we're adding 5, then 6, then 7, and then 8, and so on. There's a very simple pattern that explains how these numbers are changing from one row to the next.

So wouldn't it be convenient if the numbers on the right side followed that same pattern? And if they did, it would make a lot of sense, at least to me, that since they start at the same number, and we're making the same changes to both sides every time we move down a row, that all the rows are going to be equal, that the left side is going to be the right side in every row. Yeah?

So we're hoping that that's a plus 1, and that's a plus 2, and that's a plus 3, and that's a plus 4, and so on. These are question marks because we don't really know whether that's true yet. So let's write down what we're hoping is true.

1 plus 2 plus up to n is supposed to equal n times n plus 1 over 2. And then the next row-- 1 plus 2 plus n plus the next term, n plus 1, is supposed to equal n plus 1 times n plus 2 divided by 2. These are two consecutive rows in our big list. And what we are saying is to get from here to there, we're definitely adding n plus 1. We're just adding this one more term over here.

So is it true that to get from here to here, we're adding n plus 1? And we can check that. Let's see. n plus 1 times n plus 2 divided by 2 minus n times n plus 1 divided by 2.

Well, we can factor in n plus 1 from both sides. That's n plus 1 times n plus 2 over 2 minus n over 2. That's looking pretty good. That's going to be n plus 1 times 2 over 2 times 1. That's just n plus 1. So our hunch about that pattern is true. To go from the n -th row to the n plus first row, we're adding n plus 1. Yeah?

So what we said earlier is true. This pattern add 1, then 2, then 3, then 4, then 5 that we saw on the left is recreated on the right. Add 1, add 2, add 3, add 4, add 5, and so on. So at least to me, this feels like an explanation that should make sense, and an explanation that we should be able to write down somehow as a mathematical proof.

Does anyone have questions about this idea? I haven't said anything about how we're going to write this. But is the idea clear? Awesome. Let's see.

So let's dig in a little more and see what exactly we accomplished over there.

So this argument here on the right that we did over here, let me write down a little more carefully what we accomplished. And it's going to look like this. If 1 plus 2 plus n equals n times n plus 1 over 2, if that's true, then 1 plus 2 plus n plus 1 equals n plus 1, n plus 2 over 2 because we just added n plus 1 to both sides.

You would also want the algebra work to show that going from this to this really is just adding n plus 1, and that's exactly what we did over there. I don't want to write it again, so I just wrote the word algebra.

Does it make sense that this is what we accomplished? If this row of our list is true, if left side equals right side, then the next row is also true. Yeah? Cool.

So let's put some notation to this. I'm going to say that P of n is the statement. By the way, this colon equals is a way that we'll commonly say P of n is defined to be this thing. So not only is it equal to this, but this line right here is the definition of P of n . So let's define P of n to be the statement that 1 plus 2 plus up to n equals n times n plus 1 over 2.

So P_n is a predicate. It's not a theorem on its own. What is a theorem is that for all n greater equals 0, P of n is true. That's what we're being asked to prove. That's just a restatement of our theorem from the top board over there.

And now what we accomplished over on that board-- actually, I want a fresh board for this. That's going to be this board up here.

What we accomplished is, first of all, we proved P of 0. That's the first line up here. And then these statements up here, what that's saying is that P_n implies P_{n+1} . So we proved P_0 implies P_1 . We proved P_1 implies P_2 , P_2 implies P_3 , P_3 -- I'll stop, I promise-- implies P_4 , and all the way down. This is what we really accomplished on these last couple boards.

But this isn't the theorem we're aiming for, right? The theorem we're aiming for asks for P of 0-- oops-- and P of 1, and P of 2, and P of 3, and P of 4, and all the way down-- we want to know that all of these are true. These implications are not what we're asking for.

However, it kind of feels like this should be enough, right? Because once we know P of 0, and we know P of 0 implies P of 1, then modus ponens, we know P of 1. And once we know P of 1, this implication tells us P of 2. Once we know P of 2, pair with this, we get P of 3. Then we get P of 4. Then all the way down. We get all of them.

I don't know how to write that yet, but that's the intuition that I have guiding this. And so it really feels like if we accomplish the left side, we should be able to conclude the right side. Do we agree? All right.

And now we have the principle of induction, which just says what we wanted to say. The left-hand side of this table implies the right-hand side of this table. If you know all of these, then you conclude all of these.

So when we're structuring a proof by induction, we must prove everything here on the left. That is our responsibility. And then everything on the right, we get for free by the principle of induction.

To be more precise the principle of induction says, if you know P of 0 and for all n greater or equal to 0, P of n implies P of $n+1$, that implies for all n greater equals 0, P of n . So that's exactly what's written in this table just in formula form.

P of 0 is this first one over here. And then all of the P_n 's imply P_{n+1} s are all of the remaining stuff. And it says, if we accomplish all of these, then we can conclude these. Yes?

AUDIENCE: Is the principle of induction just not like the axiom [INAUDIBLE]?

ZACHARY ABEL: Good question. So is the principle of induction a new axiom, or is it just a name for this intuitive fact that we have? In truth, what it's really an axiom about is how the natural numbers are constructed.

Really, what it's saying is that if you start from 0, and then add 1, add 1, add 1, add 1, add 1 to infinity, you will hit all the natural numbers. That's one way of thinking about it. But in practice, when we use the principle of induction, we're using it like this.

So yeah, you can take it as an axiom. You can take it as following from the definition of the natural numbers. We didn't really go that far down in defining our fundamentals. So we'll just leave it like this. Good question.

All right. So we have the principle of induction. We know what it's accomplishing and why. So how do we actually write using the principle of induction? Let's actually write this proof now.

And the thing to always keep in mind when writing a proof by induction is this table. We are responsible for proving everything on the left side, and then we're able to conclude everything on the right side. So the theorem that we're trying to prove should be the right side. And then we're going to do the left side to get it.

So proof by induction-- we need to say what P of n -- let's see. P of n is defined as the sum and n plus 1 over 2. Same way we defined P_n before. And we'll show that P_n is true for every n greater or equal to 0 by induction.

So the first thing we're responsible for is P_0 . That's usually called the base case. Base case-- want to show P of 0. Well, P of 0 says that 0 equals 0 times 1 over 2, which is true. So we verified P of 0. Usually, the base cases are pretty straightforward. It's the inductive step-- inductive step-- that we need to watch out for.

We now need to prove this whole thing. And we usually write it like this. Assume n is greater or equal to 0. And assume P_n is true. We want to show P_{n+1} is true.

And notice, what we just wrote is exactly the unpacked outline for this theorem. It starts with for all. So assume n . It then has an implication. So assume the left side. Let's prove the right side. The things we write when we're doing a proof by induction are not magic. They're just unpacking this. And now we can do it.

The next thing I always recommend when proving something by induction-- what is P of n ? What is P of $n+1$? I know we defined it, but I can't remember what it is. I always find it much more useful to write these things down in context.

So, in other words, assume-- we're supposed to assume P of n . So assume 1 plus 2 plus up to n equals n times n plus 1 over 2. And we want to show P_{n+1} , which says that the next row, $n+1$ times $n+2$ over 2.

All right. And just by following the proof outline and by identifying the things we're supposed to identify, now we have a very clear task in front of us. Assume this row. Use it to prove this next row. And we've already done that. We know how to do that.

We know that this thing equals that thing. And then this plus that equals that because of algebra. This is all the work we already did. And now we're seeing how the induction outline leads us to figuring out that that is the right work to do. Did that all make sense? Yes?

AUDIENCE: Does this necessarily has to start at 0? Like, if it started [INAUDIBLE].

ZACHARY ABEL: Great question. Does this have to start at 0, or could it start somewhere else? Yeah. It can start wherever we want. If we didn't know that P of 0 really made sense, and we wanted to start at P of 1, we absolutely could have started at P of 1 as our base case. And then all the implications above that.

Just keep in mind that by not mentioning P of 0, we didn't prove P of 0, and so we can't conclude it either. But wherever we started and all the implications after it, those fit together to give us the thing on the right side. Yes, that's a slightly more general form of the induction principle. Please go ahead and use that version as well. All right.

So let's do one more proof by induction. This is a really pretty problem. This seems to be a mandatory problem for any textbook on discrete math. It's one of my favorites. And here is the problem.

Which board did I leave on? This one.

So the story is like this. We have a big old courtyard subdivided into lots and lots of tiny little cells. It's going to be 2^n tall and 2^n wide. So imagine 64×64 . It's a big old checkerboard. I'm just going to draw 8×8 because I don't have that kind of time.

And now we're going to take a single cell and get rid of it. Which cell? I don't know. Doesn't matter. The question is, can we then fill the remaining cells with these L trominoes. You might know the word domino. It's two squares next to each other. It's a fun game.

This is like a domino, but it's 3 squares instead of 2, and it's in the shape of an L. So it's called the L tromino. And the question is, can we fill all of this board without that one missing cell using L trominoes and nothing else? So that is our theorem.

Theorem-- a $2^n \times 2^n$ board without one cell, without any one cell can be tiled with L trominoes. So that's what we're supposed to prove. And let's see if I can set this up quickly. Yeah, that's all right. Let's see if I can get the glare a little different. OK.

I do want to plug-in. Thank you. I always forget this in Zoom meetings, too. Please share your screen. Yes. All right. So I laser cut some toys to help us play with this puzzle.

So we're supposed to prove that a $2^n \times 2^n$ board can always be tiled with L trominoes. Here, let me put something on there for you to enjoy. Well, let's phrase it like we were going to prove it with induction because we are.

So let's say that $P(n)$ is going to be a $2^n \times 2^n$ board without one cell-- without any one cell. Doesn't matter which one you pick-- can be tiled with L's. I'm going to call them L's, not L trominoes.

And notice I kind of just wrote the same thing twice. That often happens. $P(n)$ is not our theorem. It's the predicate, which is inside of the for all that is part of the theorem. So this is really every $2^n \times 2^n$ board. So if $P(n)$ is this, then our theorem is that for all n greater or equal 1, $P(n)$ is true. All right? And now we need to prove this.

OK. Well, let's start with n equals 1. You can think about what n equals 0 means. It turns out it's true. But let's start with n equals 1.

So a $2^1 \times 2^1$ board-- that's a 2×2 board. And can you see the little crosshairs in there. It's a 2×2 grid. And we're supposed to say no matter which single piece-- no matter which single cell we remove, we can always tile what remains with these L shapes. Anyone see how to do it? Yeah. All right. I'm going to put that there. Solved.

What if the missing cell was in a different place, like there? Yep. OK. Can still solve that. We're allowed to rotate. So the 2×2 case isn't very hard. So let's keep track of our progress-- $P(1)$. Actually, I'm going to write this as the 2×2 case. Not very hard. Let's step it up a little.

What about the 4 by 4 case? And say we're trying to avoid that one cell. Well, we also need to solve it for avoiding that one cell and that one cell. We need to solve all 16 cases. Because the theorem says it doesn't matter which one you exclude. You can always tile the rest.

Well, you can play around with this. I have lots and lots of pieces to play with. You can tell I had way too much fun designing this thing. And it turns out, you can make it work.

Like this one-- this is one of the 16 examples. It turns out, if I put that one there, and then that one there and there and there and there, then I've solved it. I've solved this one out of 16 cases.

But also, it turns out it's not too difficult to generalize from there. Like, if the missing cell in one of these four positions, then I know how to do it. If it's in one of these four positions, then I also know how to do it. If it's there, I can just put that there. If it's there, I can put that there. If it's in one of these four spots, can still do it. If it's one of these four spots, can still do it.

So by looking at all of those 16 cases combined, we now have a way that can put that yellow cell anywhere we want and finish tiling the rest. Not a very nice or clean argument, but at least we did all 16 cases. So that's good enough for now. Let's call the 4 by 4 case solved for now and move on.

And I want to emphasize this is actually a key moment in this proof. We just solved the 4 by 4 case by looking at all 16 cases and showing that they're all possible. I want you to hold that deep in your heart that we have solved 4 by 4. Any 4 by 4 you ever see with a single missing cell can be tiled. Yeah?

I want you to hold that deeply in your heart and then not worry too much about how we did all the 16 cases. Just know that it's true. Are we comfortable with that for now? 4 by 4 case is fully solved. Let's move on to 8 by 8.

We have an 8 by 8. There's a single cell we're trying to avoid. Maybe it's there. Maybe it's over here. Let's leave it there. I don't care. What we're really looking for here is a generalizable method. I don't want to do 64 cases.

And in the 16 by 16 case, I don't want to do-- what is that-- 256 cases? Please don't make me. We want a generalizable process that we can use to solve every board size. And here is a cool idea that works nicely.

So first thing we're going to do. We're going to shift our perspective. This isn't an 8 by 8 board. This is for 4 by 4 boards. All right? And now let me just make this one purple for now to remember that that's the special one. We cannot move that one.

However, remember we have solved the 4 by 4 case. So if we're missing a cell somewhere in this quadrant, and missing a cell somewhere in this quadrant, and missing a cell somewhere in this quadrant, then now we just have four instances of the 4 by 4 problem. We have four 4 by 4s, each missing a single cell. And we know how to solve all of those.

So we know how to solve all of the red that remains. Unfortunately, we're not actually allowed to use these individual cells. We're supposed to cover everything with our L-shaped tiles.

But remember, doesn't matter where we put these yellow ones. Anywhere we put this in here, we know how to solve that quadrant. Anywhere we put this in here, we know how to solve that quadrant. The purple one was the one that's given to us. We're not allowed to move that.

But the next clever idea is, what if we put the yellow ones right there, and then think of them as not three individual cells, but as an L in their own right? No. There we go.

So now this is going to be our key idea. If we're trying to tile everything except the single special cell, step one-- think about it as four separate quadrants. Step two-- put an L right in the middle, directly in the middle, to make sure that each of the four quadrants is now missing one cell. Yeah?

If this initial cell that we're supposed to avoid is somewhere else, well, then we just put the L to touch the other three quadrants. We just rotate it this way instead. And once again, every single quadrant, all four quadrants, are missing a single cell.

To really drive that home, let's see. We have this shape here. That's a 4 by 4 missing a corner. This shape here is a 4 by 4 missing a corner. Here's a 4 by 4 missing a corner. Here's a 4 by 4 missing that initial cell. So we have four separate problems where if we can solve the four orange cases, we'll be able to piece them back together to form a solution.

And this is why we stored that deep in our heart. We know we can solve 4 by 4s with a cell missing. So we know all four of these orange quadrants can be filled. So we already know that this bigger shape can also be filled. So we're done.

No matter where this one started, we can do this thing. We put down the cross. We put the L in the middle. Now we have four quadrants, each missing a cell. We invoke the 4 by 4 case. And now we've solved any 8 by 8 case that's given to us. Does that make sense?

So if we believe that we can solve every 4 by 4, then we also believe that we can solve every 8 by 8. Yeah? And this keeps going. If you instead had an 8 by 8-- sorry. If you instead had a 16 by 16, you can do the same thing. Draw the crosshairs, put the L in the middle. Now, all four 8 by 8 quadrants are missing a single cell. Invoke the 8 by 8 case.

We can even see that here. Go away, webcam. So here is a 16 by 16 with one cell we're supposed to avoid. So let's cut it into four quadrants and put an L right there in the middle. Now we should think of this as four separate 8 by 8s, each missing a single cell. And we trust in our heart that we have solved the 8 by 8 case.

Don't worry too much about how we solved it. Just be confident that we solved the 8 by 8 case, that there is a way to tile each of the 8 by 8 corners. So now we can piece them together to form our 16 by 16 solution. And now it's all L tiles, and we win.

So this is the idea behind this proof that we're able to show that P_n always implies P of n plus 1. If you can tile a 2 to the n by 2 to the n missing one cell, then you can tile 2 to the n plus 1 by 2 to the n plus 1 missing a tile using our strategy of taking the 2 to the n plus 1 by 2 to the n plus 1 and cutting it into four quadrants of 2 to the n by 2 to the n . Put an L in the middle. Make sure each quadrant is missing a cell. Invoke the inductive hypothesis.

That was a little sketchy. I didn't write this down carefully. You can check the lecture notes later if you want to see this written down a little more carefully. But that is the core idea and the problem I want you to see.

There's one more thing I want to mention before we go. So I kept saying, store it deep in your heart that we've solved the 4 by 4. And then use that to believe that we've solved 8 by 8. Believe that we solved 8 by 8 in order to believe that we've solved 16 by 16.

But sometimes belief isn't quite enough. And you really want to actually run this algorithm. Well, for then, all you have to do is when we have the 16 by 16, and you break it down into 8 by 8s, well, let's actually remember how we solved the 8 by 8s. We did that by splitting them up into 4 by 4s.

And then turns out you can do the same thing and split the 4 by 4s into 2 by 2s. So now each 2 by 2 is missing a single cell. But a 2 by 2 missing a single cell is just an L. So it's all L's. And then we just recombine, reverse what we did before, and now we've tiled everything, and we're done. And that's the algorithm. Thank you so much.

[APPLAUSE]