PROFESSOR: So we come now to the topic of induction, which is a standard part of a high school curriculum and you've probably seen before, but is nevertheless worth looking at at the level that we look at things in this class. So the idea of induction can be-- one way to explain it is this way. Suppose that I plan to be assigning colors to the non-negative integers like, say, in this example, I color zero blue and one red and two blue and three red and four and five green, and it goes on somehow.

OK. Now I'm going to describe to you a coloring that I have in mind, and we'll see whether you can figure out what it is. Here are the properties that my coloring has. First of all, I've colored zero red, and I've also continued the coloring satisfying the following rule. If I have an integer that's next to a red integer, then it's red also. Any integer next to a red integer is red also. So what's my coloring? Well, you obviously realize that they're all red. They have to be. And there they are.

OK. This is actually a statement. It can be read as a statement of the rule of induction. It's kind of a self-evident axiom about numbers, but let's state it abstractly. So first of all, what induction is assuming is that you have some property of numbers. Call it red, R. And R of zero you're told, and you're also told that R of zero implies R of 1, and that R of 1 implies R of 2, and R of 2 implies R of 3, and in general R of n implies R of n plus 1, and so on.

So we've written it out this way as an infinite set of implications to emphasize that that's what the rule that I stated, that if an integer is next to a red integer then it's red, is shorthand for. It's really a shorthand for this infinite number of different implications, each of which has to hold in order for you to be able to apply the rule of induction.

Well, what can you conclude if all of these things hold? Well, then you can conclude that zero is red and one is red and two is red and n is red, and so on. OK. Now of course, there's a much more concise way to express both these antecedents above the line and the conclusion below the line using quantifiers, namely the antecedents could simply be said by two predicate formulas, R of 0, comma, for all n, R of n implies R of n plus 1. That's really a summary of what we said on the first slide, that if an integer is less than a red integer, then it's red. That is, n plus 1 is next to n. If n is red, then n plus 1 is red.

Similarly, the stuff below the line, that R of zero, R of one and so on hold is simply expressed

as for all m, R of m holds. And this is the form of the induction rule. It's read that if you've proved R of 0 and you've proved that, for every n, R of n implies R of n plus 1, then you can conclude that for every m, R of m holds, where the variables are all ranging over the non-negative integers.

By the way, notice that I used n for the variable name above the line in the antecedent, and m for the variable line below in the consequent. I can use any names that I like for bound variables, just as in when you define a procedure you can name the parameters of the procedure anything you like, because they're local variables. And I've used an m in the bottom and an n in the top just to emphasize that those variables have nothing to do with each other, which is a point that sometimes confuses students.

OK. Sometimes the rule of induction is explained in terms of dominoes. You have all these dominoes lined up next to each other. You knock one over, it knocks over the next one, and so on. If that helps you think about and remember dominoes, that's fine.

OK. Let's apply induction-- maybe one of the most basic and standard applications would be to prove a numerical identity. So let's prove one that we've actually seen before. This is the formula that we've previously proved using the well-ordering principle for the sum of a geometric-- for a geometric sum. The sum of R to the 0, R to the 1, up to R to the n. And the claim is that that's equal to R to the n plus 1 minus 1 divided by R minus 1. So this sum of n plus 1 terms actually can be expressed concisely with a fairly single-- a single, simple term. Of course, this only works if R is not 1, because I can't have the denominator be zero.

All right. How do we prove it? Well, I'm going to do the proof. And at the same time that I do the proof, I'm going to show you kind of a standard template that you can pull out and use for induction proofs. So the template, it's just an organizational method to do the proof. I'm doing the template in magenta. So that's the part that really is form, not substance. There's no math in it, it's just the structure that we're going to organize the proofs in, at least in the beginning.

So here we go. The first thing you do is tell your reader that you're going to be using proof by induction. That helps them understand what's coming. So you begin with the line proof by induction on n. Now n is not in magenta because sometimes you use different variables, and sometimes there'll be many variables in the assertion. So you need to tell the reader just which one is the one that you're going to be applying induction to.

All right. That said, the most important part of the proof, the part where there's usually a

mistake-- if there's a mistake anywhere, it's usually in the absence of the statement of an induction hypothesis, or a misguided induction hypothesis. So the next part of the template says the induction hypothesis P of n is-- and in this case, our induction hypothesis is that this equality holds. That's what we're trying to prove.

So the induction hypothesis is P of n. The objective, then, implicitly, when we're doing induction with this induction hypothesis, is to prove that for all n, P of n calls. This identity works for all non-negative integers n.

OK. Having stated the induction hypothesis, the first thing we have to do is work on the base case. That is, prove it for n equals 0. Now we're telling the reader that it's n equals 0 because sometimes it's convenient to start at n equals 1 or n equals 2, and then you're just concluding that the property holds for all n greater than or equal to 1, or however-- all n greater than or equal to wherever you started.

So we're going to start at 0, which is the standard place. And what do we have to check? We have to check that the sum on the left, when n goes to-- when n is 0, is equal to the sum-- to the formula on the right when n is 0. Well the sum on the left, when n is 0, it's really just 1, because it's going from R to the 0 to R to the 0. The R and the R squared, they're a little misleading because they're not really there when n is 0.

So the left-hand side is one, and the right-hand side is R minus 1 over R minus 1, which is 1 since R is not 1. So sure enough, it checks, and we're OK. The case n equals 0 has now been proved.

So the next thing we have to do in the template is to go to the inductive step. And that's where we assume that P of n holds. And we're allowed to use the P of n assumption in order to prove that P of n plus 1 holds, where the only thing that we know about n besides that P of n holds is that n is greater than or equal to 0. And our proof has to work for all possible n that are greater than or equal to 0.

All right. Well, now we can start doing the non-template method that has to do with the content of what we're actually trying to prove. This is what I want to prove. This is P of n plus 1. It's gotten by replacing the n's in the previous equation by n plus 1's. I'd like to have that be.

OK. How do I get to that? Well, I can assume P of n, which kind of looks like it already. It's a good head start to getting to P of n plus 1. So I'm allowed to assume that this equality holds for

n. I don't know what n is except that it's a non-negative integer, but this equality holds for n. And I have to prove that it holds for n plus 1.

Well, if you look at this, what I'm trying to prove is something about the sum that goes up to R to the n plus 1. So given this equation, I can turn the left-hand side into the sum that I'm interested in. That is, the sum of powers of R up to the n plus 1st power of R simply by adding R to the n plus 1 to both sides, an obvious strategic move, or tactical move.

OK. So doing that, I get this equality, which I've now proved from the induction hypothesis. Namely, the sum up to R to the n plus 1, which is what I'm interested in, is equal to this algebraic expression on the right-hand side. And if I'm lucky, and of course I will be, the right-hand side is going to simplify to be the target expression, with n replaced by n plus 1. So what happens is-- let's put R to the n plus 1 over this common denominator, R minus 1. And I get the second term, and then you can do a little bit of algebraic simplification, trivial, and you'll realize that, sure enough, it simplifies to R to the n plus 1 plus 1 minus 1 over R minus 1, which was exactly the equality that I was hoping to prove. So in fact, at this point we can say that we've proved P of n plus 1, and we've completed the induction proof. We're done.

OK. That is the first basic example of an induction proof. And the whole template is now visible, except maybe there should have been a QED or a Done statement.

All right. By the way, as an aside, and we already saw a little problem with this, the three dots that appeared in the sum are called an ellipsis. Plural is ellipses. And they're used where the writer is trying to tell the reader that there's an obvious pattern that the reader is expected to see, which I think is fairly clear in this case. You go-- you know, it's R to the 0, R to the 1, R to the 2, R to the 3, up to R to the n.

The difficulty is that sometimes the ellipsis can cause some confusion. For example, we had to figure out that when n is 0, the left-hand side actually just meant 1. It was just R to the 0, so the R and the R squared weren't really there. One way to really avoid that kind of fence post problem where you've shown-- in order to make clear what the pattern is, you've shown more than it may-- more of a pattern than may always be there-- is to use a precise mathematical notation where I actually tell you the pattern of the i-th term, and tell you that you should sum from i equals 0 to n. So the sigma notation is shorthand for sum, and I'm telling you that the i-th term in the sum is R to the i, and it's going to run from i equals 0 to n. So this is a sort of mathematical notation for a for loop or a do loop-- do from i equals 0 to n, add R to the i plus R

to the i to the accumulator. And the sum notation is certainly more precise. But sometimes, it's actually harder to read than simply showing you the pattern, because the pattern often is visible visually.

OK, now let me tell you a little story. And it's a made-up story, but it's kind of fun to tell. This is the familiar building, the Stata Center. And this is actually a design mock-up that the architects produced for the MIT team that was overseeing the construction and design of the building to show what the student lobby would look like, the student street. Now the story goes that part of the plan for the student street was to have a plaza that was going to be built out of unit size squares, but an uncertain number of them. There was going to be a parameter that determined the size of the square.

And the size of the square was actually going to be a power of two by a power of two made out of unit size tiles. So there would be 2 to the n times 2 to the n unit size tiles filling up this square plaza. And the plaza was to be tiled with these unit tiles, but one tile space was to be left blank so that the statue of a-- what was then the potential donor, Bill, could be placed in the middle as an incentive for him to donate funds for the completion of the building, which indeed he did.

So the puzzle, then, was put forward by the architect Frank Gehry, who many regard, after Frank Lloyd Wright, as the greatest architect of the 20th century. Gehry specified for aesthetic reasons that he wanted to the square to be tiled with L-shaped tiles that were made out of three unit squares. He thought that that would give a pretty design, and it actually does.

So here's an example of tiling the n equals 3 case, [? 2 ?] [? to the-- ?] [? 2 ?] [? cubed ?] equals 8x8 plaza with Bill in the middle. There is the 8x8 plaza tiled with these L-shaped tiles, each consisting of three unit tiles. So the question was that the exact size of the square was to be determined by other architectural considerations. So it was parametrized by n, which is going to be 2 to the n by 2 to the n. The question was, can you always find such a tiling no matter how big the square is and leave Bill in the middle.

Well, let's try to prove it by induction. The induction hypothesis-- we're trying to prove a theorem that, for any 2 to the n by 2 to the n plaza, we can make Bill and Frank happy. That is, Bill's happy when he's in the middle, and Frank is happy when the rest of the square is covered with L-shaped tiles. By the way, middle is a little bit ambiguous because there are really four middle squares. But of course, it doesn't matter which one you fill, because if you

wanted a different one you could just rotate the whole square and get any one of the four middle squares empty for the Bill statue.

So an induction proof would proceed by induction on something or other. And the obvious thing is the n that's in the statement of the theorem. And the induction hypothesis would straightforwardly be that we can tile to 2 to the n by 2 to the n plaza with Bill in the middle.

OK. The base case is n equals 0. That's a 2 to the 1 by 2 to the 2. It's a 1x1 square. OK, well, not a problem. You just put Bill in the one square, and you tile the rest with no L-shaped tiles. That fits the rules. The base case n equals 1-- n equals 0 is covered. All right.

So now we come to the double size square, the square that's of size 2 to the n plus 1 by 2 to the n plus 1. I have to tile that with Bill in the middle, but I have a fairly powerful induction hypothesis that I'm allowed to assume, namely that I can tile the half size square, the 2 to the n by 2 to the n square, and get Bill in the middle. So obviously, the double size square is made out of four half-sized squares. And so I can try to fill up the whole square that's 2 to the-- the whole full-sized square, or double-sized square, 2 to the n plus 1 by 2 to the n plus 1, by working with my ability to tile them with L-shaped tiles, leaving Bill in the middle for each of those four subsquares.

So I can assume that, and now I'm stuck really. What do I do? How do I use this ability to put Bill in the middle of each of those four quadrants in order to color-- to fill in the whole thing with N-shaped tiles? I'm stuck. And the point of this example is to show you the way to get unstuck, which is kind of unexpected.

I'm actually going to get unstuck by proving something stronger. I'm actually going to prove that we can find a tiling using L-shaped squares with Bill placed in any specified square that you like. Wherever you want to put him, I can tile the rest with L-shaped tiles and leave the specified single square blank for Bill to be inserted, for a statue of Bill to be put there.

So what's unintuitive about this is that I'm proving something stronger. It ought to be harder to prove, right? But because I'm trying to prove a conclusion that's stronger, I also have a stronger induction hypothesis to work with in conducting the proof. And the net proof actually, as you'll see, is going to be easier.

So let's do it with the stronger induction hypothesis. The theorem is, again, for any 2 to the n by 2 to the n plaza, so we can make Bill and Frank happy. Prove by induction on n. But with a

revised induction hypothesis-- I'm calling it P of n again-- which is I can tile the square with Bill anywhere. So the base case, n equals 0, is the same as before. It's just 1x1. So I put Bill in the only tile that there is, which is both the middle and the corner and everything else. And the base case doesn't change.

For the inductive step, now I have a more powerful thing that I can assume is the induction hypothesis. I can assume that, in any given square of a 2 to the n by 2 to the n-- any given tile location, unit square, of a 2 to the n by 2 to the n plaza, I can tile the rest with L-shaped squares and get Bill where I wanted him to be. And I have to use that hypothesis to show that I can get Bill anywhere that's required in a 2 to the n plus 1 by 2 to the n plus 1 square.

So suppose that we want to tile Bill in that designated arbitrary square of the 2 to the n plus 1 by 2 to the n plus 1 plaza where we happen to choose a location where Bill is in the upper right quadrant, in the upper right half size square. All right. So my hypothesis, I can fill in the purple square, that quadrant, with L-shaped tiles, leaving Bill in the place that he's supposed to be.

Well, here's the trick. With the other three, since I can tile them with Bill anywhere, I'm going to tile them with Bill in the respective corners of those three other subsquares, which meet in the center of the full-size plaza, as shown here. And having done that, now it's obvious how to fill up the whole 2 to the n plus 1 by 2 to the n plus 1 plaza, because I pull those four separate pieces together to form the full-size 2 to the n plus 1 by 2 to the n plus 1 plaza. And look, I just put an L-shaped tile in the middle to fill up those three corner Bills, and I'm done. And the proof is complete. We have just proved by induction that indeed you can tile any power of 2 by power of 2 square putting-- leaving Bill wherever you want him, and the rest filled with L-shaped tiles.

Now notice that a part of this process actually is implicitly defining a recursive procedure to actually do the tiling. If you watch the way the proof went, if I was going to write a recursive procedure to do the tiling, what I would do is say OK, you give me input n plus 1, which are the dimensions-- the specification of the dimensions of the plaza, input n plus 1-- or input n means it's 2 the n by 2 to the n. How do I do that? Well, you tell me where you want Bill to be as another parameter, and then I will call myself recursively on four half size squares. So that is, call myself to do squares with dimension parameter n minus 1 four times for each quadrant, each time specifying in those quarters where I want Bill to be. The recursive procedure will return an L-shaped tiling of those four pieces, and then I take those, fit them together, tile that middle, and I wind up with a tiling of the whole region.

So what I've just talked through is the description of a very easily written recursive procedure that would print out a picture of an L-shaped tiling given, as input, any number n. And that's, in fact, that how we got the 8x8 tiling, although we did it by hand rather than writing a program. And that's enough of two examples of basic mathematical induction.