PROFESSOR: So it's time to examine uncountable sets. And that's what we're going to do in this segment. So Cantor's question was, are all sets the same size? And he gives a definitive answer of no. Cantor's theorem, which we're about the present, will show that, in fact, there isn't any biggest infinity. For any given infinity, you can find a bigger one in a very simple way.

But let's begin by coming up with the simplest form of Cantor's diagonal argument, of how do you prove that a set is not countable? Remember, a set is countable if you can list it, possibly with repeats. So A is countable if there is a sequence, a0, a1, a2, such that every element in the set A shows up at some time or other in the list, possibly more than once. And the only things in the list are elements of A.

And we saw as an example that the finite bit strings, the finite strings of zeroes and ones, with the finite binary words, are an example of a countable set. And we claimed last time, and now we're about to prove that the difference is that if you look at the infinite bit strings-- One-way infinite. They have a beginning, and they go on infinitely right. The notation being 0, 1 to the omega, where omega is an indication of one of the symbols for a kind of infinity. And this is going to be an example of an uncountable set.

How are we going to prove that? Well, the setup for using a diagonal argument is to think about drawing a matrix. Suppose that I have some way of enumerating the infinite binary sequences, the 0, 1 to the omega. So there's a sequence s0, there's a sequence s1, there's a sequence s2. Let's lay them out as though they were the rows of a matrix.

So s0 is this infinite binary sequence, 0, 0; 1, 0; and so on. And the column labels are simply the coordinate labels for s0. So this s0, 0; s0, 1; and so on. s1 is the next infinite binary sequence in this hypothetical list. And it starts 0, 1, 1, 0, and goes on, and so on, down the line.

So the row labels are this numeration of binary sequences. And the column labels are coordinate labels. And this is a matrix that's infinite to the right an infinite down. But it definitely has an upper left corner.

So the trick is to try to find an infinite binary sequence that is not in this list, that it differs from every row. If I can do that, then I've shown that any attempt to enumerate all of the binary sequences in 0, 1 to the omega-- any sequence like s0, s1, s2 of binary sequences is missing

something, which means you can't really list it.

Well, how do you find a sequence that's not here? Well, pretty easy. You look at the first digit. And that was a 0. So you choose the first digit of the new sequence to be 1, the opposite of 0. You choose the second digit to be the opposite of the coordinate of a digit 1 of s1, and complement that.

Now here, the digit 2 of S2 is a 0. So let's make that a 1. And the next one of the diagonal is a 0, and so on. We're going to complement all of the bits along this diagonal. So I get a diagonal sequence. That's why this argument is called the diagonal argument.

Well, let's think about this diagonal sequence. It just goes on, right down the diagonal of this two dimensional infinite matrix. What we can say about it is that it differs from every row. Why is that? Well, it differs from the 15th row at the 15th position, or coordinate 15. It differs from the 99th row at coordinate 99. It's not in the matrix. It's not a row of any matrix.

And that immediately tells you it's over. Any attempt to list all of the elements in 0, 1 to the omega is going to omit a diagonal element. It's not possible to list all of them. In other words, there isn't any surjection from the non-negative integers to 0, 1 to the omega, because I've just shown you how, if you give me a surjection onto binary sequences, in effect, I'm giving you with N a 0 sequence, a first sequence, a second sequence, and so on. Then I know exactly how to find something that's not there. There can't be a surjection from N to 0, 1 to the omega. It's just not true.

And that's why we can say that 0, 1 to the omega is uncountable. Definition of countable-- or an equivalent formulation of countable, remember, is that there is a surjection from the non-negative integers to the set. Well, there isn't any. We just proved it. So N is not surj 0, 1.

By the way, it's also quite easy to say that there is a surjection from the infinite binary sequence to the non-negative integers. You could map a binary sequence to the coordinate of the first one in it. And that maps every infinite binary sequence to a non-negative integer. It hits every non-negative integer lots of times. I hit five with a sequence that starts with four zeros and a one. The only sequence that doesn't go anywhere is the all-zero sequence.

But by the way, if you check the definition of surj, it doesn't require that the function be total. Surj means that there is a function that is a surjection to N. But there doesn't have to be greater or equal one arrow out. There just has to be less than or equal to one arrow out, so it's

a function.

But of course, it's easy enough to make it total, the all-zero sequence. Just map it to 0. So now you have both the 0 and the sequence that starts-- Any sequence that starts with 1 will all map to 0.

OK. So if we remember our intuition, surj is read as greater than or equal to. So this tells us that the infinite binary sequences are a larger set, at least as large a set as the non-negative integers. And the converse is true. The non-negative integers are not at least as large as the infinite binary sequences. So we can really say that non-negative integers are strictly smaller than the set of infinite binary sequences.

Now, strictly smaller is in quotes because, again, we don't know exactly what the size of infinite sets is. All we're doing we, really, is talking about properties objections, bijections, surjections, injections.

OK. So let's make an explicit definition. I'm going to say that A strict B means that there is no surjection from A to B. So if we read A surjection B intuitively, as A greater than or equal to B, this is saying it's not true that A is greater than or equal to B. Or in ordinary language and thinking about sets, if it's not true that you're greater than or equal to B, you must be strictly less than B. So that's the motivation for the word strict.

But remember, we're talking about infinite sets. And we can't go around assuming too many properties of strict until we've proved them.

One non-trivial property, by the way, is I've defined strict that it's not true that there is a surjection from A to B, but I'm not insisting that there must be a surjection from B to A, which would be the second companion part. That is, A is not greater than or equal to B, and B is greater than or equal to A. Turns out, technically, you can prove that. If there isn't any surjection from A to B, there will be a surjection from B to A. But that's using a set theoretic argument that's not so obvious. And we don't need it.

So this is the definition of strict. A straight B means you cannot get a surjection from A to B. And we're intuitively reading it as A is strictly smaller than B. And what we've just shown then is that the non-negative integers strict 0, 1 to the omega, the infinite binary sequences.

OK, now Cantor's theorem is a wonderful generalization of this. It's a powerful generalization, but the proof is pretty much the same, although it sometimes looks a little different as it's

written up. And what Cantor's theorem says, it's just beautifully elegant and simple. It says simply that the power set is strictly better than the set. A strict power set of A, for every set A, even if A is finite.

Because remember, if A is finite, say A has n elements, then the power set of A has 2 the n elements. And you could check that even for n equals 0, n is less than or equal to the-- n is less than 2 to the n. 0 is less than 2 to the 0, which is 1. 2 is less than 2 squared, which is 4, and so on. So even for finite sets, we have A strict power set of A. But the cool thing is that it works even for infinite sets.

Let's take a look. It's a diagonal argument again. But now, I mustn't assume that A is countable. I'm not going to assume that I can really list the elements of A. But we'll think about it as though we could. Let's think about this matrix again. So suppose A is this set of elements, a, b, s, t, d, e. I'm scrambling up the alphabet on purpose, because I don't want you to get the idea that we're assuming that A is countable, that you can list all the elements of A. I'm not assuming that. But I'm just writing out a sample of elements of A.

And let's suppose that I was trying to get a surjection from A to the power set of A. So suppose I have a function f that maps each of the successive elements of A to some subset of A. So f of a is part of the power set. It's a subset of capital A. f of b is a subset of capital A, and so on.

And suppose I had a setup like this. I'm going to draw a matrix that looks like the diagonal matrix. And we're going to extract a diagonal set and discover that that diagonal set is not one of the f's. It's not f of anything, which means that f is not going to be a surjection. So let's look at it again.

So here's this matrix, where I'm labeling the columns of the matrix by the elements of A. No particular order here, but in order to draw a matrix, I have to write them down in some order. And likewise, the first row is going to be f of this element a. Well, what is in f of A? f of A is going to be a set of elements. So let's just write the elements in f of A down, under the corresponding column label.

So here's an example, where f of A has an a and no b, and has it has an s in it, no c, no d. It has an e. Likewise, f of b has an a and b, and no s or t, but it's got a c, and so on.

So I filled in this matrix by taking f of an element in A, which is supposed to be a subset of A, and writing out all of the elements in that subset, under the corresponding element of the

subset. So a b goes under a b if it's in f of c, and s goes under an s if it's in f of c. Nothing goes under a t if t is not in f of C. And that's what we're seeing here.

So I'm laying out as though I was using 0's and 1's for an infinite binary sequence. I'm laying out each of the sets that are in the range of f along this row. And now with this setup, and I can define a new set, which is not going to be an f.

How do I get that? Well, what I'm going to do is in my new set I'm not going to have any of the elements that appear on the diagonal. So if a is a member of f of A, that means that a appears in this coordinate, it's not going to be in my set. If b is in f of b, meaning that b appears in the f of b row under the column b, it's not going to be in my set.

On the other hand, s is not in f of s, because there's no s there. So I'm going to put an s there in magenta. And likewise, I'm going to stick elements in or out the opposite of whether they appear on the diagonal. And this is going to give me a set D, which is going to be my diagonal set.

So if we write this out, what we're saying is suppose that I have a function f from A to the power set of A. Then what I'm going to do is define a subset of A that's not in the range of f, namely set D, which is the set of those elements in A, such that little a is not an f of a. Namely, if an element appeared on the diagonal because an element with column label A was in the row f of a, then I left that out of my set. And if it was not in that location in the matrix, I put it in my set.

So I'm keeping all the elements that aren't on the diagonal, that's my diagonal set D. And what I know about it is that D is not in the range of f, because it differs from every possible row of the matrix. If the row is labeled with f of a as f of a, then it differs in the column a, f of a from that row. And therefore, my set D is not a row of this matrix. And that means that it's not equal to f of anything.

So I've just found that there is no f arrow into D. D's not in the range of f. That means that if I had such an f from A to the power set of A, it's not a surjection, because D is always left out. So f is not a surjection. And since, you know, f is any function from A to the power set of A, none of them are surjections, and that means there is no surjection from A to the power set of A. In other words, A strict power set of A. There is no surjection.

Now, a special case of this, of course, is that the non-negative integers are strictly smaller than

the power set of N. Of course, that's just an instance of Cantor's theorem. We're applying A being the set of non-negative integers. So there is no surjection from the non-negative integers to the subsets of non-negative integers.

Again, that means that the power set of N is an example of an uncountable set, because the definition of countable is that there would be a surjection from N to power set of N. We're saying there isn't any, so it's not countable. Not countable is usually phrased as uncountable. So the power set of N is maybe our second example of an uncountable set, the first one being the infinite sequences of binary numbers.

Now as a matter of fact, just as we had a general way to prove countability-- you can show that a set is countable if there is a surjection from a set you know is countable onto the target, then the target's countable. Take the contrapositive of that lemma, and you can say that if a set A is uncountable and there is a surjection from C to A, then C has to be uncountable. That's just the contrapositive of the previous one. If C was countable, then A would be countable. So if A is uncountable, C must be uncountable.

So this gives us, again, a nice general way to prove uncountability of sets, once I have a couple in my repertoire. Well, it means that we could have deduced that 0, 1 to the omega, that the infinite binary sequences were uncountable, because we know that there's is a bijection between the infinite binary sequences and the power set of N.

We described that bijection without knowing anything about any other properties of the infinite binary sequences in the power set of N, whether they were countable or not. But now that Cantor's theorem tells us that the power set of N is uncountable and there's a bijection, the previous lemma says, in particular, there's a surjection from 0, 1 to omega to the power set of N, which means 0, 1 to the omega is uncountable.

So what I'm illustrating then is that the proof that we used directly by a diagonal argument to figure out that 0, 1 to the omega was uncountable, it's really a special case of the more general diagonal argument that we used to prove Cantor's theorem. And we get that 0, 1 to the omega is uncountable as a consequence of Cantor's theorem about the power set of N. And so we've got two different ways then to prove that the infinite binary sequences are uncountable.

Another example of an uncountable set, it's the real numbers. And they're a cute example. Remember, we saw that the rational numbers were countable. The real numbers are

uncountable. Well, how do I prove that? I'm just going to show you a surjection from the real numbers onto the infinite binary sequences.

How am I going to do that? Well, it's a kind of stupid trick, but it works. I'm using both positive and negative reals. So let's look at some real number, and look at its binary representation, so for the moment that it's positive.

So let's look at, say, the binary representation of some number, like 3 and 1/3. So that means that if we're thinking of these as binary places, this is the 0's place, the 2's place, the 4's place. This is half's place, the quarter's place, the eighth's place. Then the binary representation of 3 and 1/3 would be 2. And then this infinite repeating-- not decimal, but [? bicipal-- ?] binary expansion, 010101. And we will examine how I know that that's a third. But take it for granted that that's what you get as the repeated fraction. You could figure that out by just doing division of 1 by 3 in binary.

Anyway, just as there's a decimal expansion of every real number, there's a binary expansion just using base 2. So here's the binary expansion of 3 and 1/3. So what I'm going to do is I'm going to map 3 and 1/3 to this binary sequence. I'm going to ignore the decimal place. Binary is not decimal place. It's a [? becimal ?] place, or binary position. And I'm just going to take this to mapping the sequence, 11010101.

And I claim that this is a surjection because you're going to hit every possible binary sequence in this way. Well, almost. Let's take a closer look.

There's a problem with mapping to things that start with 0, because let's examine that a half is 0.10000000. So I would map it to that. And it will end. But there's an ambiguity, because a half is also equal to 0.011111, just as 0.999999 is equal to 1.000000 in decimal, you get the same infinite carry issue here in binary.

So numbers that end in all ones have another way to represent the very same number by a sequence that ends in all zeroes. So how am I going to hit-- if I'm using up a half to hit this one, what's left to hit that one? Oh, how about using minus 1/2? It's there, and that's part of R, so I'm just going to map the negative numbers to the version of the expansion that starts with 0 and has an infinite number of ones. And the positive one that ends with an infinite number of zeroes. And otherwise, I'm going to map plus and minus numbers to the same place.

So this is going to give me the needed surjection from R to the infinite binary sequences. And

by our previous lemma, that implies, sure enough, that the real numbers are uncountable.