PROFESSOR: Hi. Today we begin our study of integral calculus, which had its roots back in ancient Greece, roughly 600 BC, and began with an investigation of the study of area. In a certain manner of speaking, today's lecture could be called 'Calculus Revisited Revisited', in the sense that integral calculus can be studied quite apart from the fact that differential calculus was ever invented. We'll see later in this block of material that there is a wondrous relationship between integral and differential calculus, but perhaps the best proof that integral calculus can be studied independently of differential calculus lies in the fact that the ancient Greek was doing integral calculus in 600 BC, whereas differential calculus did not begin until about 1680 AD with Sir Isaac Newton.

Well, at any rate then, what we'll call today's lecture is simply 'Two-dimensional Area'. In other words, this is how the subject began, with studying the amount of space in plane regions, two-dimensional area, as opposed to, say, as we'll talk about later, three-dimensional area, which is really a fancy word for volume, et cetera.

But enough about that for the time being. Two-dimensional area, or Calculus Revisited Revisited. And to see what's happening over here, the study of area was the forerunner of integral calculus as we now know it. It began in ancient Greece, roughly 600 BC. The branch of calculus that we've been studying up until now in our course, differential calculus, did not begin until 1680 AD. Notice, then, the interesting juxtaposition in time. In other words, pedagogically, one seems to study differential calculus before integral calculus. Chronologically, integral calculus preceded differential calculus by more than 2,000 years.

Now this is a rather beautiful study. It's one of the most aesthetic parts of elementary mathematics, namely the simplicity with which the Greek was able to tackle the sophisticated problem of finding areas in general. He started with three basically simple properties of area, so simple that most of us are willing to accept them almost intuitively. We'll call these the 'Axioms for Area'.

And simply stated, they were the following. One, that the area of a rectangle is the base times

the height. Two, if one region is contained within another, the area of the contained region is less than or equal to-- in other words, can be no greater than-- that of the containing region. Roughly speaking, the smaller the region, the smaller its area.

And finally, written slightly more formally here, the old high-school axiom for plane geometry, that the area of the whole equals the sum of the areas of the parts, that if a region 'R' is subdivided into a union of mutually exclusive pieces, then the area of the entire region is equal to the sum of the areas of the constituent parts.

And with just these three axioms, the ancient Greek invented a technique for finding areas that today is still known by its original name, the 'Method of Exhaustion'. And here, exhaustion does not refer to the fact that we get tired using the system, even though, as we shall soon see, it's quite intricate. It refers to the fact that we take the region whose area we want to find and exhaust the space by squeezing it between regions which are made up of rectangles.

Stated more precisely, given a region 'R'-- and I'll define this more precisely later-- we squeeze it between two networks of rectangles, the idea being what? That we know what the area of a rectangle is, and consequently, by the fact that the area of the whole equals the sum of the areas of the parts, if we have a rectangular network, knowing how to find the area of each rectangle, we know how to find the area of the rectangular network.

Now, rather than to wax on like this philosophically, let's tackle a specific problem. In fact, give or take a little bit, this is specifically the problem that Archimedes dealt with in finding the area of parabolic segments way back, as I say, in the 300 to 600 BC era.

As an example, consider the following. We wish to determine the area of the region 'R', 'A sub R', where the region 'R' is that region which is bounded above by the curve 'y' equals 'x squared', below by the x-axis, on the left by the y-axis, and on the right by the line 'x' equals 1. In other words, this region in here.

Now here's the way we put the squeeze on it, and we'll start this quite gradually. The first thing that we observe is that if we draw a line parallel to the x-axis through the point (1 , 1), the highest point of our region here, we construct a rectangle which contains our region 'R'. And since 'R' is contained in the rectangle, the area of the region 'R' must be less than the area of the rectangle.

But notice that this particular rectangle that we've just constructed has its base equal to 1, its

height equal to 1. Therefore by our first axiom, its area is 1 times 1, or 1. Also, since intuitively, area is a positive thing, we see just from this quick diagram that whatever the area of the region is, we now have it bounded between 0 and 1. We have an upper bound. We have a lower bound. We have this thing caught, OK?

Now, there's still a lot of space between 0 and 1. The method of exhaustion refines this idea. Namely, what we do next is we say, look at-- instead of just drawing one big rectangle like this, why don't we partition the base of our region 'R' into two equal parts? In other words, let's locate the point 1/2, which of course is midway between 0 and 1.

Now what I can do is I can circumscribe two rectangles, one of which is the rectangle whose height corresponds to the x-coordinate equaling 1/2-- that means the y-coordinate is 1/4-- and the second rectangle, the one whose x-coordinate-- the x-coordinate of the height is 1, and since 'y' equals 'x squared', the height is also 1.

You see, notice that by picking a region which is a curve which is always rising, the lowest point of each region is the point that's furthest to the left. And the highest point of each region is the point that's furthest to the right. And keeping this in mind, you see by making rectangles corresponding to the points furthest to the right in each interval, I get a rectangle which contains the portion of the region 'R' that I'm talking about. In other words, if I look at this rectangular network, my region 'R' is contained inside that. Consequently, its area is less than the area of the rectangular network.

Now, what is the area of this rectangular network? Well, the big rectangle has its base equal to 1/2 and its height equal to 1. So its area is 1/2. The small rectangle has its base equal to 1/2 and its height equal to 1/4. Since its area is the base times the height, its area is 1/8. 1/8 plus 1/2 is 5/8. In other words, the area of the rectangular network which contains 'R' is 5/8. Consequently, the area of the region 'R' must be less than 5/8.

On the other hand, notice that the lowest point in the second interval corresponds to 'x' equaling 1/2. In other words, notice that this particular rectangle, whose base is 1/2 and whose height is 1/4, that this rectangle here is inscribed in the region 'R'. Its area is 1/8. And since it's contained within the region 'R', its area must be less than the area of the region 'R'. And so we're now certain that whatever the area of the region 'R' is, it's between 1/8 and 5/8.

And by the way, notice in terms of these shaded regions, that even though this is an approximation which is too large to be the right answer, it is closer to being the right answer

than this approximation here. In other words, notice that the difference between the overestimate here and the overestimate here is this amount in here. You see, we've chopped off part of the error. But we'll talk about that, as I say, in more detail in the notes. What I'd like to do is to give you an overall view of what's happening over here.

You see, to generalize the method of exhaustion, we do what the mathematician usually does. Instead of saying, let's divide the base of the region into two equal parts, or three equal parts, or four equal parts, we say, let's divide it into 'n' equal parts. Now if we divide this base into 'n' equal parts, since the whole region-- the whole length is 0 to 1, dividing into 'n' equal parts means that my points of division will be '1/n', '2/n', '3/n', et cetera, right down to 'n/n', which is 1.

What I do now is I pick the right endpoint of each point in that partition to draw my rectangle. You see, each of these rectangles that I draw this way contains the corresponding part of my region 'R'. Consequently, the area of this rectangular network must be an upper bound for the area of my region 'R'.

Now, because it's going to be an upper bound, and because it depends on what the value of 'n' is, I will denote that upper bound by 'U sub n'. And what will 'U sub n' be? It's the sum of the areas of these circumscribed rectangles. And it's not difficult to see from this picture. Let's take a look.

Since this rectangle has its height corresponding to the x-coordinate equaling '1/n', and since the y-coordinate is the square of the x-coordinate, the height of this rectangle will be ''1/n' squared'. The base is '1/n'. So the area of that circumscribed rectangle is ''1/n' squared' times '1/n'.

Similarly, the next rectangle here has its height corresponding to the x-coordinate '2/n'. So its height is ''2/n' squared'. Keep in mind the y-coordinate is the square of the x-coordinate here. The base of all of these rectangles is '1/n'. And proceeding in this way, I finally come down to the last rectangle, whose height is what? Well, the x-coordinate is 1, so the y-coordinate is 1. To keep my form here, I'll write that as 'n/n'. Its height is therefore ''n/n' squared'. The base of that rectangle is '1/n', so the area of that last rectangle is ''n/n' squared' times '1/n'.

Now you see that's what? That's an area which is too large to be 'A sub R'. Well, let's take a look now and see what's happening here. Notice that each of these terms has an 'n cubed' in the denominator. So I can factor out the 'n cubed' term.

And what I'm left with is what? The sum of the first 'n' squares. 1 squared plus 2 squared plus 3 squared, et cetera, plus et cetera, 'n squared'. In other words, this tells me how to find 'U sub n' for any value of 'n'. And we'll talk about that some more in a little while. That's an upper bound.

In a corresponding way, I can also find a lower bound. Namely, what I'll do now is pick the smallest rectangle in each region. In other words, I'll now inscribe rectangles inside my region 'R', and therefore find a rectangular region whose area is less than that of the region 'R'.

And to do that, without going through the details, notice that essentially all I have to do is shift each of these rectangles over by one partition. Namely, notice, for example, that the smallest height in the second partition here is the one which corresponds to the height '1/n'. In other words, what I do now is, to inscribe the rectangles, I just shift everything over like this.

And leaving the details out, and letting you verify these for yourselves, I can again mimic exactly what I did before. The only difference being now, that instead of the height of the last rectangle being 'n/n', notice that it's 'n - 1' over 'n', squared. In other words, the x-coordinate is 'n - 1' over 'n'. The y-coordinate is the square of the x-coordinate. In other words, notice that this point here gives rise to the height of the lowest rectangle that can be inscribed in my last portion here. So without further ado, it turns out that 'L sub n', the lower estimate, is '1 over 'n cubed'' times the sum of the first 'n - 1' squares. OK.

Now, let's keep track of just what it is that we've done over here so far. What we've done is what? For each 'n', we have squeezed 'A sub R' between two numbers, one number being an upper approximation, one number being a lower approximation. Now notice that the 'U sub n' and the 'L sub n' are functions of 'n'. 'A sub R' is a constant. That's the area of the region 'R'.

What the method of exhaustion means is simply this. We say, look. Let's take the limit of the 'L sub n's as 'n' approaches infinity, and let's take the limit of the 'U sub n's as 'n' approaches infinity, observing that for each 'n', 'A sub R' is squeezed between these two. Consequently, since 'A sub R' is squeezed between these two for each 'n', it must be squeezed between these two when we go to the limit. In other words, whatever 'A sub R' is, it must be between these two limits.

Now here's the key step. It turns out, at least in this particular problem, that the limit of 'L sub n' as 'n' goes to infinity is the same as the limit of 'U sub n' as 'n' goes to infinity. And the best

way to see that without becoming too fancy at this stage of the game is to observe that for a given 'n', the difference between 'U sub n' and 'L sub n' is just '1/n'. And again, I'll just indicate to you from this diagram how we can see that.

Notice that back here, the way we went from 'U sub n' to 'L sub n' was we just pushed over this whole network by one unit. In other words, the rectangle that we squeezed out in going from the upper sum to the lower sum was just this particular rectangle whose height was 1 and whose base was '1/n'. In other words, the area that was kicked out in going from 'L sub n' to 'U sub n' is '1/n'. This is what this thing here says here.

This is worked out in detail in the notes. But again, I just want to go through this thing quickly so we get the idea of what's happening over here. In other words, since 'U sub n' minus 'L sub n' is '1/n', that's just another way of saying what? That the limit of this difference is just the limit of '1/n' as 'n' goes to infinity. But that's clearly 0. In other words, these two limits are the same because their difference goes to 0 in the limit.

Consequently, since these two things are the same, and 'A sub R' is caught between these two, it must be that 'A sub R' is equal to this common limit. And that's precisely what the method of exhaustion was. We squeeze what we were looking for between two estimates which converge towards each other as 'n' got large.

Now, by the way, I'm going to point out the fact that it's very, very difficult in general to find what this limit is. We're going to see that in working a few specific problems. What I thought might be informative-- without going through the details right now, I have taken the liberty of computing 'L sub n' and 'U sub n' in this problem for 'n' equals 1,000. In other words, if we divided this region into 1,000 equal parts and took the network of circumscribed rectangles and inscribed rectangles, it would turn out that the area of the inscribed rectangles would be 0.3328335, and for the circumscribed rectangles, 0.3338335.

And by the way, notice how this is borne out. Notice that the difference between these two is precisely 1/1000, and that is '1/n' with 'n' equal to 1,000. But that's not the important point here. The important point is that just looking at this decimal expansion, if I didn't know anything else, I know that 'A sub R' is caught between these two. Consequently, to two decimal places, I can be sure that 'A sub R' is 0.33. And notice that if I want more and more decimal place accuracy, especially if I have access to a desk calculator, I don't really have to compute the limit exactly. I can feed the formula into the machine and put the squeeze on and get as many

decimal place accuracy as I want.

By the way, as an aside, this is exactly what we did in high school, when we said things like let pi equal 22/7. Pi is not equal to 22/7. Among other things, 22/7 is a rational number. It's the ratio of two whole numbers. Pi is an irrational number. What people really meant was that you can't tell the difference between pi and 22/7 to two decimal places. They both begin 3.14. And consequently, if all you wanted to measure was to two decimal digits, you could use 22/7 as the approximation for pi. But if you wanted to squeeze out more places, you'd have to use a more refined approach.

By the way, let me point out that one of the reasons that I chose the curve 'y' equals 'x squared' to work with was-- you may recall that back in our discussion of mathematical induction, one of the problems that we worked on was to show how we find by induction the recipe for the sum of the first 'n' squares. And by way of review, let me recall for you the fact that the sum of the first 'n' squares was given by ''n' times 'n + 1' times '2n + 1'' over 6. OK?

In other words, going back to our recipe for 'U sub n', I can now replace 1 squared plus 2 squared plus et cetera 'n squared' by ''n' times 'n + 1' times '2n + 1'' over 6. And if I now divide through by 'n cubed' judiciously, namely canceling out one 'n' with this 'n', dividing 'n + 1' by 'n' and '2n + 1' by 'n', I get that 'U sub n' is '1/6 'n + 1' over 'n'' times ''2n + 1' over 'n''. And that can be written even more suggestively as 1/6 times '1 + '1/n'' times '2 + '1/n''.

And I can now put a very nice mathematical interpretation on this. Mainly, notice that no matter how big 'n' is, '1 + '1/n'' is bigger than 1, you see, because '1/n' is positive. And '2 + '1/n'' is bigger than 2. Consequently, whatever this is, it's bigger than 1/6 times 1 times 2, which is 1/3. In other words, for each 'n', 'U sub n' is greater than 1/3. It also happens that as 'n' gets larger, 'U sub n' gets smaller. Again, looking at this recipe, the bigger 'n' is, the bigger is our denominator, and the bigger the denominator, the smaller the fraction.

And finally, notice that as 'n' goes to infinity, 'U sub n' gets arbitrarily close to 1/3 in value, because '1/n' approaches 0 in the limit. In fact, it is 0 in the limit. So summarizing then, each 'U sub n' is bigger than 1/3. As 'n' increases, 'U sub n' decreases, and the limit of 'U sub n' as 'n' approaches infinity is 1/3.

Pictorially, what this means is that if we locate 1/3 on the number line, the 'U sub n's converge uniformly in the sense of moving steadily towards the left, towards 1/3 as the limit. Well, that's an upper squeeze. The lower squeeze comes from the fact that in a similar way, we can show

that 'L sub n' is 1/6-- and look at how close this comes to paralleling the structure of 'U sub n'-- '1 - '1/n'' times '2 - '1/n''. Minuses here instead of pluses as we had above.

Now, mimicking what we did before, notice now that for each 'n'-- since we're subtracting over here-- for each 'n', 'L sub n' is less than 1/3. But now since the fractions get smaller as 'n' gets bigger, and you're subtracting them, the difference becomes larger. In other words, now notice that each 'L sub n' is less than 1/3. As 'n' increases, 'L sub n' increases. And the limit of 'L sub n' as 'n' approaches infinity is also 1/3.

In other words, if we draw this in back here, look what the 'L sub n's are doing. They're all less than 1/3, but they move steadily towards the right as 'n' increases, pushing in on 1/3. And since 'A sub R' is always caught between these two, and these two converge relentlessly upon 1/3, it must be, by definition, if there is an area at all, that the area of the region 'R' must be exactly 1/3.

And by the way, just as a quick aside, notice how important it is that we not only have upper and lower bounds which converge. They must converge to the same value. In other words, what I'm saying is suppose all the 'U sub n's get arbitrarily close to what I call script 'L sub 1', and all the 'L sub n's get arbitrarily close to 'L2' over here. Then all I'm saying is, all we would know is that the area was someplace between 'L2' and 'L1'. We couldn't conclude that it was exactly equal to 1/3. You see, the method of exhaustion implies that all the error, all the room for doubt, is squeezed out.

By the way, there was nothing sacred about choosing 'y' equals 'x squared'. We could generalize this. And we'll do this kind of rapidly because I just want you to hear what I'm saying. This is all written out in great detail in our supplementary notes.

The idea is this. Let 'f' be any positive, continuous function on 'a', 'b' and non-decreasing. The non-decreasing part is just for the convenience of being able to locate the high and low points of each partition point conveniently. Partition 'a', 'b' into 'n' equal parts, calling the first partition point 'x sub 0', the last partition point 'x sub n', and the points in between 'x1' up through ''x 'n - 1''. In other words, we've partitioned this closed interval into 'n' equal parts.

And what we can now do is what? Pick the lowest point in each region, the highest point in each region, to form a rectangle. We can form 'U sub n' and 'L sub n'. And you see it's just what? For 'L sub n', you just pick this height, which is 'f of 'x sub 0'', 'f of a'. The base is 'delta x'. You just add these all up until you get to your last inscribed partition point. That's 'x sub 'n -

1". OK.

'U sub n' is the same thing, only now your first rectangle has its height corresponding to 'x' equals 'x1'. Consequently, the height is 'f of x1'. The base is 'delta x'. Where in each of these, 'delta x' is just what? This total length, which is 'b - a', divided into 'n' equal parts.

And to review the so-called sigma notation that we have in our course-- in this section in the textbook, and we'll have exercises on this to make sure that you are familiar with this notation, the shortcut notation for writing this sum is given by this. We add up 'f of 'x sub k'' times 'delta x'. As a subscript, 'k' is allowed to vary from 1 to 'n'. And here we form the same sum, only now the subscript 'k' varies from 0 to 'n - 1'.

At any rate, the observations are this. That 'U sub n' and 'L sub n' converge to the same limit. In fact, this isn't too hard to show. If you just algebraically subtract 'L sub n' from 'U sub n', notice that the only terms that won't cancel are the last term here. See, 'f of 'x sub n'' times 'delta x' has no counterpart here, because this ends with the subscript 'n - 1'. Similarly, there is no 0 subscript in 'U sub n', so this term won't cancel.

To make a long story short, if we just subtract 'L sub n' from 'U sub n', that difference will be 'f of 'x sub n'' times 'delta x' minus 'f of 'x sub 0'' times 'delta x'. See, that's just this difference over here. 'Delta x' is 'b - a' over 'n'. 'x sub n' was called 'b'. 'x sub 0' was called 'a'. In other words, the difference between 'U sub n' and 'L sub n' is just 'f of b' minus 'f of a' times ''b - a' over 'n''.

And the important thing to observe is that 'a', 'b', 'f of a', and 'f of b' are fixed constants. The only thing that depends on 'n' is this denominator. Consequently, as 'n' goes to infinity, since the rest of this thing is a constant, the whole term goes to 0. In other words, the difference between 'Un' and 'Ln' goes to 0 in the limit. That means that the limit of 'U sub n', as 'n' approaches infinity, equals the limit of 'L sub n' as 'n' approaches infinity. 'A sub R' is always caught between 'L sub n' and 'U sub n' by this very construction. Consequently, what this means is that the area of the region 'R' must equal this common limit.

Now by the way, this looks very hard. And it is difficult. In fact, in general, these limits are very hard to find. The point that I thought would be interesting to mention at this stage of the game is that if all we want is an estimate for the area under the curve, there are faster and better ways of doing this. You see, the beauty of this technique here is that it's a technique for

exhausting the space completely. We find the exact areas this way.

Let me give you a for instance here. Namely, let me explain to you what we mean by trapezoidal approximations. Let's take the same region, 'y' equals 'x squared'. And now, we'll divide it into two parts here. And what we'll do is we'll replace the arc of the curve by the straight line segment, the chord, that joins two points here. In other words, using the accented chalk here to illustrate this. What I'm going to do is instead of finding the area of the region 'R', I'm going to find the area of the region which has the top of 'R' replaced by these two line segments.

All I want you to observe is that this first region here is a triangle whose height is 1/4 and whose base is 1/2. Consequently, its area, being one half the base times the height, is 1/16. The second region here is a trapezoid whose bases are 1 and 1/4 and whose height is 1/2. And since the area of a trapezoid is half the sum of the bases times the height, we get that the area of this trapezoid is 5/16. Therefore, the area of the triangle plus the trapezoid is 6/16 or 3/8, which is 0.375.

Recall that 1/3, we've just seen, was the exact answer. And notice how close this approximation is to the exact answer, even with just two subdivisions over here. In fact, sparing you the details, if you now divide the base here into four equal parts, thus forming your trapezoidal approximations corresponding to what? Altitudes of 1/16, 1/4, 9/16, and 1, and adding up the areas of all these trapezoids-- by the way, this is a bit unfortunate. This is a triangle. But we can view a triangle as being a degenerate trapezoid, meaning the height here just happens to be 0 at this particular point.

But let's not worry about that. Let's just get the idea of what the trapezoidal approximation idea means. We find the area of these four trapezoids, and we use that as an approximation for the area under the curve. By the way, just coming back here for a moment, notice that because this curve is always holding water, the chord will always lie above the arc. And consequently, not only do we get an approximation using trapezoids this way that's reasonably close, but we also know by the geometry here that our approximation must be too large to be the right answer. In other words, we're going to get an over-approximation. But watch how close we come.

Leaving these details for you to verify, all I show here is that if you add up the areas of these four regions, we wind up with 44/128, which is 11/32. And 11/32 is mighty close to 1/3, being

what? 11/33. In other words, look at how close, with just four subdivisions, we get to an approximation that's good for the area under the curve, without having to put the squeeze on.

But to get the exact area, we need the squeeze. We have to push the thing from above, from below, and show that these two things that we're squeezing it between, as gruesome as it sounds, converge towards each other. The thing that we're looking for is caught between them, and consequently must be the common limit. That was the method of exhaustion as known by the ancient Greeks.

And if this seems tough to you, think of it from two points of view. One, it is tough. And secondly, if people of some 2,500 years ago were able to do this, it should be at least plausible that with a little bit of effort, we can get a good feeling for what's going on here. In fact, hopefully, the exercises to this particular unit will make this a little bit easier for you to see in action.

But at any rate-- let me just make a few asides over here. The deeper asides will be made in greater detail in our supplementary notes. Notice that when we dealt with differential calculus, it was very, very important in differential calculus to have smooth curves. For finding areas, all you need are continuous curves.

In other words, for example, we can find the area of a square. But certainly, a square has sharp corners. In other words, those corners are continuous, but they're not differentiable in terms of the language of differential calculus. The point is that even continuity can be weakened.

Let me give you a definition. 'f' is called piecewise continuous on the interval from 'a' to 'b' if and only if 'f' is continuous except at a finite number of points where it has jump discontinuities. For example, in terms of this particular diagram, notice that my curve, 'y' equals 'f of x', is discontinuous at a finite number of points, namely at 'c1' and 'c2', two points. Notice that what happens at those two points is you have just a finite jump discontinuity.

The point is that whereas this curve is not continuous, notice that since a straight line having no thickness has no area, if we replace the given curve by this one-- you see, putting in these vertical lines-- notice that this does form a closed region. And to find the area of this closed region, I can pretend it was made up of these three particular regular regions. In other words, that even if I just have jumps, since the jump does not contribute towards the area, there is no harm done when one talks about piecewise continuous in finding areas rather than

continuous.

A second aside-- well, I've actually called this one aside number one, because this definition I didn't call an aside-- but the idea is this. Notice that up until now, we were assuming that we had to form our 'U sub n's and our 'L sub n's by choosing either the left endpoint of a partition or the right endpoint. What I'd like you to see also as a generalization is that if I pick any point between these two extremes and form-- look, I'll call it 'c sub k' between these two points here-- and form this particular sum, where 'c sub k' replaces either 'x sub k' or 'x sub 'k - 1'', that this sum here also gives me the area of the region 'R'.

In other words, that whatever this sum is, it's caught between 'U sub n' and 'L sub n'. And since 'U sub n' and 'L sub n' have a common limit of 'A sub R'-- in other words, since both of these squeeze in towards 'A sub R'-- this being caught between them must also be 'A sub R'. That's what I've said over here.

And to see this thing in terms of a picture, all we're saying is that when you pick the lowest point in the interval, you get the rectangle that contributes to 'L sub n'. When you pick the highest point in the interval, you get the rectangle that contributes to 'U sub n'. Consequently, for any point between these two extremes-- call that--that's what 'c sub k' is, it's any point between these two-- pick the height that corresponds to that.

And whatever rectangle you form this way, that rectangle must have a greater area than the rectangle that went into forming 'L sub n', but a lesser area than the rectangle that went into forming 'U sub n'. And consequently, that's where this particular result comes from. Again, this is done in more detail in the notes, but I think some of these things you should hear me say out loud rather than to rely on your reading it.

And finally, one more important point. Up until now, it's been clear that since we're talking about area, our region has to lie above the x-axis. We can remove the restriction that 'f' be non-negative if we replace area by net area. In other words, all I want you to see here is that if my region happens to look something like this, notice that between 'c' and 'b', 'f of x' is negative. Consequently, when I form things of the form 'f of x' times 'delta x', 'delta x' being positive, 'f of x' being negative, is going to give me a negative result.

In other words, algebraically, if I form my summation, this portion in here will give me a positive result. This portion in here will give me a negative result. Consequently, working algebraically, what I will find is not the true area but the positive minus this amount here. In other words,

what I'll find is the net area. And again, this will all be worked out in the notes.

In addition, there'll be other refinements made in the notes, such as that these partitions don't have to be into equal parts or things like this. The important point from today's lecture is this. Observe that this entire study of area is done independently of differential calculus. In our next lecture, we are going to show a truly remarkable, a wonderful relationship between the differential calculus of before and the so-called area or integral calculus that we did today. At any rate, until next time, goodbye.