Finding the Area under a Curve

As I discuss in Chapter 9, the most fundamental meaning of a derivative is that it's a rate, a *this per that* like *miles per hour,* and that when you graph the *this* as a function of the *that* (like *miles* as a function of *hours*), the derivative becomes the slope of the function. In other words, the derivative is a rate, which on a graph appears as a slope.

It sort of works the same way with integration. The most fundamental meaning of integration is to add up. And when you depict integration on a graph, you can see the adding up process as a summing up of little bits of area to arrive at the total area under a curve. Consider Figure 13-4.

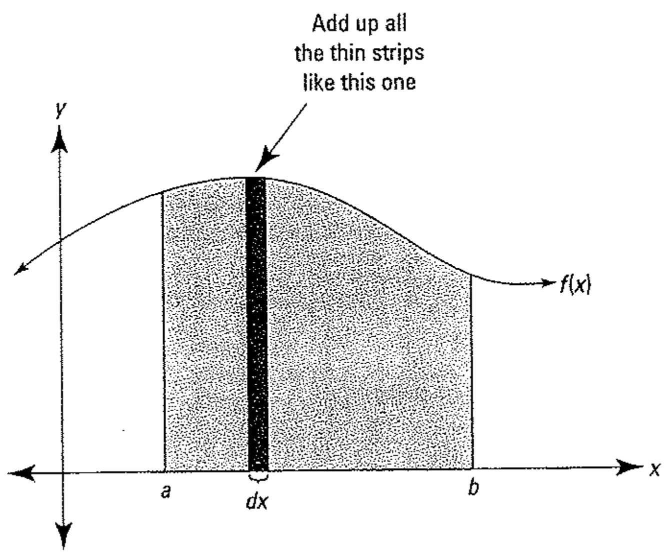

Figure 13-4: Integrating $f(x)$ from a to b means finding the area under the curve between a and b.

The shaded area in Figure 13-4 can be calculated with the following integral:

$$\int_a^b f(x)\,dx$$

Look at the thin rectangle in Figure 13-4. It has a height of $f(x)$ and a width of dx (a little bit of x), so its area (*length* times *width*, of course) is given by $f(x) \cdot dx$. The above integral tells you to add up the areas of all the narrow rectangular strips between a and b under the curve $f(x)$. As the strips get narrower and narrower, you get a better and better estimate of the area. The power of integration lies in the fact that it gives you the *exact* area by sort of adding up an infinite number of infinitely thin rectangles.

Regardless of what the tiny bits are that you're adding up — they could be little bits of distance or volume or energy (or just area) — you can represent the summation as an adding up of the areas of thin rectangular strips under a curve. If the units on both the x and y axes are, say, *feet*, then each thin rectangle measures so many feet by so many feet, and its area — *length* times *width* — is some number of *square feet.* In this case, the total area of all the rectangles gives you the area under the curve between a and b (though not to scale).

If, on the other hand, the units on the x-axis are *hours* (t) and the y-axis is labeled in *miles per hour,* then, because *rate* times *time* equals *distance,* the area of each rectangle represents an amount of distance and the total area gives you the total distance traveled during the given time interval. Or if the x-axis is labeled in *hours* (t) and the y-axis in *kilowatts* of electrical power — in which case the curve, $f(t)$, gives power usage as a function of time — then the area of each rectangular strip (*kilowatts* times *hours*) represents a number of *kilowatt-hours* of energy. In that case, the total area under the curve gives you the total number of kilowatt-hours of energy consumption between two points in time.

Figure 13-5 shows how you would do the lamp volume problem — from earlier in this chapter — by adding up areas. In this graph, the function $A(h)$ gives the cross-sectional *area* of a thin pancake slice of the lamp as a function of its height measured from the bottom of the lamp. So this time, the h-axis is labeled in *inches* (that's h as in *height* from the bottom of the lamp), and the y-axis is labeled in *square inches,* and thus each thin rectangle has a width measured in inches and a height measured in square inches. Its area, therefore, represents *inches* times *square inches,* or *cubic inches* of volume.

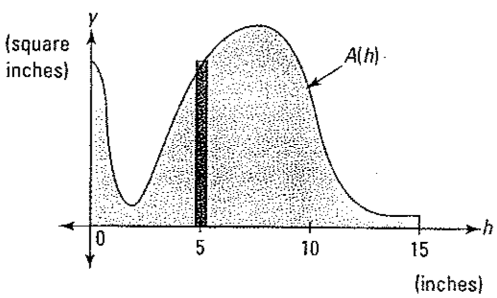

Figure 13-5: This shaded *area* gives you the *volume* of the base of the lamp in Figure 13-1.

The area of the thin rectangle in Figure 13-5 represents the *volume* of the thin pancake slice of the lamp 5 inches up from the bottom of the base. The total shaded area and thus the volume of the lamp's base is given by the following integral:

$$Volume = \underbrace{cross\text{-}sectional\ area}\times\underbrace{thickness}$$

$$V = \int_0^{15} \overbrace{A(h)}\,\overbrace{dh}$$

which means that you add up the volumes of all the thin pancake slices from 0 to 15 inches (that is, from the bottom to the top of the lamp's base), each slice having a volume given by $A(h)$ (its cross-sectional area) times dh (its height or thickness).

Dealing with Negative Area

In the examples involving volume, distance, and energy (from the previous section), you're always adding up *positive* bits of something. This is usually the case with practical problems because you can't, for instance, have a negative volume of water or use a negative number of kilowatt-hours of energy. However, you will sometimes integrate functions that go into the negatives — that's below the x-axis. Here are a few pointers for when that happens.

When using integration to calculate area, area *below* the x-axis counts as *negative* area. The total area between a and b for some curve $f(x)$ — given by the integral $\int_a^b f(x)\,dx$ — is really a *net* area where the total area below the x-axis (and above the curve) is subtracted from the total area above the x-axis (and below the curve).

Think of the x-axis as ground level, areas above the x-axis as mounds of earth, and areas below the x-axis as holes in the ground. The net area then represents the amount of earth left above ground level after you use the earth in the mounds to fill in the holes. (This net can be a negative amount.)

In Chapter 16, I show you how to calculate the total area between a curve and the x-axis where all area sections are counted as positive.

Okay, enough of this introductory stuff. In the next section, you actually calculate some areas.

Approximating Area

Before explaining how to calculate exact areas, I want to show you how to approximate areas. The approximation method is useful not only because it lays the groundwork for the exact method — integration — but because for some curves, integration is impossible, and an approximation of area is the best you can do.

Approximating area with left sums

Say you want the exact area under the curve $f(x) = x^2 + 1$ between 0 and 3. See the shaded area on the graph on the left in Figure 13-6.

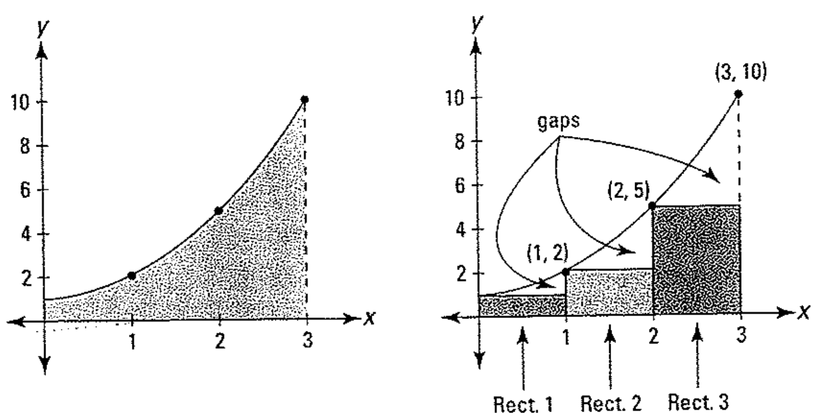

Figure 13-6: The exact area under $f(x) = x^2+1$ between 0 and 3 (left) is approximated by the area of three rectangles (right).

First, you get a rough estimate of the area by drawing three rectangles under the curve, as shown on the right in Figure 13-6, and then determining the sum of their areas.

The rectangles in Figure 13-6 represent a so-called *left sum* because the upper *left* corner of each rectangle touches the curve. Each rectangle has a width of 1 and the height of each is given by the height of the function at the rectangle's left edge. So, rectangle number 1 has a height of $f(0) = 0^2 + 1 = 1$; its area (*length* × *width* or *height* × *width*) is thus 1 × 1, or 1. Rectangle 2 has a height of $f(1) = 1^2 + 1 = 2$, so its area is 2 × 1, or 2. And rectangle 3 has a height of $f(2) = 2^2 + 1 = 5$, so its area is 5 × 1, or 5. Adding these three areas gives you a total of 1 + 2 + 5, or 8. You can see that this is an underestimate of the total area under the curve because of the three gaps between the rectangles and the curve shown in Figure 13-6.

For a better estimate, double the number of rectangles to six. Figure 13-7 shows six "left" rectangles under the curve and also how the six rectangles begin to fill up the three gaps you see in Figure 13-6.

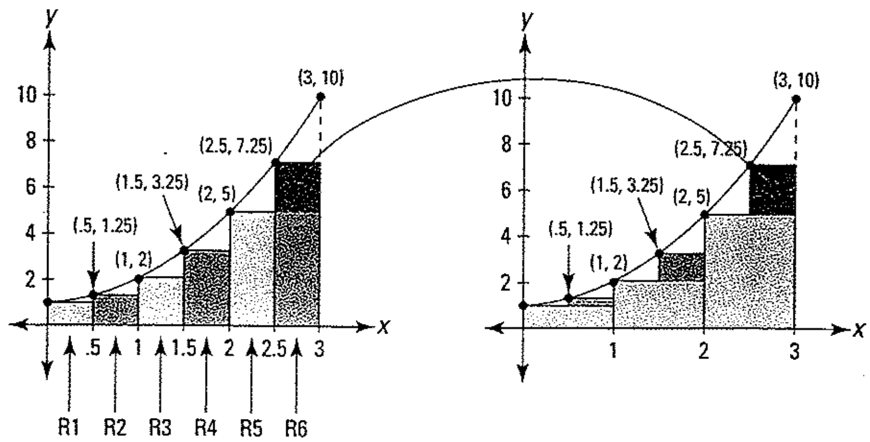

Figure 13-7: Six "left" rectangles approximate the area under $f(x) = x^2 + 1$.

See the three small shaded rectangles in the graph on the right in Figure 13-7? They sit on top of the three rectangles from Figure 13-6, and they represent how much the area estimate has improved by using six rectangles instead of three.

Now total up the areas of the six rectangles. Each has a width of 0.5 and the heights are $f(0)$, $f(0.5)$, $f(1)$, $f(1.5)$, and so on. I'll spare you the arithmetic. Here's the total: $0.5 + 0.625 + 1 + 1.625 + 2.5 + 3.625 = 9.875$. This is a better estimate, but it's still an underestimate because of the six small gaps you can see on the left graph in Figure 13-7.

Table 13-1 shows the area estimates given by 3, 6, 12, 24, 48, 96, 192, and 384 rectangles. You don't have to double the number of rectangles each time like I've done here. You can use any number of rectangles you want. I just like the doubling scheme because, with each doubling, the gaps are plugged up more and more in the way shown in Figure 13-7.

Table 13-1 Estimates of the Area under $f(x) = x^2 + 1$ Given by Increasing Numbers of "Left" Rectangles

Number of Rectangles	Area Estimate
3	8
6	9.875
12	~10.906
24	~11.445
48	~11.721
96	~11.860
192	~11.930
384	~11.965

Here's the fancy-pants formula for a left-rectangle sum:

The Left Rectangle Rule: You can approximate the exact area under a curve between a and b, $\int_a^b f(x)\,dx$, with a sum of *left* rectangles given by the following formula. In general, the more rectangles, the better the estimate.

$$L_n = \frac{b-a}{n}\left[f(x_0) + f(x_1) + f(x_2) + \ldots\ldots\ldots + f(x_{n-1})\right]$$

Where n is the number of rectangles, $\frac{b-a}{n}$ is the width of each rectangle, and the function values are the heights of the rectangles.

I better explain this formula a bit. Look back to the six rectangles shown in Figure 13-7. The width of each rectangle equals the length of the total span from 0 to 3 (which of course is 3 – 0, or 3) divided by the number of rectangles, 6. That's what the $\frac{b-a}{n}$ does in the formula.

Now, what about those xs with the subscripts? The x-coordinate of the *left* edge of rectangle 1 in Figure 13-7 is called x_0, the *right* edge of rectangle 1 (which is the same as the left edge of rectangle 2) is at x_1, the right edge of rectangle 2 is at x_2, the right edge of rectangle 3 is at x_3, and so on all the way up to the right edge of rectangle 6, which is at x_6. For the six rectangles in Figure 13-7, x_0 is 0, x_1 is 0.5, x_2 is 1, x_3 is 1.5, x_4 is 2, x_5 is 2.5, and x_6 is 3. The heights of the six left rectangles in Figure 13-7 occur at their left edges, which are at 0, 0.5, 1, 1.5, 2, and 2.5 — that's x_0 through x_5. You don't use the right edge of the last rectangle, x_6, in a left sum. That's why the list of function values in the formula stops at x_{n-1}. This all becomes clearer — cross your fingers — when you look at the formula for *right* rectangles in the next section.

Here's how to use the formula for the six rectangles in Figure 13-7:

$$L_6 = \frac{3-0}{6}\left[f(x_0) + f(x_1) + f(x_2) + f(x_3) + f(x_4) + f(x_5)\right]$$

$$= \frac{1}{2}\left[f(0) + f(0.5) + f(1) + f(1.5) + f(2) + f(2.5)\right]$$

$$= \frac{1}{2}(1 + 1.25 + 2 + 3.25 + 5 + 7.25)$$

$$= \frac{1}{2}(19.75)$$

$$= 9.875$$

Note that had I distributed the width of ½ to each of the heights after the third line in the above solution, you would have seen the sum of the areas of the rectangles — which you saw one page back. The formula just uses the shortcut of first adding up the heights and then multiplying by the width.

Whether approximating areas or finding exact areas, areas below the x-axis count as *negative*. See section "Dealing with Negative Areas" earlier in this chapter.

Approximating area with right sums

Okay, now estimate the same area under $f(x) = x^2 + 1$ from 0 to 3 with *right* rectangles. This method works just like the left sum method except that each rectangle is drawn so that its *right* upper corner touches the curve. See Figure 13-8.

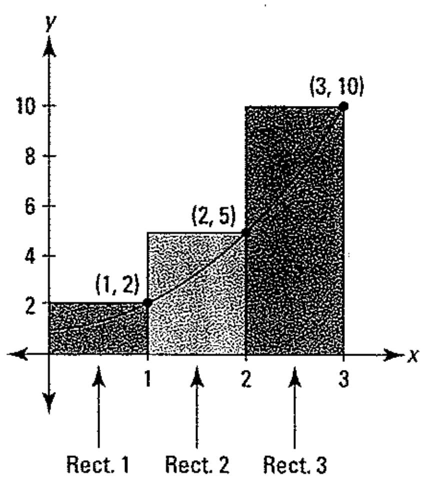

Figure 13-8: Three *right* rectangles used to approximate the area under $f(x) = x^2 + 1$

The heights of the three rectangles in Figure 13-8 are given by the function values at their *right* edges: $f(1) = 2$, $f(2) = 5$, and $f(3) = 10$. Each rectangle has a width of 1, so the areas are 2, 5, and 10, which total 17. You don't have to be a rocket scientist to see that this time you get an *over*estimate of the actual area under the curve, as opposed to the *under*estimate that you get with the left-rectangle method I detail in the previous section (more on that in a minute). Table 13-2 shows the improving estimates you get with more and more right rectangles.

Table 13-2 Estimates of the Area under $f(x) = x^2 + 1$ Given by Increasing Numbers of "Right" Rectangles

Number of Rectangles	Area Estimate
3	17
6	14.375
12	~13.156
24	~12.570
48	~12.283
96	~12.141
192	~12.070
384	~12.035

Looks like these estimates are also headed toward 12. Here's the formula for a right rectangle sum.

The Right Triangle Rule: You can approximate the exact area under a curve between a and b, $\int_a^b f(x)\,dx$, with a sum of *right* rectangles given by the following formula. In general, the more rectangles, the better the estimate.

$$R_n = \frac{b-a}{n}\left[f(x_1) + f(x_2) + f(x_3) + \ldots\ldots\ldots\ldots + f(x_n)\right]$$

where n is the number of rectangles, $\frac{b-a}{n}$ is the width of each rectangle, and the function values are the heights of the rectangles.

Now if you compare this formula to the one for a left rectangle sum (in the previous "Approximating area with left sums" section), you get the complete picture about those subscripts. The two formulas are the same except for one thing. Look at the sums of the function values in both formulas. The right sum formula has one value, $f(x_n)$, that the left sum formula doesn't have, and the left sum formula has one value, $f(x_0)$, that the right sum formula doesn't have. All the function values between those two appear in both formulas. You can get a better handle on this by comparing the three left rectangles from Figure 13-6 to the three right rectangles from Figure 13-8. Their areas and totals, which we earlier calculated, are

Three left rectangles: $1 + 2 + 5 = 8$

Three right rectangles: $2 + 5 + 10 = 17$

The sums of the areas are the same except for the left-most left rectangle and the right-most right rectangle. Both sums include the rectangles with areas 2 and 5. If you look at how the rectangles are constructed, you can see that the second and third rectangles in Figure 13-6 are the same as the first and second rectangles in Figure 13-8.

One last thing on this. The difference between the right rectangle total area (17) and the left rectangle total area (8) — that's 17 minus 8, or 9, in case you love calculus but don't have the basic subtraction thing down yet — comes from the difference between the areas of the two "end" rectangles just discussed — 10 minus 1 is also 9. All the other rectangles are a wash, no matter how many rectangles you have.

Approximating area with midpoint sums

A third way to approximate areas with rectangles is to make each rectangle cross the curve at the midpoint of its top side. A midpoint sum is a *much* better estimate of area than either a left or a right sum. Figure 13-9 shows why.

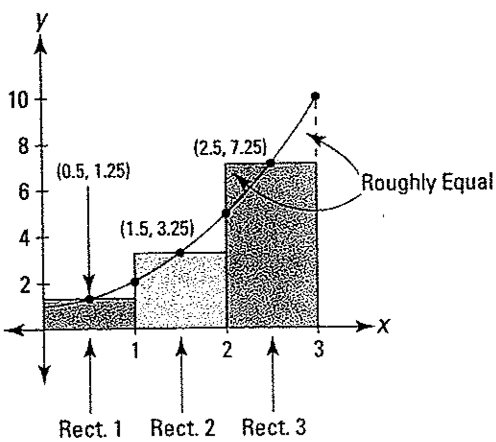

Figure 13-9: Three *midpoint* rectangles give you a much better estimate of the area under $f(x) = x^2 + 1$.

You can see in Figure 13-9 that the part of each rectangle that's above the curve looks about the same size as the gap between the rectangle and the curve. A midpoint sum produces such a good estimate because these two errors roughly cancel out each other.

For the three rectangles in Figure 13-9, the widths are 1 and the heights are $f(0.5) = 1.25$, $f(1.5) = 3.25$, and $f(2.5) = 7.25$. The total area comes to 11.75. Table 13-3 lists the midpoint sums for the same number of rectangles in Tables 13-1 and 13-2.

Table 13-3 Estimates of the Area under $f(x) = x^2 + 1$ Given by Increasing Numbers of "Midpoint" Rectangles

Number of Rectangles	Area Estimate
3	11.75
6	11.9375
12	~11.9844
24	~11.9961
48	~11.9990
96	~11.9998
192	~11.9999
384	~11.99998

If you had any doubts that the left and right sums in Tables 13-1 and 13-2 were heading to 12, Table 13-3 should dispel them. Yes, in fact, the exact area is 12 — sorry to give away the ending. And to see how much faster the midpoint approximations approach the exact answer of 12 than the left or right approximations, compare the three tables. The error with 6 midpoint rectangles is about the same as the error with 192 left or right rectangles! Here's the mumbo jumbo.

The Midpoint Rule: You can approximate the exact area under a curve between a and b, $\int_a^b f(x)\,dx$, with a sum of *midpoint* rectangles given by the following formula. In general, the more rectangles, the better the estimate.

$$M_n = \frac{b-a}{n}\left[f\left(\frac{x_0+x_1}{2}\right)+f\left(\frac{x_1+x_2}{2}\right)+f\left(\frac{x_2+x_3}{2}\right)+ \ldots\ldots\ldots + f\left(\frac{x_{n-1}+x_n}{2}\right)\right]$$

where n is the number of rectangles, $\frac{b-a}{n}$ is the width of each rectangle, and the function values are the heights of the rectangles.

All three sums — left, right, and midpoint — are called *Riemann sums* after the German mathematician G. F. B. Riemann (1826-66). Basically, any approximating sum made up of rectangles is a Riemann sum, including weird sums consisting of rectangles of unequal width. Luckily, you won't have to deal with those in this book or your calculus course.