Chapter 15

Deeper Exploration of Logs and Exponentials[1]

Not all of the models that we can use to describe real world data are based on power functions or polynomials. In fact, we saw in earlier chapters that there are many situations where exponential or logarithmic models may be needed. We also developed a way of interpreting the coefficients of these models using parameter analysis. However, parameter analysis does not give us the power needed to locate maxima and minima for such models. Only calculus tools, specifically the derivative, can do this. In this chapter, you will work with the derivatives of exponential and logarithmic functions, and you will further apply these tools to analyze models of the business world. When you have finished this chapter, you will know how to deal with many of the basic functions found in the real world. The symbolic analysis portion of this chapter will show you how, using multiplication, division and composition of models, we can build many more types of models and analyze them using calculus.

- *As a result of this chapter, students will learn*
  - √ How to use the calculus tool of derivatives to analyze models involving logarithms
  - √ How to use derivatives to analyze models involving exponentials
  - √ How compound interest works, including continuously compounded interest
- *As a result of this chapter, students will be able to*
  - √ Take derivatives of exponential functions
  - √ Take derivatives of logarithmic functions
  - √ Compute compound interest

15.1 Logarithms and their derivatives

As we have seen, there are many times when the model you develop will need to go beyond the power or polynomial models. For a multitude of reasons, the exponential and logarithmic models are the next most common models:

1. Exponentials are easy to interpret based on percent changes; thus, they can easily represent mathematically the process of accruing interest for loans or other accounting-related phenomena.

2. Logarithms are useful for dealing with some of the potential problems in modeling data, specifically the problem of non-constant variance.

3. Logarithms can be useful for simplifying many other models for analysis, since logarithms (remember the properties listed in section 12.2.1) can be used to convert many expressions involving multiplication and division into addition and subtraction problems.

These reasons alone are sufficient to justify learning how to properly use derivatives to analyze such functions. Before we get to technical, though, it's worth looking at the functions themselves and trying to figure out what we expect to happen. If we look at a graph of an exponential function, we notice immediately that the slope is always increasing. The slope is always positive, and the curve is always concave up. Thus, we expect the derivative to (a) always be positive and (b) increase as x increases. While these observations seem to tell us a lot, we have to remember that we are only looking at a small portion of the complete graph of the function, so it is possible that somewhere far from where we are looking this behavior will change. Once we have the derivative in hand, however, we can find out if this happens. (You'll have a chance to work with this in one of the problems at the end of the chapter.)

This is all in stark contrast to logarithmic functions. The graph of a logarithmic function shows more complex behavior. While it is true that the graphs seems to be always increasing, notice that the slope is decreasing as we move to the right. Thus, the logarithmic function seems to be concave down everywhere, even though it is increasing. Is it possible that somewhere far down the line the graph actually starts to decrease? We must also bear in mind that whatever we learn about one of the functions can be applied to the other, since logarithms and exponentials are inverses of each other.

The following section is devoted to learning about the derivatives of logarithmic functions. The development of this will mimic the path we took in chapter 14 to develop the derivative formulas for the power and polynomial models. Along the way we will encounter some other rules for taking derivatives: the chain rule, product rule and quotient rule. These will give us the ability to differentiate (take the derivative of) functions that are made of combinations of basic functions like logarithms and power functions. The next section will explore the exponential function and its applications to one of the most frequently used economics and business scenarios: compound interest.

15.1.1 Definitions and Formulas

Composition of Functions This is one way of making a new function from two old functions. Essentially, we take one function and "plug it into" the other function. For example, if we compose $f(x) = 2x^3$ and $g(x) = 4x-5$ we get either $h(x) = (f \circ g)(x) = f(g(x)) = 2(4x-5)^3$ or we get $k(x) = (g \circ f)(x) = 4(2x^3) - 5$ depending on the order of the composition. In general, the two orders are not the same.

Chain rule We'll be using this rule a lot. The symbolic analysis section will explain it in more detail, but the basic idea is that if you have a function composed with another function and you need the derivative of the combined object, you use the chain rule to "chain together" derivatives of each function. For example, if we start with the functions $f(x)$ and $g(x)$ above and compose them into $h(x)$ the new function h is no longer a simple power function or polynomial (although we could multiply it out into a polynomial.) But since it is composed of these simpler functions, we can still take it's derivative. In fact, the chain rule says that

$$\frac{d}{dx} f(g(x)) = \frac{df}{dg} \cdot \frac{dg}{dx}.$$

Thus $h'(x) = [df/dg][dg/dx] = [2 \cdot 3g(x)^2] \cdot [4] = 24(4x-5)^2$. A derivation and proof of the chain rule are somewhat technical; for now, think of this as a way of chaining together the derivatives so the objects which look like (but aren't really) fractions will cancel out. In the above illustration of the chain rule, the first "fraction" has the numerator we want (df) and the second "fraction" has the denominator we want (dx). Each of these "fractions" has a dg term that "cancels out" to give the derivative we want: df/dx.

Product rule The product rule allows us to take derivatives of functions that are products of simpler functions. It says that

$$\frac{d}{dx}[f(x) \cdot g(x)] = g(x) \cdot \frac{df}{dx} + f(x) \cdot \frac{dg}{dx}.$$

The proof of this rule will be given in the symbolic analysis section, and will make use of the derivative of a logarithm and the chain rule.

Quotient rule The product rule allows us to take derivatives of functions that are products of simpler functions. It says that

$$\frac{d}{dx}\left[\frac{f(x)}{g(x)}\right] = \frac{g(x)f'(x) - f(x)g'(x)}{[g(x)]^2}$$

The proof of this rule will be given in the symbolic analysis section, and will make use of the derivative of a logarithm and the chain rule.

15.1.2 Worked Examples

Example 15.1. Derivative formula for logarithmic models
In example 1 (page 318) we developed a model for the cost of electricity as a function of the number of units of electricity produced. Marginal analysis can help us to make more specific sense of this model by helping us to interpret how much each unit of electricity affects the total cost of producing the electricity. The model had the form $f(x) = A + B * \ln(x)$:

$$\text{Cost} = \text{-63,993.30} + \text{16,653.55} \cdot \text{Log(Units)}$$

Suppose that we are currently producing 500 units of electricity. How much would it cost to produce one more unit of electricity? We can put this into an Excel spreadsheet to compute it fairly easily. The results are shown below, and were obtained by setting up a formula for the difference quotient of the function, with a variable for h so that we can let h get very small. This lets us see what the instantaneous rate of change of the cost function is (this data is reproduced in the first worksheet of C15 LogDerivative.xls).

A	-63,993.30				
B	16,653.55				
X	500				
H	X+H	F(X)	F(X+H)	DF = F(X+H)-F(X)	DF/H
10	510	39501.99	39831.77	329.7840438	32.9784
1	501	39501.99	39535.26	33.27383724	33.27384
0.1	500.1	39501.99	39505.32	3.330376973	33.30377
0.01	500.01	39501.99	39502.32	0.333067669	33.30677
0.001	500.001	39501.99	39502.02	0.033307067	33.30707

From this, it seems that when current production is at 500 units, each additional unit of electricity will cost approximately \$33.31. In contrast, if are currently producing 1,000 units of electricity, the marginal cost is about \$16.65/unit.

A	-63,993.30				
B	16,653.55				
X	1000				
H	X+H	F(X)	F(X+H)	DF = F(X+H)-F(X)	DF/H
10	1010	51045.35	51211.06	165.7083324	16.57083
1	1001	51045.35	51061.99	16.64522877	16.64523
0.1	1000.1	51045.35	51047.01	1.665271738	16.65272
0.01	1000.01	51045.35	51045.51	0.166534667	16.65347
0.001	1000.001	51045.35	51045.36	0.016653542	16.65354

Now, what can this tell us about the derivative formula of a logarithmic function? Quite a lot, actually. Notice that as the production level increased (from 500 to 1000 units) the derivative (approximated by the column labeled "DF/H") decreased. Thus, we expect the derivative of a logarithmic function to be a decreasing function. This makes perfect sense when looking at the graph of a logarithmic function, since the graph "flattens out" the farther you move along the x-axis. We can repeat the same method of analysis from chapter ?? to build a table of values for $[\ln(x)]'$. If we plot these values, we get a graph much like the one below (see the second worksheet of "Ch15 LogDerivative.xls").

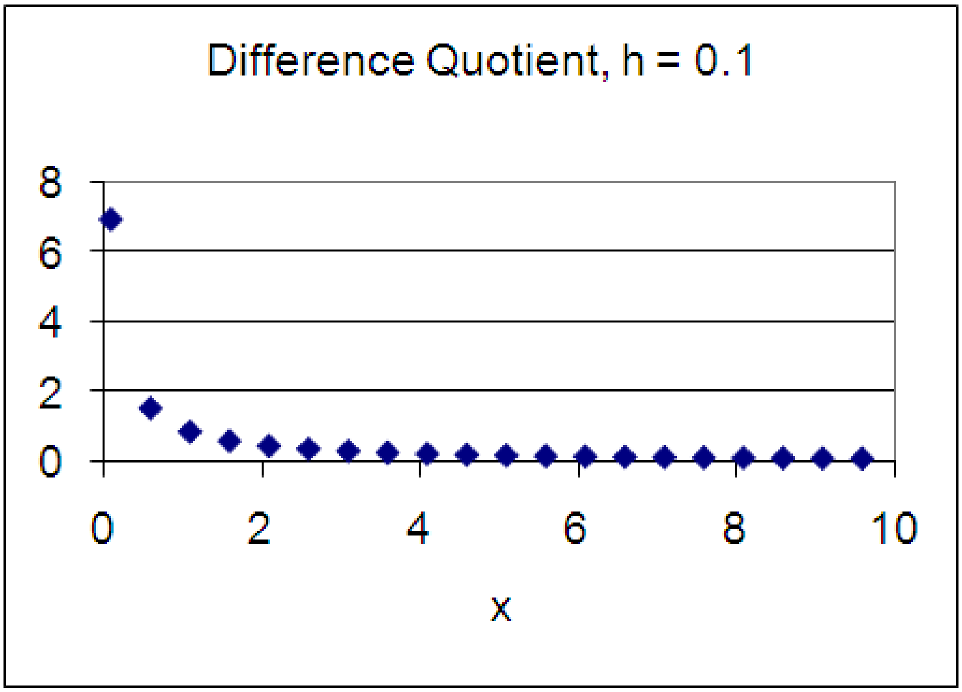

Figure 15.1: Difference quotient of a basic logarithmic function.

Notice that the difference quotient appears to be very similar to the inverse function, $f(x) = x^{-1}$. This is a power function, so we can superimpose a trend line on this data using a power function. If we do, we find remarkable agreement, even with $h = 0.1$. Reducing h will, however, quickly achieve a nearly perfect fit for the inverse function to the difference quotient. While we have not truly proven this, we can assert with some confidence that

$$\frac{d}{dx}\ln(x) = \frac{1}{x}.$$

Now, we can use this along with what we already know about derivatives to determine the derivative of a more complete logarithmic model:

$$\frac{d}{dx}\left(A + B\ln(x)\right) = \frac{d}{dx}(A) + \frac{d}{dx}\left(B\ln(x)\right) = 0 + B\frac{d}{dx}\ln(x) = \frac{B}{x}.$$

Thus, we expect that the derivative of the logarithmic function above (with $A = -63,993.30$ and $B = 16,653.55$) to be equal to $B/x = 16,653.55/x$. So when the production level is 500, the derivative should be 16,653.55/500 = 33.3071, which is extremely close to the number we estimated using the difference quotient above. If the production level is 1000, we expect the derivative to be 16,653.55/1000 = 16.65355, which is again very close to the estimates determined earlier.

Example 15.2. Derivative of a logarithmic function
Find the derivative of the function $f(x) = 3 - 2\ln(5x)$ with respect to the variable SxS.

$f'(x) = \frac{d}{dx}(3 - 2\ln(x))$	
$= \frac{d}{dx}(2) + \frac{d}{dx}(-2\ln(x))$	Using the sum rule for derivatives
$= 0 - 2\frac{d}{dx}\ln(5x)$	Derivative of a constant is zero AND derivative of a constant times a function
$= -2 \cdot \frac{1}{5x} \cdot \frac{d}{dx}(5x)$	Using the chain rule
$= -2 \cdot \frac{1}{5x} \cdot 5$	Computing the derivative of the linear function
$= -\frac{2}{x}$	Simplifying the derivative

Example 15.3. A more complex derivative
Now for the hardest example yet. Find the derivative of the compound function below:

$$h(x) = \frac{(3 + 2x + x^2)(5 + x)^4}{2 + 3x + 7x^2}.$$

There are several different paths we could take through this problem. We'll do it here by using the logarithmic derivative (one could use the chain, product and quotient rules all at once also). To do this, we take the natural logarithm of both sides and simplify the resulting mess that appears on the right hand side.

$$\begin{aligned}
\ln(h(x)) &= \ln\left[\frac{(3 + 2x + x^2)(5 + x)^4}{2 + 3x + 7x^2}\right] \\
&= \ln(3 + 2x + x^2) + \ln(5 + x)^4 - \ln(2 + 3x + 7x^2) \\
&= \ln(3 + 2x + x^2) + 4\ln(5 + x) - \ln(2 + 3x + 7x^2)
\end{aligned}$$

Taking the derivative is now a matter of using the chain rule, piece by piece. For example, we know that the derivative of the left hand side with respect to the variable x is just $h'(x)/h(x)$, where $h'(x)$ is the derivative we really want. Now we need to take the derivative of the right hand side; we'll do it in three parts, one for each term on the right hand side.

$$\begin{aligned}
\frac{d}{dx}\ln(3 + 2x + x^2) &= \frac{1}{3 + 2x + x^2} \cdot \frac{d}{dx}(3 + 2x + x^2) = \frac{2 + 2x}{3 + 2x + x^2} \\
\frac{d}{dx}[4\ln(5 + x)] &= 4 \cdot \frac{d}{dx}\ln(5 + x) = 4 \cdot \frac{1}{5 + x} \cdot \frac{d}{dx}(5 + x) = \frac{4}{5 + x} \\
\frac{d}{dx}\ln(2 + 3x + 7x^2) &= \frac{1}{2 + 3x + 7x^2} \cdot \frac{d}{dx}(2 + 3x + 7x^2) = \frac{3 + 14x}{2 + 3x + 7x^2}
\end{aligned}$$

Now we can put this all together to get

$$\frac{1}{h(x)}\frac{dh}{dx} = \frac{2+2x}{3+2x+x^2} + \frac{4}{5+x} - \frac{3+14x}{2+3x+7x^2}$$

Cross multiplying by $h(x)$ then gives us the derivative of h with respect to x

$$\frac{dh}{dx} = \left[\frac{2+2x}{3+2x+x^2} + \frac{4}{5+x} - \frac{3+14x}{2+3x+7x^2}\right] \cdot \frac{(3+2x+x^2)(5+x)^4}{2+3x+7x^2}$$

After a great deal of work, this can simplify to

$$\frac{dh}{dx} = \frac{(2+2x)(5+x)^4}{2+3x+7x^2} + \frac{4(3+2x+x^2)(5+x)^3}{2+3x+7x^2} + \frac{(3+14x)(3+2x+x^2)(5+x)^4}{(2+3x+7x^2)^2}$$

If we get a common denominator, we can further simplify this, but it doesn't really help.

15.1.3 Exploration 15A: Logs and distributions of data

Part 1. Open the data file "C15 WaitTimes.xls". This first worksheet (labeled "part 1") contains a list of 400 service times at Beef 'n Buns. Generate a histogram of the data to match the histogram below. Notice that the distribution of service times is significantly right-skewed.

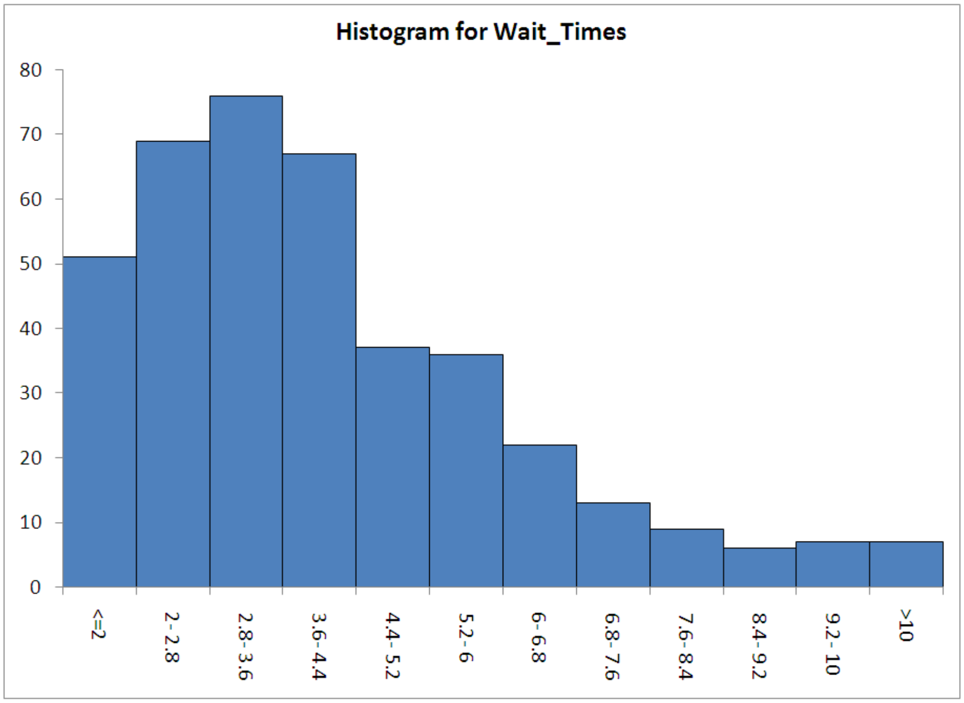

Figure 15.2: Histogram of wait times at Beef n' Buns, showing the right-skewed distribution.

One of the assumptions about linear regression involves the distribution of the data. If were to try and create a regression model to predict service times, we would find this model to have significant error, due the data's skewness. There is, however, an easy way to normalize the data in order to produce a better model. Use StatPro to create a column of wait times that has been transformed by taking the natural logarithm. Create a histogram of these logged wait times. What do you see? Under what circumstances might this be a useful tool for model building?

Part 2. The second worksheet in the file illustrates another property of logarithms. In fact, it is this property that makes the process in part 1 work. This sheet shows a graph of the natural logarithm, along with vertical and horizontal lines passing through the data points. From looking at the graph, which has points that are equally spaced in the x direction, can you explain why logarithmic functions are sometimes described as "compressing data"? Your task is to first change the x coordinates of the points (in column B) so that the change in y between successive points is the same - exactly the same. Then, use the other information in the data table and what you know about the property of logarithms to explain why this particular spacing of x values solves the problem. What other x values would work?

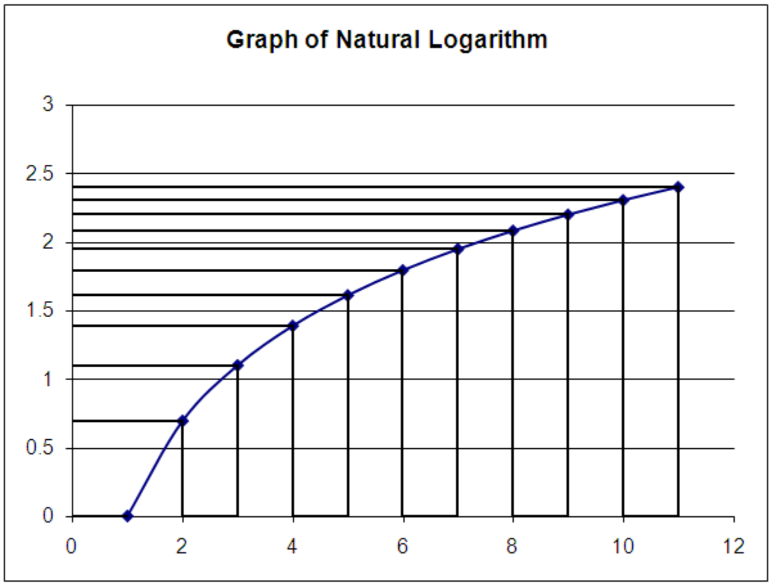

Figure 15.3: Graph of the natural logarithmic function, showing their basic properties.

15.1.4 How To Guide

Logarithmic and Log-Log plots

When you have data that spans many order of magnitude (like 1, 10, 100, 1000, 10000...) taking the logarithm of the data reduces it to a much more manageable set of numbers. For example, if we take the base-10 logarithm of each number in the preceding list, we get the numbers (0, 1, 2, 3, 4, 5...) which are must easier to use. This is the essence of many commonly used scales of measurement (the Richter scale for measuring earthquake energy and the unit of measuring sound, the decibel, are both logarithmic). This is also useful in dealing with models in which the variability in the residuals increases.

An alternate approach to actually computing the logarithm of each data point is to simply graph the data on a logarithmic scale. This is easy to do in Excel. For example, if you enter the pairs of (x, y) data points shown below and generate a standard XY (scatter) plot of the data, the graph is obviously curved, indicating a nonlinear relationship between the variables.

x	y
1	2
2	7
3	20
4	54
5	148
6	400
7	1100
8	2900
9	8100

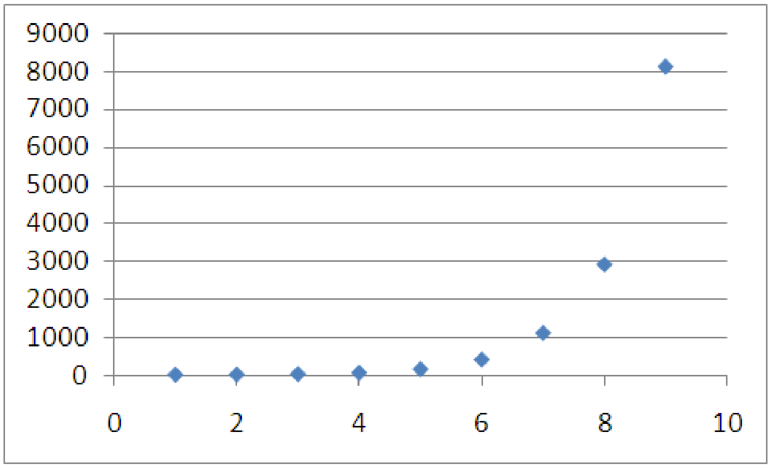

Figure 15.4: XY data showing a nonlinear relationship.

Click on the graph, then select the layout ribbon from the toolbar. Click on "axes" (not "axis titles") and select primary vertical axis. From the pull down list, choose "show axis with log scale". The graph will still display exactly the same data, but will appear to represent an almost linear relationship. This is shown on in figure 15.6. Notice that the vertical axis now looks very different. In the original graph, the evenly spaced gridlines represented an increase in the y variable of 1,000, regardless of whether you were at the top of the axis or the bottom. The spacing on the logarithmic scale, though, increases by a factor of 10 for each gridline (from 1 to 10, 10 to 100, 100 to 1000, etc.)

You can change the scale on the horizontal axis as well, letting you create log-linear, linear-log and log-log type graphs.

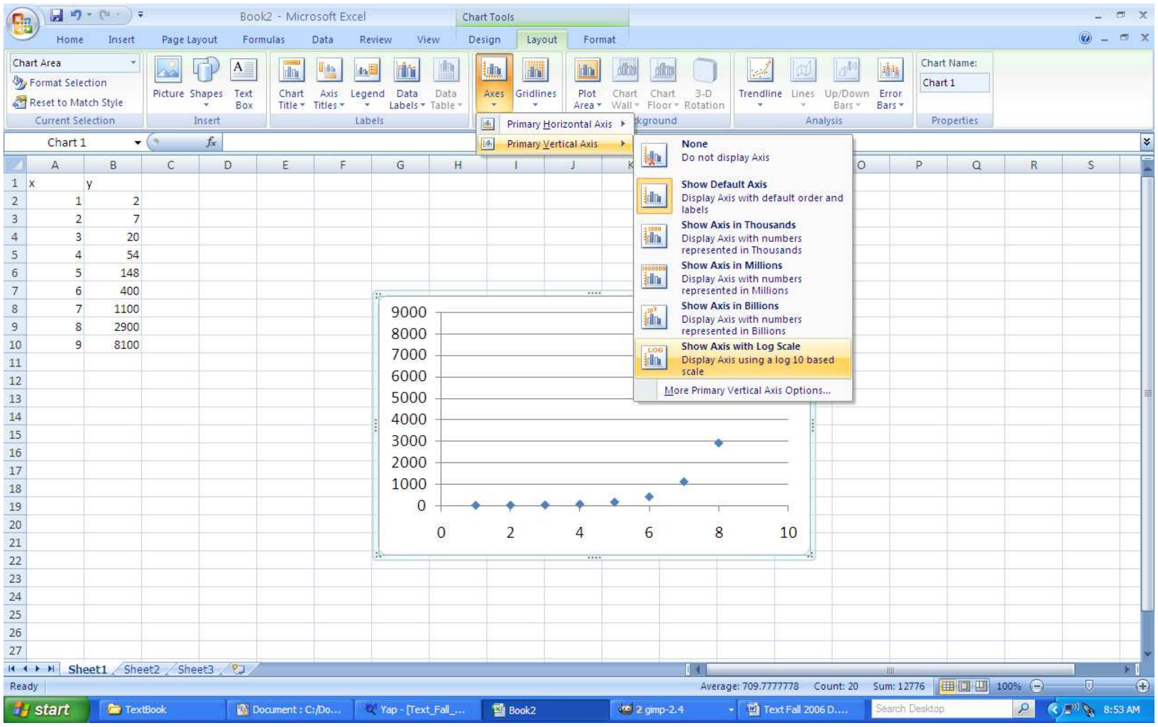

Figure 15.5: Converting the vertical axis of a graph to a logarithmic scale.

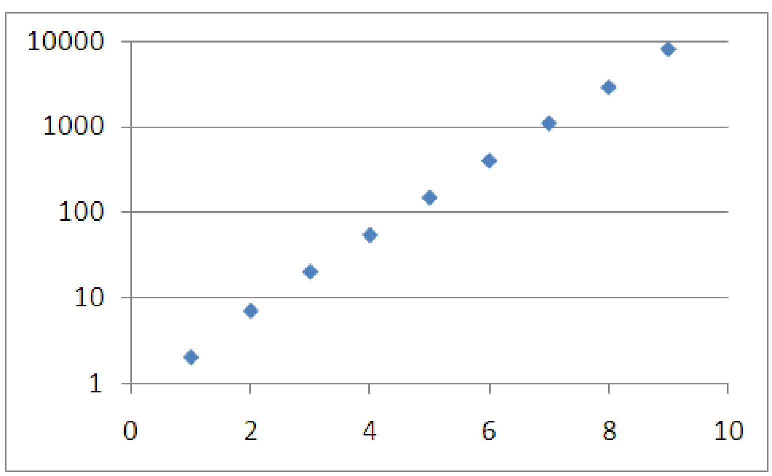

Figure 15.6: XY data showing a nonlinear relationship on a $\log(Y)$ scale.

15.2 Compound interest and derivatives of exponentials

Compound interest is one of the foundations of modern finance. The basic idea is that your investment will earn interest on the amount invested (the principal) as well as the interest itself. There are two primary versions of compound interest that we will explore in this section. The first is the easiest to make sense of, the case where there are a fixed number of times each year when the interest is computed and then added to the account. The other version is harder to understand intuitively because it involves interest being computed an infinite number of times. While it may seem that this would give you an infinite amount of money, since the interest rate for each period is infinitesimally small (it is the annual percentage rate divided by the number of compounding periods, so it is extremely small) the total amount reaches a fixed limit related to the number e.

Once we understand the basics of compound interest it can be applied to many other economic and financial concepts, such as present value and future value of an investment. The present value of an investment is the amount you would need to invest today in order to achieve a fixed level at the end of the investment period. This situation is most easily understood through the modern day phenomenon of the lottery. Most lotteries offer the winner two choices of payment: a lump sum now or small payments made over a longer period of time, say 20 years. If the winner "won" \$1 million, she would, for example, have to choose between monthly payments of \$50,000 each year for twenty years (a total of \$1 million) or a lump sum payment of \$548,811.64 right now. Ignoring all taxes, of course, which substantially change the problem under consideration, the reason the lump sum payment is so much less than the actual winnings is that you are getting it now. If you were to invest it at 3% for 20 years, you would have about \$1 million at the end, the same amount as the lottery winnings. Since the lottery company would have access to the money in the 20-year payment version, they would be earning interest on the \$1 million over that entire 20 years. But if they have to pay you all right now, they lose that interest. Thus, the present value of the \$1 million lottery winning is about \$550,000, assuming a 3% interest rate annually. We will further explore the idea of present value in the problems for this section.

15.2.1 Definitions and Formulas

Principal The amount of money initially invested or borrowed; it is the basis for computing the interest for the investment or loan.

Simple interest Simple interest is a way of computing the value of an investment based on giving interest one time only: at the very end of the investment period.

Compound interest Compound interest involves breaking the lifetime of the loan or investment into many periods. During one period, simple interest is used to compute the value of the loan or investment. During the next period, the interest for that is based not on the original principal, but on the current value of the loan including all interest from previous periods. Thus, with compound interest, you earn interest on your interest.

Continuously compounded interest This is a form of compound interest that uses, essentially, an infinite number of infinitesimally short investment periods for computing the interest. When this is done, we find that the exponential function with base e is a natural way to express the investment value.

15.2.2 Worked Examples

Example 15.4. Compound interest formulas
Suppose we were to invest an amount of principal, P, in an account that earns an interest rate r each year (this is the APR, or Annual Percentage Rate). This means that at the end of the first year, you will earn rP additional money. Thus, after one year, your account, A, has the value

$$A(1) = P + rP = P(1 + r).$$

If you were to leave the money in the account for a second year, you would earn interest not only on the principal, but also on the interest you earned the first year:

$$A(2) = P(1 + r) + P(1 + r)r = P(1 + r)(1 + r) = P(1 + r)^2.$$

What if you let the money earn interest for a third year? You would have a total of

$$A(3) = P(1 + r)^2 + P(1 + r)^2 r = P(1 + r)^2(1 + r) = P(1 + r)^3.$$

With a little work, we can show that, in general, after n years at an APR of r your principal P will earn a total of

$$A(n) = P(1 + r)^n.$$

Now, suppose that our interest is not computed annually, but is computer every month, based on the APR. This means that the actual monthly interest rate is $r/12$ and that in a single year we have 12 compounding periods. Similar logic to the previous case will tell us that after t years of compounding the interest monthly at this rate we will have

$$A(t) = P\left(1 + \frac{r}{12}\right)^{12t}$$

dollars in the account. Similarly, if we let the money be compounded n times each year, we will have an interest rate of r/n each period and a total of nt compounding periods after t years. This gives us an amount of

$$A(t) = P\left(1 + \frac{r}{n}\right)^{nt}.$$

This is, obviously, an exponential function, but with a base of $(1 + r/n)$ rather than the natural base of e. However, they are related. Consider what happens if we invest \$1 at 100% APR for one year under different compounding periods, as shown in the table below.

Schedule	Number of Periods	Total Amount
Annual	1	2
Monthly	12	2.61303529
Weekly	52	2.692596954
Daily	365	2.714567482
Hourly	8760	2.718126692
Each minute	525600	2.718279243
Each second	31536000	2.718281781
Every tenth of a second	315360000	2.71828187
Every hundredth of a second	3153600000	2.718281661

Notice that the amount of money does continue to grow, but not at the same rate. In fact, it seems that the amount of money we are earning is approaching a fixed amount. Mathematically, it is has been proven that this is the case and that the number this approaches is the number e:

> The number e is the amount of money earned in an account after investing \$1 for one year at 100% interest, compounded continuously.

Mathematicians write this fact using the limit notation:

$$\lim_{n\to\infty}\left(1+\frac{1}{n}\right)^n = e.$$

We can now use this fact to generate a formula for continuously compounded interest. First, we introduce a new variable m so that $n = r{\cdot}m$. Then we have an equivalent expression for the interest given by

$$\lim_{n\to\infty}\left(1+\frac{r}{n}\right)^n = \lim_{m\to\infty}\left(1+\frac{1}{m}\right)^{mrt} = \left[\lim_{m\to\infty}\left(1+\frac{1}{m}\right)^m\right]^{rt} = e^{rt}.$$

Thus, our formula for the amount in an account with n compounding periods changes to the following formula if we compound it continuously:

$$A(t) = Pe^{rt}.$$

Example 15.5. Derivatives of exponential functions
Now that we know about the derivatives of logarithmic functions, we can easily use the idea of a logarithmic derivative to determine the derivative of an exponential function. One of the most common exponential functions to occur in the business world relates to the future value of an investment. To get to this, though, we'll need to develop the idea of compound interest.

So, although it took us a little while to get there, and we skipped a few steps, we see that the exponential function is closely tied to the idea of compound interest. We can now ask the following. Suppose you have invested a fixed amount of money P at a fixed rate of interest r. How quickly (in time) is your money growing in value?

The question "how quickly" immediately reminds us of the idea of rates of change, so we know we are really talking about the derivative of the amount of money in the account. So, what is the derivative of the amount? We'll use our knowledge of logarithmic derivatives to help. We really want to know the derivative of $A(t)$, but we don't know the derivative of an exponential. However, the exponential function and the logarithmic function are inverses of each other, so the formula for the amount can be rewritten as

$$\ln(A(t)) = \ln\left(Pe^{rt}\right) = \ln(P) + \ln\left(e^{rt}\right) = \ln(P) + rt\ln(e) = \ln(P) + rt$$

where we have used the rules for manipulating logarithms and the fact that $\ln(e) = 1$. Now, we can take the derivative of each side of this equation, using the chain rule:

$$\frac{d}{dt}(\text{left hand side}) = \frac{d}{dt}\ln(A(t)) = \frac{1}{A(t)}\frac{dA}{dt}.$$

Now, the derivative of the right hand side is easy, since it's really a linear function (note that $ln(P)$ is a constant; it doesn't depend on the variable t with respect to which we are taking the derivative):

$$\frac{d}{dt}(\text{right hand side}) = \frac{d}{dt}\left(\ln(P) + rt\right) = \frac{d}{dt}\ln(P) + \frac{d}{dt}(rt) = 0 + r = r.$$

We can now put all this together, since we have done the same thing to both sides of the equation (namely, take the derivative with respect to t), so they are still equal to each other.

$$\frac{1}{A(t)}\frac{dA}{dt} = r \quad \Rightarrow \quad \frac{dA}{dt} = rA(t) = rPe^{rt}.$$

So, the true rate of increase of your account value is an amount of $r * \exp(rt)$ dollars per year. If you let it sit for $t = 10$ years at a rate of 2.5% your money will be increasing at a rate of $A'(10) = 0.025 \cdot P \cdot \exp(0.025 * 10) = 0.025 \cdot P \cdot \exp(0.25) = 0.032P$ dollars per year. If you had invested \$1000 initially, this would come to a growth rate of about \$32/year.

Example 15.6. Derivative of an exponential function

Find the relative rate of change of the function $g(r, t, P) = Pe^{rt}$ with respect to the variable r. The relative rate of change is just the rate of change divided by the function itself, so we have the relative rate of change as $(1/g)\cdot$ (derivative of g with respect to r).

$\frac{1}{g}\frac{\partial g}{\partial r} = \frac{1}{g}\frac{\partial}{\partial r}(Pe^{rt})$	Definition of relative rate of change, using partial derivative notation since there are several variables in the function
$= \frac{1}{g} \cdot P \cdot \frac{\partial}{\partial r}(rt)$	Derivative of a constant times a function
$= \frac{1}{g} \cdot P \cdot e^{rt} \cdot r$	Derivative of an exponential AND chain rule
$= \frac{1}{g} \cdot r \cdot g$	Derivative of a linear function
$= r$	Simplification

This means that the relative rate of change of the formula for continuously compounded interest is just equal to the interest rate itself. To understand what this means, think about the units of the rate of change with respect to r: units of dollars divided by units of interest. When we divide this by the amount (dollars) we get the relative rate of change, which is measured in 1/(interest rate). This is a relative amount, so it is like a percentage. Thus, each actual 1% increase in the interest rate (from 1% to 2% or from 5.25% to 6.25%) will increase the value of our account for a fixed amount of principal invested for a fixed period of time by r %.

Example 15.7. Application of Marginal Analysis to Business Decisions
The analysis team at Koduck has determined the following information about your current production level:

Marginal cost (MC) = $2.25/unit
Marginal Revenue (MR) = -$1.15/unit

What does this mean for Koduck?

For starters, we note that a negative value for marginal revenue means that if you increase production by 1 unit, your overall revenue (price * number sold) will actually drop. (This could be because you have already flooded the market; after all, how many pictures of water fowl can you sell in a given city?) The fact that the marginal cost is positive means that it will cost you more to make one more unit of product. Thus, it seems that increasing current production levels would not be wise: The total cost would rise and the revenue would drop, leading to lower profits. No one wants that. In fact, we should probably decrease production in order to increase profits! If we decrease production by 5 units, say, then we can expect the revenue to increase:

Change in Revenue = MR*change in production = (-$1.15/unit)*(-5 units) = $5.75.

At the same time, this would result in a decrease in cost:

Change in Cost = MC*change in production = ($2.25/unit)*(-5 units) = -$11.25.

This results in a total change in the profit of $17! It is a fact (which we will explore later) that the maximum possible profit (= revenue - cost) must happen when the marginal cost and the marginal revenue are equal. Since we can increase profits by lowering production, we must be producing more units than necessary to achieve the maximum profit.

15.2.3 Exploration 15B: Loan Amortization

In practice, the types of interest discussed in this chapter (simple, compound, and continuously compounded) are only parts of larger schemes for determining interest. One common application of simple interest is in loan amortization. The idea is that you take out a loan for a specified amount of principal, at a particular APR, for a set period of time. This time period is broken into smaller time periods (for example, a fifteen year loan for a house might be broken into monthly payment periods) and during each period you pay back some principal and some interest. However, while the total amount of each payment is generally held constant, the amount of that payment devoted to interest and principal repayment are not. In this exploration, you will construct a spreadsheet to explore the way a loan is repaid.

Suppose we take out a \$130,000 loan for a property. If the loan is at 6% interest (APR) and we pay it back monthly over a fifteen year period (180 payments) how much will we need to pay per month? Start by entering the basic information on the loan, as shown below in cells A1:B4. In cell B4, put your guess for the amount you would expect to pay. Try to be reasonable, keeping in mind that none of the interest schemes above will actually give you the amount, since the amount of interest to be paid at any one time is determined by the remaining principal on the loan. In cell E1, enter a formula to calculate the monthly interest rate (it is the APR divided by the Number of Periods in a year). Now, set up the loan amortization table headers as shown. Under "Period" enter the numbers 1, 2, 3, etc. up to 180 at 12 monthly periods per year, this will carry our loan through 15 years.

	A	B	C	D	E	F
1	Principal	\$130,000		Periodic Int	0.005	
2	APR	6.00%				
3	Num Periods	12				
4	Est Payments					
5						
6						
7						
8						
9						
10	Period	Interest	Principal	Cumulative Interest	Cumulative Principal	Remaining Principal

Figure 15.7: Setup for computing a loan amortization

Now, the interest for a particular period is easy to compute: it's just simple interest on the remaining principal balance. So, for the first period, all we need to do is multiply the periodic interest rate by the original loan amount. Once we have this, the amount of principal in the first payment is the total monthly payment minus the interest that period. The cumulative interest is just a place to track the running total on the interest we have paid and the cumulative principal tracks the total we have paid on the original loan amount; the remaining principal is the original loan minus the cumulative principal. Your formulas for the first period will probably be slightly different than the formulas for the other periods, but once you have the formulas entered, you can copy them down the table. Since the goal is to pay off the loan in 15 years (180 monthly payments) try changing the "Est Pay" amount until you find a monthly payment that leads to a balance of zero remaining principal in period 180 (cell F190).

Once you have played with this a little, you can use ”Goal Seek” to compute the actual monthly payment required to pay the loan off by a certain period of time. Try constructing a table listing different monthly payments based on changing one of the loan parameters (like the interest rate). Pay particular attention to the cumulative interest paid on the loan.

N.B.: There is a way to compute, just from the loan information, the monthly payment required. This formula, however, requires a lot of computational work, and we can get the same information by playing with it in Excel. Details of the formula will be discussed in the Symbolic Manipulation supplement for those interested. There are also automatic formulas in Excel for computing loan amortizations. If you are interested, look up the functions PMT, IPMT and PPMT.

15.2.4 How To Guide

The PMT Function

Excel can easily help you compute monthly payments on loans using the PMT function. This function requires at least three arguments, and has two optional arguments that can be used in different situations. The format of the function is shown below.

```
=PMT(Rate, NPer, PV, [FV], [Type])
```

Rate is the interest rate per period that the interest is compounded. If you are working with most loans, this will be the APR divided by 12, since the interest is compounded monthly. NPer represents the number of periods over which the loan is being computed. PV is the present value of the loan; in most cases, this is the principal, the amount that you have borrowed. The two optional arguments are FV and Type. FV is the future value of the loan; use this if you do not want the amount to be 0 at the end of the loan period. If you leave this argument out, Excel assumes that you want it to be 0. Type is either 0 or 1; 0 means that the payments are due at the end of each loan period; 1 means that payments are due at the beginning of each loan period (default is 0).

Thus, if you have a $10,000 loan, at 6% APR for 2 years, we can compute the monthly payments with the formula below (assuming payments are due at the end of each period). Notice that the number of periods is multiplied by 12 to make sure that we are computing everything in months (2 years * 12 months/year = 24 months for the number of payments).

```
=PMT(6%/12, 2*12, 10000)          (should be $443.21)
```

Notice that in two years you will pay a total of $443.21*2*12 = $10,636.95 to the loan, which is the original loan of $10,000 plus about $636 in interest. This function can also help you plan savings or retirement accounts using the FV argument. Suppose you want to know how much to save each month, starting with nothing, if you get 5% APR and want to have $250,000 in 10 years. The formula below computes the monthly payments you should make to the account (payments at beginning of each period).

```
=PMT(5%/12, 10*12, 0, 250000, 1)          (should be $1,603.29)
```

Notice that if we multiply the PMT amount by the number of periods (10*12 = 120) we get $192,394.90, indicating that over the ten years, we will pay about $192 thousand, but the interest will give us a balance of $250 thousand. Thus, we earn about $58 thousand over the savings period.

FINAL NOTE: PMT returns the answer in accounting format. This means that it is almost always a negative number displayed in red and in parentheses. This makes it easier to use in computations for loan amortization schedules, since you normally want to subtract the loan payment each period.

15.3 Homework

15.3.1 Mechanics and Techniques Problems

15.1. For each of the following functions, compute the first derivative of the function with respect to the independent variable.

1. $f(x) = 3\ln(x) + 5$
2. $h(t) = -2e^{3t}$
3. $g(s) = 5s + 3s\ln(s^2 - 4)$
4. $p(y) = 5ye^{-y^2}$
5. The logistic function $f(x) = \frac{A}{1+e^{-Bx}}$, where A and B are constants.

15.2. Find the local maxima and minima of the function given in 1d) above. Use this information to help sketch a picture of what the function looks like when plotted as $p(y)$ versus y.

15.3. The present value of an investment is the amount of money you would need to invest at a particular interest rate r for a specified period of time t in order for the investment to rise to a total value of V.

1. Assuming that there are n compounding periods per year, determine a formula for the present value of an investment.
2. Assuming that the interest is compounded continuously, determine a formula for the present value of an investment.
3. Using your formulas in a) and b) fill in the table showing the present value of a 10-year investment that has a value of \$1 million. Your table should compute this for the following range of interest rates: 1%, 2.5%, 5% and should show the results for annual compounding, monthly compounding, daily compounding and continuous compounding.

Interest Rate	Annual compounding	Monthly compounding	Daily compounding	Continuous compounding
1%				
2.5%				
5%				

15.3.2 Application and Reasoning Problems

15.4. Suppose that you are a manufacturer of widgets. At your current level of production, you have determined that each one unit increase in the production level will decrease the revenue by \$0.28. Each unit of increase in the production level leads to a drop in costs of \$0.34. Each day, your plant is improving efficiency, so each day the production level is expected to increase by 32 units. At what rate is the profit changing? Would you continue to increase the production? Why?

15.5. Prove that the exponential function of the form $y = Ae^{Bx}$ is an always increasing function of x (assuming that B is positive and A is positive). In other words, show that this function never reaches a maximum and then starts to decrease. Such functions are referred to as monotonically increasing.

15.6. Prove that the logarithmic function $y = A + B\ln(x)$ is a monotonically increasing function.

15.3.3 Memo Problem

To:	Analysis Staff
From:	Cassandra Nostradamus, CEO
Date:	May 30, 2008
Re:	Loan options

Oracular Consulting is planning to purchase $1,000,000 in computer equipment and software to upgrade the main server and our web presence. Since we do not want to reduce our liquid assets by this amount, we are considering several different loan possibilities. The terms of these loans are described below.

	Loan A	Loan B	Loan C	Loan D
APR	6%	5%	3%	2%
Number of Years	2	3	5	10
Payments per year	12	12	4	4

Analyze the four loans and provide a well-reasoned recommendation as to which loan (or loans) would be the best choice. It would certainly be nice to choose a loan that we can pay off as quickly as possible, but that may require very high monthly payments. If we are willing to pay large monthly payments, then we can take a short term for the loan, but if we need to lower the payments, we need to make a decision on some other characteristics. The three obvious ones are to compare the length of the loan, the total interest paid over the lifetime of the loan, or the monthly equivalent payments for the loan (the amount we pay each period, pro-rated to a monthly budget amount).

Attachments: None - create your own to display the results