Math 1106 – Elementary Applied Calculus
Textbook: Bittinger's Calculus and Its Applications
Lecture Notes and Suggested Exercises for Chapter 3, "Exponential and Logarithmic Functions"

Added after Day 19 class meeting, Wednesday, October 21, 2009

After finishing my lecture on Section 2.6, I continued into this chapter, talking about two types of functions you've seen before, along with the rules for deriving their slope functions (their derivatives). First, we renewed our past acquaintance with the exponential functions. These functions occur naturally in nature, particularly because they describe the way that a population of organisms reproduce and increase in number. (Can you say, "Cobb County's population is increasing exponentially!"?)

The equations for exponential functions have this basic look: $f(x) = a^x$. In other words, the input variable is used as an exponent for some fixed positive number *a* (where *a* is not equal to 1). When $a > 1$ it will be a **growth** function; if $a < 1$ then it is a decreasing exponential function, that is, a **decay** function.

During the development of the equation for the derivative of such an exponential function, we happen upon a special fixed positive number that occurs so prevalently throughout mathematics (and engineering and physics and chemistry and every other branch of science) that it has its own name: *e*. This number has a numeric value of about 2.71828, and the exponential function based on *e* has a very simple derivative rule: if $f(x) = e^x$ then $f'(x) = e^x$. Boy, that's an easy rule to remember!

Don't forget about the chain rule! If an exponential function based on *e* has an exponent that is a function of *x*, then the chain rule demands that there be a final multiplier of the derivative of that function. So, this means that if $f(x) = e^{g(x)}$ then $f'(x) = e^{g(x)} \bullet g'(x)$. That's because we have a "chain" of two functions – we could say that $f(u) = e^u$ with $u = g(x)$, and so the derivative of their composition $\frac{d}{dx} f \circ g(x) = f'(g(x)) \cdot g'(x) = e^{g(x)} \cdot g'(x)$.

For example, consider the function $f(x) = e^{4x}$; its derivative is $f(x) = 4e^{4x}$ (that factor of 4 comes from being the derivative of 4*x*).

Added after Day 20 class meeting, Monday, October 26, 2009

The topic of “inverse functions” may seem a little too theoretical for some, but when a function has a matching “inverse function,” that’s a convenient and very useful thing (not every function has a matching inverse function). And to make any headway in talking about the topic of logarithms, the best way to start is to have a remembrance of what inverse functions are.

When we’re given the defining equation for a function (for instance, $f(x) = 3x - 7$), one of our typical chores is to answer questions such as “What is the numeric value of $f(5.1)$?” (Translation: What is the numeric value of *y* when *x* = 5.1?) We ***evaluate*** the function by sticking in the prescribed value of 5.1 for *x* in the defining equation and proudly (and correctly) answer with “8.3!” Evaluating a function for a specific value of *x* is pretty straightforward, even though sometimes it would be nice to have a TI-83 handy to do the evaluation if the function is a tiny bit “ugly” with unfriendly coefficients or ridiculously high powers of *x* or if (shudder!) there are radicals involved. But in general, we “see” the way to the accomplishment of the task as simply “plugging in” the stated value for *x* and doing some arithmetic, right?

But turn the assignment around: given a specific value for the output, we are asked to ***solve*** something. “Solve $f(x) = 32.8$” can be intimidating. And no wonder! It is! In effect (using the aforementioned function) we’re being asked to find the value of *x* that makes this true: $3x - 7 = 32.8$. In truth, since this function is a simple linear one, it turns out to not be too onerous a task: just add 7 to both sides of the equation and then divide both sides by 3, and you end up with $x = {}^{39.8}\!/_{3} = 13\,{}^{4}\!/_{15}$ (I know you just *love* fractions!).

But think of what you did to “solve” the equation, and extend that idea to the general case of $f(x) = y$. In other words, solve $3x - 7 = y$ for the numeric value of for the numeric value of *x*. Now, adding 7 to both sides followed by division of both sides by 3 gives you this: $x = \frac{1}{3}y + \frac{7}{3}$. You’ve now got a matching function – the “inverse function for *f* -- that will allow you to recover the value of *x* that’s needed to plug into the function *f* that will make it evaluate out to the prescribed result of *y*. Simply swap the names of the variables, and you’ve got a new function of *x* that’s called “the inverse function for *f* “ that we name f^{-1} and write like this: $f^{-1}(x) = \frac{1}{3}x + \frac{7}{3}$. Then, when someone tells us to “solve $f(x) = 32.8$” we know that the task will be much simpler if we think of 32.8 as an x-value to be inserted into the inverse function, so that we simply evaluate $f^{-1}(32.8) = \frac{1}{3}(32.8) + \frac{7}{3} = 13\,{}^{4}\!/_{15}$. And a quick check of our answer using this answer, $f(13\,{}^{4}\!/_{15}) = 3(13\,{}^{4}\!/_{15}) - 7 = 39\,{}^{12}\!/_{15} - 7 = 32\,{}^{4}\!/_{5} = 32.8$. In effect, we’ve come “full circle” and are back where we started from.

But that's exactly the nature of two functions that are each other's "inverse function": they 'undo' each other's work. Speaking mathematically, this translates to what happens when you "compose" such functions (or "chain" them together): the composed result does nothing to the input, which means the output of the composed function is identical to the input. In terms of function notation, $f \circ f^{-1}(x) = f(f^{-1}(x)) = x$ and $f^{-1} \circ f(x) = f^{-1}(f(x)) = x$.

The most enlightening picture I can draw of that is to graph both of these functions, along with the y = x function. Notice the symmetry about the y = x line, of how the two other functions are mirror images of each other, with the y = x line being the "mirror."

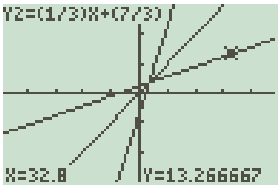

For the curious, the window settings above are

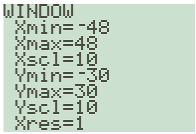

```
WINDOW
 Xmin=-48
 Xmax=48
 Xscl=10
 Ymin=-30
 Ymax=30
 Yscl=10
 Xres=1
```

Now, functions that have an inverse function must possess a special attribute: not only must they satisfy the "vertical line test" in order to be functions, they must also satisfy the "horizontal line test" so that their mirror image thru the y = x line will itself satisfy the vertical line test. One type of function that's guaranteed to have this ability to satisfy both the vertical and horizontal line tests is one that is always increasing; another type is one that is always decreasing. In other words, it never reverses its vertical direction. (The special "math-ese" word for that is "monotonic.") For a continuous function, it's both necessary as well as sufficient.

Since any exponential function is always either increasing or decreasing, it must have a matching inverse function. In particular, the exponential function $f(x) = e^x$ has as its inverse the natural logarithm function: $f^{-1}(x) = \ln(x)$. And if you thought the exponential function has an easy derivative rule, then take a look at the derivative of the natural logarithm function: $\frac{d}{dx}\ln(x) = \frac{1}{x}$. The proof of this is quite simple, utilizing the fact that the two functions are inverses of each other:

$$x \cdot \tfrac{d}{dx}\ln(x) = e^{\ln(x)} \cdot \tfrac{d}{dx}\ln(x) = \tfrac{d}{dx}e^{\ln(x)} = \tfrac{d}{dx}x = 1.$$

That may be a bit hard to swallow (or follow), but it follows from the fact that $e^{\ln(x)} = x$, which is a direct result of the inverse function relationship of $f(x) = e^x$ and $f^{-1}(x) = \ln(x)$.

Once again, don’t forget about the need to invoke the chain rule if the input is a function of *x* instead of just *x*. For example, the derivative of $g(x) = \ln(x^2 + 7x)$ is $g'(x) = \dfrac{1}{x^2 + 7x} \bullet (2x + 7) = \dfrac{2x + 7}{x^2 + 7x}$.

Now, don’t forget -- since the two functions $f(x) = e^x$ and $f^{-1}(x) = \ln(x)$ are inverses of each other, then the effect of “chaining” them together is to get the “do nothing” function $y = x$, which simply means that $e^{\ln(x)} = x$ and $\ln(e^x) = x$. That first result accentuates the statement that “a logarithm is just an exponent.”

It is also the “trick” whereby we can, in effect, ignore all the other exponential and logarithmic functions based on other positive numbers, and just concentrate on the exponential and logarithmic functions based on *e*. Imagine that *a* is a positive number (other than 1). Then, because $e^{\ln(x)} = x$, we can say that $e^{\ln(a^x)} = a^x$. But the rules for handling logarithms (see property P3 in the blue box on page 305) means that the left-hand side of that equation can be rewritten: $e^{\ln(a^x)} = e^{\ln(a)*x}$. That means we can dispense with every other exponential function, because no matter which number we use for *a*, $a^x = e^{\ln(a)*x}$.

Thus, every exponential function can be recast as one that uses *e* as its base and a particular multiple of *x* in the exponent. In fact, e^{kx} is a growth function when *k* is a positive multiplier, but it’s a decay function when the positive number *k* is preceded by a minus-sign. When speaking generically about that multiplier, by using the letter *k* to represent it, our text book adopts the convention that it will always insert an explicit negative-sign if a decay function is being talked about.

Added after Day 21 class meeting, Wednesday, October 28, 2009

Playing catchup from the flood days, I covered Sections 3.3 and 3.4 by soliciting questions about the homework for the first 4 sections of Chapter 3, and I tried to show as many of those homework questions as I could during the class and either solve them, or give broad hints as to the procedure for solving them.

This section of the book harps on the nature of the $y = e^{kx}$ function when k is a positive number. This function is an increasing function, and its slope at a particular point is proportional to the output value at that point.

That property – that an exponential function is one whose derivative is proportional to the original function – is a unique property of exponential functions. The number represented by k is the proportionality constant. It tells how swiftly the function is growing (i.e., how fast its y-value is changing) versus the change in the x-value.

The title of this section of the book incorporates this idea of proportional growth by using the notation of the German mathematician Leibniz. Remember, both Newton and Leibniz pioneered the development of calculus back in the 17th century. Back in those days, the y-value of a function was simply given the symbolic name of y.

(The $f(x)$ name came along in later centuries – a handy notation, in my estimation, but unfortunately a stumbling block to students struggling to absorb the concepts of algebra. I've found that it's difficult to teach them that $f(x)$ is just a different name for y. They keep trying to imagine that the f is a number that is multiplying x. <Sigh>.)

Well, to describe the output of the derivative of the particular function, Newton called it $\dot{y}$ (notice the tiny dot floating over the y – nowadays, we write it y' instead). But Leibniz had decided he would use the symbol $\frac{dy}{dx}$ to represent the derivative.

So, the first objective for this section of the textbook summarizes the entire notion (that an exponential function is one that is proportional to its derivative) by writing (cryptically!) this Leibniz-like equation: $\frac{dP}{dt} = kP$. I call this cryptic, because instead of using the tried and true letter y, the author has used the letter P instead. Also, instead of using the letter x to represent the input, in this case he's used the letter t. Confused yet? I hope not!

When k is a fixed (constant) positive number, we're talking about an exponential **growth** function, something that is near and dear to anyone who appreciates money. In fact, the formula for calculating the future value of a savings account that accrues continuous interest is based on this function: $A = P_0 e^{kt}$, where P_0 represents the deposited amount, t is the number of years the deposit is kept, and k represents the annual interest rate expressed as a percentage. The quantity k is also called the **rate of exponential growth**.

In general, we deal with the prototypical exponential growth function $P(t) = P_0 e^{kt}$. P_0 simply represents the initial amount (money on deposit, or a country's population, or the number of germs in a Petri dish) for some growth situation. The k stands for the growth percentage per time period (but use the decimal equivalent of a percentage – e.g., 2.5% means k is equal to the number 0.025), and the t stands for the amount of time in the standard time period. The exercises in this section are principally handled using this formula. When talking about financial matters, the t will almost always be measured in years. When using this formula for growth of populations, you'll need to pay attention to whether the time is measured in weeks, days, hours, minutes, or whatever.

In this section there is a cute handy-dandy rule for estimating how long an initial amount of money doubles in value at a given percentage rate. It's called the "Rule of 70" and it's worth remembering for future use. Divide 70 by the percentage rate (expressed as a percentage, not the decimal version) to get an estimate of how many years it would take the money deposited to double in value. Or, divide 70 by the number of years you want the amount to double in, and that will give an estimate of the percentage rate you should invest your money at. Why does this work? Well, remembering the rules for handling fractions, if you instead used the decimal version of the percentage, you'd have to be dividing it into 0.70 instead (yuck! I'd rather deal with whole numbers, thank you very much), and 0.70 is so close to the actual value of $\ln(2) \cong 0.693$. Multiply numerator and denominator by 100 does not change the value of the quotient, so multiplying ln(2) by 100 is very close to 70.

Although the title of the chapter contains the word "uninhibited", another interesting function is set forth in this section of the book: the **logistic function**. A logistic function is one that starts off exhibiting exponential growth, but at some point it "changes its mind." By this, I mean that an **inflection point** occurs (the concavity changes), and the curve continues to grow but has an upper limit below which it remains. The book very nicely shows an example of a real world situation (it's Example 8 on page 344) where a logistic function models the spread of a staph infection. Think about it -- In the early weeks after the disease starts spreading, the total number of victims grows exponentially. But eventually, the growth in total infectees has to stop growing so ferociously, and tops off below some absolute upper limit which implies that the infection has infected the maximum number of people, and there aren't any new victims to be infected.

In Section 3.4, this is just the flip-side of the previous section: it's exponential functions where the multiplier of the input variable used in the exponent is a negative multiplier. It represents the opposite of growth: decay. There are many examples of functions that exhibit such negative growth (or decay). Hot things cool off ("decay") at a rate that is proportional to the difference between their temperature and the ambient temperature. Radioactive items lose their radioactivity, which is modeled by a decay function (based on their so-called half-life). And we can derive a lot of information from "decaying" a future attained financial goal that was achieved by investing at a particular percentage rate (growth rate) back to its initial (present) value.

The decay function, in its general form, is written as $P(t) = P_0 e^{-kt}$ with an explicit minus sign (which means we will always think of the *k* as being a positive number, the "percentage rate" expressed as a decimal).

The applications given in this section about exponential decay have a very scientific flavor to them, and I know that my students have more of a business orientation. Therefore, I will generally try to concentrate on the few exercises in Section 3.4 that treat business situations. But, because the *CSI* television programs are so popular, I thought it would be intriguing to direct the attention of my students to the couple of exercises that deal with the forensic technique for determining how long it's been since an unfortunate victim met his or her demise: check out Example 8 on page 358. Unfortunately, the author's write-up of this is heavily laden with arithmetical details in which the reader can easily get lost. The most important point to recall is simply that this is another application of the idea of exponential decay, especially the result that the rate of change (the derivative) of an exponential function is always proportional to the function itself (see the discussion about "Newton's Law of Cooling" that begins at the top of page 357).

Thus, the formula for calculating the temperature of the corpse is as given at the bottom of page 358: $T = ae^{-kt} + C$. That is a bit cryptic, I know. First, the capital letter *C* represents a constant number that is simply the ambient temperature (that is, the room temperature) around the corpse, and the assumption is that the temperature of the room has not and does not change from the time of the murder until the forensic specialist finishes the work. The second mystery to clear up is what the letter *a* stands for: it is merely the difference between the standard body temperature ($98.6°F$) and the constant room temperature. Thirdly, *T(t)* stands for the output of the temperature function, where *t* represents the amount of time (in hours, usually) that have elapsed since the person was killed.

The forensic investigator arrives at the crime scene and records the room temperature, then takes and records the temperature of the corpse (don't even *think* to ask me what part of the body is used for this!). Then the investigator waits one hour and takes and records the temperature of the corpse a second time (presumably, the temperature has now decayed to a lower number than before as the body cools down to room temperature).

Well, as you can tell by trying to work through Example 6 on pages 338-359, this gets a bit gruesome with the arithmetic, because you have to first determine the value of *a* (which isn't too hard; I already told you it's the difference between the normal body temperature of the dearly-departed, $98.6°F$, and the room temperature); then you have to work out the value for *k*; and then finally you arrive at the value for *t* that represents the number of hours the body has been dead -- which is the result the investigator really wants: the number of hours that elapsed between the murder and the time the first temperature reading of the body was taken. Only, to confuse the poor student struggling to understand all this, in the textbook all of a sudden the little *t* becomes a capital *N*.

There must be a better way! And there is. You don't think the CSI investigator goes through all that garbage arithmetic, do you? Nah! She probably has a TI-83 with a little program in it, and all she has to do is input the three temperature readings and, BOOM! -- the calculator gives back the value for *N*. I've worked out the formula, and here it is:

$$N = \frac{\ln(B-C)-\ln(T_0-C)}{\ln(T_0-C)-\ln(T_1-C)} = \frac{\ln\left(\frac{B-C}{T_0-C}\right)}{\ln\left(\frac{T_0-C}{T_1-C}\right)} ,$$

where *B* stands for the normal body temperature ($98.6°F$), *C* stands for the room temperature, T_0 stands for the first temperature reading of the corpse taken by the investigator, and T_1 stands for the second reading taken exactly one hour later. Try this with the data given in Example 6 to assure yourself it works. Then try it on Exercises 32 and 33 on pages 362-363.

What I find really, really neat is that all the temperature readings can be done using a Celsius thermometer instead and the answer for *N* will be (drum roll!) the same! (Of course, in this case you have to use the normal body temperature in Celsius, which is $37°C$).

So, I've left a lot for you to ferret out for yourselves by reading the book and working through the MathXL homework. Get to work!

And as if that’s not enough, today I covered Section 3.6 about Elasticity. (Actually, I didn’t get to talk about it in the 5 o’clock class ... I’ll cover that next Monday, those of you in the 5 o’clock class!)

“There is scarcely anything in the world that some man cannot make a little worse, and sell a little more cheaply. The person who buys on price alone is this man's lawful prey.” -- John Ruskin (1819-1900)

I first read this quote in college – posted on the wall in the local Baskin-Robbins store (my favorite flavor is mint chocolate chip, what’s yours?) Despite (or maybe because of) John Ruskin’s observation about price, it is the most widely used method whereby a company attempts to influence consumers to buy its products. And, of course, the reason for wanting more people to buy your products is because increased volume will usually mean increased revenue, even if you have to drop the price to get more people to buy your product.

Mathematical models for demand (which are used to predict how many units of a product will be bought – demanded – depending on what the unit price is) usually have one important thing in common with exponential decay functions: their output decreases in value as the input variable increases in value. As a result, there is no place where the derivative of a demand function is equal to zero; so how can we use the ideas of differential calculus to our advantage in trying to find that “perfect” price?

The secret is that a demand function is only part of the ingredients for the function in which we are ***really*** interested: the revenue function. If our demand function is represented by $D(x)$ (which means that the input variable p stands for the price being charged, and so $D(x)$ represents how many units of the product will be sold — or “demanded” — at that price), then what we ***really*** want to know about is the function $R(x) = x * D(x)$. I.e., how much ***revenue*** do we get at a certain price-point? And, of course, the revenue is always equal to the demand ($D(x)$) times the price (x) – for every increase of 1 in the demand, we put one more amount of x into the cash register, right? This function will probably have a place where its derivative is equal to zero, and so calculus can get us the answer of what’s the best price to charge to maximize our revenue. (This analysis assumes that the costs are all sunk costs, which might mean that changing the price of the product will not affect the total cost we have already incurred – and this point is important, since in the final analysis we really are interested not in finding the best revenue, but the best **profit.** This is the goal of all good capitalists!)

All this is meant to lead up to the why-and-wherefore of Section 3.6: a study of the topic of *elasticity of demand*. That is, which way and how far can we stretch the price (either up or down) to beneficially affect the total revenue derived from the sale of the product?

The function $E(x)=\frac{-xD'(x)}{D(x)}$ is called the "elasticity of demand" function. This function comes about from considering the revenue function given above: $R(x)=x*D(x)$.

If you want to set the derivative of the revenue function equal to zero (to solve for the "best" price), that means we're using the product rule on the right-hand side and setting that equal to zero: $0=xD'(x)+1*D(x)$. And of course, this will be equal to zero when $D(x)=-xD'(x)$, which in turn means when $\frac{-xD'(x)}{D(x)}=1$. That means, when $E(x)=1$ then revenue is at a maximum.

So, when the *elasticity of demand* is equal to 1, the revenue will be a maximum. When $E(x)>1$, we've stretched the point ("price is elastic") and we've set the price at a point where revenue is actually on the way down (decreasing). On the other hand, when $E(x)<1$ this means that "price is inelastic" and we have flexibility to keep moving the price towards the "sweet spot" where $E(x)=1$ and revenue will keep increasing.

Clear as mud? Look at the table on page 375. The things with an elasticity value more than 1 are the things that, when price is increased, will result in less revenue being obtained by the seller; that is, they are what we call "over priced"; in fact, they are ripe for a decrease in price, because that would result in more revenue for the seller. The things with elasticity less than 1 are the things that are "under priced", so (unfortunately) we can expect that their price will increase and even more revenue will be realized by the seller. Sadly, it's the good stuff that has an inelastic demand: food, tobacco, beverages, clothing, and gasoline. Just recently, we've seen that last item undergo a significant increase in price, in fact. But has there been an appreciable decrease in demand for gasoline? Well, that's because there is another economic law at work: the law of supply, which is the other side of the economic coin. It plays a tug of war with the law of demand. But that's a topic for another chapter and another day.

And so, with the end of Chapter 3, we have completed the first part of the calculus course: the study of *differential calculus*. And not too soon: because we've only a little less than 1/3 of the semester left, and the best is yet to come: *integral calculus*! YES! On to Chapter 4!