Chapter 11

Graphical Approaches to Nonlinear Data[1]

The basic idea of this chapter is that not all relationships are linear. In fact, many of the most commonly occurring relationships come from other families of functions such as exponentials or polynomials. In this chapter, we'll explore the shapes of these different functions and learn how to control their shapes through the parameters of the model. Then, we will put our knowledge of nonlinear models to work in chapter 12 (page 347) to build and interpret nonlinear regression models.

- *As a result of this chapter, students will learn*
  - √ Know what the parameters, constants and coefficients in a model are
  - √ Know the basic shapes for each of the basic non-proportional models of interest (logarithmic/log, exponential, square, square root and reciprocal)
  - √ Logs and exponentials are inverse functions to each other
- *As a result of this chapter, students will be able to*
  - √ Select and justify a choice of non-proportional model from among several possible candidates
  - √ Choose an appropriate non-proportional model based on a scatterplot
  - √ Determine something about the parameters of a model from looking at a scatterplot
  - √ Shift the graph of a model around in order to make it better fit the data
  - √ Stretch the graph of a model in order to make it better fit the data

11.1 What if the Data is Not Proportional

Our first assumption when modeling data using regression is that the data is based on an underlying linear relationship. Such relationships are said to be proportional: if the x data increases by a certain amount, the y data increases by a fixed constant times that same amount. The fixed constant relating the x-variable changes to the y-variable changes is the called the slope of the linear model.

For many sets of data, however, the assumption of linearity is quite false. For example, the amount of electricity used in a house is related to the size of the house; larger houses are more expensive to heat or cool, so they tend to use more electricity. However, this relationship does not mean that doubling the size of the house always doubles the electricity costs. Much of the electricity use comes from lights, computers, televisions, and radios. No matter how much bigger the house, a family of four can only use so many of these devices at one time. So while the cost may increase, we might expect a more dramatic increase in electricity use when comparing a small house to a medium house, but a much less dramatic increase when comparing a medium-sized house to a large house. This implies that the slope of the model relating the electricity costs (y) to the size of the house (x) would be different for large houses than for small houses. In a linear function, this slope must be the same, regardless of the x-value being considered.

11.1.1 Definitions and Formulas

Non-proportionality Any model relating two variables (say x and y) in such a way that changes in one variable are not in a constant ratio to the changes in the second variable is said to be non-proportional. Another way of describing this is by saying that there is no constant k for which the following relation is true:

$$y_2 - y_1 = k(x_2 - x_1)$$

In the mathematical world and in the real world, most models are non-proportional.

Level-dependent Any model that is level dependent is also said to be non-proportional. The term level-dependent emphasizes that with such models, the amount that the y variable increases for a given increase in x is different if the starting point (x value or location along the horizontal axis) is moved. In other words, you can look at different x and y values and compute their differences. When we compare them, if we find that $y_2 - y_1 = k_{12}(x_2 - x_1)$ and $y_4 - y_3 = k_{34}(x_4 - x_3)$, but the k values are different, then the model is level-dependent and represents a non-proportional relationship.

Concavity Concavity is a property of non-proportional models. It refers to the amount that the graph of the model bends. If the graph bends upward, that part of the graph is said to be "concave up". If the graph bends downward in a certain area, then the graph is "concave down" in that area. Remember: concave up looks like a cup; concave down looks like a frown.

Basic function One of the six functions listed below as prototypes for fitting nonlinear data:

- linear,
- logarithmic,
- exponential,
- square,
- square root, or
- reciprocal.

In general, a function is a mathematical object that takes an input, usually in the form of a number or a set of numbers, and gives an output number. (There are other types of functions possible, but we will concentrate on functions that satisfy this definition.) For a relationship between two variables, say x and y, to be a function, it must satisfy the following statement:

Every x-value must be associated with one and only one y-value.

This means that if you draw a graph of the function, and draw a vertical line through any point on the graph, that line will only touch the graph once. This is sometimes referred to as the vertical line test. Generally, if the variable y is a function of the variable x, we write $y = f(x)$ to indicate this. If the variable y is a function of several variables (say x_1, x_2, x_3) then we write $y = f(x_1, x_2, x_3)$.

Linear function Graphs of linear functions (see figure 11.1) are straight lines. The prototypical, or base, form of a linear function that related y to x is given by $y = x$. You are more used (by this point) to seeing this in a more general form, involving two parameters, the slope and y-intercept: $y = A + Bx$. The graph of a linear function is shown below. Notice that linear functions are straight; they have no concavity at all.

Logarithmic function A logarithm (see figure 11.2) is a mathematical function very useful in scaling data that spans a large range of values, like from 1 to 1,000,000 (we will see this aspect of logarithms in a later chapter). In general, there are lots of different logarithmic functions. We will be using the natural logarithm of x as a function; this is written as $y = \ln(x)$. (Notice: **n**atural **l**ogarithm = nl $\rightarrow$ ln.) The graph of the basic natural logarithm is shown below. The basic logarithm is increasing and concave down everywhere.

The natural logarithmic function has several important properties to note. The natural log of 0 is undefined; in other words ln(0) does not exist. If $0 < x < 1$ then $\ln(x) < 0$, and $\ln(1) = 0$. This means that the point $(1, 0)$ is common to all basic log functions. This is actually a restatement of the fact that any base raised to the zero power is equal to 1.

Exponential function Exponential functions (see figure 11.3) are related to logarithmic functions. These can be written in two ways. The first form is as a base number raised

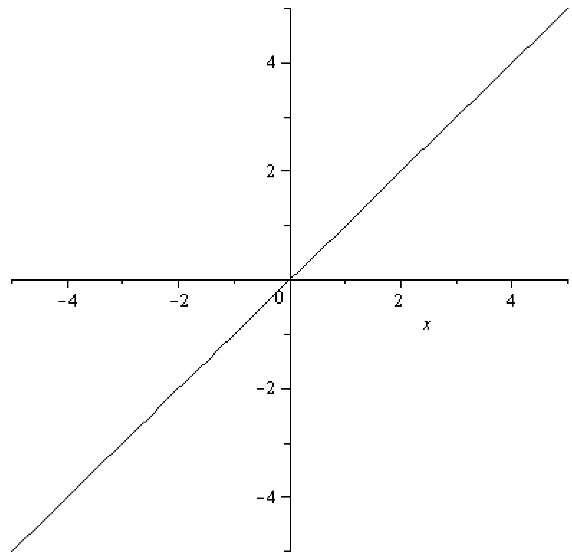

Figure 11.1: The basic linear function $y = x$.

to a variable power ($y = a^x$). The most common base to use is the number e, which is approximately 2.71828... In reality, e is an irrational number, like π. It shows up naturally in many situations, as we will see in example 4 (page 455) from chapter 15 (page 443) when examining interest rates. For now, though, the standard exponential function we will use is $y = e^x$. The second form is similar to this, but easier to type: $y = \exp(x)$. This can be read as "y is the exponential function of x" or "y equals e raised to the x power." Its graph is shown below. The basic exponential function is increasing and concave up everywhere.

In addition, since any positive number, like e, raised to a negative power is a number between 0 and 1, we know that if $-\infty < x < 0$ then $0 < e^x < 1$. Since any number raised to the zero power is 1, we also know that $e^0 = 1$ so the point $(0, 1)$ is on the graph of all basic exponential functions.

Square function You are probably familiar with the squaring function: it takes every number put into it and spits out that number raised to the second power. Thus, if we stick in the number x, we get out x^2. Thus, the basic squaring function is $y = x^2$. The graph of this function has a special name that you may have heard before: a parabola. It looks like the letter "U", centered at (0, 0). The basic squaring function is concave up everywhere. as shown in figure 11.4.

Square root function The square root function does the opposite of what the squaring function does. This function takes in a number and spits out its square root. The square root of a number is that number which, when squared, produces the number. For example, 2 is the square root of 4, since 2 × 2 is 4. The square root function is usually written as $y = \sqrt{x}$. Another way to write the function reminds us of its relationship with the squaring function: $y = x^{1/2} = x^{0.5}$. (Read this as: y is x raised

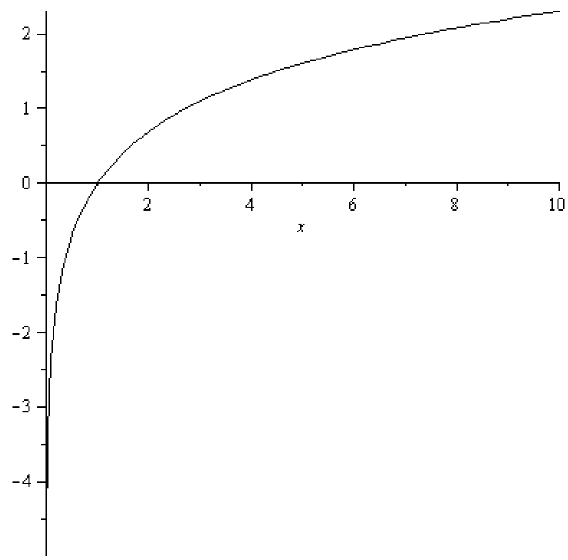

Figure 11.2: The basic logarithmic function $y = \ln(x)$.

to the one-half power or x to the 0.5 power.) The basic square root function is concave down everywhere. In figuer 11.5, the square root function is not graphed for values of x less than 0, since the square root of a negative number is an imaginary quantity.

Reciprocal function The reciprocal function takes a number and returns one divided by that number: $y = \frac{1}{x}$. This function also has an alternative form in which x is raised to a power: $y = x^{-1}$. Notice that the reciprocal function shown in figure 11.6 has several interesting features: It has different concavity on the left and the right; it does not even exist at $x = 0$ since any number divided by zero is undefined; in fact, the reciprocal function never crosses either axis.

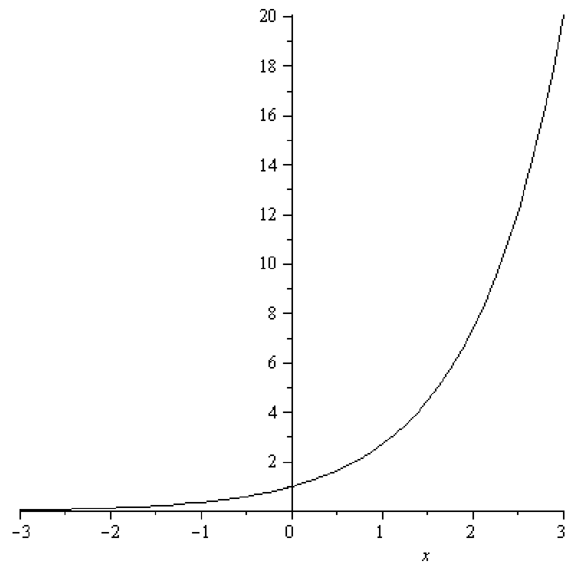

Figure 11.3: The basic exponential function $y = \exp(x)$.

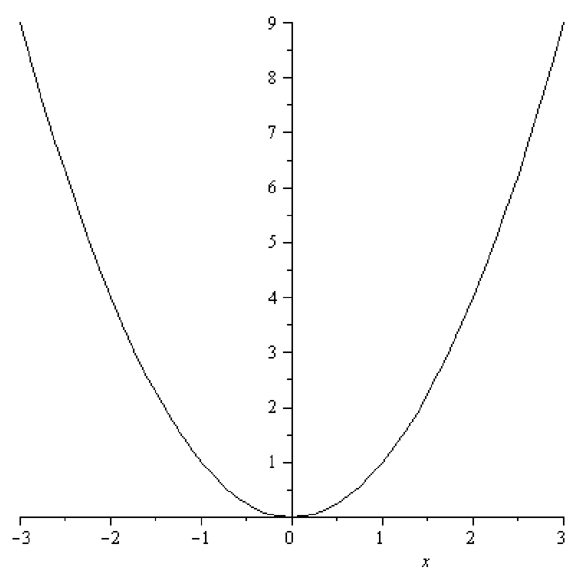

Figure 11.4: The basic squaring function $y = x^2$.

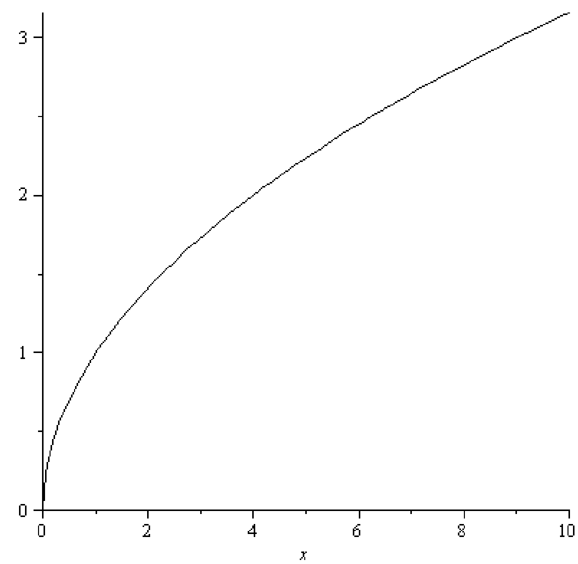

Figure 11.5: The basic square root function $y = \sqrt{x}$.

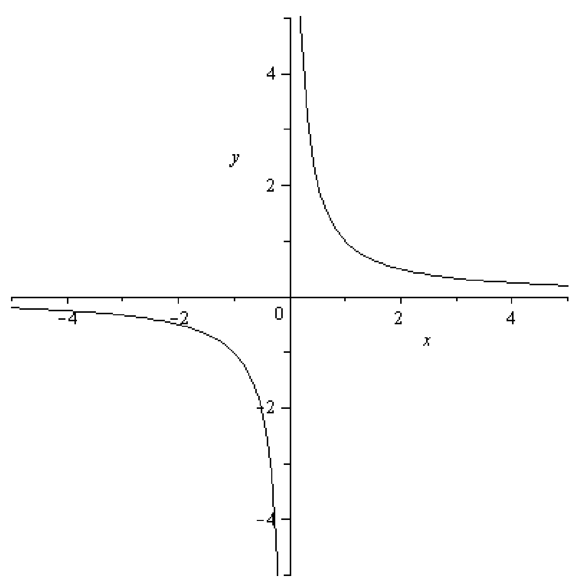

Figure 11.6: The basic reciprocal function $y = x^{-1} = \frac{1}{x}$.

11.1.2 Worked Examples

Example 11.1. Using a graph of the data to see nonlinearity
Consider the data graphed below. What can we say about it? It appears that as x increases, the y values decrease. It also looks like the data is bending upward. Mathematicians call this behavior "concave up". Let's see what happens when we apply a linear regression to these data.

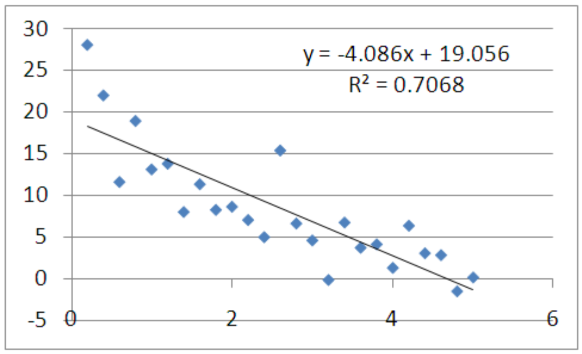

Figure 11.7: Does a linear function fit these data well?

It looks like a good candidate for fitting with a straight line, and the R^2 value is acceptable for some applications, but notice there are distinct patterns in the data, when compared to a best fit line. On the left, the data is mostly above the line, in the middle, the data is mostly below the line, and on the right, the data is mostly above the line. Patterns such as this indicate that changes in the Y data are not proportional to changes in the x data. To put this another way, the rate of change of y is not constant as a function of x, which means there is no single "slope number" that is the same for every point on the graph. Instead, these changes are level-dependent: as we move our starting point to the right, the y change for a given x change gets smaller and smaller. In a straight line, this is not the case: regardless of starting point, the y change for a given x change are the same. If the data are best represented by a linear model, we would not see any patterns in the data points when compared to the model line; the points should be spread above and below the best-fit line randomly, regardless of where along the line we are. For the graph above, though, we do see a pattern, indicating that these data are not well suited to a linear model.

Notice that R^2 by itself would not have told us the data is nonlinear, because the data is tightly clustered and has little concavity. Clearly, the more concave the data is, the worse R^2 will be for a linear fit, since lines have no concavity and cannot capture information about concavity.

Example 11.2. Comparing logarithmic models and square root models
You may have noticed that two of the functions above, the square root and the logarithmic, look very similar. Why do we need both of them? After all, the two graphs (see figure 11.8) have very similar characteristics. For example, both start off very steep for small values of x

and then flatten out as x increases. Both graphs continue to increase forever. Neither graph exists for negative values of x.

However, the graphs are actually quite different. For instance, consider the origin. The point $(0, 0)$ is a point on the square root graph (since the square root of zero is zero,) but it is not a point on the logarithmic graph. In fact, if you try to compute the natural log of zero, you will get an error, no matter what tool you use for the calculation! The logarithmic function has what is called a "vertical asymptote" at $x = 0$. This means that the graph gets very close to the vertical line $x = 0$, but never touches it. This is quite different from the square root graph which simply stops at the point $(0, 0)$. Furthermore, the square root has a horizontal intercept of $x = 0$, while the logarithmic graph crosses the x-axis at $(1, 0)$.

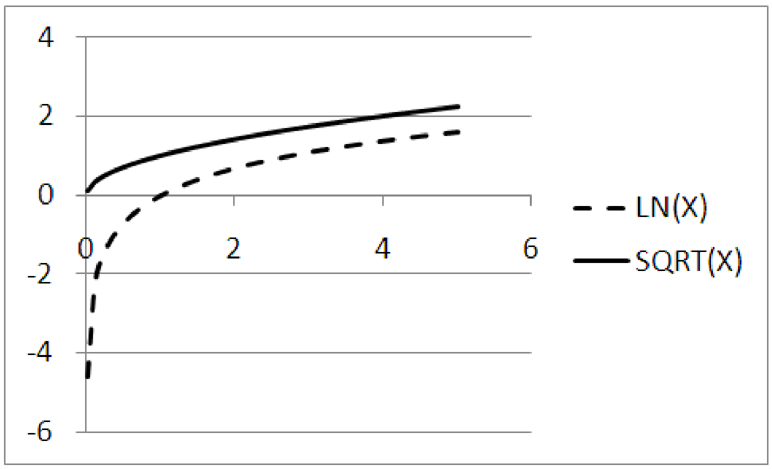

Figure 11.8: Comparison of standard log and square root functions.

You might be tempted to think that we could simply "move" the logarithmic graph over so that they both start at the same place, $(0, 0)$. Figure 11.9 shows what happens if we pick up the graph of $y = \ln(x)$ and move it to the left one unit, so that both graphs pass through $(0, 0)$. Notice that the square root graph rises sharply and then flattens, while the natural log graph rises more gradually. It also appears that the slope of the square root graph is larger and that it continues to grow larger, widening the gap between the two functions. In fact, the natural log grows so slowly that the natural log of 1,000 is only 6.9 and the natural log of 1,000,000 is 13.8! Thus, if the x-values of your data span a large range, over multiple orders of magnitude, a natural log may help scale these numbers down to a more reasonable size. This property of logs makes them useful for measuring the magnitude of an earthquake (the Richter scale) or the loudness of a sound (measured in decibels). Compare this growth to the square root function: The square root of 1,000 is about 31; the square root of 1,000,000 is 1000. This is a much larger increase than the natural log. In fact, of all the basic functions, the natural log is the slowest growing function; in a race to infinity, it will always lose.

Example 11.3. Comparing exponential models and square models

You may have also noticed that, for positive values of x, the graphs of the exponential

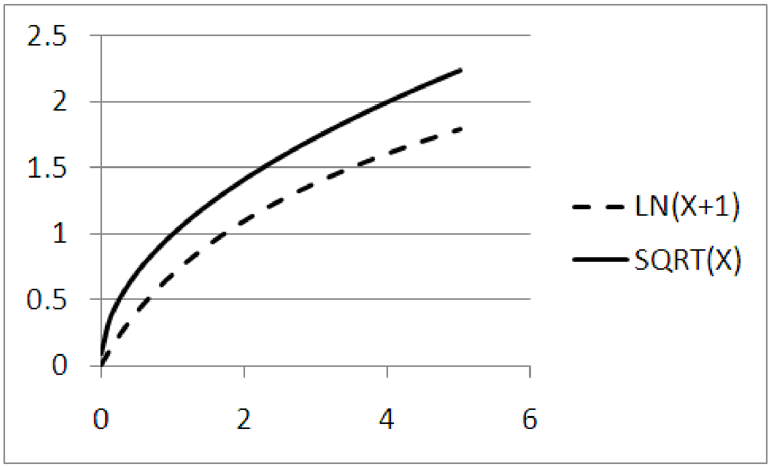

Figure 11.9: Comparison of horizontally shifted log and square root function.

function and the square function are very similar. Both are increasing. Both are growing large at a faster and faster rate, which shows in the graphs from the increasing steepness of each graph as x grows larger. This property of getting larger at an increasing rate is referred to as being "concave up." This makes the graphs bend upward, away from the x-axis so that it looks like a cup that could hold water. Both graphs also start rather flat near the origin.

Here is where the similarities end, though. The square function has a vertical and horizontal intercept at $(0, 0)$. The exponential function, on the other hand, has a vertical intercept of $y = 1$, but no horizontal intercept at all. Much like the logarithmic function (see the previous example) the exponential function has an asymptote. In this case, though, it is a horizontal asymptote at $y = 0$, rather than a vertical intercept at $x = 0$. In addition, when we look at the graphs for negative values of x, we see that the exponential function is always increasing, while the square function is decreasing for $x < 0$. This means that the square function has a minimum, or lowest, point.

These properties are also easy to see numerically from working with the functions themselves. If I take a negative number and square it, I get a positive number. Thus, $(-3)^2 = +9$, $(-2)^2 = +4$, $(-1)^2 = 1$, etc. Notice that as the negative values of x get closer to 0, the output of the square function is decreasing. For an exponential function, we notice negative exponents are really a shorthand way of writing "flip the function upside down and raise it to a positive power." Thus, to compute e^{-2}, we compute $1/e^2 \approx 1/7.3891 \approx 0.1353$. This is where the asymptotic nature of the exponential function shows through; for large negative powers, we are really computing one divided by e raised to a large positive power. Since e to a large positive power is a large positive number, one over this number is very small and close to zero.

As it turns out, the exponential function is the fastest growing of all the basic functions. In a race to infinity, it will always win.

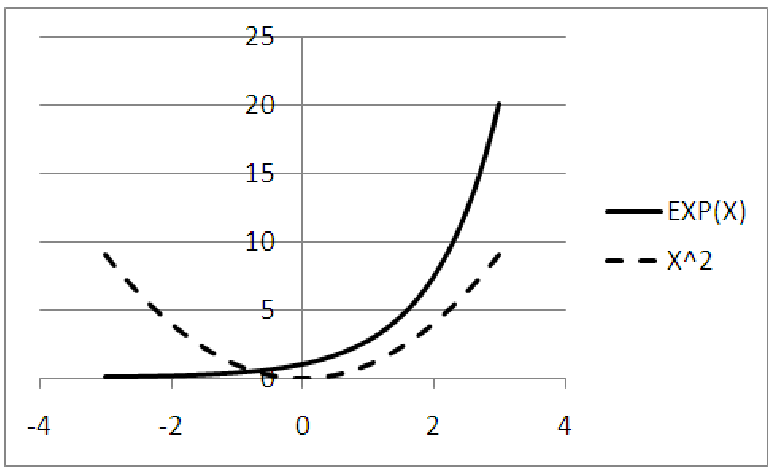

Figure 11.10: Comparison of exponential and squaring functions.

11.1.3 Exploration 11A: Developing our intuition about data that is non-proportional

Four graphs are shown below. For each, consider which of the basic functions you think would fit the data best. For your function, then describe what you think should be done to its graph in order to make it fit the data best. It might need to be shifted left or right, sifted up or down, flipped vertically or horizontally, stretched or squashed, or some combination of these.

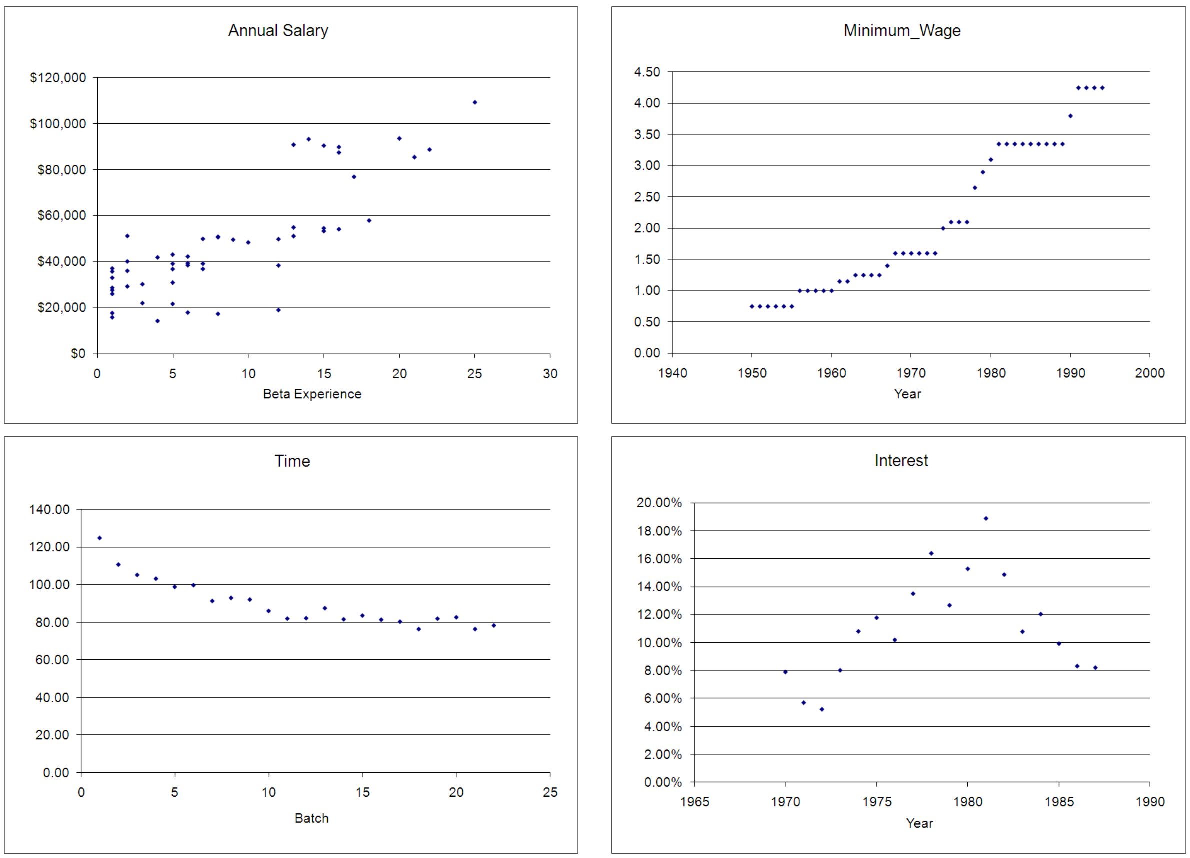

Data	Best choice for basic function	How to alter the function to fit
Annual Salary		
Minimum Wage		
Time		
Interest		

Now open the file "C11 Exploration2.xls". Test each of the possible trendlines (linear, logarithmic, exponential, power and polynomial of order 2 - do not use the moving average or higher order polynomials.) Be sure you display the equation and R^2 value for each of the possible models. Write down the equation of the best fitting model and record its R^2 in the chart below.

Data	Best Fit Trendline Equation	R^2
Annual Salary		
Minimum Wage		
Time		
Interest		

11.1.4 How To Guide

Adding Trendlines for Non-proportional Models

Excel can add trendlines for some non-proportional models to graphs. The process is virtually identical to the process used to add linear trendlines in Excel. (See the "How To Guide" for part 2 of chapter 7 (page 199) for details and screen shots.) The only difference is that in step 2 you should select the following options:

- To get an exponential fit, choose "exponential"
- To get a logarithmic fit, choose a "logarithmic"
- To get the square fit, use "Polynomial" and select "Order 2"
- It is not possible to force Excel to generate trendlines for reciprocals or square roots directly. As it turns out, these are specific cases of the more general "Power" models. However, if you add a "power" trendline to a graph, the power is one of the parameters in the model (like slope or y-intercept) so you probably will not get a power of 0.5 ($= \frac{1}{2}$ which is a square root model) or a power of -1 (for a reciprocal model). We'll talk about how to force Excel to fit these types of models with regression in the next chapter.

Computing Values of Exponentials and Logarithms

Excel uses a standard notation to compute the exponential or logarithm of a number. The notation looks a lot like the notation we have been using above:

- To compute the value of e^3, type "=EXP(3)" in a cell and hit enter.
- To get the value of e raised to whatever is in cell B2, type "=EXP(B2)"
- To compute the natural logarithm of 3, type "=LN(3)"
- To compute the natural log of the number in cell B2, type "=LN(B2)"

Note that Excel (and most calculating tools) have another logarithm function. This is the LOG(x) function. There is a slight difference between LOG(x) and LN(x). For our purposes, we will always use LN(x) when we talk about the logarithm of x.

Technical details: LOG(x) stands for the base-10 logarithm of x. LN(x) stands for the base-e logarithm of x. Essentially, when we compute a base-b logarithm of the number x we are finding the value of a so that the following equation is true: $b^a = x$. For example, since $10^2 = 100$, we know that the base-10 logarithm of 100 is 2 (i.e., $\log_{10}(100) = 2$.) Since $2^5 = 32$ we know that $\log_2(32) = 5$. Excel really only has options for base-e logarithms (LN) and base-10 logarithms (LOG). There are many other useful logarithm bases, but these are the most common, and there is a mathematical technique that relates logarithms of any two bases: $\log_b x = \frac{\ln(x)}{\ln(b)}$.

11.2 Transformations of Graphs

The basic functions introduced in the last section are very useful. With these, we can fit a lot more data than we could with just straight lines. However, we often find that the data matches the shape of one of these basic functions, but not the specific location and specific points that the basic function passes through. Consider the data shown on the left scatterplot below. The data shows the number of products (in this case motors) returned from a production line as a function of the amount of money spent on inspections in a given month (in thousands of dollars). On the right is a graph of this same data, but with the graph of the basic square function superimposed on it. They have the same shape, but are not exactly the same: The square function starts too low and rises much more quickly than the actual data, and they are not in the same place on the graph.

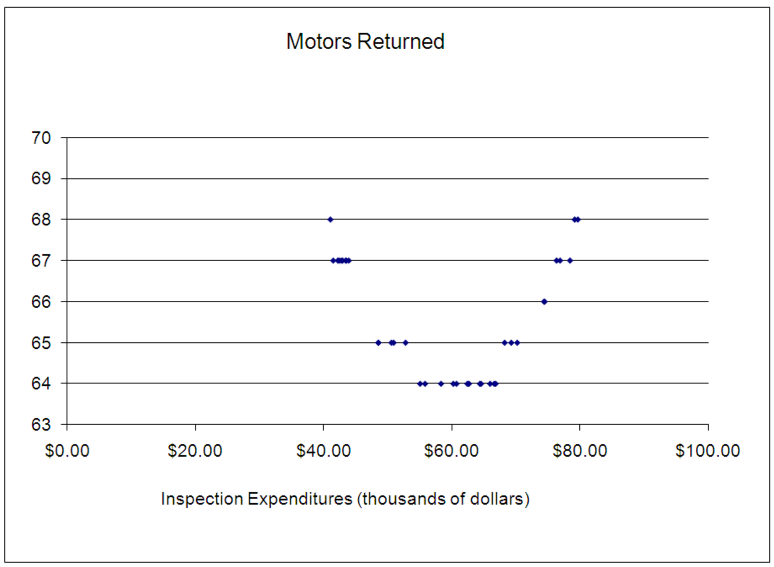

Figure 11.11: Graph of number of motors returned versus inspection expenditures.

The way to fix this is to "move the basic function around" until it fits more closely. This is very much like what we did with straight lines before: we knew we wanted a straight line, but we needed to change the slope and y-intercept until the theoretical line matched up better with the data. For the data above, it looks like we need to "squash" the square function down and move the starting point over and up. In this section, we'll explore the mathematical way of doing this. When we are done, we will have developed formulas for the basic functions that are more general and contain several parameters (almost always two). These parameters are, in general, much harder to interpret than the parameters in a linear model, but once we understand how they affect the graphs, we'll be able to out some meaning to them.

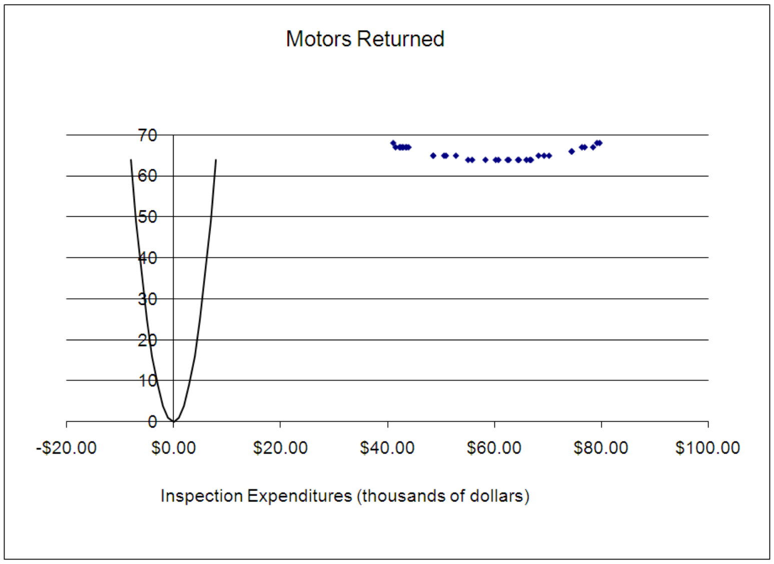

Figure 11.12: Graph of inspection expenditures with poorly fitting square model added.

11.2.1 Definitions and Formulas

Parameters A parameter is a number in the formula for a function that is constant. Changing a parameter will change the entire behavior of the function. The two parameters you are most familiar with are the slope and y-intercept of a linear function. If the slope parameter is changed, the line is more or less tilted; it may even change the direction of the tilt. If the y-intercept is changed, the graph crosses the y-axis at a different point. Most functions come in families of functions that all have the same formula, but the formula has parameters in it. Thus, linear functions of the form should really be called the "family of linear functions" since there are two parameters in the formula. To get the equation of a specific member of the family, we need to substitute in values for each of the two parameters, A and B. (Just like you need a first and last name to find a specific person in your family; you may sometimes need even more information about the person if more than one person in the family has the same name. Some functions also need more than two parameters. See quadratics below for such an example.)

Power functions This is a broad family of functions. The general form of a basic power function is where b is a number. Thus, this family includes the squaring function ($b = 2$), the square root function ($b = \frac{1}{2}$), the reciprocal function ($b = -1$), and the basic linear function ($b = 1$). This family is called the family of power functions because the independent variable, x, is always raised to a power. The shape of a power function depends on whether the power, b, is even or odd. Even power functions look something like a "U" when graphed. Odd power functions (with $b > 1$) look more like chairs: on the left they drop off; on the right they rise up high; in the middle they are relatively flat. All basic power functions pass through the origin $(0, 0)$ and the point $(1, 1)$. This is because zero raised to any power is zero and 1 raised to a power is always 1.

Polynomials A polynomial is a function made from adding together a bunch of power functions that all have whole number powers. (A whole number is a number like 5, 2, 0. Negative numbers and numbers with decimals and fractions are not allowed.) Each power function in a polynomial is multiplied by a coefficient and then they are all added together:

$$y = a_n x^n + a_{n-1} x^{n-1} + \ldots + a_2 x^2 + a_1 x^1 + a_0$$

Notice that since anything raised to the zero power is 1, there is no need to write x^0 in the last term. Each of the individual combinations of a coefficient and a power function in a polynomial is called a term. Polynomials include several well-known families of functions: the quadratics (see below) and the linear functions:

$$y = A + Bx.$$

The n in a power function gives the highest power in the polynomial. It is called the order of the polynomial. The shape of a polynomial function is highly dependent on the order of the polynomial, since this determines the leading power function in the polynomial. The following general statements can be made:

If n is even, then the polynomial function does the same thing on both sides of the y-axis: it either rises up on both sides or drops down on both sides. If n is odd, then the polynomial does the opposite on both sides: one side will rise, the other will drop. The order also determines two other properties: the maximum possible number of times the polynomial crosses the x-axis (the number of zeros) and the number of time the graph changes direction (either from increasing to decreasing or vice versa):

- Maximum number of zeros = n
- Maximum number of turning points = $n - 1$

Quadratics A quadratic function is a second-order polynomial that produces a "generalized squaring function". It is usually written in the following way:

$$y = Ax^2 + Bx + C.$$

You may have seen the famous quadratic formula. This is a formula for finding the roots of a quadratic equation. Roots are places where the function crosses the x-axis, so these points all have $y = 0$. Thus, they are solutions to the equation:

$$0 = Ax^2 + Bx + C.$$

Using the quadratic formula, we can find the x-coordinates of these crossing points:

$$x = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}$$

Excel will add quadratic trendlines to a graph; however, it refers to them by their more proper name as "polynomials of order 2".

Vertical shifting Sometimes the data we are trying to fit looks exactly like a basic function, but moved up or down. We can fix this by adding in a vertical shift to the equation. If the graph of a function has been vertically shifted, the graph has the same exact shape, only every single y-value has been increased by the same amount or every y-value has been decreased by the same amount. Effectively, this moves the entire graph of the function either up or down the y-axis. Thus, a vertical shift by k will move the y-intercept up by k units.

Horizontal shifting Sometimes, the data is moved right or left of the basic function that it is most similar to. We can compensate by adding a horizontal shift to the equation of the graph. If a graph has been horizontally shifted, the graph has been moved to the right or the left. Thus, if the graph is moved to the right h units, then the zeros of the function (if any) will all move to the right h units.

Translation This is the general term to refer to any type of shift (vertical or horizontal).

Vertical scaling It is sometimes necessary to stretch a graph out or compress the graph of a basic function so that it will match up better with the data. This can easily be done by multiplying the entire function by a scaling factor.

11.2.2 Worked Examples

Example 11.4. Vertical shift
Consider the data shown in the table below for $y = f(x)$. If we make a new function by adding the same amount, say 10, to each of these y values, then we will be creating the function $y = f(x) + 10$; each y value will be 10 more than it would be without the increase. This will result in the graph of the function being shifted up by 10 units at each data point. It's just like we picked up the graph and slid it up the y-axis 10 units.

x	$y = f(x)$	$y = f(x) + 10$
0	10	20
1	15	25
2	12	22
3	3	13
4	6	16
5	11	21
6	15	25
7	19	29
8	25	35
9	23	33
10	22	32

Example 11.5. Horizontal shift
We can also shift a graph to the left or right. In the last example, we added a value to all

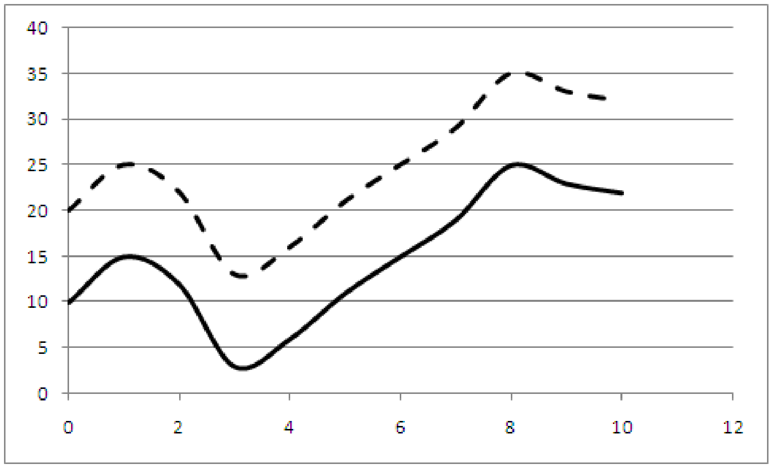

Figure 11.13: Graph of $y = f(x)$ (solid line) and $y = f(x) + 10$ (dashed line).

the y values in order to shift the graph up or down the y-axis. To shift left and right, we need to add or subtract from the x values. For example, suppose we wanted to move the graph four units to the right. The old graph would have the point (x, y) corresponding to the statement that $y = f(x)$. The new graph should have the point $(x + 4, y)$. So if the old graph had the point $(3, 3)$, the new graph should have the point $(3 + 4, 3)$ or $(7, 3)$. Here's the catch, though, the function will only give 3 for y if we plug in a value of 3 for x. We want to plug in 7 for x and get 3 out. Thus, we need to subtract 4 from each x value in order to make sure the function gives the right output. This means that to shift the function to the right 4 units, we need to plot the graph of $y = f(x - 4)$. This is shown in the data table below and the graph beside it.

x	$y = f(x)$	$y = f(x - 4)$
0	10	?
1	15	?
2	12	?
3	3	?
4	6	10
5	11	15
6	15	12
7	19	3
8	25	6
9	23	11
10	22	15
11	?	19
12	?	25
13	?	23
14	?	22

Example 11.6. Vertical scaling

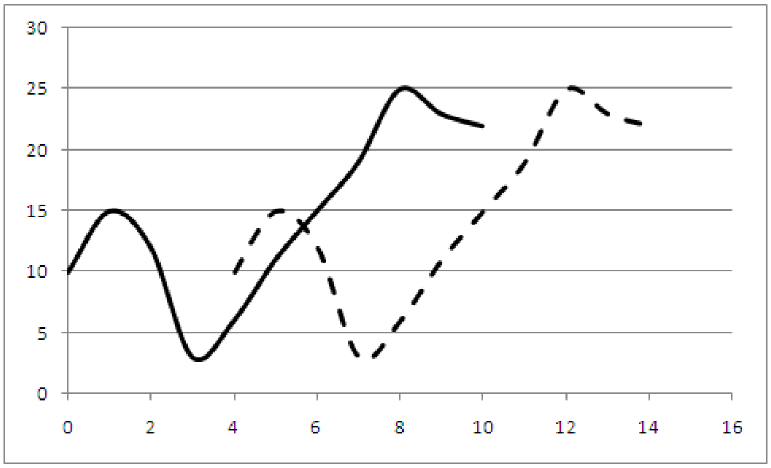

Figure 11.14: Graph of $y = f(x)$ (solid line) and $y = f(x-4)$ (dashed line).

We can also stretch the shape of a graph out. Suppose that we have a set of data that looks parabolic, so we want to use the square function. But suppose that the data contains the points $(1,1)$, $(2,8)$, $(3,18)$, and $(4,32)$. For a basic squaring function, this can't happen; the shape is right, but $2 \times 2 = 4$, not 8; $3 \times 3 = 9$, not 18; and $4 \times 4 = 16$, not 32. But notice, that each of the actual data values is simply twice what the squaring function would give. Thus, we want to graph $y = 2x^2$. This means that we should take each value of x, compute x^2, and then multiply the result by 2. This stretches the graph out to fit the data.

We can also compress a graph, squashing it down flatter instead of stretching it up taller. Suppose that the data contains the points $(1, 0.5)$, $(2, 2)$, $(3, 4.5)$ and $(4, 8)$. Each of the y values is half of what we would expect from the squaring function, so we want to graph $y = 0.5x^2$. We see that the general form to scale the graph of $y = f(x)$ is $y = a \times f(x)$ where a is a constant. The graphs below show the basic squaring function and the two functions we have just created.

But what happens if we let a be a negative number? This will simply take each of the old y values from the function and put a negative sign in front of them. This flips the graph over the x-axis, creating a mirror reflection of the original graph. Thus, the graph of $y = -f(x)$ is the same as the graph of $y = f(x)$ except that it is flipped over the x-axis. In a similar way, multiplying x by a factor can scale the graph horizontally, and negating x flips the graph horizontally over the y-axis.

Example 11.7. Combination of Shifts and Scales

Consider the graphs shown in the introduction to this section in figure 11.11 (page 325). The data for the number of motors returned as a function of inspection expenditures looks to be a basic squaring function, but shifted and scaled. Here's another look at the graph.

It looks like the graph has been shifted to the right 60 units and up 64 units. Thus, we could start by comparing the data to the graph of $y = f(x-60) + 64 = (x-60)^2 + 64$. When we do this, we find that the graph starts in the right place, but climbs too quickly. We might be tempted to simply multiple this whole thing by a constant less than one in

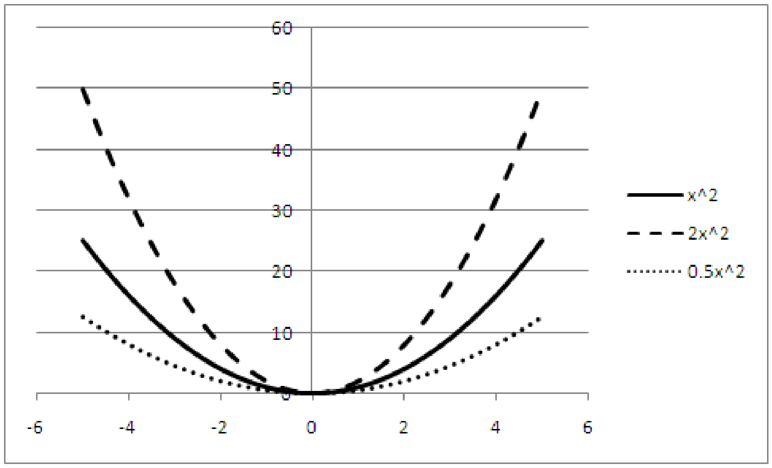

Figure 11.15: Graphs of $y = x^2$ (solid line), $y = 2x^2$ (dashed line) and $y = 0.5x^2$ (dotted line).

order to squash the graph, but this would also multiply the vertical shift by the constant, changing the starting place. We must complete the shifts and scaling in the proper order. We need to construct the fit by first shifting right, then scaling, then shifting. So, we are looking for a function of the form $y = af(x-60)+64 = a(x-60)^2+64$.

How much should we squash the graph? In other words, how big is a? The best approach here is to try a few data points. It looks like the point $(80, 68)$ is on the graph. Plugging these values in for x and y we get the following:

$$\begin{aligned} 68 &= a(80-60)^2+64 \\ 68 &= a(20)^2+64 \\ 68-64 &= a(20)^2 \\ 4 &= a(20)^2 \\ \frac{4}{20^2} &= a \\ a &= 0.01 \end{aligned}$$

Thus, the equation of the function that seems to match the data is $y = 0.01(x-60)^2+64$, where y represents the number of motors returned, and x represents the amount of money (in thousands) spent on inspection expenditures in a given month. We should check this against a few more data points, to be certain that the function is the correct one. Since the point $(75, 66)$ also appears to be on the function, we evaluate our candidate function at this x value to see if they match. At $x = 75$ our function is equal to $y = 0.01(75-60)^2+64 = 66.25$ which is very close to the value given by the data. We don't expect a perfect fit, because the data is not taken from an abstract function, but actually came from a real situation, so there will likely be some error in the best-fit function.

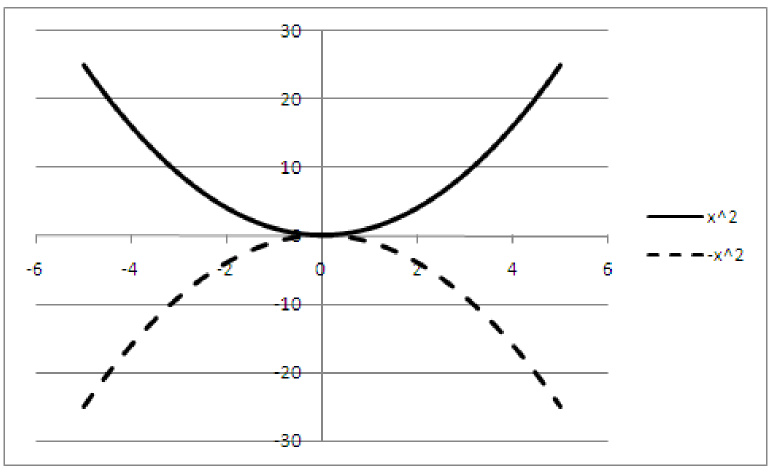

Figure 11.16: Graphs of $y = x^2$ (solid line) and $y = -x^2$ (dashed line).

11.2.3 Exploration 11B: Shifting and Scaling the Basic Models

Download and open the file C11 Exploration3.xls. The file contains several macros, so when you open the file, you you may need to click on the options button next to the security warning. Then select the option labeled "enable the content". (This is part of the security of the computer; many viruses and computer worms are hidden in macros.) When you get the file open, you should see a screen like the one shown below:

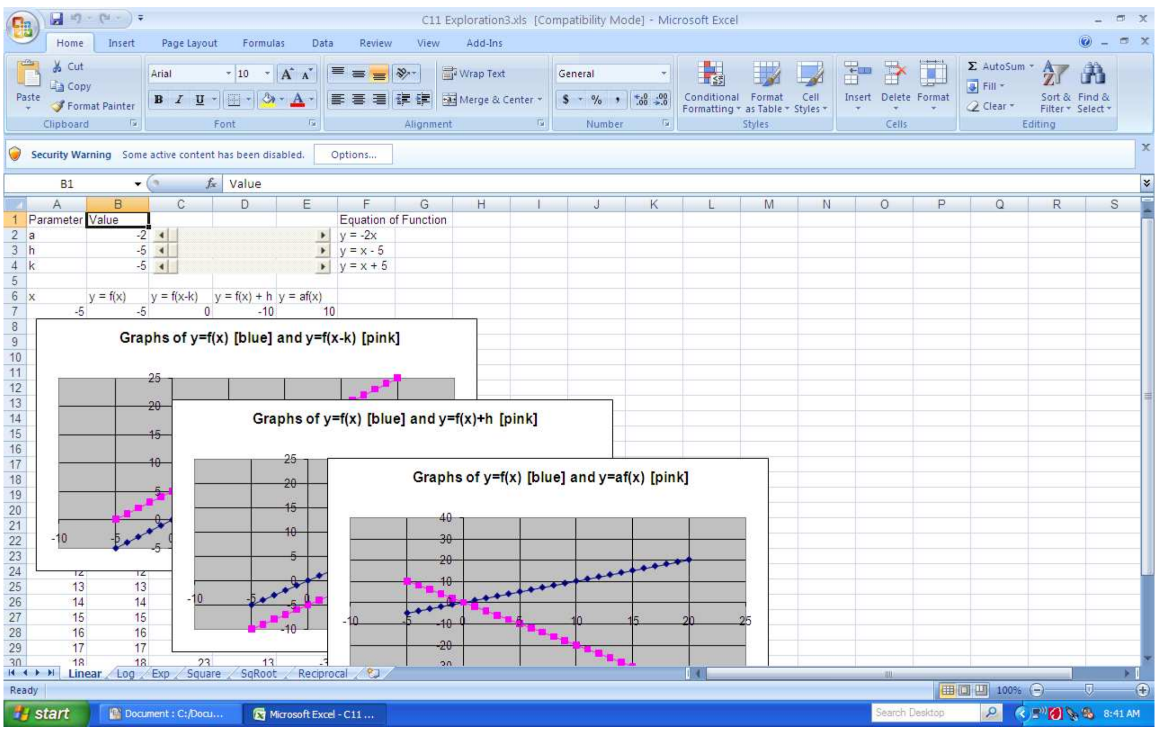

Figure 11.17: Excel file for exploring the various shifts and scalings with the basic functions.

There are six worksheets in the workbook, one for each of the basic functions we have been discussing. On each worksheet there are three slider bars and three graphs. Each graph shows the graph of the basic function itself (in blue) and one other graph (in pink). As you change the slider values, make note of how the graph of the pink function changes and how the different equations shown next to the slider bars change. To see some of the graphs, you may need to right click on them and select "Bring to Front" since Excel layers its graphs on top of each other in order to save "screen real estate". Use the worksheets and the sliders to help fill in the details about each of the functions below.

Linear Function, $f(x) = x$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x - h)$		
$y = f(x) + k$		

Logarithmic Function, $f(x) = \ln(x)$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x - h)$		
$y = f(x) + k$		

Exponential Function, $f(x) = e^x$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x - h)$		
$y = f(x) + k$		

Squaring Function, $f(x) = x^2$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x-h)$		
$y = f(x) + k$		

Square Root Function, $f(x) = \sqrt{x}$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x-h)$		
$y = f(x) + k$		

Reciprocal Function, $f(x) = \frac{1}{x}$		
Modification	Sketch	Description
$y = af(x)$		
$y = f(x-h)$		
$y = f(x) + k$		

11.2.4 How To Guide

Setting up functions in Excel for shifting and Scaling

In the how-to for chapter 7 (page 199), we introduced the idea of setting up an Excel spreadsheet to calculate a table of function values. We can use this same idea for calculating the values of a function with arbitrary shifts (horizontal and vertical) and scalings. For example, suppose we want to fit a shifted curve to a set of data that has x-values in cells A6:A25 and y-values in cells B6:B25. Let's add in some parameter values. Enter labels for each shift in A1:A3 and sample values for these shifts in B1:B3. You now have a worksheet that looks something like the one at the right.

Now we need to add in a column of values for the predicted y data, according to our formula, using the shifts and scales. Suppose that we want to use a logarithmic function to try and fit the data. So, we want to try to use the formula

y = (Vertical shift) + (Vertical Scale)*ln(X + Horizontal shift)

To do this, we enter the following formula into cell C6 and copy it down the column (Note: This is a formula for the logarithmic model we are currently working with; for other models, you will have to develop a different formula):

```
= $B$1 + $B$3*ln(A6 + $B$2)
```

Notice that we are using absolute cell references to look up the values of the parameters and compute the predicted y-values. This way, the constants will remain correct as we copy the formula down, but the x-values will change, based on which row we are in. This format will easily allow us to change the shifts and scales to try and match the actual data (in column B). A visual representation (a scatterplot) would also help, since the graph would help you see the shifts and try to move the graph of the predicted y-values closer to the actual data.

Create a scatter plot with more than one y-variable on the same axes

StatPro will not allow you to create a graph with more than one dependent variable plotted against the same independent variable. For example, in the graph below you see two y-values plotted for each x-value. To generate such graphs, you will have to use the chart wizard in Excel. The steps are outlined below.

1. Highlight all the data you wish to graph, including all the y-variables. Make sure that your data is organized with the x-variable (the independent variable) in the left-hand column.

2. Click on the Chart icon in Excel or go under "Insert" and click on "Chart"

3. Click on xy-scatter

4. Click through all the "Next"'s and click "Finish"

5. Move/resize your graph

	A	B
1	Vertical Shift	1
2	Horizontal Shift	2
3	Vertical Scaling	-3
4		
5	X	Y
6	0	5.841117
7	1	3.408326
8	2	1.682234
9	3	0.343373
10	4	-0.75056
11	5	-1.67546
12	6	-2.47665
13	7	-3.18335
14	8	-3.81551
15	9	-4.38737
16	10	-4.90944
17	11	-5.3897
18	12	-5.83434
19	13	-6.2483
20	14	-6.63553
21	15	-6.99928
22	16	-7.34223
23	17	-7.66663
24	18	-7.97439
25	19	-8.26713
26	20	-8.54625

Figure 11.18: Setting up a function with vertical and horizontal shifts, and a scaling factor.

Calculating R^2 for your fit of the data

Now, it can be very hard to see if your shifts and scales have done a good job at fitting the data, so it is helpful to add in computations to determine the value of R^2 the coefficient of determination. For this, we'll need to compute the mean of the actual y-data, the variation in the actual y-data, the residuals, and finally the R^2.

1. Compute the mean of the actual y-data.

   A good place to put the mean of the actual y-data is to put it at the bottom of the column of y-data, separated from the data by a blank row. So, in cell A28 we place a label "Y Mean" and in cell B28, we place the formula "=average(B6:B26)" to compute the mean of the y-data.

2. Compute the variation in the actual y-data.

   The variation in the y-data is the sum of the squares of the differences between each y-value and the mean of the y-data. So, we add a column to the right of the data to compute the square of these differences. In cell D5, place the label "(y-mean)ˆ2". In cell D6, we place a formula to calculate the squared difference: "=(B6-B28)ˆ2" and then copy this formula to the other cells in the column. At the bottom of the column, in cell D28, we compute the sum of these squared differences: "=sum(D6:D26)". This sum of the squared differences is the total variation in the y-data.

3. Compute the residuals (and the sum of their squares).

   We now do a similar calculation in column E. We need the squares of the differences between the actual y-data and the predicted y-data. In cell E5 we place the label

"Residualˆ2". In cell E6, we put "=(B6-C6)ˆ2" and copy this formula to the rest of column E. At the bottom of this column, in cell E28, we sum up these squared residuals: "=sum(E6:E26)".

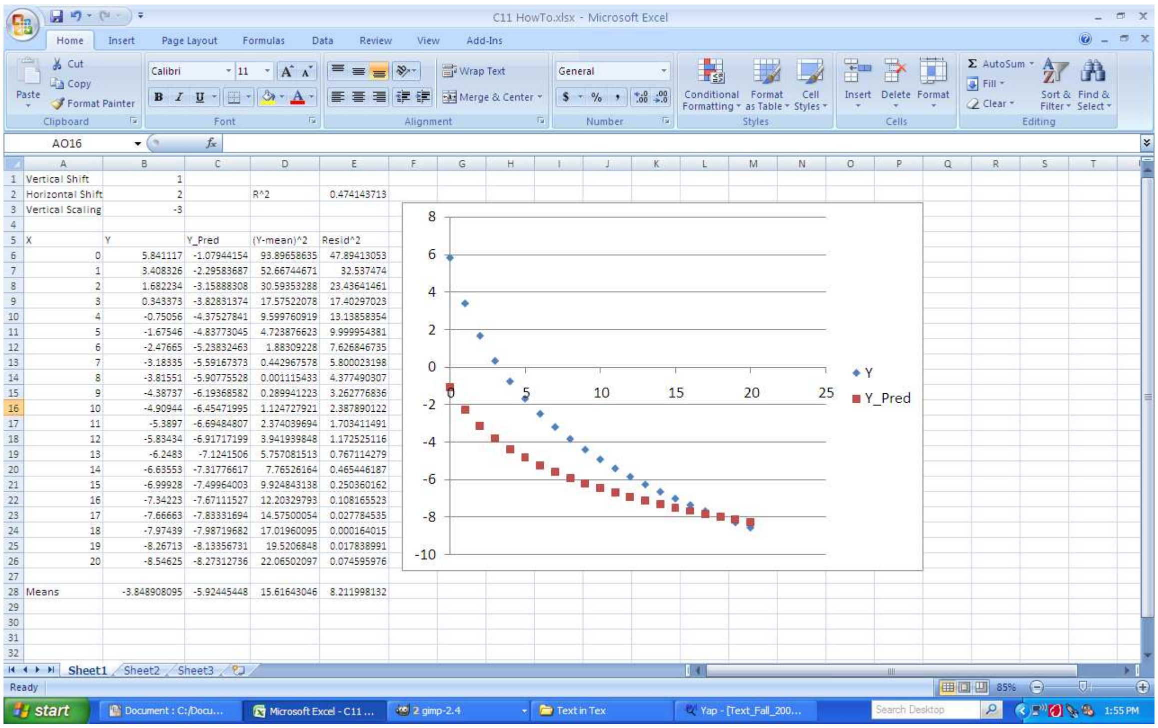

Figure 11.19: Spreadhseet for using parameters to attempt finding a best fit to a set of data.

4. Compute R^2.

   The value of R^2 can be calculated by subtracting the ratio of the sum of the squares of the residuals and the total variation in the y-data. So, in cell D2, we place the label "R^2" and in cell E2 we enter the formula "=1 - E28/D28". This will probably result in an R^2 value that is not between 0 and 1. This is because we have not actually created a "best fit" line until we do some adjustment on the parameters.

When you're done, the spreadsheet should look like figure 11.19 (you can add a graph in as well, by selecting the first three columns of data and making a scatterplot of them).

Using Goal Seek to find the best values for the shifts and scales

Once you have the spreadsheet configured to compute a set of y values based on your parameters for the shifts and scales, and have added the calculation for R^2 into the sheet to see how well your shifted and scaled curve fits the data, you can try to find the best values of the shifts and scales in one of two ways: trial and error or using goal seek. Trial and error is tedious, and can be difficult, if the data is tightly packed. Often, changing a value doesn't produce any noticeable change in the graph, making this process even harder. We recommend that you start by trial and error to get roughly the right values of the parameters, then refine the values with goal seek. (Goal seek was introduced in the computer how-to for section 7B.)

Suppose your spreadsheet is set up in the following way: You have values for three parameters (horizontal shift, vertical shift, and scale) in cells B1, B2 and B3. You have a calculated value of R^2 from the data and your best fit in cell E2. Now, there's no guarantee that the data can be "perfectly fit" so we can't try to goal seek for $R^2 = 1$. We also don't know ahead of time what the "best R^2" is, so we have to work our way up to it. Select the cell with the R^2 value in it. Then activate goal seek with "Tools/ Goal Seek". In the blank for "To value" enter a number that is a little better than the current value of R^2. For example, if the current value is 0.90, try entering 0.91. Now, the "By changing" cell should be one of the three parameters. Do this systematically: goal seek by changing the first parameter (B1) then repeat this, changing B2, then repeat changing B3, and start again at B1. Eventually, you'll get the R^2 up close to the "best possible" value. Each time, try increasing the "To value" amount a little, say 0.01 at a time. At some point, you'll know you're close to the best when the routine spits out "Goal seek may not have found a solution." When this happens, goal seek again with the same target value, but using the other parameters.

11.3 Homework

11.3.1 Mechanics and Techniques Problems

11.1. This problem deals with what happens to equations of functions and graphs of functions if you apply several different transformations, one after the other. Take $y = f(x) = x^2$ and write out each of the following functions. For each step, explain what happens to the graph in terms of how the particular change affects the appearance of the previous graph in the sequence.

Function	Written out	What happens to graph
$y = f(x) = x^2$		
$y = f(x-h)$		
$y = af(x-h)$		
$y = af(x-h) + k$		

11.2. How would the results in problem 1 be different if we changed the order to $y = a[f(x-h) + k]$?

Function	Written out	What happens to graph
$y = f(x) = x^2$		
$y = f(x-h)$		
$y = f(x-h) + k$		
$y = a[f(x-h) + k]$		

11.3. Now repeat 1 with a basic exponential function.

Function	Written out	What happens to graph
$y = f(x) = \exp(x)$		
$y = f(x - h)$		
$y = af(x - h)$		
$y = af(x - h) + k$		

11.4. Now repeat 1 with a basic logarithmic function.

Function	Written out	What happens to graph
$y = f(x) = \ln(x)$		
$y = f(x - h)$		
$y = af(x - h)$		
$y = af(x - h) + k$		

11.5. For each of the five graphs below

1. Select the best basic function to fit the data,
2. Select appropriate shifts (direction) and scaling (stretch or compress), and
3. Write down a possible equation for the graph.

11.6. Consider the data shown below in both table and graphical format.

x	y
0	2.05
5	2.69
10	3.55
15	4.23
20	4.35
24	5.08

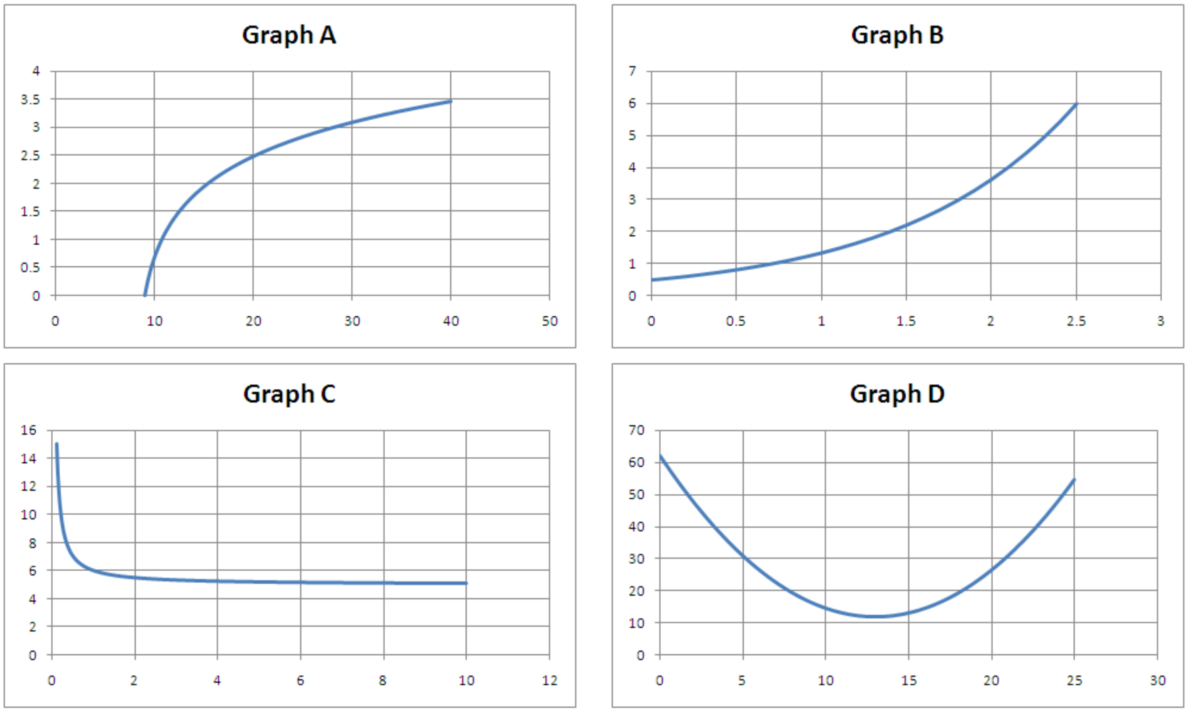

Figure 11.20: Graphs for problem 5.

1. Create a scatterplot of the data, and determine which order polynomial function (2 through 6) fits the data best. Record the results of your investigation in a table like the one shown below.

Order	Equation	R^2
2		
3		
4		
5		
6		

2. Use the parameters from your quadratic trendline (the order 2 polynomial) to manually calculate the S_e for that model. You may want to set up a spreadsheet like the one in figure 11.22.

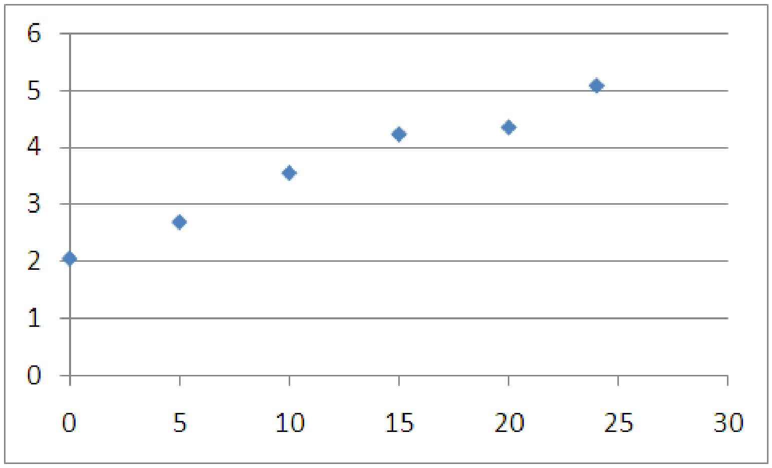

Figure 11.21: Graph of the data from problem 6.

	A	B	C	D	E	F	G	H	I
1	x	y	yhat(fitted)	Residual	Deviation		Parameters	a	
2	0	2.05						b	
3	5	2.69						c	
4	10	3.55							
5	15	4.23							
6	20	4.35					Calculations	SSR	
7	24	5.08						SSE	
8								SST	
9	ybar =>							R^2	
10								S_e	

Figure 11.22: Set up to calculate R^2 and S_e in problem 6.

11.3.2 Application and Reasoning Problems

Coming soon

11.3.3 Memo Problem

To: Analysis Staff
From: Project Management Director
Date: May 28, 2008
Re: DataCon Contract

We have received a contract from DataCon, a large data analysis provider that does general data analysis and management contracting for a wide variety of manufacturing and service sector businesses. They have subcontracted some of their business to us. They want us to fit some predictive models for four sets of data they have sent along. They want to see a best-fit Excel trendline for each data set, as well as a model that we come up with that is superimposed on both the scatterplot of the data and the best-fit trendline (See the How To Guide). DataCon management wants not only Excel's trendlines but also good fitting models we construct from shifts and scaling of basic functions because models built from basic functions are more transparent and easier to analyze than Excel's trendline models.

As usual, direct your memo to me. Include the following:

- A brief introduction
- A "Paste Special" copy into Word of your spreadsheet for each data set (there will be 4 such copies). Each spreadsheet copy will include the data set, along with your computations for R^2, and your settings for the parameters of your best-fit model.
- A copy of filled-in chart below.
- A few summary comments, including any special considerations your want to pass along about what you found.

Attachment: Data file C11 DataCon Data.XLS

Here are some suggestions for dealing with this assignment:

1. Start by fitting the best Excel trendline (don't forget to record its equation and its R^2) to a scatterplot of the data set. The table below (or one like it) will help organize the information.
2. Now try fitting your own shifted and scaled basic function on top of the scatterplot and the best-fit trendline, comparing your computed R^2 to the R^2 of the trendline.
3. You probably won't be able to better Excel's trendline (though you might!), but get as close as you reasonably can. That's all DataCon really wants or needs.

4. Here are a couple of tips:

   (a) Don't even try to do your own fit for a polynomial function (used when the scatter plot has a turn(s), etc) because Excel's polynomial fit is clear and understandable already. Your job, in this case, is to find Excel's best-fit polynomial.

   (b) If you are fitting your own exponential function, don't bother with horizontal shifts because mathematically such shifts can be absorbed by the scaling parameter.

   (c) If Excel's best trendline is a power function with a fractional power, for example, X0.42, use X0.5 for your own power function because $X^{0.5} = X^{1/2} = \sqrt{(X)}$, which is much easier to understand (remember, this is what DataCon wants).

DATA SET 1		EQUATION	R^2
	My Best Fit		
	Excel's Trendline Fit		
DATA SET 2			
	My Best Fit		
	Excel's Trendline Fit		
DATA SET 3			
	My Best Fit		
	Excel's Trendline Fit		
DATA SET 4			
	My Best Fit		
	Excel's Trendline Fit		