

The Use and Abuse of Logarithmic Axes

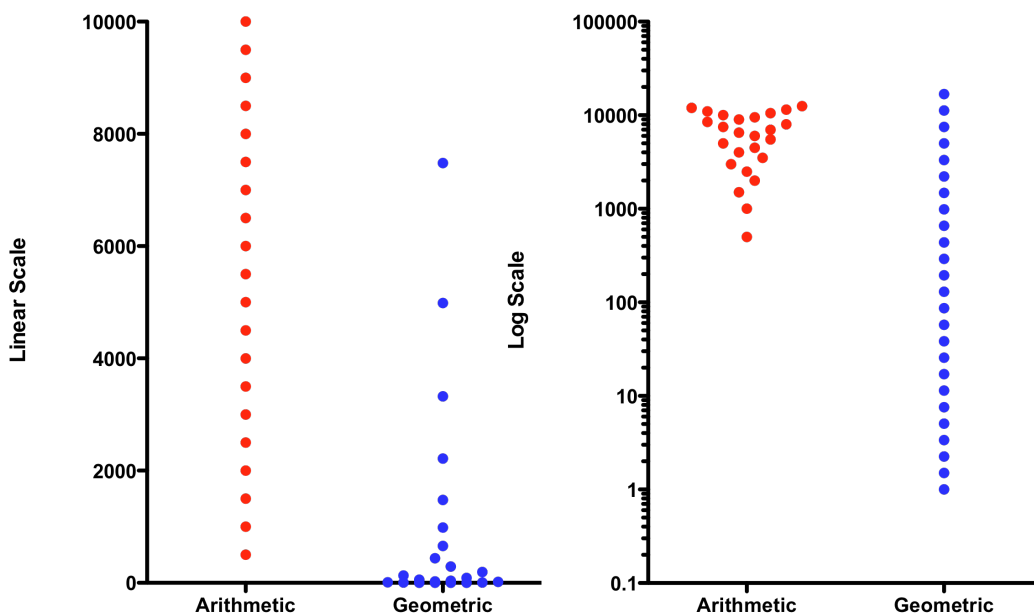
Harvey J. Motulsky
hmotulsky@graphpad.com
© 2009, GraphPad Software, Inc.
June 2009

Most graphing programs can plot logarithmic axes, which are commonly used and abused. This article explains the principles behind logarithmic axes, so you can make wise choices about when to use them and when to avoid them.

What is a logarithmic axis?

A logarithmic axis changes the scale of an axis.

The two graphs below show the same two data sets, plotted on different axes.



The graph on the left has a linear (ordinary) axis. The difference between every pair of ticks is consistent (2000 in this example).

The graph on the right has a logarithmic axis. The difference between every pair of ticks is not consistent. From the bottom tick (0.1) to the next tick is a difference of 0.9. From the top tick (100,000) down to the next highest tick (10,000) is a

difference of 90,000). What is consistent is the ratio. Each axis tick represents a value ten fold higher than the previous tick.

The red dots plot a data set with equally spaced values. Each dot represents a value with a Y value 500 higher than the dot below. The dots are equally spaced on the graph on the left, but far from equally spaced on the graph on the right. To prevent overlap, the points are jittered to the right and left so they don't overlap. The horizontal position of the red dots has no other meaning.

The blue dots represent a data set where each value represents a Y value 1.5 times higher than the one below. On the graph on the left, the lower values are almost superimposed, making it very hard to see the distribution of values (even with horizontal jittering). On the graph on the right with a logarithmic axis, the points appear equally spaced.

Interpolating between log ticks

What value is halfway between the tick for 10 and the one for 100? Your first guess might be the average of those two values, 55. But that is wrong. Values are not equally spaced on a logarithmic axis. The logarithm of 10 is 1.0, and the logarithm of 100 is 2.0, so the logarithm of the midpoint is 1.5. What value has a logarithm of 1.5? The answer is $10^{1.5}$, which is 31.62. So the value half way between 10 and 100 on a logarithmic axis is 31.62. Similarly, the value halfway between 100 and 1000 on a logarithmic axis is 316.2.

Why “logarithmic”?

In the example above, the ticks at 1, 10, 100, 1000 are equally spaced on the graph. The logarithms of 1, 10, 100 and 1000 are 0, 1, 2, 3, which are equally spaced values. Since values that are equally spaced on the graph have logarithms that are equally spaced numerically, this kind of axis is called a “logarithmic axis”.

Lingo

The term *semilog* is used to refer to a graph where one axis is logarithmic and the other isn't. When both axes are logarithmic, the graph is called a *log-log plot*.

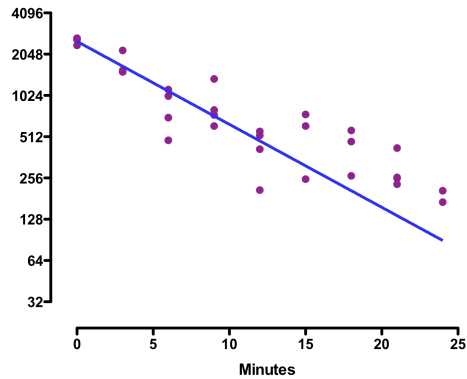
Other Bases

All the logarithms shown above are called base 10 logarithms, because the computations take 10 to some power. These are also called common logarithms.

Mathematicians prefer natural logarithms, using base e (2.7183 . . .). But they don't seem very natural to scientists, and are rarely used as an axis scale.

A base 2 logarithm is the number of doublings it takes to reach a value, and this is widely used by cell biologists and immunologists. If you start with 1 and double it four times (2, 4, 8, and 16), the result is 16, so the log base 2 of 16 is 4.

GraphPad Prism can plot log2 axes. Choose from the upper left corner of the Format Axis dialog.



Logarithmic axes cannot contain zero or negative numbers

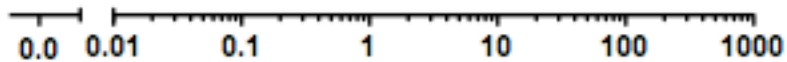
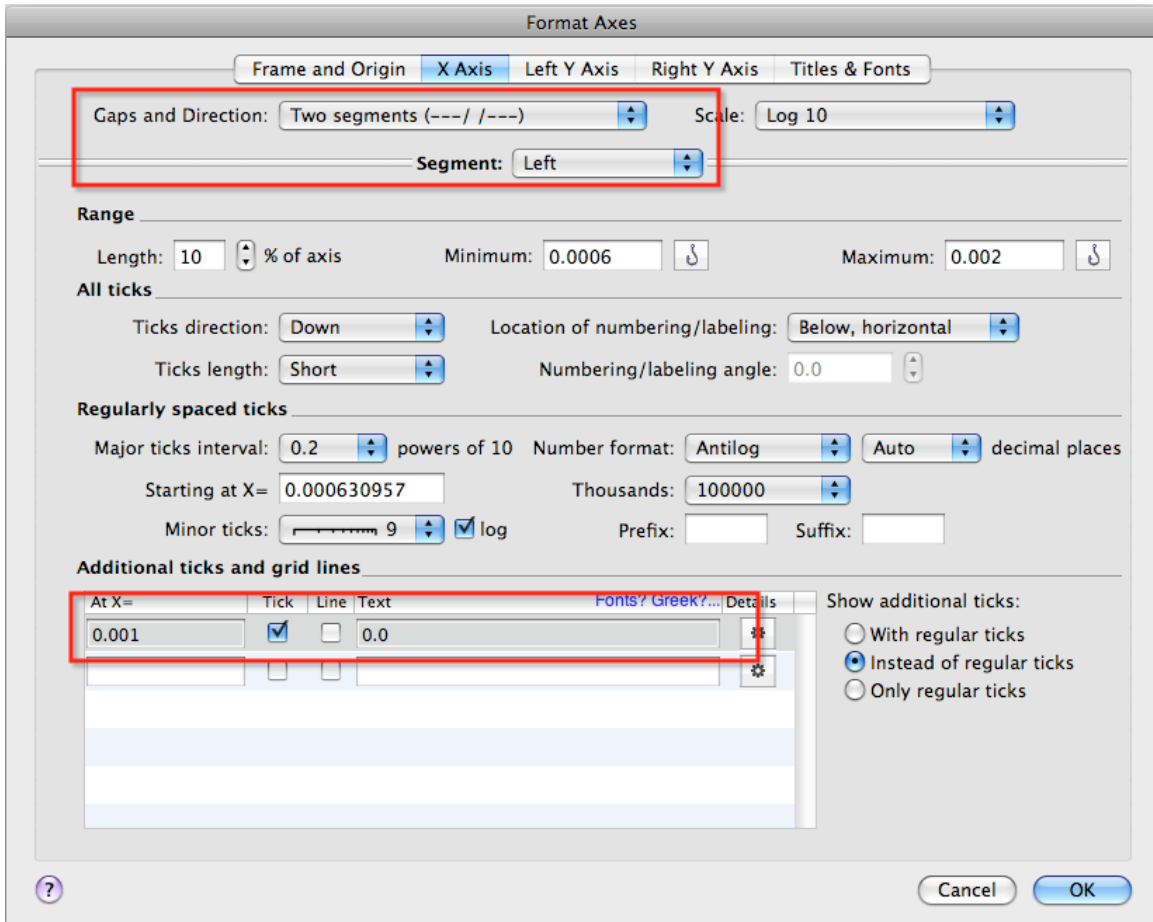
The logarithms of negative numbers and zero are simply not defined

Let's start with the fundamental definition of a logarithm. If $10^L = Z$, then L is the logarithm (base 10) of Z . If L is a negative value, then Z is a positive fraction less than 1.0. If L is zero, then Z equals 1.0. If L is greater than 0, then Z is greater than 1.0. Note that there no value of L will result in a value of Z that is zero or negative. Logarithms are simply not defined for zero or negative numbers.

Therefore a logarithmic axis can only plot positive values. There simply is no way to put negative values or zero on a logarithmic axis.

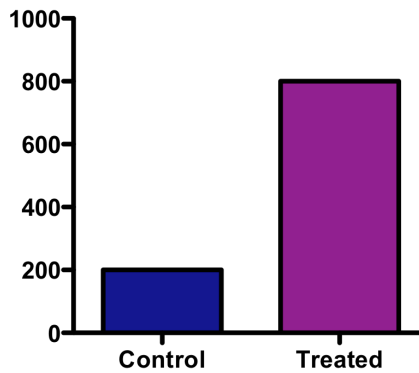
A trick to plot zero on a logarithmic axis in Prism

If you really want to include zero on a logarithmic axis, you'll need to be clever. Don't enter 0, instead enter as a small number. For example, if the smallest value in your data is 0.01, enter the zero value as 0.001. Then use the Format Axis dialog to create a discontinuous axis, and use the Additional ticks feature of Prism to label that spot on the axis as 0.0.

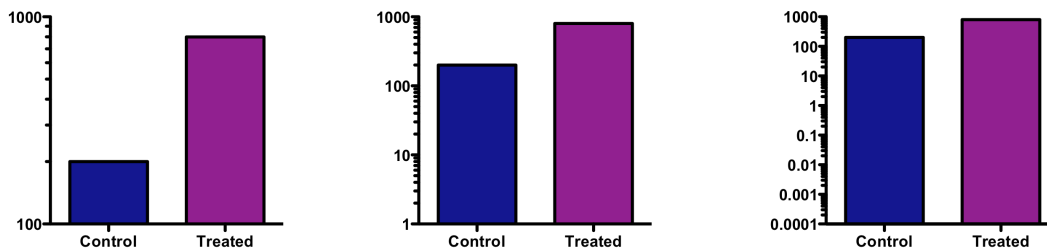


Logarithmic axes on bar graphs are misleading

The whole point of a bar graph is that the relative height of the bars tells you about the relative values plotted. In the graph below, one bar is four times as tall as the other and one value (800) is four times the other (200).



Since zero can't be shown on a log axis, the choice of a starting place is arbitrary. Accordingly, the relative height of two bars is not directly related to their relative values. The graphs below show the same data as the graph above, but with three different choices for where the axis begins. This arbitrary choice influences the relative height of the two bars, amplified in the graph on the left and minimized in the graph on the right.



If the goal is to create propaganda, a bar graph using a logarithmic axis is a great tool, as it lets you either exaggerate differences between groups or minimize them. All you have to do is carefully choose the range of your axis. Don't create bar graphs using a logarithmic axis if your goal is to honestly show the data.

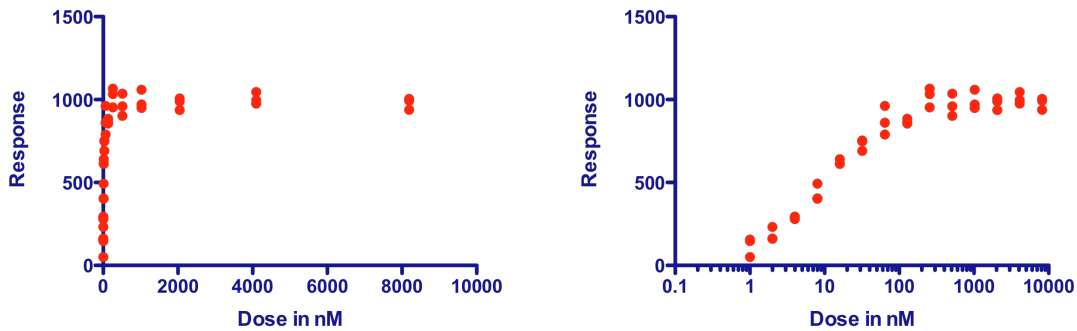
When to use a logarithmic axis

A logarithmic X axis is useful when the X values are logarithmically spaced

The X-axis usually plots the independent variable – the variable you control. If you chose X values that are constant ratios, rather than constant differences, the graph will be easier to view on a logarithmic axis.

The two graphs below show the same data. X is dose, and Y is response. The doses were chosen so each dose is twice the previous dose. When plotted with a linear axis (left) many of the values are superimposed and it is hard to see what's going on.

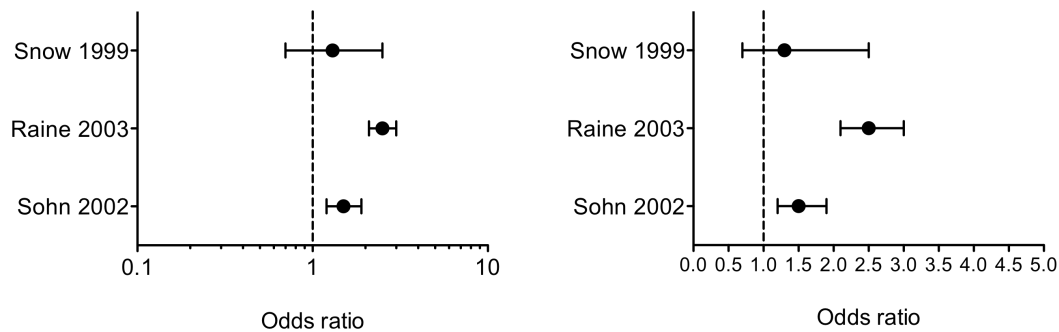
With a logarithmic axis (right), the values are equally spaced horizontally, making the graph easier to understand.



A logarithmic axis is useful for plotting ratios

Ratios are intrinsically asymmetrical. A ratio of 1.0 means no change. All decreases are expressed as ratios between 0.0 and 1.0, while all increases are expressed as ratios greater than 1.0 (with no upper limit).

On a log scale, in contrast, ratios are symmetrical. A ratio of 1.0 (no change) is halfway between a ratio of 0.5 (half the risk) and a ratio of 2.0 (twice the risk). Thus, plotting ratios on a log scale (as shown below) makes them easier to interpret. The graph below plots odds ratios from three case-control retrospective studies, but any ratio can benefit from being plotted on a log axis.

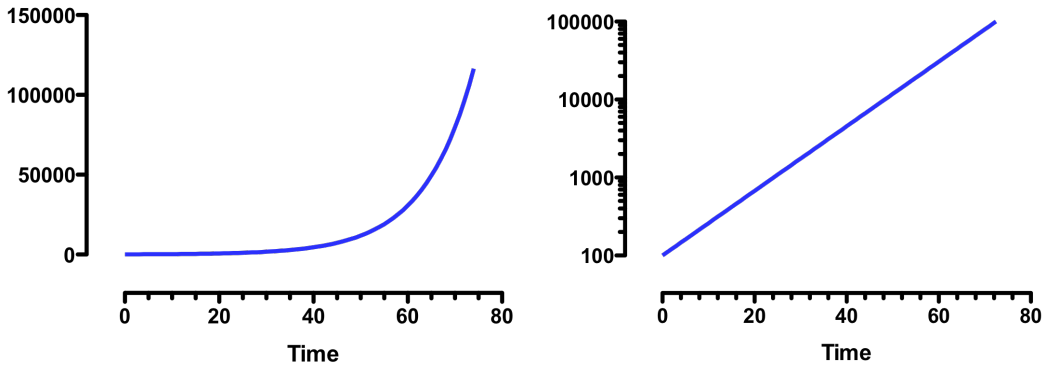


The graph above was created with the outcome (the odds ratio) plotted horizontally. Since these values are something that was determined (not something set by the experimenter), the horizontal axis is, essentially, the Y-axis even though it is horizontal.

A logarithmic axis linearizes compound interest and exponential growth

The graphs below plot exponential growth, which is equivalent to compound interest. At time = 0.0, the Y value equals 100. For each increment of 1.0 on the X

axis, the value plotted on the Y axis equals 1.1 times the prior value. This is the pattern of cell growth (with plenty of space and nutrients), and is also the pattern by which an investment (or debt) grows over time with a constant interest rate.

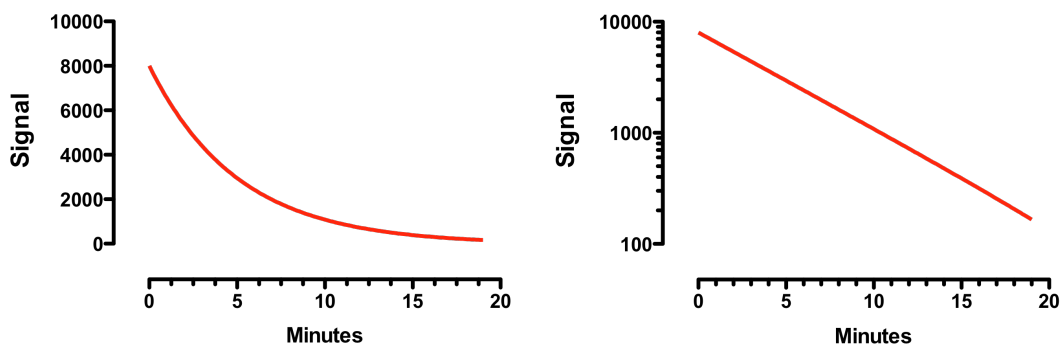


The graph on the right is identical to the one on the left, except that the Y axis has a logarithmic scale. On this scale, exponential growth appears as a straight line.

Exponential growth has a constant doubling time. For this example, the Y value (cell count, value of investment...) doubles for every time increment of 7.2657.

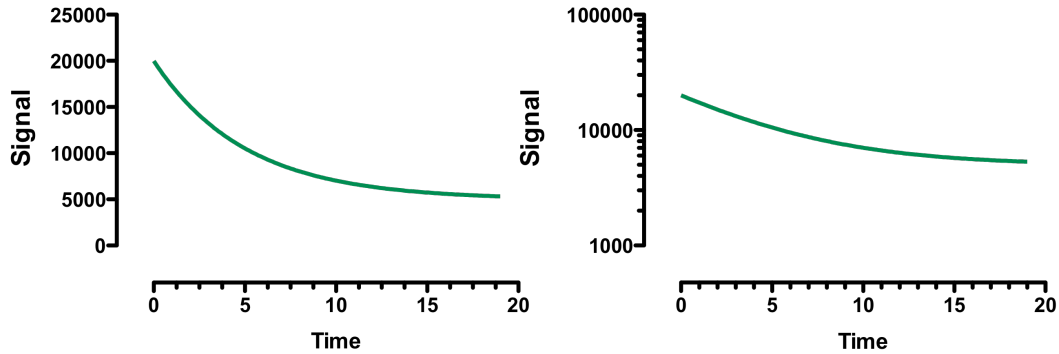
An exponential decay curve is linear on a logarithmic axis, but only when it decays to zero

The graphs below show exponential decay down to a baseline of zero. This could represent radioactive decay, drug metabolism, or dissociation of a drug from a receptor.



The half-life is consistent. For this example, the half-life is 3.5 minutes. That means that at time=3.5 minutes, the signal is half what it was at time zero. At time=7.0 minutes, the signal has cut in half again, to 25% of the original. By time = 10.5 minutes, it has decayed again to half what it was at time=7.5 minutes, which is down to 12.5% of the original value.

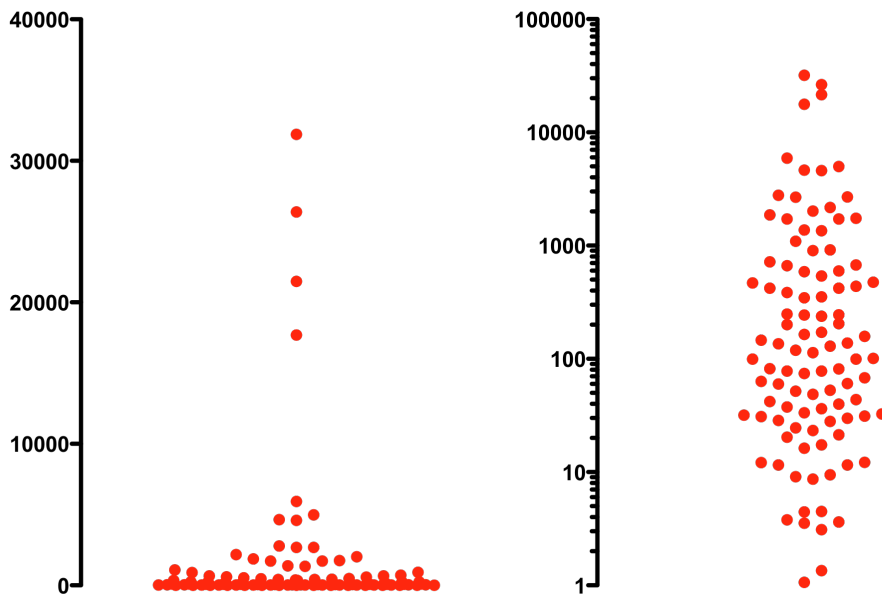
Note that when an exponential dissociation curve plateaus at a value other than zero, it will not be linear on a logarithmic axis. The graph on the right below has a logarithmic Y axis, but the exponential decay (green curve) is not a straight line.



Lognormal distributions

Plotting lognormal distributions on a logarithmic axis

The two graphs below plot the same 50 values. The graph on the left has a linear (ordinary) Y axis, while the graph on the right has a logarithmic scale.



The data are sampled from a *lognormal* distribution, which is very asymmetrical.

When shown on a linear scale (graph on the left), it is impossible to really get a sense of the distribution, since about half of the values are plotted in pile at the bottom of the graph. Another problem is that if you saw only the graph on the left, you might think the highest four values are outliers, since they seem to be so far from the others.

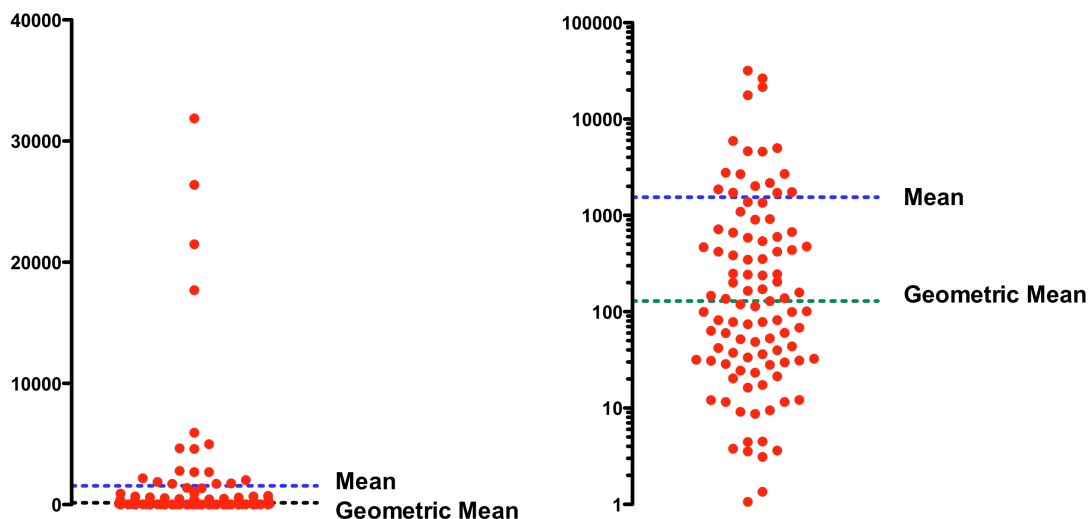
When plotted with a logarithmic axis, the distribution appears symmetrical, the highest points don't seem out of place, and you can see all the points.

These data come from a lognormal distribution. That means that the logarithms of values follow a Gaussian distribution. Plotting such values on a logarithmic plot makes the distribution more symmetrical and easier to understand.

The mean and geometric mean

The graphs below show the same data, and also show the mean and geometric mean. The geometric mean is computed by first transforming all the values to logarithms, finding the mean of those logarithms, and then reverse transforming that mean back to the original units.

With an asymmetrical distribution, the mean and geometric mean can be quite different as shown here. When plotted on a log scale (right), the mean is not near the middle of the data. That is because the points above the mean have much larger values so bring up the mean value. On a logarithmic axis, the geometric mean is near the middle of the distribution.



This graph shows that the concept of an average or mean is somewhat ambiguous when data are plotted on a logarithmic axis. Do you mean the mean of the actual values, or the mean of the logarithms (the geometric mean)? Because of this

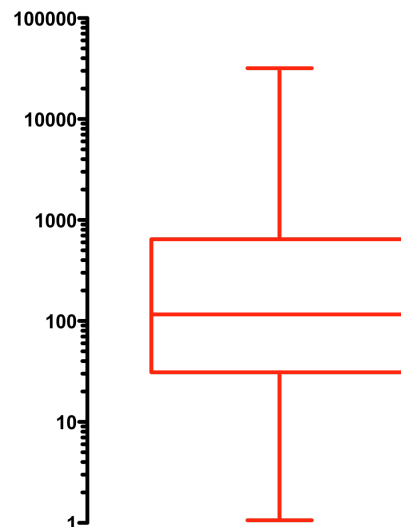
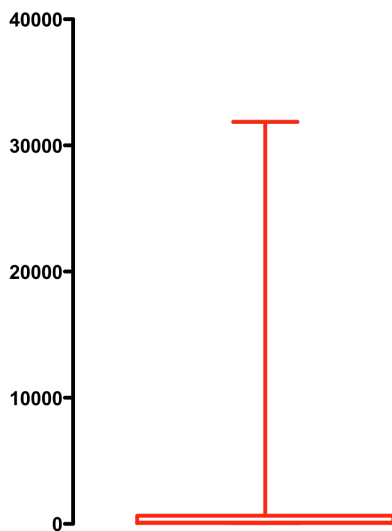
ambiguity, consider showing the median instead. For these data, the mean is 1541.3, the geometric mean is 141.2, and the median is 116.3

Displaying variability on a lognormal distributions

In the example above, the standard deviation (SD) is 4930. The range that covers a mean plus or minus one SD, therefore, extends from $1541 - 4930$ to $1541 + 4930$, which is from -3389 to 6741 . This range cannot be plotted on a logarithmic axis, because such an axis cannot include negative values

One alternative is to compute a *geometric standard deviation*, but these are not often used (much less than geometric means) and are tricky to understand.

For display purposes, if you are not going to graph every individual value as on the graphs above, a box-and-whiskers plot as shown below does a great job of showing variation. The box extends from the 25th to 75th percentiles, with a line at the median (50th percentile). The whiskers, on this graph, go down to the smallest value and up to the largest value (alternative definitions are often used). Again, the graph is easy to understand when a logarithmic axis is used (right) but is not very helpful with a linear axis.

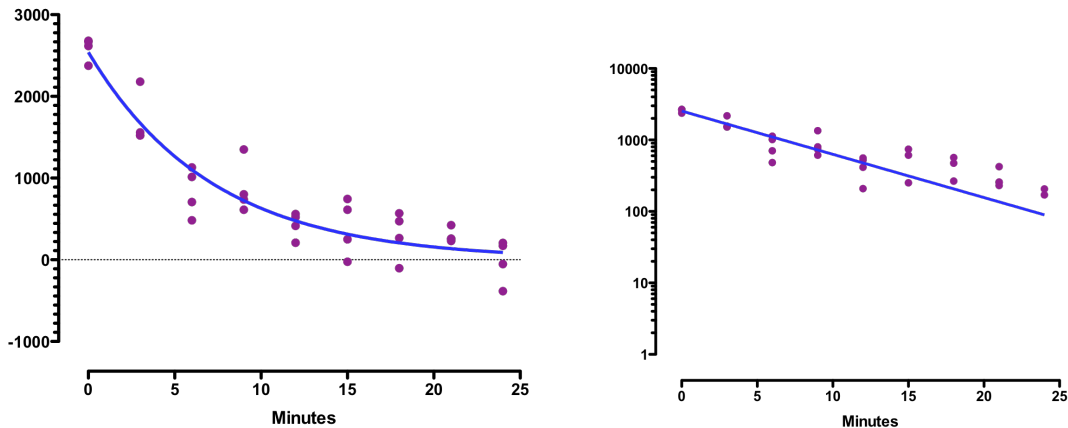


Distinguish using a logarithmic axis from plotting logarithms

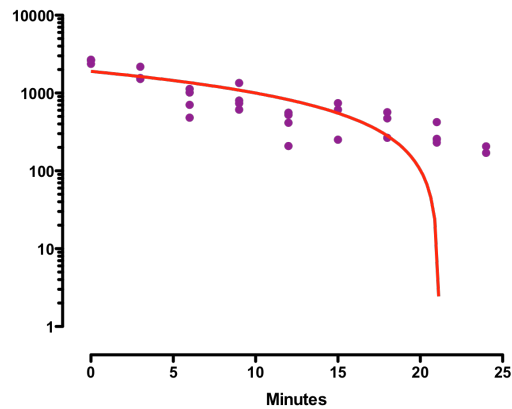
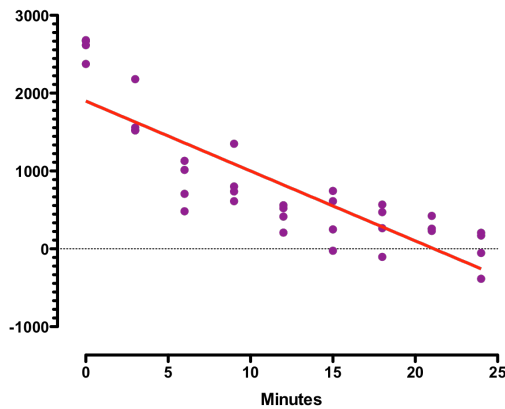
Regression fits the data, not the graph

Be careful when you fit a curve to data with a logarithmic axis. When you fit a curve with nonlinear regression (or a line with linear regression), you fit a model (equation) that defines Y as a function of X . Choosing to stretch either axis to a logarithmic scale does not change any values.

The two graphs below plot the same data. The curve in the left graph was created by using nonlinear regression to fit an exponential decay curve. That curve is also shown on the graph on the right. With a logarithmic axis, an exponential curve decaying to zero looks straight.



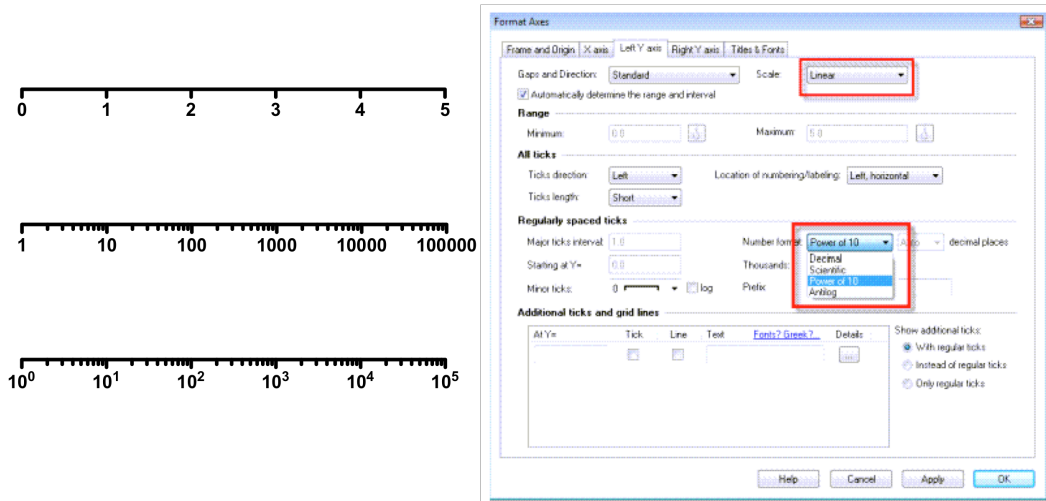
Since the graph on the right looks like a straight line, you might be tempted to fit the data with linear regression, rather than fitting an exponential decay model using nonlinear regression. The graphs below shows what happens when you fit linear regression to those data. On a linear axis (left), the linear regression line indeed looks like a line. On a logarithmic axis in contrast (right), the linear regression line appears curved. Linear regression (and nonlinear regression) fit the data, not the graph. Note that four points have negative Y values, and these are simply omitted from the graph on the right.



Use antilog or powers-of-ten numbering when plotting values that are logs

Prism's built-in dose-response equations all assume that the X values are logarithms of doses or concentrations. Entering the X values as concentrations or doses, and then stretching the axis to a logarithmic scale is not the same at all. That approach does not change the X values.

Imagine that the X values range from 1 to 10,000, and you want to fit a dose-response curve using Prism's built-in equations. Don't use a logarithmic axis. Instead, transform the values to logarithms (either before entering the data, or using Prism's Transform analysis), and plot these (which are logarithms) on a linear axis. After transforming the X values will range from 0 to 5, and the X axis will look like the top example below.



Prism lets you format the axis to show the original values (doses) rather than the transformed values (logarithms of doses). The three axes above all have the same

range (0 to 5). The first axis has decimal numbering, the second axis has antilog numbering, and the third axis uses powers-of-ten numbering. The screen shot shows how to select these on the Format Axis dialog.

Note that the bottom two axes have **not** been stretched to a logarithm scale. Only the numbers have been written in alternative formats. Even though the top of the axis is labeled 10,000 or 10^5 , the corresponding value on the graph actually is $Y=5$. A value of $Y=10,000$ would be way off scale, unless you transformed it first.