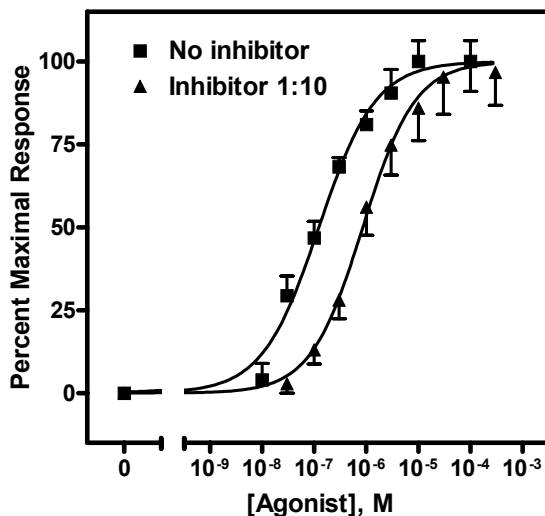


Analyzing Dose-Response Data

A dose-response curve describes the relationship between response to drug treatment and drug dose or concentration. Sensitivity to a drug acting at a specific, saturable receptor typically spans a large concentration range, so dose-response curves are usually semi-logarithmic, i.e., the amount of drug is plotted as the log of drug concentration, giving them their familiar sigmoidal shape.



This example will show you (a) how to use Prism to fit sigmoidal (also known as “logistic”) curves to your dose-response data and (b) one way to compare two dose-response curves statistically. This article includes the following techniques:

- Fitting data to a three- or four-parameter sigmoidal model
- Log transforming data

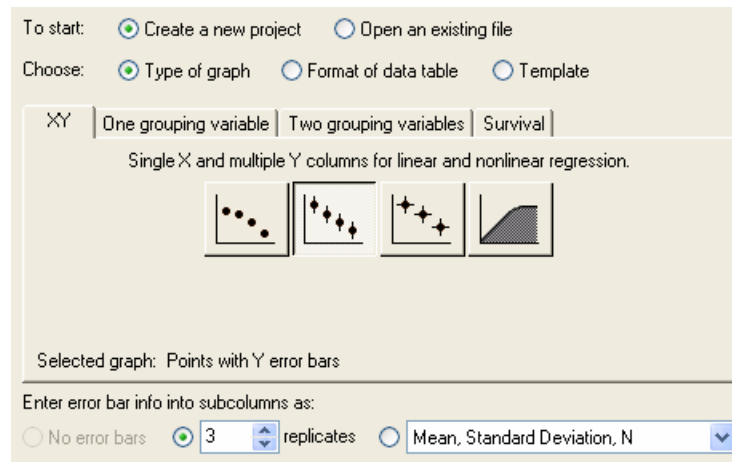
¹ Adapted from: Miller, J.R., *GraphPad Prism Version 4.0 Step-by-Step Examples*, GraphPad Software Inc., San Diego CA, 2003. *Step-by-Step Examples* is one of four manuals included with Prism 4. All are available for download as PDF files at www.graphpad.com. While the directions and figures match the Windows version of Prism 4, all examples can be applied to Apple Macintosh systems with little adaptation. We encourage you to print this article and read it at your computer, trying each step as you go. Before you start, use Prism’s **View** menu to make sure that the Navigator and all optional toolbars are displayed on your computer.

- Normalizing data
- Including a “zero” concentration control value on a log X scale
- Fixing a curve-fit parameter to a constant value
- Comparing two dose-response curves statistically
- Segmented (gapped) axes
- Displaying dose or concentration values on an X axis in different ways
- Displaying curve-fit parameters on the graph

Suppose we want to plot two dose-response curves, showing the average in vitro response to an agonist, in the presence and in the absence of an antagonist, measured in triplicate. Further, we wish to include on the graph a zero-concentration control value.

Entering Dose-Response Data

In the Welcome dialog, make the following settings:



Enter the data as shown below. The X values (agonist dose) are given as molar concentration, expressed in exponential notation. The Y values (response, in arbitrary units) go in subcolumns Y1 to Y3.

	X Values	A			B		
	Agonist	No inhibitor			Inhibitor 1:10		
	X	A:Y1	A:Y2	A:Y3	B:Y1	B:Y2	B:Y3
1	1.00e-010	0.0	0.0	0.0	0.0	0.0	0.0
2	1.00e-008	1.0	0.0	4.0			
3	3.00e-008	12.0	10.0	15.0	1.0	0.0	2.0
4	1.00e-007	19.0	18.0	22.0	5.0	3.0	6.0
5	3.00e-007	28.0	30.0	28.0	8.0	10.0	12.0
6	1.00e-006	32.0	35.0	35.0	17.0	20.0	23.0
7	3.00e-006	35.0	41.0	38.0	23.0	28.0	29.0
8	1.00e-005	39.0	43.0	44.0	27.0	31.0	34.0
9	3.00e-005				30.0	34.0	38.0
10	1.00e-004	39.0	43.0	44.0	32.0	37.0	38.0
11	3.00e-004				32.0	37.0	

Several points about this data are worth noting:

- The dosage increment (10 nM, 30 nM, 100 nM, etc) shown in the X column is typical. These values will be log-transformed, giving roughly equal horizontal spacing of data points on the dose-response curve.
- These data are a good illustration of Prism's use of only one X column. Although the X values for data sets A and B are not the same, we can combine all the X values into a single column and leave Y-value cells empty where appropriate.
- The values in row 1 represent the zero-concentration (no agonist) data. Since our X coordinates will be plotted as logarithms, and since the log of 0 is undefined, we approximate this point with an X coordinate of 1.0e-010, about 2 log units below the lowest "real" X value. Later on, we'll show how to place a zero on the X axis for this data point.
- For this example, we're assuming that the Y values do not require correction by subtraction of a baseline—the minimum Y values are already zeroes. If you wish to subtract a baseline value from each Y value in a data set, read about automatic baseline corrections in the Step-by-Step Example "Substrate-Velocity Curves and Lineweaver-Burk Plots". Later in this example, we'll demonstrate normalization of the Y values, which can also serve as a means of "baseline correction".
- You can enter the Y values in any convenient units, then transform to other units if you wish. We'll use Prism to normalize (i.e., transform to percent of maximal response) our Y values later.
- Particularly when dosing with an impure preparation, you may wish to enter X values in terms of "dilution factor" or in mass/volume (e.g., mg/mL) units. Using mass/volume units is not a problem.

Data Manipulations

This section covers the data modification necessary to produce a conventional dose-response graph, with response values normalized to span the range from 0% to 100% plotted against the logarithm of agonist concentration

Log Transforming the X Values

To convert the values in the X column to their respective logarithms, click **Analyze**. Choose **Built-in analysis**. From the **Data manipulations** category, select **Transforms**.

In the **Parameters: Transforms** dialog, choose to **Transform X values using X=Log(X)**.

At the bottom of this dialog, there is an option to **Create a new graph of the results**. But we're not ready to do that; we have one more manipulation, namely normalization, to do before making the final graph.

Prism displays a new Results sheet (partly illustrated as follows) with the log-transformed X values.

	X Values	A			B		
	Agonist	No inhibitor			Inhibitor 1:10		
	X	A:Y1	A:Y2	A:Y3	B:Y1	B:Y2	B:Y3
1	-10.000	0.000	0.000	0.000	0.000	0.000	0.000
2	-8.000	1.000	0.000	4.000			
3	-7.523	12.000	10.000	15.000	1.000	0.000	2.000
4	-7.000	19.000	18.000	22.000	5.000	3.000	6.000
5	-6.523	28.000	30.000	28.000	8.000	10.000	12.000

In a system following mass-action kinetics, the dose-response curve becomes sigmoidal only when response is plotted against the *logarithm* of dose, and Prism's built-in sigmoidal dose-response curve-fit models presuppose that your X values are in that form. But users occasionally find that their dose-response relationships are sigmoidal before log transformation. In that case, consider using the user-defined model

$$Y = \text{BOTTOM} + (\text{TOP} - \text{BOTTOM}) * X^H / (\text{EC}_{50}^H + X^H)$$

where EC₅₀ is the midpoint of the curve, and H is the Hill slope. For an example of implementing user-defined equations, see the step-by-step example “Fitting Data to User-Defined Equations”.

Normalizing the Y Values

In this section, we’ll convert the Y values in data sets A and B to a common scale. This is useful when you want to compare the shape or position (EC_{50}) of two or more curves and don’t want to be distracted by different maximum and minimum values.

With the log-transformed Results sheet in view, click **Analyze**.

In the **Analyze Data** dialog, select the **Data manipulations** category, then choose **Normalize**.

In the **Parameters: Normalize** dialog, specify how Prism will normalize the Y values. For this exercise, we’ll accept the default settings, which will define 0% and 100% as the smallest and the highest values, respectively, for each of the data sets.

Normalizing the data in this way converts it to “Percent of maximum response”. Note two points: (a) Normalization will cause the dose-response curve to extend vertically from 0 to 100 by definition, and we’ll constrain the curve fit later to ensure this. Therefore, you should be fairly confident of the accuracy of your minimum and maximum values. (b) In this example, the minimum Y values were already zeroes, so no alteration of the minimum values will be made. Note, however, that had these values been greater than 0, the normalization would effectively have produced a “baseline correction”.

Choose to present the results as **Percentages**. And since the results of this “analysis” are the values that we will finally graph, check the box to **Create a new graph of the results**. *Overlooking this checkbox is a very common mistake.* The settings are shown here:

Parameters: Normalize

How is 0% defined?

Smallest value in each data set

First value in each data set (or last, whichever is smaller)

Remove from the results page

Y = 0.0 becomes 0% for all data sets

How is 100% defined?

Largest value in each data set

Last value in each data set (or first, whichever is larger)

Remove from the results page

Y = becomes 100% for all data sets

The sum of all values in the data set (column)

Present results as

Fractions Percentages

New graph

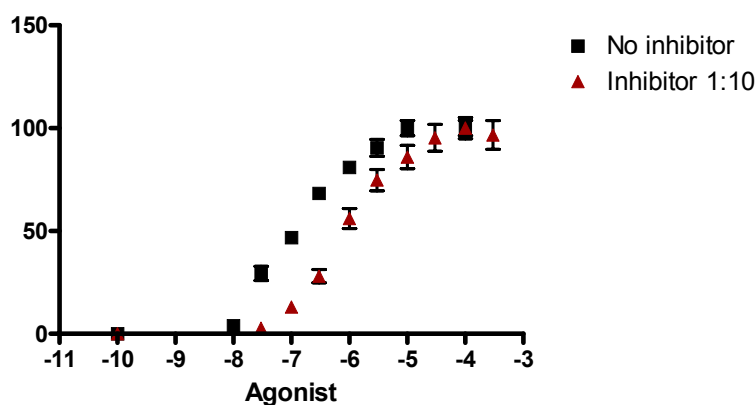
Create a new graph of the results

Help Me Decide Cancel OK

The Results sheet is displayed. Because we have replicate Y values, 0% and 100% are defined by the mean of the replicates. It is neither possible nor desirable to normalize each subcolumn separately. This is why none of the upper Y values in the table is exactly 100.

X Values	A			B		
	No inhibitor			Inhibitor 1:10		
	A:Y1	A:Y2	A:Y3	B:Y1	B:Y2	B:Y3
-10.000	0.000	0.000	0.000	0.000	0.000	0.000
-8.000	2.381	0.000	9.524			
-7.523	28.571	23.810	35.714	2.804	0.000	5.607
-7.000	45.238	42.857	52.381	14.019	8.411	16.822
-6.523	66.667	71.429	66.667	22.430	28.037	33.645
-6.000	76.190	83.333	83.333	47.664	56.075	64.486
-5.523	83.333	97.619	90.476	64.486	78.505	81.308
-5.000	92.857	102.381	104.762	75.701	86.916	95.327
-4.523				84.112	95.327	106.542
-4.000	92.857	102.381	104.762	89.720	103.738	106.542
-3.523				89.720	103.738	

Now click the yellow **Graphs** tab on the toolbar to switch to the graph of the log-transformed and normalized data.



Note that, unless you specify otherwise, Prism plots the averages of replicate data and adds error bars automatically.

Fitting a Sigmoidal (Four-Parameter Logistic) Curve

With the graph displayed, click on the **Analyze** button. From the **Curves & regression** category, select **Nonlinear regression (curve fit)**.

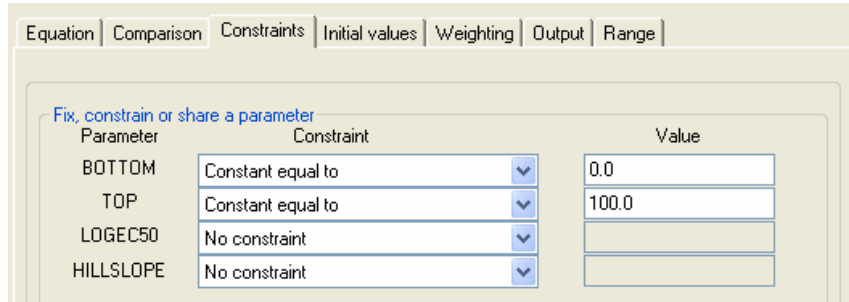
In the **Parameters: Nonlinear Regression (Curve Fit)** dialog box, choose **Classic equations**. Select **Sigmoidal dose-response (variable slope)**.

Which sigmoidal dose-response curve-fit model should you use? When you choose **Sigmoidal dose-response**, Prism assumes that the curve has “standard” steepness (a Hill slope of +/- 1) and fits only the Top, Bottom, and log EC_{50} . Other programs may refer to this model as “three-parameter logistic”. If you suspect that your curve has more or less than normal steepness, choose **Sigmoidal dose-response (variable slope)**, and Prism will also fit the Hill slope. Other programs refer to the variable-slope model as “four-parameter logistic”.

This is normally all you need do to prepare Prism to fit the curves, and if you were to click **OK** at this point, Prism would perform a four-parameter (bottom, top, log EC_{50} , and Hill slope) sigmoidal curve fit and add the curves to the graph. But we'll choose two options.

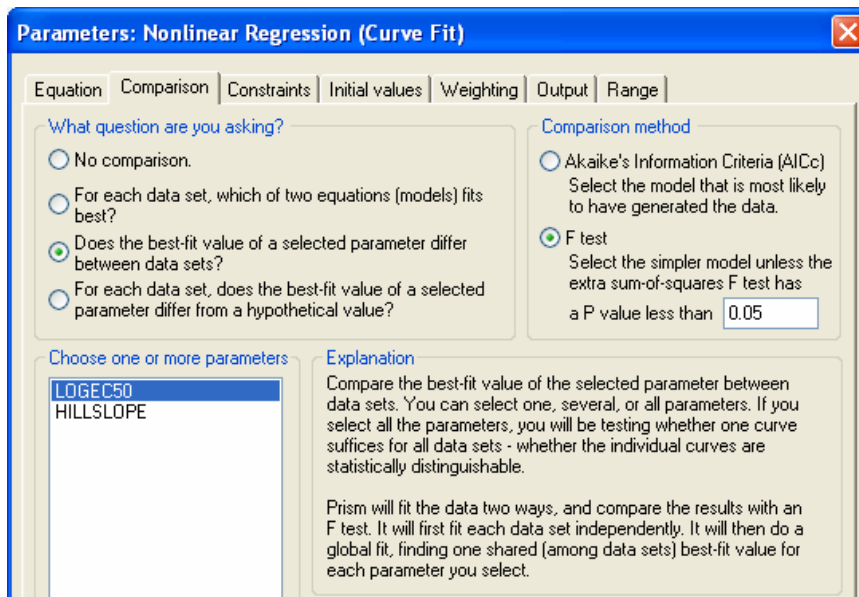
Constraining Curve-Fit Parameters

Since we normalized the original data such that the vertical range extends by definition from 0 to 100, it doesn't make sense to fit the "bottom" and the "top" of the curves. So we'll fix those parameters, leaving only the midpoint (log EC_{50}) and slope (Hill slope) of each curve to be fitted by Prism. Select the **Constraints** tab. Constrain the parameters BOTTOM and TOP to 0 and 100, respectively:



Comparing Curves Statistically

Instruct Prism to compare the fitted midpoints (log EC_{50}) of the two curves statistically. Click the **Comparison** tab and make the settings shown in the following figure:



Only two of the four parameters from the variable-slope sigmoidal model show up in the list box, since the other two parameters were constrained—therefore no longer to be fitted—in the previous section.

While you're looking at the **Comparison** page of the **Parameters: Nonlinear Regression (Curve Fit)** dialog, note the variety of comparisons available.

- **For each data set, which of two equations (models) fits best?**

Select different curve-fit equations and allow Prism to decide which better fits your data.

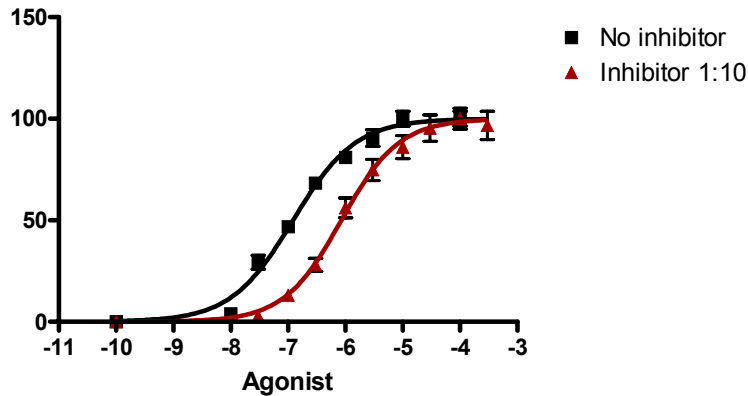
- **Does the best-fit value of a selected parameter differ between data sets?**

Compare data sets. Select (a) one or more curve-fit parameters and allow Prism to decide whether their values differ statistically among data sets, i.e., whether the curves statistically different with respect to the specified parameter(s). That's what we're doing in this example, using $\log EC_{50}$ as our criterion for comparison. Or select (b) *all* parameters and have Prism decide on the basis of the overall curve whether the data sets differ.

- **For each data set, does the best-fit value of a selected parameter differ from a hypothetical value?**

Experiment with these settings, observing how Prism changes the dialog and reading the **Explanation** frame for each option.

When you click **OK** to exit the **Parameters: Nonlinear Regression (Curve Fit)** dialog, Prism fits the curves and places them on the graph.



The Results sheet is extensive; we'll show the highlights. Select the **Table of results** subsheet in the Navigator.

Prism summarizes the result of the *F* test giving the hypotheses tested, in the first rows of the table (in the illustration, we juxtaposed the first and fourth columns by horizontal scrolling).

	Global (shared)
	Y
Comparison of Fits	
Null hypothesis	LOGEC50 same for all data sets
Alternative hypothesis	LOGEC50 different for each data set
P value	P<0.0001
Conclusion (alpha = 0.05)	Reject null hypothesis
Preferred model	LOGEC50 different for each data set
F (DFn, DFd)	163.1 (1,52)

Further down, Prism reports the values fitted under the assumptions of each hypothesis. When Prism is free to fit $\log EC_{50}$ for each data set separately (the alternate hypothesis)...

	No inhibitor	Inhibitor 1:10
	Y	Y
LOGEC50 different for each data set		
Best-fit values		
BOTTOM	0.0	0.0
TOP	100.0	100.0
LOGEC50	-6.912	-6.069
HILLSLOPE	0.7978	0.8652
EC50	1.2244e-007	8.5320e-007
Std. Error		
LOGEC50	0.04262	0.04777
HILLSLOPE	0.05848	0.07374

When Prism fits a common value for $\log EC_{50}$ to both data sets...

	No inhibitor	Inhibitor 1:10	Global (shared)
	Y	Y	Y
LOGEC50 same for all data sets			
Best-fit values			
BOTTOM	0.0	0.0	
TOP	100.0	100.0	
LOGEC50	-6.466	-6.466	-6.466
HILLSLOPE	0.6947	0.7618	
EC50	3.4198e-007	3.4198e-007	3.4198e-007
Std. Error			
LOGEC50	0.06972	0.06972	0.06972
HILLSLOPE	0.07409	0.08668	

In both cases, the fixed values of 0 and 100 for BOTTOM and TOP are included in the table; since they weren't fitted, no standard errors are shown for them. The EC_{50} shown in the table is simply the antilog of the parameter that is fitted—LOGEC50.

As we noted earlier, the null hypothesis was rejected on the basis of the F test. It appears that the experimental intervention, addition of a putative inhibitor, shifts the dose-response curve to the right.

Users frequently object to the fact that the standard error for EC_{50} is not reported. These standard error values express the uncertainty inherent in estimating the parameters by curve fitting. For the equation fitted here, the midpoint of the curve is $\log EC_{50}$, not EC_{50} . Don't try to express the standard error for EC_{50} obtained in this way using a single number—in particular, it is not the antilog of the standard error of $\log EC_{50}$. We suggest that you use the 95-percent confidence interval for EC_{50} , which Prism does report.

95% Confidence Intervals		
LOGEC50	-7.000 to -6.824	-6.167 to -5.971
HILLSLOPE	0.6773 to 0.9182	0.7139 to 1.017
EC50	1.0002e-007 to 1.4987e-007	6.8080e-007 to 1.0692e-006

Additional help for understanding your results is available when you click the Analysis Check List button.



Adjustments to the Y Axis

Click on the yellow **Graphs** tab to return to the graph. If you'd like a frame around your graph, double-click one of the axes to open the **Format Axes** dialog. Click the **General** tab, then choose **Plain Frame** in the **Frame & axes** drop-down box. We can also increase readability by adjusting the range and tick intervals for the Y axis. Select the **Left Y axis** tab. Deselect the Auto option under **Range**. Set the **Range** and **Tick options** as follows:

Range		Tick options	
<input type="checkbox"/> Auto		Major ticks: Interval: 25.0	Starting at: 0.0
Minimum: -15.0		Minor ticks: # intervals: 0	Spacing: Equal
Maximum: 115.0		All ticks: Direction: Left	Length: Short

Select the **General** tab again, then set **Origin** to **Lower Left**. This lowers the X axis to coincide with the frame around the graph.

Adjustments to the X Axis

Now we'll format the X axis, introducing a break, or gap, in the axis; labeling our zero-agonist control point; and expressing the X-axis values in different units.

Segmenting the X Axis

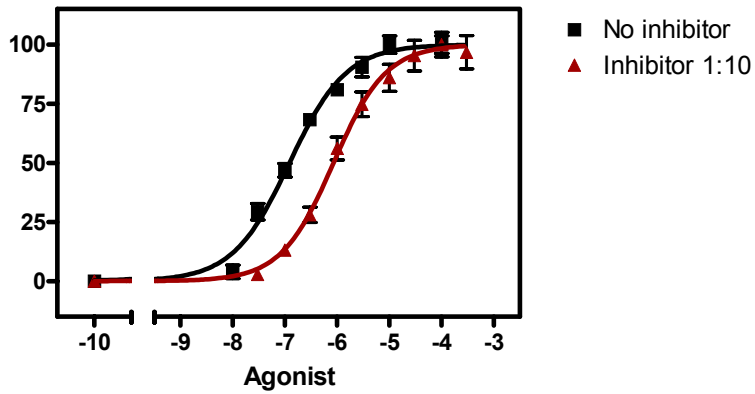
Double-click on the X axis to reopen the **Format Axes** dialog box. Choose the **X Axis** tab. Carefully set the **Appearance** and **Segment** options as follows:

Appearance	
Gaps and Direction: Two segments (---/ /---)	Scale: Linear
Segment: Left	
Range: Minimum: -10.5, Maximum: -9.5, Length: 16 % of axis	Tick options: <input type="checkbox"/> Plot only one tick on this segment, Major ticks: Interval: 1, Starting at: -10.0, Minor ticks: # intervals: 0, Spacing: Equal, All ticks: Direction: Down, Length: Short

Now reset **Segment** to **Right** and enter or verify the following settings:

Segment: Right	
Range: Minimum: -9.5, Maximum: -2.5, Length: 84 % of axis	Tick options: <input type="checkbox"/> Plot only one tick on this segment, Major ticks: Interval: 1.0, Starting at: -9.0, Minor ticks: # intervals: 0, Spacing: Equal, All ticks: Direction: Down, Length: Short

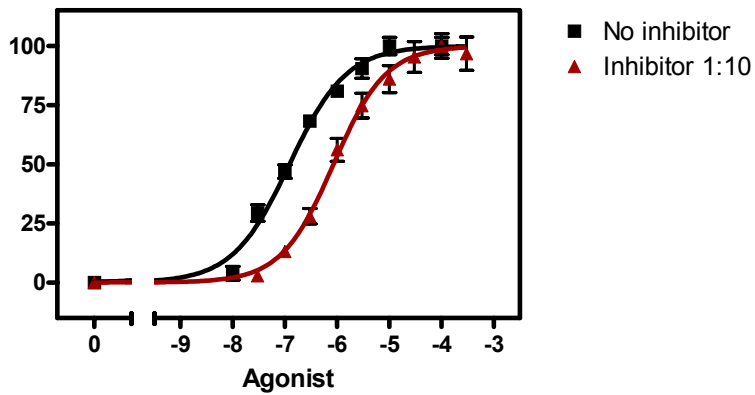
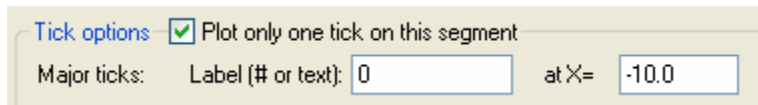
Your graph should now look like this:



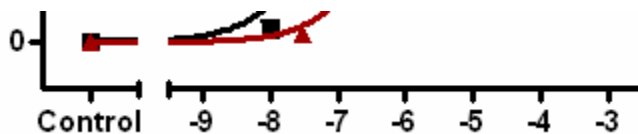
You may have to resort to some trial and error to get the look you want. One hint: Try selecting, then “dragging”, the ends of the X axis segments directly on the graph, which will adjust the value for **Length... % of total axis** automatically.

Labeling the Zero-Concentration Control Point

Now we’ll substitute the value “0” for “-10” on the left axis segment. Double-click on the left-hand segment of the X axis to reopen the **Format Axes** dialog. Verify that the **X axis** tab is selected and that **Segment** is set to **Left**. Set the checkbox to **Plot only one tick on this segment** and direct Prism to insert the label “0” at X = -10.



Note that you don’t have to use a numeric tick label where we substituted for “-10”. You could just as well label the tick “Control” or “Vehicle”.

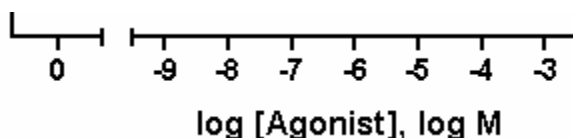


Displaying Doses or Concentrations in Different Ways

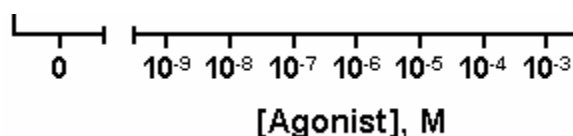
To change how Prism expresses the values on the right-hand segment of the X axis, return to the **Format Axes** dialog (**X axis** tab). Set the axis **Segment** to **Right**.

Now change the setting in the **Format** drop-down list under **Numbering or labeling**. Only two of the four choices produce usable results in this example:

Decimal:

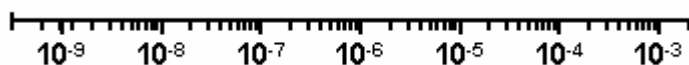


or **Power of 10:**



If you choose the **Power of 10** tick labeling option, you may also wish to add logarithmic minor ticks. In the **Format Axes** dialog (**X axis** tab), adjust **Minor ticks** and **Spacing** as shown here:

Minor ticks: # intervals: 9 Spacing: Log



Prism doesn't alter the axis *titles* to reflect the tick-labeling format—we did that manually for this illustration.

When you want to use one of the other (longer) formats, such as **Antilog**, you may find that the labels are long and overlapping each other. Try one of the following:

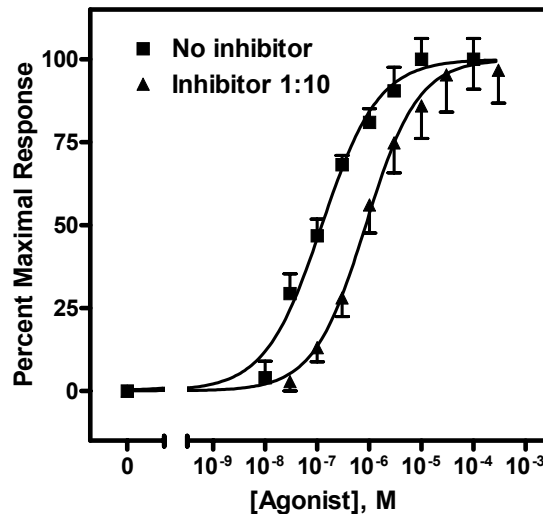
- If the crowding isn't too severe, create more room for the labels by clicking on the X axis to select it and then "dragging" the right-hand "handle" to the right.
- Enter your original data onto the data table in different units. You may be able to eliminate several place-holding zeros for tick labels in the Antilog format by entering data in micromolar or nanomolar units. Remember to modify the X-axis title if necessary, and remember that $\log EC_{50}$ and EC_{50} values reported on the curve-fit results will be in the same units that you used on the data table.
- Show the tick labels at an angle. In the **Format Axes** dialog (**X axis** tab), try these settings:

Numbering or labeling
Location: Below, angled Angled: 45

Displaying Error Bars in Only One Direction

If you'd like to prevent overlapping of closely spaced error bars, double-click on one of the symbols for the “No inhibitor” data to open the **Format Symbols and Lines** dialog. Adjust the **Error bars...Dir.** to show only the top half (**Above**) of the bar. And for this example, switch to standard deviation error bars: **Plot... Mean & SD.**

Repeat for the “Inhibitor 1:10” data, choosing this time the lower half of the bar. Your graph should now look something like this (we edited the axis titles and relocated the legends):



You wish to display some error bars associated with a data set in one direction only, and the rest in both directions. On the data table, choose the **Mean, +Error, -Error** format for the Y columns. Prism then leaves room for you to add differing positive and negative error limits. Enter zeroes on the table wherever you don't want that half of the error bar to be displayed.

Displaying Curve-Fit Parameters on the Graph

We'll now transfer some information about the best-fit parameters from the Results page directly to the graph. The simplest technique is to return to the nonlinear regression Results sheet—remember to switch back to the **Table of results** subpage—and drag your cursor to select the results of interest:

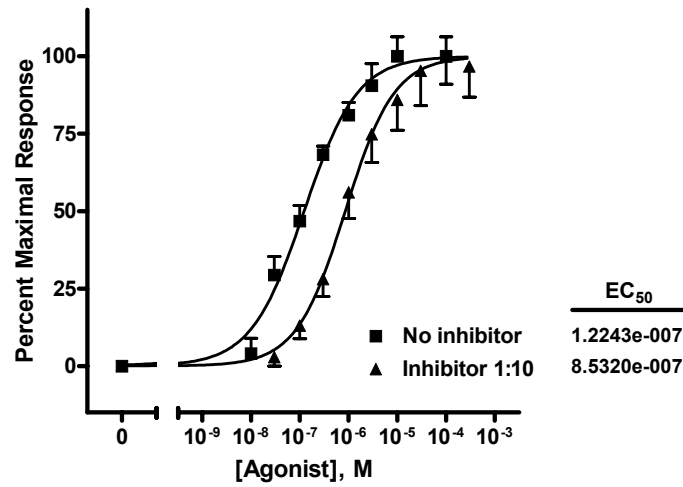
EC50	1.2240e-007	8.5320e-007
------	-------------	-------------

Choose **Edit...Copy**, switch to the graph, and choose **Edit...Paste Table**. You can now select the embedded table and drag it to a convenient open spot (in the illustration below, we removed the frame around the plot area).

With some ingenuity, you can improve the cosmetics and narrow the information down to what's important. In the illustration below, the EC_{50} values were pasted individually and situated next to the legend elements. The heading “ EC_{50} ” was generated using the text tool,

T

and the line below “ EC_{50} ” was made with the line-drawing tool.



The beauty of pasting results from the Results sheet to your graph is that, since all related steps in your project are linked, changing your data will result in all corrections being made on the graph automatically—including the pasted values for EC_{50} . Try it! This saves effort not only when you need to correct data errors, but also when you need to construct a new dose-response curve. In fact, once you've settled on the design of your dose-response curve, consider saving your entire project, for use with new data, as a template.