

Least Squares Fitting: Physics 238

In Chapter 8 of *An Introduction to Error Analysis*, by John R. Taylor, the author discusses the general theme of least-squares fitting. This is based on the normal distribution discussion in Chapter 5. Here, we will take the theoretical background as given, and show how the minimization of the least square difference leads to the standard results for some simple cases.

```
In[85]:= Clear["Global`*"]
```

Fitting a Straight Line

First, here is some sample data of {mass, angle} pairs from the torsional oscillator experiment. (See the separate importing notebook for more about entering and importing data.)

Data

```
In[86]:= SetDirectory[NotebookDirectory[]]
```

```
Out[86]:= /Users/doughera/notes/www/public_html/courses/phys238-2026/notes/Ch08
```

```
In[87]:= rawdata = Import["T0-static-20250212.csv", "CSV"]
```

```
Out[87]:= {{Mass (g), Raw Angle, Voltage}, {0, 2.98, -0.0456}, {50, 2.79, -0.4083},  
{100, 2.59, -0.7733}, {150, 2.38, -1.172}, {200, 2.17, -1.556},  
{250, 1.96, -1.958}, {300, 1.66, -2.483}, {350, 1.41, -2.726},  
{400, 1.21, -2.566}, {0, 2.98, -0.0606}, {-50, 3.2, 0.3049}, {-100, 3.4, 0.6583},  
{-150, 3.6, 1.028}, {-200, 3.81, 1.387}, {-250, 4.06, 1.812},  
{-300, 4.32, 2.237}, {-350, 4.56, 2.502}, {-400, 4.74, 2.369}, {0, 3., -0.04502}}
```

```
In[88]:= rawdata[[Range[3]]] // TableForm
```

```
Out[88]//TableForm=  
Mass (g)   Raw Angle   Voltage  
0          2.98       -0.0456  
50         2.79       -0.4083
```

```
In[89]:= data3 = Select[rawdata, VectorQ[#, NumberQ] &]
(* Select only those lines for which all elements are numbers. *)
```

```
Out[89]=
{{0, 2.98, -0.0456}, {50, 2.79, -0.4083}, {100, 2.59, -0.7733}, {150, 2.38, -1.172},
 {200, 2.17, -1.556}, {250, 1.96, -1.958}, {300, 1.66, -2.483},
 {350, 1.41, -2.726}, {400, 1.21, -2.566}, {0, 2.98, -0.0606}, {-50, 3.2, 0.3049},
 {-100, 3.4, 0.6583}, {-150, 3.6, 1.028}, {-200, 3.81, 1.387}, {-250, 4.06, 1.812},
 {-300, 4.32, 2.237}, {-350, 4.56, 2.502}, {-400, 4.74, 2.369}, {0, 3., -0.04502}}
```

For this data, we want columns 1 and 2.

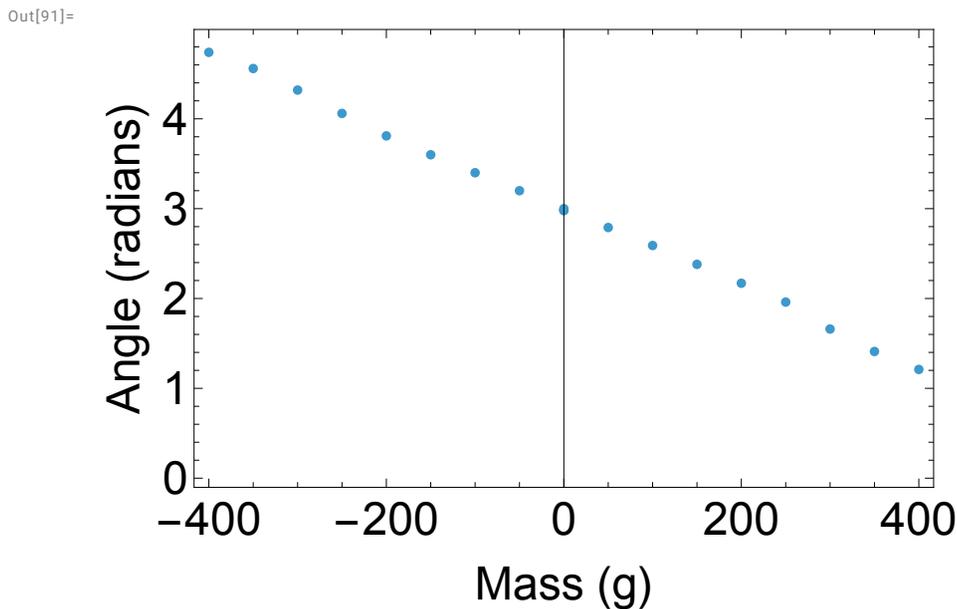
```
In[90]:= data = data3[[All, {1, 2}]]
```

```
Out[90]=
{{0, 2.98}, {50, 2.79}, {100, 2.59}, {150, 2.38},
 {200, 2.17}, {250, 1.96}, {300, 1.66}, {350, 1.41}, {400, 1.21},
 {0, 2.98}, {-50, 3.2}, {-100, 3.4}, {-150, 3.6}, {-200, 3.81},
 {-250, 4.06}, {-300, 4.32}, {-350, 4.56}, {-400, 4.74}, {0, 3.}}
```

Plot the data.

Be sure to give your axes meaningful labels

```
In[91]:= dataPlot = ListPlot[data, ImageSize -> Scaled[0.7], Frame -> True,
 LabelStyle -> Large, FrameLabel -> {"Mass (g)", "Angle (radians)"}]
```



Theoretical Curve: A straight-line fit.

Here we explore how to find the best-fit straight line by the method of least squares. First, define the target function. Use the unknown items (intercept a_0 , and slope a_1) as parameters to the function.

```
In[92]:= yfit[x_, a0_, a1_] := a0 + a1 x
```

Calculating χ^2

Define a χ^2 function for the linear fit. For given a_0 and a_1 values, compute the average of the differences squared between the data and the fit value. Within the function, the data is in $\{m, \theta\}$ pairs, so we pull out the mass of the i^{th} value with `data[[i, 1]]`, and the angle for the i^{th} value with `data[[i, 2]]`. You can use the Classroom Assistant Palette to format the sum, or you can use *Mathematica's* `Sum[]` function directly. Use whichever is easier for you to read. The denominator has the '-2' because with N data points and 2 free parameters, there are only $N-2$ degrees of freedom.

```
In[93]:= calculateChisq[data_, a0_, a1_] :=
```

$$\frac{1}{\text{Length}[\text{data}] - 2} \sum_{i=1}^{\text{Length}[\text{data}]} (\text{yfit}[\text{data}[[i, 1]], a_0, a_1] - \text{data}[[i, 2]])^2$$

(* Fancy typeset presentation *)

```
In[94]:= calculateChisq[data_, a0_, a1_] :=
```

$$\frac{\text{Sum}[(\text{yfit}[\text{data}[[i, 1]], a_0, a_1] - \text{data}[[i, 2]])^2, \{i, 1, \text{Length}[\text{data}]\}]}{(\text{Length}[\text{data}] - 2)}$$

(* Simple in-line presentation*)

Initial Explorations

This command builds an interactive window showing the data, the current fit, and the chi squared value (as the plot title). It draws vertical lines from each data point to the fit line. Move the a_0 and a_1 sliders to minimize χ^2 . *Mathematica's* "Filling" option gets confused if the data isn't sorted, so let's go ahead and sort the data by the first entry in each line.

```
In[95]:= data = SortBy[data, First]
```

```
Out[95]=
```

```
{{-400, 4.74}, {-350, 4.56}, {-300, 4.32}, {-250, 4.06}, {-200, 3.81}, {-150, 3.6},
{-100, 3.4}, {-50, 3.2}, {0, 2.98}, {0, 2.98}, {0, 3.}, {50, 2.79}, {100, 2.59},
{150, 2.38}, {200, 2.17}, {250, 1.96}, {300, 1.66}, {350, 1.41}, {400, 1.21}}
```

```
In[96]:= LinearModelFit[data, {1, x}, x]
```

```
Out[96]=
```

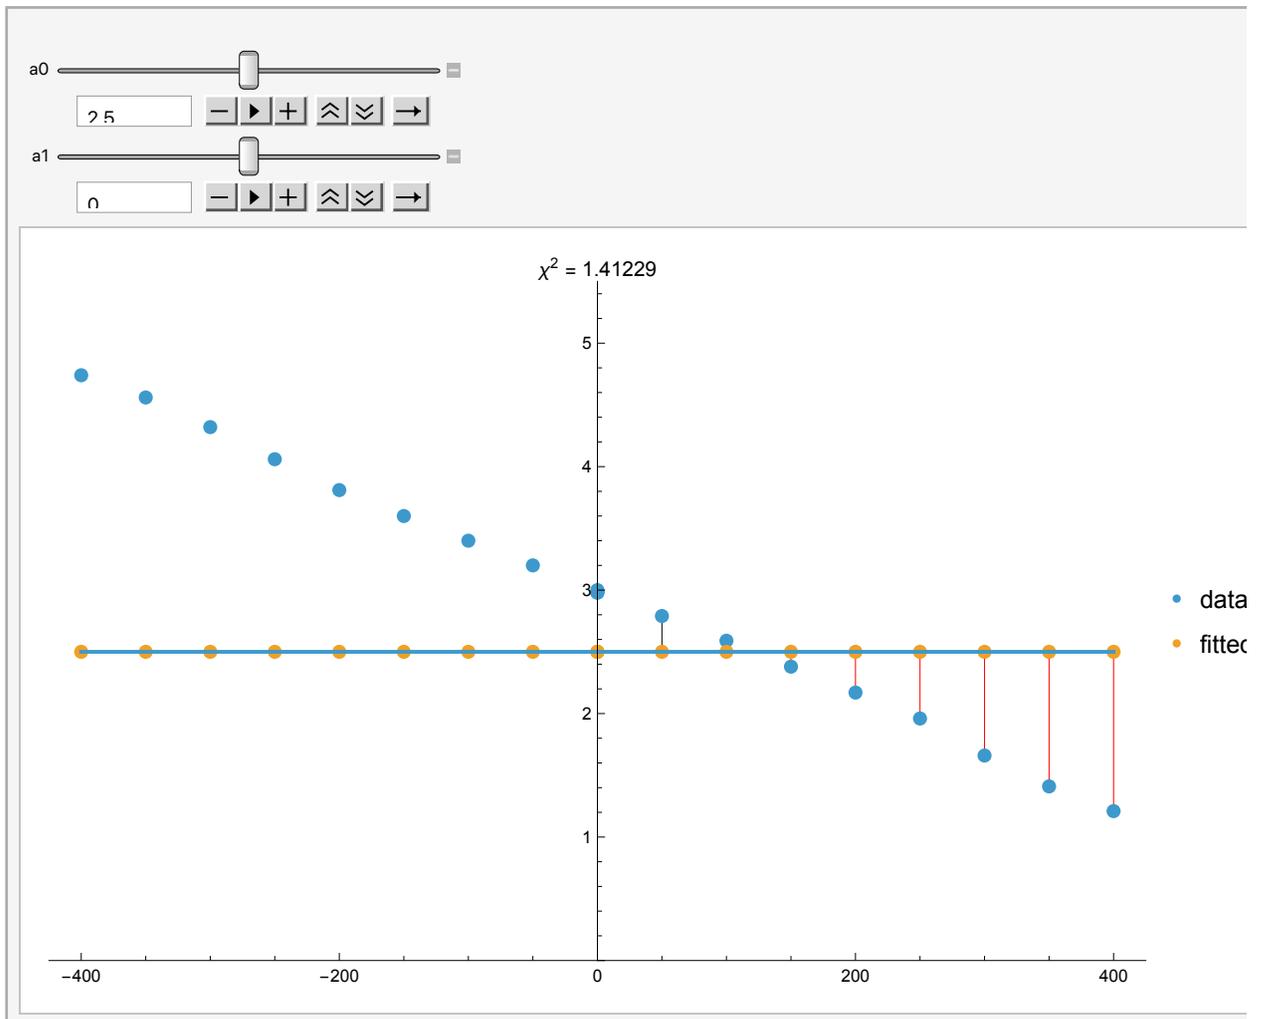
```
FittedModel[ 2.99 - 0.00436 x ]
```

```

In[97]:= Manipulate[
  fitted = Table[{data[[i, 1]], yfit[data[[i, 1]], a0, a1]}, {i, 1, Length[data]}];
  Show[{ListPlot[{data, fitted}, Filling -> {1 -> {{2}, {Red, Black}}},
    PlotRange -> {{-425, 425}, {0, 5.5}}, PlotLegends -> {"data", "fitted"}],
    Plot[yfit[x, a0, a1], {x, -400, 400}], PlotLabel ->
    StringForm[" $\chi^2 =$ `, calculateChisq[data, a0, a1], ImageSize -> Large],
    {{a0, 2.5}, 1, 4, 0.1, Appearance -> "Open"},
    {{a1, 0}, -0.01, 0.01, 0.001, Appearance -> "Open"}
]

```

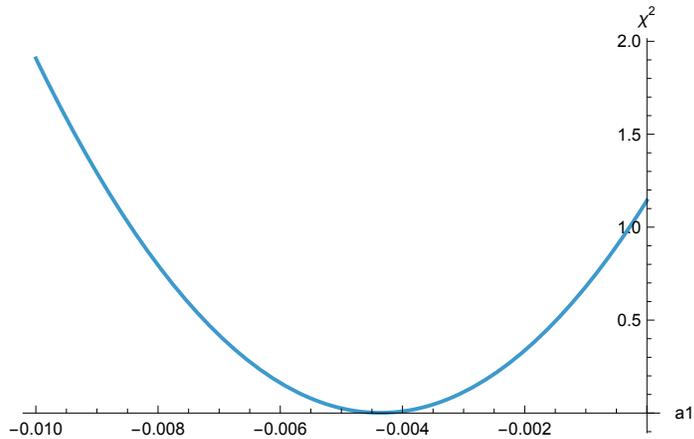
Out[97]=



Minimizing χ^2

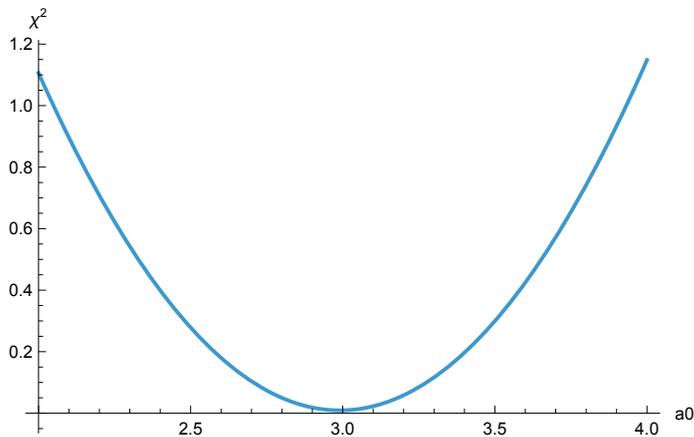
Picking the intercept $a_0 = 3$ (close to the expected value), look at how χ^2 varies with a_1 . The best value for a_1 seems to be a little less than -0.004.

```
In[98]:= Plot[calculateChisq[data, 3, a1], {a1, -0.01, 0}, AxesLabel -> {"a1", " $\chi^2$ "}]
Out[98]=
```



Picking $a_1 = -0.004$ (again, close to the expected value), look at how χ^2 varies with a_0 . The best value for the intercept is near 3.0 radians.

```
In[99]:= Plot[calculateChisq[data, a0, -0.004], {a0, 2, 4}, AxesLabel -> {"a0", " $\chi^2$ "}]
Out[99]=
```

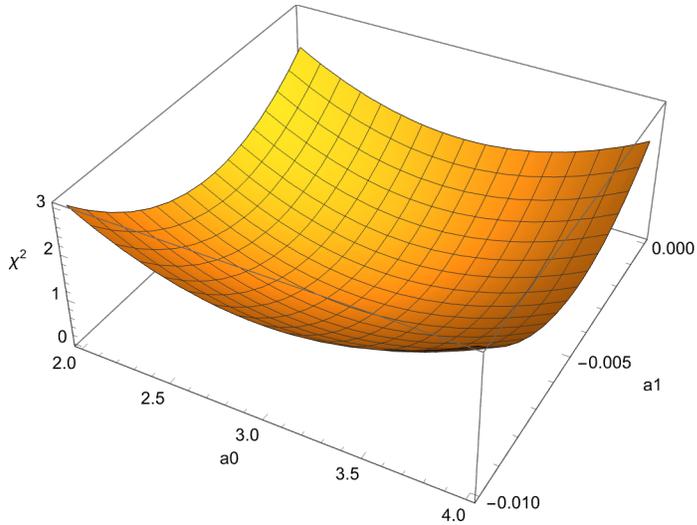


Ultimately, the “best” fit involves a two-dimensional minimization of χ^2 . Sometimes it helps to visualize this sort of thing as a 3D plot or a density plot.

In[100]:=

```
Plot3D[calculateChisq[data, a0, a1], {a0, 2, 4},
  {a1, -0.01, 0}, AxesLabel -> {"a0", "a1", " $\chi^2$ "}]
```

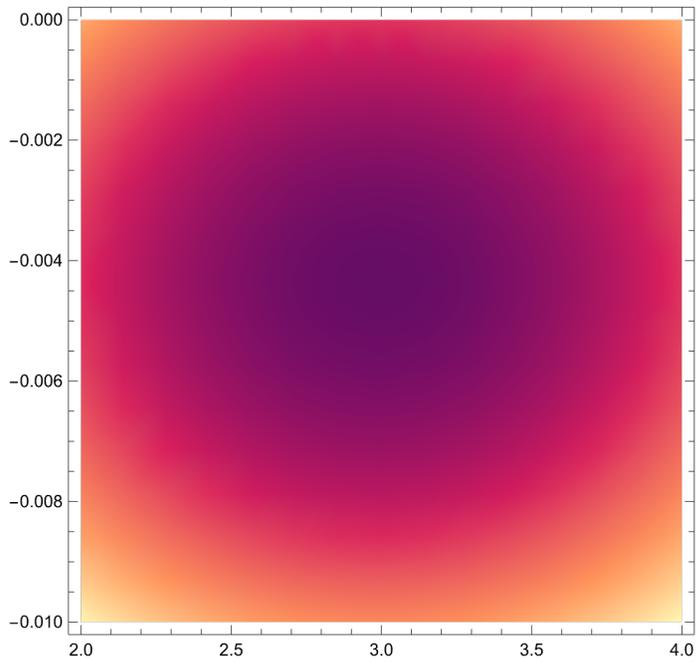
Out[100]=



In[101]:=

```
DensityPlot[calculateChisq[data, a0, a1], {a0, 2, 4}, {a1, -0.01, 0}]
```

Out[101]=



Mathematica can find the minimum here as well.

In[102]:=

```
result = FindMinimum[calculateChisq[data, a0, a1], {a0, a1}]
```

Out[102]=

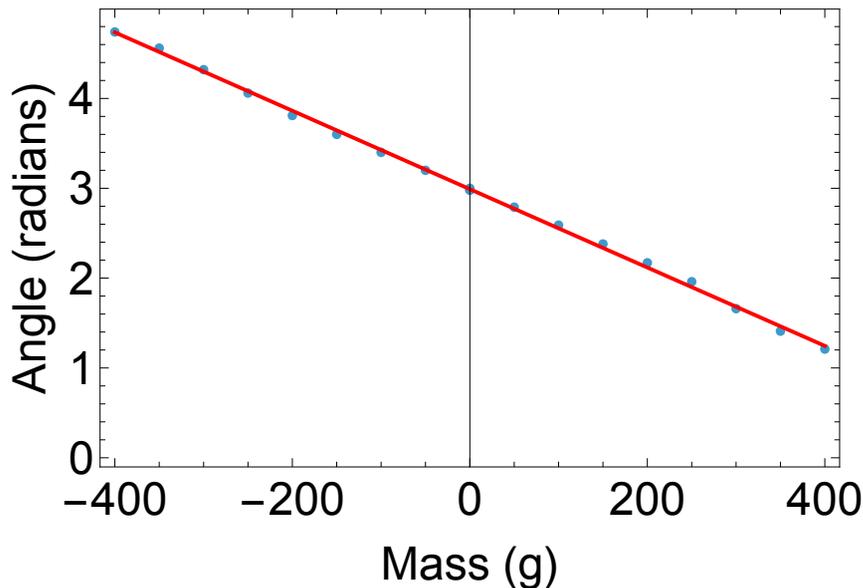
```
{0.00135759, {a0 -> 2.99053, a1 -> -0.00436275}}
```

Finally, here is the best fit curve.

In[103]:=

```
Show[{
  dataPlot,
  Plot[yfit[x, a0, a1] /. result[[2]], {x, -400, 400}, PlotStyle -> Red]
}]
```

Out[103]=



Linear Model Fit

Mathematica can do this minimization automatically. The `LinearModelFit[]` function searches the "parameter space" for the minimum value of χ^2 . It reports a "Confidence Interval" that reflects the curvature of χ^2 -- how much you can vary either parameter without making χ^2 too large.

One other interesting thing to note: Note how our χ^2 space has a "valley" (dark purple in the Density-Plot) where you can explore what happens if you change a_0 but don't change a_1 , and vice-versa. This is reflected in the `LinearModelFit[]` report. Look at the "Correlation Matrix." It tells you, in essence, how well correlated each parameter is with the other. The 1's along the diagonal mean a_0 is perfectly correlated with a_0 (obviously) and a_1 is perfectly correlated with a_1 . The off-diagonal elements tell you that a_0 and a_1 are mostly uncorrelated -- increasing one doesn't really affect the other. This isn't always true for fits. Sometimes, a change in one parameter can be partially compensated for by a change in another. In those cases, they are not completely independent parameters, and the Correlation Matrix elements are further from zero.

```
In[104]:= fit = LinearModelFit[data, {1, m}, m]
```

```
Out[104]= FittedModel[ 2.99 - 0.00436 m ]
```

See the on-line help for more information on dealing with the results from LinearModelFit. Here are some examples of things you can do with it.

```
In[105]:= fit["BestFit"]
```

```
Out[105]= 2.99053 - 0.00436275 m
```

```
In[106]:= fit["BestFitParameters"]
```

```
Out[106]= {2.99053, -0.00436275}
```

```
In[107]:= fit["ParameterConfidenceIntervalTable"]
```

```
Out[107]=
```

	Estimate	Standard Error	Confidence Interval
1	2.99053	0.00845294	{2.97269, 3.00836 }
m	-0.00436275	0.0000364825	{-0.00443972, -0.00428577 }

You can extract the values from that table, though it's clunky.

```
In[108]:= fit["ParameterConfidenceIntervalTableEntries"] // TableForm
```

```
Out[108]//TableForm=
```

2.99053	0.00845294	2.97269 3.00836
-0.00436275	0.0000364825	-0.00443972 -0.00428577

For example, here is the slope and its uncertainty (the slope is the second entry in the table):

```
In[109]:= {slope,  $\delta$ slope} = fit["ParameterConfidenceIntervalTableEntries"][[2, {1, 2}]]
```

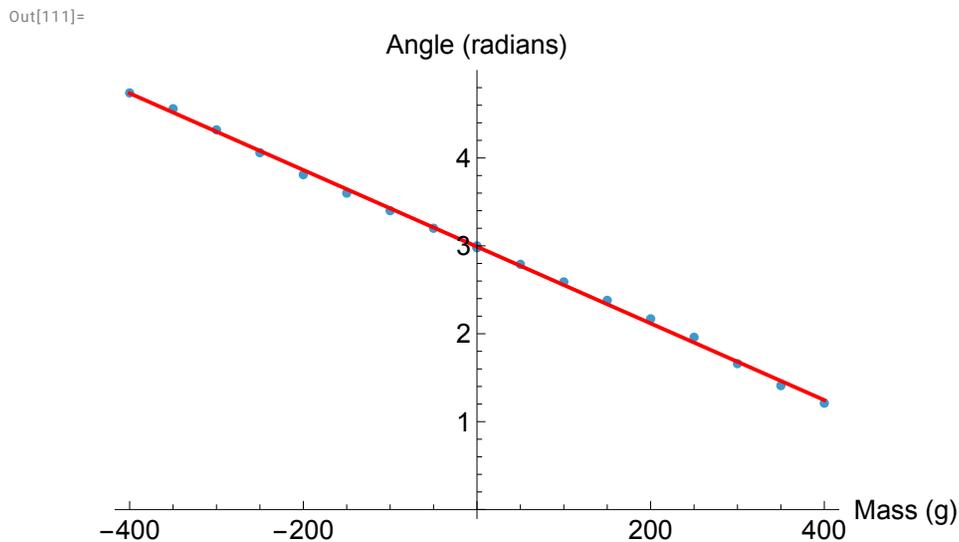
```
Out[109]= {-0.00436275, 0.0000364825}
```

```
In[110]:= fit["CorrelationMatrix"] // MatrixForm
```

```
Out[110]//MatrixForm=
```

$$\begin{pmatrix} 1. & -1.3594 \times 10^{-32} \\ -1.45388 \times 10^{-32} & 1. \end{pmatrix}$$

```
In[111]:= Show[{ListPlot[data], Plot[fit[m], {m, -400, 400}, PlotStyle -> Red]},
  LabelStyle -> Larger, AxesLabel -> {"Mass (g)", "Angle (radians)"},
  ImageSize -> Scaled[0.7]]
```



```
In[112]:= fit["EstimatedVariance"]
```

```
Out[112]= 0.00135759
```

This is the same as our χ^2 .

```
In[113]:= {a0, a1} = fit["BestFitParameters"]
```

```
Out[113]= {2.99053, -0.00436275}
```

```
In[114]:= calculateChisq[data, a0, a1]
```

```
Out[114]= 0.00135759
```

Quality of Fit

Mathematica gives you the uncertainties in the parameters, but you should also check whether the fit is reasonable. In particular, do the differences between the data points and the fit make sense? Is the size believable? *Mathematica* will give you the “Estimated Variance”, which is the same as our manually-calculated χ^2 . That is, it is the average of the square difference between each point and its corresponding fit. (Actually it divides by $N-2$, since there are 2 degrees of freedom used up by the two fit parameters, a_0 and a_1 .)

```
In[115]:= fit["EstimatedVariance"] (* Again, note this is the same as our  $\chi$ -squared. *)
```

```
Out[115]= 0.00135759
```

```
In[116]:= RMSE = Sqrt[fit["EstimatedVariance"]]
```

```
Out[116]= 0.0368455
```

The square root of the estimated variance is thus the RMS (root mean square) error -- the “typical” amount by which the fitted line misses the data. It has the same units as the original y values, and should be compared to the y uncertainties. Are differences of 0.04 radians significant? For this experiment, you can easily read the scale to 0.1 radian, so differences of 0.04 radians seem plausible.

Residuals

Mathematica also lists the individual deviations, known as Residuals. These can be useful for exploring systematic trends. You will explore this idea in a separate homework assignment.

Other Fit Properties.

This is a complete list of all available properties of the fit.

```
In[117]:= fit["Properties"]
```

```
Out[117]= {AdjustedRSquared, AIC, AICc, ANOVA, BestFitAround, BestFitDataAround,
  BasisFunctions, BetaDifferences, BestFit, BestFitParameters, BIC,
  CatcherMatrix, CoefficientOfVariation, CookDistances, CorrelationMatrix,
  CovarianceMatrix, CovarianceRatios, Data, Weights, DesignMatrix,
  DurbinWatsonD, Eigenstructure, EstimatedVariance, FitDifferences,
  FitResiduals, Function, TabularFunction, FVarianceRatios, HatDiagonal,
  MeanPredictions, MeanPredictionBands, ParameterEstimates, PartialSumOfSquares,
  PredictedResponse, Properties, Response, RSquared, SequentialSumOfSquares,
  SingleDeletionVariances, SinglePredictions, SinglePredictionBands,
  StandardizedResiduals, StudentizedResiduals, VarianceInflationFactors}
```

A note about the “RSquared” value

```
In[118]:= fit["RSquared"]
```

```
Out[118]= 0.998813
```

This is close to 1.0. Does that mean you have a good fit? Not necessarily. Consider the following parabolic data set (with a small amount of random noise added)

```
In[119]:=
parab = Table[{x, 1.0 + 0.5 x + 0.3 x2 + RandomReal[{-1, 1}]}, {x, 1, 20, 2}]
```

```
Out[119]=
{{1, 1.39824}, {3, 5.04714}, {5, 11.6745}, {7, 18.8224}, {9, 29.2596},
 {11, 43.4683}, {13, 57.3534}, {15, 75.2681}, {17, 96.6059}, {19, 117.803}}
```

Let's try fitting it with a straight line

```
In[120]:=
pfit = LinearModelFit[parab, {1, x}, x]
```

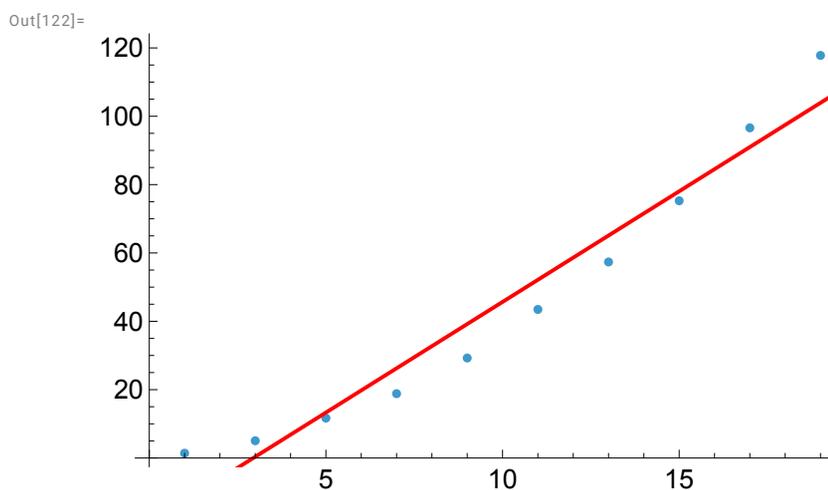
```
Out[120]=
FittedModel[-19.1 + 6.47 x]
```

```
In[121]:=
pfit["ParameterConfidenceIntervalTable"]
```

```
Out[121]=
```

	Estimate	Standard Error	Confidence Interval
1	-19.0671	6.11029	{-33.1575, -4.97673}
x	6.47372	0.52983	{5.25193, 7.69551}

```
In[122]:=
Show[{ListPlot[parab], Plot[pfit[x], {x, 0, 20}, PlotStyle → Red]},
 LabelStyle → Larger, ImageSize → Scaled[0.6]]
```



```
In[123]:=
pfit["RSquared"]
```

```
Out[123]=
0.949139
```

```
In[124]:=
Sqrt[pfit["EstimatedVariance"]]
```

```
Out[124]=
9.62483
```

The R^2 value is close to 1.0, even though we're missing a systematic trend in the data. The R^2 value does not distinguish between data points that scatter about a line and those that deviate systematically from the line. You need to look at the actual graph. Obviously, a parabola is the right fit here:

```
In[125]:= pfit2 = LinearModelFit[parab, {1, x, x^2}, x]
```

```
Out[125]= FittedModel[ 0.737 + 0.562 x + 0.296 x^2 ]
```

```
In[126]:= pfit2["ParameterConfidenceIntervalTable"]
```

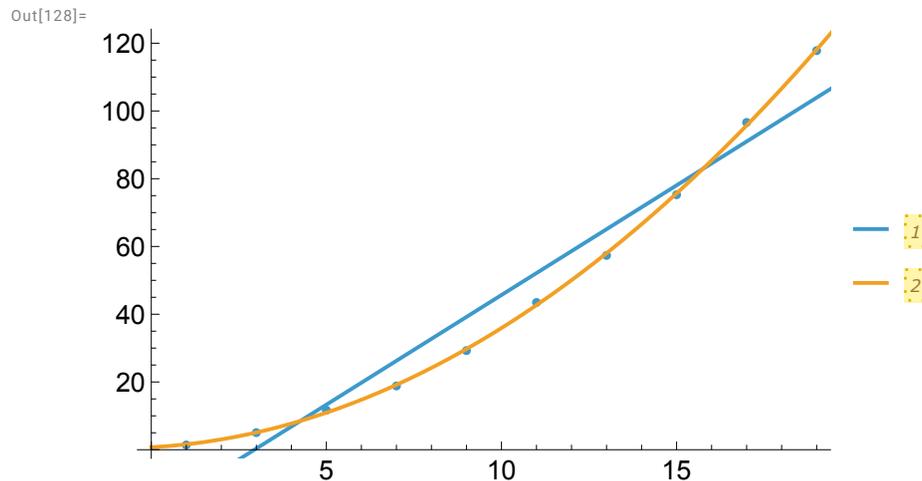
```
Out[126]=
```

	Estimate	Standard Error	Confidence Interval
1	0.736881	0.633846	{-0.761927, 2.23569 }
x	0.562086	0.147178	{0.214067, 0.910106 }
x ²	0.295582	0.00713412	{0.278712, 0.312451 }

```
In[127]:= Sqrt[pfit2["EstimatedVariance"]]
```

```
Out[127]= 0.655719
```

```
In[128]:= Show[{ListPlot[parab],
  Plot[{pfit[x], pfit2[x]}, {x, 0, 20}, PlotLegends -> Automatic]},
  LabelStyle -> Larger, ImageSize -> Scaled[0.6]]
```



Summary of Linear Fit

Read or input the data and transform into appropriate form.

```
In[129]:= rawdata = Import["T0-static-20250212.csv", "CSV"];
```

```
In[130]:= data3 = Select[rawdata, VectorQ[#, NumberQ] &]; (* 3 columns of data *)
```

```
In[131]:= data = data3[[All, {1, 2}]]; (* We only want 2 columns: mass and angle. *)
```

```
In[132]:= dataPlot = ListPlot[data, ImageSize → Scaled[0.8], Frame → True,
  LabelStyle → Large, FrameLabel → {"Mass (g)", "Angle (radians)"}];
```

```
In[133]:= fit = LinearModelFit[data, {1, m}, m]
```

```
Out[133]= FittedModel[ 2.99 - 0.00436 m ]
```

```
In[134]:= fit["BestFitParameters"]
```

```
Out[134]= {2.99053, -0.00436275}
```

```
In[135]:= fit["ParameterConfidenceIntervalTable"]
```

```
Out[135]=
```

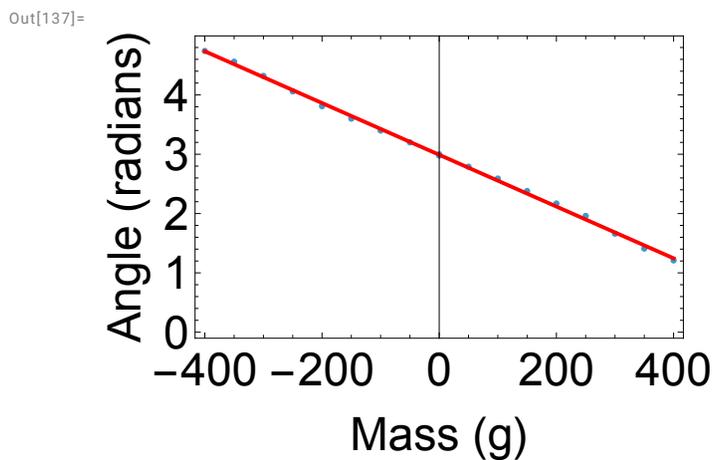
	Estimate	Standard Error	Confidence Interval
1	2.99053	0.00845294	{2.97269, 3.00836 }
m	-0.00436275	0.0000364825	{-0.00443972, -0.00428577 }

```
In[136]:= RMSE = Sqrt[fit["EstimatedVariance"]] (* Root Mean Square error *)
```

```
Out[136]= 0.0368455
```

Overlay plots

```
In[137]:= finalplot = Show[{
  dataPlot,
  Plot[fit[m], {m, -400, 400}, PlotStyle → Red]
}, ImageSize → Scaled[0.5]
]
```



Save your work

```
In[138]:= Export["T0-static-plot.pdf", finalplot]
```

```
Out[138]= T0-static-plot.pdf
```