

Regression in SPSS (Practical)

Regression practical

In this practical we will look at regressing one variable on another variable to explore the relationship between them. This work builds on the concept of correlation that we have looked at in earlier practicals but here we specify one variable as a response (or dependent) variable and look at the effect of another predictor (or independent) variable upon it.

Theories of student learning propose that a key predictor of how well an individual performs in a task is their task-related self-efficacy, that is, their confidence in their ability to perform the task well. In this example, we explore whether a measure of a student's science self-efficacy (SCIEEFF) predicts their achievement on the PISA science test (SCISCORE). The self-efficacy score was derived from students' responses to questions on how easy they would find it to perform eight science tasks on their own, such as "Identify the better of two explanations for the formation of acid rain" (see PISA datafile description for further details).

In this example **SCISCORE** is the response variable and **SCIEEFF** is the predictor variable. To begin with we will simply look at some basic summary information about the variables and plot them in a scatterplot in SPSS which is done as follows:

- Select **Descriptives** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- Copy the **Science test score[SCISCORE]** and **Science self-efficacy score[SCIEEFF]** variables into the **Variable(s)** box.
- Click on the **Options** button.
- Ensure that the **Mean, Std. deviation, Minimum** and **Maximum** options are selected only.
- Click on the **Continue** button to return to the main window.
- Click on the **OK** button to run the command.

The descriptive statistics will then appear as shown below:

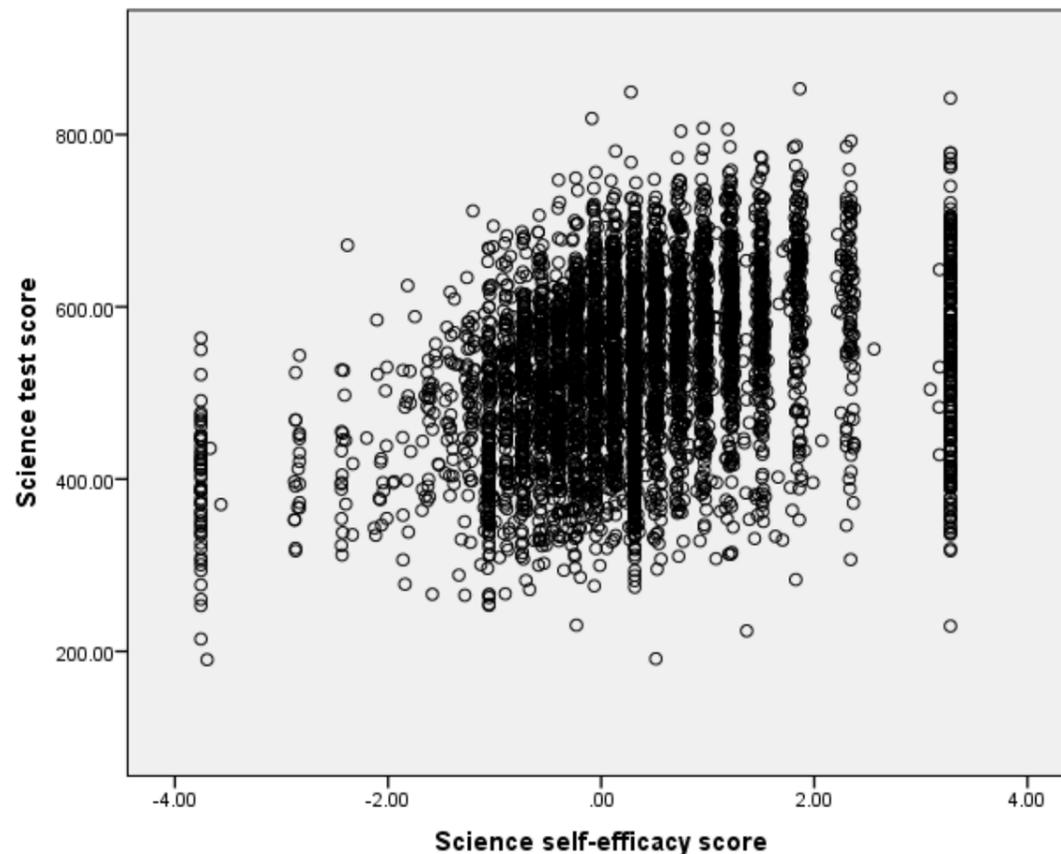
Descriptive Statistics					
	N	Minimum	Maximum	Mean	Std. Deviation
Science test score	5194	175.22	853.13	523.5944	102.57804
Science self-efficacy score	4726	-3.76	3.28	.3671	1.19427
Valid N (listwise)	4726				

Here we see a row in the table for each variable. **SCISCORE** is the response variable and takes values between 175.22 and 853.13 with a mean of 523.5944. **SCIEEFF** is the predictor variable and takes values between -3.76 and 3.28 with a mean of .3671.

We can next plot these variables against each other following instructions below:

- Select **Scatter/Dot** from the **Legacy diagnostics** available from the **Graphs** menu.
- Select **Simple Scatter** and click on Define to bring up the Simple Scatterplot window.
- Copy the **Science test score[SCISCORE]** variable into the **Y Axis** box.
- Copy the **Science self-efficacy score[SCIEEFF]** variable into the **X Axis** box.
- Click on the **OK** button.

SPSS will then draw a scatterplot of the two variables which can be seen below:



The scatterplot gives a general idea of the relationship between the two variables and we can look by eye to see if a linear relationship is suitable. In this example there appears to be a positive relationship as there are more points in the bottom-left and top-right quarters of the plot than in the top-left and bottom-right corners.

We now need to actually run the linear regression to look at if there is a significant (linear) effect of **SCIEEFF** on **SCISCORE**. This is done in SPSS as follows:

- Select **Linear** from the **Regression** submenu available from the **Analyze** menu.
- Copy the **Science test score[SCISCORE]** variable into the **Dependent** box.
- Copy the **SCIEEFF** variable into the **Independent(s)** box.
- Click on the **Statistics** button.
- On the screen appears add the tick for **Confidence Interval** to those for **Estimates** and **Model fit**.
- Click on the **Continue** button to return to the main window.
- Click on the **Save** button.
- On the screen appears select the tick for **Standardized** found under **Residuals**.
- Click on the **Continue** button to return to the main window.
- Click on the **OK** button to run the command.

The command will run and five output tables will be presented. The first of which is the **Variables Entered/Removed** table but as it is only useful when we do multiple regression we will not show it.

The next table is the **Model Summary** table.

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.341 ^a	.116	.116	94.25022

a. Predictors: (Constant), Science self-efficacy score

Here we see some fit statistics for the overall model. The statistic R here takes the value .341 and is equivalent to the Pearson correlation coefficient for a simple linear regression, that is, a regression with only one predictor variable. R square (.116) is simply the value of R squared (R multiplied by itself) and represents the proportion of variance in the response variable, **SCISCORE** explained by the predictor variables. The table also includes an adjusted R square measure which here takes value .116 and is a version of R squared that is adjusted to take account of the number of predictors (one in the case of a simple linear regression) that are in the model.

The next table is the **ANOVA** table.

ANOVA

Model		Sum of Squares	df	Mean Square	F	Sig.
1	Regression	5504631.271	1	5504631.271	619.674	.000 ^b
	Residual	41963783.385	4724	8883.104		
	Total	47468414.656	4725			

b. Predictors: (Constant), Science self-efficacy score

The ANOVA (Analysis of Variance) table is used to look at how well the predictors as a whole can account for differences in the response variable. It is used in SPSS for many models to show how the variability in the response variable is partitioned between different explanatory components (which it measures in terms of sums of squares).

Here we see in the sum of squares column that the total sum of squares, which captures the total variance in **SCISCORE** is 47468414.656. This can be split into two parts. The regression sum of squares (5504631.271) captures the variability in the values of **SCISCORE** that would be predicted on the basis of the predictor variables alone, so here on the basis of **SCIEEFF**. The residual sum of squares (41963783.385) is the variation in the dependent variable that remains unexplained by the predictor variables. The R-squared that we saw in the earlier table is related to the sum of squares - it expresses the regression (or explained) SS as a fraction of the total SS i.e. $5504631.271/47468414.656=0.116$.

These sums of squares have associated degrees of freedom (df). For the total sum of squares the df is one less than the number of observations (N - 1, here 4725) due to fitting a mean to the data. The regression sum of squares has df = 1 to account for the 1 predictor in the model. The residual df is then the difference between the total df and the regression df = 4724. The next column is the mean squares (sums of squares adjusted for dfs) which are used to construct a test statistic, F, shown in the fifth column. Here we see that F takes value 619.674 and can be used to test the null hypothesis that there is no linear relationship between the predictors and the response variable. To do this the value of F needs to be compared with an F distribution with 1 and 4724 degrees of freedom. This test results in a p value that is given in the Sig. column. Here p is quoted as .000 (reported as p < .001) which is less than the conventional .05 level used to judge significance. We therefore have sufficient evidence to reject the null hypothesis that the predictors have no effect.

The next table is the **Coefficients** table.

		Coefficients						
		Unstandardized Coefficients		Standardized Coefficients		95.0% Confidence Interval for B		
Model		B	Std. Error	Beta	t	Sig.	Lower Bound	Upper Bound
1	(Constant)	520.752	1.434		363.064	.000	517.940	523.564
	Science self-efficacy score	28.580	1.148	.341	24.893	.000	26.329	30.831

This table gives the most interesting information about the regression model. We begin with the coefficients that form the regression equation. The regression intercept (labelled Constant in SPSS) takes value 520.752 and is the value of the regression line when **SCIEEFF** takes value 0. The regression slope, the B, takes value 28.580 and is the amount by which we predict that **SCISCORE** changes for an increase in **SCIEEFF** of one unit.

Both coefficients have associated standard errors that can be used to assess their significance and also in the case of the slope, to construct a standardised coefficient. This can be seen under the Beta column and takes value .341 which represents the predicted change in **SCISCORE** in standard deviation units for an increase of one standard deviation in **SCIEEFF**. The standardised coefficient beta can be interpreted as a "unit-free" measure of effect size, one that can be used to compare the magnitude of effects of predictors measured in different units.

To test for the significance of the coefficients we need to form test statistics which are reported under the t column and are simply B / Std. Error. For the slope the t statistic is 24.893 and this value can be compared with a t distribution to test the null hypothesis that B = 0. We can see the resulting p value for the test under the Sig. column. Here p is quoted as .000 (reported as p < .001) so we have sufficient evidence to reject the null hypothesis that the slope coefficient on **SCIEEFF** is zero. We therefore have significant evidence to reject the null hypothesis that the slope is zero.

We can also check if the intercept is different from zero though this is often of less interest. For the intercept the t statistic is 363.064 and this value can be compared with a t distribution to test the null hypothesis that the intercept B = 0. We can see the resulting p value for the test under the Sig. column. Here p is quoted as .000 (reported as p < .001) so we have sufficient evidence to reject the null hypothesis that the intercept is zero. We therefore have significant evidence to reject the null hypothesis that the intercept is zero.

The final two columns give 95% confidence intervals for the coefficients, which show a lower bound of 517.940 and an upper bound of 523.564 for our estimate of the intercept.

Similarly the 95% confidence interval bounds for the **SCIEEFF** B coefficient are 26.329 and 30.831. Zero does not lie between these values, consistent with our rejection of the null hypothesis that B = 0.

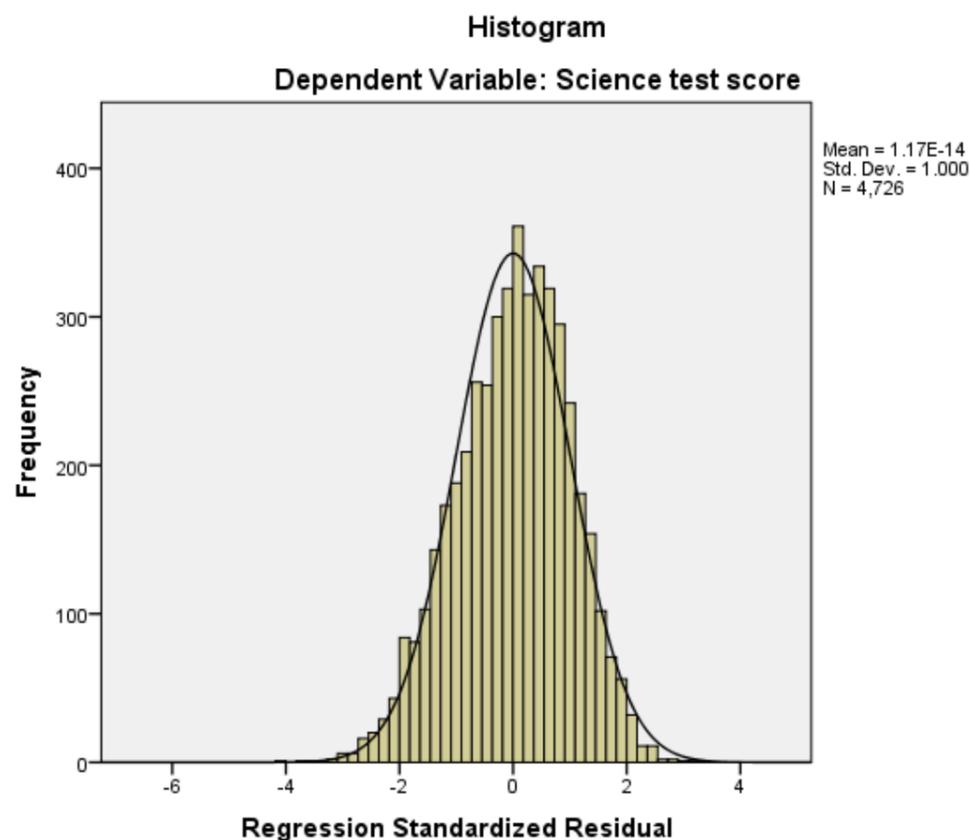
The final table is the **Residuals statistics** table.

Residuals Statistics					
	Minimum	Maximum	Mean	Std. Deviation	N
Predicted Value	413.3918	614.4227	531.2446	34.13212	4726
Residual	-385.03467	320.70566	.00000	94.24025	4726
Std. Predicted Value	-3.453	2.437	.000	1.000	4726
Std. Residual	-4.085	3.403	.000	1.000	4726

This table just summarises the predictions and residuals that come out of the regression and it is perhaps easier to look at these via plots. As we ticked the box to request that standardised residuals were saved this has resulted in an additional variable, named **ZRE_1**, being stored in the dataset at the end of the existing variables (You can see this by viewing the dataset in the SPSS Data Editor Window). We can use this variable to create some residual plots to assess the fit of the model. We will firstly plot a histogram of the residuals to check their normality which can be done in SPSS as follows:

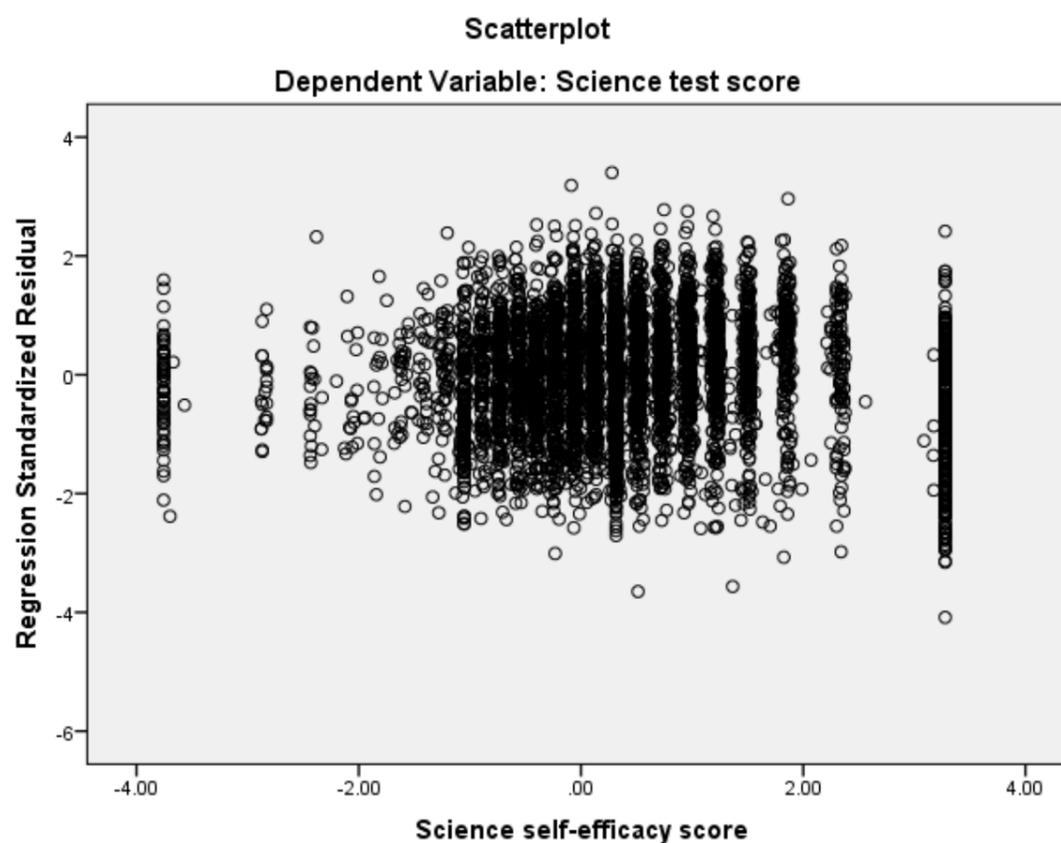
- Select **Histogram** from the **Legacy dialogs** available from the **Graphs** menu.
- Copy the **Standardized Residual [ZRE_1]** variable into the **Variable** box.

- Click on the **Display normal curve** tick box.
- Click on the **OK** button to produce the graph.



Here we hope to see the histogram of residuals roughly following the shape of the normal curve that is superimposed over them. We can also look at how the distribution of the residuals interacts with the predictor variable to check there is no relationship. We do this via a scatterplot which can be produced in SPSS as follows:

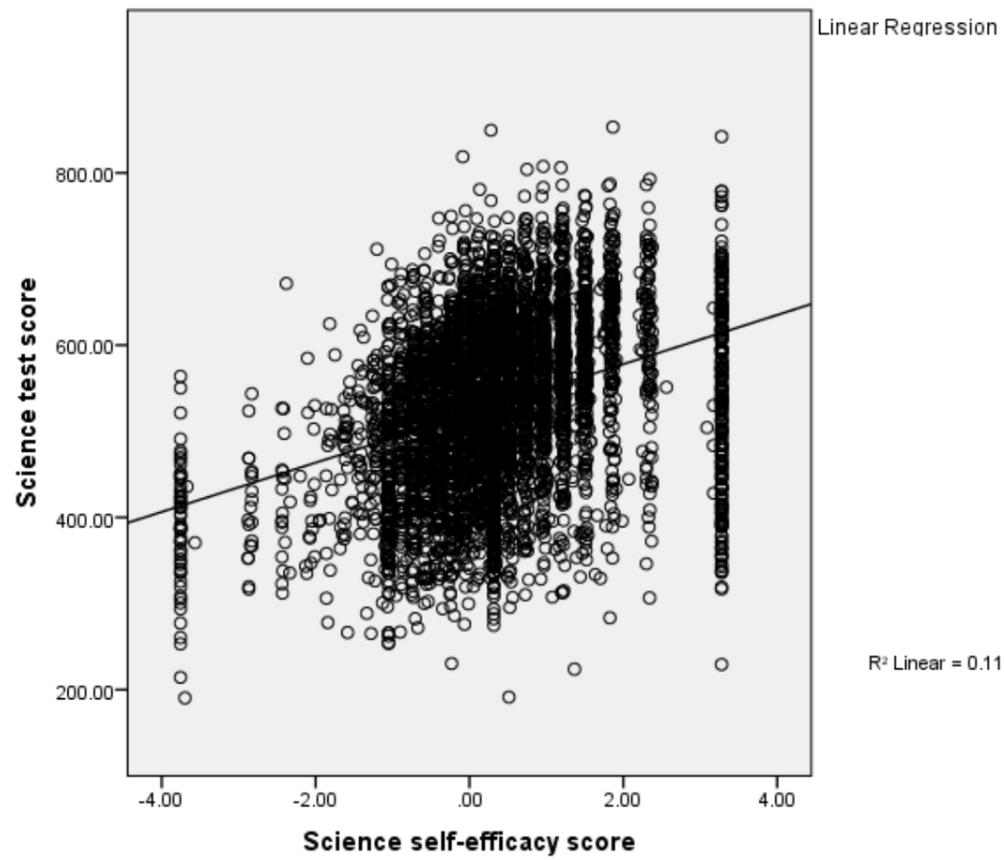
- Select **Scatter/Dot** from the **Legacy dialogs** available from the **Graphs** menu.
- Select Simple Scatter and click on Define to bring up the Simple Scatterplot window.
- Copy the **Standardized Residual [ZRE_1]** variable into the **Y Axis** box.
- Copy the **Science self-efficacy score[SCIEEFF]** variable into the **X Axis** box.
- Click on the **OK** button and the plot will appear.



Here we hope that the residuals show a random scatter when plotted against the predictor variable and also that their variability is constant across different values of the predictor variable.

Finally we might like to superimpose the regression line onto the scatterplot we drew earlier of the response against the predictor, so that the strength of the linear relationship is clear. To do this we will need to use the **Chart Editor** in SPSS so follow the following instructions:

- Locate the earlier scatterplot in the SPSS output window noting you may need to scroll up to find it.
- Double click with the left mouse button on the plot and it will pop out into a Chart Editor window.
- On the window click on the 5th button from the left on the bottom row of icons (It will say **Add Fit Line at Total** if you hover the mouse over it)
- On the **Properties** window that appears remove the tick next to **Attach label to line** as otherwise the equation is superimposed on the plot which looks untidy.
- Click on the **Close** button and the line will be added in the Chart Editor window.
- Finally close the **Chart Editor** window and the graph in the output window will now have the fixed line as shown below:



Note that the scatterplot now also contains the R-squared value which corresponds to the R-squared value we saw in the regression fit earlier.

The analysis has shown that, as expected, sense of self-efficacy in science tasks is significantly positively related to achievement in science, with around 12% of the variance in science test scores accounted for by individual differences in students' self-efficacy.