

The British Academy funded Statistical eBook grant

Professor William Browne,
Chris Charlton and Liz Washbrook
Centre for Multilevel Modelling, University of Bristol



What will we cover ?

- Background to CMM and StatJR
- Interoperability and e-Books
- British Academy grant
- Topics to be covered
- Work packages 1 – 5 progress



Background to CMM

- Cross-faculty statistical research group primarily based in Education where we are a Research Centre.
- Produce statistical software packages, MLwiN and StatJR with over 15,000 users.
- Also LEMMA online training materials with nearly 20,000 users.
- Historically research funded by the ESRC via several programme nodes to a total of more than £5M in the past 10 years
- See <http://www.bristol.ac.uk/cmm/>



Stat-JR

- A statistical package developed by the team at the Centre for Multilevel Modelling with colleagues at Southampton.
- Contains it's own (MCMC-based) estimation engine.
- System based on the idea of a suite of templates where each template performs a specific operation.
- Also allows interoperability with other software packages, so for example might have a regression template that fits regressions using various software packages.
- The initial TREE interface runs in a web browser.
- There are also newer eBook and workflow interfaces.
- Several ESRC grants have enabled Stat-JR to be written.



eBooks



+



=



An electronic book is a book-pUBLICATION in digital form. In the US more books are published online than distributed in hard copy in book shops.



Statistical (and Mathematical) eBooks

- The idea is can we incorporate statistical content into an eBook? Of course a statistical textbook is no different on paper to any other document when it comes to creating a pdf file (aside from maybe more equations!)
- The difference is in what 'enhancements' we can add and so the idea here is combining the text book with the statistics package i.e. interactive examples, allowing the user to include their own dataset etc.



Multilevel modelling with the 'tutorial' dataset

← Previous 1 2 3 4 5 Next → Go to page

Navigate through pages of eBook

- Overview
- The tutorial dataset
- Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
- Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset
 - Choosing your

Hierarchical table of contents (can be expanded / collapsed at each node)

Overview

This eBook provides a brief overview of the tutorial dataset. We are developing eBook content will appear tailored to your progress through the table of contents on the left.

EBook functionality is still being developed, so you may notice the odd thing here or there yet to be finessed (such as the large number of decimal places sometimes returned!), but we nevertheless wanted to introduce you to what we hope you find to be an interesting means of exploring statistics, and we would very much appreciate any comments you have.

Note that there may be a short delay until all available contents on a particular page are uploaded - you can keep an eye on progress either via the gauge in the top-left corner of the browser window, or by looking at the command window running in the background.

NB: if your eBook crashes, then you can reload the eBook by choosing Debug > Reload eBook from the black bar towards the top of this window. That will wipe your previous choices, I'm afraid, but it will (hopefully) breathe life back into the software!

The tutorial dataset

The **tutorial** dataset is one of the example datasets provided with the Stat-JR package (as well as with the software package MLwiN) and is summarised below. This dataset was selected from a much larger dataset of examination results from six inner London Education Authorities (school boards). A key aim of the original analysis was to establish whether some secondary schools were more 'effective' than others in promoting students' learning and development, taking account of variations in the characteristics of students when they started secondary school. The analysis then looked for factors associated with any school differences found. Thus the focus was on an analysis of examination performance after adjusting for student intake achievements.

Exploring the tutorial dataset

We'll be modelling **normexam** as the response (score) variable: as the summary below indicates, this represents the students' exam score at age 16, normalised to have an approximately standard Normal distribution.

In fact you can view the full dataset via the Resources button, which you can find in the black bar at the top of this window. In the resulting

Multilevel modelling with the 'tutorial' dataset

Finished

Summary table of tutorial dataset

Column name	n	Missing	Min	Max	Description
school	4059	0	1	65	Numeric school identifier
student	4059	0	1	198	Numeric student identifier
normexam	4059	0	-3.67	3.67	Students' exam score at age 16, normalised to have approximately a standard Normal distribution.
cons	4059	0	1	1	A column of ones. If included as an explanatory variable in a regression model, its coefficient is the intercept.
standlrt	4059	0	-2.93	3.02	Students' score at age 11 on the London Reading Test (LRT), standardised using Z-scores.
girl	4059	0	0	1	Students' gender: 0=boy; 1=girl
schgend	4059	0	1	3	School gender: 1=mixed; 2=boys' school; 3=girls' school
avslrt	4059	0	-0.76	0.64	Average LRT score in school
schav	4059	0	1	3	Average LRT score in school, coded into 3 categories: 1=bottom 25%; 2=middle 50%; 3=top 25%
vrband	4059	0	1	3	Students' score in test of verbal reasoning at age 11, coded into 3 categories: 1=top 25%; 2=middle 50%; 3=bottom 25%

Plotting variables

Here you can graphically-explore the **tutorial** dataset.

In the first two sections, below, you can produce a densityplot and XY plot, respectively; here you can re-specify your choice of variables

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
 - Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset.
 - Choosing your



Multilevel modelling with the 'tutorial' dataset

Finished

Your choice of plot

Finally, here you have more flexibility in specifying a plot of your choice. For more information on what the various options mean, please refer to the **PlotsViaR template eBook**...

Type of plot: [about](#)

...then, once you have made your choices, **your plot will appear here:**

Cross-tabulation

Here you can create a table of means and standard deviations for one variable, conditioned on another variable. The first question asks which variable to condition on: a column will be produced for each value of this variable, and so for it to be a useful guide to your data it is best if the variable you choose here consists of relatively few, discrete categories (e.g. **girl**, **schgend**, etc). If you don't want to condition on any variables, you can simply choose **cons**.

What variable do you want to condition your columns on?:

What variable do you want to produce means etc for?:

[about](#)

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
 - Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset
 - Choosing your

Multilevel modelling with the 'tutorial' dataset

← Previous 1 2 3 4 5 Next → Go to page



Finished

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
- Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset
 - Choosing your

Your choice of plot

Finally, here you have more flexibility in specifying a plot of your choice. For more information on what the various options mean, please refer to the [PlotsViaR template eBook](#)...

Which variable would you like to use to construct x-axis panel: schgend

Which variable would you like to use to construct y-axis panel: vrband

Do you want the variable name included in panel bar, or just the level: Yes

Submit

about

...then, once you have made your choices, **your plot will appear here:**

Cross-tabulation

Here you can create a table of means and standard deviations for one variable, conditioned on another variable. The first question asks which variable to condition on: a column will be produced for each value of this variable, and so for it to be a useful guide to your data it is best if the variable you choose here consists of relatively few, discrete categories (e.g. **girl**, **schgend**, etc). If you don't want to condition on any variables, you can simply choose **cons**.

What variable do you want to condition your columns on?: school

What variable do you want to produce means etc for?:



Multilevel modelling with the 'tutorial' dataset

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
- Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset
 - Choosing your

Your choice of plot

Finally, here you have more flexibility in specifying a plot of your choice. For more information on what the various options mean, please refer to the [PlotsViaR template eBook](#)...

Which variable would you like to use to construct x-axis panel:

Which variable would you like to use to construct y-axis panel:

Do you want the variable name included in panel bar, or just the level:

[about](#)

...then, once you have made your choices, **your plot will appear here:**

Cross-tabulation

Here you can create a table of means and standard deviations for one variable, conditioned on another variable. The first question asks which variable to condition on: a column will be produced for each value of this variable, and so for it to be a useful guide to your data it is best if the variable you choose here consists of relatively few, discrete categories (e.g. **girl**, **schgend**, etc). If you don't want to condition on any variables, you can simply choose **cons**.

What variable do you want to condition your columns on?:

What variable do you want to produce means etc for?:

Multilevel modelling with the 'tutorial' dataset



Finished

← Previous 1 2 3 4 5 Next → Go to page

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot
 - Cross-tabulation
- Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset
 - Choosing your

Your choice of plot


Finally, here you have more flexibility in specifying a plot of your choice. For more information on what the various options mean, please refer to the [PlotsViaR template eBook](#)...

Which variable would you like to use to construct x-axis panel:


Which variable would you like to use to construct y-axis panel:

Do you want the variable name included in panel bar, or just the level:

[about](#)



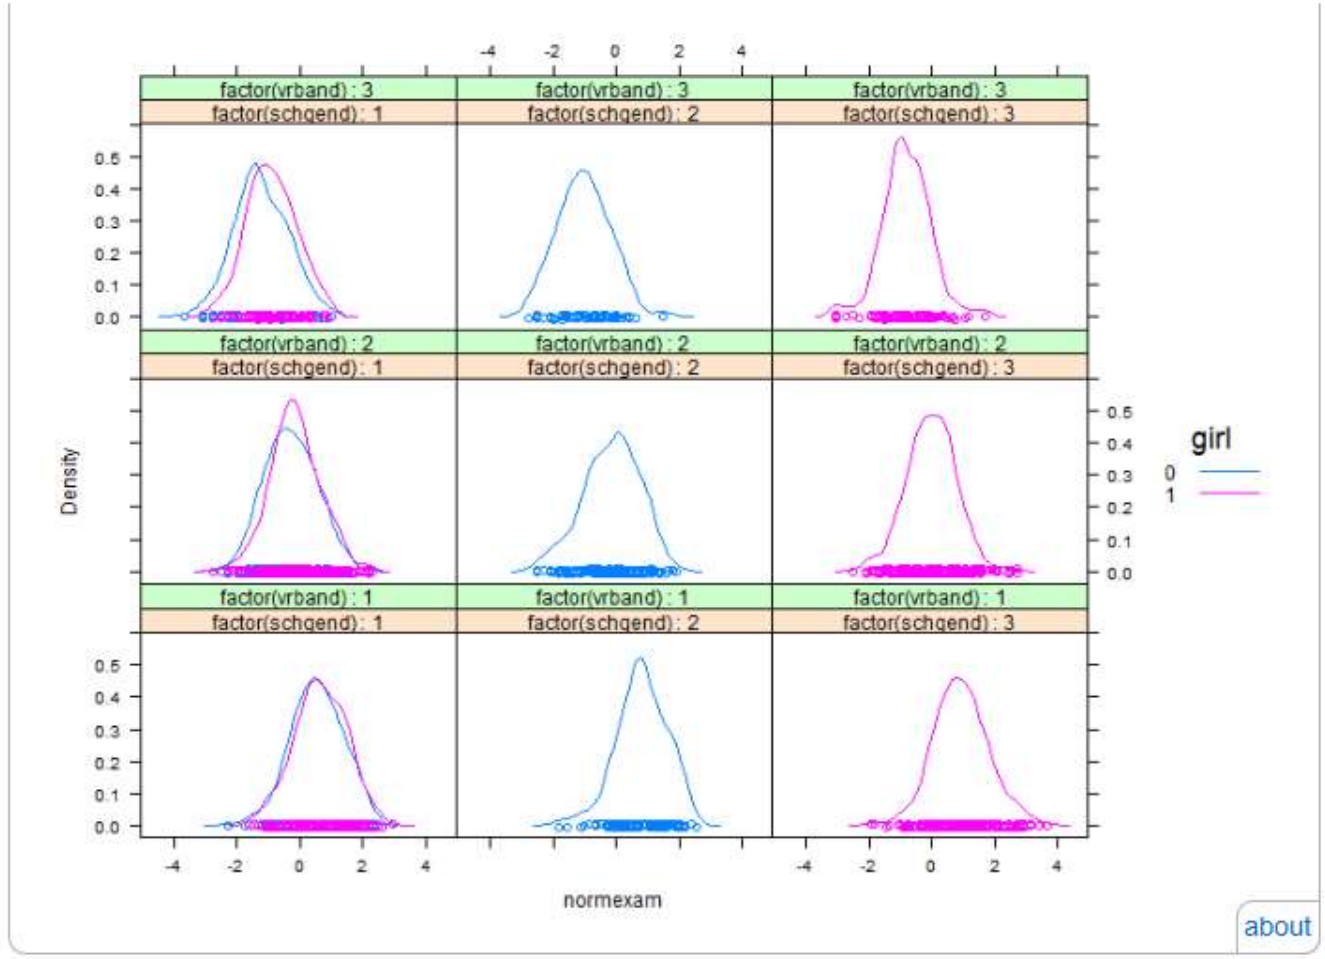
Here is the plot you requested:



Multilevel modelling with the 'tutorial' dataset

Finished

- Overview
 - The tutorial dataset
 - Exploring the tutorial dataset
 - Summary table of tutorial dataset
 - Plotting variables
 - Densityplot
 - XY plot
 - Your choice of plot**
 - Cross-tabulation
- Modelling the dataset
 - Modelling one or two levels?
 - Comparing a 1-level and 2-level model
 - Partitioning variance in a 2-level model
 - References
 - Exploring explanatory variables
 - Summary table of tutorial dataset



about

Motivation for British Academy grant

- We ran a workshop demonstrating some of the new features in StatJR attended by John MacInnes and Rich Harris.
- In current ESRC grant we have been developing Statistical Analysis Assistants (SAAs) which are interactive eBooks that assist you with your analysis.
- As a start we considered automating simple operations.
- John and Rich thought an excellent addition would be using this for teaching and automated teaching material generation.
- The initial proposal was to do everything directly in StatJR but this got switched to creating the materials to use SPSS taking advantage of interoperability.



The end product – what the *student gets*

12 sets of practical exercises (pdfs) with 3 components

1. Takes student through a particular statistical concept in detail, and how to implement it in SPSS, using a specific data example (*learning component*)
2. A worksheet that asks the student to try out their knowledge by applying the techniques to a second dataset or set of variables (*practice component*)
3. Solutions to the worksheet (*self-evaluation component*)



What the *tutor* gets

- The set of static practicals using our choice of data example (PISA data as no restrictions on access)
- Instructions to how to use the Stat-JR software to tailor the practicals to their own choice of datasets/variables
- Makes it quick and easy to
 - Create a suite of discipline-specific materials for teaching and learning
 - Produce multiple versions of worksheets (with solutions) on different substantive topics or using different data sources



Work packages

The grant has 5 work packages:

1. Work package 1 consists of choosing topics and creating a single set of static practicals with solutions
2. Work package 2 consists of extending this to allow the materials to become dynamic and work with other datasets
3. Work package 3 consists of modifying StatJR to give QM teachers tools to customise the materials
4. Work package 4 consists of complementing the practicals / solutions with concept materials (learning component)
5. Work package 5 is demonstrating the materials to the community via a workshop



Work package 1

The list of topics is finalised as:


1. Describing categorical variables (summary stats and graphs)
2. Describing continuous variables (summary stats and graphs)
3. Tabulating data
4. Checking for normality
5. Two sample t tests
6. Paired t tests
7. Non parametric tests
8. Chi-squared tests
9. Correlation
10. Linear Regression
11. ANOVA
12. Multiple Regression



Work package 2 (and 1)

- In practice we have constructed the dynamic materials first and from them used test datasets to construct static files
- At this stage we have drafts of at least parts of all the 12 practicals with a little bit of tidying required.
- In the next couple of slides we show a few screen shots to give an idea.
- Basically the practicals contain contextual text in terms of interpretation of the output but not the data context.
- When the materials are complete we intend to then construct a set of static materials using the PISA data and show how to add more data context.





variables that correspond to groups and also possible to test for group differences in several variables simultaneously.

Here however we test only one variable, V1Mass.

Below you will see instructions to perform the t test in SPSS. If you follow the instructions you will see the two tabular outputs that are embedded in the explanations below.

- Select **Compare Means** from the **Analyze** menu.
- Select **Independent-Sample T Test...** from the **Compare Means** sub-menu.
- Click on the **reset** button
- Copy the **[V1Mass]** variables into the **Test Variable(s):** box.
- Copy the **[Rep]** variable into the **Grouping Variable:** box.
- Click on the **Define Groups...** button
- Click on the **Use specified values** button
- Type **1** into the **Group 1** box.
- Type **2** into the **Group 2** box.
- Click on the **Continue** button
- Click on the **OK** button

← Instructions

The first SPSS output table contains summary statistics for all the variables considered split by group and can be seen below.

Group Statistics

	Rep	N	Mean	Std. Deviation	Std. Error Mean
V1Mass	1	60	1695.48	162.301	20.953
	2	60	1743.60	132.356	17.087

← SPSS output

The summary statistics table contains 5 columns and 1 row for each group in each variable to be tested. After the first column which contains the name of each dependent variable and group categories we next see the number of valid observations in each group, i.e. cases with a valid value of **V1Mass**. Here for the group indexed by **Rep = 1**, we have 60 observations and for **Rep = 2**, there are 60 observations. Next we see that the mean of the variable **V1Mass** for the group with **Rep = 1** is 1695.48 whilst for the group with **Rep = 2** it is 1743.6. Hence the group with **Rep = 2** has the bigger mean and the test will now establish if this distance is statistically significant.

In the next column we see the standard deviations for **V1Mass** variable in the two groups. As we will see in the next table there are two versions of the test depending on whether the variability (and therefore the standard deviations) in the two groups can be assumed equal or not. In this case the standard deviation of **V1Mass** when **Rep = 1** is 162.301 whilst for **Rep = 2** it is 132.356. So there is slightly more variability among **Rep = 1** than **Rep = 2**. But is the difference big enough to violate the assumption of equal variances? In the final column are the standard errors of the means for each group. Whilst the standard deviations measure the variability in the data the standard errors of the means measures how confident we are in the estimates of the means. As we collect more data the standard error of the mean gets smaller as we get more confident in the mean estimate and in fact the formula for the standard error of the mean = standard deviation / square root of N in this case the standard error of the mean for **V1Mass** when **Rep = 1** is 20.953 whilst for **Rep = 2** it is 17.087.

The second SPSS output table contains details of the test itself and can be seen below:

Independent Samples Test

	Levene's Test for Equality of Variances	t Test for Equality of Means								
		F	Sig.	t	df	Sig. (2-tailed)	Mean Difference	Std. Error Difference	95% Confidence Interval of the Difference	
									Lower	Upper
V1Mass	Equal variances assumed	1.547	.216	-1.790	118	.078	-48.117	27.037	-101.657	5.424

← context specific text.



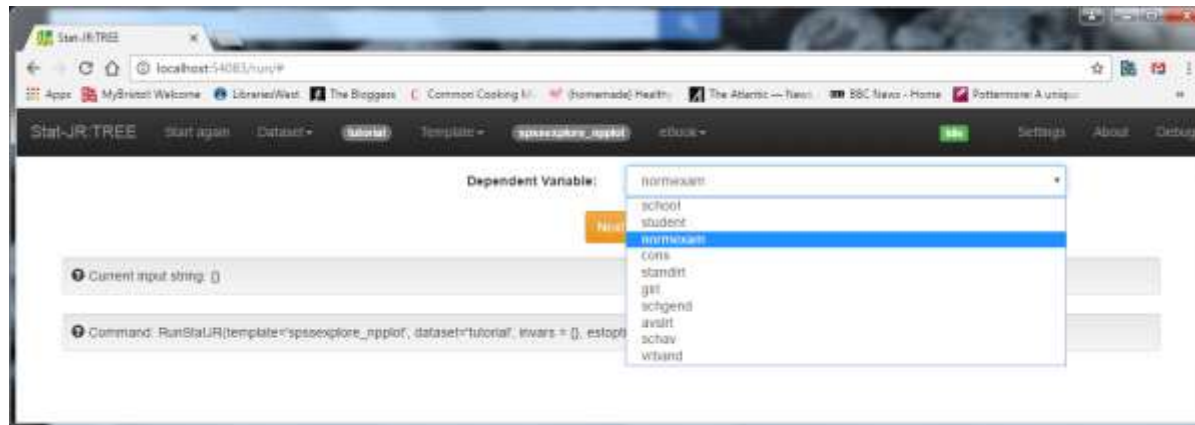
Work package 3

- The first two work packages are largely concerned with content construction whilst work package 3 involves improvements to StatJR specific to this grant. There are three main areas covered:
 - 1) *Better Interfacing with SPSS* – initially it took 30 seconds per SPSS call to use system. Now once started most practicals can be constructed in under 10 seconds.
 - 2) *Improving the eBook writer interface* – We will talk about in the next slide
 - 3) *Improving Exporting of eBooks to PDF for printing* – initially the eBook interface was great for screen display but poor for printing. This has improved to keep SPSS outputs on single pages and to ensure they are rendered appropriately



eBook writer

The first stage is to choose the appropriate template – which aligns either with a full practical or a part of a practical and to choose a dataset

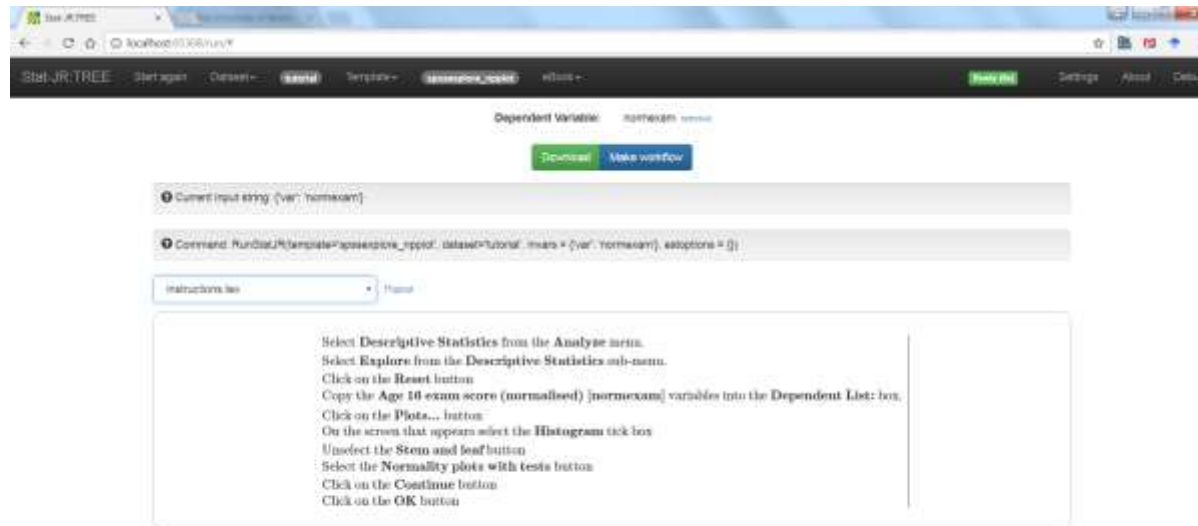


The QM teacher then chooses the particular inputs that correspond to the variables to be used in the practical.



eBook writer

StatJR then creates lots of objects including SPSS outputs, contextual text describing the outputs and blocks of instructions for using SPSS as illustrated below.

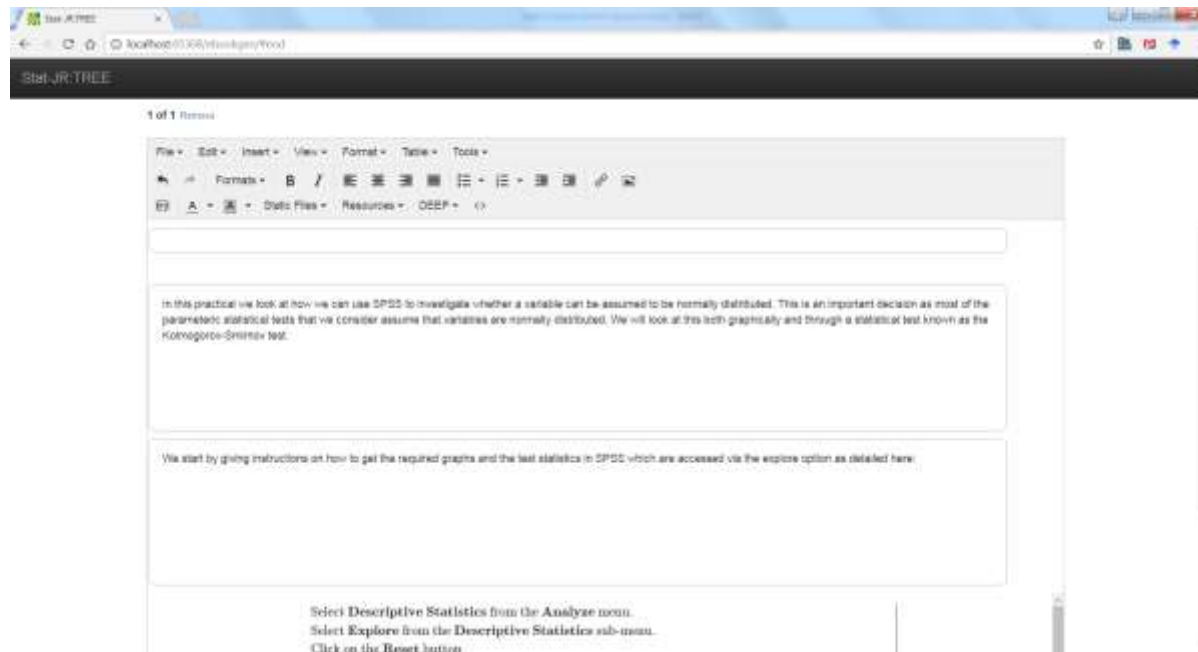


There is also a single combined output that puts these objects together



eBook writer

The QM teacher can then piece together the objects in turn as shown below:

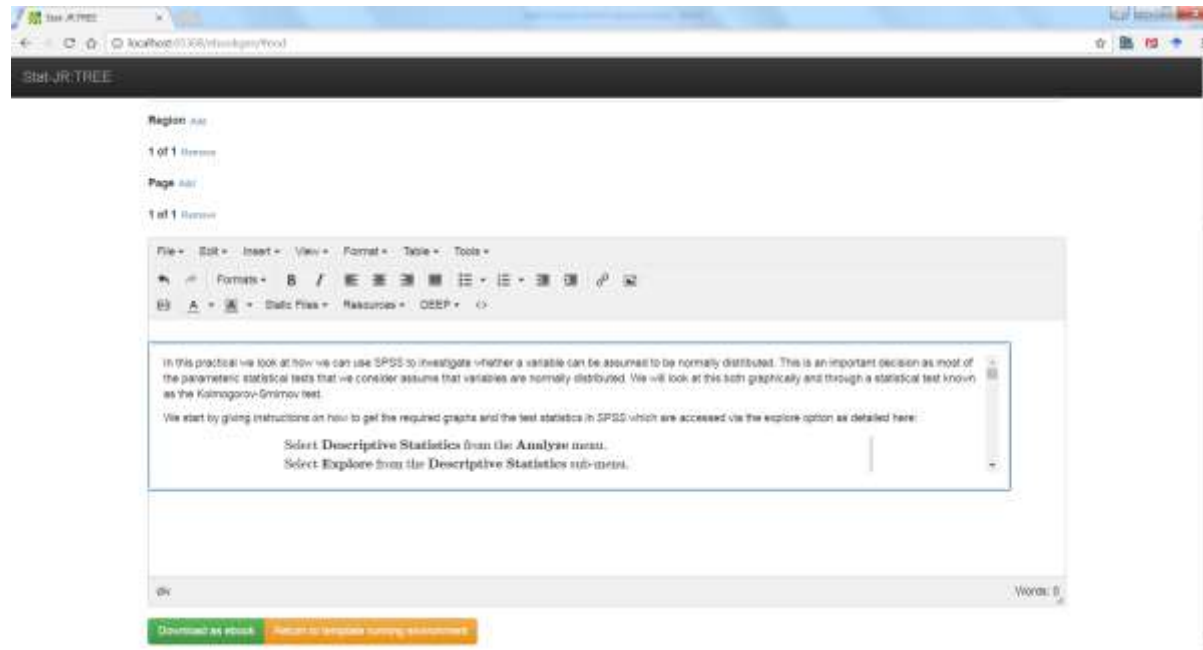


This allows them to add additional dataset specific contextual information and to construct practicals without solutions by omitting specific objects.



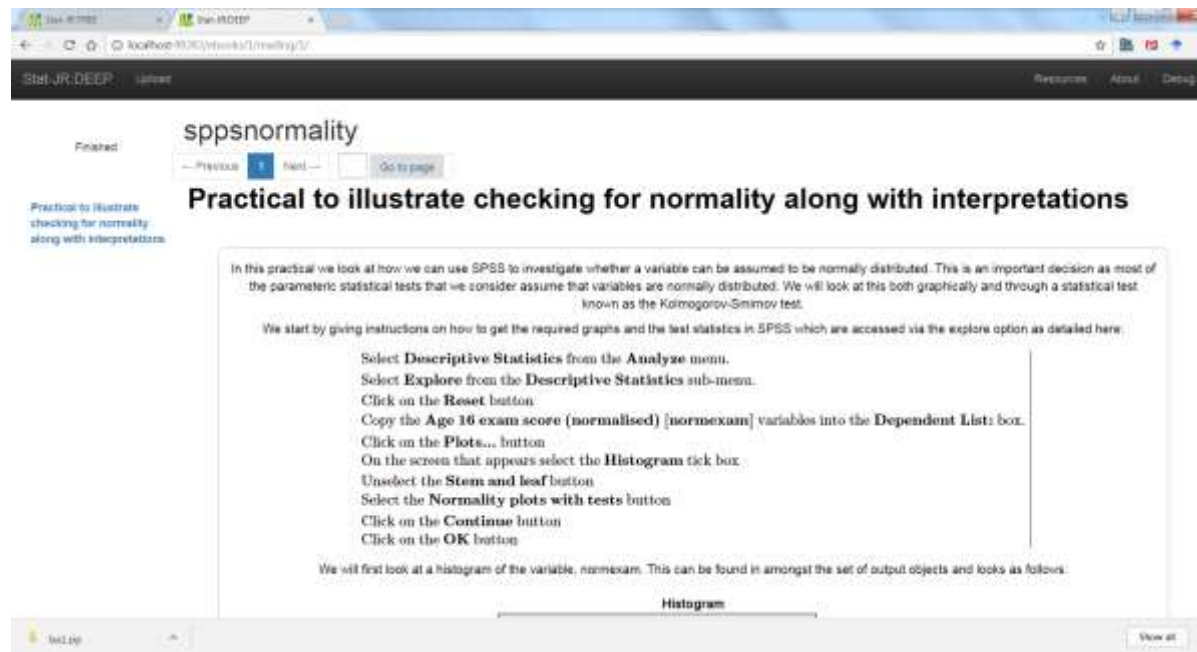
eBook writer

The other option is the instantly combined object that does the combining work for the QM teacher but is less customisable:



eBook writer

Finally in the eBook (DEEP) system we can see the final product and print to PDF file.




Current Practical 1

- Covers descriptive statistics and graphing of categorical variables
- Requires as input 2 categorical variables
- Tabulates the variables and explains percentages and missing data
- Plots bar charts for each variable separately
- Plots clustered bar charts for one variable clustering on the other
- Repeats this using percentages instead of counts



Current Practical 1 outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 6 7 8 ... 12 13 Next → Go to page

Practical 1

Descriptive statistics for categorical variables practical

Welcome to the first descriptive statistics practical in which we will look at how to investigate categorical variables in SPSS. Categorical variables can take only predefined values (or categories) and can be of two types - nominal and ordinal. For nominal variables each category has a name but there is no natural order to the categories for example categories could be colour choice where it might not make sense to describe the colour red to somehow come between blue and green in an ordering. Ordinal categories by contrast do have a natural ordering for example the categories could be the opinions on an issue with the category neutral being naturally found in an ordering between agree and disagree. In this practical we will simply look at summaries and plots of categorical variables that apply equally to both nominal and ordinal variables. In this practical we will look at two categorical variables, the first variable is **female** which has 2 categories: *No*, and *Yes*. The second variable is **freqread** which has 5 categories: *Never*, *Less than once a month*, *At least once a month*, *At least once a week*, and *Most days*. We will begin by looking at how frequent the various categories for **female** are by choosing the following in SPSS:

- Select **Frequencies** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- Copy the **Child is female(female)** variable into the **Variable(s)** box.
- Click on the **OK** button to produce the tables as shown below.

Two tables will now appear in the output window. The first is really just a count of number of observations in the dataset and how many are missing for the variable, **female**

Statistics

Child is female

N	Valid	120
	Missing	0

In this case we see there are no missing observations for the **female** variable and all 120 values are valid. In the second table we can see a list of the different categories in the data and their frequencies:

Child is female

	Frequency	Percent	Valid Percent	Cumulative Percent
Valid No	59	49.2	49.2	49.2
Yes	61	50.8	50.8	100.0
Total	120	100.0	100.0	

This second table has 5 columns which we will now describe. The first column simply gives the categories for the variable, **female** so that we can tell what each row refers to. In the second column headed *Frequencies* we get the actual numbers of occurrences of the variable and so we see that there are 59 occurrences of *No*, and 61 occurrences of *Yes*. This is a useful summary as we can compare the counts within the dataset. We might however not simply be interested in this dataset in isolation and so it is often useful to convert these counts into percentages and this is done in column 3. Here we see that we have 49.2 percent of observations are in category *No*, and 50.8 percent of observations are in category *Yes*. We have no missing data in this variable and so the fourth column simply replicates the third column as all observations are valid. Finally in the fifth column we look at cumulative percentages so that we see that 49.2 are in the first category, and unsurprisingly 100 percent of valid observations are in one of the 2 categories.

We can now repeat this for our second variable, **freqread** and to do this in SPSS we do the following:

- Select **Bar...** from the **Legacy Dialogs** submenu available from the **Graphs** menu



Current Practical 1 outputs

Stat-JR:DEEP

localhost:50517/vbooks/12/reading/1/

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 6 7 8 ... 12 13 Next → Go to page

Child is female

Here we see that each category of **freqread** is given a bar of a different colour so that one can by eye compare each category of **freqread** for different categories of **female**. So for example we see when How often child reads for pleasure takes value Never, then there are 11 observations where Child is female takes value No, and 1 observations where Child is female takes value Yes.

As with the frequency tables earlier it is often easier to look at percentages than raw counts to make comparisons and so in this case if we want to see whether the distribution of **female** is the same for different categories of variable **freqread** we can do the following in SPSS:

- Once again select **Bar...** from the **Legacy Dialogs** submenu available from the **Graphs** menu.
- Keep the choices as **Clustered and Summaries for groups of cases** before clicking on **Define**.
- Keep the **Child is female[female]** variable in the **Category Axis** box.
- Keep the **How often child reads for pleasure[freqread]** variable in the **Define Clusters by** box.
- Select **% of cases** in the **Bars Represent** choices.
- Click on the **OK** button to produce the graph as shown below.

Child is female	Never	Less than once a month	At least once a month	At least once a week	Most days
No	91.7%	8.3%	0%	0%	0%
Yes	8.3%	25.0%	37.5%	25.0%	25.0%

Here we see that now the bars represent the percentage of each category of **freqread** that are found in each category of **female**. So for example if we look again at the cases where **freqread** takes value Never, then 91.7% percent of observations have **female** taking value No, and 8.3% percent of observations have **female** taking value Yes. This ends our practical.

about



Current Practical 2

- Covers descriptive statistics and graphing of numerical variables
- Requires as input 2 numerical variables
- Shows summary statistics and frequencies (may drop this) for the variables and interprets them
- Plots histograms for each variable separately
- Plots boxplots for each variable separately and explains how they are constructed.



Current Practical 2 outputs

Finished

SPSS practicals 1 - 12 as an eBook

← Previous 1 2 3 4 5 6 7 8 9 10 11 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Practical 2

Descriptives statistics for all variables practical

In this descriptive statistics practical we will expand our investigation of variables to include continuous variables. We will look at how in SPSS we can obtain some summary statistics that describe the distribution of variables both in terms of measures of location and spread. We will also look at how we might summarise these variables graphically. We will begin by looking at how to use SPSS to get summary statistics for our first variable, **esteem**

- Select **Frequencies** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- Copy the **Rosenberg Self-Esteem Scale[esteem]** variable into the **Variable(s)** box.
- Click on the **Statistics** button to go to the statistics screen
- Here we need to select ALL the summary statistics that we are interested in looking at!
- Select **Mean, Median and Mode** from under **Central Tendency**.
- Select **Std. deviation, Variance, Range, Minimum and Maximum** from under **Dispersion**.
- Finally Select **Quartiles** from under **Percentile Values**.
- Click on the **Continue** button to return to the main window
- Click on the **OK** button to produce the tables as shown below.

The first table contains all the summary statistics that we requested for the variable as shown below:

Statistics

Rosenberg Self-Esteem Scale

N	Valid	98
	Missing	24
Mean		12.18
Median		12.00
Mode		11
Std. Deviation		1.947
Variance		3.789
Range		9
Minimum		6
Maximum		15
Percentiles	25	11.00
	50	12.00
	75	14.00



Current Practical 2 outputs

Stat-JR DEEP

localhost:50517/ebooks/12/reading/1/#ebookpage2

Stat-JR DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 **2** 3 4 5 6 7 8 ... 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Histogram

Frequency

Hours per week spent on homework (term time)

Mean = 2.24
Std. Dev. = 2.159
N = 114

- Select **Boxplot** from the **Legacy Dialogs** submenu available from the **Graphs** menu.
- We want to choose **Simple and Summaries of separate variables** from the options here.
- Next click on **Define** to set up the box plot.
- Copy the **Rosenberg Self-Esteem Scale[esteem]** variable into the **Boxes Represent:** box.
- Ignore the rest of the options and click on the **OK** button to produce the graph as shown below.




Current Practical 3

- Covers tabulations of variables
- Requires as input 2 categorical variables
- Produces a simple cross tab of the two variables and explains the numbers
- Covers percentages and shows what the 3 possible percentages in SPSS offer you – row percentages, column percentages and total percentages.



Current Practical 3 outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 6 7 8 ... 12 13 Next → Go to page

Practical 3

Tabulation practical

In this practical we will take a first look at the cross-tabulation options in SPSS. Cross tabulations are useful for summarising the relationships between pairs of categorical variables. Here we will focus on just one pair of variables, **ethnicity** and **fread**. We will begin by performing the simple cross-tabulation using the SPSS instructions shown below:

- Select **Crosstabs...** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- Copy the **Child's ethnic origin(ethnicity)** variable into the **Row(s)** box.
- Copy the **How often child reads for pleasure(fread)** variable into the **Column(s)** box.
- Click on the **OK** button to produce the table as shown below.

Child's ethnic origin * How often child reads for pleasure Crosstabulation

Count

		How often child reads for pleasure					Total
		Never	Less than once a month	At least once a month	At least once a week	Most days	
Child's ethnic origin:	White	10	7	12	24	37	90
	Asian	0	0	1	3	8	12
	Black	0	0	0	3	4	7
	Other	2	0	0	2	5	9
Total		12	7	13	32	54	110

This simple cross tabulation simply looks at the columns in the full dataset and looks at how many times each combination of our two categorical variables appears. So we can see that there are 10 occurrences where **ethnicity** takes value White and **fread** takes value Never. This is out of 90 occurrences in total where **ethnicity** takes value White and 12 occurrences in total where **fread** takes value Never.

From this table we can see that when **ethnicity** takes value White there are more instances of **fread** taking value Never than value Less than once a month. This may be because overall there are more instances of **fread** taking value Never than value Less than once a month. It can be helpful in understanding the individual cell counts to look at how their relative frequencies within a particular row or column compare to overall relative frequencies for all rows and columns. To do this we look at percentage values as well as counts and we will do this first for rows by following the instructions below:

- Select **Crosstabs...** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- The Row and Column variables should be already chosen so click on the **Cells...** button.
- Click on the **Row** tickbox found under the **Percentages** section.
- Click on the **Continue** button to return to the main window.
- Click on the **OK** button to produce the table as shown below.

Child's ethnic origin * How often child reads for pleasure Crosstabulation

	How often child reads for pleasure	
--	------------------------------------	--



Current Practical 3 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 **3** 4 5 6 7 8 ... 12 13 Next → Go to page

24) take value At least once a week and 11.1% (or 7) take value Most days.

This compares with the overall distribution of **freqread** where 10.2% (or 12) take value Never, 5.9% (or 7) take value Less than once a month, 11.0% (or 13) take value At least once a month, 27.1% (or 32) take value At least once a week and 45.8% (or 54) take value Most days. So for example we see that when **ethnicity** takes value White there is a greater percentage (11.1%) taking category Never than on average (10.2%). We can now look at columns instead of rows via the following:

- Select **Crosstabs...** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- The Row and Column variables should be already chosen so Click on the **Cells...** button
- Remove the **Row** tickbox and instead click on the **Column** tickbox under the **Percentages** section.
- Click on the **Continue** button to return to the main window.
- Click on the **OK** button to produce the table as shown below.

Child's ethnic origin * How often child reads for pleasure Crosstabulation

		How often child reads for pleasure					Total
		Never	Less than once a month	At least once a month	At least once a week	Most days	
Child's ethnic origin - White	Count	10	7	12	24	37	90
	% within How often child reads for pleasure	83.3%	100.0%	82.3%	75.0%	69.0%	76.3%
Asian	Count	0	0	1	3	8	12
	% within How often child reads for pleasure	0.0%	0.0%	7.7%	9.4%	14.0%	10.2%
Black	Count	0	0	0	3	4	7
	% within How often child reads for pleasure	0.0%	0.0%	0.0%	9.4%	7.4%	5.9%
Other	Count	2	0	0	2	5	9
	% within How often child reads for pleasure	16.7%	0.0%	0.0%	6.3%	9.3%	7.6%
Total	Count	12	7	13	32	54	118
	% within How often child reads for pleasure	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

This new table contains the same information as the last table but this time the percentages underneath each count are column rather than row percentages. These percentages represent what percentage of the observations in each column are found in each cell. So we see that when **freqread** takes value Never then for **ethnicity** 83.3% (or 10) take value White, 0% (or 0) take value Asian, 0% (or 0) take value Black and 16.7% (or 2) take value Other.

This compares with the overall distribution of **ethnicity** where 76.3% (or 90) take value White, 10.2% (or 12) take value Asian, 5.9% (or 7) take value Black and 7.6% (or 9) take value Other. So for example we see that when **freqread** takes value Never there is a greater percentage (83.3%) taking category White than on average (76.3%). Finally for completeness we can look at total percentages as follows:

- Select **Crosstabs...** from the **Descriptive Statistics** submenu available from the **Analyze** menu.
- The Row and Column variables should be already chosen so Click on the **Cells...** button
- Remove the **Column** tickbox and instead click on the **Total** tickbox under the **Percentages** section.
- Click on the **Continue** button to return to the main window.
- Click on the **OK** button to produce the table as shown below.



Current Practical 4

- Covers testing if a variable is normally distributed
- Requires just one variable as input
- Plots a histogram of the variables and interprets whether the histogram appears to be skewed or not.
- Then performs the Kolmogorov-Smirnov and Shapiro-Wilks tests and explain output
- Finally shows a QQ plot and explains what to look for in terms of normality.



Current Practical 4 outputs

Stat-JR.DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 **4** 5 6 7 8 ... 12 13 Next → Go to page

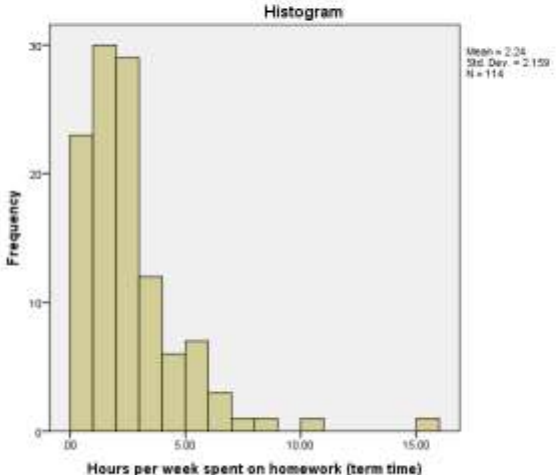
This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality**
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

The start by getting instructions on how to get the response graphs and the test statistics in SPSS. Then the instructions for the output objects are returned here:

- Select **Descriptive Statistics** from the **Analyze** menu.
- Select **Explore** from the **Descriptive Statistics** sub-menu.
- Click on the **Reset** button.
- Copy the **Hours per week spent on homework (term time)[hours_hwk]** variables into the **Dependent List:** box.
- Click on the **Plots...** button.
- On the screen that appears select the **Histogram** tick box.
- Unselect the **Stem and leaf** button.
- Select the **Normality plots with tests** button.
- Click on the **Continue** button.
- Click on the **OK** button.

We will first look at a histogram of the variable, **hours_hwk**. This can be found in amongst the set of output objects and looks as follows:



Frequency

Hours per week spent on homework (term time)

Mean = 2.24
Std. Dev. = 2.159
N = 114

Ideally for a normal distribution this histogram should look symmetric around the mean of the distribution, in this case 2.2405. This distribution appears to be significantly skewed to the right. We will next look at a statistical test to see if this backs up our visual impressions from the histogram.



Current Practical 4 outputs

Stat-JR.DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 **4** 5 6 7 8 9 10 11 12 13 Next → Go to page

Tests of Normality

	Kolmogorov-Smirnov ^a			Shapiro-Wilk		
	Statistic	df	Sig.	Statistic	df	Sig.
Hours per week spent on homework (term time)	.264	114	.000	.749	114	.000

a. Lilliefors Significance Correction

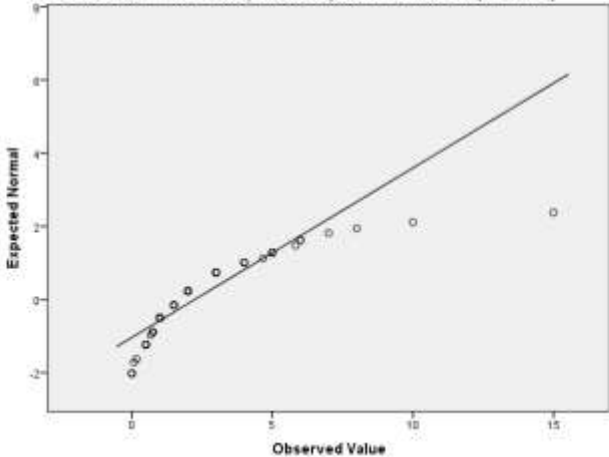
Both the Kolmogorov Smirnov and Shapiro Wilk tests produce test statistics that are used (along with a degrees of freedom parameter) to test for normality. Here we see that the Kolmogorov Smirnov statistic takes value .264 whilst the Shapiro-Wilks statistic takes value .749. Both tests have the same degrees of freedom which equals the number of data points, namely 114.

Although SPSS quotes the p value (quoted under Sig. for Kolmogorov Smirnov) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the variable follows a normal distribution. The Shapiro-Wilks p value agrees with the Kolmogorov-Smirnov p value that the null hypothesis can be rejected.

Although both these statistics tell the researcher whether the distribution followed by a variable is statistically significantly different from a normal distribution one should take care in not overinterpreting such findings. Significance will be strongly effected by the number of observations and so only a small discrepancy from normality will be deemed significant for very large sample sizes whilst very large discrepancies will be required to reject the null hypothesis for small sample sizes.

To complete our practical on checking for normality SPSS also produces a Quantile-Quantile (or QQ) plot that can be seen below.

Normal Q-Q Plot of Hours per week spent on homework (term time)




Current Practical 5

- Covers independent samples (2 samples) t tests
- Requires as input 1 response variable and 1 grouping variable (note template can do several at once)
- Practical runs the 2-sample t test in SPSS
- Then displays the summary statistics and the test tables that SPSS produces and interprets this output appropriately
- Practical also gives interpretive text as to how one would report the findings
- We intend to extend practical to include an error bar plot.



Current Practical 5 outputs



SPSS practicals 1 - 12 as an eBook

Finished

— Previous 1 2 3 4 **5** 6 7 8 ... 12 13 Next — Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

Practical 5 - Independent Samples t test

In this practical we are going to investigate how to perform a 2-sample t-test using SPSS. A 2-sample t-test is used when we have an interval or ratio level variable measured for all observations in 2 groups and we want to test if the mean of this variable is different in the two groups. The test assumes that the variable is normally distributed in both groups. To run a single test in SPSS requires that your dataset has one column containing the variable to be tested and another column of the same length with the groups identified. We will here focus on the case where the groups are identified by specific group codes although it is possible to also use a second continuous variable for groups and specify ranges of this variable that correspond to groups. It is also possible to test for group differences in several variables simultaneously.

Here however we test only one variable, **readscore**

Below you will see instructions to perform the t test in SPSS. If you follow the instructions you will see the two tabular outputs that are embedded in the explanations below.

- Select **Compare Means** from the **Analyze** menu.
- Select **Independent-Sample T Test...** from the **Compare Means** sub-menu.
- Click on the **reset** button.
- Copy the **Reading test score[readscore]** variables into the **Test Variable(s):** box.
- Copy the **Child is female[female]** variable into the **Grouping Variable:** box.
- Click on the **Define Groups...** button.
- Click on the **Use specified values** button.
- Type **0** into the **Group 1** box.
- Type **1** into the **Group 2** box.
- Click on the **Continue** button.
- Click on the **OK** button.

The first SPSS output table contains summary statistics for all the variables considered split by group and can be seen below:


Group Statistics					
	Child is female	N	Mean	Std. Deviation	Std. Error Mean
Reading test score	No	53	108.42	33.600	4.615
	Yes	55	118.04	30.338	4.091

The summary statistics table contains 5 columns and 1 row for each group in each variable to be tested. After the first column which contains the name of each dependent variable and group categories we next see the number of valid observations in each group, i.e. cases with a valid value of **readscore**. Here for the group indexed by **female = No**, we have 53 observations and for **female = Yes**, there are 55 observations. Next we see that the mean of the variable **readscore** for the group with **female = No** is 108.42 whilst for the group with **female = Yes** it is 118.04. Hence the group with **female = Yes** has the bigger mean and the test will now establish if this difference is statistically significant.

In the next column we see the standard deviations for **readscore** variable in the two groups. As we will see in the next table there are two versions of the test depending on whether the variability (and therefore the standard deviations) in the two groups can be assumed equal or not. In this case the standard deviation of **readscore** when **female = No** is 33.600 whilst for **female = Yes** it is 30.338. So there is slightly more variability among **female = No** than **female = Yes**. But is the difference big enough to violate the assumption of equal variances? In the final column are the standard errors of the means for each group. Whilst the standard deviations measure the variability in the data the standard errors of the means measures how confident we are in the estimates of the means. As we collect more data the standard error of the mean gets smaller as we get more confident in the mean estimate and in fact the formula for the standard error of the mean = standard deviation / square root of N. In this case the standard error of the mean for **readscore** when **female = No** is



Current Practical 5 outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 **5** 6 7 8 ... 12 13 Next → Go to page

among **female = No** than **female = Yes**. But is the difference big enough to violate the assumption of equal variances? In the final column are the standard errors of the means for each group. Whilst the standard deviations measure the variability in the data the standard errors of the means measures how confident we are in the estimates of the means. As we collect more data the standard error of the mean gets smaller as we get more confident in the mean estimate and in fact the formula for the standard error of the mean = standard deviation / square root of N. In this case the standard error of the mean for **readscore** when **female = No** is 4.615 whilst for **female = Yes** it is 4.091.

The second SPSS output table contains details of the test itself and can be seen below:

Independent Samples Test

		Levene's Test for Equality of Variances		t-test for Equality of Means						
		F	Sig.	t	df	Sig. (2-tailed)	Mean Difference	Std. Error Difference	95% Confidence Interval of the Difference	
									Lower	Upper
Reading test score	Equal variances assumed	1.131	.290	-1.563	106	.121	-9.621	6.156	-21.825	2.583
	Equal variances not assumed			-1.563	103.992	.122	-9.621	6.167	-21.851	2.600

The above table in fact captures two tests.

The first test is known as Levene's test and tests the assumption that the variability in the two groups is equal and if this is not the case then a slightly different t-test is performed. The two rows of numbers in the right of the table correspond to the two different versions of the test. We decide which one to report depend on the p-value (labelled Sig in SPSS) of Levene's test. The null hypothesis of Levene's test is that the variances (and SDs) in the two groups are equal in the population. When this is true, the Levene's test statistic follows a standard statistical distribution called an F distribution. Higher values of the F-statistic are associated with a lower likelihood that the sample did indeed come from a population in which the null hypothesis is true. In this case the F statistic has value 1.131 and SPSS calculates the corresponding p value for this statistic which is .290. Higher values of the F-statistic are associated with a lower likelihood that the sample did indeed come from a population in which the null hypothesis is true. This p value being greater than 0.05 does not give us strong enough evidence to reject the hypothesis of equal variances so we use the top row of numbers to the right of the table.

So if we now look at the top row of numbers we will start with the column headed mean difference. Here we see the value -9.621. If you look back at the summary statistics table this value is calculated by subtracting one mean from the other. Next to the mean difference is the standard error of the difference. This here has the value 6.156 and is calculated via a formula from the standard errors of each group and their respective sample sizes. Working back to the start of the row, the column entitled t is the statistic used in the t test and t like F is a standard statistical distribution. The t statistic is calculated simply by dividing the mean difference by its standard error so $-9.621 / 6.156 = -1.563$. Next to t is a column labelled df which stands for degrees of freedom and is a parameter used to choose the correct t distribution for the statistic. When we can assume equal variances then the degrees of freedom equal the number of observations - 2 (108) as we have used 2 degrees of freedom in estimating 2 means.

The column labelled "Sig (2-tailed)" contains a test of the null hypothesis that the means of the **readscore** variable in the two groups are the same. By default, the two-tailed test reported uses a non-directional alternative hypothesis. It gives the probability that the data in the sample came from a population in which the group means are truly equal, when either a positive or a negative difference between sample group means is evidence against that null hypothesis. To conduct a one-tailed test, in which the alternative hypothesis specifies a particular direction to the difference, we would simply halve the p-value provided by SPSS.

We can reject the null if there is sufficient evidence that the mean of Group 1 is either higher or lower than the mean of Group 2. SPSS looks Looking up the t statistic in the appropriate table gives the associated p value associated with the calculated t-statistic and degrees of freedom. Looking up the t statistic in the appropriate table gives the associated p value, in this case .121. Here we see that the p value is greater than 0.05 and therefore we cannot reject the null hypothesis that the two groups have the same means. Finally we can see the 95% confidence interval for the difference which runs from -21.825 to 2.583. Here we see it contains the value 0 backing up our failure to reject the null hypothesis.

In conclusion, we could report this to a reader as follows: Mean test scores were higher among **female = Yes** (N=55, M=118.04, SD=30.338) than **female = No** (N=53, M=108.42, SD=33.600). Levene's test was not able to reject the null hypothesis of equal variances between the two groups (F=1.131, p=.290) so an unadjusted version of the independent samples t-test was chosen. The difference in means (difference = -9.621) was not statistically significant, $t(106) = -1.563$, p=.121.

about



Current Practical 6

- Covers paired t tests
- Requires as input 2 variables to be tested via the paired test
- Practical runs the paired t test in SPSS
- Then displays the summary statistics, correlation and the test tables that SPSS produces and interprets this output appropriately
- Practical also gives interpretive text as to how one would report the findings
- We intend to extend practical to include an error bar plot.



Current Practical 6 outputs

Stat-R:DEEP

localhost:57517/ebooks/12/watmg/1/#ebookpage#

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 **6** 7 8 ... 12 13 Next → Go to page

Practical 6 - Paired t test

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test**
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

In this practical we are going to investigate how to perform a paired t-test using SPSS. A paired t-test is used when we have two interval or ratio level variable measured for all observations in a dataset and we want to test if the means of these variables are different. The test assumes that both the variables are normally distributed. To run a single test in SPSS requires that your dataset has two separate columns containing the two variables to be tested. In this situation we could perform a standard 2 sample t-test by reshaping the two variables into one long variable with an accompanying indicator column that defines which original variable each observation refers to but this would be a less efficient test as it does not take account of the paired nature of the data. Although the SPSS dialog box for the paired t test can perform several tests at once we will here focus on one test and will test the two variables, **conduct_sdq** and **hyper_sdq**.

Below you will see instructions to perform the paired t test in SPSS. If you follow the instructions you will see the three tabular outputs that are embedded in the explanations below.

- Select **Compare Means** from the **Analyze** menu.
- Select **Paired-Samples T Test...** from the **Compare Means** sub-menu.
- Click on the **Reset** button.
- Copy the **SDQ Conduct Disorder Sub-scale[conduct_sdq]** variable into the **Variable1:** box for Pair 1.
- Copy the **SDQ Hyperactivity Sub-scale[hyper_sdq]** variable into the **Variable2:** box for Pair 1.
- Click on the **OK** button.

The first SPSS output table contains summary statistics for the two variables to be compared and can be seen below:

	Mean	N	Std. Deviation	Std. Error Mean
Pair 1 SDQ Conduct Disorder Sub-scale	.60	119	1.323	.121
SDQ Hyperactivity Sub-scale	2.22	119	2.611	.239

The summary statistics table contains 5 columns and 1 row for each of the two variable to be tested. After the first column which contains the name of each variable, next we see that the mean of variable **conduct_sdq** is .60 whilst the mean of variable **hyper_sdq** is 2.22. Hence the variable **hyper_sdq** has the bigger mean and the t test will now establish if this difference is statistically significant. We next see the number of valid observations for each variable, i.e. cases with valid values for both **conduct_sdq** and **hyper_sdq**. Here we have 119 valid observations for both variables.

In the next column we see the standard deviations for **conduct_sdq** and **hyper_sdq**. In this case the standard deviation of **conduct_sdq** is 1.323 whilst for **hyper_sdq** it is 2.611. So there is slightly more variability for **hyper_sdq** than **conduct_sdq**. In the final column are the standard errors of the means for each group. Whilst the standard deviations measure the variability in the data the standard errors of the means measures how confident we are in the estimates of the means. As we collect more data the standard error of the mean gets smaller as we get more confident in the mean estimate and in fact the formula for the standard error of the mean = standard deviation / square root of N. In this case the standard error of the mean for **conduct_sdq** is .121 whilst for **hyper_sdq** it is .239.


The second SPSS output table contains information on the correlation between the two variables to be compared and can be seen below:

	N	Correlation	Sig.
Pair 1 SDQ Conduct Disorder Sub-scale & SDQ Hyperactivity Sub-scale	119	.627	.000

The correlation between two variables is a single number that describes how related they are to each other. It is represented by a correlation coefficient which is a numerical value to describe the correlation. Correlations lie between -1 and +1 with a positive value meaning that in general that large values of the first variable are more likely to be observed with large values of the second variable and conversely small values



Current Practical 6 outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 **6** 7 8 ... 12 13 Next → Go to page

The second of 33 output table contains information on the correlation between the two variables to be compared and can be seen below.

Paired Samples Correlations

	N	Correlation	Sig.
Pair 1. SDQ Conduct Disorder Sub-scale & SDQ Hyperactivity Sub-scale	119	.627	.000

The correlation between two variables is a single number that describes how related they are to each other. It is represented by a correlation coefficient which is a numerical value to describe the correlation. Correlations lie between -1 and +1 with a positive value meaning that in general that large values of the first variable are more likely to be observed with large values of the second variable and conversely small values of the first variable are more likely to be observed with small values of the second variable. In the case of a negative correlation the opposite is true and large values of the first variable are more likely to be observed with small values of the second variable and conversely small values of the first variable are more likely to be observed with large values of the second variable. A correlation of 0 means there is no (linear) relationship between the variables. We will cover correlations in more detail elsewhere but here SPSS is giving out a form of correlation known as a Pearson correlation. Here we see that the correlation between **conduct_sdq** and **hyper_sdq** is .627 it is helpful to look at the correlation between the two variables here as typically a paired t-test is more useful than a 2-sample t-test when there is a positive correlation between the two variables as is the case here. SPSS also gives out a p value which describes whether the correlation is statistically significantly different from zero. Here we see that the p value is less than 0.05 and therefore we can reject the null hypothesis that the correlation is zero.

The third SPSS output table contains details of the t test itself and can be seen below:

Paired Samples Test

	Paired Differences						t	df	Sig. (2-tailed)
	Mean	Std. Deviation	Std. Error Mean	95% Confidence Interval of the Difference					
				Lower	Upper				
Pair 1. SDQ Conduct Disorder Sub-scale - SDQ Hyperactivity Sub-scale	-1.622	2.058	.169	-1.996	-1.248	-8.595	118	.000	

The above table describes the paired t test. If we now look at the row of numbers we will start with the column headed mean (underneath paired differences). Here we see the value -1.622. If you look back at the summary statistics table this value is calculated by subtracting one mean from the other. Next to the mean is the standard deviation (of the differences) which has value 2.058. If we have two positively correlated variables then this standard deviation will typically be smaller than the standard deviations of the two variables. Next up is the standard error of the mean (of the differences). This here has the value .169 and is simply the standard deviation divided by the square root of the sample size. Moving forwards two columns, the column entitled t is the statistic used in the t test and t is a standard statistical distribution. The t statistic is calculated simply by dividing the mean difference by its standard error so $-1.622 / .169 = -8.595$. Next to t is a column labelled df which stands for degrees of freedom and is a parameter used to choose the correct t distribution for the statistic. Here the degrees of freedom equal the number of observations - 1 (118) as we have used 1 degrees of freedom in estimating the mean difference.

The column labelled "Sig (2-tailed)" contains a test of the null hypothesis that the means of the two variables (**conduct_sdq** and **hyper_sdq**) are the same. By default, the two-tailed test reported uses a non-directional alternative hypothesis. It gives the probability that the data in the sample came from a population in which the variable means are truly equal, when either a positive or a negative difference between sample means is evidence against that null hypothesis. To conduct a one-tailed test, in which the alternative hypothesis specifies a particular direction to the difference, we would simply halve the p-value provided by SPSS.

We can reject the null if there is sufficient evidence that the mean of **conduct_sdq** is either higher or lower than the mean of **hyper_sdq**. SPSS looks up the t statistic in the appropriate table gives the associated p value associated with the calculated t-statistic and degrees of freedom. Looking up the t statistic in the appropriate table gives the associated p value, in this case .000. Here we see that the p value is less than 0.05 and therefore we can reject the null hypothesis that the two groups have the same means. Finally we can see the 95% confidence interval for the difference which runs from -1.996 to -1.248. Here we see it does not contain the value 0 backing up our rejection of the null hypothesis.

In conclusion, we could report this to a reader as follows: Mean values were compared for 2 variables with sample size 119. The mean was higher for variable **hyper_sdq** (M=2.22, SD=2.611) than for variable **conduct_sdq** (M=.60, SD=1.323). The difference in means (difference = -1.622) was statistically significant, $t(118) = -8.595$, $p = .000$.

about




Current Practical 7a and b

- Covers the non-parametric tests for unpaired and paired data
- First sheet covers Mann Whitney test and requires same inputs as 2-sample t test
- Second sheet covers Wilcoxon signed rank test and requires same inputs as paired t tests
- Templates call the appropriate tests in SPSS and then display the output tables with relevant interpretation.



Current Practical 7a outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 5 6 7 8 9 ... 12 13 Next → Go to page

Practical 7a - The Mann Whitney test

In this practical we are going to investigate how to perform a Mann Whitney test using SPSS. A Mann-Whitney test is used when we have an interval or ratio level variable measured for all observations in 2 groups and we want to test if the distribution of this variable is different in the two groups but we are unable to assume normality in both groups. It can also be to compare an ordered categorical variable measured on two groups. It is the non-parametric equivalent of the 2-sample t-test but unlike the t-test it tests differences in the median rather than the mean. The test does not assume any distribution for the variable in either groups. To run a single test in SPSS requires that your dataset has one column containing the variable to be tested and another column of the same length with the groups identified. We will here focus on the case where the groups are identified by specific group codes although it is possible to also use a second continuous variable for groups and specify ranges of this variable that correspond to groups. It is also possible to test for group differences in several variables simultaneously.

Here however we test only one variable, **readscore**

Below you will see instructions to perform the Mann Whitney test in SPSS. If you follow the instructions you will see the two tabular outputs that are embedded in the explanations below.

- Select **Non Parametric Tests** from the **Analyze** menu.
- Select **Legacy Dialogs** from the **Non Parametric Tests** sub-menu.
- Select **2 Independent-Samples...** from the **Legacy Dialogs** sub-menu.
- Click on the **Reset** button.
- Copy the **Reading test score(readscore)** variable into the **Test Variable List:** box.
- Copy the **Child is female(female)** variable into the **Grouping Variable:** box.
- Click on the **Define Groups...** button.
- Type **0** into the **Group 1** box.
- Type **1** into the **Group 2** box.
- Click on the **Continue** button.
- Click on the **Exact...** button.
- On the screen that appears select the **Exact** button.
- Click on the **Continue** button.
- Click on the **OK** button.

The first SPSS output table contains a summary of the rankings for the 2 groups and can be seen below:

Ranks				
	Child is female	N	Mean Rank	Sum of Ranks
Reading test score	No	53	49.30	2613.00
	Yes	55	59.51	3273.00
	Total	108		

The Mann Whitney test works by firstly constructing a ranked list of the observations labelled in their two groups. It will then work from the lowest observation and give that observation rank 1 and the next rank 2 and so on right up to the largest observation which in this case will have rank 108. If there are observations with the same value then they are given the same rank that is an average of the ranks available (for example if three observations have the 9th smallest rank then rather than giving them ranks 9, 10 and 11 respectively they will each be given rank $10 (9+10+11)/3 = 10$).



Current Practical 7a outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 3 4 5 6 7 8 9 10 11 12 13 Next → Go to page

Yes	No	Total
55	53	108
3273.00	2613.00	

The Mann Whitney test works by firstly constructing a ranked list of the observations labelled in their two groups. It will then work from the lowest observation and give that observation rank 1 and the next rank 2 and so on right up to the largest observation which in this case will have rank 106. If there are observations with the same value then they are given the same rank that is an average of the ranks available (for example if three observations have the 9th smallest rank then rather than giving them ranks 9, 10 and 11 respectively they will each be given rank 10 $(9+10+11)/3 = 10$).

The statistics required for the test are therefore constructed and shown in the table. Here we see that for **female** category No we have 53 observations whose total sum of ranks is 2613.00. This results in a mean rank of 49.30. By contrast for **female** category Yes we have 55 observations whose total sum of ranks is 3273.00. This results in a mean rank of 59.51. So **female** category Yes has a larger mean rank than **female** category No and thus tends to take larger values. The Mann Whitney test will now decide on whether this difference in mean ranks is significant or not as is illustrated in the second table.

The second SPSS output table contains details of the test itself and can be seen below:

	Reading test score
Mann-Whitney U	1182.000
Wilcoxon W	2613.000
Z	-.1694
Asymp. Sig. (2-tailed)	.090
Exact Sig. (2-tailed)	.091
Exact Sig. (1-tailed)	.045
Point Probability	.000

The output here consists of test statistics and their significance as calculated in several ways. We are considering the Mann Whitney U statistic and to calculate this we need to consider the sums of the rankings and compare them with what we would expect if there two groups came from the same distribution. We consider each group in turn and work out for each group a U statistic. The formula here is the sum of the ranks - $N \times (N+1)/2$ for each group. So for **female** category No we have $U1 = 2613.00 - 53 \times (53+1)/2 = 1182.0$ and for **female** category Yes we have $U2 = 3273.00 - 55 \times (55+1)/2 = 1733.0$. Here $U1$ is less than $U2$ and we therefore use the value 1182.0 as our test statistic as shown in the table.

A related approach that uses ranks is the Wilcoxon W statistic which is quoted here and is the maximum of the two rank sums but we do not describe it here. The simplest way to use the Mann-Whitney U statistic is to convert it to a normal score by subtracting its mean and dividing by its standard error and that is done in the Z row. Here we see that $Z = -1.694$ and this can be compared with a standard normal distribution to test whether there are significant differences between the groups.

Here we see that the p value (quoted next to Asymp. Sig. (2-tailed)) is .090 which is greater than 0.05 and therefore we cannot reject the null hypothesis that the medians of the two groups are the same. The normal approximation used above is only an approximation to the p value and it is possible to construct the exact p value. This is given in the next row and we see that the exact p value is .091 whilst the asymptotic p value is .090. For completeness the table also gives a p value for a 1-sided test and a point probability but we will ignore these here.


In conclusion, we could report this to a reader as follows:

A comparison of the mean of the distribution of the variable readscore was desired for **female** categories No and Yes but due to the non-normality of the variable a Mann Whitney test was carried out. **female** category Yes ($N = 55$) has a larger mean rank (59.51) than **female** category No ($N = 53$) with mean rank (49.30) and thus tends to take larger values. The Mann Whitney U statistic is 1182.000 which results in an exact p value of .091. This is not significant and we cannot reject the null hypothesis of equal group medians.

about



Current Practical 7b outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Practical 7b - The Wilcoxon Sign Rank test

In this practical we are going to investigate how to perform a Wilcoxon test using SPSS. A Wilcoxon test is used when we have two interval or ratio level variables measured for a set of observations and we want to test if the distribution is different for the two variables but we are unable to assume normality for one or both of the variables. It can also be to compare ordered categorical variables. It is the non-parametric equivalent of the paired t-test but unlike the t-test it tests differences in the median rather than the mean. The test does not assume any distribution for the two variables. To run a single test in SPSS requires that your dataset has two column containing the two variables to be compared. It is possible to test for differences in several pairs of variables simultaneously.

Here however we test only one pair of variables, **conduct_sdq** and **hyper_sdq**

Below you will see instructions to perform the Wilcoxon test in SPSS. If you follow the instructions you will see the two tabular outputs that are embedded in the explanations below.

- Select **Non Parametric Tests** from the **Analyze** menu...
- Select **Legacy Dialogs** from the **Non Parametric Tests** sub-menu.
- Select **2 Related Samples...** from the **Legacy Dialogs** sub-menu.
- Click on the **Reset** button
- Copy the **SDQ Conduct Disorder Sub-scale[conduct_sdq]** variable into the **Test Pairs:** box.
- Copy the **SDQ Hyperactivity Sub-scale[hyper_sdq]** variable into the **Test Pairs:** box.
- Click on the **Exact...** button
- On the screen that appears select the **Exact** button
- Click on the **Continue** button
- Click on the **OK** button

The first SPSS output table contains a summary of the rankings for the 2 variables. Here observations are split into three types depending on whether the value of **conduct_sdq** is bigger than **hyper_sdq** (negative ranks), the value of **hyper_sdq** is bigger than **conduct_sdq** (positive ranks), and finally where both variables take the same value (ties). These can be seen below:

Ranks		N	Mean Rank	Sum of Ranks
SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale	Negative Ranks	6 ^a	22.75	136.50
	Positive Ranks	72 ^b	40.90	2944.50
	Ties	41 ^c		
	Total	119		

- a. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale
- b. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale
- c. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale

The Wilcoxon test works by firstly assigning a sign (or a tie) to the difference between each pair of observations. Here we have worked on **hyper_sdq - conduct_sdq** so that positive ranks are when **hyper_sdq** > **conduct_sdq**. Here we see that there are 6 negative ranks, 72 positive ranks and 41 ties.

Having worked out which observed pairs result in which sign for their difference, the magnitude (excluding the sign) of these differences is calculated and these are then ranked in order (excluding ties). We now see that



Current Practical 7b outputs

Stat-JR.DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Positive Ranks	72 ^a	40.90	2944.50
Ties	41 ^b		
Total	119		

a. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale
b. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale
c. SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale

The Wilcoxon test works by firstly assigning a sign (or a tie) to the difference between each pair of observations. Here we have worked on **hyper_sdq - conduct_sdq** so that positive ranks are when **hyper_sdq > conduct_sdq**. Here we see that there are 6 negative ranks, 72 positive ranks and 41 ties.

Having worked out which observed pairs result in which sign for their difference, the magnitude (excluding the sign) of these differences is calculated and these are then ranked in order (excluding ties). We now see that the total of the ranks for the negative differences is 136.50 resulting in a mean rank of 22.75 while the total of the ranks for the positive differences is 2944.50 resulting in a mean rank of 40.90. Here the mean of the positive ranks is larger than that for negative ranks suggesting that values for **hyper_sdq** are generally larger than for **conduct_sdq**. The Wilcoxon test will now decide whether this difference in mean ranks is significant or not as is illustrated in the second table.

The second SPSS output table contains details of the test itself and can be seen below:

	SDQ Hyperactivity Sub-scale - SDQ Conduct Disorder Sub-scale
Z	-7.042 ^a
Asymp. Sig. (2-tailed)	.000
Exact Sig. (2-tailed)	.000
Exact Sig. (1-tailed)	.000
Point Probability	.000

b. Based on negative ranks

The output here consists of test statistics and their significance as calculated in several ways. We are considering the Wilcoxon statistic which is calculated from the ranks and is not shown explicitly by SPSS but is used to calculate a Z score. Here we see that $Z = -7.042$ and this can be compared with a standard normal distribution to test whether there are significant differences between the groups.

Although SPSS quotes the p value (quoted next to Asympt. Sig. (2-tailed)) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the two groups have the same medians. The normal approximation used above is only an approximation to the p value and it is possible to construct the exact p value. This is given in the next row and we see that the exact p value is .000 whilst the asymptotic p value is .000. The exact p value agrees with the asymptotic p value that the null hypothesis can be rejected. For completeness the table also gives a p value for a 1-sided test and a point probability but we will ignore these here.

In conclusion, we could report this to a reader as follows:

A comparison of the mean of the distribution of the variables **conduct_sdq** and **hyper_sdq** was desired but due to the non-normality of the variables a Wilcoxon signed rank test was carried out. The mean of the positive ranks is larger than that for negative ranks suggesting that values for **hyper_sdq** are generally larger than for **conduct_sdq**. The Wilcoxon Signed rank test results in a Z statistic of -7.042 which results in an exact p value of less than 0.001. This is significant and we can reject the null hypothesis of equal medians for the 2 variables.

about



Current Practical 8

- Covers the chi-squared test and follows on from the tabulations in practical 3
- Requires as input 2 categorical variables
- Tabulates the variables and explains expected counts (note can also cover percentages)
- Performs the chi-squared test on the data and interprets the output
- Needs finishing by adding some reporting guidelines



Current Practical 8 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Practical 8 - The Chi-squared test

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test**
- Practical 9
- Practical 10
- Practical 11
- Practical 12

In this practical we are going to investigate how to perform a chi-squared test using SPSS. A chi-squared test is used when we have two categorical variables measured for all observations in a dataset and we want to test if the variables are related or independent. Independent here means that the category observed for one variable does not depend on the category observed for the other variable. To run a single test in SPSS requires that your dataset has two separate columns containing the two (categorical) variables to be tested.

Below you will see instructions to perform the chisquared test in SPSS. For the chi-squared test we will use SPSS crosstabs options that we have also looked at in our tabulation practical. If you follow the instructions you will see the three tabular outputs that are embedded in the explanations below.

- Select **Descriptive Statistics** from the **Analyze** menu.
- Select **Crosstabs...** from the **Descriptive Statistics** sub-menu.
- Click on the **reset** button.
- Copy the **female** variable into the **Row(s):** box.
- Copy the **freqread** variable into the **Column(s):** box.
- Click on the **Statistics...** button.
- In the window that appears click on the **Chi-square** tickbox.
- Click on the **Continue** button.
- Click on the **Cells...** button.
- In the window that appears:
- Under counts click on the **Expected** tick box to include it.
- Click on the **Continue** button.
- Click on the **OK** button.

The first table looks at which of the observations have non-missing values for both the two variables to be considered.

Case Processing Summary

	Cases					
	Valid		Missing		Total	
	N	Percent	N	Percent	N	Percent
Child is female * How often child reads for pleasure	118	99.3%	2	1.7%	120	100.0%

Here we see that there are 120 observations of which 2 are missing resulting in 118 that can be used in the test.

The second tabular output contains the cross-tabulation of the two variables. Here different levels of **female** are allocated to different rows whilst different values of **freqread** are allocated to different columns and each observation adds to the count of a particular cell in the table.

Child is female * How often child reads for pleasure Crosstabulation

	How often child reads for pleasure				Total
	Never	Less than once a month	At least once a month	At least once a week	
Female					
Male					



Current Practical 8 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test**
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Here we see that there are 120 observations of which 2 are missing resulting in 118 that can be used in the test.

The second tabular output contains the cross-tabulation of the two variables. Here different levels of **female** are allocated to different rows whilst different values of **freqread** are allocated to different columns and each observation adds to the count of a particular cell in the table.

Child is female * How often child reads for pleasure Crosstabulation

			How often child reads for pleasure					Total
			Never	Less than once a month	At least once a month	At least once a week	Most days	
Child is female	No	Count	11	5	8	16	19	59
		Expected Count	6.0	3.5	6.5	16.0	27.0	59.0
Yes		Count	1	2	5	16	35	59
		Expected Count	6.0	3.5	6.5	16.0	27.0	59.0
Total		Count	12	7	13	32	54	118
		Expected Count	12.0	7.0	13.0	32.0	54.0	118.0

Here we see the observed counts for the different combinations of **female** and **freqread**. So for example there are 11 observations where **female** is No and **freqread** is Never. This is out of a total of 59 observations where **female** is No and 12 observations where **freqread** is Never. Under the model of independence we expect to see 6.0 observations where **female** is No and **freqread** is Never. This means there are 5.0 more observations where **female** is No and **freqread** is Never than expected. We can look at similar differences between the other observed and expected counts and the chi-squared test looks at whether these differences are simply chance or statistically significant.

The third tabular output contains the information on test statistics for testing the hypothesis of no association / relationship between the two variables.

Chi-Square Tests

	Value	df	Asymptotic Significance (2-sided)
Pearson Chi-Square	15.052 ^a	4	.005
Likelihood Ratio	18.591	4	.002
Linear-by-Linear Association	14.825	1	.000
N of Valid Cases	118		

a. 2 cells (.00%) have expected count less than 5. The minimum expected count is 3.50.

SPSS gives several tests for significance and we will first focus on the Pearson Chi-square test. This test begins by forming the Pearson test statistic which asymptotically is formed from the observed and expected cell counts. For each cell the difference between the observed and expected counts is found and squared. This positive number is then divided by the expected count to account for different sizes of cells. Having constructed this value for each cell these are summed across all cells to give our test statistic which here is 15.052. This statistic follows a chi-squared distribution under the null hypothesis with degrees of freedom equal to (rows-1)(columns-1) which in this case equals 4. The statistic is then compared with the appropriate chi-squared distribution and this results in an asymptotic (2-sided) p value which has value .005. Here we see that the p value is less than 0.05 and therefore we can reject the null hypothesis that the two variables are independent and there is therefore some relationship between the variables.

The line headed Linear-by-Linear Association refers to a test that considers the categories of the two variables to be ordered and we will not consider this test for this example.

In conclusion, we could report this to a reader as follows: To be written

about



Current Practical 9

- Covers correlations (Pearson, Spearman and Kendall's tau)
- Requires two variables to be compared (numerical or ordered)
- First plots a scatterplot to show relationship
- Also plots histograms and performs normality tests and QQ plots to aid with choice of correlations
- Then performs each of the three correlations in turn
- Could be tweaked to only do one depending on earlier output .



Current Practical 9 outputs

Star-JR.DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 **10** 11 12 13 Next → Go to page

Practical 9

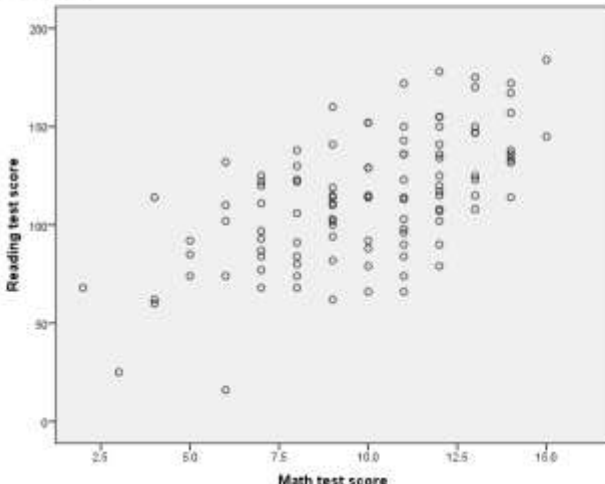
Correlation practical

In this practical we will investigate whether there is a relationship between two variables by looking how correlated they are.

To do this we firstly need to choose two variables and we will begin by simply plotting them so in SPSS:

- 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 **Reading test score[readscore]** variable into the **Y Axis** box.
- Copy the **Math test score[mathscore]** 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 shows a positive correlation between Math test score (X-axis, 0 to 16.0) and Reading test score (Y-axis, 0 to 200). The data points are represented by small circles, showing a clear upward trend where higher math scores generally correspond to higher reading scores.



Current Practical 9 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

We will now finally turn our attention to the main topic of this practical which is the calculation of the correlation between our two variables. SPSS offers several correlation coefficients and we will consider these here in turn. All three are available through the Analyse->Correlate->Bivariate option in SPSS.

- Select **Bivariate...** from the **Correlate** option available from the **Analyse** menu.
- Copy the **Reading test score[readscore]** and the **Math test score[mathscore]** variables into the **Variables** box.
- Click on the **Options** button and Select the **Means and Standard deviations** tick box.
- Click on the **Continue** button to return to main window.
- Click on the **OK** button.

The correlation command will produce two output tables. The first table which we show below simply gives means and standard deviations for the two variables we are comparing.

Descriptive Statistics

	Mean	Std. Deviation	N
Reading test score	113.31	32.195	108
Math test score	9.90	2.855	108

In the next table we see the correlation matrix for the variables we are considering.

Correlations

	Reading test score	Math test score
Reading test score	Pearson Correlation	.610**
	Sig. (2-tailed)	.000
	N	108
Math test score	Pearson Correlation	.610**
	Sig. (2-tailed)	.000
	N	108

** . Correlation is significant at the 0.01 level (2-tailed).

The correlate option can be used for more than two variables simultaneously and will then give all correlations hence the output table is in this matrix format. The table contains three numbers for each possible correlation (including the correlations of variables with themselves which always takes the value 1). For each correlation there is an estimate of the correlation, an accompanying p value and a sample size on which the correlation has been calculated. Here we are interested in the Pearson correlation between **readscore** and **mathscore** which can be found in two places in the table - either in the row for **readscore** and column for **mathscore** or the row for **mathscore** and column for **readscore**.

In this case the correlation takes value .610. This represents a large positive correlation. The correlation is given in the table, along with a significance value and a sample size which in this case is 108. This is the number of observations in which both **readscore** and **mathscore** were observed.

We can test if this correlation is significantly different from zero which will depend on (i) the magnitude of the correlation and (ii) the number of observations on which the correlation is based.

Although SPSS quotes the p value (quoted under Sig. (2-tailed)) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the correlation is 0.



Current Practical 10

- Covers simple linear regression
- Requires as input 2 numerical variables
- Shows summary statistics for both variables with interpretation
- Next does a scatterplot to look at relationship
- Then runs the regression in SPSS and interprets the tabular outputs
- Shows residuals and a histogram of them and a scatterplot of residuals against fitted values
- Finishes with a scatterplot this time with the regression line superimposed.



Current Practical 10 outputs

Stat-JR-DEEP

localhost:50517/ebooks/12/reading/1/#ebookpage11

Stat-JR-DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Practical 10

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. In this practical we will consider **readscore** as our response variable and **mathscore** as our 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 **Reading test score[readscore]** and **Math test score[mathscore]** 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:

	N	Minimum	Maximum	Mean	Std. Deviation
Reading test score	108	16	184	113.31	32.195
Math test score	108	2	15	9.90	2.855
Valid N (listwise)	108				

Here we see a row in the table for each variable. **readscore** is the response variable and takes values between 16 and 184 with a mean of 113.31. **mathscore** is the predictor variable and takes values between 2 and 15 with a mean of 9.90. We can next plot these variables against each other following instructions below:

- Select **ScatterDot** 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 **Reading test score[readscore]** variable into the **Y Axis** box.
- Copy the **Math test score[mathscore]** 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:



Current Practical 10 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 **11** 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

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

Model	Sum of Squares	df	Mean Square	F	Sig.
1. Regression	41268.951	1	41268.951	62.819	.000 ^a
Residual	69636.345	106	656.947		
Total	110905.296	107			

a. Predictors: (Constant), Math test score

The ANOVA table is used to look at the significance of the regression model as a whole and is used in SPSS for many models to show the significance of different terms in the model. An Analysis of Variance (ANOVA) basically partitions the variability in the data (which it measures in terms of sums of squares) into different pieces.

Here we see in the sum of squares column that the total sum of squares is 110905.296 and can be split into a sum of squares explained by the regression (41268.951) and that not explained which is known as the residual sum of squares (69636.345). Interestingly R-squared that we saw in the earlier table can be calculated by looking at the ratio of the regression SS to the total SS i.e. $41268.951/110905.296=0.372$.

These sums of squares have associated degrees of freedom (df) with for the total sum of squares the df being the number of observations - 1 (107) due to fitting a mean to the data. The regression sum of squares has $df = 1$ to account for 1 predictor. The residual df is then the difference between the total df and the regression $df = 106$. The next column is the mean squares (sums of squares adjusted for dfs) which are used to construct a test statistic, F which is in the next column. Here we see that F takes value 62.819 and can be used to test the null hypothesis that there is no significant regression. To do this it needs to be compared with an F distribution with 1 and 106 degrees of freedom. This test results in a p value that is given in the Sig. column. Although SPSS quotes the p value (quoted under Sig.) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that there is no regression.

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

Model	Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B	
	B	Std. Error	Beta			Lower Bound	Upper Bound
1. (Constant)	45.216	8.939		5.058	.000	37.494	62.939
Math test score	6.880	.360	.610	7.926	.000	5.159	8.601

This table gives the most interesting information about the regression model. We begin with the coefficients that form the regression equation. The regression intercept takes value 45.216 and is the value of the regression line when **mathscore** takes value 0. The regression slope takes value 6.880 and is the amount by which we predict that **readscore** changes for an increase of 1 in **mathscore**.

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 .610 which represents the predicted change in **readscore** for an increase of 1 standard deviation in **mathscore**.

To test for the significance of the coefficients we need to form test statistics which are reported under the t column and are simply $B / \text{Std. Error}$. For the slope the t statistic is 7.926 and this value can be compared with a t distribution to test the null hypothesis that the slope is 0. We can see the resulting p value for the test under the Sig. column. Although SPSS quotes the p value (quoted under Sig.) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. 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 5.058 and this value can be compared with a t distribution to test the null hypothesis that the intercept is 0. We can see the resulting p value for the test under the Sig. column. Although SPSS quotes the p value (quoted under Sig.) as 0 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the intercept is zero.



Current Practical 10 outputs

Stat-JR-DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Std. Predicted Value	-2.767	1.767	.000	1.000	108
Std. Residual	-2.750	2.063	.000	.995	108

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 requested that standardized residuals were saved this has resulted in an additional variable being stored in the dataset named **ZRE_1** at the end of the existing variables. We can use this variable to create some residuals plot 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 diagnostics** 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 as shown below.

Histogram
Dependent Variable: Reading test score

Mean = 3.20E-10
Std. Dev. = 0.995
N = 108

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 diagnostics** available from the **Graphs** menu.
- Select **Simple Scatter** and click on **Define** to bring up the **Simple Scatterplot** window.



Current Practical 10 outputs

Finished

SPSS practicals 1 - 12 as an eBook

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

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 would like to superimpose the regression line onto the scatterplot we drew earlier of the response against the predictor. 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 click on the red x to close the **Chart Editor** window and the graph in the output window will now have the forced 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



Current Practical 11

- Covers Analysis and Variance (ANOVA)
- Requires as input 1 numerical response variable and 1 categorical predictor
- Starts with a boxplot of the response for each category
- Next shows the equivalent error bar plot
- Shows descriptives for each category
- Performs the ANOVA and shows output tables with interpretation
- Performs multiple comparison tests and shows an estimated marginal means plots.
- Finishes by identifying homogeneous subsets.



Current Practical 11 outputs



SPSS practicals 1 - 12 as an eBook

Finished


← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Practical 11

In this practical we will investigate how we model the influence of a categorical predictor on a continuous response. In this case we are interested in the effect of **freqread** on our response variable **readscore**. **freqread** has 5 categories: Never, Less than once a month, At least once a month, At least once a week, Most days. We will begin to look at this relationship graphically and look at the distribution of **readscore** separately for each category and a good way to do this is via a box plot. To do this in SPSS:

- Select **Boxplot** from the **Legacy Dialogs** submenu of the **Graphs** menu.
- Select **Simple** and **Summaries for groups of cases** and click on the **Define** button.
- Transfer the **Reading test score[readscore]** variable to the **Variable** box.
- Transfer the **How often child reads for pleasure[reqread]** variable to the **Category Axis** box.
- Click on the **OK** button.

This will produce a table detailing the numbers of valid observations which we do not show here and then a plot with a box for each category as shown below:



Recall that a boxplot uses the 1st and 3rd quartiles to form the box with the median presented as a line in the middle of the box. In this case the medians are for category Never=89.00, Less than once a month=114.00, At least once a month=114.00, At least once a week=100.00 and for Most days=120.00 and we see that category Most days has the highest median whilst category Never has the lowest median.



Current Practical 11 outputs

Stat-JR.DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11**
- Practical 12

- Select **Error Bar** from the **Legacy Dialogs** submenu of the **Graphs** menu.
- Select **Simple and Summaries for groups of cases** as for the boxplot and click on the **Define** button.
- Transfer the **Reading test score[readscore]** variable to the **Variable** box.
- Transfer the **How often child reads for pleasure[freeread]** variable to the **Category Axis** box.
- Click on **OK**.

The graph will appear as shown below:

How often child reads for pleasure	Mean Reading Test Score	95% CI Lower Bound	95% CI Upper Bound
Never	85.0	58.0	112.0
Less than once a month	110.0	78.0	140.0
At least once a month	105.0	88.0	125.0
At least once a week	110.0	100.0	120.0
Most days	123.0	115.0	132.0

Here we are now plotting the means rather than medians and the error bars represent 95% confidence intervals for the different groups. The highest mean is 123.04 for category Most days whilst the lowest mean is 85.04 for category Never. Several of the confidence intervals (Never vs Less than once a month, Less than once a month vs At least once a month, At least once a month vs At least once a week, At least once a week vs Most days) do not overlap so we anticipate that **freeread** will have a significant effect on **readscore**. We have mentioned summary statistics here (medians and means) and we can access these via the **Explore** option as follows:

- Choose **Explore** from **Descriptives** within **Analyze**.
- Add **Reading test score[readscore]** to **Dependent list**.
- Add **How often child reads for pleasure[freeread]** to **Factor list**.
- Click on **OK**.



Current Practical 11 outputs

Star-JR DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Reading test score

F	df1	df2	Sig.
.658	4	101	.628

Tests the null hypothesis that the error variance of the dependent variable is equal across groups.

The Levene's test is used to test one of the underlying assumptions of the ANOVA which is the homogeneity of variances i.e. that the residual variances are equal in each group. This test requires a test statistic that has value here .658 and under the hypothesis of equal variances follows an F distribution with 4 and 101 degrees of freedom where 4 is the number of categories - 1 and 101 is the number of observations - the number of categories. Here we see the p value is .628 which is greater than 0.05 and therefore we cannot reject the Null hypothesis and so we are able to assume equal variances and proceed with the ANOVA. The ANOVA itself is described by the ANOVA table given below.

Next Tests for Between Subject

Tests of Between-Subjects Effects

Reading test score

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	12873.636 ^a	4	3218.409	3.322	.013
Intercept	702543.583	1	702543.583	725.161	.000
fread	12873.636	4	3218.409	3.322	.013
Error	97849.807	101	968.810		
Total	1472159.000	106			
Corrected Total	110723.443	105			

a. R Squared = .116 (Adjusted R Squared = .081).

The above table gives all the information required for us to decide if **fread** is a significant predictor of **readscore**. We will here go through the table column by column to explain how the ANOVA works. SPSS gives rather a lot of rows in its ANOVA tables largely because it also allows one to test the intercept term which we are less interested in here. So we will begin with the Type III Sum of Squares (SS) column and for the row Corrected total we see the value 110723.443. This value is calculated by for each observation taking the value for **readscore** and subtracting the mean of **readscore**. These values are then squared and their sum is the value 110723.443 we see in the table. (Note the value 1472159.000 in the Total row is calculated similarly but without subtracting the mean from each observation). Next the value in row **fread** is calculated by working out the mean of **readscore** for each category in **fread**. We take these category means and subtract the overall mean from them and again square and sum them to give the value 12873.636. The Corrected Model row you will see in the one way ANOVA simply repeats the **fread** row. Finally the Error row sum of squares is calculated by subtracting the **fread** SS value from the Corrected total SS value i.e. $97849.807 = 110723.443 - 12873.636$, which is effectively the sums of squares not explained by **fread**. For **fread** to be a significant predictor of **readscore** we hope its SS value is large relative to the Error SS but these numbers are based on different sample sizes and so we first need to adjust them to reflect this so the next column is the degrees of freedom (df) column. Here we see we have 106 total degrees of freedom which represents the number of observations but 105 corrected total df as we lose one by estimating the mean. For **fread** we have 4 df which is the number of categories - 1 again losing one as if we knew the mean and all but one of the category means we could calculate the last one. Finally the Error df is 101 which again is calculated by subtraction i.e. $101 = 105 - 4$. We next use the df to adjust the SS into Mean Squares (MS) and so MS for **fread** is $SS \text{ for } \text{fread} \text{ divided by } df \text{ for } \text{fread}$ which means $3218.409 = 12873.636 / 4$. Similarly for the Error column we have an MS of 968.810 . These two mean squares are now on the same scale and so we can look at their relative sizes by taking their ratio so $F = 3.322 = 3218.409 / 968.810$. This test statistic follows an F distribution with 4 and 101 degrees of freedom and equates to a p value of .013 given in the Sig. column. This p value is less than 0.05 and so we can reject the null hypothesis and we find that **fread** is a significant predictor of **readscore**.

Our next step after establishing whether or not there are significant differences in **readscore** for the different categories of **fread** is to do a post-hoc multiple comparisons test. We selected three different tests and we will look at these in turn. The first we will look at is Tukey's HSD (honest significant difference).



Current Practical 11 outputs

Stat-JR DEEP Upload Resources About Debug

Finished

SPSS practicals 1 - 12 as an eBook

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann-Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11**
- Practical 12

Finally we look at the Bonferroni procedure with the results as shown below.

Multiple Comparisons

Reading test score
Bonferroni

(i) How often child reads for pleasure	(j) How often child reads for pleasure	Mean Difference (I-J)	Std. Error	Sig.	95% Confidence Interval	
					Lower Bound	Upper Bound
Never	Less than once a month	-.2407	16.109	1.000	-70.30	22.16
	At least once a month	-.2058	13.987	1.000	-60.72	19.56
	At least once a week	-.2475	12.478	.500	-60.56	11.06
Less than once a month	Most days	-.3854*	11.852	.016	-72.56	-4.52
	Never	24.07	16.109	1.000	-22.16	70.30
	At least once a month	3.49	14.592	1.000	-38.38	45.37
At least once a month	At least once a week	-.68	13.153	1.000	-38.43	37.07
	Most days	-14.47	12.561	1.000	-50.52	21.58
	Never	20.58	13.987	1.000	-19.56	60.72
At least once a week	Less than once a month	-3.49	14.592	1.000	-45.37	38.38
	At least once a month	-4.17	10.446	1.000	-34.15	25.81
	Most days	-17.96	9.090	.067	-45.77	9.85
Most days	Never	24.75	12.478	.500	-11.06	60.56
	Less than once a month	.88	13.153	1.000	-37.07	38.43
	At least once a month	4.17	10.446	1.000	-25.81	34.15
At least once a month	Most days	-13.79	7.347	.034	-34.88	7.30
	Never	38.54*	11.852	.016	4.52	72.56
	Less than once a month	14.47	12.561	1.000	-21.58	60.52
At least once a week	At least once a month	17.96	9.090	.067	-9.85	45.77
	At least once a week	13.79	7.347	.034	-7.30	34.88

Based on observed means.
The error term is Mean Square(Error) = 993.810.
*. The mean difference is significant at the .05 level.

The Bonferroni procedure is very conservative and of the three selected we will generally get the least pairs that are statistically significantly different. We see that the following pairs Most days Never are significantly different.

The options we have chosen also give us a marginal means plot as shown below.



Current Practical 11 outputs

Stat-JR DEEP Upload Resources About Debug

localhost:57517/ebooks/12/writing/1/#bookpage12

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11**
- Practical 12

Estimated Mean

How often child reads for pleasure

In general a marginal means plot will show the means for each category adjusting for any other variables in the model but here we are simply fitting a one way ANCOVA and so we will simply see the means for each group plotted together so for example we see the mean for category Never is 84.50 and for category Less than once a month is 108.57.

Finally we have asked SPSS to find homogeneous subsets within our categories. These are groups of categories that cannot be separated (using in this case Tukey's HSD) as statistically different within the category.

Reading test score

Tukey HSD^{a,b,c}

How often child reads for pleasure	N	Subset	
		1	2
Never	8	84.50	
At least once a month	13	105.08	105.08
Less than once a month	7	108.57	108.57
At least once a week	28	109.25	109.25
Most days	50		123.04
Sig.		.295	.002

Means for groups in homogeneous subsets are displayed.
 Based on observed means.
 The error term is Mean Square(Error) = 963.810.
 a. Uses Harmonic Mean Sample Size = 12.480.
 b. The group sizes are unequal. The harmonic mean of the group sizes is used. Type I error levels are not guaranteed.
 c. Alpha = .05.

In this case we see there are 2 identified subsets with categories Never, At least once a month, Less than once a month, At least once a week in subset 1; At least once a month, Less than once a month, At least once a week, Most days in subset 2.

about




Current Practical 12

- Covers an introduction to multiple regression
- Requires as input 1 response variable and 2 predictor (numerical) variables
- Performs separate regressions for the 2 predictors
- Next performs a multiple regression for the 2 predictors together.
- Then shows how SPSS allows regressions to be done as a group
- Finishes with residuals, plotting a histogram and the residuals against each predictor in turn.



Current Practical 12 outputs



SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

Practical 12

Multiple Regression practical

In this practical we will look at regressing two different predictor variables individually on a response, followed by a model containing both of them. We will also look at a second approach to doing this. This work builds on the earlier simple linear regression practical.

We start by running the first linear regression to look at if there is a significant (linear) effect of **mathscore** on **readscore**. This is done in SPSS as follows:

- Select **Linear** from the **Regression** submenu available from the **Analyze** menu.
- Copy the **Reading test score[readscore]** variable into the **Dependent** box.
- Copy the **Math test score[mathscore]** 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 **OK** button to run the command.

SPSS will produce several tabular outputs which we described in detail in the regression practical. Here we will focus on only the model summary and coefficients tables that can be seen below:

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.610 ^a	.372	.366	25.631

a. Predictors: (Constant), Math test score

Here we see some fit statistics for the overall model. The statistic R here takes the value .610 and is equivalent to the Pearson correlation coefficient for a simple linear regression. R squared (.372) is simply the value of R squared (R multiplied by itself) and represents the proportion of variance in the response variable, **readscore** explained by **mathscore**. The table also includes an adjusted R square measure which here takes value .366 and is a version of R squared that is adjusted to take account of the number of predictors (one in the case of this simple linear regression) that are in the model. We next look at the coefficients table which is shown below.

Coefficients

Model	Unstandardized Coefficients		Standardized Coefficients		Sig.	95.0% Confidence Interval for B	
	B	Std. Error	Beta	t		Lower Bound	Upper Bound
1. (Constant)	45.216	8.939		5.058	.005	27.494	62.939
Math test score	6.880	.868	.610	7.926	.000	5.159	8.601

This table gives the most interesting information about the regression model. We begin with the coefficients that form the regression equation. The regression intercept takes value 45.216 and is the value of the regression line when **mathscore** takes value 0. The regression slope takes value 6.880 and is the amount by which we predict that **readscore** changes for an increase of 1 in **mathscore**.

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



Current Practical 12 outputs

Stat-JR:DEEP Upload Resources About Debug

SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 **13** Next → Go to page

The model and coefficients tables can be seen below:

Model Summary

Model	R	R Square	Adjusted R Square	Std. Error of the Estimate
1	.642 ^a	.412	.400	24.979

a. Predictors: (Constant), Math test score, hours per week spent on homework (term time)

This time we see some fit statistics for the multiple regression with both **mathscore** and **hours_hwk**. The statistic R here takes the value .642. R squared (.412) represents the proportion of variance in the response variable, **readscore** explained by the multiple regression. This time the adjusted R square measure takes value .400 which we can compare with .366 for just **mathscore** and .045 for just **hours_hwk**. We next look at the coefficients table which is shown below:

Coefficients

Model	Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B	
	B	Std. Error	Beta			Lower Bound	Upper Bound
1 (Constant)	37.975	9.465		4.012	.000	19.197	56.753
Hours per week spent on homework (term time)	2.485	1.109	.241	3.146	.002	1.287	5.683
Math test score	6.816	.974	.598	7.797	.000	5.082	8.551

This time the coefficients that form the regression equation are as follows: The regression intercept takes value 37.975 while the regression slope for **hours_hwk** takes value 2.485 and the slope for **mathscore** takes value 6.816. These have changed from 3.383 and 6.680 respectively when the variables are fitted individually.

This time there are two standardised slopes with the slope for **hours_hwk** taking value .241 and the slope for **mathscore** taking value .598.

For **hours_hwk** the slope has t statistic 3.146 and this value can be compared with a t distribution to test the null hypothesis that the slope is 0. We can see the resulting p value for the test under the Sig. column. The p value (quoted under Sig.) is .002 which is less than 0.05. We therefore have significant evidence to reject the null hypothesis that the slope is zero.

For **mathscore** the slope has t statistic 7.797 and this value can be compared with a t distribution to test the null hypothesis that the slope is 0. We can see the resulting p value for the test under the Sig. column. Although SPSS quotes the p value (quoted under Sig.) as .000 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the slope is zero.

For the intercept here the t statistic is 4.012 and this value can be compared with a t distribution to test the null hypothesis that the intercept is 0. We can see the resulting p value for the test under the Sig. column. Although SPSS quotes the p value (quoted under Sig.) as .000 it is not exactly 0 and is in fact simply smaller than 0.001 as SPSS is quoting the first 3 decimal places. We therefore have significant evidence to reject the null hypothesis that the intercept is zero.

The final two columns give confidence intervals for the coefficients and so we are 95 percent confident that the intercept takes a value between 19.197 and 56.753.

Similarly we are 95 percent confident that the slope for **hours_hwk** takes a value between 1.287 and 5.683. Here we see the confidence interval does not contain 0 which corresponds to the fact we could reject the null hypothesis that the slope was 0.

Finally we are 95 percent confident that the slope for **mathscore** takes a value between 5.082 and 8.551. Here we see the confidence interval does not contain 0 which corresponds to the fact we could reject the null hypothesis that the slope was 0.

Finally we will show how to run two of the regression models in one go and build up the regression in blocks. This is done in SPSS as follows:



Current Practical 12 outputs

Stat-JR:DEEP Upload Resources About Debug

Finished

SPSS practicals 1 - 12 as an eBook

← Previous 1 2 ... 6 7 8 9 10 11 12 **13** Next → Go to page

This document contains the outputs from the first 12 practicals in the British Academy project.

- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12

Here we see the model summaries for the first and third regression models earlier i.e. we fit a model with just **mathscore** and then a second model where we introduce **hours_hwk**.

Coefficients

Model	Unstandardized Coefficients		Standardized Coefficients	t	Sig.	95.0% Confidence Interval for B	
	B	Std. Error	Beta			Lower Bound	Upper Bound
1 (Constant)	46.215	9.488		4.871	.000	27.394	65.036
Math test score	6.784	.912	.595	7.439	.000	4.975	8.593
2 (Constant)	37.975	9.485		4.012	.000	19.197	56.753
Math test score	6.816	.874	.598	7.797	.000	5.082	8.551
Hours per week spent on homework (term time)	3.485	1.108	.241	3.148	.002	1.287	5.683

Similarly we have the model coefficients for the first and third models from earlier in one combined table.

Having selected standardised residuals we get an additional table, the **Residuals statistics** table.

Residuals Statistics

	Minimum	Maximum	Mean	Std. Deviation	N
Predicted Value	61.91	54.61	114.18	20.628	103
Residual	-66.357	58.936	.000	24.634	103
Std. Predicted Value	-2.534	2.445	.000	1.000	103
Std. Residual	-2.667	2.369	.000	.990	103

This table just summarises the predictions and residuals that come out of the final regression and it is perhaps easier to look at these via plots. As we requested that standardized residuals were saved this has resulted in an additional variable being stored in the dataset named **ZRE_1** at the end of the existing variables. We can use this variable to create some residuals plot 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 diagnostics** 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.

This will produce the graph as shown below.

Histogram

Dependent Variable: Reading test score

Mean = 3.22E-16
Std. Dev. = 0.990
N = 103



Current Practical 12 outputs

Stat-JR:DEEP Upload Resources About Debug

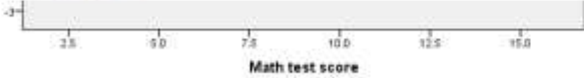
SPSS practicals 1 - 12 as an eBook

Finished

← Previous 1 2 ... 6 7 8 9 10 11 12 13 Next → Go to page

▼ This document contains the outputs from the first 12 practicals in the British Academy project.

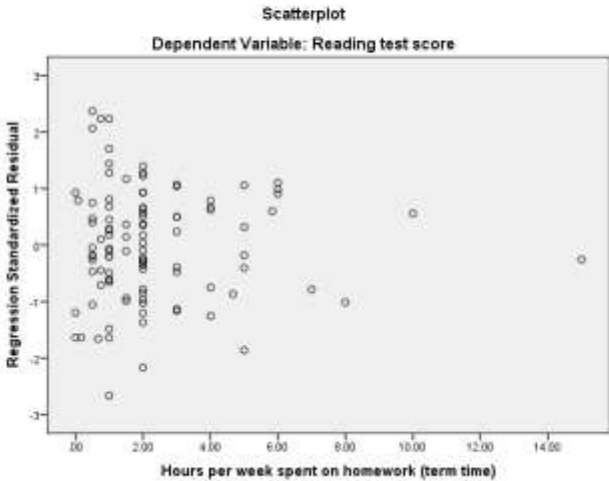
- Practical 1
- Practical 2
- Practical 3
- Practical 4 - Checking for normality
- Practical 5 - Independent Samples t test
- Practical 6 - Paired t test
- Practical 7a - The Mann Whitney test
- Practical 7b - The Wilcoxon Sign Rank test
- Practical 8 - The Chi-squared test
- Practical 9
- Practical 10
- Practical 11
- Practical 12



Here we hope not to see any pattern where there was more variability in the residuals for particular values of **Math test score[mathscore]**. We can repeat this plot for **Hours per week spent on homework (term time)[hours_hwk]** as follows:

- 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.
- Remove the **Math test score[mathscore]** variable from the **X Axis** box.
- Copy the **Hours per week spent on homework (term time)[hours_hwk]** variable into the **X Axis** box.
- Click on the **OK** button.

This will produce the graph as shown below:



Scatterplot
Dependent Variable: Reading test score

Regression Standardized Residual

Hours per week spent on homework (term time)

Note that again we hope not to see any pattern in the residuals.

about



Work package 4

- The original plan in the grant was to construct concept materials using StatJR to supplement the students learning.
- An example of such a concept eBook is shown overleaf and we have others for summary statistics and other statistical tests.
- Given the switched focus to SPSS we propose to integrate the conceptual material within the learning component of each practical (so that conceptual understanding and software skills are developed side-by-side)



SAA 2

Finished

Statistical Analysis Assistant
(Mark 2 - Chi-squared
edition)

Checking for an Association between two categorical variables

Checking for an Association between two categorical variables

You will be presented below with the choice of categorical variables to choose. Having chosen them you will then get the output to your analysis

First categorical variable:

[about](#)

Second categorical variable:

[about](#)

To do a chi-squared test we start by tabulated observed counts and totals:

Observed	cscat=0.0	cscat=1.0	cscat=2.0	Total
nsucc=0.0	188	1559	303	2050
nsucc=1.0	139	1536	440	2115
Total	327	3095	743	4165

SAA 2

Finished

[← Previous](#)
[1](#)
[2](#)
[Next →](#)

Statistical Analysis Assistant
(Mark 2 - Chi-squared
edition)

**Checking for an
Association between two
categorical variables**

To do a chi-squared test we start by tabulated observed counts and totals:

Observed	cscat=0.0	cscat=1.0	cscat=2.0	Total
nsucc=0.0	188	1559	303	2050
nsucc=1.0	139	1536	440	2115
Total	327	3095	743	4165

We can therefore work out the expected counts from the margins of the observed data

And so we expect

$$E(\text{cscat}=0.0, \text{nsucc}=0.0) = \text{Total cscat}=0.0 * \text{Total nsucc}=0.0 / \text{grand total} = 327 * 2050 / 4165 = 160.95$$

$$E(\text{cscat}=1.0, \text{nsucc}=0.0) = \text{Total cscat}=1.0 * \text{Total nsucc}=0.0 / \text{grand total} = 3095 * 2050 / 4165 = 1523.35$$

$$E(\text{cscat}=2.0, \text{nsucc}=0.0) = \text{Total cscat}=2.0 * \text{Total nsucc}=0.0 / \text{grand total} = 743 * 2050 / 4165 = 365.7$$

$$E(\text{cscat}=0.0, \text{nsucc}=1.0) = \text{Total cscat}=0.0 * \text{Total nsucc}=1.0 / \text{grand total} = 327 * 2115 / 4165 = 166.05$$

$$E(\text{cscat}=1.0, \text{nsucc}=1.0) = \text{Total cscat}=1.0 * \text{Total nsucc}=1.0 / \text{grand total} = 3095 * 2115 / 4165 = 1571.65$$

$$E(\text{cscat}=2.0, \text{nsucc}=1.0) = \text{Total cscat}=2.0 * \text{Total nsucc}=1.0 / \text{grand total} = 743 * 2115 / 4165 = 377.3$$

So the table of expected counts is

Expected	cscat=0.0	cscat=1.0	cscat=2.0	Total
nsucc=0.0	160.95	1523.35	365.7	2050.0
nsucc=1.0	166.05	1571.65	377.3	2115.0
Total	327.0	3095.0	743.0	4165.0

We next look at differences between what we observe and expect in each cell. We square these values so that every difference is positive and scale by the expected counts so that more frequently expected cells aren't overly influential. So for example for cscat=0.0, nsucc=0.0 $(O-E)^2/E = (188-160.95)^2/160.95=4.55$.

This statistic is shown in tabular form below

SAA 2

Finished

← Previous 1 **2** Next → Go to page

Statistical Analysis Assistant
(Mark 2 - Chi-squared
edition)
**Checking for an
Association between two
categorical variables**

Use the table of expected counts to

Expected	cscat=0.0	cscat=1.0	cscat=2.0	Total
nsucc=0.0	160.95	1523.35	365.7	2050.0
nsucc=1.0	166.05	1571.65	377.3	2115.0
Total	327.0	3095.0	743.0	4165.0

We next look at differences between what we observe and expect in each cell. We square these values so that every difference is positive and scale by the expected counts so that more frequently expected cells are not overly influential. So for example for cscat=0.0, nsucc=0.0 $(O-E)^2/E = (188-160.95)^2/160.95=4.55$. This statistic is shown in tabular form below

$(O-E)^2/E$	cscat=0.0	cscat=1.0	cscat=2.0
nsucc=0.0	4.55	0.83	10.75
nsucc=1.0	4.41	0.81	10.42

The test statistic for a chi-squared test is found by summing the values of this table so

$$\text{Chisq}=4.55+0.83+10.75+4.41+0.81+10.42=31.77$$

This is compared with a chi-squared table with degrees of freedom = (number of columns - 1)x(number of rows - 1) =

$$(2-1)x(3-1)=2$$

Looking up the chi-squared table the value for $P=0.05$ is 5.99 and for $P=0.01$ = 9.21

as $31.77 > 9.21$ our P value is less than 0.01 and we have strong evidence to reject the null hypothesis (at the $P=0.01$) level

The p-value is in fact less than 0.0001

Work package 5

- For this work package we intend to run a workshop to demonstrate the system and get feedback.
- The original timetable for this is month 21 or roughly Xmas time but we decided that presenting today would instead capture the same audience.
- We have also demonstrated aspects of the software to John's group in Edinburgh who were enthusiastic and discussed the software and topics with colleagues at Bristol, Exeter and Cardiff.



Questions

????????????

