MedCalc®

for Windows

statistics for
biomedical research

software manual

MedCalc Software
Broekstraat 52
B-9030 Mariakerke
Belgium

http://www.medcalcsoftware.com/

Version 11.6

MEDCALC®

MedCalc Software
Broekstraat 52
B-9030 Mariakerke
Belgium

mailto:info@medcalcsoftware.com
http://www.medcalcsoftware.com/

CONTENTS

INTRODUCTION

Program installation

System requirements

To run MedCalc, you need a PC with 2000, XP, Windows Vista or Windows 7 (32-bit or 64-bit versions), with at least 512 MB RAM and about 20 Megabyte free space on the hard disk.

MedCalc installation

- If you have downloaded the software from the Internet, you should locate the setup file **medcalcsetup.exe** on your hard disk and execute (double-click) it.
- If you have a program CD, locate the setup file **medcalcsetup.exe** on the CD and execute it.
- When installation is complete, you start MedCalc by clicking the *Start* button and point to *Programs*, next click MedCalc.

Note: To install MedCalc you must be logged on to your computer with administrator privileges.

Sample data

At first run, MedCalc creates a folder named "MedCalc" with a subfolder "Sample files" in the "(My) Documents" map. The "Sample files" folder contains the following MedCalc sample files:

Bland and Altman plot.mc1
Clustered graphs.mc1
Control chart.mc1
Cox regression - Gallstones.mc1
Data for ROC curve analysis.mc1
Dates.mc1
Logistic regression.mc1
Meta-analysis Continuous data.mc1
Meta-analysis Odds ratio.mc1
Patients - sample data.mc1
Serial measurements.mc1
Survival curves.mc1
Youden.mc1

Registration

When you start MedCalc for the first time a dialog box appears with the following options:

Buy now: to connect to the MedCalc web site and order a product key.

Enter product key: to enter your user name and product key. You only have to enter your user name and product key once. The next time you start MedCalc, the program will not ask you this information again. If you do not have a product key, you can purchase one from the MedCalc web site (http://www.medcalc.org/). If you are a registered user and you have lost your product key, use the Recover product key form http://www.medcalc.org/productkey/. Alternatively, you can contact MedCalc Software (mailto:info@medcalcsoftware.com), and we will email you your user name and product key.

Free trial: You can run the software, without any feature limitations, during 15 days without registration.

Auto-update

Important: the auto-update feature is only available if you have a legal, non-network MedCalc license.

Step 1

When you start MedCalc, the program will check whether a new version of the software is available from the MedCalc website.

When the program can successfully connect to the MedCalc website, and an update is available, the following dialog box is displayed:

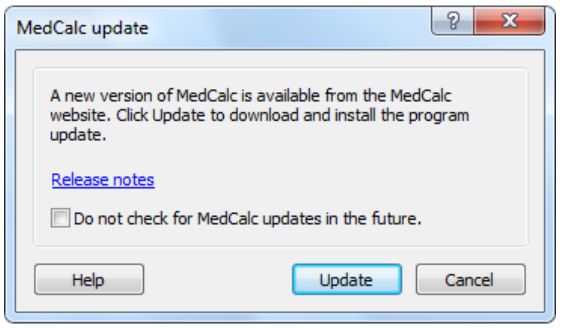

MedCalc updates are free of charge and updating to the latest version is always recommended.

If you don't want to install the update, click Cancel and MedCalc will remind you of the update after a user-defined number of days (see Options, p. 54). If you do not want MedCalc to automatically check for updates in the future, select the option *Do not check for MedCalc updates in the future*. You can always re-enable this option in MedCalc's Options dialog box (p. 54).

Step 2

If you click Update, the update will be downloaded from the MedCalc website:

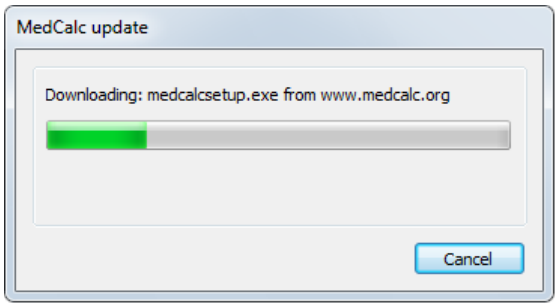

You can interrupt and cancel the download by clicking the Cancel button.

Step 3

When the update could successfully be downloaded, the update software is launched automatically.

Privacy

When MedCalc checks for updates, no personal information (name, product key) is being sent to the MedCalc web server. The software will only retrieve the most recent version number from the website and compare it with the version number of your MedCalc copy.

What's new

If you are already familiar with MedCalc, you can read about the latest changes and additions in MedCalc by selecting *What's new* in the MedCalc *Help* menu.

It is possible that new features are available in the software, which are not described in this manual. You can find information on these new features in the MedCalc Help file (see *On line help*, p. 4), or on the MedCalc web site (see *MedCalc on the Web*, p. 210). Updated versions of the manual in PDF format can be downloaded from the MedCalc web site http://www.medcalc.org.

Regional settings support

MedCalc supports regional differences as entered in the *Regional settings* dialog box in the *Control panel*.

The following are taken from the Windows settings:

- **Decimal symbol**: the character used to separate the decimal digits from whole numbers.
- **List separator**: the symbol used to separate elements in a list, e.g. arguments in spreadsheet functions, or fields when exporting data as a text file.

  E.g. when the list separator is a comma, the arguments in the RANDNORM function are separated by a comma: RANDNORM(m,s). If the list separator is a semicolon, the arguments are separated by a semicolon: RANDNORM(m;s).
- The list separator is also used to separate the fields or variables when you export spreadsheet data.
- **The list separator cannot be equal to the decimal separator!**
- If MedCalc finds the list separator to be equal to the decimal separator in the Windows settings, then it will use a semicolon as the list separator when the decimal separator is a comma.
- **Date format**: MM.DD.YY, DD.MM.YY or YY.MM.DD (see *Date functions*, p. 218).

You can use different characters for the Decimal symbol and List separator in the Options panel, see Options, Regional settings, p. 55.

The MedCalc menu bar

After starting the program, the MedCalc program window appears, with a menu bar at the top of the screen containing the following headings:

- File — Disk input and output of spreadsheet data files, printing, quit MedCalc
- Edit — Cut, copy, paste, find, insert, ...
- View — Open Contents bar, spreadsheet, notes editor, set view characteristics...
- Format — Character formatting, spreadsheet & column format, graph formatting, …
- Tools — Sort data, edit variables list, options, …
- Statistics — Statistical analysis of spreadsheet data, t-tests, Wilcoxon tests, comparison of groups, method comparison, ROC curve analysis, etc.
- Graphs — Statistical graphs, Kaplan-Meier survival curve, Control chart
- Tests — Statistical tests on tabulated or summarized data (useful when you do not have the raw data available in the spreadsheet): difference between means, standard deviations, percentages, correlation coefficients, relative risk, odds ratio
- Sampling — Calculation of sample sizes
- Window — Rearrange windows or activate a specific window
- Help — Get help and information.

On line help

At any moment during working with MedCalc *on line help* is available, e.g. information on a menu item, explanations on a dialog box, etc. After you press function key *F1* the program displays a context-sensitive help text.

You can also call the help function by selecting *Contents and Index* in the *Help* menu.

Help in dialog boxes

To get comprehensive help on the dialog box, click the Help button.

To get help on any item in the dialog box: click the question mark ? in the title bar of the dialog box, and then click an item in the dialog box.

- You can print or copy the information in a pop-up window by right clicking inside the pop-up window, and then clicking **Print Topic** or **Copy**.
- To close the pop-up window, click inside it.
- You can also get help on an item by right-clicking it, and then clicking **What's This?**

ScreenTips & quick help

When the mouse is moved over one of the buttons in a toolbar, a short explanation appears in a small popup window.

While you are making a selection in the menu, a description of the highlighted command is displayed in the status bar at the bottom of the window.

The spreadsheet data window

In MedCalc, data are entered in a spreadsheet. You open the spreadsheet window by selecting *Spreadsheet* in the *View* menu, or selecting *Data* in the *Contents bar*.

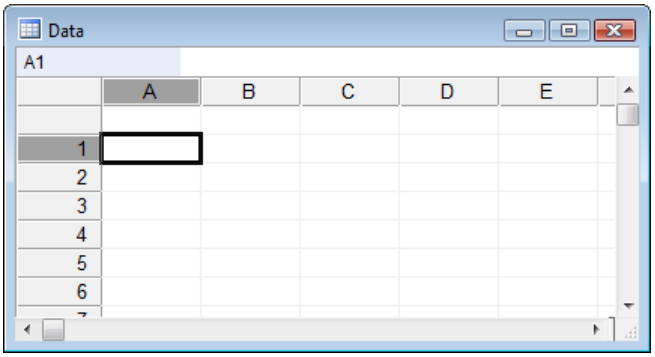

One cell (rectangle) of the spreadsheet is highlighted. This rectangle is called the *cell pointer*.

The first character(s) of this cell address (A) is the column indicator and the next number is the row number (row 1 is the first row number). Above row 1, there is an additional fixed row where you can enter a column heading, i.e. a name for the variable for which data will be entered in this column.

The column and row of the cell pointer are also called the current column and row.

The cell pointer can be moved with the arrow and/or cursor keys, or by means of the mouse. When you click on a cell, the cell pointer jumps to this cell. You can browse in the spreadsheet window by clicking on the right and bottom border of the spreadsheet window, or by using the *Page Up* and *Page Down* and other cursor or arrow keys.

By pressing the *Ctrl+Home* key you move the cell pointer to the first data cell in the spreadsheet: cell A1, and with *Ctrl+End* to the last cell used in the spreadsheet.

You can fix a number of columns in the spreadsheet by using the *Split window* command (see p. 208):

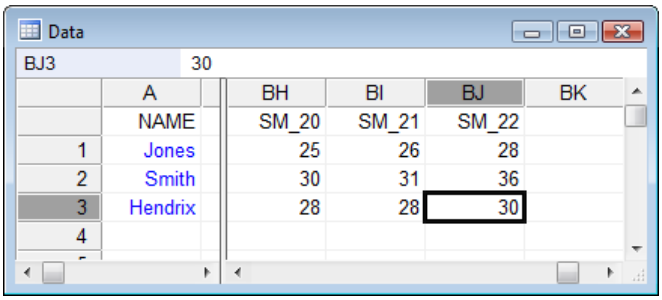

Total number of columns and rows in the spreadsheet

The default number of rows in the MedCalc spreadsheet is 16382, and the default number of columns is 520. However, you can configure MedCalc to contain a different number of rows and columns. The number of rows and columns available in the spreadsheet can be set in the *Options* dialog box, described on page 54.

How to enter data

Data for the different variables are entered in different columns of the spreadsheet. All data for a single subject or case are entered in one row in the spreadsheet. In the top row of the columns you can enter the names of the variables.

A variable name should not include any spaces. If necessary, you can use the underscore character _ to separate words, e.g. GRADE_A. Also the following characters cannot be used in a variable's name:

- + / * = < > ^ ( ) $ " ' : , .

In addition, the variable name must not start with a number and must be different from reserved words such as TRUE, FALSE, ROW and COLUMN.

The variable name should also not be equal to the address of a spreadsheet cell such as A1, S1, AB35, IL6, etc.

In order to enter the variable name LENGTH in the top row of column A, you first position the mouse pointer on this cell, and click the left mouse button. The cell pointer is now located in this cell, and you can type the variable's name on the keyboard. Next, you press the ↓ key to actually store the name in the computer's memory and move the cell pointer to cell A1. You can now enter the first value 171 for the variable LENGTH in this cell, followed by pressing the ↓ key to go to the next cell.

The data are not stored in the computer's memory until you have pressed the Enter key, or have moved the cell pointer to another cell.

When you want to change or *edit* the content of a cell, place the cell pointer on this cell and press function key *F2*. You can now move the cursor in the cell's text by means of the arrow keys ← and → and make the necessary changes. After you have made the changes, press the *Enter* key or move the cell pointer to another cell to store the new cell content in memory.

See notes on p. 218 on how to enter dates.

While you are entering data you must, from time to time, save your data on the disk. When you save the file for the first time then select the command *Save as* in the *File* menu, and next enter the file name in the file selector box described on page 8. After you have given a name to your data file, you can select the *Save* command in the *File* menu to save your data under the same file name.

If you do not save your data on the disk, your work will be lost when:

- ***you exit the program***
- ***switch of the computer***
- ***there is a power failure or a system or hardware malfunction.***

The program automatically formats the numbers that you enter according to the number of decimals selected in the *Format spreadsheet* box (see p. 40).

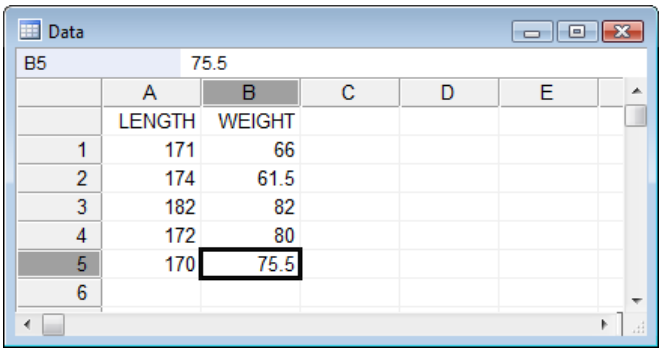

If you want a different number of decimals for a particular column, then select *Column* in the *Format* menu and enter the new number of decimals. In this dialog box you can also specify a different width for the column. You can also specify that particular columns only contain text data (option “Text format”). The latter option is particularly useful when the column contains text data that should not be interpreted as a number, e.g. dates: 05/07/1956 must be displayed as such, not as 0.0003652

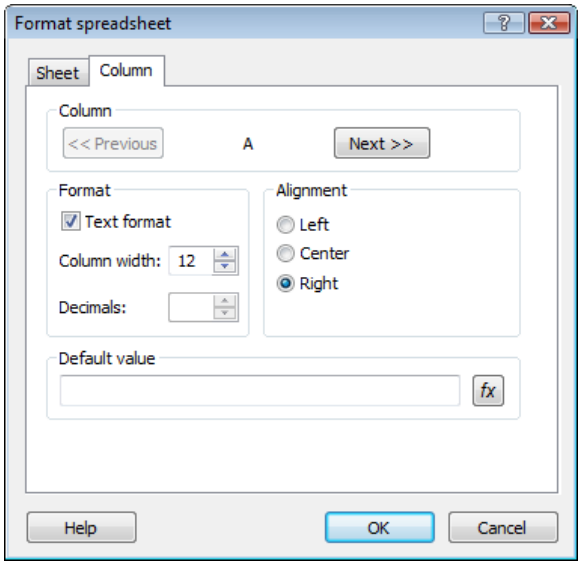

A variable may either be *categorical* or *numerical*. Categorical variables may either consist of numeric or alphanumeric (string) data. A numerical variable always consists of numerical data.

Categorical or qualitative variable

Categorical or qualitative data may either be entered as numbers or as text *strings*. A text string consists of one or more alphanumeric characters, placed in quotation marks. The program will consider every expression that cannot be interpreted as a number to be a string, even if it is not placed between quotation marks, e.g. “Green”, yellow, “10”. Distinction can be made between *Nominal* and *Ordinal* data:

Nominal data : a classification without obvious order, e.g. blood group, male/female.

Ordinal data : ordered categorical data, e.g. endometriosis stage, varicocele grade.

A variable that can only have 2 values is also called a dichotomous variable, for example pregnant/not pregnant, male/female.

In MedCalc, it is often useful to code categorical data with numerical values: 0 and 1, or 1 2 3 4, etc.

Numerical or quantitative variable

A numerical variable consists of numbers, for example 0, 25 or 3.258, or expressions that can be interpreted as a number, e.g. LOG(25) or SQRT(VAR1) where VAR1 is a variable’s name containing numerical data.

Numerical data may either be *continuous* or *discrete*.

Continuous data : numbers that can theoretically assume any value between two given values; usually *measurements*, for example: the height of a person.

Discrete data : data that are not continuous (and may have a limited number of values), usually *counts*, for example: the number of children in a family.

You can easily convert a numerical variable into a categorical variable using the Create groups tools (see p. 47-49) or the CATEGORISE function (see p. 219) or IF function (see p. 220).

Missing values

The data for all variables of one case (patient, sample) are entered on one row in the spreadsheet. When for one variable you do not know the value (or entry) for the case, you leave the corresponding cell blank and do not enter any data in this cell.

As a rule, the program will ignore an entry for a numeric variable when it is unable to interpret this entry as a number.

When text is entered in a cell for a numeric variable, the program will not take this case into account for calculations (it will not substitute the text value by a zero).

The following are recognized as numbers:

5.4	
LOG(36.5)	
HEIGHT/WEIGHT	(when 'Height' and 'Weight' are correctly defined variables)
HEIGHT/100	

The following are not recognized as numbers and are ignored for calculations:

5,8	
'4.6	
LOG(CONC)	(when 'CONC' is not a correctly defined variable or in case the variable 'CONC' has a zero, negative or missing value)
SQRT(-9)	(error!)
1/ HEIGHT	(when 'HEIGHT' is not a correctly defined variable or in case 'HEIGHT equals zero)

Data checking

After having entered the data, you should carefully check the data to ensure that they have been entered correctly.

Sometimes erroneous data input will become apparent when looking at the data range in the summary statistics report (e.g. maximum value of 78 for pH), or when plotting box-and-whisker plots, dot plots or scatter diagrams for the different variables. You should check clear outliers since they may indicate incorrect data entry, or they may result from a technical failure in measurement or from a study protocol violation. Only for such plausible reason you may exclude a value from further analysis, and not simply because a value is the smallest or largest. If there is no evidence of such a mistake then the value must remain unaltered.

You can locate any value in the spreadsheet using the *Find* procedure (p. 33).

You can exclude outliers from further calculations by using the *Exclude* command (p. 42).

How to save data

When you want to save the data, select the command *Save as* in the *File* menu. The program will display the following dialog box:

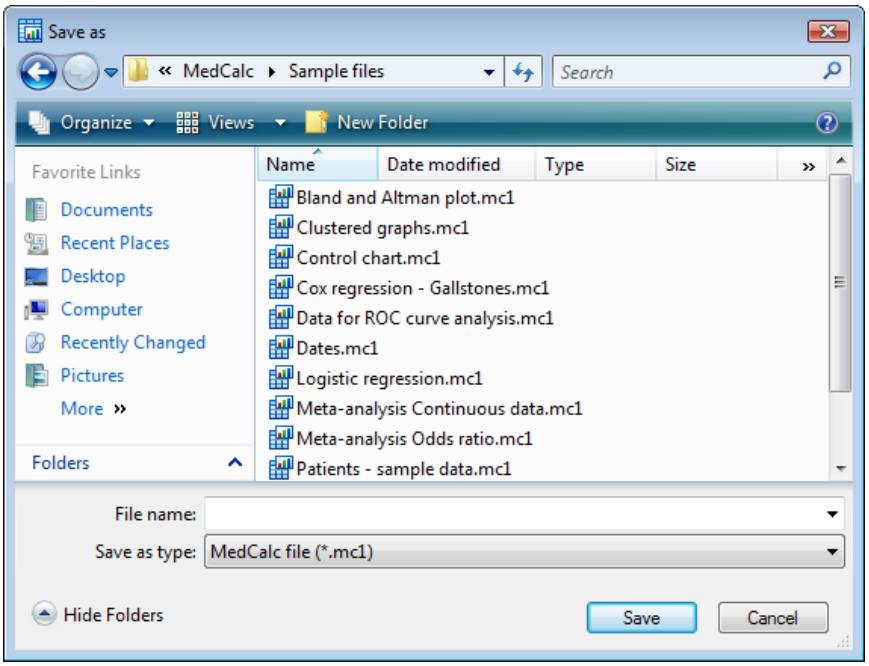

- **Save in**: select the directory where you want to save the data file.
- **File Name**: select or type the name of the file for the data. This box lists files with the file name extension selected in the *Save as type* box. You can select a file by clicking on a file name.
- **Save as type**: select the type of file you want to see in the file name list box.

When the correct file name is entered, click the SAVE button.

Statistics

After you have entered data in the spreadsheet, it is advised to save the data on disk. This is done by selecting the command *Save as* in the *File* menu. You will have to enter a name for your data file in the *File selector box* (see p. 8).

To load a data file from disk, select the command *Open* in the *File* menu. Next, select the name of the file in the *File selector box*. For example, you can select the data file *Patients - sample data.mc1*. This file contains data on length and weight for a number of persons.

When you want to obtain summary statistics for the variable *Weight*, select the command *Summary statistics* in the *Statistics* menu. The following dialog box appears on the screen.

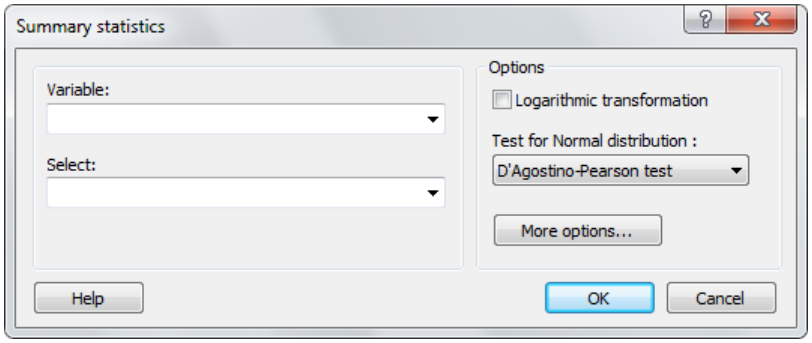

Click the button to obtain a list of variables.

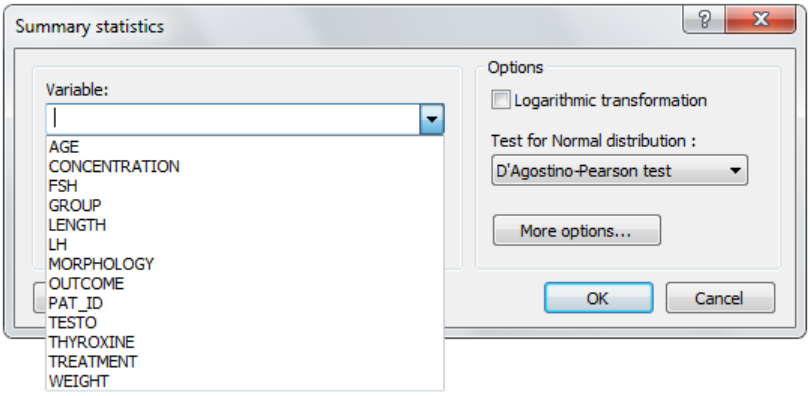

From this list you can select a variable by clicking on the variable's name.

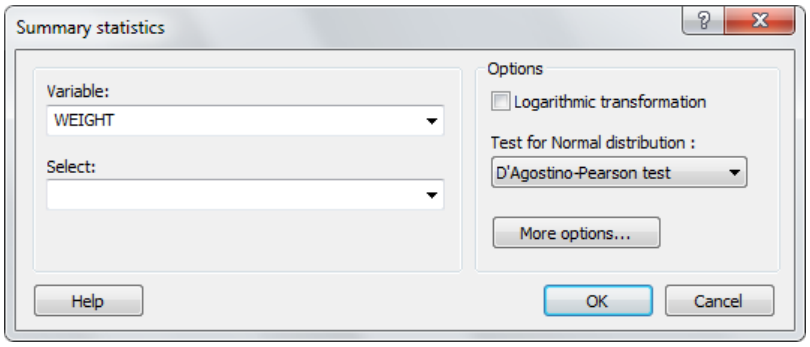

If the variable requires a logarithmic or square root, or any other mathematical transformation, then you can enter a formula in the *variable* field:

SQRT(WEIGHT)
LOG(LENGTH)

or you can enter a formula combining different variables, e.g.

WEIGHT/LENGTH

By doing so, new variables will be added to the variables list: SQRT(WEIGHT), WEIGHT/LENGTH, etc..

For an overview of mathematical operators and functions available in MedCalc, refer to pages 212 and 213.

Optionally, you may also enter a selection criterion in the *Select* field of the dialog box, in order to include only a selected subgroup of cases in the statistical analysis.

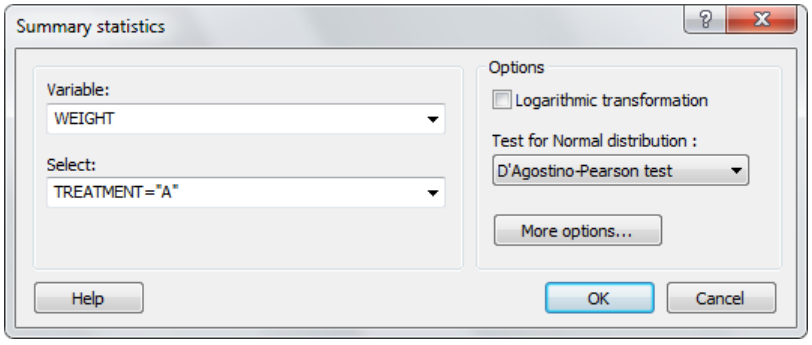

The *Select* field may contain a combination of different criteria, using the AND and OR functions (see also p. 219):

AND(LENGTH>160,LENGTH<170)

These expressions may make use of all mathematical and other spreadsheet functions as described in a next part of this manual (p. 212).

After you have entered a selection criterion in a dialog box (in order to include only a selected subgroup of cases in the statistical analysis), this selection criterion will be 'remembered' by the program and will be selectable in the *Select* list.

After you have clicked the OK button, the following statistics are calculated and displayed in the *results* window:

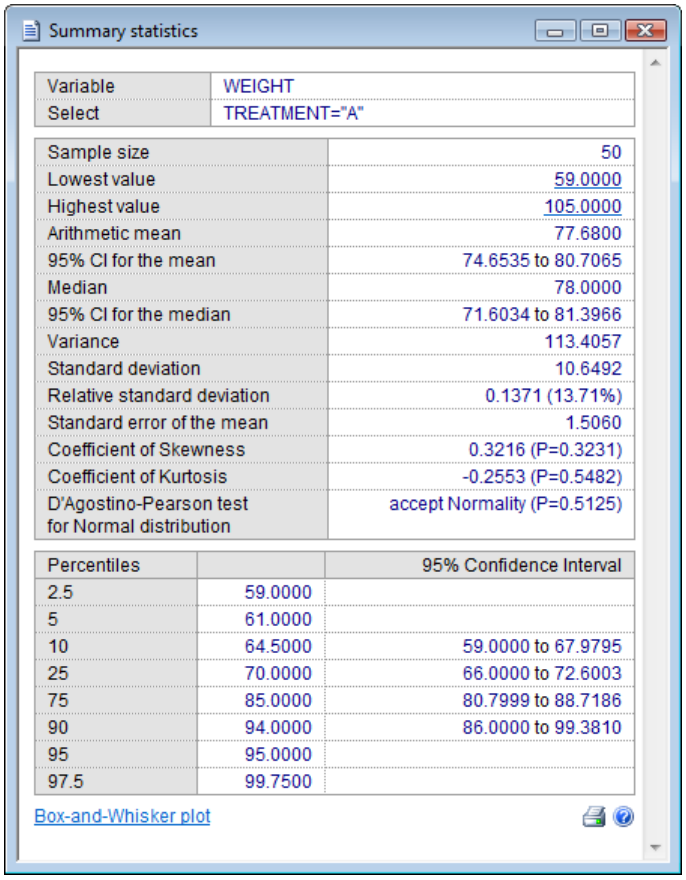

Summary statistics

Variable	WEIGHT
Select	TREATMENT="A"

Sample size	50
Lowest value	59.0000
Highest value	105.0000
Arithmetic mean	77.6800
95% CI for the mean	74.6535 to 80.7065
Median	78.0000
95% CI for the median	71.6034 to 81.3966
Variance	113.4057
Standard deviation	10.6492
Relative standard deviation	0.1371 (13.71%)
Standard error of the mean	1.5060
Coefficient of Skewness	0.3216 (P=0.3231)
Coefficient of Kurtosis	-0.2553 (P=0.5482)
D'Agostino-Pearson test for Normal distribution	accept Normality (P=0.5125)

Percentiles		95% Confidence Interval
2.5	59.0000	
5	61.0000	
10	64.5000	59.0000 to 67.9795
25	70.0000	66.0000 to 72.6003
75	85.0000	80.7999 to 88.7186
90	94.0000	86.0000 to 99.3810
95	95.0000	
97.5	99.7500	

Box-and-Whisker plot

For more details about the displayed statistics, see p. 57.

Note that when you change the data in the spreadsheet, this will not automatically cause a recalculation of the displayed statistics, but you have to repeat the statistical analysis for the new data.

To facilitate this, use the *Named tests and graphs* feature: by clicking the  button you can save the analysis so you can recall the analysis later, possibly with new data. After you have clicked this button the graph is added to the *Named tests and graphs* list in the Contents bar, with a default name, which you can edit in the Contents bar (see p. 38).

Graphs

When you want to create a graph, proceed in a way similar as for any other statistical procedure. As an example, you will create a histogram for the variable Weight of the data file you have loaded in the previous section.

To obtain the histogram, you select the command *Histogram* in the *Statistics* menu.

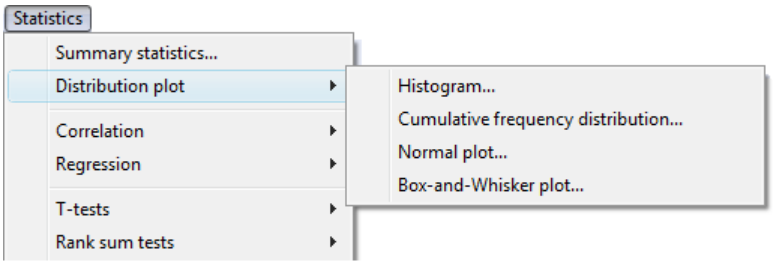

A dialog box is displayed similar to the one for *Summary statistics*. Again, enter the name of a variable and optionally a selection criterion. If you have previously entered this selection criterion in e.g. the dialog box for summary statistics, then this it be selectable in the *Select* list (click the ▾ button).

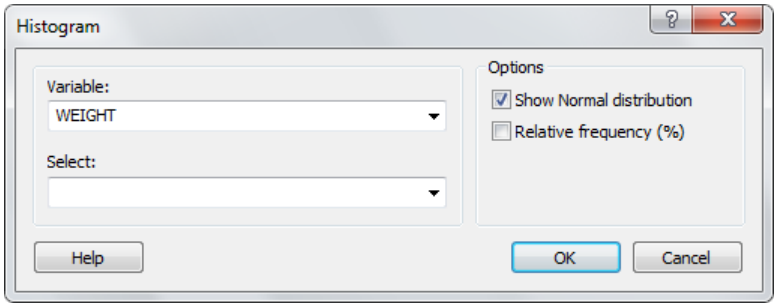

When you have identified the variable in the dialog box, click the OK button to proceed. The program will display the following box:

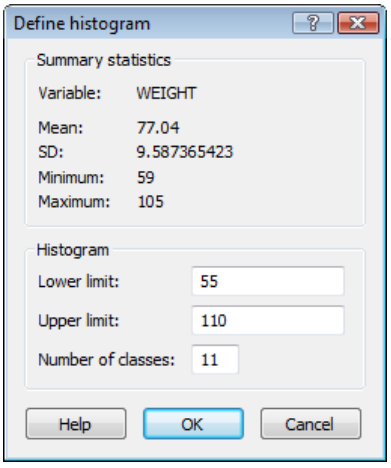

After you have clicked OK the histogram will be displayed in a new window:

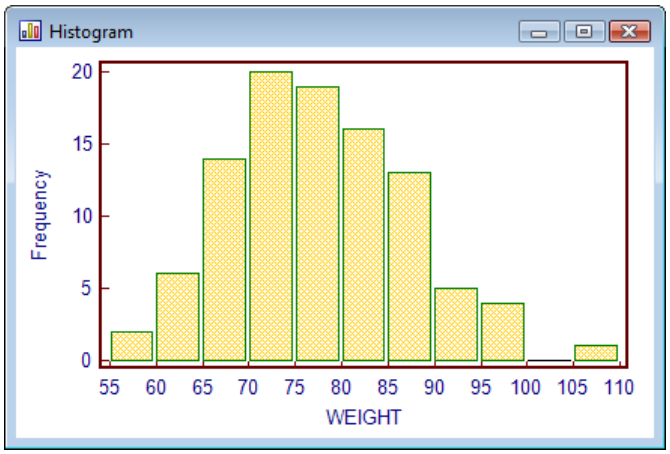

To get statistical information on the data represented in the graph (sample size, etc.):

- click the right mouse button in the graph window:

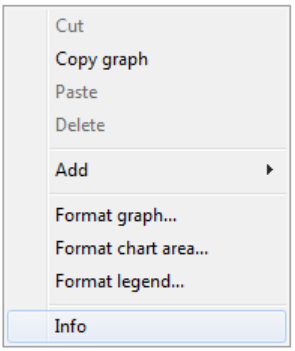

- select the *Info* command in the shortcut menu.

By clicking the  button you can save the graph so you can recreate the graph later, possibly with new data. After you have clicked this button the graph is added to the *Named tests and graphs* list in the Contents bar, with a default name, which you can edit in the Contents bar (see p. 38).

Graph formatting

Click the  button in the Formatting toolbar to format the graph. The effects of the changes in the left panels (displayed below) of the dialog box can be previewed in the right panel:

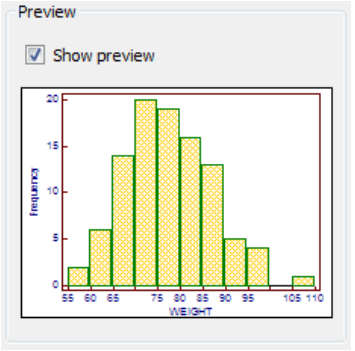

Scheme

In the graph colors box you can select a color for every object used in the graph.

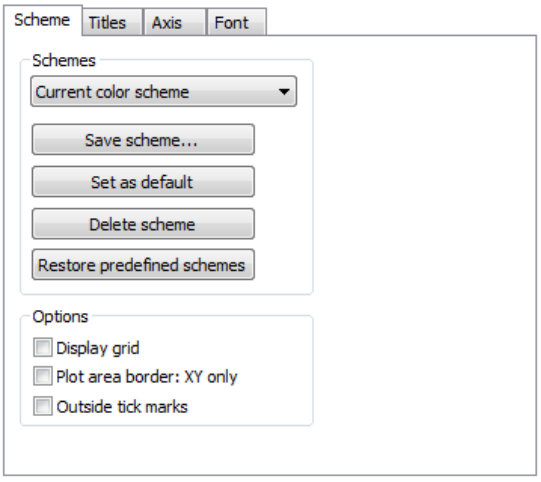

Schemes

In this box you can select a predefined or user-defined color scheme.

- *Save scheme*: "Save scheme..." can be used to save the currently defined color selections as a named scheme for later use.
- *Set as default*: When you click "Set as default", the current color scheme will be saved and used as the default color scheme for new graphs. The color schemes also include the Fills, Lines and Marker style and color selections.
- *Delete scheme*: Click this button to delete the currently selected color scheme.
- *Restore predefined schemes*: Click this button to restore the different predefined color schemes.

Options

- *Display grid*: displays gridlines in the graph.
- *Plot area border: XY only*: when this option is selected, lines for the X and Y axes are drawn. If this option is not selected, the plot area is enclosed in a rectangle.
- *Outside tickmarks*: select this option to have the tickmarks drawn on the outside of the graph frame.

Titles

In this box you can edit the main text sections (title and axes) of the graph.

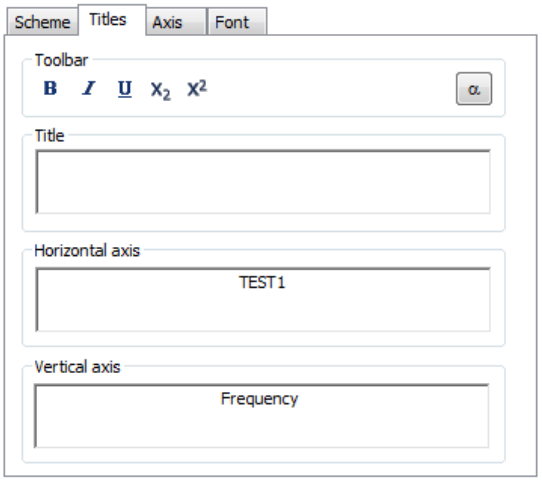

Axis

In this box you define the scaling of the X and/or Y-axis by entering the following information:

- *Minimum*: this is the lowest value that appears on the utmost left of the X-axis or bottom of the Y-axis;
- *Maximum*: this is the highest value that appears on the utmost right of the X-axis or top of the Y-axis;
- *Increment*: the increment value between two major grid lines;
- *Minor tickmarks*: the number of minor tickmarks between two major increments.

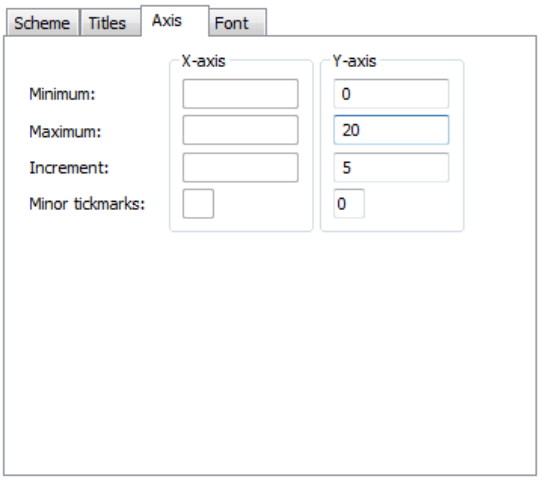

In MedCalc, the default increments are 1, 2 or 5 times a power of ten, yielding about 6 to 10 major tickmarks. If you are preparing a graph for publication, you are advised (Tukey, 1977) to limit the number of major tickmarks to about 2 or 3, so you will have to increase the value for increment, preferably also steps of 1, 2 or 5 times a power of 10.

If a logarithmic transformation of the data was selected in the graph's dialog box, then the information for the lowest value (*Minimum*) and *Increment* must be logarithmic values as well (e.g. for *Minimum* enter -2 for the value 0.01, or 0 for the value 1, 3 for the value 1000, etc.). *Increment* is then the exponent increment and should preferentially remain equal to 1.

After clicking OK, the graph will be replotted using the new axis scaling. However, the program will ignore a new axis scaling if this new scaling would cause data to fall outside the full axis range.

Font

In this box you can select the font, font style and font size used in the graph. The font *Sample* box (in which you can enter text freely) ignores the selected font size. The effect of font size can be seen in the Preview box.

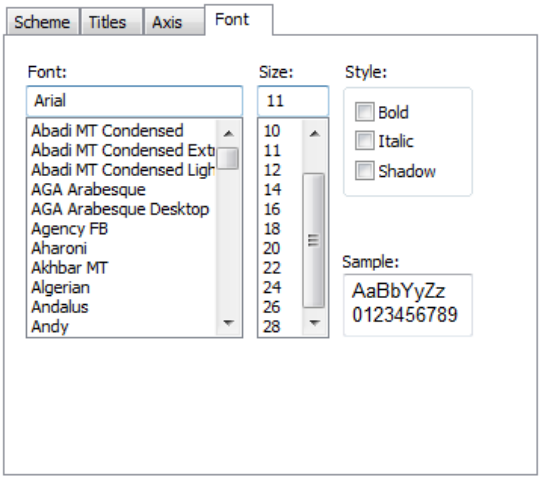

Format graph components

To format different graph components, right-click on the component and select "Format..." in the popup menu.

Example

In a box-and-whisker plot, right-click on a marker.

Select "Format marker..."

Select the desired marker attributes in the dialog box.

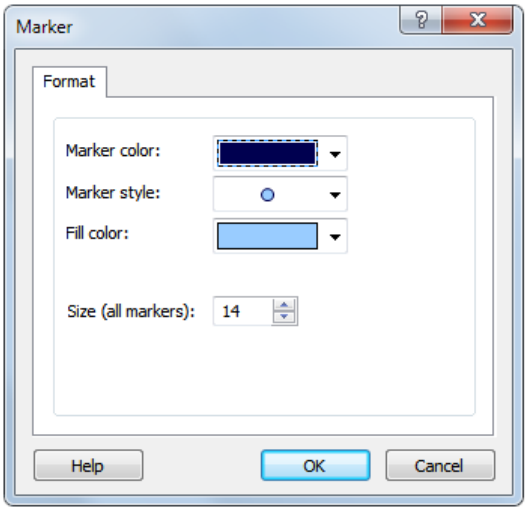

Add graphical objects

In all MedCalc graphs, you can add additional text boxes, lines, rectangles, ellipses and connectors.

Text box

Example:

P<0.05

To add a text box in the graph: right-click in the graph and in the popup menu select "Add" and "Text box":

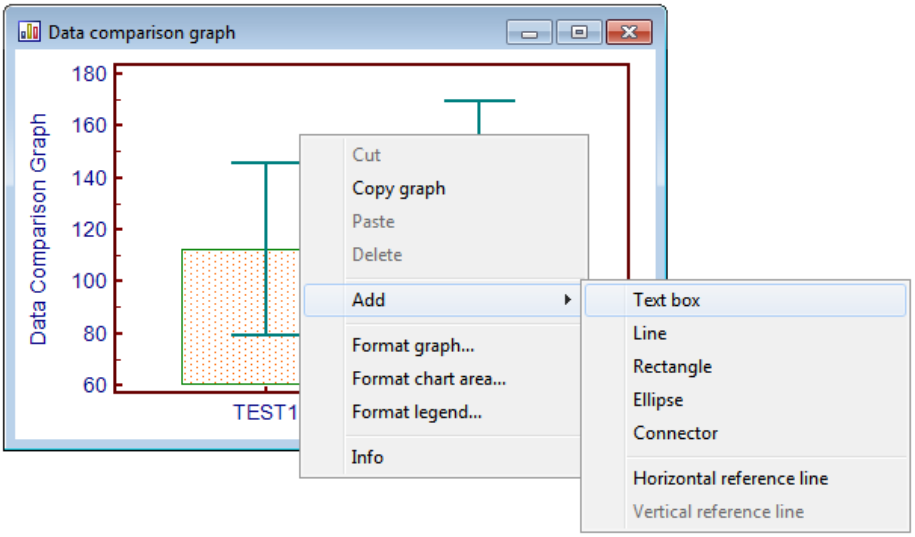

Next click in the graph to set the new text box position.

To edit the text box characteristics, double click on the text box or right-click on the text box and select "Format text box...":

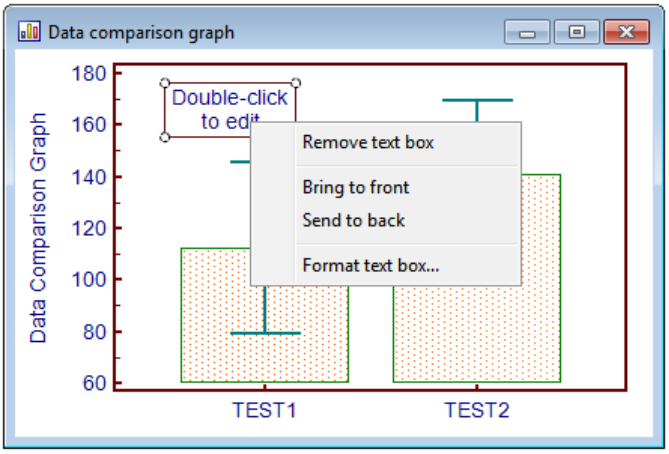

Line

Example:

To add a line in the graph: right-click in the graph and in the popup menu select "Add" and "Line".

Next click in the graph and drag the mouse to create the line.

To edit the text box characteristics, double click on the text box or right-click on the text box and select "Format line object...".

Rectangle

Example:

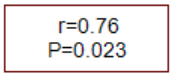

To add a rectangular frame in the graph: right-click in the graph and in the popup menu select "Add" and "Rectangle".

Next click in the graph and drag the mouse to create the rectangle.

To edit the line characteristics, double click on the text box or right-click on the text box and select "Format rectangle...".

Ellipse

Example:

To add an ellipse in the graph: right-click in the graph and in the popup menu select "Add" and "Ellipse".

Next click in the graph and drag the mouse to create the ellipse.

To edit the ellipse characteristics, double click on the text box or right-click on the text box and select "Format ellipse...".

Connector

A connector is commonly used to indicate statistical significance of differences between groups.

Example:

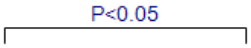

To add a connector in the graph: right-click in the graph and in the popup menu select "Add" and "Connector".

Next click in the graph and drag the mouse to create the connector.

To edit the (optional) text attached to the connector, to rotate the connector or to edit its characteristics, double click on the text box or right-click on the text box and select "Format connector...".

Selecting graphical objects

To select a graphical object:

- place the mouse pointer over the object and click the left mouse button
- small circles at the corners of the object (grab handles) indicate that the object is selected
- to add other objects to the selection, press the Ctrl key while clicking on the other objects.

Moving, resizing and deleting graphical objects

Graphical objects can be moved from their original location or can be changed in size or shape.

- select one or more objects
- click on one of the objects and move the mouse to move the selected objects, or use the arrow keys on the keyboard
- click on one of the graph grab handles (the small circles at the corners of the object) and move the mouse to resize the selected objects
- press the Delete key to delete the selected objects.

Reference lines

Reference lines are vertical or horizontal lines in a graph, corresponding with user-defined values on the x-axis and y-axis respectively.

Each graph can contain up to 16 vertical and horizontal reference lines.

Add reference line

Method 1

Right click in the graph and select "Add" and "Horizontal reference line" or "Vertical reference line".

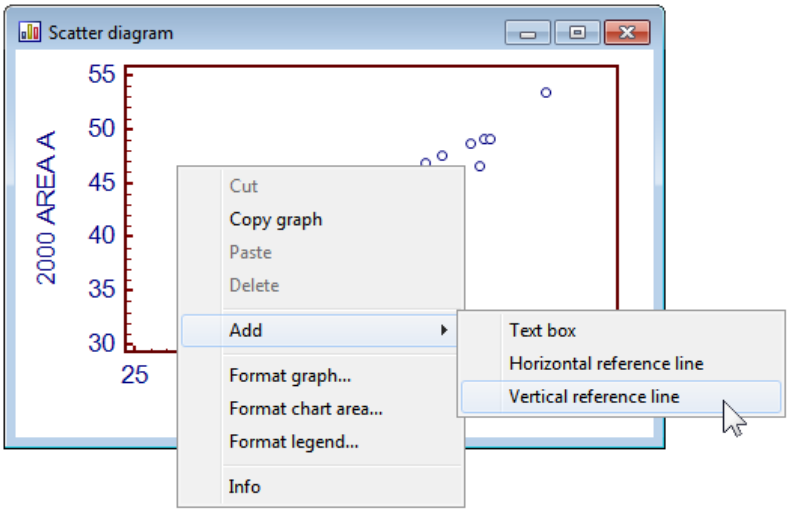

Method 2

Click in the margin of the graph to create a reference line and drag it into the graph plot area.

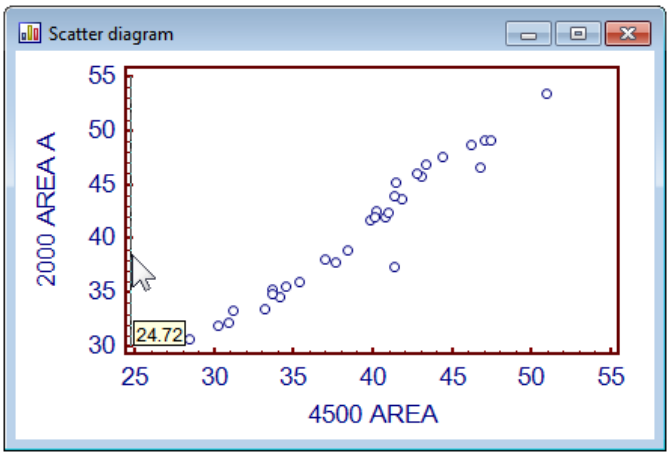

While you are dragging the reference line in the plot area, the corresponding value is displayed next to the reference line.

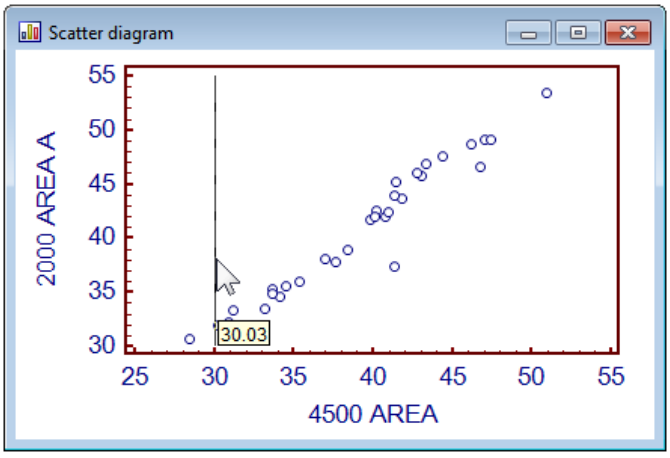

Format reference line

Right-click on the reference line and select "Format reference line" in the popup menu.

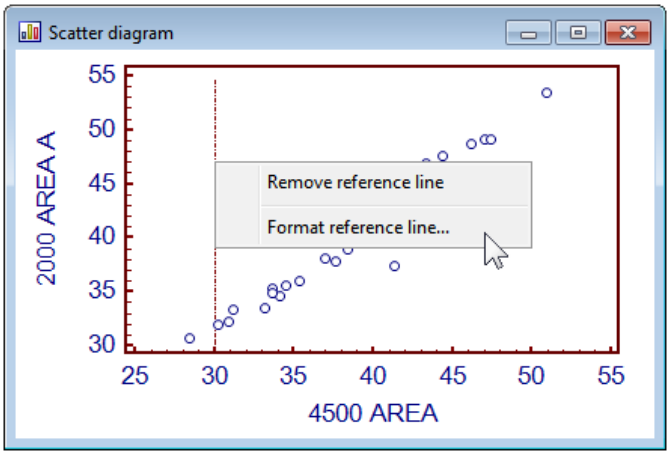

In the dialog box you can select line color, style and width, and enter an exact value for the reference line:

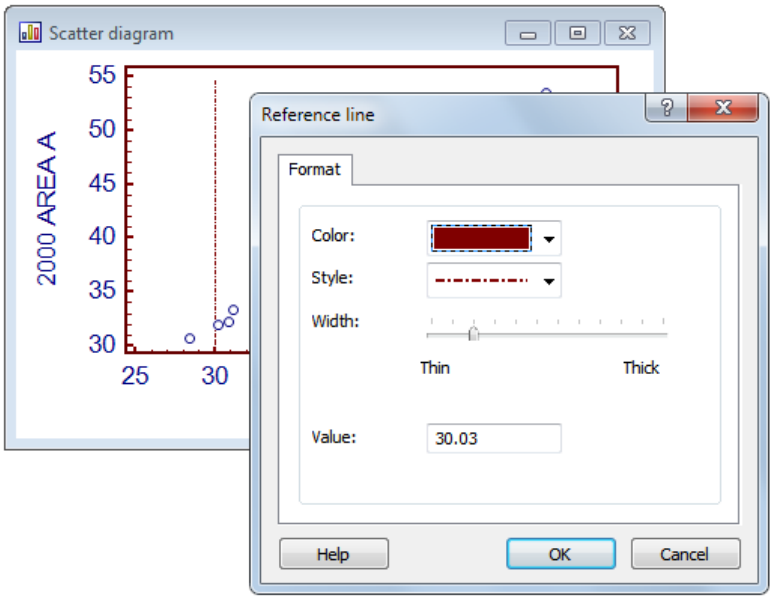

Note: all reference lines in one graph have the same color style and width.

Remove reference line

Method 1

Right-click on the reference line and select "Remove reference line" in the popup menu.

Method 2

Drag the reference line into the graph's margins.

Export metafile

When a graph window is active, you can select the *Export* command in the *File* menu to save the picture in Windows metafile format. Windows metafiles can be imported in almost every Windows word processor, drawing or presentation program (Microsoft Word, PowerPoint, etc.).

You can also transfer a graph from MedCalc to another Windows application without saving the file intermediately by selecting the *Copy* command in the *File* menu, activate the other application and there select the *Paste* command.

F7 - Repeat key

In order to repeat the last statistical analysis, you can press function key *F7* or click  . The last statistical dialog box will be redisplayed with all fields containing the same entries. For example, after you have completed the dialog box for *Correlation coefficient*, you may want to repeat the same analysis with a minor change in the dialog box, e.g. a different equation. You just need to press *F7* when the results of the first analysis are displayed without having to select anything in the menu.

Function key *F7* can also be used to enter the data from a previous dialog box in a new dialog box. For example, after you have created a scatter diagram, you want to calculate the correlation coefficient for the same variables. In order to obtain this, select *Correlation coefficient* in the *Statistics* menu. When the dialog box is displayed, you can press function key *F7* to enter the same variables as in the previous dialog box for the scatter diagram.

F7 is available to exchange information among the following ("compatible") dialog boxes:

- Summary statistics, Histogram, Cumulative distribution, Box-and-whisker plot, Normal plot, Reference interval
- Correlation, Rank correlation, Regression, Scatter diagrams
- Multiple regression, Logistic regression, Multiple variables graphs, Clustered multiple variables graphs
- T-tests, Wilcoxon tests, Data comparison graphs.
- One-way analysis of variance, Kruskal-Wallis test, Multiple comparison graphs, Clustered multiple comparison graphs.
- Bland & Altman plot, Passing & Bablok regression, Mountain plot.
- Frequency table, Frequency charts, Inter-rater agreement (Kappa), Fisher's exact test, McNemar test
- ROC curve and Comparison of ROC curves

Notes editor

Notes editor

The MedCalc *Notes editor* is an editor that you can use to edit short texts. These texts are associated with the spreadsheet data, and when you load or save the MedCalc data file, the notes are loaded or saved with it. When you want to save the text in a separate file, you have to use the *Export text* command.

The MedCalc notes offer for example the possibility to add explanations about codes used in the spreadsheet.

To create a new Notes document, right click **Notes** in the **Contents bar** (see p. 38) and select **New notes** in the shortcut menu:

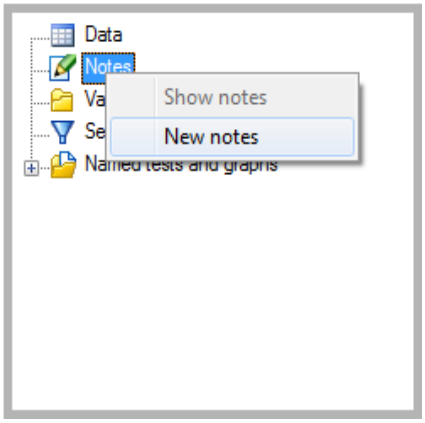

A new notes document is listed in the Contents bar with a default name "New notes". Right-click the New notes item to obtain the following shortcut menu:

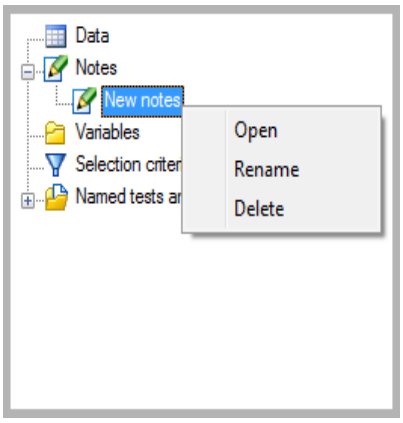

You can now open the notes document and start editing notes.

Text block functions

Deleting, moving and copying text blocks is possible by means of the commands in the *Edit* menu, or the corresponding buttons in the toolbar.

Export text

The text entered in the text window is always saved in the MedCalc data file. Therefore, after you have entered some text, select *Save* or *Save as* in the *File* menu to save the text, together with the data.

The *Export text* command can be used to save the text as a separate text file, which can be imported in a word processor.

FILE MENU

New

Select *New* if you want to close the data file, clear all data in the spreadsheet, freeing up the memory for new data. If you have already entered some data and the latest additions or changes have not been saved, then the program will display an appropriate warning. In this case you will have the option to save the data (*Yes*), or clear the data without saving the data (*No*), or cancel the *New* command, keeping the existing data in memory (*Cancel*).

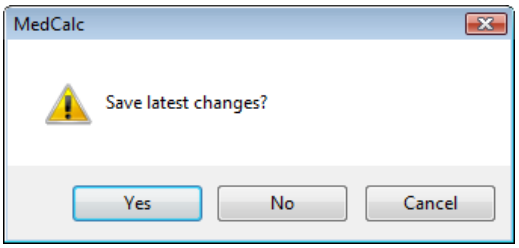

After having cleared all data in memory, the program will display a blank spreadsheet window so you can start entering new data.

Open

Button: 

Shortcut: Ctrl+O

If you want to retrieve a data file that you previously have saved on disk, select the *Open* command in the menu.

In the file selector box (described on p. 8) you select the name of the file to be read into memory. Together with the spreadsheet data the associated notes are loaded (see p. 20) and the options that are also included in the MedCalc data file.

If you want to read a MedCalc data file without erasing existing data in memory, e.g. to combine two files into one, then select the *Add file* command (see p. 26).

Supported data file types

MedCalc support the following file formats:

- MedCalc files (*.mc1)
- Excel files (*.xls, *xlsx, *.xlsm)
- SPSS files (*.sav)
- Stata files (*.dta)
- DBase files (*.dbf)
- Lotus files (*.wk1)
- SYLK (*.slk) and DIF (*.dif) files
- Text files (*.txt, *.csv, *.prn)

MedCalc files (*.mc1)

When you open a MedCalc data file (which have file extension MC1) MedCalc will lock the open file, preventing other programs and users access to the file. To close and unlock the file, you select the *New* command.

From time to time the MedCalc file format is revised. MedCalc will automatically convert old files as you open and save them. However, previous MedCalc versions may not be able to read the data files written by some newer versions of the software. It is advised that you always download and install the latest version of MedCalc (visit http://www.medcalc.org).

Excel files (*.xls, *xlsx, *.xlsm)

To import an Excel worksheet file, select the file type *Excel Worksheet*.

If the Excel file contains more than 1 worksheet, then the program will display a dialog box in which you can select the worksheet you want to import.

- MedCalc will convert all formulas in the Excel spreadsheet into their calculated values.
- The import procedure will import data files from Microsoft Excel version 2.1 and higher.
- Excel 2007 files (*.xlsx) are only supported on computers with
- Windows 7
- Windows Vista
- Windows XP with Service Pack 2 1 and Microsoft .NET Framework 3.0
- Windows Server 2003 with Service Pack 1 and Microsoft .NET Framework 3.0

SPSS files (*.sav)

When you open an SPSS data file (with file extension *sav*), MedCalc will present a *Select variables* dialog box in which you select the names of the variables to be imported in MedCalc.

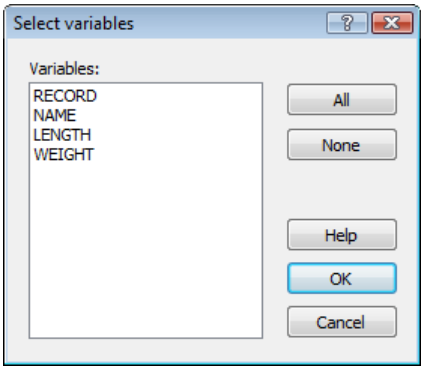

The available *variables* in the file are listed in the box at the left side of the dialog box. If you want to import all variables, i.e. all data in the file, then click the ALL button. If you do not want to import all variables, then select the names of the variables you require by clicking the variables' names. The names of the selected variables are displayed in reverse. To unselect a selected variable, click it again. By clicking the NONE button, you can cancel the selections.

When you have finished selecting the variables, click the OK button.

DBase files (*.dbf)

The file type *DBase file* must be used when you want to import a database file (with file name extension DBF) created by DBase III+ and compatible programs.

After you have selected the file name in the *File selector box*, select the names of the fields to be imported in MedCalc using the *Select variables* dialog box (see *Import SPSS file*).

After you have finished selecting the variables, click the OK button.

In DBase files *missing values* are often coded by the maximum value possible in the corresponding field (column). E.g. in a field with a maximum number of characters (field width) of 3, the number *999* will indicate a missing value. The number *9999* will be used to code missing values in a field with width of 4 characters, etc. After you have imported the DBase file, you must check whether this is the case. If it is, you can use the *Find & Replace* procedure (see p. 34) to clear the cells containing these missing values.

E.g. after you have imported the DBase file, you can open the spreadsheet window and inspect the data. When you see any values *999* in a column with a width of 3 characters, then this almost certainly means that missing values for this variable are indeed coded as *999*. You can place the cell pointer in the top row of this column, select the *Find & replace* option in the *Edit* menu, and enter the following in the dialog box:

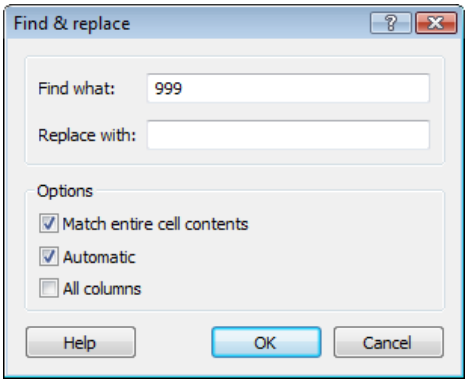

The *Replace with* field is left blank because in MedCalc missing values are indicated by an empty, blank, cell. After clicking the OK button, all *999* codes will be erased in the column.

If missing values in the DBase file were coded with a different number, for example *0*, you can take a similar approach, but you will have to be careful not to confuse the genuine *0* values with the missing value code *0*.

Lotus files (*.wk1)

Lotus files (*.wk1) are data files created by Lotus123 (version 2) and Symphony.

If the Lotus spreadsheet contains formulas that are not supported by MedCalc, it is recommended that you use the Lotus commands *Range Value* (/RV) to convert the formulas into values, before you import the file in MedCalc.

SYLK (*.slk) and DIF (*.dif) files

MedCalc can import both SYLK (Symbolic Link Format) and DIF (Data Interchange format) files.

Text files (*.txt, *.csv, *.prn)

In a text file, one text line contains data on one case, and every text line ends with a carriage return character (ASCII code 13).

In a *formatted text file*, the different fields (variables) are separated by spaces so that, when the file is opened with a text editor (such as Notepad), the variables appear as distinct columns, e.g.

```
LENGTH    WEIGHT
   171        66
   174      61.5
   182        82
   172        80
   179      72.0
   177     78.50
```

In a *delimited text file*, the data are separated by commas, and text is placed between single or double quotation marks, e.g.

```
"LENGTH","WEIGHT"
171,66
174,61.5
182,82
172,80
179,72.0
177,78.50
```

To read such a file select the *Text file* type in the dialog box and next you select the file name. The program will first check the text file to determine whether it is a formatted or delimited text file. The result is displayed in a new dialog box.

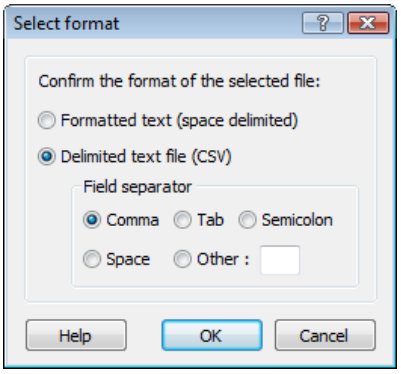

You must confirm the file type by clicking the OK button. If the program would have selected the wrong file format, you can choose the correct file format by selecting the corresponding button. Next click the OK button to proceed.

Finally, you will have to select the fields or variables to be imported in the *Select variables* dialog box (see *Import SPSS file*). If no name is available for a particular variable in the text file, then MedCalc will give the name FIELD*xx* to the variables, where *xx* is the number of the field (or variable) in the file.

Save

Button: 

Shortcut: Ctrl+S or F12

The *Save* command is used to save the data on disk.

If the data have not been saved in the native MedCalc file format (files with extension MC1) before, then MedCalc will present the File selector box and suggest the MedCalc data file format. It is recommended to save in this format.

A MedCalc data file contains the following information:

- spreadsheet data
- variables list
- selection criteria list
- named graphs and tests
- column width, precision, fill column specifications
- text entered in the *notes* window
- spreadsheet window position and size
- date of creation of file
- original Excel, Lotus, DBase or Text file name, if the current data have been imported from such data files
- other settings

MedCalc file format revisions

From time to time the MedCalc file format is revised. MedCalc will automatically convert old files as you open and save them. However, previous MedCalc versions may not be able to read the data files written by some newer versions of the software. It is advised that you always download and install the latest version of MedCalc (visit www.medcalc.org).

Save as

When you want to save the data as a new file, or save a file with a new file name, select the command *Save as*. Next, you enter a new name for the spreadsheet data in the file selector box (see p. 8).

MedCalc also allows saving the data in different file formats: Excel file, Lotus file, SYLK (Symbolic Link), DIF (Data Interchange Format) or Delimited text file (CSV).

Excel 2007 files (*.xlsx) are only supported on computers with

- Windows 7
- Windows Vista
- Windows XP with Service Pack 2 1 and Microsoft .NET Framework 3.0
- Windows Server 2003 with Service Pack 1 and Microsoft .NET Framework 3.0

Add file

If you want to read a MedCalc data file without first erasing the data in memory, then use one of the *Add file* commands.

- *Merge cells*: the cells in the file will replace existing cells in memory.
- *Append rows*: the data in the file will be placed in the open rows located under the rows already used in the spreadsheet. This option is useful to combine two files containing data on the same variables and with the same column headings.
- *Add columns*: the data will be placed in the columns at the right side of the columns already used in the spreadsheet. This option can be used to add the variables in the file to the variables already present in memory.

Export

Shortcut: F10

The *Export* command is used to export the contents of the top window. The top window can either be the spreadsheet data window and then the data can be exported as a plain text file.

If the top window is a statistics results window, these results can be exported as a text file. Finally, when the top window is a graph window, the graph can be exported as a Windows metafile.

Export data

When you want to export the data to another program, first open the spreadsheet window, or activate it and bring it on top of possible other open windows. When the spreadsheet window is on top, you can select *Export data* in order to save the spreadsheet data into one of the following file formats: MedCalc file, Excel file, Lotus file, SYLK (Symbolic Link), DIF (Data Interchange Format) or Delimited text file (CSV).

Excel 2007 files (*.xlsx) are only supported on computers with

- Windows Vista, Windows 7
- Windows XP with Service Pack 2 and Microsoft .NET Framework 3.0
- Windows Server 2003 with Service Pack 1 and Microsoft .NET Framework 3.0

Export results

When the top window is a statistical text results window, then the *Export results* command can be used to save the text of the results window as a separate text file. This text file can be imported in any word processor program.

Export graph

When the top window is a graph window, you can select *Export graph* in order to save the graph in Windows metafile format. This file can then be imported in a drawing or presentation program (e.g. PowerPoint) for further editing and enhancement.

Alternatively, you can export the graph as a *Device Independent Bitmap* (BMP), GIF format (*Graphics Interchange Format*), PCX (PC Paintbrush), TIF, PNG or JPG format.

To control the height and width of the saved image, click the "Options" button in the "Export graph as" dialog box.

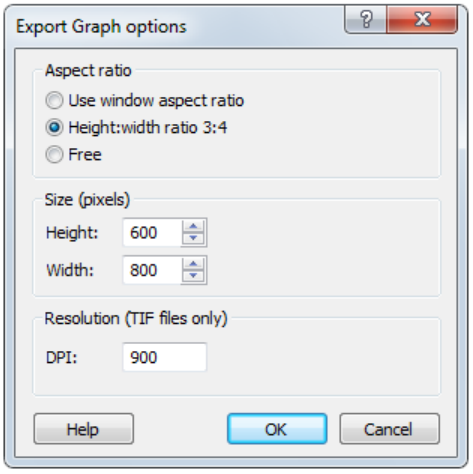

Aspect ratio

Select the aspect ratio of the image.

- **Use window aspect ratio**: the height:width ratio of the image will be the same as the height:width ratio of the graph window as displayed on the screen.
- **Height:width ratio 3:4**: the height:width ratio of the image will be 3:4
- **Free**: you can select the height and width independently.

Size (pixels)

Select the height and width, in pixels, of the image saved to disk.

Export notes

When the top window is the notes editor window, then the *Export notes* command can be used to save the notes as a separate text file. This text file can be imported in any word processor program.

Page setup

Select *Page setup* to select the page orientation (portrait, landscape), and print margins for the active document (spreadsheet, text results or notes).

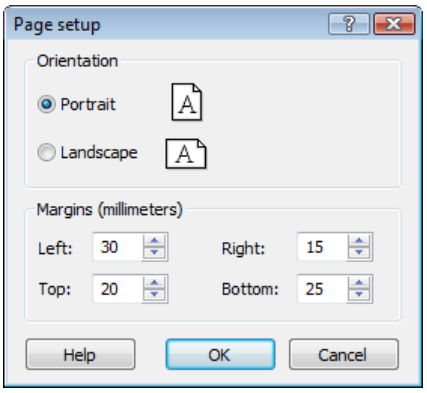

Print

Button: 

Shortcut: Ctrl+P

F9

The *Print* command is used to print the contents of the top window: either the spreadsheet, a statistical text results window, graph window, or the notes window.

After you have selected the *Print data*, *Print Notes* or *Print results* command, the program will display the following (or similar) dialog box:

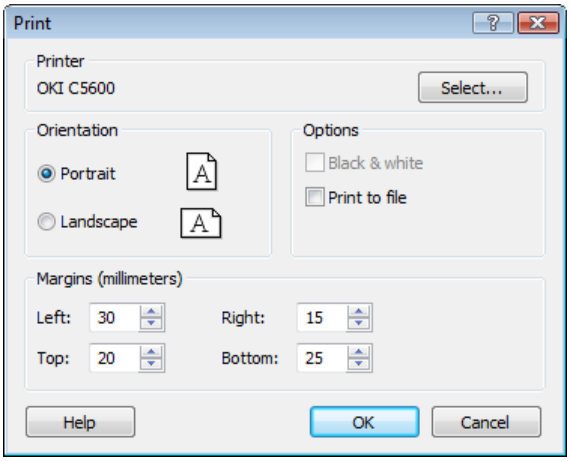

In this dialog box, you can select the printer by clicking the SELECT button.

Select the *Black & white* option in case you want to print in black and white only. When you select the *Print to file* option, the output will be redirected to a file and not to the printer.

You can select the page orientation (portrait, landscape), and print margins.

Print data

When you want to print the data, you open the spreadsheet window, or activate it and bring it on top of possible other open windows. Next, select *Print data* in the *File* menu or alternatively, press *Ctrl+P*.

The program will automatically split the columns and rows in order to fit the printed output on the paper. The numbering of the printed pages is indicated in the following diagram (the outer rectangle represents the complete spreadsheet):

Page 1-1	Page 2-1	Page 3-1
Page 1-2	Page 2-2	Page 3-2
Page 1-3	Page 2-3	Page 3-3
Page 1-4	Page 2-4	Page 3-4

Print results

Results of statistical calculations, summary statistics, correlation and regression are displayed in the *results* text window. When you want to obtain a printed copy of these results, select the *Print results* command, or press *Ctrl+P*.

Print graph

When the top window is a graph window, then you can select *Print graph* to print the graph.

In the *Preview* area of the *Print graph* dialog box, you can size and position the graph on the paper.

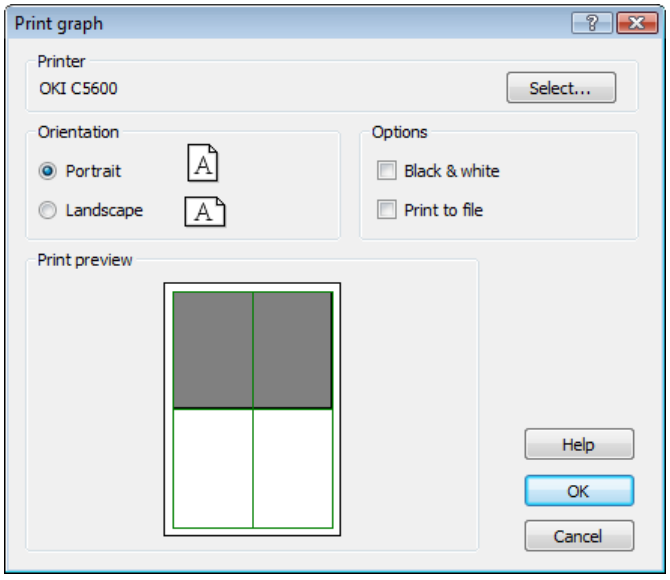

Print notes

When you want to obtain a printed copy of the text entered in the *Notes editor* window, select the *Print notes* command, or press *Ctrl+P*.

Properties

The *Properties* box displays the following information about the data and file present in memory:

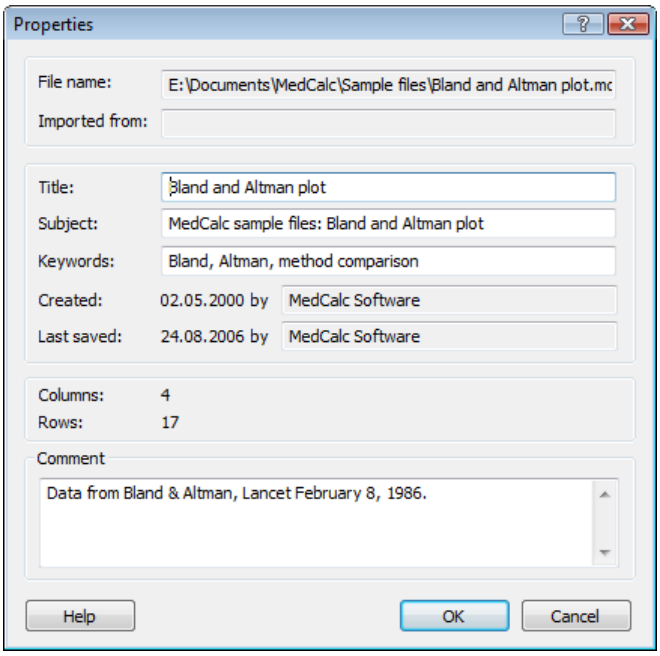

- **File name**: the name of the file loaded from disk, or NONAME.MC1 when no file has been loaded;
- **Imported from**: if the file was created by importing a Excel, Lotus, DBase or Text file, then the name of this original file is displayed;
- **Title**: the title you want to use when searching for this file. The title can be different from the file name;
- **Subject**: the subject of the file. Use this property to group similar files together, so you can search for all files that have the same subject.;
- **Keywords**: list of keywords describing the file's content, and you want to use when you search for this file.
- **Created**: the date the MedCalc file was created, or today's date if no MedCalc file has been loaded from disk;
- **Last saved**: the date when the file was saved;
- **Columns**: the number of columns used in the MedCalc spreadsheet (minimal 1);
- **Rows**: the number of rows used in the spreadsheet.
- **Comments**: comments associated with this file. Can be used when you search for this file.

Exit

Shortcut: Alt+F4

The *Exit* command must be selected to stop the program. Be sure to have saved all data before selecting this command and returning to the operating system.

However, if the latest changes have not been saved, then the program will display an appropriate warning. In this case you will have the option to save the data (Yes), or exit the program without saving the data (*No*), or cancel the *Exit* command and stay in the program (*Cancel*).

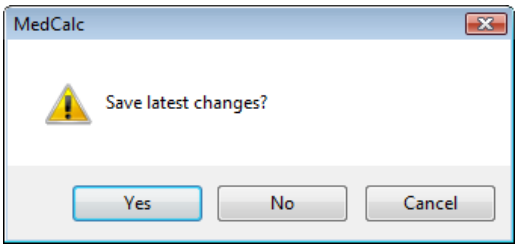

Double-clicking the MedCalc application window's Control-menu box is the same as choosing the *Exit* command from the *File* menu.

EDIT MENU

Undo

Button: 

Shortcut: Ctrl+Z

With the *Undo* command you can undo the last edit operation in the spreadsheet or notes editor window.

Cut

Button: 

Shortcut: Ctrl+X

With this command you can *remove the selection* from the window and place it on the clipboard. The selection can be a text, a text in a cell, a cell or a range of cells.

The clipboard is a temporary area in memory for data that you cut or copy. From the clipboard, you can paste the cut or copied data to another location or to another document.

To completely remove a row or column in the spreadsheet, and shift other rows and columns to fill the space, use the *Remove* command in the *Edit* menu (p. 36).

Copy

Button: 

Shortcut: Ctrl+C

With this command you can *copy the selection* from the window and place it on the clipboard. The selection can then be pasted at a new location. In the spreadsheet window, the selection can be a text, a text in a cell, a cell or a range of cells.

Spreadsheet data can be transferred to other programs by selecting the range of cells you want to copy, followed by selecting the *Copy* command. Next activate the other program and select *Paste* in the *Edit* menu.

You can also copy the MedCalc graphs to the clipboard. Next you can open a drawing or presentation program (PowerPoint) and paste the graph into it. By default, the graph is copied onto the clipboard with a height/width ratio of 2/3. If you want to use the height/width ratio of the graph window, you must select the option “Use window aspect ratio” in the *Options* box (see p. 56).

Paste

Button: 

Shortcut: Ctrl+V

The Paste command is used to paste the contents of the clipboard onto the active window.

If the *Cut-Copy-Paste* commands are used to move of copy a range of spreadsheet cells, then choosing this command pastes the contents of copied or cut cells into the area indicated by the cell pointer where the cell pointer is the upper-left corner of the paste area.

The *Paste* command can also be used to transfer data from Excel to MedCalc. First activate Excel and select the range of cells you want to transfer, followed by selecting the *Copy* command. Next activate MedCalc, open the spreadsheet window, position the cell pointer on the target location and select *Paste* in the *Edit* menu.

Delete

Shortcut: Del

With the *Delete* command you can clear the contents of a cell or a selected range of cells, a selected row or column. The *Delete* command does not affect the clipboard.

To completely remove a row or column in the spreadsheet, and shift other rows and columns to fill the space, use the *Remove* command in the *Edit* menu (p. 36).

Select all

Shortcut: Ctrl+A

This command is used to select the entire contents of the current window (does not affect the clipboard).

Find

Shortcut: Ctrl+F

The *Find* command is used to search for a specific string or data entry, starting from the current cell pointer position.

First the current column is searched through, starting at the current cell pointer position, from top to bottom. If no match is found, you are offered the choice to continue the search in the next columns.

The search specification can contain the ? and * wildcards. The ? wildcard indicates that the character on this position can be ignored and, * indicates that the rest of the cell can be ignored. For example: Jones matches Jon* and Jo?es.

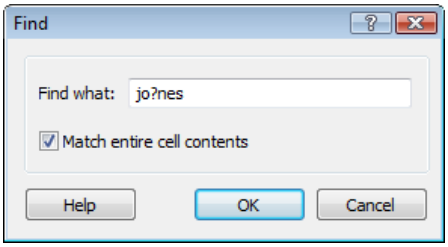

If the option *Show formulas* is selected and checked (see p. 41), then the search string is compared with the literal or raw contents of the cell, otherwise it is compared with the cells' calculated values. E.g. if cell B10 contains SIN(90), resulting in 1.00, and the search string is 'SIN*', then cell B10 will be ignored when *Show formulas* is OFF.

FIND does not distinguish uppercase from lowercase characters.

Options

- *Match entire cell contents*: Searches for an exact and complete match of characters that are specified in the Find what box.

Find & replace

Shortcut: Ctrl+H

The *Find & replace* procedure can be used to replace certain values in the spreadsheet with another value, or clear the cells containing these values. The value to be located and to be replaced is entered in the *Find* field of the dialog box. The new value is entered in the *Replace with* field:

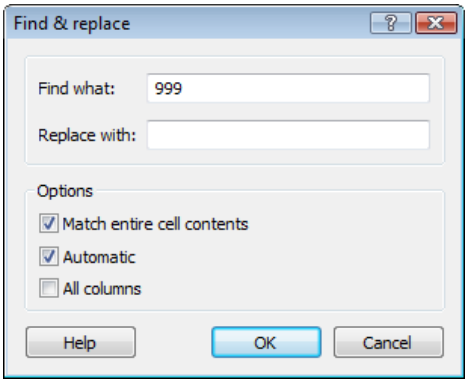

Options

- *Match entire cell contents*: Searches for an exact and complete match of characters that are specified in the Find what box.
- *Automatic*: replacements will be performed automatically. If this option is not selected, then the program will display a dialog box before the old cell content is replaced with the new text.
- *All columns*: all columns will be searched through. If this option is not selected, then the find & replace procedure will be limited to the current column (the column containing the cell pointer).

You do not have to enter any text in the *Replace with* input field. In this case, the cells containing the *find* text, will be cleared.

Refer to the *Find* command for some explanation concerning the search wildcards * and ?, and the way MedCalc performs text comparison and treats the search specification.

Remember that searching always starts from the current cell pointer position. So when you want to search an entire column, you must put the cell pointer in the top cell of this column.

Go to cell

Shortcut: Ctrl+G

Allows to quickly move the spreadsheet cellpointer to any cell in the spreadsheet. This dialog also appears when you double-click on the cell address in the MedCalc spreadsheet formula bar.

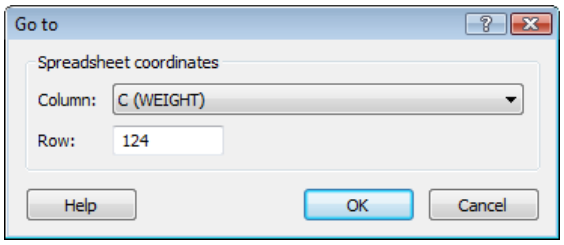

Required input

- **Column**: select a spreadsheet column.
- **Row**: enter a spreadsheet row.

Fill

Fill down

Shortcut: Ctrl+D

Use this command to copy the contents of the topmost cells of a selected range into the cells below.
Example:
Before:

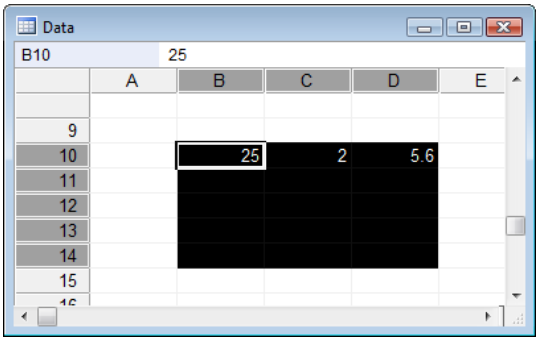

After:

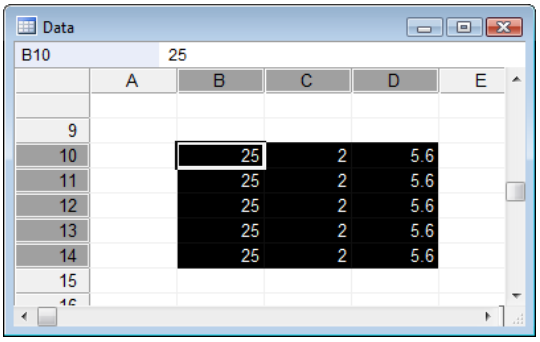

Fill right

Shortcut: Ctrl+R

Use this command to copy the contents of the leftmost cells of a selected range into the cells to the right.
Example:
Before:

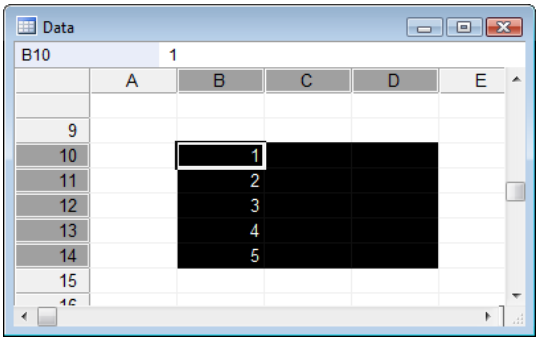

After:

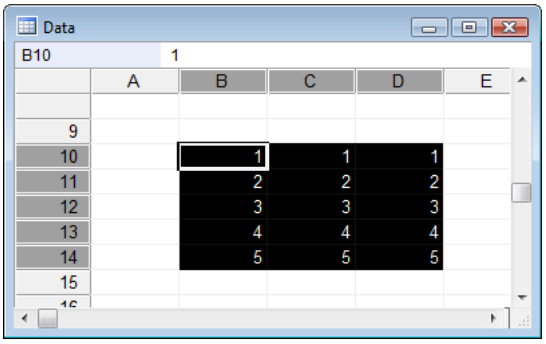

Fill series

Shortcut: Ctrl+L

Use this command to fill the selected range of cells with a series of numbers, characters, or dates. The content of the first cell in the selected range is used as the starting value for the series. Fill details can be given in the subsequent *Fill series dialog box*.

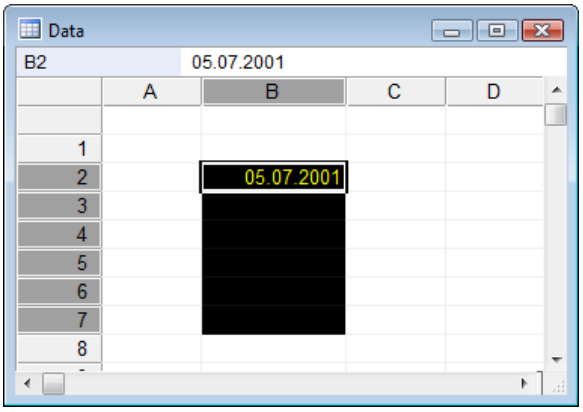

In this dialog box you enter:

- **Start value**: Enter the start value for the series. The start value must be a number, character or properly formatted date.
- **Step value**: Enter a positive or negative number to indicate the amount by which you want a series to increase or decrease.
- **Step unit (dates)**: Specify whether the series of dates will increase by days, weeks, months, or years. Available only when creating a date series.

Insert - Remove

To insert or remove the spreadsheet row or column where the cell pointer is located, select the corresponding command in the *Edit* menu.

Examples:

- When the cell pointer is located in row 20 and you select the command *Insert row* in the menu, then all rows starting with row 20 will be moved down and a new row will be inserted at this position. So row 20 will be moved to row 21, row 21 to 22, etc., and row 20 will be empty.
- When the cell pointer is located in column B and you select the command *Remove column* in the menu, then column B will be removed and all columns starting with column C will be moved to the left. So the data in column B will be lost, and column C will be moved to column B, column D to C, etc.

Transpose

With this command you can switch the orientation of data, either in a selected area, or in the complete spreadsheet.

Data from the top row of the area to transpose will appear in the left column of the transposed area, and data from the left column will appear in the top row.

When you want to transpose only a particular area of the spreadsheet, you first select this area with the mouse. In this case, the number of rows and the number of columns of the selected area must be equal (while selecting, this can be checked in the status bar).

VIEW MENU

Spreadsheet

Button: 

This command is used to open a spreadsheet window so you can enter new data, or examine the data already entered in the spreadsheet.

The active cell in the spreadsheet is indicated by the cell pointer. The cell pointer can be moved to another cell by pressing the arrow or cursor keys (see Appendix A for an overview of control keys and special key functions).

When you want to put new data in a cell, move the cell pointer to the cell and enter the data on the keyboard.

To change the content of a cell, move the cell pointer to this cell and press function key *F2*. The cell content will be displayed at the top of the window. You can now change the cell, and edit the data. The ← and → arrow keys can be used to move the cursor in the cell (just like editing a text in a word processor). The *Backspace* key is used to delete the character at the left side of the cursor, and the *Delete* key deletes the character at the right side. If desired, you can erase the entire cell by pressing the *Backspace* key repeatedly. You can use the *Insert* key to toggle between *Overwrite* and *Insert* editing mode.

The new content of the cell is stored in memory by moving the cell pointer up or down. By pressing the *Enter* key you can store the new or edited contents of a cell without moving the cell pointer.

To copy the contents of one cell to another cell, place the cell pointer on the first cell (the source cell), and select *Copy* in the *Edit* menu (you can obtain a shortcut menu by clicking the right side mouse button). By doing so, the content of this cell is copied into a buffer in memory. Now you move the cell pointer to the second cell (the target cell) and select *Paste* in the *Edit* menu. As a result, the content of the buffer is copied into this second cell.

If you want to copy a range of cells, you first select the cells with the mouse, and proceed as described above.

Show formulas

When the option *Show formulas* is checked then the spreadsheet will display the formulas that you have entered, otherwise the results of the formulas are displayed.

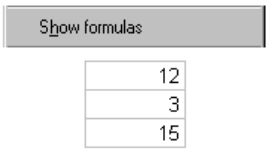

The option *Show formulas* is not checked: the program displays the results.

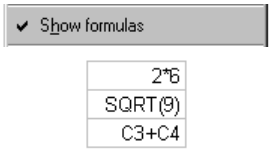

The option *Show formulas* is checked: the program displays the formulas.

If you have to enter a lot of data (on a slow computer), or when you want to see the possible relationship between cells, then it is convenient to select this option.

Show gridlines

Select this option if you want gridlines to be displayed in the spreadsheet window. When the command has a check mark next to it, the *Show gridlines* option is on.

Contents bar

Button: 

The contents bar displays the contents of the MedCalc data file in a treeview structure.

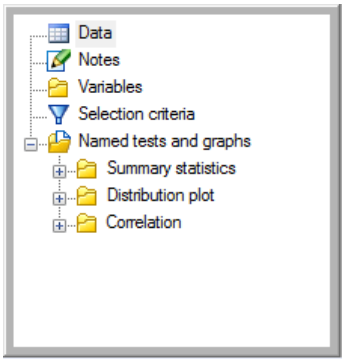

Select **Data** to open the MedCalc spreadsheet if you want to enter or edit data.

Select **Notes** if you want to write comments or notes to be saved in the MedCalc data file.

Select **Variables** to edit the Variables list. You can delete variables, create derived variables, and give descriptive labels to the variables.

Select **Selection criteria** to edit the Selection criteria list. You can create and delete selection criteria, and give descriptive labels.

In the **Named test and graphs** section you can select a named graph, which you can restore, rename or delete (right-click an item to obtain a shortcut menu).

Named tests and graphs

You can save a statistical analysis or graph in the MedCalc data file by clicking the  button in the toolbar. By doing so, the graph is added to the *Named tests and graphs list* in the *Contents* bar. A default name is given to the graph, which you can edit in the Contents bar.

This feature is very useful if you would want to recreate the graph later, possibly with new data.

Toolbars

Select the corresponding command when you want to show or hide the following toolbars:

- Standard toolbar
- Formatting toolbar

Toolbar options

- Merged toolbars: merge the two toolbars into one (all buttons will be displayed on one single row)

Check marks are displayed in the menu when the corresponding toolbar/option is selected.

Status bar

This command displays or hides the status bar at the bottom of the MedCalc main window, which shows information about the current state of the program.

A check mark is displayed in the menu when the Status bar is displayed.

Full screen

In full screen view, the MedCalc application window is maximized and status bar as well as toolbars are removed. You can restore these by using the appropriate commands in the menu (see above).

MedCalc keeps a separate set of toolbar options for Full screen and Normal view mode.

A check mark is displayed in the menu when the program operates in Full screen view mode. Click **Full screen** again to return to normal view mode.

FORMAT MENU

Character

Use this command to select a font for the top spreadsheet, notes, results or graph window. Alternatively, you can use the font and font size selectors in the formatting toolbar:

Increase font size

Button: 

Select this command to increase the font size for the top spreadsheet, notes, results or graph window.

Decrease font size

Button: 

Select this command to decrease the font size for the top spreadsheet, notes, results or graph window.

Format spreadsheet

Button: 

In the *Format spreadsheet* dialog box, you can select the following formatting options:

Sheet tab

In the Format spreadsheet box, you can select the following options:

- **Default column width and decimals**: The default column width and global number of decimals affect those columns for which a specific width and number of decimals has not been specified.

Column tab

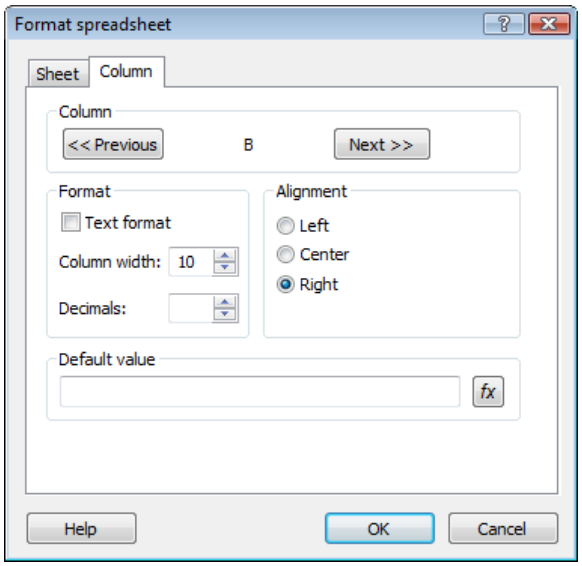

- **Text format**: When you select the option Text format, all cells in this column will be treated as text, even when a number or formula is entered. The cells will be displayed exactly as entered.

  This option is particularly useful when the column contains a text string that should not be interpreted as a number, e.g. date: 05/07/1956 will be displayed as such, not as 0.0003652 (= 5 divided by 7 divided by 1956). (see How to enter dates)
- **Width, decimals**: You can define the column's width and number of decimals.

  After you have specified a specific width or number of decimals for a column, then changing the global width or global number of decimals (in the *Format spreadsheet* dialog box, p. 40) will have no effect on the cells of this column. To reset the width and/or number of decimals of a column to the global width and/or global number of decimals, you must blank the corresponding input field and click OK.
- **Alignment**: You can select whether cell content should be centered or left or right justified in the column (default is right justified).
- **Default value**: The default value or formula for every cell in this column (this value or formula is only applied to empty cells or cells for which no value has been entered). Default values are displayed in a different color in the spreadsheet, see *Options* dialog box (p. 54)

Format graph

Button: 

Use this command to select different formatting options for the top graph window. Graph formatting is described on page 12.

Reset graph titles and options

Use this command to undo all user edits and reset the graph titles and options to default values and apply the default color scheme.

TOOLS MENU

Sort

The *Sort* command is used when you want to sort a range of rows of the spreadsheet according to the contents of a specific column.

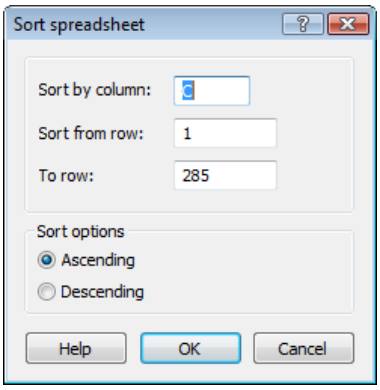

If the option *Show formulas* is not selected (see p. 41) then the spreadsheet is sorted according to the calculated values of the cells. If on the other hand the option *Show formulas* is selected then the literal content of the cell is used for sorting.

Note:

- Sorting the data in the spreadsheet is not required for any of the MedCalc statistical tests or procedures.
- The *Sort* command is only available when the MedCalc spreadsheet is open and active.

Exclude - Include

If you need to exclude data from statistical analysis, then you:

- Select the data to exclude in the spreadsheet; the selection may either be a range of cells, or one or more complete rows (cases).
- Select the *Exclude* command in the *Tools* menu.

To include data that have been excluded, you:

- Select the data to include in the spreadsheet; the selection may either be a range of cells, or one or more complete rows (cases).
- Select the *Include* command in the *Tools* menu.

Data that have been marked as excluded in the spreadsheet will be displayed in red, but you can select a different color in the *Format Spreadsheet* box (p. 40).

The *Exclude* and *Include* commands are also available in the spreadsheet shortcut menu (after right-clicking).

Note: these commands are only available when the MedCalc spreadsheet is open and active.

Fill column

With the *Fill column* command you can fill an entire column, or a range of rows in a column, with a formula.

You can select two options:

- *Convert formulas to values*: to convert the formulas to their calculated values.
- *Fill empty cells only*: to fill only empty cells with the new formula, and so saving the contents of the non-empty cells in the column.

Click the fx button to call the *Formula editor* dialog box (see Variable editor on p. 53).

In the following example, rows 1 to 100 of column A will be filled with the result of the RANDNORM(0,1) function. This function generates random numbers drawn from a Normal distribution with mean 0 and standard deviation 1.

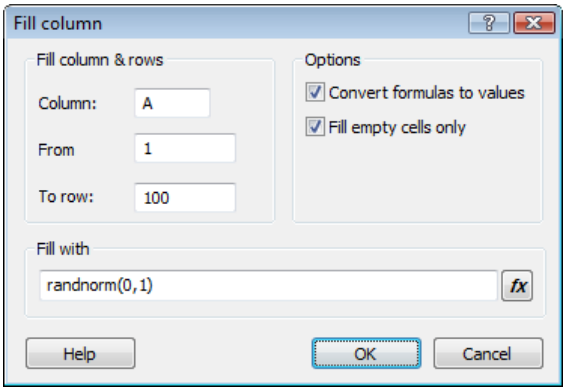

When you save the data then the formulas that you have entered in this dialog box will also be saved in the MedCalc data file, so you can easily re-use them at a later stage.

The *Fill column* command can also be used to clear a range of cells in a column, by letting the *Fill with* field blank.

Some examples of useful formulas are given in the following table.

Fill with:	**Result:**
SQRT(LENGTH)	fill with the square root of variable *LENGTH*
ROW	the cells will be filled with the row number
IF(*var*<20,"A","B")	recoding of variable *var* into two classes A and B
RAND(2)	the cells will be filled with 1 or 2 at random
RANDNORM(0,1)	fill with random numbers from a Normal distribution with mean 0 and standard deviation 1
VAR1+VAR2	fill with the sum of the variables *VAR1* and *VAR2*
	the cells will be cleared (empty *Fill with* field)

- Note: the *Fill column* command is only available when the MedCalc spreadsheet is open and active.

Stack columns

With the **Stack columns** tool you can perform a multivariate to univariate conversion by stacking several spreadsheet columns into a new column.

The procedure creates 3 new columns for data, variables and cases.

Example

Original data:

	A	B	C
	Patient	Hour12	Hour24
1	NP	3.87	3.1
2	AM	4.65	3.26
3	FS	2.34	3.18
4	CV	2.45	2.72
5			

Variables **Hour14** and **Hour24** stacked with variable **Patient** as variable with case identification:

	A	B	C	D	E	F
	Patient	Hour12	Hour24	Case	Hour	Data
1	NP	3.87	3.1	NP	Hour12	3.87
2	AM	4.65	3.26	AM	Hour12	4.65
3	FS	2.34	3.18	FS	Hour12	2.34
4	CV	2.45	2.72	CV	Hour12	2.45
5				NP	Hour24	3.1
6				AM	Hour24	3.26
7				FS	Hour24	3.18
8				CV	Hour24	2.72

This function may be useful when you have to rearrange your data for procedures such as Serial measurements analysis (p. 127), or Bland-Altman plot with multiple measurements per subject (p. 138).

Required input

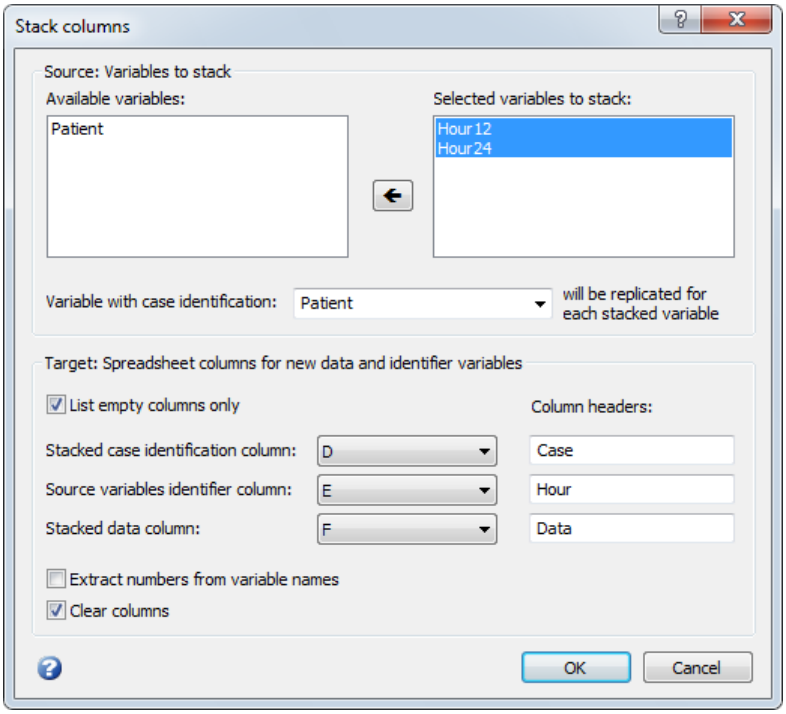

- Source: Variables to stack
  - **Variables**: select the variables of interest in the top left box and next click the right arrow button to move the selection to the Selected variables list.
  - **Variable with case identification**: a variable that will be replicated in one column for each stacked variable. If left blank, the row number will be used for case identification.
- Target: Spreadsheet columns for new data and identifier variables.
  - Option **List empty columns only**: if this option is selected, only those columns that are empty will be listed in the following selection boxes.
  - **Stacked case identification column**: a new spreadsheet column that will be filled with a case identification, replicated for each stacked variable. (column Case in the example above)

- **Source variables identifier column**: a spreadsheet column that will be filled with the variables names (column Hour in the example above)
- **Stacked data column**: a spreadsheet column that will be filled with the stacked data (column Data in the example above)
- Option **Extract numbers from variable names**: if the variable names contain numbers, like in the example Hour12 and Hour24, the numbers 12 and 14 will be extracted from the variables names and placed in the Source variables identifier column (see example below)
- Option **Clear columns**: the selected columns will be cleared prior to storing the stacked data.

Click OK to proceed.

More examples

Example with the option "Extract numbers from variable names" selected:

	A	B	C	D	E	F
	Patient	Hour12	Hour24	Case	Hour	Data
1	NP	3.87	3.1	NP	12	3.87
2	AM	4.65	3.26	AM	12	4.65
3	FS	2.34	3.18	FS	12	2.34
4	CV	2.45	2.72	CV	12	2.45
5				NP	24	3.1
6				AM	24	3.26
7				FS	24	3.18
8				CV	24	2.72

The numbers 12 and 24 are extracted from the variables names, making the variable **Hour** suitable for numerical analysis or inclusion in numerical procedures.

Example with the selection "Variable with case identification" left blank:

	A	B	C	D	E	F
	Patient	Hour12	Hour24	Case	Hour	Data
1	NP	3.87	3.1	1	12	3.87
2	AM	4.65	3.26	2	12	4.65
3	FS	2.34	3.18	3	12	2.34
4	CV	2.45	2.72	4	12	2.45
5				1	24	3.1
6				2	24	3.26
7				3	24	3.18
8				4	24	2.72

The row number is used as case identifier.

Generate random sample

Using this command you can generate a series of uniform or normal distributed random numbers.

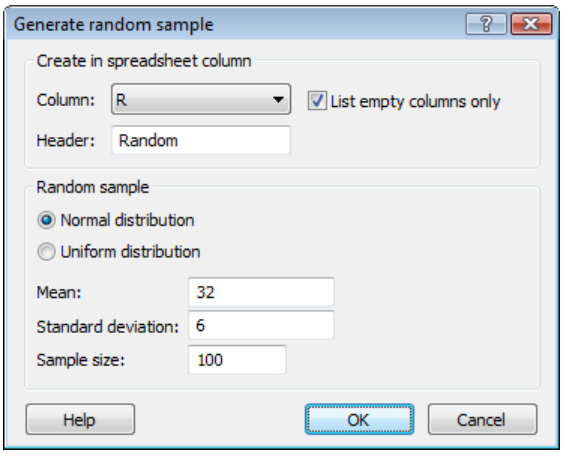

Required input

- **Column**: the column in which you want to place the random numbers. Option **List empty columns only**: if this option is selected, only empty columns are listed in the column selection box.
- **Header**: the header (top cell) for the selected column.

You can select a random sample with Normal distribution or Uniform distribution.

- For **Normal distribution** enter the desired mean and standard deviation, and the required sample size.
- For **Uniform distribution** enter the minimum and maximum values and the required sample size. If minimum and maximum values both are whole numbers, the program will generate whole random numbers only.

Click OK to proceed. The selected column in the spreadsheet is filled with the requested number of random values (all other cells in the column are cleared).

Create groups from quantiles

This tool allows to categorize a continuous variable by generating a new categorical variable with group numbers based on the quantiles of the continuous variable.

Required input

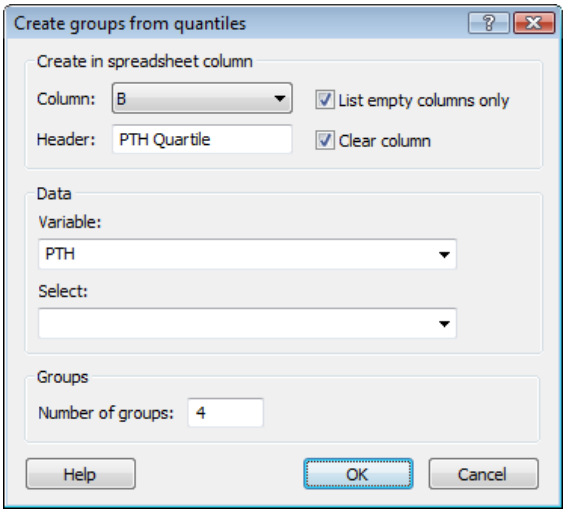

- **Column**: the column in which you want to place the group number.
- Options
  - **List empty columns only**: if this option is selected, only empty columns are listed in the column selection box.
  - **Clear column**: the selected column will be cleared prior to generating and storing the group numbers.
- **Header**: the header (top cell) for the selected column.
- **Data**: select the continuous variable and a possible selection criterion.
- **Number of groups**: the required number of groups, e.g. enter 4 to create groups based on the quartiles of the selected continuous variable.

Click OK to proceed. The selected column in the spreadsheet is filled with the group numbers (e.g. 1, 2, 3 and 4) corresponding with the quantiles of the continuous variable.

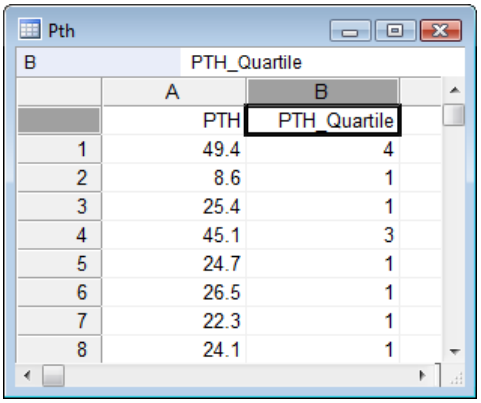

	A	B
	PTH	PTH_Quartile
1	49.4	4
2	8.6	1
3	25.4	1
4	45.1	3
5	24.7	1
6	26.5	1
7	22.3	1
8	24.1	1

Create random groups

This tool allows to assign cases to random groups.

Required input

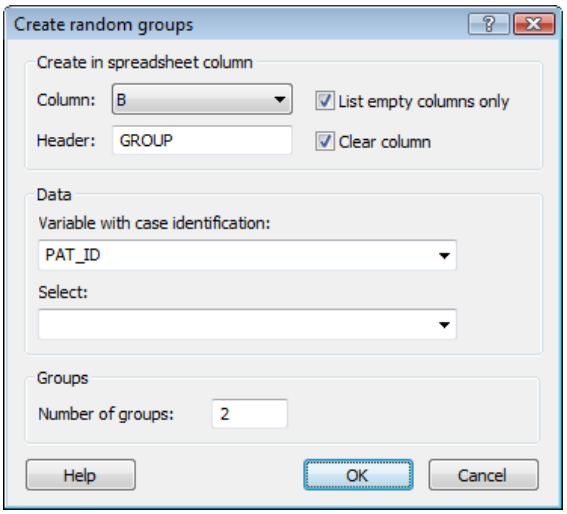

- **Column**: the column in which you want to place the group number.
- Options:
  - **List empty columns only**: if this option is selected, only empty columns are listed in the column selection box.
  - **Clear column**: the selected column will be cleared prior to generating and storing the group numbers.
- **Header**: the header (top cell) for the selected column.
- **Data**: select a variable that contains a case identification, and a possible selection criterion.
- **Number of groups**: the required number of groups, e.g. enter 2 to create 2 random groups.

Click OK to proceed. The selected column in the spreadsheet is filled with a random group number.

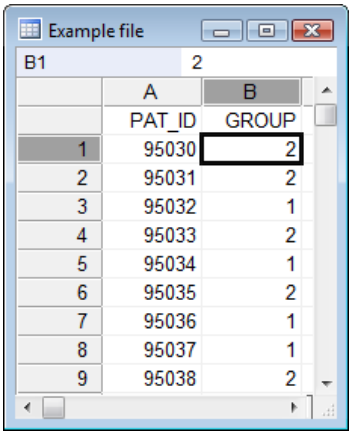

Create user-defined groups

This command allows to assign cases to groups based on a combination of user-defined criteria.

Required input

- **Column**: the column in which you want to place the group number.
- Options:
  - **List empty columns only**: if this option is selected, only those columns that are empty will be listed in the previous selection box.
- **Header**: the header (top cell) for the selected column.
- **Criteria and categories**: enter up to 7 conditions and group identifiers and a default group identifier.

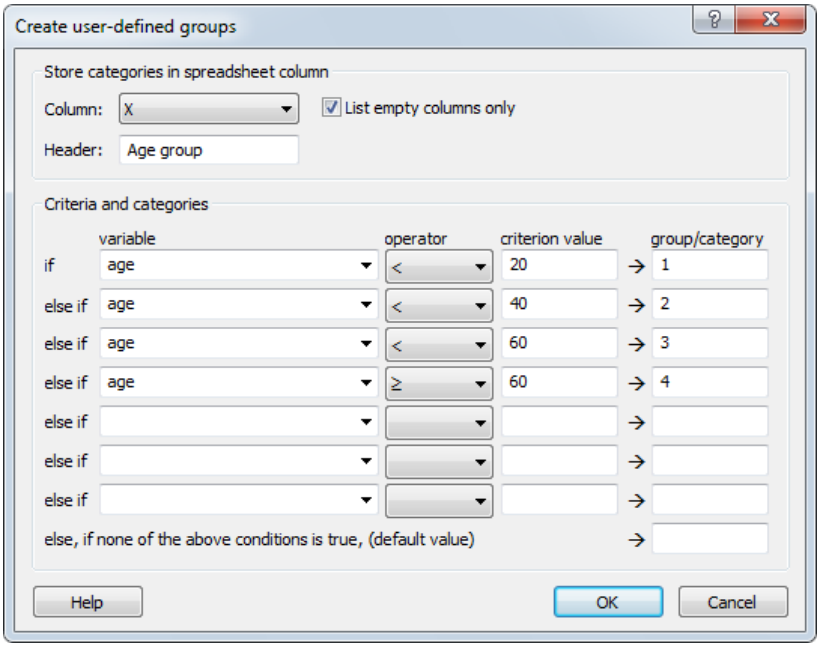

Click OK to proceed. The selected column in the spreadsheet is filled with group/category identifiers according to the different criteria. The default group identifier is used for a case when none of the selected conditions is true for that case.

Rank cases

Description

Allows to create a new variable containing the rank numbers of the data in a numeric variable.

Required input

- **Column**: the column in which you want to place the rank numbers.
- Options
  - List empty columns only: if this option is selected, only empty columns are listed in the column selection box.
  - Clear column: the selected column will be cleared prior to generating and storing the rank numbers.
- **Header**: the header (top cell) for the selected column.
- **Data**: select the numeric variable and a possible selection criterion.

Click OK to proceed. The selected column in the spreadsheet is filled with the rank numbers of the selected data.

Percentile ranks

Description

Allows to create a new variable containing the Percentile ranks of the data in a numeric variable.

Required input

- **Column**: the column in which you want to place the Percentile ranks.
- Options
  - **List empty columns only**: if this option is selected, only empty columns are listed in the column selection box.
  - **Clear column**: the selected column will be cleared prior to generating and storing the Percentile ranks.
- **Header**: the header (top cell) for the selected column.
- **Data**: select the numeric variable and a possible selection criterion.

Click OK to proceed. The selected column in the spreadsheet is filled with the Percentile ranks of the selected data.

z-scores

Description

Allows to create a new variable containing the z-scores of the data in a numeric variable.

Required input

- **Column**: the column in which you want to place the z-scores numbers.
- Options
  - List empty columns only: if this option is selected, only empty columns are listed in the column selection box.
  - Clear column: the selected column will be cleared prior to generating and storing the z-scores.
- **Header**: the header (top cell) for the selected column.
- **Data**: select the numeric variable and a possible selection criterion.
- Reference value and Standard Deviation
- Use mean and SD of sample: calculate the mean and Standard Deviation of the data in the sample and use the these in the calculations for the z-scores.
- User-defined values: use pre-specified values for reference value and Standard Deviation.
- Reference value: the pre-specified reference value.
- Standard Deviation: the pre-specified Standard Deviation.

Click OK to proceed. The selected column in the spreadsheet is filled with the z-scores of the selected data.

The z-score of every observation x is calculated as (x - reference value) / SD

Power transformation

Allows to create a new variable containing a power transformation of a numeric variable. The transformation is defined by a power parameter λ (Lambda):

$x(\lambda) = x^{\lambda}$ when $\lambda \neq 0$

$x(\lambda) = \ln(x)$ when $\lambda = 0$

Optionally, you can select the Box-Cox transformation. The Box-Cox power transformation is defined as:

$x(\lambda) = (x^{\lambda} - 1) / \lambda$ when $\lambda \neq 0$

$x(\lambda) = \ln(x)$ when $\lambda = 0$

When some of the data are negative, a shift parameter *c* needs to be added to all observations (in the formulae above x is replaced with x+*c*).

Required input

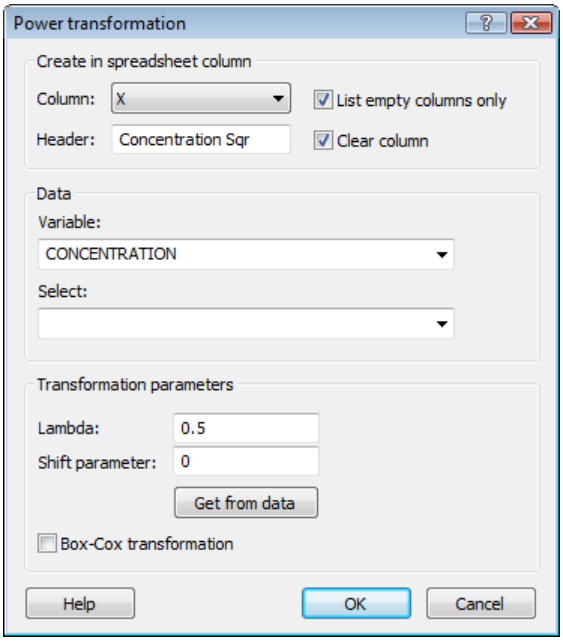

- **Column**: the column in which you want to place the transformed variable.
- **Options**
  - **List empty columns only**: if this option is selected, only empty columns are listed in the column selection box.
  - **Clear column**: the selected column will be cleared prior to generating and storing the transformed data.
- **Header**: the header (top cell) for the selected column.
- **Data**: select the numeric variable and a possible selection criterion.
- **Transformation parameters**
  - **Lambda**: the power parameter λ
  - **Shift parameter**: the shift parameter is a constant *c* that needs to be added to the data when some of the data are negative.
  - Button **Get from data**: click this button to estimate the optimal value for Lambda, and suggest a value for the shift parameter *c* when some of the observations are negative. The program will suggest a value for Lambda with 2 to 3 significant digits. It may be advantageous to manually round this value to values such as -3, -2, -1, -0.5, 0, 0.5, 1, 2 and 3 (see below).
  - Option **Box-Cox transformation**: select this option to use the Box-Cox power transformation as described above.

Click OK to proceed. The selected column in the spreadsheet is filled with the power-transformed data.

Interpretation of the power transformation

When you do not select Box-Cox transformation and the shift parameter *c* is zero then the power transformation is easy to interpret for certain values of lambda, for example:

λ = 0	logarithmic transformation
λ = 0.5	square root transformation
λ = -1	inverse transformation
λ = 1	**no** transformation!

Edit variables list

MedCalc automatically keeps track of variables that you add in the spreadsheet, while you enter data, or when you import Lotus or text files. When you save the data, the variables list is included in the file.

However, from time to time an error may occur in the list (e.g. double occurrence of a variable, or non-existing or deleted variables), and with the command *Edit variables list* you can edit the list.

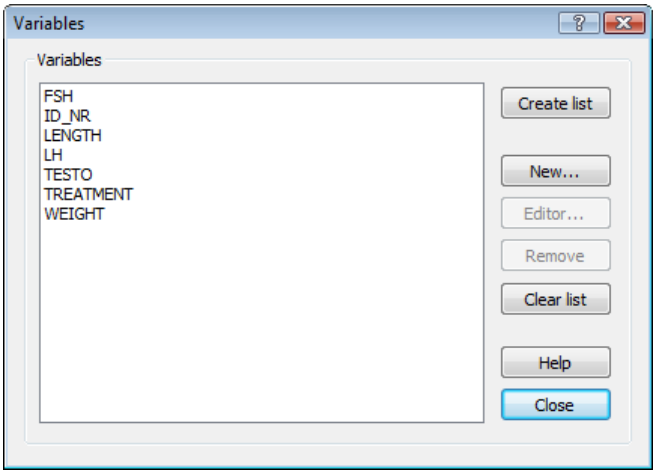

When you click the CREATE LIST button, a new list will be created based on the column headers in the spreadsheet.

Select NEW when you want to create a new derived variable in the *Variable editor* dialog box. In this dialog box, you can enter the formula for the derived variable, and/or you can enter a descriptive label for the variable.

When a variable is selected in the left box, you can click the EDITOR button to call the *Variable editor* dialog box.

You can select one or more variables in the list and next click the REMOVE button to remove the selected variables from the list.

When you click the CLEAR LIST button, the entire list will be cleared.

Variable editor

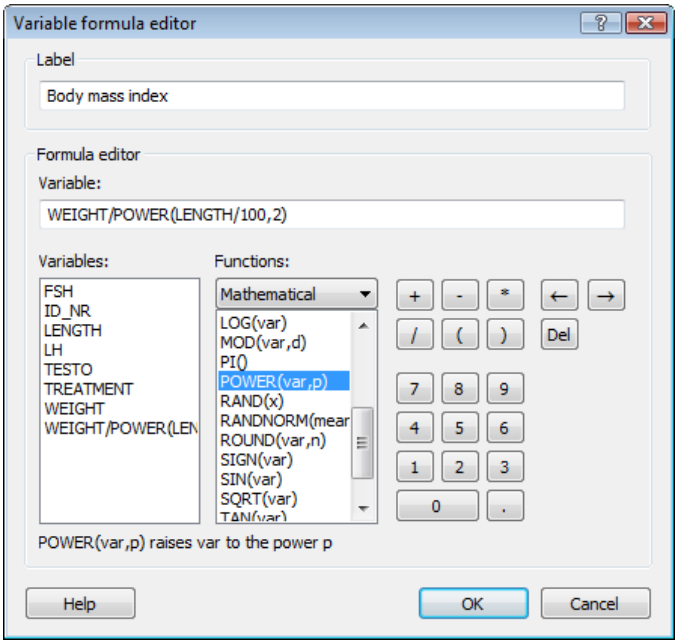

In the top box of the dialog box you can enter a descriptive label for the variable.

Next the variable is displayed, or the formula in case of a derived variable. E.g. if you have a variable named LENGTH containing the length of patients in meters, and you have another variable WEIGHT containing the weight of patients in kilograms, then you can complete the dialog box as follows to obtain a new variable containing the Body Mass Index.

Label: Body Mass Index

Formula: WEIGHT/POWER(LENGTH,2)

When entering the formula, you can select variables and functions from the respective lists.

Edit selection criteria list

After you have entered a selection criterion in a dialog box (in order to include only a selected subgroup of cases in the statistical analysis), this selection criterion will be 'remembered' by the program and put in the selection criteria list.

When you want to edit this list, select the *Edit selection criteria list*, and you obtain the following dialog box:

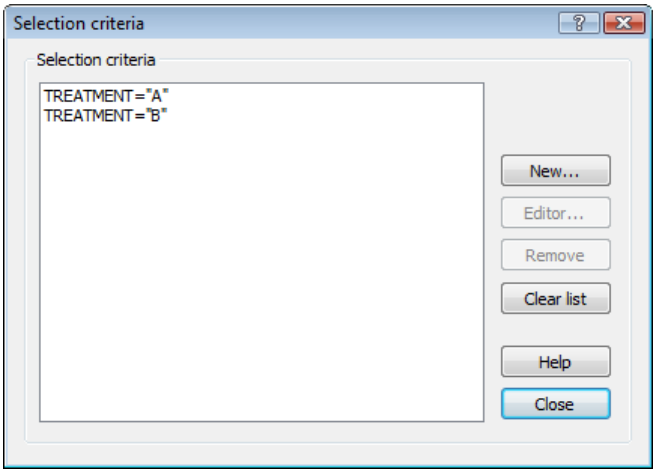

At the left side of the dialog box, the selection criteria list is displayed.

Select NEW when you want to create a new Selection criterion in the *Selection criterion editor* dialog box. In this dialog box, you can enter the formula for the new selection criterion, and/or you can enter a descriptive label for it. This editor box is similar to the *Variable editor* dialog box described on p. 53.

When a selection criterion is selected in the left box, you can click the EDITOR button to call the *Selection criterion editor* dialog box.

You can select one or more selection criteria in the list and next click the REMOVE button to remove the selected items from the list.

When you click the CLEAR LIST button, the entire list will be cleared.

Select variable for case identification

In this dialog box you can select a variable that can be used by the program to identify cases, e.g. when you click on a marker in a graph.

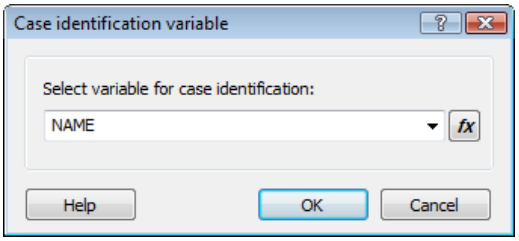

Note that you can enter a combination of variables in the dialog box, using spreadsheet functions. E.g. if the spreadsheet contains a column with header "NAME" and a column with header "FIRST_NAME", you can enter the following formula for case identification:

```
CONCAT(FIRST_NAME," ",NAME)
```

The *Formula editor* (see Variable editor on p. 53) is a useful tool to compose such a formula (click the fx button).

If no variable for case identification is selected, the program will use the spreadsheet row numbers to identify cases.

Enter key moves cell pointer…

With these commands you control the movement of the cell pointer in the spreadsheet:

- **Down:** After you press the Enter key, the cell pointer will move to the next row in the spreadsheet.
- **Right:** After you press the Enter key, the mouse will move to the next column in the spreadsheet.

Options

In the *Options* dialog box you can select several options such as the maximum number of rows and columns in the spreadsheet, default file locations, etc.

General

Recently used file list: number of entries: the number of recently used files displayed at the bottom of the File menu, so you can open the files quickly. Enter the number of files you want to display in the entries box.

Enable auto-update: select this option to enable the auto-update feature.

Check for updates every *n* days: the update frequency. If you set at 3 days, the program will check for updates every 3 days. If you set at 0 days, MedCalc will check for updates every time it is launched.

Recently used file list - number of entries: the number of recently used files displayed at the bottom of the File menu, so you can open the files quickly. Enter the number of files you want to display in the entries box.

Regional settings

Use Windows settings: select this option if you want to use the Windows settings (see Regional settings support).

Use the following alternative settings: select this option if you want to use your own settings:

Decimal symbol: enter the character to be used as decimal symbol. In the number 3.14 the character '.' is used as the decimal character; in Germany a comma is used as decimal character and Pi is written as 3,14.

List separator: enter the character to be used as the list separator. In the formula Power(2,8) the character ',' is used as the list separator. In Germany a ';' is used as list separator and the formula is written as Power(2;8).

Spreadsheet

Spreadsheet dimensions: enter the maximum number of rows that can be used in the spreadsheet. You will need to restart MedCalc before these new settings will take effect.

The default number of spreadsheet rows (which contain the cases) is 30000; the maximum number of rows is 100000.

However, to avoid memory problems (on some systems) it is advised not to set the number or rows higher than the number of cases you actually need.

The number of columns available in the MedCalc spreadsheet is 16384.

When you load a data file with more cases (rows) or variables than the number of rows and columns available in the spreadsheet, the program will display an appropriate warning.

Angle unit: the angle unit to be used in the program can be selected: radians, degrees or grades. The relation between the 3 units is as follows:

2 x Pi radians = 360 degrees = 400 grades.

Open spreadsheet window after loading/importing file: option to automatically open the spreadsheet window after a data file is read, or after data have been imported.

Enable AutoComplete for cell values: when this option is selected, MedCalc completes text entries you start to type in a column of data. If the first few letters you type match an existing entry in that column, the program fills in the remaining text for you.

Show Formula bar: select this option to show the formula bar in the spreadsheet window.

Spreadsheet colors

Numbers and text: the text color for numeric and non-numeric entries in the spreadsheet.

Default values: the color for default values as defined in Format column (p. 40).

Excluded: the color for cases or values that are excluded from calculations and statistical analysis.

Grid color: the color of the gridlines in the spreadsheet.

File locations

This tab shows the default storage location for data files, graphs, text and other items you create or use in MedCalc. Click the item you want to change, and then click Modify to set a new default location.

Text results

Colors: Select colors used in statistical text results windows: normal text color, color used for highlighting results, color for warnings and error messages, color used for background shading and line/border color.

Display small P-values in Exponential Notation: select this option to represent small P-value using exponential notation, e.g. P=0.121E-6 (P=0.000000121). This option is useful in for example the field of genomics where massive multiple testing requires very small significance levels.

Graphs

Copy graph: the **Image aspect ratio** option allows to select the width/height aspect ratio of the copied image. With the **Landscape** option the height/width ratio is 2/3 and with the **Portrait** option the height/width ratio is 3/2. With the **Graph window aspect ratio** option, the height and width of the copied image corresponds to the aspect ratio of the graph window.

Color palette: all possible colors that can be used in a graph. Click a color rectangle to modify that entry in the MedCalc color table.

STATISTICS MENU

Summary statistics

After you have selected *Summary statistics* in the menu, the *Select variable* dialog box appears on the screen, and in this dialog box you can identify the variable to be analyzed in the statistical procedure. You can also enter a selection criterion in the *Select* field, in order to include only a selected subgroup of data in the statistical procedure, as described in the *Introduction* part of this manual.

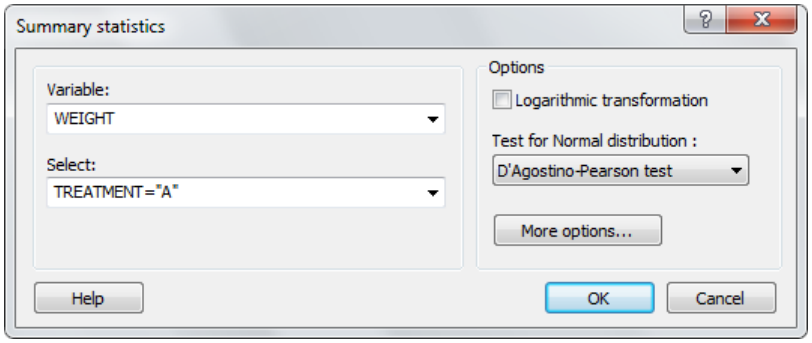

Options

- Logarithmic transformation
- Chi-square test, Kolmogorov-Smirnov test, or D'Agostino-Pearson test for Normal distribution.

  The **chi-square goodness-of-fit test** is applied to binned data (the data are put into classes) (Cochran & Snedecor, 1989) and requires a larger sample size than the **Kolmogorov-Smirnov test** (Neter et al., 1988), which is based on the greatest discrepancy between the sample cumulative distribution and the Normal cumulative distribution. The **D'Agostino-Pearson test** (Sheskin, 2004) computes a single P-value for the combination of the coefficients of Skewness and Kurtosis (see below).
- Click the **More options** button for additional options:

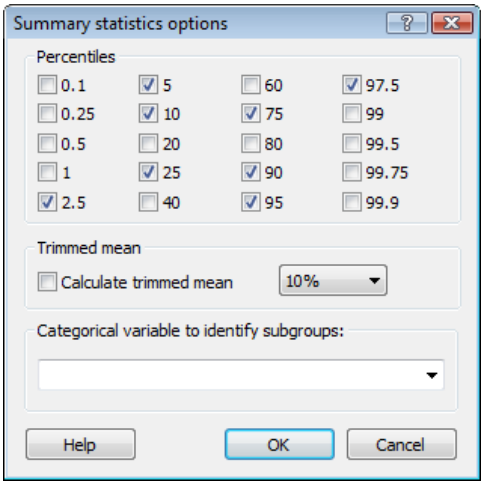

  - Percentiles: allows to select the percentiles of interest.
  - Trimmed mean: option to calculate a trimmed mean. You select the percentage of observations that will be trimmed away. For example, when you select 10% then the lowest 5% and highest 5% of observations will be dropped for the calculation of the trimmed mean.

- Subgroups: (optionally) select a categorical variable to break-up the data in several (max. 8) subgroups. Summary statistics will be given for all data and for all subgroups.

Results

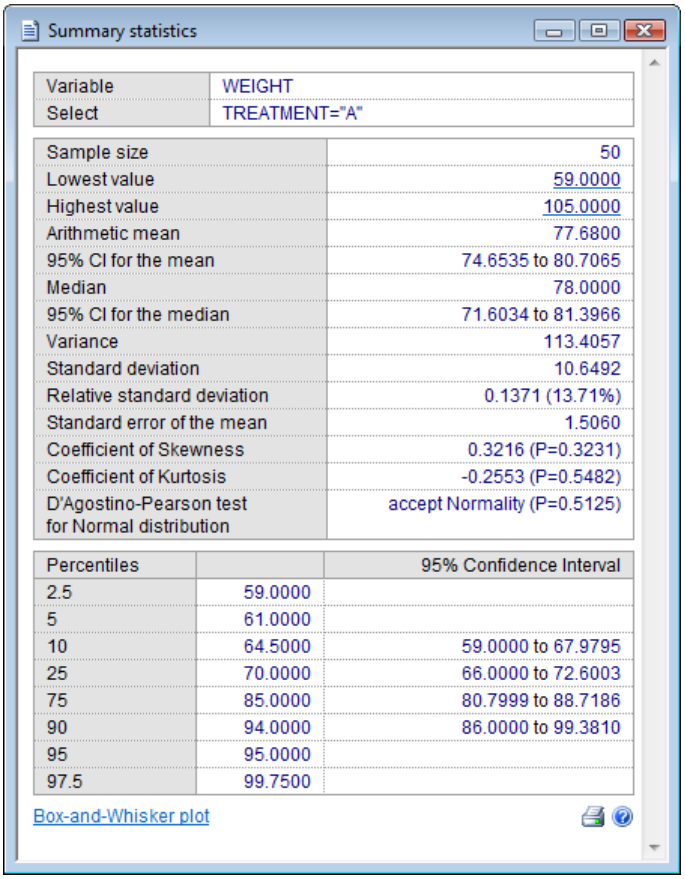

Summary statistics

Variable	WEIGHT
Select	TREATMENT="A"

Sample size	50
Lowest value	59.0000
Highest value	105.0000
Arithmetic mean	77.6800
95% CI for the mean	74.6535 to 80.7065
Median	78.0000
95% CI for the median	71.6034 to 81.3966
Variance	113.4057
Standard deviation	10.6492
Relative standard deviation	0.1371 (13.71%)
Standard error of the mean	1.5060
Coefficient of Skewness	0.3216 (P=0.3231)
Coefficient of Kurtosis	-0.2553 (P=0.5482)
D'Agostino-Pearson test for Normal distribution	accept Normality (P=0.5125)

Percentiles		95% Confidence Interval
2.5	59.0000	
5	61.0000	
10	64.5000	59.0000 to 67.9795
25	70.0000	66.0000 to 72.6003
75	85.0000	80.7999 to 88.7186
90	94.0000	86.0000 to 99.3810
95	95.0000	
97.5	99.7500	

Box-and-Whisker plot

Sample size: the number of cases ***n*** is the number of numeric entries for the variable that fulfill the selection criterion.

The **lowest value** and **highest value** of all observations (range).

Arithmetic mean: the arithmetic mean $\overline{X}$ is the sum of all observations divided by the number of observations *n*:

$$\overline{X} = \frac{\sum X}{n}$$

95% confidence interval (CI) for the mean: this is a range of values, calculated using the method described later (*see* Standard Error of the Mean), which contains the *population mean* with a 95% probability.

Median: when you have *n* observations, and these are sorted from smaller to larger, then the median is equal to the value with order number (n+1)/2. The median is equal to the 50th percentile. If the distribution of the data is Normal, then the median is equal to the arithmetic mean. The median is not sensitive to extreme values or outliers, and therefore it may be a better measure of central tendency than the arithmetic mean.

95% confidence interval (CI) for the median: this is a range of values that contains the population median with a 95% probability (Campbell & Gardner, 1988). This 95% confidence interval can only be calculated when the sample size is not too small.

Variance: the variance is the mean of the square of the differences of all values with the arithmetic mean. The variance (s^2) is calculated using the formula:

$$s^2 = \frac{\sum (X - \overline{X})^2}{n - 1}$$

Standard deviation: the standard deviation (*s* or *SD*) is the square root of the variance, and is a measure of the spread of the data:

$$s = \sqrt{\frac{\sum (X - \overline{X})^2}{n - 1}}$$

When the distribution of the observations is Normal, then 95% of all *observations* are located in the interval *mean - 1.96 SD* to *mean + 1.96 SD* (for other values see table p. 226).

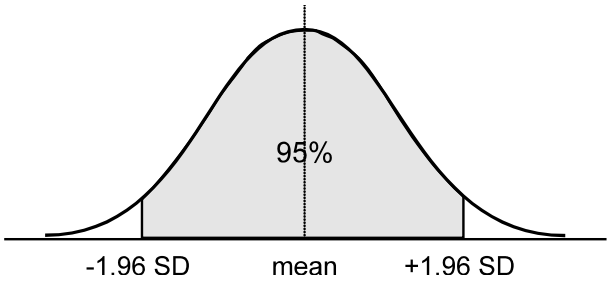

This interval should not be confused with the smaller 95% confidence interval for the mean. The interval *mean - 1.96 SD* to *mean + 1.96 SD* represents a descriptive 95% confidence range for the individual observations, whereas the 95% CI for the mean represents a statistical uncertainty of the arithmetic mean.

Relative standard deviation (RSD): this is the standard deviation divided by the mean. If appropriate, this number can be expressed as a percentage by multiplying it by 100 to obtain the *coefficient of variation*.

Standard error of the mean (SEM): is calculated by dividing the standard deviation by the square root of the sample size.

$$S.E.M. = \frac{s}{\sqrt{n}}$$

The SEM is used to calculate confidence intervals for the mean. When the distribution of the observations is Normal, or approximately Normal, then there is 95% confidence that the population mean is located in the interval $\overline{X}$ ± *t* SEM, with *t* taken from the *t*-distribution with *n-1* degrees of freedom and a *confidence* of 95% (table on p. 227). For large sample sizes, *t* is close to 1.96.

Skewness

The coefficient of Skewness is a measure for the degree of symmetry in the variable distribution. If the corresponding P-value is low (P<0.05) then the variable symmetry is significantly different from that of a Normal distribution, which has a coefficient of Skewness equal to 0 (Sheskin, 2004).

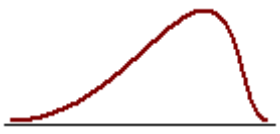

Negatively skewed distribution
or Skewed to the left
Skewness < 0

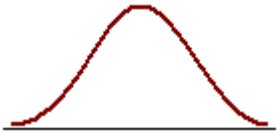

Normal distribution
Symmetrical
Skewness = 0

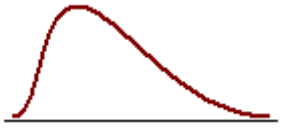

Positively skewed distribution
or Skewed to the right
Skewness > 0

Kurtosis

The coefficient of Kurtosis is a measure for the degree of peakedness/flatness in the variable distribution. If the corresponding P-value is low (P<0.05) then the variable peakedness is significantly different from that of a Normal distribution, which has a coefficient of Kurtosis equal to 0 (Sheskin, 2004).

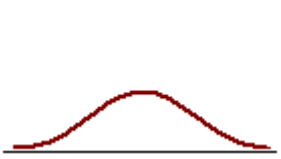

Platykurtic distribution
Low degree of peakedness
Kurtosis <0

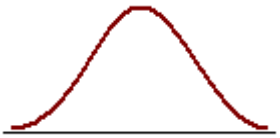

Normal distribution
Mesokurtic distribution
Kurtosis = 0

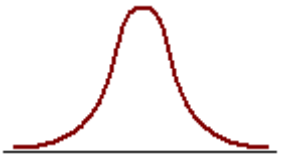

Leptokurtic distribution
High degree of peakedness
Kurtosis >0

Test for Normal distribution

The result of the test for Normal distribution is expressed as '*accept Normality*' or '*reject Normality*', with P value. If P is higher than 0.05, it may be assumed that the data have a Normal distribution and the conclusion '*accept Normality*' is displayed.

If the P value is less than 0.05, then the hypothesis that the distribution of the observations in the sample is Normal, should be rejected, and the conclusion '*reject Normality*' is displayed. In the latter case, the sample cannot accurately be described by arithmetic mean and standard deviation, and such samples should not be submitted to any parametrical statistical test or procedure, such as e.g. a t-test. To test the possible difference between not Normally distributed samples, the *Wilcoxon test* can be used, and correlation can be estimated by means of *rank correlation*.

When the sample size is small, it may not be possible to perform the selected test and an appropriate message will appear.

Percentiles (or "centiles"): when you have *n* observations, and these are sorted from smaller to larger, then the *p*-th percentile is equal to the observation with rank number:

$$R(p) = 0.5 + \frac{p \times n}{100} \qquad \text{(Lentner, 1982)}$$

When the rank number R(*p*) is a whole number, then the percentile coincides with the sample value; if R(*p*) is a fraction, then the percentile lies between the values with ranks adjacent to R(*p*) and in this case MedCalc uses interpolation to calculate the percentile.

The formula for R(*p*) is only valid when

$$\frac{1}{n} \le \frac{p}{100} \le \frac{n-1}{n}$$

E.g. the 5th and 95th percentiles can only be estimated when $n \ge 20$, since

$$\frac{1}{20} \le \frac{5}{100} \qquad \text{and} \qquad \frac{95}{100} \le \frac{20-1}{20}$$

Therefore it makes no sense to quote the 5th and 95th percentiles when the sample size is less than 20. In this case it is advised to quote the 10th and 90th percentiles, at least if the sample size is not less than 10.

The percentiles can be interpreted as follows: *p* % of the observations lie below the *p*-th percentile, e.g. 10% of the observations lie below the 10th percentile.

The 25th percentile is called the 1st quartile, the 50th percentile is the 2nd quartile (and equals the Median), and the 75th percentile is the 3rd quartile.

The numerical difference between the 25th and 75 percentile is the *interquartile range*. Within the 2.5th and 97.5th percentiles lie 95% of the values and this range is called the 95% *central range*. The 90% central range is defined by the 5th and 95th percentiles, and the 10th and 90th percentiles define the 80% central range.

Log transformation

If the option Log transformation was selected, the program will display the back-transformed results. The back-transformed mean is named the *Geometric mean*. Variance, Standard deviation and Standard error of the mean cannot be back-transformed meaningfully and are not reported.

Presentation of results

The description of the data in a publication will include the sample size and arithmetic mean. The standard deviation can be given as an indicator of the variability of the data: *the mean was 25.6 mm (SD 3.2 mm)*. The standard error of the mean can be given to show the precision of the mean: *the mean was 25.6 mm (SE 1.6 mm)*.

When you want to make an inference about the population mean, you can give the mean and the 95% confidence interval of the mean: *the mean was 25.6 (95% CI 22.4 to 28.8)*.

If the distribution of the variable is positively skewed, then a mathematical transformation of the data may be applied to obtain a Normal distribution, e.g. a logarithmic or square root transformation. After calculations you can convert the results back to the original scale. It is then useless to report the back-transformed standard deviation or standard error of the mean. Instead, you can antilog the confidence interval in case a logarithmic transformation was applied, or square the confidence interval if you have applied a square root transformation (Altman et al., 1983). The resulting confidence interval will then not be symmetrical, reflecting the shape of the distribution. If, for example, after logarithmic transformation of the data, the mean is 1.408 and the 95% confidence interval is 1.334 to 1.482, then you will antilog these statistics and report: *the mean was 25.6 mm (95% CI 21.6 to 30.3).*

If the distribution of the variable is not normal even after logarithmic or other transformation, then it is better to report the median and a percentiles range, e.g. the interquartile range, or the 90% or 95% central range: *the median was 25.6 mm (95% central range 19.6 to 33.5 mm)*. The sample size will be taken into consideration when you decide whether to use the interquartile range or the 90% or 95% central range (see p. 60) (Altman, 1980).

The precision of the reported statistics should correspond to the precision of the original data. The mean and 95% CI can be given to one decimal place more than the raw data, the standard deviation and standard error can be given with one extra decimal (Altman et al., 1983).

Finally, the summary statistics in the text or table may be complemented by a graph (see p. 70).

Outlier detection

Outlier detection is used to detect anomalous observations in sample data.

Required input

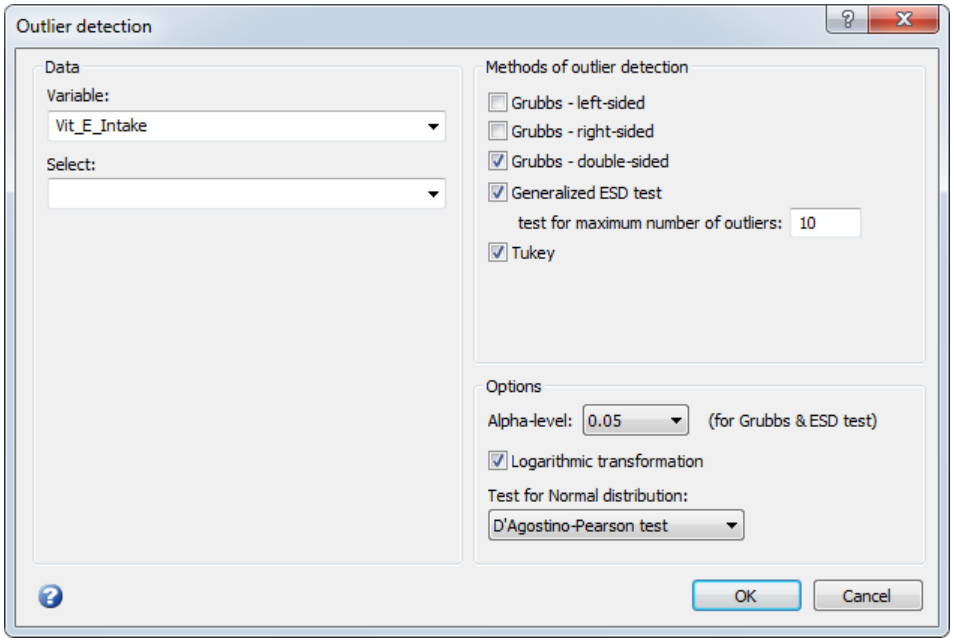

Data

- **Variable**: the name of the variable containing the data to be analyzed.
- **Select**: (optionally) a selection criterion in order to include only a selected subgroup of cases in the statistical analysis.

Methods of outlier detection:

- Grubbs - left-sided: check only the smallest value(*) (Grubbs, 1969).
- Grubbs - right-sided: check only the largest value(*) (Grubbs, 1969).
- Grubbs - double-sided: check the most extreme value at either side(*) (Grubbs, 1969).
- Generalized ESD test: the Generalized Extreme Studentized Deviate (ESD) procedure can detect multiple outliers in one step (Rosner, 1983).

  test for maximum number of outliers: enter the maximum number of outliers to detect.
- Tukey: check for multiple outliers at either side, categorized as 'outside' or 'far out' values (Tukey, 1977).

  An outside value is defined as a value that is smaller than the lower quartile minus 1.5 times the interquartile range, or larger than the upper quartile plus 1.5 times the interquartile range (the 'inner fences').

  A far out value is defined as a value that is smaller than the lower quartile minus 3 times the interquartile range, or larger than the upper quartile plus 3 times the interquartile range (the 'outer fences').

(*) The single-sided Grubbs' tests are more sensitive than the double-sided test.

Options

- Alpha level for Grubbs' and ESD test: select the alpha-level (ranging from 0.10 to 0.001), applicable only in Grubbs's test and the Generalized ESD test. With a bigger alpha-level the test will be more sensitive and outliers will more rapidly be detected; however, this may result in false-positive results.
- Logarithmic transformation: the outlier detection methods assume that the data follow an approximately normal distribution (see next option). Sometimes data should be logarithmically transformed before analysis.

  The example uses the data from Rosner (1983) on their original scale. Therefore Logarithmic transformation is performed like in the Rosner paper.
- Test for Normal distribution: see Tests for Normal distribution.

Results

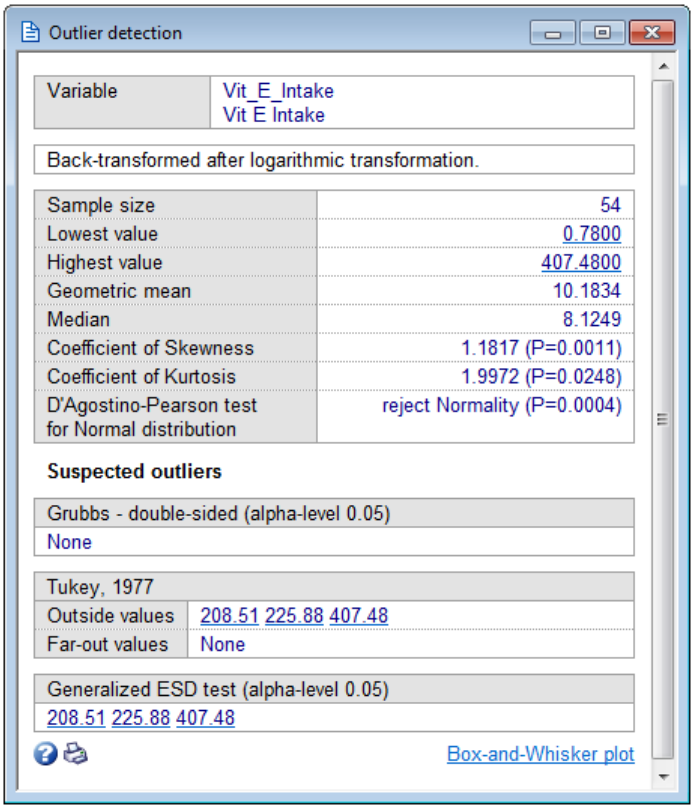

Summary statistics

- Summary statistics for the selected data are displayed (see p. 57).
- If the test for Normal distribution reports 'reject Normality' the outlier detection methods may be invalid since they assume that the data follow an approximately normal distribution. Perhaps data should have been logarithmically transformed before analysis.

  In the example, data are logarithmically transformed.

Suspected outliers

The program lists the outliers identified by the different procedures.

Grubbs' test can only be used to detect one single outlier; if you suspect there is more than one outlier you should not repeat the procedure but use the Generalized ESD test.

What to do when you have identified an outlier

Do not remove outliers automatically.

- Remove outliers only when a cause can be found for the spurious result, such as a pre-, post-, or analytical error.
- When you conclude that a pre-, post-, or analytical error is the cause of the spurious result, be aware that the same errors may exist in the other data values.
- Check the distribution of the data. Logarithmically transformed sample data may more closely follow a Normal distribution. Graph the data with and without logarithmic transformation, for example using a Box-and-Whisker plot.
- You may consider to replace the outlier value with the next highest/lowest (non-outlier) number.
- Keep the outlier but use robust or nonparametric statistical methods that do not assume that data are Normally distributed.
- Do the statistical analysis and report conclusions both with and without the suspected outlier.

In all cases, report the outliers and how you have dealt with them.

Literature

- Grubbs FE (1969) Procedures for detecting outlying observations in samples. Technometrics 11:1-21.
- Rosner B (1983) Percentage points for a generalized ESD many-outlier procedure. Technometrics 25:165-172.
- Tukey JW (1977) Exploratory data analysis. Reading, Mass: Addison-Wesley Publishing Company.

Histogram

After selecting *Histogram*, a similar dialog box is displayed as for *Summary statistics*. Enter the name of a variable and optionally a selection criterion. If you have previously entered this variable and selection criterion in the box for summary statistics, then this new variable will be selectable in the Variable list (click the button).

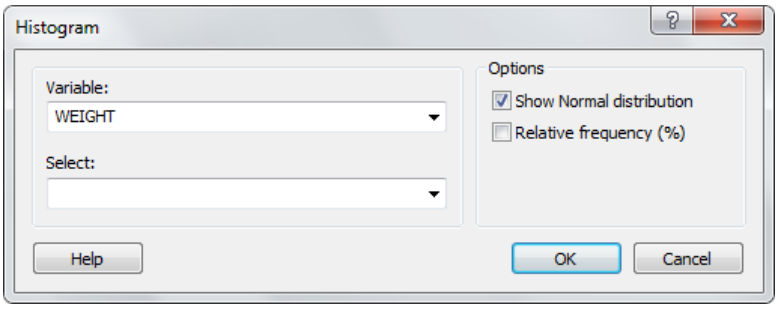

Options

- Show Normal distribution: option to have a Normal distribution curve (with Mean and Standard Deviation of the data represented in the histogram) superimposed over the histogram.
- Relative frequency (%): option to express frequencies as percentages.

After a moment (the program first collects the data and performs some calculations) the following dialog box is displayed on the screen:

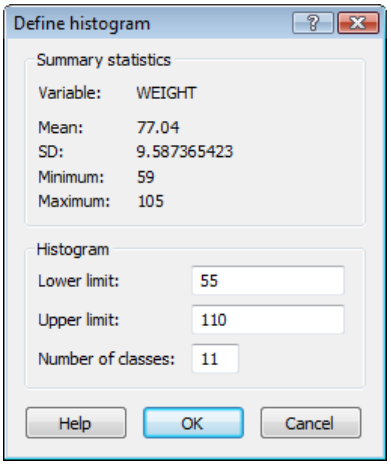

This dialog box displays:

- the mean, standard deviation, minimum and maximum value for the selected variable
- the default lower and upper limits, and the default number of classes in the histogram are displayed.

If you prefer other values than these default values, you can make the necessary changes. For *Lower* and *Upper limit*, the program will not accept values greater or less than the minimum and maximum of the variable. When you click the OK button, the program will continue with the new settings, but when you click CANCEL, the program will display the histogram with the initial default settings.

This is the histogram for the variable Weight:

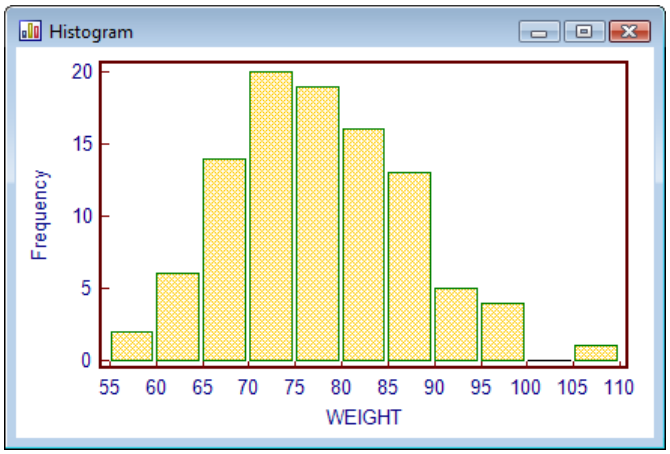

The first bar in this histogram represents the number of cases (frequency) with weight ≥ 55 and < 60. The second bar represents the number of cases with weight ≥ 60 and < 65, etc.

In the *Histogram* dialog box, you have the option (in fact it is the default setting) to superimpose a *Normal distribution* plot on the histogram.

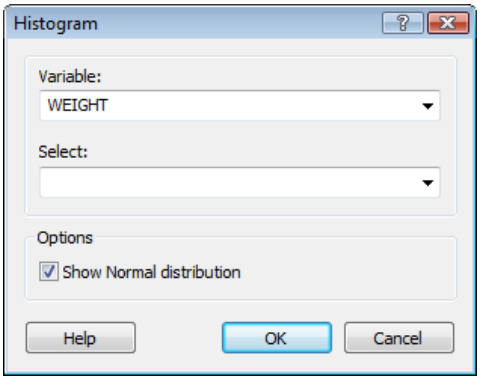

When this option is selected, a Normal distribution plot (with Mean and Standard Deviation of the data represented in the histogram) is superimposed over the histogram.

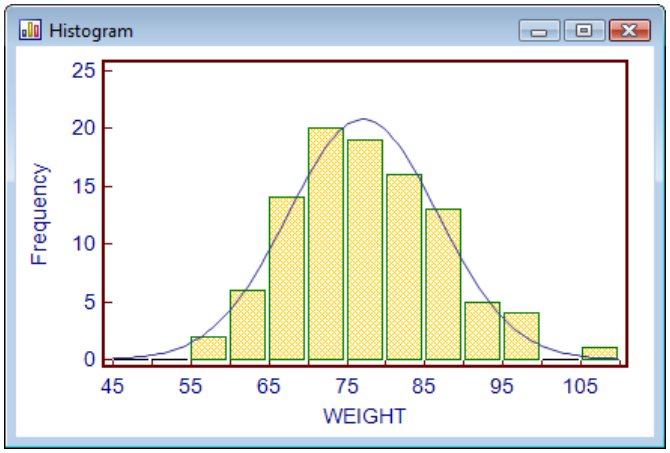

Using the histogram it can be evaluated visually whether the data are distributed symmetrically, Normally or Gaussian or whether the distribution is asymmetrical or skewed.

When the distribution is not Normal, it cannot accurately be described by mean and standard deviation, but instead the median, mode, quartiles and percentiles should be used. The latter statistics are reported in the *Summary statistics* window.

To change the titles, colors or axis scaling used in the graph, refer to page 12.

By selecting the *Export* command in the *File* menu you can export the displayed graph as a picture file (see p. 26). When you want to print the graph, press *Ctrl+P*.

Cumulative frequency distribution

To obtain a cumulative frequency distribution, select *Cumulative frequency distribution* in the menu and proceed as for *Histogram*.

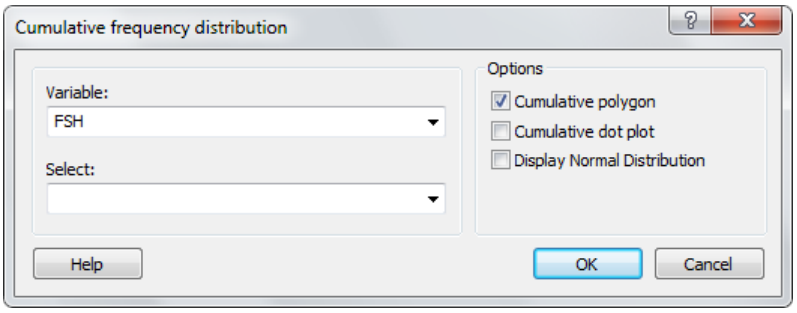

In the dialog box for *Cumulative frequency distribution*, you can either select a *Cumulative frequency polygon* or *Cumulative dot plot*, or both.

Finally you have the option to superimpose a *Normal distribution* curve on the diagram. The mean and standard deviation of this Normal distribution curve are those of the variable represented in the graph. When this option is selected the cumulative frequency distribution is plotted using a different algorithm, allowing comparison of the observed frequency distribution with the theoretical Normal distribution.

When you have completed the dialog box, click OK, or press the *Enter* key to obtain the graph.

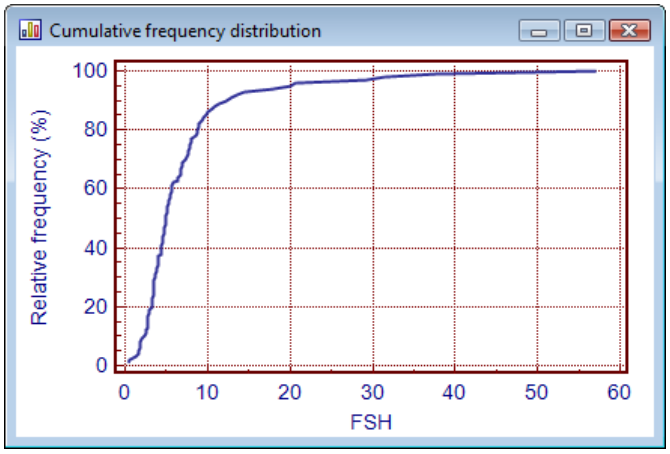

In the next figures, some combinations of the possible selections are displayed:

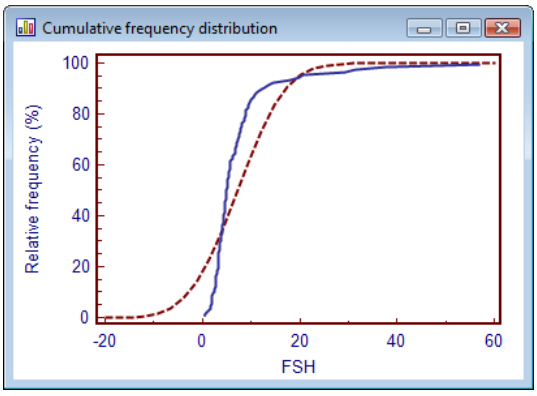

Polygon and Normal distribution

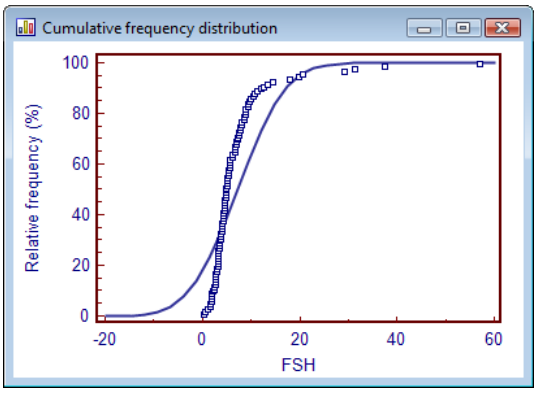

Dots and Normal distribution

Normal plot

The *Normal plot* is a graphical tool to judge the Normality of the distribution of the sample data.

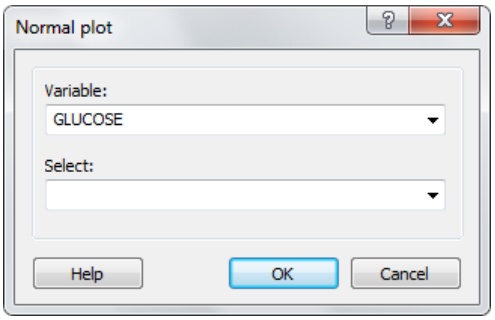

Select or enter the variable's name in the variable input field. Optionally, you may enter a selection criterion in order to include only a selected subgroup of cases in the graph. When you have completed the dialog box, click the *OK* button to proceed.

The horizontal axis of the Normal plot shows the numerical values of the observations, and the vertical axis gives the relative frequency in terms of the number of standard deviations from the mean.

A straight reference line represents the Normal distribution. If the sample data are near a Normal distribution, the data points will be near this straight line.

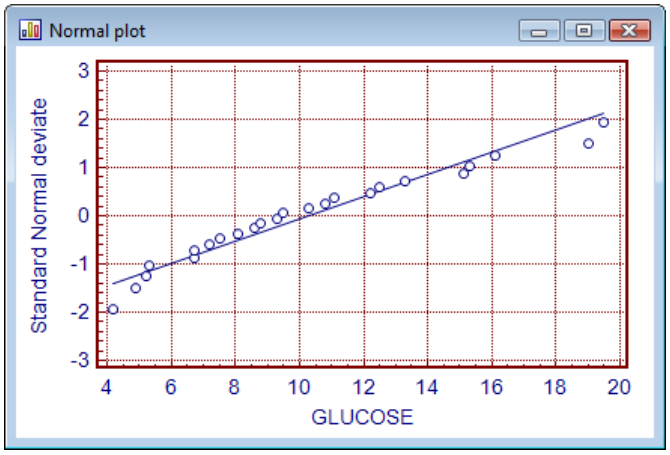

Probably a Normal distribution

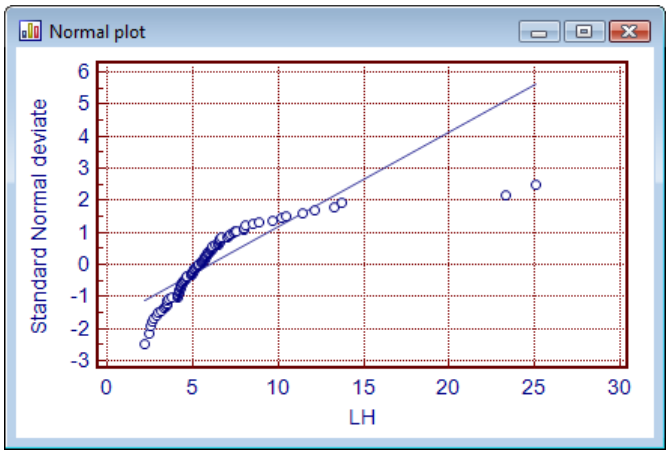

Probably not a Normal distribution

Dot plot

Description

Creates a dot plot for a single variable with different graph options such as the inclusion of a Bar, Line or Marker for mean or median, with choice of different error bars for mean (95% CI, 1 SEM, 1 SD, 2 SD, 3 SD, range) or median (95% CI, 25-75 percentiles, 10-90 percentiles, 5-95 percentiles, 2.5-97.5 percentiles, 1-99 percentiles, range), and/or Box-and-whisker plot (Tukey, 1977).

Required input

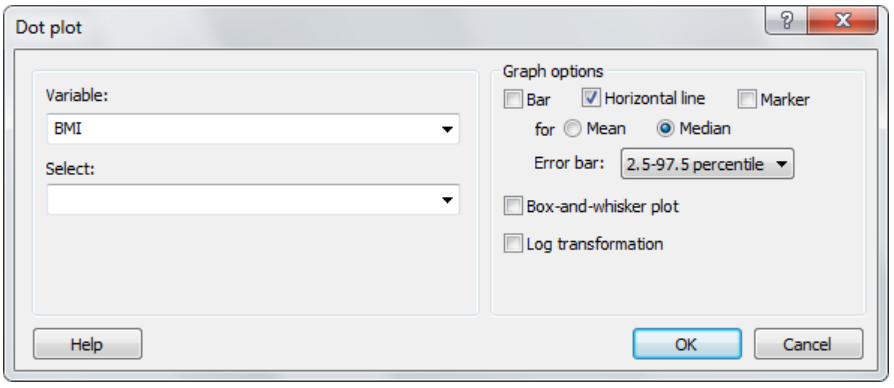

Select the variable of interest, and optionally a selection criterion to include only particular cases in the graph.

Several elements can be selected to add onto the dot plot, and some of these can be combined:

- Bar, Horizontal Line and/or Marker for mean or median

  The following error bars are available if Bars, Horizontal Line and/or Markers is selected:

  - If **mean** is selected: (none), or 95% CI for the mean, 1 SD, 2 SD, 3 SD, 1 SEM, range

    Note that 2 SEM is not in this list: when the number of cases is large, mean ± 2 SEM corresponds to the 95% confidence interval (CI) for the mean. When the number of cases is small, then the 95% CI interval is calculated as mean ± t * SEM, where t is taken from a t-table (with DF=n-1 and area A=95%).

    Although 1 SEM gives the more narrow error bar, this option is not recommended since the resulting error bar may be highly misleading, especially when the number of cases in the groups is different. Preferably the 95% CI for the mean is used for providing a valid graphical comparison of means (Pocock, 1984), or use 2 SD as an indication for the variability of the data.

- If **median** is selected: (none), or 95% CI for the median, 25-75 percentile, 10-90 percentile, 5-95 percentile, 2.5-97.5 percentiles, 1-99 percentile, range

  When the number of cases is small, it is possible that the 95% CI for the median is not defined and that it will not be displayed in the graph.

  When you use percentile ranges, take into account the number of observations: you need at least 100 observations for 1-99 percentiles, at least 20 for 5-95 percentiles, at least 10 for 10-90 percentile and at least 4 for 25-75th percentiles.

- **Box-and-Whisker plot**: see below for details.
- **Option**: logarithmic transformation of the data.

Example

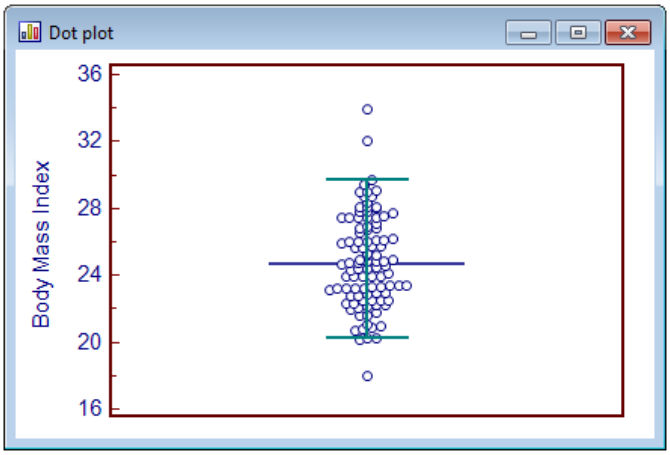

When you click an individual observation in the graph, the corresponding case is identified in a popup window (see also *Select variable for case identification* command, p. 54). If you double-click an observation, the spreadsheet window will open with the corresponding case highlighted.

Select the **Info** command in the shortcut menu that appears after right-clicking in the graph window to get detailed information on the data represented in the graph (sample size, etc.).

Box-and-whisker plot

The box-and-whisker plot (Tukey, 1977) displays a statistical summary of a variable: median, quartiles, range and, possibly, extreme values.

The *Define variable* dialog box for *Box-and-whisker plot* is similar to the one for *Summary statistics*: the name of the variable has to be entered as described on page 8:

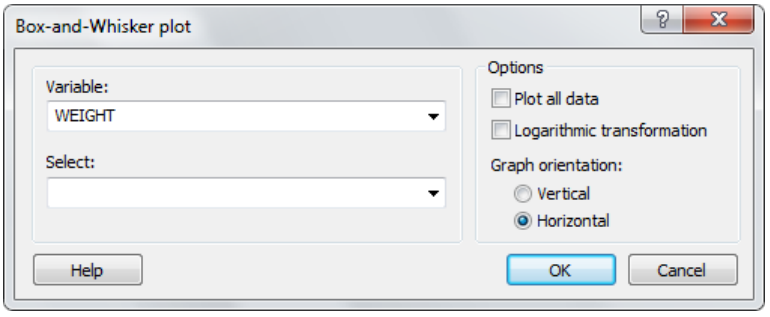

If you require a logarithmic transformation of the data (e.g. to Normalize the data), then click the LOGARITHMIC TRANSFORMATION option.

You can choose between vertical and horizontal orientation of the box-and-whisker plot.

This is the box-and-whisker plot for the variable *Weight*:

In the Box-and-whisker plot, the central box represents the values from the lower to upper quartile (25 to 75 percentile). The middle line represents the median. The horizontal line extends from the minimum to the maximum value, excluding *outside* and *far out* values, which are displayed as separate points.

- An *outside value* is defined as a value that is smaller than the lower quartile minus 1.5 times the interquartile range, or larger than the upper quartile plus 1.5 times the interquartile range (*inner fences*).
- A *far out* value is defined as a value that is smaller than the lower quartile minus 3 times the interquartile range, or larger than the upper quartile plus 3 times the interquartile range (*outer fences*). These values are plotted using a different marker, drawn in the "warning" color.

As an option, you may select to plot all individual data points. This enables you to obtain a diagram representing a statistical summary of the data without the disadvantage of concealing the real data.

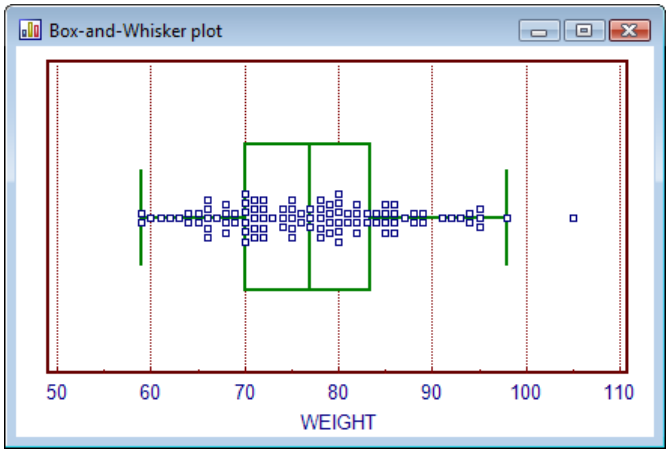

When you click an individual observation in the graph, the corresponding case is identified in a popup window (see also *Select variable for case identification* command, p. 54). If you double-click an observation, the spreadsheet window will open with the corresponding case highlighted. If the value is an outlier, you can exclude the value or the entire case from further statistical analysis by selecting the *Exclude* command in the *Tools* menu (see p. 42).

Presentation of results

The description of the data in the text or table may be complemented by a graphical representation of the data: a histogram, cumulative distribution or box-and-whisker plot. The *histogram* is not very effective to display location and spread. The *cumulative distribution* has the advantage that it makes it easy to estimate the median (or other percentile) by reading off the horizontal value at which the curve attains 50% (or other percentage) (Moses, 1987). Secondly, the plot can contain the individual observations (cumulative dot plot). Finally, the *box-and-whisker plot* may be preferable because it can combine a display of all the data together with a statistical summary.

Correlation

Correlation analysis is used to see if the values of two variables are associated. The two variables should be random samples, and should have a Normal distribution (possibly after transformation).

When you select *Correlation* in the menu, the following dialog box appears on the screen:

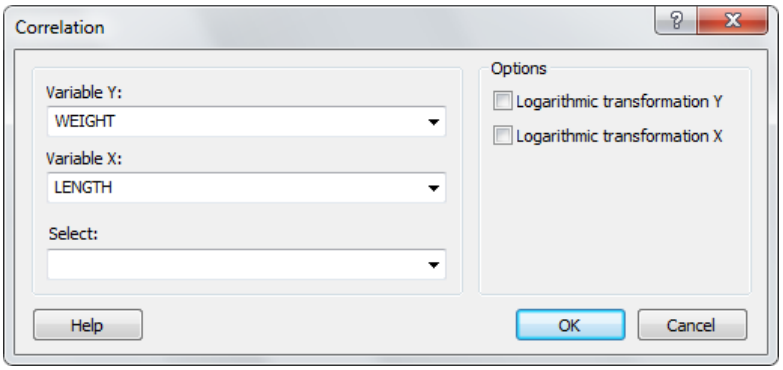

This dialog box has to be completed in a way similar to the box for summary statistics (see also p. 8), but now 2 variables must be selected. If you want to select the variables from the variables list, click the button, and select the variable in the list that is displayed. Next, you move the cursor to the *Variable X* field, and again click the button to select the variable in the list.

Finally, you can select a logarithmic transformation for one or both variable(s) in order to obtain Normal distributions.

After you click the OK button, or press the *Enter* key, you obtain the requested statistics in the *results* window:

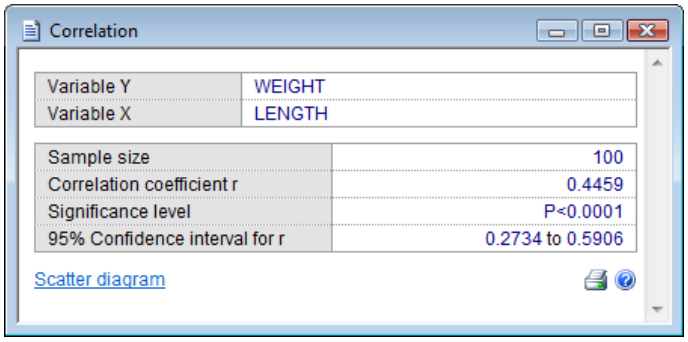

Results

Sample size: the number of data pairs **n**

Correlation coefficient with P-value: the correlation coefficient is a number between -1 and 1. In general, the correlation expresses the degree that, on an average, two variables change correspondingly.

If one variable increases when the second one increases, then there is a positive correlation. In this case the correlation coefficient will be closer to 1. For instance the height and age of children are positively correlated.

If one variable decreases when the other variable increases, then there is a negative correlation and the correlation coefficient will be closer to -1.

The P-value is the probability that you would have found the current result if the correlation coefficient were in fact zero (null hypothesis). If this probability is lower than the conventional 5% ($P<0.05$) the correlation coefficient is called statistically significant.

It is, however, important not to confuse correlation with causation. When two variables are correlated, there may or may not be a causative connection, and this connection may moreover be indirect. Correlation can only be interpreted in terms of causation if the variables under investigation provide a logical (biological) basis for such interpretation.

95% confidence interval (CI) for the correlation coefficient: this is the range of values that contains with a 95% confidence the 'true' correlation coefficient.

Presentation of results

The number of data pairs (sample size) should be reported, the correlation coefficient (two decimal places), together with the P-value and the 95% confidence interval: *the correlation coefficient was 0.45 (P<0.001, 95% CI 0.27 to 0.59).*

The relationship between two variables can be represented graphically by a scatter diagram.

Rank correlation

When the distribution of variables is not Normal, the degree of relationship between the variables can be calculated using *Rank correlation*. Instead of using the precise values of the variables, the data are ranked in order of size, and calculations are based on the differences between the ranks of corresponding values X and Y.

After selecting *Rank correlation* in the MedCalc menu, enter the names of the two variables in the following dialog box. For both variables, you can click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name. Optionally, you may also enter selection criteria in order to include only a selected subgroup of cases in the statistical analysis. Again, you can click the button to obtain a list of selection criteria already used for the current data.

MedCalc offers both Spearman's rank correlation coefficient *rho* and Kendall's *tau*.

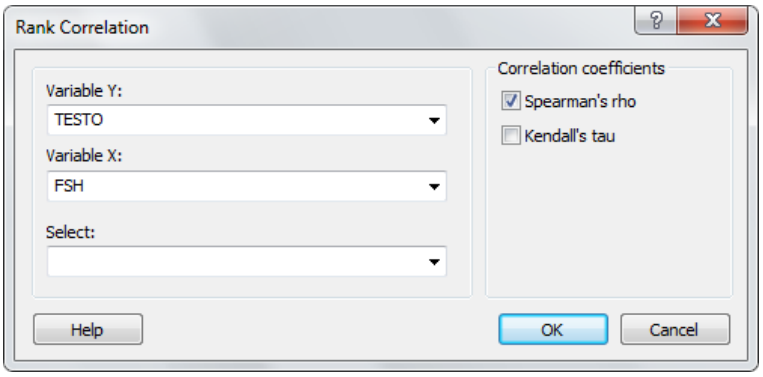

Next click the OK button, or press the *Enter* key to obtain the following statistics in the results window:

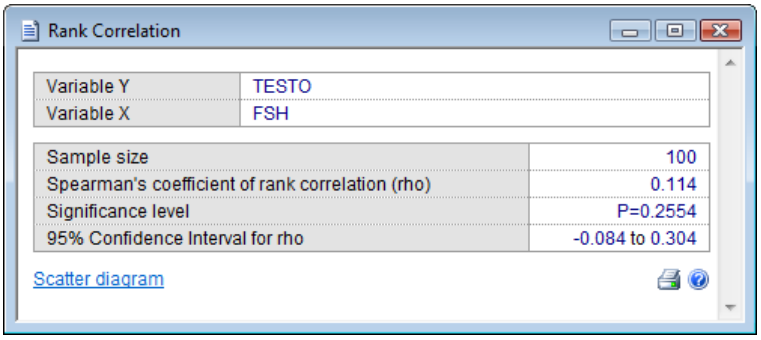

In this example the Spearman's coefficient of rank correlation *rho* is 0.114. The 95% confidence interval ranges from -0.084 to 0.304. The associated P-value is 0.255 and the conclusion therefore is that there is not a significant relationship between the two variables.

Scatter diagram

In a scatter diagram, the relation between two numerical variables is presented graphically. One variable (the variable X) defines the horizontal axis and the other (variable Y) defines the vertical axis. The values of the two variables on the same row in the data spreadsheet, give the points in the diagram.

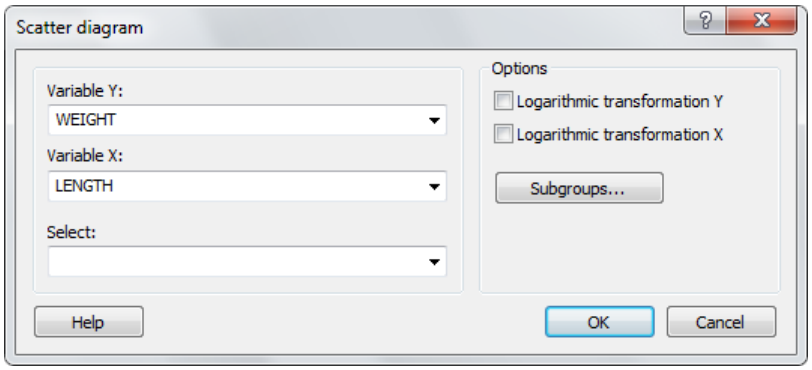

You can click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name. Optionally, you may also enter a selection criterion in order to include only a selected subgroup of cases in the graph.

You can select a logarithmic transformation for one or both variables (in this case the program will use a logarithmic scale for the corresponding axis in the graph).

Use the **Subgroups** button if you want to identify subgroups in the scatter diagram. A new dialog box is displayed in which you can select a categorical variable. The graph will display different markers for the different categories in this variable.

After you click the OK button, or press the *Enter* key you obtain the following graph:

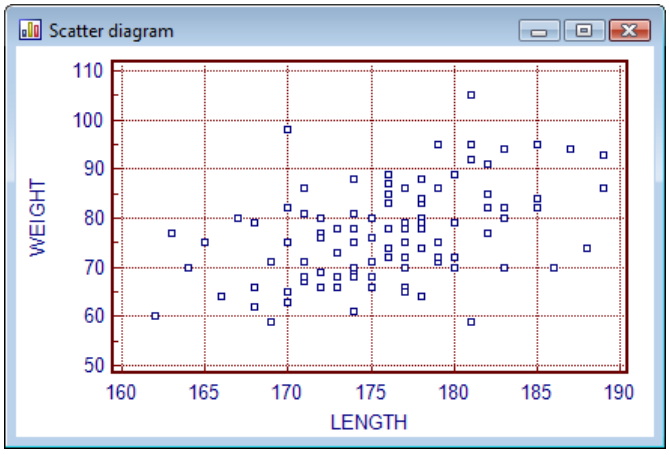

This is the same scatter diagram, but the categorical variable "Treatment" has been used to identify different **subgroups** in the graph.

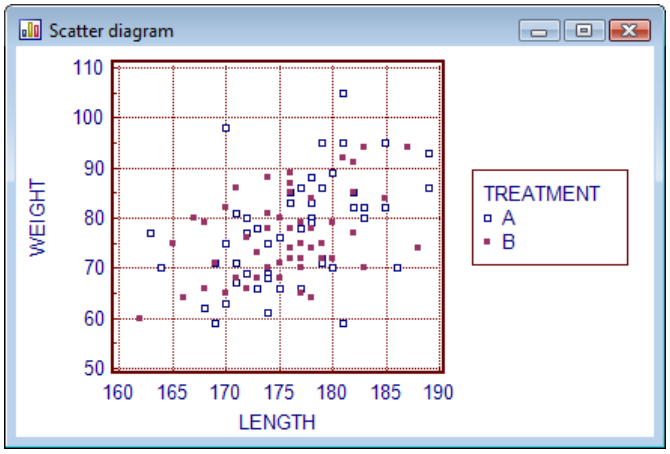

Regression

Regression is a method used to describe the relationship between two variables and to predict one variable from another (if you know one variable, then how well can you predict a second variable?).

Whereas for correlation (see p. 71) the two variables need to have a Normal distribution, in regression analysis only the dependent variable Y should have a Normal distribution. The variable X does not need to be a random sample with a Normal distribution (the values for X can be chosen by the experimenter). However, the variability of Y should be the same for each value of X.

When you select *Regression* in the menu, the following box appears on the screen:

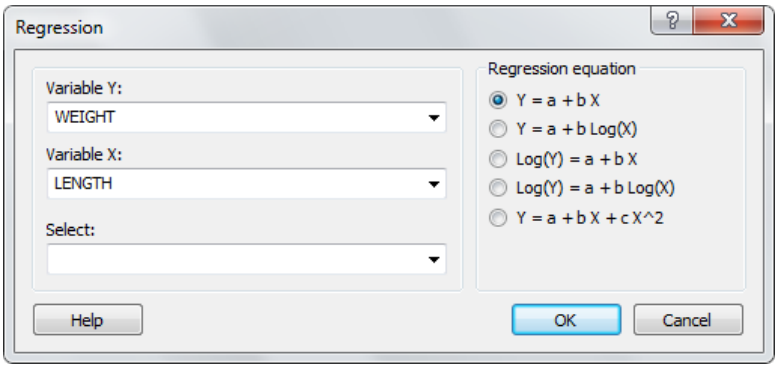

In this dialog box you identity 2 variables. If you want to select the variables from the variables list, click the button and now you can select the variable in the list. Next, you move the cursor to the *Independent X* field, and again you click the button to select the variable in the list.

Optionally, you may also enter selection criteria in order to include only a selected subgroup of cases in the statistical analysis. Again, you can click the button to obtain a list of selection criteria already used for the current data.

Finally, a regression equation (regression model, equation of approximating curve) has to be selected. The program offers a choice of 5 different equations:

Y	**=**	***a***	**+**	***b***	**X**				straight line
Y	**=**	***a***	**+**	***b***	**Log(X)**				logarithmic curve
Log(Y)	**=**	***a***	**+**	***b***	**X**				exponential curve
Log(Y)	**=**	***a***	**+**	***b***	**Log(X)**				geometric curve
Y	**=**	***a***	**+**	***b***	**X**	**+**	***c***	$\mathbf{X^2}$	quadratic regression (parabola)

where *X* represents the independent variable and *Y* the dependent variable. The coefficients *a*, *b* and *c* are calculated by the program using the method of least squares.

The following statistics will be displayed in the *results* window:

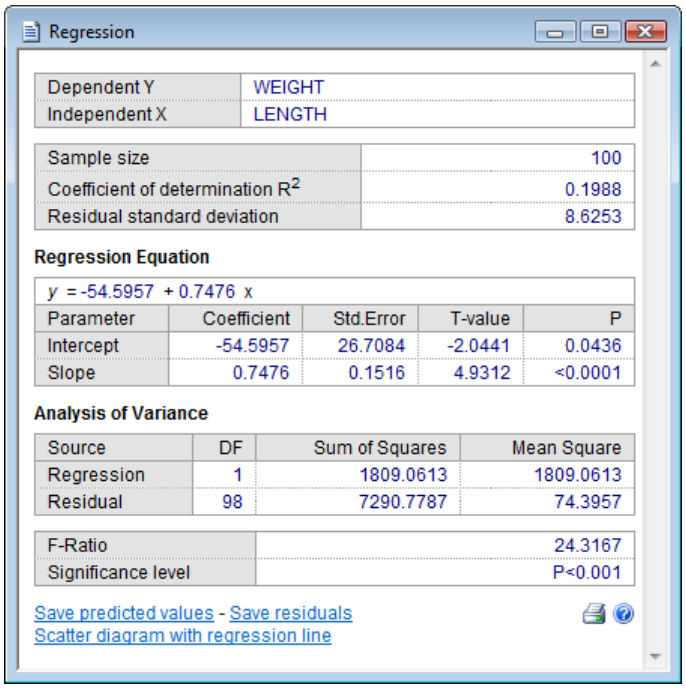

Results

Sample size: the number of data pairs **n**

Coefficient of determination: this is the proportion of the variation in the dependent variable explained by the regression model, and is a measure of the goodness of fit of the model. It can range from 0 to 1, and is calculated as follows:

$$R^2 = \frac{\text{explained variation}}{\text{total variation}} = \frac{\sum(Y_{est} - \overline{Y})^2}{\sum(Y - \overline{Y})^2}$$

where $\mathbf{Y}$ are the observed values for the dependent variable, $\overline{Y}$ is the average of the observed values and $\mathbf{Y}_{est}$ are predicted values for the dependent variable (the predicted values are calculated using the regression equation).

Residual standard deviation: the standard deviation of the residuals (residuals = differences between observed and predicted values). It is calculated as follows:

$$s_{Y.X} = \sqrt{\frac{\sum(Y - Y_{est})^2}{n - 2}}$$

The residual standard deviation is sometimes called the *Standard error of estimate* (Spiegel, 1961).

The equation of the regression curve: the selected equation with the calculated values for *a* and *b* (and for a parabola a third coefficient *c*), e.g. Y = *a* + *b* X

Next, the standard errors are given for the intercept (*a*) and the slope (*b*), followed by the t-value and the P-value for the hypothesis that these coefficients are equal to 0. If the P-values are low (e.g. less than 0.05), then you can conclude that the coefficients are different from 0.

Note that when you use the regression equation for prediction, you may only apply it to values in the range of the actual observations. E.g. when you have calculated the regression equation for height and weight for school children, this equation cannot be applied to adults.

Analysis of variance: the analysis of variance table divides the total variation in the dependent variable into two components, one which can be attributed to the regression model (labeled *Regression*) and one which cannot (labeled *Residual*). If the significance level for the F-test is small (less than 0.05), then the hypothesis that there is no (linear) relationship can be rejected.

Presentation of results

If the analysis shows that the relationship between the two variables is too weak to be of practical help, then there is little point in quoting the equation of the fitted line or curve. If you give the equation, you also report the standard error of the slope, together with the corresponding P-value. Also the residual standard deviation should be reported (Altman, 1980). The number of decimal places of the regression coefficients should correspond to the precision of the raw data.

The accompanying *scatter diagram* should include the *fitted regression line* when this is appropriate. This figure can also include the 95% confidence interval, or the 95% prediction interval, which can be more informative, or both. The legend of the figure must clearly identify the interval that is represented.

Scatter diagram & Regression line

In a scatter diagram, the relation between two numerical variables is presented graphically. One variable (the independent variable X) defines the horizontal axis and the other (dependent variable Y) defines the vertical axis. The values of the two variables on the same row in the data spreadsheet, give the points in the diagram.

The dialog box for the scatter diagram is similar to the one for *Regression* (see p. 74):

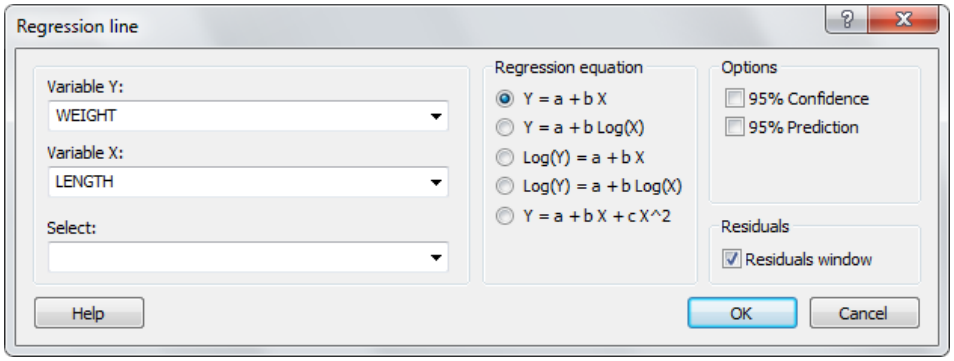

The regression curve will be drawn in the diagram. The equation of this curve is given in the *Regression* results window. When you select an equation that contains a logarithmic transformation for one or both of the variables, the program will use a logarithmic scale for the corresponding variable(s).

Finally, you can select 2 options:

- **95% Confidence**: when you select this option then two curves will be drawn parallel to the regression line. These curves represent a *95% confidence interval* for the regression line. This interval includes the true regression line with 95% probability.
- **95% Prediction**: when you select this option then two curves will be drawn parallel to the regression lines. These curves represent the *95% prediction interval* for the regression curve. The 95% prediction interval is much wider than the 95% confidence interval. For any given value of the independent variable, this interval represents the 95% probability for the values of the dependent variable.

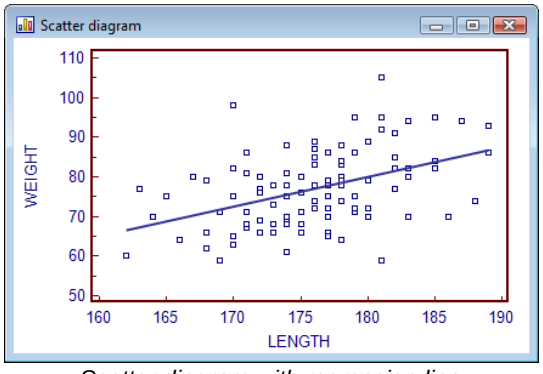

Scatter diagram with regression line

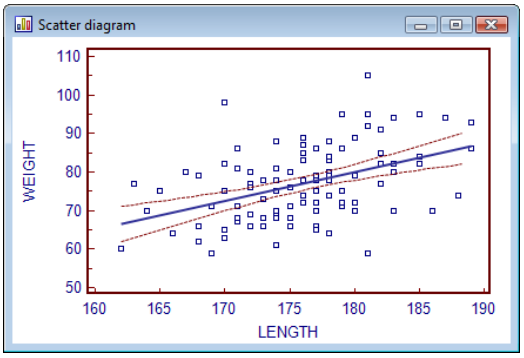

Regression line and 95% confidence interval

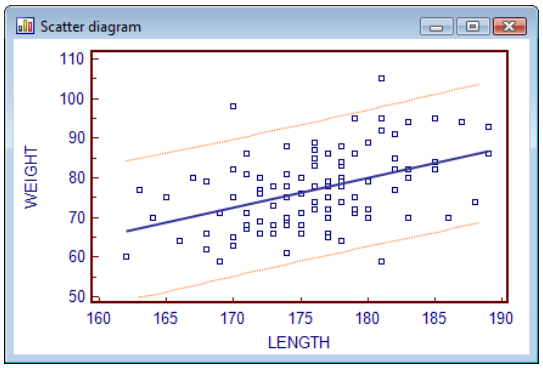

Regression line and 95% prediction interval

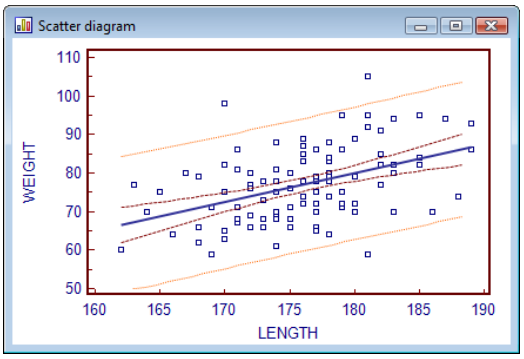

Regression line, 95% confidence interval and 95% prediction interval

When you click a point on the regression line, the program will give the x-value and the f(x) value calculated using the regression equation.

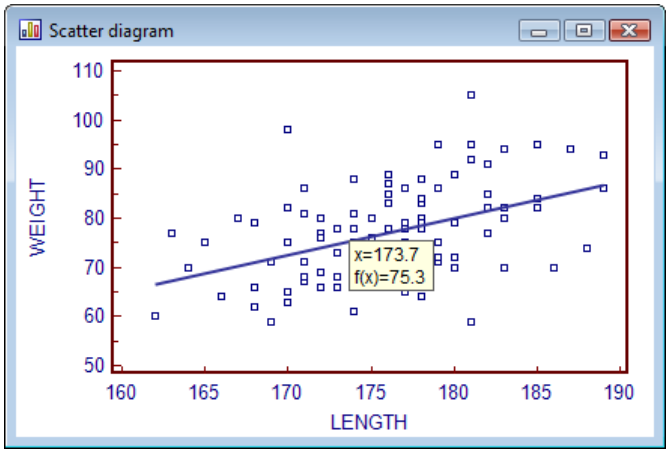

You can press *Ctrl+P* to print the scatter diagram, or function key *F10* to save the picture as file on disk. To define other titles or colors in the graph, or change the axis scaling, see p. 12.

If you want to repeat the scatter diagram, possibly to select a different regression equation, then you only have to press function key *F7*. The dialog box will re-appear with the previous entries (see p. 20).

Extrapolation

MedCalc does only show the regression line in the range of observed values. As a rule, it is not recommended to extrapolate the regression line beyond the observed range. For particular applications however, such as evaluation of stability data, extrapolation may be useful.

To allow extrapolation, right-click in the graph and select *Allow extrapolation* in the popup menu.

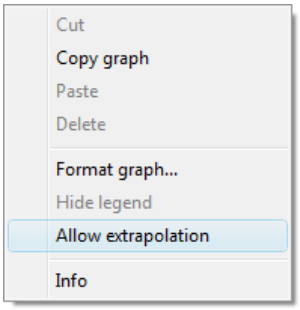

Residuals plot

When you select the option *Residuals plot* in the *Regression line* dialog box, the program will display a second window with the residuals plot. *Residuals* are the differences between the predicted values and the

observed values for the dependent variable. The residual plot allows for the visual evaluation of the goodness of fit of the selected model or equation. Residuals may point to possible outliers (unusual values) in the data or problems with the regression model. If the residuals display a certain pattern, you should consider selecting a different regression model.

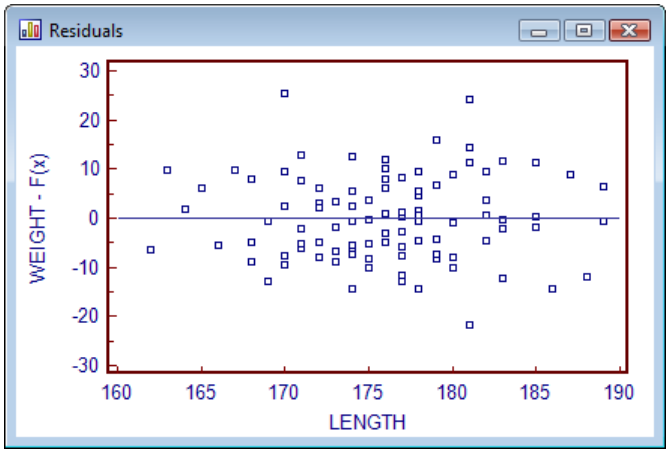

Regression equations & curves

The theoretical shape of the different regression equations is represented in the following figures (linear scaling for both dependent and independent variable). If the scatter diagram (with linear scaling) shows a pattern corresponding with one of the following curves, you should use the corresponding regression equation.

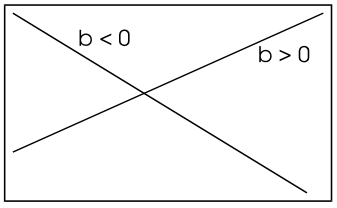

$Y = a + b\,X$

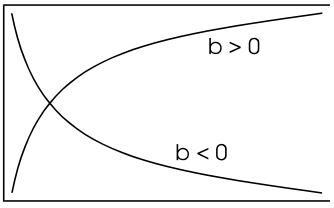

$Y = a + b\,\mathrm{Log}(X)$

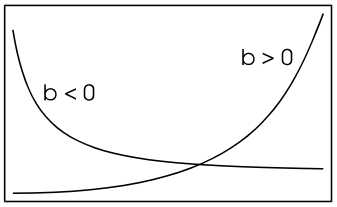

$\mathrm{Log}(Y) = a + b\,X$

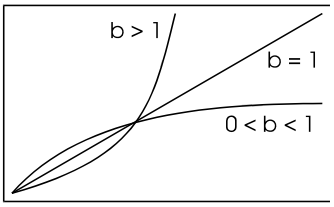

$\mathrm{Log}(Y) = a + b\,\mathrm{Log}(X)$

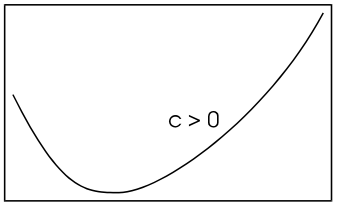

$Y = a + b\,X + c\,X^2$

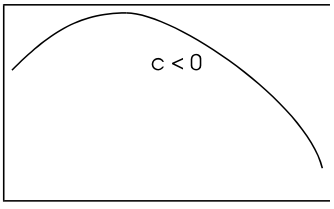

$Y = a + b\,X + c\,X^2$

Multiple regression

Multiple regression is a method used to examine the relationship between one dependent variable *Y* and one or more independent variables X_i. The regression parameters or coefficients b_i in the regression equation

$$Y = b_0 + b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k$$

are estimated using the method of least squares. In this method, the sum of squared residuals between the regression plane and the observed values of the dependent variable are minimized. The regression equation represents a (hyper)plane in a *k+1* dimensional space in which *k* is the number of independent variables X_1, X_2, X_3, ... X_k, plus one dimension for the dependent variable *Y*.

The following need to be entered in the *Multiple regression* dialog box:

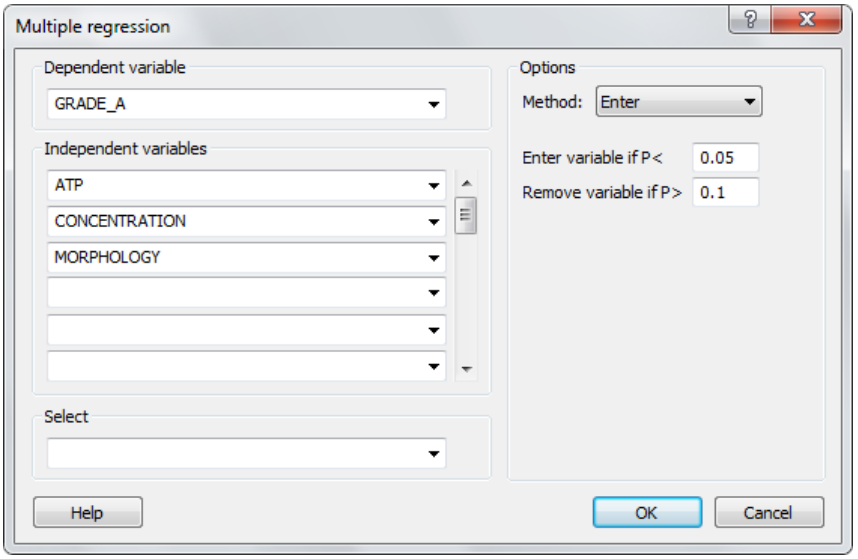

In this dialog box you first identify the *dependent* variable. For the *independent* variables you enter the names of variables that you *expect* to influence the *dependent* variable. Again, you can click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name.

Options

Method: select the way independent variables are entered into the model.

- Enter: enter all variables in the model in one single step, without checking
- Forward: enter significant variables sequentially
- Backward: first enter all variables into the model and next remove the non-significant variables sequentially
- Stepwise: enter significant variables sequentially; after entering a variable in the model, check and possibly remove variables that became non-significant.

Enter variable if P<

A variable is entered into the model if its associated significance level is less than this P-value.

Remove variable if P>

A variable is removed from the model if its associated significance level is greater than this P-value.

Results

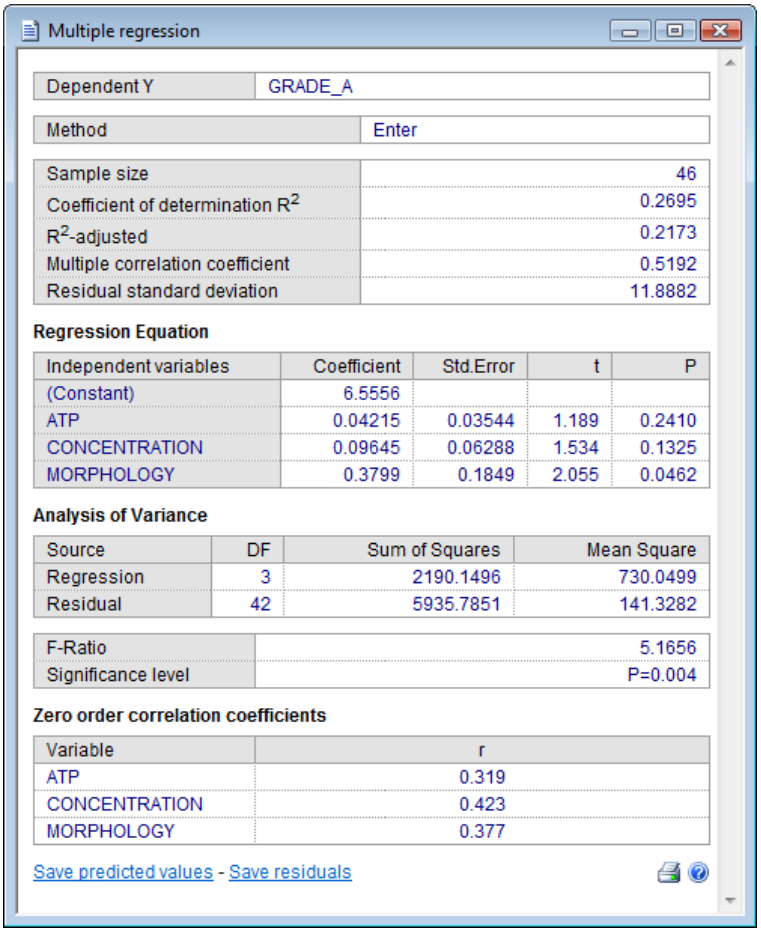

Multiple regression

Dependent Y	GRADE_A
Method	Enter

Sample size	46
Coefficient of determination R^2	0.2695
R^2-adjusted	0.2173
Multiple correlation coefficient	0.5192
Residual standard deviation	11.8882

Regression Equation

Independent variables	Coefficient	Std.Error	t	P
(Constant)	6.5556			
ATP	0.04215	0.03544	1.189	0.2410
CONCENTRATION	0.09645	0.06288	1.534	0.1325
MORPHOLOGY	0.3799	0.1849	2.055	0.0462

Analysis of Variance

Source	DF	Sum of Squares	Mean Square
Regression	3	2190.1496	730.0499
Residual	42	5935.7851	141.3282

F-Ratio	5.1656
Significance level	P=0.004

Zero order correlation coefficients

Variable	r
ATP	0.319
CONCENTRATION	0.423
MORPHOLOGY	0.377

Save predicted values - Save residuals

If you want to repeat the *Multiple regression* procedure, possibly to add or remove variables in the model, then you only have to press function key *F7*. The dialog box will re-appear with the previous entries (see p. 20).

In the *results* window, the following statistics are displayed:

Sample size: the number of data records **n**

Coefficient of determination: this is the proportion of the variation in the dependent variable explained by the regression model, and is a measure of the goodness of fit of the model. It can range from 0 to 1, and is calculated as follows:

$$R^2 = \frac{\text{explained variation}}{\text{total variation}} = \frac{\sum(Y_{est} - \overline{Y})^2}{\sum(Y - \overline{Y})^2}$$

where $\mathbf{Y}$ are the observed values for the dependent variable, $\overline{Y}$ is the average of the observed values and $\mathbf{Y}_{est}$ are predicted values for the dependent variable (the predicted values are calculated using the regression equation).

R^2-adjusted: this is the coefficient of determination adjusted for the number of independent variables in the regression model. Unlike the coefficient of determination, R^2-adjusted may decrease if variables are entered in the model that do not add significantly to the model fit.

$$R^2_{adj} = 1 - \frac{\text{unexplained variation} / (n-k-1)}{\text{total variation} / (n-1)}$$

or

$$R^2_{adj} = 1 - \frac{\sum(Y_{est} - Y)^2}{\sum(Y - \bar{Y})^2} \frac{(n-1)}{(n-k-1)}$$

Multiple correlation coefficient: this coefficient is a measure of how tightly the data points cluster around the regression plane, and is calculated by taking the square root of the coefficient of determination.

When discussing multiple regression analysis results, generally the coefficient of multiple determination is used rather than the multiple correlation coefficient.

Residual standard deviation: the standard deviation of the residuals (residuals = differences between observed and predicted values). It is calculated as follows:

$$s_{res} = \sqrt{\frac{\sum(Y - Y_{est})^2}{n-k-1}}$$

The regression equation: the different regression coefficients b_i with standard error s_{bi}, t-value and P-value. The P-value is the probability that you would have found the current result if the coefficient were equal to 0 (null hypothesis). If the P-value for one or more coefficients is less than the conventional 0.05, then these coefficients can be called statistically significant, and the corresponding independent variables exert independent effects on the dependent variable Y

Analysis of variance: the analysis of variance table divides the total variation in the dependent variable into two components, one which can be attributed to the regression model (labeled *Regression*) and one which cannot (labeled *Residual*). If the significance level for the F-test is small (less than 0.05), then the hypothesis that there is no (linear) relationship can be rejected, and the multiple correlation coefficient can be called statistically significant.

Zero order correlation coefficients: these are the simple correlation coefficients for the dependent variable *Y* and all independent variables X_i separately.

Logistic regression

Logistic regression is a technique for analyzing problems in which there are one or more independent variables that determine an outcome. The outcome is measured with a dichotomous variable (in which there are only two possible outcomes).

In logistic regression, the dependent variable is binary or dichotomous, i.e. it only contains data coded as 1 (TRUE, success, pregnant, etc.) or 0 (FALSE, "failure", "non-pregnant", etc.).

The goal of logistic regression is to find the best fitting (yet biologically reasonable) model to describe the relationship between the dichotomous characteristic of interest (dependent variable = response or outcome variable) and a set of independent (predictor or explanatory) variables. Logistic regression generates the coefficients (and its standard errors and significance levels) of a formula to predict a *logit transformation* of the probability of presence of the characteristic of interest:

$$Y = b_0 + b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k$$

where p is the probability of presence of the characteristic of interest. The logit transformation is defined as the logged odds:

$$\text{odds} = \frac{p}{1-p} = \frac{\text{probability of presence of characteristic}}{\text{probability of absence of characteristic}}$$

and

$$Y = b_0 + b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k$$

Rather than choosing parameters that minimize the sum of squared errors (like in ordinary regression), estimation in logistic regression chooses parameters that maximize the likelihood of observing the sample values.

The MedCalc dialog box for logistic regression is similar to the one for multiple regression (see p. 79). In the dialog box you first identify the *dependent* variable. Remember that the dependent variable must be binary or dichotomous, and it should only contain data coded as 0 or 1. Cases with values other than 0 or 1 for the dependent variable will be excluded from the analysis!

For the *independent* variables you enter the names of variables that you *expect* to influence the *dependent* variable.

Again, you click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name.

Options

Method: select the way independent variables are entered into the model.

- Enter: enter all variables in the model in one single step, without checking
- Forward: enter significant variables sequentially
- Backward: first enter all variables into the model and next remove the non-significant variables sequentially
- Stepwise: enter significant variables sequentially; after entering a variable in the model, check and possibly remove variables that became non-significant.

Enter variable if P<

A variable is entered into the model if its associated significance level is less than this P-value.

Remove variable if P>

A variable is removed from the model if its associated significance level is greater than this P-value.

Classification table cutoff value: a value between 0 and 1 which will be used as a cutoff value for a classification table. The classification table is a method to evaluate the logistic regression model. In this table the observed values for the dependent outcome and the predicted values (at the selected cut-off value) are cross-classified.

Categorical: click this button to identify nominal categorical variables.

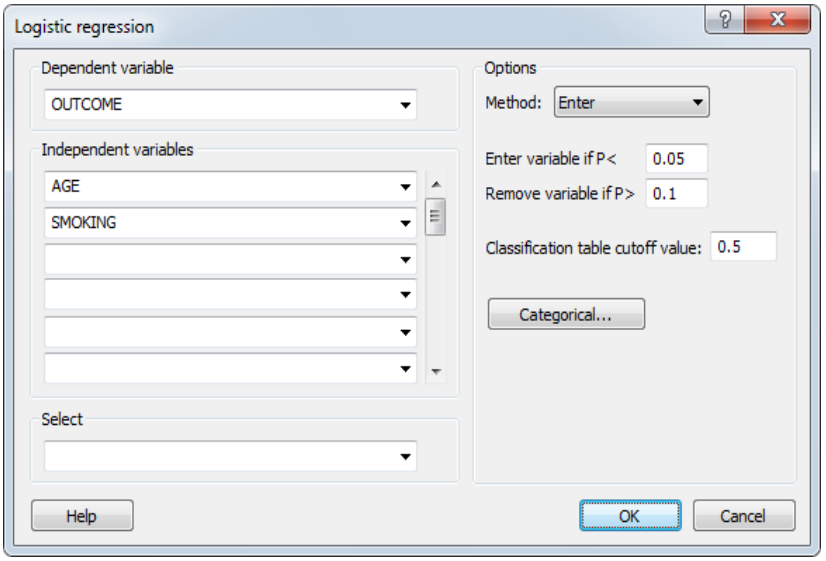

After you click OK, the following results are displayed in the results window:

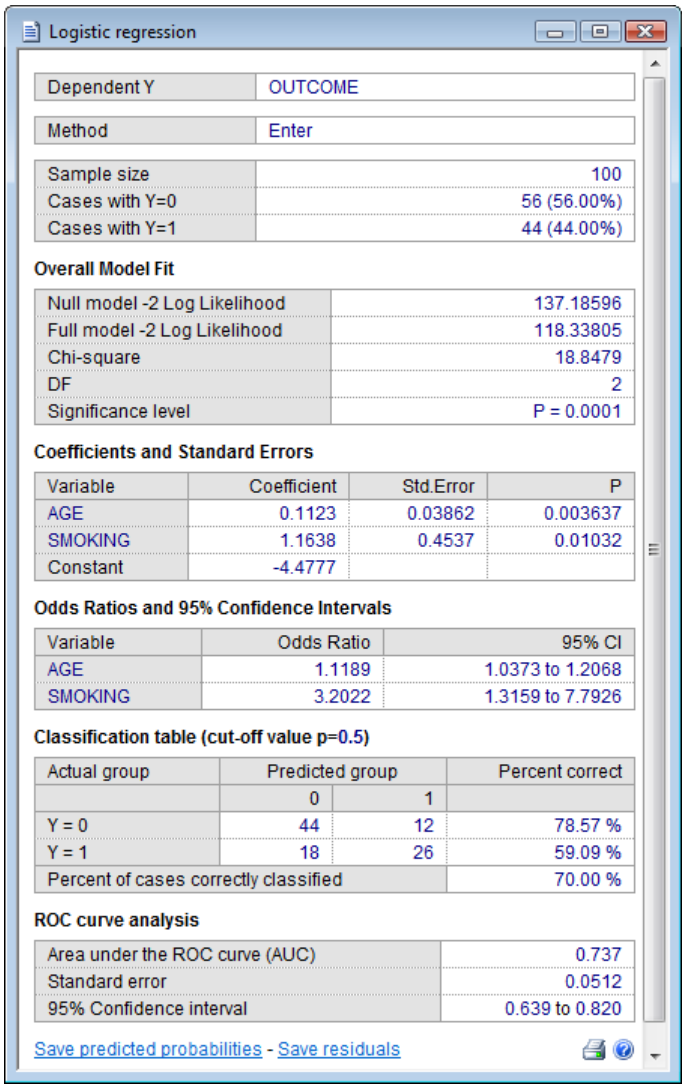

Logistic regression

Dependent Y	OUTCOME
Method	Enter

Sample size	100
Cases with Y=0	56 (56.00%)
Cases with Y=1	44 (44.00%)

Overall Model Fit

Null model -2 Log Likelihood	137.18596
Full model -2 Log Likelihood	118.33805
Chi-square	18.8479
DF	2
Significance level	P = 0.0001

Coefficients and Standard Errors

Variable	Coefficient	Std.Error	P
AGE	0.1123	0.03862	0.003637
SMOKING	1.1638	0.4537	0.01032
Constant	-4.4777		

Odds Ratios and 95% Confidence Intervals

Variable	Odds Ratio	95% CI
AGE	1.1189	1.0373 to 1.2068
SMOKING	3.2022	1.3159 to 7.7926

Classification table (cut-off value p=0.5)

Actual group	Predicted group		Percent correct
	0	1	
Y = 0	44	12	78.57 %
Y = 1	18	26	59.09 %
Percent of cases correctly classified			70.00 %

ROC curve analysis

Area under the ROC curve (AUC)	0.737
Standard error	0.0512
95% Confidence interval	0.639 to 0.820

Save predicted probabilities - Save residuals

Sample size and cases with negative and positive outcome

First the program gives sample size and the number and proportion of cases with a negative (Y=0) and positive (Y=1) outcome.

Overall model fit

The *null model* -2 Log Likelihood is given by -2 * ln(L_0) where L_0 is the likelihood of obtaining the observations if the independent variables had no effect on the outcome.

The *full model* -2 Log Likelihood is given by -2 * ln(L) where L is the likelihood of obtaining the observations with all independent variables incorporated in the model.

The difference of these two yields a Chi-Square statistic, which is a measure of how well the independent variables affect the outcome or dependent variable.

If the P-value for the overall model fit statistic is less than the conventional 0.05 then there is evidence that at least one of the independent variables contributes to the prediction of the outcome.

Regression coefficients

The regression coefficients are the coefficients b_0, b_1, b_2, ... b_k of the regression equation:

$$Y = b_0 + b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k$$

An independent variable with a regression coefficient not significantly different from 0 (P>0.05) can be removed from the regression model (press function key F7 to repeat the logistic regression procedure). If P<0.05 then the variable contributes significantly to the prediction of the outcome variable.

The logistic regression coefficients show the change (increase when $b_i>0$, decrease when $b_i<0$) in the predicted logged odds of having the characteristic of interest for a one-unit change in the independent variables.

When the independent variables X_a and X_b are dichotomous variables (e.g. Smoking, Sex) then the influence of these variables on the dependent variable can simply be compared by comparing their regression coefficients b_a and b_b.

Odds ratios with 95% CI

By taking the exponential of both sides of the regression equation as given above, the equation can be rewritten as:

$$\text{odds} = \frac{p}{1-p} = e^{b_0} \times e^{b_1 X_1} \times e^{b_2 X_2} \times e^{b_3 X_3} \times \ldots \times e^{b_k X_k}$$

It is clear that when a variable X_i increases by 1 unit, with all other factors remaining unchanged, then the odds will increase by a factor e^{b_i}. This factor e^{b_i} is the odds ratio (O.R.) for the independent variable X_i and it gives the *relative* amount by which the odds of the outcome increase (O.R. greater than 1) or decrease (O.R. less than 1) when the value of the independent variable is increased by 1 units.

E.g. The variable SMOKING is coded as 0 (= no smoking) and 1 (= smoking), and the odds ratio for this variable is 3.2. This means that in the model the odds for a positive outcome in cases that do smoke are 3.2 times higher than in cases that do not smoke.

Interpretation of the fitted equation

The logistic regression equation is:

$$\text{logit}(p) = -4.48 + 0.11 \times \text{AGE} + 1.16 \times \text{SMOKING}$$

So for 40 years old cases who do smoke logit(p) equals 1.08. Logit(p) can be back-transformed to p by the following formula:

$$p = \frac{1}{1+e^{-\text{logit}(p)}}$$

Alternatively, you can use the table on p. 229. For logit(p)=1.08 the probability p of having a positive outcome equals 0.75.

Classification table

The classification table can be used to evaluate the predictive accuracy of the logistic regression model. In this table the observed values for the dependent outcome and the predicted values (at a user defined cut-off value, for example p=0.50) are cross-classified. In our example, the model correctly predicts 70% of the cases.

ROC curve analysis

Another method to evaluate the logistic regression model makes use of ROC curve analysis. In this analysis, the power of the model's predicted values to discriminate between positive and negative cases is quantified by the Area under the ROC curve (AUC). The AUC, sometimes referred to as the c-statistic (or concordance index), is a value that varies from 0.5 (discriminating power not better than chance) to 1.0 (perfect discriminating power).

To perform a full ROC curve analysis on the predicted probabilities you can save the predicted probabilities and next use this new variable in ROC curve analysis (p. 156). The Dependent variable used in Logistic Regression then acts as the Classification variable in the ROC curve analysis dialog box.

Literature

- Pampel FC (2000) Logistic regression: A primer. Sage University Papers Series on Quantitative Applications in the Social Sciences, 07-132. Thousand Oaks, CA: Sage.

One sample t-test

Use the one sample t-test to test whether the average of observations differs significantly from a test value.

Required input

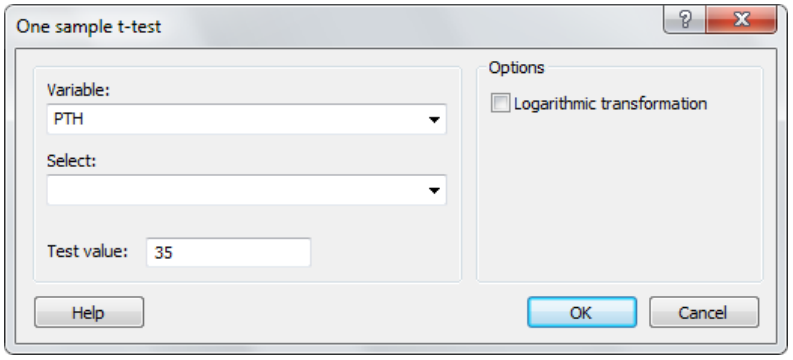

- The variable of interest. You can use the button to select variables and selection criteria.
- The test value you want to compare the sample data with.

Options

- Logarithmic transformation: select this option if the data are positively skewed.

Results

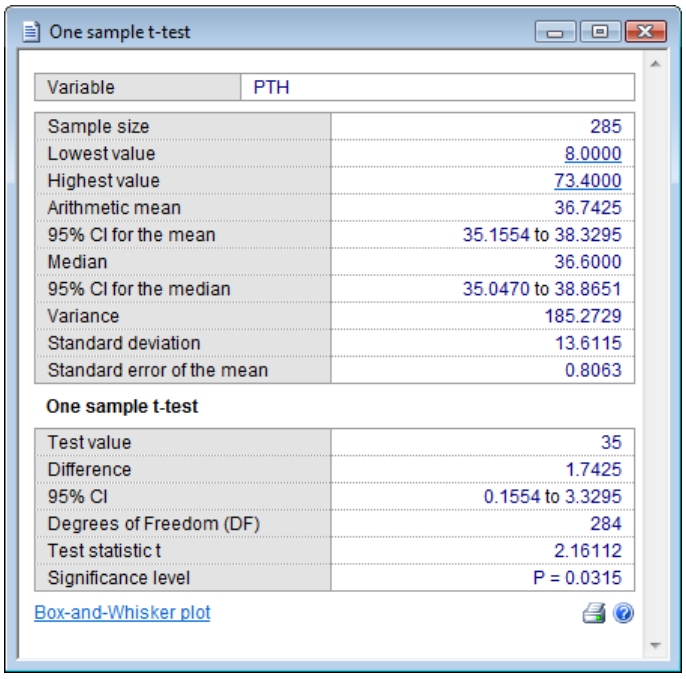

The results window first displays the summary statistics of the observations.
The One sample t-test table shows:

- The test value
- The difference Sample mean - Test value
- The 95% Confidence Interval (CI) for this difference
- The Degrees of Freedom (DF), t-statistic, and corresponding (two-sided) P-value.

If the calculated P-value is less than 0.05, the conclusion is that, statistically, the sample mean is significantly different from the test value.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of the observations, but reports the back-transformed summary statistics.

For the One sample t-test, the difference between sample mean and test value, with 95% confidence interval, are given on the log-transformed scale.

Next, the results are transformed back and the interpretation is as follows: the back-transformed difference of the logs is ratio of the sample mean / test value on the original scale.

Independent samples t-test

The independent samples t-test is used to compare the means of two independent samples.

Required input

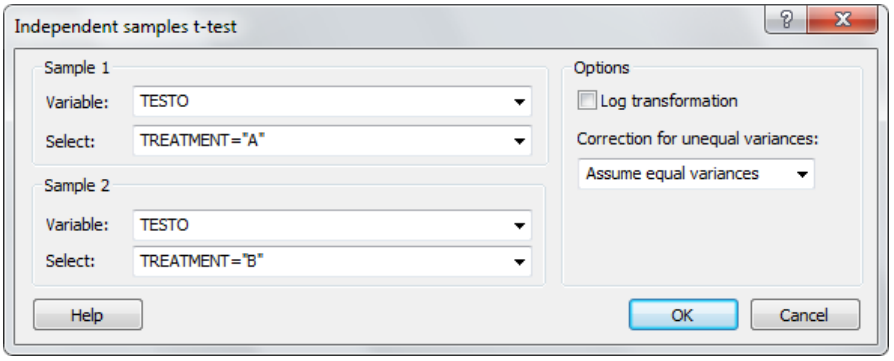

Select the variables for sample 1 and sample 2. Use the button to select variables and selection criteria in the variables list.

Options

- Logarithmic transformation: select this option if the data are positively skewed.
- Correction for unequal variances: allows to select the t-test (assuming equal variances) or the t-test corrected for unequal variances (Welch test, Armitage et al., 2002). With the option "Automatic" the software will select the appropriate test based on the F-test (comparison of variances).

Results

After you have clicked OK, the program displays the summary statistics of the two samples, followed by the statistical tests.

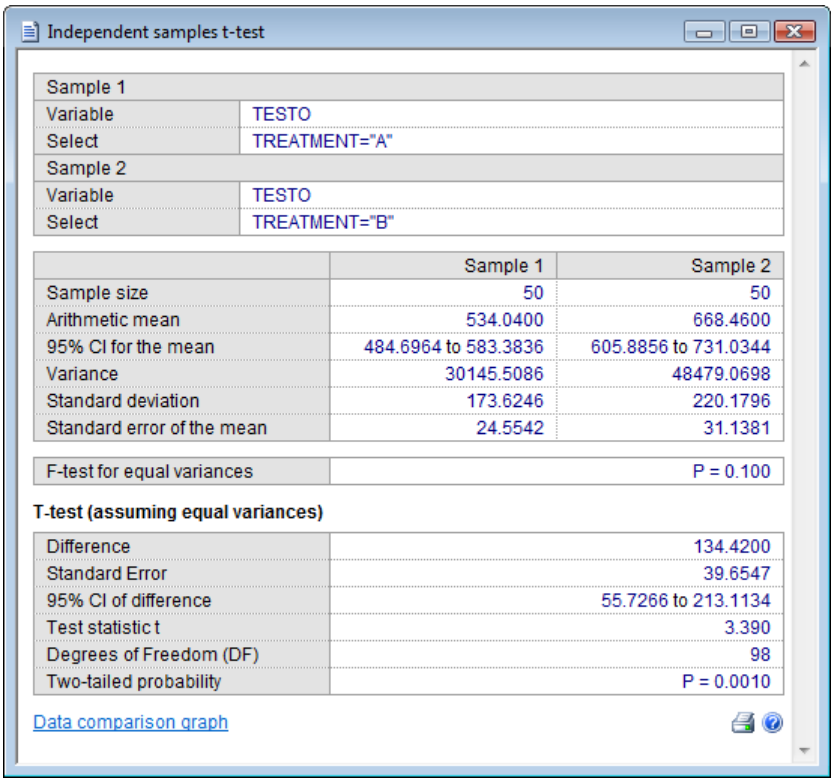

Independent samples t-test

Sample 1	
Variable	TESTO
Select	TREATMENT="A"
Sample 2	
Variable	TESTO
Select	TREATMENT="B"

	Sample 1	Sample 2
Sample size	50	50
Arithmetic mean	534.0400	668.4600
95% CI for the mean	484.6964 to 583.3836	605.8856 to 731.0344
Variance	30145.5086	48479.0698
Standard deviation	173.6246	220.1796
Standard error of the mean	24.5542	31.1381

F-test for equal variances	P = 0.100

T-test (assuming equal variances)

Difference	134.4200
Standard Error	39.6547
95% CI of difference	55.7266 to 213.1134
Test statistic t	3.390
Degrees of Freedom (DF)	98
Two-tailed probability	P = 0.0010

Data comparison graph

First an F-test is performed. If the P-value is low (P<0.05) the variances of the two samples cannot be assumed to be equal and it should be considered to use the t-test with a correction for unequal variances (Welch test) (see above).

The independent samples t-test is used to test the hypothesis that the difference between the means of two samples is equal to 0 (this hypothesis is therefore called the null hypothesis). The program displays the difference between the two means, and the 95% Confidence Interval (CI) of this difference. Next follow the test statistic t, the Degrees of Freedom (DF) and the two-tailed probability P. When the P-value is less than the conventional 0.05, the null hypothesis is rejected and the conclusion is that the two means do indeed differ significantly.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of the observations, but reports the back-transformed summary statistics.

For the t-test, the difference and 95% confidence are given, and the test is performed, on the log-transformed scale.

Next, the results of the t-test are transformed back and the interpretation is as follows: the back-transformed difference of the means of the logs is the ratio of the geometric means of the two samples (see Bland, 2000).

One-sided or two-sided tests

In MedCalc, P-values are always *two-sided* (as recommended by Fleiss, 1981, and Altman, 1991) and not *one-sided*.

- A *two-sided* (or two-tailed) P-value is appropriate when the difference between the two means can occur in both directions: it may be either negative or positive; the mean of one sample may either be smaller or larger than that of the other sample.
- A *one-sided* test should only be performed when, before the start of the study, it has already been established that a difference can only occur in one direction. E.g. when the mean of sample *A* must be more than the mean of sample *B* for reasons other than those connected with the sample(s).

Interpretation of P-values

P-values should not be interpreted too strictly. Although a significance level of 5% is generally accepted as a cut-off point for a significant versus a non-significant result, it would be a mistake to interpret a shift of P-

value from e.g. 0.045 to 0.055 as a change from significance to non-significance. Therefore the real P-values are preferably reported, P=0.045 or P=0.055, instead of P<0.05 or P>0.05, so the reader can make his own interpretation.

With regards to the interpretation of P-values as significant versus not-significant, is has been recommended to select a smaller significance level of for example 0.01 *when it is necessary to be quite certain that a difference exists before accepting it*. When a study is designed to *uncover a difference*, or when *a life-saving* drug is being studied, we should be willing to accept that there is a difference even when the P-value is as large as 0.10 or even 0.20 (Lentner, 1982). The latter states: "*The tendency in medical and biological investigations is to use too small a significance probability*".

Confidence intervals

Whereas the P-value may give information on the statistical significance of the result, the 95% confidence interval gives information to assess the clinical importance of the result.

When the number of cases included in the study is large, a biologically unimportant difference can be statistically highly significant. *A statistically significant result does not necessarily indicate a real biological difference*.

On the other hand, a high P-value can lead to the conclusion of statistically non-significant difference although the difference is clinically meaningful and relevant, especially when the number of cases is small. *A non-significant result does not mean that there is no real biological difference*.

Confidence intervals are therefore helpful in interpretation of a difference, whether or not it is statistically significant (Altman et al., 1983).

Presentation of results

It is recommended to report the results of the t-test (and other tests) not by a simple statement such as P<0.05, but by giving full statistical information, as in the following example by Gardner & Altman (1986):

> *The difference between the sample mean systolic blood pressure in diabetics and non-diabetics was 6.0 mm Hg, with a 95% confidence interval from 1.1 to 10.9 mm Hg; the t test statistic was 2.4, with 198 degrees of freedom and an associated P value of P=0.02.*

In short:

> *Mean 6.0 mm Hg, 95% CI 1.1 to 10.9; t=2.4, df=198, P=0.02*

Paired samples t-test

The *paired samples* t-test is used to test the null hypothesis that the average of the differences between a series of paired observations is zero. Observations are paired when, for example, they are performed on the same samples or subjects.

Required input

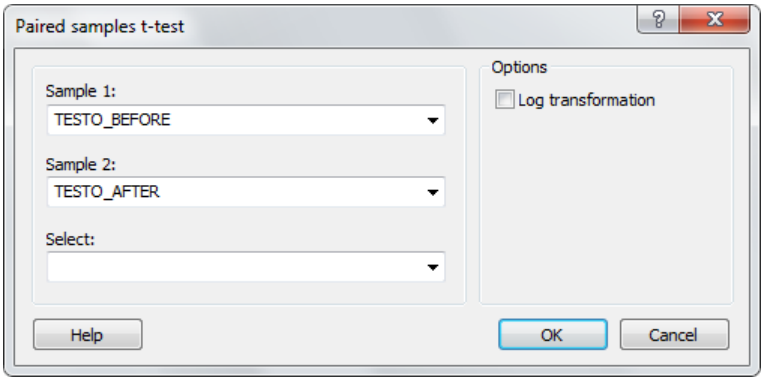

First select the variables for sample 1 and sample 2, and a possible selection criterion for the data pairs. Use the ▾ button to select variables and selection criteria in the variables list.

Options

Logarithmic transformation: select this option if the data are positively skewed.

Results

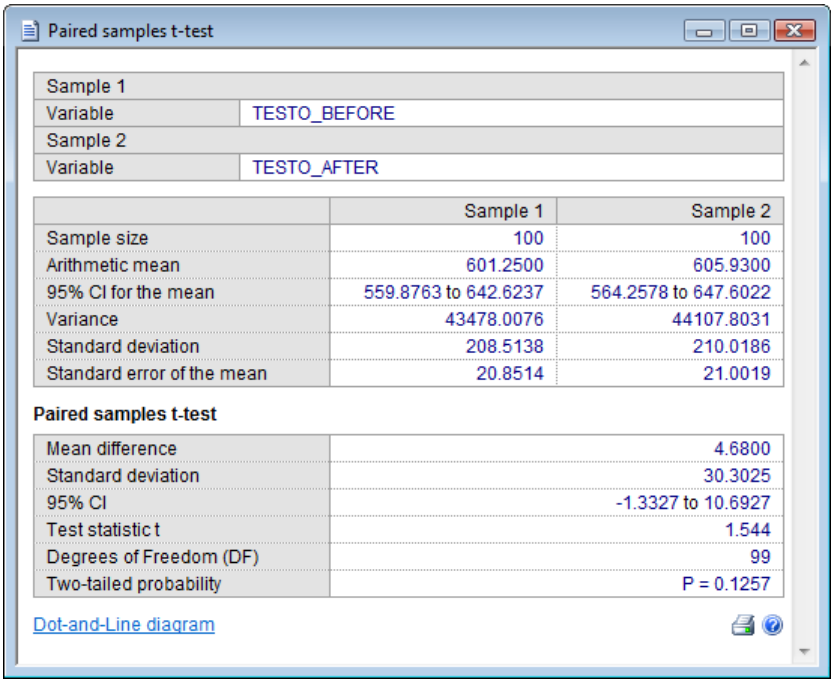

Paired samples t-test

Sample 1	
Variable	TESTO_BEFORE
Sample 2	
Variable	TESTO_AFTER

	Sample 1	Sample 2
Sample size	100	100
Arithmetic mean	601.2500	605.9300
95% CI for the mean	559.8763 to 642.6237	564.2578 to 647.6022
Variance	43478.0076	44107.8031
Standard deviation	208.5138	210.0186
Standard error of the mean	20.8514	21.0019

Paired samples t-test

Mean difference	4.6800
Standard deviation	30.3025
95% CI	-1.3327 to 10.6927
Test statistic t	1.544
Degrees of Freedom (DF)	99
Two-tailed probability	P = 0.1257

Dot-and-Line diagram

The program displays the summary statistics of the two samples followed by the mean of the differences between the paired observations, and the standard deviation of these differences, followed by a 95% confidence interval for the mean.

Note that the sample size will always be equal (only cases are included with data available for the two variables).

Next the result of the null hypothesis test is displayed. If the calculated P-value is less than 0.05, the conclusion is that the mean difference between the paired observations is statistically significantly different from 0.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of the observations, but reports the back-transformed summary statistics.

For the paired t-test, the mean difference and 95% confidence are given on the log-transformed scale.

Next, the results of the t-test are transformed back and the interpretation is as follows: the back-transformed mean difference of the logs is the geometric mean of the ratio of paired values on the original scale (see Altman, 1991).

Signed rank sum test

The **Signed rank sum test** is a test for symmetry about a test value. This test is the non-parametric alternative for the One sample t-test. It can be used when the observations are not Normally distributed.

Required input

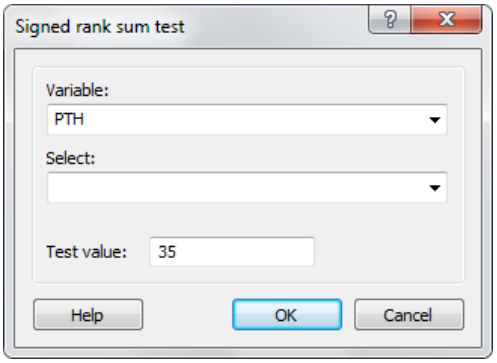

- The variable of interest. You can use the button to select variables and selection criteria.
- The test value you want to compare the sample data with.

Results

The results windows for the Signed rank sum test first displays summary statistics of the sample.

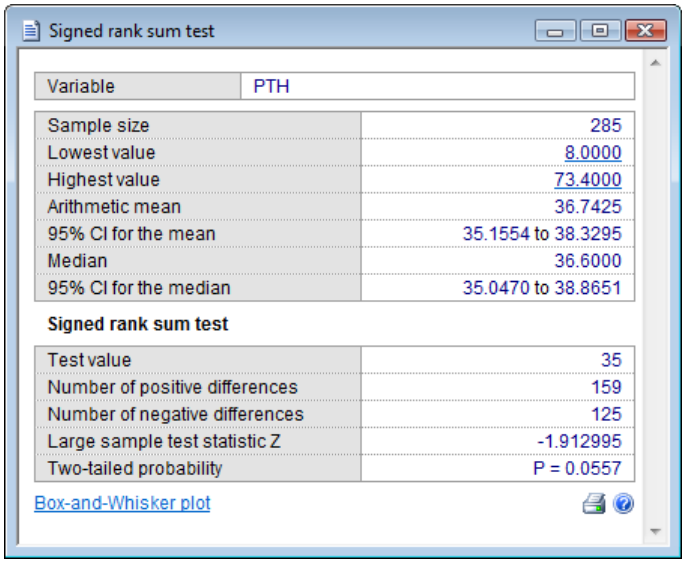

The Signed rank sum test ranks the absolute values of the differences between the sample data and the test value, and calculates a statistic on the number of negative and positive differences.

If the resulting P-value is small (P<0.05), then the sample data are not symmetrical about the test value and therefore a statistically significant difference can be accepted between the sample median and the test value.

Mann-Whitney test (independent samples)

The Mann-Whitney test is the non-parametric equivalent of the independent samples t-test. It should be used when the sample data are not Normally distributed, and they cannot be transformed to a Normal distribution by means of a logarithmic transformation.

Required input

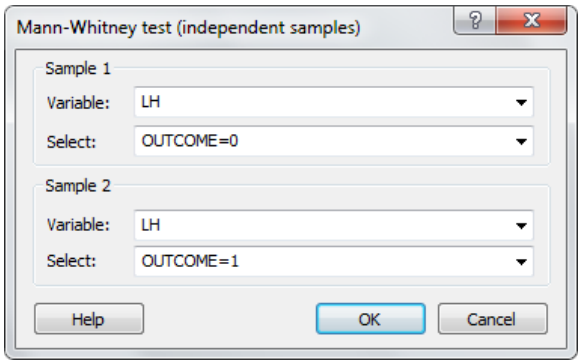

Select the variables for sample 1 and sample 2. Use the button to select variables and selection criteria in the variables list.

Results

The **Mann-Whitney test** combines and ranks the data from sample 1 and sample 2 and calculates a statistic on the difference between the sum of the ranks of sample 1 and sample 2.

If the resulting P-value is small (P<0.05) then a statistically significant difference between the two samples can be accepted.

When either or both sample sizes are large (>20) then MedCalc uses the Normal approximation (Lentner, 1982) to calculate the P-value. For small sample sizes, in the absence of ties, MedCalc calculates the exact probability (Conover, 1999).

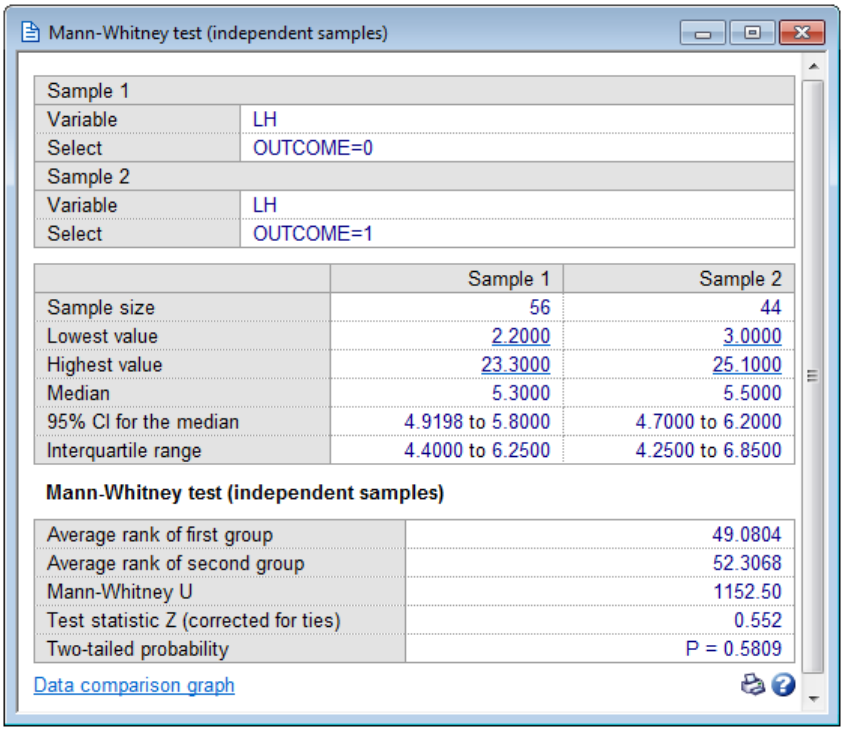

Wilcoxon test (paired samples)

The Wilcoxon test for paired samples is the non-parametric equivalent of the paired samples t-test. It should be used when the sample data are not Normally distributed, and they cannot be transformed to a Normal distribution by means of a logarithmic transformation.

Required input

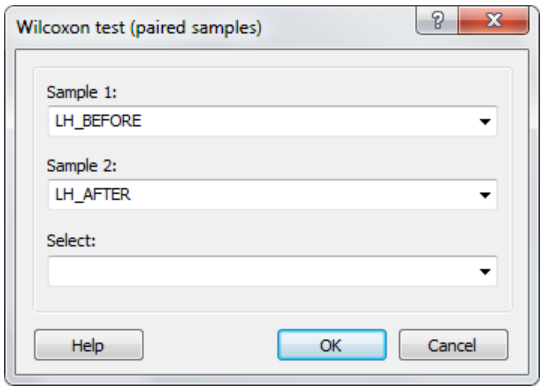

Select the variables for sample 1 and sample 2, and a possible selection criterion for the cases to be included. Use the button to select variables and selection criteria in the variables list.

Results

The **Wilcoxon test** for paired samples ranks the absolute values of the differences between the paired data in sample 1 and sample 2 and calculates a statistic on the number of negative and positive differences (differences are calculated as sample 2 - sample 1)..

If the resulting P-value is small (P<0.05) then it can be accepted that the median of the differences between the paired observations is statistically significant different from 0.

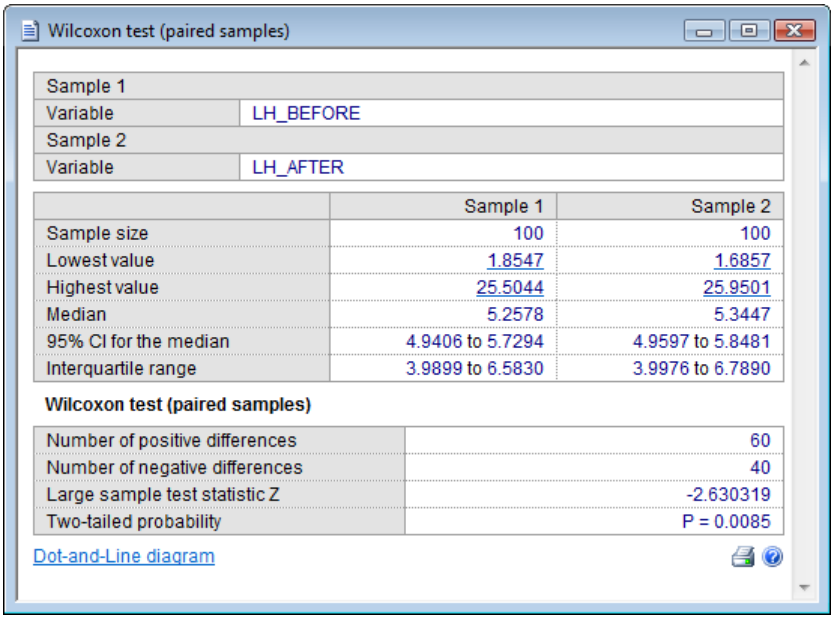

Wilcoxon test (paired samples)

Sample 1	
Variable	LH_BEFORE
Sample 2	
Variable	LH_AFTER

	Sample 1	Sample 2
Sample size	100	100
Lowest value	1.8547	1.6857
Highest value	25.5044	25.9501
Median	5.2578	5.3447
95% CI for the median	4.9406 to 5.7294	4.9597 to 5.8481
Interquartile range	3.9899 to 6.5830	3.9976 to 6.7890

Wilcoxon test (paired samples)

Number of positive differences	60
Number of negative differences	40
Large sample test statistic Z	-2.630319
Two-tailed probability	P = 0.0085

Dot-and-Line diagram

Variance ratio test (F-test)

The *Variance ratio test (F-test)* is used to compare the variances of two samples.

Required input

In the dialog box first identify the variables for sample 1 and sample 2. Use the button to select variables and selection criteria in the variables list.

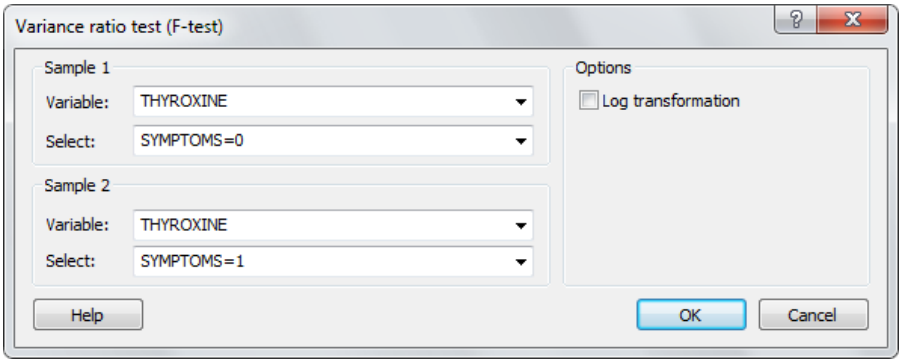

Options

Logarithmic transformation: select this option if the data are positively skewed.

Results

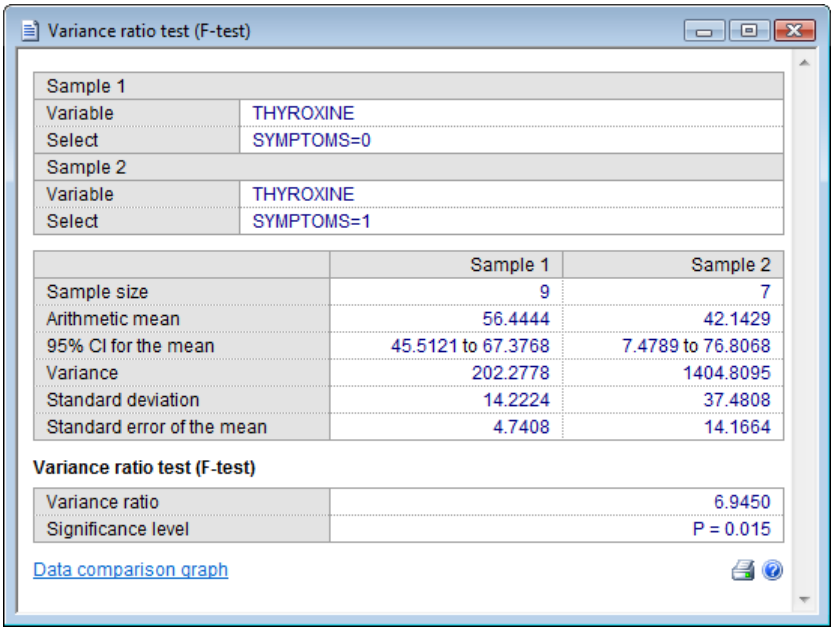

The program displays summary statistics and the variances of the two samples. Next the results of the variance ratio test and the F-statistic is given, which is the ratio of the larger variance over the smaller, with its associated (two-sided) P-value. If the P-value is less than the conventional 0.05, the null hypothesis is rejected and the conclusion is that the two variances differ significantly.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of the observations, but reports the back-transformed summary statistics.

The variance of the logs cannot be back-transformed meaningfully and therefore the variances and the variance ratio are given on the log-transformed scale.

One-way analysis of variance

One-way analysis of variance is used to test the difference between the means of several subgroups of a variable (multiple testing). For a graphical representation of ANOVA, refer to *Multiple comparison graphs* (p. 172).

Data

The following figure illustrates how data need to be entered. For ANOVA, you need one continuous variable (concentration) and one qualitative variable (grade). The data for each case are entered on one row of the spreadsheet.

The qualitative variable may contain character or numeric codes. These codes are used to break-up the data into several subgroups for the ANOVA procedure, to calculate the *Between groups* and *Within groups* variation.

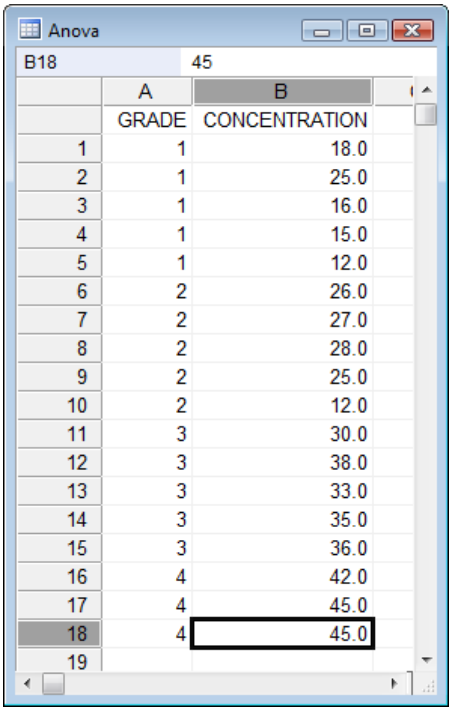

	A	B
	GRADE	CONCENTRATION
1	1	18.0
2	1	25.0
3	1	16.0
4	1	15.0
5	1	12.0
6	2	26.0
7	2	27.0
8	2	28.0
9	2	25.0
10	2	12.0
11	3	30.0
12	3	38.0
13	3	33.0
14	3	35.0
15	3	36.0
16	4	42.0
17	4	45.0
18	4	45.0
19		

Required input

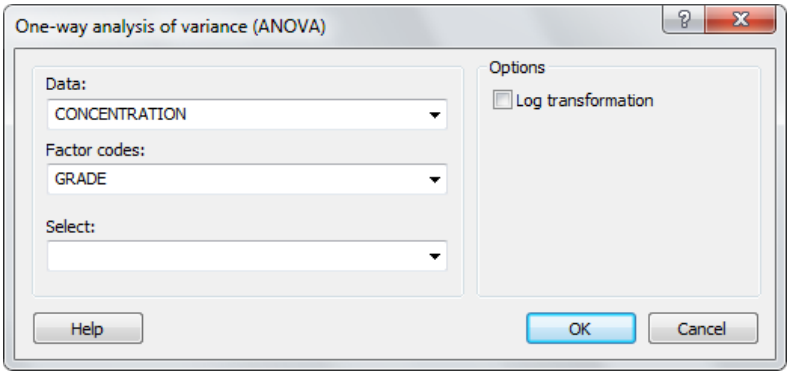

For *Data*, select a continuous variable, and for *Factor codes* the qualitative factor.

Options

Logarithmic transformation: select this option if the data of the continuous variable are positively skewed.

Results

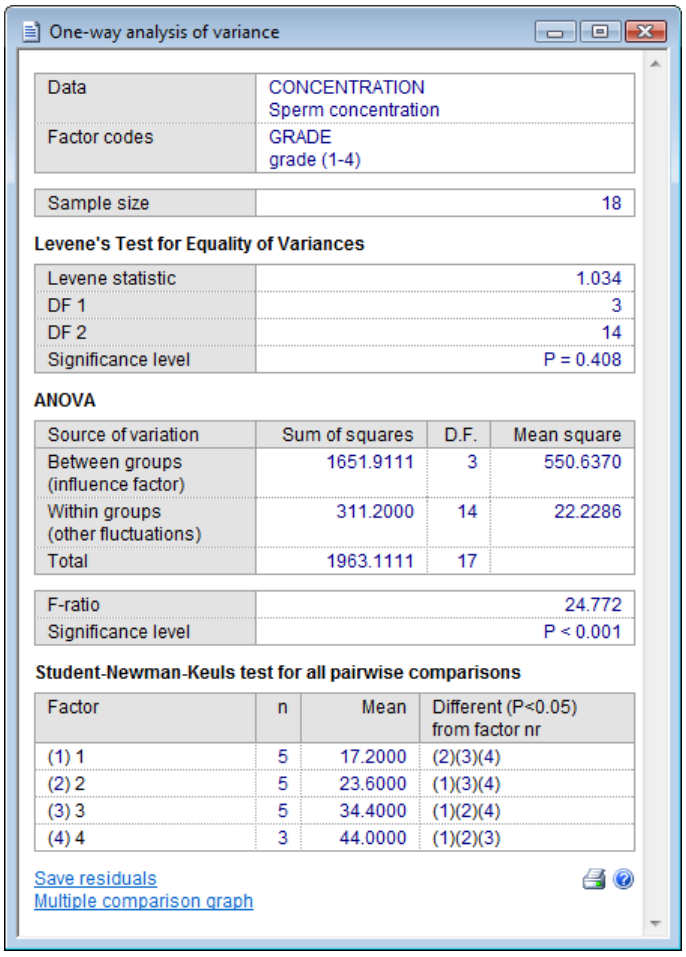

One-way analysis of variance

Data	CONCENTRATION Sperm concentration
Factor codes	GRADE grade (1-4)

Sample size	18

Levene's Test for Equality of Variances

Levene statistic	1.034
DF 1	3
DF 2	14
Significance level	P = 0.408

ANOVA

Source of variation	Sum of squares	D.F.	Mean square
Between groups (influence factor)	1651.9111	3	550.6370
Within groups (other fluctuations)	311.2000	14	22.2286
Total	1963.1111	17	

F-ratio	24.772
Significance level	P < 0.001

Student-Newman-Keuls test for all pairwise comparisons

Factor	n	Mean	Different (P<0.05) from factor nr
(1) 1	5	17.2000	(2)(3)(4)
(2) 2	5	23.6000	(1)(3)(4)
(3) 3	5	34.4000	(1)(2)(4)
(4) 4	3	44.0000	(1)(2)(3)

Save residuals
Multiple comparison graph

Levene's Test for Equality of Variances

Prior to the ANOVA test, Levene's Test for Equality of Variances is performed. If the Levene test is positive (P<0.05) then the variances in the different groups are different (the groups are not homogeneous) and you may need to apply a logarithmic transformation to the data, or use a non-parametric statistic.

Anova

The results of the ANOVA are presented in an ANOVA table, followed by the F statistic and associated P value. If the P value is less than 0.05, then you can accept the hypothesis that the means of at least two of the subgroups differ significantly.

Student-Newman-Keuls test

If the Anova test is positive (P<0.05) then MedCalc performs a *Student-Newman-Keuls te*st for pairwise comparison of subgroups.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of the dependent variable, but the different means are back-transformed and reported as the geometric means.

Two-way analysis of variance

The two-way analysis of variance is an extension to the one-way analysis of variance. There are two qualitative factors (A and B) on one dependent variable Y..

Three null hypotheses are tested in this procedure:

- factor A does not influence variable Y
- factor B does not influence variable Y
- the effect of factor A on variable Y does not depend on factor B (i.e. there is no interaction of factors A and B).

Two-way analysis of variance requires that there are data for each combination of the two qualitative factors A and B. .

Required input

Select the (continuous) dependent variable (Y) and two discrete variables for the qualitative factors (A and B) suspected to influence the dependent variable. The qualitative factors A and B may either consist of numeric or alphanumeric data. One selection criterion can also be defined in order to include only a selected subgroup of cases.

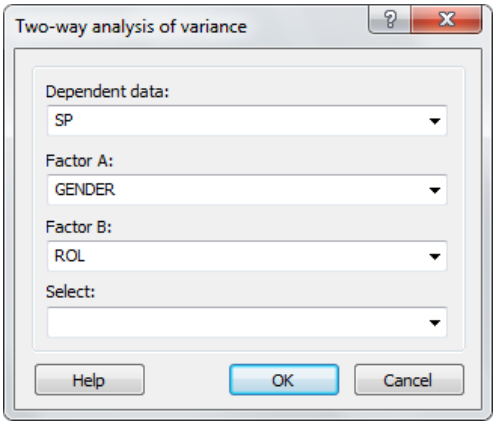

Results

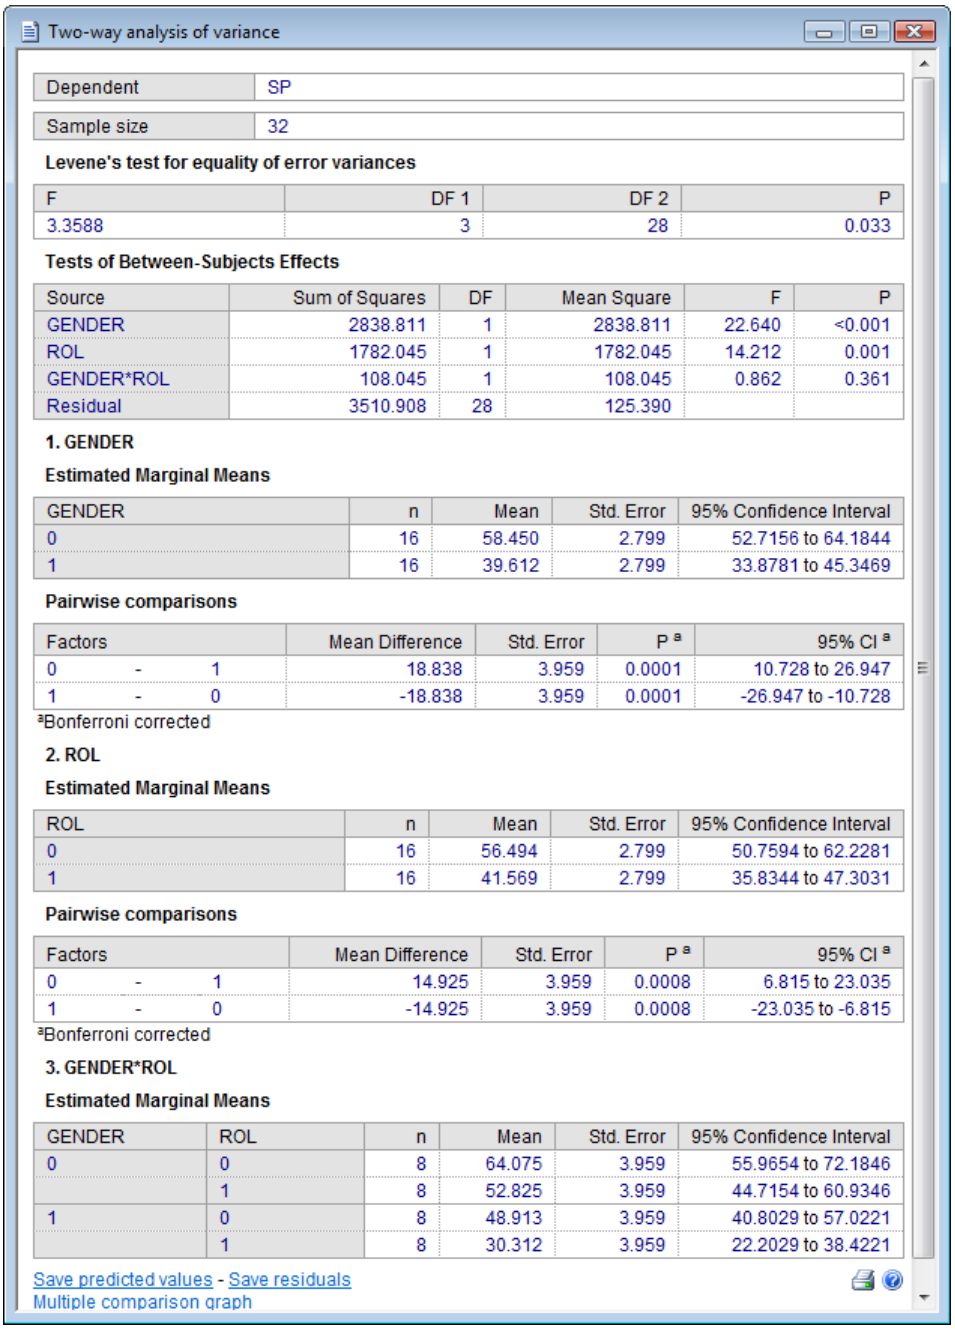

Two-way analysis of variance

Dependent	SP
Sample size	32

Levene's test for equality of error variances

F	DF 1	DF 2	P
3.3588	3	28	0.033

Tests of Between-Subjects Effects

Source	Sum of Squares	DF	Mean Square	F	P
GENDER	2838.811	1	2838.811	22.640	<0.001
ROL	1782.045	1	1782.045	14.212	0.001
GENDER*ROL	108.045	1	108.045	0.862	0.361
Residual	3510.908	28	125.390		

1. GENDER

Estimated Marginal Means

GENDER	n	Mean	Std. Error	95% Confidence Interval
0	16	58.450	2.799	52.7156 to 64.1844
1	16	39.612	2.799	33.8781 to 45.3469

Pairwise comparisons

Factors			Mean Difference	Std. Error	P [a]	95% CI [a]
0	-	1	18.838	3.959	0.0001	10.728 to 26.947
1	-	0	-18.838	3.959	0.0001	-26.947 to -10.728

[a]Bonferroni corrected

2. ROL

Estimated Marginal Means

ROL	n	Mean	Std. Error	95% Confidence Interval
0	16	56.494	2.799	50.7594 to 62.2281
1	16	41.569	2.799	35.8344 to 47.3031

Pairwise comparisons

Factors			Mean Difference	Std. Error	P [a]	95% CI [a]
0	-	1	14.925	3.959	0.0008	6.815 to 23.035
1	-	0	-14.925	3.959	0.0008	-23.035 to -6.815

[a]Bonferroni corrected

3. GENDER*ROL

Estimated Marginal Means

GENDER	ROL	n	Mean	Std. Error	95% Confidence Interval
0	0	8	64.075	3.959	55.9654 to 72.1846
	1	8	52.825	3.959	44.7154 to 60.9346
1	0	8	48.913	3.959	40.8029 to 57.0221
	1	8	30.312	3.959	22.2029 to 38.4221

Save predicted values - Save residuals
Multiple comparison graph

Levene's test for equality of variances

Prior to the ANOVA test, Levene's test for equality of variances is performed. If the Levene test is positive (P<0.05) then the variances in the groups are different (the groups are not homogeneous), and a 2-way analysis of variance is not appropriate.

Tests of Between-Subjects Effects

If the calculated P-values for the two main factors A and B, or for the 2-factor interaction is less than the conventional 0.05 (5%), then the corresponding null hypothesis is rejected, and you accept the alternative hypothesis that there is indeed a difference between groups.

When the 2-factor interaction is significant the effect of factor A is dependent on the level of factor B, and it is not recommended to interpret the means and differences between means (see below) of the main factors.

Estimated marginal means

In the following tables, the means with standard error and 95% Confidence Interval are given for all levels of the two factors. Also, differences between groups, with Standard Error, and Bonferroni corrected P-value and 95% Confidence Interval of the difference are reported.

Analysis of covariance

Description

Analysis of covariance allows to compare one variable in 2 or more groups taking into account (or to correct for) variability of other variables, called covariates.

Analysis of covariance combines one-way or two-way analysis of variance with linear regression (General Linear Model, GLM).

How to enter data

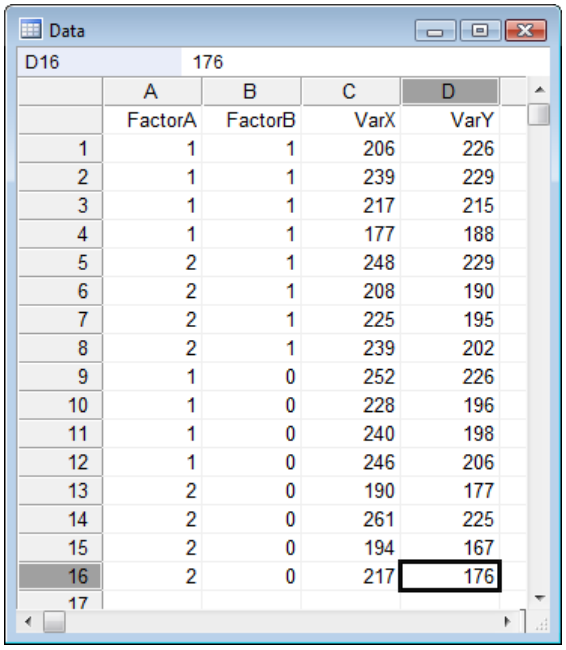

	A	B	C	D
	FactorA	FactorB	VarX	VarY
1	1	1	206	226
2	1	1	239	229
3	1	1	217	215
4	1	1	177	188
5	2	1	248	229
6	2	1	208	190
7	2	1	225	195
8	2	1	239	202
9	1	0	252	226
10	1	0	228	196
11	1	0	240	198
12	1	0	246	206
13	2	0	190	177
14	2	0	261	225
15	2	0	194	167
16	2	0	217	176
17				

In this example (data from Wildt & Ahtola, 1978) data are entered for 2 factor variables named "FactorA" and "FactorB". The variable "VarY" is the dependent variable and there is one covariate "VarX".

Required input

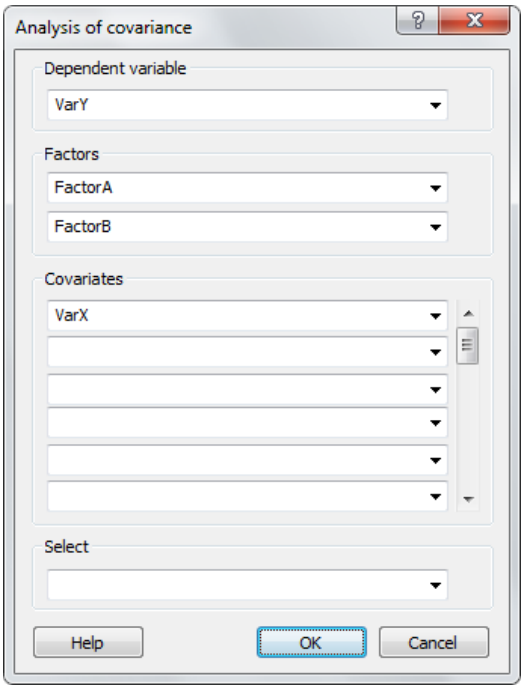

In the dialog box for ANCOVA you select:

- **Dependent variable**: the (continuous) dependent variable
- **Factors**: use one categorical variable for a one-way ANCOVA or two categorical variables for a two-way factorial ANCOVA.
- **Covariates**: one or more covariates.
- **Select**: optionally, a selection criterion in order to include only a selected subgroup of cases.

Results

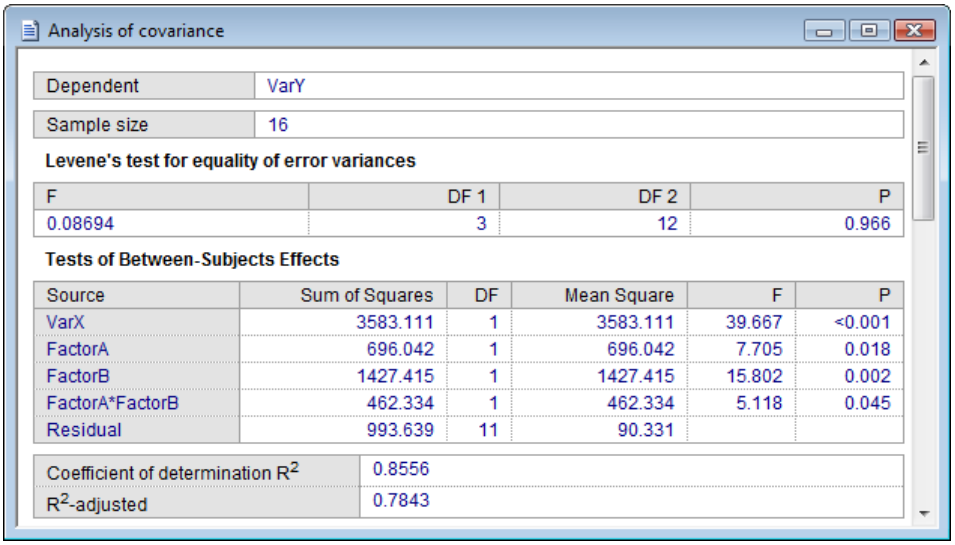

Analysis of covariance

Dependent	VarY
Sample size	16

Levene's test for equality of error variances

F	DF 1	DF 2	P
0.08694	3	12	0.966

Tests of Between-Subjects Effects

Source	Sum of Squares	DF	Mean Square	F	P
VarX	3583.111	1	3583.111	39.667	<0.001
FactorA	696.042	1	696.042	7.705	0.018
FactorB	1427.415	1	1427.415	15.802	0.002
FactorA*FactorB	462.334	1	462.334	5.118	0.045
Residual	993.639	11	90.331		

Coefficient of determination R^2	0.8556
R^2-adjusted	0.7843

Levene's test for equality of variances

Prior to the ANCOVA test, Levene's test for equality of variances is performed. If the Levene test is positive (P<0.05) then the variances in the groups are different (the groups are not homogeneous), and therefore the assumptions for ANCOVA are not met.

Tests of Between-Subjects Effects

If the calculated P-values for the two main factors A and B, or for the 2-factor interaction is less than the conventional 0.05 (5%), then the corresponding null hypothesis is rejected, and you accept the alternative hypothesis that there are indeed differences among groups.

When the 2-factor interaction (FactorA*FactorB) is significant the effect of factor A is dependent on the level of factor B, and it is not recommended to interpret the means and differences between means (see below) of the main factors.

Estimated marginal means

In the following tables, the marginal means (sometimes referred to as "corrected means") with standard error and 95% Confidence Interval are given for all levels of the two factors. Also, differences between groups, with Standard Error, and Bonferroni corrected P-value and 95% Confidence Interval of the differences are reported.

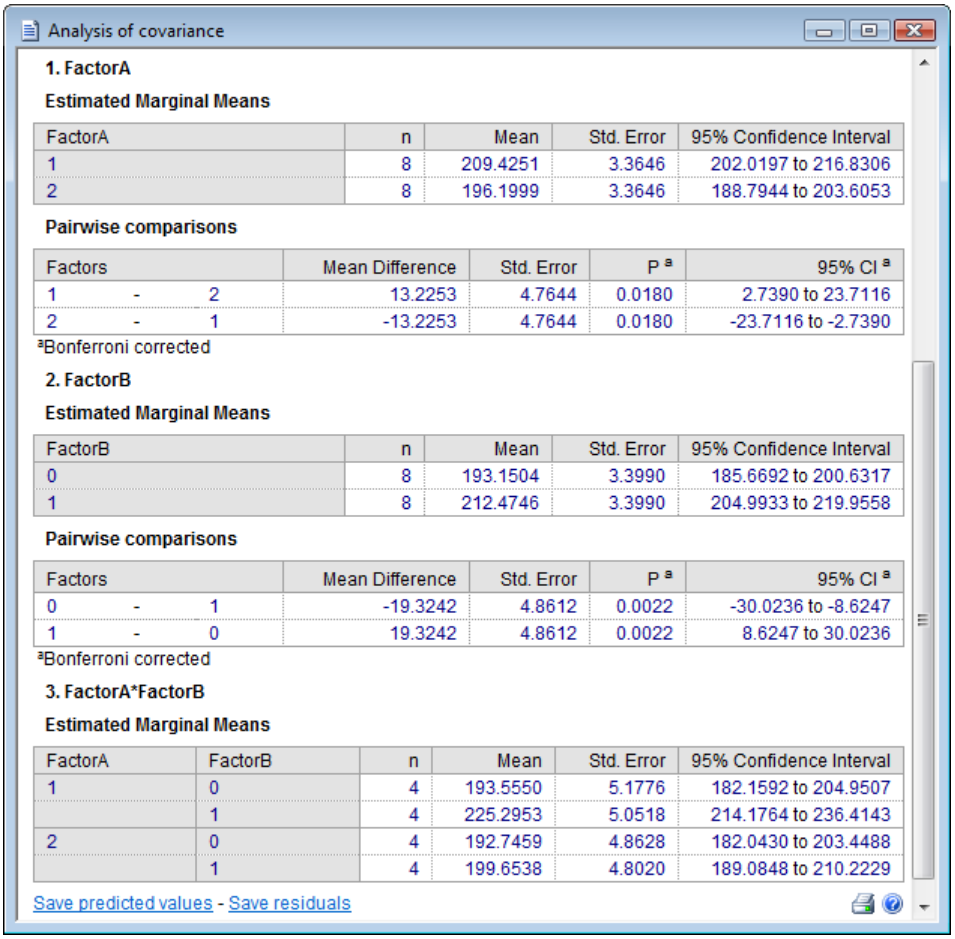

Analysis of covariance

1. FactorA

Estimated Marginal Means

FactorA	n	Mean	Std. Error	95% Confidence Interval
1	8	209.4251	3.3646	202.0197 to 216.8306
2	8	196.1999	3.3646	188.7944 to 203.6053

Pairwise comparisons

Factors			Mean Difference	Std. Error	P [a]	95% CI [a]
1	-	2	13.2253	4.7644	0.0180	2.7390 to 23.7116
2	-	1	-13.2253	4.7644	0.0180	-23.7116 to -2.7390

[a]Bonferroni corrected

2. FactorB

Estimated Marginal Means

FactorB	n	Mean	Std. Error	95% Confidence Interval
0	8	193.1504	3.3990	185.6692 to 200.6317
1	8	212.4746	3.3990	204.9933 to 219.9558

Pairwise comparisons

Factors			Mean Difference	Std. Error	P [a]	95% CI [a]
0	-	1	-19.3242	4.8612	0.0022	-30.0236 to -8.6247
1	-	0	19.3242	4.8612	0.0022	8.6247 to 30.0236

[a]Bonferroni corrected

3. FactorA*FactorB

Estimated Marginal Means

FactorA	FactorB	n	Mean	Std. Error	95% Confidence Interval
1	0	4	193.5550	5.1776	182.1592 to 204.9507
	1	4	225.2953	5.0518	214.1764 to 236.4143
2	0	4	192.7459	4.8628	182.0430 to 203.4488
	1	4	199.6538	4.8020	189.0848 to 210.2229

Save predicted values - Save residuals

General Linear Model

Since this ANCOVA procedure is an implementation of the General Linear Model (GLM), the procedure:

- reverts to one-way ANOVA when you do not specify covariates and only one factor
- reverts to a 2-way ANOVA when you specify 2 factors but no covariates
- reverts to multiple regression when you do not specify factors.

Repeated measures analysis of variance

Repeated measures analysis of variances can be used when the same parameter has been measured under different conditions on the same subjects. Subjects can be divided into different groups (Two-factor study with repeated measures on one factor) or not (Single-factor study).

A distinction is made between a **Single factor study** (without Grouping variable) or a **Two-factor study with repeated measures on one factor** (when a grouping variable is specified).

A. Single factor study

How to enter data

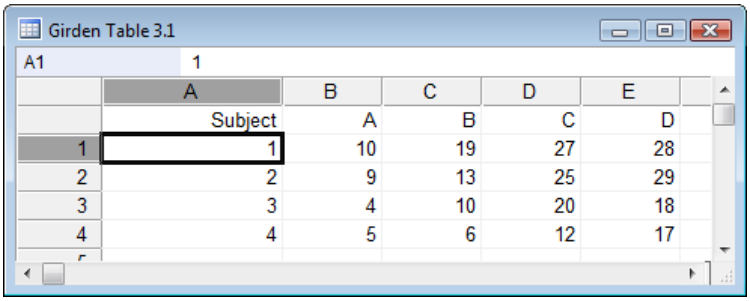

	A	B	C	D	E
	Subject	A	B	C	D
1	1	10	19	27	28
2	2	9	13	25	29
3	3	4	10	20	18
4	4	5	6	12	17

In a first column, an identification number for each cases is entered (not required). The next columns contain the data of the different measurements (example taken from Girden ER, 1992, table 3.1).

Required input

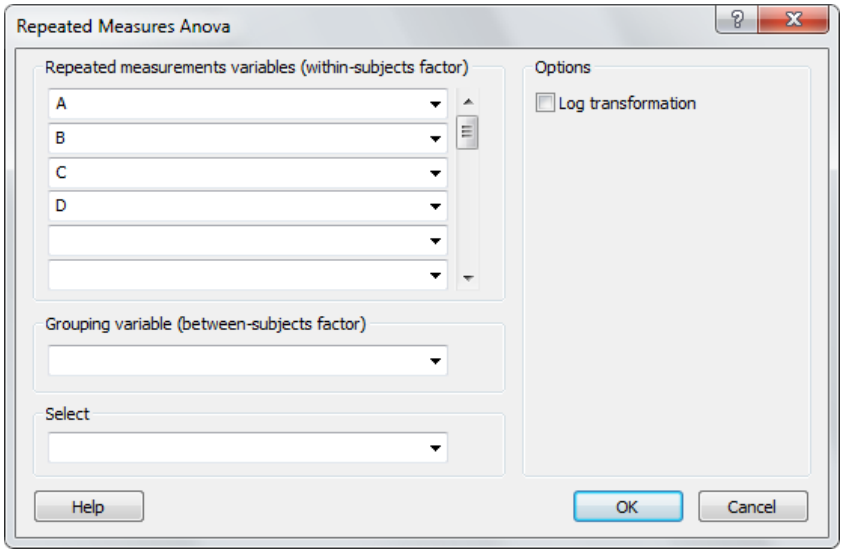

- Repeated measurements variables: the variables containing the different measurements. Note that the order in which you select the variables is important for trend analysis.
- Grouping variable: not used in a single factor study.
- Select: an optional selection criterion to include only a selected subgroup of cases.
- **Options**

  Logarithmic transformation: select this option if the data are positively skewed.

Results

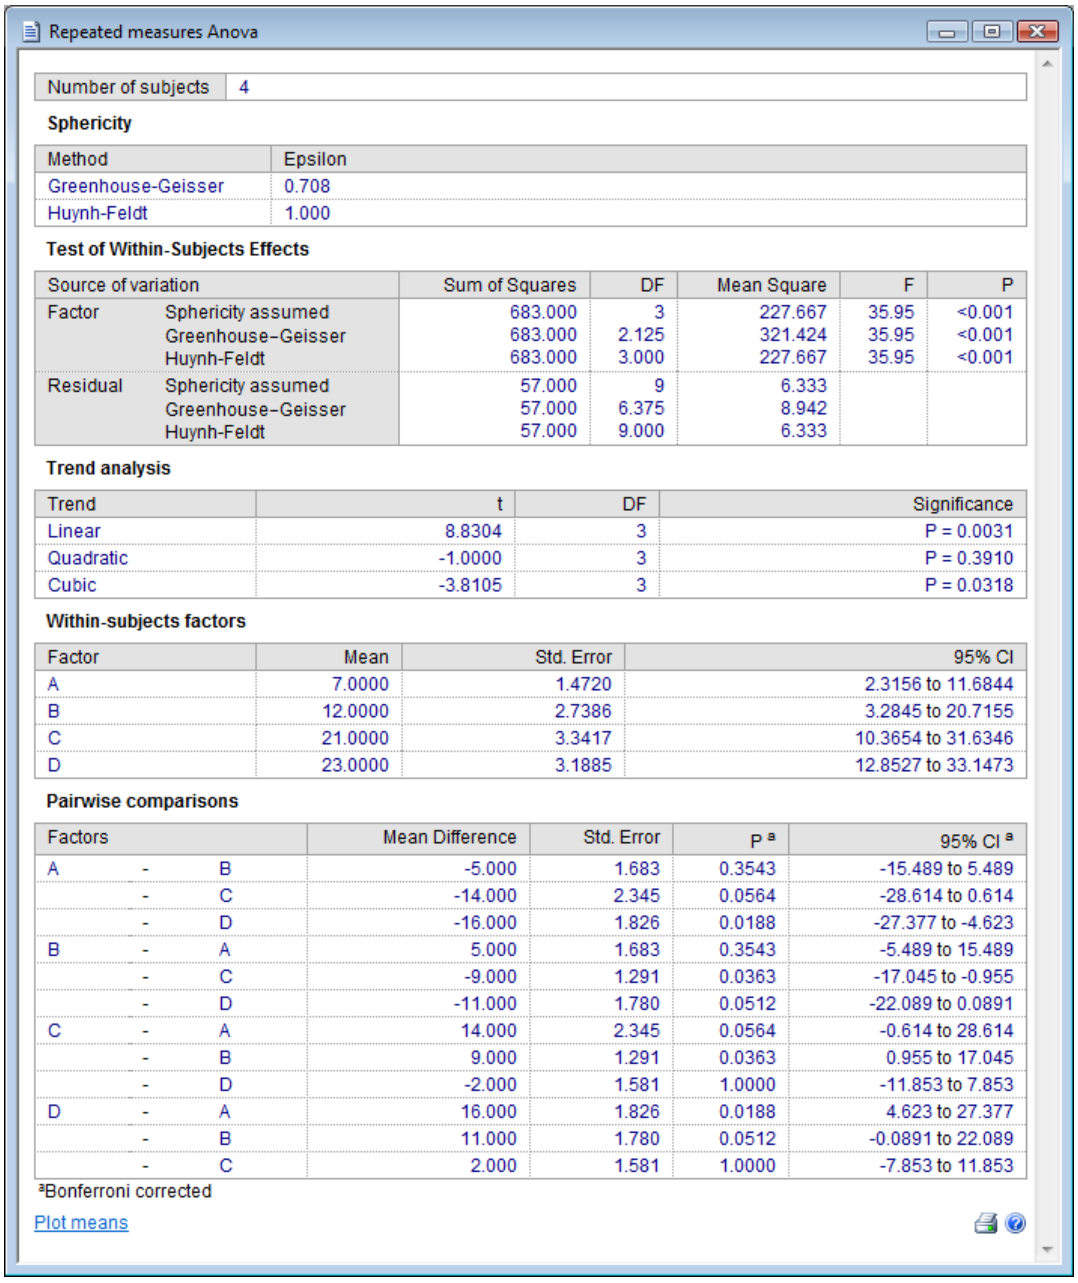

Repeated measures Anova

Number of subjects	4

Sphericity

Method	Epsilon
Greenhouse-Geisser	0.708
Huynh-Feldt	1.000

Test of Within-Subjects Effects

Source of variation		Sum of Squares	DF	Mean Square	F	P
Factor	Sphericity assumed	683.000	3	227.667	35.95	<0.001
	Greenhouse-Geisser	683.000	2.125	321.424	35.95	<0.001
	Huynh-Feldt	683.000	3.000	227.667	35.95	<0.001
Residual	Sphericity assumed	57.000	9	6.333		
	Greenhouse-Geisser	57.000	6.375	8.942		
	Huynh-Feldt	57.000	9.000	6.333		

Trend analysis

Trend	t	DF	Significance
Linear	8.8304	3	P = 0.0031
Quadratic	-1.0000	3	P = 0.3910
Cubic	-3.8105	3	P = 0.0318

Within-subjects factors

Factor	Mean	Std. Error	95% CI
A	7.0000	1.4720	2.3156 to 11.6844
B	12.0000	2.7386	3.2845 to 20.7155
C	21.0000	3.3417	10.3654 to 31.6346
D	23.0000	3.1885	12.8527 to 33.1473

Pairwise comparisons

Factors			Mean Difference	Std. Error	P [a]	95% CI [a]
A	-	B	-5.000	1.683	0.3543	-15.489 to 5.489
	-	C	-14.000	2.345	0.0564	-28.614 to 0.614
	-	D	-16.000	1.826	0.0188	-27.377 to -4.623
B	-	A	5.000	1.683	0.3543	-5.489 to 15.489
	-	C	-9.000	1.291	0.0363	-17.045 to -0.955
	-	D	-11.000	1.780	0.0512	-22.089 to 0.0891
C	-	A	14.000	2.345	0.0564	-0.614 to 28.614
	-	B	9.000	1.291	0.0363	0.955 to 17.045
	-	D	-2.000	1.581	1.0000	-11.853 to 7.853
D	-	A	16.000	1.826	0.0188	4.623 to 27.377
	-	B	11.000	1.780	0.0512	-0.0891 to 22.089
	-	C	2.000	1.581	1.0000	-7.853 to 11.853

[a]Bonferroni corrected

Plot means

The results window displays the number of subjects in the study. Note that subjects with missing values for any measurement are dropped from the analysis.

Sphericity

Sphericity refers to the equality of variances of the differences between measurements, which is an assumption of ANOVA with a repeated measures factor.

MedCalc reports the estimates (epsilon) of sphericity proposed by Greenhouse and Geisser (1958) and Huynh and Feldt (1976) (corrected by Lecoutre, 1991). The closer that epsilon is to 1, the more homogeneous are the variances of differences, and hence the closer the data are to being spherical. Both the Greenhouse-Geisser and Huynh-Feldt estimates are used as a correction factor that is applied to the degrees of freedom used to calculate the P-value for the observed value of F.

Test of Within Subjects Effects

In this table, the variation attributed to "Factor" and "Residual" variation is displayed. If the P-value next to "Factor" is low (P<0.05) it can be concluded that there is significant difference between the different measurements.

MedCalc produces two corrections based upon the estimates of sphericity by Greenhouse and Geisser (1958) and Huynh and Feldt (1976) (corrected by Lecoutre, 1991). Girden (1992) recommends that when epsilon (Greenhouse-Geisser estimate) > 0.75 then the correction according to Huynh and Feldt should be used. If epsilon < 0.75 then the more conservative correction according to Greenhouse-Geisser is preferred.

Trend analysis

The Trend analysis table shows whether the measurements show a linear or non-linear (quadratic, cubic) trend.

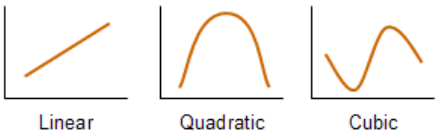

Within-subjects factors

The within-subjects factors are summarized in a table with Mean, Standard Error and 95% Confidence Interval.

Pairwise comparisons

In the Pairwise comparisons table, the different measurements are compared to each other. The mean difference with standard error, P-value, and 95% Confidence Interval of the difference is given. Bonferroni correction for multiple comparisons is applied for P-values and confidence intervals.

B. Two-factor study with repeated measures on one factor

How to enter data

Girden Table 5.1

A1: 0

	A	B	C	D	E
	Gender	Family	Luck	Internal	External
1	0	48	31	24	29
2	0	46	34	25	31
3	0	45	37	27	35
4	0	37	39	30	38
5	0	45	34	29	32
6	1	28	17	19	20
7	1	32	18	20	22
8	1	35	19	22	24
9	1	39	16	19	22
10	1	31	15	20	22

In this example the first column indicates group membership. "Male" has been coded as 0, "female" as 1. The next columns contain the data of the different measurements (example taken from Girden ER, 1992, table 5.1).

Required input

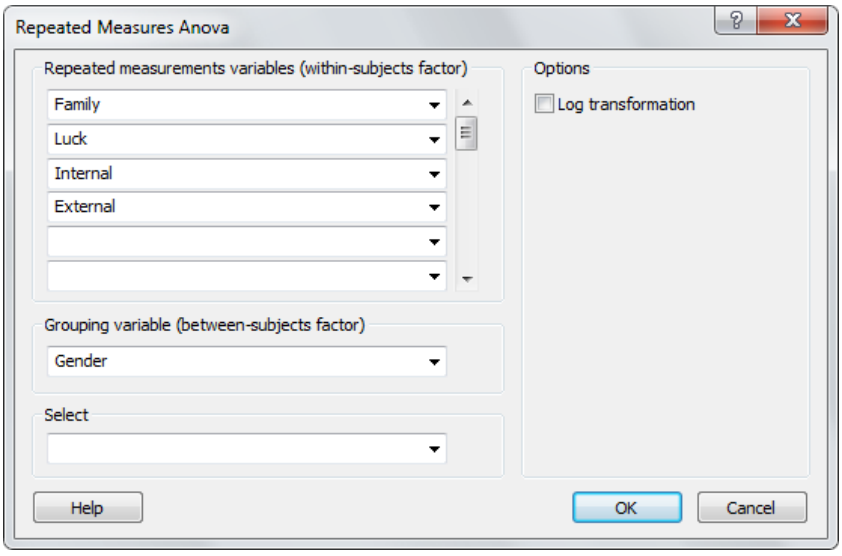

- Repeated measurements variables: the variables containing the different measurements.
- Grouping variable: a categorical variable that divides the data into groups (between-subjects factor).
- Select: an optional selection criterion to include only a selected subgroup of cases.
- **Options**

  Logarithmic transformation: select this option if the data are positively skewed.

Results

Repeated measures Anova

Number of subjects	10

Between-subjects factors (subject groups)

Gender	n
0	5
1	5
Total	10

Sphericity

Method	Epsilon
Greenhouse-Geisser	0.391
Huynh-Feldt	0.475

Test of Between-Subjects Effects

Source of variation	Sum of Squares	DF	Mean Square	F	P
Groups (Gender)	1392.400	1	1392.400	165.76	<0.001
Residual	67.200	8	8.400		

Test of Within-Subjects Effects

Source of variation		Sum of Squares	DF	Mean Square	F	P
Factor	Sphericity assumed	1336.200	3	445.400	50.52	<0.001
	Greenhouse-Geisser	1336.200	1.174	1138.414	50.52	<0.001
	Huynh-Feldt	1336.200	1.426	936.725	50.52	<0.001
Group x Factor interaction	Sphericity assumed	156.200	3	52.067	5.91	0.004
	Greenhouse-Geisser	156.200	1.174	133.079	5.91	0.033
	Huynh-Feldt	156.200	1.426	109.502	5.91	0.024
Residual	Sphericity assumed	211.600	24	8.817		
	Greenhouse-Geisser	211.600	9.390	22.535		
	Huynh-Feldt	211.600	11.412	18.542		

Plot means
Plot means clustered by groups

Between-subjects factors (subject groups)

The first table lists the different subject groups and the number of observations.

Sphericity

Sphericity refers to the equality of variances of the differences between measurements, which is an assumption of ANOVA with a repeated measures factor.

MedCalc reports the estimates (epsilon) of sphericity proposed by Greenhouse and Geisser (1958) and Huynh and Feldt (1976) (corrected by Lecoutre, 1991). The closer that epsilon is to 1, the more homogeneous are the variances of differences, and hence the closer the data are to being spherical. Both the Greenhouse-Geisser and Huynh-Feldt estimates are used as a correction factor that is applied to the degrees of freedom used to calculate the P-value for the observed value of F.

Test of Between-Subjects Effects

In this table, the variation attributed to "Groups" (between-subjects) and "Residual" variation are displayed.

- If the P-value for "Groups" is low (P<0.05) it can be concluded that there is significant difference between groups.

Test of Within-Subjects Effects

In this table, the source of variation attributed to "Factor" (within-subjects), "Group" and "Factor" interaction, "Residual" variation is displayed.

- If the P-value for "Factor" is low (P<0.05) it can be concluded that there is significant difference between measurements.
- If the P-value for "Group x Factor interaction" is low (P<0.05) it can be concluded that the difference between measurements depends on group membership.

MedCalc produces two corrections based upon the estimates of sphericity by Greenhouse and Geisser (1958) and Huynh and Feldt (1976) (corrected by Lecoutre, 1991). Girden (1992) recommends that when epsilon (Greenhouse-Geisser estimate) > 0.75 then the correction according to Huynh and Feldt should be used. If epsilon < 0.75 then the more conservative correction according to Greenhouse-Geisser is preferred.

Log transformation

If you selected the Log transformation option, the program performs the calculations on the logarithms of all measurements, but backtransforms the results to the original scale for presentation.

- In the Within-Subjects factors table, the geometric mean with its 95% Confidence is given.
- In the Pairwise comparison table, the geometric mean of the ratios of paired observations is given (which is the backtransformed mean difference of the logarithms of the paired observations).

Kruskal-Wallis test

The *Kruskal-Wallis test* (*H*-test) is an extension of the *Wilcoxon* test and can be used to test the hypothesis that a number of unpaired samples originate from the same population. In MedCalc, *Factor codes* are used to break-up the (ordinal) *data* in one variable into different sample subgroups. If the null-hypothesis, being the hypothesis that the samples originate from the same population, is rejected (P<0.05), then the conclusion is that there is a statistically significant difference between at least two of the subgroups.

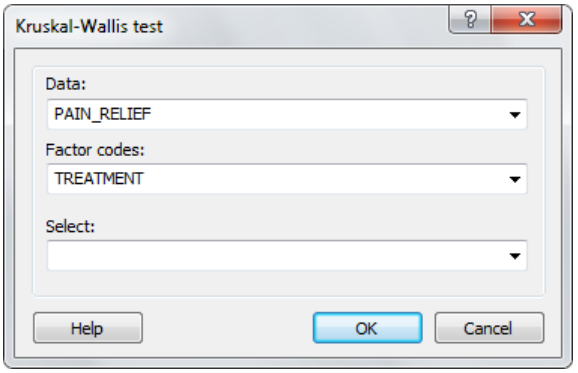

The following need to be entered in the dialog box: for *Data* select the variable containing the data, and for *Factor codes* the qualitative factor. The qualitative factor may either be character or numeric codes. These are the codes that will be used to break-up the data into several subgroups.

You can click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name. You obtain the results after you click OK.

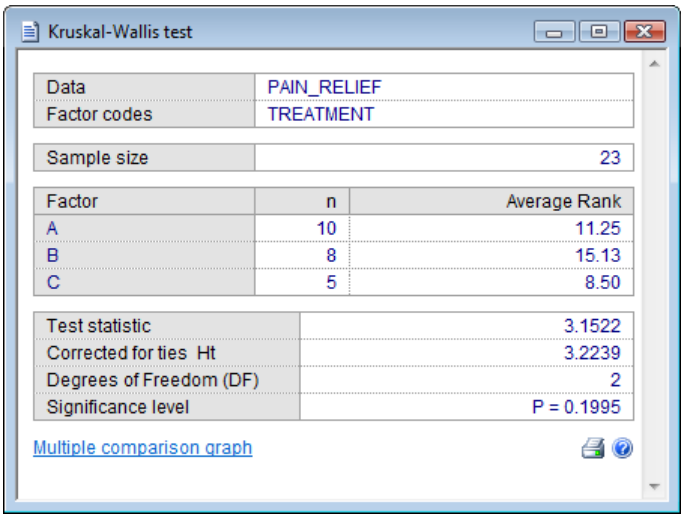

In this example, it is tested whether there is a difference of the variable *Pain_relief* for the different treatment regimens coded A, B and C in the variable *Treatment*. Pain relief was recorded on an ordinal scale from 0 to 9. Since the null-hypothesis is not rejected (P=0.1995), the conclusion is that there is no statistical significant difference between treatments.

For a graphical representation of this test, refer to *Multiple comparison graphs* (p. 172).

Post-hoc analysis

If the Kruskal-Wallis test is positive (P<0.05) then MedCalc performs a test for pairwise comparison of subgroups according to Conover, 1999.

Literature

- Conover WJ (1999) Practical nonparametric statistics, 3rd edition. New York: John Wiley & Sons.

Friedman test

The Friedman test is a non-parametric test for testing the difference between several related samples. It is an alternative for Repeated measures analysis of variances which is used when the same parameter has been measured under different conditions on the same subjects.

How to enter data

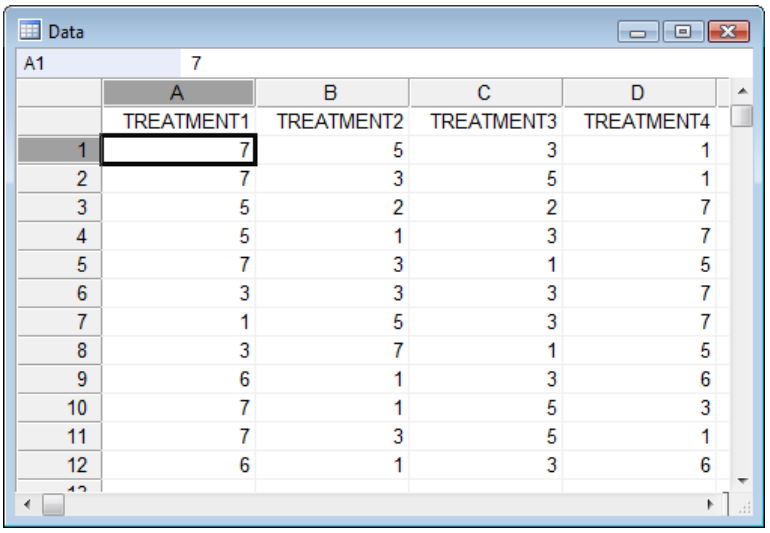

	A	B	C	D
	TREATMENT1	TREATMENT2	TREATMENT3	TREATMENT4
1	7	5	3	1
2	7	3	5	1
3	5	2	2	7
4	5	1	3	7
5	7	3	1	5
6	3	3	3	7
7	1	5	3	7
8	3	7	1	5
9	6	1	3	6
10	7	1	5	3
11	7	3	5	1
12	6	1	3	6

The columns contain the data of the different measurements (example adapted from Conover, 1999).

Required input

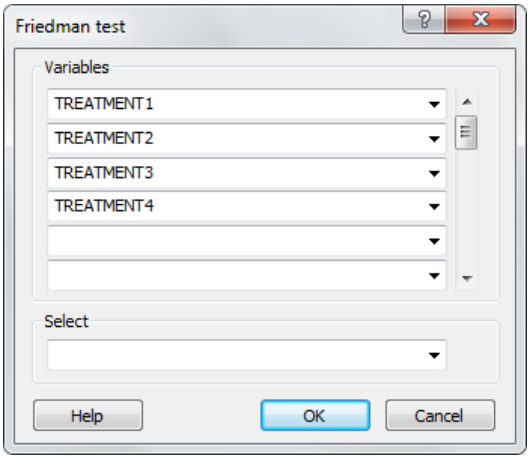

- **Variables**: the variables that contain the related observations.
- **Select**: an optional selection criterion to include only a selected subgroup of cases.

Results

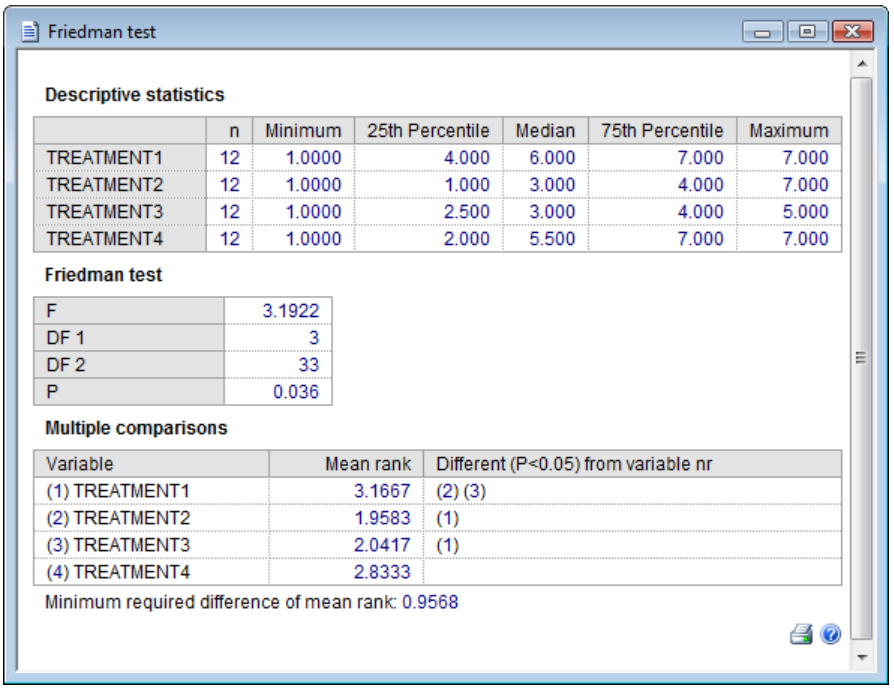

Friedman test

Descriptive statistics

	n	Minimum	25th Percentile	Median	75th Percentile	Maximum
TREATMENT1	12	1.0000	4.000	6.000	7.000	7.000
TREATMENT2	12	1.0000	1.000	3.000	4.000	7.000
TREATMENT3	12	1.0000	2.500	3.000	4.000	5.000
TREATMENT4	12	1.0000	2.000	5.500	7.000	7.000

Friedman test

F	3.1922
DF 1	3
DF 2	33
P	0.036

Multiple comparisons

Variable	Mean rank	Different (P<0.05) from variable nr
(1) TREATMENT1	3.1667	(2) (3)
(2) TREATMENT2	1.9583	(1)
(3) TREATMENT3	2.0417	(1)
(4) TREATMENT4	2.8333	

Minimum required difference of mean rank: 0.9568

Descriptive statistics

This table gives the descriptive statistics for the different variables: number of cases (n), minimum, 25th percentile, median, 75th percentile and maximum. Since the Friedman test is for related samples, cases with missing observations for one or more of the variables are excluded from the analysis, and the sample size is the same for each variable.

Friedman test

The null hypothesis for the Friedman test is that there are no differences between the variables. If the calculated probability is low (P<0.05) the null-hypothesis is rejected and it can be concluded that at least 2 of the variables are significantly different from each other.

Multiple comparisons

When the Friedman test is positive (P<0.05) then a table is displayed showing which of the variables is significantly different from which other variables.

In the example variable (1), which is TREATMENT1, is significantly different from the variables (2) and (3), which correspond to TREATMENT2 and TREATMENT3.

Literature

- Conover WJ (1999) Practical nonparametric statistics, 3rd edition. New York: John Wiley & Sons.

Frequency table & Chi-square test

The *Frequency table* procedure can be used for the following:

- to test the hypothesis that for one classification table (e.g. gender), all classification levels have the same frequency.
- to test the relationship between two classification factors (e.g. gender and profession).

In the *Frequency table* dialog box, one or two discrete variables with the classification data must be identified. Classification data may either be numeric or alphanumeric (string) values. If required, you can convert a continuous variable into a discrete variable using the CATEGORISE function (see p. 219) or IF function (see p. 220).

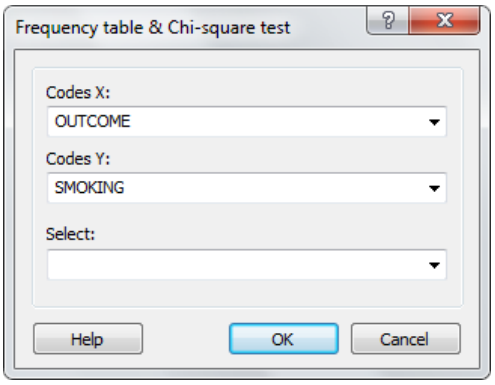

After you have completed the dialog box, click the OK button, or press the *Enter* key to obtain the frequency table with the relevant statistics.

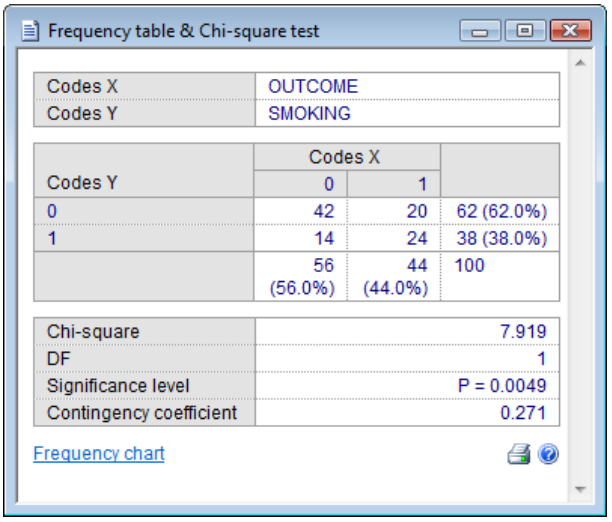

Chi-square test

- When you want to test the hypothesis that for *one single classification table* (e.g. gender), all classification levels have the same frequency, then identify only one discrete variable in the dialog form. In this case the null hypothesis is that all classification levels have the same frequency. If the calculated P-value is low (P<0.05), then you reject the null hypothesis and the alternative hypothesis that there is a significant difference between the frequencies of the different classification levels must be accepted.

  In a single classification table the *mode* of the observations is the most common observation or category (the observation with the highest frequency). A unimodal distribution has one mode; a bimodal distribution, two modes.

- When you want to study the relationship between *two classification factors* (e.g. gender and profession), then identify the two discrete variables in the dialog form. In this case the null hypothesis is that the two factors are independent. If the calculated P-value is low (P<0.05), then the null hypothesis is rejected and you accept the alternative hypothesis that there is a relation between the two factors.

Note that when the *degrees of freedom* equals 1, e.g. in case of a 2x2 table, MedCalc uses *Yates' correction for continuity*.

Chi-square test for trend

If the table has two columns and three or more rows (or two rows and three or more columns), and the categories can be quantified, MedCalc will also perform a *Chi-square test for trend*. The Cochran-Armitage test for trend (Armitage, 1955) tests whether there is a *linear* trend between row (or column) number and the fraction of subjects in the left column (or top row). The Cochran-Armitage test for trend provides a more powerful test than the unordered independence test above.

If there is no meaningful order in the row (or column) categories, then you should ignore this calculation.

Analysis of 2x2 table

If the classification table is a 2x2 table then you should **not** use the *Frequency table* and *Chi-square test* procedure when:

- the number of expected frequencies in the 2x2 table is low (in case the total number of observations is less than 20), the table should be tested using *Fisher's exact test* (see below);
- the two classification factors are not independent, or when you want to test the difference between proportions in related or paired observations (e.g. in studies in which patients serve as their own control), you must use the *McNemar test* (p. 111).

Fisher's exact test

If you have a 2x2 frequency table with small numbers of expected frequencies (in case the total number of observations is less than 20), you should not perform the *Chi-square test* as described above, but you should use *Fisher's exact test*.

In the *Fisher's exact test* dialog box, two discrete dichotomous variables with the classification data must be identified. Classification data may either be numeric or alphanumeric (string) values. If required, you can convert a continuous variable into a dichotomous variable using the CATEGORISE function (see p. 219) or IF function (see p. 220).

For example: in a study including 20 patients, 9 women and 11 men, the success of a treatment is recorded (1 = successful, 0 = no success). Is there a difference between the success rate in men and women?

The data are entered as follows in the spreadsheet:

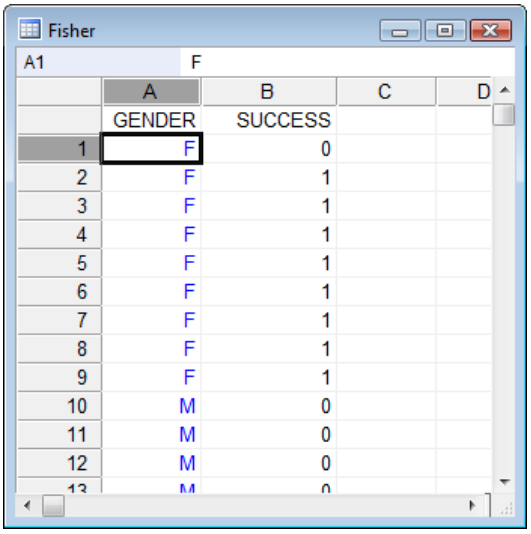

The dialog box for the Fisher's exact test is completed as follows:

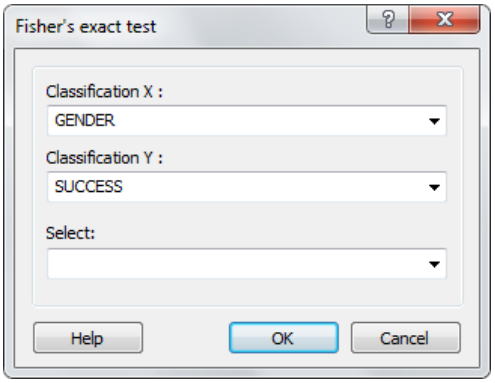

After you have completed the dialog box, click the OK button, or press the *Enter* key to obtain the frequency table with the relevant statistics:

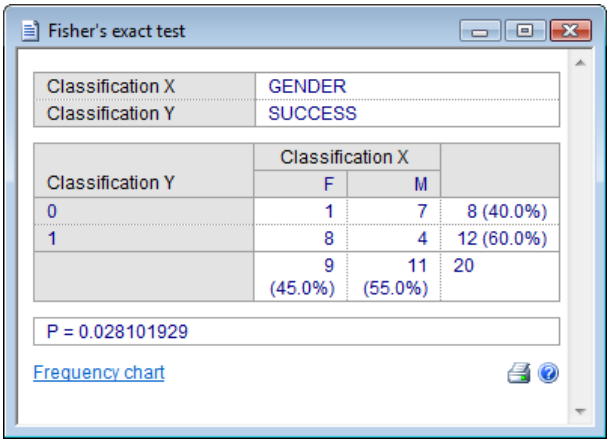

Classification table

The program displays the 2x2 classification table.

P-value

When the (two-sided) P-value (the probability of obtaining the observed result or a more extreme result) is less than the conventional 0.05, the conclusion is that there is a significant relationship between the two classification factors.

In the example P=0.028 and the conclusion therefore is that the success rate in men and women differs (or that the success rate is related to gender).

McNemar test

The *McNemar test* is a test on a 2x2 classification table when the two classification factors are dependent, or when you want to test the difference between paired proportions, e.g. in studies in which patients serve as their own control, or in studies with "before and after" design.

In the *McNemar test* dialog box, two discrete dichotomous variables with related classification data must be identified. Classification data may either be numeric or alphanumeric (string) values. If required, you can convert a continuous variable into a dichotomous variable using the CATEGORISE function (see p. 219) or IF function (see p. 220).

The variables together cannot contain more than 2 different classification values.

For example, in a study a test is performed before treatment and after treatment in 20 patients. The results of the test are coded 0 and 1. Is there a significant change in the test result before and after treatment?

The data are entered as follows in the spreadsheet:

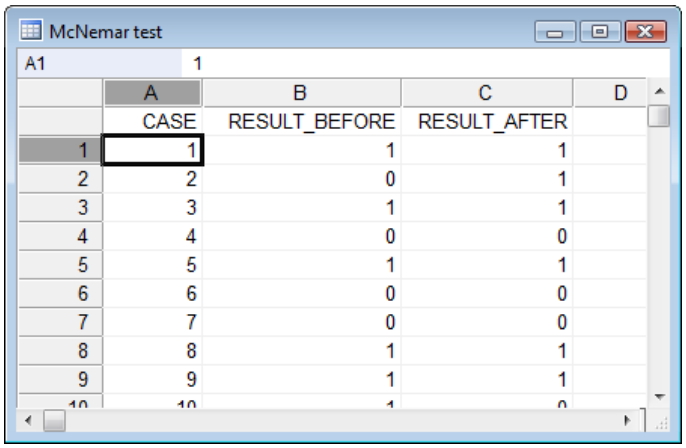

The dialog box for the McNemar test is completed as follows:

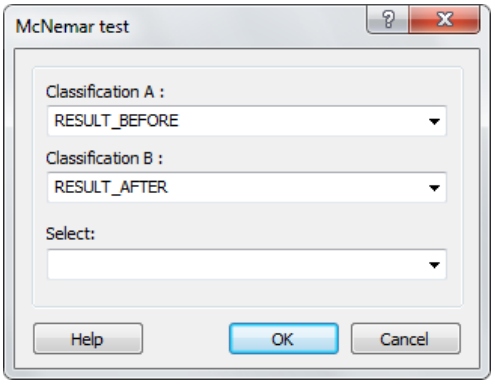

After you have completed the dialog box, click OK, or press the *Enter* key to obtain the classification table with the relevant statistics.

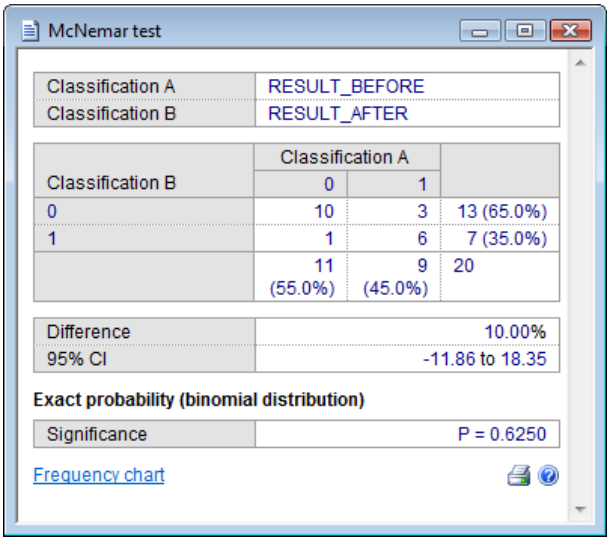

Classification table

The program displays the 2x2 classification table.

Difference and P-value

The program gives the difference between the proportions (expressed as a percentage) with 95% confidence interval.

When the (two-sided) P-value is less than the conventional 0.05, the conclusion is that there is a significant difference between the two proportions.

In the example, the difference before and after treatment is 10% with 95% CI from -11.86% to 18.35%, which is not significant (P=0.625, n =20).

Note

If the number of discordant pairs (3 + 1 in the example) is less than or equal to 25, then the two-sided P-value is based on the cumulative binomial distribution. If the number of discordant pairs is more than 25 then a chi-square approximation with Yates' correction for continuity is used.

The 95% confidence interval is calculated according to Bland, 2000.

Cochran's Q test

Cochran's Q test (Sheskin, 2004) is an extension to the McNemar test for related samples that provides a method for testing for differences between three or more matched sets of frequencies or proportions.

Example: 12 subjects are asked to perform 3 tasks. The outcome of each task is a dichotomous value, success or failure.

How to enter the data in the spreadsheet

The results are coded 0 for failure and 1 for success. In the example, subject 1 was successful in task 2 but failed tasks 1 and 3.

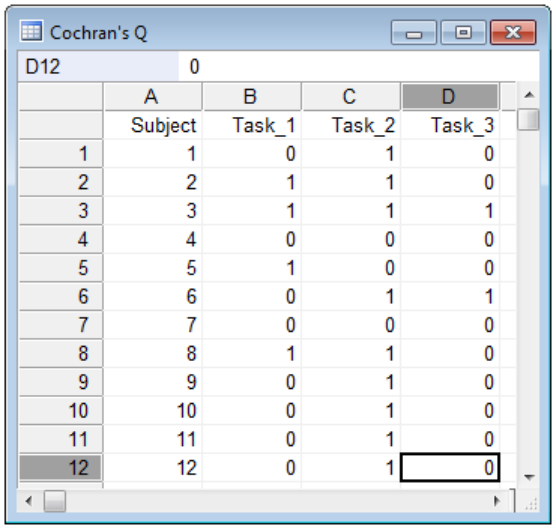

	A	B	C	D
	Subject	Task_1	Task_2	Task_3
1	1	0	1	0
2	2	1	1	0
3	3	1	1	1
4	4	0	0	0
5	5	1	0	0
6	6	0	1	1
7	7	0	0	0
8	8	1	1	0
9	9	0	1	0
10	10	0	1	0
11	11	0	1	0
12	12	0	1	0

Required input

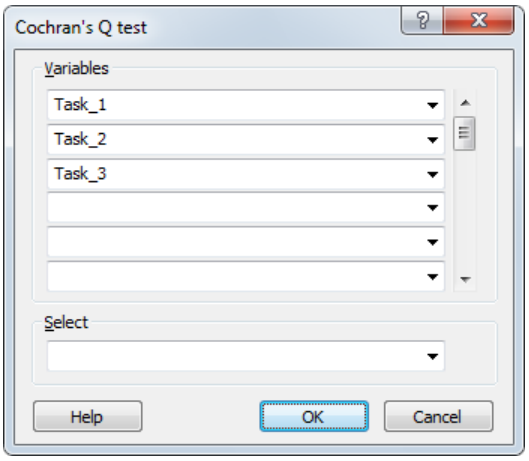

- **Variables**: the variables that contain the related observations. Data must be coded 0 to represent failure (or absence) and 1 to represent success (or presence).
- **Select**: an optional selection criterion to include only a selected subgroup of cases.

Results

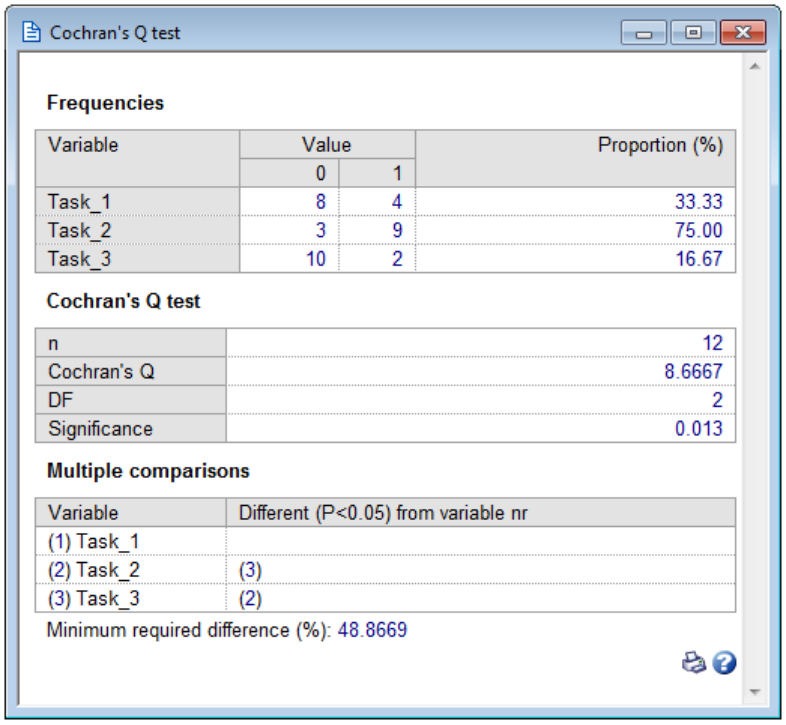

Cochran's Q test

Frequencies

Variable	Value		Proportion (%)
	0	1	
Task_1	8	4	33.33
Task_2	3	9	75.00
Task_3	10	2	16.67

Cochran's Q test

n	12
Cochran's Q	8.6667
DF	2
Significance	0.013

Multiple comparisons

Variable	Different (P<0.05) from variable nr
(1) Task_1	
(2) Task_2	(3)
(3) Task_3	(2)

Minimum required difference (%): 48.8669

Frequencies

This table gives the frequencies of the values coded 0 (meaning absence or failure) and 1 (meaning presence or success) in the different variables, the proportion (expressed as a percentage) of values coded 1

Since Cochran's Q test is for related samples, cases with missing observations for one or more of the variables are excluded from the analysis, and the number of cases is the same for each variable.

Cochran's Q test

The null hypothesis for the Cochran's Q test is that there are no differences between the variables. If the calculated probability is low (P<0.05) the null-hypothesis is rejected and it can be concluded that the proportions in at least 2 of the variables are significantly different from each other.

Multiple comparisons

When the Cochran's Q test is positive (P<0.05) then a minimum required difference for a significant difference between two proportions is calculated (Sheskin, 2004) and a table is displayed showing which of the proportions are significantly different from which other proportions.

Frequencies bar chart

Using the *Frequencies bar chart* command you can graph categorical data. In the *Frequencies bar chart* dialog box, one or two discrete variables with the classification data must be identified. Classification data may either be numeric or alphanumeric (string) values.

Graph types

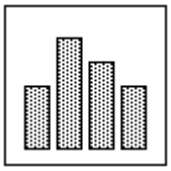

Simple column chart (one classification factor). The chart contains a single bar for each category. The height of the bars is the number of cases in the category.

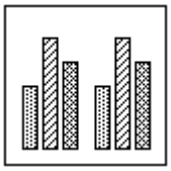

Clustered column (two classification factors). Like simple column chart, but containing a group of bars for each category in the first classification category. The height of the bars is the number of cases in the category.

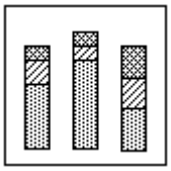

Stacked column (two classification factors). Bar segments are stacked on top of one another. There is one bar stack for each category in the first classification factor. Segments within each stack represent the contribution of categories in the second classification factor.

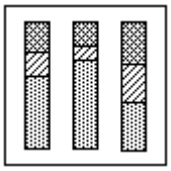

100% Stacked column (two classification factors). Bar segments are stacked on top of one another, the total equals 100%. There is one bar stack for each category in the first classification factor. Segments within each stack represent the relative contribution of categories in the second classification factor.

Kaplan-Meier survival curve

In clinical trials the investigator is often interested in the time until participants in a study present a specific event or endpoint. This event usually is a clinical outcome such as death, disappearance of a tumor, etc.

The participants will be followed beginning at a certain starting-point, and the time will be recorded needed for the event of interest to occur.

Usually, the end of the study is reached before all participants have presented this event, and the outcome of the remaining patients is unknown. Also the outcome is unknown of those participants who have withdrawn from the study. For all these cases the time of follow-up is recorded (censored data).

In MedCalc, these data can be analyzed by means of a *life-table*, or *Kaplan-Meier curve*, which is the most common method to describe survival characteristics.

In order to be able to analyze the data, you need to enter the data in the spreadsheet as follows:

- in one column, a code can be entered to assign the case to a particular group (study group - control group).
- in a second column, the survival time has to recorded

- in a third column, it must be recorded whether or not the case has reached the endpoint (by entering the code 1) or whether the time is censored, i.e. the outcome is unknown (by entering the code 0);

The order of these columns is of course not important. Also, the rows do not have to be sorted in any way.

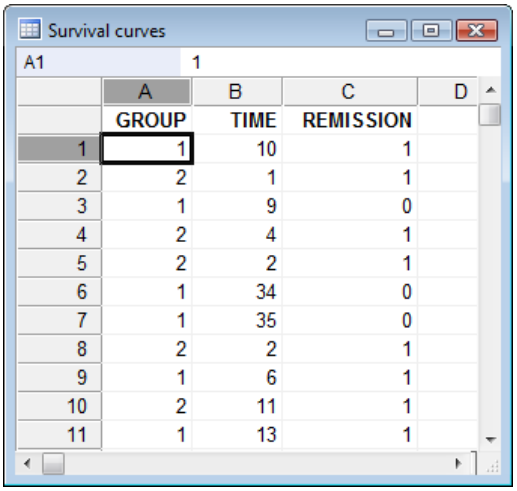

	A	B	C	D
	GROUP	TIME	REMISSION	
1	1	10	1	
2	2	1	1	
3	1	9	0	
4	2	4	1	
5	2	2	1	
6	1	34	0	
7	1	35	0	
8	2	2	1	
9	1	6	1	
10	2	11	1	
11	1	13	1	

The case in row 1 belonged to group 1, and reached the endpoint after 10 units of time. The case in row 3 also belonged to group 1 and was followed for 9 units of time. The outcome of this case is unknown (withdrawn from study, or end of study) (data from Freireich et al., Blood 1963; 21:699-716).

From these data, MedCalc can easily calculate and construct the Kaplan-Meier curve. After you have selected the *Kaplan-Meier survival curve* option in the *Graphs* menu, the following dialog box is displayed:

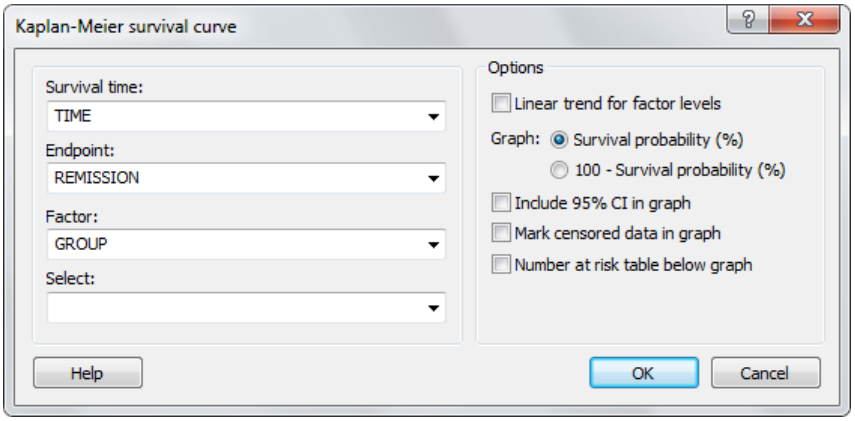

In this dialog box the following data need to be entered:

- **Survival time**

  The name of the variable containing the time to reach the event of interest, or the time of follow-up.

- **Endpoint**

  The name of a variable containing codes 1 for the cases that reached the endpoint, or code 0 for the cases that have not reached the endpoint, either because they withdrew from the study, or the end of the study was reached.

- **Factor**

  For *Factor* select a qualitative or discrete variable (grouping variable - GROUP in the example). This qualitative factor may either be character or numeric codes. These codes are used to break-up the data into several subgroups. If you want to study the effect of a continuous variable on survival time, you can convert this continuous variable into a discrete variable using the CATEGORISE function (see p. 219) or IF function (see p. 220).

  MedCalc will allow comparison of survival curves for up to 6 subgroups.

  If no *Factor* variable is selected, then MedCalc will display only one survival curve (all data are considered to belong to one group).

- **Select**

  An optional selection criterion to include only a selected subgroup of cases in the analysis.

- **Options**

  Linear trend for factor levels: Allows testing for a linear trend across levels of the factor. It is appropriate if factor levels have a natural ordering (for example, factor codes represent doses applied to different groups). Kaplan-Meier assumes that the factor levels are equally spaced.

  Graph:

  Survival probability (%): plot Survival probability (%) against time (descending curves)

  100 - Survival probability (%): plot 100 - Survival probability (%) against time (ascending curves)

  Include 95% CI in graph: Option to plot the 95% confidence interval for the survival curves.

  Mark censored data in graph: Mark censored data in the graph with a small vertical line.

  Number at risk table below graph: Shows a table below the graph with the number of subjects at risk.

When all data have been entered click the OK button, and the program will open 2 windows: one with the survival graphs and one with the results in a text window.

Graph

The survival curves are drawn as a step function, as shown in the following example:

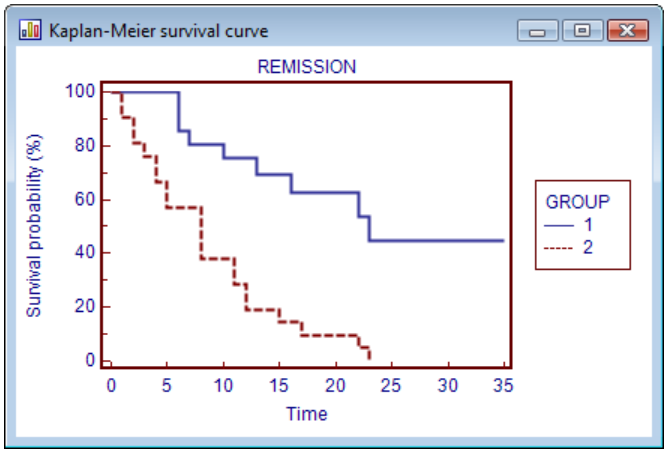

With the option “Include 95% CI in graph” selected, the graph looks like this:

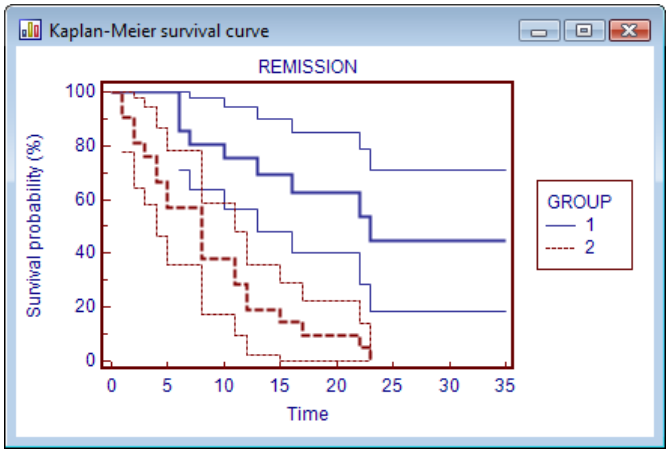

When the option "Number at risk table below graph" is selected, the result is:

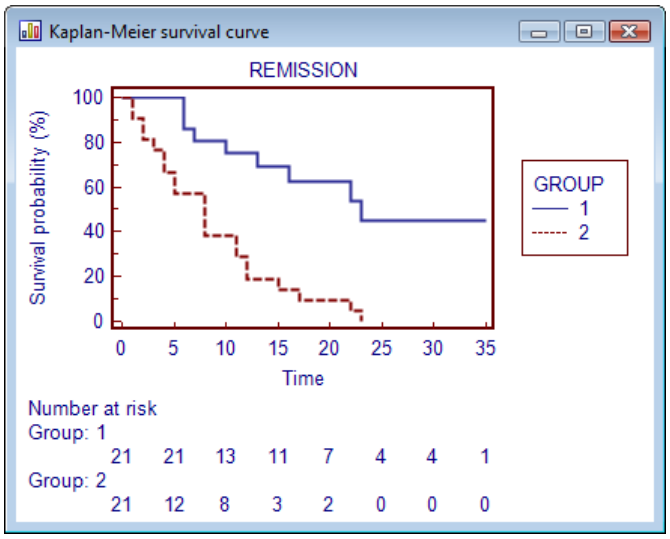

Results

Sample size, median survival time and survival proportion

In this example, the sample size was 21 cases in both groups. The median survival time was 23 units in group 1, and 8 units group 2. The median survival time is the time at which half the subjects have reached the event of interest. If the survival curve does not fall to 0.5 (50%) then the median time cannot be computed.

Next the survival proportions (with standard error) are listed for all groups for each observed time point.

Kaplan-Meier survival curve

Survival time	TIME
Endpoint	REMISSION
Factor codes	GROUP

	Factors			
	1		2	
Sample size	21		21	
Median survival	23		8	
Survival time	Survival Proportion	Standard Error	Survival Proportion	Standard Error
1	-	-	0.905	0.0641
2	-	-	0.810	0.0857
3	-	-	0.762	0.0929
4	-	-	0.667	0.103
5	-	-	0.571	0.108
6	0.857	0.0764	-	-
7	0.807	0.0869	-	-
8	-	-	0.381	0.106
9	-	-	-	-
10	0.753	0.0963	-	-
11	-	-	0.286	0.0986
12	-	-	0.190	0.0857
13	0.690	0.107	-	-
15	-	-	0.143	0.0764
16	0.627	0.114	-	-

Comparison of survival curves

When you scroll down (press the *Page Down* key), you see the result of the logrank test for the comparison between the two survival curves: 9 cases in group 1 and 21 cases in group 2 presented the outcome of interest, the chi-square statistic was 16.79 with associated P-value of less than 0.0001. The

conclusion therefore is that, statistically, the two survival curves differ significantly, or that the grouping variable has a significant influence on survival time.

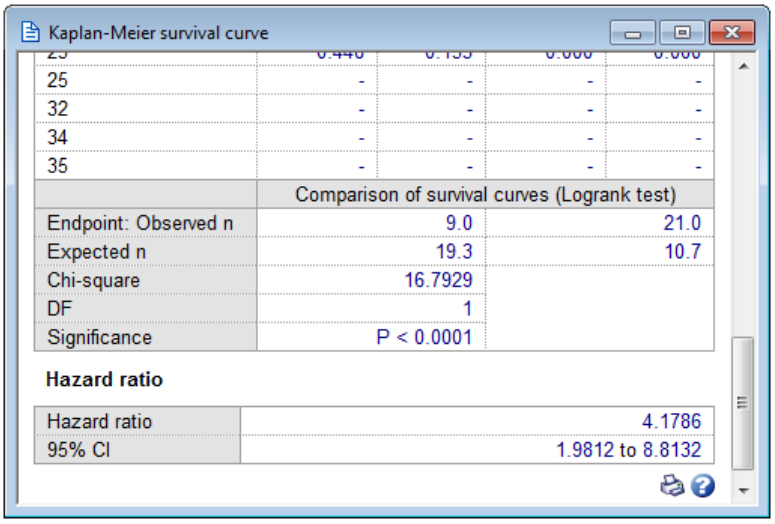

Hazard ratio

If you compare two survival curves then MedCalc also calculates the hazard ratio with its 95% confidence interval (CI). Hazard is a measure of how rapidly the event of interest occurs. The hazard ratio compares the hazards in two treatments groups.

In the example the hazard ratio is 4.1786 so that the estimated relative risk of the event of interest occurring in group 2 is 4.1786 higher than in group 1. This hazard ratio is significant different from the value 1 (corresponding to equal hazards) since the confidence interval 1.9812 to 8.8132 does not include the value 1.

Note that the computation of the hazard ratio assumes that the ratio is consistent over time, so therefore if the survival curves cross, the hazard ratio statistic should be ignored.

Logrank test for trend

If more than two survival curves are compared, and there is a natural ordering of the groups, then MedCalc can also perform the logrank test for trend. This tests the probability that there is a trend in survival scores across the groups.

Cox proportional-hazards regression

Whereas the Kaplan-Meier method with log-rank test is useful for comparing survival curves in two or more groups, Cox proportional-hazards regression allows analyzing the effect of several risk factors on survival.

The probability of the endpoint (death, or any other event of interest, e.g. recurrence of disease) is called the hazard. The hazard is modeled as:

$$H(t) = H_0(t) \times \exp(b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k)$$

where x_1 ... x_k are a collection of predictor variables and $H_0(t)$ is the baseline hazard at time *t*, representing the hazard for a person with the value 0 for all the predictor variables.

By dividing both sides of the above equation by $H_0(t)$ and taking logarithms, we obtain:

$$\ln\left(\frac{H(t)}{H_0(t)}\right) = b_1 X_1 + b_2 X_2 + b_3 X_3 + \ldots + b_k X_k$$

We call *H(t) / H₀(t)* the hazard ratio. The coefficients b_1...b_k are estimated by Cox regression, and can be interpreted in a similar manner to that of multiple logistic regression.

Suppose the covariate (risk factor) is **dichotomous,** and that the covariate is coded 1 if present and 0 if absent. Then the quantity exp(b*i*) can be interpreted as the instantaneous relative risk of an event, at any

time, for an individual with the risk factor present compared with an individual with the risk factor absent, given both individuals are the same on all other covariates.

Suppose the covariate is **continuous**, then the quantity exp(b*i*) is the instantaneous relative risk of an event, at any time, for an individual with an increase of 1 in the value of the covariate compared with another individual, given both individuals are the same on all other covariates.

Required input

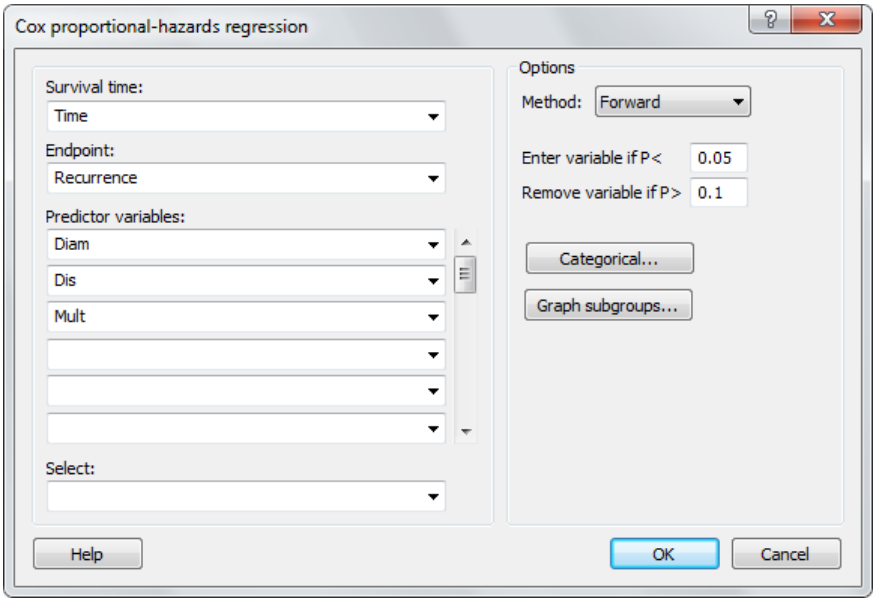

Survival time: The name of the variable containing the time to reach the event of interest, or the time of follow-up.

Endpoint: The name of a variable containing codes 1 for the cases that reached the endpoint, or code 0 for the cases that have not reached the endpoint, either because they withdrew from the study, or the end of the study was reached.

Predictor variables: Names of variables that you *expect* to predict survival time.

The Cox proportional regression model assumes that the effects of the predictor variables are constant over time. Furthermore there should be a linear relationship between the endpoint and predictor variables. Predictor variables that have a highly skewed distribution may require logarithmic transformation to reduce the effect of extreme values. Logarithmic transformation of a variable *var* can be obtained by entering LOG(*var*) as predictor variable.

Select: A selection criterion to include only a selected subgroup of cases in the graph.

Options

Method: select the way independent variables are entered into the model.

- Enter: enter all variables in the model in one single step, without checking
- Forward: enter significant variables sequentially
- Backward: first enter all variables into the model and next remove the non-significant variables sequentially
- Stepwise: enter significant variables sequentially; after entering a variable in the model, check and possibly remove variables that became non-significant.

Enter variable if P<

A variable is entered into the model if its associated significance level is less than this P-value.

Remove variable if P>

A variable is removed from the model if its associated significance level is greater than this P-value.

Categorical: click this button to identify nominal categorical variables.

Graph – Subgroups: here you can select one of the predictor variables. The graph will display different survival curves for all values in this covariate (which must be categorical, and may not contain more than 8 categories). If no covariate is selected here, then the graph will display the survival at mean of the covariates in the model.

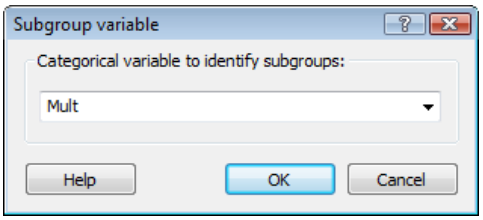

Results

In the example (taken from Bland, 2000), "survival time" is the time to recurrence of gallstones following dissolution (variable *Time*). Recurrence is coded in the variable *Recurrence* (1= yes, 0 =No). Predictor variables are *Dis* (= number of months previous gallstones took to dissolve), *Mult* (1 in case of multiple previous gallstones, 0 in case of single previous gallstones), and *Diam* (maximum diameter of previous gallstones).

Cox proportional-hazards regression

Survival time	Time
Endpoint	Recurrence

Method	Forward
Enter variable if P<	0.05
Remove variable if P>	0.1

Sample size	144

Overall Model Fit

Null model -2 Log Likelihood	339.09731
Full model -2 Log Likelihood	326.93288
Chi-square	12.1644
DF	2
Significance level	P = 0.0023

Coefficients and Standard Errors

Covariate	b	SE	P	Exp(b)	95% CI of Exp(b)
Dis	0.04292	0.01657	0.009584	1.0439	1.0107 to 1.0781
Mult	0.9635	0.3528	0.006309	2.6208	1.3173 to 5.2141

Variables not included in the model
Diam

Baseline cumulative hazard function

Time	Baseline Cumulative Hazard	At mean of Covariates	
		Cumulative Hazard	Survival
6	0.011	0.029	0.971
7	0.016	0.043	0.958
8	0.024	0.064	0.938
9	0.027	0.072	0.930
10	0.030	0.080	0.923
11	0.039	0.104	0.901
12	0.055	0.148	0.862
13	0.059	0.158	0.854
16	0.067	0.180	0.835

Overall Model Fit

The Chi-square statistic tests the relationship between time and all the covariates in the model.

Coefficients and Standard Errors

Using the *Forward* selection method, the two covariates *Dis* and *Mult* were entered in the model which significantly (0.0096 for *Dis* and 0.0063 for *Mult*) contribute to the prediction of time.

The coefficient for months for dissolution (continuous variable *Dis*) is 0.0429. Exp(b) = Exp(0.0429) is 1.0439 (with 95% Confidence Interval 1.0107 to 1.0781), meaning that for an increase of 1 month to dissolution of previous gallstones, the hazard ratio for recurrence increases by a factor 1.04. For 2 months the hazard ratio increases by a factor 1.04^2.

The coefficient for multiple gallstones (dichotomous variable *Mult*) is 0.9335. Exp(b) = Exp(0.9635) is 2.6208 (with 95% Confidence Interval 1.3173 to 5.2141), meaning that a case with previous gallstones is 2.62 more likely to have a recurrence than a case with a single stone.

Variables not included in the model

The variable *Diam* was found not to significantly contribute to the prediction of time, and was not included in the model.

Baseline cumulative hazard function

Finally, the program lists the baseline cumulative hazard $H_0(t)$, with the cumulative hazard and survival at mean of all covariates in the model.

The baseline cumulative hazard can be used to calculate the survival probability *S(t)* for any case at time t:

$$S(t) = \exp(-H_0(t) \times PI)$$

where PI is a prognostic index:

$$PI = x_1 b_1 + x_2 b_2 + x_3 b_3 + \ldots + x_k b_k$$

Graph

The graph displays the survival curves for all categories of the categorical variable *Mult* (1 in case of multiple previous gallstones, 0 in case of single previous gallstones), and for mean values for all other covariates in the model.

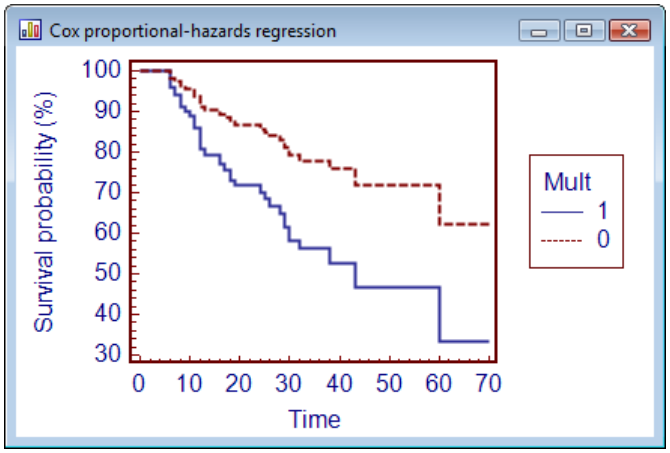

If no covariate was selected for Graph – Subgroups, or if the selected variable was not included in the model, then the graph displays a single survival curve at mean of all covariates in the model.

Meta-analysis: odds ratio

A meta-analysis integrates the quantitative findings from separate but similar studies and provides a numerical estimate of the overall effect of interest (Petrie et al., 2003).

For meta-analysis of studies with continuous measures (comparison of means), see p. 125.

Under the **fixed effects model**, t s assumed that all studies come from a common population, and that the effect size (odds ratio) s not significantly different among the different trials. This assumption is tested by the “Heterogeneity test”. If this test yields a low P-value (P<0.05), then the fixed effects model may be invalid. In this case, the **random effects model** may be more appropriate, in which both the random variation within the studies and the variation between the different studies is incorporated.

MedCalc uses the Mantel-Haenszel method for calculating the weighted summary Odds ratio under the fixed effects model. Next the heterogeneity statistic is incorporated to calculate the summary Odds ratio under the random effects model (DerSimonian and Laird).

Required input

The data of different studies can be entered as follows in the spreadsheet:

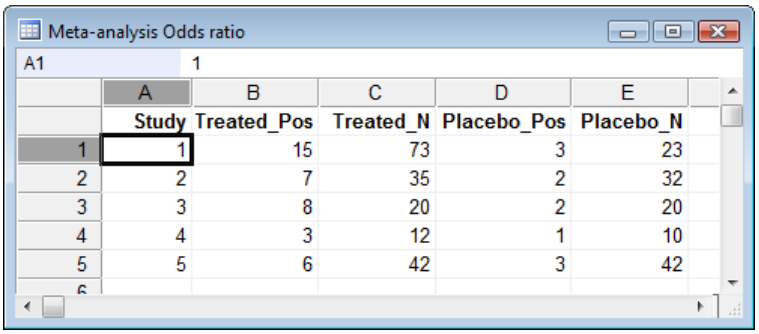

	A	B	C	D	E
	Study	Treated_Pos	Treated_N	Placebo_Pos	Placebo_N
1	1	15	73	3	23
2	2	7	35	2	32
3	3	8	20	2	20
4	4	3	12	1	10
5	5	6	42	3	42

In this example, in a first study 73 cases were treated with an active substance and of these, 15 had a positive outcome. 23 cases received a placebo and 3 of these had a positive outcome. On the next rows of the spreadsheet follow the data of 4 other studies.

The dialog box for “Meta-analysis: odds ratio” can then be completed as follows:

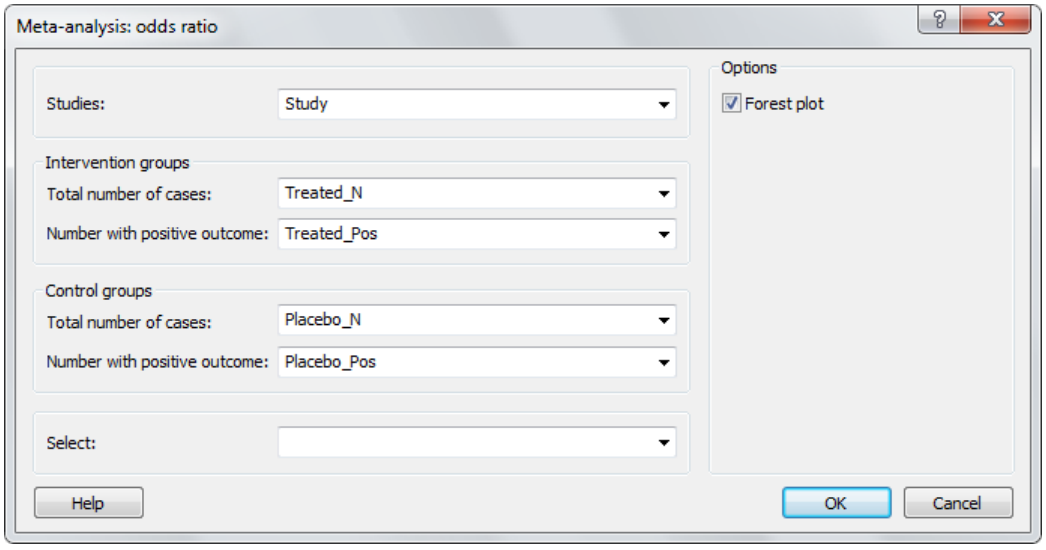

Results

Meta-analysis: odds ratio

Variable for studies	Study
1. Intervention groups	
Variable for total number of cases	Treated_N
Variable for number of positive cases	Treated_Pos
2. Control groups	
Variable for total number of cases	Placebo_N
Variable for number of positive cases	Placebo_Pos

Study	Intervention	Controls	Odds	95% CI
1	15/73	3/23	1.724	0.452 to 6.583
2	7/35	2/32	3.750	0.717 to 19.599
3	8/20	2/20	6.000	1.082 to 33.275
4	3/12	1/10	3.000	0.260 to 34.576
5	6/42	3/42	2.167	0.504 to 9.312
Total (fixed effects)	39/182	11/127	2.806	1.363 to 5.778
Total (random effects)	39/182	11/127	2.781	1.347 to 5.744

Test for heterogeneity

Q	1.5056
DF	4
Significance level	P = 0.8256

The program lists the results of the individual studies: number of positive cases, total number of cases, the odds ratio with 95% CI.

The Total Odds ratio with 95% CI is given both for the *Fixed effects model* and the *Random effects model*. If the value 1 is not within the 95% CI, then the Odds ratio is statistically significant at the 5% level (P<0.05).

The random effects model will tend to give a more conservative estimate (i.e. with wider confidence interval), but the results from the two models usually agree where there is no heterogeneity. If the test of heterogeneity is statistically significant (P<0.05) then more emphasis should be placed on the random effects model.

Graph

The results of the different studies, with 95% CI, and the overall effect with 95% CI are shown in the following graph (called a *forest plot*):

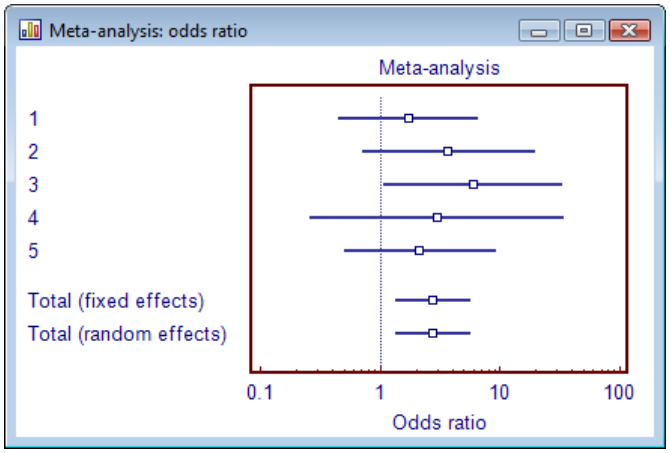

Note that the Odds ratios with 95% CI are drawn on a logarithmic scale.

Meta-analysis: continuous measure

For meta-analysis of studies with a continuous measure (comparison of means between treated cases and controls), MedCalc uses the **Hedges g** statistic as a formulation for the standardized mean difference under the fixed effects model. Next the heterogeneity statistic is incorporated to calculate the summary standardized mean difference under the random effects model (DerSimonian and Laird) (see also: Meta-analysis: odds ratio, p. 122).

The standardized mean difference Hedges g is the difference between the two means divided by the pooled standard deviation, with an adjustment for small sample bias.

Required input

The data of different studies can be entered as follows in the spreadsheet:

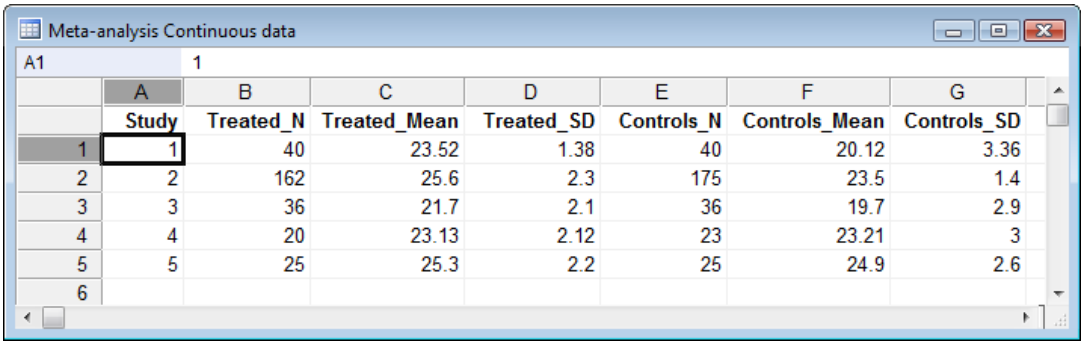

	A	B	C	D	E	F	G
	Study	Treated_N	Treated_Mean	Treated_SD	Controls_N	Controls_Mean	Controls_SD
1	1	40	23.52	1.38	40	20.12	3.36
2	2	162	25.6	2.3	175	23.5	1.4
3	3	36	21.7	2.1	36	19.7	2.9
4	4	20	23.13	2.12	23	23.21	3
5	5	25	25.3	2.2	25	24.9	2.6
6							

In this example, in a first study 40 cases were treated and the mean of the parameter of interest was 23.52 with a standard deviation of 1.38. In 23 control cases the mean was 20.12 with standard deviation of 3.36. On the next rows of the spreadsheet follow the data of 4 other studies.

The dialog box for “Meta-analysis: continuous measure” can then be completed as follows:

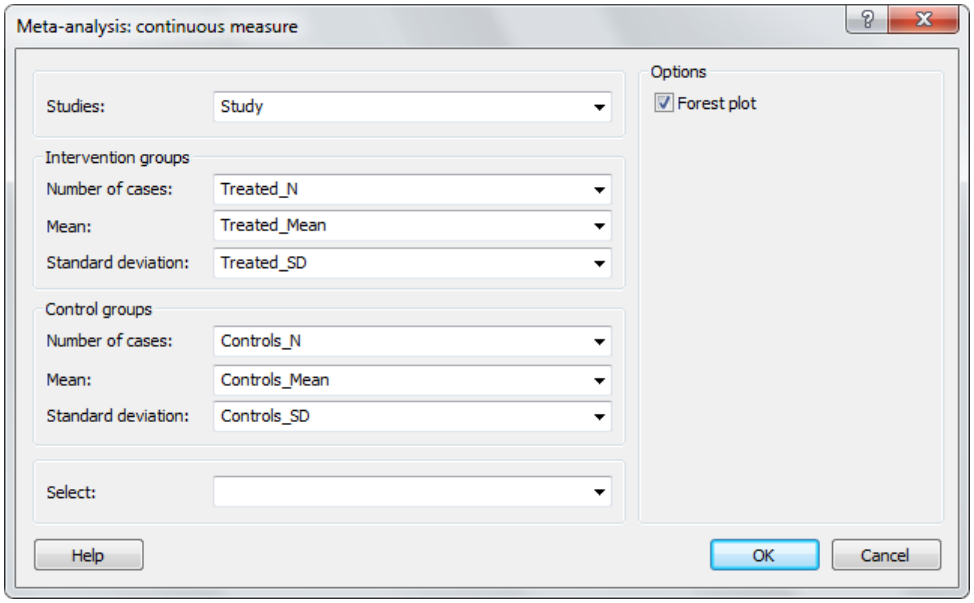

Results

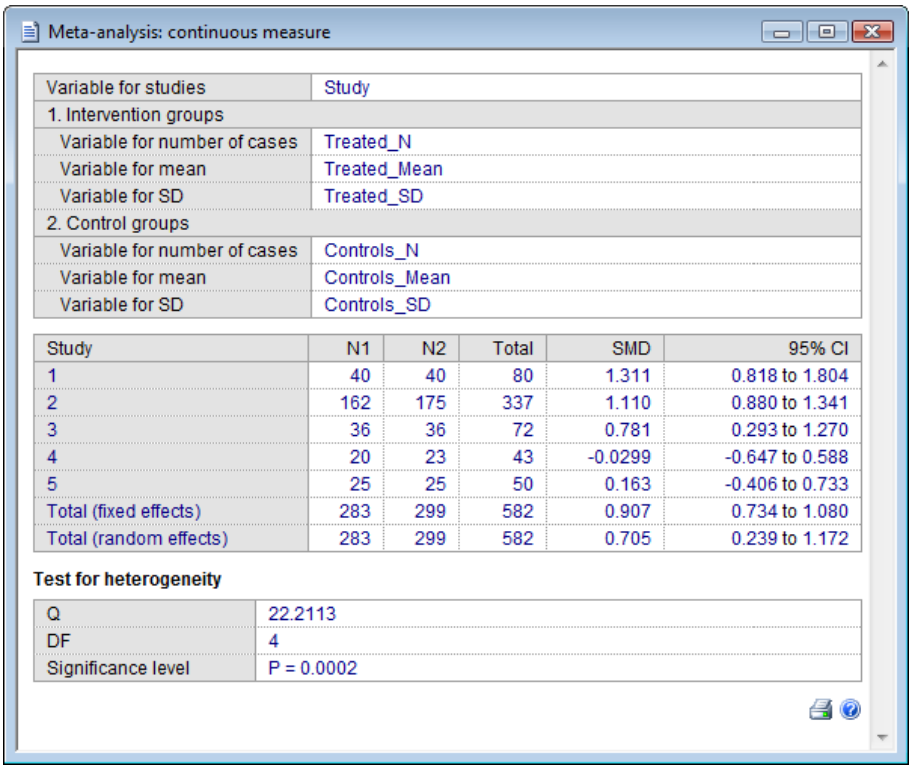

Meta-analysis: continuous measure

Variable for studies	Study
1. Intervention groups	
Variable for number of cases	Treated_N
Variable for mean	Treated_Mean
Variable for SD	Treated_SD
2. Control groups	
Variable for number of cases	Controls_N
Variable for mean	Controls_Mean
Variable for SD	Controls_SD

Study	N1	N2	Total	SMD	95% CI
1	40	40	80	1.311	0.818 to 1.804
2	162	175	337	1.110	0.880 to 1.341
3	36	36	72	0.781	0.293 to 1.270
4	20	23	43	-0.0299	-0.647 to 0.588
5	25	25	50	0.163	-0.406 to 0.733
Total (fixed effects)	283	299	582	0.907	0.734 to 1.080
Total (random effects)	283	299	582	0.705	0.239 to 1.172

Test for heterogeneity

Q	22.2113
DF	4
Significance level	P = 0.0002

The program lists the results of the individual studies: number of positive cases, total number of cases, the standardized mean difference (SMD) with 95% CI.

The Total standardized mean difference with 95% CI is given both for the *Fixed effects model* and the *Random effects model*.

If the value 0 is not within the 95% CI, then the SMD is statistically significant at the 5% level (P<0.05).

Cohen's rule of thumb for interpretation of the SMD statistic is: a value of 0.2 indicates a small effect, a value of 0.5 indicates a medium effect and a value of 0.8 or larger indicates a large effect.

The random effects model will tend to give a more conservative estimate (i.e. with wider confidence interval), but the results from the two models usually agree where there is no heterogeneity. If the test of heterogeneity is statistically significant (P<0.05), like in the example, then more emphasis should be placed on the random effects model.

Graph

The results of the different studies, with 95% CI, and the overall standardized mean difference with 95% CI are shown in the following *forest plot*:

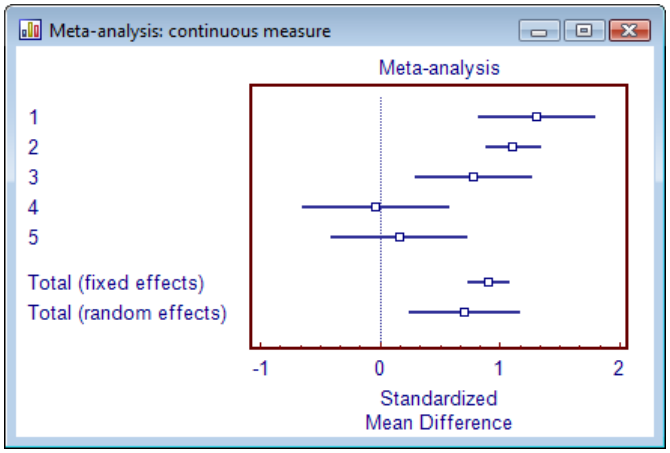

Serial measurements

Description

When serial measurements have been performed, e.g. in a glucose tolerance test, it is often useful to summarize the data on every subject into one or more summarizing measures, such as:

- height of peak (maximum value)
- time to reach peak
- difference last-first value
- area under the curve.

These summary measures can be obtained, and analyzed statistically, with the *Serial measurements* command. The program will create a graph for the observations, and perform an appropriate statistical test to evaluate whether the summary measure differs between the different subgroups.

Data input

While for any other statistical procedure in MedCalc, the data of one case have to be entered on one single row of the spreadsheet, for serial measurements MedCalc requires the data for all cases to be entered in one single column (*Value* in the example). Other columns in the spreadsheet must contain:

- a variable for the time: time should be entered as numbers and should be expressed in the same unit (e.g. hours, days); time intervals can be unequal
- a categorical variable that identifies cases
- an (optional) categorical variable that identifies the group a case belongs to.

In the example, case 1 belongs to groups "Malabsorption" and the value 0.08 was measured at time 0, value 5.70 at time 30 and 3.22 at time 45.

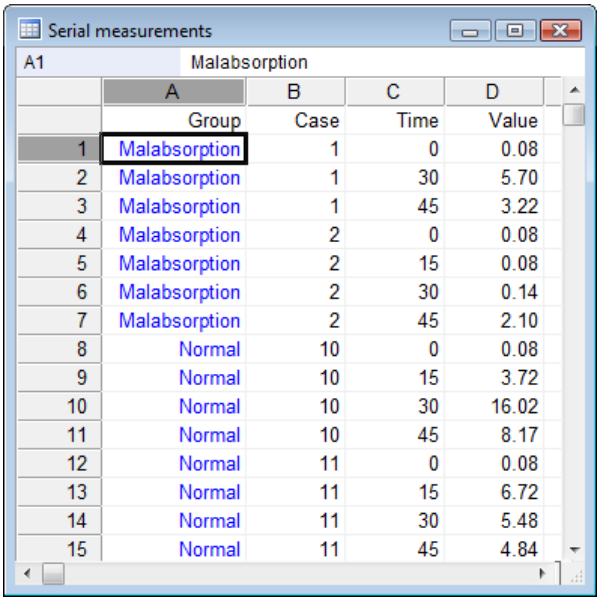

	A	B	C	D
	Group	Case	Time	Value
1	Malabsorption	1	0	0.08
2	Malabsorption	1	30	5.70
3	Malabsorption	1	45	3.22
4	Malabsorption	2	0	0.08
5	Malabsorption	2	15	0.08
6	Malabsorption	2	30	0.14
7	Malabsorption	2	45	2.10
8	Normal	10	0	0.08
9	Normal	10	15	3.72
10	Normal	10	30	16.02
11	Normal	10	45	8.17
12	Normal	11	0	0.08
13	Normal	11	15	6.72
14	Normal	11	30	5.48
15	Normal	11	45	4.84

The dialog box

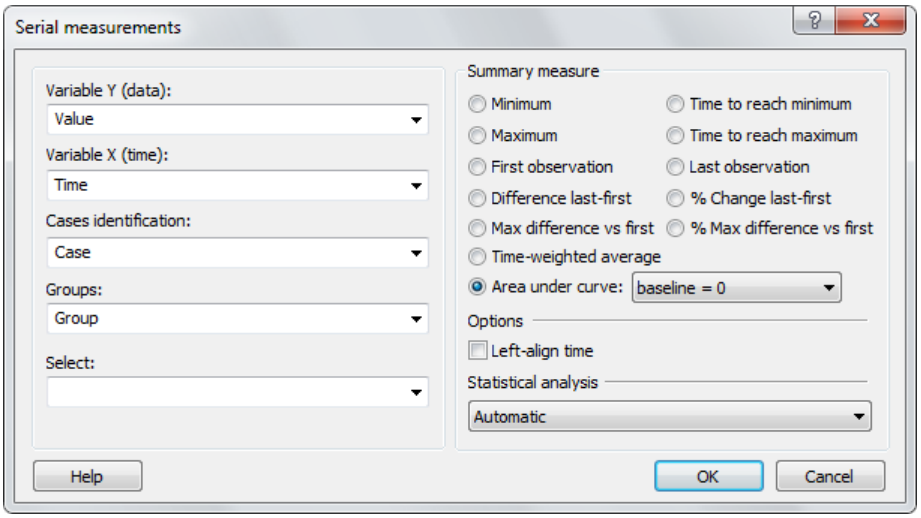

Variable Y (data): this variable contains the serial measurements of all cases.

Variable X (time): this variable contains the time of the different measurements.

Case identification: a categorical variable containing case identification data.

Groups: a categorical variable containing group identification data.

Summary measure: here you select the **summary measure of interest** for statistical analysis

- Minimum, and Time to reach minimum
- Maximum, and Time to reach maximum
- First observation, Last observation
- Difference Last-First observation
- % Change Last-First observation, calculated as 100x(Last-First)/First

  Note that in order to calculate this summary measure, the first observation should never be equal to zero.
- Maximum difference with first observation
- % Maximum difference with first observation, calculated as 100x(Max difference/First)

  Note that in order to calculate this summary measure, the first observation should never be equal to zero.
- Time-weighted average is calculated as the case AUC (baseline 0) divided by its total time interval (time of last observation minus time of first observation). If there is only one observation, this value is taken as the time-weighted average.
- Area under the curve (AUC)

  The area under the curve can be calculated in 3 different ways, depending on what is taken as the baseline value: 0, first observation, or minimum (when the first observation is taken as the baseline value, the area under the curve can be a negative number).

  Note that for the AUC, MedCalc requires that for all cases the first and last observations are set at the same time.

Options

- Left-align time: left-align time by subtracting the first time value of a case from all other time values of that case. As a result, the start time of each case is set at 0.

Statistical analysis

- **Automatic**: let MedCalc decide how to analyze the data. If this option is selected, MedCalc will analyze the summary statistics in the different groups and perform a test for Normal distribution. If the data have a Normal distribution, the software will perform a parametric test. If the data do not have a Normal distribution, the software will attempt a Logarithmic transformation. If the data have a Normal distribution after Log transformation, a statistical test will be performed on the log-transformed data; if not, a non-parametric test will be used on the non-transformed data.

- **Parametric test**: assume the summary measure has a Normal distribution and use a parametric test
- **Parametric test after Log transformation**: perform a Logarithmic transformation on the summary measure and then do a parametric test.
- **Non-parametric test**: use a non-parametric test.

For a parametric test, MedCalc uses the t-test when there are 2 groups or One-way analysis of variance (ANOVA) when there are more than 2 groups.

For a non-parametric test, MedCalc uses the Mann-Whitney when there are 2 groups or the Kruskal-Wallis test when there are more than 2 groups.

Results

The results window displays the statistics (sample size, mean, SD, etc.) for the selected summary measure in the different subgroups.

Next, a statistical test is performed to test the hypothesis that there is no difference between the subgroups. If the resulting P-value is less than 0.05, the conclusion is that there is a significant difference of the summary measure between the different subgroups.

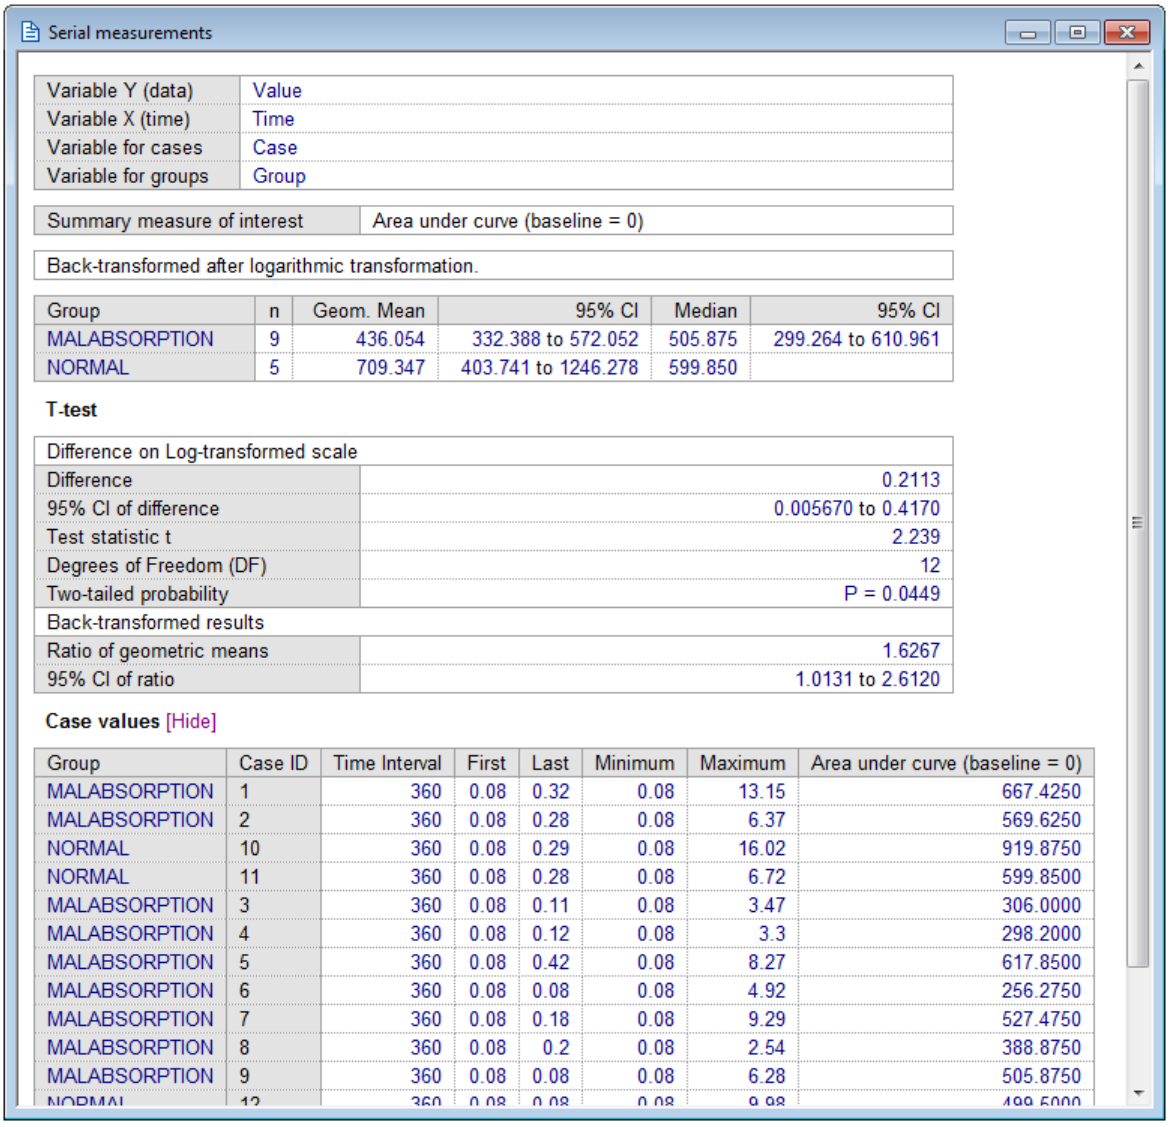

Serial measurements

Variable Y (data)	Value
Variable X (time)	Time
Variable for cases	Case
Variable for groups	Group

Summary measure of interest	Area under curve (baseline = 0)

Back-transformed after logarithmic transformation.

Group	n	Geom. Mean	95% CI	Median	95% CI
MALABSORPTION	9	436.054	332.388 to 572.052	505.875	299.264 to 610.961
NORMAL	5	709.347	403.741 to 1246.278	599.850	

T-test

Difference on Log-transformed scale	
Difference	0.2113
95% CI of difference	0.005670 to 0.4170
Test statistic t	2.239
Degrees of Freedom (DF)	12
Two-tailed probability	P = 0.0449
Back-transformed results	
Ratio of geometric means	1.6267
95% CI of ratio	1.0131 to 2.6120

Case values [Hide]

Group	Case ID	Time Interval	First	Last	Minimum	Maximum	Area under curve (baseline = 0)
MALABSORPTION	1	360	0.08	0.32	0.08	13.15	667.4250
MALABSORPTION	2	360	0.08	0.28	0.08	6.37	569.6250
NORMAL	10	360	0.08	0.29	0.08	16.02	919.8750
NORMAL	11	360	0.08	0.28	0.08	6.72	599.8500
MALABSORPTION	3	360	0.08	0.11	0.08	3.47	306.0000
MALABSORPTION	4	360	0.08	0.12	0.08	3.3	298.2000
MALABSORPTION	5	360	0.08	0.42	0.08	8.27	617.8500
MALABSORPTION	6	360	0.08	0.08	0.08	4.92	256.2750
MALABSORPTION	7	360	0.08	0.18	0.08	9.29	527.4750
MALABSORPTION	8	360	0.08	0.2	0.08	2.54	388.8750
MALABSORPTION	9	360	0.08	0.08	0.08	6.28	505.8750
NORMAL	12	360	0.08	0.08	0.08	9.98	499.5000

Graph

The Serial measurements diagram plots the serial data consecutively, for every case.

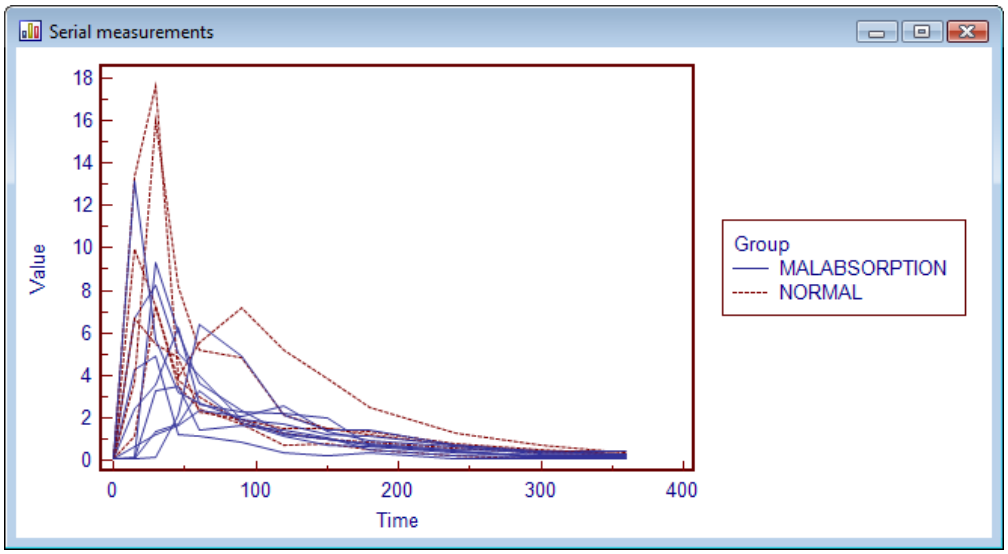

If there are many cases, you may want to create different graphs for each subgroup. To create a serial measurements graph for all cases in the "normal" group, enter the following in the *Select* box:

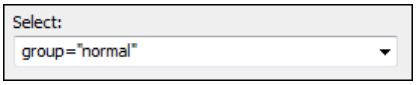

which gives the following graph:

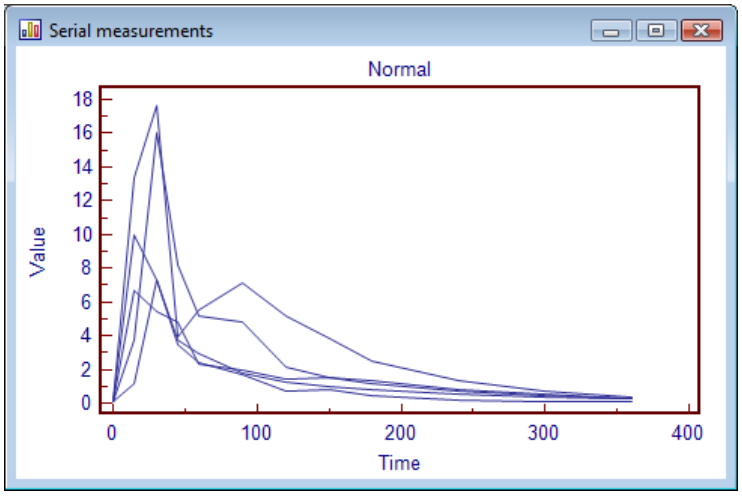

For the "Malabsorption" group, you enter:

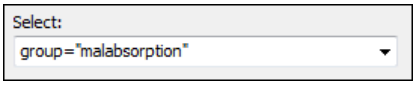

which gives the following graph:

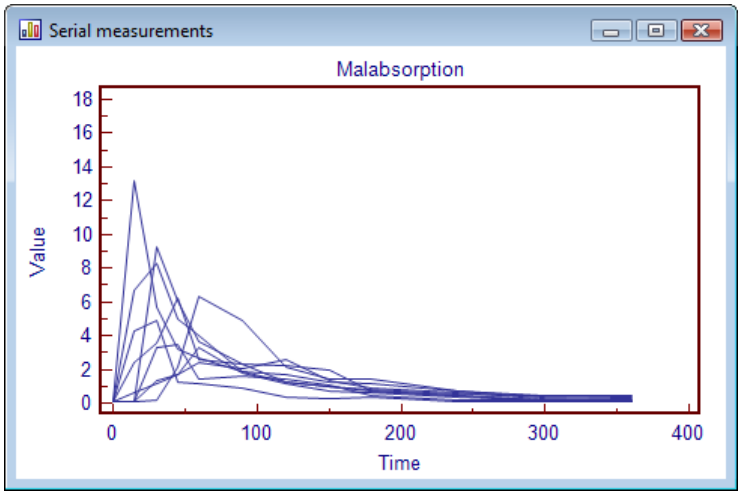

Literature

- Bland M (2000) An introduction to medical statistics, 3rd ed. Oxford: Oxford University Press.
- Mathews JNS, Altman DG, Campbell MJ, Royston P (1990) Analysis of serial measurements in medical research. British Medical Journal, 300, 230-235.

Reference interval

A Reference interval (Reference range, Normal range) can be calculated using the following two methods: (a) using the Normal distribution, (b) using a non-parametrical percentile method, and (c) optionally a "robust method" as described in the CLSI Guidelines C28-A3.

In the dialog box you identify the variable of interest. You can click the button to obtain a list of variables. In this list you can select a variable by clicking the variable's name. You can also enter a selection criterion in the *Select* field, in order to include only a selected subgroup of cases in the statistical procedure, as described in the *Introduction* part of this manual.

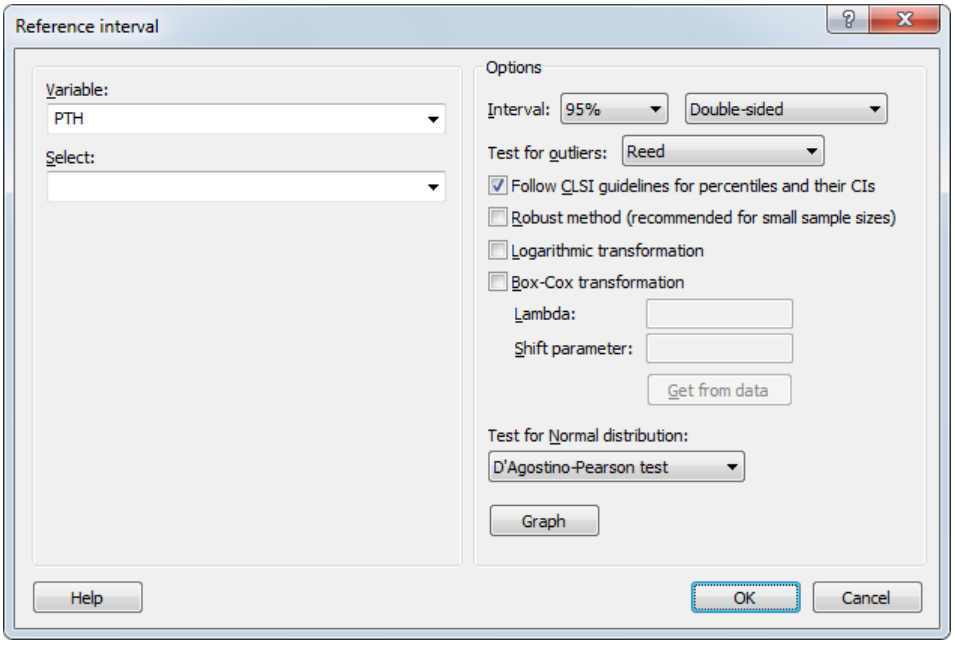

Options

- Reference interval: select a **90%, 95% or 99%** reference interval. A 95% interval is the most usual and preferred setting.
- **Double sided, left or right sided**:

  Select *Double sided* when there is both a lower and upper limit of normality (both low and high values are suspicious).

  Select *Left sided* when there is only a lower limit of normality and no upper limit of normality (only low values are suspicious).

  Select *Right sided* when there is only an upper limit of normality and no lower limit of normality (only high values are suspicious).

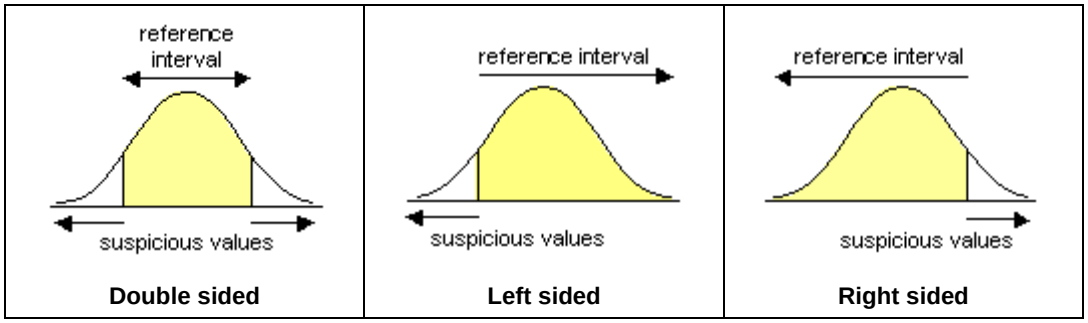

- **Test for outliers**: select the method based on Reed et at. (1971) or Tukey (1977) to automatically check the data for outliers (alternatively select none for no outlier testing). The method by Reed et at. will test only the minimum and maximum observations; the Tukey test can identify more values as outliers. The tests will create a list of possible outlying observations, but these will not automatically be excluded from the analysis. The possible outliers should be inspected by the investigator who can decide to exclude the values (see Exclude & Include, p. 42).
- Follow CLSI guidelines for percentiles and their CIs: select this option to follow the NCCLS and Clinical and Laboratory Standards Institute (CLSI) guidelines C28-A2 and C28-A3 for estimating percentiles and their 90% confidence intervals. In these guidelines, percentiles are calculated as the observations corresponding to rank r=p*(n+1). Also for the 90% confidence intervals of the reference limits the CLSI guidelines are followed and conservative confidence intervals are calculated using integer ranks (and therefore the confidence intervals are at least 90% wide).

  If you do not select this option, MedCalc calculates percentiles as the observations corresponding to rank r= p*n+0.5 (Lentner, 1982), and calculates a less conservative and more precise confidence interval using an iterative method.
- **Robust method**: select this option to calculate the reference limits with the "robust method" (CLSI Guidelines C28-A3). Recommended for smaller sample sizes (less than 120).
- Logarithmic transformation
- **Box-Cox transformation**: this will allow to perform a Box-Cox transformation with the following parameters:
  - **Lambda**: the power parameter λ
  - **Shift parameter**: the shift parameter is a constant c that needs to be added to the data when some of the data are negative.
  - Button **Get from data**: click this button to estimate the optimal value for Lambda, and suggest a value for the shift parameter c when some of the observations are negative. The program will suggest a value for Lambda with 2 to 3 significant digits.

  The Box-Cox transformation is defined as follows:

  $x(\lambda) = ((x+c)^{\lambda} - 1) / \lambda$ when $\lambda \neq 0$

  $x(\lambda) = \ln(x+c)$ when $\lambda = 0$

  When you perform a Box-Cox transformation, MedCalc will automatically transform the data with the selected parameters and will back-transform the results to the original scale for presentation.
- Chi-square test, Kolmogorov-Smirnov test, or D'Agostino-Pearson **test for Normal distribution** (see p. 57).
- **Graph**: graph option (see below).

Results

The results window for *Reference interval* displays the following information:

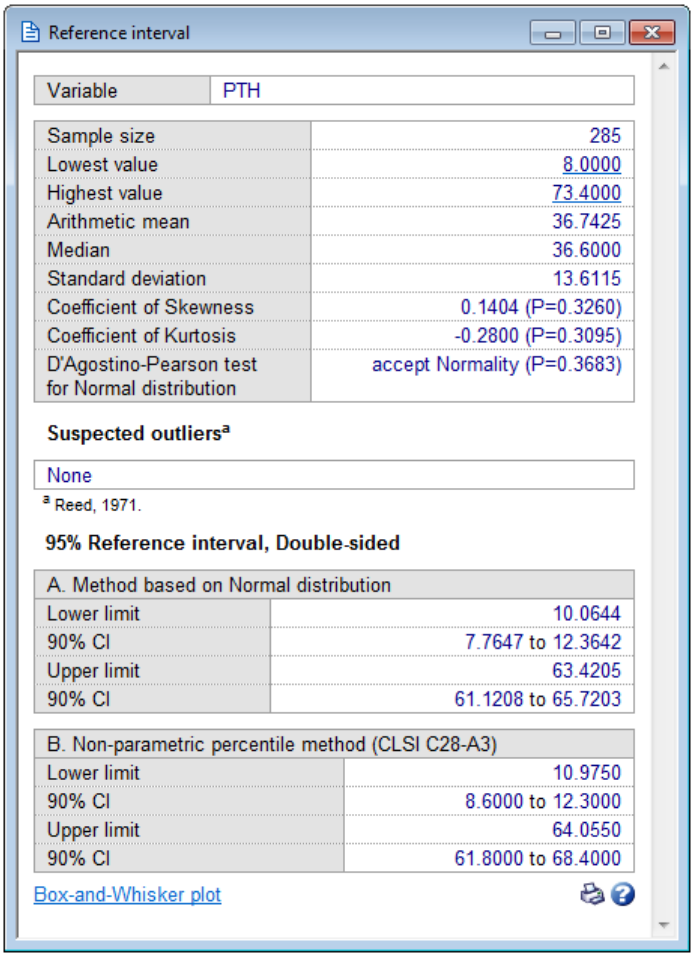

Reference interval

Variable	PTH

Sample size	285
Lowest value	8.0000
Highest value	73.4000
Arithmetic mean	36.7425
Median	36.6000
Standard deviation	13.6115
Coefficient of Skewness	0.1404 (P=0.3260)
Coefficient of Kurtosis	-0.2800 (P=0.3095)
D'Agostino-Pearson test for Normal distribution	accept Normality (P=0.3683)

Suspected outliers[a]

None

[a] Reed, 1971.

95% Reference interval, Double-sided

A. Method based on Normal distribution	
Lower limit	10.0644
90% CI	7.7647 to 12.3642
Upper limit	63.4205
90% CI	61.1208 to 65.7203

B. Non-parametric percentile method (CLSI C28-A3)	
Lower limit	10.9750
90% CI	8.6000 to 12.3000
Upper limit	64.0550
90% CI	61.8000 to 68.4000

Box-and-Whisker plot

- **Sample size**: the number of cases N is the number of numerical entries for the variable that fulfill the selection criterion.
- **Range**: the lowest and highest value of all observations.
- **Arithmetic mean**: the arithmetic mean is the sum of all observations divided by the number of observations.
- **Median**: when you have 100 observations, and these are sorted from smaller to larger, then the median is equal to the middle value. If the distribution of the data is Normal, then the median is equal to the arithmetic mean.
- **Standard Deviation**: the standard deviation is the square root of the variance. When the distribution of the observations is Normal, then 95% of observations are located in the interval Mean ± 2SD.
- **Skewness**: degree of symmetry of the sample distribution (see p. 59).
- **Kurtosis**: degree of peakedness/flatness of the sample distribution (see p. 60).
- **Test for Normal Distribution**: The result of this test is expressed as '*accept Normality*' or '*reject Normality*', with P value.

  If P is higher than 0.05, it may be assumed that the data follow a Normal distribution and the conclusion '*accept Normality*' is displayed.

  If P is less than 0.05, then the hypothesis that the distribution of the observations in the sample is Normal, should be rejected, and the conclusion '*reject Normality*' is displayed.

Log transformation

If the option Log transformation was selected, the program will display the back-transformed results. The back-transformed mean is named the *Geometric mean*. The Standard deviation cannot be back-transformed meaningfully and is not reported.

Suspected outliers

The program produces a list of possible outliers, detected by the methods based on Reed et at. (1971) or Tukey (1977). The method by Reed et at. tests only the minimum and maximum observations; the Tukey test can identify more values as outliers. Note that this does not automatically exclude any values from the analysis. The observations should be further inspected by the investigator who can decide to exclude the values. Click on the listed values (which are displayed as hyperlinks) to show the corresponding data in the spreadsheet (see Exclude & Include, p. 42).

Reference interval

The program will give the 90, 95 or 99% Reference interval, double sided or left or right sided only, as selected in the dialog box.

The reference interval is calculated using 3 different methods: (a) using the Normal distribution, (b) using a non-parametrical percentile method, and (c) optionally a "robust method" as described in the CLSI Guidelines C28-A3.

90% Confidence Intervals are given for the reference limits.

> For the robust method the confidence intervals are estimated with the bootstrap method using 5000 replications. When sample size is very small and/or the sample contains too many equal values, it may be impossible to calculate the CIs.

The results from the Normal distribution method are not appropriate when the Test for Normal distribution (see above) fails. If sample size is large (120 or more) the CLSI C28-A3 guideline recommends the percentile method and for smaller sample sizes the "robust method" is recommended.

The minimal sample size of 120 for the percentile method is the minimum number required to calculate 90% Confidence Intervals for the reference limits. A higher number of cases is required to achieve more reliable reference limits with more narrow 90% Confidence Intervals.

Graph

Click the **Graph** button in the dialog box shown above to obtain the following Reference Interval Graph box:

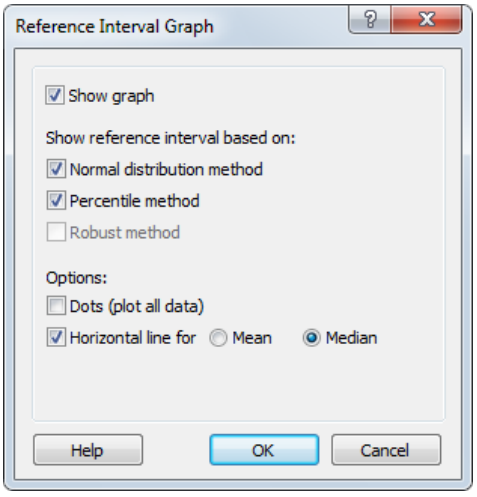

This results in the following graph:

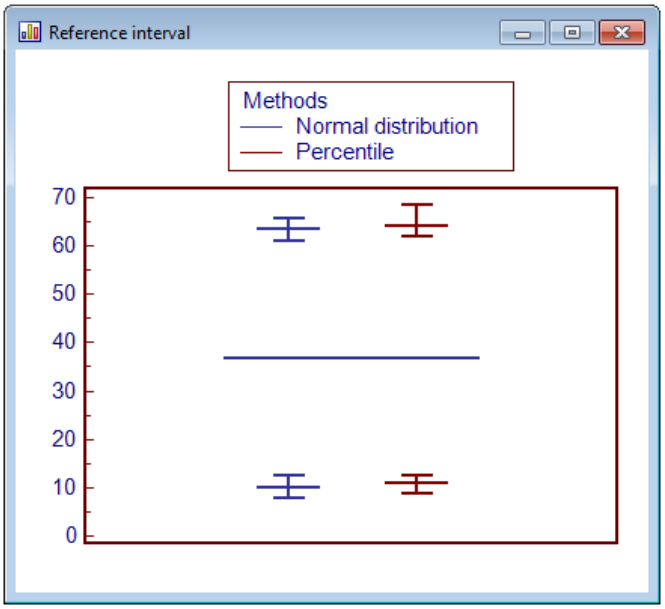

Bland-Altman plot

The Bland-Altman plot (Bland & Altman, 1986 and 1999), or difference plot, is a graphical method to compare two measurements techniques. In this graphical method the differences (or alternatively the ratios) between the two techniques are plotted against the averages of the two techniques. Alternatively (Krouwer, 2008) the differences can be plotted against one of the two methods, if this method is a reference or "gold standard" method.

Horizontal lines are drawn at the mean difference, and at the limits of agreement, which are defined as the mean difference plus and minus 1.96 times the standard deviation of the differences.

After you have selected *Bland-Altman plot* in the menu, enter the variables for the two different techniques in the following dialog box:

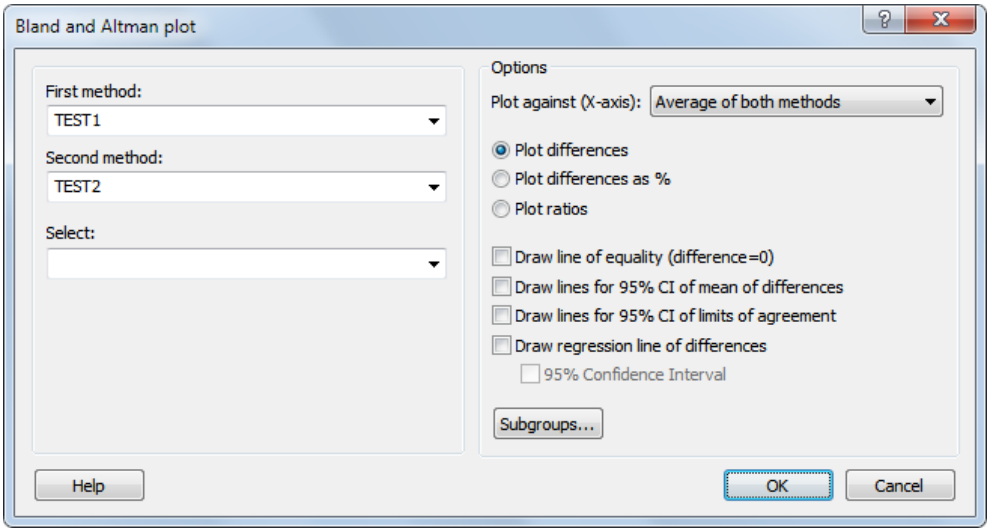

You can select the following variations of the Bland & Altman plot (see Bland & Altman, 1999; Krouwer, 2008):

- **Plot against (X-axis)**

  In the original Bland & Altman plot (Bland & Altman, 1986) the differences* between the two methods are plotted against the averages of the two methods (**recommended**, Bland & Altman, 1995).

Alternatively, you can choose to plot the differences* against one of the two methods, if this is a reference or "gold standard" method (Krouwer, 2008). Finally, you can also plot the differences* against the geometric mean of both methods.

* or ratios when this option is selected (see below).

- **Plot differences**

  This is the default option corresponding to the methodology of Bland & Altman, 1986.

- **Plot differences as %**

  When selecting this option the differences will be expressed as percentages of the values on the axis (i.e. proportionally to the magnitude of measurements). This option is useful when there is an increase in variability of the differences as the magnitude of the measurement increases.

- **Plot ratios**

  When this option is selected then the ratios of the measurements will be plotted instead of the differences (avoiding the need for log transformation). This option as well is useful when there is an increase in variability of the differences as the magnitude of the measurement increases. However, the program will give a warning when either one of the two techniques includes zero values.

Options

- Click the **Subgroups** button if you want to identify subgroups in the scatter diagram. A new dialog box is displayed in which you can select a categorical variable. The graph will use different markers for the different categories in this variable.
- **Draw line of equality**: useful for detecting a systematic difference.
- **Draw lines for 95% CI of mean of differences***: the 95% Confidence Interval of the mean difference illustrates the magnitude of the systematic difference. If the line of equality is not in the interval, there is a significant systematic difference.
- **Draw lines for 95% CI of limits of agreement**: shows lines for the 95% confidence interval for both the upper and lower limits of agreement.
- **Draw regression line of differences*** versus averages: this regression line may help to detect a proportional difference. The regression parameters are shown in the graph's info panel. Optionally, you can select to show the 95% confidence interval of this regression line.

* or ratios when this option was selected.

Use the **Subgroups** button if you want to identify subgroups in the scatter diagram. A new dialog box is displayed in which you can select a categorical variable. The graph will display different markers for the different categories in this variable.

After clicking the OK button, you obtain the following graph:

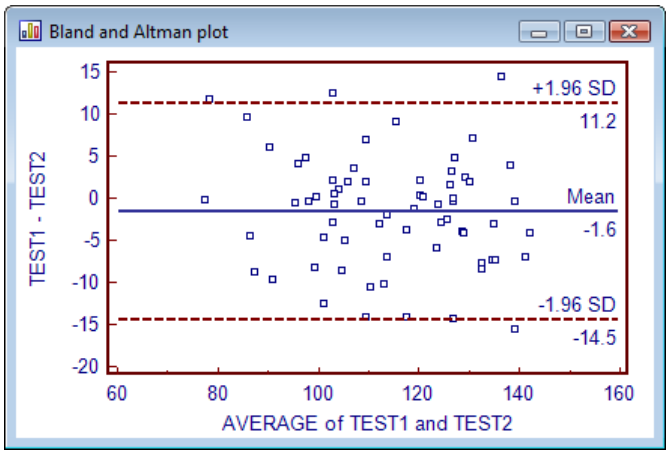

This graph displays a scatter diagram of the differences plotted against the averages of the two measurements. Horizontal lines are drawn at the mean difference, and at the limits of agreement, which are defined as the mean difference plus and minus 1.96 times the standard deviation of the differences.

To get more statistical information, right-click in the graph window and select the **Info** option in the popup menu. A new window appears inside the graph window:

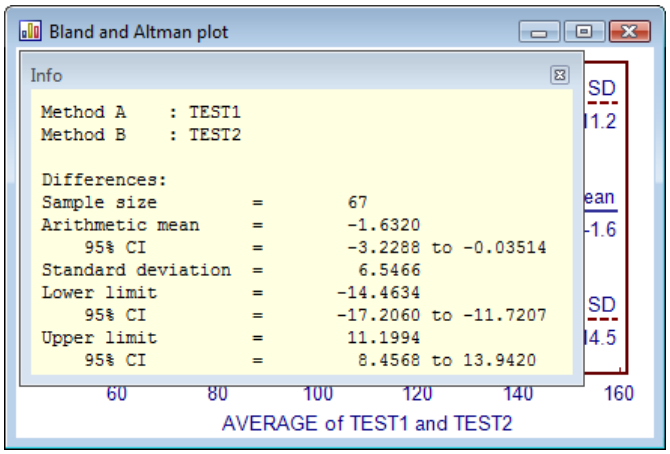

The Bland & Altman plot is useful to reveal a possible relationship between the differences and the averages (examples 1 & 2), to look for any systematic bias (example 3) and to identify possible outliers. If there is a consistent bias, it can be adjusted for by subtracting the mean difference from the new method.

If the differences within mean ± 1.96 SD are not *clinically* important, the two methods may be used interchangeably.

Some typical situations are shown in the following examples.

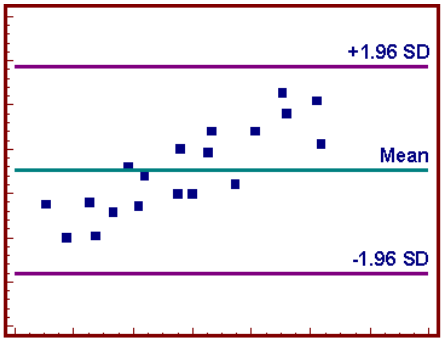

Example 1: Case of a proportional error.

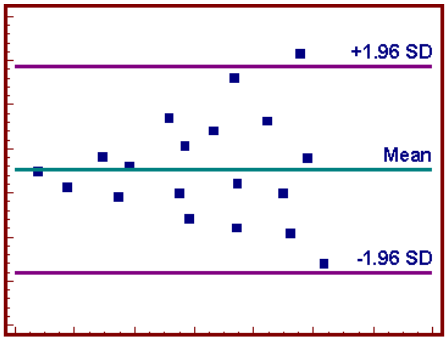

Example 2: Case where the variation of at least one method depends strongly on the magnitude of measurements.

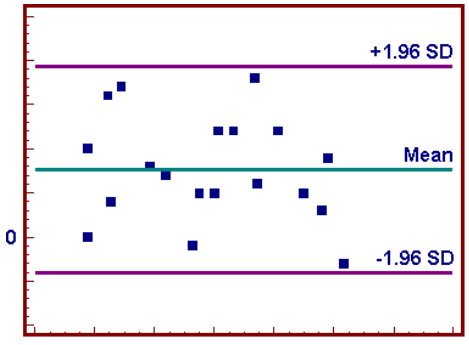

Example 3: Case of an absolute systematic error.

Repeatability

The Bland and Altman plot may also be used to assess the *repeatability* of a method by comparing repeated measurements using one single method on a series of subjects. The graph can then also be used to check whether the variability or precision of a method is related to the size of the characteristic being measured.

Since for the repeated measurements the same method is used, the mean difference should be zero. Therefore the *Coefficient of Repeatability* (CR) can be calculated as 1.96 (≈2) times the standard deviations of the differences between the two measurements (d_2 and d_1):

$$CR = 1.96 \times \sqrt{\frac{\sum(d_2 - d_1)^2}{n-1}}$$

This coefficient can be read from the Bland & Altman plot by subtracting the mean difference from the upper 95% limit of agreement, but can also be calculated using *Summary statistics*. E.g. if the names of the variables for 2 repeated measurements for FSH concentration are FSH1 and FSH2, then you define a new variable as FSH2-FSH1 and calculate the summary statistics for it. By multiplying the calculated standard deviation by 2 you obtain the coefficient of repeatability.

Bland-Altman plot with multiple measurements per subject

Use this command to create a Bland-Altman plot for method comparison when there is more than one measurement per subject with each laboratory method.

How to enter data

This procedure requires that you have your data organized like illustrated in the following example (data from Bland & Altman, 2007):

	A	B	C
	Subject	RV	IC
1	1	7.83	6.57
2	1	7.42	5.62
3	1	7.89	6.9
4	1	7.12	6.57
5	1	7.88	6.35
6	2	6.16	4.06
7	2	7.26	4.29
8	2	6.71	4.26
9	2	6.54	4.09
10	3	4.75	4.71
11	3	5.24	5.5
12	3	4.86	5.08

There is one column for subject identification (**Subject**) and one column for the measurements for each method (**RV** and **IC**).

If you have your data organized in a different format, such as the data for the multiple measurements in different columns, you can use the Stack columns tool to reorganize your data (see p. 43).

Required input

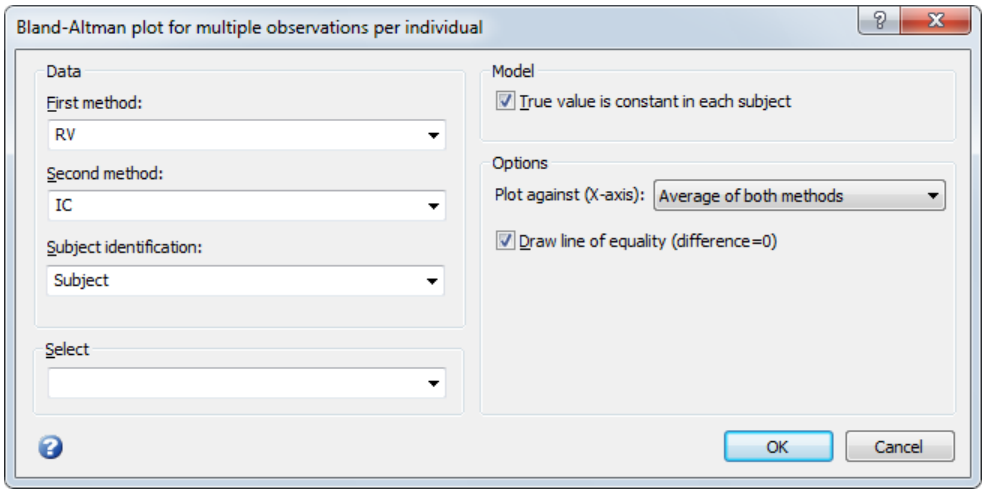

- **Data**

  First method, Second method: Select the variables for the two techniques you want to compare.

  Subject identification: Select the variable that contains the subject identification.

- **Model**

  True value is constant in each subject: Select this option if the true value is constant in each subject (e.g. with both methods several measurements were performed on the same sample).

  In the True value is constant in each subject model (see Bland & Altman, 2007) there is only one marker for each subject in the graph. The marker size is relative to the number of observations for the subject. The number of markers is equal to the number of subjects.

  In the alternative model, where the **True value varies**, there is one marker for each observation pair.

- **Options**

  Plot against (X-axis):

  In the original Bland-Altman plot (Bland & Altman, 1986) the differences between the two methods are plotted against the averages of the two methods (**recommended**, Bland & Altman, 1995).

  Alternatively, you can choose to plot the differences against one of the two methods, if this is a reference or "gold standard" method (Krouwer, 2008). Finally, you can also plot the differences against the geometric mean of both methods.

  Draw line of equality: useful for detecting a systematic difference.

Graph

This is the graph in the **True value is constant in each subject** model:

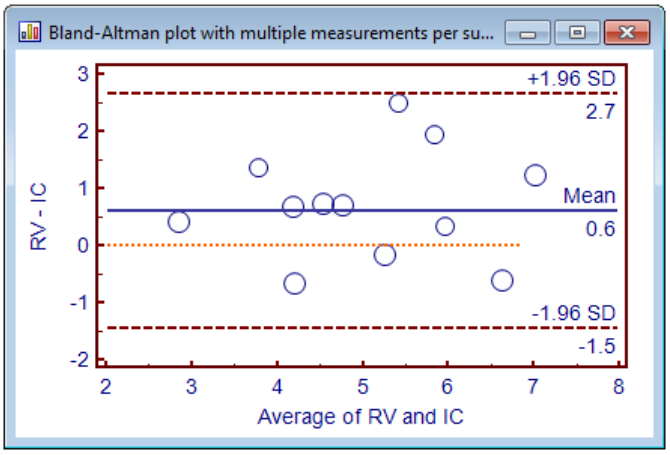

In the **True value is constant in each subject** model (see Bland & Altman, 2007) there is only one marker for each subject in the graph, and the marker size is relative to the number of observations for the subject. The number of markers is equal to the number of subjects.

In the alternative model, where the **True value varies**, there is one marker for each observation pair:

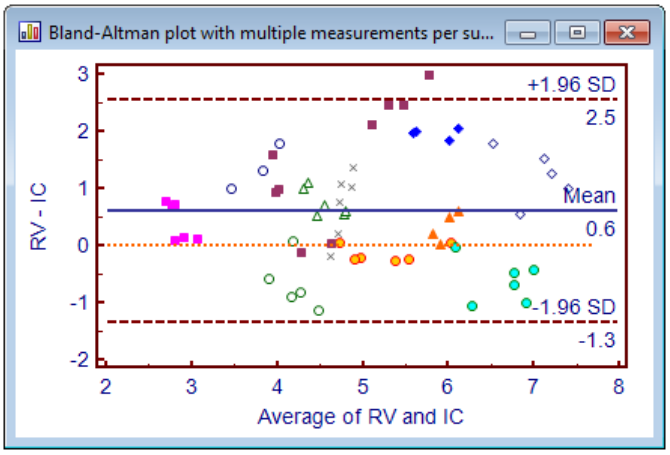

Mountain plot

A *mountain plot* (or "folded empirical cumulative distribution plot") is created by computing a percentile for each ranked difference between a new method and a reference method. To get a folded plot, the following transformation is performed for all percentiles above 50: percentile = 100 - percentile. These percentiles are then plotted against the differences between the two methods (Krouwer & Monti, 1995).

The mountain plot is a useful complementary plot to the Bland & Altman plot. In particular, the mountain plot offers the following advantages:

- It is easier to find the central 95% of the data, even when the data are not Normally distributed.
- Different distributions can be compared more easily.

In the dialog box for *Mountain plot* you enter the variables for the two or three techniques you want to compare.

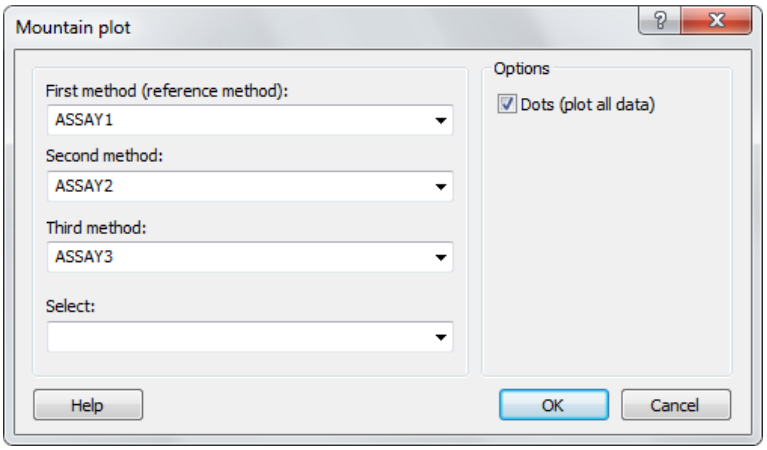

You can enter variables for two or three laboratory assays.

In case of three assays, the second and third assay will be compared with the first reference assay.

You can select an option to include all the data points in the graph. This option is very useful to identify outliers.

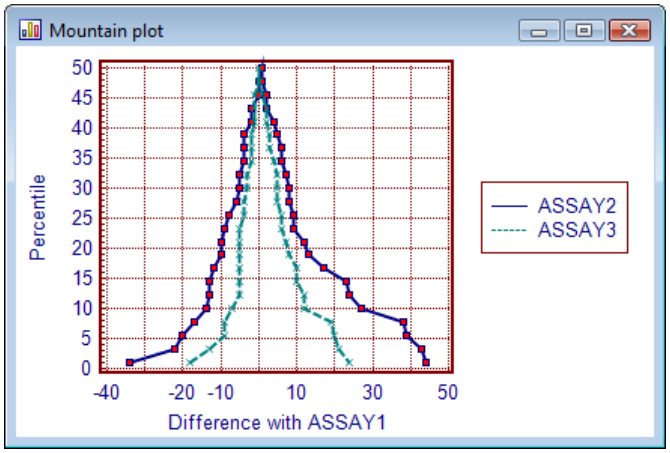

The mountain plot provides information about the distribution of the differences between methods.

If two assays are unbiased with respect to each other, the mountain will be centered over zero. Long tails in the plot reflect large differences between the methods.

In the example, the median of the differences is close to zero for both methods Assay2 and Assay3 as compared with Assay 1. The differences of Assay 1 with Assay3 tend to be smaller than the differences of Assay 1 with Assay2. Therefore Assay3 corresponds better with Assay1 than Assay2 does.

When you click an individual observation in the graph, the corresponding case is identified in a popup window (see also *Select variable for case identification* command, p. 54).

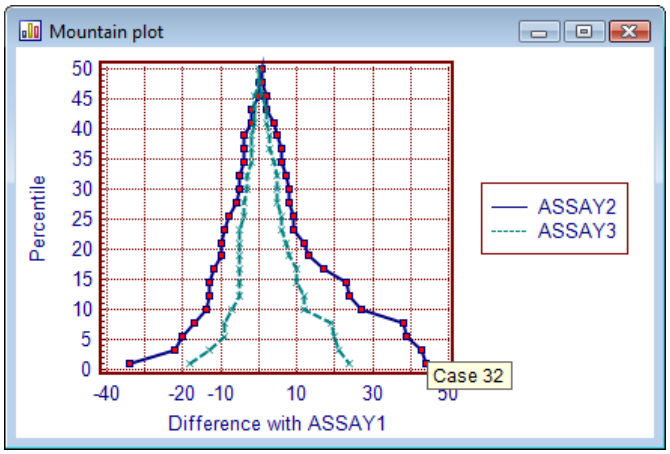

If you double-click an observation, the spreadsheet window will open with the corresponding case highlighted.

Select *Info* in the shortcut menu (which appears after right-clicking in the graph window) to get precise information on sample size, median, minimum and maximum and the most important percentiles of the distribution(s).

Deming regression

Whereas the ordinary linear regression method assumes that only the Y measurements are associated with random measurement errors, the Deming method takes measurement errors for both methods into account.

Required input

Select the variables for the two techniques you want to compare.

For each of both techniques you can either enter 2 variables (which contain repeated measurements) or you can enter only one variable, in which case you will have to enter an already established Coefficient of Variation (expressed as a percentage).

As an option, you can create 2 graphs:

- A scatter diagram with the regression line
- The *residuals* plot.

Use the **Subgroups** button if you want to identify subgroups in the scatter diagram and residuals plot. A new dialog box is displayed in which you can select a categorical variable. The graph will display different markers for the different categories in this variable.

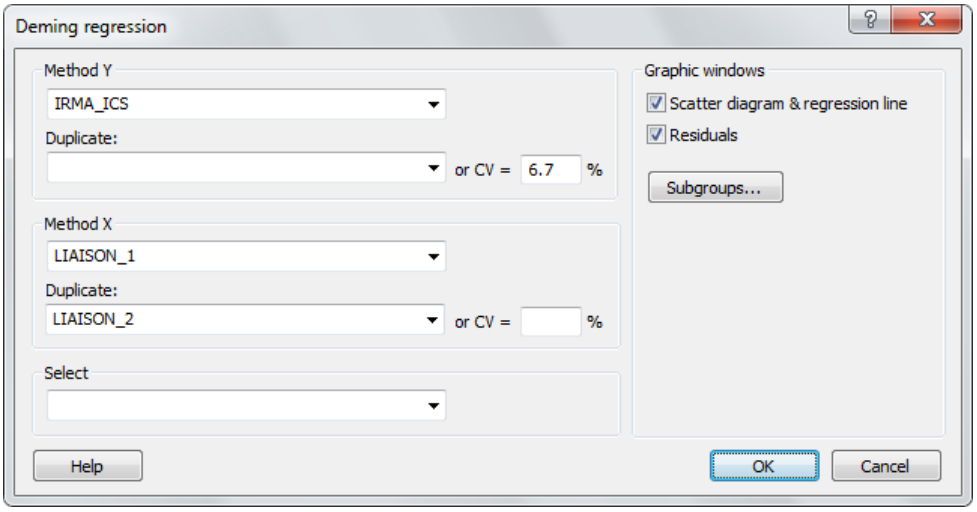

Results

The results are displayed in the following text window:

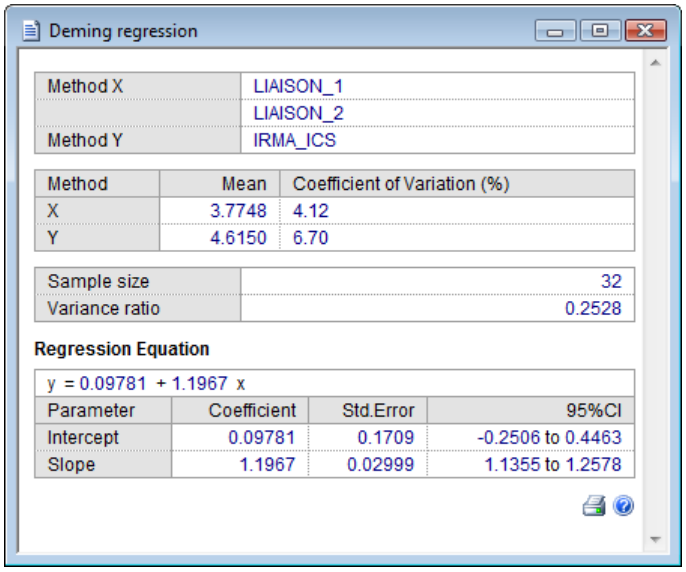

- Mean and Coefficient of Variation (%) for both methods
- Sample size: the number of (selected) data pairs
- Variance ratio: this is the ratio of the measurement errors of X and Y.
- The regression equation, Intercept and Slope with 95% confidence interval

  The Intercept and Slope are calculated according to Combleet & Gochman, 1979. The standard errors and confidence intervals are estimated using the jackknife method (Armitage et al., 2002).

  The 95% confidence interval for the Intercept can be used to test the hypothesis that A=0. This hypothesis is accepted if the confidence interval for A contains the value 0. If the hypothesis is rejected, then it is concluded that A is significant different from 0 and both methods differ at least by a constant amount.

  The 95% confidence interval for the Slope can be used to test the hypothesis that B=1. This hypothesis is accepted if the confidence interval for B contains the value 1. If the hypothesis is rejected, then it is concluded that B is significant different from 1 and there is at least a proportional difference between the two methods.

Graphs

The first graph window displays the scatter diagram with regression line (solid line) and identity line (x=y, dotted line)

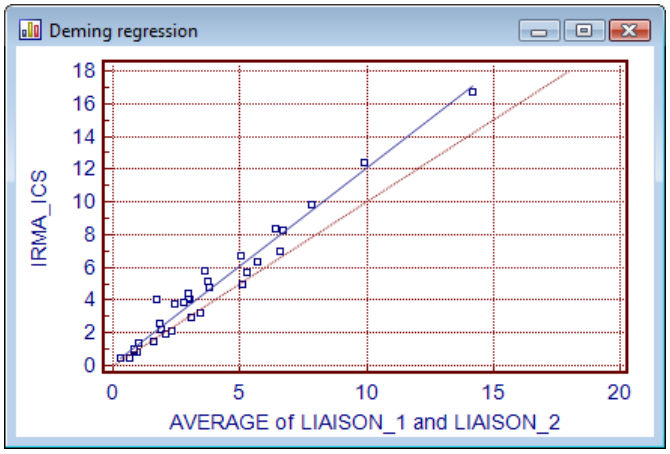

The second graph window displays the residuals:

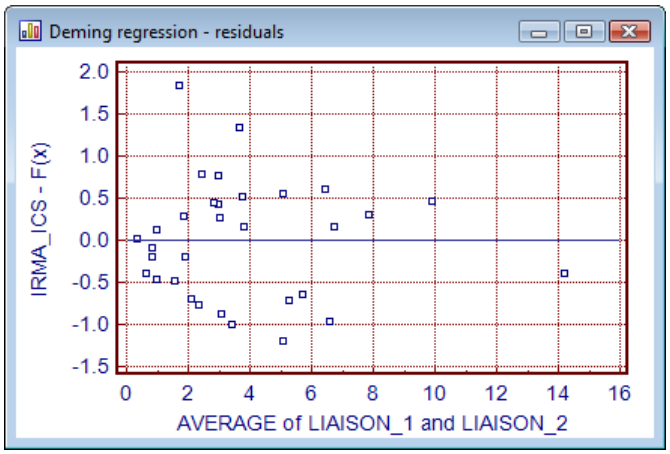

Passing & Bablok regression

Passing & Bablok (1983) have described a linear regression procedure with no special assumptions regarding the distribution of the samples and the measurement errors. The result does not depend on the assignment of the methods (or instruments) to X and Y. The slope B and intercept A are calculated with their 95% confidence interval. These confidence intervals are used to determine whether there is only a chance difference between B and 1 and between A and 0.

Required input

In the dialog box you enter the variables for the two techniques you want to compare.

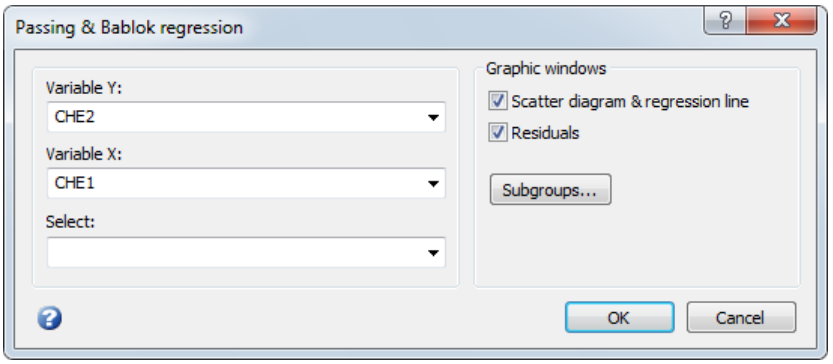

As an option, you can create 2 graphs:

- A scatter diagram with the regression line, the confidence interval for the regression line and identity line
- The *residuals* plot.

Use the **Subgroups** button if you want to identify subgroups in the scatter diagram and residuals plot. A new dialog box is displayed in which you can select a categorical variable. The graph will display different markers for the different categories in this variable.

Results

When you have completed the dialog box, click the OK button to proceed. The following results will be displayed in a text window.

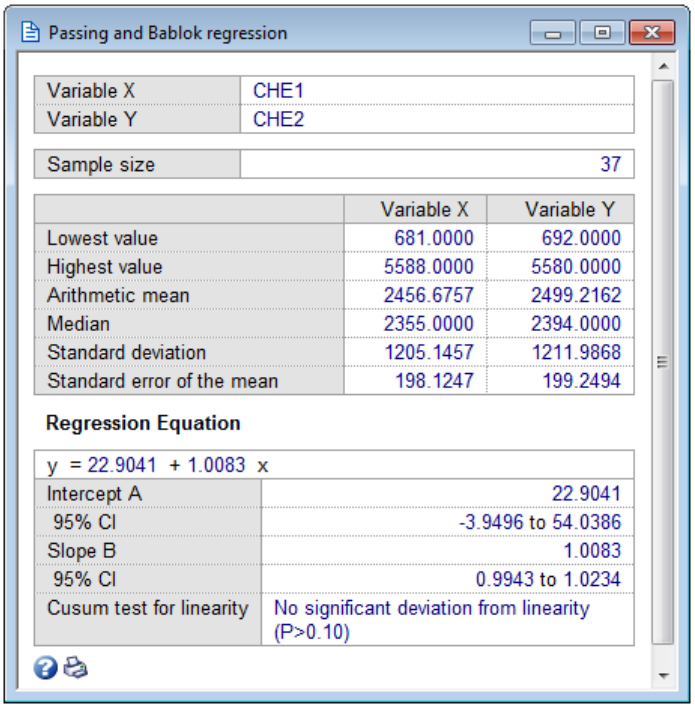

- **Sample size**: the number of (selected) data pairs
- **Summary statistics** for both variables: lowest and highest value, mean, median, standard deviation and standard error of the mean
- **The regression equation**: the regression equation with the calculated values for A and B according to Passing & Bablok (1983)
- **Intercept A and slope B with 95% confidence interval**:

  The 95% confidence interval for the intercept A can be used to test the hypothesis that A=0. This hypothesis is accepted if the confidence interval for A contains the value 0. If the hypothesis is rejected, then it is concluded that A is significant different from 0 and both methods differ at least by a constant amount.

  The 95% confidence interval for the slope B can be used to test the hypothesis that B=1. This hypothesis is accepted if the confidence interval for B contains the value 1. If the hypothesis is rejected, then it is concluded that B is significant different from 1 and there is at least a proportional difference between the two methods.
- **Test for linearity**: the test for linearity is used to evaluate how well a linear model fits the data.

Graphs

This is the scatter diagram with the regression line (solid line), the confidence interval for the regression line (dashed lines) and identity line (x=y, dotted line):

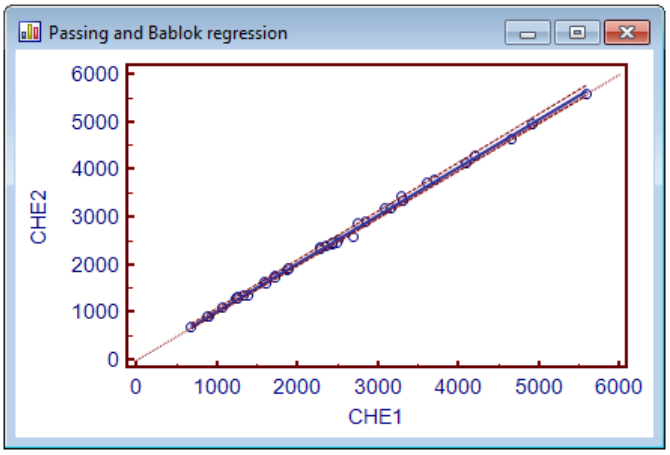

A second graph window shows the residuals:

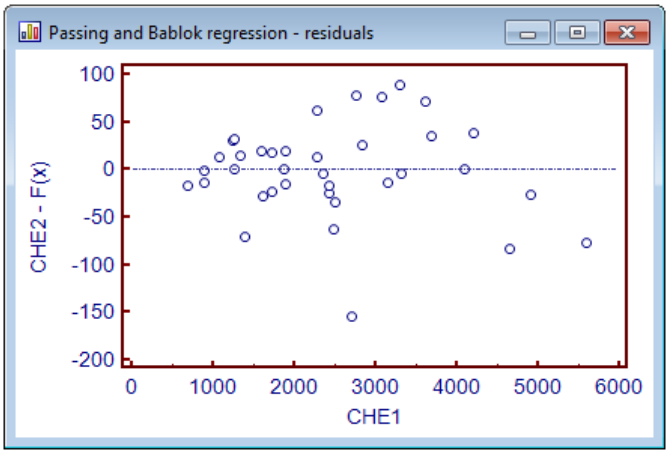

Coefficient of variation from duplicate measurements

The calculation of the Standard Deviation (SD) and Coefficient of Variation (CV) from duplicate measurements made on a number of different subjects or materials is used to determine the reproducibility of the measurements as an alternative to making a large number of observations on a single subject or material to calculate the SD and CV directly (Jones & Payne 1997).

Required input

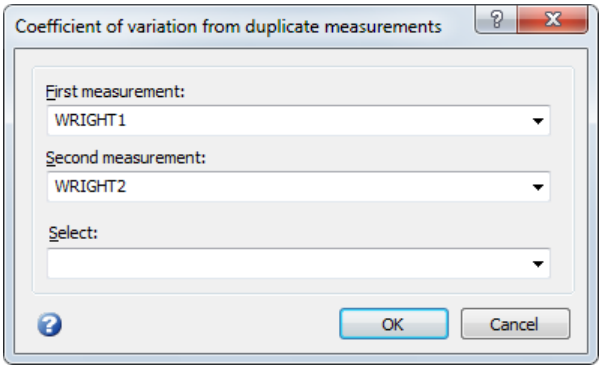

In the dialog box you select the variables that contain the data for the two measurements.

Results

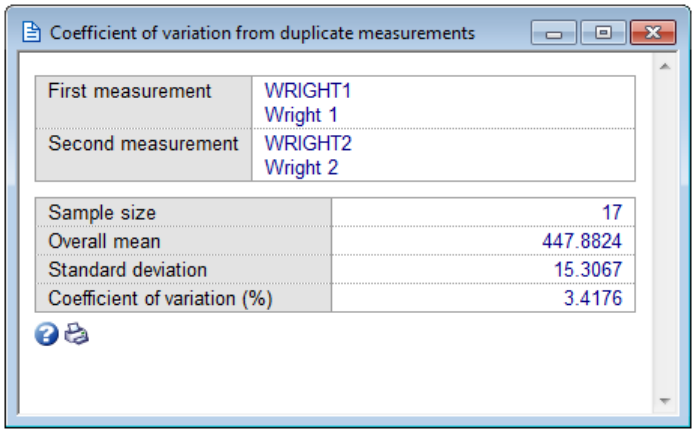

With *n* being the number of data pairs and x_1 and x_2 duplicate measurements, the SD, Mean and CV are given by:

$$SD = \sqrt{\frac{\sum (x_1 - x_2)^2}{2n}}$$

$$Mean = \frac{\sum (x_1 + x_2)}{2n}$$

$$CV\ (\%) = 100 \times \frac{SD}{Mean}$$

Intraclass correlation coefficient

The Intraclass Correlation Coefficient (ICC) is a measure of the reliability of measurements or ratings.

For the purpose of assessing inter-rater reliability and the ICC, two or preferably more raters rate a number of study subjects.

A distinction is made between two study models: (1) each subject is rated by a different and random selection of a pool of raters, and (2) each subject is rated by the same raters.

In the first model, the ICC is always a measure for Absolute agreement; in the second model a choice can be made between two types: Consistency when systematic differences between raters are irrelevant, and Absolute agreement, when systematic differences are relevant.

For example: the paired ratings (2,4), (4,6) and (6,8) are in perfect agreement, with a consistency of 1.0, but with an absolute agreement of 0.6667.

How to enter data

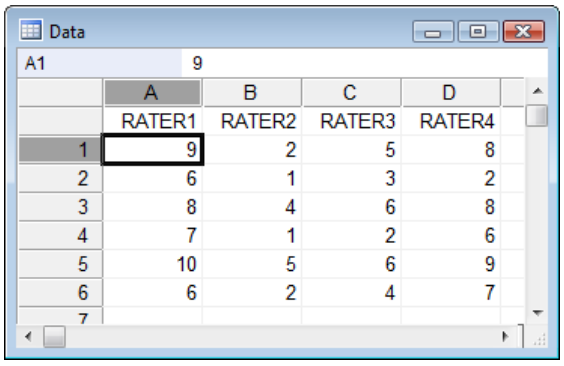

In this example (taken from Shrout PE & Fleiss JL, 1979) data are available for 4 raters on 6 subjects. The data for each subject are entered in the 4 columns.

If not all subjects are rated by the same 4 raters, the data are still entered in 4 columns, the order of which then being unimportant.

Required input

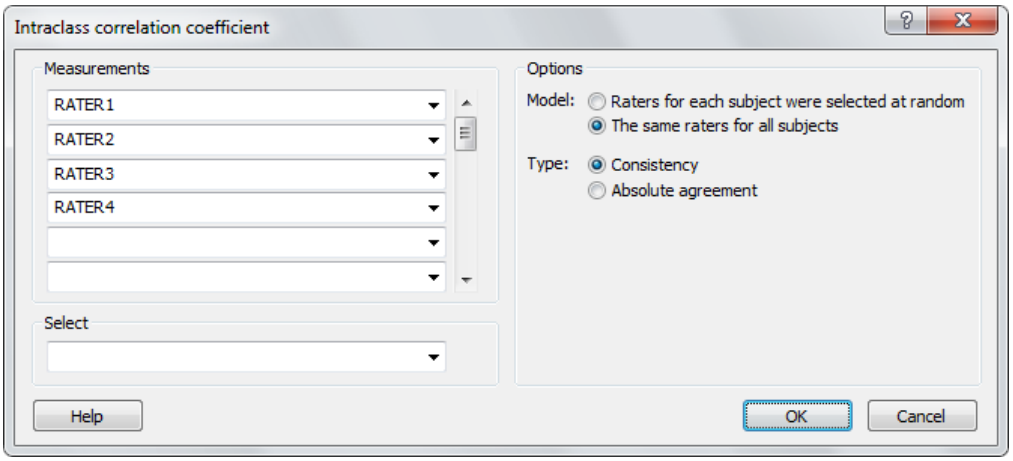

- Measurements: variables that contain the measurements of the different raters.
- Select: an optional selection criterion to include only a selected subgroup of cases.
- **Options**
  - **Model**
    - Raters for each subject were selected at random: the raters were not the same for all subjects, a random selection or raters rated each subject.
    - The same raters for all subjects: all subjects were rated by the same raters.
  - **Type**
    - Consistency: systematic differences between raters are irrelevant.
    - Absolute agreement: systematic differences are relevant

Results

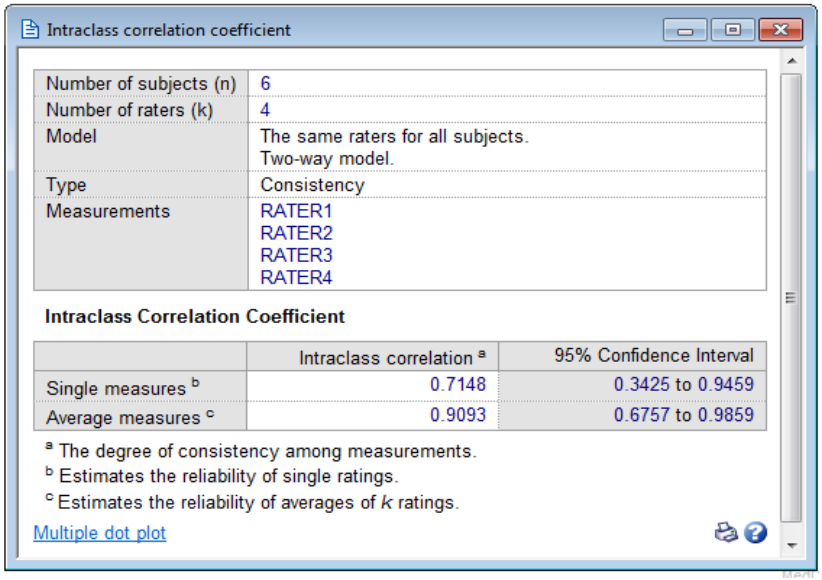

The Intraclass correlation coefficient table reports two coefficients with their respective 95% Confidence Interval.

- **Single measures**: this ICC is an index for the reliability of the ratings for one, typical, single rater.
- **Average measures**: this ICC is an index for the reliability of different raters averaged together. This ICC is always higher than the Single measures ICC.

Literature

- McGraw KO, Wong SP (1996) Forming inferences about some intraclass correlation coefficients. Psychological Methods, 1:30-46. (Correction: 1:390).
- Shrout PE, Fleiss JL (1979) Intraclass correlations: uses in assessing rater reliability. Psychological Bulletin, 86:420-428.

Concordance correlation coefficient

The concordance correlation coefficient (Lin, 1989) evaluates the degree to which pairs of observations fall on the 45° line through the origin.

Required input

Select the variables for the two techniques you want to compare.

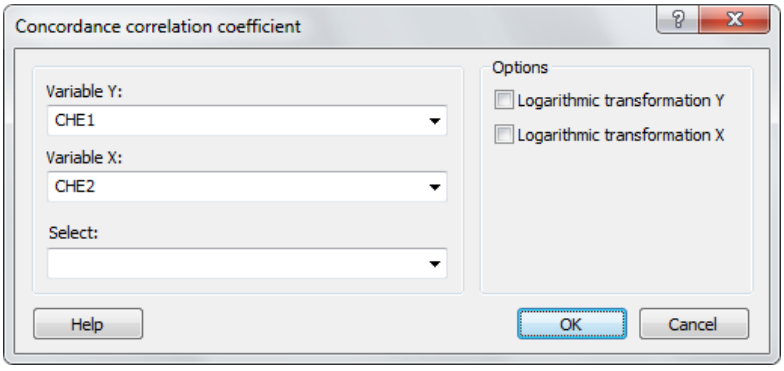

Results

The concordance correlation coefficient ρ_c contains a measurement of precision ρ and accuracy C_b:

$$\rho c = \rho\ Cb$$

where

- ρ is the Pearson correlation coefficient, which measures how far each observation deviates from the best-fit line, and is a **measure of precision**, and
- C_b is a bias correction factor that measures how far the best-fit line deviates from the 45° line through the origin, and is a **measure of accuracy**.

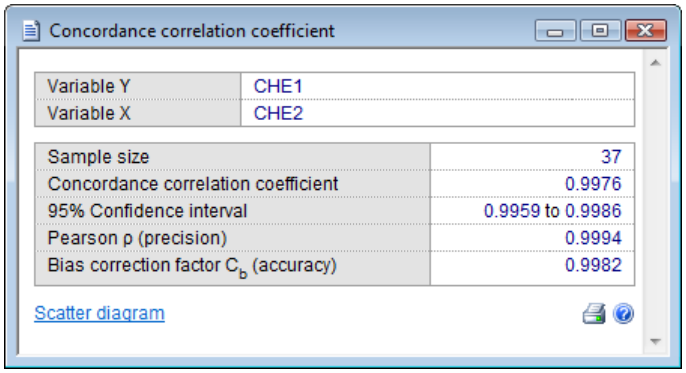

Literature

- Lin L.I-K (1989) A concordance correlation coefficient to evaluate reproducibility. Biometrics 45:255–268.
- Lin L.I-K (2000) A note on the concordance correlation coefficient. Biometrics 56:324–325.

Inter-rater agreement (kappa)

Inter-rater agreement is used to evaluate the agreement between two classifications (nominal or ordinal scales).

Agreement is quantified by the *Kappa* (***K***) statistic (Cohen, 1960; Fleiss et al., 2003). ***K*** is 1 when there is perfect agreement between the classification system; ***K*** is 0 when there is no agreement better than chance; and ***K*** is negative when agreement is worse than chance.

Data input

If you have the data already organized in a table, you can use the *Inter-rater agreement* command in the *Tests* menu (see p. 195).

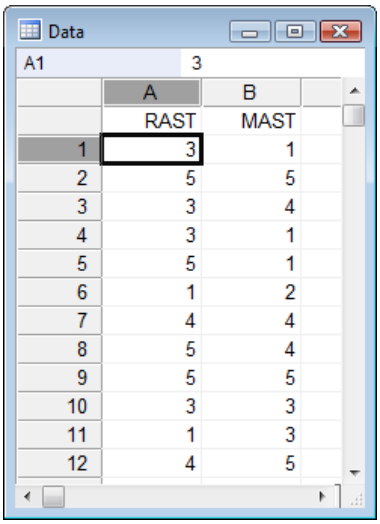

Dialog box

In the *Inter-rater agreement* dialog box, two discrete variables with the classification data from the two methods or observers must be identified. Classification data may either be numeric or alphanumeric (string) values.

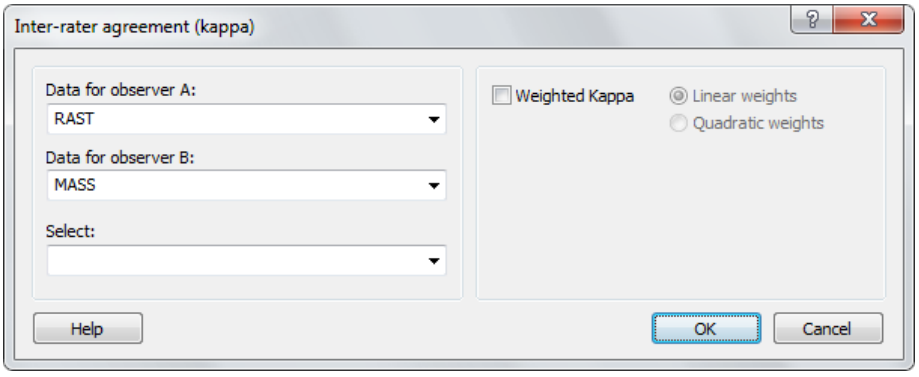

Weighted Kappa

Kappa does not take into account the degree of disagreement between observations and all disagreement is treated equally as *total* disagreement. Therefore when the categories are ordered, it is preferable to use *Weighted Kappa*, and assign different weights w_i to subjects for whom the raters differ by i categories, so that different levels of agreement can contribute to the value of Kappa.

MedCalc offers two sets of weights, called linear and quadratic. In the linear set, if there are k categories, the weights are calculated as follows:

$$w_i = 1 - \frac{i}{k-1}$$

and in the quadratic set :

$$w_i = 1 - \frac{i^2}{(k-1)^2}$$

When there are 5 categories, the weights in the linear set are 1, 0.75, 0.50, 0.25 and 0 when there is a difference of 0 (=total agreement) or 1, 2, 3 and 4 categories respectively. In the quadratic set the weights are 1, 0.937, 0.750, 0.437 and 0.

Use linear weights when the difference between the first and second category has the same importance as a difference between the second and third category, etc. If the difference between the first and second category is less important than a difference between the second and third category, etc., use quadratic weights.

Results

MedCalc calculates the inter-rater agreement statistic Kappa with 95% confidence interval.

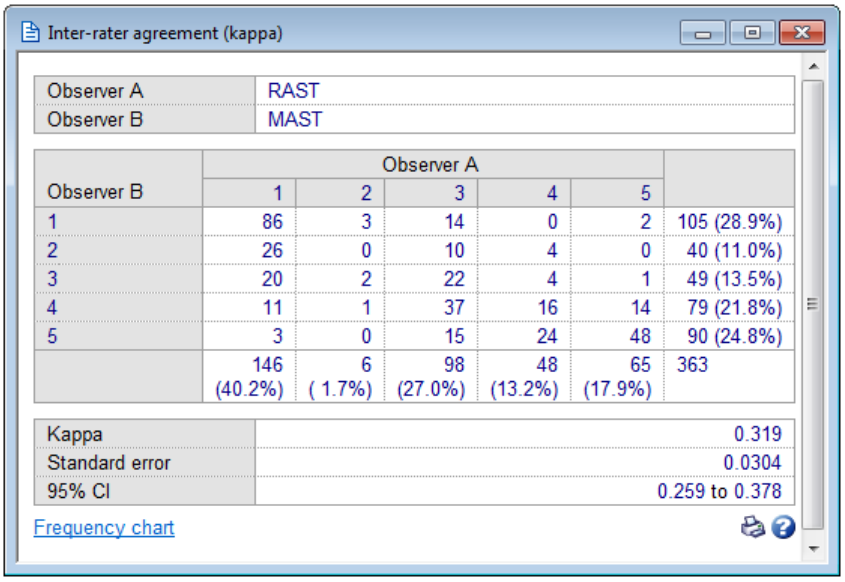

Inter-rater agreement (kappa)

Observer A	RAST
Observer B	MAST

Observer B	Observer A 1	2	3	4	5	
1	86	3	14	0	2	105 (28.9%)
2	26	0	10	4	0	40 (11.0%)
3	20	2	22	4	1	49 (13.5%)
4	11	1	37	16	14	79 (21.8%)
5	3	0	15	24	48	90 (24.8%)
	146 (40.2%)	6 (1.7%)	98 (27.0%)	48 (13.2%)	65 (17.9%)	363

Kappa	0.319
Standard error	0.0304
95% CI	0.259 to 0.378

Frequency chart

MedCalc calculates the inter-rater agreement statistic Kappa with 95% confidence interval (Fleiss et al., 2003). The Standard errors reported by MedCalc are the appropriate standard errors for testing the hypothesis that the underlying value of weighted kappa is equal to a prespecified value other than zero (Fleiss, 2003).

The ***K*** value can be interpreted as follows (Altman, 1991):

Value of ***K***	Strength of agreement
< 0.20	Poor
0.21 - 0.40	Fair
0.41 - 0.60	Moderate
0.61 - 0.80	Good
0.81 - 1.00	Very good

Cronbach's alpha

Cronbach's alpha is a statistic for investigating the internal consistency of a questionnaire (Cronbach, 1951; Bland & Altman, 1997).

How to enter data

Each question of the questionnaire results in one variable and the answers (numerically coded) are entered in the respective columns of the spreadsheet. The answers of one subject are entered on one row of the spreadsheet.

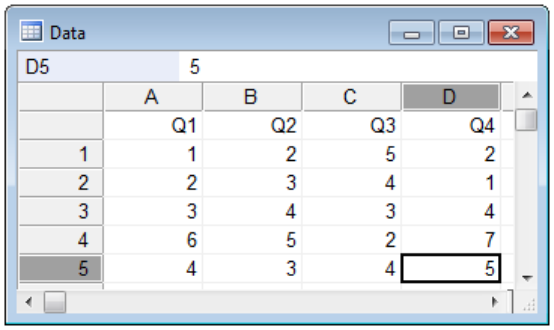

Required input

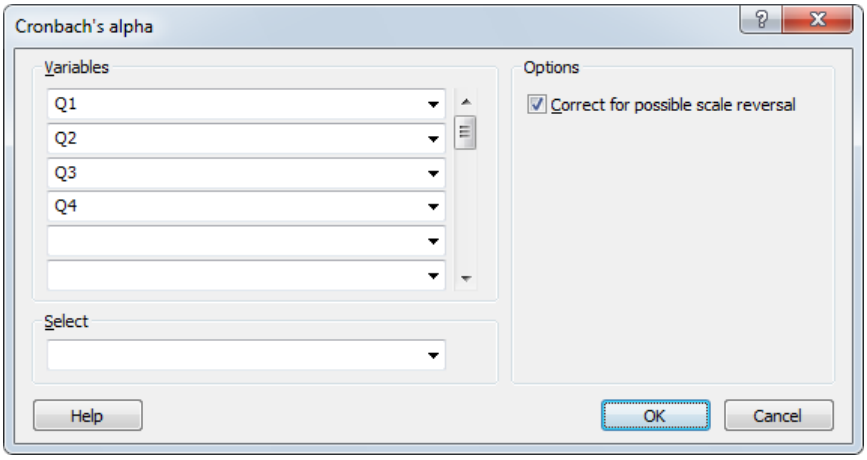

- **Variables**: the variables that contain the answers to the different questions of the questionnaire.
- **Select**: an optional selection criterion to include only a selected subgroup of subjects (rows).
- **Options**

  Correct for scale reversal: Some variables may be inversely related to other variables. When you select the option "Correct for scale reversal", MedCalc will detect these variables automatically (based on the correlation matrix) and reverse the values of those variables before analysis.

Results

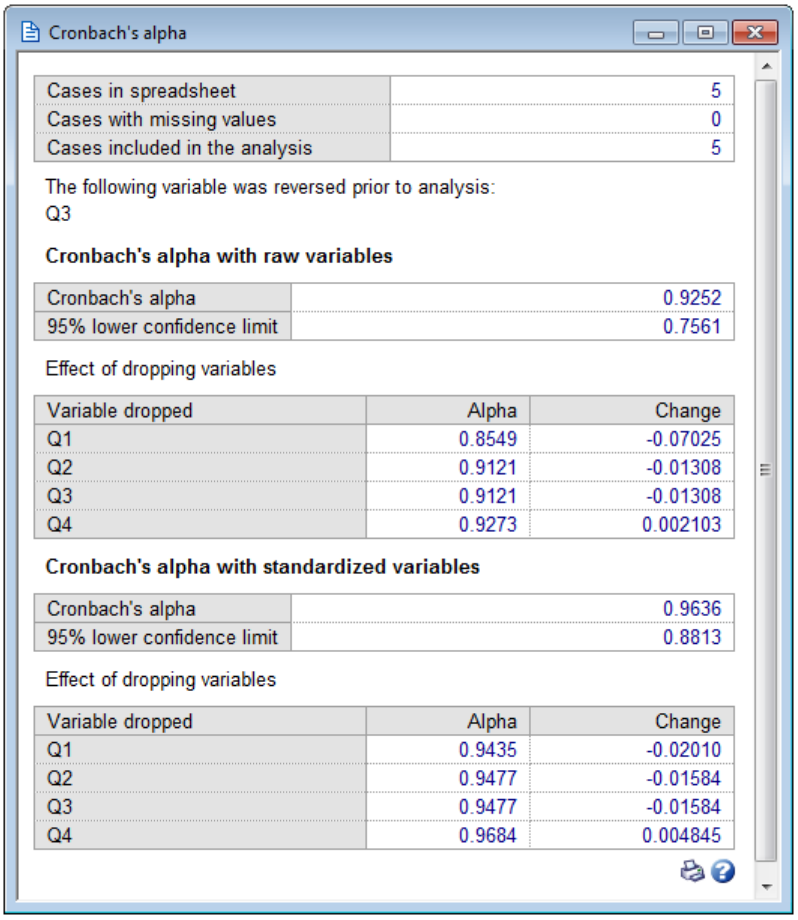

In the example, the results of Question 3 were found to be inversely related to the results of the other questions. Therefore the results of Question 3 were reversed prior to analysis.

MedCalc reports Cronbach's alpha with its lower confidence limit (Feldt, 1965).

Next, MedCalc calculates the alpha that obtained with each question in turn dropped. If the deletion of a question causes a considerable increase in alpha then you should consider dropping that question from the questionnaire.

MedCalc calculates Cronbach's alpha using the raw data and on the standardized variables (a transformation so that their mean is 0 and variance is 1). Using the "raw" data, questions that have more variability contribute more to the variability of the resulting scale; in the "standardized" form, each question gets equal weight.

For research purposes alpha should be more than 0.7 to 0.8, but for clinical purposes alpha should at least be 0.90 (Bland & Altman, 1997).

Literature

- Cronbach LJ (1951) Coefficient alpha and the internal structure of tests. Psychometrika 16:297-334.
- Bland JM, Altman DG (1997) Statistics notes: Cronbach's alpha. British Medical Journal 314:572.
- Feldt LS (1965) The approximate sampling distribution of Kuder-Richardson reliability coefficient twenty. Psychometrika 30:357-371.

Responsiveness

This allows to calculate several indices for responsiveness, which is the ability to detect any change.

Required input

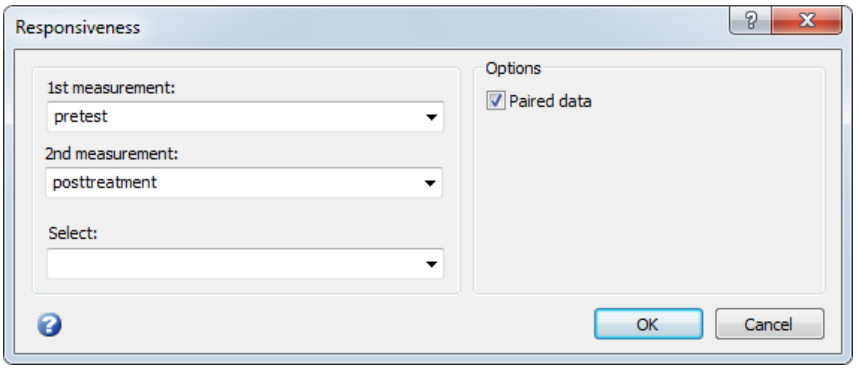

- **1st and 2nd measurement**: the variables for a 1st and 2nd measurement.
- **Select**: an optional selection criterion to include only a selected subgroup of subjects (rows).
- **Options**
  - **Paired data**: select this option when the variables for 1st and 2nd measurements contain paired data (measurements are repeated on the same subjects). If the 2 measurements are independent, unselect this option.

Results

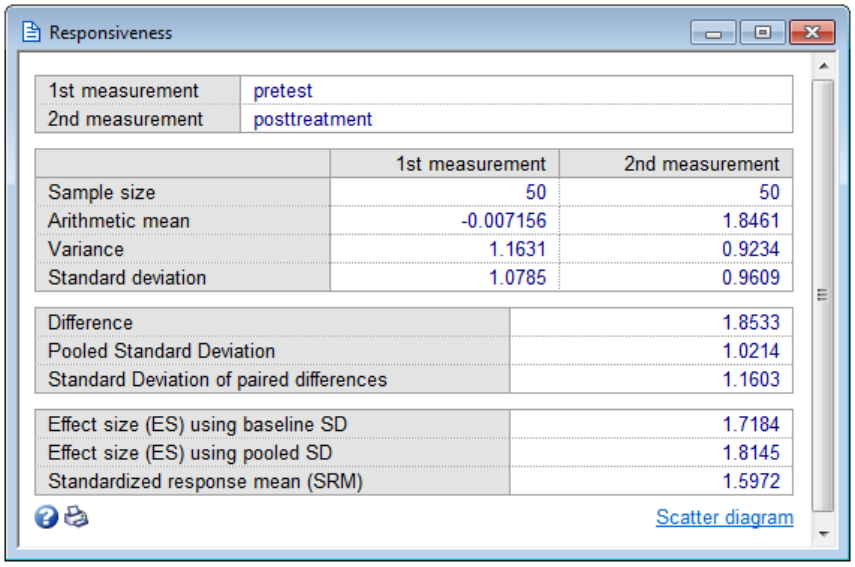

Summary statistics

The sample size, mean, variance and standard deviation of the 1st and 2nd measurement.

Difference

The average difference between the two measurements with the pooled standard deviation and (in case of paired observations) the standard deviation of the paired differences.

Indices of responsiveness

- Effect size (ES) using baseline SD: this is the average difference divided by the standard deviation of the 1st measurement.

- Effect size (ES) using pooled SD: this is the average difference divided by the pooled standard deviation of both measurements.
- Standardized response mean (SRM) : this is the average difference divided by the standard deviation of the differences between the paired measurements.

Literature

- Husted JA, Cook RJ, Farewell VT, Gladman DD (2000) Methods for assessing responsiveness: a critical review and recommendations. Journal of Clinical Epidemiology 53:459-168.
- Norman GR, Wyrwich KW, Patrick DL (2007) The mathematical relationship among different forms of responsiveness coefficients. Quality of Life Research 16:815-822.

Receiver Operating Characteristic (ROC) curve analysis

Introduction

Often the clinical researcher is confronted with the question how accurate a particular laboratory test is in identifying diseased cases. The ability of a test to discriminate diseased cases from normal cases is evaluated using Receiver Operating Characteristic (ROC) curve analysis (Metz, 1978; Zweig & Campbell, 1993). ROC curves can also be used to compare the diagnostic performance of two or more laboratory or diagnostic tests (Griner et al., 1981).

When you consider the results of a particular test in two populations, one population with a disease, the other population without the disease, you will rarely observe a perfect separation between the two groups. Indeed, the distribution of the test results will overlap, as shown in the following figure.

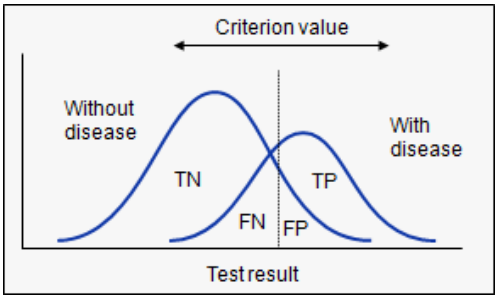

For every possible cut-off point or criterion value you select to discriminate between the two populations, there will be some cases with the disease correctly classified as positive (TP = True Positive fraction), but some cases with the disease will be classified negative (FN = False Negative fraction). On the other hand, some cases without the disease will be correctly classified as negative (TN = True Negative fraction), but some cases without the disease will be classified as positive (FP = False Positive fraction).

Table: Schematic outcomes of a test.

	Disease				
Test	**Present**	**n**	**Absent**	**n**	**Total**
Positive	True Positive	*a*	False Positive	*c*	*a* + *c*
Negative	False Negative	*b*	True Negative	*d*	*b* + *d*
Total		*a* + *b*		*c* + *d*	

Sensitivity $= \frac{a}{a+b}$	Specificity $= \frac{d}{c+d}$
Positive Likelihood Ratio $= \frac{Sensitivity}{1-Specificity}$	Negative Likelihood Ratio $= \frac{1-Sensitivity}{Specificity}$

Positive Predictive Value $= \frac{a}{a+c}$	Negative Predictive Value $= \frac{d}{b+d}$

- *Sensitivity*: probability that a test result will be positive when the disease is present (true positive rate, expressed as a percentage).
- *Specificity*: probability that a test result will be negative when the disease is not present (true negative rate, expressed as a percentage).
- *Positive likelihood ratio*: ratio between the probability of a positive test result given the *presence* of the disease and the probability of a positive test result given the *absence* of the disease, i.e.

$$\frac{\text{True Positive rate}}{\text{False Positive rate}}$$

- *Negative likelihood ratio*: ratio between the probability of a negative test result given the *presence* of the disease and the probability of a negative test result given the *absence* of the disease, i.e.

$$\frac{\text{False Negative rate}}{\text{True Negative rate}}$$

- *Positive predictive value*: probability that the disease is present when the test is positive (expressed as a percentage).
- *Negative predictive value*: probability that the disease is not present when the test is negative (expressed as a percentage).

How to enter data for ROC curve analysis

In order to perform ROC curve analysis in MedCalc you should have a measurement of interest (= the variable you want to study) and an independent diagnosis which classifies your study subjects into two distinct groups: a diseased and non-diseased group. The latter diagnosis should be independent from the measurement of interest.

In the spreadsheet, create a column DIAGNOSIS and a column for the variable of interest, e.g. TEST1. For every study subject enter a code for the diagnosis as follows: 1 for the diseased cases, and 0 for the non-diseased or normal cases. In the TEST1 column, enter the measurement of interest (this can be measurements, grades, etc. - if the data are categorical, code them with numerical values).

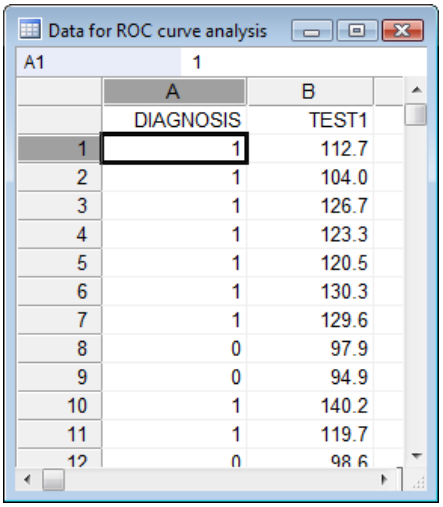

ROC curve analysis command

To obtain a ROC curve you first select the *ROC curve analysis* option in the menu and complete the following box:

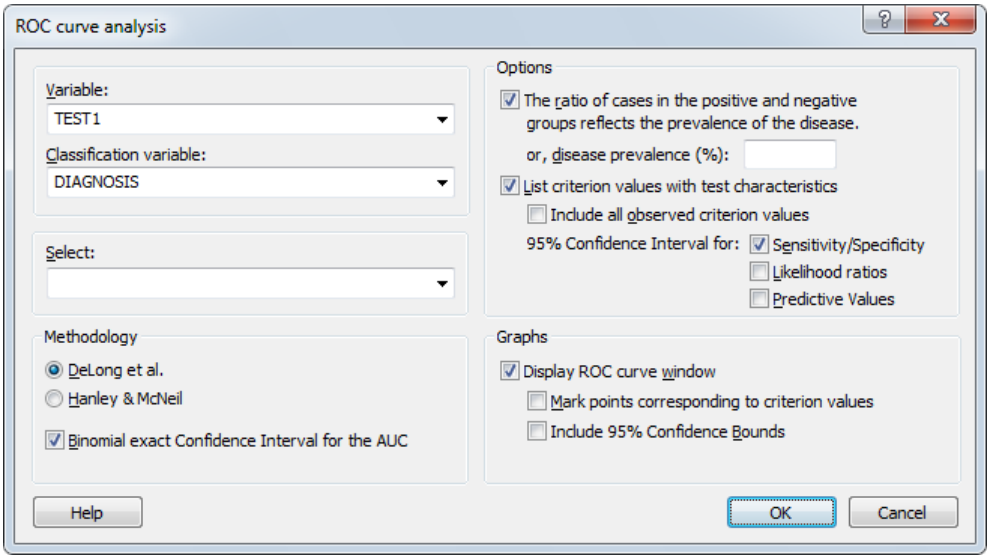

Data:

- **Variable**: identify the variable of interest.
- **Classification variable**: select or enter a dichotomous variable indicating diagnosis (0=negative, 1=positive). If diagnosis is coded differently than using the values 0 and 1, you can use the IF function to transform the codes into 0 and 1 values, e.g. IF(RESULT="pos",1,0).
- **Select**: (optionally) a selection criterion in order to include only a selected subgroup of cases (e.g. AGE>21, SEX="Male").

Methodology:

- **DeLong et al**.: use the method of DeLong et al. (1988) for the calculation of the Standard Error of the Area Under the Curve (recommended).
- **Hanley & McNeil**: use the method of Hanley & McNeil (1982) for the calculation of the Standard Error of the Area Under the Curve.
- **Binomial exact Confidence Interval for the AUC**: calculate an exact Binomial Confidence Interval for the Area Under the Curve (recommended). If this option is not selected, the Confidence Interval is calculated as AUC ± 1.96 its Standard Error.

Options:

- **Disease prevalence**: if the sample sizes in the positive and the negative group do not reflect the real prevalence of the disease, this must be indicated in the dialog box. You have then the option to enter the disease prevalence. Clinically, the disease prevalence is the same as the probability of disease being present before the test is performed. If the disease prevalence is unknown, or irrelevant for the current statistical analysis, you can ignore these fields. In this case the program will not calculate predictive values.
- **List criterion values with test characteristics:** option to create a list of criterion values corresponding with the coordinates of the ROC curve, with associated sensitivity, specificity, likelihood ratios and predictive values (if disease prevalence is known).
  - **Include all observed criterion values**: When you select this option, the program will list sensitivity and specificity for all possible threshold values. If this option is not selected, then the program will only list the more important points of the ROC curve: for equal sensitivity/specificity it will give the threshold values (criterion values) with the highest specificity/sensitivity.
- **95% Confidence Interval** for sensitivity/specificity, likelihood ratio and predictive values: select the Confidence Intervals you require.

Graphs:

- Select **Display ROC curve window** to obtain the ROC plot in a separate window.

  Options:
  - mark points corresponding to criterion values.

- display 95% Confidence Bounds for the ROC curve (Hilgers, 1991).

A few moments after you have clicked the OK button, the following appears in the results window. The report may consist of several pages of text, and press the *Page down* key to see the next pages of the report.

Results

First the program displays the number of observations in the two groups. Concerning sample size, it has been suggested that *meaningful qualitative conclusions* can be drawn from ROC experiments performed with a total of about 100 observations (Metz, 1978). A minimum of 50 cases may be required in each of the two groups, so that 1 case represents not more than 2% of the observations.

ROC curve

Variable	TEST1
Classification variable	DIAGNOSIS

Sample size	100
Positive group : DIAGNOSIS = 1	55
Negative group : DIAGNOSIS = 0	45

Disease prevalence (%)	55.0

Area under the ROC curve (AUC)	0.947
Standard Error [a]	0.0241
95% Confidence Interval [b]	0.883 to 0.982
z statistic	18.544
Significance level P (Area=0.5)	0.0001

[a] DeLong et al., 1988

[b] Binomial exact

Criterion values and coordinates of the ROC curve [Hide]

Criterion	Sensitivity	95% CI	Specificity	95% CI	+LR	-LR	+PV	-PV
>=77.3	100.00	93.5 - 100.0	0.00	0.0 - 7.9	1.00		55.0	
>94.9	100.00	93.5 - 100.0	37.78	23.8 - 53.5	1.61	0.00	66.3	100.0
>95	98.18	90.3 - 100.0	37.78	23.8 - 53.5	1.58	0.048	65.9	94.4
>102.5	98.18	90.3 - 100.0	66.67	51.0 - 80.0	2.95	0.027	78.3	96.8
>102.7	96.36	87.5 - 99.6	66.67	51.0 - 80.0	2.89	0.055	77.9	93.7
>103.2	96.36	87.5 - 99.6	73.33	58.1 - 85.4	3.61	0.050	81.5	94.3
>104	94.55	84.9 - 98.9	73.33	58.1 - 85.4	3.55	0.074	81.2	91.7
>104.5	94.55	84.9 - 98.9	75.56	60.5 - 87.1	3.87	0.072	82.5	91.9
>104.9	92.73	82.4 - 98.0	75.56	60.5 - 87.1	3.79	0.096	82.3	89.5
>108.3	92.73	82.4 - 98.0	86.67	73.2 - 94.9	6.95	0.084	89.5	90.7
>108.9 *	90.91	80.0 - 97.0	91.11	78.8 - 97.5	10.23	0.100	92.6	89.1
>110.7	81.82	69.1 - 90.9	91.11	78.8 - 97.5	9.20	0.20	91.8	80.4

The value for the *area under the ROC curve* can be interpreted as follows: an area of 0.84, for example, means that a randomly selected individual from the positive group has a test value larger than that for a randomly chosen individual from the negative group in 84% of the time (Zweig & Campbell, 1993). When the variable under study cannot distinguish between the two groups, i.e. where there is no difference between the two distributions, the area will be equal to 0.5 (the ROC curve will coincide with the diagonal). When there is a perfect separation of the values of the two groups, i.e. there no overlapping of the distributions, the area under the ROC curve equals 1 (the ROC curve will reach the upper left corner of the plot).

The *95% Confidence Interval* is the interval in which the true (population) Area under the ROC curve lies with 95% confidence.

The P-value is the probability that the sample Area under the ROC curve (0.947 in the example) is found when in fact, the true (population) Area under the ROC curve is 0.5 (null hypothesis: Area = 0.5). If P is low (P<0.05) then it can be concluded that the Area under the ROC curve is significantly different from 0.5 and that therefore there is evidence that the laboratory test does have an ability to distinguish between the two groups.

The next section of the results window lists the different selection criteria or cut-off values with their corresponding sensitivity and specificity of the test, and the positive (+LR) and negative likelihood ratio (-LR).

If the disease prevalence is known, the program also reports the positive predictive value (+PV) and negative predictive value (-PV).

When you did not select the option *Include all observed criterion values*, the program only lists the more important points of the ROC curve: for equal sensitivity (resp. specificity) it gives the threshold value (criterion value) with the highest specificity (resp. sensitivity). When you do select the option *Include all observed criterion values*, the program will list sensitivity and specificity for all possible threshold values.

The criterion value indicated with a * sign is the value corresponding with the highest average of sensitivity and specificity.

- When you select a *lower criterion value*, then the true positive fraction and sensitivity will increase. On the other hand the false positive fraction will also increase, and therefore the true negative fraction and specificity will decrease.
- When you select a *higher criterion value*, the false positive fraction will decrease with increased specificity but on the other hand the true positive fraction and sensitivity will decrease.
- If a test is used for the purpose of screening, then a cut-off value with a higher sensitivity and negative predictive value must be selected. To confirm the disease, the cases positive in the screening test will be tested again with a different test. In this second test, a high specificity and positive predictive value are required (Griner et al., 1981).

Importance of disease prevalence

Whereas *sensitivity* and *specificity*, and therefore the ROC plot, and *positive* and *negative likelihood ratio* are independent of the prevalence of the disease, *positive* and *negative predictive values* are highly dependent on the proportions of subjects who do and do not have the disease (prior probability of disease), and hence on the population studied.

Clinically, the disease prevalence is the same as the probability of disease being present before the test is performed.

If the sample sizes in the positive and the negative group do not correspond to the real prevalence of the disease, indicate this in the dialog box by deselecting the corresponding option:

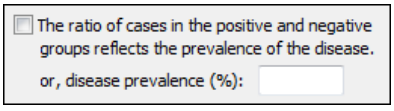

In this case the program will *not* calculate the *positive* and *negative predictive values*.

However, if you do know the disease prevalence in the population, you can enter the percentage in the dialog box:

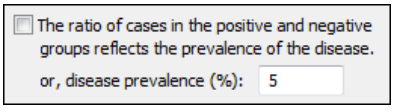

Display ROC curve

The ROC curve will be displayed in a second window when you have selected the corresponding option in the dialog box.

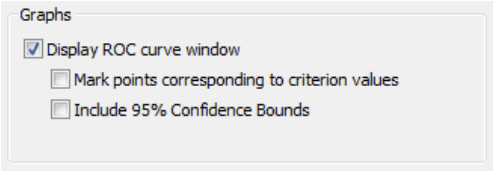

In a ROC curve the true positive rate (Sensitivity) is plotted in function of the false positive rate (100-Specificity) for different cut-off points. Each point on the ROC plot represents a sensitivity/specificity pair corresponding to a particular decision threshold. A test with perfect discrimination (no overlap in the two

distributions) has a ROC plot that passes through the upper left corner (100% sensitivity, 100% specificity). Therefore the closer the ROC plot is to the upper left corner, the higher the overall accuracy of the test (Zweig & Campbell, 1993).

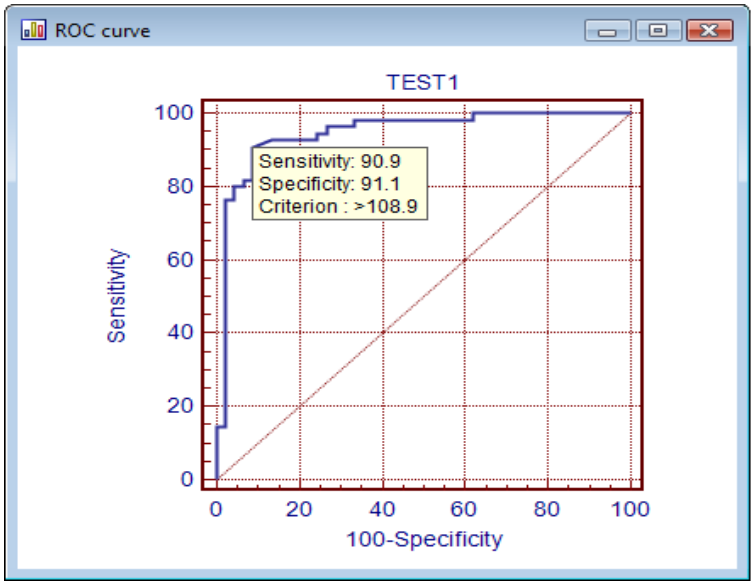

When you click on a specific point of the ROC curve, the corresponding cut-off point with sensitivity and specificity will be displayed in a small (non-printable) window.

Presentation of results

The prevalence of a disease may be different in different clinical settings. For instance the pre-test probability for a positive test will be higher when a patient consults a specialist than when he consults a general practitioner. Since positive and negative predictive values are sensitive to the prevalence of the disease, it would be misleading to compare these values from different studies where the prevalence of the disease differs, or apply them in different settings.

The data from the results window can be summarized in a table. The sample size in the two groups should be clearly stated. The table can contain a column for the different criterion values, the corresponding sensitivity (with 95% CI), specificity (with 95% CI), and possibly the positive and negative predictive value. The table should not only contain the test's characteristics for one single cut-off value, but preferably there should be a row for the values corresponding with a sensitivity of 90%, 95% and 99%, specificity of 90%, 95% and 99%, and the value corresponding with the highest accuracy (maximum sensitivity and specificity as indicated with a * mark in the results window).

With these data, any reader can calculate the negative and positive predictive value applicable in his own clinical setting when he knows the prior probability of disease (pre-test probability or prevalence of disease) in this setting, by the following formula's based on Bayes' theorem:

$$PPV = \frac{sensitivity \times prevalence}{sensitivity \times prevalence + (1-specificity) \times (1-prevalence)}$$

and

$$NPV = \frac{specificity \times (1-prevalence)}{(1-sensitivity) \times prevalence + specificity \times (1-prevalence)}$$

The negative and positive likelihood ratio must be handled with care because they are easily and commonly misinterpreted.

Interactive dot diagram

In addition to the ROC curve, MedCalc offers a second graphical display to study the accuracy of a diagnostic test, namely an *Interactive dot diagram*.

In the *Interactive dot diagram* the data of the negative and positive groups are displayed as dots on two vertical axes. Initially, a horizontal line indicates the cut-off point with the best separation (minimal false negative and false positive results) between the two groups. The corresponding test characteristics *sensitivity* and *specificity* are shown at the right side of the display.

Required input

Variable: identify the variables under study.

Classification variable: select or enter a dichotomous variable indicating diagnosis (0=negative, 1=positive). If diagnosis is coded differently than using the values 0 and 1, you can use the IF function to transform the codes into 0 and 1 values, e.g. IF(RESULT="pos",1,0).

Select: (optionally) a selection criterion in order to include only a selected subgroup of cases (e.g. AGE>21, SEX="Male").

Options: Log transformation of data.

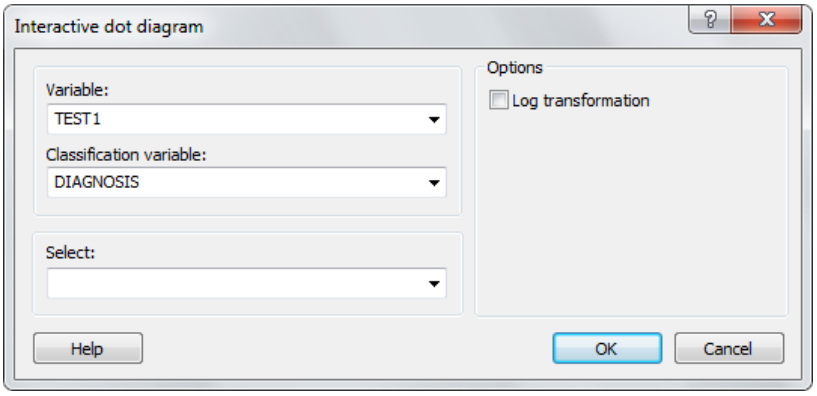

Graph

By clicking in the display, you can move the horizontal line to other values or cut-off points.

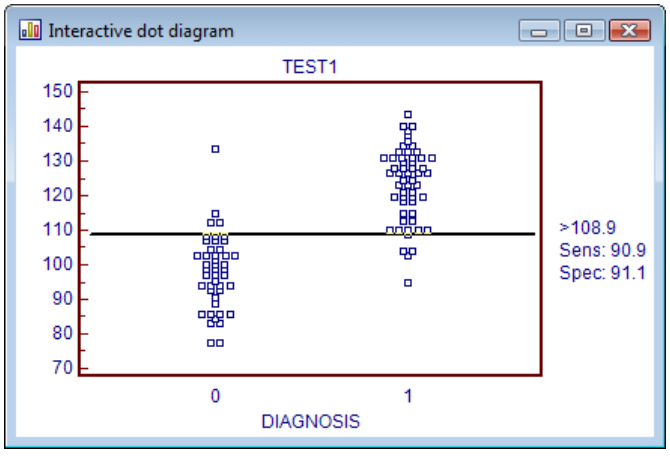

Plot versus criterion values

In this graph the sensitivity and specificity, and optionally their 95% Confidence Intervals, are plotted against the different criterion values.

Required input

Variable: identify the variables under study.

Classification variable: select or enter a dichotomous variable indicating diagnosis (0=negative, 1=positive). If diagnosis is coded differently than using the values 0 and 1, you can use the IF function to transform the codes into 0 and 1 values, e.g. IF(RESULT="pos",1,0).

Select: (optionally) a selection criterion in order to include only a selected subgroup of cases (e.g. AGE>21, SEX="Male").

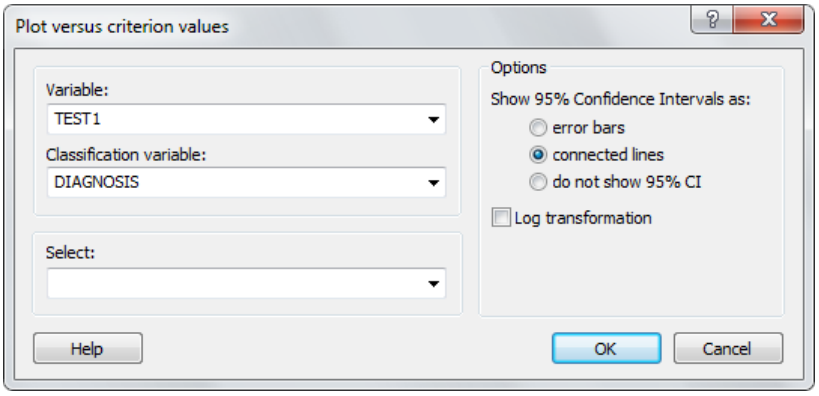

Options:

- Show 95% Confidence Intervals as:
  - **error bars**: show the 95% Confidence Interval of sensitivity and specificity as error bars
  - **connected lines**: show the 95% Confidence Interval of sensitivity and specificity as connected lines (recommended when number of criterion values is high)
  - **do not show CI**: do not show the 95% Confidence Interval of sensitivity and specificity in the graph.
- Log transformation of data.

Graph

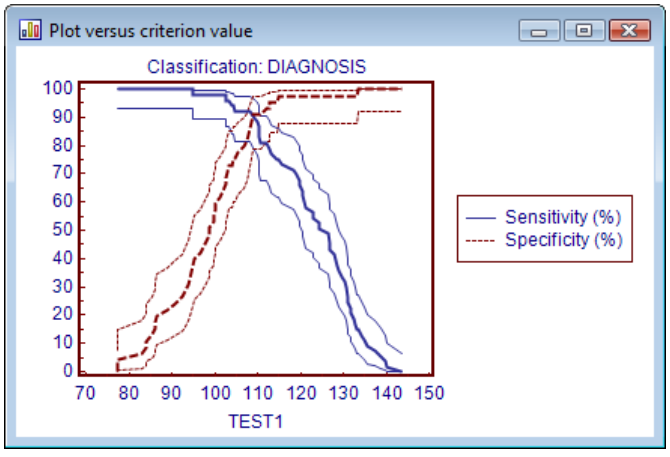

Predictive values

When you do not dispose of raw data to perform ROC curve analysis, you can still calculate positive and negative predictive values for a test when the sensitivity and specificity of the test as well as the disease prevalence (or the pretest probability of disease) are known, using Bayes' theorem (p. 161). In the *Predictive values* dialog box you enter the sensitivity, specificity and disease prevalence, all expressed as percentages. When these data are entered click TEST to see the results.

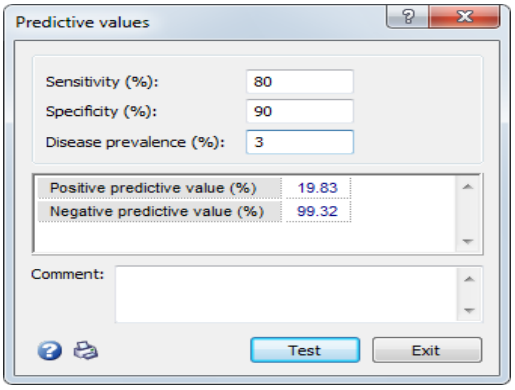

Interval likelihood ratios

Description

Allows to calculate the likelihood ratios (with 95% CI) for user-defined data intervals.

When test results have a continuous or ordinal outcome then valuable information is lost when the data are dichotomized for the calculation of sensitivity, specificity and likelihood ratios as in ROC curve analysis.

Interval likelihood ratios may be more powerful because they use more information contained in the data.

The likelihood ratio can be used to calculate the post-test probability of disease from the pre-test probability of disease.

Required input

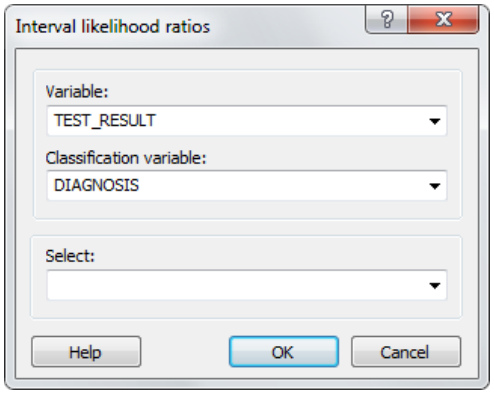

Variable: identify the variable under study.

Classification variable: select or enter a dichotomous variable indicating diagnosis (0=negative, 1=positive). If diagnosis is coded differently than using the values 0 and 1, you can use the IF function to transform the codes into 0 and 1 values, e.g. IF(RESULT="pos",1,0).

Select: (optionally) a selection criterion in order to include only a selected subgroup of cases (e.g. AGE>21, SEX="Male").

Define intervals

After some calculations, a new dialog box is displayed with suggested data intervals which you can modify.

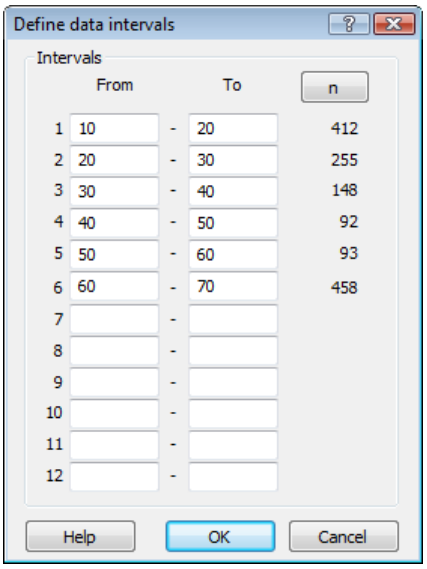

You can define up to 12 intervals. For each interval you enter the lower and upper (inclusive) boundaries.

For categorical variables, with few categories, it may suffice to enter only one number to define the "interval" as one single category.

Results

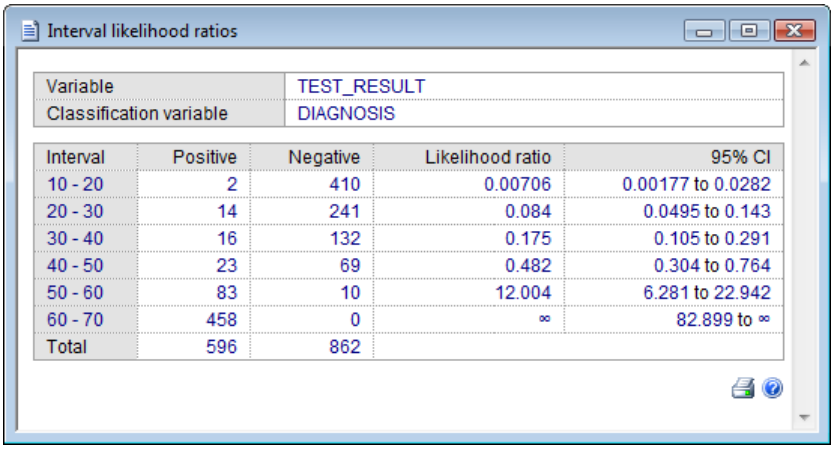

For each data interval the program reports the number of positive and negative cases in the interval and the corresponding Likelihood ratio with 95% Confidence interval.

The likelihood ratio can be used to calculate the post-test odds from the pre-test odds off disease:

post-test odds = pre-test odds x likelihood ratio

The relation between odds and probability is:

$$\text{odds} = \frac{p}{1-p} \quad \text{and} \quad p = \frac{\text{odds}}{1+\text{odds}}$$

Using these equations, you can calculate the post-test probability of disease from the pre-test probability of disease.

If, for example, the pre-test probability of disease is 0.6 then the pre-test odds is 0.6/(1-0.6) = 1.5. For a patient with test result in the interval 50-60, corresponding with a likelihood ratio of 12, the post-test odds are 1.5 x 12 = 18. The post-test probability of disease is 18/(1+18) = 0.95.

Literature

- Gardner IA, Greiner M (2006) Receiver-operating characteristic curves and likelihood ratios: improvements over traditional methods for the evaluation and application of veterinary clinical pathology tests. Veterinary Clinical Pathology, 35:8-17.

Comparison of ROC curves

Select *Comparison of ROC curves* in the menu when you want to test the statistical significance of the difference between the areas under 2 to 6 ROC curves (Hanley & McNeil, 1983).

Required input

Data

- **Variables**: identify the different variables (at least 2, maximum 6).
- **Classification variable**: select or enter a dichotomous variable indicating diagnosis (0=negative, 1=positive). If diagnosis is coded differently than using the values 0 and 1, you can use the IF function to transform the codes into 0 and 1 values, e.g. IF(RESULT="pos",1,0).
- **Select**: (optionally) a selection criterion in order to include only a selected subgroup of cases (e.g. AGE>21, SEX="Male").

Methodology

- **DeLong et al**.: use the method of Delong et al. (1988) for the calculation of the Standard Error of the Area Under the Curve (AUC) and of the difference between two AUCs (recommended).
- **Hanley & McNeil**: use the methods of Hanley & McNeil (1982, 1983) for the calculation of the Standard Error of the Area Under the Curve (AUC) and of the difference between two AUCs.
- **Binomial exact Confidence Interval for the AUC**: calculate exact Binomial Confidence Intervals for the Area Under the Curves (AUC) (recommended). If this option is not selected, the Confidence Intervals for the AUCs are calculated as AUC ± 1.96 SE (Standard Error). This option does not apply to the difference between two AUCs).

Graph

- Select **Display ROC curves window** to obtain the ROC plots in a separate graph window.

  Option:
  - mark points corresponding to criterion values.

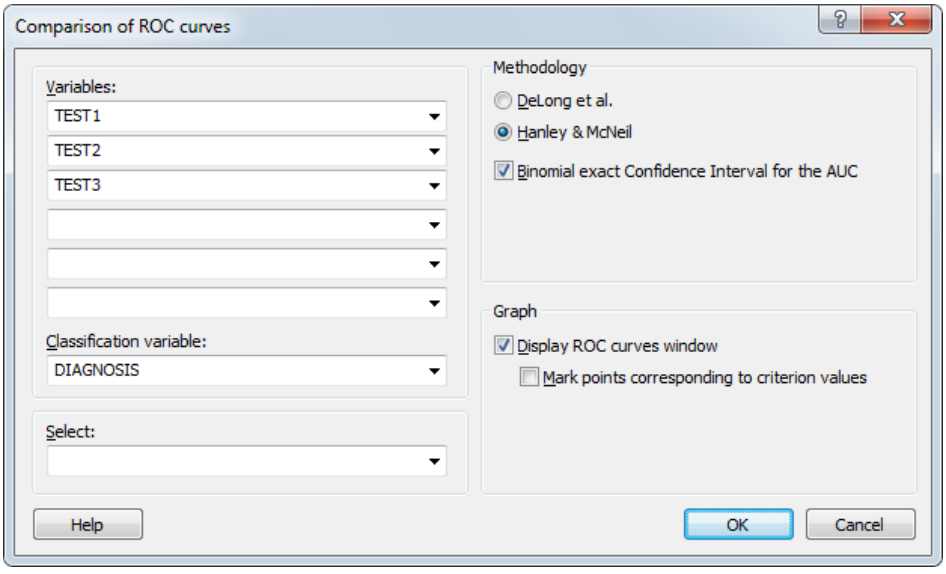

Results

The results window shows the data for the different ROC curves followed by the result of pairwise comparison of all ROC curves: the difference between the areas, the standard error, the 95% confidence

interval for the difference and P-value. If P is less than the conventional 5% (P<0.05), the conclusion is that the two compared areas are significantly different.

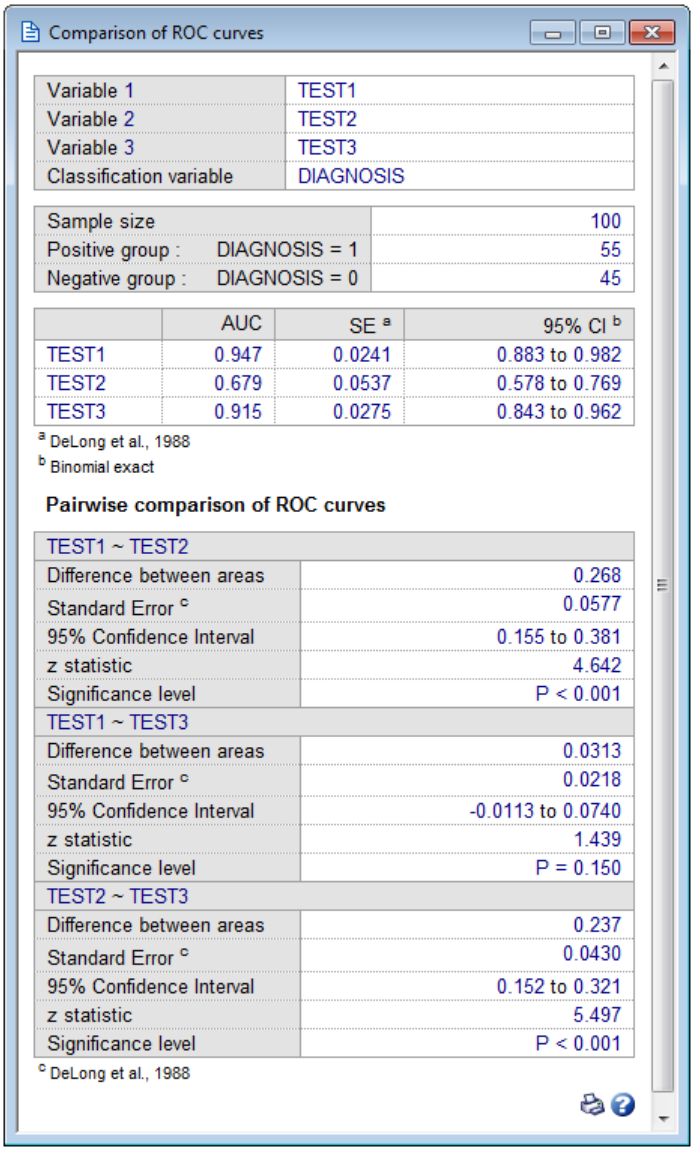

Comparison of ROC curves

Variable 1	TEST1
Variable 2	TEST2
Variable 3	TEST3
Classification variable	DIAGNOSIS

Sample size		100
Positive group :	DIAGNOSIS = 1	55
Negative group :	DIAGNOSIS = 0	45

	AUC	SE [a]	95% CI [b]
TEST1	0.947	0.0241	0.883 to 0.982
TEST2	0.679	0.0537	0.578 to 0.769
TEST3	0.915	0.0275	0.843 to 0.962

[a] DeLong et al., 1988
[b] Binomial exact

Pairwise comparison of ROC curves

TEST1 ~ TEST2	
Difference between areas	0.268
Standard Error [c]	0.0577
95% Confidence Interval	0.155 to 0.381
z statistic	4.642
Significance level	P < 0.001
TEST1 ~ TEST3	
Difference between areas	0.0313
Standard Error [c]	0.0218
95% Confidence Interval	-0.0113 to 0.0740
z statistic	1.439
Significance level	P = 0.150
TEST2 ~ TEST3	
Difference between areas	0.237
Standard Error [c]	0.0430
95% Confidence Interval	0.152 to 0.321
z statistic	5.497
Significance level	P < 0.001

[c] DeLong et al., 1988

Display Roc curves

When you have selected *Display ROC curves window* in the dialog box, the program will also open a graph window with the different ROC curves.

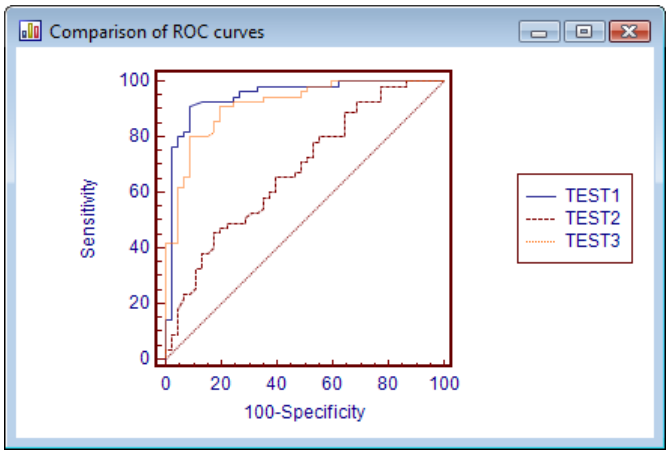

Create summary statistics table

Allows to create a publication-ready summary statistics of several variables and possible subgroups.

Required input

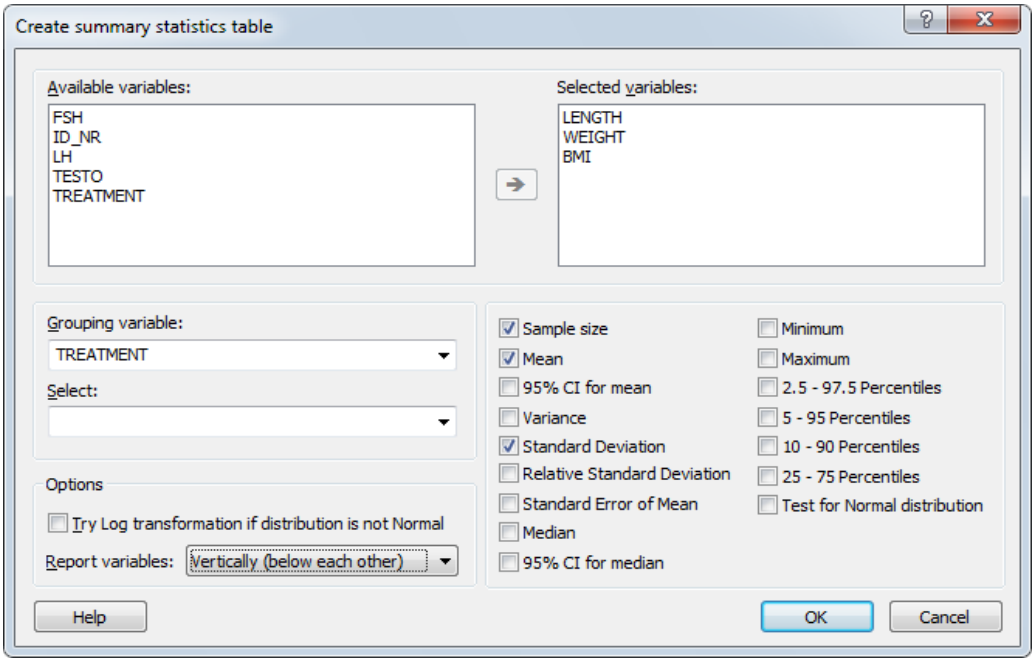

- **Variables**: select the variables of interest in the top left box and next click the right arrow button to move the selection to the Selected variables list.
- **Grouping variable**: (optionally) a categorical variable that defines subgroups.
- **Select**: an optional selection criterion.
- **Statistics**: select the statistics to be displayed in the table: Sample size, mean, 95% CI for the mean, etc.
- **Options:**
  - **Try Log Transformation if distribution is not Normal**: if this option is selected the software will perform a test for Normal distribution for all variables. If the data of a variable are found not to be Normally distributed, a Log Transformation is attempted. If after log transformation the data are found to agree better with a Normal distribution, the transformation is applied and the back-transformed results will be displayed in the table.

- **Report variables vertically or horizontally**: option to report the variables vertically (in rows below each other) or horizontally (in columns next to each other).

Results

The results window displays the summary statistics table.

With the option 'Report variables vertically':

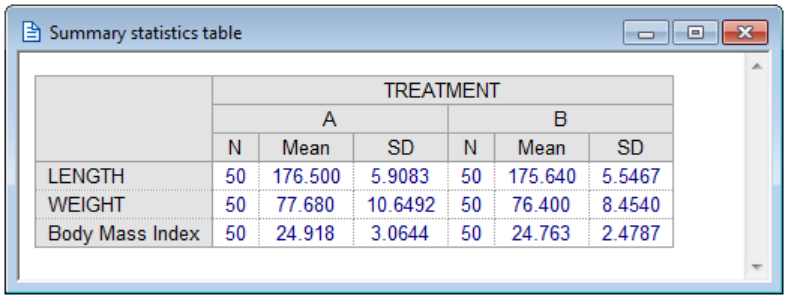

	TREATMENT					
	A			B		
	N	Mean	SD	N	Mean	SD
LENGTH	50	176.500	5.9083	50	175.640	5.5467
WEIGHT	50	77.680	10.6492	50	76.400	8.4540
Body Mass Index	50	24.918	3.0644	50	24.763	2.4787

With the option 'Report variables horizontally':

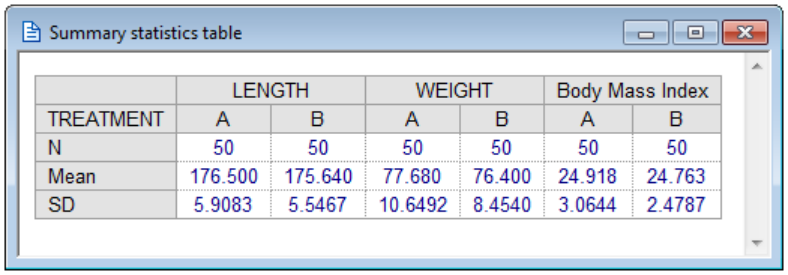

	LENGTH		WEIGHT		Body Mass Index	
TREATMENT	A	B	A	B	A	B
N	50	50	50	50	50	50
Mean	176.500	175.640	77.680	76.400	24.918	24.763
SD	5.9083	5.5467	10.6492	8.4540	3.0644	2.4787

To copy the table to Microsoft Word:

- Open your document in Microsoft Word.
- Activate the MedCalc results window.
- Select "Select all" in the Edit menu, or press Ctrl+A
- Select "Copy" in the Edit menu, or press Ctrl+C
- Activate Microsoft Word and place the cursor where you want to insert the table.
- Select "Paste" in Word's Edit menu, or press Ctrl+V

Create correlation table

This command allows to create a publication-ready table of correlation coefficients.

Required input

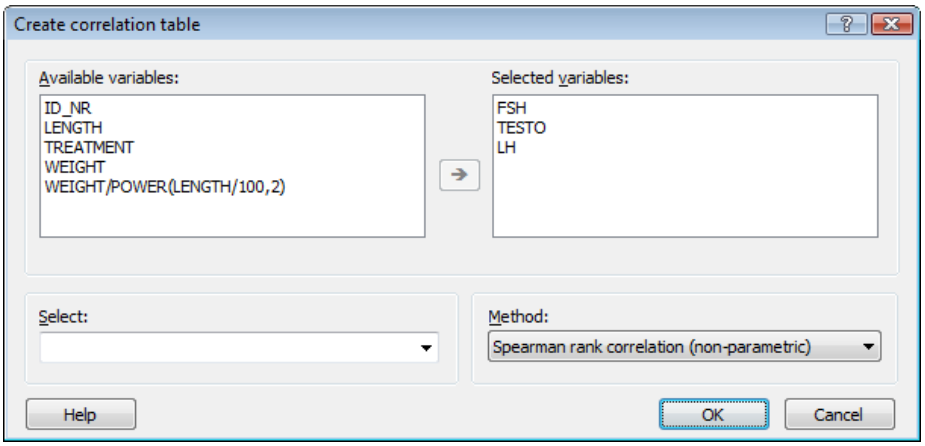

- **Variables**: select the variables of interest in the top left box and next click the right arrow button to move the selection to the Selected variables list.
- **Select**: an optional selection criterion.
- **Method**: select Pearson correlation (parametric) or Spearman rank correlation (non-parametric).

Results

The results window displays the correlation table.

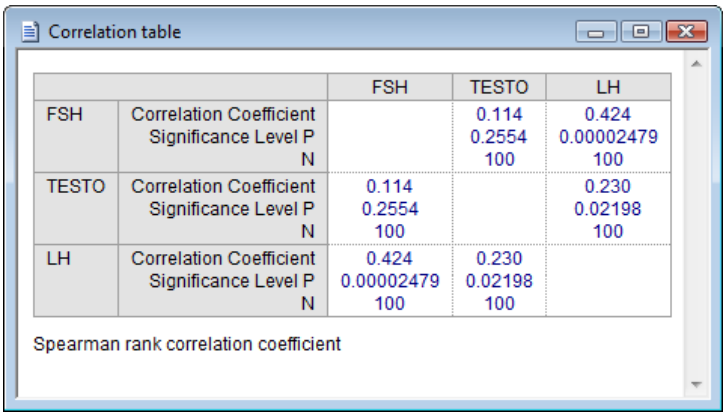

		FSH	TESTO	LH
FSH	Correlation Coefficient		0.114	0.424
	Significance Level P		0.2554	0.00002479
	N		100	100
TESTO	Correlation Coefficient	0.114		0.230
	Significance Level P	0.2554		0.02198
	N	100		100
LH	Correlation Coefficient	0.424	0.230	
	Significance Level P	0.00002479	0.02198	
	N	100	100	

To copy the table to Microsoft Word:

- Open your document in Microsoft Word.
- Activate the MedCalc results window.
- Select "Select all" in the Edit menu, or press Ctrl+A
- Select "Copy" in the Edit menu, or press Ctrl+C
- Activate Microsoft Word and place the cursor where you want to insert the table.
- Select "Paste" in Word's Edit menu, or press Ctrl+V

GRAPHS MENU

Data comparison graphs

In MedCalc, 3 graphs are available for comparison of 2 sets of data. You can select the graph type in the dialog box that appears after you have selected *Data comparison graphs* in the menu:

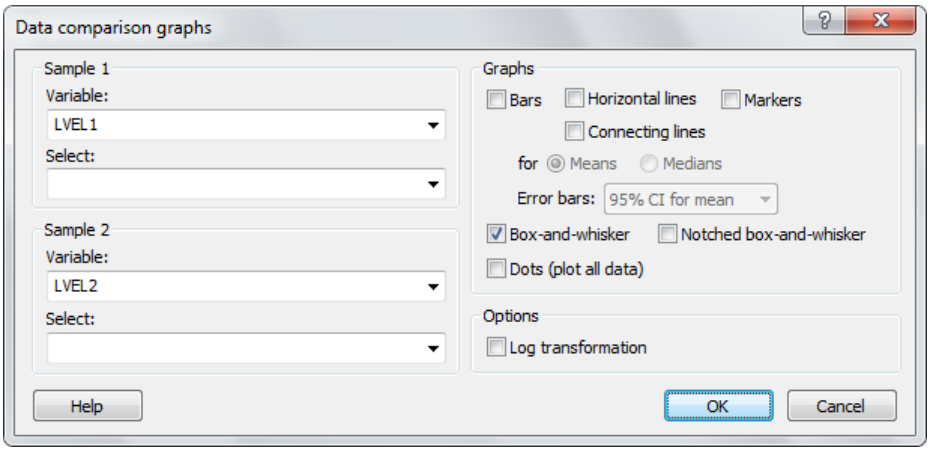

Several graphical elements can be selected to compose the graph, and some of these can be combined.

- *Bars*, *Horizontal lines*, *Markers* and/*or Connecting lines* for *means* or *medians.*
- *Error bars*: the following error bars are available if at least one of the graph types *Bars*, *Horizontal lines*, *Markers* and/*or Connecting lines* is selected:

  If *mean* is selected: (none), or 95% CI for the mean, 1 SD, 2 SD, 3 SD, 1 SEM, range

  - Note that 2 SEM is not in this list: when the number of cases is large, mean ± 2 SEM corresponds to the 95% confidence interval (CI) for the mean. When the number of cases is small, then the 95% CI interval is calculated as mean ± t * SEM, where t is taken from a t-table (with DF=n-1 and area A=95%) (see SEM on p. 59).
  - Although 1 SEM gives the more narrow error bar, this option is not recommended since the resulting error bar may be highly misleading, especially when the number of cases in the groups is different. Preferably the 95% CI for the mean is used for providing a valid graphical comparison of means (Pocock, 1984), or use 2 SD as an indication for the variability of the data.

  If *median* is selected: (none), or 95% CI for the median, 25-75 percentile, 10-90 percentile, 5-95 percentile, 1-99 percentile, range

  - When the number of cases is small, it is possible that the 95% CI for the median is not defined and that it will not be displayed in the graph.
  - When you use percentile ranges, take into account the number of observations: you need at least 100 observations for 1-99 percentiles, at least 20 for 5-95 percentiles, at least 10 for 10-90 percentile and at least 4 for 25-75th percentiles.
- The basic *Box-and-Whisker plot* (Tukey, 1977) is described on page 69.
- *Notched box-and-whisker plot*, in this variation of the box-and-whisker plot confidence intervals for the medians are provided by means of notches surrounding the medians (McGill et al., 1978). If the notches about two medians do not overlap, the medians are significantly different at a ± 95% confidence level.
- When you select *Dot plot*, all observations will be displayed in the graph.

If you want a logarithmic transformation of the values, select the LOG TRANSFORMATION option.

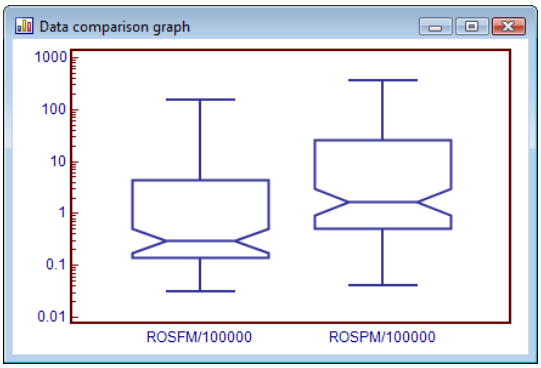

Notched box-and-whisker plots.

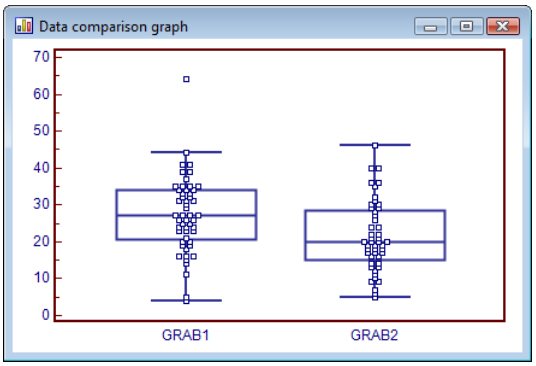

Combined box-and-whisker and dot plot.

Multiple comparison graphs

You can use *Multiple comparison graphs* to visualize the influence of a qualitative (discrete) factor on another (continuous) variable.

The graph can be composed from different elements: *Bar*, *Horizontal Lines, Markers* or *Connecting lines* for means or medians, with choice of different error bars for mean (95% CI, 1 SEM, 1 SD, 2 SD, 3 SD, range) or median (95% CI, 25-75 percentiles, 10-90 percentiles, 5-95 percentiles, 1-99 percentiles, range), *Box-and-whisker plot* (Tukey, 1977) or *Notched box-and-whisker plot* (McGill et al., 1978), and/or *Dot plot* (display all data) (see p. 171).

The following need to be entered in the dialog box: for *Data* select a continuous variable, and for *Factor codes* a qualitative factor. The qualitative factor may either be character or numeric codes. These codes are used to break-up the data into several subgroups.

When you want to use a continuous variable as the qualitative, discrete factor, you can convert the continuous data by using the CATEGORISE function (see p. 219) or IF function (see p. 220).

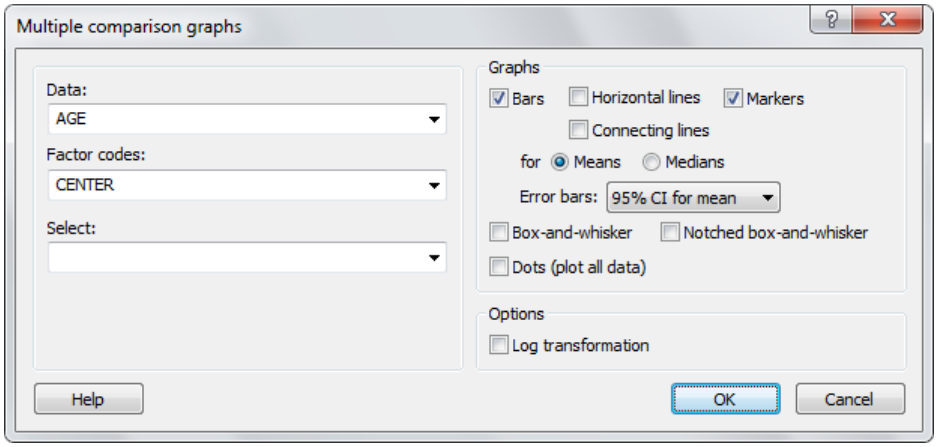

Several graphical elements can be selected to compose the graph, for a description, see p. 171.

If logarithmic transformation of the data is required, select the LOG TRANSFORMATION option.

After you have completed the form, click the OK button to obtain the graph.

Some combinations of graphs are shown in the following examples:

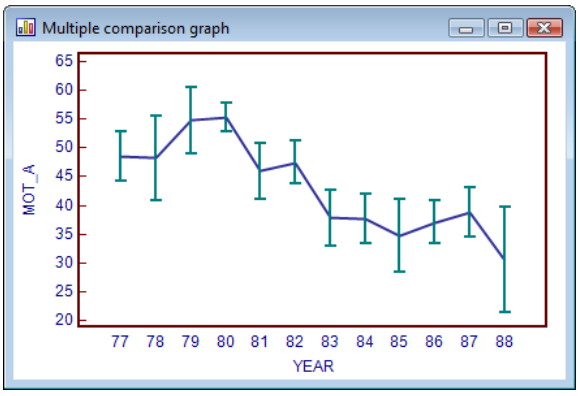

Chart with the evolution of sperm motility during 12 years. The means for every year are connected by a line. Error bars represent a 95% confidence interval for the mean.

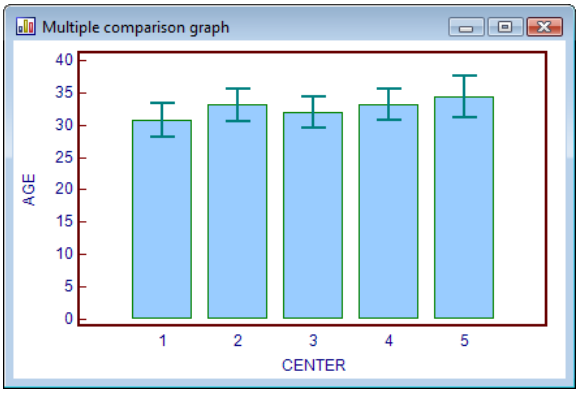

Chart with bars representing the mean age of patients in a multicenter study, and 'error bars' representing a 95% confidence interval.

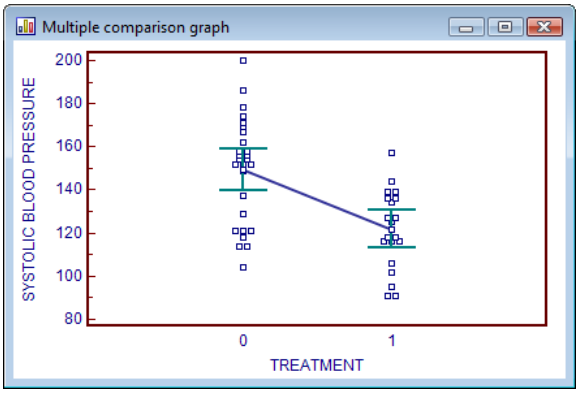

Chart with all data points for systolic blood pressure, 'error bars' representing a 95% confidence interval, and a line connecting the mean of the two treatment groups (coded 0 and 1).

Clustered multiple comparison graphs

Clustered multiple comparison graphs can be used when you want to visualize the influence of two qualitative (discrete) factors on another (continuous) variable.

The qualitative factors may either be character or numeric codes. These codes are used to break-up the data into different subgroups.

The graph can be composed from different elements: *Bar*, *Horizontal Lines, Markers* or *Connecting lines* for means or medians, with choice of different error bars for mean (95% CI, 1 SEM, 1 SD, 2 SD, 3 SD, range) or median (95% CI, 25-75 percentiles, 10-90 percentiles, 5-95 percentiles, 1-99 percentiles, range), *Box-and-whisker plot* (Tukey, 1977) or *Notched box-and-whisker plot* (McGill et al., 1978), and/or *Dot plot* (display all data) (see p. 171)

How to enter data

You need to enter data for one continuous variable (MEASUREMENT1 in the example) and 2 categorical variables (GENDER and TREATMENT in the example).

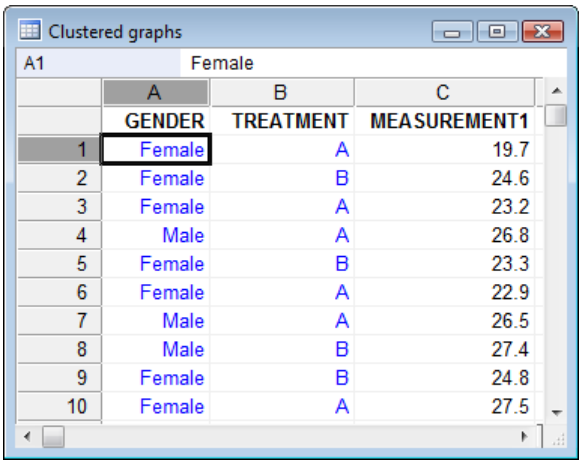

	A	B	C
	GENDER	TREATMENT	MEASUREMENT1
1	Female	A	19.7
2	Female	B	24.6
3	Female	A	23.2
4	Male	A	26.8
5	Female	B	23.3
6	Female	A	22.9
7	Male	A	26.5
8	Male	B	27.4
9	Female	B	24.8
10	Female	A	27.5

Required input

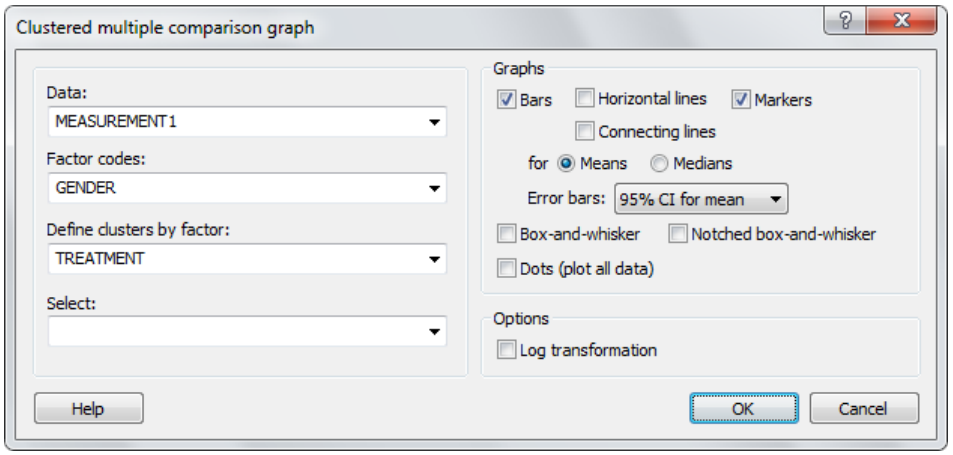

The following need to be entered in the dialog box:

- **Data**: a continuous variable that will be represented in the graph;
- **Factor codes**: a categorical variable that contains codes to break-up the data into subgroups.
- **Define clusters by factor**: a second categorical variable to make a second subdivision in the subgroups.
- **Select**: a selection criterion to include only a selected subgroup of cases in the graph.
- **Graphs**: several graphical elements can be selected to compose the graph, for a description, see p. 171.
- **Options**: Log transformation of data.

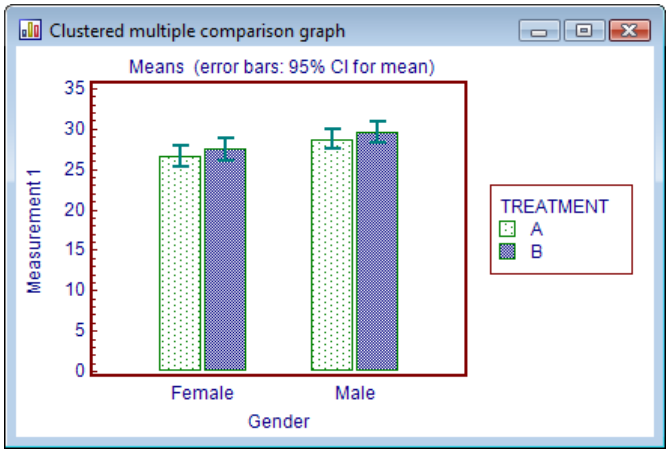

This is an example of a graph with option “Dots” selected.

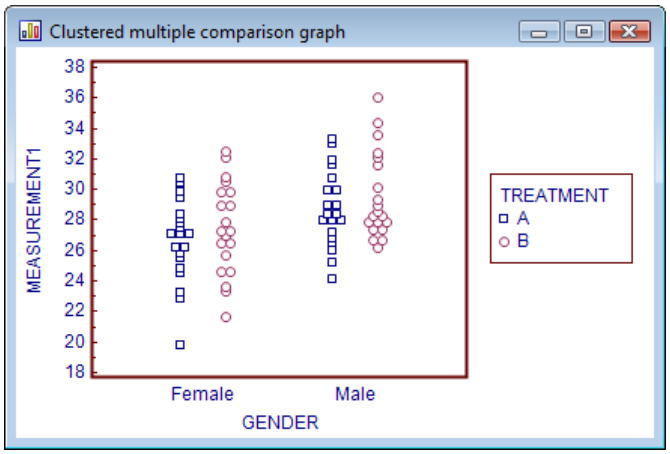

Multiple variables graphs

When you want to compare different variables, i.e. data entered in different columns of the spreadsheet, you can use *Multiple variables graphs*. In the dialog box, you select the different variables of interest, optionally followed by a selection criterion to include only a selected subgroup of cases in the graph.

Next select the different elements of the graph: *Bar*, *Horizontal Lines, Markers* or *Connecting lines* for means or medians, with choice of different error bars for mean (95% CI, 1 SEM, 1 SD, 2 SD, 3 SD, range) or median (95% CI, 25-75 percentiles, 10-90 percentiles, 5-95 percentiles, 1-99 percentiles, range), *Box-and-whisker plot* (Tukey, 1977) or *Notched box-and-whisker plot* (McGill et al., 1978), and/or *Dot plot* (display all data) (see p. 171).

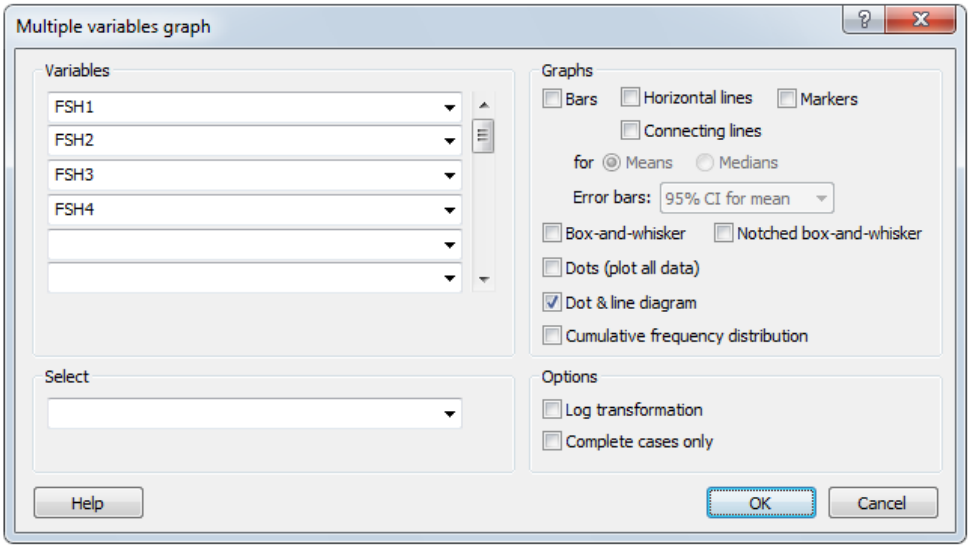

You can select the option of *logarithmic transformation* of the data, or the option to include only *complete cases* in the graph. If the latter option is selected, only cases with valid numerical data for all variables entered in the dialog box will be included in the graph.

In addition to the graphical elements as for *Data comparison graphs* described on page 171, the *Multiple variables graphs* also include the multiple *Dot and line diagram* and *Cumulative frequency distribution*.

In the *Dot and line diagram*, all observations are plotted as individual dots, and observations from the different cases (rows in the spreadsheet) are connected by a line.

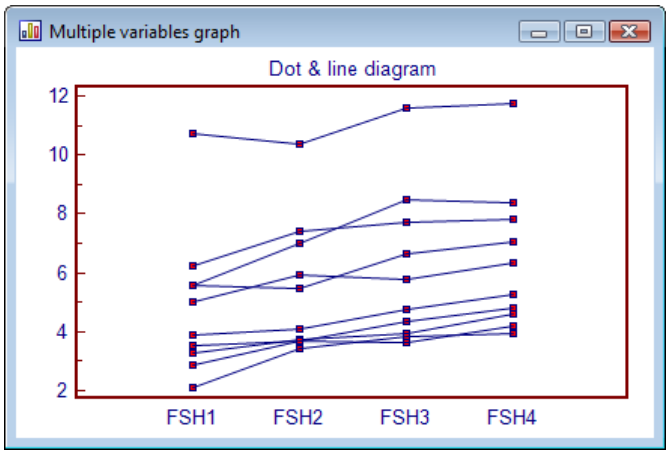

Clustered multiple variables graph

When you want to compare subgroups across different variables, you can use *Clustered multiple variables graphs*.

Clustered multiple variables graphs include: *Bar*, *Horizontal Lines, Markers* or *Connecting lines* for means or medians, with choice of different error bars for mean (95% CI, 1 SEM, 1 SD, 2 SD, 3 SD, range) or median (95% CI, 25-75 percentiles, 10-90 percentiles, 5-95 percentiles, 1-99 percentiles, range), *Box-and-whisker plot* (Tukey, 1977) or *Notched box-and-whisker plot* (McGill et al., 1978), and/or *Dot plot* (display all data) (see p. 171)

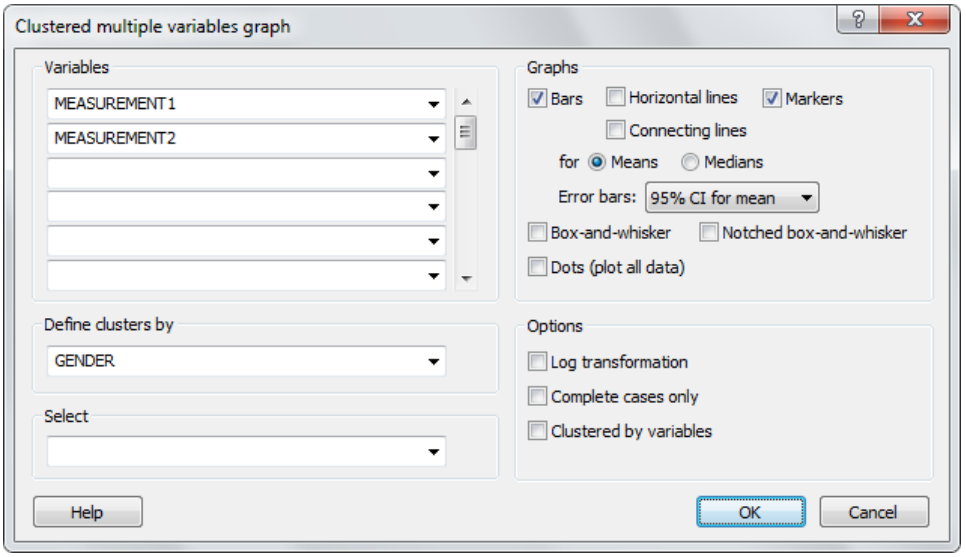

In this dialog box, you enter:

- **Variables**: the different variables of interest.
- **Define clusters by**: a categorical variable containing codes to break-up the data into subgroups.
- **Select**: a selection criterion to include only a selected subgroup of cases in the graph.
- **Graphs**: several graphical elements can be selected to compose the graph, for a description, see p. 171..
- **Options**:
  - Log transformation: Logarithmic transformation of all data,.
  - Complete cases only: Option to include only complete cases in the graph. If selected, only cases with valid numerical data for all variables selected in the dialog box will be included in the graph.

- Clustered by variables: Let clusters be defined by the Variables (X-axis will display groups defined by the Categorical variable, list of variables will appear in the Legend block).

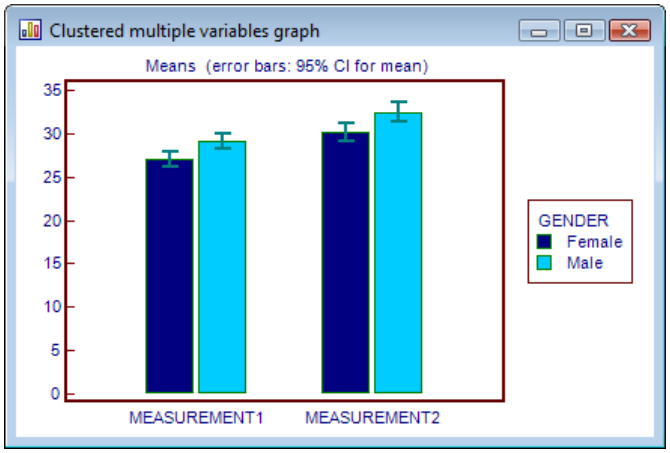

This is the same graph, but with the option “Clustered by variables” selected:

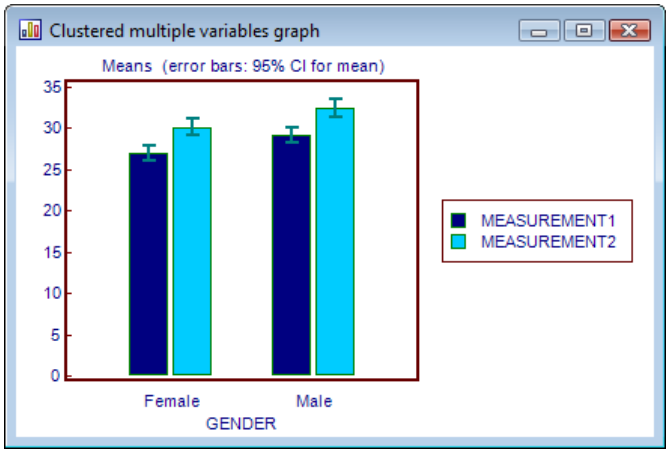

This is an example of a graph with option “Connecting lines” selected:

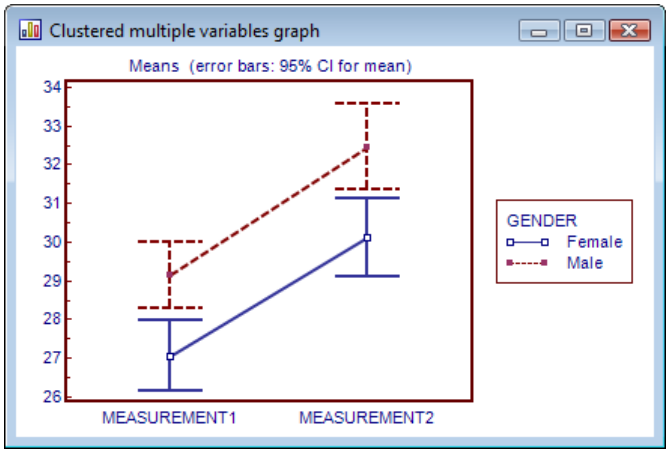

Multiple line graph

Description

Creates a graph that shows consecutive observations of different variables.

Required input

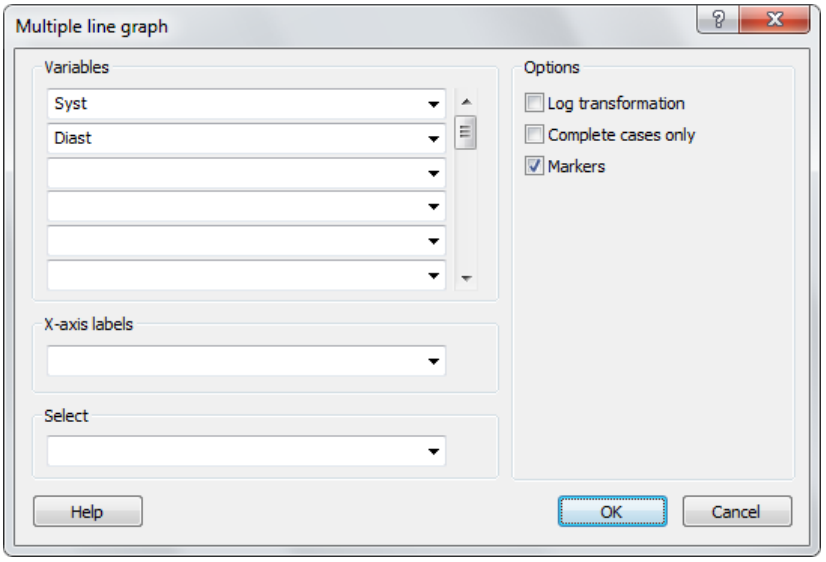

Select the variables of interest, and optionally a selection criterion to include only particular cases in the graph.

Also you can enter a variable containing labels for the X-axis (e.g. dates).

Options

- Logarithmic transformation: select this option to transform all data logarithmically
- Complete cases only: a case (spreadsheet row) will only be included in the graph when data are available for all variables.
- Markers: mark observations in the graph.

Example

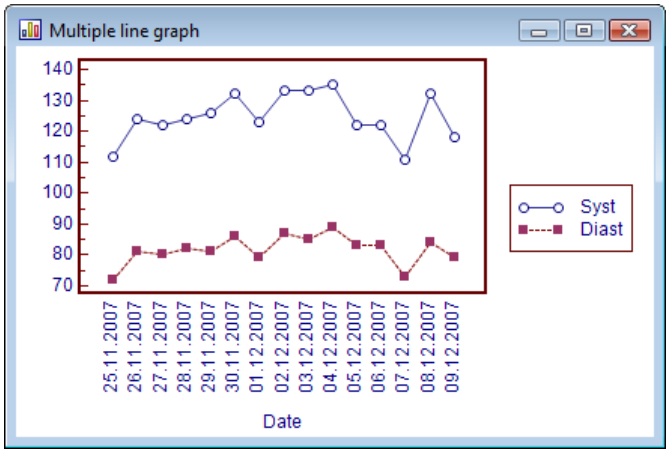

Control chart

To obtain a quality control chart select *Control chart* in the *Graphs* menu. In this chart the data are plotted consecutively, together with a line at the mean, and at -2s, +2s, -3s and +3s (s = standard deviation), i.e. at 95% and 99.7% confidence limits.

In the *Quality control* input box define the data to be represented in the quality control chart. First, enter the variable's name in the variable input field. Optionally, you may also enter a selection criterion in order to include only a subgroup of data in the graph, e.g. measurements by a particular technician. Also you can enter a variable containing labels for the X-axis (e.g. dates).

Next select the type of control limits:

- The control limits can be based on the data and in this case the program will calculate the mean and standard deviation of the selected data.

  Option: "until n =" : if the control limits must be based on the first observations only, you can enter the required number here. E.g. if you have 40 observations to be plotted in the control chart, but the control limits must be based on the first 20 observations only, you enter the number 20 here.

- You can select 'Standard' and in this case enter the Mean and Standard Deviation (SD) of the standard used.

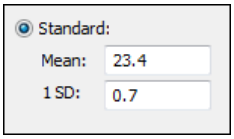

- Finally, you can enter the reference value with upper and lower control and warning limits. In this case the upper and lower control and warning limits can be asymmetrical.

Rules: see below

In the following example, a control chart is created for the technician with code 2.

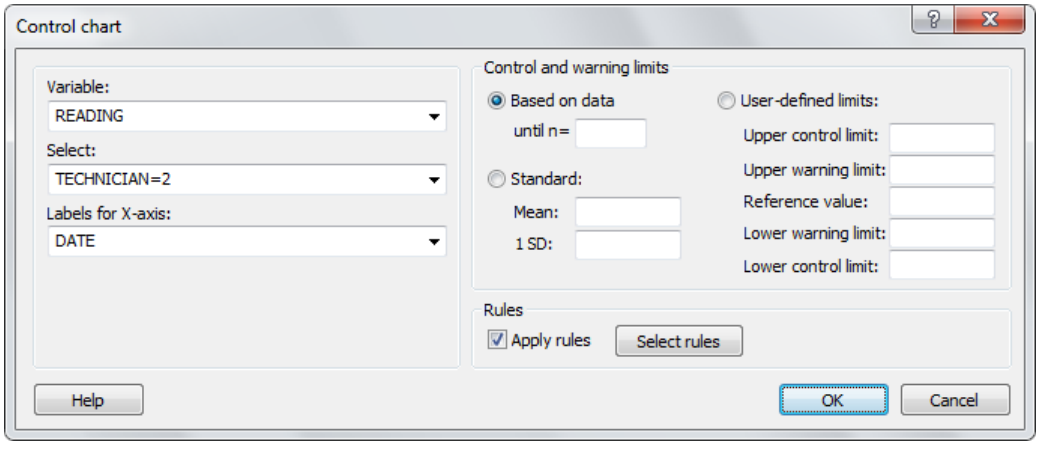

The result is displayed in the following figure:

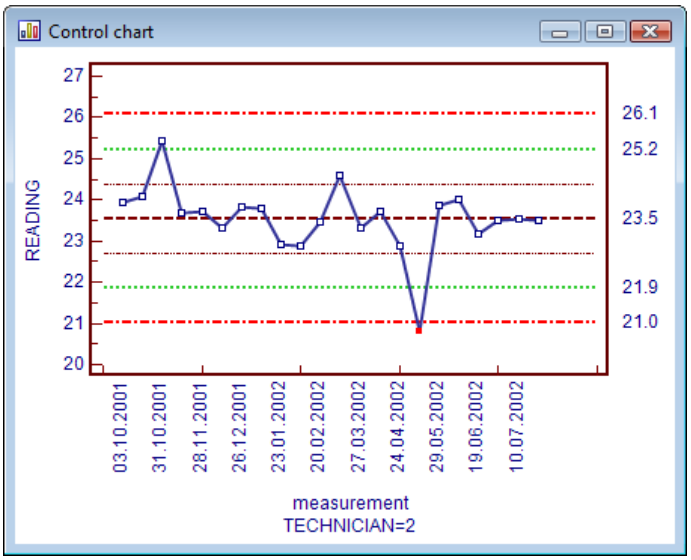

Rules

When you click the SELECT RULES button, the following dialog box is displayed, in which you can select a combination of rules to be applied in the control chart:

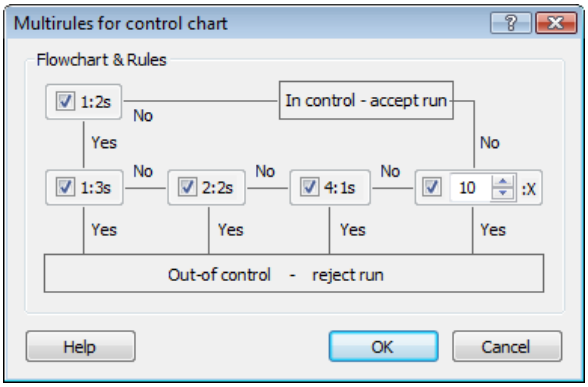

1:2S rule

If you select the 1:2S rule then the software will check all following rules only if the measurement exceeds the mean ± 2SD (or the warning limits).

If this rule is not selected, then the following rules will be checked also when the measurement does not exceed the mean ± 2SD warning limits.

This will particularly influence the 4:1S rule and 10:X rule.

1:3S rule

If the measurement exceeds the mean + 3SD or mean - 3SD, then the run is considered out of control.

This rule mainly detects random error, but may also be an indication of a large systematic error.

2:2S rule

The run is considered out of control when 2 consecutive measurements exceed the same mean + 2S or the same mean - 2S limit.

This rule detects systematic error.

4:1S rule

The run is considered out of control when 4 or more consecutive measurements exceed the same (mean + 1S) or (mean - 1S) limit.

This rule detects systematic bias.

10:X rule

The run is considered out of control when 10 or more consecutive measurements are on the same side of the mean. The software allows you to select a different value less than 10 for a higher sensitivity or more than 10 for a lower sensitivity of this rule.

This rule detects systematic bias.

Example

When the option *Apply rules* is selected in the Control chart dialog box, and the *1:3S* rule is selected in the Multirules dialog box, then when a measurement exceeds the mean + 3SD or mean - 3SD, this value will be indicated in the graph (using a marker drawn in the "warning" color). Click the value to display information on this measurement:

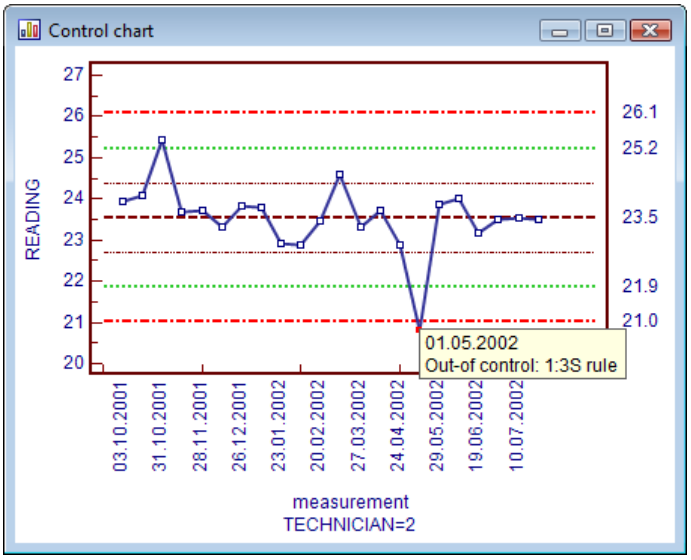

Custom control chart

When there are no data available for the variable, the program will ask whether you want to create a custom control chart, i.e. a blank control chart that you can print and on which the lab technician can mark the control measurements manually.

Youden plot

The Youden plot is a graphical method to analyze inter-laboratory data, where all laboratories have analyzed 2 samples. The plot visualizes within-laboratory variability as well as between-laboratory variability.

In medical literature you may encounter different graphs referred to as *Youden plot*.

Youden plots

1. The original Youden plot

For the original Youden plot (Youden, 1959) the two samples must be *similar and reasonably close in the magnitude of the property evaluated*.

The axes in this plot are drawn on the same scale: one unit on the x-axis has the same length as one unit on the y-axis.

Each point in the plot corresponds to the results of one laboratory and is defined by a first response variable on the horizontal axis (i.e. run 1 or product 1 response value) and a second response variable 2 (i.e., run 2 or product 2 response value) on the vertical axis.

A horizontal median line is drawn parallel to the x-axis so that there are as many points above the line as there are below it. A second median line is drawn parallel to the y-axis so that there are as many points on

the left as there are on the right of this line. Outliers are not used in determining the position of the median lines. The intersection of the two median lines is called the *Manhattan median*.

A circle is drawn that should include 95 % of the laboratories *if individual constant errors could be eliminated*.

A 45-degree reference line is drawn through the *Manhattan median*.

Interpretation

Points that lie near the 45-degree reference line but far from the Manhattan median indicate large *systematic error*.

Points that lie far from the 45-degree line indicate large *random error*.

Points outside the circle indicate large *total error*.

2. The Youden plot adapted for non-comparable samples

If two *different* products are being tested, MedCalc draws a Youden plot as described above, but the axes of the plot are not drawn on the same scale, but in this case, one standard deviation on the X-axis has the same length as one standard deviation on the y-axis.

Analogous to the 45-degree reference line in the original Youden plot, a reference line is drawn which in this case represents a constant *ratio* of the two samples.

The interpretation is the same as for the original Youden plot.

3. Other variations of the Youden plot

A common variation of the Youden plot is a scatter diagram as described above, but the circle is replaced with one or more rectangles representing 1, 2 or 3SD on both the x-axis and y-axis.

The Youden dialog box

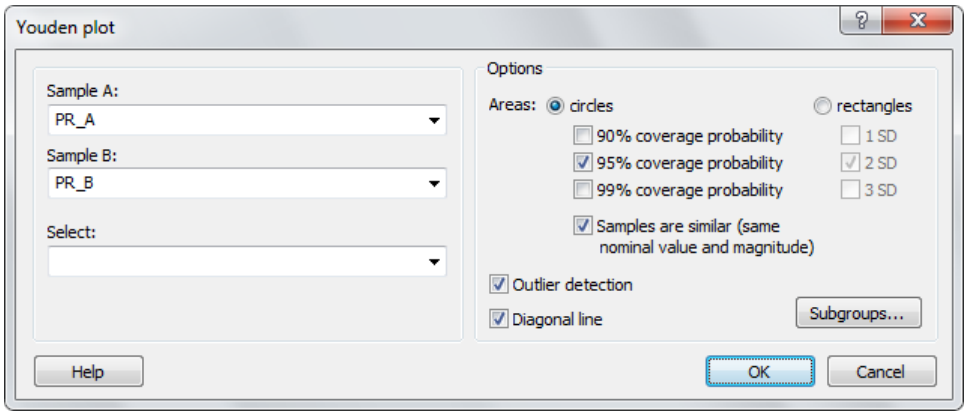

Sample A and **Sample B**: select the variables for the first and second sample.

Select: an (optional) selection criterion to include only a selected subgroup of cases in the graph.

Options

- **Areas – Circles**:

  90%, 95% or 99% Coverage probability: circles can be drawn that include 90%, 95% or 99% of the laboratories if individual constant errors could be eliminated.

  Samples are similar: select this option if the samples similar and reasonably close in the magnitude of the property evaluated.
- **Areas – Rectangles**:

  1 SD, 2SD or 3SD: draws rectangles representing 1, 2 or 3 SD on both the x-axis and y-axis.
- **Outlier detection**: MedCalc will detect outliers automatically and exclude them for calculations.
- **Diagonal line**: draws a diagonal reference line
- **Subgroups**: use the **Subgroups** button if you want to identify subgroups in the plot. A new dialog box is displayed in which you can select a categorical variable. The graph will use different markers for the different categories in this variable.

Examples

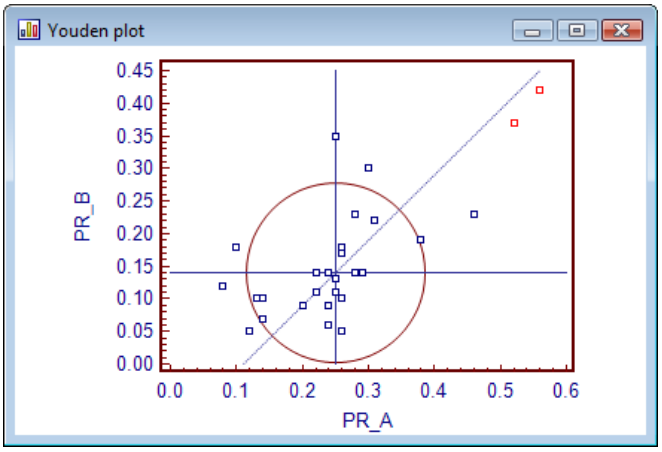

Fig. 1.

The original Youden plot. PR_A and PR_B represent similar samples. Notice the two outliers in the upper right corner of the graph.
Circle represents 95% coverage probability.

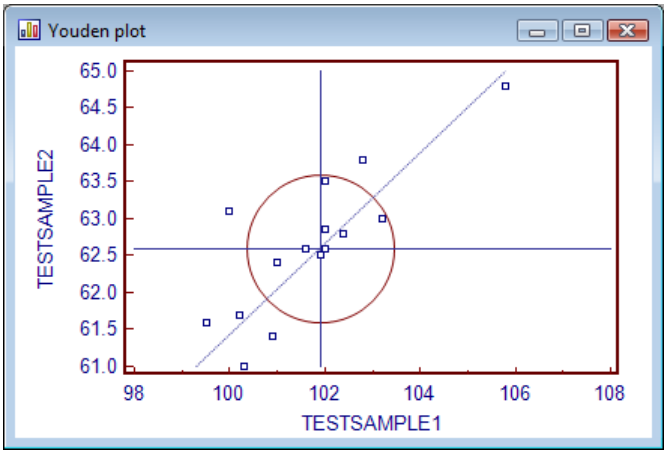

Fig 2.

Example of a Youden plot adapted for samples that are not similar.
Circle represents 95% coverage probability.

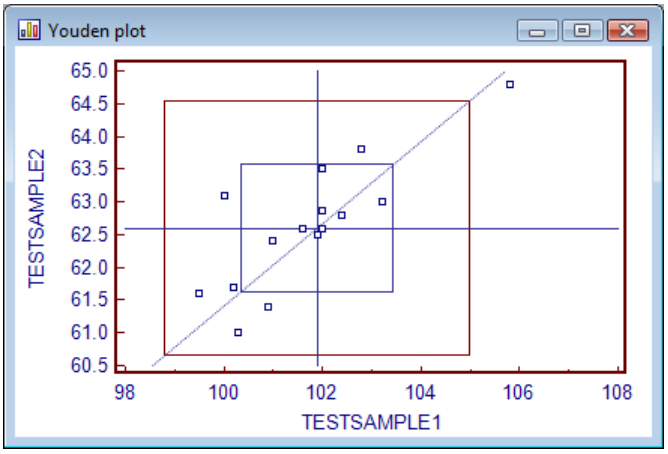

Fig. 3

Variation of Youden plot. The data are the same ad in Figure 2. Rectangles represent 1 and 2 SD.

Print all

To print multiple copies of the Youden plot, with each copy highlighting one laboratory:

- right-click in the Youden plot
- select Print all in the popup menu

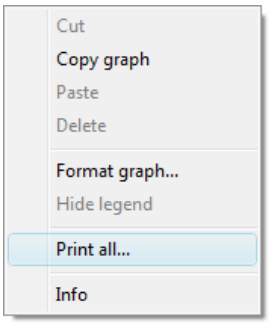
Cut
Copy graph
Paste
Delete
Format graph...
Hide legend
Print all...
Info

TESTS MENU

The *Tests* menu includes statistical tests on tabulated or summarized data. These tests are useful when you do not have the raw data available in the spreadsheet, e.g. when you want to compare e.g. two means reported in the literature, and you do not have access to the raw data.

Test for one mean

The *Test for one mean* is used to test the hypothesis that a sample mean is equal to a given mean (with unknown standard deviation) or certified value. The values are entered in the following dialog box:

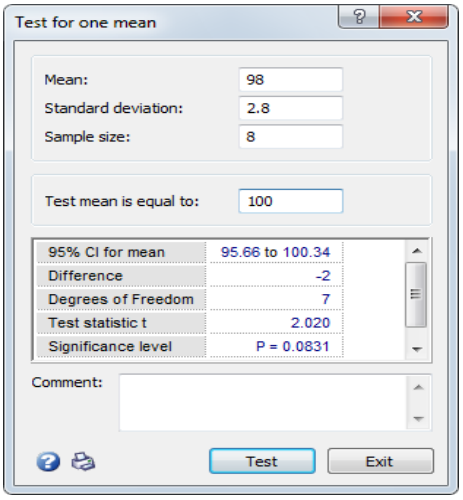

First enter the sample mean, standard deviation and sample size (n) in the dialog box. Next, in the input field *Test mean is equal to:* enter the value to compare the mean to. As an example, a sample mean of 98 with standard deviation of 2.8 and sample size 8 is compared with the certified value 100.

When all data have been entered click the TEST button.

First the program displays a 95% confidence interval for the mean. Next, the calculated t-value, degrees of freedom and probability (P) are displayed. If the P-value is less than 0.05, then the hypothesis that the mean is equal to the given value is rejected, and the alternative hypothesis that there is a significant difference between the two values can be accepted.

In the example, the calculated P-value is 0.0831 so you do not reject the hypothesis that the sample mean of 98 is equal to 100.

In the *Comment* input field you can enter a comment or conclusion that will be included on the printed report.

Test for one proportion

Description

The Test for one proportion in the Tests menu can be used to test the hypothesis that an observed proportion is equal to a pre-specified proportion.

This test is not performed on data in the spreadsheet, but on statistics you enter in a dialog box.

Required input

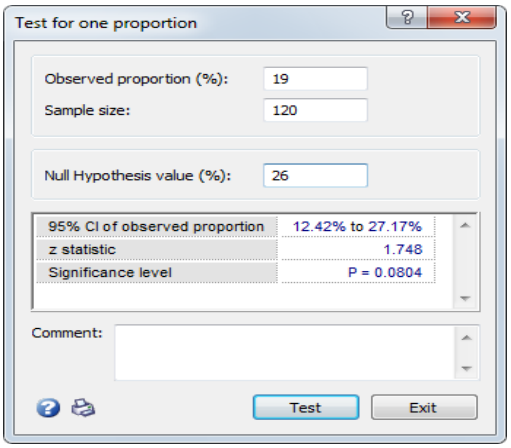

- **Observed proportion (%)**: the observed proportion, expressed as a percentage.
- **Sample size**: the sample size or total number of observations.
- **Null Hypothesis value (%)**: the pre-specified proportion (the value to compare the observed proportion to), expressed as a percentage.

When all data have been entered click the Test button.

Results

The results panel displays:

- The 95% Confidence Interval of the observed proportion.
- z statistic and associated P-value. If the P-value is less than 0.05, the hypothesis that the observed proportion is equal to the pre-specified proportion value is rejected, and the alternative hypothesis that there is a significant difference between the two proportions can be accepted.

Chi-square test

You can use the *Chi-square test* in the *Tests* menu to test the statistical significance of differences in a classification system (one-way classification) or the relationship between two classification systems (two-way classification).

To perform this Chi-square test, you must already have the data classified in a frequency table (for the test on raw data, see p. 108).

A frequency table shows the number of cases that belong simultaneously to two or more distinct categories, e.g. patients cross-classified according to both gender and age group. The data of the contingence table have to be entered in the table in the dialog form. Either a one-way classification can be used (occupying one single row or one single column), or a two-way classification table up to a 6 x 9 table.

Optionally, you can select a **Chi-square test for trend**. The Cochran-Armitage test for trend (Armitage, 1955) provides a more powerful test than the unordered test, but this test is only applicable if your classification table has 2 columns and 3 or more rows (or 2 rows and 3 or more columns), and if the data originate from ordered categories..

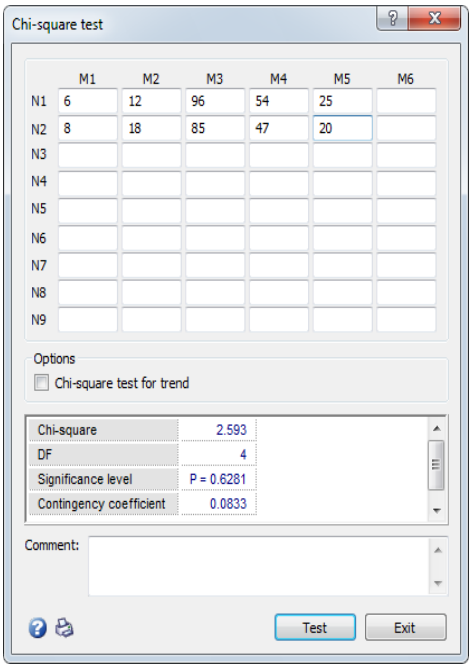

After you click the TEST button the program will automatically calculate the expected frequencies for every cell in the table, and the following results will be displayed:

- **Chi-square** with **degrees of freedom** and **P-value**. The Chi-square statistic is the sum of the squares of the differences of observed and expected frequency divided by the expected frequency for every cell:

$$\chi^2 = \sum \frac{(\text{observedcount} - \text{expectedcount})^2}{\text{expectedcount}}$$

When the *degrees of freedom* is equal to 1, e.g. in case of a 2x2 table, MedCalc uses *Yates' correction for continuity*.

If the calculated P-value is less than 0.05, then there is a statistically significant relationship between the two classifications.

- The **Contingency Coefficient** is a measure of the degree of relationship, association of dependence of the classifications in the frequency table. The coefficient is calculated as follows (**n** is the total number of cases in the table):

$$C = \sqrt{\frac{\chi^2}{\chi^2 + n}}$$

The larger the value of this coefficient, the greater the degree of association is. The maximum value of the coefficient, which is never greater than 1, is determined by the number of rows and columns in the table.

Fisher's exact test for a 2x2 table

When the number of expected frequencies in a 2x2 table is low (in case the total number of observations is less than 20), the table should be tested using Fisher's exact test. The data (representing number of cases) for the 2x2 table are entered in the dialog box.

Example: treatment A resulted in 6 successes and 1 failure (6/7 = 85.7%) whereas treatment B resulted in 2 success and 8 failures (2/10 = 20%). The following is entered in the dialog box:

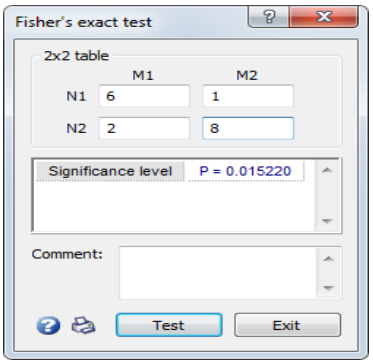

The P-value (the probability of obtaining the observed result or a more extreme result) is calculated when you click the TEST button.

The result, P=0.015 in the example, indicates that the 2 treatments gave significant different results.

In the *Comment* input field you can enter a comment or conclusion that will be included on the printed report.

McNemar test

The *McNemar test* is a test on a 2x2 classification table when the two classification factors are dependent, or when you want to test the difference between paired proportions, e.g. in studies in which patients serve as their own control, or in studies with "before and after" design.

In the example used by Bland (2000) 1319 schoolchildren were questioned on the prevalence of symptoms of severe cold at the age of 12 and again at the age of 14 years. At age 12, 356 (27%) children were reported to have severe colds in the past 12 months compared to 468 (35.5%) at age 14.

Severe colds	Severe colds at age 14		Total
at age 12	Yes	No	
Yes	212	144	356
No	256	707	963
Total	468	851	1319

Was there a significant increase of the prevalence of severe cold?

The data are entered as follows in the dialog box:

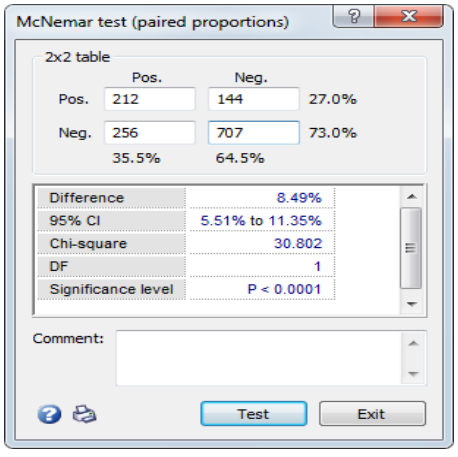

Difference and P-value

The program gives the difference between the proportions (expressed as a percentage) with 95% confidence interval.

When the (two-sided) P-value is less than the conventional 0.05, the conclusion is that there is a significant difference between the two proportions.

In the example, the difference between the prevalence at age 12 and age 14 is 8.49% with 95% CI from 5.5% to 11.4%, and is highly significant (P<0.0001).

Note

If the number of discordant pairs (144 + 256 in the example) is less than or equal to 25, then the two-sided P-value is based on the cumulative binomial distribution. If the number of discordant pairs is more than 25 then a chi-square approximation with Yates' correction for continuity is used.

The 95% confidence interval is calculated according to Bland, 2000.

Comparison of means (t-test)

When you want to compare the known mean of two independent samples, select *Means* in the *Tests* menu. In the dialog box, enter the mean, standard deviation and number of cases for both sets of data. Next you select TEST to test the statistical significance of the difference between the two means, using a Student's t-test.

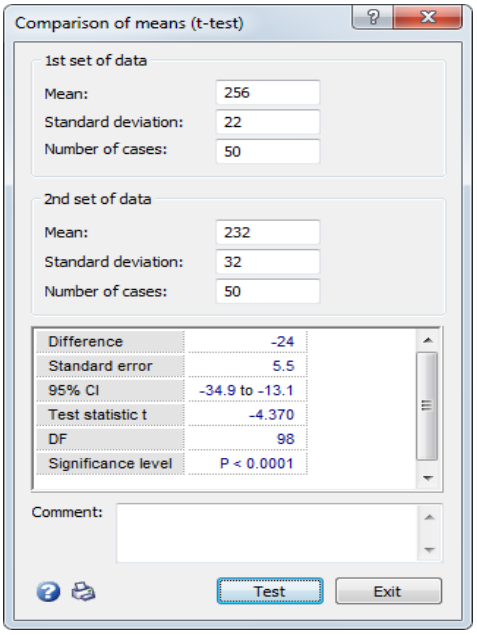

The program displays the difference between the two means, and a 95% confidence interval for this difference. Next, the P value is displayed: when this P value is less than 0.05, the conclusion is that the two means are significantly different.

In the example, the mean 256 (standard deviation 22 and sample size 50) of the first sample, is compared with the mean 232 (standard deviation 32 and sample size 50) of the second sample. The difference between the two means is -24. The 95% confidence interval for this difference ranges from -34.9 to -13.1 (note that this interval excludes the value 0). The hypothesis test results in a P-value less than 0.0001, and the conclusion therefore is that, statistically, there is a significant difference between the two means.

Comparison of standard deviations (F-test)

When you want to test the statistical significance of the difference between known *standard deviations* of two independent samples, use the *Comparison of ... Standard deviations* command in the *Tests* menu.

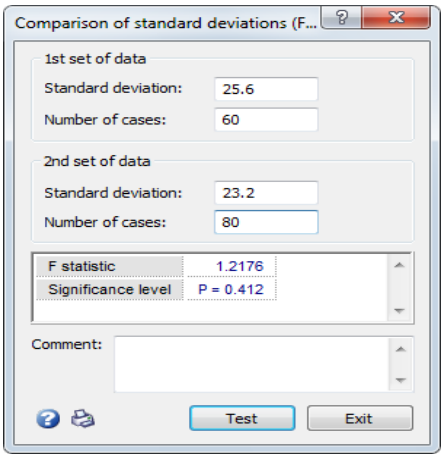

In the dialog box, enter the two standard deviations that you want to compare, and the corresponding number of cases. Next select TEST to perform the *F-test* or *variance ratio test*. In this test, the square of the standard deviations is calculated to obtain the corresponding variances. If the two variances are not significantly different, then their ratio will be close to 1.

When the calculated P value is less than 0.05 (P<0.05), the conclusion is that the two standard deviations are statistically significant different.

In the example, the standard deviation was 25.6 and sample size was 60 for the first sample, and for the second sample the standard deviation was 23.2 with sample size equal to 80. The resulting F-statistic was 1.2176 and the associated P-value was 0.412. Since P was not less than 0.05, you can conclude that there is no significant difference between the two standard deviations.

If you want to compare two known *variances*, first calculate the standard deviations, by taking the square root, and next you can compare the two standard deviations.

Comparison of correlation coefficients

When you want to test the statistical significance of the difference between two independent correlation coefficients, select *Correlation coefficients* in the *Tests* menu. In the dialog box enter the correlation coefficients and the corresponding number of cases. Next select TEST to calculate the statistical significance of the difference between the two correlation coefficients.

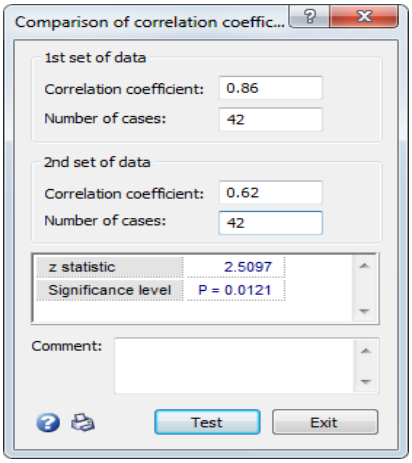

When the calculated P value is less than 0.05, the conclusion is that the two coefficients indeed are significantly different.

In the example a correlation coefficient of 0.86 (sample size = 42) is compared with a correlation coefficient of 0.62 (sample size = 42). The resulting z-statistic is 2.5097, which is associated with a P-value of 0.0140. Since this P-value is less than 0.05, it is concluded that the two correlation coefficients differ significantly.

Comparison of two proportions

When you want to compare two independent proportions (expressed as percentages), select *Comparison of proportions* in the *Tests* menu. In the dialog box, you can enter the two proportions (expressed as percentages) and the total number of cases. Next you click the TEST button to test the statistical significance of the difference between the two proportions (using a chi-square test).

For example when the prevalence of a disease in one sample of a total of 120 patients is 79%, and the prevalence of the disease in another sample of a total of 135 patients is 65%, you will enter the following in the dialog box:

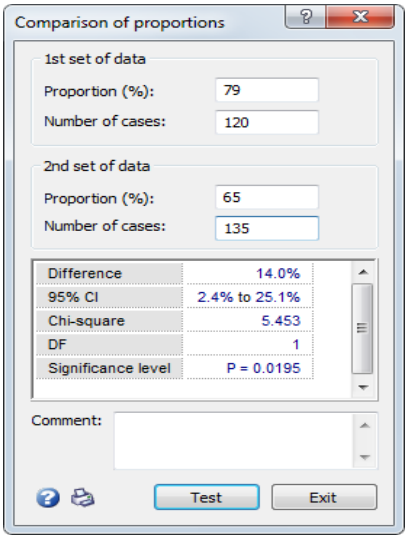

After you have selected the TEST button, the results of the test are displayed:

- the difference between the two proportions and a 95% confidence interval for this difference;
- Chi-square test (with *Yates' correction for continuity*) and P value: when this P value is less than 0.05, the conclusion is that the two proportions indeed differ significantly.

Note that when the total number of cases is less than 20, it may be more appropriate to perform Fisher's exact test (see p. 187).

Comparison of areas under independent ROC curves

Description

Allows to compare the Area under the Curve (AUC) of two independent ROC curves.

This test is not performed on data in the spreadsheet, but on statistics you enter in a dialog box.

Required **input**

Enter the **Area under the Curve (AUC)** and **Standard Error (SE)** for the two ROC curves.

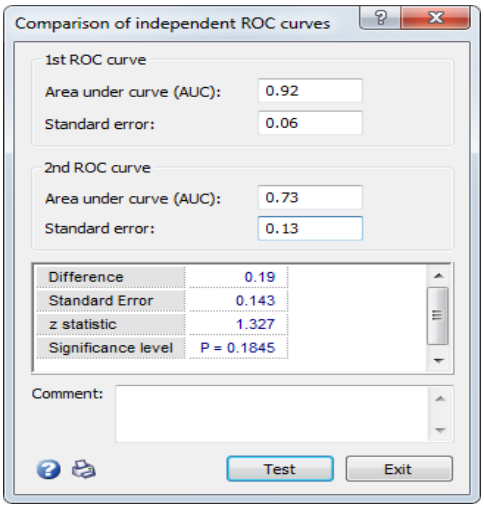

Click the Test button to test the statistical significance of the difference between the two AUCs.

Results

When the calculated P value is less than 0.05 (P<0.05), the conclusion is that the two AUCs are significantly different.

Confidence interval for a rate

Description

A rate is a ratio between a count and another measurement, for example the ratio of a number of events observed and the total number of person-years of observation.

This test is not performed on data in the spreadsheet, but on data you enter in a dialog box.

Required input

- **Numerator**: the observed number of events.
- **Denominator**: for example the total person-years.
- Option **Express result as 1:X**: when this option is selected the rate R will be displayed as 1:(1/R), e.g. the rate 10/200 equals 0.05 and can be represented as 1:20.

When all data have been entered click the Test button.

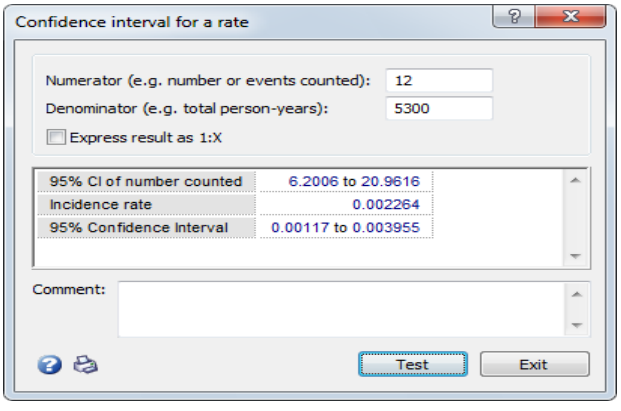

Results

The program displays:

- The Poisson 95% Confidence Interval for the number counted (the Numerator).
- The (incidence) rate.
- The 95% Confidence Interval for the incidence rate.

Literature

- Sahai H, Khurshid A (1996) Statistics in epidemiology: methods, techniques, and applications. Boca Raton, FL: CRC Press, Inc.

Comparison of two rates

Description

This procedure allows to compare the rates observed in two groups.

This test is not performed on data in the spreadsheet, but on data you enter in a dialog box.

Required input

- **Numerator**: the observed number of events in each group.
- **Denominator**: for example the total person-years for each group.
- Option **Express result as 1:X**: when this option is selected the rate R will be displayed as 1:(1/R), e.g. the rate 10/200 equals 0.05 and can be represented as 1:20.

When all data have been entered click the Test button.

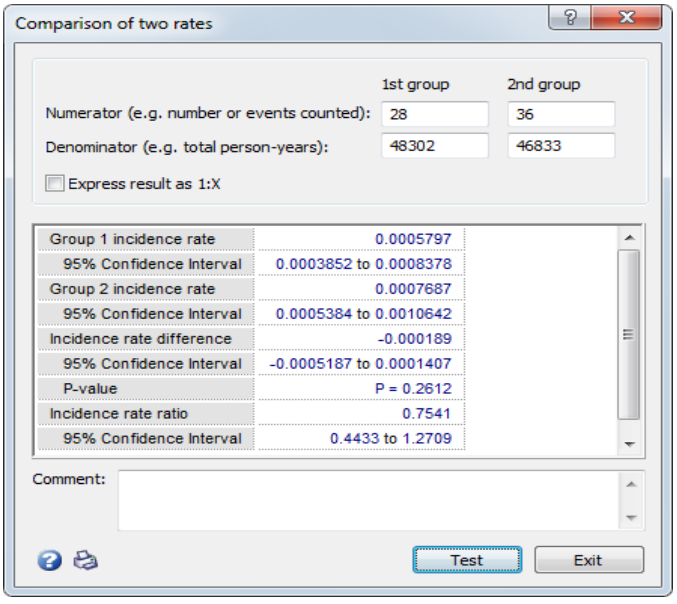

Results

The program displays:

- The (incidence) rate in the two groups with their Poisson 95% Confidence Interval.
- The difference between the two rates R2-R1 with its 95% Confidence Interval and associated P-value. If the P-value is less than 0.05 it can be concluded that there is a statistical significant difference between the two rates.
- The ratio of the two rates R1/R2 and its 95% Confidence Interval. If the value 1 is not in this interval, it can be concluded that the ratio R1/R2 is not significantly different from 1 (in which case the rates would be equal).

Literature

- Sahai H, Khurshid A (1996) Statistics in epidemiology: methods, techniques, and applications. Boca Raton, FL: CRC Press, Inc.

Relative risk

Description

This command is used to calculate a relative risk. The relative risk is the ratio of the proportions of cases having a positive outcome in two groups included in a prospective study.

In a prospective study cases are allocated to two groups and it is observed how many times the event of interest occurs.

The program calculates the relative risk and a 95% confidence interval. If the value 1 is not in the range of the confidence interval, it can be concluded that the proportions are significantly different in the two groups, and there is an increased risk in one group compared to the other.

This test is not performed on data in the spreadsheet, but on statistics you enter in a dialog box.

Required input

In the dialog box enter the number of cases with a positive and negative outcome in the two groups.

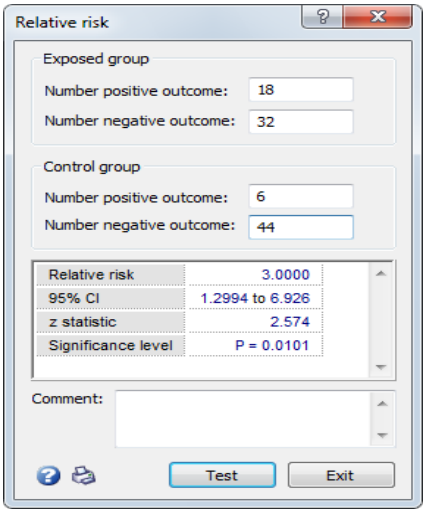

Click the Test button to perform the test.

Results

The program calculates the relative risk and a 95% confidence interval. The relative risk is the ratio of the proportions of cases having a positive outcome in the two groups. If the value 1 is not in the range of the confidence interval, it can be concluded that the proportions are significantly different in the two groups, and there is an increased risk in one group compared to the other.

In the example, there was a positive outcome in 18 cases and a negative outcome in 32 cases in a group given treatment regimen A. In a second group with treatment regimen B, 6 cases had a positive and 44 cases had a negative outcome.

The risk in the first group was 0.36 (18/50) and in the second group 0.12 (6/50). The relative risk for a positive outcome was 3.0 (0.36/0.12) with a 95% confidence interval ranging from 1.3 to 6.9. The value 1, which means equal risks in both groups, is *not* included in this interval and the conclusion is that there is a 3-fold increased risk in group A, and this increase is statistically significant at the 5% level.

Odds ratio

When you want to calculate the odds ratio in a *retrospective case-control study* (unpaired samples), select the command *Odds ratio* in the *Tests* menu.

In a retrospective study the cases with positive and negative outcome are known and they are subsequently grouped according to the occurrence of a specific characteristic.

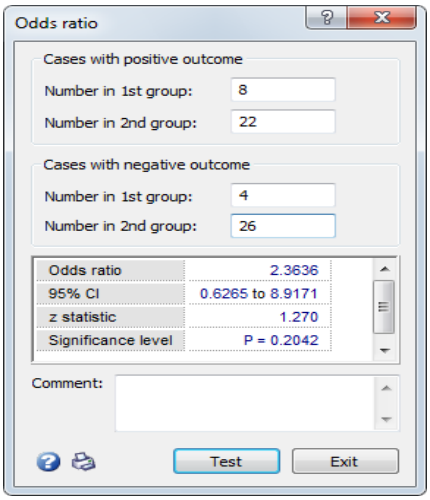

In the dialog box you enter the number of cases in the 1st and 2nd group that have a positive or negative outcome. Next, click the TEST button to perform the test.

The program will display the odds ratio, which is the ratio of the odds of the outcome in the two groups. The program also calculates a 95% confidence interval for the odds ratio. If the value 1 is not in the range of the confidence interval, then it can be concluded that there is an increased relative risk in one group compared to the other.

In the example 30 cases with a positive outcome and 30 cases with a negative outcome were selected. Next the occurrence of a particular characteristic was retrospectively investigated. For the cases with a positive outcome, 8 cases presented the characteristic and 22 did not. For the cases with a negative outcome, 4 cases presented the characteristic and 26 did not.

The resulting odds ratio is 2.4 with a 95% confidence interval ranging from 0.6 to 8.9. The value 1, which means equal odds in both groups, is included in this interval and the conclusion is that although there is a 2.4-fold increased odds of a positive outcome, this increase is not statistically significant at the 5% level.

Inter-rater agreement

Select *Inter-rater agreement* in the menu when you want to evaluate the agreement between two classification systems. If the raw data are available in the spreadsheet, use *Inter-rater agreement* in the *Statistics* menu to create the classification table and calculate Kappa (for interpretation, see p. 150).

In the dialog form you can enter the two classification systems in a 6x6 frequency table.

Select *Weighted Kappa* if the data come from an ordered scale. If the data come from a nominal scale, do not select Weighted Kappa. Use linear weights when the difference between the first and second category has the same importance as a difference between the second and third category, etc. If the difference between the first and second category is less important than a difference between the second and third category, etc., use quadratic weights.

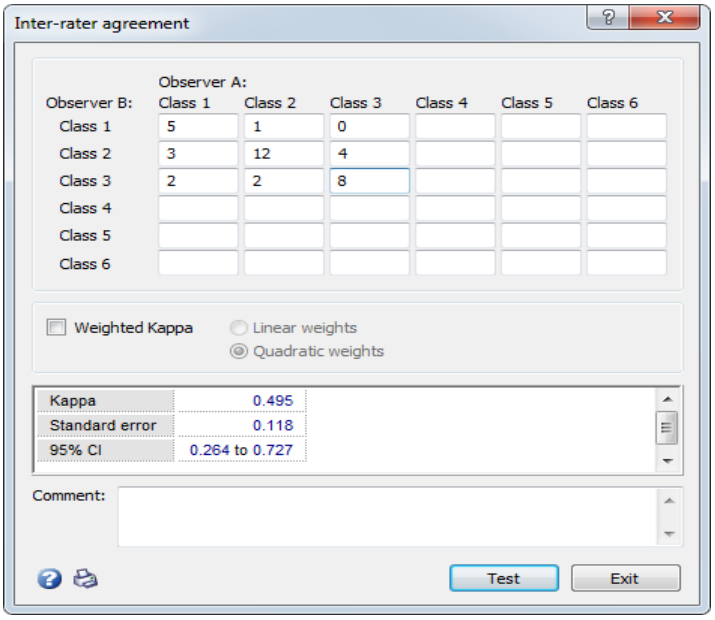

In this example, from the 6 cases that observer B has placed in class 1, observer A has placed 5 in class 1 and 1 in class 2; from the 19 cases that observer B has placed in class 2, observer A has placed 3 in class 1, 12 in class 2 and 4 in class 3; and from the 12 cases that observer B has placed in class 3, observer A has placed 2 in class 1, 2 in class 2 and 8 in class 3.

After you have entered the data, click the TEST button. The program will display the value for Kappa with its Standard Error and 95% confidence interval (CI) (Fleiss et al., 2003).

For interpretation of the Kappa statistic, see p. 150.

Diagnostic test

Description

Allows to calculate test characteristics such as sensitivity, specificity, positive and negative likelihood ratio, disease prevalence as well as positive and negative predictive power, from a 2x2 table.

Required input

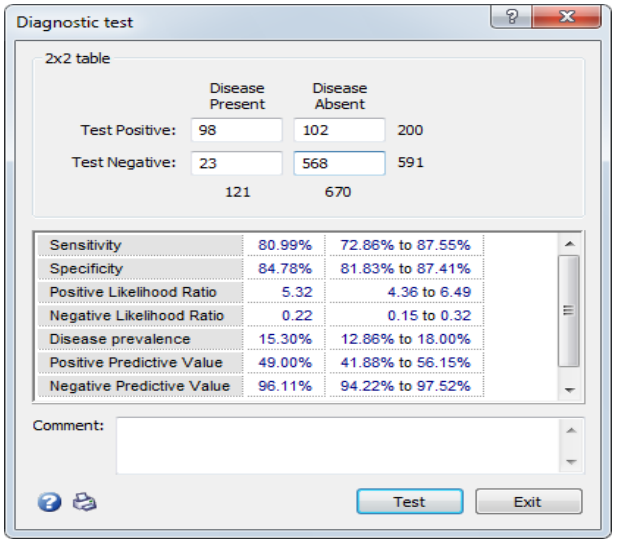

Enter the number of cases in the diseased group that test positive and negative (left column); and the number of cases in the non-diseased group that test positive and negative (right column).

Results

- *Sensitivity*: probability that a test result will be positive when the disease is present (true positive rate).
- *Specificity*: probability that a test result will be negative when the disease is not present (true negative rate).
- *Positive likelihood ratio*: ratio between the probability of a positive test result given the *presence* of the disease and the probability of a positive test result given the *absence* of the disease, i.e.

  = True positive rate / False positive rate = Sensitivity / (1-Specificity)
- *Negative likelihood ratio*: ratio between the probability of a negative test result given the *presence* of the disease and the probability of a negative test result given the *absence* of the disease, i.e.

  = False negative rate / True negative rate = (1-Sensitivity) / Specificity
- *Positive predictive value*: probability that the disease is present when the test is positive.
- *Negative predictive value*: probability that the disease is not present when the test is negative.

Sensitivity, specificity, positive and negative predictive value as well as disease prevalence are expressed as percentages for ease of interpretation.

Note

If the sample sizes in the positive (Disease present) and the negative (Disease absent) groups do not reflect the real prevalence of the disease, then the Positive and Negative predicted values cannot be estimated and you should ignore those values.

Alternatively, when the disease prevalence is known then the positive and negative predictive values can be calculated using the formula's based on Bayes' theorem:

$$PPV = \frac{sensitivity \times prevalence}{sensitivity \times prevalence + (1-specificity) \times (1-prevalence)}$$

and

$$NPV = \frac{specificity \times (1-prevalence)}{(1-sensitivity) \times prevalence + specificity \times (1-prevalence)}$$

Literature

- Gardner IA, Greiner M (2006) Receiver-operating characteristic curves and likelihood ratios: improvements over traditional methods for the evaluation and application of veterinary clinical pathology tests. Veterinary Clinical Pathology, 35:8-17.
- Griner PF, Mayewski RJ, Mushlin AI, Greenland P (1981) Selection and interpretation of diagnostic tests and procedures. Annals of Internal Medicine, 94, 555-600.
- Hanley JA, McNeil BJ (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology, 143, 29-36.
- Metz CE (1978) Basic principles of ROC analysis. Seminars in Nuclear Medicine, 8, 283-298.
- Zhou XH, NA Obuchowski, DK McClish (2002) Statistical methods in diagnostic medicine. New York: Wiley.
- Zweig MH, Campbell G (1993) Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. Clinical Chemistry, 39, 561-577.

Likelihood ratios (2xk table)

Description

Allows to calculate likelihood ratios for different test levels from a 2xk table.

When test results have a continuous or ordinal outcome then valuable information is lost when the data are dichotomized for the calculation of sensitivity, specificity and likelihood ratios as in ROC curve analysis.

Interval likelihood ratios may be more powerful because they use more information contained in the data.

The likelihood ratio can be used to calculate the post-test probability of disease from the pre-test probability of disease (see below).

Required input

Enter the number of cases in the diseased group that test positive and negative at the different test levels.

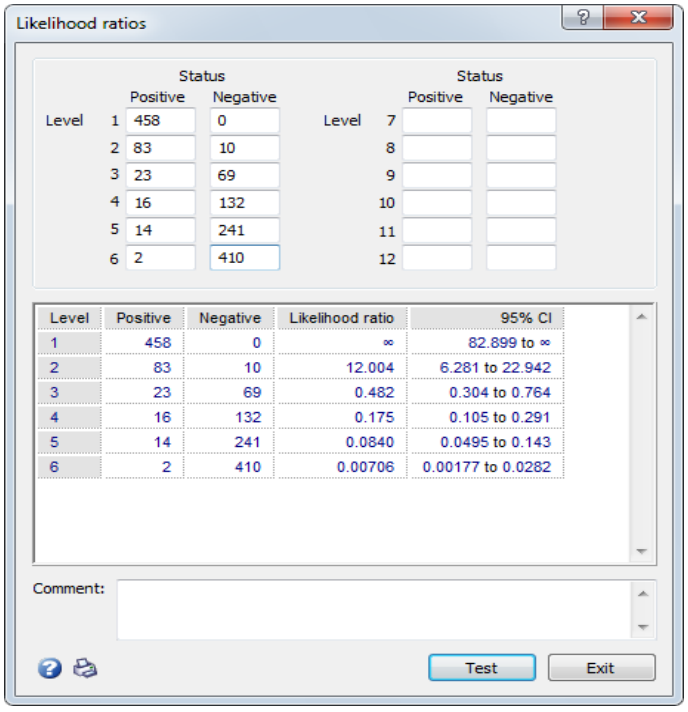

Results

For each test levels the program calculates corresponding Likelihood ratio with 95% Confidence interval.

The likelihood ratio can be used to calculate the post-test odds from the pre-test odds of disease:

post-test odds = pre-test odds x likelihood ratio

The relation between odds and probability is:

$$\text{odds} = \frac{p}{1-p} \text{ and } p = \frac{\text{odds}}{1+\text{odds}}$$

Using these equations, you can calculate the post-test probability of disease from the pre-test probability of disease.

If, for example, the pre-test probability of disease is 0.6 then the pre-test odds is 0.6/(1-0.6) = 1.5. For a case with a test result corresponding with diagnostic level 2, the likelihood ratio is 12, and the post-test odds is 1.5 x 12 = 18. The post-test probability of disease is 18/(1+18) = 0.95.

Literature

- Gardner IA, Greiner M (2006) Receiver-operating characteristic curves and likelihood ratios: improvements over traditional methods for the evaluation and application of veterinary clinical pathology tests. Veterinary Clinical Pathology, 35:8-17.

SAMPLING

Introduction

In the *Sampling* menu, you can calculate the required sample size for some common problems, taking into account the magnitude of differences and the probability to make a correct or a false conclusion.

When you perform a statistical test, you will make a *correct decision* when you

- reject a false null hypothesis, or
- accept a true null hypothesis.

On the other hand you can make *two errors*:

- you can reject a true null hypothesis, or
- you can accept a false null hypothesis.

These four situations are represented in the following table.

	Null hypothesis = TRUE	*Null hypothesis = FALSE*
Reject null hypothesis	Type I error α	Correct decision
Accept null hypothesis	Correct decision	Type II error β

For example, when you have rejected the null hypothesis in a statistical test (because P<0.05), and therefore conclude that a difference between samples exists, you can either:

- have done so correctly, and uncovered a difference where one exists;
- have rejected the null hypothesis when in fact it is true, and uncovered a difference where in fact none exits. In this case you make a Type I error. α is the (two-sided) probability of making a Type I error.

Type I error = rejecting the null hypothesis when it is true

You can avoid making a Type I error by selecting a lower significance level of the test, e.g. by rejecting the null hypothesis when P<0.01 instead of P<0.05.

On the other hand, when you accept the null hypothesis in a statistical test (because P>0.05), and conclude that there is no difference between samples, you can either:

- have correctly concluded that there is no difference;
- have accepted the null hypothesis when in fact it is false, and therefore you have failed to uncover a difference where such a difference really exists. In this case you make a Type II error. β is the probability of making a Type II error.

Type II error = accepting the null hypothesis when it is false

The power of a test is 1-β, this is the probability to uncover a difference when there really is one. For example when β is 0.10, then the power of the test is 0.90 or 90%.

Power = probability to achieve statistical significance

You can avoid making a Type II error, and increase the power of the test to uncover a difference when there really is one, mainly by increasing the *sample size*.

To calculate the required sample size, you must decide beforehand on:

- the required probability α of a Type I error, i.e. the required significance level (two-sided);
- the required probability β of a Type II error, i.e. the required power 1-β of the test;

- a quantification of the study objectives, i.e. decide what difference is biologically or clinically meaningful and worthwhile detecting.

In addition, you will sometimes need to have an idea about expected sample statistics such as e.g. the standard deviation. This can be known from previous studies.

Single mean

Description

Calculates the required sample size for the comparison of a mean with a given value.

Required input

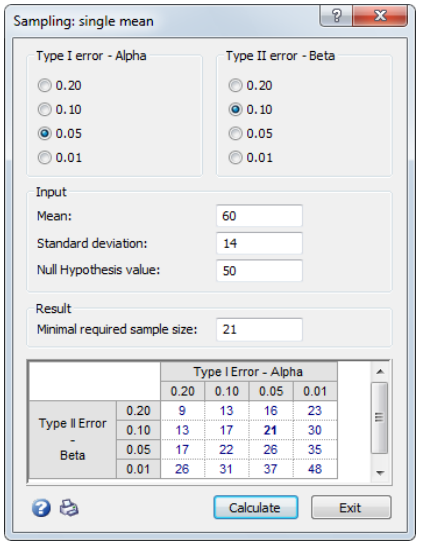

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Mean: the hypothesized mean (considered to be biologically significant different from the null hypothesis value).
- Standard deviation: hypothesized standard deviation (known from previous studies or from the literature).
- Null hypothesis value: the null hypothesis value.

Example

If for example the null hypothesis value is 50, and you consider a value of at least 60 to be significantly different, then you will enter these numbers for Null hypothesis value and Mean. From a previous study, you expect the standard deviation of the sample to be 14. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size (21 in the example).

A table shows the required sample size for different Type I and Type II Error levels.

Comparison of two paired samples

To calculate the sample size required for the comparison of two paired samples, you use the same procedure as for a Single mean.

In the dialog box, for Mean, you enter the hypothesized *mean difference*; for Standard deviation, enter the hypothesized standard deviation of the differences, and for Null hypothesis value enter the value 0.

Single proportion

Description

Calculates the required sample size for the comparison of a proportion with a given proportion.

Required input

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Proportion (%): the hypothesized proportion (considered to be biologically significant different from the null hypothesis value), expressed as a percentage.
- Null hypothesis value (%): the null hypothesis value (a proportion, expressed as a percentage).

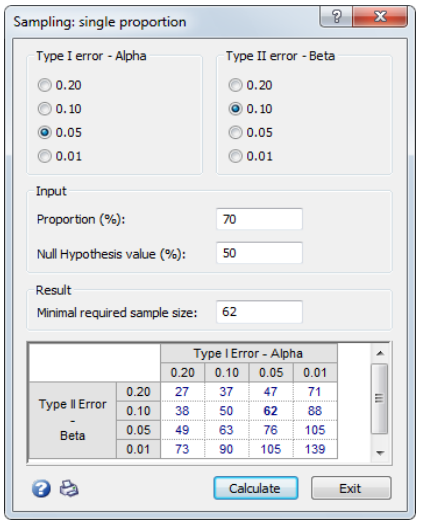

Example

For example, if the null hypothesis value is 50%, and you consider a proportion of at least 70% to be significantly different, then you enter the numbers 70 and 50 in the dialog box. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size (62 in the example).

A table shows the required sample size for different Type I and Type II Error levels.

Comparison of two means

Description

Calculates the required sample size for the comparison of two independent means.

Required input

- Type I error - alpha: the probability of making a Type I error (□-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Difference: the hypothesized difference (considered to be biologically significant).
- Standard deviation 1: hypothesized standard deviation in the first sample.
- Standard deviation 2: hypothesized standard deviation in the second sample.

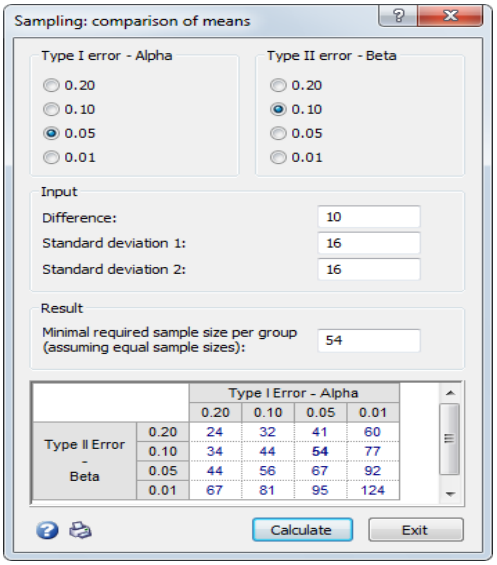

Example

For example, you are interested in detecting a difference between the sample means of at least 10. You expect the standard deviations in the two study groups to be equal to 16. Enter the value 10 for difference, and enter 16 for both standard deviations. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size (54 in the example, meaning that you will need 54 cases in each group, or a total of 108 cases).

A table shows the required sample size for different Type I and Type II Error levels.

Comparison of two paired samples

To calculate the sample size required for the comparison of two paired samples, see *Sampling for single mean*, p. 201.

Comparison of two proportions

Description

Calculates the required sample size for the comparison of two proportions.

Required input

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- First proportion (%): hypothesized proportion in the first sample.
- Second proportion (%): hypothesized proportion in the second sample (the hypothesized difference with the first proportion is considered to be biologically significant).

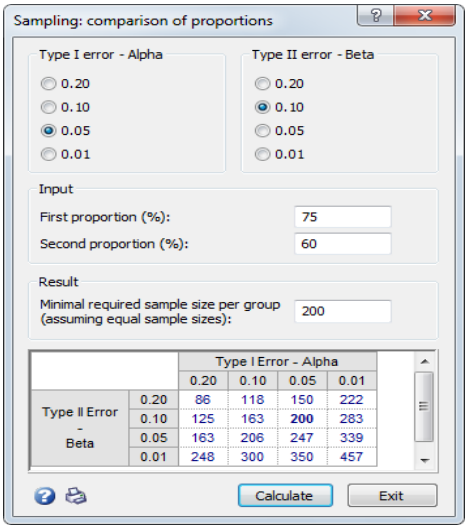

Example

For example, you are interested in detecting a difference between two proportions of at least 15. You expect the two proportions to be equal to 75 and 60, so enter these values in the dialog box. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size (200 in the example, meaning that you will need 200 cases in each group, or 400 cases in total).

A table shows the required sample size for different Type I and Type II Error levels.

Correlation coefficient

Description

Calculates the required sample size for a correlation coefficient. The sample size takes into account the required significance level and power of the test.

Required input

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.

- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Correlation coefficient: the hypothesized or anticipated correlation coefficient.

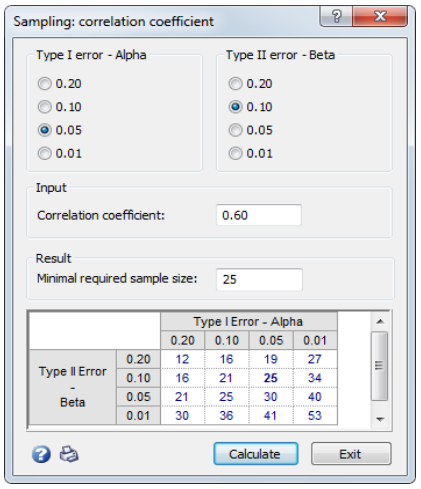

Example

For example, the correlation coefficient between two variables is thought to be 0.60. How many patients are required for this correlation coefficient to be significant different from 0.0? For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size (25 in the example, meaning that you will need 25 cases in which both variables must be measured).

A table shows the required sample size for different Type I and Type II Error levels.

Area under ROC curve

Description

Calculates the required sample size for the comparison of the area under a ROC curve with a null hypothesis value.

Required input

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Area under ROC curve: the hypothesized Area under the ROC curve (the AUC expected to be found in the study).
- Null hypothesis value: the null hypothesis AUC.

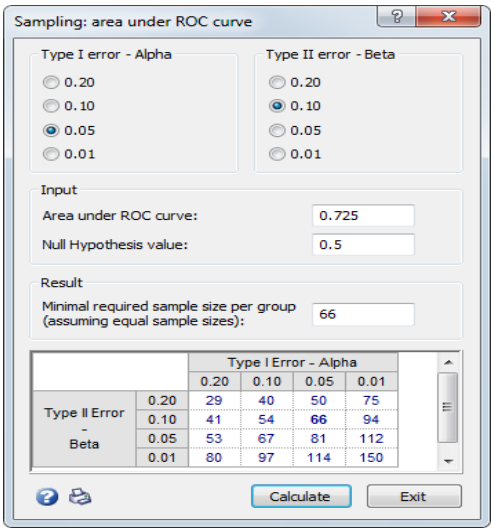

Example

If for example you want to show that the AUC of 0.725 for a particular test is significant from the null hypothesis value 0.5 (meaning no discriminating power), then you enter 0.725 for Area under ROC curve and 0.5 for Null Hypothesis value. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size. In our example this number is 66, meaning that you need 66 normal subjects (negative group) and 66 abnormal subjects (positive group), assuming equal numbers of normal and abnormal subjects. So in fact a total of 132 subjects is required.

A table shows the required sample size for different Type I and Type II Error levels.

Reference

- Hanley JA, McNeil BJ (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology, 143:29-36.

Comparison of 2 ROC curves

Description

Calculates the required sample size for the comparison of the areas under two ROC curves (derived from the same cases).

Required input

- Type I error - alpha: the probability of making a Type I error (α-level, two-sided), i.e. the probability of rejecting the null hypothesis when in fact it is true.
- Type II error - beta: the probability of making a Type II error (β-level), i.e. the probability of accepting the null hypothesis when in fact it is false.
- Area under ROC curve 1: hypothesized area for the first ROC curve.
- Area under ROC curve 2: hypothesized area for the second ROC curve.
- Correlation in positive group: the hypothesized rank correlation coefficient in the positive group (abnormal cases)
- Correlation in negative group: the hypothesized rank correlation coefficient in the negative group (normal cases)

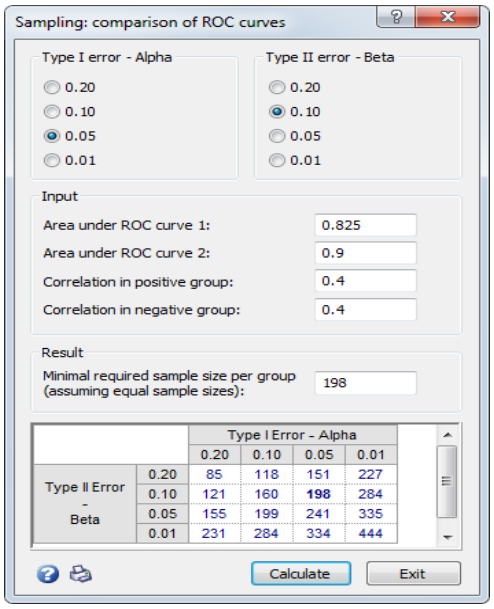

Example

For example: you are interested to show that the discriminating power of two assays (performed on the same cases), with an area under the ROC curve of 0.825 and 0.9, is significant different. From previous studies you know that the rank correlation between the two assays is 0.4 in both positive and negative cases.

You enter the values 0.825 and 0.9 for Area under ROC curve 1 and Area under ROC curve 2. Next you enter 0.4 for Correlation in positive group and Correlation in negative group. For α-level you select 0.05 and for β-level you select 0.10.

Results

After you click the CALCULATE button the program will display the required sample size. In our example this is 198, meaning that you need 198 normal subjects (negative group) and 198 abnormal subjects (positive group), assuming equal numbers of normal and abnormal subjects. So in fact a total of 396 subjects is required.

A table shows the required sample size for different Type I and Type II Error levels.

References

- Hanley JA, McNeil BJ (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology, 143:29-36.
- Hanley JA, McNeil BJ (1983) A method of comparing the areas under receiver operating characteristic curves derived from the same cases. Radiology, 148:839-843.

WINDOW MENU

Cascade

Shortcut: Shift+F5

Use this command to layer the windows in the MedCalc application window so that the title bar of each window is visible.

Tile

Shortcut: Shift+F4

With this command you can arrange the windows in the MedCalc application window side by side and in equal sizes and so that all of them are visible.

Arrange icons

This command arranges the iconized windows (spreadsheet, graphs, results) across the bottom of the MedCalc application window.

Split window – Remove split

The spreadsheet window can be divided into two panes. The two panes scroll together vertically, but can scroll independently horizontally.

To split the window into panes:

- Select the *Split window* command in the Window menu.
- Drag the *Split box* (which is displayed in the bottom right corner of the spreadsheet window) into the spreadsheet.

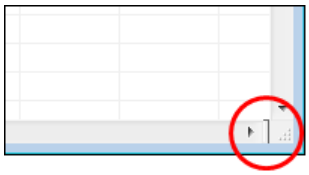

- Double-click the *Split box*.

After you have performed one of these actions, the *Split box* changes into a *Split bar* dividing the spreadsheet into two panes. You can change the position of the *Split bar* by dragging it to the left or right in the window.

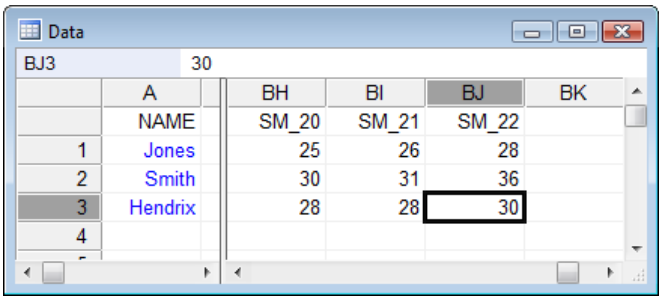

To remove the split:

- Choose *Remove split* from the *Window* menu.
- Drag the *Split box* to the utmost left or right of the spreadsheet.
- Double-click the *Split bar*.

Close all

With this command you can close all windows (spreadsheet, notes, graphs, results) in the MedCalc application window.

HELP MENU

On line help is always available by means of function key *F1*. In addition, you can select *Contents and Index* in the *Help* menu to obtain explanation on working with MedCalc.

Contents and Index

You select *Contents and Index* to browse the comprehensive on-line documentation. In Help, click one of the following tabs:

- Click the **Contents** tab to browse through topics by category.
- Click the **Index** tab to see a list of index entries, and then either type a word or scroll through the list.
- Click the **Find** tab to search for words or phrases that may be contained in a Help topic.

How to use Help

If you are new to Windows, or to the Windows Help system, select *How to use Help*. This command will activate the Windows documentation on how to use the Windows help system.

What's new

Select *What's new* in the *Help* menu to read more about what is new in your current copy of MedCalc.

MedCalc on the Web

When you select this command the system will open your default web browser (e.g. Netscape or Internet Explorer) and establish a connection with the MedCalc Web site on the Internet.

The following options are available:

- MedCalc home page: connect to the MedCalc home page for general information
- Frequently asked questions: connect to the MedCalc web site for some frequently asked questions
- On line support: connect to the MedCalc web site for specific questions and comments
- On line manual: connect to the MedCalc web site containing the complete MedCalc manual
- MedCalc update: connect to the MedCalc web site for software updates.

Register

In this dialog box you can enter your user name and product key. This will end the trial period and you can use the software without limitations.

About MedCalc

The *About MedCalc* box displays the program version number and copyright notice. Click the OK button to remove the box from the screen.

MEDCALC SPREADSHEET

Introduction

In MedCalc, as in other spreadsheets, a cell can contain a text or string entry, or it can contain a number or a formula.

An example or a string entry is a name such as e.g. Howard. When used in a formula, literal strings must be placed between quotation marks, e.g. “Howard”.

A cell can also contain a number, negative or positive, and with or without a decimal fraction, e.g. 5, -23.5, 0.03. The number of decimals displayed globally in the spreadsheet is selected in the *Format spreadsheet* box (see p.40). However, a different number of decimals may be selected for a particular column.

A formula may be a simple mathematical formula, such as e.g. SQRT(36). This formula will return the value 6, which is the square root of the number 36. Alternatively, the number 36 may be entered in a different cell of the spreadsheet, and the SQRT function can take the address of this cell as an argument, e.g. SQRT(B5).

Most of the formulas available in MedCalc work similar to their equivalents in other spreadsheet programs. When a formula begins with a cell address, this must be preceded by a = or + sign, but MedCalc formulas must not be preceded by the = character.

Operators

Arithmetic operators

The arithmetic operators, in order of precedence, are :

^	Exponentiation
-	Negation
* /	Multiplication, division
+ -	Addition, subtraction

Relational operators

Relational operators compare two numbers or two strings. The result is a logical value expressed as a number, either 1 (=TRUE), or 0 (=FALSE). The relational operators, in order of precedence, are:

=	Equality
<>	Inequality
<	Less than
>	Greater than
<=	Less than or equal to
>=	Greater than or equal to

Combination of operators

When arithmetic and relational operators are combined in one expression, the arithmetic are performed first. You have to use parentheses when you want to change the order in which the operations are performed. Operations within parentheses are performed first.

B6<5+3	Returns 1 if the contents of B6 is less than 8, otherwise this expression returns 0.
5+3*(A1>3)	Returns 8 if cell A1 contains a value greater than 3 (since A1>3 = TRUE = 1). If A1 contains 3 or a number less than 3, this expression returns 5 (A1>3 = FALSE = 0).

Relative and absolute cell addresses

If you want to use the value of a cell in a formula in another cell of the spreadsheet, then you refer to this cell by means of its cell address. This cell address consists of a column indicator and a row number, e.g. cell D14 is the cell in column D, row 14.

When a cell address includes a $ character before the column or row number, the address is considered as an absolute address. When cell C10 contains for instance the formula $B$5+$B$6, the actual addresses of these two cells are stored in memory. When the formula is copied or moved to cell C11, e.g. by inserting a row at row 9, the formula will still be $B$5+$B$6.

When a cell address does not include any $ characters, the address is considered to be a relative address. This means that the program does not store the actual address, but instead the program stores the number of columns and rows calculated relative to the cell containing the formula. When cell B5 contains the cell address B4, then the program does not store the address B4 as such, but it stores how to get from cell B5 to cell B4, in this case one row up. When the formula 2*B4 is copied from cell B5 to cell D10, the formula will be changed to 2*D9.

When cell B1 contains the formula SUM(B2:B10), and this formula is copied to cell C1, the formula will be converted to SUM(C2:C10). When it is copied to cell D51, the formula will be converted to SUM(D52:D60).

A cell address or formula may contain a combination of absolute and relative cell addressing, e.g.: A$2, $B10, $C5, SUM(B3:B20)*$A$1.

Ranges of cells are identified by two cell addresses separated by a colon. These cell addresses may also include relative and absolute cell addresses, e.g. B2:E8, $A$10:$A$25, D10:D10, etc.

When you design a spreadsheet, you should pay attention to when to use relative and when to use absolute cell addressing.

Mathematical functions

ABS Absolute value

ABS(*x*) returns the absolute value of the number *x*.

ACOS Arc cosine function

ACOS(*x*) returns the arccosine of *x*. The arccosine function is the inverse function of the cosine function and calculates the angle for a given cosine. *X* must be in the [-1..1] range. The result is a number in the range [0..180] (degrees).

ASIN Arc sine function

ASIN(*x*) returns the arcsine of *x*. The arcsine function is the inverse function of the sine function and calculates the angle for a given sine. *X* must be in the range [-1..1]. The result is a number in the range [-90..90] (degrees).

ATAN Arc tangent function

ATAN(*x*) returns the arctangent of *x*. The arctangent function is the inverse function of the tangent function and calculates the angle for a given tangent. The result is a number in the range [-90..90] (degrees).

COS Cosine function

COS(*x*) returns the cosine of an angle *x*. The result range is [-1..1].

EXP Natural exponential function

EXP(*x*) returns the natural exponential of *x*: 2.718281828^{X}. EXP is the inverse function of the LN function.

FACT Factorial function

FACT(x) returns the factorial of x. The factorial of a number *n* is equal to 1*2*3*...**n*.

FACT(5) equals 120 (=1*2*3*4*5)

INT Integer value function

INT(*x*) rounds the number *x* down to an integer.

LN Natural logarithm function

LN(*x*) returns the natural logarithm of the positive number *x* to the base *e* (*e*=2.718281828). The argument *x* must be greater than 0. LN is the inverse function of EXP.

LOG Logarithm function

LOG(*x*) returns the logarithm of the positive number *x* in the base 10. The argument *x* must be greater than 0.

MOD Modulo function

MOD(*x*,*d*) returns the remainder after *x* is divided by *d*. The result has the same sign as the divisor *d*.

PI Pi function

The function Pi() takes no argument and returns the value 3.14159265358979.

POWER Power function

POWER(*n*,*p*) returns *n* raised to the power *p*.

RAND Random number function (Uniform distribution)

RAND(*x*) returns a computer-generated random number (a) when $x \leq 1$ the result is a number between 0 and 1, or (b) when $x > 1$ then the result is a number between 1 and *x*. The function will return a different value every time the spreadsheet is recalculated.

RANDNORM Random number function (Normal distribution)

The function RANDNORM(*m*,*s*) returns a random number from a Normal distribution with mean *m* and standard deviation *s*.

ROUND Round function

ROUND(x,n) rounds the number *x* to the number of digits *n*. If *n* is negative, then *x* is rounded to the left of the decimal point.

SIGN Sign value

SIGN(*x*) returns a number that indicates the sign *x*: -1 if *x* is negative; 0 if *x* equals 0; or 1 if *x* is positive.

SIN Sine function

SIN(*x*) returns the sine of the angle *x*. The result range is [-1..1].

SQRT Square root function

SQRT(*x*) returns the square root of the positive number *x*. SQRT(*x*) = *x*^(1/2). If the argument *x* of the function is negative, then the function returns an error (missing value).

TAN Tangent function

TAN(*x*) returns the tangent of the angle *x*.

TRUNC Truncate function

TRUNC(*x*) truncates the number *x* to an integer by removing the fractional part of the number *x*.

Statistical functions

In the spreadsheet window statistical functions can be entered that perform a calculation on one or two variables or ranges of cells containing numeric values.

For all statistical functions that accept a range as argument, a list of ranges is also accepted as an argument.

E.g. SUM(A1:A3,A10,B5:B6) will calculate the sum of cells A1, A2, A3, A10, B5 and B6.

AVEDEV Average of absolute deviations

AVEDEV(*range*) returns the average of absolute deviations of the data in *range.*

AVERAGE Average

AVERAGE(*range*) computes the arithmetic mean of the contents of the cells in the specified *range*. If one of the cells in the function's range does not have a numeric value, but is empty or has a string value, then this cell will not be taken into account for calculating the average, or any of the other statistical functions.

AVERAGE(A1:B1)	returns the mean of the contents of cells A1 to B1.

CHIDIST One-tailed probability from the Chi Squared distribution

CHIDIST(*chisquared,df*) returns the probability P associated with the test statistic *chisquared* and *df* degrees of freedom.

CHIDIST(18.307,10)	returns 0.05 (rounded)

CHIINV Returns the reverse of the chi-squared distribution

CHIINV(p*,df*) returns the chi-squared value corresponding with the one-tailed P-value *p* and the specified degrees of freedom *df*. CHIINV is the inverse of the CHIDIST function.

COEFVAR Coefficient of variation

COEFVAR(*range*) returns the coefficient of variation of the data in *range.*

COUNT Count

COUNT(*range*) counts the number of cells in the specified *range* that have a numerical contents. See also the COUNTS(range) function: this function counts the number of non-empty cells.

COUNT(A1:A5)	Returns 4 if the contents of cells A1, A2,A3 and A5 is a number, and A4 is blank or contains non-numeric data.

The COUNT function has 2 variants:

COUNTNEG(range)	counts the number of negative values in *range* (< 0);
COUNTPOS(range)	counts the number of positive numbers in *range* (> 0);

COUNTS Count non-empty cells

COUNTS(*range*) counts the number of cells in *range* that are non-empty, irrespective if the cell contains a formula or a numeric or text value. This distinguishes the COUNTS function from the COUNT function that only counts cells containing a numeric value, or a formula resulting in a numeric value.

COVAR Covariance function

COVAR(range1,range2) returns the covariance (the average of the products of deviations for each data point pair).

FDIST One-tailed probability from the F distribution

FDIST(*F,v1,v2*) returns the one-tailed probability P associated with the test statistic *F* with *v1* degrees of freedom for the numerator and *v2* degrees of freedom for the denominator.

FDIST(3.291,6,40)	returns 0.01 (rounded)

FINV Inverse of the F probability distribution

FINV(*p,v1,v2*) returns the inverse of the F probability distribution where *p* is a probability associated with the F cumulative distribution, *v1* is the numerator degrees of freedom and *v2* is the denominator degrees of freedom. FINV is the inverse of the FDIST function.

GEOMEAN Geometric mean

GEOMEAN(*range*) returns the geometric mean of the data in *range.*

MAX Maximum

MAX(*range*) returns the maximum value of the contents of the cells in the specified *range*.

MAX(A1:D1,Z1)	returns the maximum value of the contents of cells A1 to D1 (A1, B1, C1, D1) and cell Z1.

MIN Minimum

MIN(*range*) returns the minimum value of the contents of the cells in the specified *range*.

NORMSDIST One-tailed probability from standardized Normal distribution

NORMSDIST(z) returns the one-tailed probability associated with the standardized Normal deviate *z*.

NORMSDIST(-1.96)	returns 0.025 (rounded)
NORMSDIST(1.96)	returns 0.975 (rounded)
NORMSDIST(0)	returns 0.5

NORMSINV Returns a Standardized Normal deviate

NORMSINV(P) returns the standardized Normal deviate z corresponding with the one-tailed probability P. P must be a value between 0 and 1 (0<P<1). NORMSINV is the inverse function of the NORMSDIST function.

NORMSINV(0.025)	returns -1.96 (rounded)
NORMSINV(0)	returns an error

SEM Standard error of the mean

SEM(*range*) returns the standard error of the mean of the data in *range*.

STDEV Standard deviation - sample

STDEV(*range*) calculates the 'sample' standard deviation of the data in *range* (divisor *n*-1).

STDEV(B3:E3)	calculates the standard deviation of the contents of cells B3, C3, D3 and E3.

SUM Sum

SUM(*range*) computes the sum of the contents of the cells in the specified *range*. The SUM function is probably the most frequently used function in any spreadsheet model.

SUM(A1:D1)	calculates the sum of the contents of cells A1, B1, C1 and D1.

The SUM function has 2 variants: SUMNEG(*range*) and SUMPOS(*range*) calculating the sum of respectively the negative and positive values in *range*.

TDIST Two-tailed probability from the Student t distribution

TDIST(t,df) returns the two-tailed probability P associated with the test statistic t and df degrees of freedom.

TDIST(3.1693,10)	returns 0.01 (rounded)

TINV Returns a value t from the Student t distribution

TINV(P,df) returns the t-value corresponding with the two-tailed P-value P and the specified degrees of freedom df. TINV is the inverse function of the TDIST function.

TINV(0.05,30)	returns 2.0423 (rounded)

VAR Variance - sample

VAR(*range*) computes the correct 'sample' variance of the data in the specified *range* (divisor *n*-1).

VAR(D8:G8)	calculates the variance of the contents of cells D8 to G8.

String functions

String functions are functions that have a non-numerical result. A string is a sequence of characters not interpreted as a number, e.g. Jones, "25".

CELL Cell text/formula function

CELL(*column,row*) returns the contents of the cell with coordinates *column* and *row* as text.

CELL("A",5)	returns the text 5*6 when cell A5 contains 5*6.

CHAR Character function

CHAR(*x*) returns the character with ANSI code *x*. The result is a string value of length 1 (if x>0).

CHAR(65.0) returns the string value 'A'

See p. 230 for a table of the ANSI character set.

CODE Character to ANSI code conversion

CODE(*str*) returns the ANSI code number of the first character of the string *str*. If the length of the string is 0, i.e. the string is empty, then the function returns 0.

CODE("Andy") returns 65.00

See p. 230 for a table of the ANSI character set.

CONCAT Concatenate strings

CONCAT(str1,str2, ...) joins two or more strings (text items) into one single string.

CONCAT("Total ","value") returns "Total value".

LEFT Left portion of string

LEFT(*str*,*n*) returns the first *n* characters of *str*. If *n* equals 0, then the LEFT function returns an empty string. If *n* is equal to or more than the length of the string *str*, the function returns the complete string.

LEFT("Position",3) returns "Pos"

LEN Length of string

LEN(*str*) returns the length of the string *str*.

LOWER Lowercase function

LOWER(str) converts the string *str* to lowercase.

LOWER('TOTAL') returns "total"

MID Middle portion of a string

MID(*str,pos,n*) returns *n* characters from *str* starting at position *pos.*

MID("statistics,3,4) returns "atis"

REPT Repeat function

REPT(*str,n*) creates a string consisting of *str* repeated *n* time.

REPT("*",5) returns "*****"

RIGHT Right portion of a string

RIGHT(*str*,*n*) returns the last *n* characters of *str*. If *n* equals 0, then the RIGHT function returns an empty string. If *n* is equal to or more than the length of the string *str*, the function returns the complete string.

STR Number to string conversion

STR(*x*,*n*) returns the numeric value *x* as a string, with *n* decimal places.

STR(25.56,1) returns the string value "25.6"

UPPER Uppercase function

UPPER(str) converts the string *str* to uppercase.

UPPER('total') returns "TOTAL"

VAL String to number conversion

VALUE(*str*) evaluates *str* as a number.

VALUE("25.0") returns 25.0
VALUE("text") returns an error

Date functions

In MedCalc, date strings, for example 5.12.72 or 5.7.1956 can be converted to serial date numbers. A serial date number is an integer number indicating the number of days since 01.01.1801. These serial date numbers can be used for comparison, sorting, arithmetic operations and statistical analysis.

A serial date number can be back-transformed to a date string using the DATE function.

The date format, i.e. the way the serial date number is converted to a string, can be chosen in the *International* dialog box of the *Control box* window that is part of the Windows operating system (see your Windows documentation).

When you want to enter date strings in the spreadsheet, it is not convenient to use the "/" character as a separator between day, month and year, because this would be interpreted as a division (10/12/88 is the number 0.0094697), unless you use quotation marks ("10/12/88" is a legal date string), or select "Text format" for the column in which you enter the dates (see p. 40). However, MedCalc allows entering dates using a dot or any non-numerical character as a separator: e.g. 10.12.88.

When you enter a date with a year less than 100, MedCalc will interpret this year as a year in the 20th century: 20.12.88 = 20.12.1988. When you enter 0 as the year number, MedCalc will interpret this as the year 1900: 20.12.00 = 20.12.1900. So when you want to enter a date situated in the 19th or 21st century, enter the year using 4 digits: 01.01.2000.

DATE Serial date number to string

DATE(*dnr*) returns the date corresponding with the serial date number *dnr* expressed as a string.

DATE(DATEVALUE(A1)+7) returns "6.9.96" when cell A1 contains the date string 30.8.96 or "30/8/96" (date format DD.MM.YY)

DATEVALUE String to serial date number

The function DATEVALUE(*str*) returns the serial date number for the date expressed in the string *str*.

DATEVALUE("10.12.88") returns 68645 if the date format is DD.MM.YY.

DATEFRAC Fractional year-number

DATEFRAC(*date*) converts *date* into a fractional year-number. The integer part of this number is the year, and the decimal fraction ranges from 0.0 to 0.99..., representing the dates 01 Jan to 31 Dec.

DATEFRAC("01.07.2000") returns 2000.5

DAY Day function

DAY(*date*) returns the day of the month of *date*. *Date* can either be a serial date number or a date string.

DAY("23.08.88") returns 23 (date format DD.MM.YY)

DAYNAME Day name function

DAYNAME(*date*) returns the name of the day of *date*. *Date* can either be a serial date number or a date string.

DAYNAME("03.12.2001") returns "Monday" (date format DD.MM.YY)

MONTH Month function

MONTH(*date*) returns the month of the year of *date*. *Date* can either be a serial date number or a date string.

MONTH("23.08.88") returns 8 (date format DD.MM.YY)

MONTHNAME Month name function

MONTHNAME(*date*) returns the name of the month of *date*. *Date* can either be a serial date number or a date string.

MONTHNAME("03.12.2001") returns "December" (date format DD.MM.YY)

TIMEVALUE String to serial time number

TIMEVALUE(str) converts a time string (in "hh:mm" format) to a serial time number. A serial time number is a number ranging from 0.00 to 1.00 representing the time 0:00 to 24:00.

TIMEVALUE("06:30") returns 0.27084

TIMEVALUE("12:00") returns 0.5

TIME Serial time number to string

TIME(*tnr*) converts the serial time number *tnr* to a time string (hh:mm).

TIME(0.5)	returns 12:00

TIMEFRAC Fractional time-number

TIMEFRAC(*time*) converts *time* into fractional time-number. The integer part of this number is the hour, and the decimal fraction ranges from 0.0 to 0.99, representing the minutes 0:00 to 0:59.

TIMEFRAC("12:30")	returns 12.5

TODAY Today's date

The function TODAY() takes no argument and returns today's date. Note that the value of this function is not fixed, so if you use this function in a file and you reload this file the next day, the function will return the new date.

TODAY()	returns 24.01.1998 if today is the 24th of January 1998 and the date format is DD.MM.YY.
DATEVALUE(TODAY())	returns the serial number of today's date.

WEEKNUM Week number function

WEEKNUM(*date*) returns the number of the week of the year of *date*. *Date* can either be a serial date number or a date string.

WEEKNUM("01.01.98")	returns 1
WEEKNUM("05.02.98")	returns 2
WEEKNUM("25.12.98")	returns 52

WEEKDAY Weekday number function

WEEKDAY(*date*) returns the weekday number of *date*. Days are numbered from Monday (=1) to Sunday (=7).

WEEKDAY("05.07.56")	returns 4, this is the 4th day of the week (=Thursday).

YEAR Year function

YEAR(*date*) returns the year of *date*. *Date* can either be a serial date number or a date string.

YEAR("23.08.88")	returns 88 (date format DD.MM.YY)
YEAR(TODAY())	returns today's year.

Logical functions

AND And function

AND(*condition1,condition2,...*) returns 1 (=TRUE) of all of its arguments are TRUE. If one of the arguments, either *condition1*, *condition2*, etc. is FALSE (or 0), the function returns 0 (=FALSE). If one of the conditions cannot be evaluated (for example because of missing data), the function returns an error.

(LENGTH>160) AND (LENGTH<170) can be written as:

AND(LENGTH>160,LENGTH<170)

CATEGORISE Categorise function

CATEGORISE("*variable*","*condition1*",*value1*,"*condition2*",*value2*,...,"*conditionN*",*valueN* [,*defaultvalue*]) recodes a variable into different categories. If "*condition1*" is true then the function returns *value1*, else if *condition2* is true then the function returns *value2*, and so on.

The first parameter is the variable name and must be placed between quotation marks. The following parameters are a serious of conditions and values. The conditions must be placed between quotation marks. The list of conditions is evaluated from left to right. If a condition is true, the value that follows is returned as the result of the function.

The last (optional) parameter of the function specifies a default value which is returned when none of the conditions is true.

CATEGORISE("VALUE","=0","Zero","<0","Negative",">0","Positive") returns the string value “Zero" when the variable VALUE contains the value 0, the string value "Negative" is returned when the variable VALUE is less than 0.

CATEGORISE("AGE",">60","old",">25,"middle-aged","young") returns the string value "old" when the variable AGE contains a value higher than 60, "middle-aged" is returned when AGE" is higher than 25. In all other cases the function returns "young".

IF If-then-else function

IF(*condition*,*x*,*y*) returns *x* if the *condition* is TRUE (=1), but returns *y* if the *condition* is FALSE (=0).

IF(C2<0,”NEG”,”POS”)	returns the string value “NEG” if the number in cell C2 is less than 0 (C2<0 = TRUE). If cell C2 contains a number equal to or more than 0, then this function returns the string value “POS”).
IF(A1>1,25,33)	returns 25 if cell A1 contains a value greater than 1. If A1 contains 1 or a value less than 1, then this expression returns 33.

This function can also be used for converting continuous data into discrete data: the IF function can be nested in order to create 3 (or more) groups. When you want to convert the variable AGE into codes for age groups of less than 30 years, 30 to 39 years and 40 or more years, you can use the following formula:

IF(AGE < 30 **,** 1 **, IF(** AGE < 40 **,** 2 **,** 3 **))**

ISNUMBER Cell-is-number function

ISNUMBER(*var*) results in the logical value 1 (=TRUE) if the calculation of *var* results in a number.

ISNUMBER(25.6)	returns 1
ISNUMBER(B2)	returns 0 (=FALSE) if cell B2 contains e.g. the string value “SMITH”

ISSTRING Cell-is-string function

ISSTRING(*var*) results in the logical value 1 (=TRUE) if the calculation of *var* does not result in a numeric value.

ISSTRING(B2)	returns 1 (=TRUE) if cell B2 contains a string value, e.g. “SMITH”
ISSTRING(25.6)	returns 0

NOT Not function

The function NOT(*x*) reverses the value of its argument *x*. If *x* is 0 or FALSE then NOT(*x*) returns 1 (= TRUE). If *x* is 1 or TRUE then NOT(*x*) returns 0 (= FALSE).

ODD Odd number

ODD(x) returns 1 (=TRUE) when *x* is an odd number, else this function returns 0 (=FALSE).

OR Or function

OR(*condition1,condition2,...*) returns 1 (=TRUE) of at least one of its arguments is TRUE (or 1). If one of the conditions cannot be evaluated (for example because of missing data), the function returns an error.

Miscellaneous functions

CHIGH Highest column number

The CHIGH() function takes no argument and returns the highest column number used in the spreadsheet.

COLUMN Spreadsheet column identifier

The identifier COLUMN refers to the column number (corresponding with variables) in the spreadsheet.

FALSE Logical constant FALSE

The logical constant FALSE corresponds to the numeric value 0.

RHIGH Highest row number

The RHIGH() function takes no argument and returns the highest row number used in the spreadsheet.

ROW Spreadsheet row identifier

The identifier ROW refers to the row number (indicating cases) in the spreadsheet (see e.g. *Fill column* dialog box, p. 43).

TRUE Logical constant TRUE

The logical constants TRUE corresponds to the numeric value 1.

APPENDIX A. CONTROL KEYS

Cascade windows	Shift+F5
Copy block (Text & Spreadsheet window)	Ctrl+C
Cut block (Text & Spreadsheet window)	Ctrl+X
Edit cell	F2
Exit MedCalc program	Alt+F4
Export file (metafile, text file, ...)	F10
Fill down	Ctrl+D
Fill right	Ctrl+R
Fill series	Ctrl+L
Find	Ctrl+F
Find and replace	Ctrl+H
Go to cell	Ctrl+G
Help	F1
Move cursor left (Text window)	←
Move cursor right (Text window)	→
Move cursor to previous line (Text window)	↑
Move cursor to previous word (Text window)	Ctrl+←
Move cursor to next line (Text window)	↓
Move cursor to next tab position (Text window)	Tab
Move cursor to next word (Text window)	Ctrl+→
On line help	F1
Open spreadsheet data file	Ctrl+O
Print window, graph, data, results	F9
Paste block (Text & Spreadsheet window)	Ctrl+V
Repeat statistics	F7
Save data	Ctrl+S or F12
Save graph as a picture file (export)	F10
Select all	Ctrl+A
Select the contents for a cell from a list of entries already in the column	Alt + ↓
Tile windows	Shift+F4

APPENDIX B. FUNCTIONS OVERVIEW

ARITHMETIC OPERATORS

^	Exponentiation
-	Negation
*	Multiplication
/	Division
+	Addition
-	Subtraction

RELATIONAL OPERATORS

=	Equality
<>	Inequality
<	Less than
>	Greater than
<=	Less than or equal to
>=	Greater than or equal to

MATHEMATICAL FUNCTIONS

ABS(*x*)	Absolute value of *x*
ACOS(*x*)	Arc cosine of *x*
ASIN(*x*)	Arc sine of *x*
ATAN(*x*)	Arc tangent of *x*
COS(*x*)	Cosine of *x*
EXP(*x*)	Natural exponential of *x*
FACT(*x*)	Factorial of *x*
INT(*x*)	Rounds *x* down to the closest integer
LN(*x*)	Natural logarithm of *x*
LOG(*x*)	Logarithm of *x* (base 10)
MOD(*x*,*d*)	Returns the remainder of the division of *x* by *d*
PI()	Returns the value 3.14159265358979
POWER(*x*,*p*)	*x* raised to the power *p* (equals *x*^*p*)
RAND(*x*)	Returns a random number ≥ 1 and ≤ *x* when x >1, or a random number ≥ 0 and < 1 when *x* ≤1
RANDNORM(*m*,*s*)	Returns a random number from a Normal distribution with mean *m* and standard deviation *s*
ROUND(*x*,*n*)	Rounds *x* to the number of digits *n*
SIGN(*x*)	Sign value
SIN(*x*)	Returns sine of *x*
SQRT(*x*)	Returns square root of *x*
TAN(*x*)	Returns tangent of *x*
TRUNC(*x*)	Removes the fractional part of *x*

STATISTICAL FUNCTIONS

AVEDEV(*range*)	Average of absolute deviations of the data in *range*
AVERAGE(*range*)	Arithmetic mean of data in *range*
CHIDIST(*chisquared*,*df*)	One-tailed probability of the chi-squared distribution
CHIINV(*p*,*df*)	Returns the inverse of the chi-squared distribution
COEFVAR(*range*)	The coefficient of variation of the data in *range*
COUNT(*range*)	Count cells containing numbers
COUNTNEG(*range*)	Counts number of negative values
COUNTPOS(*range*)	Counts number of positive numbers
COUNTS(*range*)	Counts number of non-empty cells
FDIST(*F*,*v1*,*v2*)	One-tailed probability of the F distribution
FINV(*p*,*v1*,*v2*)	Returns the inverse of the F probability distribution
GEOMEAN(*range*)	The geometric mean of the data in *range*
MAX(*range*)	Returns highest value in *range*
MIN(*range*)	Returns the lowest value in *range*
NORMSDIST(*z*)	One-tailed probability of the standardized Normal distribution
NORMSINV(*p*)	Returns a Standardized Normal deviate
SEM(*range*)	Standard error of the mean

STDEV(*range*)	Standard deviation
SUM(*range*)	Computes the sum of data in *range*
SUMNEG(*range*)	Sum of negative values
SUMPOS(*range*)	Sum of positive values
TDIST(*t*,*df*)	Two-tailed probability of the Student t distribution
TINV(*p*,*df*)	Returns the inverse of the Student t distribution
VAR(*range*)	Computes the variance

STRING FUNCTIONS

CELL(*column*,*row*)	Returns the contents of the cell with coordinates *column* and *row* as text
CHAR(*x*)	Returns character with code *x*
CODE(*str*)	Returns code of first character of the string *str*
CONCAT(*str1*,*str2*,...)	Concatenate strings
LEFT(*str*,*n*)	Returns first *n* characters of *str*
LEN(*str*)	Returns length of string *str*
LOWER(*str*)	Converts to lowercase
MID(*str*,*pos*,*n*)	returns *n* characters from *str* starting at position *pos*
REPT(*str*,*n*)	Repeat string *str n* times
RIGHT(*str*,*n*)	Returns last *n* characters of *str*
STR(*x*,*n*)	Number to string conversion
UPPER(*str*)	Converts to uppercase
VALUE(*str*)	String to number conversion

DATE FUNCTIONS

DATE(*dnr*)	Serial date number to string
DATEFRAC(*date*)	Returns a fractional year-number
DATEVALUE(*str*)	String to serial date number
DAY(*date*)	Day of *date*
DAYNAME(*date*)	Name of the day of *date*
MONTH(*date*)	Month of *date*
MONTHNAME(*date*)	Name of the month of *date*
TIME(*tnr*)	Converts the serial time number *tnr* to a time string (hh:mm)
TIMEFRAC(*time*)	Returns a fractional time-number
TIMEVALUE(*str*)	Converts a time string (in "hh:mm" format) to a serial time number
TODAY()	Today's date
WEEKNUM(*date*)	Week number of *date*
WEEKDAY(*date*)	Weekday number of *date*
YEAR(*date*)	Year of *date*

LOGICAL FUNCTIONS

AND(*cond1*,*cond2*,...)	Returns 1 of all arguments are TRUE (or 1)
CATEGORISE(..)	Categorise function
IF(*condition*,*x*,*y*)	If-then-else function
ISNUMBER(*var*)	Cell-is-number function
ISSTRING(*var*)	Cell-is-string function
NOT(*x*)	Reverses the logical value of *x*
ODD(*x*)	Returns 1 (=TRUE) when *x* is an odd number
OR(*cond1*,*cond2*,...)	Returns 1 if at least one of its arguments is TRUE (or 1)

MISCELLANEOUS

CHIGH()	Highest column number used in the spreadsheet
COLUMN	Refers to the column number
FALSE	Equals 0
RHIGH()	Highest row number used in the spreadsheet
ROW	Refers to the the row number
TRUE	Equals 1

APPENDIX C. NOTATION AND SYMBOLS

$\lvert x \rvert$	absolute value of x
α	level of hypothesis test, usually 0.05
	probability of a Type I error
β	probability of a Type II error
χ^2	value from the *Chi-squared* distribution
CI	confidence interval
CR	coefficient of repeatability
DF	degrees of freedom
F	value from the *F*-distribution
K	kappa
n or N	sample size
P	probability of hypothesis
r	Pearson correlation coefficient
r_s	Spearman rank correlation coefficient
R	coefficient of determination
RSD	relative standard deviation
Σ	Greek letter sigma, denoting *sum of.* Σx is shorthand for $\sum_{i=1}^{n} x_i$
s	standard deviation
s^2	variance
SD	standard deviation
SEM	standard error of the mean
t	value from the *t*-distribution
z	value from the *Normal* distribution

APPENDIX D. STATISTICAL TABLES

Table 1: Values of the Normal distribution

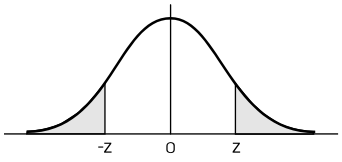

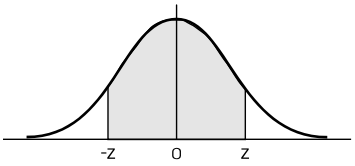

area from -∞ to -z and z to +∞	*area from -z to z*	*z*
0.001	0.999	3.291
0.005	0.995	2.807
0.01	0.99	2.576
0.05	0.95	1.960
0.10	0.90	1.645
0.20	0.80	1.282
0.50	0.50	0.675

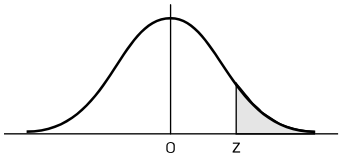

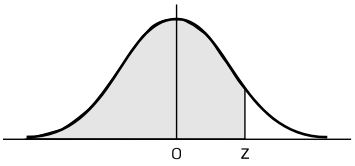

area from z to ∞	*area from -∞ to z*	*z*
0.001	0.999	3.090
0.005	0.995	2.576
0.01	0.99	2.326
0.05	0.95	1.645
0.10	0.90	1.282
0.20	0.80	0.842
0.50	0.50	0.000

Table 2: Values of the t-distribution (two-tailed)

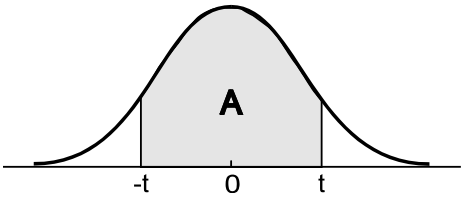

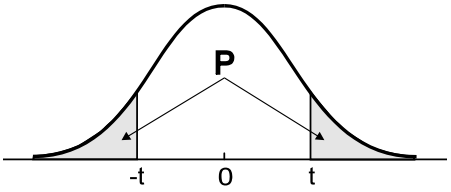

DF	A P	0.80 0.20	0.90 0.10	0.95 0.05	0.98 0.02	0.99 0.01	0.995 0.005	0.998 0.002	0.999 0.001
1		3.078	6.314	12.706	31.820	63.657	127.321	318.309	636.619
2		1.886	2.920	4.303	6.965	9.925	14.089	22.327	31.599
3		1.638	2.353	3.182	4.541	5.841	7.453	10.215	12.924
4		1.533	2.132	2.776	3.747	4.604	5.598	7.173	8.610
5		1.476	2.015	2.571	3.365	4.032	4.773	5.893	6.869
6		1.440	1.943	2.447	3.143	3.707	4.317	5.208	5.959
7		1.415	1.895	2.365	2.998	3.499	4.029	4.785	5.408
8		1.397	1.860	2.306	2.897	3.355	3.833	4.501	5.041
9		1.383	1.833	2.262	2.821	3.250	3.690	4.297	4.781
10		1.372	1.812	2.228	2.764	3.169	3.581	4.144	4.587
11		1.363	1.796	2.201	2.718	3.106	3.497	4.025	4.437
12		1.356	1.782	2.179	2.681	3.055	3.428	3.930	4.318
13		1.350	1.771	2.160	2.650	3.012	3.372	3.852	4.221
14		1.345	1.761	2.145	2.625	2.977	3.326	3.787	4.140
15		1.341	1.753	2.131	2.602	2.947	3.286	3.733	4.073
16		1.337	1.746	2.120	2.584	2.921	3.252	3.686	4.015
17		1.333	1.740	2.110	2.567	2.898	3.222	3.646	3.965
18		1.330	1.734	2.101	2.552	2.878	3.197	3.610	3.922
19		1.328	1.729	2.093	2.539	2.861	3.174	3.579	3.883
20		1.325	1.725	2.086	2.528	2.845	3.153	3.552	3.850
21		1.323	1.721	2.080	2.518	2.831	3.135	3.527	3.819
22		1.321	1.717	2.074	2.508	2.819	3.119	3.505	3.792
23		1.319	1.714	2.069	2.500	2.807	3.104	3.485	3.768
24		1.318	1.711	2.064	2.492	2.797	3.090	3.467	3.745
25		1.316	1.708	2.060	2.485	2.787	3.078	3.450	3.725
26		1.315	1.706	2.056	2.479	2.779	3.067	3.435	3.707
27		1.314	1.703	2.052	2.473	2.771	3.057	3.421	3.690
28		1.313	1.701	2.048	2.467	2.763	3.047	3.408	3.674
29		1.311	1.699	2.045	2.462	2.756	3.038	3.396	3.659
30		1.310	1.697	2.042	2.457	2.750	3.030	3.385	3.646
31		1.309	1.695	2.040	2.453	2.744	3.022	3.375	3.633
32		1.309	1.694	2.037	2.449	2.738	3.015	3.365	3.622
33		1.308	1.692	2.035	2.445	2.733	3.008	3.356	3.611
34		1.307	1.691	2.032	2.441	2.728	3.002	3.348	3.601
35		1.306	1.690	2.030	2.438	2.724	2.996	3.340	3.591
36		1.306	1.688	2.028	2.434	2.719	2.991	3.333	3.582
37		1.305	1.687	2.026	2.431	2.715	2.985	3.326	3.574
38		1.304	1.686	2.024	2.429	2.712	2.980	3.319	3.566
39		1.304	1.685	2.023	2.426	2.708	2.976	3.313	3.558
40		1.303	1.684	2.021	2.423	2.704	2.971	3.307	3.551
42		1.302	1.682	2.018	2.418	2.698	2.963	3.296	3.538
44		1.301	1.680	2.015	2.414	2.692	2.956	3.286	3.526
46		1.300	1.679	2.013	2.410	2.687	2.949	3.277	3.515
48		1.299	1.677	2.011	2.407	2.682	2.943	3.269	3.505
50		1.299	1.676	2.009	2.403	2.678	2.937	3.261	3.496

60	1.296	1.671	2.000	2.390	2.660	2.915	3.232	3.460
70	1.294	1.667	1.994	2.381	2.648	2.899	3.211	3.435
80	1.292	1.664	1.990	2.374	2.639	2.887	3.195	3.416
90	1.291	1.662	1.987	2.369	2.632	2.878	3.183	3.402
100	1.290	1.660	1.984	2.364	2.626	2.871	3.174	3.391
120	1.289	1.658	1.980	2.358	2.617	2.860	3.160	3.373
150	1.287	1.655	1.976	2.351	2.609	2.849	3.145	3.357
200	1.286	1.652	1.972	2.345	2.601	2.839	3.131	3.340
300	1.284	1.650	1.968	2.339	2.592	2.828	3.118	3.323
500	1.283	1.648	1.965	2.334	2.586	2.820	3.107	3.310
∞	1.282	1.645	1.960	2.326	2.576	2.807	3.090	3.291

Example

The mean of a sample is 128.5, SEM 6.2, sample size 32. What is the 99% confidence interval of the mean?

Degrees of freedom (DF) is n-1 = 31, t-value in column for area 0.99 is 2.744. The 99% CI is:

mean - t SEM	to	*mean + t SEM*
128.5 - 2.744 x 6.2	to	128.5 + 2.744 x 6.2
111.5	to	145.5

Table 3: Logit transformation

$$\text{logit}(p) = \ln\left[\frac{p}{1-p}\right] \quad \text{and} \quad p = \frac{1}{1+e^{-\text{logit}(p)}}$$

p	logit(p)	p	logit(p)	p	logit(p)	p	logit(p)
0.01	-4.5951	**0.26**	-1.0460	**0.51**	0.0400	**0.76**	1.1527
0.02	-3.8918	**0.27**	-0.9946	**0.52**	0.0800	**0.77**	1.2083
0.03	-3.4761	**0.28**	-0.9445	**0.53**	0.1201	**0.78**	1.2657
0.04	-3.1781	**0.29**	-0.8954	**0.54**	0.1603	**0.79**	1.3249
0.05	-2.9444	**0.30**	-0.8473	**0.55**	0.2007	**0.80**	1.3863
0.06	-2.7515	**0.31**	-0.8001	**0.56**	0.2412	**0.81**	1.4500
0.07	-2.5867	**0.32**	-0.7538	**0.57**	0.2819	**0.82**	1.5163
0.08	-2.4423	**0.33**	-0.7082	**0.58**	0.3228	**0.83**	1.5856
0.09	-2.3136	**0.34**	-0.6633	**0.59**	0.3640	**0.84**	1.6582
0.10	-2.1972	**0.35**	-0.6190	**0.60**	0.4055	**0.85**	1.7346
0.11	-2.0907	**0.36**	-0.5754	**0.61**	0.4473	**0.86**	1.8153
0.12	-1.9924	**0.37**	-0.5322	**0.62**	0.4895	**0.87**	1.9010
0.13	-1.9010	**0.38**	-0.4895	**0.63**	0.5322	**0.88**	1.9924
0.14	-1.8153	**0.39**	-0.4473	**0.64**	0.5754	**0.89**	2.0907
0.15	-1.7346	**0.40**	-0.4055	**0.65**	0.6190	**0.90**	2.1972
0.16	-1.6582	**0.41**	-0.3640	**0.66**	0.6633	**0.91**	2.3136
0.17	-1.5856	**0.42**	-0.3228	**0.67**	0.7082	**0.92**	2.4423
0.18	-1.5163	**0.43**	-0.2819	**0.68**	0.7538	**0.93**	2.5867
0.19	-1.4500	**0.44**	-0.2412	**0.69**	0.8001	**0.94**	2.7515
0.20	-1.3863	**0.45**	-0.2007	**0.70**	0.8473	**0.95**	2.9444
0.21	-1.3249	**0.46**	-0.1603	**0.71**	0.8954	**0.96**	3.1781
0.22	-1.2657	**0.47**	-0.1201	**0.72**	0.9445	**0.97**	3.4761
0.23	-1.2083	**0.48**	-0.0800	**0.73**	0.9946	**0.98**	3.8918
0.24	-1.1527	**0.49**	-0.0400	**0.74**	1.0460	**0.99**	4.5951
0.25	-1.0986	**0.50**	0.0000	**0.75**	1.0986		

APPENDIX E. ANSI CHARACTER SET

CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR	CODE	CHAR
32		64	@	96	`	128	€	160		192	À	224	à
33	!	65	A	97	a	129	•	161	¡	193	Á	225	á
34		66	B	98	b	130	‚	162	¢	194	Â	226	â
35	#	67	C	99	c	131	*ƒ*	163	£	195	Ã	227	ã
36	$	68	D	100	d	132	„	164	¤	196	Ä	228	ä
37	%	69	E	101	e	133	…	165	¥	197	Å	229	å
38	&	70	F	102	f	134	†	166	¦	198	Æ	230	æ
39	'	71	G	103	g	135	‡	167	§	199	Ç	231	ç
40	(	72	H	104	h	136	ˆ	168	¨	200	È	232	è
41	)	73	I	105	i	137	‰	169	©	201	É	233	é
42	*	74	J	106	j	138	Š	170	ª	202	Ê	234	ê
43	+	75	K	107	k	139	‹	171	«	203	Ë	235	ë
44	,	76	L	108	l	140	Œ	172	¬	204	Ì	236	ì
45	-	77	M	109	m	141	•	173	-	205	Í	237	í
46	.	78	N	110	n	142	Ž	174	®	206	Î	238	î
47	/	79	O	111	o	143	•	175	¯	207	Ï	239	ï
48	0	80	P	112	p	144	•	176	°	208	Ð	240	ð
49	1	81	Q	113	q	145	‘	177	±	209	Ñ	241	ñ
50	2	82	R	114	r	146	’	178	²	210	Ò	242	ò
51	3	83	S	115	s	147	“	179	³	211	Ó	243	ó
52	4	84	T	116	t	148	”	180	´	212	Ô	244	ô
53	5	85	U	117	u	149	•	181	µ	213	Õ	245	õ
54	6	86	V	118	v	150	–	182	¶	214	Ö	246	ö
55	7	87	W	119	w	151	—	183	·	215	×	247	÷
56	8	88	X	120	x	152	˜	184	¸	216	Ø	248	ø
57	9	89	Y	121	y	153	™	185	¹	217	Ù	249	ù
58	:	90	Z	122	z	154	š	186	º	218	Ú	250	ú
59	;	91	[	123	{	155	›	187	»	219	Û	251	û
60	<	92	\	124	|	156	œ	188	¼	220	Ü	252	ü
61	=	93	]	125	}	157	•	189	½	221	Ý	253	ý
62	>	94	^	126	~	158	ž	190	¾	222	Þ	254	þ
63	?	95	_	127	□	159	Ÿ	191	¿	223	ß	255	ÿ

LITERATURE

Altman DG (1980) Statistics and ethics in medical research. VI - Presentation of results. British Medical Journal 281:1542-1544.

Altman DG (1991) Practical statistics for medical research. London: Chapman and Hall.

Altman DG, Gardner MJ (1988) Calculating confidence intervals for regression and correlation. British Medical Journal 296:1238-1242.

Altman DG, Gore SM, Gardner MJ, Pocock SJ (1983) Statistical guidelines for contributors to medical journals. British Medical Journal 286:1489-1493.

Altman DG, Machin D, Bryant TN, Gardner MJ (Eds) (2000) Statistics with confidence, 2nd ed. BMJ Books.

Armitage P (1955) Tests for linear trends in proportions and frequencies. Biometrics 11:375–386.

Armitage P, Berry G, Matthews JNS (2002) Statistical methods in medical research. 4th ed. Blackwell Science.

Bland M (2000) An introduction to medical statistics, 3rd ed. Oxford: Oxford University Press.

Bland JM, Altman DG (1986) Statistical method for assessing agreement between two methods of clinical measurement. The Lancet i:307-310.

Bland JM, Altman DG (1995) Comparing methods of measurement: why plotting difference against standard method is misleading. The Lancet 346:1085-1087.

Bland JM, Altman DG (1997) Statistics notes: Cronbach's alpha. British Medical Journal 314:572.

Bland JM, Altman DG (1999) Measuring agreement in method comparison studies. Statistical Methods in Medical Research 8:135-160.

Bland JM, Altman DG (2007) Agreement between methods of measurement with multiple observations per individual. Journal of Biopharmaceutical Statistics. 17:571-582.

Bulpitt CJ (1987) Confidence intervals. The Lancet i: 494-497.

Campbell MJ, Gardner MJ (1988) Calculating confidence intervals for some non-parametric analyses. British Medical Journal 296:1454-1456.

CLSI (2008) Defining, establishing, and verifying reference intervals in the clinical laboratory: approved guideline - third edition. CLSI Document C28-A3. Wayne, PA: Clinical and Laboratory Standards Institute.

Cohen J (1960) A coefficient of agreement for nominal scales. Educational and Psychological Measurement 20:37-46.

Conover WJ (1999) Practical nonparametric statistics, 3rd edition. New York: John Wiley & Sons.

Cornbleet PJ, Gochman N (1979) Incorrect least-squares regression coefficients in method-comparison analysis. Clinical Chemistry 25:432-438.

Cronbach LJ (1951) Coefficient alpha and the internal structure of tests. Psychometrika 16:297-334.

Dewitte K, Fierens C, Stöckl D, LM Thienpont (2002) Application of the Bland-Altman plot for interpretation of method-comparison studies: a critical investigation of its practice. Clinical Chemistry 48:799-801.

DerSimonian R, Laird N (1986) Meta-analysis in clinical trials. Controlled Clinical Trials 7:177-188.

Feldt LS (1965) The approximate sampling distribution of Kuder-Richardson reliability coefficient twenty. Psychometrika 30:357-371.

Fleiss JL (1981) Statistical methods for rates and proportions, 2nd ed. New York: John Wiley & Sons.

Fleiss JL, Levin B, Paik MC (2003) Statistical methods for rates and proportions, 3rd ed. Hoboken: John Wiley & Sons.

Gardner IA, Greiner M (2006) Receiver-operating characteristic curves and likelihood ratios: improvements over traditional methods for the evaluation and application of veterinary clinical pathology tests. Veterinary Clinical Pathology 35:8-17.

Gardner MJ, Altman DG (1986) Confidence intervals rather than P values: estimation rather than hypothesis testing. British Medical Journal 292:746-750.

Girden ER (1992) ANOVA: repeated measures. Sage University Papers Series on Quantitative Applications in the Social Sciences, 84. Thousand Oaks, CA: Sage.

Glantz SA, Slinker BK (2001) Primer of applied regression & analysis of variance. 2nd ed. McGraw-Hill.

Greenhouse SW, Geisser S (1959) On methods in the analysis of profile data. Psychometrika 24:95-112.

Griner PF, Mayewski RJ, Mushlin AI, Greenland P (1981) Selection and interpretation of diagnostic tests and procedures. Annals of Internal Medicine 94:555-600.

Grubbs FE (1969) Procedures for detecting outlying observations in samples. Technometrics 11:1-21.

Hanley JA, McNeil BJ (1982) The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology 143:29-36.

Hanley JA, McNeil BJ (1983) A method of comparing the areas under receiver operating characteristic curves derived from the same cases. Radiology 148:839-843.

Hilgers RA (1991) Distribution-free confidence bounds for ROC curves. Methods of Information in Medicine 30:96-101.

Huitema BE (1980) The analysis of covariance and alternatives. Wiley-Interscience.

Husted JA, Cook RJ, Farewell VT, Gladman DD (2000) Methods for assessing responsiveness: a critical review and recommendations. Journal of Clinical Epidemiology 53:459-168.

Huynh H, Feldt LS (1976) Estimation of the Box correction for degrees of freedom from sample data in randomised block and split-plot designs. Journal of Educational Statistics 1:69-82.

Jones R, Payne B (1997) Clinical investigation and statistics in laboratory medicine. London: ACB Venture Publications.

Kirkwood BR, Sterne JAC (2003) Essential medical statistics, 2nd ed. Oxford: Blackwell Scientific.

Krouwer JS (2008) Why Bland-Altman plots should use *X*, not (*Y*+*X*)/2 when *X* is a reference method. Statistics in Medicine 27:778-780.

Krouwer JS, Monti KL (1995) A simple, graphical method to evaluate laboratory assays. Eur J Clin Chem Clin Biochem 33:525-527.

Lecoutre B (1991) A correction for the e approximate test in repeated measures designs with two or more independent groups. Journal of Educational Statistics 16:371-372.

Lentner C (ed) (1982) Geigy Scientific Tables, 8th edition, Volume 2. Basle: Ciba-Geigy Limited.

Lin L.I-K (1989) A concordance correlation coefficient to evaluate reproducibility. Biometrics 45:255–268.

Lin L.I-K (2000) A note on the concordance correlation coefficient. Biometrics 56:324–325.

McGill R, Tukey JW, Larsen WA (1978) Variations of box plots. The American Statistician 32:12-16.

McGraw KO, Wong SP (1996) Forming inferences about some intraclass correlation coefficients. Psychological Methods 1:30-46. (Correction: 1:390).

Mathews JNS, Altman DG, Campbell MJ, Royston P (1990) Analysis of serial measurements in medical research. British Medical Journal 300:230-235.

Metz CE (1978) Basic principles of ROC analysis. Seminars in Nuclear Medicine 8:283-298.

Moses LE (1987) Graphical methods in statistical analysis. Annual Review of Public Health 8:309-353.

NCCLS (2000) How to define and determine reference intervals in the clinical laboratory: approved guideline - second edition. NCCLS document C28-A2. Wayne, PA: NCCLS.

Neter J, Wasserman W, Whitmore GA (1988) Applied statistics. 3rd ed. Boston: Allyn and Bacon, Inc.

Neter J, Kutner MH, Nachtsheim CJ, Wasserman W (1996) Applied linear statistical models. 4th ed. McGraw-Hill.

Norman GR, Wyrwich KW, Patrick DL (2007) The mathematical relationship among different forms of responsiveness coefficients. Quality of Life Research 16:815-822.

Pampel FC (2000) Logistic regression: A primer. Sage University Papers Series on Quantitative Applications in the Social Sciences, 07-132. Thousand Oaks, CA: Sage.

Passing H, Bablok W (1983) A new biometrical procedure for testing the equality of measurements from two different analytical methods. Application of linear regression procedures for method comparison studies in Clinical Chemistry, Part I. J. Clin. Chem. Clin. Biochem. 21:709-720.

Petrie A, Bulman JS, Osborn JF (2003) Further statistics in dentistry. Part 8: systematic reviews and meta-analyses. British Dental Journal, 194:73-78.

Pocock SJ (1983) Clinical trials. A practical approach. Chichester: John Wiley & Sons.

Reed AH, Henry RJ, Mason WB (1971) Influence of statistical method used on the resulting estimate of normal range. Clinical Chemistry 17:275-284.

Rosner B (1983) Percentage points for a generalized ESD many-outlier procedure. Technometrics 25:165-172.

Sahai H, Khurshid A (1996) Statistics in epidemiology: methods, techniques, and applications. Boca Raton, FL: CRC Press, Inc.

Sheskin DJ (2004) Handbook of parametric and nonparametric statistical procedures. 3rd ed. Boca Raton: Chapman & Hall /CRC.

Shrout PE, Fleiss JL (1979) Intraclass correlations: uses in assessing rater reliability. Psychological Bulletin 86:420-428.

Snedecor GW, Cochran WG (1989) Statistical methods, 8th edition. Ames, Iowa: Iowa State University Press.

Spiegel MR (1961) Theory and problems of statistics. New York: McGraw-Hill Book Company.

Tukey JW (1977) Exploratory data analysis. Reading, Mass: Addison-Wesley Publishing Company.

Westgard JO, Barry PL, Hunt MR, Groth T (1981) A multi-rule Shewhart chart for Quality Control in Clinical Chemistry. Clinical Chemistry 27:493-501.

Wildt AR, Ahtola OT (1978) Analysis of covariance. Sage Publications.

Youden WJ (1959) Graphical diagnosis of interlaboratory test results. Industrial Quality Control 15:24-28.

Zhou XH, NA Obuchowski, DK McClish (2002) Statistical methods in diagnostic medicine. New York: Wiley.

Zweig MH, Campbell G (1993) Receiver-operating characteristic (ROC) plots: a fundamental evaluation tool in clinical medicine. Clinical Chemistry 39:561-577.

INDEX