

USING MINITAB: A SHORT GUIDE VIA EXAMPLES

The goal of this document is to provide you, the student in Math 112, with a guide to some of the tools of the statistical software package MINITAB as they directly pertain to the analysis of data you will carry out in Math 112, in conjunction with the textbook Introduction to the Practice of Statistics by David Moore and George McCabe. This guide is organized around examples. To help you get started, it includes pictures of some screens you'll see and printouts of the commands. It is neither a complete list of MINITAB capabilities nor a complete guide to all of the uses of MINITAB with this textbook, but is designed to hit the highlights and a few sticking points, so to speak, of the use of MINITAB for problems in the text, based on the Spring 1997 and Fall 1997 courses. For a brief dictionary of MINITAB functions and their EXCEL equivalents, plus more information about EXCEL, see elsewhere in this website. NOTE: This guide does update some of the MINITAB commands given in Introduction to the Practice of Statistics.

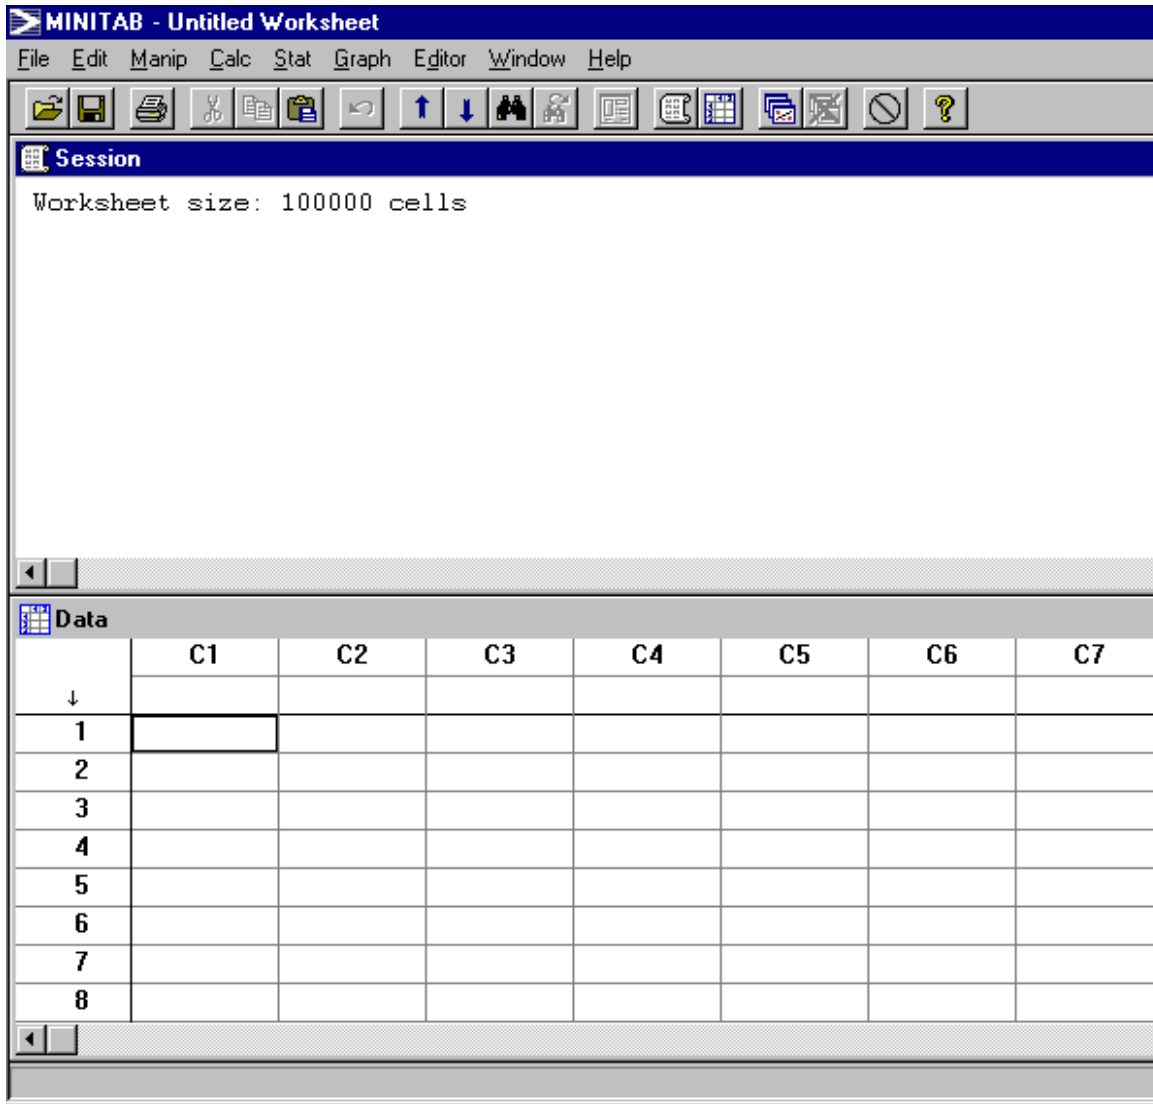
CONTENTS

1. Starting and running MINITAB on grounds
2. Describing distributions (Ch. 1 of text)
3. Scatterplots, Linear Regression, and Correlation (Ch. 2 of text)
4. Inference for Regression (Section 10.1 of text)

1. Starting and running MINITAB on grounds

- OPENING MINITAB

MINITAB should be available in any on-grounds computer lab. Begin from the Windows menu, and look for the folder marked STATISTICS (or perhaps MATHEMATICS, if there is no STATISTICS folder listed). Clicking to open this file, you should see the MINITAB icon (a blue- and white-striped arrow, labelled MINITAB). Clicking on this will open MINITAB, displaying a 'session window' on top and a 'data window' below.



- ENTERING DATA

The cursor can be placed in either window, but to begin to use MINITAB we will, of course, need to enter some data. Upon opening a new file, the data window should be ready to receive columns of data. Simply place the cursor in the first row (box) of column C1, enter the first number of your data set, and hit 'Enter'. MINITAB will then automatically move down to the second row in this column, making data easy to enter. Note that you can label your column by entering a name in the box above the first row but still below the label 'C1'. In order to have a data set to use when following the examples below, enter the following numbers in C1:

(*) 5 6.6 5.2 6.1 7.4 8.7 5.4 6.8 7.1 7.4

The data should appear in the data window as so:

The screenshot shows the MINITAB software interface. At the top is a menu bar with options: File, Edit, Manip, Calc, Stat, Graph, Editor, Window, Help. Below the menu bar is a toolbar with various icons for file operations (like Save, Print, Copy, Paste) and data manipulation (like Undo, Redo, Sort, Filter). The main area is a data grid with columns labeled C1 through C7 and rows numbered 1 through 21. The data in column C1 is as follows:

	C1	C2	C3	C4	C5	C6	C7
↓							
1	5.0						
2	6.6						
3	5.2						
4	6.1						
5	7.4						
6	8.7						
7	5.4						
8	6.8						
9	7.1						
10	7.4						
11							
12							
13							
14							
15							
16							
17							
18							
19							
20							
21							
22							

You can print out the data window at any time by locating your cursor anywhere inside the window, and clicking on the 'Print' icon below the toolbar. The printout will ignore empty columns.

- **PRINTING**

As noted immediately above, the 'Print' icon (a little printer, of course) can be used to print out the data window. Similarly, you can move the cursor to the session window to print the output placed in this window by some of the commands we'll soon explore. You can also print graph windows by clicking on the graph before attempting to print.

- **GETTING HELP FROM MINITAB**

The 'Help' command in MINITAB is very useful. Once you have pulled up a command window, say by following some of the paths below (see 'THE TOOLBAR AND COMMANDS' below), you can click on the HELP button to produce a description of the command and usually an accompanying example; these can be printed out using the 'Print' icon if you like. If you are trying to locate a feature, follow the 'Search' option after clicking on the 'Help' command on the toolbar. You may want to use this to seek out more information about the commands used below, or on topics not covered in this guide. When you begin the course and do not yet have much knowledge of statistics, the information contained in the help sheets

can be a bit overwhelming, but with time and with some basic examples in your hands (as this guide hopes to help you acquire), negotiating MINITAB by using 'Help' becomes fairly easy.

- THE TOOLBAR AND COMMANDS

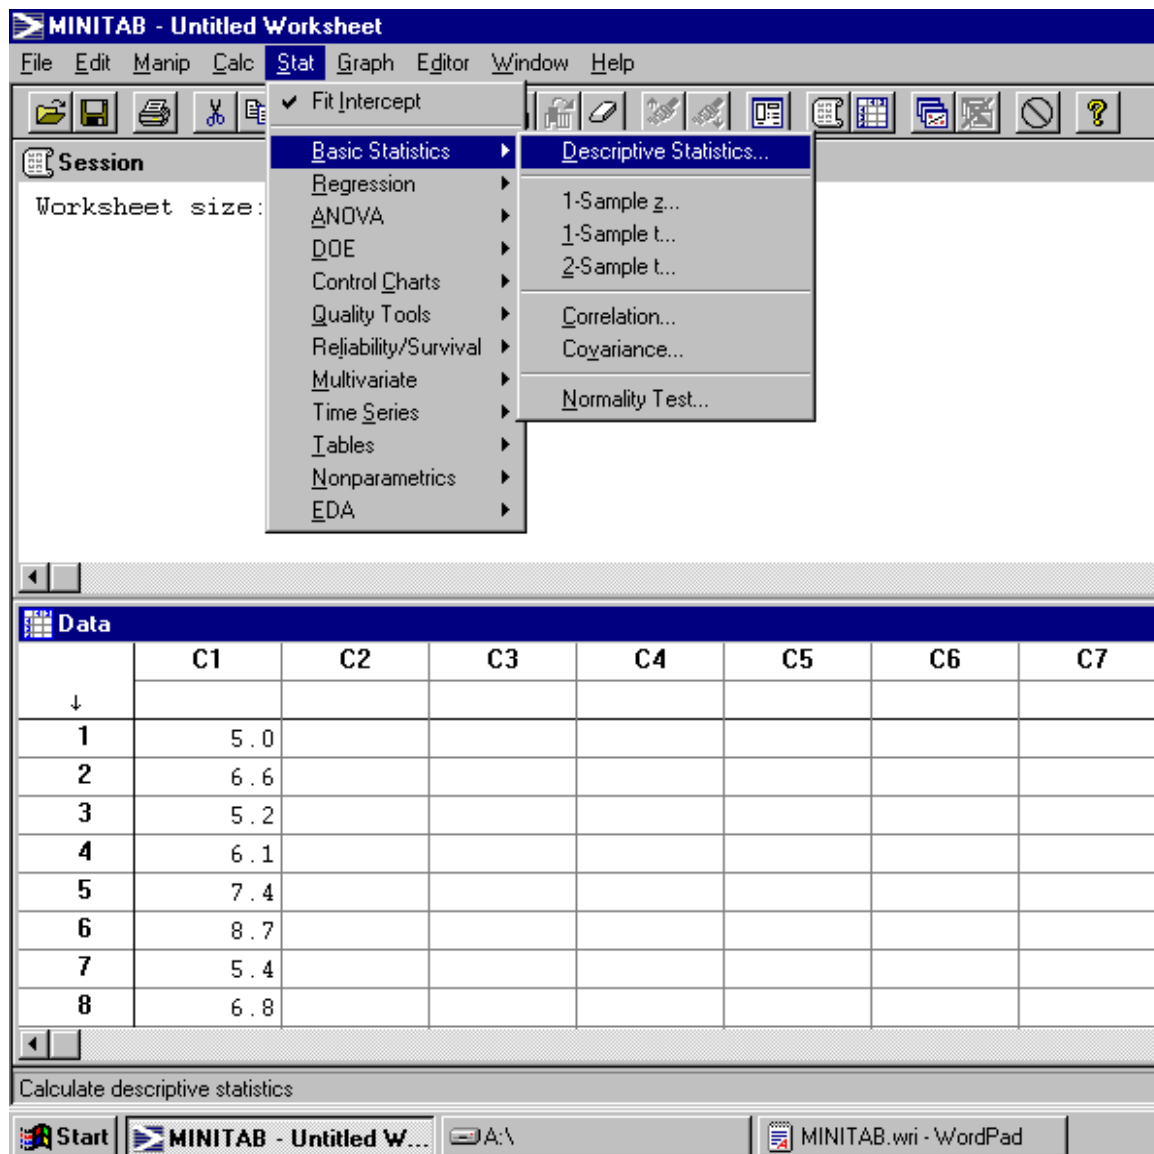
The tool bar on the top includes the headings

File Edit Manip Calc Stat Graph Editor Window Help

and using the mouse to click on each of these produces a range of options. To indicate the path of a command originating from the toolbar, we will use notation such as

Stat > Basic Statistics > Descriptive Statistics.

Check that beginning with 'Stat' you can find the 'Basic Statistics' option, and from that the 'Descriptive Statistics' selection as below:



A picture of the resulting command window is shown in the next section. (For now, hit the 'Cancel' button to close off the window without executing a command, as of yet!)

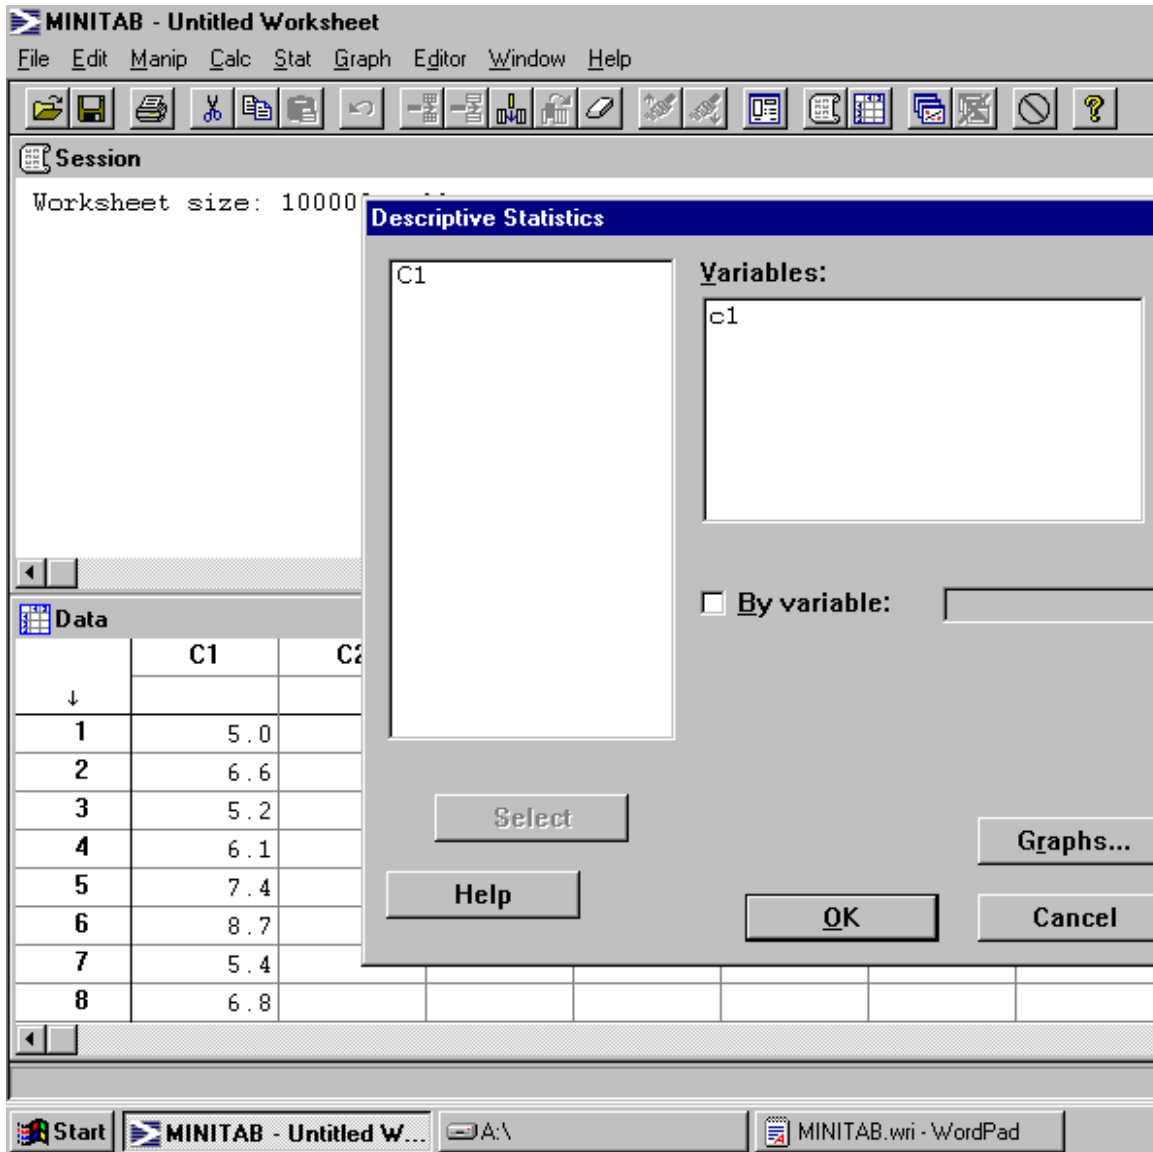
Commands can be entered in the sessions window (or by pulling up a special window for entering commands), and in some places in the text, this is the way they suggest for you to use MINITAB. In this guide we will instead, whenever possible, take advantage of the current format of MINITAB to 'mouse' along.

2. Describing Distributions (Ch. 1 of text)

NOTE: In the examples which follow, we will use the data set (*) from 'Entering Data' in Section 1 of this guide above. This data will be referred to by the label 'C1'. For additional data sets, these procedures can be repeated by substituting their corresponding column labels in place of C1, or, in most cases, the output can be produced for several different columns at the same time. We will detail the procedure for a single column only.

- DESCRIPTIVE STATISTICS (MEAN, MEDIAN, QUARTILES, STANDARD DEVIATION...)

For each column of data, you can find the mean, median, standard deviation, min, max, and first and third quartiles all in one shot by following the command path **Stat > Basic Statistics > Descriptive Statistics** and entering C1 in the 'Variables' box.



Graphs such as histograms and boxplots can be produced by selecting those options, or see those topics separately below. The output (without graphs) appears in the session window as shown below.

Descriptive Statistics

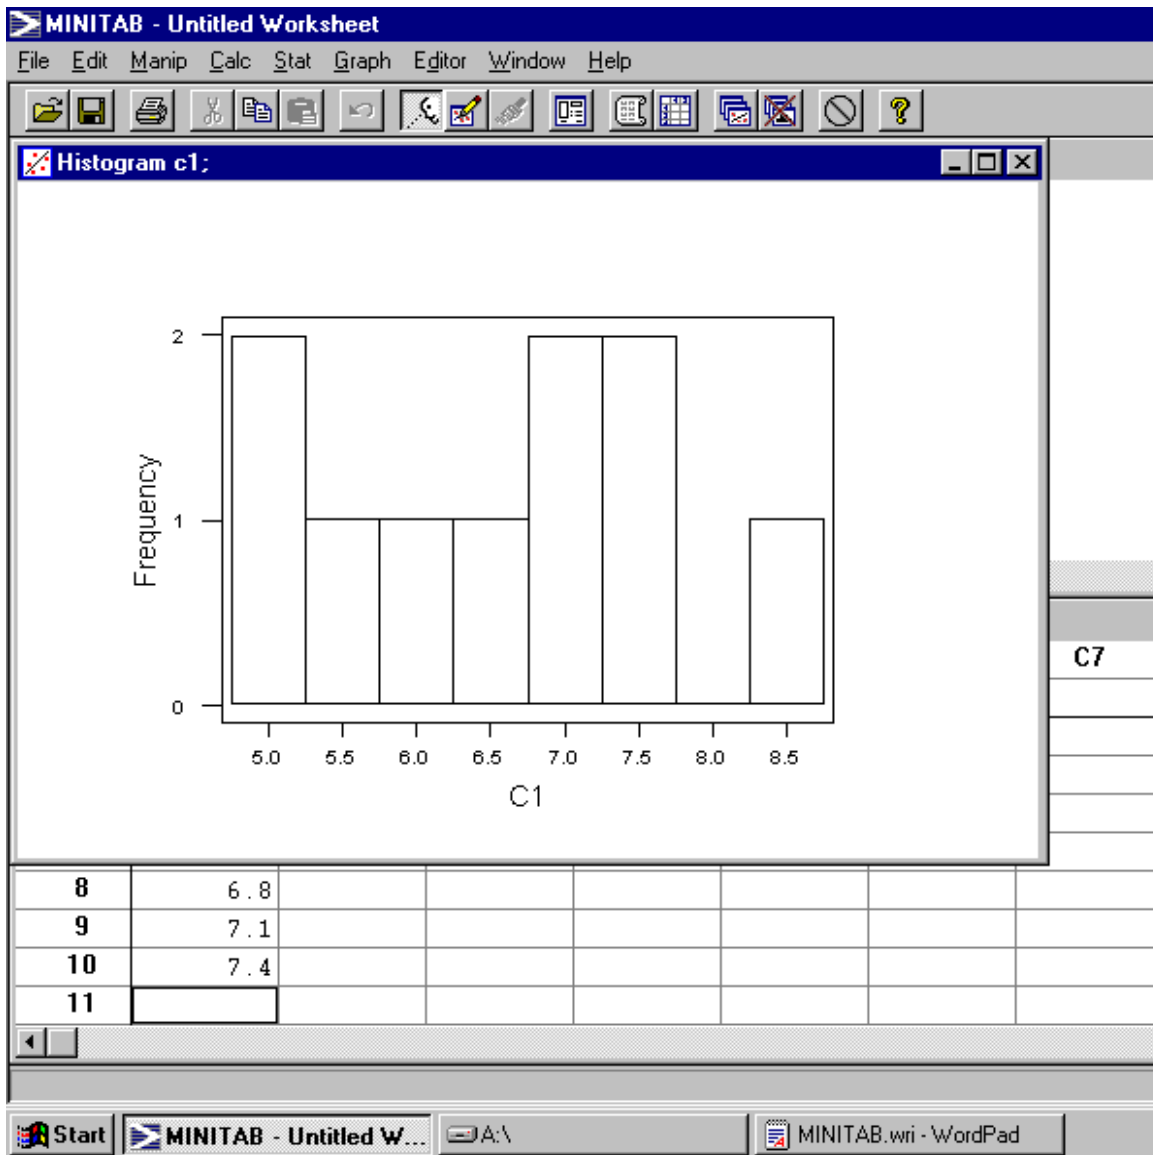
Variable	N	Mean	Median	Tr Mean	StDev	SE Mean
C1	10	6.570	6.700	6.500	1.163	0.368

Variable	Min	Max	Q1	Q3
C1	5.000	8.700	5.350	7.400

Many of these statistics (and some others such as the sum of squares and range) can also be computed separately by following **Calc > Column Statistics** and entering C1 as the 'Input variable'.

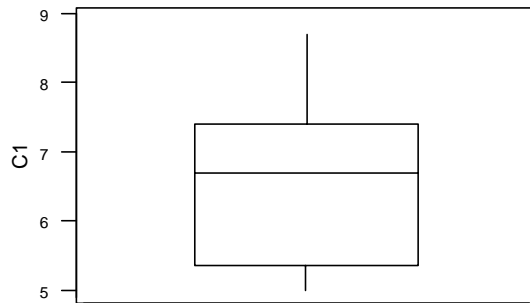
- HISTOGRAM

To produce a histogram of the data in C1, follow **Graph > Histogram**. Enter C1 in the first row under 'Graph Variable' and click 'OK'. The output is a separate graph window as below.



- **BOXPLOT**

Use **Graph > Boxplot** and enter C1 in the first row of the 'Y' column under the 'Graph variables' heading, then click on 'OK'. The output appears in a graph window.



- STEMPLOT (STEM-AND-LEAF PLOT)

Substitute C1 into the 'Variables' window appearing after following the path **Graph > Character Graph > Stem-and-Leaf**. The output appears in the session window.

Character Stem-and-Leaf Display

```
Stem-and-leaf of C1          N = 10
Leaf Unit = 0.10
```

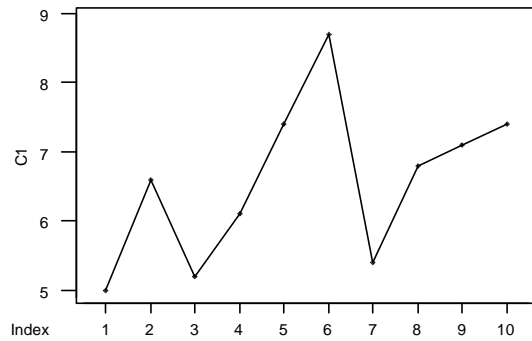
```

3      5 024
3      5
4      6 1
(2)    6 68
4      7 144
1      7
1      8
1      8 7
```

You might note that other graphs, such as histograms and boxplots, can also be created as character graphs, but they just aren't as pretty!

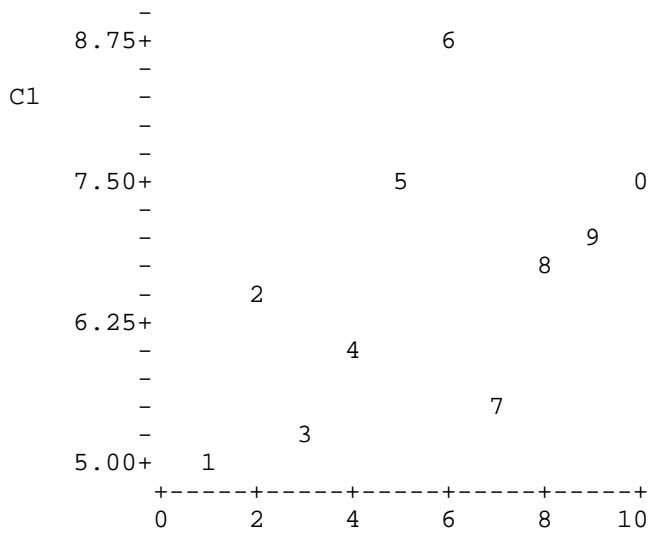
- TIME SERIES PLOT

This plot will, using the preset features, produce a time series plot labelled on the x-axis by the numbers 1, 2, 3,... placed at equally spaced intervals, and with dots connected. Check the help menu for changing this labelling. Take the path **Graph > Time Series Plot** and enter C1 in the first row of the 'Y' column under 'Graph variables', then click 'OK'.



To produce a time series plot which labels the points by their order of appearance and does not connect the dots, use the path **Graph > Character Graphs > Time Series Plot**.

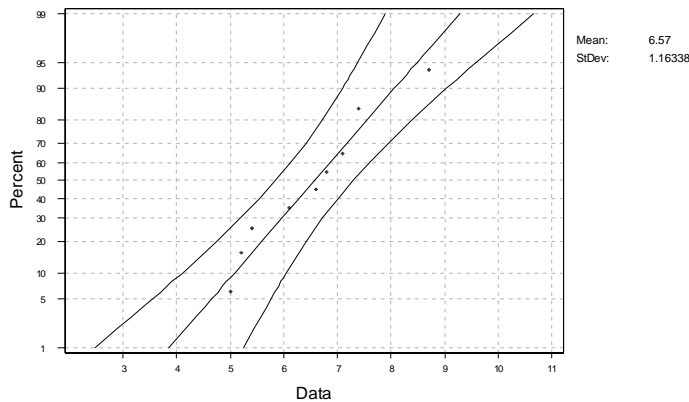
Character Multiple Time Series Plot



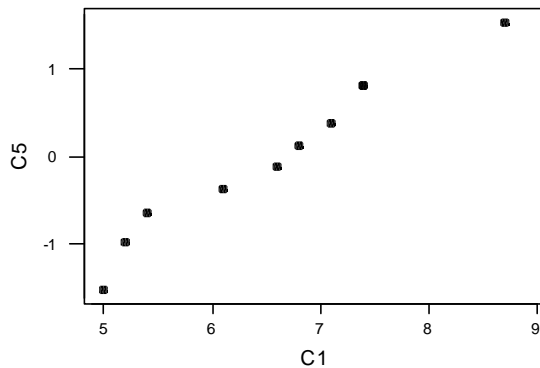
- **NORMAL QUANTILE PLOT**

As your text notes, a normal quantile plot is also called a 'normal probability plot.' To produce a plot which corresponds to the text's definition of a normal quantile plot in MINITAB, you can use the path **Graph > Probability Plot** with C1 as the variable and 'Normal' as the selection under 'Assumed distribution'. Click on 'OK'. Along with some other data, the graph appears as below:

Normal Probability Plot for C1



You'll notice that the output includes some curves which do not appear in the text's illustrations of normal quantile plots. To produce a simpler picture, without these curves, you can follow the book's suggestion and create one yourself. To do this, first enter the command **NSCORES C1 C5** in the prompt in the session window. (If you already have some data in the session window and cannot get a prompt to appear there, you can also enter this command by following the path **Edit > Command Line Editor**.) This produces a new column of data, C5. If the C5 column is already filled before you begin this process, choose another label in its place (say, if C2 or C10 is empty). Now, to get a picture just like the text's, enter the new command **PLOT C5 * C1** as above, or follow **Graph > Plot** and substitute C5 for 'Y' and C1 for 'X' in the first row of 'Graph Variables'.



As a final note, MINITAB includes other commands for producing normal probability plots and variations such as **NORMPLOT** and following the path **Stat > Basic Statistics > Normality Test**. You may wish to check with your instructor to see if some variation other than the two described above is desired.

3. Scatterplots, Linear Regression, and Correlation (Ch. 2 of text)

Note: In the examples which follow, we will use the data from Example 2.11 of the text. In this case, the 'x-variable' data is recorded as 'student' in column C1 of the data sheet, and the 'y-variable' data as 'math' in column C2. You should thus enter the following data into the data session window:

MINITAB - Untitled Worksheet

File Edit Manip Calc Stat Graph Editor Window Help

Session

Worksheet size: 100000 cells

MTB >

	C1	C2	C3	C4	C5	C6	C7
↓	student	math					
1	4595	7364					
2	4827	7547					
3	4427	7099					
4	4258	6894					
5	3995	6572					
6	4330	7156					
7	4265	7232					
8	4351	7450					
9							
10							
11							
12							

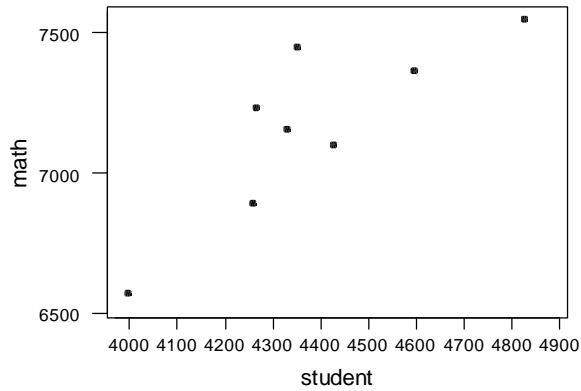
Welcome to Minitab, press F1 for help.

Note that you can clear out any previous entries in these two columns by highlighting the boxes using the mouse, and by following EDIT > CLEAR CELLS.

- SCATTERPLOT

Follow **Graph > Plot** and in the first row under 'Graph variables' enter C2 in the column for 'Y' and C1 in the column for 'X'. Hit 'OK'. As the text suggests, you can also enter the command **PLOT C2 ***

C1 in the session window (or by following the path **Edit > Command Line Editor** if you cannot get a prompt in the session window). The output, which appears in a graph window, is shown below.



- **LINEAR REGRESSION: LEAST-SQUARES REGRESSION LINE AND CORRELATION COEFFICIENT**

There are many features of MINITAB's 'Regression' command which we will want to explore. Let's begin simply by finding the equation for the least-squares regression line of 'Y' (here, 'math') on 'X' (here, 'student'). Instead of following the test's suggestion to enter commands into the session window, we will take the command path **Stat > Regression > Regression** to pull up the following window:

The screenshot shows the Minitab interface with a 'Regression' dialog box open. The dialog box has a list of variables on the left: C1 (student) and C2 (math). The 'Response' field is set to 'c2' and the 'Predictors' field is set to 'c1'. Below the list are 'Select' and 'Help' buttons. In the background, a 'Data' window shows a table with columns C1 (student) and C2 (math) and rows 1 through 12.

	C1	C2
↓	student	math
1	4595	7364
2	4827	7547
3	4427	7099
4	4258	6894
5	3995	6572
6	4330	7156
7	4265	7232
8	4351	7450
9		
10		
11		
12		

Entering C2 in the 'Response' box and C1 in the 'Predictors' box as above, and selecting 'OK' gives the following output in the session window:

Regression Analysis

The regression equation is
 $\text{math} = 2493 + 1.07 \text{ student}$

Predictor	Coef	StDev	T	P
Constant	2493	1267	1.97	0.097
student	1.0663	0.2888	3.69	0.010

S = 188.9 R-Sq = 69.4% R-Sq(adj) = 64.3%

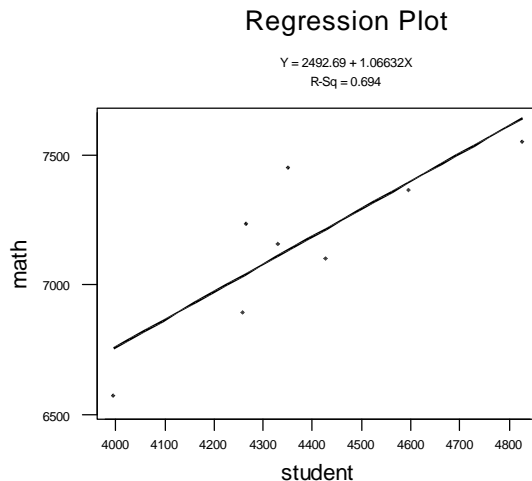
Analysis of Variance

Source	DF	SS	MS	F	P
Regression	1	486552	486552	13.63	0.010
Error	6	214209	35702		
Total	7	700762			

As you can see, the equation for the *least-squares regression line* of 'math' ('Y') on 'student' ('X') is given at the top of the output. In more detail, the slope 1.0663 (called 'b' in Ch. 2 of your text and ' β_1 ' in Ch. 10) appears in the 'Coef' column and 'student' row, while the intercept 2493 ('a' in Ch. 2 and ' β_0 ' in Ch. 10) appears in the same column, but in the 'Constant' row. The columns 'St Dev', 'T', and 'P', as well as the 'Analysis of Variance' material below them will be useful in Ch. 10, but not needed in Ch. 2. The other feature of this printout which you will need now in Ch. 2 is the *correlation coefficient*. In your text it is labelled ' r^2 ' but appears here in the printout with a capital letter as 'R-Sq'; in this case, its value is 69.4.

- FITTED LINE PLOT

To produce a picture of the least-squares regression line fitted to the scatterplot, take the path **Stat > Regression > Fitted Line Plot**. Enter C2 for 'Response (Y)' and C1 for 'Predictor (X)'. Make sure the option under 'Type of Regression Model' is 'Linear', and then click 'OK'. This will produce the plot shown below, along with the 'Regression Analysis' data as above. The equation for the regression line is given at the top of the graph, as well as the correlation coefficient.



- RESIDUAL PLOT

There are, in fact, several kinds of residual plots. Here, we will show you how to find the residual plot which corresponds to the text's use of the term. In this case, we will be plotting the residuals on the 'y-axis' and the explanatory variable values (here, the 'student' data) on the 'x-axis'. There are several ways to

create this plot. Here is the first and the most easy . Continuing from the window marked 'Regression' found in the section 'REGRESSION' above, click on the 'Graphs' button to pull up the window below:

The screenshot shows the Minitab interface with a data table and a dialog box. The data table has the following content:

	C1	C2
↓	student	math
1	4595	7364
2	4827	7547
3	4427	7099
4	4258	6894
5	3995	6572
6	4330	7156
7	4265	7232
8	4351	7450
9		
10		
11		
12		

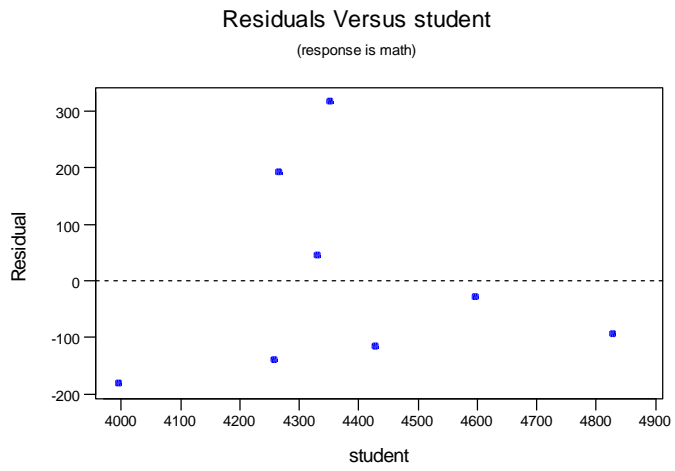
The 'Regression - Graphs' dialog box is open, showing the following settings:

- Variables: C1 student, C2 math
- Residuals for Plots: Regular, Standardized
- Residual Plots:
 - Histogram of residuals
 - Normal plot of residuals
 - Residuals versus fits
 - Residuals versus order
- Residuals versus the variables: c1

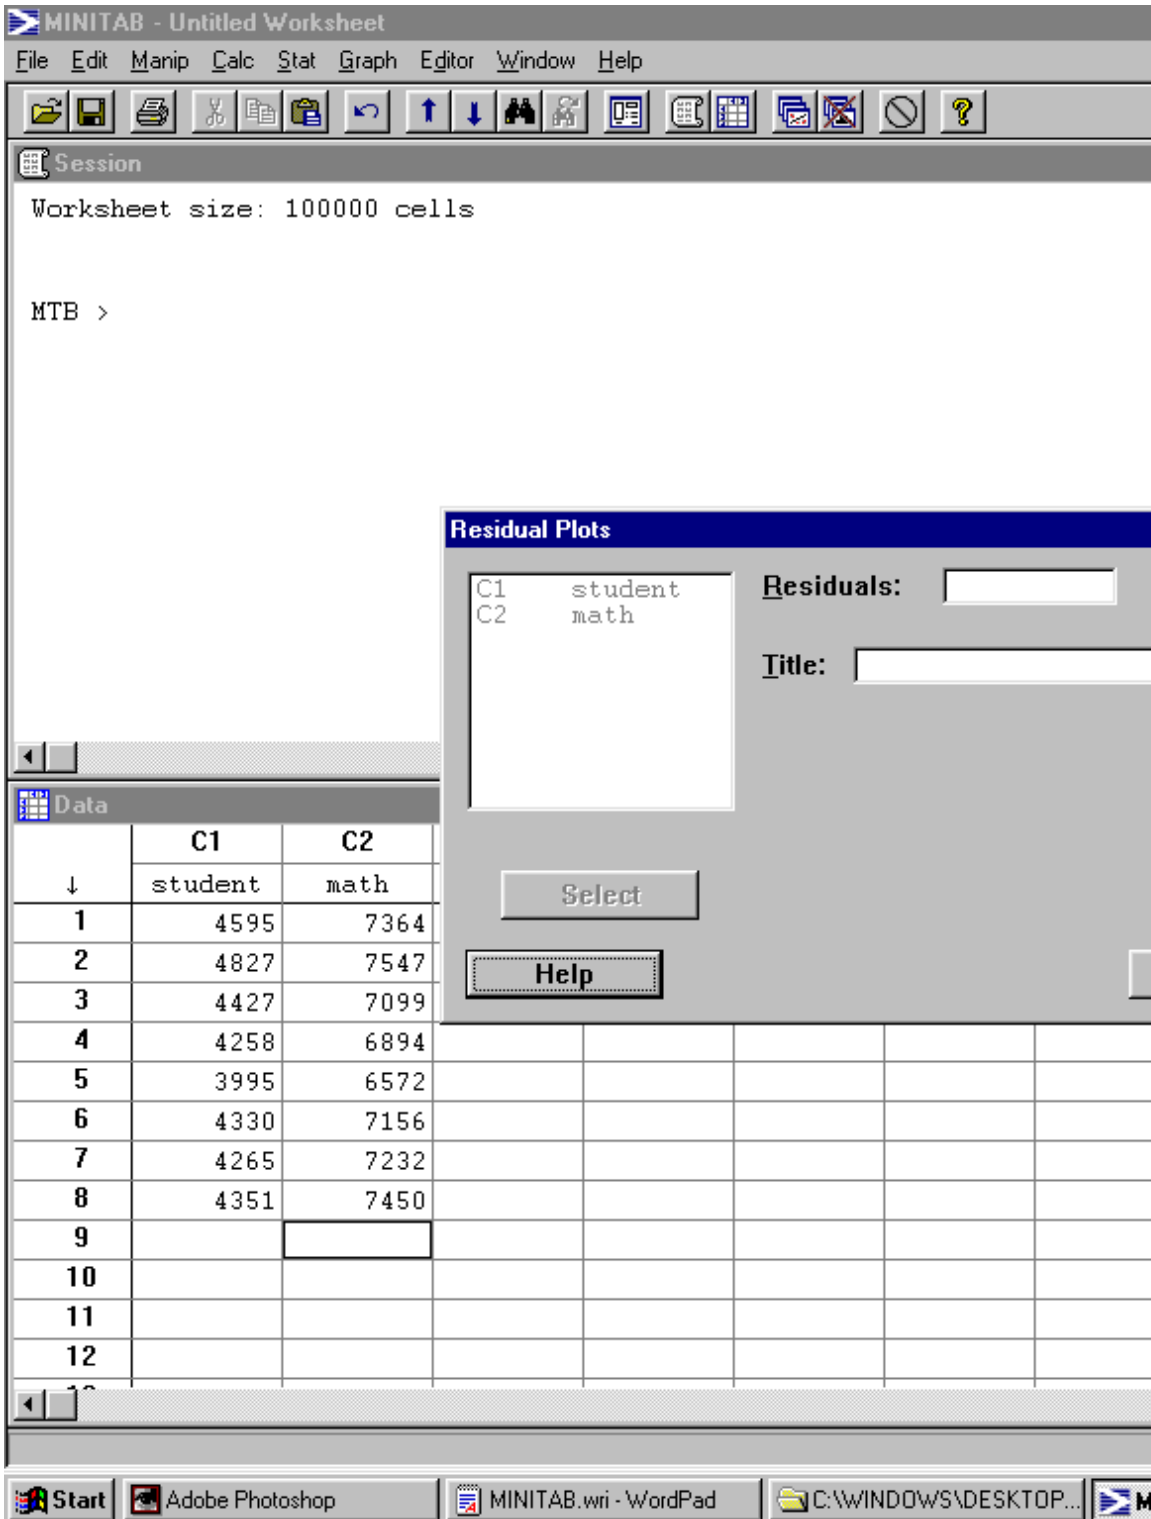
Buttons for 'Select' and 'Help' are visible at the bottom of the dialog box.

Keeping the 'Regular' selection under 'Residuals for Plots' and entering C1 under 'Residuals versus the variables' will produce (after clicking on 'OK') the graph below in a separate graph window, as well as the

'Regression Analysis' data which we examined above in 'REGRESSION'.



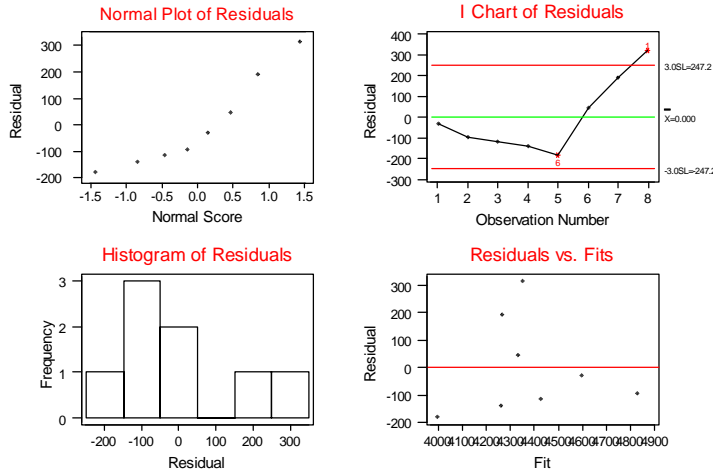
Another way to produce this residuals plot is to follow the path **Stat > Regression > Residual Plots**, which gives the window



Unfortunately, while this command path might seem to be the most natural one, you will not be able to execute this until you produce and store the residuals as a column of data. To do this, see the section on STORAGE below. Having followed the directions for storage, and assuming that the residuals are stored in column C4, then entering C4 in 'Residuals' above and entering the C1 for 'Fits' (NOT C3!) will produce a graph window containing four graphs. The residuals plot as the text defines it will be the one labelled

'Residuals vs. Fits'.

Residual Model Diagnostics



- STORAGE: PRODUCING LISTS OF RESIDUALS AND FITS

For any 'x-value' entered into the equation of the least-squares regression line, the output (labelled y^{\wedge} in your text) is the 'predicted value', also called a 'fit'. For an x-value which actually appears as an x-value used to build the regression line, the difference between the corresponding 'y-value' (the 'observed value') and the output of the equation (the 'predicted value'), is the *residual*: $\text{residual} = y - y^{\wedge}$. To store and view the predicted values (fits) and the residues for the x-values used to build the regression line, return to the 'Regression' window pictured and produced in 'REGRESSION' above (by following **Stat > Regression > Regression**). Click on the 'Storage' button. In the corresponding window, select 'Fits' and 'Residuals'. Clicking 'OK' and then clicking 'OK' again in the 'Regression' window will produce the 'Regression Analysis' data analyzed above together with the desired data as two new columns in the data window.

MINITAB - Untitled Worksheet

File Edit Manip Calc Stat Graph Editor Window Help

Session

The regression equation is
 $\text{math} = 2493 + 1.07 \text{ student}$

Predictor	Coef	StDev	T	P
Constant	2493	1267	1.97	0.097
student	1.0663	0.2888	3.69	0.010

S = 188.9 R-Sq = 69.4% R-Sq(adj) = 64.3%

Analysis of Variance

Source	DF	SS	MS	F	P
Regression	1	486552	486552	13.63	0.010
Error	6	214209	35702		
Total	7	700762			

MTB >

Data

	C1	C2	C3	C4	C5	C6	C7
↓	student	math	FITS1	RESI1			
1	4595	7364	7392.44	-28.443			
2	4827	7547	7639.83	-92.830			
3	4427	7099	7213.30	-114.301			
4	4258	6894	7033.09	-139.092			
5	3995	6572	6752.65	-180.650			
6	4330	7156	7109.87	46.132			
7	4265	7232	7040.56	191.443			
8	4351	7450	7132.26	317.740			
9							
10							
11							
12							

Thus in particular, for $x = 4258$, $\hat{y} = 7033.09$ and the residue $y - \hat{y} = 6894 - 7033.09 = -139.092$. Note that if you create the fitted line plot as above, you can also select the 'Storage' button from that window to store the fits and residuals in the same manner.

- SPECIAL OPTIONS

To set the intercept of the least-squares linear regression line equal to zero (as is done, for

example, in Exercise 9.20 of the text) click the 'Options' button in the 'Regression' window (see REGRESSION above), find the 'Fit Intercept' box and remove the check mark by clicking with the mouse. Another way to do this is simply to follow **Stat > Fit Intercept** to disengage the Fit Intercept option. In either case, when you are finished, do not forget to go back and reset the intercept for the rest of your work (and the next user's)!

To find predicted values (*fits*) for *x*-values which were not part of the original data, e.g. for student = $x = 4200$ above, simply plug in the value into the equation for the least-squares regression line found in 'REGRESSION' above, and compute by hand or use a calculator. MINITAB will carry this out for you, but it will produce much more data than you need now. (This additional data will be useful for the material in Ch. 10, discussed in Section 4 of this guide.) However, for the sake of completeness, and for later use, here is the way to ask MINITAB to produce the fit for an *x*-value: Follow **Stat > Regression > Regression** to open the 'Regression' window as pictured in 'REGRESSION' above. Click on the 'Options' button to produce the window below.

MINITAB - Untitled Worksheet

File Edit Manip Calc Stat Graph Editor Window Help

Session

Worksheet size: 100000 cells

MTB >

Regression - Options

C1 student
C2 math

Weights: |

Display

Variance inflation factors
 Durbin-Watson statistic

Prediction intervals for new observations

4200

Confidence level: 95

Storage

Fits Co
 SDs of fits Pre

Select

Help

Data

	C1	C2
↓	student	math
1	4595	7364
2	4827	7547
3	4427	7099
4	4258	6894
5	3995	6572
6	4330	7156
7	4265	7232
8	4351	7450
9		
10		
11		
12		

Start Adobe Photoshop MINITAB.wri - WordPad C:\WINDOWS\DESKTOP...

Enter the x-value for which you wish to find the predicted value, e.g. $x = 4200$ in this case, under 'Prediction intervals for new observations'. Click 'OK' here and 'OK' again in the 'Regression' window (here, you must still have columns entered for 'Response' and 'Predictor' as described previously). The 'Regression Analysis' output will now include the last line

Fit	StDev	Fit	95.0% CI	95.0% PI
6971.2	84.8	(6763.6, 7178.9)	(6464.3, 7478.2)

The fit for the value $x = 4200$ is $y^{\wedge} = 6971.2$. Unfortunately, the x -value itself is not displayed. This command can also be used to calculate the predicted values for several x -values at the same time by entering these values into a column and entering the column name in place of the single x -value above. In particular, carrying this out for the entire column of original x -values (in our case, by entering C1 instead of 4200 under 'Prediction intervals for new observations') yields the data

Fit	StDev	Fit	95.0% CI	95.0% PI
7392.4	91.0	(7169.7, 7615.2)	(6879.1, 7905.8)
7639.8	145.1	(7284.6, 7995.0)	(7056.7, 8223.0)
7213.3	68.1	(7046.6, 7380.0)	(6721.7, 7704.9)
7033.1	75.7	(6847.9, 7218.3)	(6534.9, 7531.3)
6752.6	130.0	(6434.5, 7070.8)	(6191.3, 7314.0)
7109.9	68.4	(6942.4, 7277.3)	(6618.0, 7601.7)
7040.6	74.7	(6857.6, 7223.5)	(6543.2, 7537.9)
7132.3	67.4	(6967.4, 7297.1)	(6641.3, 7623.2)

where the 'Fit' column gives the predicted values corresponding to the x -values in C1. A further discussion of this feature and its output will appear in Section 4.

4. Inference for Regression

NOTE: In this section, we employ the same sample data as in Section 3 above, namely, the data from Example 2.11 of the text. (See the note at the beginning of Section 3.)

- LEAST-SQUARES REGRESSION LINE AND POPULATION REGRESSION LINE

The *least-squares regression line* $y^{\wedge} = b_0 + b_1 x$ is the estimate for the *population regression line* $\beta_0 + \beta_1 x$. The method for computing the least-square regression line is detailed in Section 3 (see 'LINEAR REGRESSION' there). In the particular case of our sample data, the path **Stat>Regression>Regression** produced the data

The regression equation is
 $\text{math} = 2493 + 1.07 \text{ student}$

Predictor	Coef	StDev	T	P
Constant	2493	1267	1.97	0.097
student	1.0663	0.2888	3.69	0.010

S = 188.9 R-Sq = 69.4% R-Sq(adj) = 64.3%

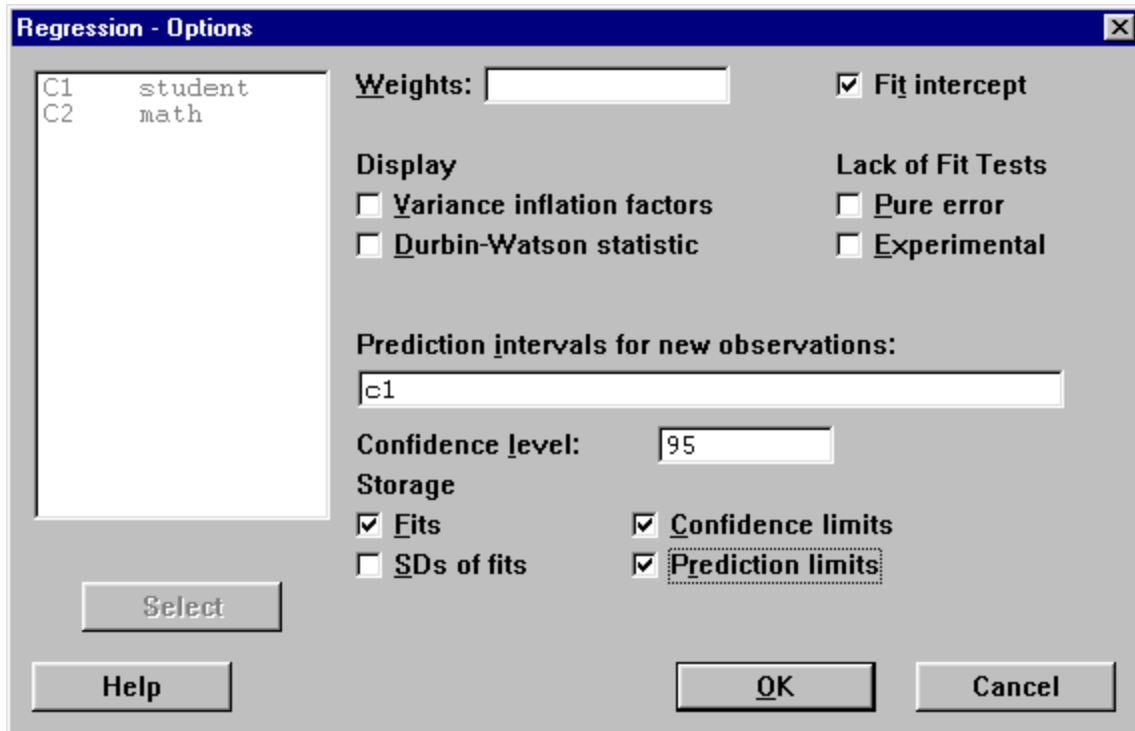
Analysis of Variance

Source	DF	SS	MS	F	P
Regression	1	486552	486552	13.63	0.010
Error	6	214209	35702		
Total	7	700762			

As we saw in Section 3, the least-squares regression line for 'math' (the y-variable) on 'student' (the x-variable) is $y^{\wedge} = 2493 + 1.07x$. More precise estimates for b_1 as 2493 and for b_0 as 1.0663 are given in the column 'Coef' (for 'Coefficient') and across from the rows 'Constant' and 'student', respectively. The remaining columns 'St Dev', 'T', and 'P' give useful information as well. The column 'St Dev' gives *estimated standard deviations* (also called *standard errors* by the text) s_{b1} and s_{b0} used to produce confidence intervals for β_1 and β_0 , respectively, with $s_{b1} = 1267$ and $s_{b0} = .2888$ for our particular example. The column 'T' gives the *t-statistics* used in significance tests for the null hypotheses $H_0 : \beta_1 = 0$ and $H_0 : \beta_0 = 0$, against a two-sided alternative. Thus, for a given row, $T = \text{Coef} / \text{St Dev}$, so that, say, the first row gives the t-statistic $t = b_1 / s_{b1}$, yielding $t = 2493/1267 = 1.97$ for our example. Finally, the column 'P' lists the *P-values* associated with the t-statistics given in the 'T' column, where again the significance test is assumed to be two-sided. Thus, the P-value for the two-sided test against the null hypothesis $H_0 : \beta_1 = 0$ is $2P(T > |t|) = 2P(T > 1.97) = 0.097$ for our example above. To find the P-value for the one-sided test against the same null hypotheses above, divide the corresponding P-value by 2.

- FITS, CONFIDENCE INTERVALS, AND PREDICTION INTERVALS

To produce the *fits*, which are the y^{\wedge} values for the observed x-values, *confidence intervals* μ^{\wedge}_y plus or minus $t^*s_{\mu^{\wedge}}$ for a mean response μ_y , and *prediction intervals* y^{\wedge} plus or minus $t^*s_{y^{\wedge}}$, follow the path **Stat>Regression>Regression** and select the 'Options' button to produce the window below:



Enter the column corresponding to the x-values (here, C1) under 'Prediction intervals for new observations'. Check the boxes marked 'Fits', 'Confidence limits', and 'Prediction limits'. Notice that the preset confidence level for these is 95%, but any other desired confidence level can be entered in the box 'Confidence level'. Click 'OK' in this window and in the one before it. Along with the regression analysis data produced above, there will also appear the desired data:

Fit	StDev	Fit	95.0% CI	95.0% PI
7392.4	91.0	(7169.7,	7615.2)	(6879.1, 7905.8)
7639.8	145.1	(7284.6,	7995.0)	(7056.7, 8223.0)
7213.3	68.1	(7046.6,	7380.0)	(6721.7, 7704.9)
7033.1	75.7	(6847.9,	7218.3)	(6534.9, 7531.3)
6752.6	130.0	(6434.5,	7070.8)	(6191.3, 7314.0)
7109.9	68.4	(6942.4,	7277.3)	(6618.0, 7601.7)
7040.6	74.7	(6857.6,	7223.5)	(6543.2, 7537.9)
7132.3	67.4	(6967.4,	7297.1)	(6641.3, 7623.2)

Notice that entering a single x-value, e.g. '7050' (for x = 'student') instead of an entire column of values (e.g. 'C1') in the window pictured above will produce a predicted value \hat{y} and associated confidence and prediction intervals as pictured below:

```
Fit  StDev  Fit      95.0% CI      95.0% PI
10010.3  773.8  ( 8116.3, 11904.3) ( 8060.6, 11959.9) XX
X denotes a row with X values away from the center
XX denotes a row with very extreme X values
```

- ANALYSIS OF VARIANCE FOR REGRESSION

The printout produced above by following the command path **Stat>Regression>Regression** includes a section labeled 'Analysis of Variance'.

Analysis of Variance

Source	DF	SS	MS	F	P
Regression	1	486552	486552	13.63	0.010
Error	6	214209	35702		
Total	7	700762			

This part of the printout matches almost exactly the analysis of variance (ANOVA) table given on page 658 of the text, with the word 'Model' in place of 'Regression'. (In addition, the printout above includes the P-value for the analysis of variance F test.) Consequently, to compare the notation of the text and that of the table above, label the row 'Regression' by 'M' (for 'Model'), the row 'Error' by 'E', and the row 'Total' by 'T'. Then reading down the sum of squares column 'SS' gives $SSM = 486552$, $SSE = 214209$, and $SST = SSM + SSE = 700762$. Similarly, reading down the mean squares column 'MS' gives $MSM = 486552$ and $MSE = 35702$. Finally, $F = MSM/MSE = 13.63$. The 'DF' column similarly gives the degrees of freedom DFM, DFE, and DFT.

Your course may or may not cover much of the material in the text's Section 9.1 on analysis of variance, but you may at least be required to read off some values from a table such as the one above to use in calculations for simple linear regression. In particular, the square r^2 of the sample correlation (see also 'INFERENCE FOR CORRELATION' below) is the fraction SSM/SST , and this, as noted above, can be pulled from the data in the ANOVA table.

- INFERENCE FOR CORRELATION

The *population correlation coefficient* ρ can be estimated using the *sample correlation* r . This value can be extracted from the data produced by following the command path **Stat>Regression>Regression** (see for example 'LEAST-SQUARES REGRESSION LINE AND POPULATION REGRESSION LINE' in this section above). Included in the output is the line

```
S = 188.9      R-Sq = 69.4%      R-Sq(adj) = 64.3%
```

Here, r^2 appears as 'R-Sq'. So, to find r , just take the positive square root of 'R-Sq', e.g., in this case take the square root of 69.4. This gives the value of r used in the test for a zero population correlation, as well as that which appears in the relationship $b_1/s_{b1} = r s_y/s_x$. Finally, $r^2 = SSM/SST$ (see 'ANALYSIS OF VARIANCE' above) so the complete printout following the path **Stat>Regression>Regression** gives two ways of finding r^2 , and hence r .

- **ADDITIONAL FEATURES**

To set the intercept of the least-squares regression line equal to zero, see the instructions under 'Special Options' at the end of Section 3 of this guide.