

Introduction to Stata
Course given at the Bank of England
Monetary Analysis
Spring/Summer 2014

Michael McMahon¹

¹ This is a version of the course and notes that I have given to MSc and PhD students in the Department of Economics at the London School of Economics (2006, 2007 – LT and MT) and at University of Warwick (2008 and 2009). It builds on earlier courses given by Martin Stewart (2004) and Holger Breinlich (2005). Any errors are my own responsibility and should you wish to contact me please email me (m.mcmahon@warwick.ac.uk).

Full Table of contents

GETTING TO KNOW STATA AND GETTING STARTED.....	5
WHY STATA?	5
WHAT STATA LOOKS LIKE	5
DATA IN STATA	6
GETTING HELP	7
<i>Manuals</i>	7
<i>Stata's in-built help and website</i>	7
<i>The web</i>	7
<i>Colleagues</i>	7
DIRECTORIES AND FOLDERS	7
GETTING DATA INTO STATA	9
READING DATA INTO STATA.....	9
<i>use</i>	9
<i>insheet</i>	9
<i>Stat/Transfer program</i>	10
<i>Manual typing</i>	10
VARIABLE AND DATA TYPES	10
<i>Indicator or data variables</i>	10
<i>Numeric or string data</i>	10
<i>Missing values</i>	11
DATABASE MANIPULATION.....	12
EXAMINING THE DATA	12
<i>List</i>	12
<i>Browse/Edit</i>	12
<i>Assert</i>	12
<i>Describe</i>	13
<i>Codebook</i>	13
<i>Summarize</i>	13
<i>Tabulate</i>	13
<i>Inspect</i>	14
<i>Graph</i>	14
SAVING THE DATASET	14
<i>Preserve and restore</i>	14
KEEPING TRACK OF THINGS	15
<i>Do-files and log-files</i>	15
<i>Labels</i>	16
<i>Notes</i>	17
<i>Review</i>	17
SOME SHORTCUTS FOR WORKING WITH STATA	17
ORGANISING DATASETS	18
<i>Rename</i>	18
<i>Recode and Replace</i>	18
<i>Keep and drop (including some notes on if-processing)</i>	18
<i>Sort</i>	19
<i>By-processing</i>	20
<i>Append and merge</i>	20
<i>Collapse</i>	21
<i>Order, aorder, and move</i>	21
CREATING NEW VARIABLES.....	22
<i>Generate, egen, replace</i>	22
<i>Converting strings to numerics and vice versa</i>	22
<i>Combining and dividing variables</i>	23
<i>Dummy variables</i>	23
<i>Lags and leads</i>	24

CLEANING THE DATA	24
<i>Fillin and expand</i>	24
<i>Interpolation and extrapolation</i>	25
PANEL DATA MANIPULATION: LONG VERSUS WIDE DATA SETS	26
<i>Reshape</i>	26
ESTIMATION	28
LINEAR REGRESSION	28
POST-ESTIMATION	31
<i>Prediction</i>	31
<i>Hypothesis testing</i>	32
<i>Extracting results</i>	33
<i>OUTREG2 – the ultimate tool in Stata/Latex or Word friendliness?</i>	34
EXTRA COMMANDS ON THE NET	34
<i>Looking for specific commands</i>	34
CONSTRAINED LINEAR REGRESSION	36
DICHOTOMOUS DEPENDENT VARIABLE	36
PANEL DATA	37
<i>Describe pattern of xt data</i>	37
<i>Summarize xt data</i>	38
<i>Tabulate xt data</i>	38
<i>Panel regressions</i>	39
TIME SERIES DATA	42
<i>Stata Date and Time-series Variables</i>	42
<i>Getting dates into Stata format</i>	43
<i>Using the time series date variables</i>	45
<i>Making Use of Dates</i>	45
<i>Time Series Tricks Using Dates</i>	46
PROGRAMMING	47
PROGRAM BASICS	47
<i>Creating or “defining” a program</i>	47
<i>Naming a program</i>	47
<i>Redefining a program</i>	48
<i>Debugging a program</i>	48
<i>Program arguments</i>	49
<i>Renaming arguments</i>	49
MACROS	50
<i>Macro contents</i>	53
<i>Manipulation of macros</i>	54
<i>Temporary objects</i>	54
LOOPING	55
<i>for</i>	55
<i>foreach</i>	56
<i>Incremental shift (number of loops is fixed)</i>	57
<i>Macro shift (number of loops is variable)</i>	58
BRANCHING	59
ADO PROGRAMMING	61
<i>Median Program -- Version #1</i>	61
<i>Median Program -- Version #2</i>	61
<i>Median Program -- Version #3</i>	62

Course Outline

At LSE, this course was run over 5 weeks while at Warwick the material was included as part of a 13 lecture course covering research methods in general. At the Bank I run the course over a single day and it is therefore not possible to cover everything – it never is with a program as large and as flexible as Stata. Therefore, I shall endeavour to take you from a position of complete novice (some having never seen the program before), to a position from which you are confident users who, through practice, can become intermediate and onto expert users.

In order to help you, the course is based around practical examples – these examples use macro data but have no economic meaning to them. They are simply there to show you how the program works.

My course website which can be accessed using the password “BOE_stata”:
http://www2.warwick.ac.uk/fac/soc/economics/staff/faculty/mcmahon/boe_stata .

The outline of the day is as follows:

Time	Activity
9.00 – 9.25	Introductions, course outline and Stata basics
9.25 – 10.30	Working through an example task list - I
10.30 – 10.45	Coffee Break
10.45 – 12.15	Working through an example task list - II
12.15 – 13.15	Lunch
13.15 – 13.45	Getting data in Stata and Database Manipulation
13.45 – 14.45	Estimation including time-series data
14.45 – 15.00	Coffee Break
15.00 – 16.15	Programming
16.15 – 16.30	Catch-up and final questions time

I am very flexible about this timetable, and I am happy to move at the pace desired by the participants. But if there is anything specific that you wish you to ask me, or material that you would like to see covered in greater detail, I am happy to accommodate these requests.

Getting to Know Stata and Getting Started

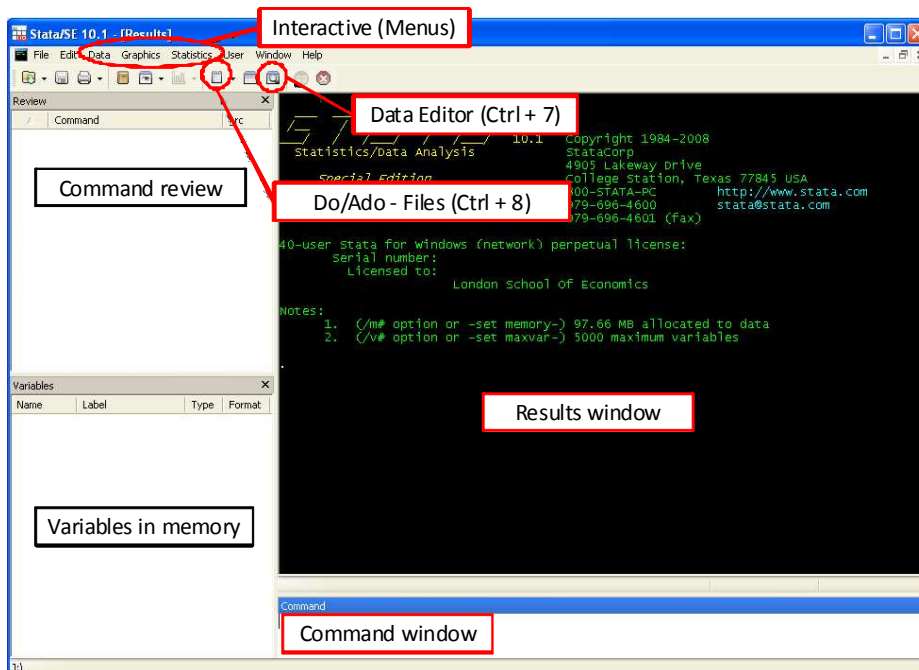
Why Stata?

There are lots of people who use Stata for their applied econometrics work. But there are also numerous people who use other packages (Eviews or Microfit for those getting started, RATS/CATS for the time series specialists, or Matlab, Gauss, or Fortran for the really advanced). So the first question that you should ask yourself is why should I use Stata?

Stata is an integrated statistical analysis packaged designed for research professionals. The official website is <http://www.stata.com/>. Its main strengths are handling and manipulating large data sets (e.g. millions of observations!), and it has ever-growing capabilities for handling panel and time-series regression analysis. The most recent version is stata12 and with each version there are improvements in computing speed, capabilities and functionality; the Bank is using stata11. It now also has pretty flexible graphics capabilities. It is also constantly being updated or advanced by users with a specific need – this means that even if a particular regression approach is not a standard feature, you can usually find someone on the web who has written a programme to carry-out the analysis and this is easily integrated with your own software.

What Stata looks like

The Stata package is located on a software server and can be started by either going through the Start menu (Start – Programs – Statistics – stata1) or by double clicking on wsestata.exe in the stata11 folder. It is possible to reconfigure the windows and even change the colour scheme!




Stata is a command-driven package. Although the newest versions also have pull-down menus from which different commands can be chosen, the best way to learn Stata is still by typing in the commands. This has the advantage of making the switch to programming much easier which will be necessary for any serious econometric work. However, sometimes the exact syntax of a command is hard to get right – in these cases, I often use the menu-commands to do it once and then copy the syntax that appears.

You can enter commands in either of three ways:

- Interactively: you click through the menu on top of the screen
- Manually: you type the first command in the command window and execute it, then the next, and so on.
- Do-file: type up a list of commands in a "do-file", essentially a computer programme, and execute the do-file.

The vast majority of your work should use do-files. If you have a long list of commands, executing a do-file once is a lot quicker than executing several commands one after another. Furthermore, the do-file is a permanent record of all your commands and the order in which you ran them. This is useful if you need to “tweak” things or correct mistakes – instead of inputting all the commands again one after another, just amend the do-file and re-run it. Working interactively is useful for “I wonder what happens if ...?” situations. When you find out what happens, you can then add the appropriate command to your do-file. To start with we’ll work interactively, and once you get the hang of that we will move on to do-files.



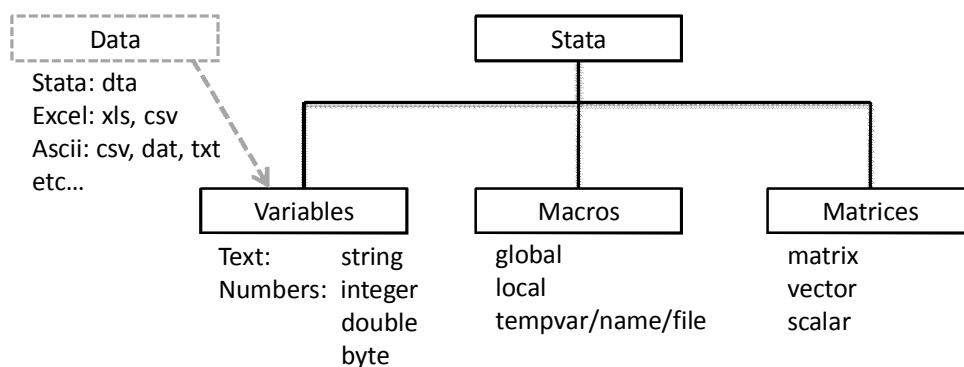
Data in Stata

Stata is a versatile program that can read several different types of data. Mainly files in its own dta format, but also raw data saved in plain text format (ASCII format). Every program you use (i.e. Excel or other statistical packages) will allow you to export your data in some kind of ASCII file. So you should be able to load all data into Stata.

When you enter the data in Stata it will be in the form of variables. Variables are organized as column vectors with individual observations in each row. They can hold numeric data as well as strings. Each row is associated with one observation, that is the 5th row in each variable holds the information of the 5th individual, country, firm or whatever information your data entails.

Information in Stata is usually and most efficiently stored in variables. But in some cases it might be easier to use other forms of storage. The other two forms of storage you might find useful are matrices and macros. Matrices have rows and columns that are not associated with any observations. You can for example store an estimated coefficient vector as a $k \times 1$ matrix (i.e. a column vector) or the variance matrix which is $k \times k$. Matrices use more memory than variables and the size of matrices is limited to 11,000, but your memory will probably run out before you hit that limit. You should therefore use matrices sparingly.

The third option you have is to use macros. Macros are in Stata what variables are in other programming languages, i.e. named containers for information of any kind. Macros come in two different flavours, local or temporary and global. Global macros stay in the system and once set, can be accessed by all your commands. Local macros and temporary objects are only created within a certain environment and only exist within that environment. If you use a local macro in a do-file it, you can only use it for code within that do-file.



Getting help

Stata is a command driven language – there are over 500 different commands and each has a particular syntax required to get any various options. Learning these commands is a time-consuming process but it is not hard. At the end of each class notes I shall try to list the commands that we have covered but there is no way we will cover all of them in this short introductory course. Luckily though, Stata has a fantastic options for getting help. In fact, most of your learning to use Stata will take the form of self-teaching by using manuals, the web, colleagues and Stata's own help function.

Manuals

The Stata manuals are available in MA – many people have them on their desks. The User Manual provides an overall view on using Stata. There are also a number of Reference Volumes, which are basically encyclopaedias of all the different commands and all you ever needed to know about each one. If you want to find information on a particular command or a particular econometric technique, you should first look up the index at the back of any manual to find which volumes have the relevant information. Finally, there is a separate Graphics Manual, panel data manual (cross-sectional time-series) and one on survey data.

Stata's in-built help and website

Stata also has an abbreviated version of its manuals built-in. Click on Help, then Contents. Stata's website has a very useful FAQ section at <http://www.stata.com/support/faqs/>. Both the in-built help and the FAQs can be simultaneously searched from within Stata itself (see menu Help>Search). Stata's website also has a list of helpful links at <http://www.stata.com/links/resources1.html>.

The web

As with everything nowadays, the web is a great place to look to resolve problems. There are numerous chat-rooms about stata commands, and plenty of authors put new programmes on their websites. Google should help you here.

Colleagues

The other place where you can learn a lot is from speaking to colleagues who are more familiar with Stata functions than you are – the Bank is littered with people who spend large parts of their days typing different commands into Stata, you should make use of them if you get stuck. You can use the user group email – see the intranet for details.

Directories and folders

Like Dos and Windows, Stata can organise files in a tree-style directory with different folders. You should use this to organise your work in order to make it easier to find things at a later date. For example, create a folder "data" to hold all the datasets you use, sub-folders for each dataset, and so on. You can use some Dos commands in Stata, including:

```
. cd "C:\Stata classes\"      - change directory to "C:\Stata classes\"
. mkdir "stata"              - creates a new directory within the current one (here, C:\Stata classes\stata)
. dir                        - list contents of directory or folder
```

Note, Stata is case sensitive, so it will not recognise the command CD or Cd. Also, quotes are only needed if the directory or folder name has spaces in it – "h:\temp\first folder" – but it's a good habit to use them all the time.

Our first few tasks:

In order to get straight into the programme, I want to do a simple series of steps that mimic the likely first experiences you would have when starting some simple data analysis work. We can then go back and talk a bit about other capabilities that Stata has, and the specifics of different commands.

The steps we will complete are the following:

1. You have, or will soon have, in your stata folder, data (in a variety of forms) from the Penn World Tables (<http://pwt.econ.upenn.edu/>)
2. Change the directory to this folder using the `cd` command.
3. Start a “do file” to keep track of your work.
4. Read the first dataset into Stata using `insheet`
5. Examine the data:
 - `browse`
 - `describe`
 - `sum`
 - look for others in the menus
6. To help your colleagues/co-authors, label the dataset and a few variables
7. Insert a note on where you got the data from
8. Rename the country code variable
9. Check to make sure that population is non-negative; generate a histogram of the population data
10. Save the dataset in Stata format in your default drive
11. Now load the second dataset into Stata
12. Check the variables that you have loaded.
13. Add this to the existing file you saved to make a larger data set covering more years – command is `append`.
14. Save this larger dataset
15. Repeat these steps to add the 3rd dataset which is already in stata format (.dta) – what is in this file?
16. Load the 4th dataset, which is already in stata format (.dta), examine these data and add this file to the existing database – here we need a `merge`.
17. Sort out the appearance of the now completed dataset.
18. Create new variables:
 - `growth`: annual growth rate of real per capita GDP
 - `high_growth` = A dummy for each year in which growth exceeds 5%
 - `EU` = dummy variable for EU membership (France, Germany, Italy, & United Kingdom)
 - `max_growth` = highest rate of growth each year across all G-7 countries
 - `Y70` = real GDP per capita in 1970 in each country
19. Run a simply linear regression of:
Growth on investment share, Y70 and population
20. Now run the same regression using country fixed effects:
 - Either, create dummies for each country and include them;
 - Or, set the panel variables (time and country) and run a fixed effects panel regression (`xtreg , fe`).

Getting Data into Stata

Reading data into Stata

There are different ways of reading or entering data into Stata:

use

If your data is in Stata format, then simply read it in as follows:

```
. use "F:\Stata classes\stata1.dta", clear
```

The `clear` option will clear the revised dataset currently in memory before opening the other one.

Or if you changed the directory already, the command can exclude the directory mapping:

```
use "stata1.dta", clear
```

insheet

If your data is originally in Excel or some other format, you need to prepare the data before reading it directly into Stata. You need to save the data in the other package (e.g. Excel) as either a csv (comma separated values) or txt (tab-delimited ASCII text) file. There are some ground-rules to be followed when saving a csv- or txt-file for reading into Stata:

- The first line in the spreadsheet should have the variable names, e.g. series/code/name, and the second line onwards should have the data. If the top row of the file contains a title then delete this row before saving.
- Any extra lines below the data or to the right of the data (e.g. footnotes) will also be read in by Stata, so make sure that only the data itself is in the spreadsheet before saving. If necessary, select all the bottom rows and/or right-hand columns and delete them.
- The variable names cannot begin with a number. If the file is laid out with years (e.g. 1980, 1985, 1990, 1995) on the top line, then Stata will run into problems. In such instances, place an underscore in front of each number (e.g. select the row and use the spreadsheet package's "find and replace" tools): 1980 becomes `_1980` and so on.
- Make sure there are no commas in the data as it will confuse Stata about where rows and columns start and finish (again, use "find and replace" to delete any commas before saving – you can select the entire worksheet in Excel by clicking on the empty box in the top-left corner, just above 1 and to the left of A).
- Some notations for missing values can confuse Stata, e.g. it will read double dots (. .) or hyphens (-) as text. Use find & replace to replace such symbols with single dots (.) or simply to delete them altogether.

Once the csv- or txt-file is saved, you then read it into Stata using the command:

```
. insheet using "F:\Stata classes\stata1.txt", clear
```

Note that if we had already changed to `F:\Stata classes\` using the `cd` command, we could simply type:

```
. insheet using "stata1.txt", clear
```

There are a few useful options for the `insheet` command ("options" in Stata are additional features of standard commands, usually appended after the command and separated by a comma – we will see many more of these). The first option is `clear` which you can use if you want to `insheet` a new file while there is still data in memory:

```
. insheet using "F:\Stata classes\stata1.txt", clear
```

Alternatively, you could first erase the data in memory using the command `clear` and then `insheet` as before.

The second option, `names`, tells Stata that the file you `insheet` contains the variable names in the first row. Normally, Stata should recognise this itself but sometimes it simply doesn't – in these cases `names` forces Stata to use the first line in your data for variable names:

```
. insheet using "F:\Stata classes\stata1.txt", names clear
```

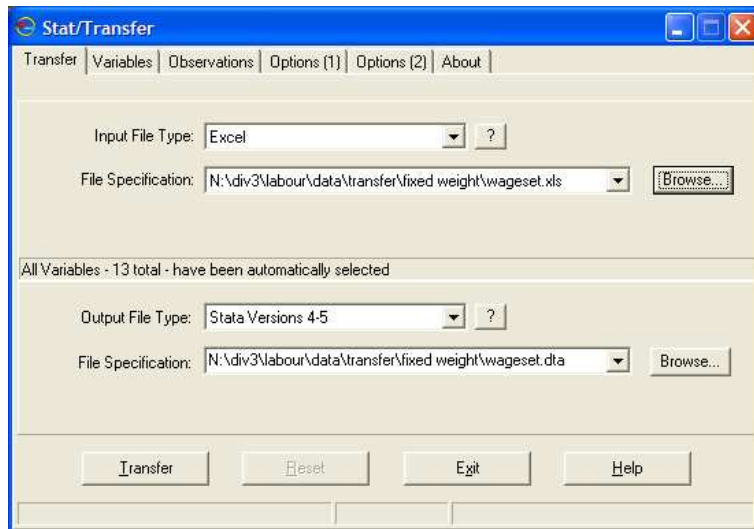
Finally, the option `delimiter("char")` tells Stata which delimiter is used in the data you want to `insheet`. Stata's `insheet` automatically recognises tab- and comma-delimited data but sometimes different delimiters are used in datasets (such as ";"):

```
. insheet using "h:\wdi-sample.txt", delimiter(";")
```

Stat/Transfer program

This is a separate package that can be used to convert a variety of different file-types into other formats, e.g. Excel into Stata or vice versa. You should take great care to examine the converted data thoroughly to ensure it was converted properly.

It is used in a very user-friendly way (see screen shot below) and is useful for changing data between lots of different packages and format.



Manual typing

The tedious last resort – if the data is not available in electronic format, you may have to type it in manually. Start the Stata program and use the edit command – this brings up a spreadsheet-like where you can enter new data or edit existing data.

This can be done directly by typing the variables into the window, or indirectly using the input command.

Variable and data types

Indicator or data variables

You can see the contents of a datafile using the `browse` or `edit` command. The underlying numbers are stored in “data variables”, e.g. the `cgdp` variable contains national income data and the `pop` variable contains population data. To know what each data-point refers to, you also need at least one “indicator variable”, in our case `countryisocode` (or `country`) and `year` tell us what country and year each particular `gdp` and population figure refers to. The data might then look as follows:

country	Countryisocode	year	pop	cgdp	openc
Canada	CAN	1990	27700.9	19653.69	51.87665
France	FRA	1990	58026.1	17402.55	43.46339
Italy	ITA	1990	56719.2	16817.21	39.44491
Japan	JPN	1990	123540	19431.34	19.81217
United Kingdom	GBR	1990	57561	15930.71	50.62695
United States	USA	1990	249981	23004.95	20.61974

This layout ensures that each data-point is on a different row, which is necessary to make Stata commands work properly.

Numeric or string data

Stata stores or formats data in either of two ways – numeric or string. Numeric will store numbers while string will store text (it can also be used to store numbers, but you will not be able to perform numerical analysis on those numbers).

Numeric storage can be a bit complicated. Underneath its Windows platform, Stata, like any computer program, stores numbers in binary format using 1's and 0's. Binary numbers tend to take up a lot of space, so Stata will try to store the data in a more compact format. The different formats or storage types are:

byte : integer between -127 and 126 e.g. dummy variable

int : integer between -32,767 and 32,766 e.g. year variable

long : integer between -2,147,483,647 and 2,147,483,646 e.g. population data

float : real number with about 8 digits of accuracy e.g. production output data

double : real number with about 16 digits of accuracy

The Stata default is “float”, and this is accurate enough for most work. However, for critical work you should make sure that your data is “double”. Note, making all your numerical variables “double” can be used as an insurance policy against inaccuracy, but with large data-files this strategy can make the file very unwieldy – it can take up lots of hard-disk space and can slow down the running of Stata. Also, if space is at a premium, you should store integer variables as “byte” or “int”, where appropriate.

String is arguably more straightforward – any variable can be designated as a string variable and can contain up to 80 characters, e.g. the variable `name` contains the names of the different countries. Sometimes, you might want to store numeric variables as strings, too. For example, your dataset might contain an indicator variable `id` which takes on 9-digit values. If `id` were stored in float format (which is accurate up to only 8 digits), you may encounter situations where different `id` codes are rounded to the same amount. Since we do not perform any calculations on `id` we could just as well store it in string format and avoid such problems.

To preserve space, only store a variable with the minimum string necessary – so the longest named `name` is “United Kingdom” with 14 letters (including the space). A quick way to store variables in their most efficient format is to use the `compress` command – this goes through every observation of a variable and decides the least space-consuming format without sacrificing the current level of accuracy in the data.

```
. compress
```

Missing values

Missing numeric observations are denoted by a single dot (`.`), missing string observations are denoted by blank double quotes (`""`).

Database Manipulation

Examining the data

It is a good idea to examine your data when you first read it into Stata – you should check that all the variables and observations are there and in the correct format.

List

As we have seen, the `browse` and `edit` commands start a pop-up window in which you can examine the raw data. You can also examine it within the results window using the `list` command – although listing the entire dataset is only feasible if it is small. If the dataset is large, you can use some options to make `list` more useable. For example, list just some of the variables:

```
. list name year GDP
+-----+
|          country   countr~e   year   pop |
+-----+-----+-----+-----+
1. |          Canada      CAN   1990   27700.9 |
2. |          France      FRA   1990   58026.1 |
3. |          Italy       ITA   1990   56719.2 |
4. |          Japan       JPN   1990   123540  |
5. |   United Kingdom    GBR   1990    57561  |
+-----+-----+-----+-----+
6. |   United States     USA   1990   249981  |
7. |          Canada      CAN   1991   28030.9 |
8. |          France      FRA   1991   58315.8 |
9. |          Italy       ITA   1991   56750.7 |
10. |          Japan       JPN   1991   123920  |
+-----+-----+-----+-----+
```

Or list just some of the observations:

```
. list in 45/49
```

Or both:

```
. list country countryisocode year pop in 45/49
+-----+-----+-----+-----+
|          country   countr~e   year   pop |
+-----+-----+-----+-----+
45. |          Italy      ITA   1997   57512.2 |
46. |          Japan      JPN   1997   126166  |
47. |   United Kingdom    GBR   1997    59014  |
48. |   United States     USA   1997   268087  |
49. |          Canada      CAN   1998    30248  |
+-----+-----+-----+-----+
```

Browse/Edit

We have already seen that `browse` starts a pop-up window in which you can examine the raw data. Most of the time we only want to view a few variables at a time however, especially in large datasets with a large number of variables. In such cases, simply list the variables you want to examine after `browse`:

```
. browse name year gdp
. browse name year gdp
```

The difference with `edit` is that this allows you to manually change the dataset.

Assert

With large datasets, it often is impossible to check every single observation using `list` or `browse`. Stata has a number of additional commands to examine data which are described in the following. A first useful command is `assert` which verifies whether a certain statement is true or false. For example, you might want to check whether all GDP values are positive as they should be:

```
assert pop>0
```

```
assert pop<0
```

If the statement is true, `assert` does not yield any output on the screen. If it is false, `assert` gives an error message and the number of contradictions.

Describe

This reports some basic information about the dataset and its variables (size, number of variables and observations, storage types of variables etc.).

```
. describe
```

Codebook

This provides extra information on the variables, such as summary statistics of numerics, example data-points of strings, and so on. Codebook without a list of variables will give information on all variables in the dataset.

```
. codebook country
```

Summarize

This provides summary statistics, such as means, standard deviations, and so on.

```
. summarize
```

Variable	Obs	Mean	Std. Dev.	Min	Max
country	0				
countryiso~e	0				
year	66	1995	3.18651	1990	2000
pop	66	98797.46	79609.33	27700.9	275423
cgdp	66	22293.23	4122.682	15930.71	35618.67
openc	66	42.54479	18.64472	15.91972	86.80463
csave	66	24.31195	5.469772	16.2536	37.80159
ki	66	23.52645	4.634476	17.00269	35.12778
grgdpch	66	1.582974	1.858131	-3.981008	5.172524

Note that `code` and `name` are string variables with no numbers, so no summary statistics are reported for them. Also, `year` is a numeric, so it has summary statistics. Additional information about the distribution of the variable can be obtained using the `detail` option:

```
. summarize, detail
```

Tabulate

This is a versatile command that can be used, for example, to produce a frequency table of one variable or a cross-tab of two variables. There are also options to get the row, column and cell percentages as well as chi-square and other statistics – check the Stata manuals or on-line help for more information.

```
. tab name
```

Name	Freq.	Percent	Cum.
Canada	10	14.29	14.29
France	10	14.29	28.57
Germany	10	14.29	42.86
Italy	10	14.29	57.14
Japan	10	14.29	71.43
United Kingdom	10	14.29	85.71
United States	10	14.29	100.00
Total	70	100.00	

Inspect

This is another way to eyeball the distribution of a variable, including as it does a mini-histogram. Also useful for identifying outliers or unusual values, or for spotting non-integers in a variable that should only contain integers.

```
. inspect cgdg
```

cgdg:				Number of Observations		
-----				Total	Integers	Non-Integers
	#	#		Negative	-	-
	#	#		Zero	-	-
	#	#		Positive	66	66
	#	#			-----	-----
	#	#	#	Total	66	66
	#	#	#	Missing	-	-
+-----					-----	
15930.71			35618.67		66	

(66 unique values)

Graph

Stata has very comprehensive graphics capabilities (type “help graph” for more details). You can graph a simple histogram with the command:

```
. graph twoway histogram cgdg
```

Or a two-way scatterplot using:

```
. graph twoway scatter cgdg pop
```

While graphs in Stata 9 and Stata 10 have the advantage of looking quite fancy, they are also very slow. Often, you just want to visualise data without actually using the output in a paper or presentation. In this case, it is useful to switch to version 7 graphics which are much faster:

```
. graph7 cgdg pop
```

Saving the dataset

The command is simply save:

```
. save "F:\Stata classes\statal.dta", replace
```

The `replace` option overwrites any previous version of the file in the directory you try saving to. If you want to keep an old version as back-up, you should save under a different name, such as “new_G7”. Note that the only way to alter the original file permanently is to save the revised dataset. Thus, if you make some changes but then decide you want to restart, just re-open the original file:

Preserve and restore

If you are going to make some revisions but are unsure of whether or not you will keep them, then you have two options. First, you can save the current version, make the revisions, and if you decide not to keep them, just re-open the saved version. Second, you can use the `preserve` and `restore` commands; `preserve` will take a “photocopy” of the dataset as it stands and if you want to revert back to that copy later on, just type `restore`.

Keeping track of things

Stata has a number of tools to help you keep track of what work you did to datasets, what's in the datasets, and so on.

Do-files and log-files

Instead of typing commands one-by-one interactively, you can type them all in one go within a do-file and simply run the do-file once. The results of each command can be recorded in a log-file for review when the do-file is finished running.

Do-files can be written in any text editor, such as Word or Notepad. Stata also has its own editor built in – click the icon along the top of the screen with the pad-and-pencil logo (although it looks like an envelope to me). Most do-files follow the following format:

```
clear
cd "c:\projects\project1\"
capture log close
log using class1.log, replace
set more off
set memory 100m
```

LIST OF COMMANDS

```
log close
```

To explain the different commands:

`clear` – clears any data currently in Stata's memory. If you try opening a datafile when one is already open, you get the error message: `no; data in memory would be lost`

`cd c:\projects\project1\` - sets the default directory where Stata will look for any files you try to open and save any files you try to save. So, if you type use `wdi-sample.dta`, Stata will look for it in this folder. If, during the session, you want to access a different directory, then just type out its destination in full, e.g. use `"c:\data\production.dta"` will look for the file in the `c:\data` folder. Note again that if you use spaces in file or directory names, you must include the file path in inverted commas.

`capture log close` – closes any log-files that you might have accidentally left open. If there were no log-file actually open, then the command `log close` on its own would stop the do-file running and give the error message: `no log file open`. Using `capture` tells Stata to ignore any error messages and keep going.

`log using class1.log, replace` – starts a log-file of all the results. The `replace` option overwrites any log file of the same name, so if you re-run an updated do-file again the old log-file will be replaced with the updated results. If, instead, you want to add the new log-file to the end of previous versions, then use the `append` option.

`set more off` – when there are a lot of results in the results window, Stata pauses the do-file to give you a chance to review each page on-screen and you have to press a key to get more. This command tells Stata to run the entire do-file without pausing. You can then review the results in the log file.

`set memory 100m` – Stata's default memory may not be big enough to handle large datafiles. Trying to open a file that is too large returns a long error message beginning: `no room to add more observations`. You can adjust the memory size to suit. First check the size of the file using the `describe` command (remember that you can use `describe` for a file that hasn't yet been read into Stata). This reports the size of the file in bytes. Then set memory just a bit bigger. Note, setting it too large can take the PC's memory away from other applications and slow the computer down, so only set it as large as necessary. For example, `describe using "c:\data\WDI-sample.dta"` reports the size of the file to be 2,730 bytes, so `set memory 1m` should be sufficient.

`log close` – closes the log file.

It is good practice to keep extensive notes within your do-file so that when you look back over it you know what you were trying to achieve with each command or set of commands. You can insert notes in two different ways:

*

Stata will ignore a line if it starts with an asterisk *, so you can type whatever you like on that line. Note, the asterisk is also useful for getting Stata to temporarily ignore commands – if you decide later to re-insert the command into your do-file, just delete the asterisk.

```
/* */
```

You can place notes after a command by inserting it inside these pseudo-parentheses, for example:

```
. use "c:\data\WDI-sample.dta", clear /* opens 1998 production data */
```

These pseudo-parentheses are also useful for temporarily blocking a whole set of commands – place `/*` at the beginning of the first command, `*/` at the end of the last, and Stata will just skip over them all.

Labels

You can put labels on datasets, variables or values – this helps to make it clear exactly what the dataset contains.

A dataset label of up to 80 characters can be used to tell you the data source, it’s coverage, and so on. This label will then appear when you describe the dataset. For example, try the following:

```
. label data "Data from Penn World Tables 6.1"
. describe
```

Variable names tend to be short – you can use up to 32 characters, but for ease of use it’s best to stick to about 8 or 10 as a maximum. This can give rise to confusion about what the variable actually represents – what exactly is `gdp` and in what units is it measured? Which is where variable labels, with a capacity of 80 characters, come in.

```
. label variable cgdp "GDP per capita in constant international dollars"
```

It can also be helpful to label different values. Imagine countries were coded as numbers (which is the case in many datasets). In this case, a tabulation may be confusing – what country does 1 represent, or 2 or 3?

```
. tabulation code
```

code	Freq.	Percent	Cum.
1	10	33.33	33.33
2	10	33.33	66.67
3	10	33.33	100.00
Total	30	100.00	

It might be better to label exactly what each value represents. This is achieved by first “defining” a label (giving it a name and specifying the mapping), then associating that label with a variable. This means that the same label can be associated with several variables – useful if there are several “yes/no/maybe” variables, for example. The label name itself can be up to 32 characters long (e.g. `countrycode`), and each value label must be no more than 80 characters long (e.g. “France” or “Italy”).

```
. label define countrycode 1 "Canada" 2 "Germany" 3 "France"
. label values code countrycode
```

Now, the tabulation should make more sense:

```
. tabulation code
```

code	Freq.	Percent	Cum.
Canada	10	33.33	33.33
Germany	10	33.33	66.67
France	10	33.33	100.00
Total	30	100.00	

see what each code represents, use `codebook` or:

```
. label list countrycode
```

```
countrycode:
```

```
1 Canada
2 Germany
3 France
```


Notes

You can also add Post-it notes to your dataset or to individual variables to, for example, remind you of the source of the data, or to remind you of work you did or intend to do on a variable.

```
. note: data from PWT
. note cgdg: This is per capita variable
```

You can also time-stamp these notes:

```
. note cgdg: TS need to add Germany to complete the G7
```

Review your notes by simply typing notes:

```
. notes
```

```
_dta:
```

```
1. data from PWT
```

```
cgdg:
```

```
1. This is per capita variable
2. 15 Feb 2006 13:01 need to add Germany to complete the G7
```

Stata will also tell you that there are notes when you use describe:

```
. describe
```

You can also delete notes. To drop all notes attached to a variable:

```
. note drop cgdg
```

To drop just one in particular:

```
. note drop cgdg in 2
```

Review

One final tool for keeping track is reviewing a list of previous commands. To see the last four, for example:

```
. #review 4
```

This is especially useful if you are working in interactive mode on a “what happens if...”. When you are happy with the sequence of commands you’ve tried, you can *review*, then cut and paste into your do-file.

Some shortcuts for working with Stata

- Most commands can be abbreviated, which saves some typing. For example: `summarize` to `sum`, `tabulate` to `tab`, `save` to `sa`. The abbreviations are noted in the Stata manuals.
- You can also abbreviate variable names when typing. This should be used with caution, as Stata may choose a variable different to the one you intended. For example, suppose you have a dataset with the variables `pop`, `popurban` and `poprural`. If you want summary statistics for `popurban`, the command `sum pop` will actually give statistics for the `pop` variable.
- Stata’s default file type is `.dta`, so you don’t need to type that when opening or saving Stata files: `sa "statal"` is the same as `sa "statal.dta"`
- You can save retyping commands or variable names by clicking on them in the review and variable windows – they will then appear in the command window. You can also cycle back and forth through previous commands using the `PageUp` and `PageDown` keys on your keyboard. Similarly, variable names can be easily entered by clicking on them in the Variables Window (bottom-left of the screen).

Organising datasets

Rename

You may want to change the names of your variables, perhaps to make it more transparent what the variable is:

```
. rename countryisocode country_code
```

Note, you can only rename one variable at a time.

Recode and Replace

You can change the values that certain variables take, e.g. suppose 1994 data actually referred to 1924:

```
. recode year 1994=1924
```

This command can also be used to recode missing values to the dot that Stata uses to denote missings. And you can recode several variables at once. Suppose a dataset codes missing population and gdp figures as -999:

```
. recode pop cgdp -999=.
```

With string variables, however, you need to use the replace command (see more on this command below):

```
. replace country="United Kingdom" if country_code == "GBR"
```

Keep and drop (including some notes on if-processing)

The original dataset may contain variables you are not interested in or observations you don't want to analyse. It's a good idea to get rid of these first – that way, they won't use up valuable memory and these data won't inadvertently sneak into your analysis. You can tell Stata to either `keep` what you want or `drop` what you don't want – the end results will be the same. For example, we can get rid of unwanted variables as follows:

```
. keep country year pop cgdp
```

or

```
. drop country_code openc csave ki
```

or

```
. drop country_code openc - gdp_growth
```

Each of these will leave you with the same set of variables. Note that the hyphen sign (-) is a useful shortcut, e.g. the first one indicates all the variables between `openc` and `gdp_growth` are to be dropped. However, you must be careful that the order of the variable list is correct, you don't want to inadvertently drop a variable that you thought was somewhere else on the list. The variable list is in the variables window or can be seen using either the `desc` or `sum` commands.

You can also drop or keep observations, such as those after or before 1995:

```
. keep if year>=1995
```

or

```
. drop if year<1995
```

Note, the different relational operators are:

`==` equal to

`!=` not equal to

`>` greater than

`>=` greater than or equal to

`<` less than

`<=` less than or equal to

Keeping observations for the years 1990 to 1995 only:

```
. keep if (year>=1990 & year<=1995)
```

or

```
. drop if (year<1990 | year>1995)
```

Or, to get really fancy, keep the observations for 1990-95 and 1997-99:

```
. keep if ((year>=1990 & year<=1995) | (year>=1997 & year<=1999))
```

Note, the different logical operators are:

```
& and
| or
~ not
! not
```

You may want to drop observations with specific values, such as missing values (denoted in Stata by a dot):

```
. drop if pop==.
```

You may want to keep observations for all countries other than those for Italy:

```
. drop if country_code!="ITA"
```

Note, with string variables, you must enclose the observation reference in double quotes. Otherwise, Stata will claim not to be able to find what you are referring to.

If you know the observation number, you can selectively keep or drop different observations. Dropping observations 1 to 10:

```
. drop if _n<=10
```

Dropping the last observation (number `_N`) in the dataset:

```
. drop if _n==_N
```

Finally, you may want to keep only a single occurrence of a specific observation type, e.g. just the first observation of each country code:

```
. keep if country[_n]~=country[_n-1]
```

or simply

```
. keep if country~=country[_n-1]
```

Stata starts at observation number one and applies the command, then moves onto observation two and applies the command again, then onto three and so on. So, starting at one `_n=1` but there is no observation `_n-1 = 0`, so the country in one cannot equal the country in zero and the observation will be kept. Moving on to two: the country in two equals the country in one (both AGO), so the observation will be dropped. Each subsequent observation with country AGO will also be dropped. When we get to an observation with a different country (which will be ALB), the two countries will be different (`AGO~=ALB`) and the observation will be kept. Thus, we will end up being left with just the first observation for each country.

Sort

From the previous example, hopefully you will have realised the importance of the order of your observations. If the country codes had started out all jumbled up, then we would have ended up with a completely different set of observations. Suppose we applied the above command to the following dataset:

Number in dataset	country	Result
1	AGO	Kept since <code>_n=0</code> does not exist
2	AGO	Dropped since <code>country==country[_n-1]</code>
3	ALB	Kept
4	ALB	Dropped
5	AGO	Kept
6	ALB	Kept
7	BEL	Kept

We would actually end up with numerous occurrences of some country codes. This shows how sorting the data first is important:

```
. sort country
```

If you wanted to make sure the observation that was kept was the earliest (i.e. 1950), then first:

```
. sort country year
```

This command first sorts the data by country, and then within each country code it sorts the data by year. This ensures that the first observation for every country (the one that is kept) will be 1950.

Note that sorting is in ascending order (A,B,C or 50, 51, 52). To sort in descending order, you need to use the `gsort` command:

```
. gsort -country
```

This gives ZWE first, then ZMB, ZAR, ZAF, YEM and so on. Note that you need to place a minus sign before every variable you want to sort in descending order. This command allows you to sort in complicated ways, e.g. to sort country codes in descending order but then years in ascending order:

```
. gsort -country year
```

By-processing

You can re-run a command for different subsets of the data using the `by` prefix. For example, to get summary statistics of population broken down by year:

```
. so year
. by year: sum pop
```

Note that you have to either sort the data first or use the `bysort` option:

```
. bysort year: sum pop
```

The `by` prefix causes the `sum` command to be repeated for each unique value of the variable `year`. The result is the same as writing a list of `sum` commands with separate `if` statements for each year:

```
. sum pop if year==1990
. sum pop if year==1991
. sum pop if year==1992
. sum pop if year==1993
. sum pop if year==1994
```

By-processing can be useful when organising your dataset. In our sort examples above we asked Stata to keep only the 1990 observations for each country. Instead of trying to make sure the data is sorted in the proper order and the `keep/drop` command is coded correctly (both of which can often be very confusing), it is much easier to:

```
. bysort country: keep if year==1990
```

That's not to say that the first methodology is entirely useless – you may have a dataset where different countries have observations for different years (so not all have 1990 data), and you may want to keep the earliest observation from each country. In such a case, you may have to revert to our earlier example.

Append and merge

You can combine different datasets into a single large dataset using the `append` and `merge` commands. `append` is used to add extra observations (rows). Suppose you have two datasets containing the G7 less Germany PWT data for different countries and/or different years. The datasets have the same variables `country` / `year` / `pop` / etc, but one dataset has data for 1970-1990 (called “`Stata2.dta`”) and the other has data for 1975-1998 (called “`stata1.dta`”).

```
. use "F:\Stata classes\stata1.dta", clear
. append using "F:\Stata classes\Stata2.dta"
. save "F:\Stata classes\G7 less Germany pwt.dta", replace
```

`append` is generally very straightforward. There is one important exception, however, if the two datasets you want to `append` have stored their variables in different formats (meaning string vs. numeric – having different numeric formats, for example byte vs. float, does not matter). In this case, Stata converts the data in the file to be appended to the format of the original file and in the process replaces all values to missing! To detect such problems while using `append`, watch out for messages like:

```
. (note: pop is str10 in using data but will be float now)
```

This indicates that a variable (here: `pop`) has been transformed from string to float – and contains all missing values now for the appending dataset (here: all years 1970-1990). It is thus very important to check via `describe` that the two files you intend to `append` have stored all variables in the same broad data categories (string/numeric). If this is not the case, you will need to transform them first (see the commands `real` and `string` below).

`merge` is used to add extra variables (columns). Suppose we now also have a second dataset containing the same indicator variables `country` / `year`, but one dataset has data for GDP per capita and other variables, and the second has data for shares in GDP per capita of consumption and investment. You must first ensure that both datasets are sorted by their common indicator variables, then use a one to one merge `merge 1:1` according to these variables.

```
. use "G7 less Germany pwt.dta", clear
. so country year
. sa "G7 less Germany pwt.dta", replace
. use "G7 extra data.dta", clear /* "master" data */
. so country year
. merge 1:1 country year using "G7 less Germany pwt.dta" /*"using" data */
. tab _merge /* 1= master, 2= using, 3= both */
```

Stata automatically creates a variable called `_merge` which indicates the results of the merge operation. It is crucial to tabulate this variable to check that the operation worked as you intended. The variable can take on the values:

1 : observations from the master dataset that did not match observations from the using dataset

2 : observations from the using dataset that did not match observations from the master dataset
3 : observations from the both datasets that matched

Ideally, all observations will have a `_merge` value of 3. However, it may be possible, for instance, that the master dataset has observations for extra countries or extra years. If so, then some observations will have a `_merge` value of 1. You should tabulate these to confirm what the extra observations refer to:

```
. tab country if _merge==1
. tab year if _merge==1
. tab _merge
```

<code>_merge</code>	Freq.	Percent	Cum.
1	31	10.95	10.95
3	252	89.05	100.00
Total	283	100.00	

`tab country if _merge==1` would allow you to look at the unmatched observations.

Finally, if you had, for example, data on companies which included the industry they are in, and separately a file with lots of industry information. You may wish to form all pairwise combinations so that any firm has industry characteristics associated with it. The command `merge m:1` or `merge 1:m` does the job. [This used to be a command called `joinby`.]

Collapse

This command converts the data into a dataset of summary statistics, such as sums, means, medians, and so on. One use is when you have monthly data that you want to aggregate to annual data:

```
. collapse (sum) monthpop, by(country year)
```

or firm-level data that you want to aggregate to industry level:

```
. collapse (sum) firmoutput, by(industry year month)
```

`by()` excludes the indicator variable that you are collapsing or summing over (`month` in the first example, `firm` in the second) – it just contains the indicator variables that you want to collapse by. Note that if your dataset contains other variables beside the indicator variables and the variables you are collapsing, they will be erased.

One possible problem that arises in the use of `collapse` is in its treatment of missings. It returns the summary statistic of missing values as zero. If, for example, when using the PWT Afghanistan (“AFG”) contains all missing values for `pop`. If you wanted to aggregate population data over time (for whatever reasons), `collapse` would report aggregate population for Afghanistan as zero, not missing. If, instead, you want aggregate population figures to be missing if any or all of the year data is missing, then use the following coding (the technicalities of it will become clearer later, after you learn how to create dummy variables):

```
. gen missing=(pop==.)
. collapse (sum) pop missing, by(countrygroup)
. replace firmoutput=. if missing>0
. rename pop aggpop
. drop missing
```

Note, if you are running this command on a large dataset, it may be worthwhile to use the `fast` option – this speeds things up skipping the preparation of a backup if the command is aborted by the user pressing `BREAK`, but this is really only useful for when you are working interactively).

Order, aorder, and move

These commands can be used to do some cosmetic changes to the order of your variable list in the variables window, e.g. if you want to have the indicator variables on top of the list. `aorder` alphabetically sorts variables and `order` brings them in a user-specified order:

```
. aorder
. order countrycode year pop
```

If you do not list certain variables after `order`, they will remain where they are. `move` is used if you simply want to swap the position of two variables, e.g. bringing `year` to the top:

```
. move year countrycode
```

Creating new variables

Generate, egen, replace

The two most common commands for creating new variables are `gen` and `egen`. We can create a host of new variables from the existing data with the `gen` command:

```
. gen realgdp=(pop*1000)*cgdg      /* real GDP in current prices */
. gen lpop=ln(pop)                 /* log population */
. gen popsq=pop^2                  /* squared population */
. gen ten=10                       /* constant value of 10 */
. gen id=_n                        /* id number of observation */
. gen total=_N                     /* total number of observations */
. gen byte yr=year-1900            /* 50,51,etc instead of 1950,1951 */
. gen str6 source="PWT6.1"        /* string variable */
. gen largeyear=year if pop>5000 & pop!=.
```

A couple of things to note. First, Stata's default data type is float, so if you want to create a variable in some other format (e.g. byte, string), you need to specify this. Second, missing numeric observations, denoted by a dot, are interpreted by Stata as a very large positive number. You need to pay special attention to such observations when using `if` statements. If the last command above had simply been `gen largeyear=year if pop>5000`, then `largeyear` would have included observations 1950–1959 for AGO, even though data for those years is actually missing.

The `egen` command typically creates new variables based on summary measures, such as sum, mean, min and max:

```
. egen totalpop=sum(pop), by(year) /* world population per year */
. egen avgpob=mean(pop), by(year) /* average country pop per year */
. egen maxpop=max(pop)             /* largest population value */
. egen countpop=count(pop)        /* counts number of non-missing obs */
. egen groupid=group(country_code) /* generates numeric id variable for countries */
```

The `egen` command is also useful if your data is in long format (see below) and you want to do some calculations on different observations, e.g. `year` is long, and you want to find the difference between 1995 and 1998 populations. The following routine will achieve this:

```
. gen temp1=pop if year==1995
. egen temp2=max(temp1), by(country_code)
. gen temp3=pop-temp2 if year==1998
. egen diff=max(temp3), by(country)
. drop temp*
```

Note that both `gen` and `egen` have `sum` options. `egen` generates the total sum, and `gen` creates a cumulative sum. The running cumulation of `gen` depends on the order in which the data is sorted, so use with caution:

```
. egen totpop=sum(pop)             /* sum total of population = single result*/
. gen cumpop=sum(pop)              /* cumulative total of population */
```

As with `collapse`, `egen` has problems with handling missing values. For example, summing up data entries that are all missing yields a total of zero, not missing (see `collapse` below for details and how to solve this problem).

The `replace` command modifies existing variables in exactly the same way as `gen` creates new variables:

```
. gen lpop=ln(pop)
. replace lpop=ln(1) if lpop==.      /* missings now ln(1)=0 */
```

Converting strings to numerics and vice versa

As mentioned before, Stata cannot run any statistical analyses on string variables. If you want to analyse such variables, you must first encode them:

```
. encode country, gen(ctyno)
. codebook ctyno /*Tells you the link with the data*/
```

This creates a new variable `ctyno`, which takes a value of 1 for CAN, 2 for FRA, and so on. The labels are automatically computed, based on the original string values – you can achieve similar results but without the automatic labels using `egen ctyno=group(country)`.

You can go in the other direction and create a string variable from a numerical one, as long as the numeric variable has labels attached to each value:

```
. decode ctyno, gen(ctycode)
```

If you wanted to convert a numeric with no labels, such as `year`, into a string, the command is:

```
. gen str4 yearcode=string(year)
```

And if you have a string variable that only contains numbers, you can convert them to a numeric variable using:

```
. gen yearno=real(yearcode)
```

This last command can be useful if a numeric variable is mistakenly read into Stata as a string. You can confirm the success of each conversion by:

```
. desc country ctyno ctycode year yearcode yearno
```

Combining and dividing variables

You may wish to create a new variable whose data is a combination of the data values of other variables, e.g. joining country code and year to get AGO1950, AGO1951, and so on. To do this, first convert any numeric variables, such as `year`, to string (see earlier), then use the command:

```
. gen str7 ctyyear=country_code+yearcode
```

If you want to create a new numeric combination, first convert the two numeric variables to string, then create a new string variable that combines them, and finally convert this string to a numeric:

```
. gen str4 yearcode=string(year)
. gen str7 popcode=string(pop)
. gen str11 yearpopcode=yearcode+popcode
. gen yearpop=real(yearpopcode)
```

```
. sum yearpopcode yearpop displays the result
```

To divide up a variable or to extract part of a variable to create a new one, use the `substr` function. For example, you may want to reduce the `year` variable to 70, 71, 72, etc. either to reduce file size or to merge with a dataset that has `year` in that format:

```
. gen str2 yr=substr(yearcode,3,2)
```

The first term in parentheses is the string variable that you are extracting from, the second is the position of the first character you want to extract (–X–), and the third term is the number of characters to be extracted (–XX). Alternatively, you can select your starting character by counting from the end (2 positions from the end instead of 3 positions from the start):

```
. gen str2 yr=substr(yearcode,-2,2)
```

Things can get pretty complicated when the string you want to divide isn't as neat as `yearcode` above. For example, suppose you have data on city population and that each observation is identified by a single variable called `code` with values such as "UK London", "UK Birmingham", "UK Cardiff", "Ireland Dublin", "France Paris", "Germany Berlin", "Germany Bonn", and so on. The `code` variable can be broken into `country` and `city` as follows:

```
. gen str10 country=substr(code,1,strpos(code," ")-1)
. gen str10 city=trim(substr(code, strpos(code," "),11))
```

The `strpos()` function gives the position of the second argument in the first argument, so here it tells you what position the blank space takes in the `code` variable. The `country` substring then extracts from the `code` variable, starting at the first character, and extracting a total of $3-1=2$ characters for UK, $8-1=7$ characters for Ireland and so on. The `trim()` function removes any leading or trailing blanks. So, the `city` substring extracts from the `code` variable, starting at the blank space, and extracting a total of 11 characters including the space, which is then trimmed off. Note, the `country` variable could also have been created using `trim()`:

```
. gen str10 country=trim(substr(code,1,strpos(code," ")))
```

Dummy variables

You can use `generate` and `replace` to create a dummy variable as follows:

```
. gen largepop=0
. replace largepop=1 if (pop>=5000 & pop!=.)
```

Or you can combine these in one command:

```
. gen largepop=(pop>=5000 & pop~=.)
```

Note, the parenthesis are not strictly necessary, but can be useful for clarity purposes. You may want to create a set of dummy

variables, for example, one for each country:

```
. tab country, gen(cdum)
```

This creates a dummy variable `cdum1` equal to 1 if the country is "CAN" and zero otherwise, a dummy variable `cdum2` if the country is "FRA" and zero otherwise, and so on up to `cdum7` for "USA". You can refer to this set of dummies in later commands using a wild card, `cdum*`, instead of typing out the entire list.

Lags and leads

To generate lagged population in the G7 dataset:

```
. so countrycode year  
. by countrycode: gen lagpop=pop[_n-1] if year==year[_n-1]+1
```

Processing the statement country-by-country is necessary to prevent data from one country being used as a lag for another, as could happen with the following data:

country	Year	pop
AUS	1996	18312
AUS	1997	18532
AUS	1998	18751
AUT	1950	6928
AUT	1951	6938
AUT	1952	6938

The `if` argument avoids problems when there isn't a full panel of years – if the dataset only has observations for 1950, 1955, 1960-1998, then lags will only be created for 1961 on. A lead can be created in similar fashion:

```
. so country year  
. by country: gen leadpop=pop[_n+1] if year==year[_n+1]-1
```

Cleaning the data

This section covers a few techniques that can be used to fill in gaps in your data.

Fillin and expand

Suppose you start with a dataset that has observations for some years for one country, and a different set of years for another country:

country	Year	pop
AGO	1960	4816
AGO	1961	4884
ARG	1961	20996
ARG	1962	21342

You can "rectangularize" this dataset as follows:

```
. fillin country year
```

This creates new missing observations wherever a county-year combination did not previously exist:

country	Year	pop
AGO	1960	4816
AGO	1961	4884
AGO	1962	.
ARG	1960	.
ARG	1961	20996
ARG	1962	21342

It also creates a variable `_fillin` that shows the results of the operation; 0 signifies an existing observation, and 1 a new one. If no country had data for 1961, then the fillin command would create a dataset like:

country	Year	pop
AGO	1960	4816
AGO	1962	.
ARG	1960	.
ARG	1962	21342

So, to get a proper “rectangle”, you would first have to ensure that at least one observation with `year=1961` exists:

```
. expand 2 if _n==1
. replace year=1961 if _n==_N
. replace pop=. if _n==_N
```

`expand 2` creates 2 observations identical to observation number one (`_n==1`) and places the additional observation at the end of the dataset, i.e. observation number `_N`. As well as recoding the year in this additional observation, it is imperative to replace all other data with missing values – the original dataset has no data for 1961, so the expanded dataset should have missings for 1961. After this has been done, you can now apply the `fillin` command to get a complete “rectangle”.

These operations may be useful if you want to estimate missing values by, for example, extrapolation. Or if you want to replace all missing values with zero or some other amount.

Interpolation and extrapolation

Suppose your population time-series is incomplete – as with some of the countries in the PWT (e.g. STP which is Sao Tome and Principe). You can linearly interpolate missing values using:

```
. so country
. by country: ipolate pop year, gen(ipop)
```

country	Year	pop
STP	1995	132
STP	1996	135.29
STP	1997	.
STP	1998	141.7
STP	1999	144.9
STP	2000	148

Note, first of all, that you need to interpolate by country, otherwise Stata will simply interpolate the entire list of observations irrespective of whether some observations are for one country and some for another. The first variable listed after the `ipolate` command is the variable you actually want to interpolate, the second is the dimension along which you want to interpolate. So, if you believe population varies with time, you can interpolate along the time dimension. You then need to specify a name for a new variable that will contain all the original and interpolated values – here `ipop`. You can use this cleaned-up version in its entirety in subsequent analysis, or you can select values from it to update the original variable, e.g. to clean values for STP only:

```
. replace pop=ipop if country=="STP"
```

Linear extrapolation can be achieved with the same command, adding the `epolate` option, e.g. to extrapolate beyond 2000:

```
. so country
. by country: ipolate pop year, gen(ipop) epolate
```

Note, however, that Stata will fail to interpolate or extrapolate if there are no missing values to start with. No 2001 or 2002 observations actually exist, so Stata will not actually be able to extrapolate beyond 2000. To overcome this, you will first have to create blank observations for 2001 and 2002 using `expand` (alternatively, if these observations exist for other countries, you can rectangularise the dataset using `fillin`).

Panel Data Manipulation: Long versus Wide data sets

Reshape

Datasets may be laid out in wide or long formats. Suppose we keep population data for 1970-75 only:

```
. keep country country_code year pop  
. keep if year<=1975
```

In long format, this looks like:

country	country_code	year	Pop
Canada	CAN	1970	21324
Canada	CAN	1971	21962.1
Canada	CAN	1972	22219.6
Canada	CAN	1973	22493.8
Canada	CAN	1974	22808.4
Canada	CAN	1975	23142.3
France	FRA	1970	52040.8
France	FRA	1971	52531.8
France	FRA	1972	52993.1
France	FRA	1973	53420.5
France	FRA	1974	53771
France	FRA	1975	54016

And the same data in wide format looks like:

countryisocode	pop1970	pop1971	pop1972	pop1973	pop1974	pop1975	country
CAN	21324	21962.1	22219.6	22493.8	22808.4	23142.3	Canada
FRA	52040.8	52531.8	52993.1	53420.5	53771	54016	France
GBR	55632	55928	56097	56223	56236	56226	United Kingdom
GER	77709	78345	78715	78956	78979	78679	Germany
ITA	53821.9	54073.5	54381.3	54751.4	55110.9	55441	Italy
JPN	103720	104750	106180	108660	110160	111520	Japan
USA	205089	207692	209924	211939	213898	215981	United States

The vast majority of Stata commands work best when the data is in long format. In any case, to convert from long to wide:

```
. reshape wide pop, i(country_code) j(year)
```

or from wide to long:

```
. reshape long pop, i(country_code) j(year)
```

The variable(s) immediately behind `long` or `wide` is the one that contains the data we want to reshape (the “data variable”, in our case `pop`). Note that in the `reshape long` case, Stata will reshape all variables that start with the letters you put behind `long`. Here, there are actually six of them (`pop1970-pop1975`, all starting with `pop`). The `i()` specifies the variable(s) whose unique values denote a logical observation in wide format. In our case, this is `country`. It uniquely identifies every data entry in wide format (here: `pop`). The `j()` specifies the variable whose unique values denote a sub-observation, in our case `year`. That is, within every group of countries, `year` uniquely identifies observations. In long format, `i()` and `j()` together completely identify each observation.

If there are more than two indicator variables in wide format, then be careful to include the correct list in `i()`. For example, if there were also an `agegroup` indicator variable, so that `pop` actually referred to population in a given age group, then we could reshape the data from `country / agegroup / year / pop` to `country / agegroup / pop1960 / pop1961 / etc` using:

```
. reshape wide pop, i(country agegroup) j(year)
```

If there is more than one data variable, first drop the variables you are not interested in, and then make sure to include the full list you are interested in reshaping within the command:

```
. reshape wide pop cgdp pi, i(country) j(year)
```

This will create new variables `pop1970-1975`, `cgdp1970-1975` and `pi1970-1975`. Note if you had not dropped all other variables beforehand, you would get an error message. For example, if you had forgotten to delete `cc`:

```
. cc not constant within country  
. Type "reshape error" for a listing of the problem observations.
```

As Stata suggests, “reshape error” will list all observations for which country does not uniquely identify observations in wide format (here, these are actually all observations!). More generally, any variable that varies across both i() and j() variables either needs to be dropped before reshape wide or be included in the data variable list. Intuitively, Stata would not know where to put the data entries of such variables once year has gone as an identifier.

We could also have reshaped the original long data to have the country variable as wide:

```
. reshape wide pop, i(year) j(country) string
```

Note, you need to specify the string option when j() is a string variable. Browsing the resulting data:

year	popCAN	popFRA	popGBR	popGER	popITA	popJPN	popUSA
1970	21324	52040.8	55632	77709	53821.9	103720	205089
1971	21962.1	52531.8	55928	78345	54073.5	104750	207692
1972	22219.6	52993.1	56097	78715	54381.3	106180	209924
1973	22493.8	53420.5	56223	78956	54751.4	108660	211939
1974	22808.4	53771	56236	78979	55110.9	110160	213898
1975	23142.3	54016	56226	78679	55441	111520	215981

To create variables named CANpop / FRApop / GBRpop instead of popCAN/popFRA/popGBR, use:

```
. reshape wide @pop, i(year) j(country) string
```

The @ is useful when, for example, you start with a dataset that has the dimension you want to reshape written the “wrong” way around. Suppose you are given a dataset with country / youngpop / oldpop. You can reshape the pop variable to long to give country / agegroup / pop using:

```
. reshape long @pop, i(country) j(agegroup) string
```

Estimation

We now move on from the manipulation of databases to the more exciting material of running regressions. In this tutorial, we shall use data from Prof. Caselli's "Accounting for Cross-Country Income Differences" which is forthcoming in Handbook of Economic Growth. There is a link to these data on my website. But before we start looking at the basics of regression commands, let us look in Stata help for the details on estimation. The basic information is:

- There are many different models of estimation. The main commands include `regress`, `logit`, `logistic`, `sureg`.
- Most have a similar syntax:

command varlist [weight] [if exp] [in range] [, options]

- 1st variable in the varlist is the dependent variable, and the remaining are the independent variables.
- You can use Stata's syntax to specify the estimation sample; you do not have to make a special dataset.
- You can, at any time, review the last estimates by typing the estimation command without arguments.
- The `level()` option to indicate the width of the confidence interval. The default is `level(95)`.

Once you have carried out your estimation, there are a number of post-estimation commands that are useful:

- You can recall the estimates, VCM, standard errors, etc...;
- You can carry out hypothesis testing => `test` (Wald tests), `testnl` (non-linear Wald tests), `lrtest` (likelihood-ratio tests), `hausman` (Hausman's specification test);
- You can use Stata's `predict` command, which does predictions and residual calculations.

The rest of the estimation notes will use Prof. Francesco Caselli's database from his handbook of Economic Growth paper which contains real GDP and other variables on up to 105 countries. This data is downloadable in .dta form from his website or using a link on the course webpage. Download this data to the relevant directory on your computer and save it as you wish to call it. I have called it "Caselli_handbook.dta".

Linear regression

Stata can do most of fancy regressions (and most of which we will not talk about in these classes). Just so that you know the main ones, here is an abbreviated list of other regression commands that may be of interest:

<code>anova</code>	analysis of variance and covariance
<code>cnreg</code>	censored-normal regression
<code>heckman</code>	Heckman selection model
<code>intreg</code>	interval regression
<code>ivreg</code>	instrumental variables (2SLS) regression
<code>newey</code>	regression with Newey-West standard errors
<code>prais</code>	Prais-Winsten, Cochrane-Orcutt, or Hildreth-Lu regression
<code>qreg</code>	quantile (including median) regression
<code>reg</code>	ordinary least squares regression
<code>reg3</code>	three-stage least squares regression
<code>rreg</code>	robust regression (NOT robust standard errors)
<code>sureg</code>	seemingly unrelated regression
<code>svyheckman</code>	Heckman selection model with survey data
<code>svyintreg</code>	interval regression with survey data
<code>svyivreg</code>	instrumental variables regression with survey data
<code>svyregress</code>	linear regression with survey data
<code>tobit</code>	tobit regression
<code>treatreg</code>	treatment effects model
<code>truncreg</code>	truncated regression
<code>xtabond</code>	Arellano-Bond linear, dynamic panel-data estimator
<code>xtintreg</code>	panel data interval regression models
<code>xtreg</code>	fixed- and random-effects linear models
<code>xtregar</code>	fixed- and random-effects linear models with an AR(1) disturbance
<code>xttobit</code>	panel data tobit models

We will focus on this is the most basic form of linear regression. *regress* fits a model of *depvar* on *varlist* using linear regression. The *help regress* command will bring up the following instructions for using *regress*.

regress *depvar* [*varlist*] [*weight*] [*if exp*] [*in range*] [, *level*(#) *beta* *robust* *cluster*(*varname*) *score*(*newvar*) *hc2* *hc3* *hascons* *noconstant* *tsscons* *noheader* *eform*(*string*) *depname*(*varname*) *mse1* *plus*]

Looking in the bottom of this help file will explain the options as follows:

Options

level(#) specifies the confidence level, in %, for confidence intervals of the coefficients; see help level.

beta requests that normalized beta coefficients be reported instead of confidence intervals. *beta* may not be specified with *cluster*().

robust specifies that the Huber/White/sandwich estimator of variance is to be used in place of the traditional calculation. *robust* combined with *cluster*() further allows observations which are not independent within cluster (although they must be independent between clusters). See [U] 23.14 Obtaining robust variance estimates.

cluster(*varname*) specifies that the observations are independent across groups (clusters) but not necessarily independent within groups. *varname* specifies to which group each observation belongs; e.g., *cluster*(*personid*) in data with repeated observations on individuals. *cluster*() can be used with *pweights* to produce estimates for unstratified cluster-sampled data, but see help *svyregress* for a command especially designed for survey data. Specifying *cluster*() implies *robust*.

score(*newvar*) creates a new variable for the scores from the equation in the model. The new variable contains each observation's contribution to the score; see [U] 23.15 Obtaining scores.

hc2 and **hc3** specify an alternative bias correction for the robust variance calculation. *hc2* and *hc3* may not be specified with *cluster*(). *hc2* uses $u_j^2/(1-h_j)$ as the observation's variance estimate. *hc3* uses $u_j^2/(1-h_j)^2$ as the observation's variance estimate. Specifying either *hc2* or *hc3* implies *robust*.

Hascons indicates that a user-defined constant or its equivalent is specified among the independent variables. Some caution is recommended when using this option as resulting estimates may not be as accurate as they otherwise would be. Use of this option requires "sweeping" the constant last, so the moment matrix must be accumulated in absolute rather than deviation form. This option may be safely specified when the means of the dependent and independent variables are all "reasonable" and there are not large amounts of collinearity between the independent variables. The best procedure is to view *hascons* as a reporting option -- estimate with and without *hascons* and verify that the coefficients and standard errors of the variables not affected by the identity of the constant are unchanged. If you do not understand this warning, it is best to avoid this option.

noconstant suppresses the constant term (intercept) in the regression.

tsscons forces the total sum of squares to be computed as though the model has a constant; i.e., as deviations from the mean of the dependent variable. This is a rarely used option that has an effect only when specified with *nocons*. It affects only the total sum of squares and all results derived from the total sum of squares.

noheader, **eform**(), **depname**(), **mse1**, and **plus** are for ado-file writers; see [R] *regress*.

As described above, most estimation commands will follow this type of syntax but the available options will differ and so you should check the relevant help files if you wish to use these approaches. Of course, Stata has a number of defaults and so you don't need to include any options if you don't wish to change the default (though it is always good to figure out what the default is!)

Lets start with a very simple regression of GDP per worker (*y*) on capital-output ratio (*k*).

```
. regress y k
```

Source	SS	df	MS			
Model	2.5465e+10	1	2.5465e+10	Number of obs =	104	
Residual	2.3380e+09	102	22921482.3	F(1, 102) =	1110.99	
				Prob > F =	0.0000	
				R-squared =	0.9159	
				Adj R-squared =	0.9151	
Total	2.7803e+10	103	269936187	Root MSE =	4787.6	

y	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
k	.3319374	.0099587	33.33	0.000	.3121844	.3516904
_cons	4720.016	617.1018	7.65	0.000	3495.998	5944.035

There are a few points to note here:

- The first variable listed after the `regress` (or `reg` for short) command is the dependent variable, and all subsequently listed variables are the independent variables.
- Stata automatically adds the constant term or intercept to the list of independent variables (use the `noconstant` option if you want to exclude it).
- The top-left corner gives the ANOVA decomposition of the sum of squares in the dependent variable (Total) into the explained (Model) and unexplained (Residual).
- The top-right corner gives the statistical significance results for the model as a whole.
- The bottom section gives the results for the individual independent variables.

The `regress` command can be used with the `robust` option for estimating the standard errors using the Huber-White sandwich estimator (to correct the standard errors for heteroscedasticity):

```
. regress y k, robust
```

Regression with robust standard errors						
				Number of obs =	104	
				F(1, 102) =	702.15	
				Prob > F =	0.0000	
				R-squared =	0.9159	
				Root MSE =	4787.6	

y	Coef.	Robust Std. Err.	t	P> t	[95% Conf. Interval]	
k	.3319374	.0125268	26.50	0.000	.3070905	.3567842
_cons	4720.016	506.2807	9.32	0.000	3715.811	5724.222

The coefficient estimates are exactly the same as in straightforward OLS, but the standard errors take into account heteroscedasticity. Note, the ANOVA table is deliberately suppressed as it is no longer appropriate in a statistical sense.

Sometimes you also want to allow for more general deviations from the iid-assumption on the error term. The option `cluster(group)` allows for arbitrary correlation within specified groups (see Wooldridge, “Econometrics of Cross-Section and Panel Data”, chapter 4, for more details and limitations of this approach). For example, you might think that in a panel of countries, errors are correlated across time but independent across countries. Then, you should cluster standard errors on countries. In our example, we do not have a time dimension so clustering on country yields the same results as the robust option (which is a special case of the cluster option):

```
. regress y k, cluster(country)
```

Stata comes with a large amount of regression diagnostic tools, such as tests for outliers, heteroskedasticity in the errors etc. A good survey is available at <http://www.ats.ucla.edu/stat/stata/webbooks/reg/chapter2/statareg2.htm>. We will focus on two useful tools for detecting influential observations and looking at partial correlations. The first tool is the command `lvr2plot` (read leverage-versus-residual squared plot). This is not available after the `robust` option is used so lets revert back to the original regression:

```
. regress y k
. lvr2plot, mlabel(country)
```

This plots the leverages of all observations against their squared residuals (the option `mlabel` labels points according to the variable listed in brackets behind it). Leverage tells you how large the influence of a single observation on the estimated coefficients is. Observations with high values could potentially be driving the results obtained (especially if they also have a large squared residual) so we should check whether excluding them changes anything.

The second command is `avplot` (added-variable plot) which graphs the partial correlation between a specified regressor and the dependent variable. For this not to be simply the fitted values, we should add another variable such as human capital (`h`). Formally

```
. regress y k h
. avplot k, mlabel(country)
```

For some very basic econometrics which also comes with the necessary Stata commands, see http://www.cas.lancs.ac.uk/short_courses/notes/stata/session5.pdf for model diagnostics.

Post-estimation

Once you have done your regression, you usually want to carry out some extra analysis such as forecasting or hypothesis testing. Here is a list of the most useful post-estimation commands:

Command	Description
<code>adjust</code>	Tables of adjusted means and proportions
<code>estimates</code>	Store, replay, display, ... estimation results
<code>hausman</code>	Hausman's specification test after model fitting
<code>lincom</code>	Obtain linear combinations of coefficients
<code>linktest</code>	Specification link test for single-equation models
<code>lrtest</code>	Likelihood-ratio test after model fitting
<code>mfx</code>	Marginal effects or elasticities after estimation
<code>nlcom</code>	Nonlinear combinations of estimators
<code>predict</code>	Obtain predictions, residuals, etc. after estimation
<code>predictnl</code>	Nonlinear predictions after estimation
<code>suest</code>	Perform seemingly unrelated estimation
<code>test</code>	Test linear hypotheses after estimation
<code>testnl</code>	Test nonlinear hypotheses after estimation
<code>vce</code>	Display covariance matrix of the estimators

Prediction

A number of predicted values can be obtained after all estimation commands, such as `reg`, `cnsreg`, `logit` or `probit`. The most important are the predicted values for the dependent variable and the predicted residuals. For example, suppose we run the basic regression again:

```
. regress y k h
. predict y_hat          /* predicted values for dependent var */
. predict r, residual    /* predicted residuals */
```

Stata creates new variables containing the predicted values, and these variables can then be used in any other Stata command, e.g. you can graph a histogram of the residuals to check for normality.

If we run a selected regression (e.g. just using OECD countries) and then wish to know how well this regression fits, we could run the following commands:

```
regress y k h if oecd==1

predict y_hat_oecd if oecd==1
predict r_oecd     if oecd==1, residual
```

The `if` statements are only necessary if you are running the analysis on a subset of dataset currently loaded into Stata. If you want to make out-of-sample predictions, just drop the `if` statements in the `predict` commands.

```
predict y_hat_oecd_full
predict r_oecd_full, residual
```

Hypothesis testing

The results of each estimation automatically include for each independent variable a t-test (for linear regressions) and a z-test (for regressions such as logit or probit) on the null hypothesis that the “true” coefficient is equal to zero. You can also perform an F-test or χ^2 test on this hypothesis using the `test` command:

```
. regress y k h y1985 ya
. test y1985 /*since Stata defaults to comparing the listed terms to zero, you can simply use the variable*/

( 1) y1985 = 0

      F( 1,    63) =   15.80
      Prob > F =   0.0002
```

The F-statistic with 1 numerator and 63 denominator degrees of freedom is 15.80. The p -value or significance level of the test is basically zero (up to 4 digits at least), so we can reject the null hypothesis even at the 1% level – `y1985` is significantly different from zero. Notice that, since the F -distribution with 1 numerator degree of freedom is identical to the t -distribution, so the F-test result is the same as the square of the t-test result in the regression. Also the p -values associated with each test agree.

You can perform any test on linear hypotheses about the coefficients, such as:

```
. test y1985=0.5          /* test coefficient on y1985 equals 0.5 */
. test y1985 h           /* test coefficients on y1985 & h jointly zero */
. test y1985+h=-0.5     /* test coefficients on y1985 & h sum to -0.5 */
. test y1985=h          /* test coefficients on y1985 & h are the same */
```

With many Stata commands, you can refer to a list of variables using a hyphen, e.g. `desc k- ya` gives descriptive statistics on `exp`, `ya` and every other variable on the list between them. However, the `test` command interprets the hyphen as a minus, and gets confused because it thinks you are typing a formula for it to test. If you want to test a long list of variables, you can use the `testparm` command (but remember to use the `order` command to bring the variables in the right order first)

```
. order k h y1985 ya

. testparm k-ya

( 1) k = 0
( 2) h = 0
( 3) y1985 = 0
( 4) ya = 0

      F( 4,    63) =   370.75
      Prob > F =    0.0000
```


Extracting results

We have already seen how the `predict` command can be used to extract predicted values from Stata's internal memory for use in subsequent analyses. Using the `generate` command, we can also extract other results following a regression, such as estimated coefficients and standard errors:

```
regress y k h y1985 ya, robust
gen b_cons=_b[_cons]      /* beta coefficient on constant term */
gen b_k=_b[k]             /* beta coefficient on GDP60 variable */
gen se_k=_se[k]          /* standard error */
```

You can `tabulate` the new variables to confirm that they do indeed contain the results of the regression. You can then use these new variables in subsequent Stata commands, e.g. to create a variable containing t-statistics:

```
. gen t_k=b_k/se_k
or, more directly:
. gen t_k=_b[k]/_se[k]
```

Stata stores extra results from estimation commands in `e()`, and you can see a list of what exactly is stored using the `ereturn list` command:

```
. regress y k h y1985 ya, robust
. ereturn list

. ereturn list

scalars:
      e(N) = 68
    e(df_m) = 4
    e(df_r) = 63
      e(F) = 273.7198124833108
    e(r2) = .9592493796249692
  e(rmse) = 3451.985251440704
    e(mss) = 17671593578.3502
    e(rss) = 750720737.0983406
  e(r2_a) = .9566620386487768
    e(ll) = -647.8670640006279
  e(ll_0) = -756.6767273270843

macros:
    e(depvar) : "y"
    e(cmd) : "regress"
  e(predict) : "regres_p"
    e(model) : "ols"
  e(vcetype) : "Robust"
```

```
matrices:
      e(b) : 1 x 5
      e(V) : 5 x 5
```

```
functions:
    e(sample)
```

`e(sample)` is a useful tool to have. Earlier we ran the following commands:

```
regress y k h if oecd==1
predict y_hat_oecd if oecd==1
```

but in the event that the “if” statement is complex, we may wish to simply tell Stata to predict using the same sample it used in the regression. We can do this using the `e(sample)`:

```
predict y_hat_oecd if e(sample)
```

`e(N)` stores the number of observations, `e(df_m)` the model degrees of freedom, `e(df_r)` the residual degrees of freedom,

`e(F)` the F-statistic, and so on. You can extract any of these into a new variable:

```
. gen residualdf=e(df_r)
```

And you can then use this variable as usual, e.g. to generate p-values:

```
. gen p_k=tprob(residualdf,t_k)
```

The `tprob` function uses the two-tailed cumulative Student's t-distribution. The first argument in parenthesis is the relevant degrees of freedom, the second is the t-statistic.

In fact, most Stata commands – not just estimation commands – store results in internal memory, ready for possible extraction. Generally, the results from other commands are stored in `r()`. You can see a list of what exactly is stored using the `return list` command, and you can extract any you wish into new variables:

```
. sum y
```

Variable	Obs	Mean	Std. Dev.	Min	Max
y	105	18103.09	16354.09	630.1393	57259.25

```
. return list
```

scalars:

```
      r(N) = 105
r(sum_w) = 105
r(mean)  = 18103.08932466053
r(Var)   = 267456251.2136306
r(sd)    = 16354.08973968379
r(min)   = 630.1392822265625
r(max)   = 57259.25
r(sum)   = 1900824.379089356
```

```
. gen mean_y=r(mean)
```

Note that the last command will give exactly the same results as `egen mean_y=mean(y)`.

OUTREG2 – the ultimate tool in Stata/Latex or Word friendliness?

There is a tool which will automatically create excel, word or latex tables or regression results and it will save you loads of time and effort. It formats the tables to a journal standard and was originally just for word (outreg) but now the updated version will also do tables for latex also.

However, it does not come as a standard tool and so before we can use it, we must learn how to install extra ado files (not to be confused with running our own do files).

Extra commands on the net

Looking for specific commands

If you are trying to perform an exotic econometric technique and cannot find any useful command in the Stata manuals, you may have to programme in the details yourself. However, before making such a rash move, you should be aware that, in addition to the huge list of commands available in the Stata package and listed in the Stata manuals, a number of researchers have created their own extra commands. These extra commands range from the aforementioned exotic econometric techniques to mini time-saving routines. For example, the command `outreg`.

You need to first locate the relevant command and then install it into your copy of Stata. The command can be located by trying different searches, e.g. to search for a command that formats the layout of regression results, I might search for words like “format” or “table”:

```
. search format regression table
```

Keyword search

Keywords: format regression table

Search: (1) Official help files, FAQs, Examples, SJs, and STBs

Search of official help files, FAQs, Examples, SJs, and STBs

FAQ Can I make regression tables that look like those in journal articles?
..... UCLA Academic Technology Services
5/01 <http://www.ats.ucla.edu/stat/stata/faq/outreg.htm>

STB-59 sg97.3 Update to formatting regression output
(help outreg if installed) J. L. Gallup
1/01 p.23; STB Reprints Vol 10, p.143
small bug fixes

STB-58 sg97.2 Update to formatting regression output
(help outreg if installed) J. L. Gallup
11/00 pp.9--13; STB Reprints Vol 10, pp.137--143
update allowing user-specified statistics and notes, 10%
asterisks, table and column titles, scientific notation for
coefficient estimates, and reporting of confidence interval
and marginal effects

STB-49 sg97.1 Revision of outreg
(help outreg if installed) J. L. Gallup
5/99 p.23; STB Reprints Vol 9, pp.170--171
updated for Stata 6 and improved

STB-46 sg97 Formatting regression output for published tables
(help outreg if installed) J. L. Gallup
11/98 pp.28--30; STB Reprints Vol 8, pp.200--202
takes output from any estimation command and formats it as
in journal articles

(end of search)

You can read the FAQ by clicking on the blue hyperlink. This gives some information on the command. You can install the command by first clicking on the blue command name (here `sg97.3`, the most up-to-date version) and, when the pop-up window appears, clicking on the install hyperlink. Once installed, you can create your table and then use the command `outreg` as any other command in Stata. The help file will tell you the syntax.

However, I mentioned `outreg2` and this has not appeared here, so I may need to update more.

But we know that `outreg2` exists so how do we find it to install? Well, type `outreg2` into google to convince yourself that it exists. Then type:

```
search outreg2, net
```

Web resources from Stata and other users
(contacting <http://www.stata.com>)

1 package found (Stata Journal and STB listed first)

```
-----  
outreg2 from http://fmwww.bc.edu/RePEc/bocode/o  
'OUTREG2': module to arrange regression outputs into an illustrative table  
/ outreg2 provides a fast and easy way to produce an illustrative / table  
of regression outputs. The regression outputs are produced / piecemeal and  
are difficult to compare without some type of / rearrangement. outreg2
```

(click here to return to the previous screen)

(end of search)

Click on the blue link and follow instructions to install the ado file and help.

Constrained linear regression

Suppose the theory predicts that certain the coefficients should be identical. We can estimate a regression model where we constrain the coefficients to be equal to each other. To do this, first define a constraint and then run the `cnsreg` command:

```
. constraint define 1 rev=assass /* constraint is given the number 1 */
. cnsreg gr6085 lgdp60 sec60 prim60 gcy rev assass pi60 if year==1990, constraint(1)
```

```
Constrained linear regression                                Number of obs =      100
                                                           F( 6, 93) = 12.60
                                                           Prob > F      = 0.0000
                                                           Root MSE     = 1.5025
```

```
( 1) - rev + assass = 0
```

gr6085	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
lgdp60	-1.617205	.2840461	-5.69	0.000	-2.181264 -1.053146
sec60	.0429134	.0122297	3.49	0.001	.0184939 .0673329
prim60	.0352023	.007042	5.00	0.000	.0212183 .0491864
gcy	-.0231969	.017786	-1.30	0.195	-.0585165 .0121226
rev	-.2335536	.2877334	-0.81	0.419	-.804935 .3378279
assass	-.2335536	.2877334	-0.81	0.419	-.804935 .3378279
pi60	-.0054616	.0024692	-2.21	0.029	-.0103649 -.0005584
_cons	12.0264	2.073177	5.80	0.000	7.909484 16.14332

Notice that the coefficients for `REV` and `ASSASS` are now identical, along with their standard errors, t-stats, etc. We can define and apply several constraints at once, e.g. constrain the `LGDP60` coefficient to equal `-1.5`:

```
. constraint define 2 lgdp60=-1.5
. cnsreg gr6085 lgdp60 sec60 prim60 gcy rev assass pi60 if year==1990, constraint(1 2)
```

Dichotomous dependent variable

When the dependent variable is dichotomous (zero/one), you can run a Linear Probability Model using the `regress` command. You may also want to run a `logit` or a `probit` regression. The difference between these three models is the assumption that you make about the probability distribution of the latent dependent variable (LPM assumes an identity function, Logit a logistic distribution function, and Probit a normal distribution function).

For the sake of trying out these commands, let us “explain” why a country is an OECD member using a logit regressions:

```
. logit OECD lrgdpl if year==1990
```

```
Iteration 0: log likelihood = -63.180951
Iteration 1: log likelihood = -37.103818
...
Iteration 6: log likelihood = -21.991401
Iteration 7: log likelihood = -21.99139
```

```
Logit estimates                                Number of obs =      135
                                                LR chi2(1)      =       82.38
                                                Prob > chi2     =       0.0000
Log likelihood = -21.99139                    Pseudo R2      =       0.6519
```

OECD	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
lrgdpl	4.94118	1.119976	4.41	0.000	2.746067 7.136292
_cons	-47.38448	10.7335	-4.41	0.000	-68.42176 -26.3472

Panel Data

If you are lucky enough to have a panel dataset, you will have data on n countries/people/firms/etc, over t time periods, for a total of $n \times t$ observations. If t is the same for each country/person/firm then the panel is said to be balanced; but for most things Stata is capable of working out the optimal/maximum dataset available. There are a few things to note before using panel data commands:

1. Panel data should be kept in long form (with separate person and time variables). However, sometimes your data may be in wide form and needs to be converted to long form using the `reshape` command (see class 2).
2. You have to declare your data a panel. One way to do this is using the command `iis` and `tss`. Another is using the `tsset` command. To do this, you need two indicator variables, indicating the unit (`iss`) and time (`tss`) dimensions of your panel. In our case, these are simply `year` and `country`. Note that panel dimensions cannot be string variables so you should first encode `country` (see last week). Once you have done this, use the `tsset` command:

```
. encode country, gen(country_no)
. tsset country_no year
```

You are now free to use Stata's panel data commands, although I will only make use of a few main ones (**bolded**):

xtdes	Describe pattern of xt data
xtsum	Summarize xt data
xttab	Tabulate xt data
xtdata	Faster specification searches with xt data
xtline	Line plots with xt data
xtreg	Fixed-, between- and random-effects, and population-averaged linear models
xtregar	Fixed- and random-effects linear models with an AR(1) disturbance
xtgls	Panel-data models using GLS
xtpcse	OLS or Prais-Winsten models with panel-corrected standard errors
xtcrhh	Hildreth-Houck random coefficients models
xtivreg	Instrumental variables and two-stage least squares for panel-data models
xtabond	Arellano-Bond linear, dynamic panel data estimator
xttobit	Random-effects tobit models
xtintreg	Random-effects interval data regression models
xtlogit	Fixed-effects, random-effects, & population-averaged logit models
xtprobit	Random-effects and population-averaged probit models
xtcloglog	Random-effects and population-averaged cloglog models
xtpoisson	Fixed-effects, random-effects, & population-averaged Poisson models
xtnbreg	Fixed-effects, random-effects, & population-averaged negative binomial models
xtgee	Population-averaged panel-data models using GEE

Describe pattern of xt data

`xtdes` is very useful to see if your panel is actually balanced or whether there is large variation in the number of years for which each cross-sectional unit is reporting.

```
. xtdes

country_no: 1, 2, ..., 168          n =          168
   year: 1950, 1951, ..., 2000      T =           51
      Delta(year) = 1; (2000-1950)+1 = 51
      (country_no*year uniquely identifies each observation)

Distribution of T_i:   min      5%      25%      50%      75%      95%      max
                    51       51       51       51       51       51       51
```


G7	Overall		Between		Within
	Freq.	Percent	Freq.	Percent	Percent
0	8211	95.83	161	95.83	100.00
1	357	4.17	7	4.17	100.00
Total	8568	100.00	168	100.00	100.00

(n = 168)

Panel regressions

xtreg is a generalisation of the regress commands. As with the summary data above, we can make use of the information in the cross-section (between) and also in the time-series (within). Also, as per your econometrics training, Stata allows you to run fixed-effects (fe), random effects (re) and between estimators using xtreg. More complicated estimation (such as Arellano-Bond) have specific xt estimation commands.

Fixed Effects Regression

Fixed effects regression controls for unobserved, but constant, variation across the cross-sectional units. It is equivalent to including a dummy for each country/firm in our regression. Let us use the xtreg command with the fe option:

```
. xtreg grgdpch gdp60 openk kc kg ki, fe
```

```
Fixed-effects (within) regression      Number of obs      =      5067
Group variable (i): country_no        Number of groups   =      112

R-sq:  within = 0.0164                  Obs per group: min =      2
      between = 0.2946                  avg =      45.2
      overall = 0.0306                  max =      51

F(4,4951)                             =      20.58
corr(u_i, Xb) = -0.4277                 Prob > F           =      0.0000
```

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	(dropped)				
openk	-.0107672	.0042079	-2.56	0.011	-.0190166 - .0025178
kc	-.0309774	.0089545	-3.46	0.001	-.0485322 - .0134225
kg	-.0733306	.0147568	-4.97	0.000	-.1022604 - .0444007
ki	.1274592	.0178551	7.14	0.000	.0924552 .1624631
_cons	4.425707	.7451246	5.94	0.000	2.964933 5.886482
sigma_u	1.6055981				
sigma_e	6.4365409				
rho	.05858034	(fraction of variance due to u_i)			

```
F test that all u_i=0:      F(111, 4951) =      1.82      Prob > F = 0.0000
```

Notice that gdp60, the log of GDP in 1960 for each country, is now dropped as it is constant across time for each country and so is subsumed by the country fixed-effect.

Between Effects

We can now use the xtreg command with the be option. This is equivalent to running a regression on the dataset of means by cross-sectional identifier. As this results in loss of information, between effects are not used much in practice.

```
. xtreg grgdpch gdp60 openk kc kg ki, be
```

```
Between regression (regression on group means)  Number of obs      =      5067
Group variable (i): country_no                  Number of groups   =      112

R-sq:  within = 0.0100                  Obs per group: min =      2
      between = 0.4575                  avg =      45.2
      overall = 0.0370                  max =      51

F(5,106)                                       =      17.88
```

sd(u_i + avg(e_i.))= 1.277099 Prob > F = 0.0000

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	-.5185608	.1776192	-2.92	0.004	-.8707083 -.1664134
openk	.0008808	.0029935	0.29	0.769	-.0050541 .0068156
kc	-.0151328	.009457	-1.60	0.113	-.0338822 .0036166
kg	-.0268036	.0149667	-1.79	0.076	-.0564765 .0028693
ki	.1419786	.0213923	6.64	0.000	.0995662 .184391
_cons	4.657591	1.587533	2.93	0.004	1.510153 7.80503

Random Effects

The command for a linear regression on panel data with random effects in Stata is `xtreg` with the `re` option. Stata's random-effects estimator is a weighted average of fixed and between effects.

```
. xtreg grgdpch gdp60 openk kc kg ki, re
```

```
Random-effects GLS regression                      Number of obs                      =                      5067
Group variable (i): country_no                      Number of groups                      =                      112

R-sq:    within = 0.0143                              Obs per group: min =                      2
          between = 0.4235                                                              avg =                      45.2
          overall = 0.0389                                                              max =                      51

Random effects u_i ~ Gaussian                      Wald chi2(5)                      =                      159.55
corr(u_i, X)                      = 0 (assumed)                      Prob > chi2                      =                      0.0000
```

grgdpch	Coef.	Std. Err.	z	P> z	[95% Conf. Interval]
gdp60	-.5661554	.1555741	-3.64	0.000	-.8710751 -.2612356
openk	-.0012826	.0024141	-0.53	0.595	-.0060142 .003449
kc	-.0270849	.0061971	-4.37	0.000	-.039231 -.0149388
kg	-.0506839	.0101051	-5.02	0.000	-.0704895 -.0308783
ki	.1160396	.0127721	9.09	0.000	.0910067 .1410725
_cons	6.866742	1.239024	5.54	0.000	4.4383 9.295185
sigma_u	.73122048				
sigma_e	6.4365409				
rho	.01274156	(fraction of variance due to u_i)			

Choosing Between Fixed and Random Effects

Choosing between FE and RE models is usually done using a Hausman test, and this is easily completed in Stata using the `Hausman` command. To run a Hausman test we need to run the RE and FE models and save the results using the `store` command. We then instruct Stata to retrieve the 2 sets of results and carry-out the test.

For example, using the same estimates as above, we can write the following in our do file:

```
xtreg grgdpch gdp60 openk kc kg ki, fe
estimates store fe
```

```
xtreg grgdpch gdp60 openk kc kg ki, re
estimates store re
```


hausman fe re

	---- Coefficients ----			
	(b)	(B)	(b-B)	sqrt(diag(V_b-V_B))
	fe	re	Difference	S.E.
openk	-.0107672	-.0012826	-.0094846	.0034465
kc	-.0309774	-.0270849	-.0038924	.0064637
kg	-.0733306	-.0506839	-.0226467	.0107541
ki	.1274592	.1160396	.0114196	.0124771

b = consistent under Ho and Ha; obtained from xtreg
B = inconsistent under Ha, efficient under Ho; obtained from xtreg

Test: Ho: difference in coefficients not systematic

chi2(4) = (b-B)'[(V_b-V_B)^(-1)](b-B)
= 20.15
Prob>chi2 = 0.0005

As described in the results, the null hypothesis is that there is no difference in the coefficients estimated by the efficient RE estimator and the consistent FE estimator. If there is no difference, then use the RE estimator – i.e. if the P-value is insignificant and the Prob>chi2 larger than .05. Otherwise, you should use FE, or one of the other solutions for unobserved heterogeneity.

Time series data

Stata has a very particular set of functions that control time series commands. But in order to use these commands, you must ensure that you tell Stata. As with the panel data commands above, we can do this using the `tsset` command – data must be sorted by the time series (or with panel data, by the panel data variable and then the date variable). For example:

```
sort datevar
```

```
tsset datevar
```

OR

```
sort panelvar datevar
```

```
tsset panelvar datevar
```

Once you have done this, you are free to use the time series commands – I present a selection of these below (type `help time` for the full list):

```
tsset      Declare a dataset to be time-series data
tsfill     Fill in missing times with missing observations in time-series data
tsappend   Add observations to a time-series dataset
tsreport   Report time-series aspects of a dataset or estimation sample

arima      Autoregressive integrated moving-average models
arch       Autoregressive conditional heteroskedasticity (ARCH) family of estimators

tssmooth_ma    Moving-average filter
tssmooth_nl    Nonlinear filter

corrgram      Tabulate and graph autocorrelations
xcorr         Cross-correlogram for bivariate time series
dfuller       Augmented Dickey-Fuller unit-root test
pperron       Phillips-Perron unit-roots test
archlm        Engle's LM test for the presence of autoregressive conditional heteroskedasticity

var           Vector autoregression models
svar          Structural vector autoregression models
varbasic      Fit a simple VAR and graph impulse-response functions
vec           Vector error-correction models

varsoc        Obtain lag-order selection statistics for VARs and VECMs
varstable     Check the stability condition of VAR or SVAR estimates
vecrank       Estimate the cointegrating rank using Johansen's framework

irf create    Obtain impulse-response functions and FEVDs
vargranger    Perform pairwise Granger causality tests after var or svar

irf graph     Graph impulse-response functions and FEVDs
irf cgraph    Combine graphs of impulse-response functions and FEVDs
irf ograph    Graph overlaid impulse-response functions and FEVDs
```

All of these can be implemented where appropriate by using the help function, manuals and internet resources (or colleagues know-how).

Stata Date and Time-series Variables

However, one of the issues with time series in Stata, and something that particularly challenges new users of Stata, is the data format used in the program. Therefore, I below provide some more advanced notes on this specialist topic.

The key thing is that there are 2 possible types of entry – date entries (which work in general for storing dates in Stata) and time-series entries (which are useful when we are not using daily data). Stata stores dates as the number of elapsed periods since January 1, 1960. When using a data-set that is not daily data, we want to use Stata's time-series function rather than the date function – the reason is that the dates for quarterly data will be about 3 months apart but the number of days between them will

vary so telling Stata to go from Q1 to Q2 will involve changing the date from (for example) January 1st to April 1st – which is either 90 days or 91 days depending on whether it is a leap-year. Obviously our life would be easier if we could just tell Stata that one entry is Q1, and the other entry is Q2. For example, if we want to take first-differences between quarters, or even more tricky if we wanted to take seasonal differences – Q1 minus Q1 from previous year.

Therefore when we have a variable that identifies the time-series elements of a dataset, we must tell Stata what type of data we are using – is it daily, weekly, monthly, quarterly, half-yearly or yearly. Therefore, if you use daily data it will be the number of elapsed days since January 1st 1960 (which is therefore zero), but if you use quarterly data, it is the number of elapsed quarters since 1960 Q1. The following table explains the different formats –:

There is a format for each of these time periods:

Format	Description	Beginning	+1 Unit	+2 Units	+3 Units
%td	daily	01jan1960	02jan1960	03Jan1960	04Jan1960
%tw	weekly	week 1, 1960	week 2, 1960	week 3, 1960	week 4, 1960
%tm	monthly	Jan, 1960	Feb, 1960	Mar, 1960	Apr, 1960
%tq	quarterly	1st qtr, 1960	2nd qtr, 1960	3rd qtr, 1960	4th qtr, 1961
%th	half-yearly	1st half, 1960	2nd half, 1960	1st half, 1961	2nd half, 1961
%ty	yearly	1960	1961	1962	1963

Obviously, what you tell Stata here is highly important; we will see how to convert our data into Stata dates in a moment, but for now assume that we have a Stata date for January 1, 1999 – this is an elapsed date of 14245 (the number of days since January 1st 1960). If we were to use this number as different types of time-series data, then there would be very different outcomes as shown in the following table:

Daily	Weekly	Quarterly	Half-yearly	Yearly
%td	%tw	%tq	%th	%ty
01 Jan 1999	2233 W50	5521 Q2	9082 H2	-

These dates are so different because the elapsed date is actually the number of weeks, quarters, etc., from the first week, quarter, etc of 1960. The value for %ty is missing because it would be equal to the year 14,245 which is beyond what Stata can accept.

Therefore if we have a date format of 14245, but we want this to point to quarterly data, then we would need to convert it using special Stata functions. These functions translate from %td dates:

wofd(varname)	daily to weekly
mofd(varname)	daily to monthly
qofd(varname)	daily to quarterly
yofd(varname)	daily to yearly

Looking up in help can also show how to convert numbers between other formats.

Getting dates into Stata format

This section covers how we get an existing date or time variable into the Stata format for dates – from here we can rewrite it as quarterly, monthly, etc... using the above commands. There are 3 different considerations depending on how your existing “date variable” is set up:

1. Date functions for single string variables

For example, your existing date variable is called raw_date and is of the form “20mar1999” – then it is said to be a single string variable (the string must be easily separated into its components so strings like “20mar1999” and “March 20, 1999” are acceptable). If you have a string like “200399”, we would need to convert it to a numeric variable first and then use technique 3 below.

To convert the raw_date variable to a daily time-series date, we use the command:

```
gen daily=date(raw_date, "dmy")
```

The “dmy” portion indicates the order of the day, month and year in the variable; so if the variable was of the form values been coded as “March 20, 1999” we would have used “mdy” instead.

The year must have 4 digits or else it returns missing values – therefore if the original date only has two digits, we place

the century before the "y.":

```
gen daily=date(raw_date, "dm19y")
```

Or, if we have non-daily dates, we can use the following functions:

```
weekly(stringvar, "wy")
```

```
monthly(stringvar, "my")
```

```
quarterly(stringvar, "qy")
```

```
halfyearly(stringvar, "hy")
```

```
yearly(stringvar, "y")
```

For example, if our data is 2002Q1, then

```
gen quarterly= quarterly(raw_data, "yq")
```

will get our elapsed quarters since 1960 Q1.

2. Date functions for partial date variables

If there are separate variables for each element of the date; for example:

month	day	year
7	11	1948
1	21	1952
11	2	1994
8	12	1993

We can use the `mdy()` function to create an elapsed Stata date variable. The month, day and year variables must be numeric. Therefore we can write:

```
gen mydate = mdy(month,day,year)
```

Or, with quarterly data, we would use the "yq()" function:

```
gen qtr=yq(year,quarter)
```

All of the functions are:

```
mdy(month,day,year)      for daily data
```

```
yw(year, week)          for weekly data
```

```
ym(year,month)          for monthly data
```

```
yq(year,quarter)        for quarterly data
```

```
yh(year,half-year)      for half-yearly data
```

3. Converting a date variable stored as a single number

As discussed above, if you have a single numeric variable, we need to first convert it into its component parts in order to use the `mdy` function. For example, imagine the variable is of the form `yyyymmdd` (for example, 19990320 for March 20 1999); now we need to split it into year, month and day as follows:

```

gen year = int(date/10000)

gen month = int((date-year*10000)/100)

gen day = int((date-year*10000-month*100))

gen mydate = mdy(month,day,year)

```

In each case the `int(x)` command returns the integer obtained by truncating `x` towards zero.

Using the time series date variables

Once we have the date variable in Stata elapsed time form, it is not the most intuitive to work with. For example, here is how a new variable called `stata_date` will look by using the command

```
gen stata_date = mdy(month,day,year)
```

month	day	year	stata_date
7	11	1948	-4191
1	21	1952	-2902
8	12	1993	12277
11	2	1994	12724

Therefore to display the `stata_date` in a more user-friendly manner, we can use the `format` command as follows:

```
format stata_date %d
```

This means that `stata_date` will now be displayed as:

month	day	year	stata_date7	11	1948	11jul1948
1	21	1952	21jan1952			
8	12	1993	12aug1993			
11	2	1994	02nov1994			

It is possible to use alternatives to `%d`, or to use `%td` to display elapsed dates in numerous other ways – in fact, we can control everything about the display. For example if I had instead written:

```
format stata_date %dM_d,_CY
```

Then we would get:

month	day	year	stata_date7	11	1948	July 11, 1948
1	21	1952	January 21, 1952			
8	12	1993	August 12, 1993			
11	2	1994	November 2, 1994			

See `help dfmt` for more details.

Making Use of Dates

If we want to use our dates in an `if` command, we have a number of options:

1. Exact dates
We have a selection of functions `d()`, `w()`, `m()`, `q()`, `h()`, and `y()` to specify exact daily, weekly, monthly, quarterly, half-yearly, and yearly dates respectively. For example:

```
reg x y if w(1995w9)

sum income if q(1988-3)
```

```
tab gender if y(1999)
```

2. A date range

If you want to specify a range of dates, you can use the `tin()` and `twithin()` functions:

```
reg y x if tin(01feb1990,01jun1990)
```

```
sum income if twithin(1988-3,1998-3)
```

The difference between `tin()` and `twithin()` is that `tin()` includes the beginning and end dates, whereas `twithin()` excludes them. Always enter the beginning date first, and write them out as you would for any of the `d()`, `w()`, etc. functions.

Time Series Tricks Using Dates

Often in time-series analyses we need to "lag" or "lead" the values of a variable from one observation to the next. Or we need to take differences or seasonal differences. One way is to generate a whole bunch of variables which represent the lag or the lead, the difference, etc... But if we have many variables, this can take up a lot of memory.

You should use the `tsset` command before any of the "tricks" in this section will work. This has the added advantage that if you have defined your data as a panel, Stata will automatically re-start any calculations when it comes to the beginning of a new cross-sectional unit so you need not worry about values from one panel being carried over to the next.

- Lags and Leads

These use the `L.varname` (to lag) and `F.varname` (to lead) commands. Both work the same way:

```
reg income L.income
```

This regresses `income(t)` on `income(t-1)`

If you wanted to lag `income` by more than one time period, you would simply change the `L.` to something like `"L2."` or `"L3."` to lag it by 2 and 3 time periods respectively.

- Differencing

Used in a similar way, the `D.varname` command will take the first difference, `D2.varname` will take the double difference (difference in difference), etc... For example:

Date	income	D.income	D2.income
02feb1999	120	.	.
02mar1999	130	10	.
02apr1999	145	15	5

- Seasonal Differencing

The `S.varname` command is similar to the `D.varname`, except that the difference is always taken from the current observation to the `nth` observation: In other words: `S.income=income(t)-income(t-1)` and `S2.income=income(t)-income(t-2)`

Date	income	S.income	S2.income
02feb1999	120	.	.
02mar1999	130	10	.
02apr1999	145	15	25

Programming

Program Basics

Creating or “defining” a program

A program contains a set of commands and is activated by a single command. A do-file is essentially one big program – it contains a list of commands and is activated by typing:

```
. do "statal.do"
```

You can also create special programs within a do-file, especially useful when you have a set of commands that are going to be used repetitively. The use of these programs will initially be demonstrated interactively, but they are best used within a do-file.

Suppose you want to create new variables that contain the average values (across countries and years) of some of the underlying variables in the dataset and at the same time display on screen these averages. No single Stata command will do this for you, but there are a couple of ways you can combine separate Stata commands to reach your goal. The most efficient method is:

```
. egen mean_kc=mean(kc)
. tab mean_kc
```

mean_kc	Freq.	Percent	Cum.
72.53644	8,568	100.00	100.00
Total	8,568	100.00	

```
. egen mean_kg=mean(kg)
. tab mean_kg
```

mean_kg	Freq.	Percent	Cum.
20.60631	8,568	100.00	100.00
Total	8,568	100.00	

The tasks are the same for each variable you are interested in. To avoid repetitive typing or repetitive cutting and pasting, you can create your own program that combines both tasks into a single command (note, in what follows the first inverted comma or single-quote of `1' is on the top-left key of your keyboard, the second inverted comma is on the right-hand side on the key with the @ symbol, and inside the inverted commas is the number one, not the letter L):

```
program define means
egen mean_`1'=mean(`1')
tab mean_`1'
end
```

You have now created your own Stata command called `mean`, and the variable you type after this new command will be used in the program everywhere there is a `1'. For example, `mean kg` will use `kg` everywhere there is a `1'. This command can now be applied to any variable you wish:

```
. means ki
```

mean_ki	Freq.	Percent	Cum.
15.74088	8,568	100.00	100.00
Total	8,568	100.00	

Naming a program

You can give your program any name you want as long as it isn't the name of a command already in Stata, e.g. you cannot name it `summarize`. Actually, you can create a program called `summarize`, but Stata will simply ignore it and use its own

summarize program every time you try using it. To check whether Stata has already reserved a particular name:

```
. which sum
built-in command: summarize

. which means
command means not found as either built-in or ado-file
r(111);
```

Redefining a program

You may want to change your program in some way, such as altering a line, or adding or dropping a line. For example, the `tabulate` command displays more than just the single number we are interested in. We can provide a more user-friendly result using the `display` command, where everything in double-quotes (") is interpreted as straightforward text and anything not in double-quotes is interpreted as something in Stata memory, such as a variable name or results of a previous command (e.g. `e(_b)` or `e(_se)` from a `regress` command):

```
. display "Mean of kg = " mean_kg
Mean of kg = 21.15341
```

Note, the value of `mean_kg` that is displayed is that of the first observation (`_n=1`). In our example, the value just so happens to be the same for all observations so we don't care whether it is displaying the first, tenth or one-hundredth observation. However, this will not be so in many other examples, so care needs to be taken when using this command in this way.

We can redefine our `mean` program by replacing the `tabulate` command with this `display` command. To do so, we must first drop the old `mean` program (placing `capture` before the command avoids Stata tripping up if there is no program called `mean` defined in the first place – useful for preventing Stata crashing in the middle of a long do-file):

```
capture program drop means
program define means
egen mean_`1' = mean(`1')
display "Mean of `1' = " mean_`1'
end
```

Now we need to drop the existing `mean_kc` and `mean_kg` variables and re-run the commands to get:

```
. means kc
Mean of kc = 72.536438

. means kg
Mean of kg = 20.606308
```

Debugging a program

Your program may crash out half-way through for some reason:

```
. means kc
mean_kc already defined
r(110);
```

Here, Stata tells us the reason why the program has crashed – you are trying to create a new variable called `mean_kc` but there is an old variable already called that. Our `mean` program is a very simple one, so we can figure out very quickly that the problem arises with the first line, `egen mean_`1' = mean(`1')`. However, with more intricate programs, it is not always so obvious where the problem lies. This is where the `set trace` command comes in handy. This command traces the execution of the program line-by-line so you can see exactly where it trips up. Because the trace details are often very long, it is usually a good idea to `log` them for review afterwards.

```
. log using "debug.log", replace
. set more off
. set trace on
. means kg
. set trace off
. log close
```


Program arguments

Our `means` command was defined to handle only one argument – ``1'`. It is possible to define more complicated programs to handle several arguments, ``1'`, ``2'`, ``3'`, and so on. These arguments can refer to anything you want – variable names, specific values, strings of text, command names, if statements, and so on. For example, we can define a program that displays the value of a particular variable (argument ``1'`) for a particular country (argument ``2'`) and year (argument ``3'`):

```
capture program drop show
program define show
    tempvar obs
    quietly gen `obs'=`1' if (countryisocode=="`2'" & year=="`3'")
    so `obs'
    display "`1' of country `2' in `3' is: " `obs'
end
```

To see this in action:

```
. show pop USA 1980
pop of country USA in 1980 is: 227726

. show pop FRA 1990
pop of country FRA in 1990 is: 58026.102
```

Some things to note about this program:

- Line 1 creates a temporary variable that will exist while the program is running but will be automatically dropped once the program has finished running. Once this `tempvar` has been defined, it must be referred to within the special quotes (``'`), just as with the arguments.
- Line 2 starts with `quietly`, which tells Stata to suppress any onscreen messages resulting from the operation of the command on this line.
- Make sure to properly enclose any string arguments within double-quotes. ``2'` will contain a string of text, such as ARG or FRA, so when ``2'` is being used in a command it should be placed within double-quotes (`"`).
- Line 3 ensures that observation number one will contain the value we are interested in. Missings are interpreted by Stata as arbitrarily large, so when the data is sorted in ascending order our value will be at the top of the list, ahead of these missings.

As we have seen, use of strings can cause a bit of a headache. A further complication may arise if the argument itself is a string containing blank spaces, such as United States instead of USA. Stata uses blank spaces to number the different arguments, so if we tried `show kg United States 1980`, Stata would assign `kg` to ``1'`, `United` to ``2'`, `States` to ``3'` and `1980` to ``4'`. The way to get around this is to enclose any text containing important blank spaces within double-quotes – the proper command then would be:

```
. show kg "United States" 1980
```

Renaming arguments

Using ``1'`, ``2'`, ``3'`, and so on can be confusing and prone to error. It is possible to assign more meaningful names to the arguments at the very beginning of your program so that the rest of the program is easier to create. Make sure to continue to include your new arguments within the special quotes (``'`):

```
capture program drop show
program define show
args var cty yr
tempvar obs
quietly gen `obs'=`var' if countryisocode=="`cty'" & year=="`yr'"
so `obs'
display "`var' of country `cty' in `yr' is: " `obs'
end
```

```
show kg USA 1980

kg of country USA in 1980 is: 13.660507
```

Macros

A Stata macro is different to an Excel macro. In Excel, a macro is like a recording of repeated actions which is then stored as a mini-program that can be easily run – this is what a do file is in Stata. Macros in Stata are the equivalent of variables in other programming languages. A macro is used as shorthand – you type a short macro name but are actually referring to some longer name or string of characters. For example, you may use the same list of independent variables in several regressions and want to avoid retyping the list several times. Just assign this list to a macro. Using the PWT dataset:

```
. local varlist gdp60 openk kc kg ki
. regress grgdpch `varlist' if year==1990
```

Source	SS	df	MS			
Model	694.520607	5	138.904121	Number of obs =	111	
Residual	2175.67123	105	20.7206784	F(5, 105) =	6.70	
				Prob > F =	0.0000	
				R-squared =	0.2420	
				Adj R-squared =	0.2059	
				Root MSE =	4.552	
Total	2870.19184	110	26.0926531			

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
gdp60	-1.853244	.6078333	-3.05	0.003	-3.058465	-.6480229
openk	-.0033326	.0104782	-0.32	0.751	-.0241088	.0174437
kc	-.0823043	.0356628	-2.31	0.023	-.153017	-.0115916
kg	-.0712923	.0462435	-1.54	0.126	-.1629847	.0204001
ki	.2327257	.0651346	3.57	0.001	.1035758	.3618757
_cons	16.31192	5.851553	2.79	0.006	4.709367	27.91447

```
. regress grgdpch `varlist' if year==1980
```

Source	SS	df	MS			
Model	880.685302	5	176.13706	Number of obs =	111	
Residual	8130.51957	105	77.4335197	F(5, 105) =	2.27	
				Prob > F =	0.0524	
				R-squared =	0.0977	
				Adj R-squared =	0.0548	
				Root MSE =	8.7996	
Total	9011.20487	110	81.9200443			

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]	
gdp60	-.2969159	1.143709	-0.26	0.796	-2.564679	1.970847
openk	.0023893	.0218479	0.11	0.913	-.0409311	.0457097
kc	-.1349823	.0518524	-2.60	0.011	-.237796	-.0321686
kg	-.1363929	.0845697	-1.61	0.110	-.304079	.0312932
ki	-.1307708	.1124651	-1.16	0.248	-.3537683	.0922267
_cons	16.98343	9.885368	1.72	0.089	-2.617433	36.58429

Macros are of two types – local and global. Local macros are “private” – they will only work within the program or do-file in which they are created. Thus, for example, if you are using several programs within a single do-file, using local macros for each means that you need not worry about whether some other program has been using local macros with the same names – one program can use `varlist` to refer to one set of variables, while another program uses its `varlist` to refer to a completely different set of variables. Global macros are “public” – they will work in all programs and do files – `varlist` refers to exactly the same list of variables irrespective of the program that uses it. Each type of macro has its uses, although local macros are the most commonly used type.

Just to illustrate this, let’s work with an example. The program `reg1` will create a local macro called `varlist` and will also use that macro. The program `reg2` will not create any macro, but will try to use a macro called `varlist`. Although `reg1` has a macro by that name, it is local or private to it, so `reg2` cannot use it:

```
program define reg1
local varlist gdp60 openk kc kg ki
reg grgdpch `varlist' if year==1990
end
```

```
. reg1
```

Source	SS	df	MS	Number of obs =	111
Model	694.520607	5	138.904121	F(5, 105) =	6.70
Residual	2175.67123	105	20.7206784	Prob > F =	0.0000
				R-squared =	0.2420
				Adj R-squared =	0.2059
Total	2870.19184	110	26.0926531	Root MSE =	4.552

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	-1.853244	.6078333	-3.05	0.003	-3.058465 - .6480229
openk	-.0033326	.0104782	-0.32	0.751	-.0241088 .0174437
kc	-.0823043	.0356628	-2.31	0.023	-.153017 -.0115916
kg	-.0712923	.0462435	-1.54	0.126	-.1629847 .0204001
ki	.2327257	.0651346	3.57	0.001	.1035758 .3618757
_cons	16.31192	5.851553	2.79	0.006	4.709367 27.91447

```
. capture program drop reg2
```

```
. program define reg2  
1. reg grgdpch `varlist' if year==1990  
2. end
```

```
. reg2
```

Source	SS	df	MS	Number of obs =	129
Model	0	0	.	F(0, 128) =	0.00
Residual	4008.61956	128	31.3173404	Prob > F =	.
				R-squared =	0.0000
				Adj R-squared =	0.0000
Total	4008.61956	128	31.3173404	Root MSE =	5.5962

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
_cons	.9033816	.492717	1.83	0.069	-.0715433 1.878306

Now, suppose we create a global macro called `varlist` – it will be accessible to all programs. Note, local macros are enclosed in the special quotes (``'`), global macros are prefixed by the dollar sign (`$`).

```
. global varlist gdp60 openk kc kg ki
```

```
. capture program drop reg1
```

```
. program define reg1  
1. local varlist gdp60 openk kc kg ki  
2. reg grgdpch `varlist'  
3. reg grgdpch $varlist  
4. end
```

```
. capture program drop reg2
```

```
. program define reg2  
1. reg grgdpch $varlist  
2. reg grgdpch `varlist'  
3. end
```

```
. reg1
```

Source	SS	df	MS	Number of obs =	5067
Model	8692.09731	5	1738.41946	F(5, 5061) =	41.21
Residual	213498.605	5061	42.1850632	Prob > F =	0.0000
Total	222190.702	5066	43.859199	R-squared =	0.0391
				Adj R-squared =	0.0382
				Root MSE =	6.495

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	-.5393328	.1268537	-4.25	0.000	-.7880209 -.2906447
openk	-.0003768	.0020639	-0.18	0.855	-.0044229 .0036693
kc	-.0249966	.0055462	-4.51	0.000	-.0358694 -.0141237
kg	-.0454862	.0089808	-5.06	0.000	-.0630924 -.02788
ki	.1182029	.011505	10.27	0.000	.0956481 .1407578
_cons	6.344897	1.045222	6.07	0.000	4.295809 8.393985

Source	SS	df	MS	Number of obs =	5067
Model	8692.09731	5	1738.41946	F(5, 5061) =	41.21
Residual	213498.605	5061	42.1850632	Prob > F =	0.0000
Total	222190.702	5066	43.859199	R-squared =	0.0391
				Adj R-squared =	0.0382
				Root MSE =	6.495

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	-.5393328	.1268537	-4.25	0.000	-.7880209 -.2906447
openk	-.0003768	.0020639	-0.18	0.855	-.0044229 .0036693
kc	-.0249966	.0055462	-4.51	0.000	-.0358694 -.0141237
kg	-.0454862	.0089808	-5.06	0.000	-.0630924 -.02788
ki	.1182029	.011505	10.27	0.000	.0956481 .1407578
_cons	6.344897	1.045222	6.07	0.000	4.295809 8.393985

```
. reg2
```

Source	SS	df	MS	Number of obs =	5067
Model	8692.09731	5	1738.41946	F(5, 5061) =	41.21
Residual	213498.605	5061	42.1850632	Prob > F =	0.0000
Total	222190.702	5066	43.859199	R-squared =	0.0391
				Adj R-squared =	0.0382
				Root MSE =	6.495

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
gdp60	-.5393328	.1268537	-4.25	0.000	-.7880209 -.2906447
openk	-.0003768	.0020639	-0.18	0.855	-.0044229 .0036693
kc	-.0249966	.0055462	-4.51	0.000	-.0358694 -.0141237
kg	-.0454862	.0089808	-5.06	0.000	-.0630924 -.02788
ki	.1182029	.011505	10.27	0.000	.0956481 .1407578
_cons	6.344897	1.045222	6.07	0.000	4.295809 8.393985

Source	SS	df	MS	Number of obs =	5621
Model	0	0	.	F(0, 5620) =	0.00
Residual	250604.237	5620	44.5915012	Prob > F =	.
Total	250604.237	5620	44.5915012	R-squared =	0.0000
				Adj R-squared =	0.0000
				Root MSE =	6.6777

grgdpch	Coef.	Std. Err.	t	P> t	[95% Conf. Interval]
_cons	2.069907	.0890675	23.24	0.000	1.8953 2.244513

As you will see, Stata runs two fully specified regressions in the first case but only one in the last case since again, the program `reg2` does not recognize ``varlist'`.

Macro contents

We introduced macros by showing how they can be used as shorthand for a list of variables. In fact, macros can contain practically anything you want – variable names, specific values, strings of text, command names, if statements, and so on. Note, we were actually using macros implicitly earlier in the class. When we created the programs `mean` and `show`, the arguments (e.g. `pop ARG 1980`) were passed to the programs via local macros (``1'`, ``2'`, ``3'`). These local macros contained variables (`kg`) and specific values (`ARG` and `1980`). Some other examples of what macros can contain:

Text

Text is usually contained in double quotes ("`"`) though this is not necessary for macro definitions:

```
. local ctynome "United States"
```

gives the same result as

```
. local ctynome United States
```

A problem arises whenever your macro name follows a backslash (`\`). Whenever this happens, Stata ignores the first single quote (`'`) of the macro name and so fails to properly load the macro:

```
. local filename PWT.dta
. use "F:\Stata classes\`filename'"
```

```
invalid `''
r(198);
```

To get around this problem, use double backslashes (`\\`) instead of a single one:

```
. use "F:\Stata classes\\`filename'"
```

Statements

Using macros to contain statements is essentially an extension of using macros to contain text. For example, if we define the local macro:

```
. local year90 "if year==1990"
```

then,

```
. reg grgdpc $varlist `year90'
```

is the same as:

```
. reg grgdpc gdp60 openk kc kg ki if year==1990
```

Note that when using `if` statements, double quotes become important again. For simplicity, consider running a regression for all countries whose codes start with "B". First, I define a local macro and then use it in the `reg` command:

```
. local ctynome B
. reg grgdpc gdp60 openk kc kg ki if substr(country,1,1)=="`ctynome'"
```

Although it does not matter whether I define `ctynome` using double quotes or not, it is important to include them in the `if`-statement since the variable `country` is string. The best way to think about this is to do what Stata does: replace ``ctynome'` by its content. Thus, `substr(country,1,1)=="`ctynome'"` becomes `substr(country,1,1)=="B"`. Omitting the double quotes would yield `substr(country,1,1)==B` which as usual results in an error message (since the results of the `substr`-operation is a string).

Numbers and expressions

```
. local i=1
. local result=2+2
```

Note, when the macro contains explicitly defined numbers or equations, an equality sign must be used. Furthermore, there must be no double-quotes, otherwise Stata will interpret the macro contents as text:

```
. local problem="2+2"
```

Thus, the `problem` macro contains the text `2+2` and the `result` macro contains the number 4. Note that as before we could also have assigned “2+2” to `problem` while omitting the equality sign. The difference between the two assignments is that assignments using “=” are evaluations, those without “=” are copy operations. That is, in the latter case, Stata simply copies “2+2” into the macro `problem` while in the former case it evaluates the expression behind the “=” and then assigns it to the corresponding macro. In the case of strings these two ways turn out to be equivalent. There is one subtle difference though: evaluations are limited to string lengths of 244 characters (80 in Intercooled Stata) while copy operations are only limited by available memory. Thus, it is usually safer to omit the equality sign to avoid parts of the macro being secretly cut off (which can lead to very high levels of confusion ...)

While a macro can contain numbers, it is essentially holding a string of text that can be converted back and forth into numbers whenever calculations are necessary. For this reason, macros containing numbers are only accurate up to 13 digits. When precise accuracy is crucial, scalars should be used instead:

```
. scalar root2=sqrt(2)
. display root2
1.4142136
```

Note, when you call upon a macro, it must be contained in special quotes (e.g. `display `result'`), but this is not so when you call upon a scalar (e.g. `display root2` and not `display `root2'`).

Manipulation of macros

Contents of macros can be changed by simply redefining a macro. For example, if the global macro `result` contains the value “2+2” typing:

```
. global result "2+3"
```

overwrites its contents. If you want to drop a specific macro, use the `macro drop` command:

```
. macro drop year90
```

To drop all macros in memory, use `_all` instead of specific macro names. If you want to list all macros Stata has saved in memory instead (including a number of pre-defined macros), type:

```
. macro list
```

or

```
. macro dir
```

Macro names starting with an underscore (“_”) are local macros, the others are global macros. Similarly, to drop or list scalars, use the commands `scalar drop` and `scalar list` (or `scalar dir`) respectively.

Temporary objects

Besides in macros and variables, Stata can also store information in so-called temporary variables which are often used in longer programmes:

- `tempvar` assigns names to the specified local macro names that may be used as temporary variable names in a dataset (we have already seen this type earlier on). When the program or do-file concludes, any variables with these assigned names are dropped:

```
program define temporary
tempvar logpop
gen `logpop'=log(pop)
sum pop if `logpop'>=8
end
```

```
. temporary
(2721 missing values generated)
```

Variable	Obs	Mean	Std. Dev.	Min	Max
pop	4162	43433.71	126246.4	2989	1258821

Since the `tempvar logcgdp` is dropped at the end of the program, trying to access it later on yields an error message:

```
. sum pop if `logpop'>=8
```

```
>8 invalid name
```

- `tempname` assigns names to the specified local macro names that may be used as temporary scalar or matrix names. When the program or do-file concludes, any scalars or matrices with these assigned names are dropped. This command is used more rarely than `tempvar` but can be useful if you want to do matrix-algebra in Stata subroutines (see the Stata User Guide [U], p. 220 for an example).
- `tempfile` assigns names to the specified local macro names that may be used as names for temporary files. When the program or do-file concludes, any datasets created with these assigned names are erased. For example, try the following programme:

```
program define temporary2
    tempfile cgd
    keep country year cgd
    save "`cgd'"
    clear
    use "`cgd'"
    sum year
end

. temporary2
file C:\DOCUME~1\Michael\LOCALS~1\Temp\ST_0c000012.tmp saved
```

Variable	Obs	Mean	Std. Dev.	Min	Max
year	8568	1975	14.72046	1950	2000

This saves the variables `country` `year` `cgd` in a temporary file that is automatically erased as soon as the programme terminates (check this by trying to reload “`cgd`” after termination of the programme “`temporary`”).

Looping

There are a number of techniques for looping or repeating commands within your do-file, thus saving you laborious retyping or cutting and pasting. These techniques are not always mutually exclusive – you can often use one or more different techniques to achieve the same goal. However, it is usually the case that one technique is more suitable or more efficient in a given instance than the others. Therefore, it is best to learn about each one and then choose whichever is most suitable when you come across a looping situation.

for

`for-processing` allows you to easily repeat Stata commands. As an example, we can use the PWT dataset and create the mean of several variables all at once:

```
. for varlist kc ki kg: egen mean_X=mean(X)

-> egen mean_kc=mean(kc)
-> egen mean_ki=mean(ki)
-> egen mean_kg=mean(kg)
```

The `egen` command is repeated for every variable in the specified `varlist`, with the `X` standing in for the relevant variable each time (note, instead of typing out a long `varlist`, you could e.g. use `varlist kc-ki` to signify every variable listed between `kc` and `kg`, inclusive). You can see in the variables window that our three desired variables have been created.

```
for varlist kc ki kg: display "Mean of X = " mean_X

-> display `"'Mean of kc = "' mean_kc
Mean of kc = 72.536438

-> display `"'Mean of ki = "' mean_ki
Mean of ki = 15.740885

-> display `"'Mean of kg = "' mean_kg
Mean of kg = 20.606308
```

The onscreen display includes both the individual commands and their results. To suppress the display of the individual commands, use the `noheader` option:

```
for varlist kc ki kg, noheader: display "Mean of X = " mean_X
```

```
Mean of kc = 72.536438
Mean of ki = 15.740885
Mean of kg = 20.606308
```

To suppress both the individual commands and their results, you need to specify `quietly` before `for`. The example we have used above repeats commands for a list of existing variables (`varlist`). You can also repeat for a list of new variables you want to create (`newlist`):

```
. for newlist ARG FRA USA : gen Xpop=pop if countryisocode=="X" & year==1995
```

It is also possible to repeat for a list of numbers (`numlist`) or any text you like (`anylist`). For example, suppose we wanted to append several similarly named data files to our existing dataset:

```
. for numlist 1995/1998: append using "F:\Stata classes\dataX.dta"
```

Note, the full file name `F:\Stata classes\dataX.dta` must be enclosed in double quotes, otherwise Stata will get confused and think the backslash `\` is a separator belonging to the `for` command:

```
. for numlist 1995/1998: append using F:\Stata classes\dataX.dta
-> append using F:
file F: not found
r(601);
```

It is possible to nest several loops within each other. In this case, you need to specify the name of the macro Stata uses for the list specified after `for` (in above examples, Stata automatically used `"X"`):

```
. for X in varlist kg cgdp: for Y in numlist 1990/1995: sum X if year==Y
```

It is also possible to combine two or more commands into a single `for`-process by separating each command with a backslash `\`:

```
. for varlist kg cgdp, noheader: egen mean_X=mean(X) \ display "Mean of X = " mean_X
```

If the list of commands you want to repeat is very long and/or complicated, it may be worthwhile using `for` in conjunction with a custom-made program containing your list of commands:

```
capture program drop mean
program define mean
quietly egen mean_`1'=mean(`1')
display mean_`1'
end
```

```
for varlist kg cgdp: mean X
```

foreach

We know that it is possible to combine several commands into a single `for`-process. This can get quite complicated if the list of commands is quite long, but we saw how you can overcome this by combining `for` with a custom-made program containing your list of commands. The `foreach` command does the same thing without the need for creating a separate program:

```
foreach var in kg cgdp {
    egen mean_`var'=mean(`var')
    display "Mean of `var' = " mean_`var'
}
```

```
Mean of kg = 20.606308
Mean of cgdp = 7.4677978
```


With the `foreach...in` command, `foreach` is followed by a macro name that you assign (e.g. `var`) and `in` is followed by the list of arguments that you want to loop (e.g. `kg cgdp`). This command can be easily used with variable names, numbers, or any string of text – just as `for` (in fact, `foreach` officially replaces `for` from version 8 onwards though `for` continues to work).

While this command is quite versatile, it still needs to be redefined each time you want to execute the same list of commands for a different set of arguments. For example, the program above will display the mean of `kg` and `cgdp`, but suppose that later on in your do-file you want to display the means of some other variables – you will have to create a new `foreach` loop. One way to get around this is to write the `foreach` loop into a custom-made program that you can then call on at different points in your do-file:

```
capture program drop mean
program define mean
foreach var of local 1 {
    egen mean_`var'=mean(`var')
    display "Mean of `var' = " mean_`var'
}
end

. mean "kg cgdp"
Mean of kg = 20.606308
Mean of cgdp = 7.4677978

. mean "ki pop"
Mean of ki = 15.740885
Mean of pop = 31252.467
```

This method works, but can be quite confusing. Firstly, `of local` is used in place of `in`. Secondly, reference to the local macro ``1'` in the first line does not actually use the single quotes we are used to. And thirdly, the list of arguments after the executing command must be in double quotes (so that everything is passed to the macro ``1'` in a single go). For these reasons, it can be a good idea to use `foreach` only when looping a once-off list. A technique called macro shift can be used when you want to loop a number of different lists (see later).

Incremental shift (number of loops is fixed)

You can loop or repeat a list of commands within your do-file using the `while` command – as long as the `while` condition is true, the loop will keep on looping. There are two broad instances of its use – the list of commands are to be repeated a fixed number of times (e.g. 5 loops, one for each year 1980-84) or the number of repetitions may vary (e.g. maybe 5 loops for a list of 5 years, or maybe 10 loops for a list of 10 years). We will look first at the incremental shift technique for a fixed number of loops. We can see how it works using the following very simple example:

```
local i=1
while `i'<=5 {
    display "loop number " `i'
    local i=`i'+1
}

loop number 1
loop number 2
loop number 3
loop number 4
loop number 5
```

The first command defines a local macro that is going to be the loop increment – it can be seen as a counter and is set to start at 1. It doesn't have to start at 1, e.g. if you are looping over years, it may start at 1980.

The second command is the `while` condition that must be satisfied if the loop is to be executed. This effectively sets the upper limit of the loop counter. At the end of the `while` command is an open bracket `{` that signifies the start of the looped or repeated set of commands. Everything between the two brackets `{ }` will be executed each time you go through the `while` loop. The final command before the close bracket `}` increases or increments the counter, readying it to go through the loop again (as long as the `while` condition is still satisfied). In actuality, it is redefining the local macro ``i'` – which is why there are no single quotes on the left of the equality but there are on the right. The increase in the counter does not have to be unitary, e.g. if you are using bi-annual data you may want to fix your increment to 2. All the looped commands within the brackets are defined in terms of the local macro ``i'`, so in the first loop everywhere there is an ``i'` there will now be a 1, in the second loop a 2, and so on.

To see a more concrete example, we will create a program to display the largest per capita GDP each year for every year 1980-84:

```
capture program drop maxcgdp
program define maxcgdp
    local i=1980
    while `i'<=1984 {
        quietly sum cgdp if year==`i'
        display `i' " " r(max)
        local i=`i'+1
    }
end

maxcgdp
1980 9.4067564
1981 9.5102491
1982 9.5399132
1983 9.6121063
1984 9.7101154
```

Macro shift (number of loops is variable)

The incremental shift technique used a fully defined counter with a fixed start (1980), end (1984) and increment (1 year). You type a single command (`maxrgdp1`) to execute the program that loops over this fully defined counter. However, this technique cannot be used if the required replications are not so neatly definable, e.g. you want to repeat a set of commands for 1980, 1984, 1986 and 1995, or you want to repeat the commands for 1980-84 and 1990-94. Instead, you write a program that is executed by the command and a list of arguments that represent the required replications (e.g. `maxrgdp1 1980 1984 1986 1995`). Stata will allocate the first argument to local macro ``1'`, the second to local macro ``2'`, and so on. Thus, you need to shift through each of these arguments or local macros in order to shift through the required replications. A simple example of how this works:

```
capture program drop displayno
program define displayno
while "`1'"~="" {
    display `1'
    macro shift
}
end

. displayno 1 2 4 8 10
1
2
4
8
10
. displayno 77 90876 8
77
90876
8
```

The command `macro shift` is used here instead of the counter increment device – it shifts the contents of local macros one place to the left; ``1'` disappears and ``2'` becomes ``1'`, ``3'` becomes ``2'`, and so on. So, in the example above, ``1'` initially contained the number 77, ``2'` contained 90876 and ``3'` contained 8. The looped commands are in terms of ``1'` only so the first replication uses the number 77. The `macro shift` command then shifts ``2'` into the ``1'` slot, so the second replication uses the number 90876. Similarly, the third replication uses the number 8.

The `while` command at the start of the loop ensures that it will keep on looping until the local macro ``1'` is empty, i.e. it will work as long as `"`1'"` is not an empty string `"`. This is similar to the `while` command in the incremental shift technique, but here the loop is defined in terms of ``1'` instead of ``i'` and it is contained in double quotes. The use of double quotes is a convenient way to ensure the loop continues until the argument or macro ``1'` contains nothing – it has nothing to do with whether the arguments are strings of text or numbers.

For a more realistic application of this technique, we can revisit our `maxcgdp` program:

```

capture program drop maxcgdp
program define maxcgdp
    while "`1'"~="" {
        quietly sum cgdp if year==`1'
        display `1' " " r(max)
        macro shift
    }
end

. maxcgdp 1983 1991 1999
1983 9.6121063
1991 10.137326
1999 10.699246

```

Note that, essentially, the only things that have changed are the format of the `while` command, the format of the shifting mechanism and the way in which the local macro in the loop is defined (``1'` instead of ``i'`).

The macro `shift` technique is commonly used to shift through variables rather than actual values. For example:

```

capture program drop means
program define means
    while "`1'"~="" {
        tempvar mean
        quietly egen `mean'=mean(`1')
        display "Mean of `1' = " `mean'
        macro shift
    }
end

```

Now, we can display the mean of a single variable:

```

means kg
Mean of kg = 20.606308

```

or of a list of variables:

```

. mean kg pop cgdp
Mean of kg = 20.606308
Mean of pop = 31252.467
Mean of cgdp = 7.4677978

```

Branching

Branching allows you to do one thing if a certain condition is true, and something else when that condition is false. For example, suppose you are doing some sort of analysis year-by-year but you want to perform different types of analyses for the earlier and later years. For simplicity, suppose you want to display the minimum per capita GDP for the years to 1982 and the maximum value thereafter:

```

capture program drop minmaxcgdp
program define minmaxcgdp
    local i=1980
    while `i'<=1984 {
        if `i'<=1982 {
            local function min
        }
        else {
            local function max
        }
        quietly sum cgdp if year==`i'
        display `i' " " r(`function')
        local i=`i'+1
    }
end

. minmaxcgdp
1980 5.518826
1981 5.9500399

```

```
1982 6.0682001
1983 9.6121063
1984 9.7101154
```

The structure of this program is almost identical to that of the `maxcgdp` program created earlier. The only difference is that `egen` in line 6 is now a `min` or `max` function depending on the `if/else` conditions in lines 3 and 4.

It is very important to get the brackets `{}` correct in your programs. Firstly, every `if` statement, every `else` statement, and every `while` statement must have their conditions fully enclosed in their own set of brackets – thus, if there are three condition with three open brackets `{`, there must also be three close brackets `}`. Secondly, nothing except comments in `/* */` should be typed after a close bracket, as Stata automatically moves on to the next line when it encounters a close bracket. Thus, Stata would ignore the `else` condition if you typed:

```
. if `i'<=1982 {local function min} else {local function max}
```

Thirdly, it is necessary to place the brackets and their contents on different lines, irrespective of whether the brackets contain one or more lines of commands. Finally, it is possible to embed `if/else` statements within other `if/else` statements for extra levels of complexity, so it is crucial to get each set of brackets right. Suppose you want to display the minimum per capita GDP for 1980 and 1981, the maximum for 1982 and 1983, and the minimum for 1984:

```
capture program drop minmaxrgdpl
program define minmaxrgdpl
    local i=1980
    while `i'<=1984 {
        if `i'<=1981 {
            local function min
        }
        else {
            if `i'<=1983 {
                local function max
            }
            else {
                local function min
            }
        }
        quietly sum cgdp if year==`i'
        display `i' " " r(max)
        local i=`i'+1
    }
end
```

```
. minmaxcgdp
1980 5.518826
1981 5.9500399
1982 9.5399132
1983 9.6121063
1984 6.1164122
```

One final thing to note is that it is important to distinguish between the conditional `if`:

```
. sum cgdp if cgdp>8
and the programming if:
if cgdp >8 {
    sum cgdp
}
```

The conditional `if` summarizes all the observations on `cgdp` that are greater than 8. The programming `if` looks at the first observation on `cgdp` to see if it is greater than 8, and if so, it executes the `sum cgdp` command, i.e. it summarizes *all* observations on `cgdp` (try out the two commands and watch the number of observations).

ADO programming

ADO programming involves setting up user-defined programmes which will be stored in the memory of Stata and then can be retrieved as a command whenever you use Stata. For example, *regress* is used as an ado file. While I think it is pretty advanced to start programming your own complex ado files, some very simple files might be of use. The following simple examples come from <http://www.ats.ucla.edu/stat/stata/stat130/median.htm>. They are 3 programs of increasing complexity (and therefore flexibility) to take the median of a series or set of series.

Median Program -- Version #1

Basic program to deal with one variable.

```
program define median1
  version 6
  sort `1'
  quietly count if `1' ~= .
  local n = r(N)
  local mid = int(`n'/2)
  local odd = mod(`n',2)

  if `odd' {
    local median = `1'[`mid'+1]
  }
  else {
    local median = (`1'[`mid'] + `1'[`mid'+1])/2
  }

  display "Median of `1' = `median'"
end
```

Median Program -- Version #2

Multiple variables and saves results in return list.

```
program define median2, rclass
  display
  display in green " Variable          N      Median"
  display in green "-----"

  while "`1'" ~= "" {
    quietly count if `1' ~= .
    local n = r(N)
    local i = int(`n'/2)
    local odd = mod(`n',2)
    sort `1'
    if `odd' {
      local median = `1'[`i'+1]
    }
    else {
      local median = (`1'[`i'] + `1'[`i'+1])/2
    }
    display in yellow %9s "`1'" %9.0f `n' %10.2f `median'
    macro shift
  }
  return local Mdn = `median'
  return local N = `n'
end
```

Median Program -- Version #3

Allows for multiple series and for “if” and “in” statements.

```
program define median3, rclass
  syntax varlist [if] [in]
  tokenize `varlist'
  preserve
  marksample touse
  display
  display in green " Variable           N      Median"
  display in green "-----"

  while "`1'" ~= "" {
    quietly keep if `touse'
    quietly count
    local n = r(N)
    local i = int(`n'/2)
    local odd = mod(`n',2)
    sort `1'
    if `odd' {
      local median = `1'[`i'+1]
    }
    else {
      local median = (`1'[`i'] + `1'[`i'+1])/2
    }
    display in yellow %9s "`1'" %9.0f `n' %10.2f `median'
    macro shift
  }
  return local Mdn = `median'
  return local N = `n'
end
```

These programmes can be written (separately) in the do file editor and then saved as .ado files. You should save them in the Stata directory which contains ado file updates, under the “m” folder. Then whenever you type:

```
median1 variablename
```

it will calculate the median for you.