









The Workflow of Data Analysis: Principles and Practice





HERBORD



The Workflow of Data Analysis: Principles and Practice

J. SCOTT LONG
Department of Sociology
Indiana University–Bloomington



A Stata Press Publication
StataCorp LP
College Station, Texas





Copyright © 2009 by StataCorp LP
All rights reserved. First edition 2009

Published by Stata Press, 4905 Lakeway Drive, College Station, Texas 77845
Typeset in L^AT_EX 2_ε
Printed in the United States of America
10 9 8 7 6 5 4 3 2 1

ISBN-10: 1-59718-047-5
ISBN-13: 978-1-59718-047-4

No part of this book may be reproduced, stored in a retrieval system, or transcribed, in any form or by any means—electronic, mechanical, photocopy, recording, or otherwise—without the prior written permission of StataCorp LP.

Stata is a registered trademark of StataCorp LP. L^AT_EX 2_ε is a trademark of the American Mathematical Society.



THE

To Valerie

THE



BE



BO

BO



THE

BY





HERBORD



Contents

List of tables	xxi
List of figures	xxiii
Preface	xxv
1 Introduction	1
1.1 Replication: The guiding principle for workflow	2
1.2 Steps in the workflow	3
1.2.1 Cleaning data	4
1.2.2 Running analysis	4
1.2.3 Presenting results	4
1.2.4 Protecting files	4
1.3 Tasks within each step	5
1.3.1 Planning	5
1.3.2 Organization	5
1.3.3 Documentation	5
1.3.4 Execution	6
1.4 Criteria for choosing a workflow	6
1.4.1 Accuracy	6
1.4.2 Efficiency	6
1.4.3 Simplicity	7
1.4.4 Standardization	7
1.4.5 Automated	7
1.4.6 Usability	7
1.4.7 Scalability	8
1.5 Changing your workflow	8

1.6	How the book is organized	8
2	Planning, organizing, and documenting	11
2.1	The cycle of data analysis	13
2.2	Planning	14
2.3	Organization	18
2.3.1	Principles for organization	18
2.3.2	Organizing files and directories	19
2.3.3	Creating your directory structure	20
	A directory structure for a small project	21
	A directory structure for a large, one-person project	22
	Directories for collaborative projects	23
	Special-purpose directories	25
	Remembering what directories contain	27
	Planning your directory structure	28
	Naming files	29
	Batch files	30
2.3.4	Moving into a new directory structure (advanced topic)	30
	Example of moving into a new directory structure	31
2.4	Documentation	34
2.4.1	What should you document?	36
2.4.2	Levels of documentation	37
2.4.3	Suggestions for writing documentation	38
	Evaluating your documentation	39
2.4.4	The research log	39
	A sample page from a research log	40
	A template for research logs	42
2.4.5	Codebooks	43
	A codebook based on the survey instrument	43
2.4.6	Dataset documentation	44
2.5	Conclusions	45

<i>Contents</i>	ix
3 Writing and debugging do-files	47
3.1 Three ways to execute commands	47
3.1.1 The Command window	48
3.1.2 Dialog boxes	49
3.1.3 Do-files	49
3.2 Writing effective do-files	50
3.2.1 Making do-files robust	51
Make do-files self-contained	51
Use version control	53
Exclude directory information	53
Include seeds for random numbers	55
3.2.2 Making do-files legible	55
Use lots of comments	56
Use alignment and indentation	57
Use short lines	58
Limit your abbreviations	61
Be consistent	63
3.2.3 Templates for do-files	63
Commands that belong in every do-file	63
A template for simple do-files	66
A more complex do-file template	66
3.3 Debugging do-files	68
3.3.1 Simple errors and how to fix them	68
Log file is open	68
Log file already exists	68
Incorrect command name	69
Incorrect variable name	69
Incorrect option	70
Missing comma before options	70
3.3.2 Steps for resolving errors	70

Step 1: Update Stata and user-written programs	70
Step 2: Start with a clean slate	71
Step 3: Try other data	72
Step 4: Assume everything could be wrong	72
Step 5: Run the program in steps	72
Step 6: Exclude parts of the do-file	73
Step 7: Starting over	74
Step 8: Sometimes it is not your mistake	74
3.3.3 Example 1: Debugging a subtle syntax error	75
3.3.4 Example 2: Debugging unanticipated results	77
3.3.5 Advanced methods for debugging	81
3.4 How to get help	82
3.5 Conclusions	82
4 Automating your work	83
4.1 Macros	83
4.1.1 Local and global macros	84
Local macros	84
Global macros	85
Using double quotes when defining macros	85
Creating long strings	85
4.1.2 Specifying groups of variables and nested models	86
4.1.3 Setting options with locals	88
4.2 Information returned by Stata commands	90
Using returned results with local macros	92
4.3 Loops: foreach and forvalues	92
The foreach command	94
The forvalues command	95
4.3.1 Ways to use loops	95
Loop example 1: Listing variable and value labels	96
Loop example 2: Creating interaction variables	97

<i>Contents</i>	xi
Loop example 3: Fitting models with alternative measures of education	98
Loop example 4: Recoding multiple variables the same way	98
Loop example 5: Creating a macro that holds accumulated information	99
Loop example 6: Retrieving information returned by Stata	100
4.3.2 Counters in loops	101
Using loops to save results to a matrix	102
4.3.3 Nested loops	104
4.3.4 Debugging loops	105
4.4 The include command	106
4.4.1 Specifying the analysis sample with an include file	107
4.4.2 Recoding data using include files	107
4.4.3 Caution when using include files	109
4.5 Ado-files	110
4.5.1 A simple program to change directories	111
4.5.2 Loading and deleting ado-files	112
4.5.3 Listing variable names and labels	112
Version 1	113
Version 2	114
Version 3	116
Version 4	116
4.5.4 A general program to change your working directory	117
4.5.5 Words of caution	118
4.6 Help files	118
4.6.1 nmlabel.hlp	118
4.6.2 help me	121
4.7 Conclusions	122
5 Names, notes, and labels	123
5.1 Posting files	123
5.2 A dual workflow for data management and statistical analysis	125

5.3	Names, notes, and labels	127
5.4	Naming do-files	127
5.4.1	Naming do-files to re-create datasets	128
5.4.2	Naming do-files to reproduce statistical analysis	128
5.4.3	Using master do-files	129
	Master log files	131
5.4.4	A template for naming do-files	132
	Using subdirectories for complex analyses	133
5.5	Naming and internally documenting datasets	134
	Never name it final!	135
5.5.1	One time only and temporary datasets	135
5.5.2	Datasets for larger projects	136
5.5.3	Labels and notes for datasets	136
5.5.4	The datasignature command	137
	A workflow using data signature	138
	Changes datasignature does not detect	139
5.6	Naming variables	141
5.6.1	The fundamental principle for creating and naming variables	141
5.6.2	Systems for naming variables	142
	Sequential naming systems	143
	Source naming systems	143
	Mnemonic naming systems	144
5.6.3	Planning names	144
5.6.4	Principles for selecting names	145
	Anticipate looking for variables	145
	Use simple, unambiguous names	146
	Try names before you decide	148
5.7	Labeling variables	149
5.7.1	Listing variable labels and other information	149
	Changing the order of variables in your dataset	152

- 5.7.2 Syntax for label variable 153
- 5.7.3 Principles for variable labels 154
 - Beware of truncation 154
 - Test labels before you post the file 155
- 5.7.4 Temporarily changing variable labels 155
- 5.7.5 Creating variable labels that include the variable name . . . 156
- 5.8 Adding notes to variables 158
 - 5.8.1 Commands for working with notes 159
 - Listing notes 159
 - Removing notes 160
 - Searching notes 160
 - 5.8.2 Using macros and loops with notes 160
 - 5.8.3 Using value labels 161
 - 5.8.4 Creating value labels is a two-step process 162
 - Step 1: Defining labels 162
 - Step 2: Assigning labels 162
 - Why a two-step system? 162
 - Removing labels 163
 - 5.8.5 Principles for constructing value labels 163
 - 1) Keep labels short 163
 - 2) Include the category number 164
 - 3) Avoid special characters 166
 - 4) Keeping track of where labels are used 167
 - 5.8.6 Cleaning value labels 168
 - 5.8.7 Consistent value labels for missing values 169
 - 5.8.8 Using loops when assigning value labels 169
- 5.9 Using multiple languages 171
 - 5.9.1 Using label language for different written languages 172
 - 5.9.2 Using label language for short and long labels 172
- 5.10 A workflow for names and labels 174

Step 1: Plan the changes	174
Step 2: Archive, clone, and rename	175
Step 3: Revise variable labels	175
Step 4: Value labels	175
Step 5: Verify the changes	176
5.10.1 Step 1: Check the source data	176
Step 1a: List the current names and labels	176
Step 1b: Try the current names and labels	178
5.10.2 Step 2: Create clones and rename variables	180
Step 2a: Create clones	181
Step 2b: creating rename commands	181
Step 2c: rename variables	182
5.10.3 Step 3: Revising variable labels	183
Step 3a: Creating variable label commands	183
Step 3b: revising variable labels	184
5.10.4 Step 4: Revising value labels	185
Step 4a: List the current labels	186
Step 4b: Create label define commands to edit	187
Step 4c: Revise labels and add them to dataset	190
5.10.5 Step 5: Checking the new names and labels	192
5.11 Conclusions	192
6 Cleaning your data	193
6.1 Importing data	194
6.1.1 Data formats	194
ASCII data formats	194
Binary-data formats	196
6.1.2 Ways to import data	197
Stata commands to import data	197
Using other statistical packages to export data	199
Using a data conversion program	199



6.1.3	Verifying data conversion	199
	Converting the ISSP 2002 data from Russia	200
6.2	Verifying variables	206
6.2.1	Values review	207
	Values review of data about the scientific career	208
	Values review of data on family values	211
6.2.2	Substantive review	212
	What does time to degree measure?	212
	Examining high-frequency values	214
	Links among variables	215
	Changes in survey questions	220
6.2.3	Missing-data review	221
	Comparisons and missing values	221
	Creating indicators of whether cases are missing	223
	Using extended missing values	224
	Verifying and expanding missing-data codes	225
	Using include files	232
6.2.4	Internal consistency review	233
	Consistency in data on the scientific career	234
6.2.5	Principles for fixing data inconsistencies	236
6.3	Creating variables for analysis	237
6.3.1	Principles for creating new variables	237
	New variables get new names	237
	Verify that new variables are correct	238
	Document new variables	239
	Keep the source variables	239
6.3.2	Core commands for creating variables	240
	The generate command	240
	The clonevar command	240
	The replace command	241

6.3.3	Creating variables with missing values	242
6.3.4	Additional commands for creating variables	244
	The recode command	244
	The egen command	245
	The tabulate, generate() command	247
6.3.5	Labeling variables created by Stata	248
6.3.6	Verifying that variables are correct	249
	Checking the code	250
	Listing variables	250
	Plotting continuous variables	251
	Tabulating variables	253
	Constructing variables multiple ways	254
6.4	Saving datasets	255
	Steps in saving a dataset	255
6.4.1	Selecting observations	256
	Deleting cases versus creating selection variables	256
6.4.2	Dropping variables	257
	Selecting variables for the ISSP 2002 Russian data	258
6.4.3	Ordering variables	258
6.4.4	Internal documentation	259
6.4.5	Compressing variables	259
6.4.6	Running diagnostics	260
	The codebook, problem command	260
	Checking for unique ID variables	262
6.4.7	Adding a datasignature	264
6.4.8	Saving the file	264
6.4.9	After a file is saved	265
6.5	Extended example of preparing data for analysis	265
	Creating control variables	265
	Creating binary indicators of positive attitudes	268



<i>Contents</i>	xvii
Creating four-category scales of positive attitudes	271
6.6 Merging files	273
6.6.1 Match merging	273
Sorting the ID variable	274
6.6.2 One-to-one merging	275
Combing unrelated datasets	275
Forgetting to match merge	277
6.7 Conclusions	278
7 Analyzing data and presenting results	279
7.1 Planning and organizing statistical analysis	279
7.1.1 Planning in the large	280
7.1.2 Planning in the middle	281
7.1.3 Planning in the small	283
7.2 Organizing do-files	283
7.2.1 Using master do-files	284
7.2.2 What belongs in your do-file?	286
7.3 Documentation for statistical analysis	287
7.3.1 The research log and comments in do-files	287
7.3.2 Documenting the provenance of results	288
Captions on graphs	290
7.4 Analyzing data using automation	290
7.4.1 Locals to define sets of variables	291
7.4.2 Loops for repeated analyses	292
Computing t tests using loops	292
Loops for alternative model specifications	294
7.4.3 Matrices to collect and print results	295
Collecting results of t tests	295
Saving results from nested regressions	298
Saving results from different transformations of articles	300
7.4.4 Creating a graph from a matrix	302

7.4.5	Include files to load data and select your sample	303
7.5	Baseline statistics	304
7.6	Replication	305
7.6.1	Lost or forgotten files	305
7.6.2	Software and version control	306
7.6.3	Unknown seed for random numbers	306
	Bootstrapped standard errors	306
	Letting Stata set the seed	307
	Training and confirmation samples	308
	Using a global that is not in your do-file	310
7.7	Presenting results	310
7.7.1	Creating tables	311
	Using spreadsheets	311
	Regression tables with esttab	313
7.7.2	Creating graphs	315
	Colors, black, and white	316
	Font size	318
7.7.3	Tips for papers and presentations	318
	Papers	318
	Presentations	319
7.8	A project checklist	320
7.9	Conclusions	320
8	Protecting your files	323
8.1	Levels of protection and types of files	324
8.2	Causes of data loss and issues in recovering a file	326
8.3	Murphy's law and rules for copying files	329
8.4	A workflow for file protection	330
	Part 1: Mirroring active storage	330
	Part 2: Offline backups	332
8.5	Archival preservation	335

<i>Contents</i>	xix
8.6 Conclusions	337
9 Conclusions	339
A How Stata works	341
A.1 How Stata works	341
Stata directories	341
The working directory	342
A.1.1 Working on a network	343
A.2 Customizing Stata	345
A.2.1 Fonts and window locations	345
A.2.2 Commands to change preferences	345
Options that can be set permanently	345
Options that need to be set each session	347
A.2.3 profile.do	347
Function keys	348
A.3 Additional resources	348
References	351
Author index	355
Subject index	357





HERBORD





Tables

3.1	Command abbreviations used in the book	63
7.1	Example table using <code>esttab</code>	315
8.1	Issues related to backup from the perspective of a data analyst . . .	325



HERBIBOON



Figures

2.1	Data analysis cycle	13
2.2	Spreadsheet plan of a directory structure	29
2.3	Sample page from a research log	41
2.4	Workflow research log template	42
2.5	Survey instrument used for the SGC-MHS Study	44
2.6	Data registry spreadsheet	45
4.1	Viewer window displaying <code>help nmlabel</code>	119
4.2	Viewer window displaying <code>help me</code>	122
5.1	The dual workflow of data management and statistical analysis	125
5.2	The dual workflow of data management and statistical analysis, version 2	130
5.3	Sample spreadsheet	145
6.1	A fixed-format ASCII file	195
6.2	A free-format ASCII file with variable names	195
6.3	A binary file in Stata format	196
6.4	Computed descriptive statistics	201
6.5	Computed frequency distributions	201
6.6	Stat/Transfer dialog box	202
6.7	Stat/Transfer case selection dialog box	203
6.8	Combined missing values	205
6.9	Aspects of data verification	206
6.10	Thumbnails of two-way graphs	220
6.11	Spreadsheet of possible reasons for missing data	227

6.12	Spreadsheet with variables that require similar data processing grouped together	228
6.13	Merging datasets	275
7.1	Example of a research log	288
7.2	Example showing the recording of the origin of each number within a paper	289
7.3	Addition of the name of the do-file, location within it, person who ran it, and the data	289
7.4	Spreadsheet with pasted text	312
7.5	Convert text to columns wizard	312
7.6	Spreadsheet with header information added	313
7.7	Colored graph in black and white	317
7.8	Graph with small, hard to read text and a graph with readable text	318
8.1	Levels of protection for files	324
8.2	A two-part workflow for protecting files used in data analysis	330

Preface

This book is about methods that allow you to work efficiently and accurately when you analyze data. Although it does not deal with specific statistical techniques, it discusses the steps that you go through with any type of data analysis. These steps include planning your work, documenting your activities, creating and verifying variables, generating and presenting statistical analyses, replicating findings, and archiving what you have done. These combined issues are what I refer to as the *workflow of data analysis*. A good workflow is essential for replication of your work, and replication is essential for good science.

My decision to write this book grew out of my teaching, researching, consulting, and collaborating. I increasingly saw that people were *drowning* in their data. With cheap computing and storage, it is easier to create files and variables than it is to keep track of them. As datasets have become more complicated, the process of managing data has become more challenging. When consulting, much of my time was spent on issues of data management and figuring out what had been done to generate a particular set of results. In collaborative projects, I found that problems with workflow were multiplied. Another motivation came from my work with Jeremy Freese on the package of Stata programs known as SPost (Long and Freese 2006). These programs were downloaded 10,000 times last year, and we were contacted by hundreds of users. Responding to these questions showed me how researchers from many disciplines organize their data analysis and the ways in which this organization can break down. When helping someone with what appeared to be a problem with an SPost command, I often discovered that the problem was related to some aspect of the user's workflow. When people asked if there was something they could read about this, I had nothing to suggest.

A final impetus for writing the book came from Bruce Fraser's *Real World Camera Raw with Adobe Photoshop CS2* (2005). A much touted advantage of digital photography is that you can take a lot of pictures. The catch is keeping track of thousands of pictures. Imaging experts have been aware of this issue for a long time and refer to it as *workflow*—keeping track of your work as it flows through the many stages to the final product. As the amount of time I spent looking for a particular picture became greater than the time I spent taking pictures, it was clear that I needed to take Fraser's advice and develop a workflow for digital imaging. Fraser's book got me thinking about data analysis in terms of the concept of a workflow.

After years of gestation, the book took two years to write. When I started, I thought my workflow was very good and that it was simply a matter of recording what I did. As writing proceeded, I discovered gaps, inefficiencies, and inconsistencies in what I did.

Sometimes these involved procedures that I knew were awkward, but where I never took the time to find a better approach. Some problems were due to oversights where I had not realized the consequences of the things I did or failed to do. In other instances, I found that I used multiple approaches for the same task, never choosing one as the best practice. Writing this book forced me to be more consistent and efficient. The advantages of my improved workflow became clear when revising two papers that were accepted for publication. The analyses for one paper were completed before I started the workflow project, whereas the analyses for the other were completed after much of the book had been drafted. I was pleased by how much easier it was to revise the analyses in the paper that used the procedures from the book. Part of the improvement was due to having better ways of doing things. Equally important was that I had a consistent and documented way of doing things.

I have no illusions that the methods I recommend are the best or only way of doing things. Indeed, I look forward to hearing from readers who have suggestions for a better workflow. Your suggestions will be added to the book's web site. However, the methods I present work well and avoid many pitfalls. An important aspect of an efficient workflow is to find one way of doing things and sticking with it. Uniform procedures allow you to work faster when you initially do the work, and they help you to understand your earlier work if you need to return to it at a latter time. Uniformity also makes working in research teams easier since collaborators can more easily follow what others have done. There is a lot to be said in favor of having established procedures that are documented and working with others who use the same procedures. I hope you find that this book provides such procedures.

Although this book should be useful for anyone who analyzes data, it is written within several constraints. First, Stata is the primary computing language because I find Stata to be the best, general purpose software for data management and statistical analysis. Although nearly everything I do with Stata can be done in other software, I do not include examples from other packages. Second, most examples use data from the social sciences, because that is the field in which I work. The principles I discuss, however, apply broadly to other fields. Finally, I work primarily in Windows. This is not because I think Windows is a better operating system than Mac or Linux, but because Windows is the primary operating system where I work. Just about everything I suggest works equally well in other operating systems, and I have tried to note when there are differences.

I want to thank the many people who commented on drafts or answered questions about some aspect of workflow. I particularly thank Curtis Child, Nadine Reibling, and Tait Runnfeldt Medina whose detailed comments greatly improved the book. I also thank Alan Acock, Art Alderson, Myron Gutmann, Patricia McManus, Eliza Pavalko, Jack Thomas, Leah VanWey, Rich Watson, Terry White, and Rich Williams for talking with me about workflow. David M. Drukker at StataCorp answered many of my questions. His feedback made it a better book and his friendship made it more fun to write. Some of the material in this book grew out of research funded by NIH Grant Number R01TW006374 from the Fogarty International Center, the National Institute

of Mental Health, and the Office of Behavioral and Social Science Research to Indiana University–Bloomington. Other work was supported by an anonymous foundation and The Bayer Group. I gratefully acknowledge support provided by the College of Arts and Sciences at Indiana University.

Without the unintended encouragement from my dear friend Fred, I would not have started the book. Without the support of my dear wife Valerie, I would not have completed it. Long overdue, this book is dedicated to her.



A word about fonts, files, commands, and examples

The book uses standard Stata conventions for typography. Things printed in a typewriter-style typeface are Stata commands and options. For example, `use mydata, clear`. Italics indicate information that you should add. For example, `use dataset-name, clear` indicates that you should substitute the name of your dataset. When I provide the syntax for a command, I generally show only some of the options. For full documentation, you can type `help command-name` or check the reference manual. Manuals are referred to with the usual Stata conventions. For example, [R] `logit` refers to the `logit` entry in the *Base Reference Manual* and [D] `sort` refers to the `sort` entry in the *Data Management Reference Manual*.



The book includes many examples that I encourage you to try as you read. If the name of a file begins with `wf` you can download that file. I use (file: *filename.do*) to let you know the name of the do-file that corresponds to the example being presented. With few exceptions (e.g., some ado-files), if the name of a file does not begin with `wf` (e.g., `science2.dta`), the file is not available for download. To download the examples, you must be in Stata and connected to the Internet. There are two Workflow packages for Stata 10 (`wf10-part1` and `wf10-part2`) and two for Stata 9 (`wf09-part1` and `wf10-part1`). To find and install the packages, type `findit workflow`, choose the packages you need, and follow the instructions. Although two packages are needed because of the large number of examples, I refer to them simply as the Workflow package. Before trying these examples, be sure to update your copy of Stata as described in [GS] **20 Updating and extending Stata—Internet functionality**. Additional information related to the book is located at <http://www.indiana.edu/~jslsoc/workflow.htm>.

Bloomington, IN
August 2008

Scott Long



HERBORD



1 Introduction

This book is about methods for analyzing your data effectively, efficiently, and accurately. I refer to these methods as the *workflow of data analysis*. Workflow involves the entire process of data analysis including planning and documenting your work, cleaning data and creating variables, producing and replicating statistical analyses, presenting findings, and archiving your work. You already have a workflow, even if you do not think of it as such. This workflow might be carefully planned or it might be ad hoc. Since workflow for data analysis is rarely described in print or formally taught, researchers often develop their workflow in reaction to problems they encounter and from informal suggestions from colleagues. For example, after you discover two files with the same name but different content, you might develop procedures (i.e., a workflow) for naming files. Too often, good practice in data analysis is learned inefficiently through trial and error. Hopefully, my book will shorten the learning process and allow you to spend more time on what you really want to do.

Reactions to early drafts of this book convinced me that both beginners and experienced data analysts can benefit from a more formal consideration of how they do data analysis. Indeed, when I began this project, I thought that my workflow was pretty good and that it was simply a matter of writing down what I routinely do. I was surprised and pleased by how much my workflow improved as a result of thinking about these issues systematically and from exchanging ideas with other researchers. Everyone can improve their workflow with relatively little effort. Even though changing your workflow involves an investment of time, you will recoup this investment by saving time in later work and by avoiding errors in your data analysis.

Although I make many specific suggestions about workflow, most of the things that I recommend can be done in other ways. My recommendations about the best practice for a particular problem are based on my work with hundreds of researchers and students from all sectors of employment and from fields ranging from chemistry to history. My suggestions have worked for me and most have been refined with extensive use. This is not to say that there is only one way to accomplish a given task or that I have the best way. In Stata, as in any complex software environment, there are a myriad of ways to complete a task. Some of these work only in the limited sense that they get a job done but are error prone or inefficient. Among the many approaches that work well, you will need to choose your preferred approach. To help you do this, I often discuss several approaches to a given task. I also provide examples of ineffective procedures since seeing the consequences of a misguided approach can be more effective than hearing about the virtues of a better approach. These examples are all real, based on mistakes I made (and

I have made lots) or mistakes I encountered when helping others with data analysis. You will have to choose a workflow that matches the project at hand, the tools you have, and your temperament. There are as many workflows as there are people doing data analysis, and there is no single workflow that is ideal for every person or every project. What is critical is that you consider the general issues, choose your own procedures, and stick with them unless you have a good reason to change them.

In the rest of this chapter, I provide a framework for understanding and evaluating your workflow. I begin with the fundamental principle of replicability that should guide every aspect of your workflow. No matter how you proceed in data analysis, you must be able to justify and reproduce your results. Next I consider the four steps involved in all types of data analysis: preparing data, running analysis, presenting results, and preserving your work. Within each step there are four major tasks: planning the work, organizing your files and materials, documenting what you have done, and executing the analysis. Since there are alternative approaches to accomplish any given aspect of your work, what makes one workflow better than another? To answer this question, I provide several criteria for evaluating the way you work. These criteria should help you decide which procedures to use, and they motivate many of my recommendations for best practice that are given in this book.

1.1 Replication: The guiding principle for workflow

Being able to reproduce the work you have presented or published should be the cornerstone of any workflow. Science demands replicability and a good workflow facilitates your ability to replicate your results. How you plan your project, document your work, write your programs, and save your results should anticipate the need to replicate. Too often researchers do not worry about replication until their work is challenged. This is not to say that they are taking shortcuts, doing shoddy work, or making decisions that are unjustified. Rather, I am talking about taking the steps necessary so that all the good work that has been done can be easily reproduced at a later time. For example, suppose that a colleague wants to expand upon your work and asks you for the data and commands used to produce results in a published paper. When this happens, you do not want to scramble furiously to replicate your results. Although it might take a few hours to dig out your results (many of mine are in notebooks stacked behind my file cabinets), this should be a matter of retrieving the records, not trying to remember what it was you did or discovering that what you documented does not correspond to what you presented.

Think about replication throughout your workflow. At the completion of each stage of your work, take an hour or a day if necessary to review what you have done, to check that the procedures are documented, and to confirm that the materials are archived. When you have a draft of a paper to circulate, review the documentation, check that you still have the files you used, confirm that the do-files still run, and double check that the numbers in your paper correspond to those in your output. Finally, make sure that all of this is documented in your research log (discussed on page 37).

If you have tried to replicate your own work months after it was completed or tried to reproduce another author's results using only the original dataset and the published paper, you know how difficult it can be to replicate something. A good way to understand what is required to replicate your work is to consider some of the things that can make replication impossible. Many of these issues are discussed in detail later in the book. First, you have to find the original files, which gets more difficult as time passes. Once you have the files, are they in formats that can be analyzed by your current software? If you can read the file, do you know exactly how variables were constructed or cases were selected? Do you know which variables were in each regression model? Even if you have all of this information, it is possible that the software you are currently using does not compute things exactly the same way as the software you used for the original analyses. An effective workflow can make replication easier.

A recent example, illustrates how difficult it can be to replicate even simple analyses. I collected some data that were analyzed by a colleague in a published paper. I wanted to replicate his results to extend the analyses. Due to a drive failure, some of his files were lost. Neither of us could reproduce the exact results from the published paper. We came close, but not close enough. Why? Suppose that 10 decisions were made in the process of constructing the variables and selecting the sample for analysis. Many of these decisions involve choices between options where neither choice is incorrect. For example, do you take the square root of publications or the log after adding .5? With 10 such decisions, there are $2^{10} = 1,024$ different outcomes. All of them will lead to similar findings, but not exactly the same findings. If you lose track of decisions made in constructing your data, you will find it very difficult to reproduce what you have done. By the way, remarkably another researcher who was using these data discovered the secret to reproducing the published results.

Even if you have the original data and analysis files, it can be difficult to reproduce results. For published papers, it is often impossible to obtain the original data or the details on how the results were computed. Freese (2007) makes a compelling argument for why disciplines should have policies that govern the availability of information needed to replicate results. I fully support his recommendations.

1.2 Steps in the workflow

Data analysis involves four major steps: cleaning data, performing analysis, presenting findings, and saving your work. Although there is a logical sequence to these steps, the dynamics of an effective workflow are flexible and highly dependent upon the specific project. Ideally, you advance one step at a time, always moving forward until you are done. But, it never works that way for me. In practice, I move up and down the steps depending on how the work goes. Perhaps I find a problem with a variable while analyzing the data, which takes me back to cleaning. Or my results provide unexpected insights, so I revise my plans for analysis. Still, I find it useful to think of these as distinct steps.

1.2.1 Cleaning data

Before substantive analysis begins, you need to verify that your data are accurate and that the variables are well named and properly labeled. That is, you clean the data. First, you must bring your data into Stata. If you received the data in Stata format, this is as simple as a single `use` command. If the data arrived in another format, you need to verify that they were imported correctly into Stata. You should also evaluate the variable names and labels. Awkward names make it more difficult to analyze the data and can lead to mistakes. Likewise, incomplete or poorly designed labels make the output difficult to read and lead to mistakes. Next verify that the sample and variables are what they should be. Do the variables have the correct values? Are missing data coded appropriately? Are the data internally consistent? Is the sample size correct? Do the variables have the distribution that you would expect? Once these questions are resolved, you can select the sample and construct new variables needed for analysis.

1.2.2 Running analysis

Once the data are cleaned, fitting your models and computing the graphs and tables for your paper or book are often the simplest part of the workflow. Indeed, this part of the book is relatively short. Although I do not discuss specific types of analysis, later I talk about ways to ensure the accuracy of your results, to facilitate later replications, and to keep track of your do-files, data files, and log files regardless of the statistical methods you are using.

1.2.3 Presenting results

Once the analyses are complete, you want to present them. I consider several issues in the workflow of presentation. First, you need to move the results from your Stata output into your paper or presentation. An efficient workflow can automate much of this work. Second, you need to document the provenance of all findings that you present. If your presentation does not preserve the source of your results, it can be very difficult to track them down later (e.g., someone is trying to replicate your results or you must respond to a reviewer). Finally, there are a number of simple things that you can do to make your presentations more effective.

1.2.4 Protecting files

When you are cleaning your data, running analyses, and writing, you need to protect your files to prevent loss due to hardware failure, file corruption, or unintentional deletions. Nobody enjoys redoing analyses or rewriting a paper because a file was lost. There are a number of simple things you can do to make it easier to routinely save your work. With backup software readily available and the cost of disk storage so cheap, the hardest parts of making backups is keeping track of what you have. Archiving is distinct from backing up and more difficult since it involves the long-term preservation

of files so that they will be accessible years into the future. You need to consider if the file formats and storage media will be accessible in the future. You must also consider the operating system you use (it is now difficult to read data stored using the CP/M operating system), the storage media (can you read 5 1/4" floppy disks from the 1980s or even a ZIP disk from a few years ago?), natural disasters, and hackers.

1.3 Tasks within each step

Within each of the four major steps, there are four primary tasks: planning your work, organizing your materials, documenting what you do, and executing the work. While some tasks are more important within particular steps (e.g., organization while planning), each task is important for all steps of the workflow.

1.3.1 Planning

Most of us spend too little time planning and too much time working. Before you load data into Stata, you should draft a plan of what you want to do and assess your priorities. What types of analyses are needed? How will you handle missing data? What new variables need to be constructed? As your work progresses, periodically reassess your plan by refining your goals and analytic strategy based on the work you have completed. A little planning goes a long way, and I almost always find that planning saves time.

1.3.2 Organization

Careful organization helps you work faster. Organization is driven by the need to find things and to avoid duplication of effort. Good organization can prevent you from searching for lost files or, worse yet, having to reconstruct them. If you have good documentation about what you did, but you cannot find the files used to do the work, little is gained. Organization requires you to think systematically about how you name files and variables, how you organize directories on your hard drive, how you keep track of which computer has what information (if you use more than one computer), and where you store research materials. Problems with organization show up when you have not been working on a project for a while or when you need something quickly. Throughout the book, I make suggestions on how to organize materials and discuss tools that make it easier to find and work with what you have.

1.3.3 Documentation

Without adequate documentation, replication is virtually impossible, mistakes are more likely, and work usually takes longer. Documentation includes a research log that records what you do and codebooks that document the datasets you create and the variables they contain. Complete documentation also requires comments in your do-files and

labels and notes within data files. Although I find writing documentation to be an onerous task, certainly the least enjoyable part of data analysis, I have learned that time spent on documentation can literally save weeks of work and frustration later. Although there is no way to avoid time spent writing documentation, I can suggest things that make documenting your work faster and more effective.

1.3.4 Execution

Execution involves carrying out specific tasks within each step. Effective execution requires the right tools for the job. A simple example is the editor used to write your programs. Mastering a good text editor can save you hours when writing your programs and will lead to programs that are better written. Another example is learning the most effective commands in Stata. A few minutes spent learning how to use the `recode` command can save you hours of writing `replace` commands. Much of this book involves selecting the right tool for the job. Throughout my discussion of tools, I emphasize standardizing tasks and automating them. The reason to standardize is that it is generally faster to do something the way you did it before than it is to think up a new way to do it. If you set up templates for common tasks, your work becomes more uniform, which makes it easier to find and avoid errors. Efficient execution requires assessing the trade-off between investing the time in learning a new tool, the accuracy gained by the new tools, and the time you save by being more efficient.

1.4 Criteria for choosing a workflow

As you work on the various tasks in each step of your workflow, you will have choices of different ways to do things. How do you decide which procedure to use? In this section, I consider several criteria for evaluating your current workflow and choosing from among alternative procedures for your work.

1.4.1 Accuracy

Getting the correct answer is the *sine qua non* of a good workflow. Oliveira and Stewart (2006, 30) make the point very well, “If your program is not correct, then nothing else matters.” At each step in your work, you must verify that your results are correct. Are you answering the question you set out to answer? Are your results what you wanted and what you think they are? A good workflow is also about making mistakes. Invariably, mistakes will happen, probably a lot of them. Although an effective workflow can prevent some errors, it should also help you find and correct them quickly.

1.4.2 Efficiency

You want to get your analyses done as quickly as possible, given the need for accuracy and replicability. There is an unavoidable tension between getting your work done and

the need to work carefully. If you spend so much time verifying and documenting your work that you never finish the project, you do not have a viable workflow. On the other hand, if you finish by publishing incorrect results, both you and your field suffer. You want a workflow that gets things done as quickly as possible without sacrificing the accuracy of your results. A good workflow, in effect, increases the time you have to do your work, without sacrificing the accuracy of what you do.

1.4.3 Simplicity

A simpler workflow is better than a more complex workflow. The more complicated your procedures, the more likely you will make mistakes or abandon your plan. But what is simple for one person might not be simple for another. Many of the procedures that I recommend involve programming methods that may be new to you. If you have never used a loop, you might find my suggestion of using a loop much more complex than repeating the same commands for multiple variables. With experience, however, you might decide that loops are the simplest way to work.

1.4.4 Standardization

Standardization makes things easier because you do not have to repeatedly decide how to do things and you will be familiar with how things look. When you use standardized formats and procedures, it is easier to see when something is wrong and ensure that you do things consistently the next time. For example, my do-files all use the same structure for organizing the commands. Accordingly, when I look at the log file, it is easier for me to find what I want. Whenever you do something repeatedly, consider creating a template and establishing conventions that become part of your routine workflow.

1.4.5 Automated

Procedures that are automated are better because you are less likely to make mistakes. Entering numbers into your do-file by hand is more error prone than using programming tools to transfer the information automatically. Typing the same list of variables multiple times in a do-file makes it easy to create lists that are supposed to be the same but are not. Again automation can eliminate this problem. Automation is the backbone for many of the methods recommended in this book.

1.4.6 Usability

Your workflow should reflect the way you like to work. If you set up a workflow and then ignore it, you do not have a good workflow. Anything that increases the chances of maintaining your workflow is helpful. Sometimes it is better to use a less efficient approach that is also more enjoyable. For example, I like experimenting with software and prefer taking longer to complete a task while learning a new program than getting

things done quicker the old way. On the other hand, I have a colleague who prefers using a familiar tool even if it takes a bit longer to complete the task. Both approaches make for a good workflow since they compliment our individual styles of work.

1.4.7 Scalability

Some ways of work are fine for small jobs but do not work well for larger jobs. Consider the simple problem of alphabetizing ten articles by author. The easiest approach is to lay the papers on a table and pick them up in order. This works well for ten articles but is dreadfully slow with 100 or 1,000 articles. This issue is referred to as *scalability*—how well do procedures work when applied to a larger problem? As you develop your workflow, think about how well the tools and practices you develop can be applied to a larger project. An effective workflow for a small project where you are the only researcher might not be sustainable for a large project involving many people. Although you can visually inspect every case for every variable in a dataset with 25 measures of development in 80 countries, this approach does not work with the National Longitudinal Survey that has thousands of cases and thousands of variables. You should strive for a workflow that adapts easily to different types of projects. Few procedures scale perfectly. As a consequence you are likely to need different workflows for projects of different complexities.

1.5 Changing your workflow

This book has hundreds of suggestions. Decide which suggestions will help you the most and adapt them to the way you work. Suggestions for minor changes to your workflow can be adopted at any time. For example, it only takes a few minutes to learn how to use `notes`, and you can benefit from this command almost immediately. Other suggestions might require major changes to how you work and should only be made when you have the time to fully integrate them into your work. It is a bad idea to make major changes when a deadline is looming. On the other hand, make sure you find time to improve your workflow. Time spent improving your workflow should save time in the long run and improve the quality of your work. An effective workflow is something that evolves over time, reflecting your experience, changing technology, your personality, and the nature of your current research.

1.6 How the book is organized

This book is organized so that it can be read front to back by someone wanting to learn about the entire workflow of data analysis. I also wanted it to be useful as a reference for people who encounter a problem and who want a specific solution. For this purpose, I have tried to make the index and table of contents extensive. It is also useful to understand the overall structure of this book before you proceed with your reading.

Chapter 2 – Planning, organizing, and documenting your work discusses how to plan your work, organize your files, and document what you have done. Avoid the temptation of skipping this chapter so that you can get to the “important” details in later chapters.

Chapter 3 – Writing and debugging do-files discusses how do-files should be used for almost all of your work in Stata. I provide information on how to write more effective do-files and how to debug programs that do not work. Both beginners and advanced users should find useful information here.

Chapter 4 – Automating Stata is an introduction to programming that discusses how to create macros, run loops, and write short programs. This chapter is not intended to teach you how to write sophisticated programs in Stata (although it might be a good introduction), rather it discusses tools that all data analysts should find useful. I encourage every reader to study the material in this chapter before reading chapters 5–7.

Chapter 5 – Names and labels discusses both principles and tools for creating names and labels that are clear and consistent. Even if you have received data that is labeled, you should consider improving the names and labels in the dataset. This chapter is long and includes a lot of technical details that you can skip until you need them.

Chapter 6 – Cleaning data and constructing variables discusses how to check whether your data is correct and how to construct new variables and verify that they were created correctly. At least 80% of the work in data analysis involves getting the data ready, so this chapter is essential.

Chapter 7 – Analyzing, presenting, and replicating results discusses how to keep track of the analyses used in presentations and papers, issues to consider when presenting your results, and ways to make replication simpler.

Chapter 8 – Saving your work discusses how to back up and archive your work. This seemingly simple task is often frustratingly difficult and involves subtle problems that can be easy to overlook.

Chapter 9 – Conclusions draws general conclusions about workflow.

Appendix A reviews how the Stata program operates; considers working with a networked version of Stata, such as that found in many computer labs; explains how to install user-written programs, such as the *Workflow* package; and shows you how to customize the way in which Stata works.

Additional information about workflow, including examples and discussion of other software, is available at <http://www.indiana.edu/~jslsoc/workflow.htm>.