

# Maximum Likelihood Estimation with Stata

Fifth Edition

JEFFREY PITBLADO  
*StataCorp LLC*

BRIAN POI  
*Poi Consulting LLC*

WILLIAM GOULD  
*StataCorp LLC*



**stata**® *Press*

A Stata Press Publication  
StataCorp LLC  
College Station, Texas



Copyright © 1999, 2003, 2006, 2010, 2024 StataCorp LLC  
All rights reserved. First edition 1999  
Second edition 2003  
Third edition 2006  
Fourth edition 2010  
Fifth edition 2024

Published by Stata Press, 4905 Lakeway Drive, College Station, Texas 77845

Typeset in L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub>

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

Print ISBN-10: 1-59718-411-X

Print ISBN-13: 978-1-59718-411-3

ePub ISBN-10: 1-59718-412-8

ePub ISBN-13: 978-1-59718-412-0

Library of Congress Control Number: 2023947523

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 LLC.

Stata, **stata**, Stata Press, Mata, **mata**, and NetCourse are registered trademarks of StataCorp LLC.

Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations.

NetCourseNow is a trademark of StataCorp LLC.

L<sup>A</sup>T<sub>E</sub>X 2<sub>ε</sub> is a trademark of the American Mathematical Society.

Other brand and product names are registered trademarks or trademarks of their respective companies.

# Contents

	List of tables	xiii
	List of figures	xv
	Preface to the fifth edition	xvii
	Versions of Stata	xix
	Notation and typography	xxi
<b>1</b>	<b>Theory and practice</b>	<b>1</b>
1.1	The likelihood-maximization problem . . . . .	3
1.2	Likelihood theory . . . . .	5
1.2.1	All results are asymptotic . . . . .	9
1.2.2	Likelihood-ratio tests and Wald tests . . . . .	10
1.2.3	The outer product of gradients variance estimator . . . . .	11
1.2.4	Robust variance estimates . . . . .	12
1.3	The maximization problem . . . . .	14
1.3.1	Numerical root finding . . . . .	14
	Newton's method . . . . .	14
	The Newton–Raphson algorithm . . . . .	16
1.3.2	Quasi-Newton methods . . . . .	18
	The BHHH algorithm . . . . .	19
	The DFP and BFGS algorithms . . . . .	19
1.3.3	Numerical maximization . . . . .	20
1.3.4	Numerical derivatives . . . . .	21
1.3.5	Numerical second derivatives . . . . .	25
1.4	Monitoring convergence . . . . .	26

<b>2</b>	<b>Estimation with <code>mlexp</code></b>	<b>29</b>
2.1	Syntax . . . . .	29
2.2	Normal linear regression . . . . .	30
2.3	Initial values . . . . .	31
2.4	Restricted parameters . . . . .	33
2.5	Robust standard errors . . . . .	35
2.6	The probit model . . . . .	36
2.7	Specifying derivatives . . . . .	38
2.8	Additional estimation features . . . . .	40
2.9	Wrapping up . . . . .	42
<b>3</b>	<b>Introduction to <code>ml</code></b>	<b>43</b>
3.1	The probit model . . . . .	43
3.2	Normal linear regression . . . . .	46
3.3	Robust standard errors . . . . .	48
3.4	Weighted estimation . . . . .	49
3.5	Other features of <code>method-gf0</code> evaluators . . . . .	49
3.6	Limitations . . . . .	50
<b>4</b>	<b>Overview of <code>ml</code></b>	<b>53</b>
4.1	The terminology of <code>ml</code> . . . . .	53
4.2	Equations in <code>ml</code> . . . . .	54
4.3	Likelihood-evaluator methods . . . . .	62
4.4	Tools for the <code>ml</code> programmer . . . . .	65
4.5	Common <code>ml</code> options . . . . .	65
4.5.1	Subsamples . . . . .	66
4.5.2	Weights . . . . .	67
4.5.3	OPG estimates of variance . . . . .	68
4.5.4	Robust estimates of variance . . . . .	69
4.5.5	Survey data . . . . .	70
4.5.6	Constraints . . . . .	71
4.5.7	Choosing among the optimization algorithms . . . . .	71

<i>Contents</i>	vii	
4.6	Maximizing your own likelihood functions . . . . .	75
4.7	Appendix: More about scalar parameters . . . . .	76
<b>5</b>	<b>Method lf</b>	<b>79</b>
5.1	The linear-form restrictions . . . . .	80
5.2	Examples . . . . .	81
5.2.1	The probit model . . . . .	81
5.2.2	Normal linear regression . . . . .	83
5.2.3	The Weibull model . . . . .	85
5.3	The importance of generating temporary variables as doubles . . . . .	87
5.4	Problems you can safely ignore . . . . .	89
5.5	Nonlinear specifications . . . . .	90
5.6	The advantages of lf in terms of execution speed . . . . .	91
<b>6</b>	<b>Methods lf0, lf1, and lf2</b>	<b>95</b>
6.1	Comparing these methods . . . . .	95
6.2	Outline of evaluators of methods lf0, lf1, and lf2 . . . . .	96
6.2.1	The todo argument . . . . .	97
6.2.2	The b argument . . . . .	97
	Using mlevel to obtain equation and free-parameter values . . . . .	99
6.2.3	The lnfj argument . . . . .	101
6.2.4	Arguments for scores . . . . .	102
6.2.5	The H argument . . . . .	103
	Using mlmatsum to define H . . . . .	105
6.2.6	Aside: Stata's scalars . . . . .	107
6.3	Summary of methods lf0, lf1, and lf2 . . . . .	110
6.3.1	Method lf0 . . . . .	110
6.3.2	Method lf1 . . . . .	111
6.3.3	Method lf2 . . . . .	113
6.4	Examples . . . . .	115
6.4.1	The probit model . . . . .	115
6.4.2	Normal linear regression . . . . .	117

6.4.3	The Weibull model . . . . .	124
<b>7</b>	<b>Methods d0, d1, and d2</b>	<b>129</b>
7.1	Comparing these methods . . . . .	129
7.2	Outline of method d0, d1, and d2 evaluators . . . . .	130
7.2.1	The todo argument . . . . .	131
7.2.2	The b argument . . . . .	131
7.2.3	The lnf argument . . . . .	132
	Using lnf to indicate that the likelihood cannot be calculated	133
	Using mlsum to define lnf . . . . .	134
7.2.4	The g argument . . . . .	136
	Using mlvecsum to define g . . . . .	136
7.2.5	The H argument . . . . .	138
7.3	Summary of methods d0, d1, and d2 . . . . .	139
7.3.1	Method d0 . . . . .	139
7.3.2	Method d1 . . . . .	142
7.3.3	Method d2 . . . . .	144
7.4	Panel-data likelihoods . . . . .	146
7.4.1	Calculating lnf . . . . .	148
7.4.2	Calculating g . . . . .	152
7.4.3	Calculating H . . . . .	156
	Using mlmatbysum to help define H . . . . .	156
7.5	Other models that do not meet the linear-form restrictions . . . . .	164
<b>8</b>	<b>Debugging likelihood evaluators</b>	<b>171</b>
8.1	ml check . . . . .	171
8.2	Using the debug methods . . . . .	173
8.2.1	First derivatives . . . . .	175
8.2.2	Second derivatives . . . . .	185
8.3	ml trace . . . . .	188
<b>9</b>	<b>Setting initial values</b>	<b>191</b>
9.1	ml search . . . . .	192

9.2	ml plot . . . . .	195
9.3	ml init . . . . .	197
<b>10</b>	<b>Interactive maximization</b>	<b>201</b>
10.1	The iteration log . . . . .	201
10.2	Pressing the Break key . . . . .	202
10.3	Maximizing difficult likelihood functions . . . . .	204
<b>11</b>	<b>Final results</b>	<b>207</b>
11.1	Graphing convergence . . . . .	207
11.2	Redisplaying output . . . . .	209
<b>12</b>	<b>Writing do-files to maximize likelihoods</b>	<b>213</b>
12.1	The structure of a do-file . . . . .	213
12.2	Putting the do-file into production . . . . .	214
<b>13</b>	<b>Writing ado-files to maximize likelihoods</b>	<b>217</b>
13.1	Writing estimation commands . . . . .	217
13.2	The standard estimation-command outline . . . . .	219
13.3	Outline for estimation commands using ml . . . . .	220
13.4	Using ml in noninteractive mode . . . . .	221
	13.4.1 Parsing with help from <code>_get_diopts</code> . . . . .	222
13.5	Advice . . . . .	223
	13.5.1 Syntax . . . . .	224
	13.5.2 Estimation subsample . . . . .	226
	13.5.3 Parsing with help from <code>mlopts</code> . . . . .	230
	13.5.4 Weights . . . . .	233
	13.5.5 Constant-only model . . . . .	234
	13.5.6 Initial values . . . . .	238
	13.5.7 Storing results in <code>e()</code> . . . . .	241
	13.5.8 Displaying ancillary parameters . . . . .	241
	13.5.9 Exponentiated coefficients . . . . .	243
	13.5.10 Offsetting linear equations . . . . .	245
	13.5.11 Program properties . . . . .	247

<b>14</b>	<b>Writing ado-files for survey data analysis</b>	<b>251</b>
14.1	Program properties . . . . .	251
14.2	Writing your own predict command . . . . .	254
<b>15</b>	<b>Mata-based likelihood evaluators</b>	<b>257</b>
15.1	Introductory examples . . . . .	257
15.1.1	The probit model . . . . .	257
15.1.2	The Weibull model . . . . .	260
15.2	Evaluator function prototypes . . . . .	262
	Method-lf evaluators . . . . .	263
	lf-family evaluators . . . . .	263
	d-family evaluators . . . . .	264
15.3	Utilities . . . . .	265
	Dependent variables . . . . .	266
	Obtaining model parameters . . . . .	266
	Summing individual or group-level log likelihoods . . . . .	267
	Calculating the gradient vector . . . . .	267
	Calculating the Hessian . . . . .	268
15.4	Random-effects linear regression . . . . .	269
15.4.1	Calculating lnf . . . . .	270
15.4.2	Calculating g . . . . .	271
15.4.3	Calculating H . . . . .	272
15.4.4	Results at last . . . . .	273
15.5	Ado-file considerations . . . . .	276
<b>16</b>	<b>Mata's moptimize() function</b>	<b>279</b>
16.1	Introductory examples . . . . .	280
16.1.1	The probit model . . . . .	280
16.1.2	The Weibull model . . . . .	283
16.2	Restricting the estimation sample . . . . .	286
16.2.1	Using moptimize_init_touse() . . . . .	286
16.2.2	Not using moptimize_init_touse() . . . . .	287



16.3	Estimation preliminaries . . . . .	288
16.3.1	Weights . . . . .	289
16.3.2	Panel data and clusters . . . . .	289
16.3.3	Survey data . . . . .	290
16.3.4	Initial values . . . . .	290
16.4	Estimation . . . . .	291
16.5	Results . . . . .	292
16.5.1	Displaying results . . . . .	292
16.5.2	Retrieving results . . . . .	293
16.5.3	Storing results in <code>e()</code> . . . . .	293
16.6	Estimation commands . . . . .	294
16.6.1	Common maximization options . . . . .	295
16.6.2	Initial values . . . . .	296
16.6.3	Constraints . . . . .	299
16.7	Regression redux . . . . .	302
<b>17</b>	<b>Other examples</b>	<b>309</b>
17.1	The logit model . . . . .	309
17.2	The probit model . . . . .	311
17.3	Normal linear regression . . . . .	313
17.4	The Weibull model . . . . .	316
17.5	The Cox proportional hazards model . . . . .	319
17.6	The random-effects regression model . . . . .	322
17.7	The seemingly unrelated regression model . . . . .	325
17.8	A bivariate Poisson regression model . . . . .	338
17.8.1	A bivariate Poisson distribution . . . . .	338
17.8.2	Bivariate Poisson regression . . . . .	347
17.8.3	Discussion . . . . .	356
17.9	Epilogue . . . . .	357
<b>A</b>	<b>Syntax of <code>mlexp</code></b>	<b>359</b>
<b>B</b>	<b>Syntax of <code>ml</code></b>	<b>365</b>

<b>C</b>	<b>Syntax of moptimize()</b>	<b>391</b>
<b>D</b>	<b>Likelihood-evaluator checklists</b>	<b>419</b>
D.1	Method lf . . . . .	419
D.2	Method d0 . . . . .	420
D.3	Method d1 . . . . .	421
D.4	Method d2 . . . . .	423
D.5	Method lf0 . . . . .	426
D.6	Method lf1 . . . . .	427
D.7	Method lf2 . . . . .	429
<b>E</b>	<b>Listing of estimation commands</b>	<b>433</b>
E.1	The logit model . . . . .	433
E.2	The probit model . . . . .	435
E.3	The normal model . . . . .	437
E.4	The Weibull model . . . . .	440
E.5	The Cox proportional hazards model . . . . .	442
E.6	The random-effects regression model . . . . .	444
E.7	The seemingly unrelated regression model . . . . .	448
E.8	A bivariate Poisson regression model . . . . .	455
	<b>References</b>	<b>463</b>
	<b>Author index</b>	<b>467</b>
	<b>Subject index</b>	<b>469</b>

*(Pages omitted)*

# Preface to the fifth edition

*Maximum Likelihood Estimation with Stata, Fifth Edition* is written for researchers in all disciplines who need to compute maximum likelihood estimators that are not available as prepackaged routines. To get the most from this book, you should be familiar with Stata, but you will not need any special programming skills, except in chapters 13 and 14, which detail how to take an estimation technique you have written and add it as a new *command* to Stata. No special theoretical knowledge is needed either, other than an understanding of the likelihood function that will be maximized.

Like the rest of Stata, the tools one uses to implement maximum likelihood estimators in Stata have undergone many enhancements over the years, and a new version of this book reflecting those changes is warranted. The core of the book continues to focus on the `ml` suite of commands. We have also added a new chapter for the `mlexp` command, which is useful not only for pedagogical and prototyping purposes but also for implementing relatively simple estimators with zero programming. For those who are familiar with Mata, Stata's matrix programming language, we have also included a new chapter describing the `moptimize()` suite of functions for implementing maximum likelihood estimators entirely within Mata.

Chapter 1 provides a general overview of maximum likelihood estimation theory and numerical optimization methods, with an emphasis on the practical implications of each for applied work. Chapter 2 covers the `mlexp` command for implementing relatively simple estimators with no programming skills required. Chapter 3 is an introduction to the `ml` command, which provides substantially more flexibility than `mlexp` and can be used to implement arbitrarily complex maximum-likelihood estimators. Chapter 4 is an overview of the `ml` command and the notation used throughout the rest of the book. Chapters 5–11 detail, step by step, how to use Stata to maximize user-written likelihood functions. Chapter 12 describes how to package all the user-written code in a do-file so that it can be conveniently reapplied to different datasets and model specifications. Chapter 13 details how to structure the code in an ado-file to create a new Stata estimation command. Chapter 14 shows how to add survey estimation features to existing `ml`-based estimation commands.

Chapters 15 and 16 are more advanced and show how to use Mata to implement maximum likelihood estimators. Chapter 15 shows how to write your likelihood evaluator in Mata while continuing to use the `ml` command to specify your model, maximize the likelihood function, and report results. Chapter 16 shows how to implement an estimator using Mata's `moptimize()` function and bypass `ml` altogether.

Chapter 17, the final chapter, provides examples. For a set of estimation problems, we derive the log-likelihood function, show the derivatives that make up the gradient and Hessian, write one or more likelihood-evaluation programs, and so provide a fully functional estimation command. We use the estimation command to fit the model to a dataset. An estimation command is developed for each of the following:

- Logit and probit models
- Linear regression
- Weibull regression
- Cox proportional hazards model
- Random-effects linear regression for panel data
- Seemingly unrelated regression
- Bivariate Poisson regression

Appendices contain full syntax diagrams for all the `m1` subroutines, useful checklists for implementing each maximization method, and program listings of each estimation command covered in chapter 17.

We acknowledge William Sribney as one of the original developers of `m1` and the principal author of the first edition of this book.

College Station, TX  
October 2023

Jeffrey Pitblado  
Brian Poi  
William Gould

# Versions of Stata

This book was written for Stata 18. Regardless of what version of Stata you are using, verify that your copy of Stata is up to date and obtain any free updates; to do this, enter Stata, type

```
. update query
```

and follow the instructions.

Having done that, if you are still using a version older than 18—such as Stata 16—you will likely run into compatibility issues with some of the code and examples in this book. In that case, you should purchase an upgrade to Stata before continuing. We will assume that you are running Stata 18 or perhaps an even newer version.

All the programs in this book follow the outline

```
program myprog
  version 18
  ...
end
```

Because Stata 18 is the current release of Stata at the time this book was written, we write `version 18` at the top of our programs. You could omit the line, but we recommend that you include it because Stata is continually being developed and sometimes details of syntax change. Placing `version 18` at the top of your program tells Stata that, if anything has changed, you want the version 18 interpretation.

Coding `version 18` at the top of your programs ensures they will continue to work in the future.

What about programs you write in the future? Perhaps the here and now for you is Stata 19 or Stata 20. Using this book, should you put `version 18` at the top of your programs, or should you put `version 19` or `version 20`? Probably, you should substitute the more modern version number. The only reason you would not want to make the substitution is that the syntax of `ml` itself has changed and you want to use the version of syntax described in this book.

Anyway, if you are using a version more recent than 18, type `help whatsnew` to see a complete listing of what has changed. That will help you decide what to code at the top of your programs: unless the listing clearly states that `ml`'s syntax has changed, substitute the more recent version number.

*(Pages omitted)*

## 3 Introduction to ml

`ml` is the Stata command to implement maximum likelihood (ML) estimators that cannot be handled by `mlexp` and to write your own commands that perform ML estimation. Obtaining ML estimates requires the following steps:

1. Derive the log-likelihood function from your probability model.
2. Write a program that calculates the log-likelihood values and, optionally, its derivatives. This program is known as a likelihood evaluator.
3. Identify a particular model to fit using your data variables and the `ml model` statement.
4. Fit the model using `ml maximize`.

This chapter illustrates steps 2, 3, and 4 using the probit model for dichotomous (0/1) variables and the linear regression model assuming normally distributed errors.

In this chapter, we fit our models explicitly, handling each coefficient and variable individually. New users of `ml` will appreciate this approach because it closely reflects how you would write down the model you wish to fit on paper; and it allows us to focus on some of the basic features of `ml` without becoming overly encumbered with programming details. We will also illustrate this strategy's shortcomings so that once you become familiar with the basics of `ml` by reading this chapter, you will want to think of your model in a slightly more abstract form, providing much more flexibility.

In the next chapter, we discuss `ml`'s probability model parameter notation, which is particularly useful when, as is inevitably the case, you decide to change some of the variables appearing in your model. If you are already familiar with `ml`'s  $\theta$ -parameter notation, you can skip this chapter with virtually no loss of continuity with the rest of the book.

Chapter 17 contains the derivations of log-likelihood functions (step 1) for models discussed in this book.

### 3.1 The probit model

In section 2.6, we fit a probit model to predict whether a car is foreign or domestic based on its weight and price using `auto.dta`. Here we illustrate how to fit that model using `ml`. Recall that our statistical model is



$$\begin{aligned}\pi_j &= \Pr(\text{foreign}_j \mid \text{weight}_j, \text{price}_j) \\ &= \Phi(\beta_1 \text{weight}_j + \beta_2 \text{price}_j + \beta_0)\end{aligned}$$

where we use the subscript  $j$  to denote observations and  $\Phi(\cdot)$  denotes the standard normal distribution function. The log likelihood for the  $j$ th observation is

$$\begin{aligned}\ln \ell_j &= \begin{cases} \ln \Phi(\beta_1 \text{weight}_j + \beta_2 \text{price}_j + \beta_0) & \text{if } \text{foreign}_j = 1 \\ 1 - \ln \Phi(\beta_1 \text{weight}_j + \beta_2 \text{price}_j + \beta_0) & \text{if } \text{foreign}_j = 0 \end{cases} \\ &= \begin{cases} \ln \Phi(\beta_1 \text{weight}_j + \beta_2 \text{price}_j + \beta_0) & \text{if } \text{foreign}_j = 1 \\ \ln \Phi(-\beta_1 \text{weight}_j - \beta_2 \text{price}_j - \beta_0) & \text{if } \text{foreign}_j = 0 \end{cases} \quad (3.1)\end{aligned}$$

where we used the fact that  $1 - \Phi(w) = \Phi(-w)$ .

With our log-likelihood function in hand, we write a program to evaluate it:

```

----- begin myprobit_gf0.ado -----
program myprobit_gf0
  version 18
  args todo b lnfj
  tempvar xb
  quietly generate double `xb' = `b'[1,1]*weight + `b'[1,2]*price + ///
    `b'[1,3]
  quietly replace `lnfj' = ln(normal(`xb')) if foreign == 1
  quietly replace `lnfj' = ln(normal(-1*`xb')) if foreign == 0
end
----- end myprobit_gf0.ado -----

```

We named our program `myprobit_gf0.ado`, but you could name it anything you want as long as it has the extension `.ado`. The name without the `.ado` extension is what we use to tell `ml model` about our likelihood function. We added `gf0` to our name to emphasize that our evaluator is a general-form problem and that we are going to specify no (0) derivatives. We will return to this issue when we use the `ml model` statement.

Our program accepts three arguments. The first, `todo`, we can safely ignore for now. In later chapters, when we discuss other types of likelihood-evaluator programs, we will need that argument. The second, `b`, contains a row vector containing the parameters of our model ( $\beta_0$ ,  $\beta_1$ , and  $\beta_2$ ). The third argument, `lnfj`, is the name of a temporary variable that we are to fill in with the values of the log-likelihood function evaluated at the coefficient vector `b`. Our program then created a temporary variable to hold the values of  $\beta_1 \text{weight}_j + \beta_2 \text{price}_j + \beta_0$ . We created that variable to have storage type `double`; we will discuss this point in greater detail in the next chapter, but for now you should remember that when coding your likelihood evaluator, you must create temporary variables as `doubles`. The last two lines replace `lnfj` with the values of the log likelihood for `foreignj` equal to 0 and 1, respectively. We could have used Stata's `cond()` function to combine those two statements into one, but using two statements is arguably clearer. Because `b` and `lnfj` are arguments passed to our program and `xb`

is a temporary variable we created with the `tempvar` commands, their names are local macros that must be dereferenced using left- and right-hand single quote marks to use them; see [U] **18.7 Temporary objects**.

The next step is to identify our model using `ml model`. To that end, we type

```
. sysuse auto
(1978 automobile data)
. ml model gf0 myprobit_gf0 (foreign = weight price)
```

We first loaded in our dataset, because `ml model` will not work without the dataset in memory. Next we told `ml` that we have a `method-gf0` likelihood evaluator named `myprobit_gf0`, our dependent variable is `foreign`, and our independent variables are `weight` and `price`. In subsequent chapters, we examine all the likelihood-evaluator types; `method-gf0` (general form) evaluator programs most closely follow the mathematical notation we used in (3.1) and are therefore perhaps easiest for new users of `ml` to grasp, but we will see that they have disadvantages as well. This general-form evaluator simply receives a vector of parameters and a variable into which the observations' log-likelihood values are to be stored.

The final step is to tell `ml` to maximize the likelihood function and report the coefficients:

```
. ml maximize
Initial:      Log likelihood = -51.292891
Alternative:  Log likelihood = -45.055272
Rescale:     Log likelihood = -45.055272
Iteration 0:  Log likelihood = -45.055272
Iteration 1:  Log likelihood = -20.770386
Iteration 2:  Log likelihood = -18.023563
Iteration 3:  Log likelihood = -18.006584
Iteration 4:  Log likelihood = -18.006571
Iteration 5:  Log likelihood = -18.006571

                                Number of obs =      74
                                Wald chi2(2)  =   14.09
                                Prob > chi2   =  0.0009

Log likelihood = -18.006571
```

foreign	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
weight	-.003238	.0008643	-3.75	0.000	-.004932	-.0015441
price	.000517	.0001591	3.25	0.001	.0002052	.0008287
_cons	4.921935	1.330065	3.70	0.000	2.315055	7.528816

You can verify that we would obtain identical results using `probit`:

```
. probit foreign weight price
```

This example was straightforward because we had only one equation and no auxiliary parameters. Next we consider linear regression with normally distributed errors.

## 3.2 Normal linear regression

Now suppose we want to fit a linear regression of `turn` on `length` and `headroom`:

$$\text{turn}_j = \beta_1 \text{length}_j + \beta_2 \text{headroom}_j + \beta_3 + \epsilon_j$$

where  $\epsilon_j$  is an error term. If we assume that each  $\epsilon_j$  is independent and identically distributed as a normal random variable with mean zero and variance  $\sigma^2$ , we have what is often called normal linear regression; and we can fit the model by ML. As derived in section 17.3, when we assume homoskedasticity, we can write the log likelihood for the  $j$ th observation in terms of  $\phi(z)$ , the standard normal density function, as

$$\ln \ell_j = \ln \phi \left( \frac{\text{turn}_j - \beta_1 \text{length}_j - \beta_2 \text{headroom}_j - \beta_3}{\sigma} \right) - \ln \sigma$$

There are four parameters in our model:  $\beta_1$ ,  $\beta_2$ ,  $\beta_3$ , and  $\sigma$ , so we will specify our `ml model` statement so that our likelihood evaluator receives a vector of coefficients with four columns. As a matter of convention, we will use the four elements of that vector in the order we just listed so that, for example,  $\beta_2$  is the second element and  $\sigma$  is the fourth element. Our likelihood-evaluator program is

```

----- begin mynormal1_gf0.ado -----
program mynormal1_gf0
  version 18
  args todo b lnfj
  tempvar xb
  quietly generate double `xb' = `b'[1,1]*length +          ///
                                `b'[1,2]*headroom + `b'[1,3]
  quietly replace `lnfj' = ln(normalden((turn - `xb')/`b'[1,4])) - ///
                          ln(`b'[1,4])
end
----- end mynormal1_gf0.ado -----

```

In our previous example, when we typed

```
. ml model gf0 myprobit_gf0 (foreign = weight price)
```

`ml` knew to create a coefficient vector with three elements because we specified two right-hand-side variables, and by default `ml` includes a constant term unless we specify the `noconstant` option, which we discuss in the next chapter. How do we get `ml` to include a fourth parameter for  $\sigma$ ? Perhaps the most transparent solution is to type

```
. ml model gf0 mynormal1_gf0 (turn = length headroom) /sigma
```

The notation `/sigma` tells `ml` to include a fourth element in our coefficient vector and to label it `sigma` in the output. In chapter 4, we will see other ways to specify parameters like  $\sigma$  and have them labeled slightly differently. Having identified our model, we can now maximize the log-likelihood function:

```

. ml maximize
Initial:      Log likelihood =      -<inf> (could not be evaluated)
Feasible:    Log likelihood =   -8418.567
Rescale:     Log likelihood =  -327.16314
Rescale eq:  Log likelihood = -215.53986
Iteration 0: Log likelihood = -215.53986 (not concave)
Iteration 1: Log likelihood = -213.33272 (not concave)
Iteration 2: Log likelihood = -211.10519 (not concave)
Iteration 3: Log likelihood = -209.60577 (not concave)
Iteration 4: Log likelihood = -207.93771 (not concave)
Iteration 5: Log likelihood = -206.43844 (not concave)
Iteration 6: Log likelihood = -205.19618 (not concave)
Iteration 7: Log likelihood = -204.11373 (not concave)
Iteration 8: Log likelihood = -203.00329 (not concave)
Iteration 9: Log likelihood = -202.1822 (not concave)
Iteration 10: Log likelihood = -201.42449 (not concave)
Iteration 11: Log likelihood = -200.64468 (not concave)
Iteration 12: Log likelihood = -199.9014 (not concave)
Iteration 13: Log likelihood = -199.18937 (not concave)
Iteration 14: Log likelihood = -198.48172 (not concave)
Iteration 15: Log likelihood = -197.78556 (not concave)
Iteration 16: Log likelihood = -197.10597 (not concave)
Iteration 17: Log likelihood = -196.43819 (not concave)
Iteration 18: Log likelihood = -195.78002 (not concave)
Iteration 19: Log likelihood = -195.13253 (not concave)
Iteration 20: Log likelihood = -194.4956 (not concave)
Iteration 21: Log likelihood = -193.86829 (not concave)
Iteration 22: Log likelihood = -193.2503 (not concave)
Iteration 23: Log likelihood = -192.64164 (not concave)
Iteration 24: Log likelihood = -192.0421 (not concave)
Iteration 25: Log likelihood = -191.45135 (not concave)
Iteration 26: Log likelihood = -190.86936 (not concave)
Iteration 27: Log likelihood = -190.29601 (not concave)
Iteration 28: Log likelihood = -189.73109 (not concave)
Iteration 29: Log likelihood = -189.17464 (not concave)
Iteration 30: Log likelihood = -188.62662
Iteration 31: Log likelihood = -167.31438 (backed up)
Iteration 32: Log likelihood = -163.31365
Iteration 33: Log likelihood = -163.18798
Iteration 34: Log likelihood = -163.18765
Iteration 35: Log likelihood = -163.18765

Number of obs =      74
Wald chi2(2) = 219.18
Prob > chi2 = 0.0000
Log likelihood = -163.18765

```

turn	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
length	.1737846	.0134739	12.90	0.000	.1473762	.2001929
headroom	-.1542078	.3546293	-0.43	0.664	-.8492685	.5408529
_cons	7.450477	2.19735	3.39	0.001	3.143751	11.7572
/sigma	2.195259	.1804492	12.17	0.000	1.841585	2.548933

The point estimates match those we obtain from typing

```
. regress turn length headroom
```

The standard errors differ by a factor of  $\sqrt{71/74}$  because of the same degree-of-freedom adjustment we discussed in section 2.5.

### 3.3 Robust standard errors

Robust standard errors are commonly reported nowadays along with linear regression results because they allow for correct statistical inference even when the tenuous assumption of homoskedasticity is not met. Cluster-robust standard errors can be used when related observations' errors are correlated. Obtaining standard errors with most estimation commands is trivial: you just specify the option `vce(robust)` or `vce(cluster id)`, where *id* is the name of a variable identifying groups. Using our previous regression example, you might type

```
. regress turn length headroom, vce(robust)
```

For the evaluator functions we have written so far, both of which have been method `gf0`, obtaining robust or cluster-robust standard errors is no more difficult than with other estimation commands. To refit our linear regression model, obtaining robust standard errors, we type

```
. ml model gf0 mynormal1_gf0 (turn = length headroom) /sigma, vce(robust)
. ml maximize, nolog
Initial:      Log pseudolikelihood =    -<inf> (could not be evaluated)
Feasible:    Log pseudolikelihood =  -8418.567
Rescale:     Log pseudolikelihood = -327.16314
Rescale eq:  Log pseudolikelihood = -215.53986

                                     Number of obs =      74
                                     Wald chi2(2)  = 298.85
                                     Prob > chi2   = 0.0000

Log pseudolikelihood = -163.18765
```

turn	Robust		z	P> z	[95% conf. interval]	
	Coefficient	std. err.				
length	.1737846	.0107714	16.13	0.000	.152673	.1948961
headroom	-.1542078	.2955882	-0.52	0.602	-.73355	.4251344
_cons	7.450477	1.858003	4.01	0.000	3.808857	11.0921
/sigma	2.195259	.2886184	7.61	0.000	1.629577	2.760941

`ml model` accepts `vce(cluster id)` with method-`gf0` evaluators just as readily as it accepts `vce(robust)`.

Being able to obtain robust standard errors just by specifying an option to `ml model` should titillate you. When we discuss other types of evaluator programs, we will see that in fact there is a lot of work happening behind the scenes to produce robust standard errors. With method-`gf0` evaluators (and other linear-form evaluators), `ml` does all the work for you.

### 3.4 Weighted estimation

Stata provides four types of weights that the end-user can apply to estimation problems. Frequency weights, known as `fweights` in the Stata vernacular, represent duplicated observations; instead of having five observations that record identical information, `fweights` allow you to record that observation once in your dataset along with a frequency weight of 5, indicating that observation is to be repeated a total of five times. Analytic weights, called `awweights`, are inversely proportional to the variance of an observation and are used with group-mean data. Sampling weights, called `pweights`, denote the inverse of the probability that an observation is sampled and are used with survey data where some people are more likely to be sampled than others. Importance weights, called `iweights`, indicate the relative “importance” of the observation and are intended for use by programmers who want to produce a certain computation.

Obtaining weighted estimates with method-gf0 likelihood evaluators is the same as with most other estimation commands. Suppose that in `auto.dta`, `rep78` is actually a frequency weight variable. To obtain frequency-weighted estimates of our probit model, we type

```
. ml model gf0 myprobit_gf0 (foreign = weight price) [fweight = rep78]
. ml maximize
Initial:      Log likelihood = -162.88959
Alternative:  Log likelihood = -159.32929
Rescale:     Log likelihood = -156.55825
Iteration 0: Log likelihood = -156.55825
Iteration 1: Log likelihood = -72.414357
Iteration 2: Log likelihood = -66.82292
Iteration 3: Log likelihood = -66.426129
Iteration 4: Log likelihood = -66.424675
Iteration 5: Log likelihood = -66.424675

                                Number of obs =    235
                                Wald chi2(2)  =   58.94
                                Prob > chi2   =  0.0000

Log likelihood = -66.424675
```

foreign	Coefficient	Std. err.	z	P> z	[95% conf. interval]	
weight	-.0027387	.0003576	-7.66	0.000	-.0034396	-.0020379
price	.0004361	.0000718	6.07	0.000	.0002953	.0005768
_cons	4.386445	.5810931	7.55	0.000	3.247523	5.525366

Just like with obtaining robust standard errors, we did not have to do anything to our likelihood-evaluator program. We just added a weight specification, and `ml` did all the heavy lifting to make that work. You should be impressed. Other evaluator types require you to account for weights yourself, which is not always a trivial task.

### 3.5 Other features of method-gf0 evaluators

In addition to easily obtaining robust standard errors and weighted estimates, method-gf0 likelihood evaluators provide several other features. By specifying the `svy` option

to `ml model`, you can obtain results that take into account the complex survey design of your data. Before using the `svy` option, you must first `svyset` your data; see [U] **26.19 Survey data**.

You can restrict the estimation sample by using `if` and `in` conditions in your `ml model` statement. Again, `method-gf0` evaluators require you to do nothing special to make them work. See [U] **11 Language syntax** to learn about `if` and `in` qualifiers.

## 3.6 Limitations

We have introduced `ml` using `method-gf0` evaluators because they align most closely with the way you would write the likelihood function for a specific model. However, writing your likelihood evaluator in terms of a particular model with prespecified variables severely limits your flexibility.

For example, say that we had a binary variable `good` that we wanted to use instead of `foreign` as the dependent variable in our probit model. If we simply change our `ml model` statement to read

```
. ml model gf0 myprobit_gf0 (good = weight price)
```

the output from `ml maximize` will label the dependent variable as `good`, but the output will otherwise be unchanged! When we wrote our likelihood-evaluator program, we hardcoded in the name of the dependent variable. As far as our likelihood-evaluator program is concerned, changing the dependent variable in our `ml model` statement did nothing.

When you specify the dependent variable in your `ml model` statement, `ml` stores the variable name in the global macro `$ML_y1`. Thus a better version of our `myprobit_gf0` program would be

```
----- begin myprobit_gf0_good.ado -----
program myprobit_gf0_good
  version 18
  args todo b lnfj
  tempvar xb
  quietly generate double `xb' = `b'[1,1]*weight + `b'[1,2]*price + ///
                                `b'[1,3]
  quietly replace `lnfj' = ln(normal(`xb')) if $ML_y1 == 1
  quietly replace `lnfj' = ln(normal(-1*`xb')) if $ML_y1 == 0
end
----- end myprobit_gf0_good.ado -----
```

With this change, we can specify dependent variables at will.

Adapting our program to accept an arbitrary dependent variable was straightforward. Unfortunately, making it accept an arbitrary set of independent variables is much more difficult. We wrote our likelihood evaluator assuming that the coefficient vector `'b'` had three elements, and we hardcoded the names of our independent variables

in the likelihood-evaluator program. If we were hell-bent on making our method-`gf0` evaluator work with an arbitrary number of independent variables, we could examine the column names of `'b'` and deduce the number of variables, their names, and even the number of equations. In the next chapter, we will learn a better way to approach problems using `ml` that affords us the ability to change regressors without having to modify our evaluator program in any way.