

lrtest — Likelihood-ratio test after estimation

Description	Quick start	Menu	Syntax
Options	Remarks and examples	Stored results	Methods and formulas
References	Also see		

Description

`lrtest` performs a likelihood-ratio test of the null hypothesis that the parameter vector of a statistical model satisfies some smooth constraint. To conduct the test, both the unrestricted and the restricted models must be fit using the maximum likelihood method (or some equivalent method), and the results of at least one must be stored using [estimates store](#).

`lrtest` also supports composite models. In a composite model, we assume that the log likelihood and dimension (number of free parameters) of the full model are obtained as the sum of the log-likelihood values and dimensions of the constituting models.

Quick start

Likelihood-ratio test that the coefficients for `x2` and `x3` are equal to 0

```
logit y x1 x2 x3
estimates store full
logit y x1 if e(sample)
estimates store restricted
lrtest full restricted
```

Display additional information, including AIC and BIC

```
lrtest full restricted, stats
```

Likelihood-ratio test that the coefficients for `x1` and `x3` are equal

```
constraint 1 x1=x3
logit y x1 x2 x3, constraints(1)
estimates store constrained
lrtest full constrained
```

Compare stored estimates `full` with the last model run

```
lrtest full
```

Menu

Statistics > Postestimation

Syntax

```
lrtest modelspec1 [modelspec2] [, options]
```

*modelspec*₁ and *modelspec*₂ specify the restricted and unrestricted model in any order. *modelspec*_# is *name* | . | (*namelist*)

name is the name under which estimation results were stored using `estimates store` (see [R] [estimates store](#)), and “.” refers to the last estimation results, whether or not these were already stored. If *modelspec*₂ is not specified, the last estimation result is used; this is equivalent to specifying *modelspec*₂ as “.”.

If *namelist* is specified for a composite model, *modelspec*₁ and *modelspec*₂ cannot have names in common; for example, `lrtest (A B C) (C D E)` is not allowed because both model specifications include C.

<i>options</i>	Description
<code>stats</code>	display statistical information about the two models
<code>dir</code>	display descriptive information about the two models
<code>df(#)</code>	override the automatic degrees-of-freedom calculation; seldom used
<code>force</code>	force testing even when apparently invalid

`collect` is allowed; see [U] [11.1.10 Prefix commands](#).

Options

`stats` displays statistical information about the unrestricted and restricted models, including the information indices of Akaike and Schwarz.

`dir` displays descriptive information about the unrestricted and restricted models; see `estimates dir` in [R] [estimates store](#).

`df(#)` is seldom specified; it overrides the automatic degrees-of-freedom calculation.

`force` forces the likelihood-ratio test calculations to take place in situations where `lrtest` would normally refuse to do so and issue an error. Such situations arise when one or more assumptions of the test are violated, for example, if the models were fit with `vce(robust)`, `vce(cluster clustvar)`, or `pweights`; when the dependent variables in the two models differ; when the null log likelihoods differ; when the samples differ; or when the estimation commands differ. If you use the `force` option, there is no guarantee as to the validity or interpretability of the resulting test.

Remarks and examples

[stata.com](http://www.stata.com)

The standard way to use `lrtest` is to do the following:

1. Fit either the restricted model or the unrestricted model by using one of Stata’s estimation commands and then store the results using `estimates store name`.
2. Fit the alternative model (the unrestricted or restricted model) and then type `‘lrtest name .’`. `lrtest` determines for itself which of the two models is the restricted model by comparing the degrees of freedom.

Often, you may want to store the alternative model with `estimates store name2`, for instance, if you plan additional tests against models yet to be fit. The likelihood-ratio test is then obtained as `lrtest name name2`.

Remarks are presented under the following headings:

Nested models
Composite models

Nested models

`lrtest` may be used with any estimation command that reports a log likelihood, including `heckman`, `logit`, `poisson`, `stcox`, and `streg`. You must check that one of the model specifications implies a statistical model that is *nested within* the model implied by the other specification. Usually, this means that both models are fit with the same estimation command (for example, both are fit by `logit`, with the same dependent variables) and that the set of covariates of one model is a subset of the covariates of the other model. Second, `lrtest` is valid only for models that are fit by maximum likelihood or by some equivalent method, so it does not apply to models that were fit with probability weights or clusters. Specifying the `vce(robust)` option similarly would indicate that you are worried about the valid specification of the model, so you would not use `lrtest`. Third, `lrtest` assumes that under the null hypothesis, the test statistic is (approximately) distributed as χ^2 . This assumption is not true for likelihood-ratio tests of “boundary conditions”, such as tests for the presence of overdispersion or random effects (Gutierrez, Carter, and Drukker 2001).

▷ Example 1

We have data on infants born with low birthweights along with the characteristics of the mother (Hosmer, Lemeshow, and Sturdivant 2013; see also [R] `logistic`). We fit the following model:

```
. use https://www.stata-press.com/data/r18/lbw
(Hosmer & Lemeshow data)
. logistic low age lwt i.race smoke ptl ht ui
Logistic regression                                Number of obs =   189
                                                    LR chi2(8)      =   33.22
                                                    Prob > chi2    = 0.0001
                                                    Pseudo R2     = 0.1416
Log likelihood = -100.724
```

low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
age	.9732636	.0354759	-0.74	0.457	.9061578	1.045339
lwt	.9849634	.0068217	-2.19	0.029	.9716834	.9984249
race						
Black	3.534767	1.860737	2.40	0.016	1.259736	9.918406
Other	2.368079	1.039949	1.96	0.050	1.001356	5.600207
smoke	2.517698	1.00916	2.30	0.021	1.147676	5.523162
ptl	1.719161	.5952579	1.56	0.118	.8721455	3.388787
ht	6.249602	4.322408	2.65	0.008	1.611152	24.24199
ui	2.1351	.9808153	1.65	0.099	.8677528	5.2534
_cons	1.586014	1.910496	0.38	0.702	.1496092	16.8134

Note: `_cons` estimates baseline odds.

We now wish to test the constraint that the coefficients on `age`, `lwt`, `ptl`, and `ht` are all zero or, equivalently here, that the odds ratios are all 1. One solution is to type

```
. test age lwt ptl ht
( 1) [low]age = 0
( 2) [low]lwt = 0
( 3) [low]ptl = 0
( 4) [low]ht = 0
      chi2( 4) = 12.38
      Prob > chi2 = 0.0147
```

This test is based on the inverse of the information matrix and is therefore based on a quadratic approximation to the likelihood function; see [R] [test](#). A more precise test would be to refit the model, applying the proposed constraints, and then calculate the likelihood-ratio test.

We first save the current model:

```
. estimates store full
```

We then fit the constrained model, which here is the model omitting `age`, `lwt`, `ptl`, and `ht`:

```
. logistic low i.race smoke ui
Logistic regression                               Number of obs = 189
                                                    LR chi2(4) = 18.80
                                                    Prob > chi2 = 0.0009
                                                    Pseudo R2 = 0.0801

Log likelihood = -107.93404
```

low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
race						
Black	3.052746	1.498087	2.27	0.023	1.166747	7.987382
Other	2.922593	1.189229	2.64	0.008	1.316457	6.488285
smoke						
ui	2.419131	1.047359	2.04	0.041	1.035459	5.651788
_cons	.1402209	.0512295	-5.38	0.000	.0685216	.2869447

Note: `_cons` estimates baseline odds.

That done, `lrtest` compares this model with the model we previously stored:

```
. lrtest full .
Likelihood-ratio test
Assumption: . nested within full
      LR chi2(4) = 14.42
      Prob > chi2 = 0.0061
```

Let’s compare results. `test` reported that `age`, `lwt`, `ptl`, and `ht` were jointly significant at the 1.5% level; `lrtest` reports that they are significant at the 0.6% level. Given the quadratic approximation made by `test`, we could argue that `lrtest`’s results are more accurate.

`lrtest` explicates the assumption that, from a comparison of the degrees of freedom, it has assessed that the last fit model (`.`) is nested within the model stored as `full`. In other words, `full` is the unconstrained model and `.` is the constrained model.

The names in “(Assumption: `.` nested in `full`)” are actually links. Click on a name, and the results for that model are replayed.

◀

Aside: The `nestreg` command provides a simple syntax for performing likelihood-ratio tests for nested model specifications; see [R] [nestreg](#). In the [previous example](#), we fit a full logistic model, used `estimates store` to store the full model, fit a constrained logistic model, and used `lrtest` to report a likelihood-ratio test between two models. To do this with one call to `nestreg`, use the `lrtable` option.

□ Technical note

`lrttest` determines the degrees of freedom of a model as the rank of the (co)variance matrix $e(V)$. There are two issues here. First, the *numerical* determination of the rank of a matrix is a subtle problem that can, for instance, be affected by the scaling of the variables in the model. The rank of a matrix depends on the number of (independent) linear combinations of coefficients that sum exactly to zero. In the world of numerical mathematics, it is hard to tell whether a very small number is really nonzero or is a real zero that happens to be slightly off because of roundoff error from the finite precision with which computers make floating-point calculations. Whether a small number is being classified as one or the other, typically on the basis of a threshold, affects the determined degrees of freedom. Although Stata generally makes sensible choices, it is bound to make mistakes occasionally. The moral of this story is to make sure that the calculated degrees of freedom is as you expect before interpreting the results.

□

□ Technical note

A second issue involves `regress` and related commands such as `anova`. Mainly for historical reasons, `regress` does not treat the residual variance, σ^2 , the same way that it treats the regression coefficients. Type `estat vce` after `regress`, and you will see the regression coefficients, not $\hat{\sigma}^2$. Most estimation commands for models with ancillary parameters (for example, `streg` and `heckman`) treat all parameters as equals. There is nothing technically wrong with `regress` here; we are usually focused on the regression coefficients, and their estimators are uncorrelated with $\hat{\sigma}^2$. But, formally, σ^2 adds a degree of freedom to the model, which does not matter if you are comparing two regression models by a likelihood-ratio test. This test depends on the difference in the degrees of freedom, and hence being “off by 1” in each does not matter. But, if you are comparing a regression model with a larger model—for example, a heteroskedastic regression model fit by `arch`—the automatic determination of the degrees of freedom is incorrect, and you must specify the `df(#)` option.

□

▷ Example 2

Returning to the low-birthweight data in [example 1](#), we now wish to test that the coefficient on `2.race` (black) is equal to that on `3.race` (other). The base model is still stored under the name `full`, so we need only fit the constrained model and perform the test. With z as the index of the logit model, the base model is

$$z = \beta_0 + \beta_1 \text{age} + \beta_2 \text{lwt} + \beta_3 \text{2.race} + \beta_4 \text{3.race} + \dots$$

If $\beta_3 = \beta_4$, this can be written as

$$z = \beta_0 + \beta_1 \text{age} + \beta_2 \text{lwt} + \beta_3 (\text{2.race} + \text{3.race}) + \dots$$

We can fit the constrained model as follows:

```
. constraint 1 2.race = 3.race
. logistic low age lwt i.race smoke ptl ht ui, constraints(1)
Logistic regression                                Number of obs =   189
                                                    Wald chi2(7)  =  25.17
Log likelihood = -100.9997                          Prob > chi2   =  0.0007
( 1)  [low]2.race - [low]3.race = 0
```

	low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]
age		.9716799	.0352638	-0.79	0.429	.9049649 1.043313
lwt		.9864971	.0064627	-2.08	0.038	.9739114 .9992453
race						
Black		2.728186	1.080207	2.53	0.011	1.255586 5.927907
Other		2.728186	1.080207	2.53	0.011	1.255586 5.927907
smoke		2.664498	1.052379	2.48	0.013	1.228633 5.778414
ptl		1.709129	.5924776	1.55	0.122	.8663666 3.371691
ht		6.116391	4.215585	2.63	0.009	1.58425 23.61385
ui		2.09936	.9699702	1.61	0.108	.8487997 5.192407
_cons		1.309371	1.527398	0.23	0.817	.1330839 12.8825

Note: **_cons** estimates baseline odds.

Comparing this model with our original model, we obtain

```
. lrtest full .
Likelihood-ratio test
Assumption: . nested within full
LR chi2(1) = 0.55
Prob > chi2 = 0.4577
```

By comparison, typing `test 2.race=3.race` after fitting our base model results in a significance level of 0.4572. Alternatively, we can first store the restricted model, here using the name `equal`. Next, `lrtest` is invoked specifying the names of the restricted and unrestricted models (we do not care about the order). This time, we also add the option `stats` requesting a table of model statistics, including the model selection indices AIC and BIC.

```
. estimates store equal
. lrtest equal full, stats
Likelihood-ratio test
Assumption: equal nested within full
LR chi2(1) = 0.55
Prob > chi2 = 0.4577
Akaike's information criterion and Bayesian information criterion
```

Model	N	ll(null)	ll(model)	df	AIC	BIC
equal	189	.	-100.9997	8	217.9994	243.9334
full	189	-117.336	-100.724	9	219.448	248.6237

Note: BIC uses N = number of observations. See **[R] IC note**.

Composite models

lrtest supports composite models; that is, models that can be fit by fitting a series of simpler models or by fitting models on subsets of the data. Theoretically, a composite model is one in which the likelihood function, $L(\theta)$, of the parameter vector, θ , can be written as the product

$$L(\theta) = L_1(\theta_1) \times L_2(\theta_2) \times \cdots \times L_k(\theta_k)$$

of likelihood terms with $\theta = (\theta_1, \dots, \theta_k)$ a partitioning of the full parameter vector. In such a case, the full-model likelihood $L(\theta)$ is maximized by maximizing the likelihood terms $L_j(\theta_j)$ in turn. Obviously, $\log L(\hat{\theta}) = \sum_{j=1}^k \log L_j(\hat{\theta}_j)$. The degrees of freedom for the composite model is obtained as the sum of the degrees of freedom of the constituting models.

► Example 3

As an example of the application of composite models, we consider a test of the hypothesis that the coefficients of a statistical model do not differ between different portions (“regimes”) of the covariate space. Economists call a test for such a hypothesis a *Chow test*.

We continue the analysis of the data on children of low birthweight by using logistic regression modeling and study whether the regression coefficients are the same among the three races: white, black, and other. A likelihood-ratio Chow test can be obtained by fitting the logistic regression model for each of the races and then comparing the combined results with those of the model previously stored as full. Because the full model included dummies for the three races, this version of the Chow test allows the intercept of the logistic regression model to vary between the regimes (races).

```
. logistic low age lwt smoke ptl ht ui if 1.race, nolog
Logistic regression                               Number of obs =    96
                                                    LR chi2(6)      =   13.86
                                                    Prob > chi2    =  0.0312
Log likelihood = -45.927061                       Pseudo R2      =  0.1311
```

	low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]
age		.9869674	.0527757	-0.25	0.806	.8887649 1.096021
lwt		.9900874	.0106101	-0.93	0.353	.9695089 1.011103
smoke		4.208697	2.680133	2.26	0.024	1.20808 14.66222
ptl		1.592145	.7474264	0.99	0.322	.6344379 3.995544
ht		2.900166	3.193537	0.97	0.334	.3350554 25.1032
ui		1.229523	.9474768	0.27	0.789	.2715165 5.567715
_cons		.4891008	.993785	-0.35	0.725	.0091175 26.23746

Note: **_cons** estimates baseline odds.

```
. estimates store white
```

8 lrtest — Likelihood-ratio test after estimation

```
. logistic low age lwt smoke ptl ht ui if 2.race, nolog
Logistic regression
Log likelihood = -12.654157
Number of obs = 26
LR chi2(6) = 10.12
Prob > chi2 = 0.1198
Pseudo R2 = 0.2856
```

low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
age	.8735313	.1377846	-0.86	0.391	.6412332	1.189983
lwt	.9747736	.016689	-1.49	0.136	.9426065	1.008038
smoke	16.50373	24.37044	1.90	0.058	.9133647	298.2083
ptl	4.866916	9.33151	0.83	0.409	.1135573	208.5895
ht	85.05605	214.6382	1.76	0.078	.6049308	11959.27
ui	67.61338	133.3313	2.14	0.033	1.417399	3225.322
_cons	48.7249	169.9216	1.11	0.265	.0523961	45310.94

Note: **_cons** estimates baseline odds.

```
. estimates store black
. logistic low age lwt smoke ptl ht ui if 3.race, nolog
Logistic regression
Log likelihood = -37.228444
Number of obs = 67
LR chi2(6) = 14.06
Prob > chi2 = 0.0289
Pseudo R2 = 0.1589
```

low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]	
age	.9263905	.0665386	-1.06	0.287	.8047407	1.06643
lwt	.9724499	.015762	-1.72	0.085	.9420424	1.003839
smoke	.7979034	.6340585	-0.28	0.776	.1680885	3.787586
ptl	2.845675	1.777944	1.67	0.094	.8363053	9.682908
ht	7.767503	10.00537	1.59	0.112	.6220764	96.98826
ui	2.925006	2.046473	1.53	0.125	.7423107	11.52571
_cons	49.09444	113.9165	1.68	0.093	.5199275	4635.769

Note: **_cons** estimates baseline odds.

. estimates store other

We are now ready to perform the likelihood-ratio Chow test:

```
. lrtest (full) (white black other), stats
Likelihood-ratio test
Assumption: full nested within (white, black, other)
LR chi2(12) = 9.83
Prob > chi2 = 0.6310
Akaike's information criterion and Bayesian information criterion
```

Model	N	ll(null)	ll(model)	df	AIC	BIC
full	189	-117.336	-100.724	9	219.448	248.6237
white	96	-52.85752	-45.92706	7	105.8541	123.8046
black	26	-17.71291	-12.65416	7	39.30831	48.11499
other	67	-44.26039	-37.22844	7	88.45689	103.8897

Note: BIC uses N = number of observations. See **[R] IC note**.

We cannot reject the hypothesis that the logistic regression model applies to each of the races at any reasonable significance level. By specifying the **stats** option, we can verify the degrees of freedom of the test: $12 = 7 + 7 + 7 - 9$. We can obtain the same test by fitting an expanded model with interactions between all covariates and race.


```
. logistic low race##c.(age lwt smoke ptl ht ui)
```

```
Logistic regression
```

```
Number of obs = 189
LR chi2(20) = 43.05
Prob > chi2 = 0.0020
Pseudo R2 = 0.1835
```

```
Log likelihood = -95.809661
```

	low	Odds ratio	Std. err.	z	P> z	[95% conf. interval]
race						
Black		99.62137	402.0829	1.14	0.254	.0365434 271578.9
Other		100.3769	309.586	1.49	0.135	.2378638 42358.38
age						
		.9869674	.0527757	-0.25	0.806	.8887649 1.096021
lwt						
		.9900874	.0106101	-0.93	0.353	.9695089 1.011103
smoke						
		4.208697	2.680133	2.26	0.024	1.20808 14.66222
ptl						
		1.592145	.7474264	0.99	0.322	.6344379 3.995544
ht						
		2.900166	3.193537	0.97	0.334	.3350554 25.1032
ui						
		1.229523	.9474768	0.27	0.789	.2715165 5.567715
race#c.age						
Black		.885066	.1474079	-0.73	0.464	.638569 1.226714
Other		.9386232	.0840486	-0.71	0.479	.7875366 1.118695
race#c.lwt						
Black		.9845329	.0198857	-0.77	0.440	.9463191 1.02429
Other		.9821859	.0190847	-0.93	0.355	.9454839 1.020313
race#c.smoke						
Black		3.921338	6.305992	0.85	0.395	.167725 91.67917
Other		.1895844	.1930601	-1.63	0.102	.025763 1.395113
race#c.ptl						
Black		3.05683	6.034089	0.57	0.571	.0638301 146.3918
Other		1.787322	1.396789	0.74	0.457	.3863582 8.268285
race#c.ht						
Black		29.328	80.7482	1.23	0.220	.1329492 6469.623
Other		2.678295	4.538712	0.58	0.561	.0966916 74.18702
race#c.ui						
Black		54.99155	116.4274	1.89	0.058	.8672471 3486.977
Other		2.378976	2.476124	0.83	0.405	.309335 18.29579
_cons						
		.4891008	.993785	-0.35	0.725	.0091175 26.23746

Note: **_cons** estimates baseline odds.

```
. lrtest full .
```

```
Likelihood-ratio test
```

```
Assumption: full nested within .
```

```
LR chi2(12) = 9.83
```

```
Prob > chi2 = 0.6310
```

Applying `lrtest` for the full model against the model with all interactions yields the same test statistic and p -value as for the full model against the composite model for the three regimes. Here the specification of the model with interactions was convenient, and `logistic` had no problem computing the estimates for the expanded model. In models with more complicated likelihoods, such as Heckman's selection model (see [R] [heckman](#)) or complicated survival-time models (see [ST] [streg](#)), fitting the models with all interactions may be numerically demanding and may be much more time consuming than fitting a series of models separately for each regime.

Given the model with all interactions, we could also test the hypothesis of no differences among the regions (races) by a Wald version of the Chow test by using the `testparm` command; see [R] [test](#).

```
. testparm race#c.(age lwt smoke ptl ht ui)
( 1) [low]2.race#c.age = 0
( 2) [low]3.race#c.age = 0
( 3) [low]2.race#c.lwt = 0
( 4) [low]3.race#c.lwt = 0
( 5) [low]2.race#c.smoke = 0
( 6) [low]3.race#c.smoke = 0
( 7) [low]2.race#c.ptl = 0
( 8) [low]3.race#c.ptl = 0
( 9) [low]2.race#c.ht = 0
(10) [low]3.race#c.ht = 0
(11) [low]2.race#c.ui = 0
(12) [low]3.race#c.ui = 0

      chi2( 12) =      8.24
      Prob > chi2 =    0.7663
```

We conclude that, here, the Wald version of the Chow test is similar to the likelihood-ratio version of the Chow test.

◀

Stored results

`lrtest` stores the following in `r()`:

Scalars

<code>r(p)</code>	<i>p</i> -value for likelihood-ratio test
<code>r(df)</code>	degrees of freedom
<code>r(chi2)</code>	LR test statistic

Programmers wishing their estimation commands to be compatible with `lrtest` should note that `lrtest` requires that the following results be returned:

<code>e(cmd)</code>	name of estimation command
<code>e(ll)</code>	log likelihood
<code>e(V)</code>	variance-covariance matrix of the estimators
<code>e(N)</code>	number of observations

`lrtest` also verifies that `e(N)`, `e(ll_0)`, and `e(depvar)` are consistent between two noncomposite models.

Methods and formulas

Let L_0 and L_1 be the log-likelihood values associated with the full and constrained models, respectively. The test statistic of the likelihood-ratio test is $LR = -2(L_1 - L_0)$. If the constrained model is true, LR is approximately χ^2 distributed with $d_0 - d_1$ degrees of freedom, where d_0 and d_1 are the model degrees of freedom associated with the full and constrained models, respectively (Greene 2018, 554–555).

`lrtest` determines the degrees of freedom of a model as the rank of `e(V)`, computed as the number of nonzero diagonal elements of `invsym(e(V))`.

References

- Greene, W. H. 2018. *Econometric Analysis*. 8th ed. New York: Pearson.
- Gutierrez, R. G., S. L. Carter, and D. M. Drukker. 2001. [sg160: On boundary-value likelihood-ratio tests](#). *Stata Technical Bulletin* 60: 15–18. Reprinted in *Stata Technical Bulletin Reprints*, vol. 10, pp. 269–273. College Station, TX: Stata Press.
- Hosmer, D. W., Jr., S. A. Lemeshow, and R. X. Sturdivant. 2013. *Applied Logistic Regression*. 3rd ed. Hoboken, NJ: Wiley.
- Raciborski, R. 2015. Spotlight on irt. *The Stata Blog: Not Elsewhere Classified*. <http://blog.stata.com/2015/07/31/spotlight-on-irt/>.
- Tauchmann, H. 2023. [lgrgtest: Lagrange multiplier test after constrained maximum-likelihood estimation](#). *Stata Journal* 23: 386–401.

Also see

- [R] [test](#) — Test linear hypotheses after estimation
- [R] [testnl](#) — Test nonlinear hypotheses after estimation
- [R] [nestreg](#) — Nested model statistics

Stata, Stata Press, and Mata are registered trademarks of StataCorp LLC. Stata and Stata Press are registered trademarks with the World Intellectual Property Organization of the United Nations. StataNow and NetCourseNow are trademarks of StataCorp LLC. Other brand and product names are registered trademarks or trademarks of their respective companies. Copyright © 1985–2023 StataCorp LLC, College Station, TX, USA. All rights reserved.



For suggested citations, see the FAQ on [citing Stata documentation](#).