

Title

vce_option — Variance estimators

<p>Syntax Methods and formulas</p>	<p>Description Also see</p>	<p>Options</p>	<p>Remarks and examples</p>
--	---------------------------------	----------------	-----------------------------

Syntax

estimation_cmd ... [, *vce*(*vcetype*) ...]

<i>vcetype</i>	Description
<hr/>	
Likelihood based	
<code>oim</code>	observed information matrix (OIM)
<code>opg</code>	outer product of the gradient (OPG) vectors
Sandwich estimators	
<code>robust</code>	Huber/White/sandwich estimator
<code>cluster</code> <i>clustvar</i>	clustered sandwich estimator
Replication based	
<code>bootstrap</code> [, <i>bootstrap_options</i>]	bootstrap estimation
<code>jackknife</code> [, <i>jackknife_options</i>]	jackknife estimation

Description

This entry describes the `vce()` option, which is common to most estimation commands. `vce()` specifies how to estimate the variance–covariance matrix (VCE) corresponding to the parameter estimates. The standard errors reported in the table of parameter estimates are the square root of the variances (diagonal elements) of the VCE.

Options

SE/Robust

`vce(oim)` is usually the default for models fit using maximum likelihood. `vce(oim)` uses the observed information matrix (OIM); see [R] [ml](#).

`vce(opg)` uses the sum of the outer product of the gradient (OPG) vectors; see [R] [ml](#). This is the default VCE when the `technique(bhhh)` option is specified; see [R] [maximize](#).

`vce(robust)` uses the robust or sandwich estimator of variance. This estimator is robust to some types of misspecification so long as the observations are independent; see [U] [20.21 Obtaining robust variance estimates](#).

If the command allows `pweights` and you specify them, `vce(robust)` is implied; see [U] 20.23.3 Sampling weights.

`vce(cluster clustvar)` specifies that the standard errors allow for intragroup correlation, relaxing the usual requirement that the observations be independent. That is, the observations are independent across groups (clusters) but not necessarily within groups. *clustvar* specifies to which group each observation belongs, for example, `vce(cluster personid)` in data with repeated observations on individuals. `vce(cluster clustvar)` affects the standard errors and variance–covariance matrix of the estimators but not the estimated coefficients; see [U] 20.21 Obtaining robust variance estimates.

`vce(bootstrap [, bootstrap_options])` uses a bootstrap; see [R] bootstrap. After estimation with `vce(bootstrap)`, see [R] bootstrap postestimation to obtain percentile-based or bias-corrected confidence intervals.

`vce(jackknife [, jackknife_options])` uses the delete-one jackknife; see [R] jackknife.

Remarks and examples

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

Remarks are presented under the following headings:

Prefix commands

Passing options in vce()

Prefix commands

Specifying `vce(bootstrap)` or `vce(jackknife)` is often equivalent to using the corresponding prefix command. Here is an example using `jackknife` with `regress`.

```
. use http://www.stata-press.com/data/r13/auto
(1978 Automobile Data)
. regress mpg turn trunk, vce(jackknife)
(running regress on estimation sample)
Jackknife replications (74)
-----|----- 1 -----|----- 2 -----|----- 3 -----|----- 4 -----|----- 5
.....|..... 50
.....
Linear regression                               Number of obs   =          74
                                                Replications    =          74
                                                F( 2,          73) =         66.26
                                                Prob > F        =         0.0000
                                                R-squared       =         0.5521
                                                Adj R-squared   =         0.5395
                                                Root MSE      =         3.9260
```

mpg	Jackknife		t	P> t	[95% Conf. Interval]	
	Coef.	Std. Err.				
turn	-.7610113	.150726	-5.05	0.000	-1.061408	-.4606147
trunk	-.3161825	.1282326	-2.47	0.016	-.5717498	-.0606152
_cons	55.82001	5.031107	11.09	0.000	45.79303	65.84699

```
. jackknife: regress mpg turn trunk
(running regress on estimation sample)
```

```
Jackknife replications (74)
```

```
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
      1      2      3      4      5
.....
.....
```

```
Linear regression
```

Number of obs	=	74
Replications	=	74
F(2, 73)	=	66.26
Prob > F	=	0.0000
R-squared	=	0.5521
Adj R-squared	=	0.5395
Root MSE	=	3.9260

mpg	Coef.	Jackknife Std. Err.	t	P> t	[95% Conf. Interval]	
turn	-.7610113	.150726	-5.05	0.000	-1.061408	-.4606147
trunk	-.3161825	.1282326	-2.47	0.016	-.5717498	-.0606152
_cons	55.82001	5.031107	11.09	0.000	45.79303	65.84699

Here it does not matter whether we specify the `vce(jackknife)` option or instead use the `jackknife` prefix.

However, `vce(jackknife)` should be used in place of the `jackknife` prefix whenever available because they are not always equivalent. For example, to use the `jackknife` prefix with `clogit` properly, you must tell `jackknife` to omit whole groups rather than individual observations. Specifying `vce(jackknife)` does this automatically.

```
. use http://www.stata-press.com/data/r13/clogitid
. jackknife, cluster(id): clogit y x1 x2, group(id)
(output omitted)
```

This extra information is automatically communicated to `jackknife` by `clogit` when the `vce()` option is specified.

```
. clogit y x1 x2, group(id) vce(jackknife)
(running clogit on estimation sample)
Jackknife replications (66)
-----|-----|-----|-----|-----|-----|-----|-----|-----|-----|
      1      2      3      4      5
.....
.....
```

```
Conditional (fixed-effects) logistic regression
```

Number of obs	=	369
Replications	=	66
F(2, 65)	=	4.58
Prob > F	=	0.0137
Pseudo R2	=	0.0355

```
Log likelihood = -123.41386
```

```
(Replications based on 66 clusters in id)
```

y	Coef.	Jackknife Std. Err.	t	P> t	[95% Conf. Interval]	
x1	.653363	.3010608	2.17	0.034	.052103	1.254623
x2	.0659169	.0487858	1.35	0.181	-.0315151	.1633489

Passing options in `vce()`

If you wish to specify more options to the bootstrap or jackknife estimation, you can include them within the `vce()` option. Below we request 300 bootstrap replications and save the replications in `bsreg.dta`:

```
. use http://www.stata-press.com/data/r13/auto
(1978 Automobile Data)
. regress mpg turn trunk, vce(bootstrap, nodots seed(123) rep(300) saving(bsreg))
```

```
Linear regression                Number of obs   =       74
                                Replications   =       300
                                Wald chi2(2)     =      127.28
                                Prob > chi2      =       0.0000
                                R-squared        =       0.5521
                                Adj R-squared    =       0.5395
                                Root MSE     =       3.9260
```

mpg	Observed Coef.	Bootstrap Std. Err.	z	P> z	Normal-based [95% Conf. Interval]	
turn	-.7610113	.1361786	-5.59	0.000	-1.027916	-.4941062
trunk	-.3161825	.1145728	-2.76	0.006	-.540741	-.0916239
_cons	55.82001	4.69971	11.88	0.000	46.60875	65.03127

```
. bstat using bsreg
```

```
Bootstrap results                Number of obs   =       74
                                Replications   =       300
```

```
command: regress mpg turn trunk
```

	Observed Coef.	Bootstrap Std. Err.	z	P> z	Normal-based [95% Conf. Interval]	
turn	-.7610113	.1361786	-5.59	0.000	-1.027916	-.4941062
trunk	-.3161825	.1145728	-2.76	0.006	-.540741	-.0916239
_cons	55.82001	4.69971	11.88	0.000	46.60875	65.03127

Methods and formulas

By default, Stata's maximum likelihood estimators display standard errors based on variance estimates given by the inverse of the negative Hessian (second derivative) matrix. If `vce(robust)`, `vce(cluster clustvar)`, or `pweights` is specified, standard errors are based on the robust variance estimator (see [U] [20.21 Obtaining robust variance estimates](#)); likelihood-ratio tests are not appropriate here (see [SVY] [survey](#)), and the model χ^2 is from a Wald test. If `vce(opg)` is specified, the standard errors are based on the outer product of the gradients; this option has no effect on likelihood-ratio tests, though it does affect Wald tests.

If `vce(bootstrap)` or `vce(jackknife)` is specified, the standard errors are based on the chosen replication method; here the model χ^2 or F statistic is from a Wald test using the respective replication-based covariance matrix. The t distribution is used in the coefficient table when the `vce(jackknife)` option is specified. `vce(bootstrap)` and `vce(jackknife)` are also available with some commands that are not maximum likelihood estimators.

Also see

[R] **bootstrap** — Bootstrap sampling and estimation

[R] **jackknife** — Jackknife estimation

[XT] *vce_options* — Variance estimators

[U] **20 Estimation and postestimation commands**