

vce_options — Variance estimators

Description Syntax Options Remarks and examples
 Methods and formulas Reference Also see

Description

This entry describes the *vce_options*, which are common to most xt estimation commands. Not all the options documented below work with all xt estimation commands; see the documentation for the particular estimation command. If an option is listed there, it is applicable.

The `vce()` option 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.

Syntax

estimation_cmd ... [, *vce_options* ...]

<i>vce_options</i>	Description
<code>vce(oim)</code>	observed information matrix (OIM)
<code>vce(opg)</code>	outer product of the gradient (OPG) vectors
<code>vce(robust)</code>	Huber/White/sandwich estimator
<code>vce(cluster <i>clustvar</i>)</code>	clustered sandwich estimator
<code>vce(bootstrap [, <i>bootstrap_options</i>])</code>	bootstrap estimation
<code>vce(jackknife [, <i>jackknife_options</i>])</code>	jackknife estimation
<code>nmp</code>	use divisor $N - P$ instead of the default N
<code>scale(x2 dev phi #)</code>	override the default scale parameter; available only with population-averaged models

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] [20.22 Obtaining robust variance estimates](#).

If the command allows `pweights` and you specify them, `vce(robust)` is implied; see [U] [20.24.3 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 to say, 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] 20.22 **Obtaining robust variance estimates**.

`vce(bootstrap [, bootstrap_options])` uses a nonparametric 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**.

`nmp` specifies that the divisor $N - P$ be used instead of the default N , where N is the total number of observations and P is the number of coefficients estimated.

`scale(x2|dev|phi|#)` overrides the default scale parameter. By default, `scale(1)` is assumed for the discrete distributions (binomial, negative binomial, and Poisson), and `scale(x2)` is assumed for the continuous distributions (gamma, Gaussian, and inverse Gaussian).

`scale(x2)` specifies that the scale parameter be set to the Pearson χ^2 (or generalized χ^2) statistic divided by the residual degrees of freedom, which is recommended by [McCullagh and Nelder \(1989\)](#) as a good general choice for continuous distributions.

`scale(dev)` sets the scale parameter to the deviance divided by the residual degrees of freedom. This option provides an alternative to `scale(x2)` for continuous distributions and for over- or underdispersed discrete distributions.

`scale(phi)` specifies that the scale parameter be estimated from the data. `xtgee`'s default scaling makes results agree with other estimators and has been recommended by [McCullagh and Nelder \(1989\)](#) in the context of GLM. When comparing results with calculations made by other software, you may find that the other packages do not offer this feature. In such cases, specifying `scale(phi)` should match their results.

`scale(#)` sets the scale parameter to `#`. For example, using `scale(1)` in `family(gamma)` models results in exponential-errors regression (if you assume independent correlation structure).

Remarks and examples

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

When you are working with panel-data models, we strongly encourage you to use the `vce(bootstrap)` or `vce(jackknife)` option instead of the corresponding prefix command. For example, to obtain jackknife standard errors with `xtlogit`, type

```

. use https://www.stata-press.com/data/r17/clogitid
. xtlogit y x1 x2, fe vce(jackknife)
(running xtlogit on estimation sample)

Jackknife replications (66)
-----|-----|-----|-----|-----|
..... 50
.....

Conditional fixed-effects logistic regression      Number of obs   =   369
                                                    Replications    =    66
Group variable: id                               Number of groups =    66

                                                    Obs per group:
                                                    min =         2
                                                    avg =         5.6
                                                    max =         10

                                                    F(2, 65)       =   4.58
                                                    Prob > F       =   0.0137

Log likelihood = -123.41386
                                                    (Replications based on 66 clusters in id)

```

y	Coefficient	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

If you wish to specify more options to the bootstrap or jackknife estimation, you can include them within the `vce()` option. Below we refit our model requesting bootstrap standard errors based on 300 replications, we set the random-number seed so that our results can be reproduced, and we suppress the display of the replication dots.

```

. xtlogit y x1 x2, fe vce(bootstrap, reps(300) seed(123) nodots)
Conditional fixed-effects logistic regression      Number of obs   =   369
                                                    Replications    =   300
Group variable: id                               Number of groups =    66

                                                    Obs per group:
                                                    min =         2
                                                    avg =         5.6
                                                    max =         10

                                                    Wald chi2(2)   =    9.26
                                                    Prob > chi2    =   0.0097

Log likelihood = -123.41386
                                                    (Replications based on 66 clusters in id)

```

y	Observed coefficient	Bootstrap std. err.	z	P> z	Normal-based [95% conf. interval]	
x1	.653363	.307093	2.13	0.033	.0514717	1.255254
x2	.0659169	.0477384	1.38	0.167	-.0276486	.1594824

□ Technical note

To perform jackknife estimation on panel data, you must omit entire panels rather than individual observations. To replicate the output above using the `jackknife` prefix command, you would have to type

```

. jackknife, cluster(id): xtlogit y x1 x2, fe
(output omitted)

```

Similarly, bootstrap estimation on panel data requires you to resample entire panels rather than individual observations. The `vce(bootstrap)` and `vce(jackknife)` options handle this for you automatically. □

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` are specified, standard errors are based on the robust variance estimator (see [U] 20.22 **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.

Reference

McCullagh, P., and J. A. Nelder. 1989. *Generalized Linear Models*. 2nd ed. London: Chapman & Hall/CRC.

Also see

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

[R] **jackknife** — Jackknife estimation

[R] **ml** — Maximum likelihood estimation

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