

ivregress postestimation — Postestimation tools for ivregress

| | | | |
|----------------------------------|--------------------------------------|-----------------------------------|--------------------------------------|
| Description | Syntax for predict | Menu for predict | Options for predict |
| Syntax for estat | Menu for estat | Options for estat | Remarks and examples |
| Stored results | Methods and formulas | References | Also see |

Description

The following postestimation commands are of special interest after `ivregress`:

| Command | Description |
|-------------------------------|---|
| <code>estat endogenous</code> | perform tests of endogeneity |
| <code>estat firststage</code> | report “first-stage” regression statistics |
| <code>estat overid</code> | perform tests of overidentifying restrictions |

These commands are not appropriate after the `svy` prefix.

The following postestimation commands are also available:

| Command | Description |
|------------------------------------|---|
| <code>contrast</code> | contrasts and ANOVA-style joint tests of estimates |
| <code>estat summarize</code> | summary statistics for the estimation sample |
| <code>estat vce</code> | variance–covariance matrix of the estimators (VCE) |
| <code>estat (svy)</code> | postestimation statistics for survey data |
| <code>estimates</code> | cataloging estimation results |
| <code>forecast</code> ¹ | dynamic forecasts and simulations |
| <code>hausman</code> | Hausman’s specification test |
| <code>lincom</code> | point estimates, standard errors, testing, and inference for linear combinations of coefficients |
| <code>margins</code> | marginal means, predictive margins, marginal effects, and average marginal effects |
| <code>marginsplot</code> | graph the results from margins (profile plots, interaction plots, etc.) |
| <code>nlcom</code> | point estimates, standard errors, testing, and inference for nonlinear combinations of coefficients |
| <code>predict</code> | predictions, residuals, influence statistics, and other diagnostic measures |
| <code>predictnl</code> | point estimates, standard errors, testing, and inference for generalized predictions |
| <code>pwcompare</code> | pairwise comparisons of estimates |
| <code>test</code> | Wald tests of simple and composite linear hypotheses |
| <code>testnl</code> | Wald tests of nonlinear hypotheses |

¹ `forecast` is not appropriate with `svy` estimation results.

Special-interest postestimation commands

`estat endogenous` performs tests to determine whether endogenous regressors in the model are in fact exogenous. After GMM estimation, the C (difference-in-Sargan) statistic is reported. After 2SLS estimation with an unadjusted VCE, the Durbin (1954) and Wu–Hausman (Wu 1974; Hausman 1978) statistics are reported. After 2SLS estimation with a robust VCE, Wooldridge’s (1995) robust score test and a robust regression-based test are reported. In all cases, if the test statistic is significant, then the variables being tested must be treated as endogenous. `estat endogenous` is not available after LIML estimation.

`estat firststage` reports various statistics that measure the relevance of the excluded exogenous variables. By default, whether the equation has one or more than one endogenous regressor determines what statistics are reported.

`estat overid` performs tests of overidentifying restrictions. If the 2SLS estimator was used, Sargan’s (1958) and Basman’s (1960) χ^2 tests are reported, as is Wooldridge’s (1995) robust score test; if the LIML estimator was used, Anderson and Rubin’s (1950) χ^2 test and Basman’s F test are reported; and if the GMM estimator was used, Hansen’s (1982) J statistic χ^2 test is reported. A statistically significant test statistic always indicates that the instruments may not be valid.

Syntax for predict

```
predict [type] newvar [if] [in] [, statistic]
```

```
predict [type] { stub* | newvarlist } [if] [in], scores
```

| <i>statistic</i> | Description |
|------------------|-------------|
|------------------|-------------|

Main

| | |
|---------------------------------------|---|
| <code>xb</code> | linear prediction; the default |
| <code>residuals</code> | residuals |
| <code>stdp</code> | standard error of the prediction |
| <code>stdf</code> | standard error of the forecast |
| <code>pr(<i>a</i>,<i>b</i>)</code> | $\Pr(a < y_j < b)$ |
| <code>e(<i>a</i>,<i>b</i>)</code> | $E(y_j a < y_j < b)$ |
| <code>ystar(<i>a</i>,<i>b</i>)</code> | $E(y_j^*), y_j^* = \max\{a, \min(y_j, b)\}$ |

These statistics are available both in and out of sample; type `predict ... if e(sample) ...` if wanted only for the estimation sample.

`stdf` is not allowed with `svy` estimation results.

where a and b may be numbers or variables; a missing ($a \geq .$) means $-\infty$, and b missing ($b \geq .$) means $+\infty$; see [U] 12.2.1 Missing values.

Menu for predict

Statistics > Postestimation > Predictions, residuals, etc.

Options for predict

Main

xb, the default, calculates the linear prediction.

residuals calculates the residuals, that is, $y_j - \mathbf{x}_j \mathbf{b}$. These are based on the estimated equation when the observed values of the endogenous variables are used—not the projections of the instruments onto the endogenous variables.

stdp calculates the standard error of the prediction, which can be thought of as the standard error of the predicted expected value or mean for the observation’s covariate pattern. This is also referred to as the standard error of the fitted value.

stdf calculates the standard error of the forecast, which is the standard error of the point prediction for 1 observation. It is commonly referred to as the standard error of the future or forecast value. By construction, the standard errors produced by **stdf** are always larger than those produced by **stdp**; see *Methods and formulas* in [R] **regress postestimation**.

pr(a,b) calculates $\Pr(a < \mathbf{x}_j \mathbf{b} + u_j < b)$, the probability that $y_j | \mathbf{x}_j$ would be observed in the interval (a, b) .

a and b may be specified as numbers or variable names; lb and ub are variable names;

pr(20,30) calculates $\Pr(20 < \mathbf{x}_j \mathbf{b} + u_j < 30)$;

pr(lb,ub) calculates $\Pr(lb < \mathbf{x}_j \mathbf{b} + u_j < ub)$; and

pr(20,ub) calculates $\Pr(20 < \mathbf{x}_j \mathbf{b} + u_j < ub)$.

a missing ($a \geq .$) means $-\infty$; **pr(.,30)** calculates $\Pr(-\infty < \mathbf{x}_j \mathbf{b} + u_j < 30)$;

pr(lb,30) calculates $\Pr(-\infty < \mathbf{x}_j \mathbf{b} + u_j < 30)$ in observations for which $lb \geq .$

and calculates $\Pr(lb < \mathbf{x}_j \mathbf{b} + u_j < 30)$ elsewhere.

b missing ($b \geq .$) means $+\infty$; **pr(20,.)** calculates $\Pr(+\infty > \mathbf{x}_j \mathbf{b} + u_j > 20)$;

pr(20,ub) calculates $\Pr(+\infty > \mathbf{x}_j \mathbf{b} + u_j > 20)$ in observations for which $ub \geq .$

and calculates $\Pr(20 < \mathbf{x}_j \mathbf{b} + u_j < ub)$ elsewhere.

e(a,b) calculates $E(\mathbf{x}_j \mathbf{b} + u_j | a < \mathbf{x}_j \mathbf{b} + u_j < b)$, the expected value of $y_j | \mathbf{x}_j$ conditional on $y_j | \mathbf{x}_j$ being in the interval (a, b) , meaning that $y_j | \mathbf{x}_j$ is truncated.

a and b are specified as they are for **pr()**.

ystar(a,b) calculates $E(y_j^*)$, where $y_j^* = a$ if $\mathbf{x}_j \mathbf{b} + u_j \leq a$, $y_j^* = b$ if $\mathbf{x}_j \mathbf{b} + u_j \geq b$, and $y_j^* = \mathbf{x}_j \mathbf{b} + u_j$ otherwise, meaning that y_j^* is censored. a and b are specified as they are for **pr()**.

scores calculates the scores for the model. A new score variable is created for each endogenous regressor, as well as an equation-level score that applies to all exogenous variables and constant term (if present).

Syntax for estat

Perform tests of endogeneity

```
estat endogenous [varlist] [, _lags(#) forceweights forcenonrobust]
```

Report “first-stage” regression statistics

```
estat firststage [, all forcenonrobust]
```

Perform tests of overidentifying restrictions

```
estat overid [ , lags(#) forceweights forcenonrobust ]
```

Menu for estat

Statistics > Postestimation > Reports and statistics

Options for estat

Options for `estat` are presented under the following headings:

Options for estat endogenous

Options for estat firststage

Options for estat overid

Options for estat endogenous

`lags`(#) specifies the number of lags to use for prewhitening when computing the heteroskedasticity- and autocorrelation-consistent (HAC) version of the score test of endogeneity. Specifying `lags`(0) requests no prewhitening. This option is valid only when the model was fit via 2SLS and an HAC covariance matrix was requested when the model was fit. The default is `lags`(1).

`forceweights` requests that the tests of endogeneity be computed even though `aweight`s, `pweight`s, or `iweight`s were used in the previous estimation. By default, these tests are conducted only after unweighted or frequency-weighted estimation. The reported critical values may be inappropriate for weighted data, so the user must determine whether the critical values are appropriate for a given application.

`forcenonrobust` requests that the Durbin and Wu–Hausman tests be performed after 2SLS estimation even though a robust VCE was used at estimation time. This option is available only if the model was fit by 2SLS.

Options for estat firststage

`all` requests that all first-stage goodness-of-fit statistics be reported regardless of whether the model contains one or more endogenous regressors. By default, if the model contains one endogenous regressor, then the first-stage R^2 , adjusted R^2 , partial R^2 , and F statistics are reported, whereas if the model contains multiple endogenous regressors, then Shea's partial R^2 and adjusted partial R^2 are reported instead.

`forcenonrobust` requests that the minimum eigenvalue statistic and its critical values be reported even though a robust VCE was used at estimation time. The reported critical values assume that the errors are independent and identically distributed (i.i.d.) normal, so the user must determine whether the critical values are appropriate for a given application.

Options for estat overid

`lags(#)` specifies the number of lags to use for prewhitening when computing the heteroskedasticity- and autocorrelation-consistent (HAC) version of the score test of overidentifying restrictions. Specifying `lags(0)` requests no prewhitening. This option is valid only when the model was fit via 2SLS and an HAC covariance matrix was requested when the model was fit. The default is `lags(1)`.

`forceweights` requests that the tests of overidentifying restrictions be computed even though `aweight`s, `pweight`s, or `iweight`s were used in the previous estimation. By default, these tests are conducted only after unweighted or frequency-weighted estimation. The reported critical values may be inappropriate for weighted data, so the user must determine whether the critical values are appropriate for a given application.

`forcenonrobust` requests that the Sargan and Basman tests of overidentifying restrictions be performed after 2SLS or LIML estimation even though a robust VCE was used at estimation time. These tests assume that the errors are i.i.d. normal, so the user must determine whether the critical values are appropriate for a given application.

Remarks and examples

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

Remarks are presented under the following headings:

estat endogenous
estat firststage
estat overid

estat endogenous

A natural question to ask is whether a variable presumed to be endogenous in the previously fit model could instead be treated as exogenous. If the endogenous regressors are in fact exogenous, then the OLS estimator is more efficient; and depending on the strength of the instruments and other factors, the sacrifice in efficiency by using an instrumental-variables estimator can be significant. Thus, unless an instrumental-variables estimator is really needed, OLS should be used instead. `estat endogenous` provides several tests of endogeneity after 2SLS and GMM estimation.

► Example 1

In [example 1](#) of [\[R\] ivregress](#), we fit a model of the average rental rate for housing in a state as a function of the percentage of the population living in urban areas and the average value of houses. We treated `hsngval` as endogenous because unanticipated shocks that affect rental rates probably affect house prices as well. We used family income and region dummies as additional instruments for `hsngval`. Here we test whether we could treat `hsngval` as exogenous.

```
. use http://www.stata-press.com/data/r13/hsng
(1980 Census housing data)
. ivregress 2sls rent pcturban (hsngval = faminc i.region)
(output omitted)
. estat endogenous

Tests of endogeneity
Ho: variables are exogenous
Durbin (score) chi2(1)          = 12.8473   (p = 0.0003)
Wu-Hausman F(1,46)             = 15.9067   (p = 0.0002)
```

Because we did not specify any variable names after the `estat endogenous` command, Stata by default tested all the endogenous regressors (namely, `hsngval`) in our model. The null hypothesis of the Durbin and Wu–Hausman tests is that the variable under consideration can be treated as exogenous. Here both test statistics are highly significant, so we reject the null of exogeneity; we must continue to treat `hsngval` as endogenous.

◀

The difference between the Durbin and Wu–Hausman tests of endogeneity is that the former uses an estimate of the error term’s variance based on the model assuming the variables being tested are exogenous, while the latter uses an estimate of the error variance based on the model assuming the variables being tested are endogenous. Under the null hypothesis that the variables being tested are exogenous, both estimates of the error variance are consistent. What we label the Wu–Hausman statistic is Wu’s (1974) “ T_2 ” statistic, which Hausman (1978) showed can be calculated very easily via linear regression. Baum, Schaffer, and Stillman (2003, 2007) provide a lucid discussion of these tests.

When you fit a model with multiple endogenous regressors, you can test the exogeneity of a subset of the regressors while continuing to treat the others as endogenous. For example, say you have three endogenous regressors, `y1`, `y2`, and `y3`, and you fit your model by typing

```
. ivregress depvar ... (y1 y2 y3 = ...)
```

Suppose you are confident that `y1` must be treated as endogenous, but you are undecided about `y2` and `y3`. To test whether `y2` and `y3` can be treated as exogenous, you would type

```
. estat endogenous y2 y3
```

The Durbin and Wu–Hausman tests assume that the error term is i.i.d. Therefore, if you requested a robust VCE at estimation time, `estat endogenous` will instead report Wooldridge’s (1995) score test and a regression-based test of exogeneity. Both these tests can tolerate heteroskedastic and autocorrelated errors, while only the regression-based test is amenable to clustering.

▷ Example 2

We refit our housing model, requesting robust standard errors, and then test the exogeneity of `hsngval`:

```
. use http://www.stata-press.com/data/r13/hsng
(1980 Census housing data)
. ivregress 2sls rent pcturban (hsngval = faminc i.region), vce(robust)
(output omitted)
. estat endogenous
Tests of endogeneity
Ho: variables are exogenous
Robust score chi2(1)          = 2.10428 (p = 0.1469)
Robust regression F(1,46)    = 4.31101 (p = 0.0435)
```

Wooldridge’s score test does not reject the null hypothesis that `hsngval` is exogenous at conventional significance levels ($p = 0.1469$). However, the regression-based test does reject the null hypothesis at the 5% significance level ($p = 0.0435$). Typically, these two tests yield the same conclusion; the fact that our dataset has only 50 observations could be contributing to the discrepancy. Here we would be inclined to continue to treat `hsngval` as endogenous. Even if `hsngval` is exogenous, the 2SLS estimates are still consistent. On the other hand, if `hsngval` is in fact endogenous, the OLS estimates would not be consistent. Moreover, as we will see in our discussion of the `estat overid` command, our additional instruments may be invalid. To test whether an endogenous variable can be treated as exogenous, we must have a valid set of instruments to use to fit the model in the first place!

◀

Unlike the Durbin and Wu–Hausman tests, Wooldridge’s score and the regression-based tests do not allow you to test a subset of the endogenous regressors in the model; you can test only whether all the endogenous regressors are in fact exogenous.

After GMM estimation, `estat endogenous` calculates what Hayashi (2000, 220) calls the C statistic, also known as the difference-in-Sargan statistic. The C statistic can be made robust to heteroskedasticity, autocorrelation, and clustering; and the version reported by `estat endogenous` is determined by the weight matrix requested via the `wmatrix()` option used when fitting the model with `ivregress`. Additionally, the test can be used to determine the exogeneity of a subset of the endogenous regressors, regardless of the type of weight matrix used.

If you fit your model using the LIML estimator, you can use the `hausman` command to carry out a traditional Hausman (1978) test between the OLS and LIML estimates.

estat firststage

For an excluded exogenous variable to be a valid instrument, it must be sufficiently correlated with the included endogenous regressors but uncorrelated with the error term. In recent decades, researchers have paid considerable attention to the issue of instruments that are only weakly correlated with the endogenous regressors. In such cases, the usual 2SLS, GMM, and LIML estimators are biased toward the OLS estimator, and inference based on the standard errors reported by, for example, `ivregress` can be severely misleading. For more information on the theory behind instrumental-variables estimation with weak instruments, see Nelson and Startz (1990); Staiger and Stock (1997); Hahn and Hausman (2003); the survey article by Stock, Wright, and Yogo (2002); and Angrist and Pischke (2009, chap. 4).

When the instruments are only weakly correlated with the endogenous regressors, some Monte Carlo evidence suggests that the LIML estimator performs better than the 2SLS and GMM estimators; see, for example, Poi (2006) and Stock, Wright, and Yogo (2002) (and the papers cited therein). On the other hand, the LIML estimator often results in confidence intervals that are somewhat larger than those from the 2SLS estimator.

Moreover, using more instruments is not a solution, because the biases of instrumental-variables estimators increase with the number of instruments. See Hahn and Hausman (2003).

`estat firststage` produces several statistics for judging the explanatory power of the instruments and is most easily explained with examples.

► Example 3

Again building on the model fit in [example 1](#) of [\[R\] ivregress](#), we now explore the degree of correlation between the additional instruments `faminc`, `2.region`, `3.region`, and `4.region` and the endogenous regressor `hshgval`:

```
. use http://www.stata-press.com/data/r13/hshg
(1980 Census housing data)
. ivregress 2sls rent pcturban (hshgval = faminc i.region)
(output omitted)
. estat firststage
```

First-stage regression summary statistics

| Variable | R-sq. | Adjusted R-sq. | Partial R-sq. | F(4,44) | Prob > F |
|----------|--------|----------------|---------------|---------|----------|
| hshgval | 0.6908 | 0.6557 | 0.5473 | 13.2978 | 0.0000 |

Minimum eigenvalue statistic = 13.2978

Critical Values # of endogenous regressors: 1
 Ho: Instruments are weak # of excluded instruments: 4

| | 5% | 10% | 20% | 30% |
|-----------------------------------|-------|-------|-------|------|
| 2SLS relative bias | 16.85 | 10.27 | 6.71 | 5.34 |
| 2SLS Size of nominal 5% Wald test | 10% | 15% | 20% | 25% |
| LIML Size of nominal 5% Wald test | 24.58 | 13.96 | 10.26 | 8.31 |
| | 5.44 | 3.87 | 3.30 | 2.98 |

To understand these results, recall that the first-stage regression is

$$\text{hsngval}_i = \pi_0 + \pi_1 \text{pcturban}_i + \pi_2 \text{faminc} + \pi_3 \text{2.region} + \pi_4 \text{3.region} + \pi_5 \text{4.region} + v_i$$

where v_i is an error term. The column marked “R-sq.” is the simple R^2 from fitting the first-stage regression by OLS, and the column marked “Adjusted R-sq.” is the adjusted R^2 from that regression. Higher values purportedly indicate stronger instruments, and instrumental-variables estimators exhibit less bias when the instruments are strongly correlated with the endogenous variable.

Looking at just the R^2 and adjusted R^2 can be misleading, however. If `hsngval` were strongly correlated with the included exogenous variable `pcturban` but only weakly correlated with the additional instruments, then these statistics could be large even though a weak-instrument problem is present.

The partial R^2 statistic measures the correlation between `hsngval` and the additional instruments after *partialling out* the effect of `pcturban`. Unlike the R^2 and adjusted R^2 statistics, the partial R^2 statistic will not be inflated because of strong correlation between `hsngval` and `pcturban`. [Bound, Jaeger, and Baker \(1995\)](#) and others have promoted using this statistic.

The column marked “F(4, 44)” is an F statistic for the joint significance of π_2 , π_3 , π_4 , and π_5 , the coefficients on the additional instruments. Its p -value is listed in the column marked “Prob > F”. If the F statistic is not significant, then the additional instruments have no significant explanatory power for `hsngval` after controlling for the effect of `pcturban`. However, [Hall, Rudebusch, and Wilcox \(1996\)](#) used Monte Carlo simulation to show that simply having an F statistic that is significant at the typical 5% or 10% level is not sufficient. [Stock, Wright, and Yogo \(2002\)](#) suggest that the F statistic should exceed 10 for inference based on the 2SLS estimator to be reliable when there is one endogenous regressor.

`estat firststage` also presents the [Cragg and Donald \(1993\)](#) minimum eigenvalue statistic as a further test of weak instruments. [Stock and Yogo \(2005\)](#) discuss two characterizations of weak instruments: first, weak instruments cause instrumental-variables estimators to be biased; second, hypothesis tests of parameters estimated by instrumental-variables estimators may suffer from severe size distortions. The test statistic in our example is 13.30, which is identical to the F statistic just discussed because our model contains one endogenous regressor.

The null hypothesis of each of Stock and Yogo’s tests is that the set of instruments is weak. To perform these tests, we must first choose either the largest relative bias of the 2SLS estimator we are willing to tolerate or the largest rejection rate of a nominal 5% Wald test we are willing to tolerate. If the test statistic exceeds the critical value, we can conclude that our instruments are not weak.

The row marked “2SLS relative bias” contains critical values for the test that the instruments are weak based on the bias of the 2SLS estimator *relative to* the bias of the OLS estimator. For example, from past experience we might know that the OLS estimate of a parameter β may be 50% too high. Saying that we are willing to tolerate a 10% relative bias means that we are willing to tolerate a bias of the 2SLS estimator no greater than 5% (that is, 10% of 50%). In our rental rate model, if we

are willing to tolerate a 10% relative bias, then we can conclude that our instruments are not weak because the test statistic of 13.30 exceeds the critical value of 10.22. However, if we were willing to tolerate only a relative bias of 5%, we would conclude that our instruments are weak because $13.30 < 16.85$.

The rows marked “2SLS Size of nominal 5% Wald test” and “LIML Size of nominal 5% Wald test” contain critical values pertaining to Stock and Yogo’s (2005) second characterization of weak instruments. This characterization defines a set of instruments to be weak if a Wald test at the 5% level can have an actual rejection rate of no more than 10%, 15%, 20%, or 25%. Using the current example, suppose that we are willing to accept a rejection rate of at most 10%. Because $13.30 < 24.58$, we cannot reject the null hypothesis of weak instruments. On the other hand, if we use the LIML estimator instead, then we can reject the null hypothesis because $13.30 > 5.44$.



□ Technical note

Stock and Yogo (2005) tabulated critical values for 2SLS relative biases of 5%, 10%, 20%, and 30% for models with 1, 2, or 3 endogenous regressors and between 3 and 30 excluded exogenous variables (instruments). They also provide critical values for worst-case rejection rates of 5%, 10%, 20%, and 25% for nominal 5% Wald tests of the endogenous regressors with 1 or 2 endogenous regressors and between 1 and 30 instruments. If the model previously fit by `ivregress` has more instruments or endogenous regressors than these limits, the critical values are not shown. Stock and Yogo did not consider GMM estimators.



When the model being fit contains more than one endogenous regressor, the R^2 and F statistics described above can overstate the relevance of the excluded instruments. Suppose that there are two endogenous regressors, Y_1 and Y_2 , and that there are two additional instruments, z_1 and z_2 . Say that z_1 is highly correlated with both Y_1 and Y_2 but z_2 is not correlated with either Y_1 or Y_2 . Then the first-stage regression of Y_1 on z_1 and z_2 (along with the included exogenous variables) will produce large R^2 and F statistics, as will the regression of Y_2 on z_1 , z_2 , and the included exogenous variables. Nevertheless, the lack of correlation between z_2 and Y_1 and Y_2 is problematic. Here, although the order condition indicates that the model is just identified (the number of excluded instruments equals the number of endogenous regressors), the irrelevance of z_2 implies that the model is in fact not identified. Even if the model is overidentified, including irrelevant instruments can adversely affect the properties of instrumental-variables estimators, because their biases increase as the number of instruments increases.

▷ Example 4

`estat firststage` presents different statistics when the model contains multiple endogenous regressors. For illustration, we refit our model of rental rates, assuming that both `hshgval` and `faminc` are endogenously determined. We use `i.region` along with `popden`, a measure of population density, as additional instruments.

```
. ivregress 2sls rent pcturban (hsngval faminc = i.region popden)
(output omitted)
. estat firststage
Shea's partial R-squared
```

| Variable | Shea's Partial R-sq. | Shea's Adj. Partial R-sq. |
|----------|-------------------------|------------------------------|
| hsngval | 0.3477 | 0.2735 |
| faminc | 0.1893 | 0.0972 |

```
Minimum eigenvalue statistic = 2.51666
```

```
Critical Values # of endogenous regressors: 2
Ho: Instruments are weak # of excluded instruments: 4
```

| | | | | |
|-----------------------------------|-------|------|------|------|
| 2SLS relative bias | 5% | 10% | 20% | 30% |
| | 11.04 | 7.56 | 5.57 | 4.73 |
| 2SLS Size of nominal 5% Wald test | 10% | 15% | 20% | 25% |
| | 16.87 | 9.93 | 7.54 | 6.28 |
| LIML Size of nominal 5% Wald test | 4.72 | 3.39 | 2.99 | 2.79 |

Consider the endogenous regressor `hsngval`. Part of its variation is attributable to its correlation with the other regressors `pcturban` and `faminc`. The other component of `hsngval`'s variation is peculiar to it and orthogonal to the variation in the other regressors. Similarly, we can think of the instruments as predicting the variation in `hsngval` in two ways, one stemming from the fact that the predicted values of `hsngval` are correlated with the predicted values of the other regressors and one from the variation in the predicted values of `hsngval` that is orthogonal to the variation in the predicted values of the other regressors.

What really matters for instrumental-variables estimation is whether the component of `hsngval` that is orthogonal to the other regressors can be explained by the component of the predicted value of `hsngval` that is orthogonal to the predicted values of the other regressors in the model. Shea's (1997) partial R^2 statistic measures this correlation. Because the bias of instrumental-variables estimators increases as more instruments are used, Shea's adjusted partial R^2 statistic is often used instead, as it makes a degrees-of-freedom adjustment for the number of instruments, analogous to the adjusted R^2 measure used in OLS regression. Although what constitutes a "low" value for Shea's partial R^2 depends on the specifics of the model being fit and the data used, these results, taken in isolation, do not strike us as being a particular cause for concern.

However, with this specification the minimum eigenvalue statistic is low. We cannot reject the null hypothesis of weak instruments for either of the characterizations we have discussed.

◀

By default, `estat firststage` determines which statistics to present based on the number of endogenous regressors in the model previously fit. However, you can specify the `all` option to obtain all the statistics.

□ Technical note

If the previous estimation was conducted using `aweight`s, `pweight`s, or `iweight`s, then the first-stage regression summary statistics are computed using those weights. However, in these cases the minimum eigenvalue statistic and its critical values are not available.

If the previous estimation included a robust VCE, then the first-stage F statistic is based on a robust VCE as well; for example, if you fit your model with an HAC VCE using the Bartlett kernel and four lags, then the F statistic reported is based on regression results using an HAC VCE using the Bartlett kernel and four lags. By default, the minimum eigenvalue statistic and its critical values are not displayed. You can use the `forcenonrobust` option to obtain them in these cases; the minimum eigenvalue statistic is computed using the weights, though the critical values reported may not be appropriate. □

estat overid

In addition to the requirement that instrumental variables be correlated with the endogenous regressors, the instruments must also be uncorrelated with the structural error term. If the model is overidentified, meaning that the number of additional instruments exceeds the number of endogenous regressors, then we can test whether the instruments are uncorrelated with the error term. If the model is just identified, then we cannot perform a test of overidentifying restrictions.

The estimator you used to fit the model determines which tests of overidentifying restrictions `estat overid` reports. If you used the 2SLS estimator without a robust VCE, `estat overid` reports Sargan's (1958) and Basman's (1960) χ^2 tests. If you used the 2SLS estimator and requested a robust VCE, Wooldridge's robust score test of overidentifying restrictions is performed instead; without a robust VCE, Wooldridge's test statistic is identical to Sargan's test statistic. If you used the LIML estimator, `estat overid` reports the Anderson–Rubin (1950) likelihood-ratio test and Basman's (1960) F test. `estat overid` reports Hansen's (1982) J statistic if you used the GMM estimator. Davidson and MacKinnon (1993, 235–236) give a particularly clear explanation of the intuition behind tests of overidentifying restrictions. Also see Judge et al. (1985, 614–616) for a summary of tests of overidentifying restrictions for the 2SLS and LIML estimators.

Tests of overidentifying restrictions actually test two different things simultaneously. One, as we have discussed, is whether the instruments are uncorrelated with the error term. The other is that the equation is misspecified and that one or more of the excluded exogenous variables should in fact be included in the structural equation. Thus a significant test statistic could represent either an invalid instrument or an incorrectly specified structural equation.

► Example 5

Here we refit the model that treated just `hsngval` as endogenous using 2SLS, and then we perform tests of overidentifying restrictions:

```
. ivregress 2sls rent pcturban (hsngval = faminc i.region)
(output omitted)
. estat overid

Tests of overidentifying restrictions:
Sargan (score) chi2(3) = 11.2877 (p = 0.0103)
Basman chi2(3)      = 12.8294 (p = 0.0050)
```

Both test statistics are significant at the 5% test level, which means that either one or more of our instruments are invalid or that our structural model is specified incorrectly.

One possibility is that the error term in our structural model is heteroskedastic. Both Sargan's and Basman's tests assume that the errors are i.i.d.; if the errors are not i.i.d., then these tests are not valid. Here we refit the model by requesting heteroskedasticity-robust standard errors, and then we use `estat overid` to obtain Wooldridge's score test of overidentifying restrictions, which is robust to heteroskedasticity.

```
. ivregress 2sls rent pcturban (hsngval = faminc i.region), vce(robust)
(output omitted)
. estat overid
Test of overidentifying restrictions:
Score chi2(3)          =   6.8364   (p = 0.0773)
```

Here we no longer reject the null hypothesis that our instruments are valid at the 5% significance level, though we do reject the null at the 10% level. You can verify that the robust standard error on the coefficient for `hsngval` is more than twice as large as its nonrobust counterpart and that the robust standard error for `pcturban` is nearly 50% larger. ◀

□ Technical note

The test statistic for the test of overidentifying restrictions performed after GMM estimation is simply the sample size times the value of the objective function $Q(\beta_1, \beta_2)$ defined in (5) of [R] `ivregress`, evaluated at the GMM parameter estimates. If the weighting matrix \mathbf{W} is optimal, meaning that $\mathbf{W} = \text{Var}(z_i u_i)$, then $Q(\beta_1, \beta_2) \overset{A}{\sim} \chi^2(q)$, where q is the number of overidentifying restrictions. However, if the estimated \mathbf{W} is not optimal, then the test statistic will not have an asymptotic χ^2 distribution.

Like the Sargan and Basman tests of overidentifying restrictions for the 2SLS estimator, the Anderson–Rubin and Basman tests after LIML estimation are predicated on the errors’ being i.i.d. If the previous LIML results were reported with robust standard errors, then `estat overid` by default issues an error message and refuses to report the Anderson–Rubin and Basman test statistics. You can use the `forcenonrobust` option to override this behavior. You can also use `forcenonrobust` to obtain the Sargan and Basman test statistics after 2SLS estimation with robust standard errors. □

By default, `estat overid` issues an error message if the previous estimation was conducted using `aweight`s, `pweight`s, or `iweight`s. You can use the `forceweights` option to override this behavior, though the test statistics may no longer have the expected χ^2 distributions.

Stored results

After 2SLS estimation, `estat endogenous` stores the following in `r()`:

Scalars

| | |
|-----------------------------|---|
| <code>r(durbin)</code> | Durbin χ^2 statistic |
| <code>r(p_durbin)</code> | p -value for Durbin χ^2 statistic |
| <code>r(wu)</code> | Wu–Hausman F statistic |
| <code>r(p_wu)</code> | p -value for Wu–Hausman F statistic |
| <code>r(df)</code> | degrees of freedom |
| <code>r(wudf_r)</code> | denominator degrees of freedom for Wu–Hausman F |
| <code>r(r_score)</code> | robust score statistic |
| <code>r(p_r_score)</code> | p -value for robust score statistic |
| <code>r(hac_score)</code> | HAC score statistic |
| <code>r(p_hac_score)</code> | p -value for HAC score statistic |
| <code>r(lags)</code> | lags used in prewhitening |
| <code>r(regF)</code> | regression-based F statistic |
| <code>r(p_regF)</code> | p -value for regression-based F statistic |
| <code>r(regFdf_n)</code> | regression-based F numerator degrees of freedom |
| <code>r(regFdf_r)</code> | regression-based F denominator degrees of freedom |

After GMM estimation, `estat endogenous` stores the following in `r()`:

| | |
|---------------------|---------------------------------------|
| Scalars | |
| <code>r(C)</code> | C χ^2 statistic |
| <code>r(p_C)</code> | p -value for C χ^2 statistic |
| <code>r(df)</code> | degrees of freedom |

`estat firststage` stores the following in `r()`:

| | |
|-------------------------------|--|
| Scalars | |
| <code>r(mineig)</code> | minimum eigenvalue statistic |
| Matrices | |
| <code>r(mineigcv)</code> | critical values for minimum eigenvalue statistic |
| <code>r(multiresults)</code> | Shea's partial R^2 statistics |
| <code>r(singleresults)</code> | first-stage R^2 and F statistics |

After 2SLS estimation, `estat overid` stores the following in `r()`:

| | |
|---------------------------|---|
| Scalars | |
| <code>r(lags)</code> | lags used in prewhitening |
| <code>r(df)</code> | χ^2 degrees of freedom |
| <code>r(score)</code> | score χ^2 statistic |
| <code>r(p_score)</code> | p -value for score χ^2 statistic |
| <code>r(basmann)</code> | Basmann χ^2 statistic |
| <code>r(p_basmann)</code> | p -value for Basmann χ^2 statistic |
| <code>r(sargan)</code> | Sargan χ^2 statistic |
| <code>r(p_sargan)</code> | p -value for Sargan χ^2 statistic |

After LIML estimation, `estat overid` stores the following in `r()`:

| | |
|------------------------------|--|
| Scalars | |
| <code>r(ar)</code> | Anderson–Rubin χ^2 statistic |
| <code>r(p_ar)</code> | p -value for Anderson–Rubin χ^2 statistic |
| <code>r(ar_df)</code> | χ^2 degrees of freedom |
| <code>r(basmann)</code> | Basmann F statistic |
| <code>r(p_basmann)</code> | p -value for Basmann F statistic |
| <code>r(basmann_df_n)</code> | F numerator degrees of freedom |
| <code>r(basmann_df_d)</code> | F denominator degrees of freedom |

After GMM estimation, `estat overid` stores the following in `r()`:

| | |
|---------------------------|--|
| Scalars | |
| <code>r(HansenJ)</code> | Hansen's J χ^2 statistic |
| <code>r(p_HansenJ)</code> | p -value for Hansen's J χ^2 statistic |
| <code>r(J_df)</code> | χ^2 degrees of freedom |

Methods and formulas

Methods and formulas are presented under the following headings:

Notation
estat endogenous
estat firststage
estat overid

Notation

Recall from [R] `ivregress` that the model is

$$\mathbf{y} = \mathbf{Y}\boldsymbol{\beta}_1 + \mathbf{X}_1\boldsymbol{\beta}_2 + \mathbf{u} = \mathbf{X}\boldsymbol{\beta} + \mathbf{u}$$

$$\mathbf{Y} = \mathbf{X}_1\boldsymbol{\Pi}_1 + \mathbf{X}_2\boldsymbol{\Pi}_2 + \mathbf{V} = \mathbf{Z}\boldsymbol{\Pi} + \mathbf{V}$$

where \mathbf{y} is an $N \times 1$ vector of the left-hand-side variable, N is the sample size, \mathbf{Y} is an $N \times p$ matrix of p endogenous regressors, \mathbf{X}_1 is an $N \times k_1$ matrix of k_1 included exogenous regressors, \mathbf{X}_2 is an $N \times k_2$ matrix of k_2 excluded exogenous variables, $\mathbf{X} = [\mathbf{Y} \ \mathbf{X}_1]$, $\mathbf{Z} = [\mathbf{X}_1 \ \mathbf{X}_2]$, \mathbf{u} is an $N \times 1$ vector of errors, \mathbf{V} is an $N \times p$ matrix of errors, $\boldsymbol{\beta} = [\boldsymbol{\beta}_1 \ \boldsymbol{\beta}_2]$ is a $k = (p + k_1) \times 1$ vector of parameters, and $\boldsymbol{\Pi}$ is a $(k_1 + k_2) \times p$ vector of parameters. If a constant term is included in the model, then one column of \mathbf{X}_1 contains all ones.

estat endogenous

Partition \mathbf{Y} as $\mathbf{Y} = [\mathbf{Y}_1 \ \mathbf{Y}_2]$, where \mathbf{Y}_1 represents the p_1 endogenous regressors whose endogeneity is being tested and \mathbf{Y}_2 represents the p_2 endogenous regressors whose endogeneity is not being tested. If the endogeneity of all endogenous regressors is being tested, $\mathbf{Y} = \mathbf{Y}_1$ and $p_2 = 0$. After GMM estimation, `estat endogenous` refits the model treating \mathbf{Y}_1 as exogenous using the same type of weight matrix as requested at estimation time with the `wmatrix()` option; denote the Sargan statistic from this model by J_e and the estimated weight matrix by \mathbf{W}_e . Let $\mathbf{S}_e = \mathbf{W}_e^{-1}$. `estat endogenous` removes from \mathbf{S}_e the rows and columns corresponding to the variables represented by \mathbf{Y}_1 ; denote the inverse of the resulting matrix by \mathbf{W}'_e . Next `estat endogenous` fits the model treating both \mathbf{Y}_1 and \mathbf{Y}_2 as endogenous, using the weight matrix \mathbf{W}'_e ; denote the Sargan statistic from this model by J_c . Then $C = (J_e - J_c) \sim \chi^2(p_1)$. If one simply used the J statistic from the original model fit by `ivregress` in place of J_c , then in finite samples $J_e - J$ might be negative. The procedure used by `estat endogenous` is guaranteed to yield $C \geq 0$; see Hayashi (2000, 220).

Let $\hat{\mathbf{u}}_c$ denote the residuals from the model treating both \mathbf{Y}_1 and \mathbf{Y}_2 as endogenous, and let $\hat{\mathbf{u}}_e$ denote the residuals from the model treating only \mathbf{Y}_2 as endogenous. Then Durbin's (1954) statistic is

$$D = \frac{\hat{\mathbf{u}}'_e \mathbf{P}_{Z\mathbf{Y}_1} \hat{\mathbf{u}}_e - \hat{\mathbf{u}}'_c \mathbf{P}_Z \hat{\mathbf{u}}_c}{\hat{\mathbf{u}}'_e \hat{\mathbf{u}}_e / N}$$

where $\mathbf{P}_Z = \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'$ and $\mathbf{P}_{Z\mathbf{Y}_1} = \mathbf{Z} \ \mathbf{Y}_1'(\mathbf{Z} \ \mathbf{Y}_1)^{-1}[\mathbf{Z} \ \mathbf{Y}_1]'$. $D \sim \chi^2(p_1)$. The Wu–Hausman (Wu 1974; Hausman 1978) statistic is

$$WH = \frac{(\hat{\mathbf{u}}'_e \mathbf{P}_{Z\mathbf{Y}_1} \hat{\mathbf{u}}_e - \hat{\mathbf{u}}'_c \mathbf{P}_Z \hat{\mathbf{u}}_c) / p_1}{\{\hat{\mathbf{u}}'_e \hat{\mathbf{u}}_e - (\hat{\mathbf{u}}'_e \mathbf{P}_{Z\mathbf{Y}_1} \hat{\mathbf{u}}_e - \hat{\mathbf{u}}'_c \mathbf{P}_Z \hat{\mathbf{u}}_c)\} / (N - k_1 - p - p_1)}$$

$WH \sim F(p_1, N - k_1 - p - p_1)$. Baum, Schaffer, and Stillman (2003, 2007) discuss these tests in more detail.

Next we describe Wooldridge's (1995) score test. The nonrobust version of Wooldridge's test is identical to Durbin's test. Suppose a robust covariance matrix was used at estimation time. Let $\hat{\mathbf{e}}$ denote the sample residuals obtained by fitting the model via OLS, treating \mathbf{Y} as exogenous. We then regress each variable represented in \mathbf{Y} on \mathbf{Z} ; call the residuals for the j th regression \hat{r}_j , $j = 1, \dots, p$. Define $\hat{k}_{ij} = \hat{e}_i \hat{r}_{ij}$, $i = 1, \dots, N$. We then run the regression

$$\mathbf{1} = \theta_1 \hat{\mathbf{k}}_1 + \dots + \theta_p \hat{\mathbf{k}}_p + \epsilon$$

where $\mathbf{1}$ is an $N \times 1$ vector of ones and ϵ is a regression error term. $N - \text{RSS} \sim \chi^2(p)$, where RSS is the residual sum of squares from the regression just described. If instead an HAC VCE was used at estimation time, then before running the final regression we prewhiten the $\widehat{\mathbf{k}}_j$ series by using a VAR(q) model, where q is the number of lags specified with the `lags()` option.

The regression-based test proceeds as follows. Following Hausman (1978, 1259), we regress \mathbf{Y} on \mathbf{Z} and obtain the residuals $\widehat{\mathbf{V}}$. Next we fit the augmented regression

$$\mathbf{y} = \mathbf{Y}\beta_1 + \mathbf{X}_1\beta_2 + \widehat{\mathbf{V}}\gamma + \epsilon$$

by OLS regression, where ϵ is a regression error term. A test of the exogeneity of \mathbf{Y} is equivalent to a test of $\gamma = \mathbf{0}$. As Cameron and Trivedi (2005, 276) suggest, this test can be made robust to heteroskedasticity, autocorrelation, or clustering by using the appropriate robust VCE when testing $\gamma = \mathbf{0}$. When a nonrobust VCE is used, this test is equivalent to the Wu–Hausman test described earlier. One cannot simply fit this augmented regression via 2SLS to test the endogeneity of a subset of the endogenous regressors; Davidson and MacKinnon (1993, 229–231) discuss a test of $\gamma = \mathbf{0}$ for the homoskedastic version of the augmented regression fit by 2SLS, but an appropriate robust test is not apparent.

estat firststage

When the structural equation includes one endogenous regressor, `estat firststage` fits the regression

$$\mathbf{Y} = \mathbf{X}_1\pi_1 + \mathbf{X}_2\pi_2 + \mathbf{v}$$

via OLS. The R^2 and adjusted R^2 from that regression are reported in the output, as well as the F statistic from the Wald test of $H_0: \pi_2 = \mathbf{0}$. To obtain the partial R^2 statistic, `estat firststage` fits the regression

$$\mathbf{M}_{\mathbf{X}_1}\mathbf{y} = \mathbf{M}_{\mathbf{X}_1}\mathbf{X}_2\xi + \epsilon$$

by OLS, where ϵ is a regression error term, ξ is a $k_2 \times 1$ parameter vector, and $\mathbf{M}_{\mathbf{X}_1} = \mathbf{I} - \mathbf{X}_1(\mathbf{X}_1'\mathbf{X}_1)^{-1}\mathbf{X}_1'$; that is, the partial R^2 is the R^2 between \mathbf{y} and \mathbf{X}_2 after eliminating the effects of \mathbf{X}_1 . If the model contains multiple endogenous regressors and the `all` option is specified, these statistics are calculated for each endogenous regressor in turn.

To calculate Shea's partial R^2 , let \mathbf{y}_1 denote the endogenous regressor whose statistic is being calculated and \mathbf{Y}_0 denote the other endogenous regressors. Define $\widetilde{\mathbf{y}}_1$ as the residuals obtained from regressing \mathbf{y}_1 on \mathbf{Y}_0 and \mathbf{X}_1 . Let $\widehat{\mathbf{y}}_1$ denote the fitted values obtained from regressing \mathbf{y}_1 on \mathbf{X}_1 and \mathbf{X}_2 ; that is, $\widehat{\mathbf{y}}_1$ are the fitted values from the first-stage regression for \mathbf{y}_1 , and define the columns of $\widehat{\mathbf{Y}}_0$ analogously. Finally, let $\widetilde{\widetilde{\mathbf{y}}}_1$ denote the residuals from regressing $\widehat{\mathbf{y}}_1$ on $\widehat{\mathbf{Y}}_0$ and \mathbf{X}_1 . Shea's partial R^2 is the simple R^2 from the regression of $\widetilde{\widetilde{\mathbf{y}}}_1$ on $\widehat{\mathbf{y}}_1$; denote this as R_S^2 . Shea's adjusted partial R^2 is equal to $1 - (1 - R_S^2)(N - 1)/(N - k_Z + 1)$ if a constant term is included and $1 - (1 - R_S^2)(N - 1)/(N - k_Z)$ if there is no constant term included in the model, where $k_Z = k_1 + k_2$. For one endogenous regressor, one instrument, no exogenous regressors, and a constant term, R_S^2 equals the adjusted R_S^2 .

The Stock and Yogo minimum eigenvalue statistic, first proposed by Cragg and Donald (1993) as a test for underidentification, is the minimum eigenvalue of the matrix

$$\mathbf{G} = \frac{1}{k_Z} \widehat{\Sigma}_{\mathbf{V}\mathbf{V}}^{-1/2} \mathbf{Y}'\mathbf{M}'_{\mathbf{X}_1} \mathbf{X}_2(\mathbf{X}_2'\mathbf{M}_{\mathbf{X}_1}\mathbf{X}_2)^{-1}\mathbf{X}_2'\mathbf{M}_{\mathbf{X}_1} \mathbf{Y}\widehat{\Sigma}_{\mathbf{V}\mathbf{V}}^{-1/2}$$

where

$$\widehat{\Sigma}_{\mathbf{V}\mathbf{V}} = \frac{1}{N - k_Z} \mathbf{Y}' \mathbf{M}_Z \mathbf{Y}$$

$\mathbf{M}_Z = \mathbf{I} - \mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1}\mathbf{Z}'$, and $\mathbf{Z} = [\mathbf{X}_1 \ \mathbf{X}_2]$. Critical values are obtained from the tables in [Stock and Yogo \(2005\)](#).

estat overid

The [Sargan \(1958\)](#) and [Basmann \(1960\)](#) χ^2 statistics are calculated by running the auxiliary regression

$$\widehat{\mathbf{u}} = \mathbf{Z}\boldsymbol{\delta} + \mathbf{e}$$

where $\widehat{\mathbf{u}}$ are the sample residuals from the model and \mathbf{e} is an error term. Then Sargan's statistic is

$$S = N \left(1 - \frac{\widehat{\mathbf{e}}'\widehat{\mathbf{e}}}{\widehat{\mathbf{u}}'\widehat{\mathbf{u}}} \right)$$

where $\widehat{\mathbf{e}}$ are the residuals from that auxiliary regression. Basmann's statistic is calculated as

$$B = S \frac{N - k_Z}{N - S}$$

Both S and B are distributed $\chi^2(m)$, where m , the number of overidentifying restrictions, is equal to $k_Z - k$, where k is the number of endogenous regressors.

Wooldridge's (1995) score test of overidentifying restrictions is identical to Sargan's (1958) statistic under the assumption of i.i.d. and therefore is not recomputed unless a robust VCE was used at estimation time. If a heteroskedasticity-robust VCE was used, Wooldridge's test proceeds as follows. Let $\widehat{\mathbf{Y}}$ denote the $N \times k$ matrix of fitted values obtained by regressing the endogenous regressors on \mathbf{X}_1 and \mathbf{X}_2 . Let \mathbf{Q} denote an $N \times m$ matrix of excluded exogenous variables; the test statistic to be calculated is invariant to whichever m of the k_2 excluded exogenous variables is chosen. Define the i th element of $\widehat{\mathbf{k}}_j$, $i = 1, \dots, N$, $j = 1, \dots, m$, as

$$k_{ij} = \widehat{q}_{ij} u_i$$

where \widehat{q}_{ij} is the i th element of $\widehat{\mathbf{q}}_j$, the fitted values from regressing the j th column of \mathbf{Q} on $\widehat{\mathbf{Y}}$ and \mathbf{X}_1 . Finally, fit the regression

$$\mathbf{1} = \theta_1 \widehat{\mathbf{k}}_1 + \dots + \theta_m \widehat{\mathbf{k}}_m + \boldsymbol{\epsilon}$$

where $\mathbf{1}$ is an $N \times 1$ vector of ones and $\boldsymbol{\epsilon}$ is a regression error term, and calculate the residual sum of squares, RSS. Then the test statistic is $W = N - \text{RSS}$. $W \sim \chi^2(m)$. If an HAC VCE was used at estimation, then the $\widehat{\mathbf{k}}_j$ are prewhitened using a VAR(p) model, where p is specified using the `lags()` option.

The Anderson–Rubin (1950), AR, test of overidentifying restrictions for use after the LIML estimator is calculated as $\text{AR} = N(\kappa - 1)$, where κ is the minimal eigenvalue of a certain matrix defined in [Methods and formulas](#) of `[R] ivregress`. $\text{AR} \sim \chi^2(m)$. (Some texts define this statistic as $N \ln(\kappa)$ because $\ln(x) \approx (x - 1)$ for x near one.) Basmann's F statistic for use after the LIML estimator is calculated as $B_F = (\kappa - 1)(N - k_Z)/m$. $B_F \sim F(m, N - k_Z)$.

Hansen's J statistic is simply the sample size times the value of the GMM objective function defined in (5) of `[R] ivregress`, evaluated at the estimated parameter values. Under the null hypothesis that the overidentifying restrictions are valid, $J \sim \chi^2(m)$.

References

- Anderson, T. W., and H. Rubin. 1950. The asymptotic properties of estimates of the parameters of a single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics* 21: 570–582.
- Angrist, J. D., and J.-S. Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton, NJ: Princeton University Press.
- Basmann, R. L. 1960. On finite sample distributions of generalized classical linear identifiability test statistics. *Journal of the American Statistical Association* 55: 650–659.
- Baum, C. F., M. E. Schaffer, and S. Stillman. 2003. Instrumental variables and GMM: Estimation and testing. *Stata Journal* 3: 1–31.
- . 2007. Enhanced routines for instrumental variables/generalized method of moments estimation and testing. *Stata Journal* 7: 465–506.
- Bound, J., D. A. Jaeger, and R. M. Baker. 1995. Problems with instrumental variables estimation when the correlation between the instruments and the endogenous explanatory variable is weak. *Journal of the American Statistical Association* 90: 443–450.
- Cameron, A. C., and P. K. Trivedi. 2005. *Microeconometrics: Methods and Applications*. New York: Cambridge University Press.
- Cragg, J. G., and S. G. Donald. 1993. Testing identifiability and specification in instrumental variable models. *Econometric Theory* 9: 222–240.
- Davidson, R., and J. G. MacKinnon. 1993. *Estimation and Inference in Econometrics*. New York: Oxford University Press.
- Durbin, J. 1954. Errors in variables. *Review of the International Statistical Institute* 22: 23–32.
- Hahn, J., and J. A. Hausman. 2003. Weak instruments: Diagnosis and cures in empirical econometrics. *American Economic Review Papers and Proceedings* 93: 118–125.
- Hall, A. R., G. D. Rudebusch, and D. W. Wilcox. 1996. Judging instrument relevance in instrumental variables estimation. *International Economic Review* 37: 283–298.
- Hansen, L. P. 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50: 1029–1054.
- Hausman, J. A. 1978. Specification tests in econometrics. *Econometrica* 46: 1251–1271.
- Hayashi, F. 2000. *Econometrics*. Princeton, NJ: Princeton University Press.
- Judge, G. G., W. E. Griffiths, R. C. Hill, H. Lütkepohl, and T.-C. Lee. 1985. *The Theory and Practice of Econometrics*. 2nd ed. New York: Wiley.
- Nelson, C. R., and R. Startz. 1990. The distribution of the instrumental variable estimator and its t ratio when the instrument is a poor one. *Journal of Business* 63: S125–S140.
- Poi, B. P. 2006. Jackknife instrumental variables estimation in Stata. *Stata Journal* 6: 364–376.
- Sargan, J. D. 1958. The estimation of economic relationships using instrumental variables. *Econometrica* 26: 393–415.
- Shea, J. S. 1997. Instrument relevance in multivariate linear models: A simple measure. *Review of Economics and Statistics* 79: 348–352.
- Staiger, D. O., and J. H. Stock. 1997. Instrumental variables regression with weak instruments. *Econometrica* 65: 557–586.
- Stock, J. H., J. H. Wright, and M. Yogo. 2002. A survey of weak instruments and weak identification in generalized method of moments. *Journal of Business and Economic Statistics* 20: 518–529.
- Stock, J. H., and M. Yogo. 2005. Testing for weak instruments in linear IV regression. In *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, ed. D. W. K. Andrews and J. H. Stock, 80–108. New York: Cambridge University Press.
- Wooldridge, J. M. 1995. Score diagnostics for linear models estimated by two stage least squares. In *Advances in Econometrics and Quantitative Economics: Essays in Honor of Professor C. R. Rao*, ed. G. S. Maddala, P. C. B. Phillips, and T. N. Srinivasan, 66–87. Oxford: Blackwell.
- Wu, D.-M. 1974. Alternative tests of independence between stochastic regressors and disturbances: Finite sample results. *Econometrica* 42: 529–546.

Also see

[R] [ivregress](#) — Single-equation instrumental-variables regression

[U] [20 Estimation and postestimation commands](#)