

A general approach to testing for autocorrelation

Christopher F Baum & Mark E Schaffer

Boston College/DIW Berlin Heriot–Watt University/CEPR/IZA

Stata Conference, New Orleans, July 2013

Introduction

Testing for autocorrelation in a time series is a common task for researchers working with time-series data.

We present a new Stata command, `actest`, which generalizes our earlier `ivactest` (Baum, Schaffer, Stillman, *Stata Journal* 7:4, 2007) and provides a more versatile framework for autocorrelation testing.

Background

The standard Q test statistic, Stata's `wntestq` (Box and Pierce, 1970), refined by Ljung and Box (1978), is applicable for univariate time series under the assumption of strictly exogenous regressors.

Breusch (1978) and Godfrey (1978) in effect extended the B-P-L-B approach (Stata's `estat bgodfrey`, B-G) to test for autocorrelation in models with weakly exogenous regressors.

Although these tests are more general and much more useful than tests that consider only the AR(1) alternative, such as the Durbin–Watson statistic, the B-P-L-B and B-G tests have important limitations.

The B-P-L-B and Breusch–Godfrey tests are not applicable:

- when serial correlation up to order q is expected to be present, so they cannot test for serial correlation at orders $q + 1, q + 2 \dots$ for $q > 0$
- when the model contains endogenous regressors and is thus estimated by IV or IV-GMM
- in the context of overlapping data, as we often encounter in the financial markets
- in the presence of conditional heteroskedasticity in the error process

Cumby and Huizinga (1992) provide a framework that extends the implementation of the Q statistic to deal with these limitations. Their test also allows for testing for autocorrelation of order $(q + 1)$ through $(q + s)$, where under the null hypothesis there may be autocorrelation of order q or less in the form of $MA(q)$. Their test may also be applied in the context of panel data.

The Baum–Schaffer–Stillman `ivreg2` package, as described in *Stata Journal* (2007), contains the `ivactest` command, which implements the Cumby–Huizinga (C-H) test after OLS, IV, IV-GMM and LIML estimation.

We present an enhanced and extended command, `actest`, for the testing of autocorrelation in the errors of OLS, IV, IV-GMM and LIML estimates for a single time series, including testing for autocorrelation at specific lag orders.

We demonstrate the relationship between the C-H test, developed for the large- T setting, and the test for $AR(p)$ in a large- N setting, developed by Arellano and Bond (1991) and implemented by Roodman as `abar` for application to a single residual series. Our `actest` command may also be applied in the panel context, and reproduces results of the `abar` test in a variety of settings.

Earlier tests for multiple orders of autocorrelation

The first tests for autocorrelation, based on the alternative of an $AR(1)$ model of the error process, only considered that possible departure from independence. From a pedagogical standpoint, such a test is dangerous, as a failure to reject may be taken as a clean bill of health, implying the absence of serial correlation: which it is not.

The Box–Pierce portmanteau (or Q) test, developed in 1970, may be applied to a univariate time series, and is often considered to be a general test for ‘white noise’: thus its name in Stata, `wntestq`. The test implemented by that command is the refinement proposed by Ljung and Box (1978), implementing a small-sample correction.

However, if the portmanteau test is applied to a set of regression residuals, the regressors in the model are assumed to be strictly exogenous and homoskedastic.

For illustration, we compute the Q statistic for one lag, and illustrate its computation via `actest`. The `bp` option specifies the Q test, and `small` indicates that the Ljung–Box form of the statistic, with its small sample correction, is to be computed. Without the `small` option, the original Box–Pierce statistic will be computed.

```
. wntestq air, lags(1)
Portmanteau test for white noise
-----
Portmanteau (Q) statistic = 132.1415
Prob > chi2(1)           = 0.0000
```

```
. actest air, lags(1) bp small
Cumby–Huizinga test for autocorrelation
H0: variable is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=0 (serially uncorrelated) HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) =132.142	0.0000	1	Chi-sq(1) =132.142	0.0000

Test requires conditional homoskedasticity

As you can see from the output, `actest` automatically displays a test statistic for all specified lags, as well as a test for each lag order. In the single-lag case, these are identical. The null hypothesis is that the variable tested is a moving average process of order q : $MA(q)$. By default, $q = 0$, implying white noise. The alternatives considered is that serial correlation is present in that range of lags, or for that specified lag.

For a single lag, the Ljung–Box portmanteau statistic is identical to the Cumby–Huizinga (C-H) test statistic. We may also apply each test for a range of lag orders:

```
. wntestq air, lags(4)
Portmanteau test for white noise
-----
Portmanteau (Q) statistic =    427.7387
Prob > chi2(4)           =         0.0000
```

```
. actest air, lags(4) bp small
Cumby–Huizinga test for autocorrelation
H0: variable is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=0 (serially uncorrelated) HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) =132.142	0.0000	1	Chi-sq(1) =132.142	0.0000
1 - 2	Chi-sq(2) =245.646	0.0000	2	Chi-sq(1) =113.505	0.0000
1 - 3	Chi-sq(3) =342.675	0.0000	3	Chi-sq(1) = 97.029	0.0000
1 - 4	Chi-sq(4) =427.739	0.0000	4	Chi-sq(1) = 85.064	0.0000

Test requires conditional homoskedasticity

For the range of lags 1–4, the C-H statistic is identical to the Ljung–Box Q reported by `wntestq`. The right-hand panel also indicates that serial correlation is present at each lag. Those findings cannot be produced by the B-P-L-B test, as its null hypothesis assumes the absence of autocorrelation at all lags.

The Breusch–Godfrey test, developed independently by those two authors in 1978 publications, is meant to be applied to a set of regression residuals under the assumption of weakly exogenous, or predetermined, regressors. Although its implementation in official Stata as `estat bgodfrey` classifies it as a post-estimation command, it may be applied to a single time series by regressing that series on a constant:

```
. qui reg air
. estat bgodfrey, lags(1)
Breusch–Godfrey LM test for autocorrelation
```

lags (p)	chi2	df	Prob > chi2
1	130.900	1	0.0000

H0: no serial correlation

In this case, the regressor (the units vector) is of course strictly exogenous.

Our `actest` also functions as a post-estimation command, so that if no *varname* is specified, it operates on the residual series of the last estimation command:

```
. actest, lags(1)
Cumby-Huizinga test for autocorrelation
H0: disturbance is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) =130.900	0.0000	1	Chi-sq(1) =130.900	0.0000

Test allows predetermined regressors/instruments
 Test requires conditional homoskedasticity

The `actest` statistic is identical to that produced by the B-G test.

The advantage of the B-G test over tests for $AR(1)$ is that it may be applied to test a null hypothesis over a range of lag orders:

```
. estat bgodfrey, lags(4)  
Breusch–Godfrey LM test for autocorrelation
```

lags (p)	chi2	df	Prob > chi2
4	132.364	4	0.0000

H0: no serial correlation

We may reproduce the B–G test results with `actest` for the same number of lags:

```
. actest, lags(4)
Cumby–Huizinga test for autocorrelation
H0: disturbance is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) =130.900	0.0000	1	Chi-sq(1) =130.900	0.0000
1 - 2	Chi-sq(2) =131.954	0.0000	2	Chi-sq(1) = 40.202	0.0000
1 - 3	Chi-sq(3) =132.208	0.0000	3	Chi-sq(1) = 22.708	0.0000
1 - 4	Chi-sq(4) =132.364	0.0000	4	Chi-sq(1) = 15.970	0.0001

Test allows predetermined regressors/instruments
 Test requires conditional homoskedasticity

The `actest` statistic for the range of lags 1–4 is identical to the B-G statistic. Note that on the right-hand panel, the null for each specific lag is that the process is $MA(lag - 1)$ rather than $MA(lag)$.

This hypothesis cannot be tested by B-G, as under its null hypothesis there is no autocorrelation at any lag order. It makes no sense to test for autocorrelation, say, at the 4th lag while assuming that it is not present at any lower lag order.

However, the B-P-L-B and B-G tests, and the C-H test in its default form, are all based upon conditional homoskedasticity of the error process. We can relax this assumption in `actest` by specifying the **robust** option:

```
. actest, lags(4) robust
Cumby–Huizinga test for autocorrelation
H0: disturbance is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 55.852	0.0000	1	Chi-sq(1) = 55.852	0.0000
1 - 2	Chi-sq(2) = 59.940	0.0000	2	Chi-sq(1) = 20.886	0.0000
1 - 3	Chi-sq(3) = 63.790	0.0000	3	Chi-sq(1) = 13.761	0.0002
1 - 4	Chi-sq(4) = 65.304	0.0000	4	Chi-sq(1) = 10.526	0.0012

```
Test allows predetermined regressors/instruments
Test robust to heteroskedasticity
```

The test for lag orders 1–4 again strongly rejects the null of independence in the series, as does the test at each individual lag.

The Cumby–Huizinga test in perspective

In each of these examples, we have performed a test on a univariate time series. Each test may be applied to the residuals of a nontrivial regression model under the assumption of strict exogeneity (B-P-L-B), or weakly exogenous or predetermined regressors (B-G):

```
. qui reg air time
. qui predict double airhat, residual
. wntestq airhat, lags(4)
Portmanteau test for white noise


---


Portmanteau (Q) statistic = 107.6173
Prob > chi2(4)           = 0.0000
```

To reproduce these results with `actest`, we must also employ the `strict` option to specify that the regressors are assumed to be strictly exogenous:

```
. actest, lags(4) bp small strict
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags $>q$

H0: $q=0$ (serially uncorrelated) HA: s.c. present at range specified			H0: $q=0$ (serially uncorrelated) HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 77.958	0.0000	1	Chi-sq(1) = 77.958	0.0000
1 - 2	Chi-sq(2) = 90.266	0.0000	2	Chi-sq(1) = 12.308	0.0005
1 - 3	Chi-sq(3) = 91.425	0.0000	3	Chi-sq(1) = 1.159	0.2816
1 - 4	Chi-sq(4) = 107.617	0.0000	4	Chi-sq(1) = 16.192	0.0001

Test requires strictly exogenous regressors/instruments

Test requires conditional homoskedasticity

The `actest` statistic for lag orders 1–4 is identical to the Q statistic.

The B-P-L-B statistic can be modified to allow for predetermined but not strictly exogenous regressors. As Hayashi shows in his textbook (p. 146–147), this requires ‘predeterminedness’ and a strong form of conditional homoskedasticity: that the expectation of the error conditioned on both its own history and the history of the regressors is zero, and that the expectation of the squared error under the same conditioning is σ_u^2 . Hayashi calls this the ‘modified Box–Pierce Q ’, and shows that it is asymptotically equivalent to the B-G test statistic.

We demonstrate the equivalency between B-G and `actest`:

```
. estat bgodfrey, lags(4)
```

Breusch–Godfrey LM test for autocorrelation

lags (p)	chi2	df	Prob > chi2
4	95.947	4	0.0000

H0: no serial correlation

```
. actest, lags(4)
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags $>q$

H0: $q=0$ (serially uncorrelated) HA: s.c. present at range specified			H0: $q=\text{specified lag}-1$ HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 76.740	0.0000	1	Chi-sq(1) = 76.740	0.0000
1 - 2	Chi-sq(2) = 94.492	0.0000	2	Chi-sq(1) = 5.936	0.0148
1 - 3	Chi-sq(3) = 95.007	0.0000	3	Chi-sq(1) = 0.511	0.4748
1 - 4	Chi-sq(4) = 95.947	0.0000	4	Chi-sq(1) = 7.218	0.0072

Test allows predetermined regressors/instruments

Test requires conditional homoskedasticity

The equivalency also holds for predetermined, or weakly exogenous, regressors in this $AR(2)$ model:

```
. qui reg air L(1/2).air
. estat bgodfrey, lags(4)
```

Breusch–Godfrey LM test for autocorrelation

lags (p)	chi2	df	Prob > chi2
4	15.506	4	0.0038

H0: no serial correlation

```
. actest, lags(4)
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags $>q$

H0: $q=0$ (serially uncorrelated) HA: s.c. present at range specified			H0: $q=\text{specified lag}-1$ HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 5.409	0.0200	1	Chi-sq(1) = 5.409	0.0200
1 - 2	Chi-sq(2) = 7.979	0.0185	2	Chi-sq(1) = 2.578	0.1084
1 - 3	Chi-sq(3) = 12.490	0.0059	3	Chi-sq(1) = 1.755	0.1853
1 - 4	Chi-sq(4) = 15.506	0.0038	4	Chi-sq(1) = 7.387	0.0066

Test allows predetermined regressors/instruments

Test requires conditional homoskedasticity

Beyond the ability to compute a robust version of the test, allowing for conditional heteroskedasticity, the C-H framework also allows us to consider the null hypothesis as allowing for serial correlation at some lower lag order. To illustrate, we assume that the error process is $MA(2)$ under the null:

```
. actest, lags(3 4)
```

```
Cumby–Huizinga test for autocorrelation
```

```
H0: disturbance is MA process up to order q
```

```
HA: serial correlation present at specified lags >q
```

H0: disturbance is $MA(q)$, $q=2$ HA: s.c. present at range specified			H0: $q=\text{specified lag}-1$ HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
3 - 3	Chi-sq(1) = 1.755	0.1853	3	Chi-sq(1) = 1.755	0.1853
3 - 4	Chi-sq(2) = 9.046	0.0109	4	Chi-sq(1) = 7.387	0.0066

```
Test allows predetermined regressors/instruments
```

```
Test requires conditional homoskedasticity
```

For lag orders 3–4, the null that the residuals are $MA(2)$ rather than $MA(4)$ is rejected.

The ability to allow for lower-order serial correlation implies that the C-H test framework is considerably more flexible than those of the earlier tests, which all assume no serial correlation under the null: in `actest` terms, that $q = 0$. Allowing for $MA(q)$ under the null requires the use of a kernel-robust VCE, which is the truncated kernel with bandwidth set to q .

Hayashi points out that the truncated kernel is a natural kernel to use when the autocorrelation dies out at a predetermined lag q , obviating the need for large- T asymptotics when the bandwidth increases with T . As in Baum–Schaffer–Stillman’s `ivreg2`, the default when using the `kernel()` or `bw()` options is to compute an AC-robust VCE. To compute a HAC VCE, the `robust` option should also be specified.

Another useful feature of the C-H testing framework is that it can give us greater insight into the form of dependence in the error process. For instance, if we looked at this regression with the B-G test:

```
. qui reg investment L(1/4).income
. estat bgodfrey, lags(1/8)
Breusch–Godfrey LM test for autocorrelation
```

lags (p)	chi2	df	Prob > chi2
1	64.511	1	0.0000
2	64.601	2	0.0000
3	64.641	3	0.0000
4	64.641	4	0.0000
5	65.147	5	0.0000
6	65.438	6	0.0000
7	65.750	7	0.0000
8	66.566	8	0.0000

H0: no serial correlation

We conclude that there is serious autocorrelation at all lag lengths.

However, consider the right-hand panel of the equivalent `actest` results:

```
. actest, lags(8)
```

```
Cumby–Huizinga test for autocorrelation
```

```
H0: disturbance is MA process up to order q
```

```
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 64.511	0.0000	1	Chi-sq(1) = 64.511	0.0000
1 - 2	Chi-sq(2) = 64.601	0.0000	2	Chi-sq(1) = 18.232*	0.0000
1 - 3	Chi-sq(3) = 64.641	0.0000	3	Chi-sq(1) = 9.245*	0.0024
1 - 4	Chi-sq(4) = 64.641	0.0000	4	Chi-sq(1) = 5.799	0.0160
1 - 5	Chi-sq(5) = 65.147	0.0000	5	Chi-sq(1) = 3.211	0.0731
1 - 6	Chi-sq(6) = 65.438	0.0000	6	Chi-sq(1) = 1.402	0.2364
1 - 7	Chi-sq(7) = 65.750	0.0000	7	Chi-sq(1) = 0.388	0.5335
1 - 8	Chi-sq(8) = 66.566	0.0000	8	Chi-sq(1) = 0.002	0.9617

```
Test allows predetermined regressors/instruments
```

```
Test requires conditional homoskedasticity
```

```
* Eigenvalues adjusted to make matrix positive semidefinite
```

In this case, we can see that although the joint tests of lag orders all soundly reject, the test for $MA(4)$ vs. $MA(5)$ cannot reject at the 95% level. This suggests that the conclusion of the B-G test is being strongly influenced by the clear autocorrelation at lags 1–4, and might lead us to including more lags than necessary in a HAC estimator of the VCE.

The C-H framework also relaxes the assumption of predetermined regressors, as in an IV context, the requirement for predeterminedness is applied to the *instruments* rather than the *regressors*. To illustrate:

```
. webuse lutkepohl, clear
(Quarterly SA West German macro data, Bil DM, from Lutkepohl 1993 Table E.1)
. qui ivreg2 investment (income = L(1/2).income)
. actest, lags(3)
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags >q

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 68.947	0.0000	1	Chi-sq(1) = 68.947	0.0000
1 - 2	Chi-sq(2) = 69.029	0.0000	2	Chi-sq(1) = 21.716	0.0000
1 - 3	Chi-sq(3) = 69.182	0.0000	3	Chi-sq(1) = 12.362	0.0004

Test allows predetermined regressors/instruments

Test requires conditional homoskedasticity

We may also conduct this test as heteroskedasticity-robust:

```
. actest, lags(3) robust
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags >q

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 37.402	0.0000	1	Chi-sq(1) = 37.402	0.0000
1 - 2	Chi-sq(2) = 37.616	0.0000	2	Chi-sq(1) = 12.128	0.0005
1 - 3	Chi-sq(3) = 37.631	0.0000	3	Chi-sq(1) = 7.003	0.0081

Test allows predetermined regressors/instruments

Test robust to heteroskedasticity

In both forms of the test, the null hypothesis is overwhelmingly rejected.

Cumby–Huizinga test vs. Arellano–Bond test

Arellano and Bond (1991) introduced a test for autocorrelation in dynamic panel data (DPD) estimation for the fixed T , large N context supported by `xtabond` et al. Roodman (2004, 2012) implemented this as a standalone test, `abar`.

The A-B test was originally devised for DPD models, in which there is $AR(1)$ (actually, $MA(1)$) present in the differenced errors by construction, the presence of significant $AR(2)$ is a diagnostic test of the validity of the instruments, complementary to the standard Sargan–Hansen test of overidentifying restrictions. In this context, the null allows for $AR(1)$ ($q > 0$ in C-H terms) while testing for $AR(2)$, $AR(3)$...

Roodman noted that this test is ‘quite general in its applicability—more general than `dwstat`, `durbina`, `bgodfrey` and `xtserial`.’¹ and that ‘it can be applied to linear GMM regressions in general, and thus to the special cases of OLS and 2SLS.’

The `abar` test may be applied after `regress`, `ivreg`, `ivregress` and `ivreg2` (Baum–Schaffer–Stillman) for the homoskedastic, robust, and cluster-robust forms of those commands, as well as regressions with HAC VCEs estimated by `newey`, `newey2` (Roodman), `ivregress` and `ivreg2`. It may be applied to fixed-effects models estimated with these commands, but is not appropriate for fixed effects models with fixed- T large- N asymptotics (Wooldridge, MIT Press, pp. 310–311).

¹`abar` help file

The A-B test implemented in `abar` for linear DPD models is essentially the C-H test in the panel context. Thus, there is general equivalence between the `abar` test applied to an OLS, 2SLS or IV-GMM regression and the C-H test implemented by `actest`.

Some differences exist, as `actest` defaults to an AC-robust version of the C-H test, supporting the truncated kernel. Whereas `abar` reports tests of serial correlation at individual lag orders only, `actest` also reports tests at ranges of lag orders. Unless explicitly specified in the original estimation, the `abar` test assumes no autocorrelation under the null ($q = 0$).

In contrast, when testing for a particular lag q , `actest` allows for autocorrelation (in the form of $MA(q - 1)$) at lower lag orders. This default behavior of `abar` makes it less attractive than `actest`, as we may often want to accommodate autocorrelation at a lower lag order and not assume its absence.

We illustrate this equivalence by testing for $AR(1)$... $AR(4)$ with `abar` following an OLS regression with classical VCE. The `abar` statistic is standard Normal under the null, so we convert its reported results to χ^2 tests.

```
. qui reg investment income
. abar, lags(4)
Warning: The Arellano-Bond test is only valid for time series only if they are
> ergodic.
Arellano-Bond test for AR(1): z = 8.33 Pr > z = 0.0000
Arellano-Bond test for AR(2): z = 7.43 Pr > z = 0.0000
Arellano-Bond test for AR(3): z = 6.77 Pr > z = 0.0000
Arellano-Bond test for AR(4): z = 5.73 Pr > z = 0.0000
. forv i=1/4 {
2. loc ar`i' = r(ar`i')^2 * e(N)/e(df_r)
3. loc dia "`dia' `ar`i'`"
4. }
. di "`dia'"
70.90516463704326 56.45707136113094 46.89961653711498 33.59539738062808
```


In order to compare with `actest`, we use the `q0` option, which specifies that no serial correlation is assumed under the null for individual lag-order tests:

```
. actest, lags(4) q0
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags $>q$

H0: $q=0$ (serially uncorrelated) HA: s.c. present at range specified			H0: $q=0$ (serially uncorrelated) HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 70.905	0.0000	1	Chi-sq(1) = 70.905	0.0000
1 - 2	Chi-sq(2) = 70.973	0.0000	2	Chi-sq(1) = 56.457	0.0000
1 - 3	Chi-sq(3) = 71.083	0.0000	3	Chi-sq(1) = 46.900	0.0000
1 - 4	Chi-sq(4) = 71.881	0.0000	4	Chi-sq(1) = 33.595	0.0000

Test allows predetermined regressors/instruments

Test requires conditional homoskedasticity

The `abar` statistics are equal to those in the right-hand panel above.

We may also employ the test after robust or HAC estimation of the VCE. In the latter case, we employ the Bartlett kernel as is used in the Newey–West HAC estimator. Bandwidth=5 in `ivreg2` terms implies four lags in the kernel.

```
. qui ivreg2 investment income, robust kernel(bartlett) bw(5)
. abar, lags(4)
Warning: The Arellano-Bond test is only valid for time series only if they are
> ergodic.
Arellano-Bond test for AR(1): z =    3.01   Pr > z = 0.0026
Arellano-Bond test for AR(2): z =    2.86   Pr > z = 0.0042
Arellano-Bond test for AR(3): z =    2.72   Pr > z = 0.0065
Arellano-Bond test for AR(4): z =    2.43   Pr > z = 0.0151
. forv i=1/4 {
2.     loc ar`i' = r(ar`i')^2
3.     loc dia "`dia' `ar`i'`"
4. }
. di "`dia'"
9.041350773390105 8.191274390789095 7.403328244433819 5.9082939073147
```

In order to compare with `actest`, we add the `kernel()` and `bw()` options to indicate that the test should be computed in a HAC context:

```
. actest, lags(4) q0 robust kernel(bartlett) bw(5)
Cumby–Huizinga test for autocorrelation
H0: disturbance is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=0 (serially uncorrelated) HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 9.041	0.0026	1	Chi-sq(1) = 9.041	0.0026
1 - 2	Chi-sq(2) = 9.380	0.0092	2	Chi-sq(1) = 8.191	0.0042
1 - 3	Chi-sq(3) = 9.628	0.0220	3	Chi-sq(1) = 7.403	0.0065
1 - 4	Chi-sq(4) = 9.643	0.0469	4	Chi-sq(1) = 5.908	0.0151

```
Test allows predetermined regressors/instruments
Test robust to heteroskedasticity
```

The `abar` statistics are equal to those in the right-hand panel above.

We may also apply the test in a panel context using pooled OLS:

```
. webuse abdata, clear
. qui reg n w k, clu(id)
. abar, lags(3)
Arellano-Bond test for AR(1): z =    5.92   Pr > z = 0.0000
Arellano-Bond test for AR(2): z =    5.76   Pr > z = 0.0000
Arellano-Bond test for AR(3): z =    5.62   Pr > z = 0.0000
. forv i=1/3 {
  2.      loc ar`i' = r(ar`i')^2
  3.      loc dia "`dia' `ar`i''"
  4. }
. di "`dia'"
35.01428083793625 33.12713487950008 31.63010894000373
```

The `abar` test in this context is robust to within-panel autocorrelation.

In order to compare with `actest`, we add the `cluster()` option to indicate that the test should be computed in a cluster-robust context:

```
. actest, lags(3) clu(id)
Cumby–Huizinga test for autocorrelation
H0: disturbance is MA process up to order q
HA: serial correlation present at specified lags >q
```

H0: q=0 (serially uncorrelated) HA: s.c. present at range specified			H0: q=specified lag-1 HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 35.014	0.0000	1	Chi-sq(1) = 35.014	0.0000
1 - 2	Chi-sq(2) = 56.630	0.0000	2	Chi-sq(1) = 33.127	0.0000
1 - 3	Chi-sq(3) = 58.136	0.0000	3	Chi-sq(1) = 31.630	0.0000

Test allows predetermined regressors/instruments

Test robust to heteroskedasticity and within-cluster autocorrelation

The `abar` statistics are equal to those in the right-hand panel above.

The tests are also equivalent in this panel setting when we estimate a model fit to first differences via instrumental variables (IV-GMM) or LIML with a cluster-robust VCE. We illustrate IV-GMM:

```
. qui ivreg2 D.n (D.w D.k = D(1/2).(w k)), noco gmm2s clu(id)
. abar, lags(3)
Arellano-Bond test for AR(1): z =    3.97   Pr > z = 0.0001
Arellano-Bond test for AR(2): z =    1.81   Pr > z = 0.0705
Arellano-Bond test for AR(3): z =    0.42   Pr > z = 0.6739
. forv i=1/3 {
  2.         loc ar`i` = r(ar`i`)^2
  3.         loc dia "`dia' `ar`i`'"
  4. }
. di "`dia'"
15.74212072538034 3.271292893080344 .1771184678625001
```

The same test statistics and p -values for each lag order are produced by `actest`:

```
. actest, lags(3) clu(id)
```

Cumby–Huizinga test for autocorrelation

H0: disturbance is MA process up to order q

HA: serial correlation present at specified lags $>q$

H0: $q=0$ (serially uncorrelated) HA: s.c. present at range specified			H0: $q=\text{specified lag}-1$ HA: s.c. present at lag specified		
lags	chi2	p-val	lag	chi2	p-val
1 - 1	Chi-sq(1) = 15.742	0.0001	1	Chi-sq(1) = 15.742	0.0001
1 - 2	Chi-sq(2) = 16.547	0.0003	2	Chi-sq(1) = 3.271	0.0705
1 - 3	Chi-sq(3) = 16.661	0.0008	3	Chi-sq(1) = 0.177	0.6739

Test allows predetermined regressors/instruments

Test robust to heteroskedasticity and within-cluster autocorrelation

In the panel context, we are considering whether `actest` should also accept residuals produced by `areg`, as in that framework the partialled-out fixed effects can be treated as predetermined, so that application of the C-H test is straightforward.

actest syntax

The current version of `actest` may be employed with a *varname*, in which case that variable is tested; otherwise, it is assumed that an appropriate estimation command has been previously executed, and the residuals from that command are to be tested.

As we have demonstrated, `actest` implements a number of options that allow it to match, the results of a number of other tests for autocorrelation. These include:

- `lags (numlist)` : specifies the lag orders at which autocorrelation is to be tested. If a single value, tested up to that value. If a *numlist*, tested for that range of lags, assuming autocorrelation at lower lag orders under the null.
- `strictexog`: regressors in prior estimation are assumed to be strictly exogenous, as they are in B-P-L-B tests.

- `q0`: for single lag-order tests, null hypothesis specifies no autocorrelation ($q = 0$).
- `bp`: perform the Box–Pierce test.
- `small`: perform the Ljung–Box variant of the Box–Pierce test, with small-sample correction.
- `robust`: make test robust to arbitrary heteroskedasticity in the error process.
- `cluster(varlist)`: make test cluster-robust to specified variable(s): two-way clustering is supported, as in `ivreg2`.

- `bw (#)` : make test robust to arbitrary autocorrelation, using the specified bandwidth in kernel estimator. This is the appropriate VCE to use, in conjunction with the default truncated kernel, when you know the degree of autocorrelation under the null. This is the case for overlapping data, where a given $MA(q)$ process is induced.
- `kernel (string)` : make test robust to arbitrary autocorrelation, using specified kernel (per choices in `ivreg2`. Caution: generally, the default truncated kernel will be appropriate for HAC-robustified tests.
- `psd (string)` : some kernel-robust VCEs are not guaranteed to produce positive semidefinite VCEs in finite samples. Default behavior: replace negative eigenvalues with absolute values, per Stock and Watson, *Econometrica*, 2008. With the `psd(psd0)` option, negative eigenvalues are replaced with zeros.

Some housekeeping details:

- `actest` essentially supersedes `ivactest`, as described in Baum–Schaffer–Stillman, *Stata Journal* 7:4. Users of the earlier routine should `ssc install actest`.
- `actest` makes use of the same Mata object library, `livreg2.mlib`, used by the Baum–Schaffer–Stillman `ivreg2` package in its recent versions, providing access to all VCE options in `ivreg2` such as two-way clustering. The library will be automatically installed with `actest`.
- Some relevant references: Cumby and Huizinga, *Econometrica* 60:1, 1992, 185–195; Cumby and Huizinga, NBER Technical Working Paper 90, 1990; Arellano and Bond, *Review of Economic Studies* 58:2, 1991, 277–297; Roodman, *Stata Journal* 12:4, 2012, 766–767.