Improved tests for Granger noncausality in panel data

Jiaqi Xiao
University of Birmingham
Birmingham, U.K.
jxx963@outlook.com

Artūras Juodis
University of Amsterdam
Amsterdam, The Netherlands
a.juodis@uva.nl

Yiannis Karavias
University of Birmingham
Birmingham, U.K.
i.karavias@bham.ac.uk

Vasilis Sarafidis
BI Norwegian Business School
Oslo, Norway
vasilis.sarafidis@bi.no

Jan Ditzen
Free University of Bozen-Bolzano
Bozen, Italy
jan.ditzen@unibz.it

Abstract. In this article, we introduce the `xtgrangert` command, which implements the panel Granger noncausality testing approach developed by Juodis, Karavias, and Sarafidis (2021, *Empirical Economics* 60: 93–112). This test offers superior size and power performance to existing tests, which stem from the use of a pooled estimator that has a faster $\sqrt{NT}$ convergence rate. The test has several other useful properties: it can be used in multivariate systems; it has power against both homogeneous and heterogeneous alternatives; and it allows for cross-section dependence and cross-section heteroskedasticity.

Keywords: st0706, xtgrangert, xtgrangert postestimation, panel data, Granger causality, Nickell bias, heterogeneous panels, half-panel jackknife, cross-section dependence

1   Introduction

Predictive (Granger) causality and feedback is an important aspect of applied time-series and longitudinal panel-data analysis. Granger (1969) developed a statistical concept of causality between two or more time-series variables, according to which a variable x "Granger-causes" a variable y if the variable y can be better predicted using past values of both x and y rather than using solely past values of y. The concept of "Granger causality" has been widely adopted in economics, medicine, chemistry, physics, biology, engineering, and beyond.

Granger causality is also useful when the data consist of multiple time series, as in the case of panel data. Methods on testing for Granger causality using panel-data models are very well cited and widely available in standard econometric software. Prominent

examples include the generalized method of moments (GMM) approach of Holtz-Eakin, Newey, and Rosen (1988), which is valid for homogeneous panels with a few time-series observations (T), and the methods of Dumitrescu and Hurlin (2012) and Emirmahmutoglu and Kose (2011), suitable for heterogeneous, large-T panels. The GMM approach of Holtz-Eakin, Newey, and Rosen (1988) has been implemented in Stata by Abrigo and Love (2016) with the command `pvargranger`, whereas the method of Dumitrescu and Hurlin (2012) is available in both EViews and Stata; see, for example, the command `xtgcause` by Lopez and Weber (2017).

Recently, Juodis, Karavias, and Sarafidis (2021) developed a new method for testing the null hypothesis of no Granger causality, which is valid in models with homogeneous or heterogeneous coefficients. The novelty of their approach lies in the fact that under the null hypothesis, the Granger-causality parameters equal zero, and thus they are homogeneous. This allows the use of a pooled fixed effects-type estimator for these parameters only, which guarantees a $\sqrt{NT}$ convergence rate, where N denotes the number of cross-sectional units in the panel and T denotes the number of time-series observations in the panel.[1] To account for the so-called Nickell bias of the pooled estimator, their testing procedure makes use of the half-panel jackknife (HPJ) method of Dhaene and Jochmans (2015). The resulting approach works very well under circumstances that are empirically relevant: many cross-section units, a moderate time dimension, heterogeneous nuisance parameters, and high persistence.

The method of Juodis, Karavias, and Sarafidis (2021) has a number of advantages relative to existing approaches. In particular, the GMM approach of Holtz-Eakin, Newey, and Rosen (1988) is not appealing when T is (even moderately) large. This is due to the well-known problem of using too many instruments, which often renders the usual GMM-based inference highly inaccurate; see, for example, Bun and Sarafidis (2015) and remark 8 in Juodis and Sarafidis (2022). Moreover, when feedback based on past own values is heterogeneous (that is, the autoregressive parameters vary across individuals), inferences may not be valid even asymptotically. On the other hand, while the method of Dumitrescu and Hurlin (2012) accommodates heterogeneous slopes under both null and alternative hypotheses, their test statistic is theoretically justified only for sequences where $N/T^2 \to 0$. This implies that when T is sufficiently smaller than N, that is, $T << N$, this method can suffer from substantial size distortions. In an extended Monte Carlo experiment, Juodis, Karavias, and Sarafidis (2021) show that their method outperforms the method of Dumitrescu and Hurlin (2012) in terms of power.

The present article introduces a new command, `xtgrangert`, that implements the Granger noncausality test of Juodis, Karavias, and Sarafidis (2021). The command reports the Wald test statistic and its p-value, the null and the alternative hypotheses, and regression results with respect to the HPJ bias-corrected pooled estimator. The command offers options for both manual and automatic lag-length selection, using a Bayesian information criterion (BIC). The command further allows for cross-sectional dependence and cross-sectional heteroskedasticity in the errors. Finally, the command

---

1. The autoregressive parameters and intercepts (fixed effects) are still allowed to be heterogeneous.

can test for Granger causality in equations with multiple relevant variables.[2] The panel must be balanced.

Notably, by construction `xtgrangert` is computationally faster than `xtgcause`, especially so when N is relatively large. This is because the former is based on a single, pooled regression, whereas the latter runs N individual regressions and retrieves N individual-specific Wald test statistics, which are subsequently averaged over i.[3]

The `xtgrangert` command is applied to a real dataset from the U.S. banking industry, where we perform Granger noncausality tests to examine the type of temporal relation between profitability, cost inefficiency, and asset quality. Our results show that past values of inefficiency contain information that helps to predict profitability, while this is not the case for asset quality.

The remainder of the article is organized as follows. Section 2 briefly outlines the Wald test approach developed by Juodis, Karavias, and Sarafidis (2021). Section 3 describes the syntax of the `xtgrangert` command. Section 4 illustrates the command using a real dataset. Section 5 concludes.

2   A bias-corrected test for Granger noncausality

We consider the following linear dynamic panel-data model,

$$y_{i,t} = \phi_{0,i} + \sum_{p=1}^{P} \phi_{p,i} y_{i,t-p} + \sum_{p=1}^{P} \beta_{p,i} x_{i,t-p} + \varepsilon_{i,t} \tag{1}$$

for $i = 1, \dots, N$ and $t = 1, \dots, T$. Without loss of generality and for ease of exposition, $x_{i,t}$ is assumed to be a scalar. The parameters $\phi_{0,i}$ denote the individual-specific effects, $\varepsilon_{i,t}$ are the errors, $\phi_{p,i}$ denote the heterogeneous autoregressive coefficients, $p = 1, \dots, P$, and $\beta_{p,i}$ are the heterogeneous feedback coefficients, or Granger-causality parameters.

The restriction that the number of lags of $y_{i,t}$ is the same as that of $x_{i,t}$ has the benefit of a minimal computational cost when it comes to lag-length selection. Such restriction is also imposed by `xtgcause` and `pvargranger`.

The null hypothesis that $x_{i,t}$ does not Granger-cause $y_{i,t}$ can be formulated as a set of linear restrictions on the parameters in (1):

$$H_0 \colon \beta_{p,i} = 0, \quad \text{for all } i \text{ and } p$$

The alternative hypothesis is

$$H_1 \colon \beta_{p,i} \neq 0 \quad \text{for some } i \text{ and } p$$

---

2. The command does not consider Granger-causal relations that exist only more than one period ahead; see, for example, Dufour and Renault (1998).

3. To provide some indication of the likely computational gains of our method, in the application of this article ($N = 450$, $T = 56$), we note that when the maximum number of lags equals 5, `xtgrangert` requires roughly 1 second to test the null hypothesis, whereas `xtgcause` takes about 33 seconds.

Failure to reject the null hypothesis can be interpreted as $x_{i,t}$ not Granger-causing $y_{i,t}$.[4] The same applies when $x_{i,t}$ consists of multiple relevant variables and is a $k \times 1$ vector of regressors.

The main feature of the above setup, utilized in the Granger noncausality test proposed by Juodis, Karavias, and Sarafidis (2021), is that under the null hypothesis, $\beta_{p,i} = 0$, for all i and p. In other words, the model is homogeneous in the feedback coefficients. This allows the use of a pooled estimator for $\{\beta_{p,i}\}_{i=1}^{N}$. Pooled estimators have a faster rate of convergence, $\sqrt{NT}$, which means that they benefit from a larger value of both N and T. However, they are subject to the so-called Nickell bias. Juodis, Karavias, and Sarafidis (2021) propose that this bias is corrected using the HPJ method of Dhaene and Jochmans (2015). Although bias corrections have been previously shown to reduce the power of tests (Karavias and Tzavalis 2016, 2017), Juodis, Karavias, and Sarafidis (2021) demonstrate that this test has very good power in empirically relevant scenarios.

The above arguments are demonstrated as follows. Rewrite (1) as

$$y_{i,t} = \mathbf{z}'_{i,t}\boldsymbol{\phi}_i + \mathbf{x}'_{i,t}\boldsymbol{\beta}_i + \varepsilon_{i,t} \tag{2}$$

where $\mathbf{z}_{i,t} = (1, y_{i,t-1}, \ldots, y_{i,t-p})'$, $\mathbf{x}_{i,t} = (x_{i,t-1}, \ldots, x_{i,t-p})'$, $\boldsymbol{\phi}_i = (\phi_{0,i}, \ldots, \phi_{p,i})'$, and $\boldsymbol{\beta}_i = (\beta_{1,i}, \ldots, \beta_{p,i})'$. Stacking (2) over time yields

$$\mathbf{y}_i = \mathbf{Z}_i\boldsymbol{\phi}_i + \mathbf{X}_i\boldsymbol{\beta}_i + \boldsymbol{\varepsilon}_i$$

where $\mathbf{y}_i = (y_{i,1}, \ldots, y_{i,T})'$, $\mathbf{Z}_i = (\mathbf{z}_{i,1}, \ldots, \mathbf{z}_{i,T})'$, $\mathbf{X}_i = (\mathbf{x}_{i,1}, \ldots, \mathbf{x}_{i,T})'$, and $\boldsymbol{\varepsilon}_i = (\varepsilon_{i,1}, \ldots, \varepsilon_{i,T})'$. Under the null hypothesis, $\boldsymbol{\beta}_i = \boldsymbol{\beta} = \mathbf{0}$. The pooled least-squares estimator of $\boldsymbol{\beta}$ is defined as follows,

$$\widehat{\boldsymbol{\beta}} = \left(\sum_{i=1}^{N} \mathbf{X}'_i \mathbf{M}_{\mathbf{Z}_i} \mathbf{X}_i \right)^{-1} \left(\sum_{i=1}^{N} \mathbf{X}'_i \mathbf{M}_{\mathbf{Z}_i} \mathbf{y}_i \right) \tag{3}$$

where $\mathbf{M}_{\mathbf{Z}_i} = \mathbf{I}_T - \mathbf{Z}_i \left(\mathbf{Z}'_i \mathbf{Z}_i \right)^{-1} \mathbf{Z}'_i$. Fernández-Val and Lee (2013) show that under general conditions, and as $N, T \to \infty$ with $N/T \to \kappa^2 \in [0; \infty)$, we have

$$\sqrt{NT} \left(\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}_0 \right) \to \mathbf{J}^{-1} N \left(-\kappa \mathbf{B}, \mathbf{V} \right)$$

where $\mathbf{J} = \text{plim}_{N,T \to \infty} (NT)^{-1} \sum_{i=1}^{N} \mathbf{X}'_i \mathbf{M}_{\mathbf{Z}_i} \mathbf{X}_i$, $\mathbf{V}$ denotes the variance–covariance matrix, and $\mathbf{B}$ is the bias arising from the fact that N and T are of the same order.

To remove the bias of the pooled estimator, we use the HPJ estimator of Dhaene and Jochmans (2015), which is defined as follows:

$$\widetilde{\boldsymbol{\beta}} = \widehat{\boldsymbol{\beta}} + \left\{ \widehat{\boldsymbol{\beta}} - \frac{1}{2} \left(\widehat{\boldsymbol{\beta}}_{1/2} + \widehat{\boldsymbol{\beta}}_{2/1} \right) \right\} = \widehat{\boldsymbol{\beta}} + T^{-1}\widehat{\mathbf{B}}$$

---

4. Obviously, nontrivial power of the test requires that there exist sufficiently many individuals with nonzero coefficients.

where $\widehat{\boldsymbol{\beta}}_{1/2}$ is the estimator in (3) calculated using all units but only the first half of the time-series observations and $\widehat{\boldsymbol{\beta}}_{2/1}$ is the estimator in (3) calculated using all units but only the second half of the available time-series observations.

The bias-corrected estimator then forms the basis of a Wald test for Granger noncausality. In particular, under mild regularity assumptions reported in Juodis, Karavias, and Sarafidis (2021), as $N, T \to \infty$ with $N/T \to \kappa^2 \in [0, \infty)$, we have

$$\widehat{W}_{\mathrm{HPJ}} = NT\widetilde{\boldsymbol{\beta}}' \left(\widehat{\mathbf{J}}^{-1}\widehat{\mathbf{V}}\widehat{\mathbf{J}}^{-1} \right)^{-1} \widetilde{\boldsymbol{\beta}} \to \chi^2(P)$$

where $\widehat{\mathbf{J}} = (NT)^{-1} \sum_{i=1}^{N} \mathbf{X}_i'\mathbf{M}_{\mathbf{Z}_i}\mathbf{X}_i$.

When the errors are assumed to be homoskedastic along both time and cross-sectional dimensions, then

$$\widehat{\mathbf{V}} = \widehat{\sigma}^2\widehat{\mathbf{J}}$$

with the variance estimator given by

$$\widehat{\sigma}^2 = \frac{1}{N(T-1-P)-P} \sum_{i=1}^{N} \left(\mathbf{y}_i - \mathbf{X}_i\widehat{\boldsymbol{\beta}} \right)' \mathbf{M}_{\mathbf{Z}_i} \left(\mathbf{y}_i - \mathbf{X}_i\widehat{\boldsymbol{\beta}} \right) \tag{4}$$

On the other hand, if the errors are cross-sectionally heteroskedastic,

$$\widehat{\mathbf{V}} = \frac{1}{N(T-1-P)-P} \sum_{i=1}^{N} \mathbf{X}_i'\mathbf{M}_{\mathbf{Z}_i}\widehat{\boldsymbol{\varepsilon}}_i\widehat{\boldsymbol{\varepsilon}}_i'\mathbf{M}_{\mathbf{Z}_i}\mathbf{X}_i \tag{5}$$

The model in (1) can allow for weak cross-section dependence as in Sarafidis and Wansbeek (2012) and Dumitrescu and Hurlin (2012). Under weak cross-sectional dependence, the HPJ estimator $\widetilde{\boldsymbol{\beta}}$ remains consistent, but $\widehat{\mathbf{V}}$ in the above equations is not. In this case, an estimator for $\widehat{\mathbf{V}}$ is obtained by using the pairs bootstrap as in Gonçalves and Kaffo (2015). Unreported Monte Carlo simulations show that this approach works well in finite samples.

3   The xtgrangert command

3.1   Syntax

xtgrangert *depvar* $\big[$ *indepvars* $\big]$ $\big[$ *if* $\big]$ $\big[$ *in* $\big]$ $\big[$, lags(#) maxlags(#) het nodfc
    <u>boot</u>strap$\big[$(#*reps*, seed(*seed*))$\big]$ sum $\big]$

Data must be xtset before using xtgrangert. The panel must be balanced.

3.2 Options

lags(*#*) specifies the number of lags of dependent and independent variables to be
added to the regression. The default is lags(1). The lags of the dependent variable
are partialed out.

maxlags(*#*) specifies the upper bound of lags. The BIC is used to select the number of
lags that provides the best model fit. lags() and maxlags() cannot be used at the
same time.

het allows for cross-sectional heteroskedasticity.

nodfc does not apply the degrees-of-freedom correction in (4) and (5). This option is
mostly useful under cross-sectional heteroskedasticity.

bootstrap $\left[\,(\#reps,\ \text{seed}(seed))\,\right]$ specifies a bootstrap variance estimator in the HPJ
Wald statistic that allows for cross-sectional dependence and uses a custom seed and
#reps replications. By default, 100 replications are used based on the current seed.
This is useful in the presence of weak cross-sectional dependence.

sum presents results on the sum of the estimated feedback coefficients. This option can
be useful when $P > 1$.

3.3 Stored results

xtgrangert stores the following in e():

Scalars
    e(N)   number of units
    e(T)   number of time periods
    e(p)   number of lags
    e(BIC)   BIC value
    e(W_HPJ)   Wald statistic
    e(pvalue)   *p*-value for the HPJ Wald test

Matrices
    e(b_HPJ)   HPJ coefficient estimator
    e(Var_HPJ)   variance–covariance matrix of the HPJ estimator
    e(b_Sum_HPJ)   sum of the HPJ estimates of the feedback coefficients
    e(Var_Sum_HPJ) variance of the sum of the HPJ estimators

3.4 Postestimation command

predict can be used after xtgrangert. The syntax for predict is

predict *newvar* $\left[\,if\,\right]$ $\left[\,in\,\right]$ $\left[\,,\ \underline{\text{res}}\text{iduals xb}\,\right]$

residuals calculates the residuals.

xb calculates the linear prediction on the partialed-out variables.

4   Example

4.1   Estimation of the determinants of banks' capital adequacy ratios

To illustrate the `xtgrangert` command, we perform Granger noncausality tests and examine the type of temporal relation between profitability, cost efficiency, and asset quality in the U.S. banking industry. We draw a random sample of 450 U.S. bank-holding companies, each one observed over 56 time periods, namely, 2006:Q1–2019:Q4. The data are publicly available, and they have been downloaded from the Federal Deposit Insurance Corporation website.[5]

We focus on the following model,

$$\mathrm{ROA}_{i,t} = \phi_{0,i} + \sum_{p=1}^{P} \phi_{p,i} \mathrm{ROA}_{i,t-p} + \sum_{p=1}^{P} \beta_{1,p,i} \texttt{inefficiency}_{i,t-p}$$

$$+ \sum_{p=1}^{P} \beta_{2,p,i} \texttt{quality}_{i,t-p} + \varepsilon_{i,t}$$

for $i = 1, \ldots, N(= 450)$ and $t = P + 1, \ldots, T(= 56)$.

$\mathrm{ROA}_{i,t}$ stands for the "return on assets" and is used as a measure of profitability; in particular, it is defined as annualized net income expressed as a percentage of average total assets. $\texttt{inefficiency}_{i,t-p}$ presents a measure of cost inefficiency, which has been constructed from a stochastic cost frontier model using a translog function form.[6] Finally, $\texttt{quality}_{i,t-p}$ represents the quality of banks' assets and is computed as the total amount of loan-loss provisions expressed as a percentage of assets. Thus, a higher level of loan-loss provisions indicates lower quality.

We start by testing whether the pair of `inefficiency` and `quality` Granger-causes ROA. Then, we consider univariate tests by modeling ROA as a function of `inefficiency` and `quality` separately. Throughout, we allow for a maximum of four lags of the dependent variable and the covariates. The following results are obtained:

---

5. See https://www.fdic.gov/.
6. See section 5 in Juodis, Karavias, and Sarafidis (2021) for more details.

```
. use xtgrangert_example

. xtset cert time
Panel variable: cert (strongly balanced)
 Time variable: time, 1 to 56
         Delta: 1 unit

. xtgrangert roa inefficiency quality, maxlags(4) het
Juodis, Karavias and Sarafidis (2021) Granger non-causality Test
--------------------------------------------------------------------------
Number of units= 450                    Obs. per unit (T) = 55
Number of lags = 1                      BIC          = -34257.34
--------------------------------------------------------------------------
JKS non-causality test

H0: Selected covariates do not Granger-cause roa.
H1: H0 is violated.

HPJ Wald test  : 30.2387
p-value        : 0.0000
--------------------------------------------------------------------------
BIC selection:
    lags = 1, BIC = -34257.336*
    lags = 2, BIC = -33371.195
    lags = 3, BIC = -32727.595
    lags = 4, BIC = -32715.923
--------------------------------------------------------------------------
              Results for the Half-Panel Jackknife estimator
          Cross-sectional heteroskedasticity-robust variance estimation
```

	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
inefficiency						
L1.	.2562039	.0572807	4.47	0.000	.1439358	.368472
quality						
L1.	-.0162294	.0444754	-0.36	0.715	-.1033996	.0709409

As we can see, the null hypothesis that cost inefficiency and asset quality do not Granger-cause profitability is rejected at the 5% level of significance. The optimal number of lags equals 1 according to the BIC.[7] The option `het` requests computing cross-sectional heteroskedasticity-robust standard errors.

In addition to the Wald test statistic, the command also reports regression results with respect to the HPJ bias-corrected pooled estimator. The regression output above indicates that the test outcome may be driven by `inefficiency`. To shed some light on this issue, we test for Granger noncausality for each variable separately using univariate tests. We obtain the following output:

---

7. Assuming that the maximum number of lags is 4, we tested the residuals for remaining serial correlation of order up to 3 using the community-contributed command `xtqptest` by Wursten (2018). We did not find evidence of residual serial correlation (p-value = 0.089). The commands for getting these results are `predict epsilonres, residuals` and `xtqptest epsilonres, lags(3) force`.

```
. xtgrangert roa inefficiency, maxlags(4) het

Juodis, Karavias and Sarafidis (2021) Granger non-causality Test
--------------------------------------------------------------------------------
Number of units= 450                        Obs. per unit (T) = 55
Number of lags = 1                          BIC           = -33295.8
--------------------------------------------------------------------------------
JKS non-causality test

H0: inefficiency does not Granger-cause roa.
H1: inefficiency does Granger-cause roa for at least one panelvar.

HPJ Wald test  : 24.3174
p-value        : 0.0000
--------------------------------------------------------------------------------
BIC selection:
    lags = 1, BIC = -33295.799*
    lags = 2, BIC = -32170.227
    lags = 3, BIC = -31112.604
    lags = 4, BIC = -30724.676
--------------------------------------------------------------------------------
              Results for the Half-Panel Jackknife estimator
        Cross-sectional heteroskedasticity-robust variance estimation
```

	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
inefficiency						
L1.	.2549723	.0517052	4.93	0.000	.1536319	.3563127

```
. xtgrangert roa quality, maxlags(4) het

Juodis, Karavias and Sarafidis (2021) Granger non-causality Test
--------------------------------------------------------------------------------
Number of units= 450                        Obs. per unit (T) = 55
Number of lags = 1                          BIC           = -33816.06
--------------------------------------------------------------------------------
JKS non-causality test

H0: quality does not Granger-cause roa.
H1: quality does Granger-cause roa for at least one panelvar.

HPJ Wald test  : 0.2090
p-value        : 0.6476
--------------------------------------------------------------------------------
BIC selection:
    lags = 1, BIC = -33816.061*
    lags = 2, BIC = -32649.24
    lags = 3, BIC = -31479.433
    lags = 4, BIC = -30607.217
--------------------------------------------------------------------------------
              Results for the Half-Panel Jackknife estimator
        Cross-sectional heteroskedasticity-robust variance estimation
```

	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
quality						
L1.	-.0201426	.0440637	-0.46	0.648	-.1065059	.0662207

The output on the top (bottom) corresponds to the Granger noncausality univariate test of the relationship between profitability and cost inefficiency (asset quality). The null hypothesis that `inefficiency` does not Granger-cause ROA is rejected at the 5%

level of significance. This implies that past values of `inefficiency` contain information that helps to predict ROA over and above the information contained in past values of ROA. On the other hand, one fails to reject the null hypothesis that `quality` does not Granger-cause ROA.

To illustrate further options of the `xtgrangert` command, we split the sample into two groups according to their asset size, where the partitioning is determined based on the kmeans clustering algorithm available in Stata. Subsequently, we test for Granger noncausality for the smallest banks in the sample, using data from 2011:Q1 onward, which corresponds to quarter 1 of the first year following the enactment of the Dodd–Frank Wall Street Reform and Consumer Protection Act of 2010.[8] We obtain the following results:

```
. xtgrangert roa inefficiency quality, maxlags(4) het sum, if cluster==2 &
> time>20
Juodis, Karavias and Sarafidis (2021) Granger non-causality Test
-------------------------------------------------------------------------------
Number of units= 183                      Obs. per unit (T) = 34
Number of lags = 2                        BIC          = -10307.93
-------------------------------------------------------------------------------
JKS non-causality test

H0: Selected covariates do not Granger-cause roa.
H1: H0 is violated.

HPJ Wald test  : 36.3572
p-value        : 0.0000
-------------------------------------------------------------------------------
BIC selection:
    lags = 1, BIC = -10249.44
    lags = 2, BIC = -10307.934*
    lags = 3, BIC = -9788.8299
    lags = 4, BIC = -9963.9685
-------------------------------------------------------------------------------
        Sum of Half-Panel Jackknife coefficients across lags (lags>1)
        Cross-sectional heteroskedasticity-robust variance estimation
```

	Coefficient	Std. err.	z	P>|z|	[95% conf. interval]	
inefficiency	.4906756	.2405474	2.04	0.041	.0192113	.9621398
quality	−.1765458	.1235961	−1.43	0.153	−.4187897	.0656981

As before, the null hypothesis that `inefficiency` and `quality` do not Granger-cause ROA is rejected at the 5% level of significance. Note that the optimal number of lags equals 2. The option `sum` requests reporting the sum of the lags of the regression coefficients for each variable.

---

8. The Dodd–Frank Act is a U.S. federal law enacted during 2010, aiming "to promote the financial stability of the United States by improving accountability and transparency in the financial system, to end 'too big to fail', to protect the American taxpayer by ending bailouts, to protect consumers from abusive financial services practices, and for other purposes"; see https://www.cftc.gov/LawRegulation/DoddFrankAct/index.htm. In a nutshell, the Dodd–Frank Act has instituted a new failure-resolution regime, which seeks to ensure that losses resulting from bad decisions by managers are absorbed by equity and debt holders, thus potentially reducing moral hazard.

5   Conclusions

`xtgrangert` implements the Granger noncausality test of Juodis, Karavias, and Sarafidis (2021). The command reports the Wald test statistic and its *p*-value, the null and the alternative hypotheses, and regression results with respect to the HPJ bias-corrected pooled estimator. The command offers options for both manual and automatic lag-length selection, using a BIC. Moreover, the command allows for cross-section dependence and cross-section heteroskedasticity in the errors.

The Granger causality testing approach developed by Juodis, Karavias, and Sarafidis (2021) is computationally fast and widely applicable. In the presence of cross-section dependence, some caution must be exercised with interpreting the results. This is because we do not provide a formal proof that the bootstrap methodology used in this article is asymptotically valid in this case. The same issue applies to other notable contributions in the literature; see, for example, Dumitrescu and Hurlin (2012). While the Monte Carlo simulations are encouraging, a formal investigation on the properties of estimators and the bootstrap is an interesting direction for future research.

A special case of strong cross-section dependence is the additive time effect of the two-way fixed effects model, which is frequently used in the literature. Typically, this time effect is removed by transforming the data to deviations from their cross-sectional means for each time period. This approach is valid if the panel is assumed to be homogeneous across i. After the transformation, the Juodis, Karavias, and Sarafidis (2021) test can be applied without using the bootstrap. If, however, the panel is heterogeneous, then the demeaning no longer removes the additive time effect.

6   Acknowledgments

7   Programs and supplemental materials

To install a snapshot of the corresponding software files as they existed at the time of publication of this article, type

```
. net sj 23-1
. net install st0706      (to install program files, if available)
. net get st0706          (to install ancillary files, if available)
```

The `xtgrangert` command also is available on the Statistical Software Components Archive and can be installed directly in Stata with the command

```
. ssc install xtgrangert
```

8 References

Abrigo, M. R. M., and I. Love. 2016. Estimation of panel vector autoregression in Stata. *Stata Journal* 16: 778–804. https://doi.org/10.1177/1536867X1601600314.

Bun, M. J. G., and V. Sarafidis. 2015. Dynamic panel data models. In *The Oxford Handbook of Panel Data*, ed. B. H. Baltagi, 76–110. Oxford: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199940042.013.0003.

Dhaene, G., and K. Jochmans. 2015. Split-panel jackknife estimation of fixed-effect models. *Review of Economic Studies* 82: 991–1030. https://doi.org/10.1093/restud/rdv007.

Dufour, J.-M., and E. Renault. 1998. Short run and long run causality in time series: Theory. *Econometrica* 66: 1099–1125. https://doi.org/10.2307/2999631.

Dumitrescu, E.-I., and C. Hurlin. 2012. Testing for Granger non-causality in heterogeneous panels. *Economic Modelling* 29: 1450–1460. https://doi.org/10.1016/j.econmod.2012.02.014.

Emirmahmutoglu, F., and N. Kose. 2011. Testing for Granger causality in heterogeneous mixed panels. *Economic Modelling* 28: 870–876. https://doi.org/10.1016/j.econmod.2010.10.018.

Fernández-Val, I., and J. Lee. 2013. Panel data models with nonadditive unobserved heterogeneity: Estimation and inference. *Quantitative Economics* 4: 453–481. https://doi.org/10.3982/QE75.

Gonçalves, S., and M. Kaffo. 2015. Bootstrap inference for linear dynamic panel data models with individual fixed effects. *Journal of Econometrics* 186: 407–426. https://doi.org/10.1016/j.jeconom.2015.02.017.

Granger, C. W. J. 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* 37: 424–438. https://doi.org/10.2307/1912791.

Holtz-Eakin, D., W. Newey, and H. S. Rosen. 1988. Estimating vector autoregressions with panel data. *Econometrica* 56: 1371–1395. https://doi.org/10.2307/1913103.

Juodis, A., Y. Karavias, and V. Sarafidis. 2021. A homogeneous approach to testing for Granger non-causality in heterogeneous panels. *Empirical Economics* 60: 93–112. https://doi.org/10.1007/s00181-020-01970-9.

Juodis, A., and V. Sarafidis. 2022. An incidental parameters free inference approach for panels with common shocks. *Journal of Econometrics* 229: 19–54. https://doi.org/10.1016/j.jeconom.2021.03.011.

Karavias, Y., and E. Tzavalis. 2016. Local power of fixed-T panel unit root tests with serially correlated errors and incidental trends. *Journal of Time Series Analysis* 37: 222–239. https://doi.org/10.1111/jtsa.12144.

———. 2017. Local power of panel unit root tests allowing for structural breaks. *Econometric Reviews* 36: 1123–1156. https://doi.org/10.1080/07474938.2015.1059722.

Lopez, L., and S. Weber. 2017. Testing for Granger causality in panel data. *Stata Journal* 17: 972–984. https://doi.org/10.1177/1536867X1801700412.

Sarafidis, V., and T. Wansbeek. 2012. Cross-sectional dependence in panel data analysis. *Econometric Reviews* 31: 483–531. https://doi.org/10.1080/07474938.2011.611458.

Wursten, J. 2018. Testing for serial correlation in fixed-effects panel models. *Stata Journal* 18: 76–100. https://doi.org/10.1177/1536867X1801800106.

About the authors

Jiaqi Xiao was a full time student at the B.Sc. in Mathematical Economics and Statistics at the Birmingham Business School in the University of Birmingham.

Artūras Juodis is an assistant professor of economics at the Amsterdam School of Economics in the University of Amsterdam.

Yiannis Karavias is an associate professor in financial economics at the Birmingham Business School in the University of Birmingham.

Vasilis Sarafidis is a professor of economics at the Department of Economics in the BI Norwegian Business School.

Jan Ditzen is an assistant professor in econometrics at the Free University of Bozen-Bolzano.