

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

A portmanteau test for serial correlation in a linear panel model

Jochmans, Koen; Verardi, Vincenzo

Published in:
Stata Journal

DOI:
[10.1177/1536867x20909695](https://doi.org/10.1177/1536867x20909695)

Publication date:
2020

Document Version
Publisher's PDF, also known as Version of record

[Link to publication](#)

Citation for published version (HARVARD):

Jochmans, K & Verardi, V 2020, 'A portmanteau test for serial correlation in a linear panel model', *Stata Journal*, vol. 20, no. 1, pp. 149-161. <https://doi.org/10.1177/1536867x20909695>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A portmanteau test for serial correlation in a linear panel model

Koen Jochmans
University of Cambridge
Cambridge, UK
kj345@cam.ac.uk

Vincenzo Verardi
Université de Namur
Namur, Belgium
vverardi@unamur.be

Abstract. We introduce the command `xtserialpm` to perform the portmanteau test developed in Jochmans (2019, Cambridge Working Papers in Economics No. 1993, University of Cambridge, Faculty of Economics). The procedure tests for serial correlation of arbitrary form in the errors of a linear panel model after estimation of the regression coefficients by the within-group estimator. The test is designed for short panels and can deal with general missing-data patterns. The test is different from the related portmanteau test of Inoue and Solon (2006, *Econometric Theory* 22: 835–851), which is performed by `xtistest` (Wursten, 2018, *Stata Journal* 18: 76–100), in that it allows for heteroskedasticity. In simulations documented below, `xtserialpm` is found to provide a more powerful test than `xthrttest` (Wursten 2018), which performs the test for first-order autocorrelation of Born and Breitung (2016, *Econometric Reviews* 35: 1290–1316). We also provide comparisons with `xtistest` and `xtserial` (Drukker, 2003, *Stata Journal* 3: 168–177). These tests perform well under stationarity but break down under even mild forms of heteroskedasticity.

Keywords: `st0592`, `xtserialpm`, heteroskedasticity, fixed-effects model, portmanteau test, serial correlation, short panel data, unbalanced panel

1 Introduction

Consider panel data on an outcome $y_{i,t}$ and a set of covariates $\mathbf{x}_{i,t}$, where $i = 1, \dots, N$ and $t = 1, \dots, T$. The data are independent across groups i but may be dependent within groups. The workhorse specification to analyze such data is the regression model

$$y_{i,t} = \mathbf{x}_{i,t}^{\top} \boldsymbol{\beta} + v_{i,t}, \quad v_{i,t} = \alpha_i + \varepsilon_{i,t}$$

where α_i is a latent individual effect and $\varepsilon_{i,t}$ is an idiosyncratic disturbance whose mean is normalized to zero. These disturbances are taken to be mean independent of the regressors and the individual effects but are otherwise allowed to be (conditionally) heteroskedastic and correlated within each group i . Our aim is to test whether the $\varepsilon_{i,t}$ are correlated within groups. Although we do not make it explicit in the notation, everything to follow applies to settings where the panel is unbalanced (possibly with gaps) provided that missingness is at random. As such, the test discussed below is suitable for data with a group structure; the number of observations on a given group need not be the same, and the observations need not be collected over time. Examples are data on student test scores stratified by classroom or data on individual members of households.

The command `xtserialpm`, which we introduce in this article, performs a test for the (multivariate) null of no correlation at any order. The alternative is that at least one error pair is correlated. Thus, `xtserialpm` performs a portmanteau test. The test performed was developed in Jochmans (2019). This test is different from the portmanteau test of Inoue and Solon (2006), which is implemented in `xtistest` (Wursten 2018), and is robust to heteroskedasticity. This is important because requiring the errors to have a constant variance within each group is often unrealistic. One situation where heteroskedasticity will arise is when the error process is not in its steady state; this is typical in short panels. A second situation is where errors are conditionally heteroskedastic and some of the regressors are nonstationary. An example here would be a wage regression where the regressors include such characteristics as age, tenure and experience, and number of children, all of which are nonstationary.

The portmanteau paradigm is to be contrasted with an approach that tests against a specific alternative. Using a portmanteau test is of interest if no strong stand can be taken on the particular form of correlation that should serve as the alternative. This is relevant in many panel-data applications, especially when the observations for a given group do not have a natural ordering (such as time, for example). On the other hand, there are cases where attention may be limited to first-order autocorrelation patterns. In such a case, `xtserial` (Wooldridge 2010, 319–320; Drukker 2003) or its heteroskedasticity-robust version `xthrttest` (Born and Breitung 2016; Wursten 2018) may be of use.¹ While tests against specific alternatives may have poor power if the alternative is ill chosen, they have the advantage that the dimension of the null hypothesis is independent of the sample size. A portmanteau test, on the other hand, necessarily has a null whose dimension grows with the length of the panel, T . Moreover, `xtserialpm` and `xtistest` are not well suited for panels where $T(T-1)/2$ is not small relative to N .

We introduce the test that is the subject of this article in section 2. We give the syntax of the command `xtserialpm`, which implements the test, in section 3, and we provide an example in section 4. We give the results of a simulation study in section 5. The Monte Carlo analysis compares the performance of `xtserialpm` with `xtistest`, `xtserial`, and `xthrttest` in various settings. While `xtistest` and `xtserial` are competitive under homoskedasticity, they are unreliable under heteroskedasticity. Although `xthrttest` is designed to be size correct, it is found to have poor power. Moreover, it is virtually unable to detect most violations from the null, even those for which it was designed.

2 The test statistic

The presence of the group-level effect α_i complicates the construction of a test based on the data in levels. The approach followed in Jochmans (2019) is to test the null that the difference between all pairwise within-group correlations are zero. There are $T(T-1)/2$ covariances, so

1. Under homoskedasticity, `xtqptest` (Born and Breitung 2016; Wursten 2018) can also be used to test for correlation up to a fixed order.

$$q := \frac{T(T-1)}{2} - 1 = \frac{(T+1)(T-2)}{2}$$

linearly independent differences. There are many ways of selecting q such differences. How they are chosen is irrelevant in practice because each will deliver numerically the same test statistic. A convenient way for notational purposes is as follows. Let Δ denote the first-differencing operator; that is, $\Delta v_{i,t} = v_{i,t} - v_{i,t-1}$. Then, we test the null $\mathbf{H}_0 : \mathbb{E}(v_{i,t'} \Delta v_{i,t}) = 0$ for all t and each $t' \leq t-2$ and $t' = t+1$ against the alternative $\mathbf{H}_1 : \mathbb{E}(v_{i,t'} \Delta v_{i,t}) \neq 0$ for some t and $t' \leq t-2$ or $t' = t+1$.

An exercise in adding up shows that this indeed involves q moments. The rationale for them comes from the observation that

$$\begin{aligned} \mathbb{E}(v_{i,t'} \Delta v_{i,t}) &= \mathbb{E}(\varepsilon_{i,t'} \Delta \varepsilon_{i,t}) + \mathbb{E}(\alpha_i \Delta \varepsilon_{i,t}) \\ &= \mathbb{E}(\varepsilon_{i,t'} \Delta \varepsilon_{i,t}) \\ &= \mathbb{E}(\varepsilon_{i,t'} \varepsilon_{i,t}) - \mathbb{E}(\varepsilon_{i,t'} \varepsilon_{i,t-1}) \end{aligned}$$

which is indeed the difference between two covariances. The main transition here uses $\mathbb{E}(\alpha_i \Delta \varepsilon_{i,t}) = \mathbb{E}\{\alpha_i \mathbb{E}(\Delta \varepsilon_{i,t} | \alpha_i)\} = 0$, which follows from iterated expectations and the assumption that $\mathbb{E}(\varepsilon_{i,t} | \alpha_i) = 0$.

The q restrictions that make up our null can be written compactly as

$$\mathbb{E}(\mathbf{\Upsilon}_i^\top \mathbf{D} \mathbf{v}_i) = \mathbf{0}$$

where we have introduced the $(T-1) \times q$ matrix

$$\mathbf{\Upsilon}_i := \left(\begin{array}{cccccc|ccc} 0 & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & v_{i,3} & 0 & \cdots & 0 \\ v_{i,1} & 0 & 0 & 0 & & 0 & & 0 & 0 & v_{i,4} & & \vdots \\ 0 & v_{i,1} & v_{i,2} & 0 & & 0 & & 0 & \vdots & & \ddots & 0 \\ \vdots & & & \ddots & & & & \vdots & 0 & 0 & & v_{i,T} \\ 0 & 0 & 0 & 0 & \cdots & v_{i,1} & \cdots & v_{i,T-2} & 0 & 0 & \cdots & 0 \end{array} \right)$$

and the $T \times 1$ vector $\mathbf{v}_i := (v_{i,1}, \dots, v_{i,T})^\top$ and write \mathbf{D} for the $(T-1) \times T$ matrix first-difference operator; so $\mathbf{D} \mathbf{v}_i = (\Delta v_{i,2}, \dots, \Delta v_{i,T})^\top$, for example. The left block of the matrix $\mathbf{\Upsilon}_i$ is reminiscent of the instrument matrix for the generalized method of moments estimator of dynamic panel models (see, for example, Arellano [2003, 88–89]). The right block does not appear there, because it would not provide valid moment conditions in that context. The null can be tested using a minimum-distance statistic in a sample version of the moment restrictions as soon as three observations per group are available. Note that the dimension of the null grows with T . As such, the approach is designed for short panels, where q is small compared with N .

To make the test operational, we need to replace the unobserved $v_{i,t}$ with an estimator. For this, an estimator of β is needed. `xtserialpm` uses the within-group least-squares estimator (as computed by `xtreg`, `fe`),

$$\mathbf{b} := \left(\sum_{i=1}^N \mathbf{X}_i^\top \mathbf{M} \mathbf{X}_i \right)^{-1} \sum_{i=1}^N \mathbf{X}_i^\top \mathbf{M} \mathbf{y}_i$$

where we have collected all observations for a given group in $\mathbf{y}_i := (y_{i,1}, \dots, y_{i,T})^\top$ and $\mathbf{X}_i := (\mathbf{x}_{i,1}, \dots, \mathbf{x}_{i,T})^\top$ and \mathbf{M} denotes the usual $T \times T$ projection matrix that transforms observations into deviations from within-group means. Given \mathbf{b} , the residuals

$$u_{i,t} := y_{i,t} - \mathbf{x}_{i,t}^\top \mathbf{b}$$

can be used as estimators of the $v_{i,t}$.

We then define

$$\mathbf{s}_i := \mathbf{U}_i^\top \mathbf{D} \mathbf{u}_i - \left(\sum_{j=1}^N \mathbf{U}_j^\top \mathbf{D} \mathbf{X}_j \right) \left(\sum_{j=1}^N \mathbf{X}_j^\top \mathbf{M} \mathbf{X}_j \right)^{-1} \mathbf{X}_i^\top \mathbf{M} \mathbf{u}_i$$

where the matrix \mathbf{U}_i is the sample version of Υ_i constructed using the residuals $u_{i,t}$ in place of the unobservable $v_{i,t}$; that is,

$$\mathbf{U}_i := \left(\begin{array}{cccccccc|ccc} 0 & 0 & 0 & 0 & \cdots & 0 & \cdots & 0 & u_{i,3} & 0 & \cdots & 0 \\ u_{i,1} & 0 & 0 & 0 & & 0 & & 0 & 0 & u_{i,4} & & \vdots \\ 0 & u_{i,1} & u_{i,2} & 0 & & 0 & & 0 & \vdots & & \ddots & 0 \\ \vdots & & & \ddots & & & & \vdots & 0 & 0 & & u_{i,T} \\ 0 & 0 & 0 & 0 & \cdots & u_{i,1} & \cdots & u_{i,T-2} & 0 & 0 & \cdots & 0 \end{array} \right)$$

and we have introduced $\mathbf{u}_i := (u_{i,1}, \dots, u_{i,T})^\top$. The test statistic for our null can then be written as the quadratic form

$$\mathbf{s}^\top \mathbf{V}^{-1} \mathbf{s}$$

where $\mathbf{s} := \sum_{i=1}^N \mathbf{s}_i$ and \mathbf{V} is an estimator of the variance of \mathbf{s} . `xtserialpm` uses the uncentered estimator

$$\mathbf{V} := \sum_{i=1}^N \mathbf{s}_i \mathbf{s}_i^\top$$

as the default. Use of a centered variance estimator is available as an option. In the simulations reported on below, we found that use of the centered estimator is power enhancing but comes at the expense of size distortion in small samples.

Our test statistic has an interpretation that explains its form. Because the second part of \mathbf{s}_i sums to zero, we have

$$\mathbf{s} = \sum_{i=1}^N \mathbf{s}_i = \sum_{i=1}^N \mathbf{U}_i^\top \mathbf{D} \mathbf{u}_i$$

This is a sample version of the moments we are aiming to test. The second part of \mathbf{s}_i is present to ensure that \mathbf{V} is a consistent estimator of the variance of \mathbf{s} . Moreover, the naive variance estimator $\sum_{i=1}^N (\mathbf{U}_i^\top \mathbf{D} \mathbf{u}_i)(\mathbf{U}_i^\top \mathbf{D} \mathbf{u}_i)^\top$ ignores the fact that the statistic is constructed with residuals rather than (unobservable) errors and will generally not be consistent.

Under the null,

$$\mathbf{s}^\top \mathbf{V}^{-1} \mathbf{s} \xrightarrow{d} \chi_q^2$$

as $N \rightarrow \infty$. A test of our null then amounts to comparing the test statistic with the quantiles of the χ_q^2 distribution. Large values are evidence against the null of no serial correlation. This test is consistent against any alternative except the one where all covariances are constant.² Asymptotic power results and calculations for special cases are provided in Jochmans (2019).

3 The `xtserialpm` command

`xtserialpm` is a standalone command that one can run without first running `xtreg`. You must `xtset` your data prior to executing `xtserialpm`. Unbalanced panel data are allowed.

The command has the following syntax:

```
xtserialpm depvar [indepvars] [if] [in] [, center noisily]
```

`center` returns the test statistic computed with a centered covariance matrix as discussed above.

`noisily` displays the preliminary within-group estimator. The output is the same as that produced by `xtreg, fe`.

Running the command produces a table with the value of the test statistic and the associated p -value. The layout of the table mimics the layout of the table produced by `xtserial`.

`xtserialpm` stores the following in `r()`:

Scalars

<code>r(stat)</code>	test statistic
<code>r(df)</code>	degrees of freedom of its limit distribution
<code>r(p)</code>	p -value of the test

Help is available by typing `help xtserialpm`.

2. This is due to the presence of the fixed effects, and the same is true for all other available tests of serial correlation in a panel setting.

4 Example

We use the data from the illustration in Drukker (2003). The following extract loads the data:

```
. webuse nlswork
(National Longitudinal Survey. Young Women 14-26 years of age in 1968)
. xtset idcode year
    panel variable:  idcode (unbalanced)
    time variable:  year, 68 to 88, but with gaps
                   delta: 1 unit
```

The portmanteau test is computed as

```
. xtserialpm ln_wage c.age##c.age ttl_exp c.tenure##c.tenure i.south if year<=70
Jochmans portmanteau test for within-group correlation in panel data.
H0: no within-group correlation
    Chi-sq( 2) =      25.658
    Prob > Chi-sq =      0.0000
```

The result provides strong evidence for the presence of serial correlation in the errors.

To compute the test statistic using a centered covariance matrix estimator, use the `center` option as follows:

```
. xtserialpm ln_wage c.age##c.age ttl_exp c.tenure##c.tenure i.south if
> year<=70, center
Jochmans portmanteau test for within-group correlation in panel data.
H0: no within-group correlation
    Chi-sq( 2) =      26.180
    Prob > Chi-sq =      0.0000
```

The test statistic is slightly larger, and our initial conclusion is unaltered.

To perform the test of Inoue and Solon (2006) in this example, we first generate residuals from the within-group regression by typing

```
. quietly xtreg ln_wage c.age##c.age ttl_exp c.tenure##c.tenure i.south if
> year<=70, fe
. predict u, residuals
(441 missing values generated)
```

We then perform the test on these residuals:

```
. keep if year<=70
(24,241 observations deleted)
. xtistest u, lags(all)
```

By default, `xtistest` checks for correlation (in the within-group residuals) up to the second order only.^{3,4} Here `xtistest` is used with the `lags()` option set to `all` so that the command yields the portmanteau test as originally introduced in Inoue and Solon (2006).

The output of the test is

```
Inoue and Solon (2006) LM-test on variables u
Panelvar: idcode
Timevar: year
```

Variable	IS-stat	p-value	N	maxT	balance?
u	159.44	0.000	2206	3	gaps

```
Notes: Under H0, LM ~ chi2((T-1)(T-2)/2)
H0: No auto-correlation of any order.
Ha: Auto-correlation of some order.
```

The same conclusion regarding our null is reached.

5 Simulations

We provide size and power comparisons between `xtserialpm`, `xtistest`, `xthrtest`, and `xtserial`. We consider different specifications for the errors and provide results for different panel dimensions. In all cases, outcomes were generated with fixed effects drawn from a standard normal and with two regressors—the first standard normal and the second zero or one according to the toss of a fair coin—each with a coefficient set to unity. From the time-series literature, we consider alternative specifications where the errors follow a first-order autoregressive or first-order moving-average process. Both `xtserial` and `xthrtest` were designed specifically to detect such forms of serial correlation. It is straightforward to concoct specifications of the error process where these tests will not be able to detect any deviation from the null.

-
- For any choice of lag, `xtistest` still delivers a portmanteau test, albeit based on fewer moment restrictions, and not a test for serial correlation up to a given order. Note, furthermore, that serial correlation in the within-group residuals need not be most pronounced at low orders even if the true errors (in levels) are most strongly correlated at small lags. A lower choice for `lag` need not improve the power of the test.
 - If interest lies only in testing that the first $T' < T$ covariances are the same, `xtserialpm` can be applied without modification to the subpanel obtained by dropping all cross sections $t > T'$. This will deliver a valid test.

Our first set of results concerns first-order autoregressive error processes of the form

$$\varepsilon_{i,t} = \rho \varepsilon_{i,t-1} + \eta_{i,t} \quad t = 2, \dots, T \quad |\rho| < 1$$

where the innovations $\eta_{i,t}$ are independent standard normal. Here the null corresponds to $\rho = 0$. In (A1), we draw the initial value $\varepsilon_{i,1}$ from its steady-state distribution. This implies that the error process is strictly stationary (and hence homoskedastic). In (A2), we set $\varepsilon_{i,1} = 0$ for all groups. This introduces time-series heteroskedasticity for any value of the autoregressive coefficient. Moreover, we have

$$\mathbb{E}(\varepsilon_{i,1}^2) = 0, \quad \mathbb{E}(\varepsilon_{i,2}^2) = 1, \quad \mathbb{E}(\varepsilon_{i,3}^2) = 1 + \rho^2, \quad \mathbb{E}(\varepsilon_{i,4}^2) = 1 + \rho^2 + \rho^4$$

and so on. The heteroskedasticity is mild but present both under the null and the alternative.

We present simulation results for panels with $N = 100$ and $T \in \{3, 6, 9\}$. This corresponds to $q \in \{2, 14, 35\}$, which, relative to N , can be considered small, moderate, and large. Results are reported in figure 1 by means of power plots (as obtained over 10,000 replications). For each test, the power curve plots the rejection frequency of the test against the value for $\rho \in (-1, 1)$ that was used to generate the data. A test is size correct if its rejection frequency under the null equals its size (here set to 5%, the level at which the horizontal axis is set). For a given alternative, the rejection frequency is the complement of the probability of making a type II error. Hence, given two tests that are both size correct, the one with a higher rejection frequency is superior. The plots in figure 1 contain the power curves for the portmanteau tests `xtserialpm` (full) and `xtistest` (dashed) as well as for the tests targeted to detect autocorrelation at first-order, `xthrttest` (dotted) and `xtserial` (dashed-dotted). Note that `xthrttest` requires $T \geq 4$ and so is absent from the plots for $T = 3$; $T \geq 3$ suffices for the other three tests.




Figure 1. Power against first-order autoregressive alternatives; power is reported for `xtserialpm` (solid line), `xtistest` (dashed line), `xthrttest` (dotted line), and `xtserial` (dashed-dotted line)

The left plots shows that, under homoskedasticity, all tests control size well. The command `xtserialpm` performs best here, which is not surprising given that this is the ideal setting for this test. Both the portmanteau tests do well and are roughly equally able to reject the null when it is false. `xtserialpm` does better in the shortest panel, while `xtistest` does better in the longest panel. Both observations arise from the fact that `xtserialpm` tests more moment conditions. `xthrttest` lacks power against most alternatives. Although its ability to detect violations from the null improves somewhat with the length of the panel, it is uniformly outperformed by all other tests even in the longest panel considered here.

The right plots show the impact of time-series heteroskedasticity on all the tests. Being robust to heteroskedasticity, both `xtserialpm` and `xthrttest` continue to be size correct. Moreover, the heteroskedasticity improves the power of `xtserialpm` relative to the stationary case, especially for $T = 3$. `xthrttest`, on the other hand, continues to struggle to detect any violation of the null. Both `xtserial` and `xtistest` are now severely size distorted, with their probability of a type I error far exceeding 5%. Because $|\rho| < 1$, the error process is mean reverting and so will become stationary as $t \rightarrow \infty$. Moreover, the errors become homoskedastic for large values of t . This explains why the performance of `xtserial` improves as T grows. Of course, no such improvement occurs for `xtistest`. On the other hand, we stress that the properties of `xtserial` and `xtistest` would not improve if instead N would increase.

Our second set of simulation results involves moving-average processes of order one; that is,

$$\varepsilon_{i,t} = \eta_{i,t} + \theta\eta_{i,t-1} \quad t = 1, \dots, T \quad \theta \in (-\infty, +\infty)$$

where the innovations $\eta_{i,t}$ are again independent standard normal. The null corresponds to $\theta = 0$. In (B1), we draw the initial value $\eta_{i,0}$ from the standard normal distribution, again implying stationarity. In (B2), we set $\eta_{i,0} = 0$. This leads to heteroskedasticity under the alternative but not under the null. This is different from (A2). Here, because errors are homoskedastic under the null, all tests will remain size correct. Note that heteroskedasticity is limited to the first observation, $\varepsilon_{i,1}$, whose variance is equal to 1; $\varepsilon_{i,2}, \dots, \varepsilon_{i,T}$ all have variance $1 + \theta^2$.

Figure 2 provides the power curves for these two specifications. We plot power against the first-order autoregressive coefficient (under stationarity), ρ . This coefficient is one to one with θ in the sense that

$$\theta = \frac{1 + \sqrt{1 - 4\rho^2}}{2\rho}$$

when $\rho \neq 0$ and $\theta = 0$ if $\rho = 0$. Note that $-(1/2) \leq \rho \leq (1/2)$. The main conclusions from the autoregressive specifications carry over. Both `xtserial` and `xtistest` do well under homoskedasticity but have erratic power patterns under heteroskedasticity. `xthrttest`, although size correct, continues to be incapable of detecting any violation from the null. `xtserialpm` performs well in all specifications and, as such, yields the most reliable test.




Figure 2. Power against first-order moving-average specifications; power is reported for `xtserialpm` (solid line), `xtistest` (dashed line), `xthrttest` (dotted line), and `xtserial` (dashed-dotted line)

6 Conclusion

We have introduced the command `xtserialpm` to test for arbitrary patterns of serial correlation in the errors of a fixed-effects regression model fit from short panel data. Contrary to the existing portmanteau test performed by `xtistest`, it is robust to heteroskedasticity. Both tests are designed for micropanels. For macropanels, where T is not small relative to N , only tests against specific alternatives can properly control size. Such tests are implemented in `xtserial` and `xthrttest`. Simulation evidence shows that even mild forms of heteroskedasticity make the properties of `xtistest` and `xtserial` break down. Unfortunately, heteroskedasticity in short panels is the rule rather than

the exception. Further, `xthrtest`, although size correct under heteroskedasticity, is far less powerful than `xtserialpm` even when the alternative under question is characterized by well-pronounced dependence at first order. The conclusion from the theory and simulation evidence presented here is that, when heteroskedasticity is suspected, only `xtserialpm` will provide a suitable test when $T(T-1)/2$ is small compared with N . On the other hand, only `xthrtest` is guaranteed to be size correct when $T(T-1)/2$ is large relative to N . However, it may not pick up violations from the null if they do not occur at first order. Furthermore, even if they do, our simulations show that the test needs the sample size to be substantial to safeguard against type II errors with reasonable probability.

7 Acknowledgments

Koen Jochmans gratefully acknowledges financial support from the European Research Council through Starting Grant n° 715787. Vincenzo Verardi gratefully acknowledges financial support from the FNRS. Comments and suggestions by Helmut Farbmacher and an anonymous referee were much appreciated.

8 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 20-1
. net install st0592      (to install program files, if available)
. net get st0592         (to install ancillary files, if available)
```

9 References

- Arellano, M. 2003. *Panel Data Econometrics*. Oxford: Oxford University Press.
- Born, B., and J. Breitung. 2016. Testing for serial correlation in fixed-effects panel data models. *Econometric Reviews* 35: 1290–1316. <https://doi.org/10.1080/07474938.2014.976524>.
- Drukker, D. M. 2003. Testing for serial correlation in linear panel-data models. *Stata Journal* 3: 168–177. <https://doi.org/10.1177/1536867X0300300206>.
- Inoue, A., and G. Solon. 2006. A portmanteau test for serially correlated errors in fixed effects models. *Econometric Theory* 22: 835–851. <https://doi.org/10.1017/S0266466606060385>.
- Jochmans, K. 2019. Testing for correlation in error-component models. Cambridge Working Papers in Economics No. 1993, University of Cambridge, Faculty of Economics. <https://doi.org/10.17863/CAM.47807>.

Wooldridge, J. M. 2010. *Econometric Analysis of Cross Section and Panel Data*. 2nd ed. Cambridge, MA: MIT Press.

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

Koen Jochmans is a reader in the Faculty of Economics and a fellow of Trinity Hall at the University of Cambridge.

Vincenzo Verardi is an associate researcher of the FNRS.