UNIVERSITY OF CAMBRIDGE
Faculty of Economics

Cambridge Working Papers in Economics

Cambridge Working Papers in Economics: 1886

A PORTMANTEAU TEST FOR CORRELATION IN SHORT PANELS

Koen Jochmans

3 January 2018

Inoue and Solon (2006, A Portmanteau test for serially correlated errors in fixed effects models, Econometric Theory 22, 835{851) presented an elegant approach to test for serial correlation of arbitrary form in fixed-effect models for short panel data. Their approach requires the choice of a regularization parameter that may severely affect the power of the test and for which no optimal selection rule is available. We present a modified version of their test that uses strictly more information and does not require any regularization parameter. Monte Carlo simulations are provided to illustrate the power gains of our procedure.

A PORTMANTEAU TEST FOR CORRELATION IN SHORT PANELS

KOEN JOCHMANS*

UNIVERSITY OF CAMBRIDGE

THIS VERSION: DECEMBER 21, 2018

Abstract

Inoue and Solon (2006, A Portmanteau test for serially correlated errors in fixed effects models, *Econometric Theory* 22, 835–851) presented an elegant approach to test for serial correlation of arbitrary form in fixed-effect models for short panel data. Their approach requires the choice of a regularization parameter that may severely affect the power of the test and for which no optimal selection rule is available. We present a modified version of their test that uses strictly more information and does not require any regularization parameter. Monte Carlo simulations are provided to illustrate the power gains of our procedure.

Keywords: fixed effects, panel data, statistical power, serial correlation, test.

*Address: University of Cambridge, Faculty of Economics, Austin Robinson Building, Sidgwick Avenue, Cambridge CB3 9DD, U.K. E-mail: kj345@cam.ac.uk. Support of the European Research Council through grant n$^{\text{o}}$ 715787 (MiMo) is gratefully acknowledged.

1 Introduction

In panel data models with fixed effects testing whether errors are serially correlated is complicated by the need to estimate the unit-specific intercepts, especially in short panels. Tests against (first-order) autoregressive errors have been developed by Baltagi and Wu (1999), Wooldridge (2002, pp. 275), and Wooldridge (2002, pp. 282–283) and Drukker (2003). Inoue and Solon (2006) proposed an elegant Portmanteau test, i.e., a procedure to test against serial correlation of arbitrary form. Such an approach is desirable 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 unit do not have a natural ordering such as time. Indeed, the Monte Carlo work in Inoue and Solon (2006) aptly illustrates the gain in terms of power of using a Portmanteau test.

The test of Inoue and Solon (2006) involves the choice of a regularization parameter. This choice affects the power of the test, possibly quite dramatically. The way in which it does so, and to what extent, depends on the particular alternative under consideration. Hence, choosing the regularization parameter amounts to tacking a stand on what type of alternative one wishes to best arm oneself against. This is at odds with the Portmanteau paradigm.

We show below that working with this regularization parameter is both inefficient and unnecessary. A modified Portmanteau test is obtained that uses more information and does not involve any regularization parameter. Like the original test statistics of Inoue and Solon (2006), ours is straightforward to implement, and it can be adapted to handle unbalanced panel data in the same way. Some numerical results on power are provided to support our claims.

The tests discussed here are designed to be applicable to short panels. A Portmanteau test for long panels was developed by Okui (2009).

2 Testing for serial correlation

Suppose we have $N \times T$ panel data where, for each randomly sampled unit $i = 1, \ldots, N$, we observe the vector of outcomes $\boldsymbol{y}_i := (y_{i1}, \ldots, y_{iT})'$ along with the $T \times K$ matrix of covariates $\boldsymbol{X}_i : (\boldsymbol{x}_{i1}, \ldots, \boldsymbol{x}_{iT})'$. The standard fixed-effect model for such data specifies that

$$\boldsymbol{y}_i = \boldsymbol{X}_i \boldsymbol{\beta} + c_i \, \boldsymbol{\iota}_T + \boldsymbol{\varepsilon}_i,$$

where c_i is unit i's fixed effect, $\boldsymbol{\iota}_T$ is the T-vector of ones, and $\boldsymbol{\varepsilon}_i := (\varepsilon_{i1}, \ldots, \varepsilon_{iT})'$ is a vector of shocks. The latter are taken to be mean-independent of the covariates and the fixed effect. The standard specification (see, e.g., Wooldridge 2002) additionally assumes that these latter shocks are homoskedastic and serially uncorrelated. Our interest here lies in testing this assumption.

2.1 Inoue and Solon's (2006) Portmanteau test

Using the Lagrange-Multiplier (LM) principle, Inoue and Solon (2006) proposed a simple test for the null

$$\boldsymbol{\Sigma} := E(\boldsymbol{\varepsilon}_i \boldsymbol{\varepsilon}_i') = \sigma^2 \, \boldsymbol{I}_T, \tag{2.1}$$

where $\boldsymbol{I}_T$ is the $T \times T$ identity matrix and σ^2 is an unknown positive constant. We assume throughout that $\boldsymbol{\Sigma}$ is positive definite. To describe their approach let $\boldsymbol{M} := \boldsymbol{I}_T - \frac{1}{T}(\boldsymbol{\iota}_T \boldsymbol{\iota}_T')$ be the matrix that transforms observations into deviations from within-group means and let

$$\hat{\boldsymbol{\beta}} := \left(\sum_{i=1}^{N} \boldsymbol{X}_i' \boldsymbol{M} \boldsymbol{X}_i \right)^{-1} \left(\sum_{i=1}^{N} \boldsymbol{X}_i' \boldsymbol{M} \boldsymbol{y}_i \right)$$

be the fixed-effect least-squares estimator of $\boldsymbol{\beta}$. Under standard regularity conditions, $\hat{\boldsymbol{\beta}} \xrightarrow{p} \boldsymbol{\beta}$ as $N \to \infty$ (and, here and later, T is held fixed). Consequently, the least-squares residuals, $\hat{\boldsymbol{e}}_i := \boldsymbol{M}\boldsymbol{y}_i - \boldsymbol{M}\boldsymbol{X}_i \hat{\boldsymbol{\beta}}$, consistently estimate the demeaned shocks $\boldsymbol{e}_i := \boldsymbol{M}\boldsymbol{\varepsilon}_i$. Now,

$$\boldsymbol{\Omega} := E(\boldsymbol{e}_i \boldsymbol{e}_i') = \boldsymbol{M} \boldsymbol{\Sigma} \boldsymbol{M},$$

which, under the null (2.1) equals $\sigma^2 \boldsymbol{M}$. An (unconstrained) plug-in estimator and an estimator under the null take the form

$$\hat{\boldsymbol{\Omega}} := \frac{1}{N}\sum_{i=1}^{N} \hat{\boldsymbol{e}}_i \hat{\boldsymbol{e}}_i', \qquad \check{\boldsymbol{\Omega}} := \frac{1}{N}\sum_{i=1}^{N} \frac{\hat{\boldsymbol{e}}_i' \hat{\boldsymbol{e}}_i}{T-1} \boldsymbol{M},$$

respectively. A natural way to test (2.1), then, is to evaluate whether their difference can be considered large (in magnitude) under the null. Note that the within-group operation implies that $\boldsymbol{\iota}_T' \boldsymbol{e}_i = 0$ and $\boldsymbol{\iota}_T' \hat{\boldsymbol{e}}_i = 0$ for all i. Hence, $\boldsymbol{\Omega}$ and the estimators $\hat{\boldsymbol{\Omega}}$ and $\check{\boldsymbol{\Omega}}$ all have rank $T-1$.

It will prove fruitful to introduce the shorthand

$$\hat{\boldsymbol{v}}_i := \operatorname{vec}\left(\hat{\boldsymbol{e}}_i \hat{\boldsymbol{e}}_i' - \frac{\hat{\boldsymbol{e}}_i' \hat{\boldsymbol{e}}_i}{T-1} \boldsymbol{M} \right).$$

This allows to write

$$\bar{\hat{\boldsymbol{v}}} := \frac{1}{N}\sum_{i=1}^{N} \hat{\boldsymbol{v}}_i = \operatorname{vec}(\hat{\boldsymbol{\Omega}} - \check{\boldsymbol{\Omega}}), \qquad \hat{\boldsymbol{V}} := \frac{1}{N}\sum_{i=1}^{N} \hat{\boldsymbol{v}}_i \hat{\boldsymbol{v}}_i',$$

The reduced-rank of the covariance matrix estimators implies that $\hat{\boldsymbol{V}}$, too, is singular. To sidestep this issue Inoue and Solon (2006) work with a $(T-1)\times(T-1)$ submatrix of $\hat{\boldsymbol{\Omega}}$ and $\check{\boldsymbol{\Omega}}$ by dropping their nth row and column, where n is to be chosen. Any such submatrix involves

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

distinct covariances. To state their test statistic introduce the $T^2 \times m$ selection matrix $\boldsymbol{\Delta}_n := \partial \operatorname{vec}(\boldsymbol{\Omega}) / \partial \operatorname{vech}(\boldsymbol{\Omega}_{-n})'$, where $\boldsymbol{\Omega}_{-n}$ is the submatrix obtained on deleting the nth row and column from $\boldsymbol{\Omega}$. The test statistic of Inoue and Solon (2006) can then be written as

$$\mathrm{LM}_n := N\, \bar{\hat{\boldsymbol{v}}}' \boldsymbol{\Delta}_n (\boldsymbol{\Delta}_n' \hat{\boldsymbol{V}} \boldsymbol{\Delta}_n)^{-1} \boldsymbol{\Delta}_n' \bar{\hat{\boldsymbol{v}}},$$

for chosen $n \in \{1, \ldots, T\}$.

Inoue and Solon (2006) derived the large-sample properties of this test statistic under standard conditions collected in Assumption 1.

Assumption 1.

(a) $\boldsymbol{X}_i$ *and* $\boldsymbol{\varepsilon}_i$ *are i.i.d. and have finite fourth-order moments;*

(b) $E(\boldsymbol{\varepsilon}_i|\boldsymbol{X}_i, c_i) = \mathbf{0}$*;*

(c) $E(\boldsymbol{X}_i'\boldsymbol{M}\boldsymbol{X}_i)$ *is non-singular;*

(d) $T \geq 3$.

In the sequel, we let $\boldsymbol{V}$ be the probability limit of $\hat{\boldsymbol{V}}$, $\boldsymbol{C}$ denotes any symmetric $T \times T$ matrix with zero diagonal and

$$\boldsymbol{\delta} := \mathrm{vec}\left(\boldsymbol{MCM} - \frac{\mathrm{tr}(\boldsymbol{MC})}{T-1}\boldsymbol{M}\right),$$

with its dependence on $\boldsymbol{C}$ left implicit. The following is a restatement of Theorem 1 in Inoue and Solon (2006).

Theorem 1. *Let Assumption 1 hold and suppose that the matrix* $\boldsymbol{\Delta}_n'\boldsymbol{V}\boldsymbol{\Delta}_n$ *is non-singular.*

(a) Under the null (2.1),

$$\mathrm{LM}_n \xrightarrow{d} \chi^2_m.$$

(b) Under a sequence of local alternatives of the form $\boldsymbol{\Sigma} = \sigma^2\boldsymbol{I}_T + \boldsymbol{C}/\sqrt{N}$,

$$\mathrm{LM}_n \xrightarrow{d} \chi^2_m(\gamma_n),$$

where $\chi^2_m(\gamma_n)$ *is a non-central* χ^2*-distribution with* m *degrees of freedom and non-centrality parameter* $\gamma_n := \boldsymbol{\delta}'\boldsymbol{\Delta}_n(\boldsymbol{\Delta}_n'\boldsymbol{V}\boldsymbol{\Delta}_n)^{-1}\boldsymbol{\Delta}_n'\boldsymbol{\delta}$.

The theorem shows that (i) the number of restrictions being tested is m; and (ii) that the power of the test depends on the choice of n. Inoue and Solon (2006) provide a discussion on (ii) and illustrate it in their Monte Carlo work.

2.2 A new Portmanteau test

Note that $\hat{\boldsymbol{v}}_i$ can be seen as a plug-in estimator of

$$\boldsymbol{v}_i := \mathrm{vec}\left(\boldsymbol{e}_i\boldsymbol{e}_i' - \frac{\boldsymbol{e}_i'\boldsymbol{e}_i}{T-1}\boldsymbol{M}\right).$$

Now,

$$E(\boldsymbol{v}_i) = \operatorname{vec}\left(\boldsymbol{\Omega} - \frac{\sigma^2}{T}\boldsymbol{M}\right)$$

and so the null (2.1) can be equivalently stated as the collection of T^2 moment conditions

$$E(\boldsymbol{v}_i) = \mathbf{0}.$$

A test based on LM_n can be understood as a joint test of a subset of m moment conditions

$$E(\boldsymbol{\Delta}_n' \boldsymbol{v}_i) = \mathbf{0}.$$

The matrix $\boldsymbol{\Delta}_n$ serves to resolve the fact that some of the moment conditions are linear combinations of the others and, hence, redundant. This selection is too severe, however, as there are

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

linearly-independent moment conditions and $r - m = T - 2$. Thus, LM_n ignores useful information.

When $T = 3$, for example, LM_n tests a single moment condition. For $n = 1, 2, 3$, respectively, it checks whether

$$(\hat{\boldsymbol{\Omega}})_{2,3} - (\check{\boldsymbol{\Omega}})_{2,3}, \quad (\hat{\boldsymbol{\Omega}})_{1,3} - (\check{\boldsymbol{\Omega}})_{1,3}, \quad \text{or} \quad (\hat{\boldsymbol{\Omega}})_{1,2} - (\check{\boldsymbol{\Omega}})_{1,2},$$

is large relative to its standard deviation. The first and third of these test statistics look at first-order covariances while the second involves a second-order covariance. If serial dependence is most pronounced at first order, LM_2 would have considerably more difficulty to pick up violations from (2.1) than would LM_1 and LM_3. However, any pair of these three moments are linearly independent, and so we may equally test them jointly. While a joint test need not be more powerful than the most powerful of the LM_n tests (as a function of n), one would expect it to perform at least as well as the worst of the LM_n tests and typically considerably better.

Combining all available moments also avoids the need to choose the regularization parameter n, making the implementation of the test immune to data snooping. A simple way to proceed is to retain all those moments that relate to the lower diagonal part of the matrix $\hat{\boldsymbol{\Omega}} - \check{\boldsymbol{\Omega}}$ except for the lower-left entry, i.e., $(\hat{\boldsymbol{\Omega}} - \check{\boldsymbol{\Omega}})_{T,1}$.[1] If we let $\overline{\mathrm{vech}}(\boldsymbol{\Omega})$ denote the operator that returns these lower-diagonal entries for matrix $\boldsymbol{\Omega}$ then we can write our test statistic as

$$\mathrm{LM} := N\, \bar{\hat{\boldsymbol{v}}}' \boldsymbol{\Delta}(\boldsymbol{\Delta}'\hat{\boldsymbol{V}}\boldsymbol{\Delta})^{-1}\boldsymbol{\Delta}'\bar{\hat{\boldsymbol{v}}}, \qquad \boldsymbol{\Delta} := \partial \mathrm{vec}(\boldsymbol{\Omega})/\partial \overline{\mathrm{vech}}(\boldsymbol{\Omega})'.$$

The following theorem follows along the same lines as Theorem 1 of Inoue and Solon (2006).

Theorem 2. *Let Assumption 1 hold and suppose that the matrix $\boldsymbol{\Delta}'\boldsymbol{V}\boldsymbol{\Delta}$ is non-singular.*

(a) Under the null (2.1)*,*

$$\mathrm{LM} \xrightarrow{d} \chi^2_r.$$

(b) Under a sequence of local alternatives of the form $\boldsymbol{\Sigma} = \sigma^2 \boldsymbol{I}_T + \boldsymbol{C}/\sqrt{N}$,

$$\mathrm{LM} \xrightarrow{d} \chi^2_r(\gamma),$$

where $\chi^2_r(\gamma)$ is a non-central χ^2-distribution with r degrees of freedom and non-centrality parameter $\gamma := \boldsymbol{\delta}'\boldsymbol{\Delta}(\boldsymbol{\Delta}'\boldsymbol{V}\boldsymbol{\Delta})^{-1}\boldsymbol{\Delta}'\boldsymbol{\delta}$.

To contrast with Theorem 1 note that Theorem 2 involves testing strictly more moment conditions and that the limit distributions under both the null and the alternative are independent of a regularization parameter.

3 Monte Carlo illustration

To illustrate the potential gains from using LM over LM_n we conducted a simulation experiment. Random samples of size $N = 100$ where drawn from the multivariate normal distribution $N(\boldsymbol{0}, \boldsymbol{\Sigma})$. This corresponds to a simple model with no regressors and all fixed effects set to zero.

We report results for two types of configurations for $\boldsymbol{\Sigma}$ and for $T \in \{3, 4\}$. The first configuration has dependence between the first and second observation but not between the others, i.e.,

$$\boldsymbol{\Sigma} = \begin{pmatrix} 1 & \rho & 0 \\ \rho & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \qquad \boldsymbol{\Sigma} = \begin{pmatrix} 1 & \rho & 0 & 0 \\ \rho & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix},$$

for $T = 3$ and $T = 4$, respectively. Here, ρ is the correlation between the first and second measurement. The second configuration considered is a stationary autoregressive process, i.e.,

$$\boldsymbol{\Sigma} = \frac{1}{1-\rho^2} \begin{pmatrix} 1 & \rho & \rho^2 \\ \rho & 1 & \rho \\ \rho^2 & \rho & 1 \end{pmatrix}, \qquad \boldsymbol{\Sigma} = \frac{1}{1-\rho^2} \begin{pmatrix} 1 & \rho & \rho^2 & \rho^3 \\ \rho & 1 & \rho & \rho^2 \\ \rho^2 & \rho & 1 & \rho \\ \rho^3 & \rho^2 & \rho & 1 \end{pmatrix},$$

where ρ is the autoregressive parameter. Under the null all covariance matrices reduce to the identity matrix of appropriate dimension.

Figure 1 contains the sampling distributions of the test statistics under the null for $T = 3$. From top left to bottom right the plots relate to LM_1, LM_2, LM_3, and finally, LM. The relevant asymptotic distribution for each statistic is also given in each plot. For the first three statistics this is the χ^2_1 distribution. For the LM statistic it is the χ^2_2 distribution. For each of the statistics the asymptotic approximation is quite satisfactory in light of the small sample size.

The plots in Figure 2 provide the power functions for all the tests and all configurations as a function of ρ. The upper two plots concern the non-stationary configuration. The lower two plots deal with the autoregressive specification. Plots on the left contain power functions for $T = 3$. Plots on the right give power functions for $T = 4$. The power functions of the LM_n tests are dashed and dashed-dotted lines while the power function of the LM test is marked by a full line. In our designs the power functions of several of the LM_n tests overlap so that only two distinct curves are visible in each plot. The plots also include a dashed horizontal line at the chosen significance level of 5%.

The plots clearly show that the choice of the regularization parameter n can have quite

Figure 1: Sampling distributions under the null

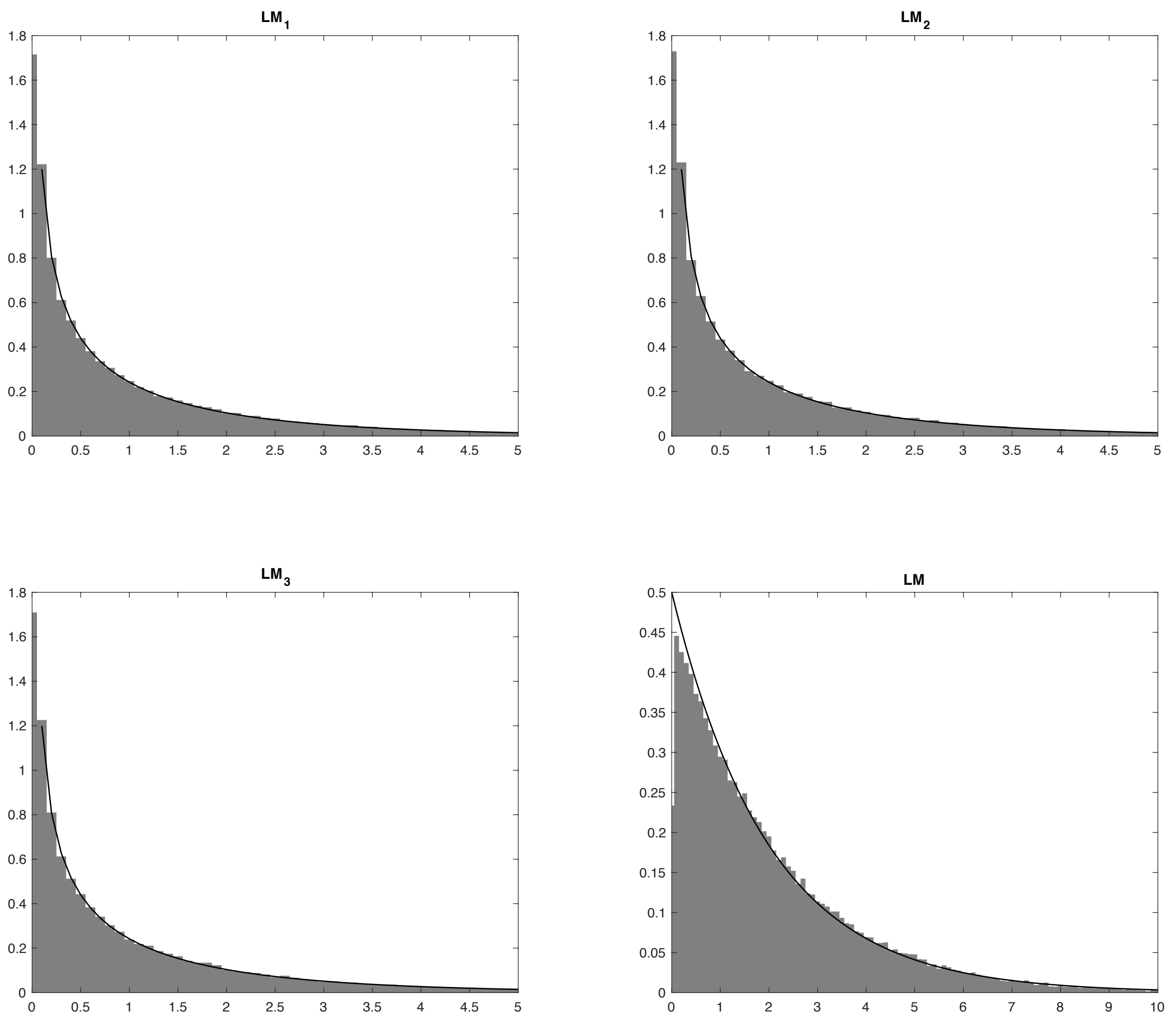

Sampling distributions computed over 100,000 Monte Carlo replications. Reference distribution is χ^2_1 for LM_1, LM_2, and LM_3 and χ^2_2 for LM.

Figure 2: Power functions

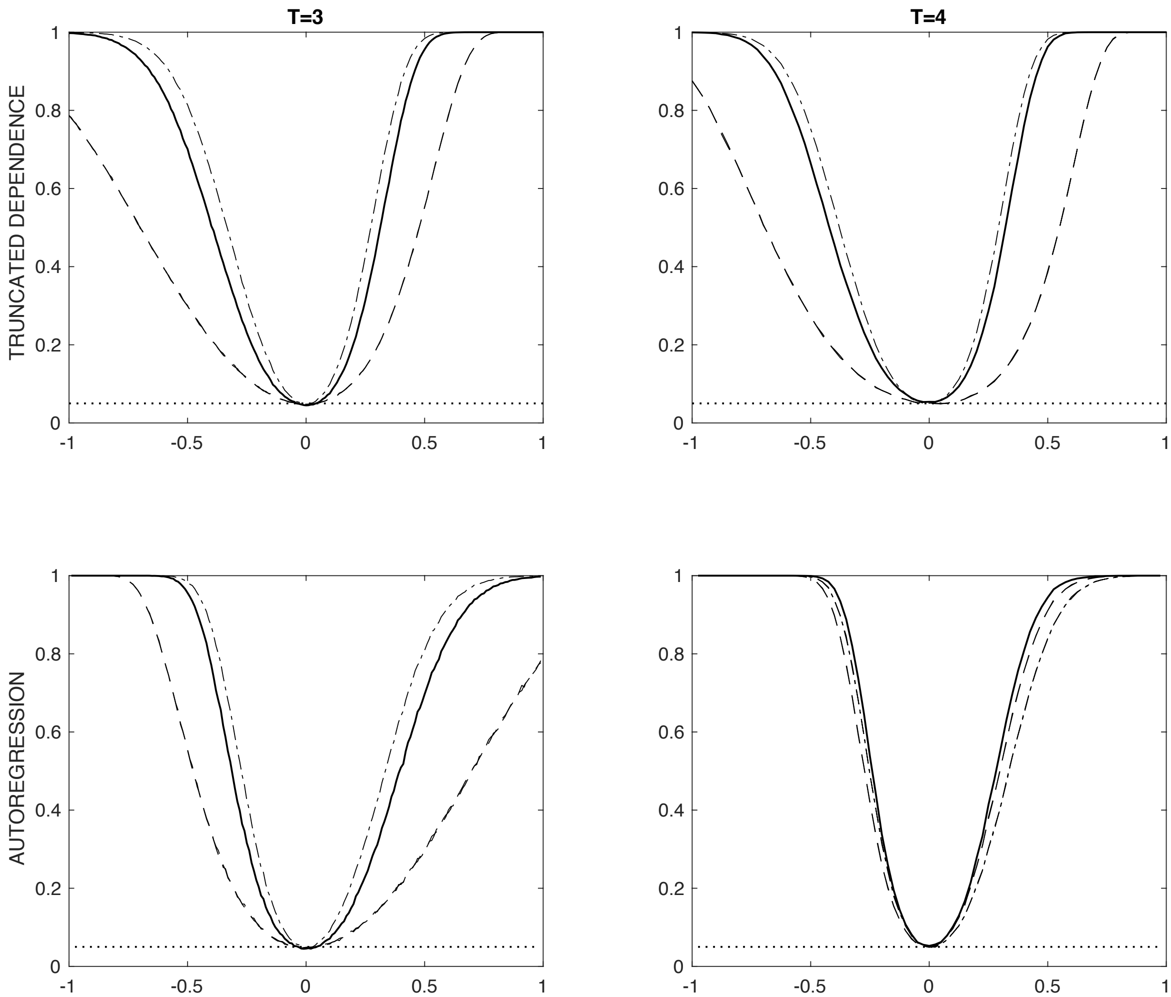

Power functions computed over 100,000 Monte Carlo replications. Theoretical size is 5%.

dramatic effects on power. For example, in the upper-left plot LM_3 picks up violations from the null much more frequently than do LM_1 and LM_2 while in the lower-left plot both LM_1 and LM_3 have much lower power than LM_2. The plots on the right concern four LM_n tests, and each one of the two functions plotted corresponds to two of these tests. The joint LM test always does considerably better than the worst of the LM_n tests. Moreover, in all but the last design it is only marginally less powerful than the strongest of the LM_n tests. When $T = 4$ the LM test is actually the most powerful against autoregressive alternatives of all the tests considered, and this uniformly over all ρ in $(-1, 1)$; the difference with the most powerful LM_n test is, nonetheless, small.

Notes

[1]Other choices are possible. All give numerically the same test statistic.

References

Baltagi, B. H. and P. X. Wu (1999). Unequally spaced panel data regressions with AR(1) disturbances. *Econometric Theory 15*, 814–823.

Drukker, D. M. (2003). Testing for serial correlation in linear panel-data models. *Stata Journal 3*, 168–177.

Inoue, A. and G. Solon (2006). A Portmanteau test for serially correlated errors in fixed effects models. *Econometric Theory 22*, 835–851.

Okui, R. (2009). Testing serial correlation in fixed effects regression models based on asymptotically unbiased autocorrelation estimators. *Mathematics and Computers in Simulation 79*, 2897–2909.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.