cemfi

working paper
2302

Highly Irregular Serial Correlation Tests

Dante Amengual
Xinyue Bei
Enrique Sentana

May 2023

CENTRO DE ESTUDIOS
MONETARIOS Y FINANCIEROS
Casado del Alisal 5, 28014 Madrid, Spain
www.cemfi.es

CEMFI Working Paper No. 2302
May 2023

Highly Irregular Serial Correlation Tests

Abstract

We develop tests for neglected serial correlation when the information matrix is repeatedly singular under the null. Specifically, we consider white noise against a multiplicative seasonal AR model, and a local-level model against a nesting UCARIMA one. Our proposals, which involve higher-order derivatives, are asymptotically equivalent to the likelihood ratio test but only require estimation under the null. Remarkably, we show that our proposed tests effectively check that certain autocorrelations of the observations are 0, so their asymptotic distribution is standard. We conduct Monte Carlo exercises that study their finite sample size and power properties, comparing them to alternative approaches.

JEL Codes: C22, C32, C52, C12.

Keywords: Generalized extremum tests, higher-order identifiability, likelihood ratio test.

Dante Amengual
CEMFI
amengual@cemfi.es

Xinyue Bei
Duke University
xinyue.bei@duke.edu

Enrique Sentana
CEMFI
sentana@cemfi.es

Acknowledgement

We would like to thank Gabriele Fiorentini as well as the audience at the XIII Workshop in Time Series Econometrics (Zaragoza, March 2023) for useful comments and suggestions. Of course, the usual caveat applies. Amengual and Sentana gratefully acknowledge financial support from the Spanish Ministry of Science and Innovation through grant PID2021-128963NB-I00.

1 Introduction

The econometric literature on serial correlation tests, which can be traced back at least to Durbin and Watson (1950, 1951) and the references therein, is vast. Given that Lagrange Multiplier (LM) tests only require estimation of the model parameters under the null, following Breusch (1978) and Godfrey (1978a,b), they became the preferred choice for neglected serial correlation tests in econometric applications. In addition to computational considerations, which continue to be very relevant for resampling procedures, two other important advantages of LM tests are that (i) rejections provide a clear indication of the specific directions along which modelling efforts should focus, and (ii) they are often easy to interpret as moment tests, so they remain informative for alternatives they are not designed for. Furthermore, under standard regularity conditions, they are asymptotically equivalent to the Likelihood ratio (LR) and Wald tests under the null and sequences of local alternatives, and thus they share their optimality properties.

One of those standard regularity conditions is a full rank information matrix of the unrestricted model parameters evaluated under the null. However, Fiorentini and Sentana (2016) highlighted some examples of neglected serial correlation tests in which this condition does not hold despite the fact that the model parameters are locally identified both under the null and the alternative hypotheses. To tackle this problem, they applied the "extremum" tests proposed by Lee and Chesher (1986), thereby obtaining asymptotic chi-square distributions under the null. As is well known, Lee and Chesher (1986) studied the restrictions that the null imposes on higher-order optimality conditions. Sometimes, the second derivative suffices, but it might be necessary to study the third or even higher-order ones. They proved the asymptotic equivalence between their extremum tests and the corresponding LR tests under the null and sequences of local alternatives in unrestricted contexts. Using earlier results by Cox and Hinkley (1974), this equivalence intuitively follows from the fact that the extremum tests can often be re-interpreted as standard LM tests of a suitable transformation of the parameters such that the new information matrix is no longer singular. In contrast, Wald tests are extremely sensitive to reparametrization under these circumstances.

Importantly, though, the nullity of the information matrix of the alternative model under the null is assumed to be 1 in all the aforementioned references. The purpose of this paper is to develop tests for neglected serial correlation asymptotically equivalent to the LR test in some highly irregular situations in which the nullity of the information matrix is two or higher. To do so, we rely on the generalized extremum tests (GET) we have proposed in a companion paper – see Amengual, Bei and Sentana (2023). For illustrative purposes, we use as examples two classes of time series models very popular among practitioners: the multiplicative seasonal ARIMA models put forward by Box and Jenkins (1970), and the UCARIMA models, which constitute the basis of the "structural time series" models studied by Harvey (1989) (see Lippi and Reichlin (1992) for an insightful comparison of some important characteristics of these two models).

We show that our proposed tests effectively check that certain autocorrelations of the observations are 0, which in turn implies that their asymptotic distribution is standard. This is somewhat remarkable because GET statistics typically have unusual asymptotic distributions (see e.g. Amengual, Bei and Sentana (2022)).

We conduct Monte Carlo exercises that study the finite sample size and power properties of our proposal and compare it to other tests for neglected serial correlation. We find that our suggested parametric bootstrap procedures yield very reliable test sizes for the small sample sizes typically encountered in empirical applications to macroeconomic data. In addition, we confirm the power superiority of our tests over their competitors. Finally, we also confirm their substantial computational superiority over the corresponding LR tests, which require the maximization over the entire parameter space of an unrestricted log-likelihood function which is extremely flat around its maximum when the null hypothesis is true. These computational advantages are particularly pertinent for computing the bootstrap critical values mentioned above.

The rest of the paper is organized as follows. In Sections 2 and 3, we derive our proposed tests for the two aforementioned examples and study both their asymptotic properties and their finite sample ones. Next, we present our conclusions in Section 4, relegating proofs and some additional results to the appendix.

2 Multiplicative seasonal ARIMA models

Box and Jenkins (1970) introduced the popular multiplicative seasonal ARIMA model to capture the autocorrelation of series with strong seasonal patterns, such as their famous airline passenger example. The serial dependence structure of these models is perfectly understood, and the same is true of the properties of the maximum likelihood estimators (MLE) of their parameters in normal circumstances. Moreover, LM tests for neglected serial correlation in such models have been readily available for several decades.

However, what it is far less known is that in some cases, the standard regularity conditions that guarantee the asymptotic validity of such tests do not hold. Next, we showcase the difficulties involved by means of a rather simple example.

2.1 The test statistic

Suppose that after taking regular and seasonal differences of an observed time series, a researcher would like to formally assess the need for a more complicated dependence structure. Specifically, assuming the data is observed at the quarterly frequency, one of the alternatives that she might consider is the following Ar(2)-Sar(2) process:

$$(1-\vartheta_1 L)(1-\vartheta_2 L)(1-\vartheta_3 L^4)(1-\vartheta_4 L^4)(y_t-\varphi_M)=\varepsilon_t, \tag{1}$$

with $E(\varepsilon_t)=0$ and $V(\varepsilon_t)=\varphi_V$, where $y_t=\Delta\Delta_4 x_t$ and x_t denotes the original data, so that $H_0:\boldsymbol{\vartheta}=\mathbf{0}$, with $\boldsymbol{\vartheta}=(\vartheta_1,\vartheta_2,\vartheta_3,\vartheta_4)'$.

As usual, non-linear least squares estimation coincides with Gaussian ML, so that the criterion function will be

$$-\frac{T}{2}\ln(2\pi)-\frac{T}{2}\ln\varphi_V-\sum_{t=1}^{T}\frac{[y_t-\mu_t(\varphi_M,\boldsymbol{\vartheta})]^2}{2\varphi_V}, \tag{2}$$

where the conditional mean under the alternative is

$$\begin{aligned}\mu_t(\varphi_M,\vartheta)=&\varphi_M+(\vartheta_1+\vartheta_2)(y_{t-1}-\varphi_M)-\vartheta_1\vartheta_2(y_{t-2}-\varphi_M)+(\vartheta_3+\vartheta_4)(y_{t-4}-\varphi_M)\\&-(\vartheta_1+\vartheta_2)(\vartheta_3+\vartheta_4)(y_{t-5}-\varphi_M)+\vartheta_1\vartheta_2(\vartheta_3+\vartheta_4)(y_{t-6}-\varphi_M)\\&-\vartheta_3\vartheta_4(y_{t-8}-\varphi_M)+(\vartheta_1+\vartheta_2)\vartheta_3\vartheta_4(y_{t-9}-\varphi_M)-\vartheta_1\vartheta_2\vartheta_3\vartheta_4(y_{t-10}-\varphi_M).\end{aligned}$$

The model parameters under the null are φ_M and φ_V, whose restricted MLEs coincide with the sample mean and variance (with denominator T) of y_t. Moreover, the MLEs of the parameters of the alternative model, which also include $\boldsymbol{\vartheta}$, usually converge to their true values at the standard $\sqrt{T}$ rate.

However, as we shall formally prove below, the information matrix of model (1) evaluated at $\boldsymbol{\vartheta}=\mathbf{0}$ has two zero eigenvalues because

$$\frac{\partial l_t}{\partial\vartheta_1}-\frac{\partial l_t}{\partial\vartheta_2}=0 \tag{3}$$

and

$$\frac{\partial l_t}{\partial\vartheta_3}-\frac{\partial l_t}{\partial\vartheta_4}=0, \tag{4}$$

which makes this testing problem a highly irregular one.[1]

As we show in the proof of Proposition 1, we can find a suitable reparametrization

$$(\varphi_M,\varphi_V,\vartheta_1,\vartheta_2,\vartheta_3,\vartheta_4)\mapsto(\phi_M,\phi_V,\theta_{i1},\theta_{i2},\theta_{u1},\theta_{u2})$$

[1] This irregularity is particularly relevant for Wald tests, which are extremely sentitive to reparametrizations in this context. For example, Fiorentini and Paruolo (2009) found that the rate of convergence of a sequential Cochrane-Orcutt-type estimator of what is effectively the product of the first two autocorrelations of y_t is T rather than $T^{\frac{1}{2}}$ or $T^{\frac{1}{4}}$ when $\vartheta_1=\vartheta_2=0$ in a non-seasonal version of model (1) in which $\vartheta_3=\vartheta_4=0$.

that isolates the singularity in the last two parameters in such a way that the first derivatives of the log-likelihood function corresponding to θ_{u1} and θ_{u2} are both 0, where $\boldsymbol{\theta}_i' = (\theta_{i1}, \theta_{i2})$ contains the parameters of the alternative model that are first-order identified while $\boldsymbol{\theta}_u' = (\theta_{u1}, \theta_{u2})$ refers to those that are first-order underidentified but second-order identified.

Fortunately, the assumptions of Theorem 1 in Amengual, Bei and Sentana (2023) apply to the second derivatives

$$\frac{\partial^2 l_t}{(\partial \theta_{u1})^2} = \frac{2\,(y_t - \phi_M)\,(y_{t-2} - \phi_M)}{\phi_V}, \tag{5}$$

$$\frac{\partial^2 l_t}{\partial \theta_{u1} \partial \theta_{u2}} = 0, \tag{6}$$

and

$$\frac{\partial^2 l_t}{(\partial \theta_{u2})^2} = \frac{2\,(y_t - \phi_M)\,(y_{t-8} - \phi_M)}{\phi_V}, \tag{7}$$

because the asymptotic covariance matrix of

$$\frac{\partial l_t}{\partial \phi_M}, \frac{\partial l_t}{\partial \phi_V}, \frac{\partial l_t}{\partial \theta_{i1}}, \frac{\partial l_t}{\partial \theta_{i2}}, \theta_{u1}^2 \frac{\partial^2 l_t}{(\partial \theta_{u1})^2} + \theta_{u2}^2 \frac{\partial^2 l_t}{(\partial \theta_{u2})^2} + 2\theta_{u1}\theta_{u2} \frac{\partial^2 l_t}{\partial \theta_{u1} \partial \theta_{u2}}$$

scaled by $\sqrt{T}$ has full rank for any $(\theta_{u1}, \theta_{u2}) \neq (0, 0)$, which allows us to obtain the following result:

Proposition 1

$$LR_T = GET_T + O_p(T^{-\frac{1}{4}}),$$

under H_0, where LR_T is the likelihood ratio statistic based on (2), and

$$GET_T = T(\hat{r}_{1T}^2 + \hat{r}_{4T}^2 + \hat{r}_{2T}^2 \mathbf{1}[\hat{r}_{2T} \geq 0] + \hat{r}_{8T}^2 \mathbf{1}[\hat{r}_{8T} \geq 0]), \tag{8}$$

where $\mathbf{1}[.]$ is the usual indicator function and

$$\hat{r}_{jT} = \frac{1}{T} \sum\nolimits_t \frac{(y_t - \tilde{\phi}_M)(y_{t-j} - \tilde{\phi}_M)}{\tilde{\phi}_V},$$

with $\tilde{\phi}_M = T^{-1} \sum_t y_t$ and $\tilde{\phi}_V = T^{-1} \sum_t (y_t - \tilde{\phi}_M)^2$.

Therefore, the GET_T statistic is simply focusing on the first two regular sample autocorrelations and the first two seasonal ones, which is very intuitive in view of (1). Given that these estimated autocorrelations are asymptotically independent under the null, the asymptotic distribution of (8) will be a mixture of χ_2^2, χ_3^2 and χ_4^2 with weights $\frac{1}{4}$, $\frac{1}{2}$ and $\frac{1}{4}$, respectively.[2]

Furthermore, we can show that a test of white noise against the multiplicative AR(k)-SAR(k_s) model

$$\prod\nolimits_{j=1}^{k} (1 - \vartheta_j L) \prod\nolimits_{j=k+1}^{k+k_s} (1 - \vartheta_j L^4)(y_t - \varphi_M) = \varepsilon_t$$

[2] The partially one-sided nature of the test arises from the multiplicative nature of the alternative, which forces the roots to be always real. Additive alternatives, which allow for complex roots too, give rise to two-sided tests.

for $k \geq 3$ or $k_s \geq 3$ will numerically coincide with the statistic in (8). The rationale is as follows. When the null is true, we can prove that the MLE of an additive AR(3) is such that all three roots of the lag polynomial are real with probability tending to 0, unless one of the roots is forced to be 0. Consequently, the LR for multiplicative AR(3) is asymptotically equivalent to the LR for AR(2), and the same applies to the corresponding GETs. Perhaps less surprisingly, we can also show that we would obtain exactly the same test statistic if we considered multiplicative MA alternatives instead.

Finally, it is important to mention that our proposed test, which is based on sample autocorrelations, is numerically invariant to affine transformations of the observed series y_t. Effectively, this means that its finite sample distribution is pivotal with respect to $\boldsymbol{\varphi} = (\varphi_M, \varphi_V)'$. Therefore, one can estimate the sample mean and variance of y_t, and apply our test directly to the standardized series as if they were the observed variables.

2.2 Simulation evidence

Next, we study the finite sample size and power properties of the testing procedures we introduced in the previous subsection by means of several extensive Monte Carlo exercises. Given that no nuisance parameters are effectively involved under the null, we can set the unconditional mean and variance of the innovation ε_t to 0 and 1, respectively, both under the null and alternative hypotheses without any loss of generality. We also estimate φ_M and φ_V with the sample mean and variance, respectively, which effectively impose the null.

As alternative hypotheses we consider the covariance stationary models

$$(1 - .1L - .1L^2 - .1L^3 - .1L^4)y_t = \varepsilon_t \quad (H_{a_1})$$

and

$$(1 - .4L)(1 + .4L)(1 - .4L^4)(1 + .4L^4)y_t = \varepsilon_t \quad (H_{a_2}),$$

for which the first, second, fourth and eighth autocorrelation coefficients in the population are (0.14,0.14,0.14,0.03) and (0,0.16,0.03,0.16). Note that two of the roots of the first process are complex conjugates, while our test is designed for the case of real roots.

We approximate the exact finite sample distribution using 10,000 simulated samples under the maintained hypothesis that the y_t's are *i.i.d.* as standard normals. In fact, we could thus obtain "exact" critical values for any sample size by increasing the number of simulations. Alternatively, one could consider a non-parametric bootstrap procedure that randomly draws with replacement from the observations, which would eliminate any time series dependence while allowing for any marginal distribution. Either way, we do not need to take into account the sensitivity of the critical values to $\tilde{\varphi}$ because the test statistics are numerically invariant to the values of these estimators.

In Table 1 we compare the results of our test for $T = 100$ (top) and $T = 400$ (bottom) with three alternative procedures: LM-AR(1) and LM-SAR(4), which denote standard LM tests based on the score of an AR(1) and a Wallis (1972)-style seasonal AR(4), respectively, and a moment test based on the first two regular sample autocorrelations and the first two seasonal ones (MT), which is effectively the two-sided version of (8), whose asymptotic distribution is χ^2_4 under the null.

In turn, the last six columns present the rejection rates at the 1%, 5% and 10% levels for the two alternatives. The behavior of the different test statistics is in accordance with expectations. In particular, our proposal is the most powerful for H_{a_2}, which is not very surprising given that it is designed to direct power against such multiplicative alternatives with real roots. But it is also the top performer for H_{a_1} even though the process has two complex roots, which is perhaps not entirely surprising in view of the positive value of the relevant population autocorrelations.

The scatterplot in Figure 1 visually illustrates the asymptotic equivalence under the null between LR_T and GET_T statistics stated in Proposition 1, with the Gaussian rank correlation coefficients[3] between them being 0.932 and 0.986 across Monte Carlo samples of size $T = 100$ and 400, respectively. Finally, our results also indicate indicate that the LR takes 755 (921) seconds of CPU time for 10,000 samples of length 100 (400), while computing GET only requires 0.20 (0.24) seconds, respectively, which makes a huge difference in the calculation of the bootstrap critical values.

3 UCARIMA models

These popular unobserved component models assume that the observed time series are the superposition of two or more latent ARIMA time series models, whose parameters can be estimated by maximizing the Gaussian log-likelihood function of the observed data, which can be readily obtained either as a by-product of the Kalman filter prediction equations or from Whittle's (1962) frequency domain asymptotic approximation. Once the parameters have been estimated, filtered values of the unobserved components can be extracted by means of the Kalman smoother or its Wiener-Kolmogorov counterpart. These estimation and filtering issues are well understood (see e.g. Harvey (1989) for a textbook treatment).

[3]The Gaussian rank correlation coefficient between two variables is the usual Pearson correlation coefficient between the Gaussian scores of those variables, which are obtained by applying the inverse Gaussian cumulative distribution function transform to the ranks of the observations on each variable divided by $n+1$ (see Amengual, Sentana and Tian (2022)). Like the Spearman correlation coefficient, the Gaussian one is less sensitive to outliers than the Pearson one.

In contrast, tests that assess the correct specification of the parametric ARIMA models for the underlying components are far less well studied, even though the various outputs of an UCARIMA model could be misleading under misspecified dynamics. As mentioned in the introduction, Fiorentini and Sentana (2016) provided a thorough discussion of such tests, highlighting the popular local level model as an example in which the LM test cannot be computed in the usual way because the information matrix of the alternative model is sometimes singular under the null. Unfortunately, the extremum tests of Lee and Chesher (1986) cannot be applied when the nullity of the information matrix is two or more. Next, we study a simple example of this situation.

3.1 The test statistic

Undoubtedly, the local level model is the most popular UCARIMA model among practitioners. It assumes that

$$\begin{aligned} x_t &= z_t + u_t, & (9) \\ \Delta z_t &= f_t, & (10) \\ u_t &= v_t, & (11) \\ \begin{pmatrix} f_t \\ v_t \end{pmatrix} | I_{t-1} &\sim N\left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_f^2 & 0 \\ 0 & \sigma_v^2 \end{pmatrix}\right], & (12) \end{aligned}$$

where f_t and v_t follow two univariate white noise processes orthogonal at all leads and lags, and σ_f^2 and σ_v^2 are both strictly positive to exclude degenerate cases. Thus, the observed series is simply a random walk plus noise, whose first differences $y_t = \Delta x_t$ follow an MA(1) process with coefficient

$$\beta_y = \frac{1}{2}\left(\sqrt{q^2 + 4q} - 2 - q\right) < 0, \quad (13)$$

where $q = \sigma_f^2/\sigma_v^2$ denotes the positive but bounded signal to noise ratio, and residual variance

$$\sigma_a^2 = -\sigma_v^2/\beta_y.$$

As is well known, this model justifies the popular Exponentially Weighted Moving Average (EWMA) prediction rule, which has proved remarkably successful in many applications ranging from macro time series to volatility forecasts. However, EWMA predictions become suboptimal if (10) or (11) are dynamically misspecified, so it makes sense to test them against some more general alternatives.

To illustrate the issues that may arise, we consider the following nesting model:

$$\left.\begin{array}{c} (1 - \psi_1 L - \psi_2 L^2)\Delta z_t = f_t \\ (1 - \alpha L)u_t = v_t \end{array}\right\} \quad (14)$$

in which the "signal" z_t follows an an ARIMA(2,1,0) process while the "noise" u_t a stationary AR(1) process. As a result, the null hypothesis of interest is $H_0 : \alpha = \psi_1 = \psi_2 = 0$.

Once again, we can formally prove that the nullspace of the information matrix of the parameters of model (14) evaluated under the null is 2 because the first-derivatives of the log-likelihood function corresponding to ψ_1 and ψ_2 are linear combinations of the ones corresponding to σ_f^2, σ_v^2 and α. As a result, we show in the proof of Proposition 2 that we can find a suitable reparametrization

$$(\sigma_f^2, \sigma_v^2, \alpha, \psi_1, \psi_2) \mapsto (\sigma_f^{2\dagger}, \sigma_v^{2\dagger}, \alpha^\dagger, \psi_1^\dagger, \psi_2^\dagger)$$

that isolates the singularity in the last two parameters in such a way that the first-derivatives of the log-likelihood function corresponding to $\psi_1^\dagger$ and $\psi_2^\dagger$ are both 0.

Like Fiorentini and Sentana (2016), we can explicitly relate this singularity to the identification conditions for UCARIMA models in Hotta (1989). Specifically, although model (14) is generally identified, it is locally equivalent around the null to the following model:

$$\left.\begin{array}{c} \Delta z_t = (1 - \psi_1 L - \psi_2 L^2) f_t \\ u_t = (1 - \alpha L) v_t \end{array}\right\}, \tag{15}$$

in the sense that the (absolute value of the) scores and information matrices are identical when H_0 holds. Unlike model (14), which generates the autocorrelation structure of a restricted ARMA(3,3) for y_t, model (15) generates the autocorrelation structure of an unrestricted MA(2), which depends on three parameters only, namely the two MA coefficients plus the variance of the reduced form innovations. In contrast, model (15) depends on five parameters, namely ψ_1, ψ_2 and α together with σ_f^2 and σ_v^2, which means that the MA(2) reduced form can only identified a manifold of dimension 2 of the structural parameters.

In addition, we show that after the aforementioned reparametrization,

$$\frac{\partial^2 l_t}{(\partial \psi_1^\dagger)^2} = 0$$

and

$$\frac{\partial^3 l_t}{(\partial \psi_1^\dagger)^3} = 0,$$

while

$$\frac{\partial^2 l_t}{(\partial \psi_2^\dagger)^2} \neq 0,$$

which means that these two parameters have different degrees of identification. Fortunately, the assumptions of the more general Theorem 2 in Amengual, Bei and Sentana (2023) apply, allowing us to obtain the following result:

Proposition 2

$$LR_T = GET_T + O_p(T^{-\frac{1}{8}}),$$

under H_0, where LR_T is the corresponding likelihood ratio statistic, and

$$GET_n = \begin{pmatrix} \tilde{r}_{2T} & \tilde{r}_{3T} & \tilde{r}_{4T} \end{pmatrix} \mathcal{V}^{-1}_{\boldsymbol{\rho}_a\boldsymbol{\rho}_a} \begin{pmatrix} \tilde{r}_{2T} \\ \tilde{r}_{3T} \\ \tilde{r}_{4T} \end{pmatrix},$$

with

$$\tilde{r}_{jT} = \frac{\sum_t y_t y_{t-j}}{\sum_t y_t^2}$$

and

$$\mathcal{V}_{\boldsymbol{\rho}_a\boldsymbol{\rho}_a} = \lim_{T\to\infty} V\left[\sqrt{T}\begin{pmatrix} \tilde{r}_{2T} \\ \tilde{r}_{3T} \\ \tilde{r}_{4T} \end{pmatrix}\right]. \tag{16}$$

Therefore, both LR_T and GET_T are effectively testing that the second, third and fourth autocorrelations of y_t are 0. This result is not entirely surprising in view of the fact that y_t follows an MA(1) model under the null and an ARMA(3,3) under the alternative. Unlike what happened in the model discussed in section 2, though, the sample autocorrelations are no longer asymptotically independent under the null, so we need their asymptotic covariance matrix, which is particularly simple to obtain in the frequency domain using the expressions in Appendix B.1, as we explain in the proof of Proposition 2.

3.2 Simulation evidence

To assess the size properties of our proposed test, we generate 10,000 samples of lengths $T = 100$ and $T = 400$ of the local level model (9)-(12). Under the null, we simulate Gaussian shocks with $\sigma_f^2 = 1$ and $\sigma_v^2 = 0.5$, so that the signal to noise ratio is neither too small nor too high.

We compute GET_T using its spectral version (A3) with the information matrix (A2) estimated using (B7) after computing the periodogram using the fast Fourier transform. It is also important to emphasize that the LR_T statistic requires the estimation of model (14). For the reasons described in the introduction, this is a non-trivial numerical task. To increase the chances that we obtain the correct unrestricted ML estimates, we maximize the spectral log-likelihood of model (14) starting from the true values of the parameters in each design.

Although our main interest lies in the GET_T and LR_T statistics in Proposition 2, we also consider the following two moment tests for comparison purposes:

1. no second-order serial correlation in y_t,

2. no second- or third-order serial correlation in y_t.

Importantly, in computing these moment tests, we use the relevant elements of (A2) to obtain the adjusted asymptotic covariance matrix of the second and third sample autocovariances.

Unlike what happens in the multiplicative seasonal ARIMA model in section 2, the finite sample distribution of GET_T and LR_T is not pivotal with respect to the value of the signal to noise ratio q, even though both statistics are numerically invariant to the scale of Δy_t. For that reason, we conduct a parametric bootstrap procedure whereby for each of those 10,000 simulated samples, we simulate another $NB-1$ samples in which we set σ_f^2 to 1 without loss of generality and $(1+q)^{-1}$ to its estimated value, so that we can automatically compute size-adjusted rejection rates, as forcefully argued by Horowitz and Savin (2000).[4]

We present the rejection rates under the null for the tests at the 10%, 5% and 1% in the first three columns of Table 2 for samples of length 100 (top) and 400 (bottom).[5] As can be seen, all the testing procedures have reasonable size in both cases, which is reassuring for macro applications.[6]

Next, we simulate and estimate 10,000 samples of the same length of the following two alternative data generation processes (DGPs):

$$\left.\begin{array}{c}(1+0.5L+0.4L^2)\Delta z_t = f_t \\ (1-0.5L)u_t = v_t\end{array}\right\} \quad (H_{a_1})$$

and

$$\left.\begin{array}{c}(1-0.1L+0.5L^2)\Delta z_t = f_t \\ (1+0.5L)u_t = v_t\end{array}\right\} \quad (H_{a_2})$$

with the same σ_f^2 and σ_v^2 as in the null hypothesis. The first four autocorrelation coefficients of these processes in the population are (-0.32,-0.19,0.15,-0.04) and (-0.42,0.03,-0.15,0.15), respectively.

The corresponding rejection rates, which we report in the last six columns of Table 2, indicate that the behavior of the different test statistics is in accordance with expectations. For both alternatives, the GET and LR tests are more powerful than the competitors. Interestingly, our proposal is the most powerful for H_{a_2} while it has slightly less power than LR for H_{a_1}.

As in the first example, the scatterplot in Figure 2 visually illustrates the asymptotic equivalence under the null between LR_T and GET_T in Proposition 2, with the Gaussian rank correlation coefficients between the GET and LR test statistics across Monte Carlo samples of size

[4] In fact, the bounded support of $(1+q)^{-1}$ allows us to compute a table of "exact" critical values for a fine grid of values of this reduced-form MA coefficient before running the actual simulations (see Appendix D.1 in Amengual and Sentana (2015) for details). The same procedure works if we replace $(1+q)^{-1}$ by either β_y in (13) or the first-order autocorrelation of y_t, which are both between 0 and -1, but it is trickier to apply to q because this parameter can take any positive real value in the sample.

[5] Given the number of Monte Carlo replications, the 95% asymptotic confidence intervals for the Monte Carlo rejection probabilities under the null are (.80,1.20), (4.57,5.43) and (9.41,10.59) at the 1%, 5% and 10% levels.

[6] In a very small fraction of the samples of size $T = 100$ simulated under the null (0.62%), we encountered the "pile-up" problem associated to a positive first-order sample autocorrelation for y_t. In contrast, this never happened under either of the alternatives, or indeed when $T = 400$.

$T = 100$ and 400 generated under the null being 0.743 and 0.807, respectively, reflecting the slower rate of convergence. The simulation results also indicate that the LR takes 1250 (1763) seconds of CPU time for 10,000 samples of length 100 (400) while computing GET only requires 4.5 (5.5) seconds, respectively, which once again makes a huge difference in the calculation of the bootstrap critical values.

4 Conclusions

We characterize the singularity of the information matrix of a multiplicative seasonal Ar model à la Box and Jenkins under the null of white noise, as well as of a trend plus signal Ucarima model that nests the popular local level process. Using the generalization in Amengual, Bei and Sentana (2023) of the extremum-type tests in Lee and Chesher (1986) to models in which the nullity of the information matrix under the null hypothesis is strictly larger than one, we explain how to obtain an LM-type test based on higher-order derivatives which is asymptotically equivalent to the LR despite said singularity but only requires estimation under the null. This is particularly relevant for resampling-based inference because the fact that several log-likelihood derivatives are 0 under the null implies that the LR requires the estimation of all the parameters that appear under the alternative in a model whose log-likelihood function is extremely flat.

Our proposed dynamic specification tests are simple to implement and even simpler to interpret. And although some of our theoretical derivations make extensive use of frequency domain methods for time series, we provide a simple time domain interpretation of the statistics, so that empirical researchers who are not familiar with spectral analysis can still apply them easily.

We conduct Monte Carlo exercises that study the finite sample size and power properties of our proposals and compare them to alternative approaches. We find that our suggested parametric bootstrap procedures work very well, and that our tests have more power than alternative procedures. We also find that the computational advantages of our GET procedures relative to the LR ones are very substantial.

In the two examples that we consider the model parameters are only identified up to higher-order when the null is true. As a result, a local power analysis of our proposed tests would necessarily involve sequences of those parameters converging to zero at unusually low rates. Nevertheless, given that in both cases our test statistics have χ^2-like asymptotic distributions under the null, they would approximately follow non-central χ^2 distributions in large samples if we ignore inequality constraints. Finding exact expressions for the non-centrality parameters constitutes an interesting avenue for further research.

References

Amengual, D., Bei, X. and Sentana, E. (2022): "Normal but skewed?, *Journal of Applied Econometrics* 37, 1295–1313.

Amengual, D., Bei, X. and Sentana, E. (2023): "Hypothesis tests with a repeatedly singular information matrix", mimeo, CEMFI.

Amengual, D. and Sentana, E. (2015): "Is a normal copula the right copula?", CEMFI Working Paper 1504.

Amengual, D., Sentana, E. and Tian, Z. (2022): "Gaussian rank correlation and regression", chapter 11 in A. Chudik, C. Hsiao and A. Timmermann (eds.) *Essays in honor of M. Hashem Pesaran: Panel modeling, micro applications and econometric methodology*, *Advances in Econometrics* 43B, 269–306, Emerald.

Breusch, T.S. (1978): "Testing for autocorrelation in dynamic linear models", *Australian Economic Papers* 17, 334–355.

Box, G. and Jenkins, G. (1970): *Time series analysis: forecasting and control*. Holden-Day.

Cox, D. and Hinkley, D. (1974): *Theoretical statistics*, Chapman and Hall.

Dunsmuir, W. (1979): "A central limit theorem for parameter estimation in stationary vector time series and its application to models for a signal observed with noise", *Annals of Statistics* 7, 490-506.

Dunsmuir, W. and Hannan, E.J. (1976): "Vector linear time series models", *Advances in Applied Probability* 8, 339-364.

Durbin, J. and Watson, G.S. (1950): "Testing for serial correlation in least squares regression, I", *Biometrika* 37, 409–428.

Durbin, J. and Watson, G.S. (1951): "Testing for serial correlation in least squares regression, II", *Biometrika* 38, 159–179.

Fiorentini, G. and Paruolo, P. (2009): "Testing residual autocorrelation in autoregressive processes with a zero root", mimeo, Università di Firenze.

Fiorentini, G. and Sentana, E. (2016): "Neglected serial correlation tests in UCARIMA models", *SERIEs* 7, 121–178.

Fuller, W.A. (1995): *Introduction to time series analysis*, 2nd ed., Wiley.

Godfrey, L.G. (1978a): "Testing against general autoregressive and moving average error models when the regressors include lagged dependent variables", *Econometrica* 46, 1293–1301.

Godfrey, L.G. (1978b): "Testing for higher order serial correlation in regression equations when the regressors include lagged dependent variables", *Econometrica* 46, 1303–10.

Harvey, A.C. (1989): *Forecasting, structural models and the Kalman filter*, Cambridge University Press.

Horowitz, J. and Savin, N.E. (2000): "Empirically relevant critical values for hypothesis tests: a bootstrap approach", *Journal of Econometrics* 95, 375–389.

Hotta, L.K. (1989): "Identification of unobserved components models", *Journal of Time Series Analysis* 10, 259–270.

Lee, L.F. and Chesher, A. (1986): "Specification testing when score test statistics are identically zero", *Journal of Econometrics* 31, 121–149.

Lippi, M. and Reichlin, L. (1992): "On persistence of shocks to economic variables: a common misconception", *Journal of Monetary Economics* 29, 87–93.

Wallis, K.F. (1972): "Testing for fourth order autocorrelation in quarterly regression equations", *Econometrica* 40, 617–636.

Whittle, P. (1962): "Gaussian estimation in stationary time series", *Bulletin of the International Statistical Institute* 39, 105–129.

Appendices

A Proofs

For the sake of brevity, we have not included below a detailed verification that the multiplicative seasonal Arima model and the Ucarima models that we consider satisfy all the assumptions required for the application of Theorem 1 and 2 in Amengual, Bei and Sentana (2023), respectively.

Proof of Proposition 1

The scores evaluated under the null will be

$$\begin{aligned}
\frac{\partial l_t}{\partial \varphi_M} &= \frac{y_t - \varphi_M}{\varphi_V}, \\
\frac{\partial l_t}{\partial \varphi_V} &= \frac{(y_t - \varphi_M)^2 - \varphi_V}{2\varphi_V}, \\
\frac{\partial l_t}{\partial \vartheta_1} &= \frac{\partial l_t}{\partial \vartheta_2} = \frac{(y_t - \varphi_M)(y_{t-1} - \varphi_M)}{\varphi_V} \quad \text{and} \\
\frac{\partial l_t}{\partial \vartheta_3} &= \frac{\partial l_t}{\partial \vartheta_4} = \frac{(y_t - \varphi_M)(y_{t-4} - \varphi_M)}{\varphi_V},
\end{aligned}$$

which immediately imply (3) and (4), thereby confirming that the nullity of the information matrix is 2.

Consider the reparametrization from the original set of parameters $\boldsymbol{\varrho} = (\varphi_M, \varphi_V, \vartheta_1, \vartheta_2, \vartheta_3, \vartheta_4)'$ to a different set $\boldsymbol{\rho} = (\phi_M, \phi_V, \theta_{i1}, \theta_{i2}, \theta_{u1}, \theta_{u2})'$ defined by

$$\begin{aligned}
\varphi_M &= \phi_M, \\
\varphi_V &= \phi_V, \\
\vartheta_1 &= \theta_{i1} - \theta_{u1}, \\
\vartheta_2 &= \theta_{u1}, \\
\vartheta_3 &= \theta_{i2} - \theta_{u2} \quad \text{and} \\
\vartheta_4 &= \theta_{u2}.
\end{aligned}$$

The corresponding first-order derivatives under the equivalent hypothesis $H_0 : \theta_{i1} = \theta_{u1} = \theta_{i2} = \theta_{u2} = 0$ are

$$\begin{aligned}
\frac{\partial l_t}{\partial \theta_{i1}} &= \frac{(y_t - \phi_M)(y_{t-1} - \phi_M)}{\phi_V}, \\
\frac{\partial l_t}{\partial \theta_{i2}} &= \frac{(y_t - \phi_M)(y_{t-4} - \phi_M)}{\phi_V}
\end{aligned}$$

$$\frac{\partial l_t}{\partial \theta_{u1}} = 0, \quad \text{and}$$
$$\frac{\partial l_t}{\partial \theta_{u2}} = 0.$$

In turn, the second-order derivatives involving θ_{u1} and θ_{u2} are given in (5), (6) and (7).

Let $\theta_{u1} = \eta \upsilon_1$ and $\theta_{u2} = \eta \upsilon_2$ with $\upsilon_1^2 + \upsilon_2^2 = 1$ and consider the simplified null hypothesis $H_0 : \eta = 0$ for fixed values of υ_1 and υ_2. In this context, the only relevant quantity associated to η is

$$\frac{\partial^2 l_t}{\partial \eta^2} = 2\upsilon_1^2 \frac{(y_t - \phi_M)(y_{t-2} - \phi_M)}{\phi_V} + 2\upsilon_2^2 \frac{(y_t - \phi_M)(y_{t-8} - \phi_M)}{\phi_V}.$$

Moreover, given that

$$E\left(\frac{\partial l_t}{\partial \boldsymbol{\phi}} \frac{\partial l_t}{\partial \boldsymbol{\theta}_i'}\right) = \mathbf{0} \quad \text{and} \quad E\left[\frac{\partial l_t}{\partial \boldsymbol{\phi}} vech'\left(\frac{\partial^2 l_t}{\partial \boldsymbol{\theta}_u \partial \boldsymbol{\theta}_u'}\right)\right] = \mathbf{0}$$

under the null, we can ignore the parameter uncertainty in estimating ϕ_M and ϕ_V, at least asymptotically.

In this context, the GET statistic will be given by

$$GET_T = \sup_{||\boldsymbol{\upsilon}||=1} T^{-1}[S'_{\boldsymbol{\theta}_i}(\tilde{\boldsymbol{\phi}}, \mathbf{0}), \mathcal{H}_\eta(\tilde{\boldsymbol{\phi}}, 0, \boldsymbol{\upsilon})]\mathcal{V}^{-1}(\tilde{\boldsymbol{\phi}}, \boldsymbol{\upsilon})[S'_{\boldsymbol{\theta}_i}(\tilde{\boldsymbol{\phi}}, \mathbf{0}), \mathcal{H}_\eta(\tilde{\boldsymbol{\phi}}, 0, \boldsymbol{\upsilon})]', \tag{A1}$$

where

$$S_{\boldsymbol{\theta}_i}(\boldsymbol{\rho}) = [S_{\theta_{i1}}(\boldsymbol{\rho}), S_{\theta_{i2}}(\boldsymbol{\rho})]',$$
$$\mathcal{H}_\eta(\boldsymbol{\phi}, \eta, \boldsymbol{\upsilon}) = \sum_{t=1}^{T} \frac{\partial^2 l_t}{\partial \eta^2}, \quad \text{and}$$
$$\mathcal{V}(\boldsymbol{\phi}, \boldsymbol{\upsilon}) = Var\{T^{-1/2}[S'_{\boldsymbol{\theta}_i}(\boldsymbol{\phi}, \mathbf{0}), \mathcal{H}_\eta(\boldsymbol{\phi}, 0, \boldsymbol{\upsilon})]'|\boldsymbol{\phi}, \mathbf{0}\}.$$

Importantly, the sup-type statistic (A1) can be computed analytically in this example. Specifically, straightforward algebra shows that

$$GET_T = T \sup_{||\boldsymbol{\upsilon}||\neq 0} \left\{\tilde{r}_1^2 + \tilde{r}_4^2 + \frac{(\upsilon_1^2\tilde{r}_2 + \upsilon_2^2\tilde{r}_8)^2}{\upsilon_1^4 + \upsilon_2^4} 1[\upsilon_1^2\tilde{r}_2 + \upsilon_2^2\tilde{r}_8 \geq 0]\right\}.$$

In addition, we can show that the value of $\boldsymbol{\upsilon}$ that maximizes the above expression will be proportional to the vector

$$(\sqrt{\tilde{r}_2 \mathbf{1}\,[\tilde{r}_2 \geq 0]}, \sqrt{\tilde{r}_8 \mathbf{1}\,[\tilde{r}_8 \geq 0]})$$

if $\tilde{r}_2 \geq 0$ or $\tilde{r}_8 \geq 0$, and to $(1, 1)$ otherwise, which confirms (8). □

Proof of Proposition 2

We can use expression (B4) in Appendix B.1 to compute the spectral approximation to the log-likelihood function of model (14) with $g_{yy}(\omega; \boldsymbol{\theta})$ given in (B8) and $\boldsymbol{\theta} = (\sigma_f^2, \sigma_v^2, \alpha, \psi_1, \psi_2)'$.

To simplify the notation, let us define the vector

$$\mathbf{C}(\omega) = \frac{2\pi I_{yy}(\omega) - g_{yy}(\omega;\boldsymbol{\gamma})}{g_{yy}^2(\omega;\boldsymbol{\gamma})} \begin{bmatrix} 1 \\ \cos(\omega) \\ \cos(2\omega) \\ \cos(3\omega) \\ \cos(4\omega) \end{bmatrix},$$

which corresponds to the contribution of frequency ω to the spectral score of an MA(4) model parametrized in terms of its unconditional variance and first four autocovariances, say $\boldsymbol{\gamma} = (\gamma_0, \gamma_1, \gamma_2, \gamma_3, \gamma_4)$, evaluated at $\gamma_2 = \gamma_3 = \gamma_4 = 0$, as can be immediately seen from (B5). Importantly, $g_{yy}(\omega;\boldsymbol{\theta}) = g_{yy}(\omega;\boldsymbol{\gamma})$ for all ω under the locally equivalent null hypotheses $H_0 : \alpha = \psi_1 = \psi_2 = 0$ and $H_0 : \gamma_2 = \gamma_3 = \gamma_4 = 0$ when both σ_f^2 and σ_v^2 are strictly positive.

Therefore, we can write the contribution of frequency ω to the spectral score as

$$\begin{aligned}
\frac{\partial l_t}{\partial \sigma_f^2} &= \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \end{pmatrix} \mathbf{C}(\omega) \\
\frac{\partial l_t}{\partial \sigma_v^2} &= \begin{pmatrix} 2 & -2 & 0 & 0 & 0 \end{pmatrix} \mathbf{C}(\omega) \\
\frac{\partial l_t}{\partial \alpha} &= \begin{pmatrix} -2\sigma_f^2 & 4\sigma_f^2 & -2\sigma_f^2 & 0 & 0 \end{pmatrix} \mathbf{C}(\omega) \\
\frac{\partial l_t}{\partial \psi_1} &= \begin{pmatrix} 0 & 2\sigma_f^2 & 0 & 0 & 0 \end{pmatrix} \mathbf{C}(\omega) \\
\frac{\partial l_t}{\partial \psi_2} &= \begin{pmatrix} 0 & 0 & 2\sigma_f^2 & 0 & 0 \end{pmatrix} \mathbf{C}(\omega)
\end{aligned}$$

We can immediately notice that the last two elements of this score belong to the linear span of the first three.

To isolate those singularities, we conduct a two-step reparametrization as follows. First, we consider

$$\begin{aligned}
\sigma_f^2 &= \sigma_f^{2\diamond} - 2\sigma_f^{2\diamond}\psi_1^\diamond + \sigma_f^{2\diamond}(\psi_1^\diamond)^2 - 2\sigma_f^{2\diamond}\psi_2^\diamond, \\
\sigma_v^2 &= \sigma_v^{2\diamond} + \sigma_f^{2\diamond}\psi_1^\diamond + 2\sigma_f^{2\diamond}\psi_2^\diamond, \\
\alpha &= \alpha^\diamond + \frac{\sigma_f^{2\diamond}}{\sigma_v^{2\diamond}}(\psi_1^\diamond)^2 + \frac{\sigma_f^{2\diamond}}{\sigma_v^{2\diamond}}\psi_2^\diamond, \\
\psi_1 &= \psi_1^\diamond, \\
\psi_2 &= \psi_2^\diamond,
\end{aligned}$$

and then

$$
\begin{aligned}
\sigma_f^{2\diamond} &= \sigma_f^{2\dagger} - \sigma_v^{2\dagger}(\psi_1^\dagger)^3, \\
\sigma_v^{2\diamond} &= \sigma_f^{2\dagger} + \frac{1}{2}\sigma_v^{2\dagger}(\psi_1^\dagger)^3, \\
\alpha^\diamond &= \alpha^\dagger - \frac{(\sigma_f^{2\dagger})^2 + 2\sigma_v^{2\dagger}\sigma_f^{2\dagger}}{2(\sigma_v^{2\dagger})^2}(\psi_1^\dagger)^3, \\
\psi_1^\diamond &= \psi_1^\dagger, \\
\psi_2^\diamond &= \psi_2^\dagger - \frac{(\psi_1^\dagger)^2}{2}.
\end{aligned}
$$

As a consequence, the relevant derivatives after reparametrization will be

$$
\begin{aligned}
\frac{\partial l}{\partial \sigma_f^{2\dagger}} &= \frac{\partial l}{\partial \sigma_f^2} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \end{pmatrix}\mathbf{C}(\omega) \\
\frac{\partial l}{\partial \sigma_v^{2\dagger}} &= \frac{\partial l}{\partial \sigma_v^2} = \begin{pmatrix} 2 & -2 & 0 & 0 & 0 \end{pmatrix}\mathbf{C}(\omega) \\
\frac{\partial l}{\partial \alpha^\dagger} &= \frac{\partial l}{\partial \alpha} = \begin{pmatrix} -2\sigma_f^2 & 4\sigma_f^2 & -2\sigma_f^2 & 0 & 0 \end{pmatrix}\mathbf{C}(\omega)
\end{aligned}
$$

and

$$
\frac{\partial l}{\partial \psi_1^\dagger} = \frac{\partial l}{\partial \psi_2'} = \frac{\partial^2 l}{(\partial \psi_1')^2} = \frac{\partial^3 l}{(\partial \psi_1')^3} = 0.
$$

In addition, straightforward calculations deliver

$$
\frac{\partial^2 l}{(\partial \psi_2^\dagger)^2} = \mathbf{C}'(\omega)\begin{bmatrix} 2\sigma_f^2(\sigma_v^2 - 2\sigma_f^2)/\sigma_v^2 \\ 8\sigma_f^4/\sigma_v^2 \\ -8\sigma_f^2 \\ -4\sigma_f^4/\sigma_v^2 \\ 4\sigma_f^2 \end{bmatrix},
$$

$$
\frac{\partial^2 l}{\partial \psi_1^\dagger \partial \psi_2^\dagger} = \mathbf{C}'(\omega)\begin{bmatrix} -2\sigma_f^4/\sigma_v^2 \\ -2\sigma_f^2(\sigma_v^2 - 2\sigma_f^2)/\sigma_v^2 \\ -2\sigma_f^2(\sigma_f^2 + 2\sigma_v^2)/\sigma_v^2 \\ 4\sigma_f^2 \\ 0 \end{bmatrix},
$$

and

$$
\frac{\partial^4 l}{(\partial \psi_1^\dagger)^4} = \mathbf{C}'(\omega)\begin{bmatrix} 6\sigma_f^2(4\sigma_f^4 + 14\sigma_f^2\sigma_v^2 + 9\sigma_v^4)/\sigma_v^4 \\ -24\sigma_f^2(2\sigma_f^4 + 7\sigma_f^2\sigma_v^2 + 2\sigma_v^4)/\sigma_v^4 \\ 24\sigma_f^2(\sigma_f^4 + 4\sigma_f^2\sigma_v^2 + 2\sigma_v^4)/\sigma_v^4 \\ -12\sigma_f^4/\sigma_v^4 \\ -12\sigma_f^2 \end{bmatrix}
$$

Thus, in the notation of Theorem 2 in Amengual, Bei and Sentana (2023), we have

$$
\mathcal{S}_{\boldsymbol{\theta},T} = \sum_{t=1}^{T} \mathbf{C}(\omega_t),
$$

where $\omega_t = 2\pi t/T$ (for $0 = 1, \ldots, T-1$) are the usual Fourier frequencies,

$$\lambda_{\boldsymbol{\theta}}(\boldsymbol{\theta}) = \begin{pmatrix} -2\sigma_f^2\alpha - 4\sigma_f^2(\psi_2)^2 - \frac{2\sigma_f^2(\sigma_f^2+2\sigma_v^2)}{\sigma_v^2}\psi_1\psi_2 + \frac{\sigma_f^2(\sigma_f^4+4\sigma_f^2\sigma_v^2+2\sigma_v^4)}{\sigma_v^4}\psi_1^4 \\ -\frac{2\sigma_f^4}{\sigma_v^2}\psi_2^2 + 4\sigma_f^2\psi_1\psi_2 - \frac{\sigma_f^4}{2\sigma_v^4}\psi_1^4 \\ 2\sigma_f^2\psi_2^2 - 12\sigma_f^2\psi_1^4 \end{pmatrix},$$

and

$$\Lambda_T = \left\{\sqrt{T}\lambda_{\boldsymbol{\theta}}\right\} \to \Lambda = \mathbb{R}^3.$$

The interpretation of $\mathbf{C}(\omega_t)$ as a spectral log-likelihood score allows us to obtain the asymptotic variance of $\mathcal{S}_{\boldsymbol{\theta},n}$ suitably scaled by $T^{-\frac{1}{2}}$ using expression (B6). Specifically, if we partition the autocovariances into $\boldsymbol{\gamma}_n = (\gamma_0, \gamma_1)$ and $\boldsymbol{\gamma}_a = (\gamma_2, \gamma_3, \gamma_4)$, then we will have that

$$T^{-\frac{1}{2}}\mathcal{S}_{\boldsymbol{\theta},T} \xrightarrow{d} N\left[\begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \begin{pmatrix} \mathcal{I}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_n} & \mathcal{I}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a} \\ \mathcal{I}'_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a} & \mathcal{I}_{\boldsymbol{\gamma}_a\boldsymbol{\gamma}_a} \end{pmatrix}\right], \tag{A2}$$

with the different elements evaluated at $\boldsymbol{\gamma}_a = \mathbf{0}$.

On this basis, we obtain

$$\begin{aligned} GET_T &= \left[T^{-\frac{1}{2}}\sum_{t=1}^{T}\mathbf{C}'(\omega_t)\right]\begin{pmatrix} -\mathcal{I}^{-1}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_n}\mathcal{I}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a} \\ \mathbf{I}_3 \end{pmatrix}\left(\mathcal{I}_{\boldsymbol{\gamma}_a\boldsymbol{\gamma}_a} - \mathcal{I}'_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a}\mathcal{I}^{-1}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_n}\mathcal{I}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a}\right)^{-1} \\ &\quad \times\begin{pmatrix} -\mathcal{I}'_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_a}\mathcal{I}^{-1}_{\boldsymbol{\gamma}_n\boldsymbol{\gamma}_n} & \mathbf{I}_3 \end{pmatrix}\left[T^{-\frac{1}{2}}\sum_{t=1}^{T}\mathbf{C}'(\omega_t)\right]. \end{aligned} \tag{A3}$$

Given that the restricted MLEs for σ_f^2 and σ_v^2 are such that in large samples the estimated model will perfectly match the sample variance and first autocovariance of y_t with probability approaching 1, the first two components of $\mathcal{S}_{\boldsymbol{\theta},T}$ evaluated at $\tilde{\boldsymbol{\theta}}_T$ will be 0, which in turn implies that GET_T is effectively testing that the second, third and fourth autocovariances of y_t are simultaneously 0 on the basis of their sample counterparts, but taking into account the sampling uncertainty in estimating those autocovariances when the true process is the local level model (9)-(12).

Finally, applying the delta method to go from the autocovariances γ_j $(j = 0, ..., 4)$ to the autocorrelations $\rho_j = \gamma_j/\gamma_0$ $(j = 1, ..., 4)$ delivers the expressions in the statement of the proposition. □

B Additional results

B.1 Maximum likelihood estimation in the frequency domain

Henceforth, we assume that y_t is a covariance stationary series, which may require taking first or seasonal differences of the observations, as in the examples in sections 2 and 3.

Let

$$I_{yy}(\omega) = \frac{1}{2\pi T}\sum_{t=1}^{T}\sum_{s=1}^{T}(y_t-\mu)(y_s-\mu)e^{-i(t-s)\omega}$$

denote the periodogram of y_t and $\omega_j = 2\pi j/T$ (for $j = 0, \dots T-1$) the usual Fourier frequencies. If we assume that the spectral density $g_{yy}(\omega;\boldsymbol{\theta})$ is not zero at any of those frequencies, the so-called Whittle (1962)'s (discrete) spectral approximation to the log-likelihood function is

$$-\frac{T}{2}\ln(2\pi) - \frac{1}{2}\sum_{j=0}^{T-1}\ln|g_{yy}(\omega_j;\boldsymbol{\theta})| - \frac{1}{2}\sum_{j=0}^{T-1}\frac{2\pi I_{yy}(\omega_j)}{g_{yy}(\omega_j;\boldsymbol{\theta})}. \tag{B4}$$

The MLE of μ, which only enters through $I_{yy}(\omega)$, is the sample mean, so in what follows we focus on demeaned variables. In turn, the score with respect to all the remaining parameters is

$$\frac{\partial l_t}{\partial\boldsymbol{\theta}} = \frac{1}{2}\sum_{j=0}^{T-1}\frac{\partial g_{yy}(\omega_j;\boldsymbol{\theta})}{\partial\boldsymbol{\theta}}M(\omega_j;\boldsymbol{\theta})m(\omega_j;\boldsymbol{\theta}), \tag{B5}$$

$$m(\omega;\boldsymbol{\theta}) = 2\pi I_{yy}(\omega) - g_{yy}(\omega;\boldsymbol{\theta}),$$

$$M(\omega;\boldsymbol{\theta}) = g_{yy}^{-2}(\omega;\boldsymbol{\theta}).$$

The information matrix is block diagonal between μ and the elements of $\boldsymbol{\theta}$, with the (1,1)-element being $g_{yy}(0)$ and the (2,2)-block

$$\mathbf{Q}(\boldsymbol{\theta}) = \frac{1}{4\pi}\int_{-\pi}^{\pi}\frac{\partial g_{yy}(\omega_j;\boldsymbol{\theta})}{\partial\boldsymbol{\theta}}M(\omega;\boldsymbol{\theta})\left\{\frac{\partial g_{yy}(\omega_j;\boldsymbol{\theta})}{\partial\boldsymbol{\theta}}\right\}^{*}d\omega, \tag{B6}$$

where * denotes the conjugate transpose of a matrix. A consistent estimator will be provided either by the outer product of the score or by

$$\boldsymbol{\Phi}(\boldsymbol{\theta}) = \frac{1}{2}\sum_{j=0}^{T-1}\frac{\partial g_{yy}(\omega_j;\boldsymbol{\theta})}{\partial\boldsymbol{\theta}}M(\omega_j;\boldsymbol{\theta})\left\{\frac{\partial g_{yy}(\omega_j;\boldsymbol{\theta})}{\partial\boldsymbol{\theta}}\right\}^{*}. \tag{B7}$$

In fact, by selecting an artificially large value for T in (B7), one can approximate (B6) to any desired degree of accuracy. In addition, the univariate nature of y_t implies that both $g_{yy}(\omega_j;\boldsymbol{\theta})$ and its derivatives are real.

Formal results showing the strong consistency and asymptotic normality of the resulting ML estimators of dynamic latent variable models under suitable regularity conditions were provided by Dunsmuir (1979), who generalized earlier results for VARMA models by Dunsmuir and Hannan (1976). These authors also show the asymptotic equivalence between time and frequency domain MLEs.

B.2 The autocorrelation structure of the UCARIMA model

We can derive the autocovariance structure of the $y_t = \Delta x_t$ by the usual inverse Fourier transformation

$$\gamma_{yy}(k) = cov(y_t, y_{t-k}) = \int_{-\pi}^{\pi} e^{i\omega k} g_{yy}(\omega) d\omega$$

after exploiting that $g_{yy}(\omega)$ is the sum of the spectral densities of the signal and noise components, $s_t = \Delta z_t$ and $n_t = \Delta u_t$, respectively, which are cross-sectionally uncorrelated at all leads and lags. Specifically, we know that

$$y_t = \frac{1}{1 - \psi_1 L - \psi_2 L^2} f_t + \frac{1 - L}{1 - \alpha L} u_t = s_t + n_t,$$

where the first component, s_t, is an AR(2) process while the second component, n_t, is an ARMA(1,1) with a unit root on the MA part.

Thus,

$$\begin{aligned} g_{yy}(\omega; \boldsymbol{\theta}) &= = g_{ss}(\omega; \boldsymbol{\theta}) + g_{nn}(\omega; \boldsymbol{\theta}) \\ &= \frac{\sigma_f^2}{(1 - \psi_1 e^{-i\omega} - \psi_2 e^{-2i\omega})(1 - \psi_1 e^{i\omega} - \psi_2 e^{2i\omega})} + \frac{(1 - e^{-i\omega})(1 - e^{i\omega})\sigma_v^2}{(1 - \alpha e^{-i\omega})(1 - \alpha e^{i\omega})} \\ &= \frac{\sigma_f^2}{(1 + \psi_1^2 + \psi_2^2) - 2\psi_1(1 - \psi_2)\cos\omega - 2\psi_2 \cos 2\omega} + \frac{2(1 - \cos\omega)\sigma_v^2}{(1 + \alpha^2) - 2\alpha\cos\omega}. \quad \text{(B8)} \end{aligned}$$

However, the expressions for $\gamma_{yy}(k)$ are somewhat easier to obtain in the time domain as the sum of the autocovariances of the two underlying components.

The autocovariances of the AR(2) process for the signal are given by the usual Yule-Walker recursion

$$\gamma_{ss}(k) = \psi_1 \gamma_s(k - 1) + \psi_2 \gamma_s(k - 2), \quad \text{(B9)}$$

with initial conditions

$$\begin{aligned} \gamma_{ss}(0) &= \left(\frac{1 - \psi_2}{1 + \psi_2}\right) \frac{\sigma_f^2}{(1 - \psi_2)^2 - \psi_1^2}, \\ \gamma_{ss}(1) &= \left(\frac{\psi_1}{1 + \psi_2}\right) \frac{\sigma_f^2}{(1 - \psi_2)^2 - \psi_1^2}, \end{aligned}$$

which yields

$$\begin{aligned} \gamma_{ss}(2) &= \frac{\psi_1^2 + \psi_2(1 - \psi_2)}{1 - \psi_2} \gamma_s(0), \\ \gamma_{ss}(3) &= \frac{\psi_1[\psi_1^2 + \psi_2(2 - \psi_2)]}{1 - \psi_2} \gamma_s(0), \quad \text{and} \\ \gamma_{ss}(4) &= \frac{\psi_1[\psi_1^3 + \psi_1\psi_2(3 - \psi_2)] + \psi_2^2(1 - \psi_2)}{1 - \psi_2} \gamma_s(0). \end{aligned}$$

To find the solution for general k, it is convenient to find the roots of the characteristic equation (B9), which are given by

$$\begin{aligned}\delta_1 &= \frac{1}{2}\psi_1 + \frac{1}{2}\sqrt{\psi_1^2 + 4\psi_2}\\ \delta_2 &= \frac{1}{2}\psi_1 - \frac{1}{2}\sqrt{\psi_1^2 + 4\psi_2}\end{aligned}$$

When the roots are different (real or complex), the autocorrelation of order k will be given by

$$\gamma_{ss}(k) = \frac{\delta_1^{k+1}(1-\delta_2^2) - \delta_2^{k+1}(1-\delta_1^2)}{(\delta_1 - \delta_2)(1+\delta_1\delta_2)}\gamma_s(0).$$

Applying L'Hôpital's rule, this simplifies to

$$\gamma_{ss}(k) = \left[1 + k\frac{(1-\delta^2)}{(1+\delta^2)}\right]\delta^k\gamma_s(0)$$

when the two roots are equal, which happens for $\psi_2 = -\psi_1^2/4$ (see e.g. Fuller (1995)).

In turn, the autocovariances of the ARMA(1,1) process for the noise will be

$$\begin{aligned}\gamma_{nn}(0) &= \sigma_v^2\left[1 + \frac{(\alpha-1)^2}{1-\alpha^2}\right] = \frac{2\sigma_v^2}{\alpha+1},\\ \gamma_{nn}(1) &= \sigma_v^2\left[(\alpha-1) + \frac{(\alpha-1)^2\alpha}{1-\alpha^2}\right] = \frac{(\alpha-1)\sigma_v^2}{\alpha+1},\end{aligned}$$

and

$$\gamma_{nn}(k) = \frac{\alpha^{k-1}(\alpha-1)\sigma_v^2}{\alpha+1}.$$

Finally,

$$\gamma_{yy}(k) = \gamma_{ss}(k) + \gamma_{nn}(k).$$

C Tables and figures

Table 1: Monte Carlo rejection rates (in %) under the alternative hypotheses for the white noise versus multiplicative seasonal AR test.

	Alternative hypotheses					
	H_{a_1}			H_{a_2}		
	1%	5%	10%	1%	5%	10%
	$T = 100$					
GET	29.3	48.2	57.7	26.5	51.9	64.2
LR	15.3	33.3	44.2	20.8	46.6	59.3
LM-AR(1)	16.8	32.3	40.9	3.1	10.8	17.0
LM-SAR(4)	15.4	31.7	41.6	3.8	11.6	18.4
MT	27.0	44.3	53.3	22.0	43.7	55.0
	$T = 400$					
GET	87.4	94.7	96.8	92.1	97.9	99.1
LR	81.0	92.3	95.3	92.2	98.0	99.1
LM-AR(1)	60.1	76.9	84.1	3.7	10.8	17.7
LM-SAR(4)	57.3	78.2	86.0	4.9	13.7	22.1
MT	85.3	93.3	95.8	89.6	96.7	98.2

Notes: Results based on 10,000 samples. The mean and variance parameters φ_M and φ_V are estimated under the null using their sample analogs. GET is computed as defined in section 2.1. DGPs: the true unconditional mean and the variance of the innovations are set to 0 and 1, respectively, under both the null and alternative hypotheses. As for the alternative hypotheses,

$$(1 - 0.1L - 0.1L^2 - 0.1L^3 - 0.1L^4)y_t = \varepsilon_t \quad (H_{a_1})$$

and

$$(1 - 0.4L)(1 + 0.4L)(1 - 0.4L^4)(1 + 0.4L^4)y_t = \varepsilon_t \quad (H_{a_2}).$$

LM-AR(1) and LM-SAR(4) denote the Lagrange multiplier tests based on the score of an AR(1) and a seasonal SAR(4), respectively. MT refers to the two-sided version of GET. Finite sample critical values are computed using a parametric bootstrap procedure.

Table 2: Monte Carlo rejection rates (in %) under the null and alternative hypotheses for the local level model versus the UCARIMA model (14) test.

	Null hypothesis			Alternative hypotheses H_{a_1}			Alternative hypotheses H_{a_2}		
	1%	5%	10%	1%	5%	10%	1%	5%	10%
$T = 100$									
GET	1.1	5.3	10.4	5.8	21.7	35.5	8.6	22.7	33.1
LR	1.6	5.4	10.3	13.5	31.0	44.0	7.3	18.3	27.8
2^{nd} autocorrelation	1.2	5.3	10.2	4.4	13.7	23.6	2.2	7.2	13.7
2^{nd} & 3^{rd} autocorrelation	1.3	5.5	10.0	5.3	19.7	31.4	2.1	7.8	13.2
$T = 400$									
GET	1.0	5.0	10.0	67.0	87.0	93.7	51.5	72.7	83.0
LR	1.0	5.2	10.2	77.3	91.0	95.4	31.9	50.6	63.0
2^{nd} autocorrelation	1.0	5.0	10.0	30.3	56.0	69.1	2.4	8.9	15.3
2^{nd} & 3^{rd} autocorrelation	1.0	5.0	10.1	52.8	76.8	85.8	2.5	8.4	14.2

Notes: Results based on 10,000 samples. The local level parameters σ_f^2 and σ_u^2 are estimated under the null. GET is computed as defined in section 3.1. DGPs: We simulate Gaussian shocks with $\sigma_f^2 = 1$ and $\sigma_v^2 = 0.5$ under both the null and the alternatives. Alternative hypotheses:

$$\left.\begin{array}{c}(1+0.5L+0.4L^2)\Delta z_t = f_t \\ (1-0.5L)u_t = v_t\end{array}\right\} \; (H_{a_1}) \quad \text{and} \quad \left.\begin{array}{c}(1-0.1L+0.5L^2)\Delta z_t = f_t \\ (1+0.5L)u_t = v_t\end{array}\right\} \; (H_{a_2}).$$

2^{nd} autocorrelation (2^{nd} & 3^{rd} autocorrelation) denote the moment test of no second-order (no second- or third-order) serial correlation in y_t. Finite sample critical values are computed using a parametric bootstrap procedure.

Figure 1: Alignment of GET and LR under the null under null for the white noise versus multiplicative seasonal AR test.

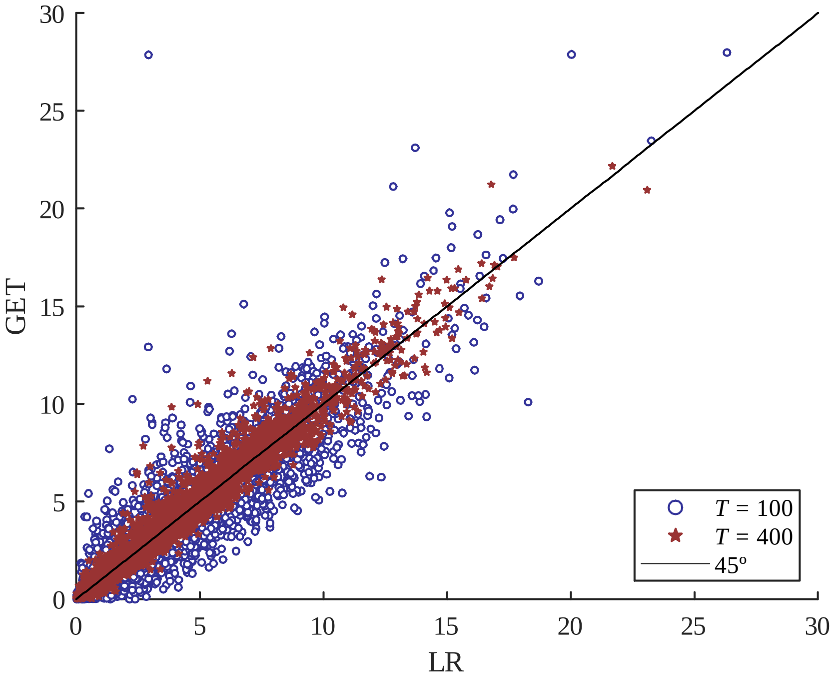

Notes: Scatter plots of the GET_T and LR_T test statistics. Results based on $10,000$ simulated samples of size T of $y \sim i.i.d.$ Gaussian. GET is computed as explain in section 2.1. The true mean and variance of the simulated data are set to 0 and 1, and the elements of $\boldsymbol{\varphi}$ are estimated using the sample mean and variance, respectively.

Figure 2: Alignment of GET and LR under the null for the local level model versus the UCARIMA model (14).

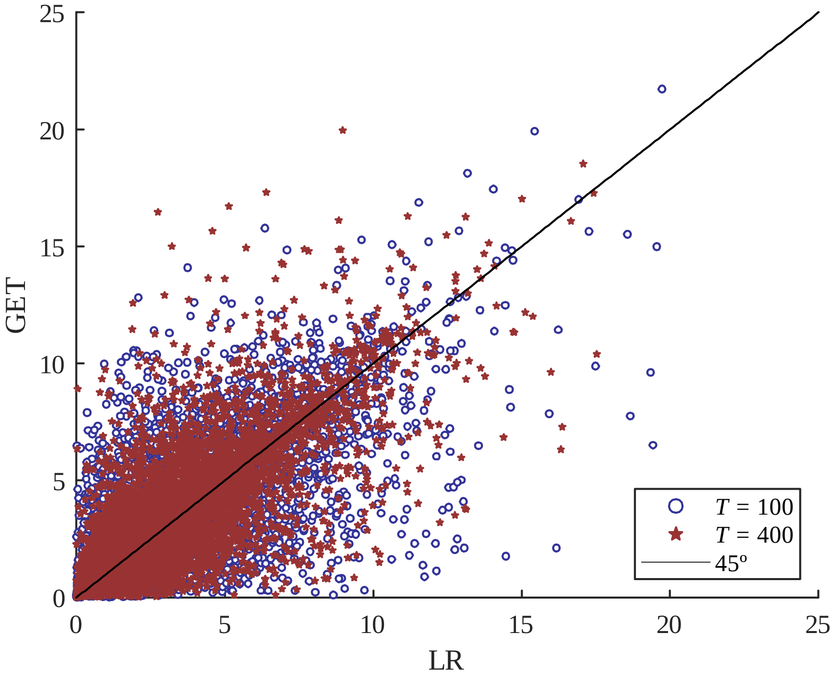

Notes: Scatter plots of the GET_T and LR_T test statistics. Results based on $10,000$ simulated samples of size T of the model under the null with Gaussian shocks with $\sigma_f^2 = 1$ and $\sigma_v^2 = 0.5$. GET is computed as explained in section 3.1.