

Semiparametric Lag Dependent Functions

Erik Lindström

Division of Mathematical Statistics
Centre for Mathematical Sciences
Lund University, Box 118
SE-22100 Lund, Sweden

Abstract

This paper introduces a suite of tests for linear and general non-linear serial dependence. The problem is complex as the number of variations of realistic non-linear alternatives is very large.

The alternative model is defined in terms of penalized truncated polynomial splines, making the approximation capable of accurately approximating a large class of linear and non-linear processes.

The asymptotic distribution for the test statistic is derived, and we show using Monte Carlo simulations that one test is equivalent to the Ljung-Box test, other linear tests corresponds to ordinary or partial sample autocorrelation while the non-linear versions are capable of detecting non-linear effects, even when other tests fail to do so.

Mathematics Subject Classification: 62M10, 62M07, 60G10

Keywords: time series, hypothesis testing, Stationary processes

1 Introduction

Dependence tests are the starting point for most statistical studies. Popular linear tests for serial dependence autocorrelation are [19, 22, 24, 20]. It is often necessary to account for linear dependence prior to testing for non-linear dependence.

Non-linear tests for dependence include extensions of the Ljung-Box test, see [21], Volterra expansions, see e.g. [13, 30] and generalizations based on polynomial and exponential functions, see [4, 8] and references therein. Series expansions are rarely optimal for global approximation of functions, cf. Runge's phenomenon, which is why threshold models, cf. [28, 29] have been popular, see [10, 15]. Even more general alternatives are Neural Net based tests, see [27] and Lag Dependent Functions (LDFs), see [23, 18], that is based local

polynomial regression. Yet other alternative are spectral tests and tests based on the probability integral transform, cf. [25, 17].

We derive linear, non-linear and strictly non-linear tests for serial dependence. Penalized splines are used to approximate a general time series model. This class is often used for non-linear regression problems as splines can accurately approximate a huge class of functions without the limitations of polynomial expansions, cf. [26, 7]. The explanatory properties of penalized splines are closely related to local polynomial regression, see [26, 7, 16], making the proposed tests similar in spirit to [23], hence the title of the paper.

The advantage of the proposed tests in this paper compared to the tests in [23] is twofold: The first advantage is that the scope of the proposed method is broader as the test is not restricted to homoscedastic data. The second advantage is computational efficiency, as only a single regression is being run. We also present approximate confidence bands, thereby eliminating the need for bootstrap methods.

The remainder of the paper is organized as follows. Section 2 explains how semi-parametric methods can be used to approximate a general, alternative model and introduces several tests. These are evaluated in Section 3 using Monte Carlo simulations and Section 4 concludes.

2 Dependence testing

The sequence of observations is throughout the paper assumed be generated from an ergodic and strictly stationary process $\{Y_n\}_{n=1}^N$. This section will derive tests for linear and non-linear dependence.

It is practically impossible to use general non-linear autoregressive model for testing due to the "curse of dimensionality". A commonly used approximation is the additive model, see [9, 26], that approximates a K -dimensional function with a sum of 1-dimensional functions, depending of lagged values of the process, or cross-terms between lagged values.

Splines can approximate almost arbitrary functions. Regression splines are easy to use but will often overfit data. Overfitting is avoided by using penalized splines. We used a q th order truncated polynomial basis spline

$$Y_n = \psi_0 + \sum_{k=1}^K \left(\psi_k Y_{n-k} + \sum_{j=2}^q \alpha_{kj} Y_{n-k}^j + \sum_{l=1}^p \beta_{kl} (Y_{n-k} - \kappa_l)_+^q \right) + \xi_n, \quad (1)$$

where $\text{Var}(\xi) = \Sigma$, p is the number of knots and $\{\kappa_l\}_{l=1..p}$ are the positions of the knots. The knots are placed according to the percentiles of the data in this paper. It has been claimed that the number of knots is unimportant as long as then number is large, see [26] but recent research has found that this is not entirely, and that penalized splines are asymptotically equivalent to either

regression splines or smoothing splines, see [5]. They found that it is optimal to let the number of knots grow slower than the number of observations for penalized splines defined in terms of truncated polynomial splines.

All spline models in this paper can be written in matrix notation

$$Y = X\psi + B_\alpha\alpha + B_\beta\beta + \xi = Z\theta + \xi, \quad (2)$$

where X is the linear part, B_α contains higher order polynomial terms and B_β contains all the spline terms from the basis expansion. Penalized splines differ from regression splines as the spline coefficients are restricted. The optimization problem for estimating penalized parameters is given by

$$\hat{\theta} = \arg \min_{\theta} (Y - Z\theta)^T \Sigma^{-1} (Y - Z\theta), \quad (3)$$

$$\text{subject to } \theta^T D \theta \leq c, \quad (4)$$

where c is some constant and $D = \text{diag}(\mathbf{0}_{1+qK}, \mathbf{1}_{pK})$ is a diagonal matrix with ones on the diagonal elements if the parameter is to be penalized and zero otherwise. The equivalent optimization problem can be written as

$$\hat{\theta} = \arg \min_{\theta} ((Y - Z\theta)^T \Sigma^{-1} (Y - Z\theta) + \gamma^2 \theta^T D \theta) \quad (5)$$

where $\gamma^2 = \gamma^2(c)$ is a constant tuning the size of the penalty. The estimate, written in matrix notation, is then given by

$$\hat{\theta} = (Z^T \Sigma^{-1} Z + \gamma^2 D)^{-1} Z^T \Sigma^{-1} Y, \quad (6)$$

while the associated projection matrix S_γ is given by

$$S_\gamma = Z (Z^T \Sigma^{-1} Z + \gamma^2 D)^{-1} Z^T \Sigma^{-1}. \quad (7)$$

2.1 Hypothesis testing

F -tests are used for hypothesis testing when the model is defined in terms of regression splines. We argue, cf. [9, 26] that these tests also can be used for penalized splines under certain conditions.

The standard F -test is derived for independent, homoscedastic innovations, being defined as ratios of χ^2 variables. The estimated innovations $\hat{\xi} = Y - Z\hat{\theta}$ must therefore be normalized. We define the *Normalized Residual Sum of Squares* as

$$NRSS = \hat{\xi}^T \Sigma^{-1} \hat{\xi} \xrightarrow{d} \chi^2(\cdot) \quad (8)$$

It follows that this quadratic form converge in distribution to a χ^2 distributed random variable. It can be shown, see [9], that the equivalent residual degrees of freedom is given by

$$df = N - 2 \text{trace}(S_\gamma) + \text{trace}(\Sigma S_\gamma^T \Sigma^{-1} S_\gamma), \quad (9)$$

where N is the number of observations.

Remark 2.1. *A computationally efficient implementation of Equation (9) is obtained by computing*

$$\Gamma = (Z^T \Sigma^{-1} Z + \gamma^2 D)^{-1} Z^T \Sigma^{-1} Z, \quad (10)$$

$$df = n - 2\text{trace}(\Gamma) + \text{trace}(\Gamma\Gamma). \quad (11)$$

This is shown by using $\text{trace}(AB) = \text{trace}(BA)$. The computational gain can be quite substantial, as the dimension of the Γ matrix is $d \times d$ where d is the dimension of the parameter vector, while the naive implementation would involve multiplications of $N \times N$ matrices, where N is the number of observations.

We define the test statistic analogously with the standard F -test, comparing the *Normalized Residual Sum of Squares* for the restricted model (*Res*) and the general model (*Gen*), normalized by the residual variance. The test statistic is given by

$$f = \frac{\left(\hat{\xi}_{Res}^T \Sigma^{-1} \hat{\xi}_{Res} - \hat{\xi}_{Gen}^T \Sigma^{-1} \hat{\xi}_{Gen} \right) / (df_{Res} - df_{Gen})}{\hat{\xi}_{Gen}^T \Sigma^{-1} \hat{\xi}_{Gen} / df_{Gen}}. \quad (12)$$

It has been proposed that this statistic is approximately $f \in F(df_{Res} - df_{Gen}, df_{Gen})$, see e.g. [9]. The F -test has been criticised for being inaccurate by [9] who suggested a second order correction of the test. Another study, [6], found that the distribution for asymptotic test statistic for linear splines is non-standard while the asymptotic distribution for quadratic (and higher order) splines is regular. This paper uses cubic splines ($q = 3$), as they are almost regular for large samples, see Section 3.1 for a simulation study.

The test is influenced by the smoothing factor γ^2 . The parameter is not identifiable under the null hypothesis for the tests that we propose, but it will still influence the power. Choosing γ^2 too small will increase the risk for overfitting, while choosing it too large will result in a polynomial model. Our simulations show, see Section 3.1, that the size of tests is correct over a large range of values for γ^2 as long as the truncated polynomial basis is based on higher order polynomials. A heuristic approach would be to pick γ^2 such that some smoothing is obtained, but not too much. Another approach is to estimate the parameter, using AIC, GVC or REML methods, see [11, 14].

A computationally expensive alternative to the F -test is the bootstrap methodology in [23] as penalized splines still are magnitudes faster than fitting a local polynomial regression. The bootstrap methodology in [12] are specifically designed for computationally efficient testing of penalized splines.

2.2 Tests for serial dependence

The framework can be used to derive many tests, including standard tests for autocorrelation, partial autocorrelation and portmanteau tests. We also

introduce several previously unknown tests; see Section 2.2.6, 2.2.8 and 2.2.9. The tests in Section 2.2.1-2.2.3 correspond to autocorrelation and generalizations, the partial dependence tests in Section 2.2.4-2.2.6 include partial autocorrelation while the portmanteau tests in Section 2.2.7-2.2.9 include Box-Pierce/Ljung-Box type tests and non-linear generalizations. The parameters in this section are the parameters defined in Equation (1).

2.2.1 Linear dependence (SACF)

Fitting a linear model at lag k , i.e. estimating ψ_0 and ψ_k and testing

$$H_0 : \psi_k = 0, \quad (13)$$

$$H_A : \psi_k \neq 0. \quad (14)$$

using Equation (12) and (9) will result in an estimate and a confidence band for the sample autocorrelation at lag k .

2.2.2 Non-linear dependence (LDF)

The framework can be used for testing non-linear dependence at lag k , similar to lag dependent functions (LDFs), see [23]. Fitting a non-linear model at lag k , i.e. estimating the $\psi_k, \alpha_{kj}, j = 2, \dots, q$ and $\beta_{kl}, l = 1, \dots, p$ and testing

$$H_0 : \psi_k = 0, \alpha_{kj} = 0, \beta_{kl} = 0 \quad (15)$$

$$H_A : \text{some } \psi_k \neq 0, \alpha_{kj} \neq 0 \text{ or } \beta_{kl} \neq 0. \quad (16)$$

will result in an estimate and a confidence band for the lag dependent function at lag k .

2.2.3 Strictly non-linear dependence (NLDF)

Lag dependent functions do not distinguish between linear and non-linear dependence. This can be done by using strictly non-linear lag dependent functions (NLDF), see [23]. This is done within the proposed framework by fitting a non-linear model at lag k , estimating $\psi_k, \alpha_{kj}, j = 2, \dots, q$ and $\beta_{kl}, l = 1, \dots, p$ and testing

$$H_0 : \alpha_{kj} = 0, \beta_{kl} = 0, \quad (17)$$

$$H_A : \text{some } \alpha_{kj} \neq 0 \text{ or } \beta_{kl} \neq 0. \quad (18)$$

This will result in an estimate and a confidence band for the strictly non-linear lag dependent function at lag k .

2.2.4 Partial linear dependence (PACF)

Extending the linear dependence test to partial linear dependence is done by testing for additional linear dependence at lag K by estimating ψ_k , $k = 0, \dots, K$ and testing

$$H_0 : \psi_K = 0, \quad (19)$$

$$H_A : \psi_K \neq 0. \quad (20)$$

This test is essentially the partial autocorrelation (PACF). The test resembles Section 2.2.1 but the residual degrees of freedom is different as more parameters are estimated.

2.2.5 Partial non-linear dependence (PLDF)

We compute the partial lag dependent function, i.e. testing for partial (linear and non-linear) dependence, at lag K by fitting a fully non-linear model for lag $k = 1, \dots, K$ and testing

$$H_0 : \psi_K = 0, \alpha_{Kj} = 0, \beta_{Kl} = 0, j = 2, \dots, q, l = 1, \dots, p, \quad (21)$$

$$H_A : \text{some } \psi_K \neq 0, \alpha_{Kj} \neq 0 \text{ or } \beta_{Kl} \neq 0. \quad (22)$$

This test is useful for finding the appropriate number of lags in a forward selection procedure.

2.2.6 Partial strictly non-linear dependence (New)

It is also possible to test for partial strictly non-linear dependence by testing a $(K - 1)$ th order non-linear model and linear terms at lag K against a K th order fully non-linear model. This is achieved by testing

$$H_0 : \alpha_{Kj} = 0, \beta_{Kl} = 0, j = 2, \dots, q, l = 1, \dots, p, \quad (23)$$

$$H_A : \text{some } \alpha_{Kj} \neq 0 \text{ or } \beta_{Kl} \neq 0. \quad (24)$$

2.2.7 A linear portmanteau test

The null hypothesis in the Box-Pierce, Ljung-Box and Monti tests is that data is *iid*, the alternative is that the data generating process is a linear process. We can test this joint hypothesis by fitting an $\text{AR}(K)$ process, and testing

$$H_0 : \psi_k = 0, k = 1 \dots K, \quad (25)$$

$$H_A : \text{some } \psi_k \neq 0. \quad (26)$$

The resulting test is performing similarly to the Box-Ljung test as we will see in Section 3.2.

Example 2.2. *The Ljung-Box test can either be applied to data or to residuals. The test statistic for residuals, when an ARMA(a, b) was fitted to data is approximately $\chi^2(K - a - b)$ where K is the number of lags tested. By comparison, the test statistic for comparing an AR($a + b$) to a larger AR(K) model in our framework is approximately $F(K - a - b, N - K)$ which is equivalent to $\chi^2(K - a - b)$ for large samples.*

2.2.8 Non-linear portmanteau (new)

Another alternative would be to test if data is *iid* against the alternative that the data generating process is a general, non-linear process, represented by a semi-parametric NLAR(K) process. This can be tested using the hypothesis

$$H_0 : \psi_k = 0, \alpha_{kj} = 0, \beta_{kl} = 0, \quad k = 1, \dots, K, \quad j = 2, \dots, q, \quad l = 1, \dots, p, \quad (27)$$

$$H_A : \text{some } \psi_k \neq 0, \alpha_{kj} \neq 0 \text{ or } \beta_{kl} \neq 0. \quad (28)$$

This test is less powerful than the linear portmanteau test for purely linear structures, but the test is also capable of detecting non-linear structures.

2.2.9 Strictly non-linear portmanteau (new)

It is also possible to test for strictly non-linear dependence. The hypothesis would then be to compare a semi-parametric NLAR(K) model to a linear AR(K) model, testing

$$H_0 : \alpha_{kj} = 0, \beta_{kl} = 0, \quad k = 1 \dots K, \quad j = 2, \dots, q, \quad l = 1, \dots, p, \quad (29)$$

$$H_A : \text{some } \alpha_{kj} \neq 0 \text{ or } \beta_{kl} \neq 0. \quad (30)$$

The test can be used for testing residuals from a linear model for additional non-linear autodependence.

3 Monte Carlo simulations

We evaluate the proposed tests on simulated data, applying the tests to several linear and non-linear models. We focus on power, but compute size when possible.

3.1 Size study

We have validated the size by simulating $N_{sim} = 10\,000$ independent samples from

$$y_n = m + kx_n + e_n, \quad (31)$$

where x_n are known variables (drawn from $U(0, 1)$) and e_n are pairwise independent standard Gaussian random variables. Four different samples sizes were used, $N_{obs} = 50, 200, 500$ or $1\ 000$, where the first corresponds to the sample size in [6]. We test for non-linear effects, by defining

$$H_0 : \alpha_j = 0, \beta_l = 0, \text{ for } j = 2, \dots, q, l = 1, \dots, p \tag{32}$$

$$H_A : \text{some } \alpha_j \text{ or } \beta_l \neq 0 \tag{33}$$

The results, when varying γ^2 are reported in Figure 1. We see that the truncated linear spline basis is biased (upwards and downwards, depending on γ^2), while higher order (≥ 3) splines are practically unbiased for large samples (an approximate confidence interval is given by $I_{\hat{p}} \approx (0.045, 0.055)$). The size study makes us confident to use penalized splines for medium sized and large samples.

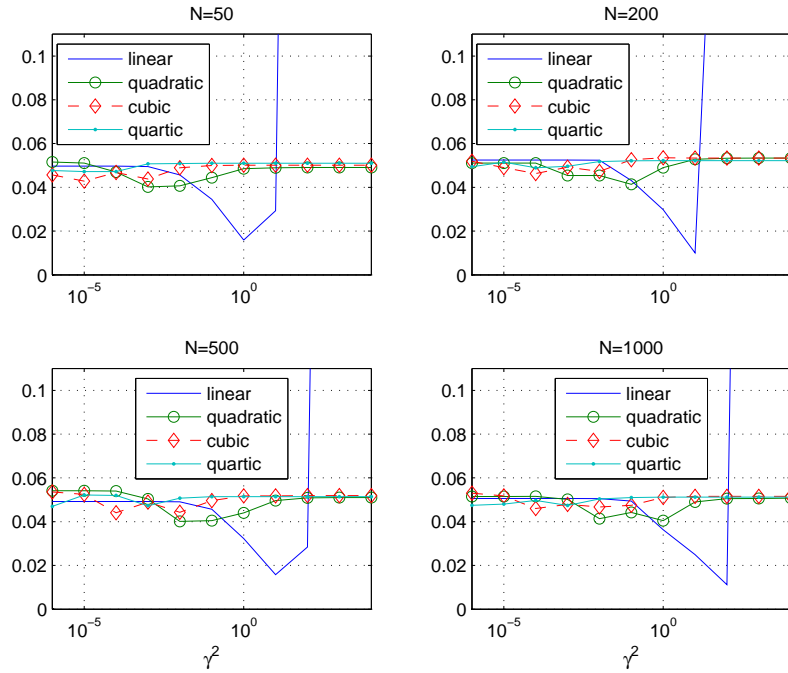


Figure 1: Estimated probability of rejecting the null hypothesis when the null hypothesis (linear regression) is true. The tests are performed at the 5% percent level.

3.2 AR(1) model

Linear models are well known, and several tests exist for testing linear dependence. We have focused on the Box-Ljung test as our benchmark. Let $\{Y\}$ be

an AR(1)-process

$$Y_n = a_1 Y_{n-1} + e_t,$$

where a_1 is being varied from $a_1 = -0.7$ to $a_1 = 0.7$ and e_t are *iid* standard Gaussian random variables. We apply the tests in Section 2.2.7, 2.2.8 and 2.2.9 and compare them with the Ljung-Box test. The simulations are based on 1 000 independent simulations. All tests are based on 5 lags (to emulate the situation when we do not know the true model order), and performed at a 5% significance level. We used $p = 10$ knots for the truncated polynomial spline basis for each lag and the penalty parameter was chosen as $\gamma^2 = 1$. Our sample sizes are almost an order of magnitude larger than e.g. [6] who reported that F-tests often are biased. We believe that the bias they found is related to small sample effects.

N=200 The statistical power of the tests when $N = 200$ observations are used is presented in Figure 2. There is little difference between the Box-Ljung test and our linear test of dependence. The non-linear test is less powerful than the linear tests while the strictly non-linear tests fails to detect any dependence (rejects the null hypothesis in approximately 4% of the simulations) regardless of the value of a_1 . Some bias is found for the non-linear tests in these simulations.

N=1 000 Increasing the sample size to $N = 1\,000$ observations decreases the bias and increases the power, see Figure 2. The strictly non-linear test is now detecting non-linear terms in approximately 5% of the simulations, as is expected. The Box-Ljung test and our linear test were almost identical in terms of power, but our test suite was also able to test for other types of dependence including non-linearities.




Figure 2: Testing for white noise when the true model is an AR(1) process, $Y_n = a_1 Y_{n-1} + e_n$. The graph on the left presents the power when $N = 200$ observations were used and the graph on the right presents the power when $N = 1\,000$ observations were used.

3.3 Logistic AR model

The test suite was also applied to a non-linear AR(1), driven by Gaussian additive noise. The model, used by e.g. [23] is given by

$$Y_n = \frac{1}{1 + \exp(-5Y_{n-1} + 2.5)} + 0.25e_n, \quad (34)$$

where e_n is a sequence of *iid* standard Gaussian random variables. 10 000 independent trajectories, each consisting of $N = 1\,000$ observations was simulated from the model. The linear and non-linear portmanteau tests, both using 5 lags and a 5% significance level were applied to each series. Both tests detected dependence in all simulations.

We proceeded with applying partial non-linear tests (see Section 2.2.5) to pinpoint the size of the model, comparing a constant to an NLAR(1), NLAR(1) to NLAR(2) etc. at the 5% level. The results from the partial tests are presented in Table 1, indicating that an NLAR(1) is the preferred model.

Percentage of rejections of the null hypothesis (PLDF)					
H_0 Model	NLAR(0)	NLAR(1)	NLAR(2)	NLAR(3)	NLAR(4)
Rejection (%)	100	4.5	4.2	4.2	4.1

Table 1: Percentage of tests rejecting the null hypothesis (NLAR(k)) in favour of a larger model (NLAR($k+1$)) testing for partial non-linear dependence. The data generating process is a NLAR(1) process and the tests were performed at the 5% level.

The final test was comparing linear first order autoregressive model against a general first order autoregressive model using the strictly non-linear test, see Section 2.2.3. The purely linear model was rejected in all simulations.

3.4 Non-linear AR model - humps

The logistic AR process did not present a challenge for any of the tests we tried. A more complex model was defined from the Matlab `humps.m` function. It is given by

$$h(x) = \frac{1}{(x - 0.3)^2 + 0.01} + \frac{1}{(x - 0.9)^2 + 0.04} - 6. \quad (35)$$

A scaled version of this function is used to define a non-linear process, leading to

$$Y_n = ah(Y_{n-1}) + e_n \quad (36)$$

where a is a small number that we use as a design parameter. Two realizations, using $a = 0.01$ and $a = 0.1$ is presented in Figure 3, together with the estimated sample autocorrelations.




Figure 3: Two realizations from the non-linear (humps) autoregressive process with different parameters (top/bottom). The right panel presents the estimated autocorrelations.

It is hard to distinguish these realizations from white noise, but the realization using $a = 0.1$ is clearly non-Gaussian due to the spikes.

Formal testing for linear dependence was not done as the autocorrelation is practically zero for all lags. We have instead used the following non-linear alternative models:

Quadratic $Y_n = c_0 + c_1 Y_{n-1} + c_2 Y_{n-1}^2 + e_n$, testing $H_0 : c_2 = 0$.

Cubic $Y_n = c_0 + c_1 Y_{n-1} + c_2 Y_{n-1}^2 + c_3 Y_{n-1}^3 + e_n$, testing $H_0 : c_2 = 0, c_3 = 0$.

Cubic+exp $Y_n = c_0 + c_1 Y_{n-1} + c_2 Y_{n-1}^2 + c_3 Y_{n-1}^3 + c_4 Y_{n-1} \exp(-|Y_{n-1}|) + e_n$, testing $H_0 : c_2 = c_3 = c_4 = 0$

Strictly non-linear LDF at lag 1 (PS) defined using penalized splines, cf. Section 2.2.9. The tests were run using $\gamma^2 = 10^{-10}, 10^{-8}, \dots, 10^2$ and $p = 30$ knots. The optimal γ^2 , computed using AIC ranges from about 10^2 or larger when $a = 0$ to $10^{-4} - 10^{-7}$ for $a \geq 0.02$.

The quadratic, cubic and cubic+exp alternative models were used by e.g. [4]. All tests were applied to 1 000 independent trajectories, each consisting of 100 observations. The tests were performed at the 5% confidence level. The resulting statistical power is presented in Figure 4.




Figure 4: Probability of rejecting the linear model in favour of a non-linear model when the a parameter is being varied. The tests were performed with the alternative model being a quadratic function, a cubic function, a cubic + exponential function or a penalized truncated polynomial spline basis (PS).

We see that the purely polynomial models are struggling to find non-linear structures, and that the penalized spline models with the largest penalty is marginally outperforming them. The combination of the cubic polynomial and the exponential-type term is somewhat more successful while the more weakly penalized spline based tests are quite successful at finding the non-linear structure in the data. The size ($a = 0$) is reasonable for all tests.

3.5 Non-linear GARCH model

The GARCH model, see [1], is one of the most popular models in econometrics, due to the simplicity of the model and the predictive performance. The standard model is given by

$$X_t = \sigma_t z_t, \quad z_t \in \mathcal{N}(0, 1), \tag{37}$$

$$\sigma_t^2 = \omega + \alpha X_{t-1}^2 + \beta \sigma_{t-1}^2. \tag{38}$$

Several extensions have been suggested, see [2] including non-parametric versions. We have studied a switching volatility GARCH in, [3], defined as

$$\sigma_t^2 = \omega + \alpha X_{t-1}^2 + (\beta_+ \mathbf{1}_{\{X_{t-1} > 0\}} + \beta_- \mathbf{1}_{\{X_{t-1} \leq 0\}}) \sigma_{t-1}^2, \tag{39}$$

$$= \omega + \alpha X_{t-1}^2 + f(X_{t-1}) \sigma_{t-1}^2 \tag{40}$$

The rationale for replacing the β -parameter with a function that depends on the sign of the lagged return, X_{t-1} is that good news and bad news have

different influence on the conditional variance. The heteroscedastic framework was used to test if $f(X_{t-1})$ is constant, approximating equation (40) by

$$Y_t = X_t^2 = \omega + \alpha X_{t-1}^2 + f(X_{t-1})\sigma_{t-1}^2 + \nu_t \tag{41}$$

$$\approx \omega + \alpha X_{t-1}^2 + \left(\psi_0 + \sum_{j=1}^q \theta_j X_{t-1}^j + \sum_{l=1}^p \omega_l (X_{t-j} - \kappa_l)_+^q \right) \sigma_{t-1}^2 + \nu_t, \tag{42}$$

where $\nu_t = X_t^2 - \sigma_t^2 = \sigma_t^2(z_t^2 - 1)$ is heteroscedastic white noise, and σ_t^2 was estimated using a standard GARCH(1,1) model. We test for non-linearities by defining

$$H_0 : \theta_j = 0, \omega_l = 0, j = 1 \dots q, l = 1 \dots p, \tag{43}$$

$$H_A : \theta_j \neq 0 \text{ or } \omega_l \neq 0. \tag{44}$$

Data was generated using $\omega = 5, \alpha = 0.2, \beta_- = 0.75$ while β_+ was varied between 0.1 and 0.8. The test was evaluated by simulating 1 000 independent trajectories from the non-linear GARCH model, each consisting of 1 000 observation for every value of β_+ . The results, computed for three different values of γ^2 , is presented in Figure 5. AIC suggests that γ^2 should vary from 10^4 for small values of β_+ to approx 10^1 for larger values.




Figure 5: Probability of rejecting the standard GARCH-process when the data generating process is a switching GARCH with parameter β_+ .

It can be seen that the test is able to detect the non-linear structure in the switching GARCH process, despite the fact that the non-linear term is discontinuous. It can also be seen that the size is reasonable.

4 Conclusion

This paper presents a framework for testing serial dependence. Nine different tests were derived, including linear, non-linear and strictly non-linear versions. The tests are based on penalized splines, making the non-linear versions capable of fitting very general structures without losing too much power. Testing is done using approximate F -tests. The size was correct for large samples when higher order splines are used, cf. Section 3.1, 3.2, 3.4 and 3.5 and are computationally much more efficient than simulation or bootstrap based tests.

The Monte Carlo simulations show that the linear versions are equal in power to the Ljung-Box test for large samples while the non-linear versions found dependence where other tests had difficulties.

References

- [1] T. Bollerslev. Generalized Autoregressive Conditional Heteroscedasticity. *Journal of Econometrics*, 31:307–327, 1986.
- [2] T. Bollerslev. Glossary to ARCH (GARCH). In Tim Bollerslev, Jeffrey R Russel, and Mark Watson, editors, *Volatility and Time Series Econometrics: Essays in Honour of Robert F. Engle*. Oxford University Press, 2010.
- [3] P. Bühlmann and A. J. McNeil. An algorithm for nonparametric garch modelling. *Computational Statistics & Data Analysis*, 40(4):665–683, October 2002.
- [4] J.L. Castle and D.F. Hendry. A low-dimension portmanteau test for non-linearity. *Journal of Econometrics*, 2010.
- [5] G. Claeskens, T. Krivobokova, and J.D. Opsomer. Asymptotic properties of penalized spline estimators. *Biometrika*, 96(3):529, 2009.
- [6] C. Crainiceanu and D. Ruppert. Likelihood ratio tests for goodness-of-fit of a nonlinear regression model. *Journal of Multivariate Analysis*, 91(1):35–42, 2004.
- [7] J. Friedman, R. Tibshirani, and T. Hastie. *The elements of statistical learning: Data mining, inference, and prediction*. Springer-Verlag New York, 2009.
- [8] C.W.J. Granger and T. Teräsvirta. *Modelling nonlinear economic relationships*. Oxford University Press, USA, 1993.
- [9] T. J. Hastie and R. J. Tibshirani. *Generalized Additive Models*. Chapman and Hall, 1990.

- [10] K.C. Hung, S.H. Cheung, W.S. Chan, and L.X. Zhang. On a robust test for SETAR-type nonlinearity in time series analysis. *Journal of Forecasting*, 2009.
- [11] C.M. Hurvich, J.S. Simonoff, and C.L. Tsai. Smoothing parameter selection in nonparametric regression using an improved akaike information criterion. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(2):271–293, 2002.
- [12] G. Kauermann, G. Claeskens, and J.D. Opsomer. Bootstrapping for penalized spline regression. *Journal of Computational and Graphical Statistics*, 18(1):126–146, 2009.
- [13] D. Keenan. A tukey nonadditivity-type test for time series nonlinearity. *Biometrika*, 72(1):39, 1985.
- [14] T. Krivobokova, C.M. Crainiceanu, and G. Kauermann. Fast adaptive penalized splines. *Journal of computational and graphical statistics*, 17(1), 2008.
- [15] G. Li and W.K. Li. Testing a linear time series model against its threshold extension. *Biometrika*, 98(1):243, 2011.
- [16] Y. Li and D. Ruppert. On the asymptotics of penalized splines. *Biometrika*, 95(2):415, 2008.
- [17] E. Lindström. Model validation for diffusion processes using generalized gaussian residuals. In *First Brazilian Conference on Statistical Modelling in Insurance and Finance*, 2003.
- [18] Erik Lindström, Jonas Ströjby, and Stefan Ingi Adalbjörnsson. Non-Linear Portmanteau Tests. In *Proceedings IFAC System Identification*, volume 15, 2009.
- [19] G. M. Ljung and G. E. P. Box. On a Measure of Lack of Fit in Time Series Models. *Biometrika*, 65:297–303, 1978.
- [20] E. Mahdi and A. I. McLeod. Improved multivariate portmanteau test. *Journal of Time Series Analysis*, 2011. doi: 10.1111/j.1467-9892.2011.00752.x.
- [21] A. I. McLeod and W. K. Li. Diagnostic checking arma time series models using squared residual autocorrelations. *Journal of Time Series Analysis*, 4:269–273, 1983.
- [22] A.C. Monti. A proposal for a residual autocorrelation test in linear models. *Biometrika*, 81(4):776, 1994.

- [23] H. Aa. Nielsen and H. Madsen. A Generalization of some Classical Time Series Tools. *Computational Statistics and Data Analysis*, 37(1):13–31, 2001.
- [24] D. Peña and J. Rodríguez. A powerful portmanteau test of lack of fit for time series. *Journal of the American Statistical Association*, 97(458):601–610, 2002.
- [25] M. Rosenblatt. Remarks on a Multivariate Transformation. *Annals of Mathematical Statistics*, 23:470–472, 1952.
- [26] D. Ruppert, M. P. Wand, and R. J. Carroll. *Semiparametric Regression*. Cambridge University Press, 2003.
- [27] T. Teräsvirta, C. F. Lin, and C. Granger. Power of the neural network linearity test. *Journal of Time Series Analysis*, 14(2):209–220, 1993.
- [28] H. Tong. *Non-linear time series: a dynamical system approach*, volume 6. Oxford University Press, 1993.
- [29] H. Tong. Threshold models in time series analysis 30 years on. *Statistics and Its Interface*, 2010.
- [30] R.S. Tsay. Nonlinearity tests for time series. *Biometrika*, 73(2):461, 1986.

Received: October, 2012