



Lund University
Department of Statistics

Threshold Detection in Autoregressive Non-linear Models

Jacob Järås and Azadeh M.Gishani

Master Thesis: 15 ECTS
Supervisor: Peter Gustafsson
Spring 2010

Abstract

In this paper we fit non-linear models. We build Threshold Autoregressive (TAR) and Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models and estimate the parameters associated to the models, e.g. the threshold for the TAR model. The TAR and the GARCH model concept are applied to simulated data and to three empirical datasets, two River flow time series and one Blowfly data set. We observe significant non-linear effects from the tests for the three empirical time series. Two different TAR models fit to the Blowfly data. We fit ultimately TAR models to the river data sets. The fitted AR-GARCH model does not give satisfactory results for the three empirical data sets.

Keywords: Non-linear Time Series, Threshold Autoregressive Model, ARCH, GARCH.

Contents

1	Introduction	3
1.1	Background	3
1.2	Problem	3
1.3	Purpose	3
2	The Data Sets	4
3	Models	4
3.1	Non-Linear Models	4
3.1.1	The TAR and the SETAR Models	5
3.1.2	TGSP Algorithm	6
3.1.3	The ARCH and GARCH Models	6
3.2	Model Specification	7
3.2.1	Model Diagnostics	7
3.2.2	Calculations	8
3.2.3	Overfitting and Parameter Redundancy	9
3.2.4	Nonlinearity Tests	9
4	Applications	11
4.1	Simulated Series	11
4.1.1	Simulated TAR Model	11
4.1.2	Simulated GARCH Model	14
4.2	Real Series	17
4.2.1	The Blowfly Data	17
4.2.2	The River Flow in Northern Scotland	20
5	Conclusions and Discussions	35
5.1	Conclusions	35
5.2	Discussions	35
A	MATLAB routine for TGSP	38
B	Matlab routine for simulating a TAR model	39

1 Introduction

1.1 Background

In the field of time series there are many linear processes that can be fitted by the autoregressive moving average (ARMA) models. However, there are several time series which do not exhibit a linear behaviour. These processes cannot be well fitted by the common and popular ARMA models.

To adequately fit these non-linear time series, other more complicated models have to be taken into account that has the ability to capture the dynamics of the series more precisely.

Although many tests exist that has the capability of detecting the presence of nonlinearity there isn't any single dominating test that is able to pin down the exact type of nonlinearity. Rather can the test available, suggest that some form of nonlinearity structure is present in the data and perhaps aim towards a general class of nonlinear models that should be examined closer.

1.2 Problem

Since the category of non-linearity includes such a wide range of different shapes that a time series can take, the primary problem is that of correctly specify or approximate the right kind of non-linearity and then fit a model to it. There are many different models to choose from when it comes to non-linearity and choosing the best one may be a matter of taste. Two analysts may choose different models for the same data.

The choice of model depends not only on the specific series under study as well as on the characteristics of the model under consideration but also on the analysts experience and background. The model that agrees the most with the features of the series being analyzed should be applied and studied more thorough.

The models we are looking at in this thesis are the Threshold Autoregressive (TAR) and the General Autoregressive Conditional Heteroskedasticity (GARCH) models. The model that agrees the most with the features of the series being analyzed should be applied and studied more thorough.

1.3 Purpose

The purpose of the thesis is to compare the characteristics and the performance of different models developed to handle non-linearity and heteroskedasticity (time varying variance). We use simulated TAR and GARCH data as well as three real TAR time series data. We then apply the TAR and the GARCH models to the simulated and the real data and evaluate the findings. In specific, we seek to find out if possible thresholds is present in the data.

2 The Data Sets

In the first part of the thesis, we simulate some TAR and GARCH processes and apply the TAR and ARCH/GARCH models to the simulated data. By simulations we have the benefit of observing the performance of each model to different kinds of data. In the second part we apply the models to real data. The data sets used is the Blowfly data set and two river flow data set taken from UK's centre of Ecology & Hydrology.

The A.J Nicholsons Blowfly data is a frequently used data in population ecology. It is the data derived from the experiment which observes every other day, the population change of 1000 blowflies caged in, with a restricted amount of food being available daily.¹ See Tong [12].

The data sets used in the second part are obtained from the home page of UK's centre of Ecology & Hydrology.² These data sets come from the measure station Easter Turnaig which monitors the Oykel river flow which is operated by the Scottish Environment Protection Agency (SEPA). The second data set comes from the Comber measure station which monitors Enler river flow. The entire series measure from late 1977 till 2008 and includes some 11500 observations. For our particular analysis of the Oykel and the Enler River, we consider only the measured series from 2005 to 2008, two samples of totally 1461 observations. The flow of the rivers is measured in cubic metres per second (also known as cumecs).

3 Models

3.1 Non-Linear Models

Among models for non-linearity we consider and further investigate are the TAR model, introduced by Tong [11] and the ARCH model first presented by Engle [8] and further developed by Bollerslev [2] to become the GARCH model. It should be noted that there exist many different extensions and development of the above mentioned models.

For fitting the models we use Matlab and R. In R, we have downloaded the package (TSA) which is a TAR model extension written by Chan [5]. We have also used a Matlab-R link that establishes a connection between the two programs making it available to run commands to R from within Matlab. We use a significance level of 5% throughout the analysis.

¹The blowfly data set can be accessed at:<http://robjhyndman.com/TSDL/data/blowfly.dat>

²The river data sets can be found at: http://www.ceh.ac.uk/data/nrfa/river_flow_data.html

3.1.1 The TAR and the SETAR Models

Threshold Autoregressive (TAR) models are commonly referred to as piecewise linear models or regime-switching models. They consist of k , AR(p) parts where one process change to another according to the value of an observed variable, a threshold. Once the series cross the threshold value, the process takes on another value. TAR models are usually referred to as TAR (k,p) where p is the autoregressive order in each regime. Since there can be different numbers of autoregressive orders in each regime, the model is sometimes denoted TAR(k).

The Self-Exciting Threshold Autoregressive (SETAR) model, first introduced by Tong [11], is a special case of the TAR model. Here, the movements between the regimes are controlled or governed by a variable called threshold just as in the TAR model with the difference that the threshold of a SETAR model is Self-Exciting. This means that, unlike the TAR model, where the threshold is assumed to be an exogenous variable, the threshold variable of a SETAR model is a certain lagged value of the series itself, an endogenous variable.

To estimate a time series with an assumed TAR model behavior it is essential to know the value of the threshold parameter in the series. If the value of the threshold parameter is known, the estimation of the TAR model is readily available. Then for example a two regime TAR model can be represented by equation (1). The observations are sorted according to if y_{t-1} is greater or smaller than the threshold value and each regime can be estimated as a common AR process.

$$Y_t = I_t[\alpha_{10} + \sum_{i=1}^p \alpha_{1i}Y_{t-i}] + (1 - I_t)[\alpha_{20} + \sum_{i=1}^r \alpha_{2i}Y_{t-i}] + \epsilon_t \quad (1)$$

where the error term ϵ_t is a white-noise process and I_t is an indicator function such as

$$I_t = 1 \text{ if } Y_{t-1} > \tau \text{ and } I_t = 0 \text{ if } Y_{t-1} \leq \tau$$

and τ is the threshold variable that separates the two regimes. The more general form also allows different error terms variances across regimes :

$$Y_t = \begin{cases} \alpha_{10} + \alpha_{11}Y_{t-1} + \dots + \alpha_{1p}Y_{t-p} + \epsilon_{1t} & \text{if } Y_{t-1} > \tau \\ \alpha_{20} + \alpha_{21}Y_{t-1} + \dots + \alpha_{2r}Y_{t-r} + \epsilon_{2t} & \text{if } Y_{t-1} \leq \tau \end{cases} \quad (2)$$

If the switch of the process takes more than one period to occur, then its appropriate to select the delay parameter d which solely is the length of the lag of the threshold variable, the size of the lag. The delay parameter is a positive integer. It may be estimated along with the other parameters of the model. As with the estimation of the threshold value, mentioned earlier, the estimation procedure of the delay parameter of a TAR model is usually done for each potential value of the d . The model that yields the smallest residual sum of squares makes the most consistent estimate of the delay parameter.

3.1.2 TGSP Algorithm

In most situations, the threshold value is unknown and needs to be estimated along with the other parameters in the model. One way of finding the unknown threshold value is the method proposed by Chan [4] mentioned in Enders [7]. The suggested approach is based on the idea that the threshold value is one of the elements of the series itself. We now can consider all of the elements of the series as a potential threshold value and fit a model of the form (1) or (2) to each value of y_t . Only the middle 70 or 80% of the series is tested for the threshold. The reason for this restriction is to keep a satisfactory amount of observations on each regime when estimating the threshold and the other parameters in the TAR model.

We specify the regimes and estimate a model in each regime by least squares regression. For each of the potential threshold the residual sum of squares is saved and we choose the threshold value that is corresponding to the model with the least sum of squared residuals. Putting it in a different way, it could be explained that the threshold model seeks to divide the observations in such a way that each regime has a homoskedastic error process.

Here the algorithm for finding the threshold is called the Threshold Grid Search Process (TGSP), and a simple routine for it in MATLAB can be found in Appendix A.

3.1.3 The ARCH and GARCH Models

Autoregressive Conditional Heteroskedasticity (ARCH), first introduced by Engle [8], and Generalized Autoregressive Conditional Heteroskedasticity (GARCH), first introduced by Bollerslev [2], models are models whose primary field of interest is in the financial sector. They were originally developed to capture the dynamics of the conditional volatility. Bollerslevs further development of Engle ARCH model lead to GARCH model. The main difference from the ARCH model is that the GARCH allow for the error terms to be modelled as an ARMA process. Two important characteristics within the sector of financial time series are fat tails and volatility clustering, which can be captured and explained by one of the members of the GARCH family models.

If we estimate a series with an ARMA model, and the conditional variance of this series is not a constant and it is predicted with an AR(p) process of , then we have a time series with ARCH effect; consider the following AR(1) process:

$$y_t = a_0 + a_1 y_{t-1} + \epsilon_t \quad (3)$$

where ϵ_t is the error process.

$$Var(y_{t+1}|y_t) = E_t(\epsilon_{t+1})^2 \quad (4)$$

The estimated conditional variance is:

$$\hat{\epsilon}_t^2 = \alpha_0 + \alpha_1 \hat{\epsilon}_{t-1}^2 + \alpha_2 \hat{\epsilon}_{t-2}^2 + \dots + \alpha_q \hat{\epsilon}_{t-q}^2 + \nu_t \quad (5)$$

where ν_t is a white-noise process.

Now consider that the error process of the fitted ARMA model is:

$$\epsilon_t = \nu_t \sqrt{h_t} \tag{6}$$

where the variance of ν_t process, σ_ν^2 , is one.

$$h_t = \alpha_0 + \sum_{i=1}^q \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^p \beta_i h_{t-i} \tag{7}$$

Then we have GARCH(p,q).

3.2 Model Specification

Model selection is a set of steps designed to basically make the process of model selection faster and easier. By following the below steps, misspecifications will much easier be detected than if not following the steps, since the advancement of the models follow the order of the steps.

1. Make a visual inspection of the data by plotting. This is to help determine the shape of the potential nonlinearity. Structural breaks as well as outliers will deviate from the pattern.
2. The second step is to fit the best linear model possible to the data. Then running a set of tests will help evaluate if the residuals are serially correlated.
3. Step three is to test for non-linearity. We consider Tsay (Cryer & Chan [6]), Keenan (Cryer & Chan [6]) and The Likelihood Ratio Test (Cryer & Chan [6]). Other test exists as well to alarm the presence of a non-linear behaviour, such as the RESET test, the McLeod-Li test and various Lagrange multiplier tests. See Enders [7] and Cryer & Chan [6].
4. In case non-linearity is found, then the appropriate form of non-linearity should be determined. As there are many kinds of non-linearity, there is no general way.
5. The chosen non-linear model should make a better fit to the data than the linear model, and all parameters should be significant.

3.2.1 Model Diagnostics

The other very important part of building a model, or even one of the most important ones, is that of validating the residuals. In order for a model to be correctly specified, there shouldn't be any remaining structure in the residuals after having fitted a

model. If there is, predictions from the model may be unreliable and misleading. That would mean that the model to some degree failed to capture or explain the underlying generating mechanism of the data.

On the other hand, if the model is correctly specified, then the remaining residuals should be approximately independently identically normally distributed, in other words white noise.

We assess the goodness of fit for the models and, if necessary, suggest appropriate modifications in case of poor fit.

3.2.2 Calculations

The raw residuals from the TAR model are defined by subtracting the fitted value from the data (Cryer & Chan 2008), where the t -th fitted value is the estimated conditional mean of Y_t given past values of Y 's; the residuals, $\hat{\epsilon}_t$, are calculated as:

$$\begin{aligned} \hat{\epsilon}_t = Y_t - \{ & \hat{\Phi}_{1,0} + \hat{\Phi}_{1,1}Y_{t-1} + \dots + \hat{\Phi}_{1,p}Y_{t-p} \} I(Y_{t-d} \leq \hat{\tau}) \\ & - \{ \hat{\Phi}_{2,0} + \hat{\Phi}_{2,1}Y_{t-1} + \dots + \hat{\Phi}_{2,q}Y_{t-q} \} I(Y_{t-d} > \hat{\tau}) \end{aligned} \quad (8)$$

Standardised residuals are calculated to overcome the problem of different variances of the raw residuals. The standardized residuals \hat{e}_t are obtained by normalization, which is, dividing each of the raw residuals by their standard deviation.

$$\hat{e}_t = \frac{\hat{\epsilon}_t}{\hat{\sigma}_1 I(Y_{t-d} \leq \hat{\tau}) + \hat{\sigma}_2 I(Y_{t-d} > \hat{\tau})} \quad (9)$$

Other measures of spread we will use in the evaluation section is the Root Mean Squared Error, RMSE.

$$RMSE = \sqrt{\frac{1}{\tau} \sum_{t=T+1}^{T+\tau} (h_t^2 - \sigma_t^2)^2} \quad (10)$$

where τ is the number of predictions, h_t^2 is the real variance and σ_t^2 is the predicted variance. The statistics give information about the spread of the residuals from a model fit. It can be described as being the standard deviation of the residuals, if unbiased.

Another test used to check residuals is the Jarque-Bera test

$$JB = \frac{n}{6} \left(S^2 + \frac{1}{4}K^2 \right) \quad (11)$$

where n is the number of observations, S is the sample Skewness, a measure of the asymmetry and K the sample kurtosis, a measure of the flatness.

This test uses the different moments of the sample to describe various characteristics of the distribution. For example, the first moment is the mean, the second moment is the variance, and third and fourth moment is the Skewness and Kurtosis respectively. See Jarque & Bera [1].

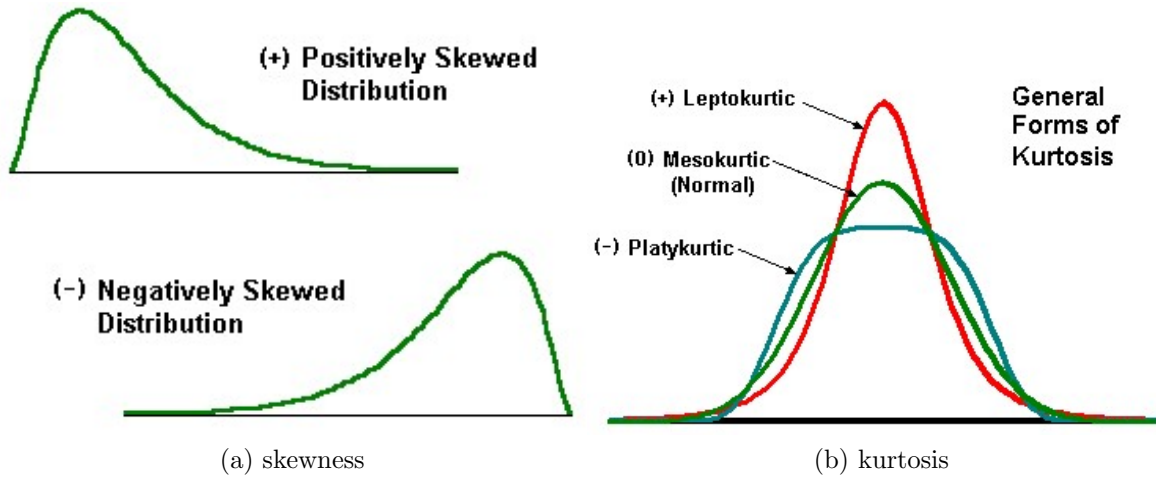


Figure 1: Visualisation of skewness and kurtosis

The normal distribution has skewness and kurtosis of both being equal to zero.

3.2.3 Overfitting and Parameter Redundancy

One way of assessing the performance of a model that has been specified is to fit another model to the same data but with an additional coefficient. For instance, if we have a TAR(2, 1, 1) model we may want to check if the TAR(2, 2, 1), TAR(2, 1, 2) or TAR(2, 2, 2) make a better one than our original model. The point of this procedure, called over fitting, is to see if the additional coefficient(s) are significantly different from zero as well as to see if the original estimates of the coefficients in the TAR(2, 1, 1) model change significantly due to the extra added coefficient. An indication of the overall fit of the model by this procedure allows us to make modifications if necessary.

3.2.4 Nonlinearity Tests

Portmanteau tests are residual based test that establish the degree of correlation a variable has on itself on various time intervals. They operate under the null hypothesis of residual independence. The alternative hypothesis is then that they are dependent.

We make use of various tests to assess if the series are nonlinear. The following tests are used: ARCH/GARCH test, Ljung-Box Q test, Keenan test, Tsay test and the likelihood ratio test for threshold non-linearity. In addition, we look at the ACF of the

residuals (squared as well) and consider the Jarque-Bera test for normality.

The ARCH/GARCH test we are mentioning here is the Lagrange Multiplier test suggested by Engle [8] for pretesting a series whether ARCH model is appropriate for it or not. The first step on ARCH/GARCH test is to have the residuals of an OLS regression; then the test statistic is the Lagrange Multiplier test statistic TR^2 which is derived from the regression of the squared residuals on a constant and p lags; TR^2 , where T is the number of squared residuals included in the regression and R^2 is the sample multiple correlation coefficient, has χ_p^2 distribution. A built-in MATLAB routine is used for this test. For further reading on Lagrange Multiplier test and the ARCH/GARCH test, see Enders [7] and Engle [8].

The Ljung-Box Q-test is a typical portmanteau test which checks if the residual of an estimated ARMA (p, q) model is white-noise or not. The test statistic is written as:

$$Q = T(T + 2) \sum_{k=1}^s \frac{r_k^2}{T - k} \quad (12)$$

The Q-statistic mentioned above has a distribution with s degrees of freedom. See Enders [7].

The null hypothesis, H_0 , of the likelihood ratio test for threshold non-linearity is that the fitted model to the series is an AR(p) model and the alternative hypothesis, H_1 , is that the fitted model to the series is a TAR(k) model with autoregressive order p in each regime. For example if, we suspect TAR non-linearity we estimate an AR(1) model as the H_0 and a TAR(2; 1, 1) as H_1 . The likelihood ratio test statistic of threshold non-linearity (TLRT) is

$$T_n = (n - p) \log \left\{ \frac{\hat{\sigma}^2(H_0)}{\hat{\sigma}^2(H_1)} \right\} \quad (13)$$

Where $\hat{\sigma}^2(H_0)$ is the sum of squared residuals of the fitted AR(p) model to the series and $\hat{\sigma}^2(H_1)$ is the sum of squared residuals of the TAR(k) model fitted to the data. The distribution of the test statistic is not a standard distribution. See Chan [3] and Tong [13].

Keenan's test is based on approximating a non-linear stationary series by Volterra expansion, which is:

$$y_t = \mu + \sum_{\mu=-\infty}^{\infty} \theta_{\mu} \epsilon_{t-\mu} + \sum_{\nu=-\infty}^{\infty} \sum_{\mu=-\infty}^{\infty} \theta_{\mu\nu} \epsilon_{t-\mu} \epsilon_{t-\nu} \quad (14)$$

where ϵ_t is a white noise process.

The null hypothesis of the Keenan test considers that the coefficients of the double sum part are zero and so the series is linear. The notations for Keenan test, TLRT test and Tsay test is from Cryer & Chan [6].

The Tsay test is an expansion of Keenans test which allows for non-linearity detection in a more general way. The point we should consider about Keenan and Tsay tests is that they detect quadratic non-linearity and they cannot be adequate for TAR models since they are linear in each regime. See Cryer & Chan [6].

4 Applications

4.1 Simulated Series

4.1.1 Simulated TAR Model

In this part we simulate a TAR model by help of a MATLAB routine presented in appendix II. The algorithm used for simulation is a simple straight forward way of doing it. The simulated model is:

$$Y_t = \begin{cases} 0.6y_{t-1} + \epsilon_t & \text{if } y_{t-1} \leq -0.2 \quad (\text{Lower regime}) \\ -0.9y_{t-1} + \epsilon_t & \text{if } y_{t-1} > -0.2 \quad (\text{Upper regime}) \end{cases} \quad (15)$$

where ϵ_t has mean zero and variance 0.25. This series is a simple TAR(2; 1, 1) model in which the two regimes can be distinguished easily. Figure 2 shows the plot of the series versus the first lag values, which gives an overview of the two different regimes separated by the threshold, here represented by the dotted red line.




Figure 2: Plot of the series versus its first lag values

Figure 3 shows the series versus time. In plot (a) and (b), we graph the two separated series which build the TAR model; in the plot (c) it can be seen that the TAR(1) model is not similar to any of the series which are building it; in plot (d) the elements of the

upper and lower regimes of the TAR model are separated and it can be observed how the TAR model uses the both series according to the threshold value to build the TAR procedure.

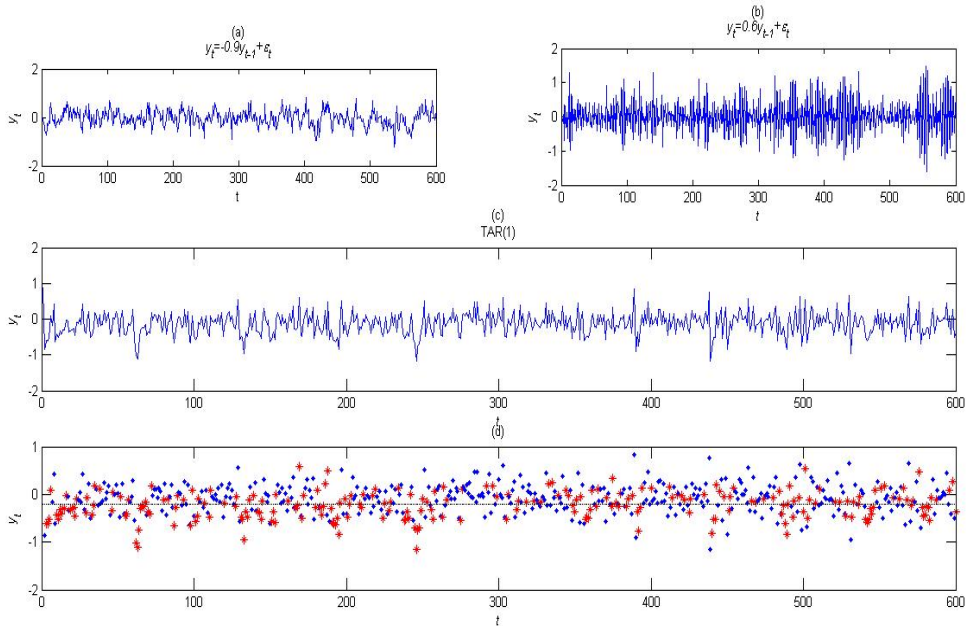


Figure 3: Simulated TAR model; In plot (d) the red stars show the lower regime, and the blue dots the upper regime.

Figure 4 shows the sample autocorrelation (SACF) and the sample partial autocorrelation (SPACF) for the simulated TAR(2, 1, 1) data, for the squared data, and for the third power of the simulated series; although the SACF and PACF of the series itself can not give much information on the non-linearity, the SACF of the higher powers of the series can point out the non-linear correlation of the data. Here, the SACF of the squared series point out correlation of the lags of the squared data which is an indication of presence of non-linearity.

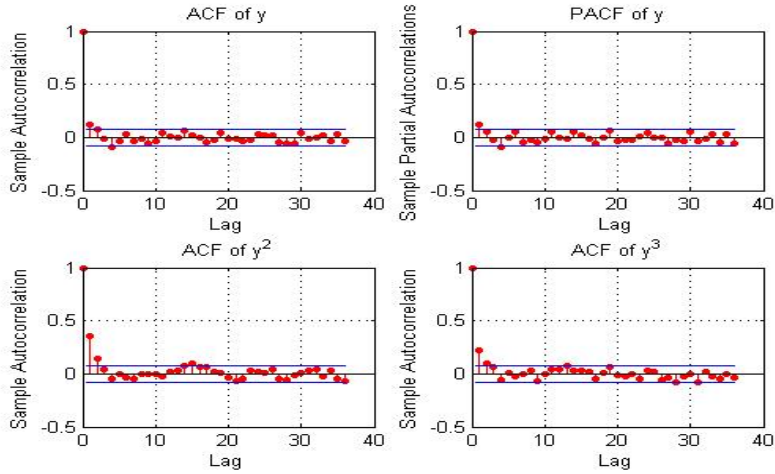


Figure 4: The sampled ACF and PACF for simulated TAR models

Now we examine the data for possible serial correlation and non-linearity. Ideally, these tests are able to pin down a suitable model that fits the data well. See Table 1.

Table 1: Results of pre-test for non-linearity of simulated TAR(2; 1, 1) data

Tests	p-value	test statistics	critical value
ARCH/GARCH Test			
Using up to lag 15	0.000	94.2181	24.9958
Using up to lag 20	0.000	97.0089	31.4104
Using up to lag 25	0.000	98.1159	37.6525
LBQ Test			
Using up to lag 15	0.0047	32.9864	24.9958
Using up to lag 20	0.0007	46.3696	31.4104
Using up to lag 25	0.0022	49.8536	37.6525

It's clear that some kind of dependence is present, both the test for ARCH effect and LB-Q test indicate that some form of heteroskedasticity in the data is present. Further, the tests for detecting non-linearity all have p-values less than 0.000. Keenan, Tsay and TLRT test suggest that the data is non-linear resulting in the test-statistics being 174.7, 70.74 and 28.79 respectively. Since the tests indicate heteroskedasticity; we try to fit an ARCH or GARCH model to the data, which in this case a GARCH model is not significant but an ARCH model is.

The estimated model with ARCH effect for the simulated data is:

$$y_t = -0.10294 + \epsilon_t \quad (16)$$

where the conditional variance is:

$$\sigma_t^2 = 0.077297 + 0.20838\epsilon_{t-1}^2 \quad (17)$$

In Table 2 the results of the fitted ARCH model is presented and it shows that the estimated ARCH model in equation (16) is significant.

Table 2: Result of fitted ARCH model to the simulated TAR(2; 1, 1)

Parameter	Value	Standard Error	T-statistic	P-value
C	-0.102	0.0143	-7.161	0.000
K	0.077	0.007	10.33	0.000
ARCH(1)	0.208	0.0778	2.677	0.011

AIC:297.572

4.1.2 Simulated GARCH Model

In this section, we try to fit a TAR model to a simulated GARCH model. The following set of simulated data with GARCH(2,1) effect is considered:

$$y_t = 0.35 + \epsilon_t \quad (18)$$

Where the conditional variance is:

$$\sigma_t^2 = 0.2004 + 0.1167\sigma_{t-1}^2 + 0.2369\sigma_{t-2}^2 + 0.4097\epsilon_{t-1}^2 \quad (19)$$

Figure 5 shows the plotted GARCH(2,1) model in equation (18) and its error term ϵ_t is called innovations which is plotted in the second plot of Figure 5. The innovations has a conditional variance which is AR(2) process in equation (19) and is shown in the third plot of Figure 5.

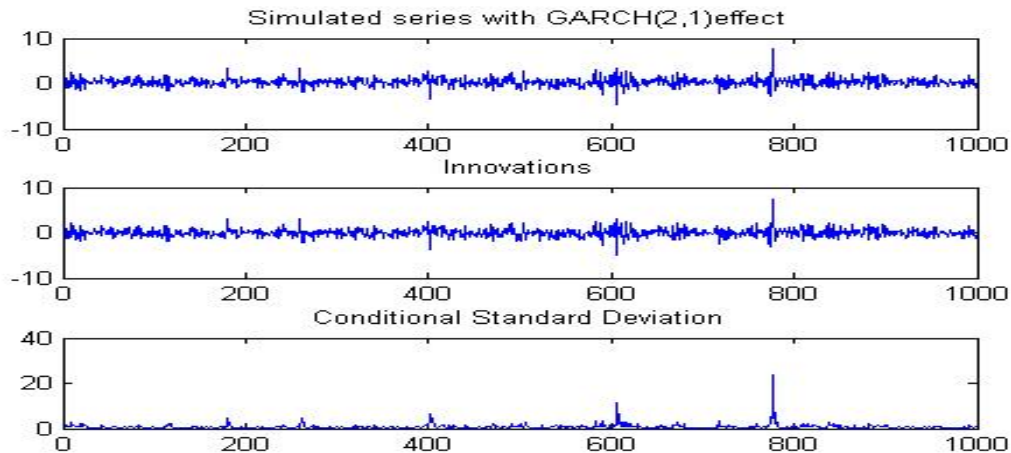


Figure 5: The simulated GARCH model

TGSP is used to find the threshold value for the simulated GARCH series; Figure 6 shows the sum of squared residuals(SSR) of the fitted models for the possible threshold values and the estimated threshold by this method for the series with GARCH effect is -0.2552 and the corresponding model for the threshold is:

$$Y_t = \begin{cases} 0.0533y_{t-1} + e_{t1} & \text{if } y_{t-1} \leq -0.2552 \quad (\text{Lower regime}) \\ 0.0945y_{t-1} + e_{t2} & \text{if } y_{t-1} > -0.2552 \quad (\text{Upper regime}) \end{cases} \quad (20)$$




Figure 6: The sum of squared residuals of the TAR models fitted to the GARCH data

But as it can be seen in Figure 6 the estimated Sum of Squared Residuals (SSR) of the series are fluctuating a lot and it can be interpreted as either the GARCH model has several thresholds or the SSR build a new non-linear process. The other point which should be noticed is that since we have not considered the distribution of the threshold value and the confidence interval for the threshold we can not choose this model as a significant model, even if there might exist significant thresholds, because the algorithm is unclear about the significance about the thresholds.

One of the possible tools to check the goodness of the fit for the model is that the residuals of the fitted model should be independent and uncorrelated, and the independence shall not be effected by transformations such as squaring. As we can see in Figure 7, the residuals of the upper regime are correlated, and this can be an indicator of lack of fit of the model.

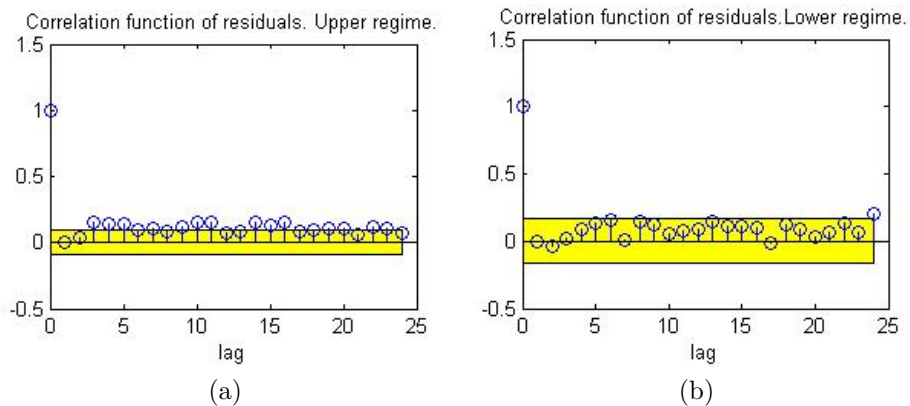


Figure 7: Correlation function of residuals for upper and lower regimes

If we move further on and check the squared residuals we observe that the squared residuals are not white noise and therefore the TAR model does not catch any of the ARCH/GARCH characteristics of the series. See Figure 8.

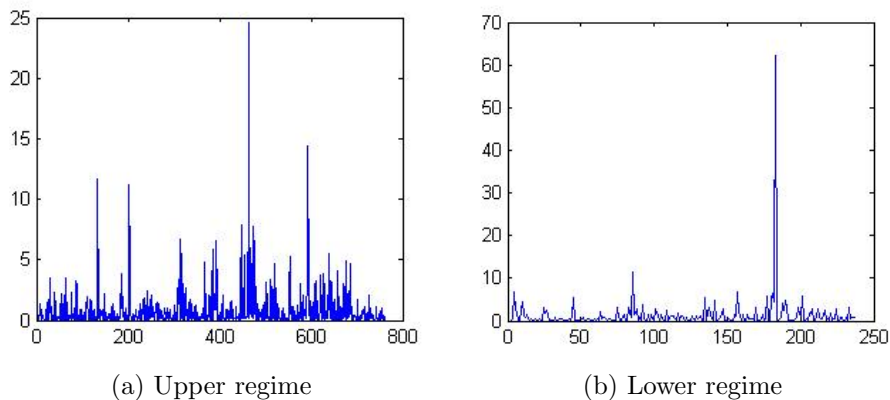


Figure 8: Squared residuals plotted for lower and upper regimes of the estimated TAR model

4.2 Real Series

4.2.1 The Blowfly Data

The plot of the blowfly data, presented in Figure 9 shows that either the original series or the transformed one, doesn't follow a regular linear series.




Figure 9: The Blowfly data

As it can be seen in Table 3, the tests of non-linearity show that the Blowfly population does not follow a linear model, and the likelihood ratio test for threshold nonlinearity (TLRT test) shows that a TAR model can be fitted to the data.

Table 3: Results of Blowfly data pre-test for non-linearity

Tests	Original		Logarithm	
	p-value	test statistics	p-value	test statistics
ARCH/GARCH Test				
Using up to lag 15	0.178	19.970	1.000	1.250
Using up to lag 20	0.355	21.730	0.842	13.760
Using up to lag 25	0.425	25.660	0.963	13.900
LBQ Test				
Using up to lag 15	0.007	31.790	0.000	48.99
Using up to lag 20	0.000	63.110	0.000	113.5
Using up to lag 25	0.000	68.240	0.000	116.7
Keenan Test	0.000	48.800	0.000	29.70
Tsay Test	0.000	16.07	0.000	10.87
TLRT Test	0.000	27.26	0.000	89.59

In Tong [12], it is suggested that the linear model applied in the article lacks in fit because it fails to catch two very important biological features, namely that there is a threshold due to food limitation and that there is a delay due to development time of the population till the flies mature.

By using these arguments we fit a TAR model to the data. Here, we try to fit TAR models to both transformed ³ and untransformed data, and we get TAR(2; 1, 2) and TAR(2; 2, 2) models respectively, with threshold one at point 914.113, $\log_{10}914.113 = 2.961$, and one at 4404.

Table 4: Result of TAR model fit for logarithm transformed Blowfly data

Model: TAR(2,1,2)								
Delay parameter: 8								
Estimated Threshold: 2.961								
The threshold is searched from the 15 percentile to the 85 percentile of all data.								
The estimated threshold is the 17.6 percentile of all data.								
AIC: -319.8								
Coef.	Lower Regime				Upper Regime			
	Est.	Std.Err	t-stats	p-value	Est.	Std.Err	t-stats	p-value
Intercept	2.706	0.183	14.81	0.000	0.263	0.074	3.551	0.004
Lag1	0.264	0.051	5.147	0.000	1.261	0.057	22.07	0.000
Lag2	0.000	0.000	0.000	0.000	-0.343	0.059	-5.817	0.00
Nr of Observations is 64				Nr of Observations is 298				
Residual Std.Err = 0.143			R = 0.999		Residual Std.Err = 0.154		R = 0.998	
F-stat (df = 2,60) = 19998			p-value = 0		F-stat (df=3,288) = 46149		p-value = 0	

The first threshold can be explained as the point where the population starts increasing more rapidly as a consequence of high population of mature flies:

$$y_t^* = \begin{cases} 2.706 + 0.264y_{t-1}^* + e_{t1}^* & \text{if } y_{t-1}^* \leq 2.961 \\ 0.263 + 1.261y_{t-1}^* - 0.343y_{t-2}^* + e_{t2}^* & \text{if } y_{t-1}^* > 2.961 \end{cases} \quad (21)$$

where e_{t1}^* and e_{t2}^* are the error processes with specifications in Table 4. Here $y_{t-i}^* = \log_{10}^{y_{t-i}}$ for $i = 0, 1, 2$.

But the second threshold can be interpreted as the point where lack of food can cause a noticeable decline in the process of population increase:

$$y_t = \begin{cases} 852.9 + 0.905y_{t-1} - 0.122y_{t-2} + e_{t1} & \text{if } y_{t-1} \leq 4404 \\ 105.9 + 1.281y_{t-1} - 0.351y_{t-2} + e_{t2} & \text{if } y_{t-1} > 4404 \end{cases} \quad (22)$$

where e_{t1} and e_{t2} are the error processes with specifications in Table 5.

³Logarithm in base 10 is used for the transformation.

4.2.2 The River Flow in Northern Scotland

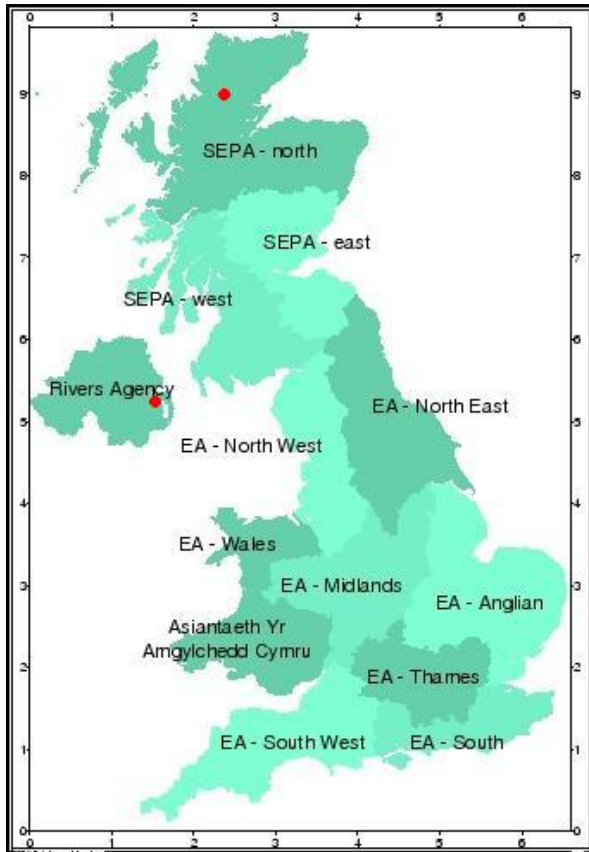


Figure 10: UK, Scotland and Ireland

Here we apply the TAR model concept to the two rivers described earlier. The catchment area for Oykel, the upper red dot, is 330.7 km^2 . For Enler River, the catchment area is significantly smaller, 61.8 km^2 . Here, catchment area refers to drainage basin which is an area where rainwater and melting water from snow and ice drains downhill towards a river, lake, reservoir, sea or ocean. The drainage basin includes all land and waterways available that conveys the water to the specific point in that drainage basin to where the water gathers up, the river of Oykel or Enler in this case. One drainage basin is separated from another one by the topography, settlement, mountains, forest etc.

To the left side, a map over the UK, Scotland and Ireland shows the position (red dot) on the top centre of the Eastern Turnaig Oykel river flow and Comber station that measure the river of Enler close to Ireland. This region and perhaps specifically Scotlands climate is to a large degree defined by the Gulfstream

which passes just north of Scotland. Due to the Gulfstreams position, Scotlands climate is mild and warmer than other places with similar positioning in the world. In general, the climate could be described as being very changeable. Rain followed by sun or vice versa is rather common throughout the day as well as wide variations in climate over short distances. The months that have the lowest precipitations are May and June.

The Oykel River

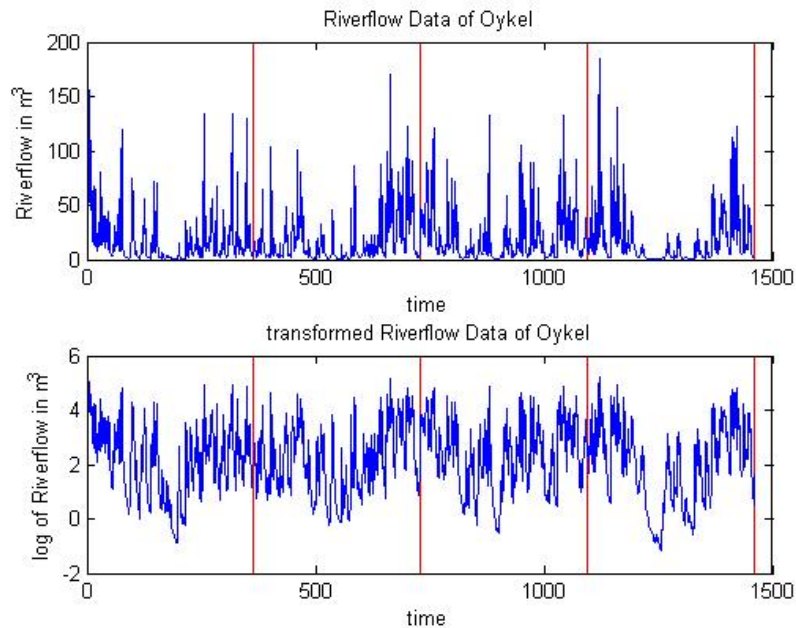


Figure 11: Oykel riverflow data taransformed and untransformed

The time series is plotted in its original, untransformed shape and in its transformed shape. We see that the original series exhibit great variations, such as cycles, on a yearly basis. Further, the cycles seem to indicate very low activity around the months of May to June and, except for these months, the rivers behaviour is very temperamental the rest of the year.

The behaviour of the Oykel River and other rivers in this part of Scotland are probably greatly influenced by the catchment area through which they flow and the capacity of the catchment area is to a large extent influenced by the climate at a given time of year. We assume that a very intensive rainfall will result in an inability of the catchment area to absorb all the water and thus lead to some sort of surface runoff, where water flows more or less rapidly, pending on the topography of the area, towards the natural gather point to, e.g. the river. In that case, it would mean that the rainfall exceeded some sort of critical point to where the soil could handle, absorb all the water. The catchment capacity is then full and the response of the soil accounting for all the water may result in surface runoffs, and probably a greater chance of flooding and erosions.

In the opposite case, when no or very little precipitation has occurred for a period of time, droughts would be more probable to occur. The event of passing the threshold isn't longer that probable because the capacity of the soil to absorb the water will have drained by some kind of constant demand.

This diversity of climate at different time points of the year as well as the complexity of the structure of water flows in the catchment area may point towards the need to apply a TAR model to the data for the reasons described. Firstly, the climate in this part of Scotland is very erratic, with very high peaks and low downs. Secondly, if the level of the river is triggered by a threshold, the soils inability to absorb water, it seems natural to consider a TAR model with this feature inbuilt.

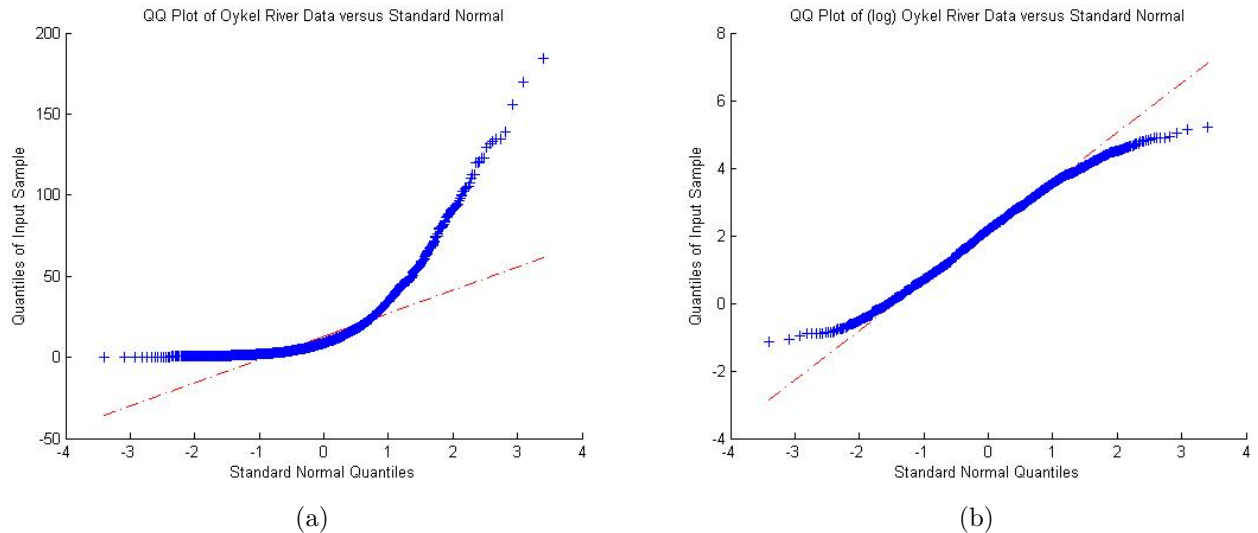


Figure 12: Q-Q plot of Oykel river data, raw and transformed

We perform a logarithmic transformation to the series in order to attempt to:

- a) Make a better fit to our model.
- b) Normalize the data in respect to the variance, e.g. make it more homoskedastic.
- c) To increase the importance of small values and to reduce the impact of large values.

The transformation of the logarithm of the river data result in the skewness and the kurtosis to change drop drastically from 2.55 and 11.19 to -0.12 and 2.33, respectively. We expect this to have a large effect in the analysis, improving the fit and normalize the data, although the remaining data display a leptokurtic shape.

Following the steps in the model specification part, we start with the simplest model. We have already seen the plot of the data. Next step is to have a look at the SACF and the SPACF.

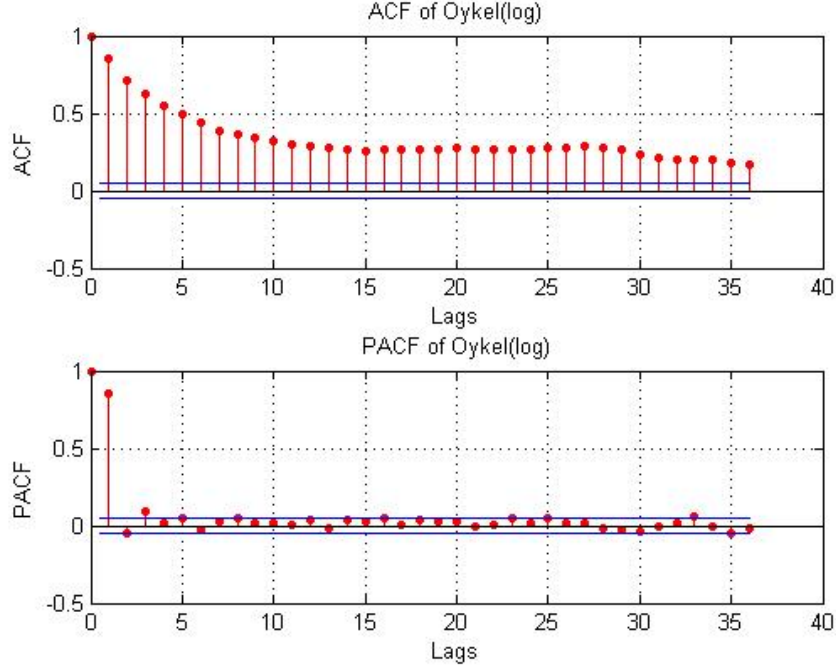


Figure 13: Sampled Autocorrelation and Sampled Partial Autocorrleataion of the log-transformed Oykel river data

We see that it is an Autoregressive process. We examine the ACF as well as the PACF plots of the logarithmic Oykel river flow series and get a rough visual estimate of the number of parameters we shall include in the ARMA model. Initially one can say one or perhaps two or even three AR parameters. The number of MA parameters will initially be set to zero, however this may change later on.

Table 6: Results of the fitted ARMA models to Oykel river data

	Model 1	Model 2	Model 3	Model 4	Model 5
ARMA(p,q)	(1,0)	(2,0)	(3,1)	(4,0)	(3,2)
AIC	3041	3040	3023	3030	3018
SSR	682.7	680.9	670.7	673.8	667.5
RMSE	0.684	0.683	0.679	0.680	0.677

We tried to fit many different ARMA models to the data and some even including MA parameters. We also considered the AR(2) but a closer examination of the ACF of the residuals revealed serial correlation at different lag intervals so this model was rejected. The model with the best fit in respect to AIC, RMSE and SSR is the last estimated model. It is an ARMA(3, 0) model.

However, it is possible to drop the Moving Average term without a too big of a loss in goodness-of-fit.

Balancing between these two models, ARMA(3, 1) and AR(3), is perhaps a matter of personal preference. However, we choose the AR(3) for a couple of reasons, one being that the model containing only three coefficients is more parsimonious, that is, it is simpler to interpret than the other one that contain four parameters. Other reasons are that the residuals of the AR(3) model are as good as those from the ARMA(3, 1) model. From this perspective, it seems more reasonable to choose the AR(3) above. We find no evidence of ARCH/GARCH effects, and we decide to leave the ARMA model as that.

Table 7: Results of the fitted ARMA(3,0) model to Oykel river data

Coefficients	Estimates	Std. Error	Test Statistics	p-value
Intercept	0.293	0.036	8.214	0.000
AR1	0.904	0.026	34.71	0.000
AR2	-0.140	0.035	-3.994	0.000
AR3	0.097	0.026	3.723	0.000
RMSE=0.680 SSR = 674.1 AIC = 3028				

Table 8: Results of pre-test for non-linearity of Oykel river data

Tests	p-value	test statistics	critical value
ARCH/GARCH Test			
Using up to lag 15	0.867	9.19	24.99
Using up to lag 20	0.951	10.7	31.41
Using up to lag 25	0.976	13.0	37.65
LBQ Test			
Using up to lag 15	0.838	9.69	24.99
Using up to lag 20	0.769	15.1	31.41
Using up to lag 25	0.802	18.8	37.65

The diagnostics checking of the residuals and the overall fit look well. We do not suspect non-linearity at this time. The hypothesis of GARCH effect is rejected. Our residuals look like noise. However, not to be surprised later on, we do test for non-linearity. The tests are carried on from the AR(3) in Table 9.

Table 9: Results of test for non-linearity of Oykel river data

Tests	p-value	test statistics
Keenan Test	0.000	105.1
Tsay Test	0.000	16.27
TLRT		
Setting p=3, d=1	0.000	112.30
Setting p=3, d=2	0.005	33.53
Setting p=3, d=3	0.008	21.64

Here we see the result from the test for non-linearity. We see that they all indicate the presence of some kind of non-linearity at 5% significance level. Keenan and Tsay test indicate that quadratic nonlinearity is present, TLRT that we may have a potential threshold model with two regions.

As we at this point only consider GARCH and TAR models to model non-linearity, we are probably at risk of misspecification. We have ruled out the GARCH since diagnostics did not support this kind of non-linearity in the variance. We decide to estimate a TAR(k) model and in the following step we enable a grid search to find potential thresholds in the series.




Figure 14: Oykel river flow, raw and transformed plot of potential thresholds vs SSR

The threshold grid search procedure, TGSP, described earlier, plots the results of the TAR models of the middle 70% observations against their SSR to find potential thresholds. The threshold corresponding to the smallest residual sum of squares is chosen. We add that when searching for the threshold in the untransformed data we did not

yield any satisfying result, no through or dip was encountered in the plot that indicated a possible threshold.

The second figure of the logarithmic transformation does indicate several potential thresholds. Interpretation is not straightforward without a confidence interval; a sudden drop of a large magnitude could suggest that there exist several thresholds. The first around $y = 0.8$, the second around $y = 1.2$, a third around $y = 1.5$, and the fourth around $y=2.6$.

We first fit a TAR (2, 1, 1) model with delay set to 1 to the Oykel River flow data set. Then we fit another model of form TAR (2, 2, 1), $d=1$ and continue in this fashion by adding and removing parameters until we are satisfied. Our criterion for goodness-of-fit is the Akaike criterion in combination with SSR, RMSE, numbers of parameters etc. Simply put, we look for the most parsimonious model.

Table 10: Results of the fitted TAR models to transformed Oykel river data

TAR(k)	Model 1 (2; 1, 1)	Model 2 (2; 2, 1)	Model 3 (2; 3,1)	Model 4 (2; 1, 2)	Model 5 (2; 2, 2)
Threshold	0.827	1.497	1.497	0.8207	1.497
Observation	268;1193	479; 982	481; 908	268; 1193	479;982
Delay	1	1	1	1	1
AIC	2983	2932	2930	2975	2928
SSR	665.1	636.0	636.5	660.9	634.4
RMSE	0.675	0.661	0.661	0.673	0.660
LB-Q test (p-values)					
Using up to lag 15	0.049	0.089	0.093	0.614	0.408
Using up to lag 20	0.082	0.103	0.108	0.612	0.352
Using up to lag 25	0.151	0.206	0.209	0.677	0.480
ARCH test (p-values)					
Using up to lag 15	0.859	0.943	0.943	0.871	0.929
Using up to lag 20	0.937	0.971	0.970	0.938	0.963
Using up to lag 25	0.972	0.981	0.980	0.969	0.974
J-B test (p-values)	0.000	0.000	0.000	0.000	0.000

In addition to the models in Table 10, we also estimated various models with different delay parameters. None of them resulted in a better fit. Examining the outcomes of the results in Table 10 closer, it seems like the models with an estimated threshold of 1.5 all have better goodness of fit compared to the models that were estimated with a threshold of 0.8. The models above are all rather similar in terms of model fit. We also fit another model, a TAR(2; 2, 3), and let it represent the final choice of model. The main reason for choosing the TAR model below is that the diagnostics yielded the best result.

Table 11: Result of the best fitted TAR model to the transformed Oykel river data

Model: TAR(2,2,3)								
Delay parameter: 1								
Estimated Threshold: 1.497								
The threshold is searched from the 15 percentile to the 85 percentile of all data.								
The estimated threshold is the 32.9 percentile of all data.								
AIC: 2918 SSR: 629.4 RMSE: 0.657								
Coef.	Lower Regime				Upper Regime			
	Est.	Std.Err	t-stats	p-value	Est.	Std.Err	t-stats	p-value
Const.	0.202	0.039	0.189	0.000	0.739	0.085	8.727	0.000
Lag1	1.906	0.118	16.089	0.000	0.767	0.032	24.33	0.000
Lag2	-0.823	0.101	-8.115	0.000	-0.137	0.036	-3.806	0.000
Lag3	—	—	—	—	0.086	0.028	3.079	0.002
Nr of Observations is 481					Nr of Observations is 980			
Residual Std.Err = 0.603			R = 0.741		Residual Std.Err = 0.684		R = 0.946	
F-stat (df = 3,476) = 453.28			p-value = 0		F-stat (df=4, 975) = 4238		p-value = 0	

From the output of the selected model we firstly start off with transforming the threshold back to its original scale. The threshold translates into circa 4.47 cubic metres of water.

This can be interpreted as being the critical level of what the catchment basin can support in terms of water storage in the soil before the excess water beyond the threshold of 4.47 cubic meters start responding with surface runoffs. If the total amount of water, including melted snow and ice, rain and hail exceeds that level, the predicted level of water will pass the threshold and the consequences result in a higher probability of floods, erosions and landslides since the water will have no way to go.

As seen in the time series plot of the Oykel river flow, once the threshold is exceeded, the response of the rivers behaviour is very episodic and erratic. The upper regime is to a larger extent more influenced by the previous periods. This is perhaps because of when more persistent rain periods occur, the water does not disappear as fast as in periods of less rain.

The storage capacity of the catchment gets more and more filled up and as the rain falls over the entire area, the first water reaches the river within a short period of time. Water that falls further away has to travel longer, through different layers of soil, and the effect of all these factors may result in a longer delay.

On the other hand, when the amount of water lie under the threshold, the predicted water response is fairly predictable, not varying too much. During the dry month, with lower precipitation, it seems like whenever the threshold is crossed, the catchment response doesnt collapse.

In between the heavy rainfalls during the dry middle months, the soil is able to drain the water falling into the catchment area and keep the storage of water emptier than during rainy periods. This is because of the soils constant demand for water is higher, probably due to vegetation, underlying canals etc, than the total amount of precipitation in the area so the threshold is not being crossed. In the lower regime, the present amount of cumecs is greatly influenced by the two previous periods. Warmer weather and evaporation makes the water level less dependent to previous precipitation.

$$Y_t^* = \begin{cases} 0.202 + 1.906y_{t-1}^* - 0.823y_{t-2}^* + e_{t1}^* & \text{if } y_{t-1}^* \leq 1.497 \\ 0.739 + 0.767y_{t-1}^* - 0.137y_{t-2}^* + 0.086y_{t-3}^* + e_{t2}^* & \text{if } y_{t-1}^* > 1.497 \end{cases} \quad (23)$$

where e_{t1}^* and e_{t2}^* are the error processes with specifications in Table 11. Here $y_{t-i}^* = \log_{10}^{y_{t-i}}$ for $i = 0, 1, 2, 3$.

The Enler river

Another river we will have a look at is the river of Enler. The series consists of 1461 observations here as well. In between 27/8-2008 to 3/9-2008, a total of 8 observations were missing from the dataset. Missing data is probably unavoidable in the long run, and the deciding upon a strategy for handling them will aid in the analysis. When faced with missing data there are a few options available. Dependent on the dataset in use, its size and characteristics, there are a few things to do.

One can simply ignore them and delete them, which lead to loss of information in one or another way.

The most common thing to do is to use various imputation methods. Imputation is a replacement method where the missing value is substituted in on or another way.

There is a wide range of different ways to do the imputation. One can replace the missing value with the mean of the entire dataset. This approach is likely to be good if the data set doesn't display too much variety. We apply an imputation technique called hot deck imputation, Tsikriktsis [14]. The hot deck imputation method is simple; here we substitute the missing data by real data from a similar case of the original set for which the data isn't missing.

For instance, for this dataset we consider replacing the missing data with real data from periods with similar record, here from September to February. Since the frequency of missing data only is about 0.5 % we do not expect imputation in this case to be a potential threat to our analysis. We assume that the values to be replaced have no negative effect in the analysis.

Table 12: Result of Hot Deck Imputation

Day	Month	Year	Value
27	8	2008	0.604
28	8	2008	0.656
29	8	2008	0.696
30	8	2008	0.709
31	8	2008	0.762
1	9	2008	0.845
2	9	2008	0.961
3	9	2008	1.010

We begin our analysis by plotting the river in its untransformed and transformed shape. This river has a much lower water flow in cumecs than does Oykel River. One very noticeable feature of the river that struck us is that the middle 4-6 months are rather calm, with year 2007 as an exception. When the precipitation does occur, the rivers response seems irregular and somewhat sporadic throughout the year, almost clustering up. It should be remembered that the river is quite small, about one tenth of the size of Oykel in terms of cumecs. The catchment area is also a lot smaller, being about 5 times smaller than that of the Oykel.

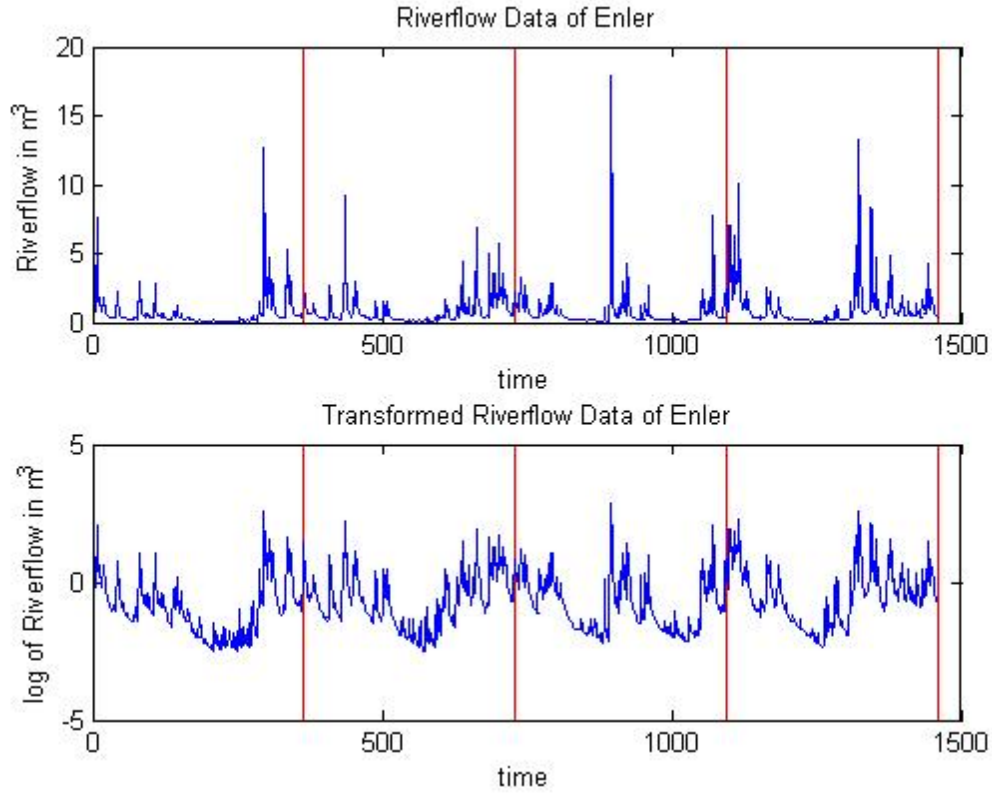


Figure 15: Enler river data, original and transformed data

First impression of this series that struck us is the resemblance of a GARCH model. Later on, we shall see if the GARCH model fit the series.

As in the previous analysis of Oykel, after having inspected ACF and PACF, the initial autoregressive order to include is about three. We evaluate as before, various combinations of model settings and select the best fitting linear model.

Table 13: Results of log of the Enler River

	Model 1	Model 2	Model 3
ARMA(p, q)	(1, 0)	(2, 0)	(3, 0)
AIC	1680	1678	1669
SSR	269.0	268.1	265.9
RMSE	0.429	0.428	0.427

The model we select is the first one. It seems reasonable, having only one parameter and have as much explanatory power as the other ones. The residuals are checked for disturbances, we check the squared residuals as well. From the tests, we can see that

there may be something fishy at lag 25 for the LB-Q test, which should be investigated further. A visual inspection could reveal more than the numerical test, so we plot the residuals from the AR(1) model.

Table 14: Test of residuals for ARCH/GARCH effect

Tests	p-value	test statistics	critical value
ARCH/GARCH Test			
Using up to lag 15	0.210	19.07	24.99
Using up to lag 20	0.393	21.07	31.41
Using up to lag 25	0.287	28.44	37.65
LBQ Test			
Using up to lag 15	0.082	23.08	24.99
Using up to lag 20	0.080	29.40	31.41
Using up to lag 25	0.020	41.51	37.65

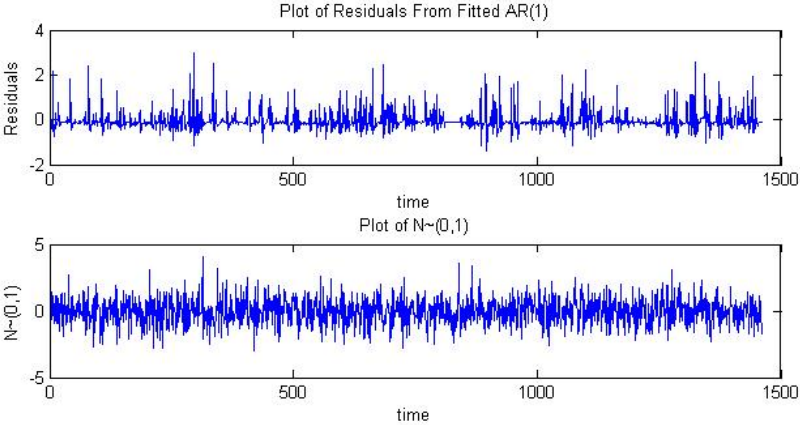


Figure 16: Residuals of Enler river data from AR(1)

An instant impression tells us that some kind of structure is present. The residuals do not seem normally distributed, rather clustering up sporadically. We suspect that arch effect or some other mechanism is in present in the data. However though, there is little support for arch effect. We will estimate an AR(1)-GARCH(1, 1) and see if we can catch the underlying dynamics.

Table 15: Result of AR(1)-GARCH(1, 1) fit to Enler river data

Parameter	Estimate	Sts.Error	Test Statistic	p-value
C	-0.054	0.020	-2.635	0.013
AR(1)	0.923	0.012	74.88	0.000
K	0.026	0.004	6.266	0.000
GARCH(1)	0.799	0.029	27.35	0.000
ARCH(1)	0.057	0.010	5.300	0.000
AIC = 1661				

A model taking the form AR(2)-GARCH(1,1) was also estimated. Its AIC is 1649. When dividing the Innovations with the Sigmas, there should be a white noise if the model is correct for the data; unfortunately it is not, residuals from this model are not normally distributed as well. More GARCH coefficients were estimated but the residuals didn't yield normality.




Figure 17: Innovations divided by Sigmas from AR(1)-GARCH(1, 1)

The structure does not resemble normal homoskedastic process. We conclude that the AR-GARCH models do not work for this particular data set. Having ruled out AR-GARCH, the only one remaining to test and estimate is the TAR model. The tests for non-linearity indicate that quadratic nonlinearity is present as well as a possible threshold alternative. The AR model under the null hypothesis is tested against the alternative, a TAR model with autoregressive order one in each regime and delay parameter equal to one.

Table 16: Results of test for non-linearity

Test	p-value	Test-stat
Keenan Test	0.000	62.45
Tsay Test	0.000	37.38
TLRT		
Setting p=1, d=1	0.000	50.7840

With the threshold search algorithm enabled, we search for thresholds in both the raw

and transformed data. As seen in the plots, no visible threshold is present, just a slope. A threshold would look like a single through in the plot.

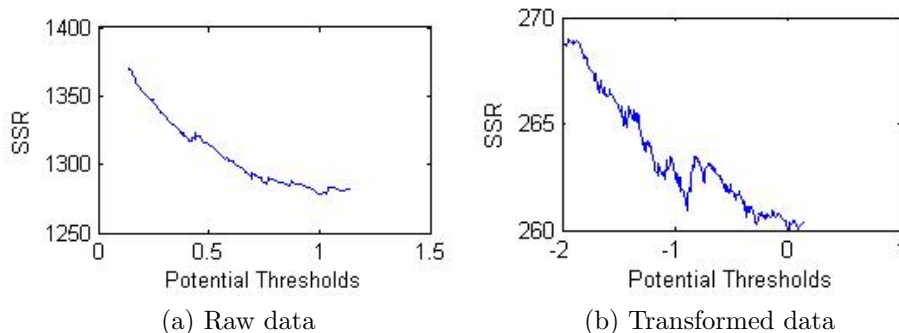


Figure 18: Plot of potential thresholds vs SSR, original and transformed Enler river flow data. There is some indications of a potential threshold at the value $y = -0.955$.

Table 17: Result of TAR model fit

Model: TAR(2, 1, 1)								
Delay parameter: 1								
Estimated Threshold: -0.955								
The threshold is searched from the 15 percentile to the 85 percentile of all data.								
The estimated threshold is the 44.8 percentile of all data.								
AIC: 1584 SSR: 129.9 RMSE: 0.269								
	Lower Regime				Upper Regime			
Coef.	Est.	Std.Err	t-stats	p-value	Est.	Std.Err	t-stats	p-value
Const.	-0.090	0.055	-1.611	0.1075	-0.045	0.017	-2.702	0.007
Lag1	00.917	0.033	27.6922	0.0000	0.789	0.023	33.54	0.000
Nr of Observations is 665					Nr of Observations is 796			
Residual Std.Err = 0.354			R = 0.954		Residual Std.Err = 0.473			R = 0.587
F-stat (df=2, 651)=6969			p-value = 0		F-stat (df=2, 804)=572			p-value = 0

Since we have no support for the GARCH, we conclude that the best model would be the linear AR(1) or the piecewise linear model TAR(2, 1, 1) $d=1$. The residuals of the TAR model are still subject to serial correlation at certain lags, but less than from the AR(1) model.

In deciding which model to choose, considering that both models are subject to serial correlation, and that the AR(1) model is simpler than the TAR model perhaps the AR seems as a better choice. One aspect that the TAR model captures, which is interesting in this case, is the threshold opposed to the AR. The TAR model allows us to more

precisely inference on the rivers behaviour so we choose this one. The regimes are separated by the threshold of 0.385 cumecs.

When the rivers flow is below the threshold, it is easier to predict. The soil has more capacity to store the water and rain doesnt have a very large impact on the river flow. When the river flow exceeds this threshold and enters another state, the behavior of the river change. In this new state, drastic changes are expected from whenever precipitation occurs and the rivers response becomes much harder to predict.

5 Conclusions and Discussions

5.1 Conclusions

After performing the diagnostic tests and modelling procedures, there were some results obtained for the simulated data sets and for real data sets.

For the simulated data, we were successful in fitting a GARCH model to the simulated TAR series and a TAR model to the simulated GARCH series. Although it was possible to fit another model to the real data one should pay careful attention about the choice of model before finally making a selection. Numerical measures of goodness of fit is without hesitation of utmost importance when fitting a model, but there are other criterias to consider as well, such as the relevance of the model to the data to be fitted.

There were two different TAR models fitted to the Blowfly data while there was no ARCH effect detected and therefore no ARCH/GARCH model could be fitted to the data.

To a certain degree the result of the Oykel and the Enler data indicate that the threshold models may be appropriate for modelling and analysing these two rivers.

For the Oykel, the TAR model performed slightly better even though the AR(3) model was simpler. Any GARCH effects were not supported for.

The second river, Enler, was also tested for thresholds. Here, we encountered ARCH effects but when investigating it closer, the fitted AR-GARCH model didn't give satisfactory result. Here we ultimately compared the AR and the TAR model and finally choose the simple TAR model which performed better than the other models.

It's realistic to think of the cumecs in the river as a result of different degrees of rain precipitation. Depending on many factors, such as layers of soil, the distance that the water has to travel to the river, the catchment will respond differently, even to a similar precipitation depending on if the previous time periods where subject to rainfall or not. Depending on the current state, once the crossing of the threshold is made, the catchment enters a new state, to where its now unstable and the probability of landslides, erosions and floods increases.

5.2 Discussions

A potential question in this inference for the Blowfly data is whether the upper regime of the logarithmic threshold and the lower regime of the untransformed threshold collapse on each other so that there is a three regime TAR model.

We should also take into consideration that the population dynamic cannot be well explained with a simple TAR model, because it is quite difficult to choose the delay and threshold which best explains the data and perhaps a more complicated function of threshold and delay parameter is needed for modeling.

Another way to go further into modeling of the population data sets such as Blowfly data is to consider a dynamic threshold variable in which the switching threshold value while going from lower regime to upper regime is different from the threshold value of going back from upper regime to lower regime.

While fitting a TAR model to the data, we need to estimate the unknown threshold value and consequently we need to find the confidence interval for the threshold. The problem we face for calculating the critical values of the threshold variable is that we should know the distribution of the threshold variable which is not a standard distribution. There are bootstrap methods for finding the distribution, and a solution similar to finding the critical values for maximum likelihood ratio tests for finding the critical values. See Enders [7], Hansen [10] and Joliffe & Jones & Garthwaite [9].

In the fitting of TAR models to rivers, it would be interesting to estimate a three regime TAR model that has the dynamic to model whether the catchment are in drought, medium rain precipitation or intensive rain precipitation. For this, we would have to engage in the challenge to program an algorithm, more complicated, that would need further restrictions.

The Oykel River, in Northern Scotland showed evidence of multiple regimes. This may point towards a three-regime TAR model, where the lower regime describe the river system in times of drought, a middle regime that describe the river system in periods of medium precipitation and finally the upper regime that would describe the river system in times of heavy rain downfall. A three regime TAR model would more accurately describe the behaviour of the rivers respond to different degrees of rainfall, or seasons of the year, of the catchment area.

References

- [1] Jarque C. Bera A. Efficient tests for normality, heteroskedasticity and serial independence of regression residuals: Monte carlo evidence. *Economics Letter*, 7:313–318, 1981.
- [2] T. Bollerslev. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrics*, 31:307–327, 1986.
- [3] K. S. Chan. Percentage points of likelihood ratio tests for threshold autoregression. *Journal of the Royal Statistical Society*, 53:691–696, 1991.
- [4] K. S. Chan. Consistency and limiting distribution of the least squares estimator of a threshold autoregressive model. *The Annals of Statistics*, 21:520–533, 1993.
- [5] K. S. Chan. *TSA: Time Series Analysis*, 2008. R package version 0.97.
- [6] Chan K.S. Cryer, J. *Time Series Analysis [electronic resource]: With Applications in R*. Springer Science + Business Media, New York, 2008.
- [7] W. Enders. *Applied Econometric Time Series*. John Wiley Sons, third edition, 2004.
- [8] Robert F Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, July 1982.
- [9] Joliffe I. Jones B. Garthwaite, P. *Statistical Inference*. Prentice Hall, London, 1995.
- [10] B. E. Hansen. Inference in tar models. 1996. Forthcoming in *Studies in Non-linear Dynamics and Econometrics*.
- [11] H. Tong. *Threshold Models in Non-linear Time Series Analysis*. Springer, New York, 1983.
- [12] H. Tong. Non-linear time series modelling in population biology: A preliminary case study. *Lecture Notes in Control Information Science*, 106, 1988.
- [13] H. Tong. *Non-Linear Time Series :A Dynamical System Approach*. Oxford University Press, Clarendon New York, 1990.
- [14] N. Tsikriktsis. A review of techniques for treating missing data in om survey research. *Operations Management*, 24:53–62, 2005.

A MATLAB routine for TGSP

```
function[tau b minr y1 y2]=tauest(y)
% Gets the time series as input and estimates the threshold value
% tau <-- the threshold value
% b <-- the coefficients of the AR models estimated in each regime
% minr <-- the residuals corresponding to the estimated threshold
% y1 <-- the elements of the series in lower regime
% y2<-- the elements of the series in upper regime

[a ind]=sort(y);
mini=ceil(.15*length(a));
maxi=ceil(.85*length(a));
z1=[];z2=[];
residuals=ones(1,maxi-mini+1)*1000000;
size(residuals);
for i=mini:maxi
    for j=1:length(y)-1
        if y(j)<a(i)
            z1=[z1;y(j+1) y(j)];
        else
            z2=[z2;y(j+1) y(j)];
        end
    end
    [b1(i,:) bint1 r1]=regress(z1(:,1),[ones(size(z1(:,2))) z1(:,2)]);
    [b2(i,:) bint2 r2]=regress(z2(:,1),[ones(size(z2(:,2))) z2(:,2)]);
    z1=[];
    z2=[];
    residuals(1,i-mini+1)=sum(r1.^2)+sum(r2.^2);
end
min(residuals(1:end));
[stres ind2]=sort(residuals(:));
minr=residuals(ind2(1));
tauind=ind2(1)+mini-1;
tau=a(tauind);
b=[b1(tauind,:);b2(tauind,:)];
figure
plot(a(mini:maxi),residuals(:))
hold on
xlabel(' Sorted \ity_t');
ylabel('Residuals')
hold off
y1=y(y<tau);
```

```

y2=y(y>tau);
end

```

B Matlab routine for simulating a TAR model

```

function[y]=tar(p1,p2,tau,d,n)
% p1: arma coefficient vector for the lower regime
%(The first value should be the constant)
% p2: arma coefficeint vector for upper regime(The first value should be
% the constant)
% tau: threshold value
% d: the delay parameter
% n: number of simulations wanted
p1n=length(p1);
p2n=length(p2);
p=max(p1n,p2n);
y=[rand(d,1);zeros(n-d,1)];
lagy=zeros(n,p);
lagy=[ones(n,1),lagmatrix(y,[1:p])];
lagy(isnan(lagy))=0;
figure
hold on
for j=d+1:n
    if y(j-d)<tau
        y(j)=sum(lagy(j,1:p1n)*p1')+normrnd(0,.25);
        lagy=[ones(n,1),lagmatrix(y,1:p)];
        lagy(isnan(lagy))=0;
        plot(j,y(j),'*')
    else
        y(j)=sum(lagy(j,1:p2n)*p2')+normrnd(0,.25);
        lagy=[ones(n,1),lagmatrix(y,[1:p])];
        lagy(isnan(lagy))=0;
        plot(j,y(j),'r*')
    end
end
end
hold off
figure
y=y;
y2=y.^2;
y3=y.^3;
save y y;
load y;

```



```

figure
subplot(2,1,1); plot(y); hold on; plot(1:length(y),tau,'r');
title('yt series')
subplot(2,2,4); plot(sort(y));hold on; plot(1:length(y),tau,'r');
title('yt series sorted ascending')
subplot(2,2,3); plot(tau,y,'r'); hold on;
for i=d+1:length(y)
hold on
plot(y(i-d),y(i));
end
hold off
figure
subplot(2,2,1); autocorr(y,36,0,1.96); title('ACF of y')
subplot(2,2,2); parcorr(y,36,0,1.96); title('PACF of y')
subplot(2,2,3); autocorr(y2,36,0,1.96); title('ACF of y^{2}')
subplot(2,2,4); autocorr(y3,36,0,1.96); title('ACF of y^{3}')
end

```