Monitoring Structural Change in Variance, with an Application to European Nominal Exchange Rate Volatility

Frédéric Carsoule*
Tinbergen Institute
Erasmus University Rotterdam

Philip Hans Franses
Econometric Institute
Erasmus University Rotterdam

June 23, 1999

Econometric Institute Report 9925/A

Abstract

In this paper we propose a sequential testing approach for a structural change in the variance of a time series, which amounts to a procedure with a controlled asymptotic size as we repeat the test. Our approach builds on that taken in Chu, Stinchcombe & White (1996) for structural change in the parameters of a linear regression model. We provide simulation evidence to examine the empirical size and power of our procedure. We apply our approach to 14 weekly observed European exchange rates for 1985-1998 and we find ample evidence for the presence of structural changes in nominal exchange rate volatility, where generally a reduction of volatility is found.

Keywords: Structural change, variance, monitoring procedure, exchange rate volatility

*Corresponding author: Tinbergen Institute, Erasmus University Rotterdam, P.O.Box 1738, NL-3000 DR ROTTERDAM, phone: +31 10 4088931, carsoule@few.eur.nl

1 Introduction

An important empirical question in international finance concerns the stability of nominal or real exchange rate volatility. Particulary, it is of interest to investigate whether exchange rate volatility differs across various monetary regimes. Examples of relevant studies are Baxter & Stockman (1989), Eichengreen (1988) and Stockman (1983), among others. In this paper we also focus on the stability of exchange rate volatility. We narrow our focus to nominal rates, as they appear to govern the behavior of real rates, see for example Mussa (1986). We consider European exchange rates (all against the German Mark) for weekly data from 1985-1998 in order to examine if the European Monetary Union (EMU) had an effect on country-specific exchange rate volatility.

Before we turn to the empirical investigation, we first propose a novel sequential testing method that is relevant for our purposes. We need this new method for the two following reasons. Several tests for a structural change in the variance at some point in time assume prior knowledge of its location. For our empirical problem, however, it will be very difficult (if not impossible) to pinpoint the exact date that the EMU has become relevant for each specific country. One possibility is now to sequentially apply tests for structural change to a range of possible breakpoints. The major drawback of this approach, however is that the size of the test can not be controlled. Therefore, in this paper we decide to build on the approach taken in Chu et al. (1996) by adapting their method to our situation, that is, to monitoring structural change in the variance. This results in a procedure with controlled asymptotic size as we repeat the test. Given its importance we decide to discuss this method first and to evaluate its empirical performance in substantial detail.

The outline of our paper is as follows. In Section 2, we propose our approach to monitoring a structural change in the variance. The null hypothesis concerns an iid sequence with constant variance. In Section, 3 we evaluate the empirical size and power of our method. As it is sometimes found that exchange rates can be described by Autoregressive Conditional Heteroskedasticity (ARCH) type models, we propose an alternative test that can handle ARCH behavior, and we evaluate its empirical size using Monte Carlo simulations to see if the empirical size gets much affected. In Section 4, we apply our method to 14 European nominal exchange rates. Our main finding is that volatility seems to have decreased for most exchange rates. We conclude our paper in Section 5 with some remarks.

2 Monitoring structural change in variance

In this section we propose a method for monitoring structural change in variance. We first outline the monitoring procedure in section 2.1, and next we discuss the specific test we use in this procedure in section 2.2. Finally we examine how our method should be adapted in case the data display simple ARCH patterns.

2.1 The monitoring procedure

We consider the problem of testing whether a parameter in a statistical model remains constant over time. Most tests in the literature concerning change point problems are retrospective tests, that is given a set of observations, it is the aim to decide if a change has occured within a span of data see for example Andrews (1993) and Ploberger, Kramer & Kontrus (1989), among many others. In contrast a monitoring procedure for structural change allows us to check parameter constancy, any time new data become available. The purpose of the monitoring procedure described in Chu et al. (1996) is to detect a change in a parameter that can occur at any time. It is assumed that one has m observations, called the historical data set. Then we gather data and we want to be able to check at any moment if the model estimated for the historical data set fits the new data. As it is notified in Chu et al. (1996), repeating a retrospective test of structural change each time we get new observations leads us to reject the hypothesis of constancy with a large probability, even if no change actually occurs. Consequently, such a procedure will have a poor size.

A way to avoid this problem is to apply a sequential approach. Such a procedure can be described as follows. Given that m is fixed, we define, for every $n \geq m$ a statistic, T_n, called the detecting statistic or the detector, such that, under the hypothesis of no change, T_n may cross a boundary function $g(n/m)$, for some $n \geq m$, with small probability (0.05 or 0.10). Conversely, we expect T_n to cross $g(n/m)$ with a large probability in case of true change.

More precisely, the procedure can be viewed as a stopping rule. Let $X_1, \ldots, X_n$ be the sequence of observations available at time n and let ψ_n the σ-field generated by $X_1, \ldots, X_n$. We define the stopping time τ as

$$\tau = \inf\{n \geq m, \ |T_n| > g(m,n)\}. \tag{1}$$

This τ is adapted to the filtration, $\psi = \{\psi_n\}_{n \geq 1}$. The monitoring procedure will be stopped as soon as the process T_n is crossing g and τ will therefore take a finite value. If H_0 is the null hypothesis that a parameter remains constant over time, and if H_1 is the alternative that a change occurs for some $n \geq m$, we want that

$$P\{\tau < \infty | H_0\} = \alpha \tag{2}$$

with for example $\alpha = 0.05$ or 0.10, and

$$P\{\tau < \infty | H_1\} = 1 \tag{3}$$

In order to guarantee the requirement (2), which ensures a good size of the procedure, we use Theorem (3.4) in Chu et al. (1996). This theorem does not give an exact relation for a fixed m, but an approximation based on weak convergence, for m large enough. Denote S_n the process such that $T_n = m^{-1/2} S_n$. We assume that the sequence of processes $\{m^{-1/2} S_{[mt]},\ t \in [0,\infty)\}$ converges in law to a standard Wiener process, W. Then for a convenient function g, under certain conditions, the above theorem states that $P\{\tau < \infty | H_0\}$ is approximately equal to the probability that the absolute value of a Wiener process crosses at least once the path of the function g. More precisely, under H_0

$$P(|S_n| \geq \sqrt{m} g(n/m),\ \text{for some } n \geq m) \simeq P(|W(t)| \geq g(t),\ \text{for some } t \geq 1) \tag{4}$$

We will use $g(t) = [t(a^2 + \ln t)]^{1/2}$ as the boundary function, see also Chu et al. (1996). This function satisfies the conditions of the theorem and for $a = 7.78$ (6.25), α in (2) will be equal to 0.05 (0.10).

2.2 CUSUM of Squares Test

We wish to monitor the stability of the variance of a time series. Consider first the case that there is no time-dependence between the observations. Let $\{X_n\}_{n\geq 1}$ be an independent sequence of normal random variables with mean 0 and variance σ_n^2. First we assume that

$$\sigma_1^2 = \cdots = \sigma_m^2 \equiv \sigma_0^2.$$

Next, observing new data, we want to detect if a change occurs in the variance. That is, we want to test the null hypothesis

$$H_0 : \sigma_n^2 = \sigma_0^2, \quad n = m+1, \ldots$$

against the alternative

$$H_1 : \sigma_n^2 \text{ changes at some } n \geq m+1.$$

The statistic T_n used for our monitoring procedure is a Cumulative Sum of Squares as proposed in Brown, Durbin & Evans (1975). Inclan & Tiao (1994) use a centered version of this statistic for a retrospective test on a change of the variance of an iid sequence of centered random variables.

Denote $C_k = \sum_{i=1}^k X_i^2$ as the cumulative sum of squares of the sequence $\{X_t\}$, and denote the centered (for $n \geq m$) cumulative sum of squares as

$$D_n = \frac{C_n}{C_m} - \frac{n-2}{m-2}.$$

The process mD_n is equal to $\sum_{i=1}^n \{X_i^2/\hat{\sigma}^2 - 1\} - 2\ (n-m)/(m-2)$, where $\hat{\sigma}^2$ is the Maximum Likelihood Estimator of σ^2 based on the historical dataset

$X_1, \cdots, X_m$, that is $\hat{\sigma}^2 = m^{-1} \sum_{i=1}^{m} X_i^2$. Denote $S_n = 2^{-1/2} m D_n$. Then

$$\begin{aligned} S_n = & \ 2^{-1/2} \sum_{i=1}^{n} \left(\frac{X_i^2}{\hat{\sigma}^2} - 1 \right) - 2^{1/2} \frac{n-m}{m-2} \\ = & \ 2^{-1/2} \sum_{i=m+1}^{n} \left(\frac{X_i^2}{\hat{\sigma}^2} - \frac{m}{m-2} \right) \end{aligned} \tag{5}$$

Inclan & Tiao (1994) show that under H_0,

$$m^{-1/2} \ S_{[mt]} \Rightarrow W^0(t), \tag{6}$$

where "$\Rightarrow$" denotes weak convergence and where $W^0(t) = W(t) - t\, W(1)$, with W is a standard Wiener process. They prove this result for $t \in [0,1]$. However, it naturally can be extended to $[0, \infty)$.

The asymptotic process itself is not a Wiener process, but a similar approximation as the approximation in (4) can be deduced using

$$\left\{ W^0(t) \ , \quad t \geq 1 \right\} \stackrel{d}{=} \left\{ (t-1) W(t/(t-1)) \ , \quad t \geq 1 \right\},$$

where "$\stackrel{d}{=}$" denotes equality in law. Denote $b(t) = t/(t-1)$ and $a(t) = t-1$. We have therefore

$$\begin{aligned} & P \left\{ |S_n| \geq \sqrt{m} (g \circ b)(n/m) \times a(n/m) \ , \quad \text{for some } n \geq m \right\} \\ & \quad \simeq P \left\{ |W(t)| \geq g(t) \ , \quad \text{for some } t \geq 1 \right\} \end{aligned} \tag{7}$$

As $b(t) = t/(t-1)$, $b(n/m) = n/(n-m)$ and $a(t) = t-1$, $a(n/m) = (n-m)/m$. Next we have

$$\begin{aligned} & g(t) = [t(a^2 + \log t)]^{1/2}, \text{ then} \\ & g(b(n/m)) = g(n/(n-m)) = \left[n/(n-m)(a^2 + \log(n/(n-m))) \right]^{1/2} \end{aligned}$$

and given these results, we obtain

$$\sqrt{m} g(b(n/m)) a(n/m) = \sqrt{(n-m)/m} \sqrt{n} (a^2 + \log(n/(n-m)))^{1/2}. \tag{8}$$

Finally, we obtain the approximation

$$\begin{aligned} & P \left\{ |S_n| \geq \sqrt{(n-m)/m} \sqrt{n} (a^2 + \log(n/(n-m)))^{1/2}, \text{ for some } n \geq m \right\} \\ & \quad \simeq P \left\{ |W(t)| \geq g(t) \ , \quad \text{for some } t \geq 1 \right\} \end{aligned}$$

It is this approximation which we will use in our empirical work below.

2.3 Autoregressive Conditional Heteroskedasticity

As we aim to apply our method to financial time series, consider now a model where the conditional variance depends on the past. In its basic form, the sequence $\{X_n\}$ can be described as $X_n = Y_n\sqrt{h_n}$, where $\{Y_n\}$ is an iid sequence of standard normal variables and h_n is the conditional variance of Y_n, $h_n = \omega_n + \alpha X_{n-1}^2$, with $\omega_n > 0$ and $0 < \alpha < 1$. This model is called an ARCH(1) model, see Engle (1982). The parameter α remains constant over time. The intercept in the conditional variance equation depends on n but we assume that it is constant up to time m,

$$\omega_1 = \cdots = \omega_m \equiv \omega_0.$$

We now aim to detect if a change occurs in ω when new data are observed, that is, we want to test the new null hypothesis

$$H_0 : \omega_n = \omega_0, \quad n = m+1, \ldots$$

against the alternative

$$H_1 : \omega_n \text{ changes at some } n \geq m+1.$$

We will use the same threshold as for the iid case but we now have to use a different detector that takes into account that the variance at time n depends of the square of the previous observation X_{i-1}^2.

The parameters α and ω are estimated from the historical data set, $X_1, \ldots, X_m$ by the maximum likelihood estimators, $\hat{\alpha}$ and $\hat{\omega}$. Denote by $\hat{h}_i$ the estimate of the conditional variance at time i, that is $\hat{h}_i = \hat{\omega} + \hat{\alpha} X_{i-1}^2$. The detecting statistic is then defined by

$$S_n = 2^{-1/2} \left(m^{-1} \sum_{i=1}^{m} \frac{1}{\hat{h}_i^2} \right)^{-1/2} \sum_{i=m+1}^{n} \frac{1}{\hat{h}_i} \left(\frac{X_i^2}{\hat{h}_i} - 1 \right) \tag{9}$$

We will reject the null hypothesis as soon as the process S_n crosses the boundary given in (8).

Before we apply our monitoring procedure for a structural change in variance to weekly exchange rate data, we first consider its empirical performance in the next section using Monte Carlo simulations.

3 Monte Carlo evidence

In this section, we present results from simulation experiments to examine the finite sample behavior of the monitoring procedure. Size and power are first investigated for an independent sequence of normal variables. Next the size and power of the test is examined for an ARCH(1) process.

3.1 Finite sample size for an iid sequence

We generate 5000 replications of an iid sequence drawn from a standard normal distribution and we compute the frequency that the detecting statistic crosses at least once the boundary functions defined by (8). The statistic S_n, as defined in (5), has mean zero but its distribution is not symmetric as its median is negative. Unreported simulation results show that introducing a negative bias provides a better size. In case of a decrease of the variance, the finite sample power will obviously be improved and in case of an increase, it will not be substantially reduced. Therefore, instead of using S_n defined by (5), we use $\tilde{S}_n = S_n - \sqrt{2}m/(m-2)$ which is equivalent to $1/\sqrt{2}\sum_{i=1}^{n}(X_i^2/\hat{\sigma}^2 - m/(m-2))$.

We consider asymptotic sizes of 0.05 and 0.10, which correspond with choosing $a = 7.78$ and $a = 6.25$ in (8). We compare the performance of the test under the null hypothesis for $m = 25, 50, 75, 100, 200, 500$ and 800. The approximation (4) says that the probability of crossing holds for some $n \geq m$. We calculate the frequency of rejection of the null hypothesis for some n lying between m and q, where q is 2,4,6,9 and 19 times m. The results are summarized in Table 1.

From Table 1, we can draw the following conclusion. For asymptotic size of 0.05, the empirical size gets close to the asymptotic size for $m \geq 200$ and $q \geq 6m$. Interestingly, for asymptotic size of 0.10, the monitoring procedure is a little conservative for large m and q.

3.2 Finite sample power

To study the finite sample power of the test procedure, we simulate a shift in the variance of the sequence of independent random variables at time $t = m \times 1.1$ in the same way as is done in Chu et al. (1996). As the distribution of the process S_n is not symmetric, we examine the two cases where the variance decreases and increases. Denote by σ_1^2 the variance of X_i after the change. We consider the case $\sigma_1^2 = 0.5$ and the case $\sigma_1^2 = 2$. The results are reported in Table 2.

The finite sample power results in Table 2 confirm the asymmetry in the distribution of S_n. Generally, when the variance decreases the power is smaller, then when it increases (with a similar sized break). For $m \geq 100$ and $q \geq 4m$ the empirical power is close to 1. Hence, for practical purposes, it seems useful to consider such values for m and q.

Figures 1 and 2 present the delay needed by the procedure to detect the change, that is, the delay needed by the detecting statistic to cross the boundary function. Here we just consider the case of the boundary function corresponding with a critical value of 10%. The average and the median are computed, given that the procedure detects the change. As before, a change of the variance is simulated at time $t = m \times 1.1$ The delay of detection will depend on the size of the change. We consider the cases where after the change the variance is less than 1 with the new variance varying between 0.01 and 0.7, and where the variance is larger than 1 with the new variance varying between 1.2 and 8. We consider $m = 100$ and 500 and $q = 19m$. The x-axis of the graphs gives the variance after the change and the y-axis the ratio of the delay of detection and m.

The graphs in Figures 1 and 2 can be interpreted as follows. For example, consider the upper panel of Figure 1, where a value 0.5 on the x-axis corresponds with values of 1.6 and 0.9 on the y-axis. The value 0.5 means that if the variance of the observations shifts from 1 to 0.5, the average number of observations needed for detection is $1.6m$ and the median number is $0.9m$, given that the procedure detects it. In this special case, from Table 2, we can see that for a shift from 1 to 0.5, m=100, asymptotic size of 0.10 and a monitoring horizon of $19m$, the probability that the test procedure rejects the null hypothesis is 0.97.

3.3 Finite sample size for an ARCH(1) process

Finally, we consider the test when the test statistic is modified to account for ARCH(1) behavior, see (9). The data generating process used for the simulation is

$$X_i = Y_i\sqrt{h_i}, \text{ where } h_i = 0.6 + 0.4X_{i-1}^2,$$

and Y_i is an iid sequence of standard normal variable. As in the iid case, we will not use S_n as defined in (9), but instead we use

$$\tilde{S}_n = 2^{-1/2}\left(m^{-1}\sum_{i=1}^{n}\frac{1}{\hat{h}_i^2}\right)^{-1/2}\sum_{i=1}^{n}\frac{1}{\hat{h}_i}\left(\frac{X_i^2}{\hat{h}_i} - \frac{m}{m-2}\right).$$

Once again, the use of this statistic will improve the finite sample size.

We compute the test for $m = 75, 100, 200$ and 500, and for the same monitoring horizon as for the iid case. The number of replications is 2500. The results are reported in Table 3.

From the entries of Table 3, we notice that the ARCH(1)-corrected monitoring procedure seems to have reasonable size for $m \geq 200$ and $q \geq 4m$. When m gets very large, the procedure results in a conservative test.

In order to investigate the power of the monitoring procedure in the ARCH case, we consider again a change in the parameter ω at time $t = m \times 1.1$. We examine the case where the ω after the change is equal to half the initial value, that is $\omega = 0.3$, and the case where it is two times the initial value, ω=1.2. The results are reported in Table 4.

The results in Table 4 show that the empirical power is moderately large to close to unity for $m \geq 100$ and $q \geq 4m$. Notice again that the test is more powerful in case of an increase in the variance than it is for a decrease.

4 European exchange rate volatility

In this section we apply our monitoring procedure for a structural change in the variance to 14 weekly observed nominal European exchange rates (all against the German mark). In Section 4.1, we provide some details on the data and the research method. In Section 4.2, we apply our monitoring procedure and discuss the main empirical results.

4.1 The data and research method

We consider 14 nominal European exchange rates. The first 8 exchange rates concern currencies that (partly) constitute the Euro currency. There are the Finnish Markka, the Dutch Guilder, the Belgian Franc, the French Franc, the Italian Lire, the Spanish Peseta, the Portuguese Escudo and the Austrian Shilling. The relevant data for these currencies (that is, their returns) appear in the upper left panel of Figures 3 to 10. The next 6 exchange rates are not part of the Euro. These are the UK Pound, the Danish Kroner, the Swedish Kroner, the Norwegian Kroner, the Greek Drachme and the Swiss Franc. The relevant data appear in Figures 11 to 16. For each exchange rate, we calculate the weekly returns as the first differences of the logged exchange rates.

Our sample contains weekly data that cover 1985 to 1998, which amounts to 726 observations. During this period, the European Monetary Union was founded. Although the currency of this EMU, that is, the Euro, was officially introduced in January 1999, way before that date it was known that the EMU would result in a single currency. Central banks in the European countries involved likely adjusted their policies given this knowledge, although it is uncertain when they started to do so exactly. Hence, it would be difficult to have a precise indication of a week or period which may correspond to a potential structural break in country-specific policies of central banks. Of course, the hypothesis is that the founding of the EMU and the predictable introduction of the Euro would have led to a reduction in the variance of European exchange rates. For the countries that are not part of the EMU one may expect that the EMU also would lead to a variance reduction, although perhaps to a lesser extent.

In the Figures 3 to 16, we will present the following graphs. As we set m= 100, the first graph in the upper left panel contains 626 weekly exchange rate returns. We decide to set m at 100 as our simulation results suggest that this choice leads to good size and power properties. We experimented with $m = 75$, but this resulted in rather fragile results, and with $m = 150$, but this led to approximately similar outcomes as obtained for $m = 100$. The monitoring procedure is based on setting $q = 6m$. The first graph of the exchange rates themselves immediately indicates for many series that sometimes there are substantial outliers.

In the second graph, in the upper right panel, we give the results for the monitoring method as discussed in Section 2.2. From various graphs we can see that outliers have a significant impact on the CUSUM test statistic. Therefore, in the bottom two panels of each graph we give the CUSUM test results when we use methods that may make the test more robust to those outliers. The first, in the bottom left panel, is the CUSUM test statistic when corrected for ARCH(1) patterns. The second, in the bottom right panel, is the monitoring test when applied to outlier-corrected data. The procedure for this correction is the following. We calculate the Median Absolute Deviation (MAD) for all the data points. Next, we delete those observations for which the genuine data point is 4 times the MAD. The numbers of observations that are deleted are given in Table 5.

4.2 Empirical results

The graphs in the Figures 3 to 16 lead to the following conclusions, where we first deal with the Euro countries.

When we do take care of the (obvious) outliers in the Finnish Markka, the CUSUM test suggests that there was a significant increase in the variance. However, when these outliers are set at the MAD value, the bottom right panel of Figure 3 suggests that was no change in the volatility. For the Dutch Guilder we observe for each of the three methods that there was a significant level shift downwards. A similar conclusion can be drawn for the Belgian Franc in Figure 5. When we take care of outliers, the 10% critical value gets hit towards the end of the sample. The graphs for the French Franc suggest that the variance for the period after 1985-1986 (which roughly corresponds with $m = 100$) differs from that for those two years. However, if we take care of outliers, the evidence for such a volatility change seems to disappear. In contrast to these four series, Figure 7 shows that the Italian Lire experienced a significant increase in volatility. This conclusion holds for all three methods, that is, whether we take care of ARCH(1) behavior and of outliers or not. For the Spanish Peseta in Figure 8 we observe that the CUSUM test tends to indicate that there was a significant increase in volatility. However, when we correct for outliers, or when we do not take count of ARCH, this evidence is not very strong. This is in contrast to the results in Figure 9 for the Portuguese Escudo, which clearly indicate a significant reduction in variance. Finally, for the Austrian Shilling one hardly needs a statistical test to indicate that there was a reduction in volatility. Indeed, even a casual look at Figure 10 suggests such a reduction for the last few years. Reassuringly, our monitoring procedure confirms this visual evidence.

In sum, for the Euro countries we generally find (except for Italy) that nominal exchange rate volatility experienced structural shifts downwards. In other words, the last few years have witnessed a significant decrease in the variance. This seems to confirm our conjecture.

We now turn to analyzing the volatility of exchange rates of countries that do not participate in the Euro. The results in Figure 11 for the UK Pound show that all CUSUM tests point towards a reduction in volatility (that is, when comparing 1987-1998 with 1985 and 1986). In contrast, for the Danish Kroner all tests suggest that there is no structural change in volatility, although there is a tendency downwards toward the end of the sample. The test results in Figure 13 for the Swedish Kroner all less easy to interpret. First, there is an indication that there is a reduction in variance in 1987-1988. Next, there is a large outlier, which seems to precede a period with larger volatility. Indeed, the basic and the ARCH(1)-corrected CUSUM test indicate this increase in variance. It seems that the outlier removal method results in deleting too many observations that are informative for the rather obvious upward shift in variance. The Norwegian Kroner, for which the test results are depicted in Figure 14, seems to display a significant reduction in volatility. Finally, for the Greek Drachme we would conclude that there is reduction in volatility, although deleting 43 observations seems to lead to different conclusions. For the Swiss Franc we are tempted to conclude that there is some evidence for a variance reduction in 1988 and 1989, but that this does not seem to hold for the rest of the sample.

5 Conclusion

In this paper we proposed a monitoring procedure to examine a structural break in the variance of an iid process or an ARCH(1) process. We showed through simulations that our procedure has good empirical size and power properties. We applied our procedure to 14 weekly European exchange rates (against the German mark). Our main finding was that all rates (except for Italy) have witnessed a significant decrease in variance or appeared to have a stable variance. The evidence for a variance decrease was most convincing for currencies that constitute the Euro.

One potentially fruitful area where our method can be useful concerns the construction of linear or nonlinear (G)ARCH models. For some exchange rates, we could reject the constancy of the conditional variance. As it is well known that neglecting such a shift can lead to a substantial estimation bias in the (G)ARCH parameters, see for example Lamoureux & Lastrapes (1990), our method can be used as a simple diagnostic tool for the adequacy of a (G)ARCH model. Additionally, if our method suggests a variance shift during a certain period of time, it may be useful to specify nonlinear GARCH models in which parameters experience sudden or smooth changes.

Table 1: Empirical size for an iid sequence

Asymptotic size 0.05

m	25	50	75	100	200	500	800
q							
$2m$	6.3	5.6	5.2	4.9	4.6	4.4	4.3
$4m$	7.3	6.4	5.6	5.5	5	4.8	4.7
$6m$	7.7	6.7	5.8	5.7	5.2	4.9	4.8
$9m$	8	6.8	6	5.9	5.3	5	5
$19m$	8.3	7	6.3	6	5.5	5.1	5.1

Asymptotic size 0.10

q							
$2m$	7.8	7.5	7.1	7.1	7.3	7.2	7.3
$4m$	9	8.6	7.6	7.9	8	8	8
$6m$	9.5	8.9	8.3	8.3	8.3	8.3	8.3
$9m$	9.9	9.2	8.6	8.5	8.6	8.5	8.5
$19m$	10.3	9.5	8.9	8.8	8.8	8.7	8.7

Table 2: Finite sample power of CUSUM of squares monitoring

Decrease of variance from 1 to 0.5:

Asymptotic size 0.05

m	25	50	75	100	200	500	800
q							
2m	1.5	9.3	34.3	60.2	97.9	100	100
$4m$	1.7	22.3	57.9	81.5	99.7	100	100
$6m$	2	29.7	66.7	87.4	99.9	100	100
$9m$	2.4	35.3	72.2	90.5	99.9	100	100
$19m$	3.2	42	77.9	93	99.9	100	100

Asymptotic size 0.10

q							
$2m$	3.0	26.3	57	77.9	99.3	100	100
$4m$	7.0	45.8	77	91.5	99.9	100	100
$6m$	9.8	53.7	82.7	94.4	99.9	100	100
$9m$	12.4	59.2	86.4	95.9	99.9	100	100
$19m$	15.6	65.2	89.5	97.2	99.9	100	100

Increase of variance from 1 to 2:

Asymptotic size 0.05

m	25	50	75	100	200	500	800
q							
$2m$	54.8	76.3	89	95.4	99.9	100	100
$4m$	63.9	85.6	95.2	98.5	100	100	100
$6m$	67.7	88.6	96.8	99	100	100	100
$9m$	70.2	90.4	97.6	99.4	100	100	100
$19m$	72.8	92.3	98.4	99.7	100	100	100

Asymptotic size 0.10

q							
$2m$	60.0	80.9	92.1	96.9	100	100	100
$4m$	69.3	89.4	96.9	99.1	100	100	100
$6m$	72.9	91.8	98.0	99.4	100	100	100
$9m$	75.3	93.2	98.5	99.7	100	100	100
$19m$	77.8	94.6	99.0	99.8	100	100	100

Table 3: Empirical size for an ARCH(1) process

Asymptotic size 0.05

m	75	100	200	500
q				
$2m$	10.3	8.2	5.8	3.5
$4m$	11.4	8.9	6.3	3.9
$6m$	11.9	9.1	6.5	4.1
$9m$	12.2	9.6	6.8	4.1
$19m$	12.6	10.0	6.8	4.3

Asymptotic size 0.10

q				
$2m$	13.4	10.8	8.7	7.2
$4m$	15.3	11.9	9.5	7.8
$6m$	15.7	12.4	9.8	8.1
$9m$	16.1	12.9	10.0	8.4
$19m$	16.5	13.4	10.3	8.7

Table 4: Finite sample power for an ARCH(1) process

Variance change ω=0.6 to ω=0.3:

Asymptotic size 0.05

m	75	100	200	500
q				
$2m$	32.6	47.2	88.4	100
$4m$	46.4	63.1	96.0	100
$6m$	51.7	69.2	97.3	100
$9m$	55.6	73.1	98.0	100
$19m$	60.2	77.1	98.8	100

Asymptotic size 0.10

q				
$2m$	47.5	61.9	93.6	100
$4m$	61.4	76.1	97.9	100
$6m$	66.0	80.6	98.8	100
$9m$	70.0	83.8	99.2	100
$19m$	73.6	87.1	99.4	100

Variance change ω=0.6 to ω=1.2:

Asymptotic size 0.05

m	75	100	200	500
q				
$2m$	68.2	76.5	96.2	100
$4m$	78.8	87.0	98.8	100
$6m$	82.8	89.8	99.4	100
$9m$	84.8	92.1	99.7	100
$19m$	87.4	93.7	99.8	100

Asymptotic size 0.10

q				
$2m$	74.7	82.8	97.6	100
$4m$	84.6	91.0	99.5	100
$6m$	87.3	93.4	99.8	100
$9m$	89.4	94.8	99.8	100
$19m$	91.1	96.0	99.9	100

Table 5: Number of outliers deleted by countries

UK Pound	6	Belgian Franc	38
Danish Kroner	21	French Franc	26
Swedish Kroner	17	Italian Lire	64
Norwegian Kroner	22	Spanish Peseta	35
Finnish Markka	16	Portuguese Escudo	30
Dutch Guilder	14	Greak Drachme	43
Swiss Franc	3	Austrian Shilling	44

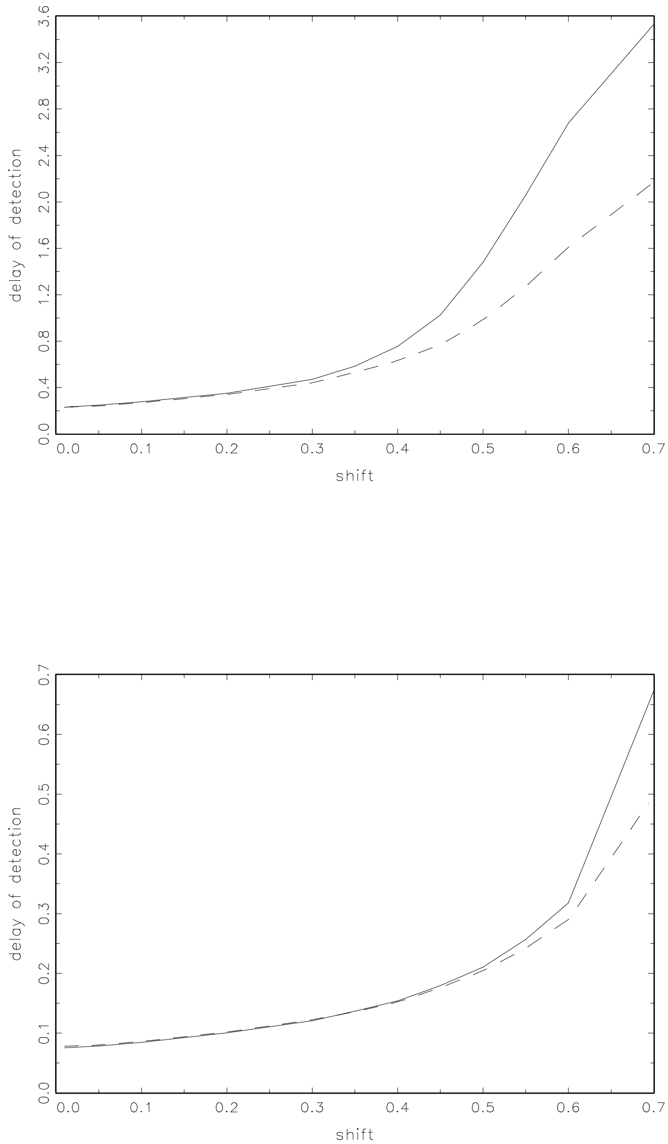

Figure 1: Delay of detection for m=100 and m=500 in case of a decrease of the variance ($q = 19m$)
straight line: average values of the delay, dashed line: median values of the delays

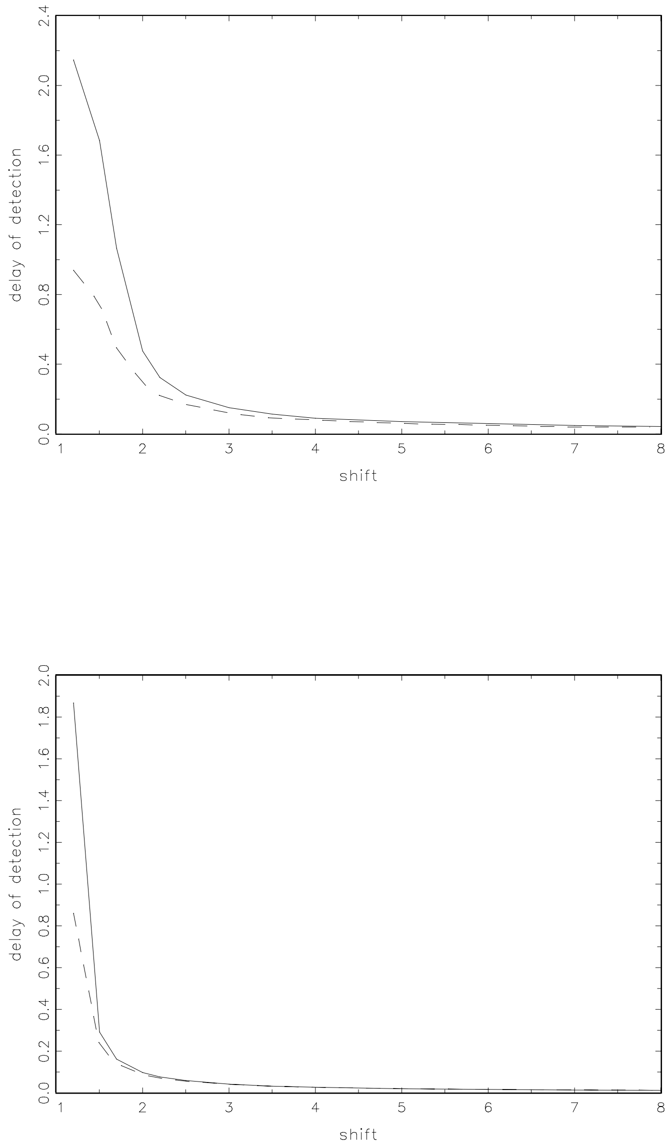

Figure 2: Delay of detection for m=100 and m=500 in case of an increase of the variance ($q = 19m$)
straight line: average values of the delay, dashed line: median values of the delays

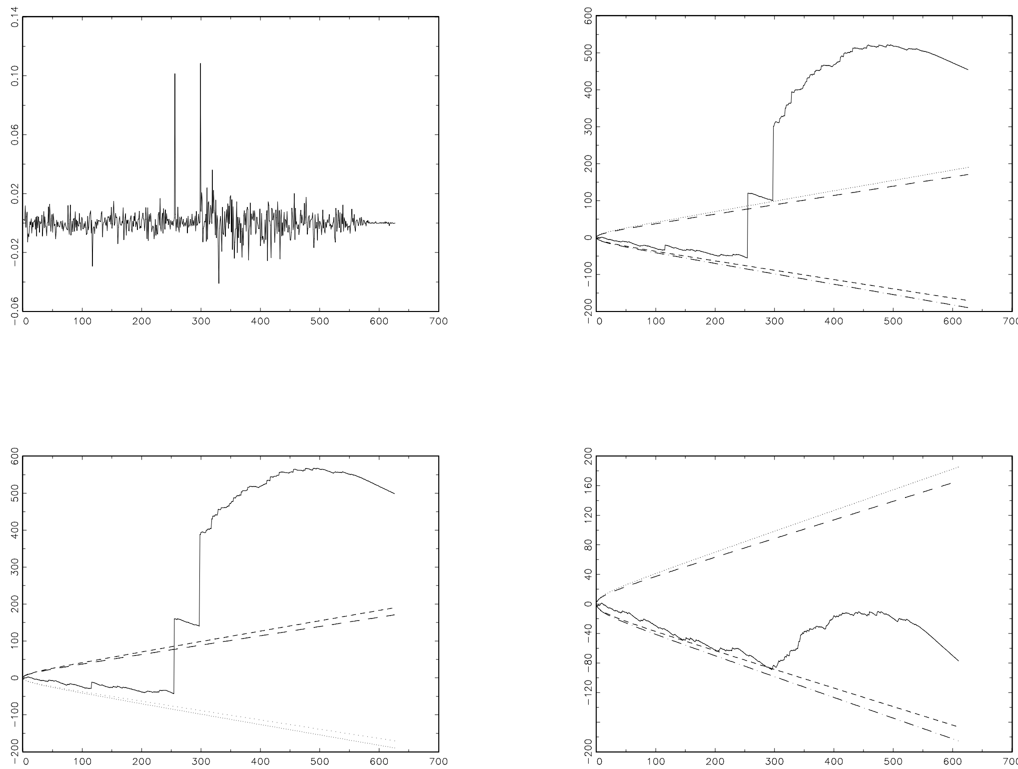

Figure 3: Finnish Markka
the dotted line corresponds with the 5% critical value, and the dashed line with 10% critical value.
upper left panel : nominal exchange rate returns
upper right panel : test without ARCH
bottom left panel : test with ARCH
bottom right panel : test while some observations have been deleted

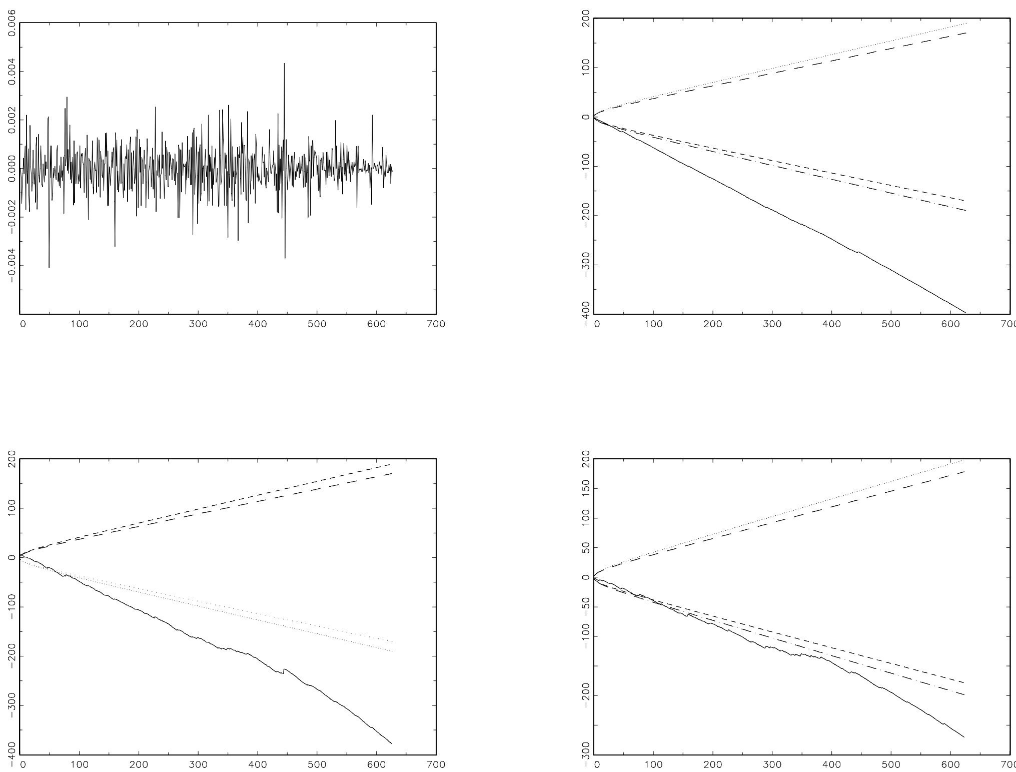

Figure 4: Dutch Guilder

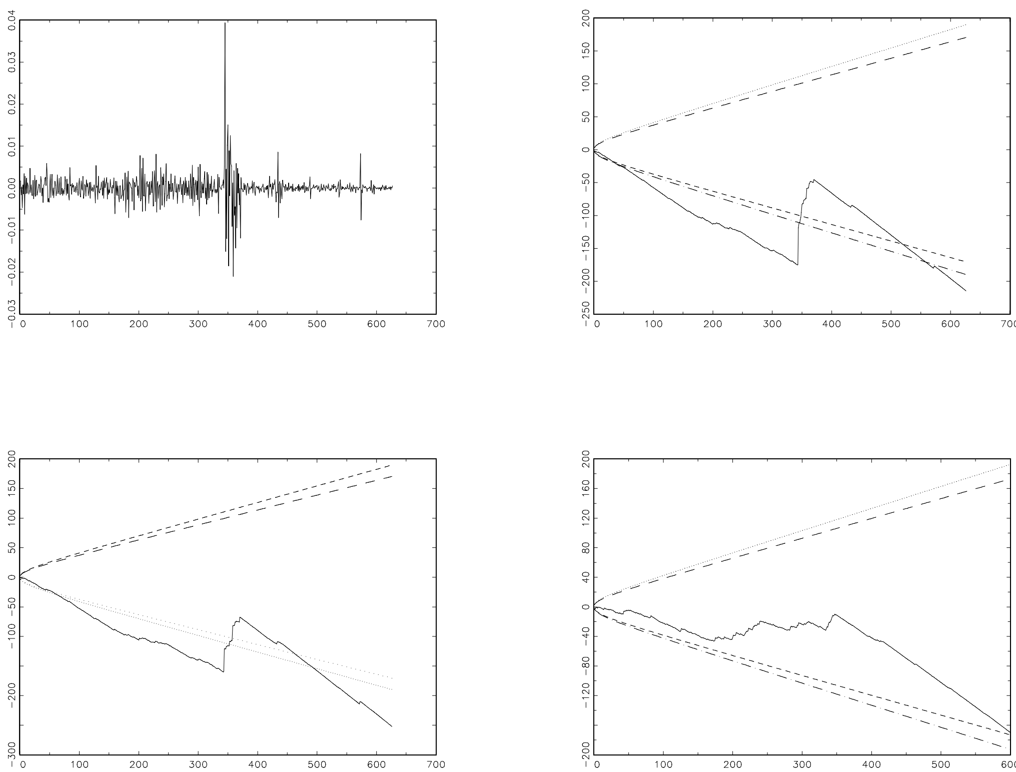

Figure 5: Belgian Franc

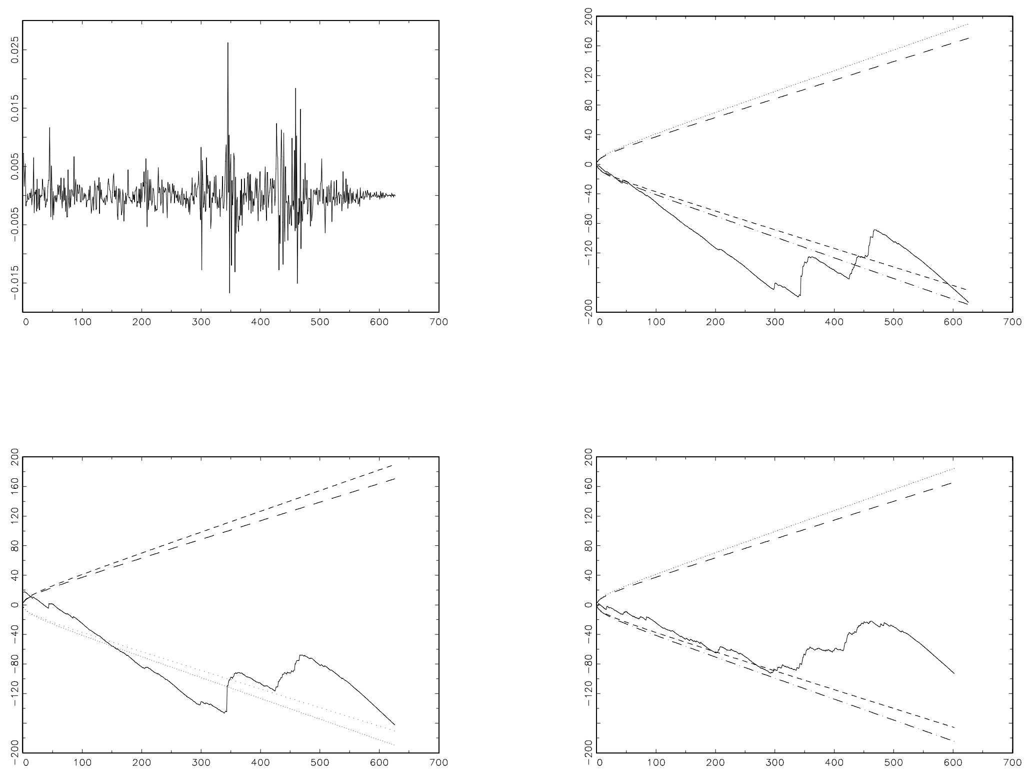

Figure 6: French Franc

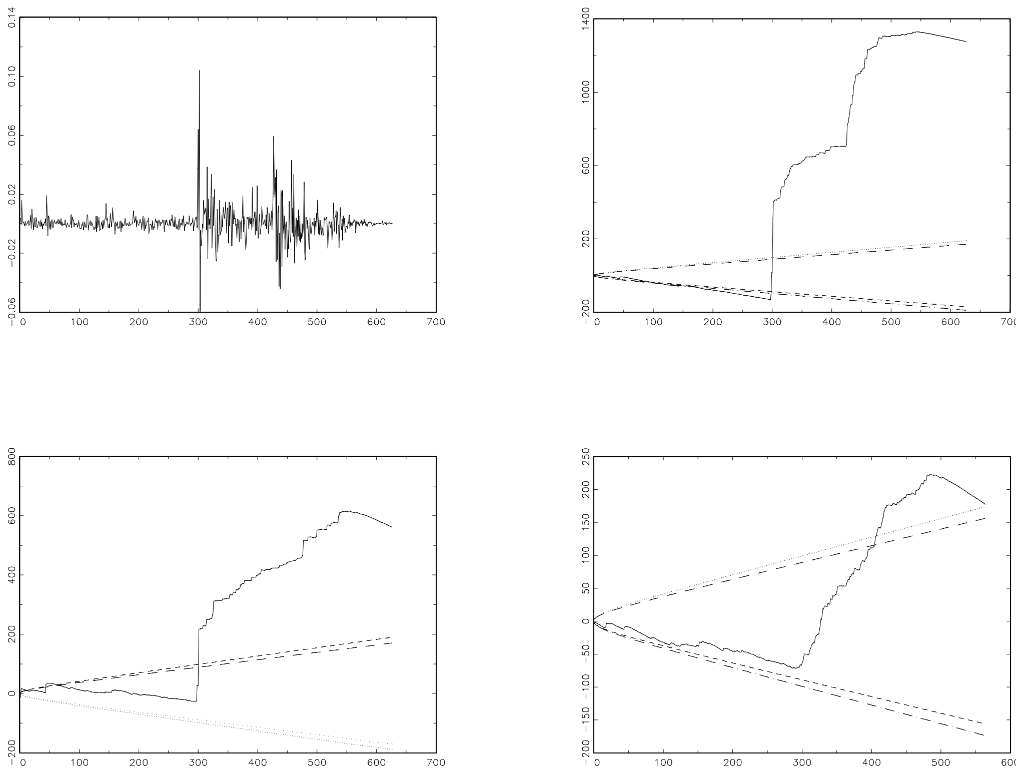

Figure 7: Italian Lire

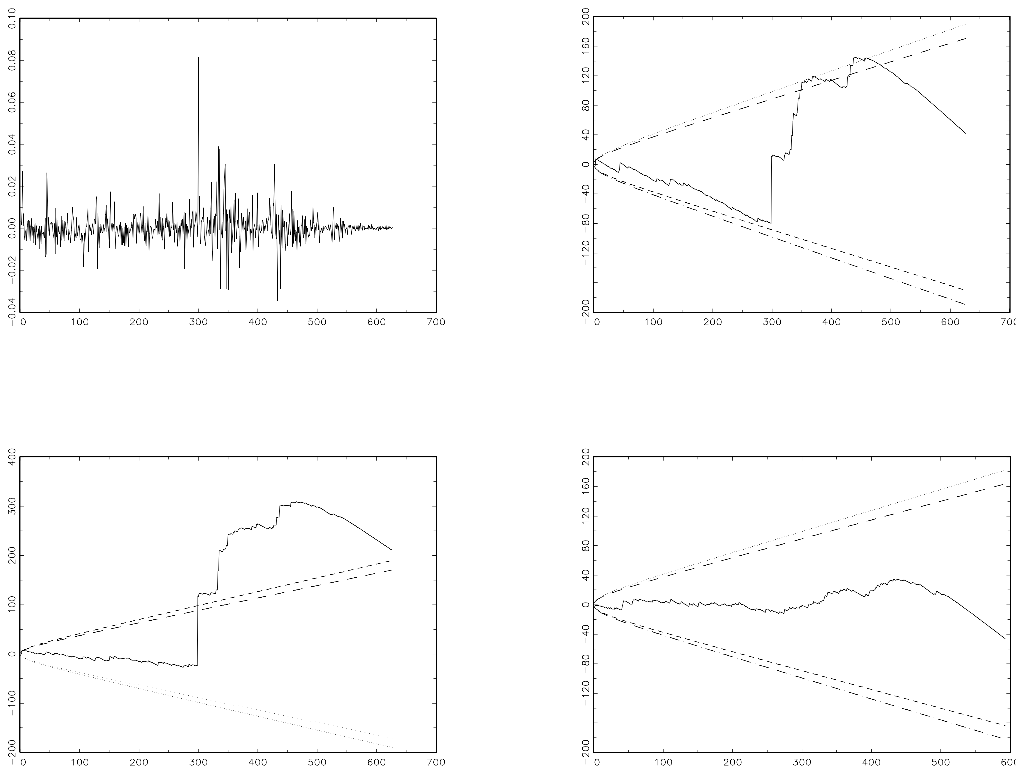

Figure 8: Spanish Peseta

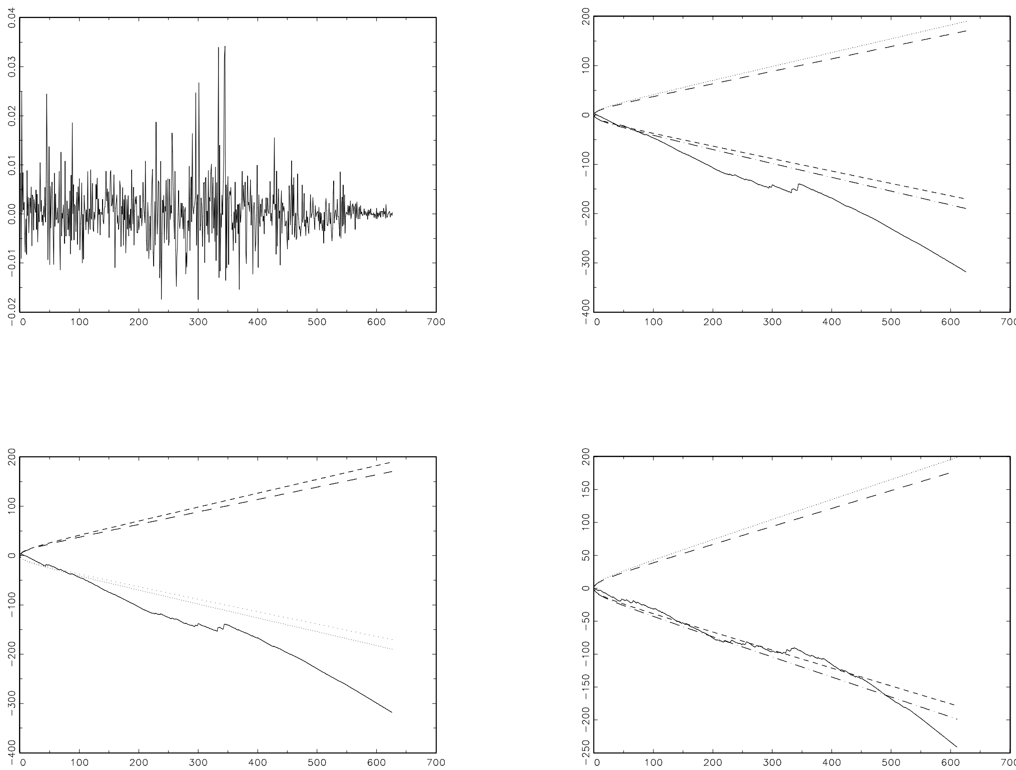

Figure 9: Portuguese Escudo

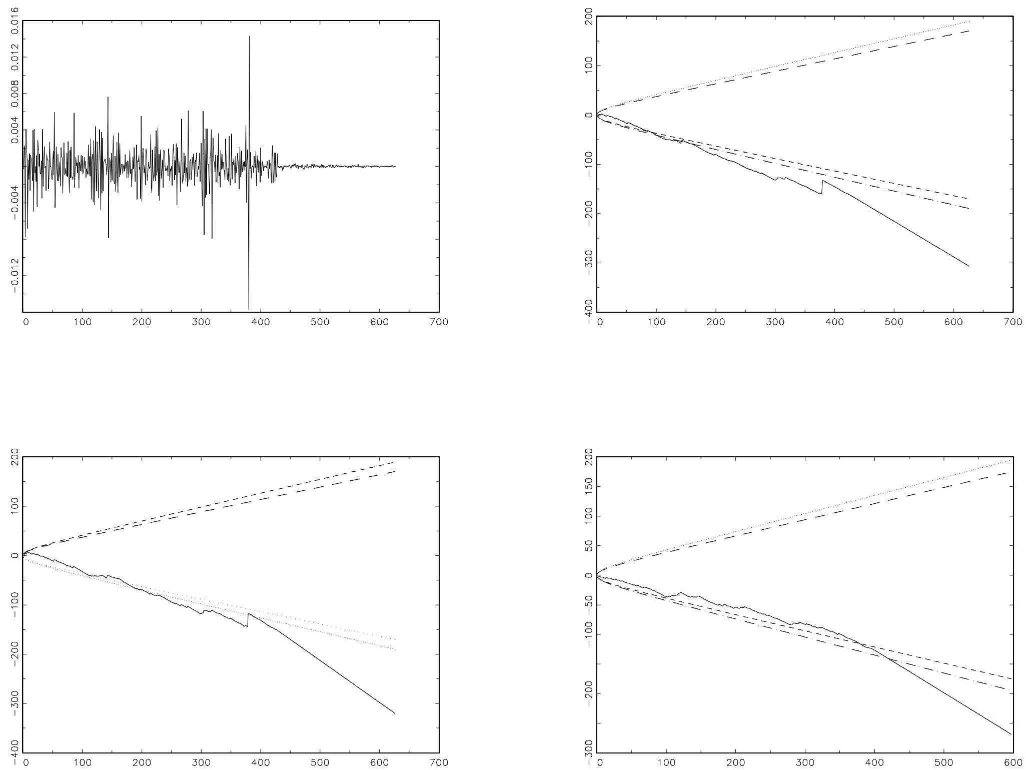

Figure 10: Austrian Shilling

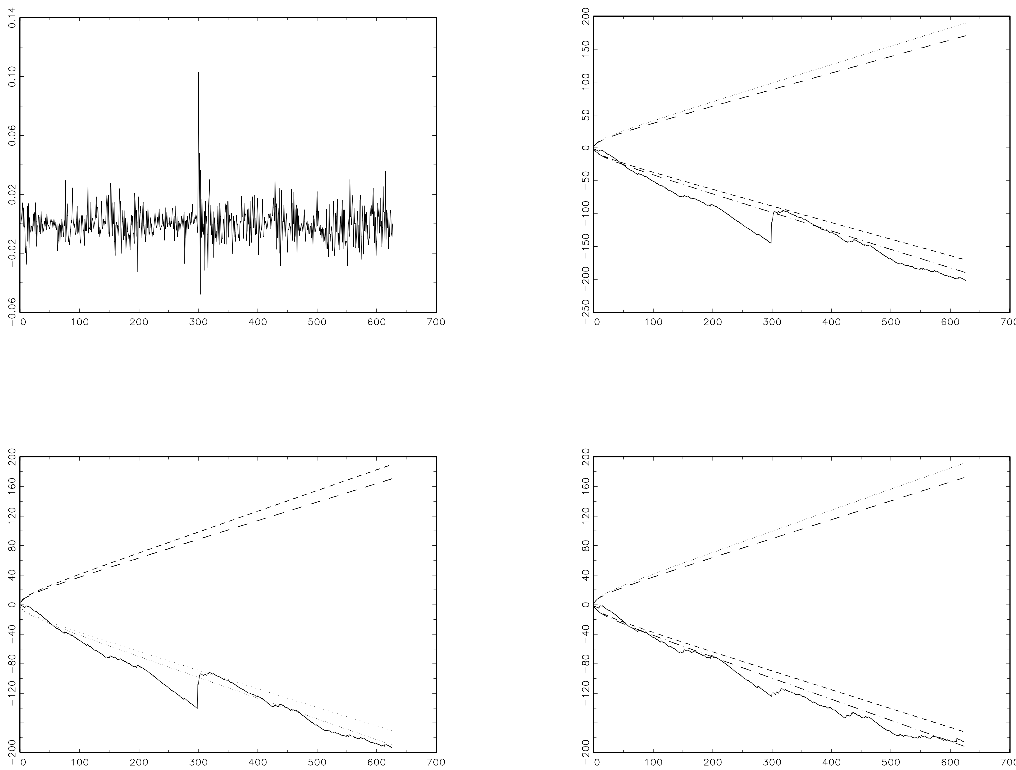

Figure 11: UK Pound

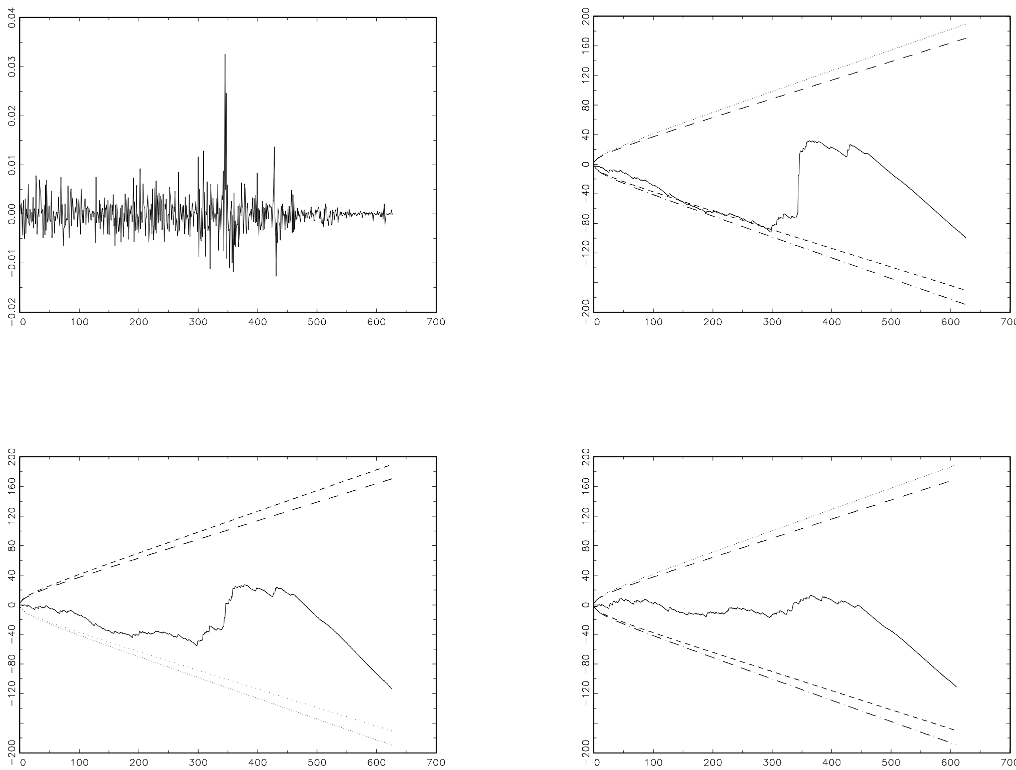

Figure 12: Danish Kroner

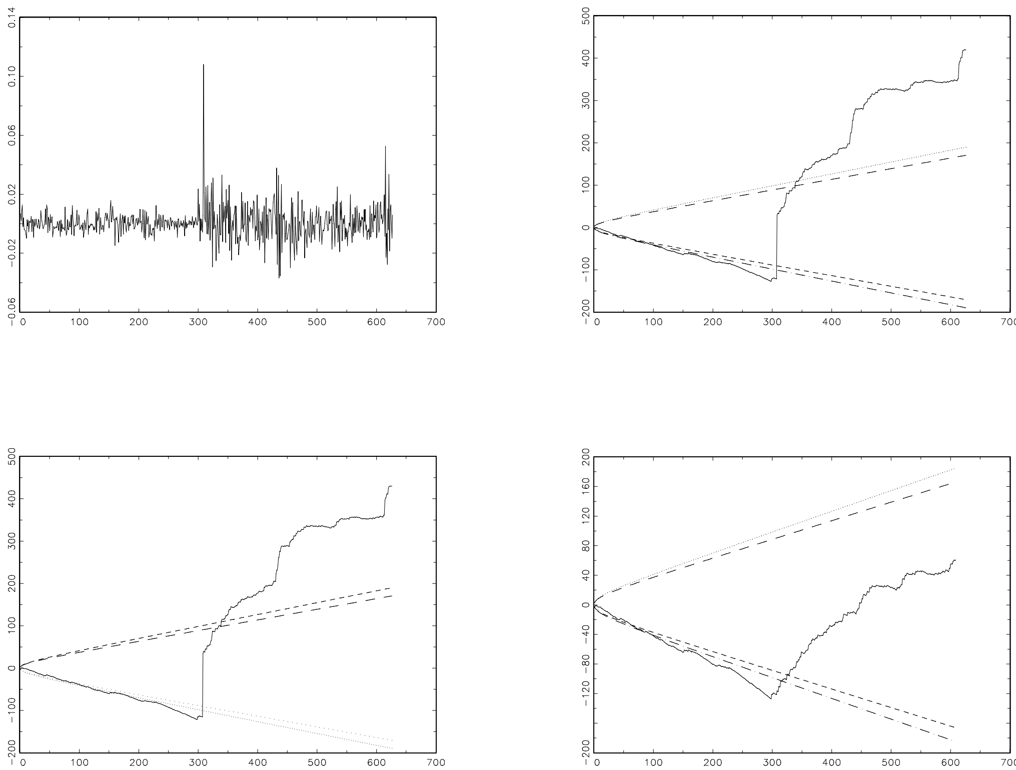

Figure 13: Swedish Kroner

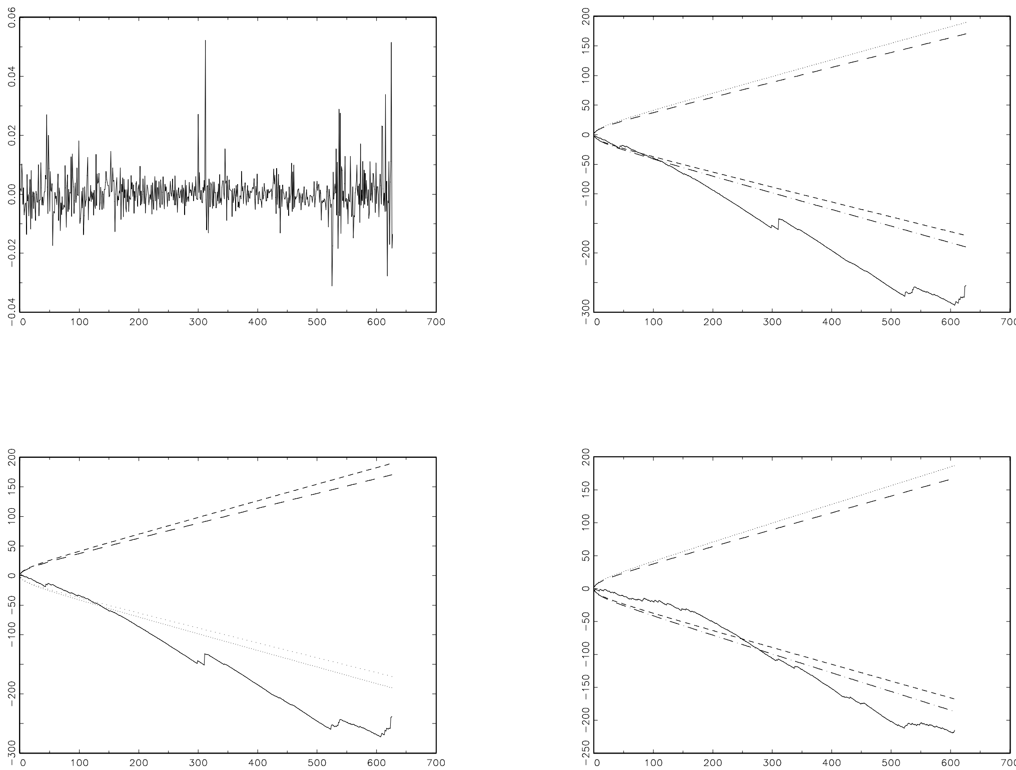

Figure 14: Norwegian Kroner

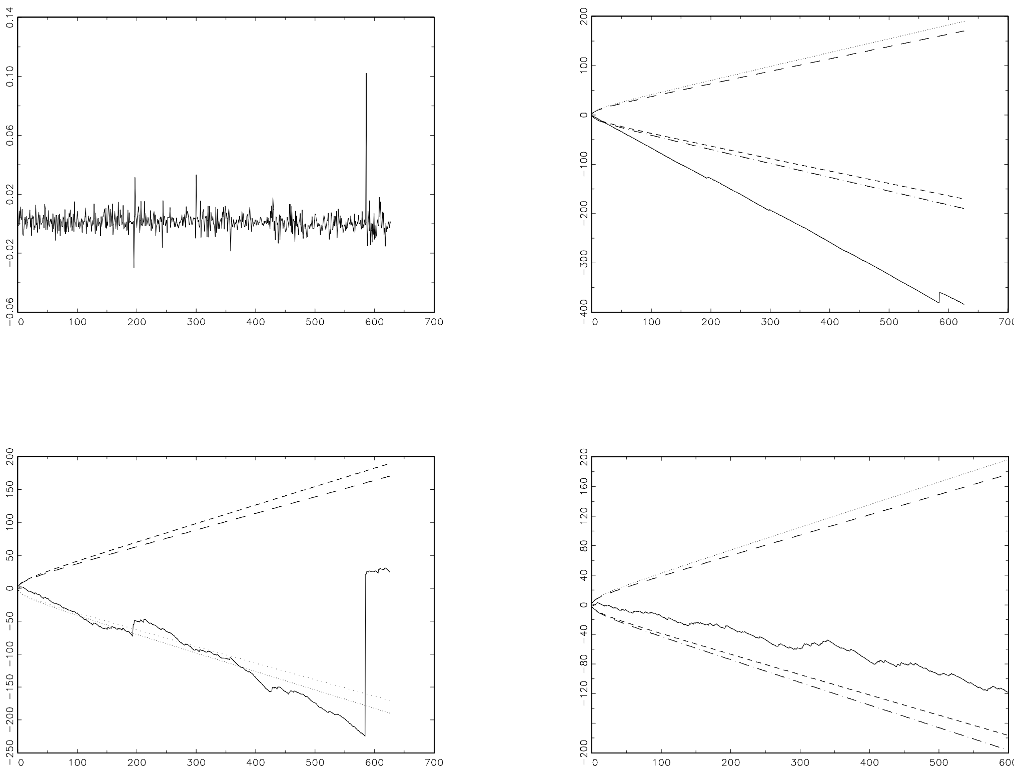

Figure 15: Greek Drachme

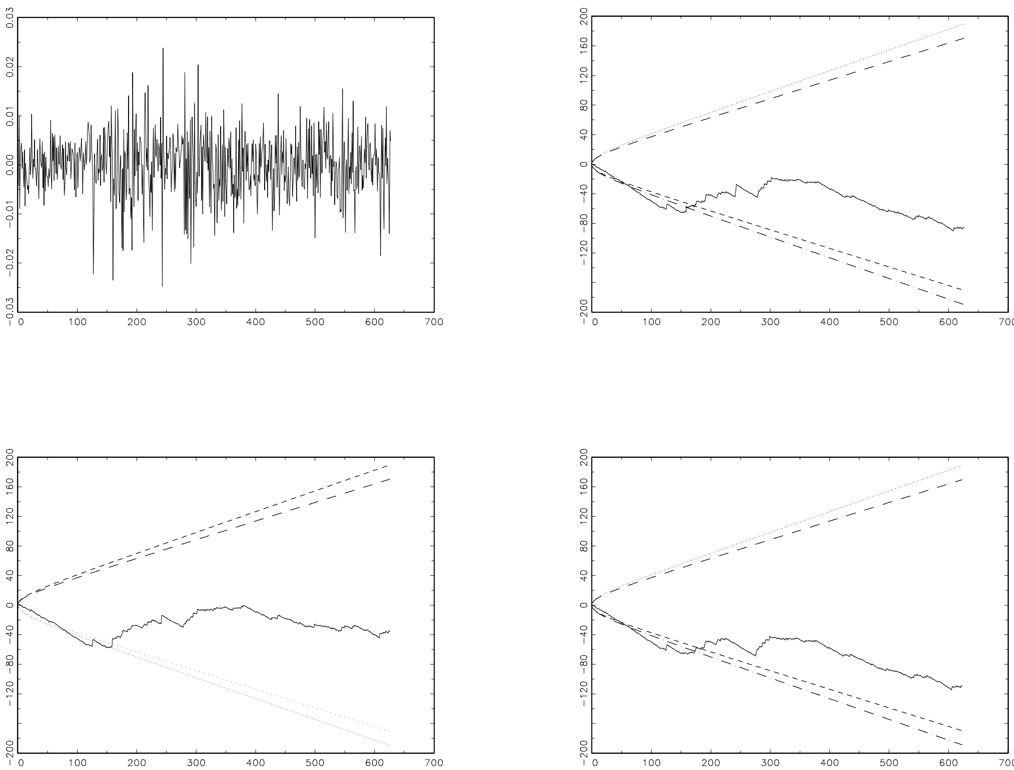

Figure 16: Swiss Franc

References

Andrews, D. W. K. (1993), 'Test for parameter instability and structural change point', *Econometrica* **61**, 821–856.

Baxter, M. & Stockman, A. C. (1989), 'Business cycles and the exchange-rate regime: Some international evidence', *Journal of Monetary Economics* **23**, 377–400.

Brown, R. L., Durbin, J. & Evans, J. M. (1975), 'Techniques for testing the constancy of regression relationships over time', *Journal of the Royal Statistical Society* **37**, 149–163.

Chu, C. S. J., Stinchcombe, M. & White, H. (1996), 'Monitoring structural change', *Econometrica* **64**, 1045–1065.

Eichengreen, B. (1988), 'Real exchange rate behavior under alternativew international monitary regimes', *European Economic Review* **32**, 363–371.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Inclan, C. & Tiao, G. C. (1994), 'Use of cumulative sums of squares for retrospective detection of changes of variance', *Journal of the American Statistical Association* **89**, 913–923.

Lamoureux, C. G. & Lastrapes, W. D. (1990), 'Persistence in variance, structural change, and the GARCH model', *Journal of Business and Economic Statistics* **8**, 225–234.

Mussa, M. (1986), 'Nominal exchange rate regimes and the behavior of exchange rates: Evidence and imlications', *Carnegie-Rochester conference Series on Public Policy* **25**, 117–214.

Ploberger, W., Kramer, W. & Kontrus, K. (1989), 'A new test for structural stability in the linear regression model', *Journal of Econometrics* **40**, 307–318.

Stockman, A. C. (1983), 'Real exchange rate under alternative nominal exchange rate systems', *Journal of International Money and Finance* **2**, 147–166.