Tinbergen Institute Discussion Paper TI98-039/4

Long Memory and Level Shifts: Re-Analyzing Inflation Rates

Charles S. BOS *
Tinbergen Institute, Erasmus University Rotterdam

Philip Hans FRANSES and Marius OOMS
Econometric Institute, Erasmus University Rotterdam

February 27th, 1998

Abstract

A key application of long memory time series models concerns inflation. Long memory implies that shocks have a long-lasting effect. It may however be that empirical evidence for long memory is caused by neglecting one or more level shifts. Since such level shifts are not unlikely for inflation, where the shifts may be caused by sudden oil price shocks, we examine whether evidence for long memory (indicated by the relevance of an ARFIMA model) in G7 inflation rates is spurious or exaggerated. Our main findings are that apparent long memory is quite resistant to level shifts, although for a few inflation rates we find that evidence for long memory disappears.

Keywords: Long memory, fractional integration, structural change, inflation

*Correspondence to Charles S. Bos, Tinbergen Institute, Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands. Email: cbos@few.eur.nl

1 Introduction

A key application of long memory time series models concerns inflation. For example, Hassler & Wolters (1995) and Baillie, Chung & Tieslau (1996) find convincing evidence for the presence of long memory characteristics in, especially, inflation rates in the G7 countries. Long memory implies that shocks have a long-lasting effect. Similar to the arguments in the literature on unit roots versus mean shifts, see Perron (1989) and Perron & Vogelsang (1992), it may however be that empirical evidence for long memory is caused by neglecting one or more level shifts. Since such level shifts are not unlikely for inflation, where these may be caused by sudden oil price shocks, we examine whether evidence for long memory (indicated by the relevance of an ARFIMA model) in G7 inflation rates is spurious or exaggerated.

The outline of the study is as follows. In Section 2 we start with a brief motivation by having a closer look at monthly US inflation, thereby extending some recent results summarized in Ooms (1996). In Section 3 we put forward the relevant theory for testing for long memory and structural level shifts. Our results build on that of Hidalgo & Robinson (1996), who showed that the Wald test is applicable to testing for breaks in a long memory model and on Cheung (1993) since we put forward LM and Wald tests. Section 4 deals with a simulation study of the practical performance of the tests. In Section 5 we apply our tests to monthly G7 inflation rates, where we assume that structural level shifts concur with substantial oil price changes. Our main findings are that apparent long memory is quite resistant to mean shifts, although for a few inflation rates we find that evidence that long memory disappears. In Section 6, we conclude our paper with several remarks.

2 A motivation

Consider the monthly US inflation rate in Figure 1. The data cover January 1957-December 1995 and concern all commodities (source: Bureau of Labor Statistics, series SA0). In the same graph, we draw straight lines that suggest that US inflation has undergone four different regimes. First, until approximately 1967, inflation is stable at a low level. Then the Vietnam war exerts its effect on prices. Inflation is higher, but still quite stable (around this higher level). At the time of the first oil crisis, inflation almost doubles, while at the same time starting to display higher variability. This period of high inflation ends approximately halfway 1981. The final subsample shows a return to earlier inflation levels, although the variability of inflation stays high.

The apparent level shifts might reflect genuine long memory properties intrinsic to inflation. However, the level shifts may also be caused by exogenous events such as the oil crisis. It may also be that the data are better described by a long memory model with mean shifts.

To examine the impact of including mean shifts, we estimate an ARFIMA model (of the type discussed in section 3 below) for U.S. inflation (adjusted for seasonal means). We allow for zero, two or four breaks in our ARFIMA model. The timing of the breaks is fixed exogenously and they correspond with shortly before and after the first oil crisis (1973:07 and 1976:07) and shortly before and after the second oil crisis (1979:01 and 1982:07). In Table 1 we give some key results concerning the fractional differencing parameter d in an ARFIMA model and the parameters for the mean shifts.

The behavior of the parameter d indicating the degree of fractional integration is interesting. Under the assumption of no mean shifts, we find clear indication of long memory, with

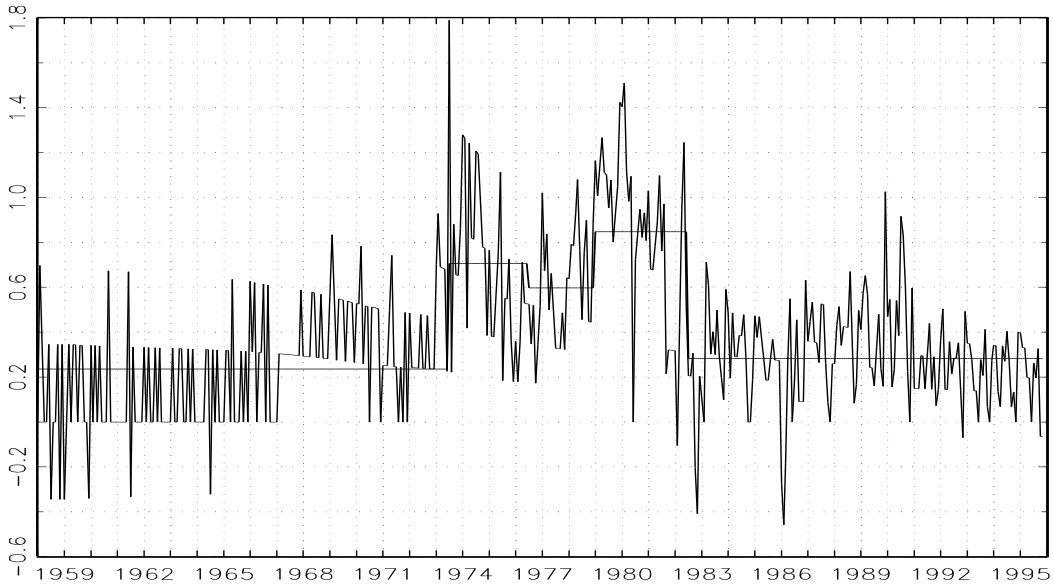

Figure 1: US inflation levels

Table 1: Estimated values of fractional parameter d and mean shifts β in an ARFIMA model for U.S. inflation

Nr. of breaks	0	2	4
$\hat{d}$	0.50 (0.05)	0.40 (0.06)	0.38 (0.06)
$\hat{\beta}_{1973:07}$		1.85 (0.35)	1.65 (0.40)
$\hat{\beta}_{1976:07}$			-0.25 (0.44)
$\hat{\beta}_{1979:01}$			0.97 (0.43)
$\hat{\beta}_{1982:07}$		-1.71 (0.35)	-2.10 (0.39)

Note: Standard errors are given in parentheses

an estimate for d even at the border of the non-stationary region. Allowing for two breaks we observe that $\hat{d}$ reduces considerably. Finally, allowing for four breaks does not seem to change much, since $\hat{d}$ obtains about the same value as in the case of two shifts.

3 Some theoretical results

The empirical results in Table 1 evoke interest in the following issues. The first concerns how one would formally address modeling breaks and fractional integration jointly. The second concern is with test statistics that are useful to examine if structural shifts in an ARFIMA model are statistically plausible. The asymptotic distribution of these statistics is then relevant, but also their small sample performance. In this section we deal with these issues, except for the simulation evidence which we postpone to Section 4.

The ARFIMA model

A fractionally integrated model aims to capture the long memory that is apparent in a time series. Where the influence of a shock in a stationary [I(0)] model disappears after a limited number of periods (depending on the short memory parameters in the autoregressive and moving average parts), and where the effect of a shock lasts forever in a unit root [I(1)] process, the fractionally integrated model [FI(d) with $d \in (0,1)$] takes up an intermediate position, see Granger & Joyeux (1980), Hosking (1981), and more recently, Baillie (1996) and Beran (1994).

The ARFIMA(p,d,q) model is written as

$$\Phi(L)(1-L)^d(z_t - \mu_z) = \Theta(L)\epsilon_t \qquad t = 1,..,T \tag{1}$$

where z_t is the time series at time t, μ_z its mean, and $\Phi(L) = 1 - \phi_1 L - .. - \phi_p L^p$ is the stable autoregressive polynomial in the lag operator L and $\Theta(L) = 1 + \theta_1 L + .. + \theta_q L^q$ the invertible moving average part ($p, q \in \{0,1,2,..\}$). $\Phi(L)$ and $\Theta(L)$ together define the short memory characteristics of the model. We assume the noise process ϵ_t to be Gaussian, with expectation zero and variance σ_ϵ^2. The long memory behavior is governed by the part $(1-L)^d$. If d is an integer, the I(d) process is non-fractional, and taking d^{th} differences of z_t leads to a (stationary) I(0) series. If $d \in (0,1)$, one says that z_t exhibits long memory behavior.

Including a level shift

To allow for a level shift, after a fraction τ ($0 < \tau < 1$) of the data, we write the observations y_t as the sum of an unobserved ARFIMA process and the term for the level shift:

$$y_t = z_t + \mu I_{\{t > \tau T\}} \tag{2}$$

The parameter μ indicates the size of the level shift in the series y_t at time τT. We define the relative level shift as

$$\beta = \frac{\mu}{\sigma_z} \tag{3}$$

with σ_z being the standard deviation of the ARFIMA process. If the level shift μ and the timing of the break τ are known, this standard deviation can be estimated directly using the empirical standard deviation of the underlying process z_t.

The extension of (2) to k breaks is straightforward. We define μ_r as the r-th shift in level, compared to the previous level, and we define the relative breaksize β_r similar to (3), where $r = 1, .., k$. When we allow for k level changes at prespecified fractions $0 < \tau_1 < .. < \tau_k < 1$, we can extend (2) to

$$z_t = y_t - \sum_{r=1}^{k} \mu_r I_{\{t>\tau_r T\}}. \tag{4}$$

The sample mean of the underlying process z_t is now defined as

$$\overline{z} = \frac{1}{T}\sum_t y_t - \frac{1}{T}\sum_t \sum_{r=1}^{k} \mu_r I_{\{t>\tau_r T\}} = \overline{y} - \frac{1}{T}\sum_{r=1}^{k} \lfloor T(1-\tau_r)\rfloor \mu_r, \tag{5}$$

where $\lfloor\cdot\rfloor$ denotes the operator to take the integer part of the argument (the entier function).

In this paper we assume we know the values of $\tau_1, .., \tau_k$. It is of course possible to endogenize the timing of the breaks, as in Andrews (1993) or Bai (1996). This would complicate matters too much for the moment, and further investigation into this subject is postponed for later research.

The spectrum of an ARFIMA model

In order to derive the following results, we assume from now on that z_t is a zero mean, stationary and invertible ARFIMA process, which is obtained from the original data by filtering out the known level shifts, and, if needed, by appropriate differencing, and subtracting the sample mean $\overline{z}$, i.e.,

$$z_t^a(\mu, \lfloor d + \frac{1}{2}\rfloor) = y_t - \sum_{r=1}^{k} \mu_r I_{\{t>\tau_r T\}} \tag{6}$$

$$z_t^b(\mu, \lfloor d + \frac{1}{2}\rfloor) = (1-L)^{\lfloor d+\frac{1}{2}\rfloor} z_t^a \tag{7}$$

$$z_t(\mu, \lfloor d + \frac{1}{2}\rfloor) = z_t^b - \overline{z^b}. \tag{8}$$

We assume that z_t in (8) can be described by an ARFIMA(p, d, q) model with d in $[-0.5, 0.5)$.

The autocovariance generating function (ACGF) of an ARFIMA(p, d, q) process is written as

$$g(z; \Psi) = \sigma_\epsilon^2 \left((1-z)(1-z^{-1})\right)^{-d} \frac{\Theta(z)\Theta(z^{-1})}{\Phi(z)\Phi(z^{-1})}, \tag{9}$$

see Harvey (1989). To save notation, we use Φ as the set of parameters $\{\phi_1, .., \phi_p\}$ in the polynomial $\Phi(L)$, Θ likewise for the MA parameters and we write Ψ as shorthand for the ARFIMA parameters $\{\Phi, d, \Theta, \sigma_\epsilon^2\}$. We use μ to denote the level shifts μ_1 to μ_r. The spectral generating function (SGF) of the ARFIMA model is given by

$$g(\lambda; \Psi) = \sigma_\epsilon^2 \left|1 - e^{i\lambda}\right|^{-2d} \frac{\left|\Theta(e^{i\lambda})\right|^2}{\left|\Phi(e^{i\lambda})\right|^2}, \tag{10}$$

leading to the power spectrum, which is used extensively in the likelihood function as

$$f(\lambda; \Psi) = \frac{1}{2\pi} g(\lambda; \Psi). \tag{11}$$

The loglikelihood

With z_t defined as in (8), the loglikelihood of the ARFIMA model is

$$\ln \mathcal{L}(y|\Phi, d, \Theta, \mu, \sigma_\epsilon^2) = -\frac{T}{2}\ln 2\pi - \frac{1}{2}\ln|\Sigma_T(\Psi)| - \frac{1}{2}z'\Sigma_T^{-1}(\Psi)z. \tag{12}$$

The covariance matrix of z is $\Sigma_T(\Psi) = [\gamma(j-l)]_{j,l=1}^T$, with $\gamma(j)$ the j-th autocovariance of the process z. The loglikelihood depends only on the level shifts through the change from observations y_t to the underlying process z_t. This dependence is not stressed in the notation.

Although it is possible to construct the exact likelihood function in the time domain (see Sowell (1992)), we use an approximation in the frequency domain following Harvey (1989). The latter procedure is computationally simpler. In practice, the problem with the calculation of the loglikelihood function is found in the covariance matrix $\Sigma_T(\Psi)$, which is a $T \times T$ matrix. Calculation of its determinant and inverse is time-consuming. Harvey (1989, section 4.3) proposes to use the following approximations

$$\ln|\Sigma_T(\Psi)| = T\ln 2\pi + \sum_{j=0}^{T-1}\ln f(\lambda_j;\Psi) \tag{13}$$

with $f(\lambda;\Psi)$ the power spectrum of the process z_t at frequency λ and $\lambda_j = 2\pi j/T$, and

$$z'\Sigma_T(\Psi)z = \sum_{j=0}^{T-1}\frac{I_z(\lambda_j;\mu)}{f(\lambda_j;\Psi)}, \tag{14}$$

where $I_z(\lambda;\mu)$ denotes the break-adjusted periodogram of z_t at frequency λ. In this notation, the dependence of the periodogram $I_z(\lambda;\mu)$ on the level shifts μ is made explicit again, see the Appendix.

When calculating the elements of equations (13)-(14), elements at frequency zero are disregarded, as advocated in Beran (1994), hence the summations start at $j = 1$. Taking all above results together leads to Whittle's approximative loglikelihood function, denoted by

$$\begin{aligned}
&\ln \mathcal{L}(y|\Phi, d, \Theta, \mu, \sigma_\epsilon^2) \\
&\quad = -\frac{T}{2}\ln 2\pi - \frac{1}{2}\ln|\Sigma_T(\Psi)| - \frac{1}{2}z'\Sigma_T^{-1}(\Psi)z \\
&\quad \approx -\frac{T}{2}\ln 2\pi - \frac{T}{2}\ln 2\pi - \frac{1}{2}\sum_{j=1}^{T-1}\ln f(\lambda_j;\Psi) - \frac{1}{2}\sum_{j=1}^{T-1}\frac{I_z(\lambda_j;\mu)}{f(\lambda_j;\Psi)} \\
&\quad = -T\ln 2\pi - \sum_{j=1}^{T^*}\delta_j \ln f(\lambda_j;\Psi) - \sum_{j=1}^{T^*}\delta_j\frac{I_z(\lambda_j;\mu)}{f(\lambda_j;\Psi)}.
\end{aligned} \tag{15}$$

Here we use the fact that for the power spectrum it holds that $f(\lambda;\Psi) = f(-\lambda;\Psi)$ and that $f(\lambda;\Psi) = f(\lambda + 2\pi;\Psi)$. Furthermore, $T^* = \lfloor T/2 \rfloor$ and weights δ_j are defined as

$$\delta_j = \begin{cases} \frac{1}{2} & \text{if } j = \frac{T}{2} = \left\lfloor \frac{T}{2} \right\rfloor \\ 1 & \text{else} \end{cases} \tag{16}$$

for $j = 1, .., T^*$. The purpose of the weighting is only to make sure that the midpoint of the range of frequencies is not used twice if T is even.

Testing for breaks

Testing whether level shifts occur corresponds with testing a linear restriction on the parameters $\xi = \{\Psi, \mu\} = \{\Phi, d, \Theta, \sigma_\epsilon^2, \mu\}$. The relevant null hypothesis and alternative hypothesis are

$$H_0 : R\xi = \mu = 0 \qquad H_1 : R\xi = \mu \neq 0, \tag{17}$$

where the alternative implies that we are testing against k breaks at prespecified moments. R is the matrix to select the parameters to be restricted from the vector ξ. In this setting, the Wald, Lagrange Multiplier (LM) and Likelihood Ratio (LR) tests can be used. In our case, the parameters under the null hypothesis of no breaks are more easily calculated than under the alternative, as the inclusion of breaks in the likelihood function would imply that for every evaluation of the likelihood, the periodogram $I_z(\lambda; \mu)$ would have to be recalculated. Thus, the LM principle is our first choice for the next section, where the simulations are presented. In Section 5, where the analysis is done for the countries of the G7, both the LM and Wald test statistics are used.

Hidalgo & Robinson (1996) have proved that the Wald test statistic on a single structural change in a long memory environment follows a χ_1^2 distribution, when using a non-parametric estimator for d and an OLS estimate (disregarding the value of $\hat{d}$ or any short memory parameters) for the mean of the series before and after the break. Their proof cannot easily be translated to a setting where an iterative generalized least squares or (approximate) maximum likelihood procedure is used. However, as parametric estimators tend to converge at least as fast as non-parametric estimators (assuming a correct model specification under the null hypothesis), this asymptotic χ^2 distribution can be expected to hold in our case too. As the Wald and LM tests are asymptotically equivalent, the results of Hidalgo & Robinson can be translated to the LM test. Finally, the extension to multiple breaks is straightforward, as long as we condition on the number and the timing of the breakpoints. Asymptotically, we expect a χ_k^2 distribution for the test on k breaks under the null hypothesis.

The calculation of the test statistics follows the familiar lines:

$$W_\mu = \hat{\mu}' \left(R \, J^{-1}(\hat{\Psi}, \hat{\mu}) R' \right)^{-1} \hat{\mu}, \tag{18}$$

$$LM_\mu = \left. \frac{\partial \ln \mathcal{L}}{\partial \xi} \right|_{\hat{\xi}_0}' J^{-1}(\hat{\Psi}_0, 0) \left. \frac{\partial \ln \mathcal{L}}{\partial \xi} \right|_{\hat{\xi}_0}, \tag{19}$$

with ξ as defined before. The $\hat{\Psi}$ and $\hat{\mu}$ are the unrestricted estimates of the parameters in the model, whereas $\hat{\xi}_0 = \{\hat{\Psi}_0, 0\}$ is the estimate of the parameters under the null hypothesis of no breaks. The $J(\Psi, \mu)$ denotes the information matrix

$$J(\Psi, \mu) = -E \begin{bmatrix} \frac{\partial^2 \ln \mathcal{L}}{\partial \Psi \partial \Psi'} & \frac{\partial^2 \ln \mathcal{L}}{\partial \Psi \partial \mu'} \\ \frac{\partial^2 \ln \mathcal{L}}{\partial \mu \partial \Psi'} & \frac{\partial^2 \ln \mathcal{L}}{\partial \mu \partial \mu'} \end{bmatrix}. \tag{20}$$

In appendix A we provide more details on the calculation of these test statistics.

4 Simulation evidence

In this section we report on some simulation evidence concerning the small sample properties of the estimators of the parameters and of the LM test for level shifts.

Two data generating processes

We generate 512 (=T) observations from an AR(1) process with $\phi = 0.8$ and from an ARFIMA$(1, d, 0)$ process with $\phi = 0.4$ and $d = 0.3$ (see appendix B for information on the generator used for the ARFIMA model). The variance σ_ϵ^2 of the disturbances is taken to be 1. Halfway the sample, we add a level shift of size $\beta = 1$, i.e., a shock of one time the standard deviation of the underlying process is added to the mean of the series after the observation at $\frac{1}{2}T$. Given the variance of the ARFIMA processes, which is $\sigma_z^2 = 2.778$ for the AR(1) model and $\sigma_z^2 = 2.357$ in the ARFIMA(1, d, 0) case, the parameter μ in (2) equals 1.667 and 1.535, respectively. For both time series, i.e., AR(1) and ARFIMA$(1, d, 0)$, with and without a level shift, we estimate the parameters of an ARFIMA$(1, d, 1)$ model. Our simulations are based on 5000 replications.

The empirical distribution of the estimators

In Figure 2 the empirical distribution functions of $\hat{\phi}, \hat{d}$ and $\hat{\theta}$ are shown for data generated according to the AR(1) process without a break. Figure 3 shows the estimation results for the same parameters, but estimated for data which include a level shift. The boxes indicate the 5%, 50% and 95% quantiles, whereas the dashed lines cross at the original parameter value.

From the first graph we can conclude that the estimation procedure leads to consistent results. The estimated medians (with the corresponding values in the DGP between brackets) are $\hat{\phi}_{50\%} = 0.79$ [0.8], $\hat{d}_{50\%} = 0.00$ [0.0] and $\hat{\theta}_{50\%} = 0.01$ [0.0]. Including a break in the data generating process (DGP) leads to Figure 3, which depicts the same estimators. In the left-hand graph we see that in a certain fraction of the simulations, inclusion of a level shift in the DGP leads to data with unit root properties ($\hat{\phi}$ close to 1). A closer examination of the estimation results shows that in these cases d is estimated around or even below zero. In the other cases, ϕ was estimated below the 'true' value of 0.8, i.e. the median decreases to 0.68. Except for cases where a unit root was found, d is estimated at a value higher than in the previous simulations. Indeed, a correlation of -0.89 between $\hat{\phi}$ and $\hat{d}$ is found. The median $\hat{d}_{50\%}$ shifts upwards to a value of 0.27. Such a value for the degree of fractional integration is in general taken as a strong indication of long memory behavior. The estimates of the parameter θ decrease somewhat, relative to the original value of zero, while their spread is higher than before. Clearly, neglecting a level shift in otherwise short memory data may lead one to believe that long memory resides in the data.

When the ARFIMA$(1, d, 0)$ model is taken as the DGP, we obtain the results as given in Figures 4 and 5. Figure 4 shows the consistency of the approximative Whittle estimator in the presence of long memory in the data. This result agrees with those reported in Hauser (1997), where encouraging results for this estimator are obtained from an extensive simulation study. The medians (with values of the DGP between brackets) found are $\hat{\phi}_{50\%} = 0.41$ [0.4], $\hat{d}_{50\%} = 0.29$ [0.3] and $\hat{\theta}_{50\%} = 0.03$ [0.0]. For the DGP where a break is included, we observe shifts in the empirical distribution of the estimators. Now, no indication of unit roots is found. The estimates of ϕ shift down to a distribution with 0.28 as a median. Estimated $\hat{d}$ values indicate even stronger long memory characteristics than found before. The parameter $\hat{\theta}$ does not change much. Again, a positive bias in $\hat{d}$ is found when a level shift is neglected.

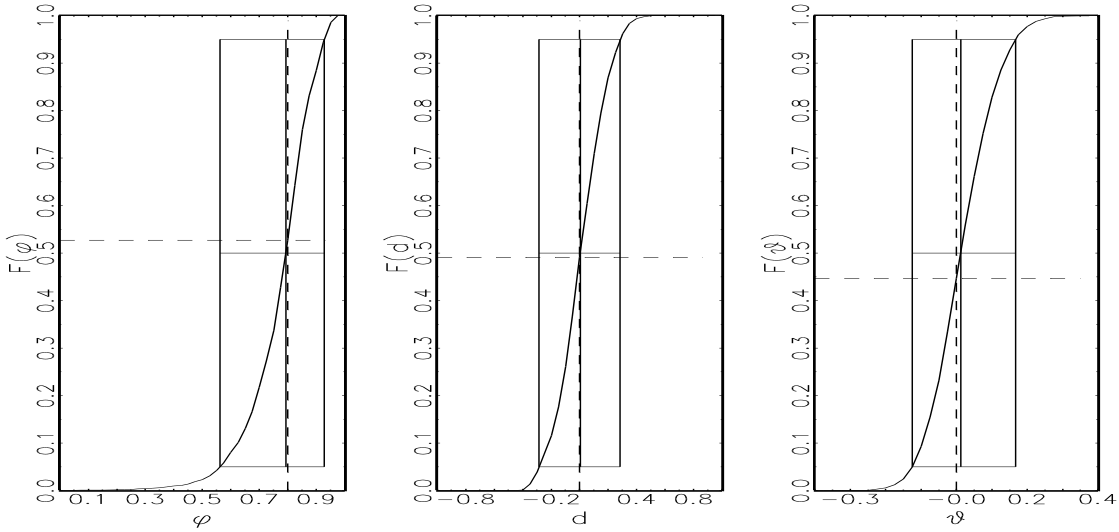

Figure 2: Estimating parameters in an ARFIMA$(1, d, 1)$ model. The DGP is an AR(1) model with $\phi = 0.8, \sigma_\epsilon^2 = 1$, without a level shift.

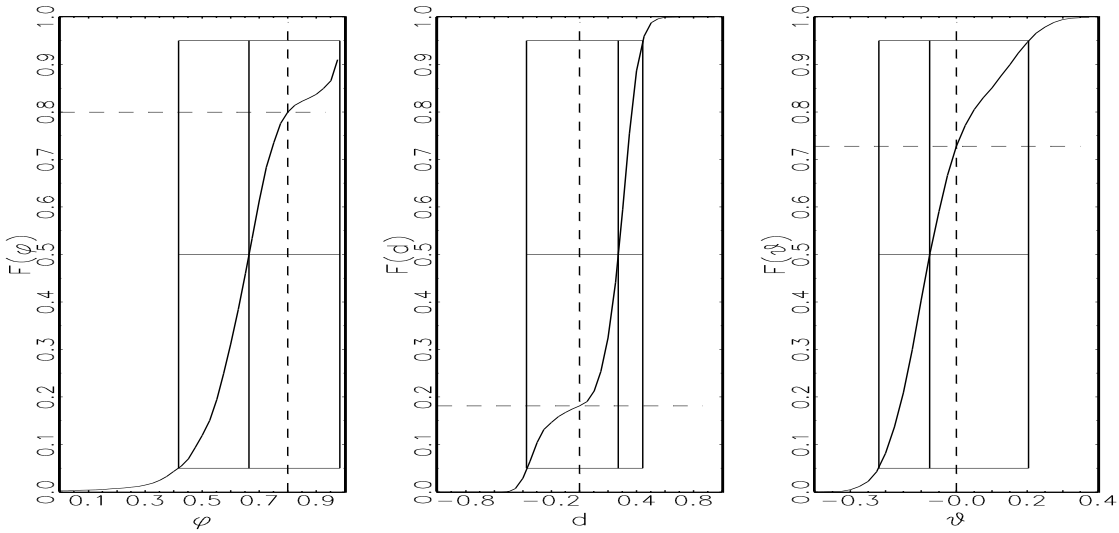

Figure 3: Estimating parameters in an ARFIMA$(1, d, 1)$ model. The DGP is an AR(1) model with $\phi = 0.8, \sigma_\epsilon^2 = 1$, with a level shift of size $\beta = 1$.

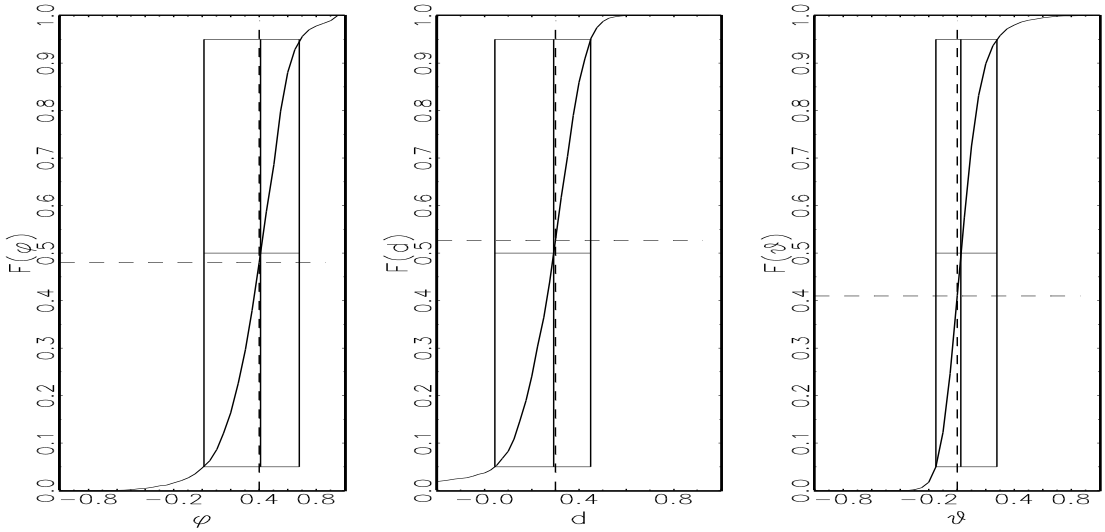

Figure 4: Estimating parameters in an ARFIMA$(1, d, 1)$ model. The DGP is an ARFIMA$(1, d, 0)$ model with $\phi = 0.4, d = 0.3, \sigma_\epsilon^2 = 1$, without a level shift.

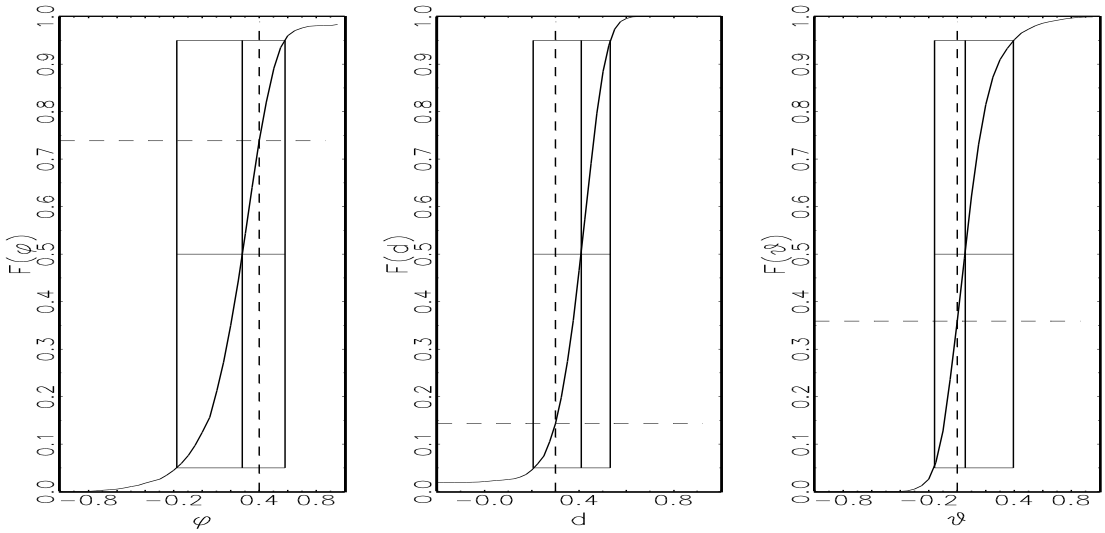

Figure 5: Estimating parameters in an ARFIMA$(1, d, 1)$ model. The DGP is an ARFIMA$(1, d, 0)$ model with $\phi = 0.4, d = 0.3, \sigma_\epsilon^2 = 1$, with a level shift of size $\beta = 1$.

The empirical distribution of the LM test statistic for a level shift

Figure 6 depicts the empirical density and distribution function of the Lagrange Multiplier test statistic, when the presence of a break is tested in the series simulated according to the AR(1) DGP. The dashed curve in the same graph indicates the χ^2_1 distribution, which is supposed to be the asymptotic distribution of the statistic. The empirical density function exhibits the same shape as the χ^2 density. However, large values of the test statistic occur too often, as indicated by the heavier tail of the empirical distribution when compared to the χ^2 distribution. The standard 5% critical value for the χ^2 distribution is 3.84. Using this critical value leads to an empirical size of the test of 14%. The empirical 5% critical value is found at a value of the test statistic of 7.80.

The effects of the inclusion of a break in the DGP on the LM test statistic are summarized in Figure 7. Although the shape of the density function still resembles a χ^2, much larger values of the test statistic are found. The standard χ^2-based critical value would lead to a (correct) rejection of the null hypothesis of no break in 82% of the cases. Using the empirical critical value of 7.80 reduces the power to 68%, which is still quite reasonable.

When the data are generated according to the ARFIMA(1, d, 0) model, the empirical size of the test (at a nominal level of 5 %) is 21% (see Figure 8). Even when generating under the null, more large values of the test statistic are found. The empirical critical value at the 95% confidence level is 12.19, which is considerably higher than the original 3.84. Finally, when the DGP is the ARFIMA(1,d,0) model with a break, we obtain the results as in Figure 9. If the critical value of 3.84 is used, the empirical power is still around 80%, as indicated by the horizontal dashed line in the graph on the right hand side of Figure 9. Using the empirical 5% critical value in this case however lowers the power to 49%. The findings on size and power are summarized in Table 2. The third column in this table reports the rejection rates for the Beran test for goodness-of-fit, advocated in Beran (1994), at a nominal level of 5%. We interpret the numbers in this column as that this test statistic does not signal important residual correlation in the models fitted to the data.

Our simulations lead to the conclusion that a neglected level shift has a substantial effect on the parameter estimates. The LM test seems to be able to detect a level shift, although the power can be low and some size distortions do occur.

5 Inflation: Long memory and level shifts

In this section we re-analyze part of the series previously used by Baillie et al. (1996). The dataset consists of the Consumer Price Indices (CPI) for the countries of the G7: Canada, France, Germany, Italy, Japan, the United Kingdom and the United States. The data for the U.S. originates from the Bureau of Labor Statistics concerning the overall price index SA0, and it ranges from January 1957 until November 1995. Indices for the other countries are extracted from Citibase. Observations on the months January 1948 until March 1990 are available. Inflation rates are constructed from the price indices by taking $y_t = 100\Delta \ln \mathrm{CPI}_t$. As the inflation rates exhibit rather erratic behavior in the first years of the sample, we only use the data starting in 1958. For the U.S., we have a sample of 455 observations, while for the other countries 387 observations are available. To account for part of the seasonality, the data y_t are first adjusted for seasonal means. The parameters in the ARFIMA models for the resulting series are estimated by optimizing the likelihood as described in Section 3. As

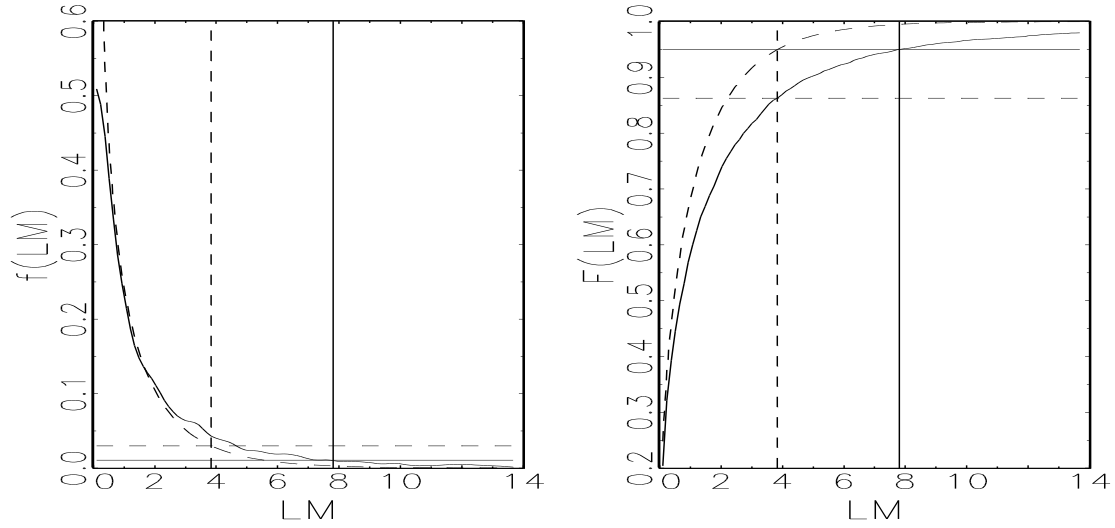

Figure 6: Empirical distribution of the LM-test statistic. The DGP is an AR(1) model with $\phi = 0.8, \sigma_\epsilon^2 = 1$, without a level shift.

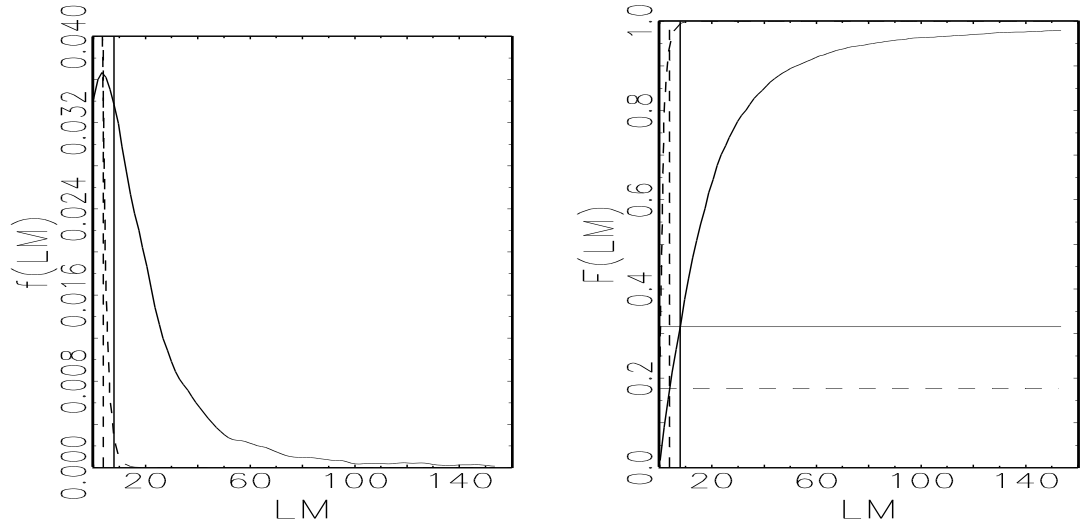

Figure 7: Empirical distribution of the LM-test statistic. The DGP is an AR(1) model with $\phi = 0.8, \sigma_\epsilon^2 = 1$, with a level shift of size $\beta = 1$.

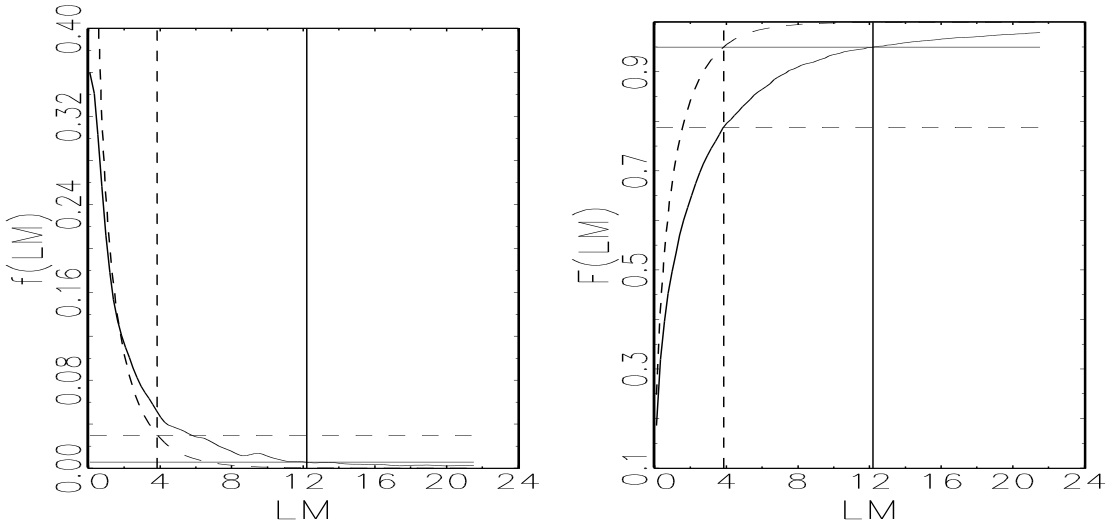

Figure 8: Empirical distribution of the LM-test statistic. The DGP is an ARFIMA$(1, d, 0)$ model with $\phi = 0.4, d = 0.3, \sigma_\epsilon^2 = 1$, without a level shift.

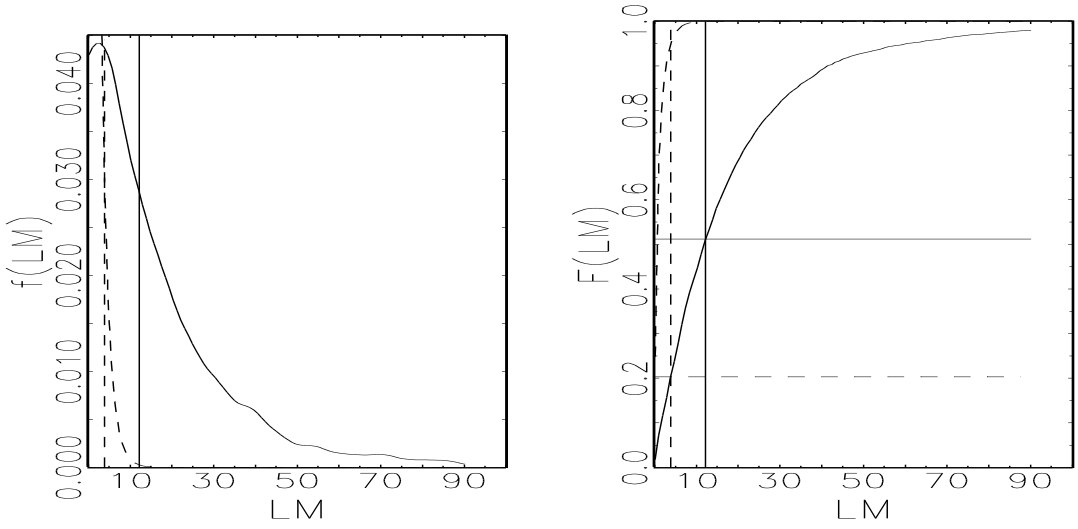

Figure 9: Empirical distribution of the LM-test statistic. The DGP is an ARFIMA$(1, d, 0)$ model with $\phi = 0.4, d = 0.3, \sigma_\epsilon^2 = 1$, with a level shift of size $\beta = 1$.

Table 2: Empirical rejection frequencies of the LM test statistic for a level shift and the Beran test for residual white noise

DGP: AR(1), $\phi = 0.8$			
β	LM > 3.84	LM > 7.80	$p_{\text{Beran}} < 0.05$
0.00	13.74	5.00	4.08
1.00	82.48	68.46	3.90
DGP: ARFIMA$(1, d, 0)$, $\phi = 0.4, d = 0.3$			
β	LM > 3.84	LM > 12.19	$p_{\text{Beran}} < 0.05$
0.00	21.22	5.00	4.54
1.00	79.72	48.78	4.26

Note: Each series consists of 512 observations. The number of replications is 5000

the estimation is done in the frequency domain, this adjustment for seasonality corresponds to putting the periodogram to zero at the seasonal frequencies (see Ooms & Hassler (1997)). The timing of the breaks corresponds with the oil price shocks (see Table 1), and is taken to be equal for all countries.

We aim to consider an ARFIMA model. Several settings with different degrees in the AR and MA polynomials are tried. Our specification search results in a model with AR parameters ϕ_1, ϕ_{12} and ϕ_{13} together with the degree of integration d and residual variance σ_ϵ^2. This model appears to capture the short and long run correlations quite well, as is indicated by the Beran (1994) test for white noise. Specifically, adding moving average parameters often leads to root cancellation, and hardly improves the residual variance.

In Table 3 the results of the estimations are presented. For each country we consider three models. First, the pure ARFIMA model is considered. The parameter $\hat{\phi}_{13}$ is significant only in the case of German inflation rates. For the other countries, the parameter can be omitted. The fractional integration parameter d is estimated around the commonly found value of 0.4. For the U.S., we find $\hat{d} = 0.501$. Theoretically, the approximative Whittle estimator is only consistent in the range of $d \in [-0.5, 0.5)$, although it is known to be little biased if the degree of integration lies just outside the stationary region. The Beran test for the absence of residual correlation, which is reported along with its corresponding p-value, does not indicate strong correlation in the residuals. The $\hat{\sigma}_\epsilon$ reported is the estimated standard deviation of the disturbances. The $\hat{\sigma}_z$ denotes the standard deviation of the ARFIMA process, which in this no break case equals the standard deviation of the process z_t itself.

Allowing for a level shift in July 1973 and July 1982 leads to parameter estimates as reported in the second panel of Table 3. Apart from the ARFIMA parameters, the sizes of the level shifts μ_r are estimated as well. Subtracting the estimated level shifts from the data leads to the underlying process z_t as in Section 3. The empirical standard deviation $\hat{\sigma}_z$ of this underlying process is used to calculate the relative break sizes $\hat{\beta}_r$ as defined in equation (3). Reported standard deviations of the estimates of β_r are calculated from the original standard deviations of the level shifts $\hat{\mu}_r$, taking $\hat{\sigma}_z$ is given. As $\hat{\sigma}_z$ is not a given, fixed parameter, the true uncertainty about the β's is likely to be somewhat larger. Significant values of $\hat{\beta}_r$ are found in all countries except for Germany and Japan. For most countries, a considerably lower degree of fractional integration is found compared with the no break case. Also, the standard deviation of the residuals $\hat{\sigma}_\epsilon$ and of the underlying process $\hat{\sigma}_z$ is smaller, as expected after inclusion of extra parameters.

We also calculate the LM and Wald test statistics for the absence of level shifts. For each country, the value of the statistic and the corresponding p-value are reported. For the calculation of the p-value, it is assumed that the statistic follows a χ_k^2 distribution, with k the hypothesized number of structural mean shifts. The LM and Wald test both point in the same direction. The hypothesis of no breaks seems to be rejected convincingly for five out of seven countries. Finally, notice that for the U.S. and the U.K., the Beran test statistic is getting worse. However, adding AR or MA components to our maintained model does not yield improvement.

The final panel of Table 3 concerns two more breaks, in July 1976 and January 1979. For Canada or France, no dramatic changes occur (as compared to the two break case). For Germany it is interesting to see that a temporally higher inflation seems to be found between 1979 and 1982. The Wald and LM tests however do not reject the null of no breaks against the alternative of four breaks. For Italy and the U.K. these extra breakpoints do not lead to a strong change. Japan, however, seems to have undergone higher inflation in the period

Table 3: ARFIMA models (with and without level shifts) for monthly inflation rates in the G7 countries

Parameter	Canada		France		Germany		Italy		Japan		UK		USA	
$\hat{\phi}_1$	-0.140	(0.067)	-0.055	(0.073)	0.033	(0.080)	-0.231	(0.060)	-0.224	(0.067)	-0.157	(0.069)	-0.268	(0.060)
$\hat{\phi}_{12}$	0.228	(0.052)	0.139	(0.053)	0.096	(0.052)	-0.093	(0.050)	0.160	(0.049)	0.145	(0.050)	0.078	(0.046)
$\hat{\phi}_{13}$					0.144	(0.051)								
$\hat{d}$	0.324	(0.049)	0.357	(0.056)	0.207	(0.064)	0.445	(0.044)	0.267	(0.056)	0.385	(0.054)	0.501	(0.050)
$\hat{\sigma}_\epsilon$	0.304	(0.011)	0.320	(0.012)	0.257	(0.009)	0.514	(0.018)	0.694	(0.025)	0.530	(0.019)	0.239	(0.008)
$\hat{\sigma}_z$	0.382		0.407		0.286		0.672		0.737		0.646		0.328	
LL	-87.971		-107.836		-22.977		-290.133		-407.146		-302.569		7.598	
Beran/p_{Beran}	0.341	0.157	0.313	0.598	0.337	0.206	0.290	0.888	0.314	0.571	0.349	0.087	0.341	0.137
Two shifts	Canada		France		Germany		Italy		Japan		UK		USA	
$\hat{\phi}_1$	-0.021	(0.077)	0.056	(0.089)	0.065	(0.086)	-0.124	(0.074)	-0.200	(0.071)	-0.095	(0.077)	-0.205	(0.068)
$\hat{\phi}_{12}$	0.224	(0.052)	0.137	(0.052)	0.087	(0.052)	-0.118	(0.051)	0.158	(0.050)	0.127	(0.051)	0.066	(0.046)
$\hat{\phi}_{13}$					0.133	(0.051)								
$\hat{d}$	0.114	(0.062)	0.186	(0.076)	0.163	(0.072)	0.274	(0.063)	0.225	(0.062)	0.276	(0.066)	0.400	(0.059)
$\hat{\sigma}_\epsilon$	0.292	(0.011)	0.313	(0.011)	0.256	(0.009)	0.502	(0.018)	0.691	(0.025)	0.521	(0.019)	0.231	(0.008)
$\hat{\beta}_{1973:07}$	1.639	(0.214)	1.244	(0.292)	0.326	(0.261)	1.554	(0.276)	0.368	(0.239)	1.289	(0.294)	1.851	(0.350)
$\hat{\beta}_{1982:07}$	-1.258	(0.235)	-1.446	(0.287)	-0.468	(0.315)	-0.956	(0.302)	-0.521	(0.268)	-0.723	(0.344)	-1.708	(0.354)
$\hat{\sigma}_z$	0.304		0.334		0.275		0.532		0.716		0.557		0.259	
LM/p_{LM}	290.144	0.000	45.950	0.000	2.650	0.266	40.727	0.000	4.936	0.085	29.972	0.000	49.863	0.000
Wald/p_{Wald}	64.957	0.000	32.207	0.000	2.957	0.228	39.890	0.000	4.357	0.113	19.557	0.000	33.786	0.000
LL	-74.028		-99.602		-21.911		-282.305		-405.319		-295.851		21.117	
Beran/p_{Beran}	0.312	0.608	0.309	0.660	0.331	0.291	0.266	0.988	0.315	0.561	0.354	0.060	0.354	0.044
Four shifts	Canada		France		Germany		Italy		Japan		UK		USA	
$\hat{\phi}_1$	-0.023	(0.077)	0.075	(0.092)	0.051	(0.087)	-0.127	(0.075)	-0.126	(0.082)	-0.111	(0.079)	-0.202	(0.068)
$\hat{\phi}_{12}$	0.221	(0.052)	0.130	(0.052)	0.093	(0.053)	-0.117	(0.051)	0.140	(0.051)	0.135	(0.052)	0.059	(0.047)
$\hat{\phi}_{13}$					0.138	(0.052)								
$\hat{d}$	0.106	(0.062)	0.161	(0.080)	0.175	(0.076)	0.274	(0.065)	0.117	(0.075)	0.297	(0.075)	0.381	(0.058)
$\hat{\sigma}_\epsilon$	0.291	(0.010)	0.312	(0.011)	0.255	(0.009)	0.502	(0.018)	0.682	(0.025)	0.520	(0.019)	0.230	(0.008)
$\hat{\beta}_{1973:07}$	1.452	(0.279)	1.277	(0.347)	0.028	(0.352)	1.403	(0.358)	0.922	(0.247)	1.086	(0.410)	1.653	(0.402)
$\hat{\beta}_{1976:07}$	0.074	(0.331)	-0.227	(0.385)	0.114	(0.401)	0.107	(0.387)	-0.923	(0.308)	0.185	(0.493)	-0.248	(0.440)
$\hat{\beta}_{1979:01}$	0.432	(0.319)	0.443	(0.386)	0.597	(0.378)	0.247	(0.368)	-0.033	(0.296)	0.173	(0.408)	0.974	(0.435)
$\hat{\beta}_{1982:07}$	-1.529	(0.275)	-1.638	(0.335)	-0.763	(0.351)	-1.117	(0.345)	-0.335	(0.252)	-0.797	(0.379)	-2.097	(0.394)
$\hat{\sigma}_z$	0.302		0.329		0.277		0.531		0.689		0.570		0.255	
LM/p_{LM}	852.230	0.000	49.741	0.000	7.099	0.131	49.586	0.000	29.441	0.000	31.151	0.000	64.879	0.000
Wald/p_{Wald}	73.176	0.000	39.512	0.000	6.154	0.188	41.029	0.000	22.897	0.000	17.873	0.001	42.317	0.000
LL	-72.667		-98.978		-20.130		-281.860		-400.906		-295.601		23.560	
Beran/p_{Beran}	0.312	0.607	0.306	0.705	0.332	0.278	0.268	0.986	0.321	0.458	0.354	0.059	0.362	0.020

Note: Standard errors between parentheses

from 1973 until 1976, which is the first oil crisis period. The breaksizes at those moments are about equal in absolute size and opposite in sign. Both tests point out that allowing for four breaks should be preferred to the assumption of no level shifts. For U.S. inflation the four breaks do seem to matter. The degree of fractional integration decreases a little further, and both LM and Wald test statistics obtain larger values, indicating strong evidence for level shifts. On the other hand, the problems indicated by the Beran test statistic for residual serial correlation increase.

Allowing for level shifts is seen to have a huge effect on the degree of fractional integration. In Canada, two breaks suffice to have the degree of integration diminish to a level that is no longer significant. In Japan, the first two breakpoints chosen do not seem to fit well the moments at which the mean level of inflation underwent a change. However, in the setting of four breakpoints, a high inflation period is neutralized, and the resulting z_t series displays no significant fractional integration. In France and Germany, $\hat{d}$ decreases to a level that usually is considered as not a strong indication of the presence of long memory, although the parameter itself is still significantly different from zero. In Italy, the estimate of d in a model with two level shifts is notably lower than in the pure ARFIMA model. This also holds true for the U.K. and U.S.

6 Conclusions

In this paper we investigated the effect a level shift can have on the apparent long memory characteristics. Especially for inflation rates, where long memory seems to exist, level shifts because of exogenous shocks may also have occurred. A framework for combining level shifts and long memory was put forward. In Section 4, a simulation study was performed to investigate the possible effects of a level shift on the estimate of the long memory parameter d in the ARFIMA model. A shock of only one standard deviation of the underlying series already could lead to the erroneous impression that long memory was present in the data. Encouraged by the results of the simulations, an investigation of the inflation rates in the countries of the G7 was performed. Where a pure ARFIMA model replicates previous results, that is, a significant value of $\hat{d}$ is found in several countries, addition of a set of level shifts did decrease the degree of fractional integration in various countries.

The results of the simulations indicated a size distortion for the LM test. In the controlled environment of a simulation, adjusting for this distortion is possible. Even if the size is controlled, the power of the test is not impressive. This is a problem that is hard to escape. Fractional integration and the occurrence of level shifts can be quite hard to distinguish in samples of medium size. Indeed, an unreported investigation using a larger sample did lead to better results for the size and power of the test, though improvement was slow. A possible way of improving the small sample results would be to work with the exact Maximum Likelihood method, as propagated by Sowell (1992).

In our empirical work, timing of the breakpoints was taken to be fixed exogenously. This may be seen as a drawback, and hence, endogenizing the breaks would be an interesting path of further research.

A Calculating the likelihood and the test statistics

In the calculation of the likelihood and the test statistics, several elements are combined. In this appendix the way the elements of these functions were calculated, are given.
In equation (11), the power spectrum was given as

$$f(\lambda;\Psi) = \frac{\sigma_\epsilon^2}{2\pi}\left|1-e^{i\lambda}\right|^{-2d}\frac{\left|\Theta(e^{i\lambda})\right|^2}{\left|\Phi(e^{i\lambda})\right|^2} \qquad (10)+(11)$$

The factors in this formula are calculated using

$$e^{-i\lambda} = \cos\lambda - i\sin\lambda \quad \Leftrightarrow \quad e^{i\lambda} = \cos\lambda + i\sin\lambda \qquad (21)$$

$$\begin{aligned}\left|1-e^{i\lambda}\right|^{-2d} &= \left|1-\cos\lambda - i\sin\lambda\right|^{-2d} \\ &= \left(1-2\cos\lambda+\cos^2\lambda+\sin^2\lambda\right)^{-d} = (2-2\cos\lambda)^{-d} \qquad (22)\end{aligned}$$

$$\left|\Theta(e^{i\lambda})\right|^2 = \left|1+\sum_{j=1}^{q}\theta_j e^{ij\lambda}\right|^2 = (1+\sum_{j=1}^{q}\theta_j\cos j\lambda)^2 + (\sum_{j=1}^{q}\theta_j\sin j\lambda)^2 \qquad (23)$$

$$\left|\Phi(e^{i\lambda})\right|^2 = \left|1-\sum_{j=1}^{p}\phi_j e^{ij\lambda}\right|^2 = (1-\sum_{j=1}^{p}\phi_j\cos j\lambda)^2 + (-\sum_{j=1}^{p}\phi_j\sin j\lambda)^2 \qquad (24)$$

Before calculating the periodogram, the data following the underlying ARFIMA process are calculated from $y_t, t = 1,..,T$, the moments of the k breaks $\tau_r, r = 1,..,k$, and the sizes of the corresponding level shifts $\mu_r, r = 1,..,k$. Assuming d lies in the stationary realm (see remarks in section 3), and reiterating a slightly adjusted version of equations (6)-(8), we have

$$z_t^a(\mu,d) = z_t^b(\mu,d) = y_t - \sum_{r=1}^{k}\mu_r I_{\{t>\tau_r T\}} \qquad (6')$$

$$z_t(\mu,d) = z_t^b - \overline{z^b} \qquad (8')$$

Of this transformed data set, the periodogram is calculated as

$$\begin{aligned}I_z(\lambda;\mu) &= \frac{1}{2\pi T}\left|\sum_{t=1}^{T} z_t e^{it\lambda}\right|^2 \\ &= \frac{1}{2\pi T}\left((\sum_{t=1}^{T} z_t\cos t\lambda)^2 + (\sum_{t=1}^{T} z_t\sin t\lambda)^2\right) \qquad (25)\end{aligned}$$

Defining $T^* = \lfloor T/2 \rfloor, \lambda_j = \frac{2\pi j}{T}$ and weights δ_j as in section 3,

$$\delta_j = \begin{cases} \frac{1}{2} & \text{if } j = \frac{T}{2} = \left\lfloor \frac{T}{2} \right\rfloor \\ 1 & \text{else} \end{cases} \qquad j = 1,..,T^* \qquad (16')$$

all elements in the likelihood function are known:

$$\begin{aligned}\ln\mathcal{L}(y|\Phi,d,\Theta,\mu,\sigma_\epsilon^2) &= \ln\mathcal{L}(y|\Psi,\mu) \\ &= -T\ln 2\pi - \sum_{j=1}^{T^*}\delta_j \ln f(\lambda_j;\Psi) - \sum_{j=1}^{T^*}\delta_j\frac{I_z(\lambda_j;\mu)}{f(\lambda_j;\Psi)} \qquad (15')\end{aligned}$$

The analytical gradient of the loglikelihood with respect to the parameters of the model was used, during the optimization routine for the model and in the calculation of the test statistics. Separating the vector of ARFIMA parameters Ψ and the level shifts μ, we find

$$\frac{\partial \ln \mathcal{L}(y|\Psi,\mu)}{\partial \psi_r} = -\sum_{j=1}^{T^*} \delta_j \frac{1}{f(\lambda_j;\Psi)} \frac{\partial f(\lambda_j;\Psi)}{\partial \psi_r} + \sum_{j=1}^{T^*} \delta_j \frac{I_z(\lambda_j;\mu)}{f(\lambda_j;\Psi)^2} \frac{\partial f(\lambda_j;\Psi)}{\partial \psi_r} \tag{26}$$

$$\frac{\partial \ln \mathcal{L}(y|\Psi,\mu)}{\partial \mu_r} = -\sum_{j=1}^{T^*} \delta_j \frac{1}{f(\lambda_j;\Psi)} \frac{\partial I_z(\lambda_j;\mu)}{\partial \mu_r} \tag{27}$$

The derivatives of the power spectrum w.r.t. the ARFIMA parameters in Ψ are

$$\frac{\partial f(\lambda;\Psi)}{\partial \phi_r} = f(\lambda;\Psi) \times \frac{-1}{\left|\Phi(e^{i\lambda})\right|^2} \frac{\partial \left|\Phi(e^{i\lambda})\right|^2}{\partial \phi_r} \tag{28}$$

$$\frac{\partial f(\lambda;\Psi)}{\partial \theta_r} = f(\lambda;\Psi) \times \frac{1}{\left|\Theta(e^{i\lambda})\right|^2} \frac{\partial \left|\Theta(e^{i\lambda})\right|^2}{\partial \theta_r} \tag{29}$$

$$\frac{\partial f(\lambda;\Psi)}{\partial d} = f(\lambda;\Psi) \times (-\ln(2-2\cos\lambda)) \tag{30}$$

$$\frac{\partial f(\lambda;\Psi)}{\partial \sigma_\epsilon} = f(\lambda;\Psi) \times \frac{2}{\sigma_\epsilon} \tag{31}$$

Two derivatives are left unspecified in the equations (28) and (29). These follow, after some tedious but simple algebra, as

$$\frac{\partial \left|\Phi(e^{i\lambda})\right|^2}{\partial \phi_r} = -2\left((1-\sum_{k=1}^{p}\phi_k\cos k\lambda)\cos r\lambda - \sum_{k=1}^{p}\phi_k \sin k\lambda \sin r\lambda\right) \tag{32}$$

$$\frac{\partial \left|\Theta(e^{i\lambda})\right|^2}{\partial \theta_r} = 2\left((1+\sum_{k=1}^{q}\theta_k\cos k\lambda)\cos r\lambda + \sum_{k=1}^{q}\theta_k \sin k\lambda \sin r\lambda\right) \tag{33}$$

The last element that was left unknown in these equations was the derivative of the periodogram of z w.r.t. the value of the breaks. This last derivative is found to be

$$\frac{\partial I_z(\lambda,\mu)}{\partial \mu_r} = \frac{1}{\pi T}\left(\left(\sum_{t=1}^{T} z_t \cos t\lambda\right)\left(\sum_{t=1}^{T}\frac{\partial z_t}{\partial \mu_r}\cos t\lambda\right) + \left(\sum_{t=1}^{T} z_t \sin t\lambda\right)\left(\sum_{t=1}^{T}\frac{\partial z_t}{\partial \mu_r}\sin t\lambda\right)\right) \tag{34}$$

with

$$\frac{\partial z_t}{\partial \mu_r} = I_{\{\tau_r T > t\}} + \frac{T-\tau_r T}{T} \approx I_{\{\tau_r T > t\}} + 1 - \tau_r \tag{35}$$

Constructing the gradient is done by combining the equations. Although it is possible to derive the analytical second derivative as well, this would become even harder. For calculation of the hessian, the numerical first derivative of the (analytical) gradient is used.

B Generating an ARFIMA process

Beran (1994) describes a method, originating from Davies & Harte (1987), to generate T observations from a stationary Gaussian model, given the autocovariances $\gamma(0), .., \gamma(T-1)$. The following steps are taken:

1. Calculate the autocovariances of the model, $\gamma(0), \gamma(1), .., \gamma(T-2), \gamma(T-1)$, and define

$$\gamma^*(j) = \begin{cases} \gamma(j) & \text{if } j \leq T-1 \\ \gamma(T-j) & \text{if } T-1 < j \leq 2T-3 \end{cases} \qquad j = 0, .., 2T-3 \tag{36}$$

2. Calculate the fourier transform of the $\gamma^*(\cdot)$,

$$g_k = \sum_{j=0}^{2T-3} \gamma^*(j) e^{ij\lambda_k} \tag{37}$$

   for $k = 0, .., 2T-3$ and $\lambda_k = \frac{2\pi k}{2T-2}$. These g_k should all result to be positive.

3. Generate random normals $U_1, .., U_{T-2}$ and $V_1, .., V_{T-2}$ all independent and with variance 1; simulate U_0 and U_{T-1} as independent of all other values and eachother, with variance 2. Define $V_0 = V_{T-1} = 0$. Construct random variables Z_k in the complex plane as

$$Z_k = \begin{cases} U_k + iV_k & k = 0, .., T-1 \\ U_{2T-k-2} - iV_{2T-k-2} & k = T, .., 2T-3 \end{cases} \tag{38}$$

4. The observations y_t are now calculated as

$$y_t = \frac{1}{\sqrt{T-1}} \sum_{k=0}^{2T-3} \sqrt{g_k} e^{i(t-1)\lambda_k} Z_k \tag{39}$$

The procedure for generating an ARFIMA(p, d, q) start with generating a ARFIMA(0,d,0) as described above, using the autocovariance function of the process as follows (Gradshteyn & Ryzhik 1965, p. 372):

$$\gamma(k) = \sigma_\epsilon^2 \frac{(-1)^k \Gamma(1-2d)}{\Gamma(k-d+1)\Gamma(1-k-d)} \tag{40}$$

for the covariances needed in (36). Then, the y_t that are generated should be used instead of the usual disturbances in a routine generating an ARMA(p, q). The result will be a series y_t' distributed as a ARFIMA(p, d, q).

References

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

Bai, J. (1996), Estimation of a change point in multiple regression models. Unpublished manuscript, M.I.T.

Baillie, R. T. (1996), 'Long memory processes and fractional integration in econometrics', *Journal of Econometrics* **73**, 5–59.

Baillie, R. T., Chung, C.-F. & Tieslau, M. A. (1996), 'Analysing inflation by the fractionally integrated ARFIMA-GARCH model', *Journal of Applied Econometrics* **11**, 23–40.

Beran, J. (1994), *Statistics for Long-Memory Processes*, New York: Chapman & Hall.

Cheung, Y.-W. (1993), 'Tests for fractional integration: A Monte Carlo investigation', *Journal of Time Series Analysis* **14**, 331–345.

Davies, R. B. & Harte, D. S. (1987), 'Tests for Hurst effect', *Biometrika* **74**, 95–102.

Gradshteyn, I. S. & Ryzhik, I. M. (1965), *Tables of Integrals, Series and Products*, London: Academic Press.

Granger, C. W. J. & Joyeux, R. (1980), 'An introduction to long-range time series models and fractional differencing', *Journal of Time Series Analysis* **1**, 15–30.

Harvey, A. C. (1989), *Forecasting, Structural Time Series Models and the Kalman Filter*, Cambridge: Cambridge University Press.

Hassler, U. & Wolters, J. (1995), 'Long memory in inflation rates: International evidence', *Journal of Business & Economic Statistics* **13**, 37–46.

Hauser, M. A. (1997), Maximum likelihood estimators for ARFIMA models: A Monte Carlo study. Unpublished manuscript.

Hidalgo, J. & Robinson, P. M. (1996), 'Testing for structural change in a long memory environment', *Journal of Econometrics* **70**, 159–174.

Hosking, J. R. M. (1981), 'Fractional differencing', *Biometrika* **68**, 165–176.

Ooms, M. (1996), Long memory and seasonality in US consumer price inflation: An empirical investigation at varying levels of aggregation. Unpublished manuscript, Erasmus University Rotterdam.

Ooms, M. & Hassler, U. (1997), 'On the effect of seasonal adjustment on the log-periodogram regression', *Economics Letters* **56**, 135–141.

Perron, P. (1989), 'The great crash, the oil price shock, and the unit root hypothesis', *Econometrica* **57**, 1361–1401.

Perron, P. & Vogelsang, T. J. (1992), 'Testing for a unit root in a time series with a changing mean: Corrections and extensions', *Journal of Business and Economic Statistics* **10**, 467–469.

Sowell, F. (1992), 'Maximum likelihood estimation of stationary univariate fractionally integrated time series models', *Journal of Econometrics* **53**, 165–188.