Detecting a Change in Inflation Persistence in the Presence of Long Memory

Evidence from a New Approach

Uwe Hassler*and Barbara Meller

Goethe University Frankfurt

First Version, February 25, 2009

Abstract

Inflation stability is an important issue for monetary policy. However, there is still no common agreement on whether or not US inflation persistence is stable. Most studies ignore the fractionally integrated nature of inflation which may lead to misspecification and incorrect test results. In the presence of fractional integration, we propose a new test that finds a break in persistence at unknown dates and establishes the significance of the break. Using this test, we find significant changes in US inflation persistence taking place in 1973 and 1980.

Keywords: Long Memory; Test for unknown Break in Fractional Integration; Change in Persistence; Inflation Dynamics

JEL classification: E31; C22

* **Corresponding Author:** Uwe Hassler, Goethe Universität Frankfurt, Grueneburgplatz 1, 60323 Frankfurt am Main, Germany. E-mail: hassler@wiwi.uni-frankfurt.de.

1 Introduction

Inflation persistence is an important issue for economists and especially for central bankers. If inflation persistence is high, a shock to the price level will increase inflation for a long time period. Inflation might even behave like a random walk in which case central banks cannot control it. On the contrary, if inflation is integrated of order zero, then inflation reverts back to its initial level fast, after a shock occurred. In that case, central bankers do not even have to act in response to a shock on inflation. Inflation persistence is thought to reveal the credibility of central banks. Low inflation persistence indicates that inflation expectations are anchored. People expect that inflation will return to the target inflation quickly after it has deviated (Blinder et al., 2008). Lastly it is important to study inflation persistence because it is at the heart of empirical tests on monetary policy rules, on the hypothesis of the natural rate of inflation and on the natural rate of unemployment (Cogley and Sargent, 2005).

Not only the level of inflation persistence is important in economic analyses but also its stability. There are several sources and consequences of a potential change in inflation persistence. Mishkin (2007) argues that (un)anchored inflation expectations were at the heart of a fall (rise) in inflation persistence. Further, anchored inflation expectations were brought about by an aggressive monetary policy, dedicated to price stability. However, a decrease in inflation persistence might not only be a good sign but could also inherit some dangers. Pivetta and Reis (2007) argue that a decrease in inflation persistence could lead to a rejection of long-run monetary neutrality which might mislead the central bank to exploit the trade-off between inflation and unemployment. O'Reilly and Whelan (2005) point to the importance of inflation persistence stability for policy analysis. In view of the Lucas' critique, they test whether inflation persistence in the Euro area has changed due to drastic changes in the monetary policy regime. If inflation persistence is not stable, O'Reilly and Whelan (2005) argue, this will affect forecasting and limit the usefulness of backward-looking models of inflation. As a consequence, if parameters were unstable, central banks should use more structural models that can cope with shifting dynamics.[1]

There are different ways to measure inflation persistence. Most often, it is

[1] They come to the conclusion, however, that this is not necessary as inflation persistence is stable.

measured as the sum of the lagged coefficients on inflation. If the sum is close to one, then inflation is thought to be very persistent. Alternatively, some studies use the largest autoregressive root or the half-life of a unit shock.[2] Yet other studies investigate inflation persistence by testing whether inflation is I(0) or I(1). Recently, Stock and Watson (2007) presented another measure of inflation persistence using a time-varying first-order integrated moving average model. They decompose the innovation variance of this model into a permanent component and a transitory component. The permanent component represents inflation persistence. Cogley and Sargent (2005) measure inflation persistence using the normalized spectrum of inflation at zero frequency. Lastly, inflation persistence can be measured by the long memory or fractional integration parameter of inflation.

Since the early evidence by Hassler and Wolters (1995) and Baillie et al. (1996) it is well established that inflation exhibits long memory. In view of this evidence, Gadea and Mayoral (2006) show that the sum of the AR coefficients and unit root tests could lead to erroneous conclusions. Similarly, the normalized spectral density of inflation at zero frequency, used by Cogley and Sargent (2005), is not a reliable measure of persistence in the presence of fractional integration, see Hauser et al. (1999). As a straight forward consequence, one also has to account for fractional integration when measuring the *change* in inflation persistence. Kumar and Okimoto (2007) argue that tests for inflation persistence stability using unit root tests or autoregressive coefficients (e.g. Pivetta and Reis (2007)) might lead to incorrect conclusions in the presence of fractional integration. Of the same tenor, Stock and Watson (2007) might draw incorrect conclusion when using Andrews' (1993) critical values to determine the significance of a beak in persistence but not accounting for fractional integration. Andrews' critical values are not likely to be applicable in the presence of fractional integration.[3]

As Kumar and Okimoto (2007), we use long memory measures to determine a change in inflation persistence. However, we do not use visual judgment of rolling window estimates or analyze two exogenously split subsample. Our paper contributes to the existing literature by deriving and applying a new test for the timing and significance of a break in fractional

[2]For a detailed overview on these different measures of inflation persistence, the reader is referred to Mishkin (2007) and Pivetta and Reis (2007).

[3]Hassler and Olivares (2008) have Monte Carlo evidence, that their supremum statistic does not obey asymptotically the distribution established by Andrews (1993) because their series exhibits long memory.

integration. The test is based on the maximum of the squared t-statistic. It accounts for fractional integration by differencing the series before applying a lag-augmented LM-test. We therefore apply our test to an I(0) series so that we can use Andrews (1993)' critical values. Another advantage of our approach is that it is able to detect a break in the long memory parameter close to the boundaries of the sample. This is because our test does not rely on the estimation of the long memory parameter before and after the break. If this was the case, we could only search breaks in a smaller range around the middle of the sample because those estimates are measured in the frequency domain and require more observations than an OLS regression. We apply our new test to US inflation rates and find that there are two significant changes in persistence in 1973 and in 1980.

The rest of the paper is organized as follows. In the next section, we will give an overview on tests against a break in long memory. In section three, we describe our new test procedure as well as its finite sample properties. Section four tests for a change in US inflation persistence and reconciles our results with the existing literature. Finally, section five concludes.

2 Review of (tests against) a break in persistence

2.1 A break in persistence

First, we recall how long memory or equivalently fractional integration is defined. Under the null hypothesis the observed time series y_t is integrated of order d,

$$\Delta^d y_t = (1-L)^d y_t = e_t = \sum_{j=0}^{\infty} c_j \varepsilon_{t-j} \ \sim \ I(0)\,, \tag{1}$$

where e_t is a stationary and invertible short memory process, ε_t is white noise and c_j are the Wold coefficients. L denotes the conventional lag operator, and fractional differences are defined through the usual binomial expansion,

$$(1-L)^d = 1 - dL - \frac{d\,(1-d)}{2} L^2 - \frac{d\,(1-d)\,(2-d)\,L^3}{6} - \cdots = \sum_{j=0}^{\infty} \pi_j L^j\,,$$

where the coefficients π_j are defined implicitly. Similarly, one may expand the inverse filter,

$$y_t = (1-L)^{-d} e_t = \sum_{j=0}^{\infty} \psi_j e_{t-j} = \sum_{j=0}^{\infty} \gamma_j \varepsilon_{t-j} \,,$$

where the impulse response function γ_j is given by convolution. Following Brockwell and Davis (1991) one can show with some constant γ

$$\gamma_j \sim \gamma \, j^{d-1} \quad \text{as } j \to \infty \,. \tag{2}$$

For $d = 1$, past shocks have a permanent effect, while for $0.5 \leq d < 1$ we observe nonstationarity with transitory shocks[4], $\gamma_j \to 0$ as $j \to \infty$; finally for $0 < d < 0.5$ the impulse responses die out fast enough to be square-summable resulting in a stationary process, but still dying out so slowly that γ_j is not summable, why we speak of long memory. In any case, d is interpreted as degree of persistence or memory parameter measuring how slowly the effect of past shocks dies out.

Working with finite samples of size T the theoretical difference operator from (1) has to be adjusted. Given a finite past starting with the first observation y_1, the infinite expansion is truncated in practice. We call the truncated differences $\Delta_t^d(y)$ and denote it by x_t for brevity,

$$x_t = \Delta_t^d(y) = \sum_{j=0}^{t-1} \pi_j y_{t-j} \,, \; t = 1, \ldots, T \,. \tag{3}$$

Next, we model a break in persistence and review the literature on tests of such break. As alternative hypothesis to (1) we model

$$y_t = \begin{cases} y_{0,t} \sim I(d)\,, & t = 1, \ldots, [\lambda T] \\ y_{1,t} \sim I(d+\theta)\,, & t = [\lambda T]+1, \ldots, T \end{cases} \tag{4}$$

with unknown potential break fraction $\lambda \in [0, 1]$. The null hypothesis may now be recast as

$$H_0 : \theta = 0 \,. \tag{5}$$

[4]Such a feature is sometimes called "mean-reversion" although Phillips and Xiao (1999) argue that this is a misnomer given the nonstationarity.

If such a sudden break is considered as too extreme in practice, there still may be a smooth transition,

$$y_{1,t} = y_{0,[\lambda T]} + \sum_{j=0}^{t-1-[\lambda T]} \psi_j \, e_{t-j} \,, \quad t = [\lambda T]+1, \ldots, T \,, \tag{6}$$

where ψ_j are the coefficients from expanding $(1-L)^{-d-\theta}$.

2.2 Testing against a break in persistence

There exists a considerable body of literature that deals with a break from an I(0) to an I(1) process, and vice versa, starting with tests pioneered by Kim (2000) and Busetti and Taylor (2004)

For general d with $|d| < 0.5$ Beran and Terrin (1996) consider the so-called Whittle estimator of d, which is a parametric frequency domain approximation to maximum likelihood [ML]. In a long memory environment the Whittle estimator has been discussed by Fox and Taqqu (1986), and it was shown to be $\sqrt{T}$-consistent following a limiting normal distribution. Beran and Terrin (1996) suggest to compute from non-overlapping subsamples,

$$\begin{array}{lll} \widehat{d}_0(\lambda) & \text{for} & t = 1, \ldots, [\lambda T] \\ \widehat{d}_1(\lambda) & \text{for} & t = [\lambda T]+1, \ldots, T \end{array}$$

where λ is varied systematically (with $0 < \pi < 0.5$),

$$\lambda \in (\pi, 1-\pi) \,.$$

The test statistic is $\sup |D(\lambda)|$ over all potential break points with the difference

$$D(\lambda) = \frac{\sqrt{T}\sqrt{\lambda\,(1-\lambda)}(\widehat{d}_0(\lambda) - \widehat{d}_1(\lambda))}{\kappa}$$

with κ from the Fisher information of ML estimation. The limiting distribution under H_0 for $|d| < 0.5$ is given by Beran and Terrin (1999) as the supremum of a normalized Brownian bridge, depending on the boundary choice of the interval $(\pi, 1-\pi)$. Following Andrews (1993) one would compute the square instead of the modulus (where the limit process is sometimes also called tied-down Bessel process),

$$\sup_{\lambda \in (\pi, 1-\pi)} D^2(\lambda) \tag{7}$$

in order to use the extensively tabulated critical values in Andrews (1993, Table I).

In applied work, it is not always practical to set up a complete parametric specification as required for Whittle estimation. Therefore, Robinson (1995) considers the so-called *local* Whittle estimator of d, which is semiparametric. Although limiting normality was established by Robinson (1995), the local Whittle estimator converges to the true parameter value only at rate m, where m denotes the bandwidth of lower order than T. Shimotsu and Phillips (2005) improve the local Whittle estimator so that the estimator is also valid for $d > 0.5$. They term this improved estimator *exact* local Whittle estimator, see also Shimotsu and Phillips (2006). Let $\widehat{\delta}_i(\lambda)$ with bandwidth m_i, $i = 0, 1$, denote the local Whittle estimators from the respective subsamples.

For *fixed* break fraction λ given exogeneously, Shimotsu (2006) establishes asymptotic independence of the estimators, such that

$$\Delta(\lambda) = \frac{2\sqrt{m_0 m_1}}{\sqrt{m_0 + m_1}}(\widehat{\delta}_0(\lambda) - \widehat{\delta}_1(\lambda))$$

has a limiting normal distribution. When the break point is unknown, one may want to compute in analogy to (7)

$$\sup_{\lambda \in (\pi, 1-\pi)} \Delta^2(\lambda)\,.$$

Hassler and Olivares (2008) have Monte Carlo evidence, however, that this supremum statistic $\sup \Delta^2(\lambda)$ does not obey asymptotically the distribution established by Andrews (1993). This suggests that the assumption of $\sqrt{T}$-consistent estimators (see Andrews, 1993, Theorem 1) is not only sufficient but also necessary, because this assumption is not met by the semiparametric local Whittle estimator. This motivates the test proposed in the next section, which relies on $\sqrt{T}$-consistent estimation.

3 A new test against a break in persistence

3.1 Derivation

Our new test is based on the regression-based Lagrange Multiplier [LM] test proposed by Demetrescu, Kuzin and Hassler [DKH] (2008). They propose a test for any prespecified value d, which they call augmented LM test [ALM] because it relies on a regression augmented by lags. The LM test for fractional integration was pioneered by Robinson (1994). It is locally most powerful against fractional alternatives. The simple regression-based version with lag-augmentation by Demetrescu et al. (2008) can approximate with growing p many AR(∞) processes with short memory.[5] The original LM tests by Robinson (1994) in the time or frequency domain do not dispose of such a simple and general correction for short memory; cf. also Tanaka (1999). In order to test whether the true order of integration is significantly different from d, DKH estimate the following regression by OLS:

$$x_t = \widehat{\phi}\, x_{t-1}^* + \widehat{a}_1 x_{t-1} + \ldots + \widehat{a}_p x_{t-p} + \widehat{\varepsilon}_t, \tag{8}$$

with

$$x_{t-1}^* = \sum_{j=1}^{t-1} \frac{x_{t-j}}{j},$$

where x_t is the difference of y_t defined in (3). The null hypothesis translates into $\phi = 0$ and can be tested with a t-type test statistic t_ϕ, which follows asymptotically a standard normal distribution. All regressors in (8) are (asymptotically) stationary, and Demetrescu et al. (2008) prove that $\widehat{\phi}$ converges at rate $\sqrt{T}$. Hence, the OLS regression satisfies the assumptions for a sup test in line with Andrews (1993).

We now turn to our new test for a break in long memory. We extend equation (8) to include a dummy variable accounting for a possible break in long memory. We propose

$$x_t = \widehat{\phi} x_{t-1}^* + \widehat{\psi}\, D_t(\lambda) x_{t-1}^* + \sum_{j=1}^{p} \widehat{a}_j\, x_{t-j} + \widehat{\varepsilon}_t \tag{9}$$

with step dummy

$$D_t(\lambda) = \begin{cases} 0\,, & t = 1, \ldots, [\lambda T] \\ 1\,, & t = [\lambda T] + 1, \ldots, T \end{cases} .$$

[5] Since the regressors are not orthogonal, Demetrescu et al. (2008) advice against data-driven lag-length selection, as the model selection step affects subsequent inference about ϕ even asymptotically, see e.g. Leeb and Pötscher (2005). Instead they propose to choose the lag length p in (8) deterministically following the rule of thumb $p = \left[4(T/100)^{1/4}\right]$.

λ is varied over the interval $[\pi, 1-\pi]$, where $\pi \in (0, 1/2)$. In our Monte Carlo simulation as well as our application, we choose $\pi = 0.15$. The candidate for a break point is determined by

$$\hat{\lambda} = \arg\max(t_{\psi}^2(\lambda)) \, .$$

Next, the significance of the break at $\hat{\lambda}$ is tested. To that end, one evaluates the series of squared t-statistics for $\psi = 0$. The supremum,

$$\sup_{\lambda \in [\pi, 1-\pi]} t_{\psi}^2(\lambda) \, ,$$

can be compared with the familiar critical values by Andrews (1993).

The power of the test under the alternative $\theta \neq 0$ will be investigated by means of computer simulations. Before doing so we address as further case the situation where the null hypothesis is true ("no break in persistence"), but the the model in (4) or (1) is misspecified in that d is not the true order of integration. Instead we assume

$$\Delta^{d+\delta} y_t = e_t \, , \quad t = 1, \ldots, T \, , \tag{10}$$

where $\delta \neq 0$ is the degree of misspecification. Hence, the differences x_t from (3) are not $I(0)$ but rather $I(\delta)$, and the test regression (9) is performed with $I(\delta)$-variables. Consequently, the critical values from Andrews (1993) do not apply, although no change in persistence occurs. The effect of incorrect differencing under H_0 will be investigated experimentally, too.

3.2 Finite sample performance

In order to investigate the finite sample performance of our test for a change in persistence, we conduct a Monte Carlo study. In the basic set-up we simulate a time series with 1000 observations. A break in persistence takes place in the middle of the sample. The process follows an $I(0)$ before and an $I(\theta)$ process after the break:

$$x_t = \begin{cases} \varepsilon_t \, , & t = 1, \ldots, [\lambda_0 T] \\ (1-L)^{-\theta} \varepsilon_t \, , & t = [\lambda_0 T] + 1, \ldots, T \end{cases} \, , \tag{11}$$

where $\lambda_0 = 0.5$ is the true break fraction, θ is the true order of integration and ε_t is randomly drawn from a standard normal distribution. The break

Table 1. Rejection of the null hypothesis for different values of θ

θ	1% CV	5% CV	10% CV	$\overline{\hat{\lambda}}$	$\sigma(\hat{\lambda})$	MSE($\hat{\lambda}$)
-0.4	100%	100%	100%	0.49	0.03	0.00
-0.3	99%	100%	100%	0.49	0.05	0.00
-0.2	77%	91%	94%	0.49	0.09	0.01
-0.1	15%	34%	46%	0.49	0.17	0.03
0	1%	5%	10%	0.49	0.23	0.05
0.1	14%	32%	46%	0.50	0.18	0.03
0.2	72%	89%	93%	0.50	0.10	0.01
0.3	99%	100%	100%	0.50	0.05	0.00
0.4	100%	100%	100%	0.50	0.03	0.00

Notes: The table shows for different value of θ how often the null hypothesis was rejected in the Monte Carlo study using the 1%, 5% and 10% critical values. The table also reports the mean of the estimated break fraction, its variance and its mean squared error. Simulation is conducted under the basic set-up: $T = 1000$, $\lambda_0 = 0.5$, fractional noise.

in our set-up is especially hard to detect because the transition from short to long memory is smooth, as described in equation (6). We run regression (9) with $p = \left[4(T/100)^{1/4}\right] = 7$ lags and determine the supremum of the squared t-statistics for the hypothesis that $\psi = 0$. Table 1 shows the result of 1000 iterations: the rejection rate of the hypothesis that there is no break using Andrews' critical values for 1%, 5% and 10% significance, the mean of the estimated break fractions, their standard deviation and their mean squared error for different values of θ. Figure 1 visualizes the power and the size of the test for different values of θ.

The simulation results in Table 1 are in line with expectations. The larger the difference in the order of integration before and after the break, the easier the break is detected and correctly allocated. In other words, the larger θ in absolute terms, the higher the rejection rate and the smaller the MSE($\hat{\lambda}$). Overall, the performance of our test in a finite sample is satisfactory. The size of the test is good: 1.10%, 4.60%, 10.00% corresponding to the critical values of the 1%, 5% and 10% significance level. The power is extremely high if the difference in the long memory parameter before and after the break is

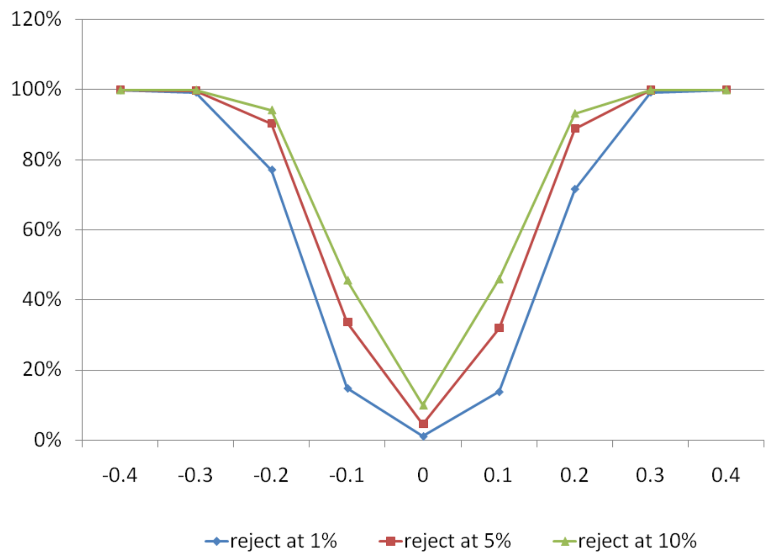

Figure 1. Basic set-up, rejection rates plotted against θ

greater than 0.3 is absolute terms. Even if the difference is only ± 0.2, the power is still high. Figure 1 depicts the symmetry of the rejection rates with respect to θ around zero.

Next, we investigate the performance of our test in the light of some variations in the simulation set-up. In Figure 2 the 5%-rejection rate of our test is plotted against θ for different values of the break fraction: $\lambda_0 \in \{0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8\}$. When the break fraction equals 0.2 or 0.8, the power of the test decreases notably if θ is smaller or equal to 0.2, in absolute terms. For all other cases, the power is high and the $MSE(\hat{\lambda})$s are comparable to the ones reported in Table 1.[6]

In Figure 3, the 5%-rejection rate is plotted against θ for different sample sizes, $T \in \{250, 500, 1000, 2000\}$. Unsurprisingly, the power of our test decreases as the sample size decreases. For $T = 250$, the test is only of limited use and for $T = 500$, the difference of the long memory parameters should be at least 0.3 for the test to have enough power.

The last variation of the simulation set-up is the distribution of the time

[6]Tables containing equivalent information as reported in Table 1 are available for all variations to the simulation set-up reported in this subsection.

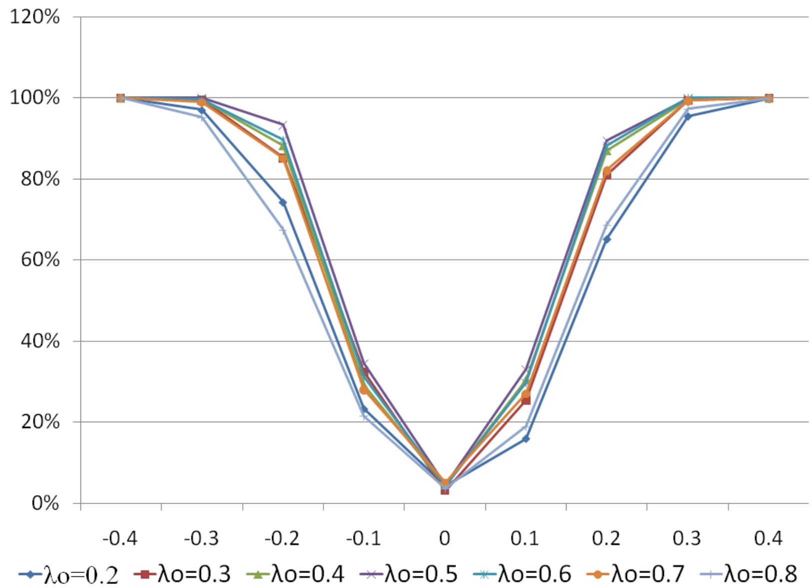

Figure 2. Different break fractions

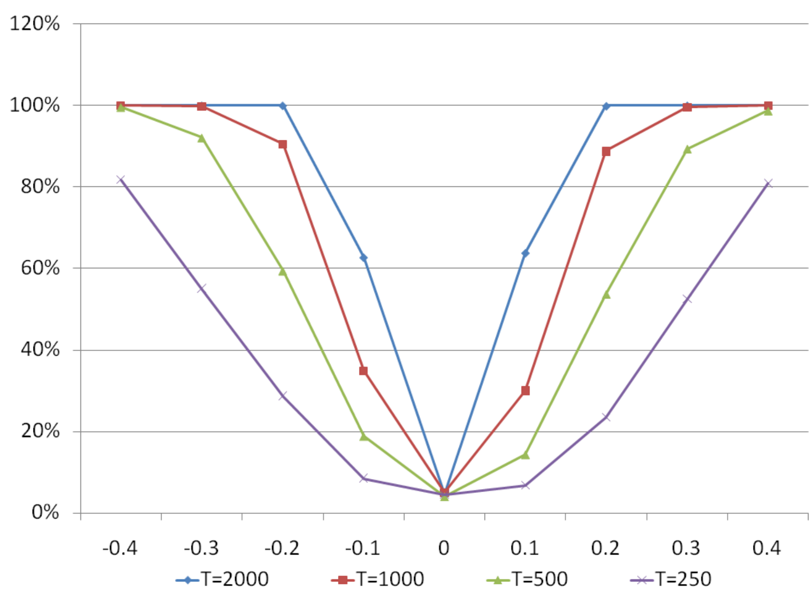

Figure 3. Different sample size

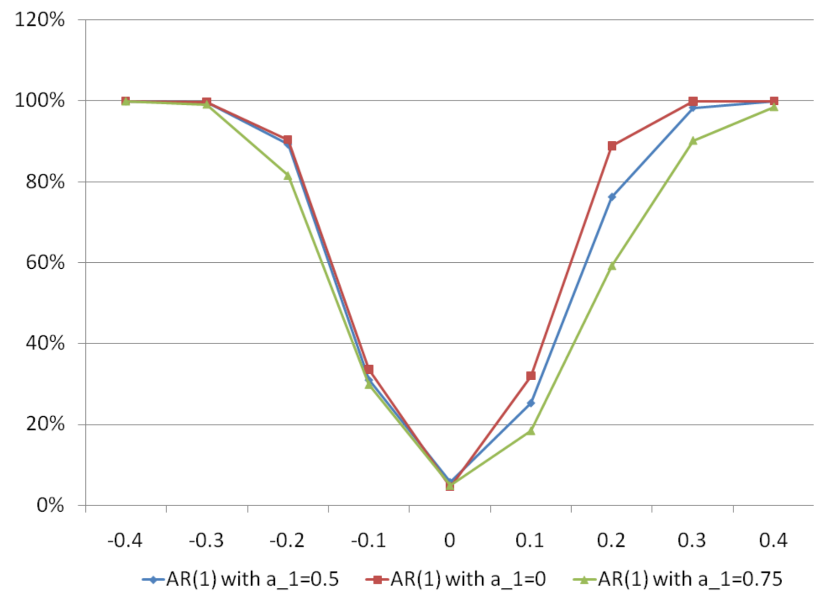

Figure 4. AR(1) process

series. In the definition of process x_t in equation (11), we replace ε_t by an AR(1) process e_t:

$$e_t = a_1 e_{t-1} + \varepsilon_t.$$

The AR(1) coefficient takes on the values 0, 0.5 and 0.75. Due to the lagged variables included in regression (9), the power of our test is only mildly affected. Only if the AR(1) coefficient is as high as 0.75, a break in persistence is not easily detected if $\theta \in [0.1, 0.2]$. The size of the test is unaffected.

Concluding, the power of our test depends especially on the difference of the order of integration before and after the break. If this difference is larger than 0.3, the power is good. Furthermore, the power is influenced by the sample size while it is almost unaffected by changes in the true break fraction or the value of the AR(1) coefficient. The size is very good throughout all simulation set-ups.

We now investigate the size of our test in case that no break occurs but the overall order of fractional integration is incorrectly specified under H_0, see (10). To that end, we simulate fractional noise processes having an unchanged order of integration $d + \delta$ with $\delta > 0$ while H_0 continues to assume $\delta = 0$. In other words, we difference the time series only by order

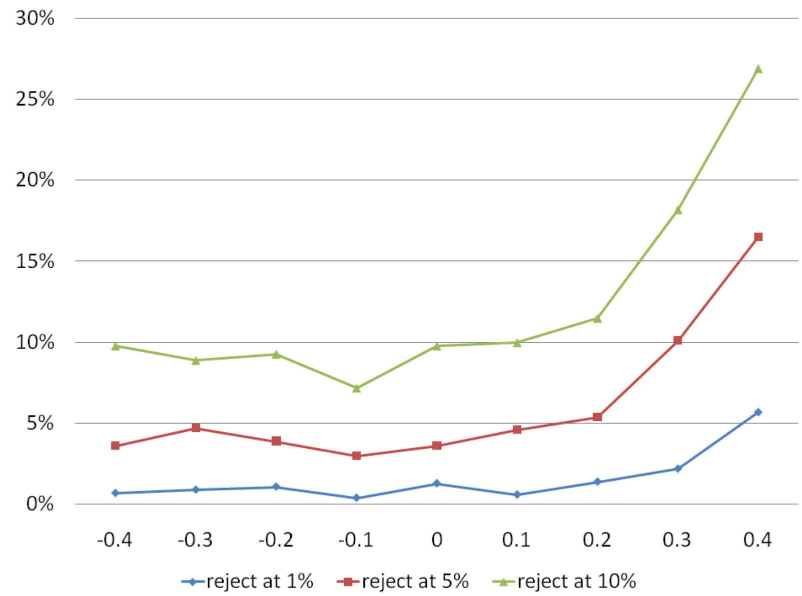

Figure 5. Size of the Test when Long Memory is Ignored

d and apply our test to the fractionally integrated process I(δ). The size of our test performs well for $\delta \leq 0.2$. However, for larger values of δ, the misspecification leads to a high rejection rate. Long memory is confounded with a break in fractional integration. Therefore, it is crucial to account for long memory before applying our test.

We conclude that our test works very well in our different settings if overall long memory is captured appropriately a priori.

4 Testing against a change in U.S. inflation persistence

4.1 Constructing inflation rates

We use monthly US CPI data collected by the Organization for Economic Cooperation and Development (OECD). The sample runs from January 1966 until June 2008 yielding 509 observations. Inflation is computed as follows:

$$\pi_t = 1200(\log(CPI_t) - \log(CPI_{t-1})).$$

We seasonally adjust inflation by subtracting seasonal means. In a next step, we test whether there is a shift in the overall mean:

$$\pi_t = \begin{cases} y_t + \mu_1 \, , & t = 1, \ldots, [\tau T] \\ y_t + \mu_2 \, , & t = [\tau T] + 1, \ldots, T \end{cases}$$

where y_t is the seasonally adjusted inflation and τ is the unknown, potential break fraction. In order to find τ, we follow Hsu (2005) and modify the exact local Whittle estimator. In a grid search over $\tau \in [0.15, 0.85]$, d is estimated while accounting for a mean shift at the same time. That break fraction for which the concentrated local Whittle function is minimized is the candidate. The candidate lies in the interval 1981/8 to 1982/7, depending on the bandwidth used in the local Whittle estimation. Next, we test H_0 : $\mu_1 = \mu_2$ using a test statistic proposed by Hidalgo and Robinson (1996):

$$HR = T^{d-0.5} \frac{\hat{\mu}_1(\tau) - \hat{\mu}_2(\tau)}{\sqrt{\Omega}}$$

where

$$\Omega = \frac{C}{d(2d+1)} \frac{\tau^{2d} + (1-\tau)^{2d} - 1}{\tau(1-\tau)}$$

$$C = G 2\Gamma(1-2d)\cos((0.5-d)\pi)$$

$$G = \frac{1}{m} \sum_{j=1}^{m} \lambda_j^{2d} I_y(\lambda_j)$$

and $I_y(\lambda_j)$ is the periodogram of inflation evaluated at the harmonic frequency, $\hat{\mu}_1(\tau)$ is the estimated mean before the break point, $\hat{\mu}_2(\tau)$ is the estimated mean after the break point and d is replaced by the order of integration estimated accounting for the given break point. We repeat the analysis for different values of the bandwidth: $m \in \{T^{0.60}, T^{0.65}, T^{0.70}, T^{0.75}\}$. The HR-statistic has a limiting normal distribution if the break fraction is given exogenously. This is not the case here as we found the break fraction using a grid search. To that end, Ω is likely to be underestimated and p-values computed under the normal distribution would be too small. However, if we do not account for a possible mean shift, our long memory parameter might be spuriously high (Hsu, 2005). Therefore, we are willing to accept the significance of a break point under normality and acknowledge that we might allow for insignificant mean shifts. We observe that the break in mean

Table 2. Estimation of the order of Integration

m	$T^{0.60} = 42$	$T^{0.65} = 57$	$T^{0.70} = 78$	$T^{0.75} = 107$
$\hat{d}$	0.44	0.34	**0.35**	0.29
	[0.31, 0.57]	[0.23, 0.45]	**[0.25, 0.44]**	[0.21, 0.37]
date	1981/8	1981/10	**1981/10**	1982/7
HR	2.07	2.65	**2.61**	3.07

Notes: $\hat{d}$ is the exact local Whittle estimate of d of the seasonally adjusted inflation which is demeaned accounting for one mean shift. The bandwidth is denoted by m. The 90% confidence interval of the estimation of d is given in parenthesis. Furthermore, the date of the mean shift and the HR-statistic is given. The statistics corresponding to the most appropriate choice of m are highlighted.

is satisfactorily significant (p-value: 0.00-0.04) and account for it when estimating the order of integration of inflation. The HR-statistic as well as the dating of the mean shift are given in the Table 2.

Furthermore, Table 2 reports the estimation results of the order of integration using different values for the bandwidth. As is well known, the appropriate choice of m is crucial in the estimation of d. If m is chosen too small, then the estimate has a great standard deviation and might be imprecise. On the other hand, choosing m too large results into a downward bias. Our estimate of d seems to stabilize for $m = T^{0.65}$ and $m = T^{0.70}$ while the estimate for $m = T^{0.75}$ already exhibits a small downward bias. In order to maximize the number of observations not leading to a downward bias, we continue to work with $m = T^{0.70}$. The order of integration of inflation in the whole sample period therefore is 0.35 implying that inflation is stationary.

We investigate whether there is a second break in mean. To that end, we allow for two mean shifts in Hsu's modified local Whittle estimation. We compute the HR-statistic for the second break given that the first break is significant. To that end, we subtract the first mean-shift from the whole series and search for the second mean-shift (dating: 1973/1-1973/2). We then compute the HR-statistic for the second mean-shift as if it was the single mean shift. The p-values of the test statistics varies between 0.27-0.37, depending on m.[7] The second break is not significant even at the 10%

[7] Alternatively to the sequential procedure we also allow for two mean shifts simulta-

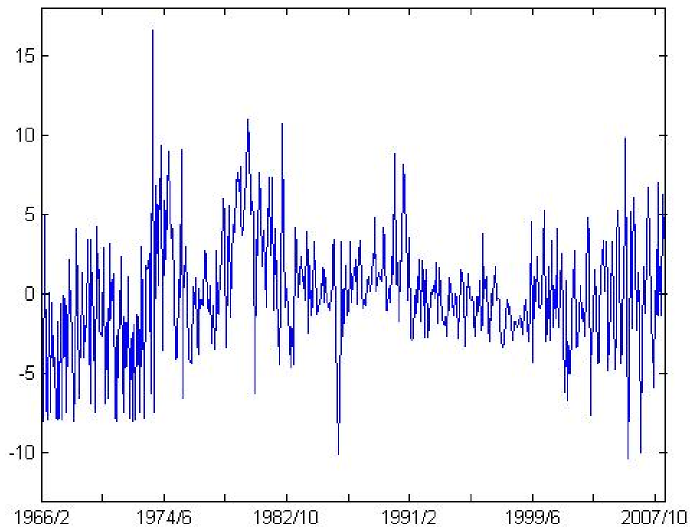

Figure 6. Adjusted Monthly US Inflation

significance level. Therefore, we only account for one shift in mean. In Figure 6 inflation accounting for seasonality and the break in mean is plotted against time.

4.2 Testing against a change in inflation persistence

We now turn to the estimation of the change in persistence. As a first step, we apply the difference filter to the mean adjusted inflation rates (y_t):

$$x_t = (1 - L)^{\hat{d}} y_t,$$

where $\hat{d}$ is the estimated order of integration of the whole sample reported in Table 2. Next, we estimate regression (9) using $p = \left[4(509/100)^{1/4}\right] = 6$ lags, for a given break fraction $\lambda \in (0.15, 0.85)$. We then compute the t-statistic for the null hypothesis $\psi = 0$. We repeat this procedure for all possible break points to find that value of λ that corresponds to the supremum of the squared t-statistics.

neously and obtain similar break points and p-values.

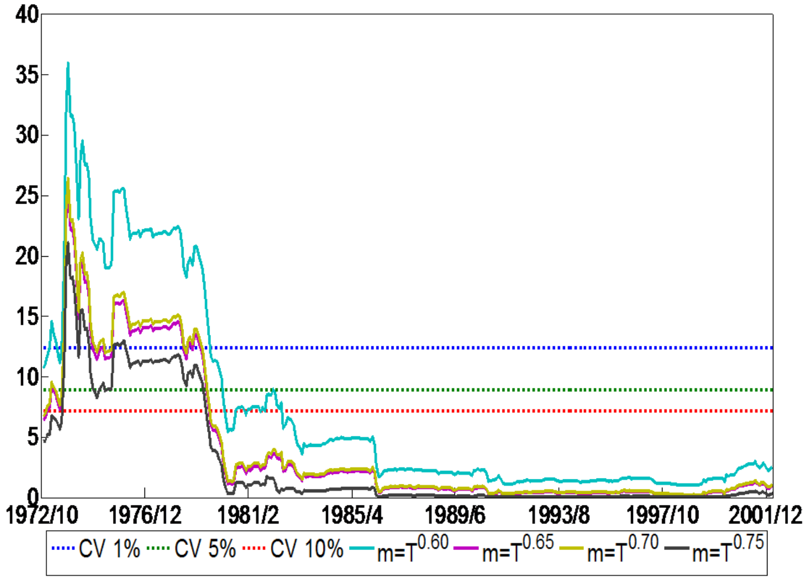

Figure 7. Evolution of t^2-statistic and CV for a search in the interval 15%-85% of the observations

The search for the break point is conducted on the middle 70% of the observations, i.e. on the time period 1972/10- 2002/3.[8] Figure 7 shows the evolution of the t^2-statistics of our test, for different definitions of m, as well as the critical values tabulated by Andrews (for $\pi = 0.15$). The t^2-statistics have their maximum value in October 1973.[9] The change in persistence is significant at a 1% significance level using Andrews' critical values.

We investigate whether there is a second change in persistence. To this end, we leave out the observations previous to the first change in persistence and repeat our analysis on the time period starting in 1973/10.[10] Again, we leave out the first and last 15% observations and search for a change in persistence in the interval 1979/4-2003/4. As can be seen in Figure 8, there is a second break point in June 1980. The second break point is not as significant as the first one. It is significant only at the 10% significance

[8] Note that the sample size is 509-6=503 because of the 6 lagged values in regression (9).

[9] Note that there is a local maximum in 1982/1.

[10] A search for a change in persistence using the sample before the first change is not suitable because the time series is so short.

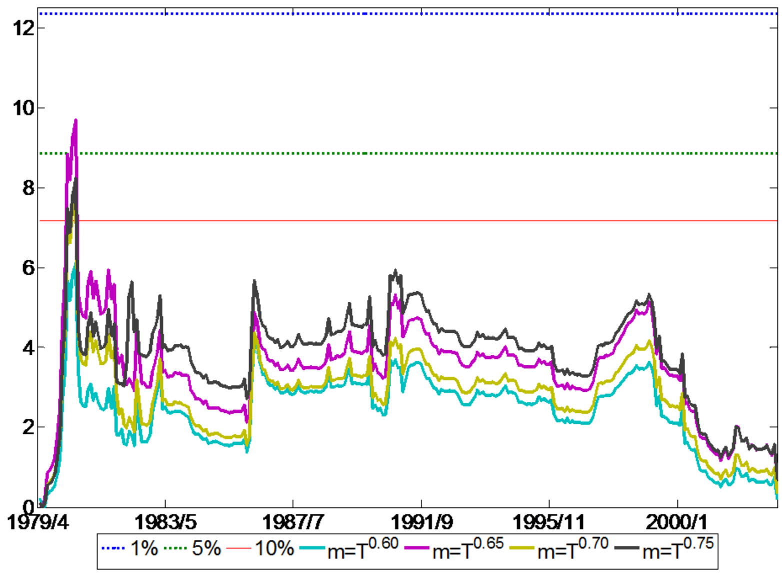

Figure 8. Evolution of t^2-statistic and CV in search for second break point

level for our choice of $m = T^{0.70}$. The supremum of the squared t-statistics is 7.82 and the 1% (5%, 10%) critical value is equal to 12.35 (8.85, 7.17).[11]. We conclude that there is a change in persistence in October 1973 and that there is somewhat weaker evidence for a second change in persistence taking place in June 1980.

One virtue of our test is that it can find a change in d even with few observations before or after the break point because it relies on OLS estimation. In contrast to that, a test on the difference between the estimate of d before and after some exogenously given break point has to have enough data points to measure the long memory parameter in the frequency domain in both periods. On this account, we were able to find the first break in persistence which already took place in 1973. However, the short time period before the break point does not allow us to reliably estimate the order of integration before the break. In order to provide an idea of whether inflation was reduced or increased after the change in persistence in 1973, Table

[11]The second change is significant at the 5% significance level if we choose $m = T^{0.65}$, significant at the 10% significance level if we choose $m = T^{0.75}$ and not significant at the 10% significance level if we choose $m = T^{0.60}$

Table 3. Estimates of d in different periods

	1966:2-2008:6	1973:11-2008:6	1980:7-2008:6	1966:2-1980:6
$\hat{d}$	0.35	0.27	0.09	0.54
	[0.25, 0.44]	[0.17, 0.37]	[-0.02, 0.20]	[0.41, 0.68]
m	78	68	59	37

Notes: The table shows exact local Whittle estimates of d for different time periods. Estimations are conducted over the whole time period, over the period after the first and after the second break and over the period until the second break. 95%-confidence intervals are given in brackets. Bandwidth $m = T^{0.70}$ is used in the estimations.

3 reports the estimated order of integration of the whole time period and of the period 1973-2008. Inflation persistence after the first break point is 0.27 and is therefore lower than inflation persistence in the whole period, 0.35. This decrease is not overly large because we are not comparing inflation persistence after the break point to inflation persistence before the break point. However, we expect that inflation persistence before the break point is higher than 0.35 and therefore the decrease was more serious. It is easier to demonstrate the change in persistence after the second break because there are enough observations to compute the order of integration in both periods. Table 3 reports the order of integration before the second break point, 0.54, as well as the order of integration after the second break point, 0.09.[12] We conclude that inflation persistence decreased after each break in persistence.

4.3 Reconciling our results with the existing literature

Several studies collect evidence that inflation persistence changed over the last four decades while others argue that inflation persistence was stable.

Pivetta and Reis (2007) conclude that inflation persistence is stable since 1965. However, they describe two periods in which inflation persistence could be suspected to be unstable. In the period 1982-1983, close to our second break, there is evidence for a "possible short-lived change" (p.29). However,

[12]The order of integration before the second break point is computed in the interval 1966-1980 including the first break point. The order of integration after the first and before the second break is probably smaller, however there are too few observations for such an estimation.

they argue that this was an exceptional period known as the Volcker disinflation and should not be confused with a structural change in the economy. When using recursive instead of rolling window estimation, the evidence vanishes. Therefore, they conclude that inflation persistence is indeed stable and the evidence for a change is only due to the short sample period. Another period that receives their attention is the period 1974-75. The confidence intervals of the largest autoregressive root as well as the sum of autoregressive coefficients tighten around this period. However, they dismiss this change point because the two persistence estimates point to a change of opposite directions. In addition, Pivetta and Reis (2007) test the hypothesis of stability of the parameters of the autoregression using the sup-Wald Likelihood Ratio test. The test points to a break in persistence in the third quarter of 1974, which is close to the dating of our first break point. However, Pivetta and Reis (2007)'s break point is not significant.

Most studies find that inflation persistence is time-variant. Our results are easiest comparable with those of Kumar and Okimoto (2007). Kumar and Okimoto (2007) argue that the test of Pivetta and Reis (2007) relying on the largest autoregressive root is insensitive to a change in the long memory of inflation and therefore the test does not detect such a break in persistence. Kumar and Okimoto (2007) test whether there is a change in the long memory of inflation using two strategies. Firstly, they ad hoc divide their sample into two periods, 1960.5-1982.4 and 1982.5-2003.4. They find a significant difference of the long memory parameter of those two subsamples. Their exogenously chosen break point is close to our second break point, 1980/6. However, our analysis shows that a change in the order of integration already takes place in 1973/10. Secondly, they plot 15-year rolling windows of long memory parameters and visually detect a break in the early 1980s. While sufficient for autoregressive parameters, we feel 15 years are insufficient to reliably estimate the order of integration.

Using a different measure of inflation persistence, other studies also find that inflation persistence in the US has decreased in the last years. Mishkin (2007) finds a decline in inflation persistence in the late 1990s using a rolling window of the coefficient sum of an AR(5) model. However, he points out that the decline of the sum of the lagged coefficients on inflation might be insignificant if modest changes in the methodology are made such as the extension of the sample period, the correction for small-sample bias or the use of a different price index.

Cogley and Sargent (2005) conclude that inflation persistence increased during the 1960s and 1970s and decreased afterwards. However, there remains uncertainty with respect to the timing and the number of peaks in their persistence measure.[13] Some evidence suggests that there was one peak in 1980 while other evidence suggests that there were two peaks of inflation persistence, either in the late 1960s and in 1980 or when the oil shocks occurred, 1973 and 1979.

Stock and Watson (2007) describe inflation persistence by the standard deviation of the permanent innovation. Their measure indicates high persistence in the 1970s to 1983 and low before and afterwards. However, one could also interpret the graphs of their persistence measure as upward trending until 1975 and downward trending afterwards. In any case, one can certainly conclude that the measure is not stable over time.

In the discussion on whether inflation persistence is stable, those studies dominate that find time-variant persistence. There is however still uncertainty about the timing of the change in persistence. In the studies mentioned, the consensus seems to be that there was a change in persistence in the early 1980s. This also fits with the Great Moderation and the start of the Volcker disinflation.

In line with most studies, we find evidence for time-varying inflation persistence. We contribute to the existing literature by testing not only for the significance but also the dating of a shift in inflation persistence accounting for fractional integration. We are confident that our test identifies reliably the timing of significant breaks in inflation persistence. Inflation persistence decreases in 1973 and 1980.

5 Concluding remarks

We proposed a new test against a break in the order of integration. The test follows Andrews (1993) in the context of a lag-augmented LM-regression. We account for long memory by differencing the time series by the appropriate order before using our test procedure. Monte Carlo simulations indicate that the power of our test depends especially on the difference of the order of integration before and after the break. If this difference is larger than 0.3, the power is very good. Furthermore, the power is influenced by the sample size while it is almost unaffected by changes in the true break fraction

[13] Cogley and Sargent (2005) measure inflation persistence using the normalized spectrum of inflation at zero frequency.

or the value of the AR(1) coefficient. The size is very good for all those specifications. However, if the long memory parameter of the whole sample is misspecified, our test is not conducted on an I(0) process. In case that the misspecification of d exceeds 0.3 the size of the test is too high and we reject the hypothesis that there is no break too often. We conclude that our test works reasonably well in our different settings if the overall long memory is captured appropriately a priori.

Using our new test, we investigate whether inflation persistence, i.e. the order of integration of inflation, in the US has changed. We find that there is a decrease in persistence in October 1973 and that there is somewhat weaker evidence for a second decrease in persistence taking place in June 1980. Those break dates are in line with changes in the monetary policy stance and macroeconomic dynamics and provide new evidence to the rather elusive literature on the stability of inflation persistence. We contribute to the existing literature by testing not only for the significance but also the dating of a shift in inflation persistence accounting for fractional integration. We leave it to future research to shed light on the stability of inflation persistence in other countries.

References

Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica 61*, 821–56.

Baillie, R. T., C.-F. Chung, and M. A. Tieslau (1996). Analysing inflation by the fractionally integrated ARFIMA-GARCH model. *Journal of Applied Econometrics 11*, 23–40.

Beran, J. and N. Terrin (1996). Testing for a change of the long-memory parameter. *Biometrika 83*, 233–233.

Beran, J. and N. Terrin (1999). Correction: Testing for a change of the long-memory parameter. *Biometrika 86*, 627–638.

Blinder, A. S., M. Ehrmann, M. Fratzscher, J. De Haan, and D.-J. Jansen (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature 46*, 910–945.

Brockwell, P. J. and R. A. Davis (1991). *Time Series: Theory and Methods*. Springer.

Busetti, F. and A. M. R. Taylor (2004). Tests of stationarity against a change in persistence. *Journal of Econometrics 123*, 33–66.

Cogley, T. and T. J. Sargent (2005). Drift and volatilities: Monetary policies and outcomes in the post WWII U.S. *Review of Economic Dynamics 8*, 262–302.

Demetrescu, M., V. Kuzin, and U. Hassler (2008). Long memory testing in the time domain. *Econometric Theory 24*, 176–215.

Fox, R. and M. Taqqu (1986). Large-sample properties of parameter estimates for strongly dependent stationary gaussian time series. *The Annals of Statistics 14*, 517–532.

Gadea, M. D. and L. Mayoral (2006). The persistence of inflation in oecd countries: A fractionally integrated approach. *International Journal of Central Banking 2*, 51–104.

Hassler, U. and M. Olivares (2008). Long memory and structural change: New evidence from German stock market returns. Unpublished.

Hassler, U. and J. Wolters (1995). Long memory in inflation rates: International evidence. *Journal of Business & Economic Statistics 13*, 37–45.

Hauser, M. A., B. M. Pötscher, and E. Reschenhofer (1999). Measuring persistence in aggregate output: ARMA models, fractionally integrated ARMA models and nonparametric procedures. *Empirical Economics 24*, 243–269.

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

Hsu, C.-C. (2005). Long memory or structural changes: An empirical examination on inflation rates. *Economics Letters 88*, 289–294.

Kim, J.-Y. (2000). Detection of change in persistence of a linear time series. *Journal of Econometrics 95*, 97–116.

Kumar, M. S. and T. Okimoto (2007). Dynamics of persistence in international inflation rates. *Journal of Money, Credit and Banking 39*, 1457–1479.

Leeb, H. and B. Pötscher (2005). Model selection and inference: Facts and fiction. *Econometric Theory 21*, 21–59.

Mishkin, F. S. (2007). Inflation dynamics. *International Finance 10*, 317–334.

O'Reilly, G. and K. Whelan (2005). Has euro-area inflation persistence changed over time? *Review of Economics and Statistics 87*, 709–720.

Phillips, P. and Z. Xiao (1999). A primer on unit root testing. *Journal of Economic Surveys 12*, 423–470.

Pivetta, F. and R. Reis (2007). The persistence of inflation in the United States. *Journal of Economic Dynamics and Control 31*, 1326–1358.

Robinson, P. (1994). Efficient tests of nonstationary hypotheses. *Journal of the American Statistical Associations 89*, 1420–1437.

Robinson, P. (1995). Gaussian semiparametric estimation of long range dependence. *Annals of Statistics 23*, 1630–1661.

Shimotsu, K. (2006). Simple (but effective) tests of long memory versus structural breaks. *Working Paper, Queen's University*.

Shimotsu, K. and P. C. Phillips (2006). Local Whittle estimation of fractional integration and some of its variants. *Journal of Econometrics 130*, 209–233.

Shimotsu, K. and P. C. B. Phillips (2005). Exact local Whittle estimation of fractional integration. *Annals of Statistics 33*, 18901933.

Stock, J. H. and M. W. Watson (2007). Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking 39*, 3–33.

Tanaka, K. (1999). The nonstationary fractional unit root. *Econometric Theory 15*, 549–583.