



TI 2007-099/4

Tinbergen Institute Discussion Paper

Long Memory Modelling of Inflation with Stochastic Variance and Structural Breaks

*C.S. Bos**

*S.J. Koopman**

M. Ooms

VU University Amsterdam

* Tinbergen Institute

Tinbergen Institute

The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam

Roetersstraat 31

1018 WB Amsterdam

The Netherlands

Tel.: +31(0)20 551 3500

Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam

Burg. Oudlaan 50

3062 PA Rotterdam

The Netherlands

Tel.: +31(0)10 408 8900

Fax: +31(0)10 408 9031

Most TI discussion papers can be downloaded at
<http://www.tinbergen.nl>.

Long memory modelling of inflation with stochastic variance and structural breaks

Charles S. Bos, Siem Jan Koopman and Marius Ooms
Tinbergen Institute Amsterdam &

Department of Econometrics, VU University Amsterdam,
De Boelelaan 1105, 1081 HV Amsterdam, NL

`cbos@feweb.vu.nl`, `s.j.koopman@feweb.vu.nl`, `mooms@feweb.vu.nl`

December 12, 2007

Abstract

We investigate changes in the time series characteristics of postwar U.S. inflation. In a model-based analysis the conditional mean of inflation is specified by a long memory autoregressive fractionally integrated moving average process and the conditional variance is modelled by a stochastic volatility process. We develop a Monte Carlo maximum likelihood method to obtain efficient estimates of the parameters using a monthly dataset of core inflation for which we consider different subsamples of varying size. Based on the new modelling framework and the associated estimation technique, we find remarkable changes in the variance, in the order of integration, in the short memory characteristics and in the volatility of volatility.

Keywords: Time varying parameters, Importance sampling; Monte Carlo simulation; Stochastic Volatility; Fractional Integration

JEL classification: C15, C32, C51, E23, E31

1 Introduction

Many monetary authorities, financial institutions, pension funds and private investors demand realistic statistical models for inflation to assess the real value of wealth, income and returns, both in the short run and in the long run. In this paper we develop a new exact maximum likelihood method to investigate the statistical properties of an inflation process with long memory and stochastic volatility.

It is well established that the statistical properties of postwar U.S. inflation underwent a number of structural breaks. There are many ways to explain persistent changes in the mean, variance and autocorrelation of U.S. inflation. The changes in the time series properties of inflation may be due to changing monetary policies (short run direct price controls in the 1950s and 1970s, long run indirect inflation controls starting in the 1980s) or to changes in the process generating price shocks. Many types of shocks have been investigated, we mention technological progress, unemployment changes, output gap disturbances, fluctuations in real unit labour costs as in Galí and Gertler (1999), oil price shocks, as in Hooker (2002), changes

in the sectoral distribution of price changes as in Ball and Mankiw (1995), trade unions as in Bowdler and Nunziata (2007), and exchange rate pass-through as in Campa and Goldberg (2005). Cecchetti, Hooper, Kasman, Schoenholtz, and Watson (2007) provide an extensive discussion of the empirical evidence of changes in the time series properties of inflation. It has been recognised that empirical models should allow for such changes in mean, volatility and persistence of inflation. In particular, the Great Inflation period of the 1970s and early 1980s shows the highest mean, the highest volatility and the highest persistence in mean and volatility. The statistical significance and the economic importance of the changes in mean and persistence are debatable but the reduction in volatility since the mid 1980s is relatively undisputed.

1.1 Empirical evidence of changes in the dynamic properties of inflation

Evans and Wachtel (1993) first investigated the fundamental changes in the persistence and conditional volatility of U.S. inflation using the Markov-Switching model of Hamilton (1990) and many studies had followed suit. Kim (1993) modelled changing volatility in inflation in a Bayesian Markov-Switching model. Alogoskoufis and Smith (1991) and Emery (1994) used parameter stability tests to show the increase in inflation persistence in the 1970s and the following decrease in the 1980s, respectively. The persistence of price inflation and wage inflation that appeared in the 1970s have been widely recognised as important stylised facts to be captured in macroeconomic models since Fuhrer and Moore (1995) published their influential article “Inflation Persistence”. There is a continuing debate about the statistical significance of inflation persistence and its stability over time. For example, Pivetta and Reis (2007) challenged the findings of changing persistence by Cogley and Sargent (2002). More recently Cogley and Sargent (2007) suggested that the change in persistence is more clearly established for the inflation gap, defined as inflation in deviation from a stochastic trend for the unconditional mean. It seems that one can only show significant changes in persistence when assuming a specific model structure. An exception is Harvey, Leybourne, and Taylor (2006) who show a significant change in the persistence in monthly U.S. inflation from $I(1)$ to $I(0)$ using a nonparametric variance ratio test and assuming a stationary variance process for the innovations. By considering a general classical frequency domain framework, encompassing simple versions of fractionally integrated processes as in Robinson (2003), local to unit root models as in Stock (1994) and unobserved component models as in Harvey (1989), Müller and Watson (2006) did not find conclusive evidence for significant changes in low-frequency behaviour of postwar U.S. inflation. The structural volatility changes in inflation have been accompanied by similar changes in a wide range of U.S. macroeconomic time series, see Kim and Nelson (1999), McConnell and Perez-Quiros (2000), Sensier and van Dijk (2004) and Kim, Nelson, and Piger (2004).

Structural explanations of the decrease in volatility and the reduction in persistence are still actively debated. For example, Benati (2006) reviews changes in inflation persistence in a wider historical and international perspective and concludes that inflation persistence might well be associated with monetary regimes, thereby questioning the structural character of inflation persistence. For example, the decrease in U.S. inflation volatility is associated with a change in monetary policy. Warne and Vredin (2006) found support for this explanation in a bivariate Structural Vector Autoregressive (SVAR) model for inflation and unemployment, where volatility breaks are captured by endogenous two-regime Markov-Switching. Primiceri (2005) captured the changing variance in a Stochastic Volatility (sv) specification of a trivari-

ate SVAR, also including an interest rate. Cogley and Sargent (2007) analysed a similar Vector Autoregressive (VAR) model, estimating stochastic volatility in inflation-gap persistence in the U.S. in addition to time variation in the VAR parameters. Sims and Zha (2006) analysed a SVAR with six variables switching between nine regimes. All these authors used seasonally adjusted data and found significant changes in U.S. inflation volatility.

Stock and Watson (2007) and Nason (2006) studied the time variation in inflation persistence using standard Autoregressive Integrated Moving Average (ARIMA) and Unobserved Components (UC) time series models. Nason (2006) used rolling exact maximum likelihood estimates for U.S. monthly inflation. Stock and Watson (2007) used quarterly U.S. data in fixed subsamples and introduced SV processes in both the nonstationary level component and the stationary component in inflation. They adopted Bayesian estimation methods using informative priors. Cecchetti, Hooper, Kasman, Schoenholtz, and Watson (2007) extended the empirical analysis of Stock and Watson (2007) to the G7 countries.

1.2 Evidence of long memory in inflation

Long memory modelling has been studied in econometrics and finance ever since Mandelbrot introduced long memory specifications for price processes, see Mandelbrot (1969) for early references and Baillie (1996) and Robinson (2003) for econometric literature reviews. In this paper we add a new model in this long tradition. Long memory models only started to become widely used in the 1980s when Geweke and Porter-Hudak (1983) developed the log periodogram regression estimator for the order of integration parameter d in the ARFIMA model of Granger and Joyeux (1980) and Hosking (1981). They applied their estimator to postwar U.S. CPI inflation data and found that inflation was integrated of an order d , $I(d)$, with d around 0.5, clearly different from the values of zero and one which were assumed in the earlier literature. This was an interesting alternative characterisation of the long-run behaviour of inflation with important economic implications. Moreover, the parameterisation provided simple tests for non-stationary values of $d > 0.5$ against stationary values $d < 0.5$ and vice versa. Moreover, the order of integration d turned out to be a crucial parameter for long-run forecasting, as it determines the rate of increase of dynamic prediction intervals at long horizons for values of $d > 0.5$, see Beran (1994, §8.7). This is important for inflation, but even more so for the consumer price level, which is often used in long-term financial contracts. Doornik and Ooms (2004) illustrated these effects

Hassler and Wolters (1995) showed that a range of international postwar inflation series could effectively be modelled with a fractional order of integration. Baillie, Chung, and Tieslau (1996) further extended the evidence on long memory in international inflation rates by estimating fractional orders of integration using approximate maximum likelihood in the time domain. Their estimates took account of autoregressive conditional heteroskedasticity, initially introduced by Engle (1982) to model heteroskedasticity of U.K. inflation. Ling and Li (1997) provided additional asymptotic theory for these ARFIMA-GARCH models. Beran and Feng (2001) and Feng, Beran, and Yu (2007) developed estimation and inference methods for the ARFIMA-GARCH model in combination with semi-parametric estimates of deterministic flexible smooth trends.

A range of inflation series for different countries has provided strong evidence of high persistence and changing volatility over time. High estimates of inflation persistence in ARMA or in ARFIMA models could partly be explained by structural shifts in the unconditional mean. These shifts may lead to overestimating the persistence and to wrong conclusions regarding

the order of integration of the stochastic part of the model. This point was strongly made by Perron (1989), who analysed the effect of structural changes in the mean on tests for the unit root hypothesis developed by Dickey and Fuller (1979). For example, allowing for seven breaks in the model for the mean of the U.S. CPI index he concluded: “After 1929 the unit root is no longer present”, cf. Perron (1989, p. 1385). In the context of ARFIMA models for international inflation rates Bos, Franses, and Ooms (1999) showed how structural changes in the mean affect estimates of the fractional integration parameter d .

A related problem for the interpretation of persistence estimates in inflation occurs when volatility is (nearly) nonstationary. Then the specification of the volatility component does influence the parameter estimates in the conditional mean part of the model. Boswijk and Klaassen (2004) discussed theoretical results and practical implications of this interaction of volatility and mean specification for tests of the $I(1)$ -hypothesis in an ARIMA-GARCH context. Based on an unobserved components model with SV, Cavaliere and Taylor (2005) analysed the effect of nonstationary volatility specifications on the limiting distribution of tests of the $I(0)$ -hypothesis. When evaluating time variation in persistence measures in inflation it is therefore crucial to investigate stability of the mean and volatility at the same time. Reliable tests for the stability of long-memory persistence in the presence of a nearly nonstationary mean and (nearly) nonstationary volatility have not been developed yet.

On the other hand, high persistence in inflation can lead to overrejections in tests of the null hypothesis of stability of the unconditional mean, even if the underlying data generating process satisfies the assumptions of the test. For example, Müller (2005) clearly showed that partial sum tests like the KPSS-test of Kwiatkowski, Phillips, Schmidt, and Shin (1992) do not work in the presence of high persistence, at least not for practical sample sizes. Nyblom and Harvey (2001) clearly indicated the limitations of similar stability tests in a parametric unobserved components context.

If the mean is modelled as a stochastic processes with occasional extremely long-lasting changes in the mean it can even be considered as a substitute for long-memory, removing the scope for an additional ARFIMA process to model deviations from this mean, see Engle and Smith (1999), Parke (1999) and Diebold and Inoue (2001). At a more basic level it is well known that one can approximate a fractionally integrated process by an ARIMA process with an MA root close to the unit circle as in Brodsky and Hurvich (1999).

In sum, there is evidence of long memory in inflation rates, but in the evaluation of this evidence one should consider the possibility of persistent changes in the mean and volatility.

1.3 Contribution of this paper

In this paper we take a new step by modelling heteroskedasticity in an ARFIMA model by means of a Stochastic Volatility (SV) model. SV models have been reviewed by Taylor (1994) and in the edited volume of Shephard (2005). Shephard (1996) extensively discussed the differences and similarities of ARCH and SV models. Taking into account many of the criticisms of existing research we provide a different view on the changing time series characteristics of inflation. We examine changes in the long-run persistence in the mean by monitoring changes in the estimate of the long-memory parameter d of the ARFIMA model. Furthermore, we investigate fluctuations in the volatility pattern, in the persistence of volatility (via the autoregressive parameter and the variance parameter of the SV component) and in short memory characteristics (via the ARMA parameters present in the ARFIMA model).

In contrast to Primiceri (2005) and Stock and Watson (2007) we use monthly time series

data and we adopt maximum likelihood methods instead of Bayesian methods. We extend the analysis of Stock and Watson (2007) by using an ARFIMA specification which allows us to circumvent the a priori choice for the order of integration. However, as we do not use an unobserved components model we can only distinguish a single source of stochastic volatility.

We adapt the Monte Carlo maximum likelihood estimation procedure of Koopman and Bos (2004) to allow its treatment within an ARFIMA model. The details are given in this paper. The evaluation of the loglikelihood function is comparatively fast, so that we are able to obtain the rolling and recursive estimates necessary for the analysis of time variation of the parameters. We label our model ARFIMA-SV which should not be confused with the long memory SV models of Breidt, Crato, and De Lima (1998) and Brockwell (2007). They model long memory in the SV component instead of long memory in the conditional mean. As we allow for long-memory and stochastic volatility we cannot use techniques as in Bai and Perron (2003) to formally establish the statistical significance of changes in parameters between subsamples. Anyhow, even in a model with constant mean Beran and Terrin (1996) and Ling (2007) showed that in order to reliably detect structural changes in the long memory parameter d , one requires considerable subsample sizes. Development of similar tests for our model are outside the scope of this paper.

The remainder of this paper is organised as follows. Section 2 introduces the ARFIMA-SV model and the maximum likelihood estimation of the fixed parameters and the volatility component. Section 3 presents empirical results for U.S. inflation. We present the estimation results for the ARFIMA-SV model and we analyse parameter (in)stability. We also use our framework to compare the results for the ARFIMA-SV model with specifications from the existing literature. Section 4 concludes.

2 Long memory model with stochastic volatility

Our model combines a long memory model as specified in ARFIMA-form by Granger and Joyeux (1980) and Sowell (1992) with a stochastic volatility model as presented by Taylor (1982). Robinson (2003) reviewed long memory models. Our method is not confined to ARFIMA models, we can consider all general linear time series models with a stationary autocorrelation function. In terms of asymptotic representations the data generating process of our approach relates to the stationary "fractional Brownian motion of Type I" in the terminology of Marinucci and Robinson (1999), in contrast to nonstationary Type II long memory specifications, which were considered inferior by Mandelbrot and Van Ness (1968). Davidson and Hashimzade (2008) discuss important differences between Type I and Type II fractional Brownian motions in the development of asymptotic theory.

We choose the ARFIMA specification as it delivers the well known ARIMA models as special cases. Moreover, the ARIMA(0,1,1) model encompasses the random walk plus noise model, or local level model, which is very effective for forecasting inflation, see Stock and Watson (1999) and Stock and Watson (2007). Sowell (1992) discussed the exact maximum likelihood estimation of ARFIMA parameters. Doornik and Ooms (2003) refined Sowell's method and introduced a computationally efficient implementation. Exact maximum likelihood avoids ad hoc decisions on the likelihood for the initial observations of the process. Because of the relatively strong correlation of distant observations this is more important for long memory ARFIMA models than for short memory ARMA models. Note that one can take first differences of the process to ensure stationarity. This is relevant for inflation data in the 1970s.

We add the SV specification to the ARFIMA model. The SV specification avoids the ad hoc decisions for the computation of the likelihood of the initial observations that one needs to make in estimating GARCH type models as introduced by Bollerslev (1986).

2.1 The ARFIMA-SV Model

The ARFIMA(p, d, q) model for a time series y_t is based on a linear Gaussian long memory process for the conditional mean as given by

$$\Phi(L)(1-L)^d(y_t - \mu) = \Theta(L)\varepsilon_t, \quad \varepsilon_t \sim \mathcal{NID}(0, \sigma^2), \quad t = 1, \dots, T, \quad (1)$$

where $\Phi(L) = (1 - \phi_1 L - \dots - \phi_p L^p)$ is the autoregressive (AR) polynomial and $\Theta(L) = (1 + \theta_1 L + \dots + \theta_q L^q)$ is the moving average (MA) polynomial in the lag operator L , with $L^k y_t = y_{t-k}$ for $k = 0, 1, \dots$ and integer orders $p \geq 0$ and $q \geq 0$. The mean of the long memory process is μ . The disturbance ε_t is serially independent and normally distributed with zero mean and variance σ^2 . The fractional difference operator $(1-L)^d$ with real d is defined by

$$(1-L)^d = \sum_{j=0}^{\infty} \delta_j L^j = \sum_{j=0}^{\infty} \binom{d}{j} (-L)^j.$$

We assume $-1 < d < 0.5$. The autoregressive moving average (ARMA) part of the model is assumed invertible and stationary. We further assume that the equations $\Theta(z) = 0$ and $\Phi(z) = 0$ for unknown z do not have common roots. Granger and Joyeux (1980) define the process of $y_t - \mu$ to be integrated of order d , denoted $I(d)$.

The mean μ can be replaced by a time-varying deterministic component μ_t and specified as in a linear model to capture deterministic variation in the time series. More specifically, we may specify the mean function as $\mu_t = X_t' \beta$ where X_t is a $k \times 1$ vector of covariates (including the constant) and β is a $k \times 1$ vector of fixed but unknown coefficients. The resulting ARFIMA model is standard and its treatment is considered by many, including Hosking (1981), Sowell (1992) and Beran (1994).

To deal with possibly time-varying volatility in the time series, the ARFIMA model is extended by replacing σ^2 in (1) for σ_t^2 which is modelled by the SV model. For this purpose, we specify $\sigma_t^2 = \exp(h_t)$ where the log-variance h_t follows a stationary Gaussian AR(1) process with zero mean, that is

$$h_{t+1} = (1 - \rho)\gamma + \rho h_t + \xi_t, \quad \xi_t \sim \mathcal{NID}(0, \sigma_\xi^2), \quad t = 1, \dots, T, \quad (2)$$

with autoregressive coefficient $0 < \rho < 1$. The constant term γ in (2) is the mean of the volatility process. The disturbances ξ_t are serially independent and normally distributed with zero mean and variance σ_ξ^2 . The disturbances ε_t in (1) and ξ_t in (2) are mutually independent for all $t, s = 1, \dots, T$. The parameter ρ captures the persistence of volatility and σ_ξ^2 measures the volatility of volatility. A selection of readings on the SV model is given by Shephard (2005).

Fast and reliable methods for the exact maximum likelihood estimation of parameters in the ARFIMA model (1) with arbitrary d have been discussed by Doornik and Ooms (2003). Koopman and Bos (2004) have developed Monte Carlo maximum likelihood estimation methods for the ARIMA-SV model, that is model (1)-(2) with integer $d = 0, 1, 2, \dots$. These methods

are based on the non-Gaussian state space framework and treated in Durbin and Koopman (2001). They allow for simultaneous estimation of the ARMA and SV parameters. In the next subsection we discuss an efficient combination of the ARFIMA and SV methods of estimation to treat the ARFIMA-SV model (1)-(2) with arbitrary d .

2.2 Likelihood estimation

Consider ARFIMA-SV model (1)-(2) for the simple case with $\mu_t = 0$, $p = 1$ and $q = 1$ and with $\Phi(L) = 1 - \phi L$ and $\Theta(L) = 1 + \theta L$. The unknown parameters are d , ϕ , θ , γ , ρ and σ_ξ^2 and they are collected in the parameter vector ψ . Given a realised value for $h = (h_1, h_2, \dots, h_T)'$, the density of the data vector $y = (y_1, y_2, \dots, y_T)'$ follows from the prediction error decomposition of the Gaussian density and is given by

$$\log p(y|h; \psi) = -\frac{T}{2} \log 2\pi - \frac{1}{2} \sum_{t=1}^T \log f_t - \frac{1}{2} \sum_{t=1}^T v_t^2 / f_t, \quad (3)$$

where v_t is the one-step ahead prediction error and f_t is its variance, see Schweppe (1965) and Brockwell and Davis (1993). Koopman and Bos (2004) used the Kalman filter to compute v_t and f_t for the ARIMA-SV model. In this paper we need to treat an ARFIMA process, so this would require an equivalent process represented in state space form. However, Chan and Palma (1998) have proven that a state space representation of the ARFIMA process with a finite dimension of the state does not exist. Therefore we resort to the algorithm of Durbin (1960), modified by Doornik and Ooms (2003) to compute v_t and f_t .

Golub and Van Loan (1996, §4.7.2) discuss the Durbin algorithm as a special implementation of the Choleski decomposition for Toeplitz matrices and argue that the algorithm is computationally efficient. For our purpose, the covariance matrix of data vector y is of Toeplitz form and needs to be decomposed. The Durbin algorithm requires explicit expressions for the autocovariance function of the time series process y_t in terms of the model parameters. Sowell (1992) derived reliable algorithms to compute the autocovariance function. Doornik and Ooms (2003) further refined these. Let Ω denote the covariance matrix of y , then the algorithm computes the f_t for $t = 1, \dots, T$ and $v = P^{-1}y$, where P is the $T \times T$ Choleski factor matrix so that $\Omega = PFP'$ with $F = \text{diag}(f_1, \dots, f_T)$ and $|P| = 1$. The Choleski factor P is not stored as this is prohibitive for large T . Only the projection coefficients for the last observation y_T are in memory at the end of the computations. Note that the Kalman Filter computes a similar Choleski factorisation, see Durbin and Koopman (2001). Therefore the Durbin-algorithm and the Kalman filter deliver the same numerical results for v_t and f_t for stationary models that can be put in state space form. It is clear that in cases where $\sigma_t^2 = \sigma^2$, σ^2 can be concentrated out of the likelihood function for y_t as it is just a scaling factor.

2.3 Importance sampling

Consider the ARFIMA-SV model with $\sigma_t^2 = \exp(h_t)$ where h_t is modelled by (2). We can express the density of y_1, \dots, y_T as given by

$$p(y; \psi) = \int p(y, h; \psi) dh = \int p(y|h; \psi) p(h; \psi) dh = \int p(v|h; \psi) p(h; \psi) dh \quad (4)$$

It is noticed that the Jacobian of the transformation from y to $v = (v_1', \dots, v_n')'$ is unity. Therefore the density of $y|h$ is the same as the density of $v|h$ and $p(v|h; \psi)$ is given by (3).

The difficulty in evaluating the integral to obtain the density of the ARFIMA-SV model is well recognised. A possible Monte Carlo technique for the evaluation of this density is importance sampling and is applied to stochastic volatility models by Danielson (1994), Shephard and Pitt (1997), Sandmann and Koopman (1998) and Durham and Gallant (2002). This approach is taken below and it involves approximating the density via averages of simulations from an approximating model.

Given a realised value of y , the evaluation of the resulting likelihood function (4) can be based on importance sampling. For this purpose, the observation density is represented by

$$p(y; \psi) = \int \frac{p(v|h; \psi)p(h; \psi)}{g(h|v; \psi)} g(h|v; \psi) dh, \quad (5)$$

where $g(h|v; \psi)$ is the importance density. The importance sampler will be based on the prediction errors v^σ of the ARFIMA model (1) with σ_t^2 kept fixed at 1. To increase the efficiency of the importance sample, the importance density $g(\cdot)$ is suitably chosen so that it is close to the density $p(\cdot)$. The importance density is constructed using an approximating linear Gaussian state space model. It can be obtained as described in Appendix A.

It follows from Bayes' rule that $g(h|v; \psi) = g(v|h; \psi)g(h; \psi)/g(v; \psi)$. Substituting in the denominator of (5), and observing that $g(h; \psi) = p(h; \psi)$, one obtains

$$p(y; \psi) = g(v; \psi) \int \frac{p(v|h; \psi)}{g(v|h; \psi)} g(h|v; \psi) dh. \quad (6)$$

Based on M simulated series $h^{(i)}$ from the importance density $g(h|v^\sigma; \psi)$, we approximate the logarithm of (6) via importance sampling and obtain the Monte Carlo estimate of the loglikelihood

$$\log p(\widehat{y}; \psi) = \log g(v^\sigma; \psi) + \log \left[\frac{1}{M} \sum_{i=1}^M \frac{p(v^\sigma|h^{(i)}; \psi)}{g(v^\sigma|h^{(i)}; \psi)} \right]. \quad (7)$$

The estimator (7) is subject to simulation error, which can be made arbitrarily small by increasing the number of simulations M . Generating drawings from the importance density is referred to as simulation smoothing. Various simulation smoothing methods exist to compute conditional draws from the approximating model $g(h|v^\sigma; \psi)$. We have used the fast and simple simulation smoother of Durbin and Koopman (2002).

The Monte Carlo estimate of the loglikelihood function in (7) is maximised with respect to ψ . The evaluations of the Monte Carlo loglikelihood for different values of ψ are based on the same set of random numbers for the sampling of $h^{(i)}$, $i = 1, \dots, M$, to ensure a smooth loglikelihood function in ψ . In our main applications we use $M = 250$, but M can often be reduced without important effects on the empirical results. Standard error estimates of the parameters are obtained from the numerical Hessian of the loglikelihood. This procedure is numerically stable and is implemented for the programming language `0x` of Doornik (2006) using the ARFIMA functions as documented in Doornik and Ooms (2003) and adopting state space functions of `SsfPack` as in Koopman, Shephard, and Doornik (1999). The programs are available upon request.

2.4 Estimation of the volatility process

The volatility vector h can be estimated by the importance sampling methods as well. In a similar way as for the construction of the likelihood function, we develop the importance

sampling estimator

$$\hat{h}_t = E(h_t|y) = \bar{\omega}^{-1} \sum_{i=1}^M \omega_i \hat{h}_t^{(i)}, \quad \omega_i = \frac{p(v^\sigma|h^{(i)}; \psi)}{g(v^\sigma|h^{(i)}; \psi)}, \quad \bar{\omega} = \sum_{i=1}^M \omega_i, \quad (8)$$

where $\hat{h}_t^{(i)}$ is the smoothed estimate of h_t based on the stationary linear Gaussian state space model (9) and (2) using the Kalman filter and associated smoother, see Appendix A. In a similar way, the variance of the smoothed volatility estimate is given by

$$\text{var}(h_t|y) = \bar{\omega}^{-1} \sum_{i=1}^M \omega_i V_t^{(i)},$$

where $V_t^{(i)}$ is the smoothed variance estimate V_t associated with $\hat{h}_t^{(i)}$.

3 Empirical study of U.S. core inflation

3.1 Data

We apply our model to a long monthly time series of inflation. We use the U.S. City Average core consumer price index of the Bureau of Labor Statistics (BLS), series CUUR0000SA0L1E. This index, which excludes the direct effect of price changes for food and energy and is further denoted by P_t . Monthly U.S. core inflation, π_t , is calculated as $\pi_t = 100\{\log(P_t) - \log(P_{t-1})\}$ and can be interpreted as the percentage price change per month. Our series start in 1965:1. The preceding years of data lack sufficient precision as the CPI index is rounded at two digits and normalised at 100 for 1982-1984. This makes the earlier data unfit for the analysis of stochastic volatility.

Table 1 reports descriptive statistics of our series. Figure 1 shows time series plots of the price index, monthly inflation, and sample autocorrelation functions of inflation and of changes in inflation. The time-varying character of mean and volatility are clear from the time series plots. There are also some possible outliers in inflation, 1980:7 showing most clearly. The autocorrelation function shows strong persistence and a seasonal pattern, so it seems advisable to build a model that allows for long range dependence and seasonal correlations. However, we cannot model all aspects of the data in great detail, as the main aim of our paper is to examine time variation in the persistence of the mean and variance of the inflation process. Efficiency of the estimate is an issue. Even simple nonlinear long memory models need a substantial sample for a satisfactory empirical identification of the parameters.

We therefore take a simple approach regarding seasonal movements. We take the mean of the seasonal component in π_t as fixed and remove it by a priori OLS regression, resulting in our dependent variable y_t . We use y_t in the remainder of our paper. For the full sample, we also estimated fixed seasonal factors within the model, resulting in seasonal estimates that were very close to the OLS estimates. We include seasonal AR parameters in the model to capture the remaining seasonality. The next section presents our main empirical estimates and diagnostic tests for ARFIMA-SV models for y_t .

The discussion of the estimation results for the ARFIMA-SV model is organised as follows. First, Section 3.2 presents the estimates of the full sample, for models with and without an SV-component. More importantly, it also reports these results for the two subsamples. Next, Section 3.3 shows the sensitivity of the results with respect to the begin point and endpoint of the subsamples.







Figure 1: (a) Monthly time series of U.S. core price index (1982-84=100), (b) Monthly core inflation adjusted for fixed seasonals, (c) Sample autocorrelation function (SACF) of core inflation up to a lag 100 months (d) SACF of changes core in inflation

3.2 Full-sample, Great-Inflation, and Post-Great Moderation results

Table 2 presents detailed estimation results for models that are restricted to be stationary and for three samples. This table illustrates the effect of the SV extension on the ARFIMA estimates and shows the presence of a structural break in the parameters. Table 3 displays the corresponding ARFIMA-SV estimates for Δy_t to show the consequence of relaxing the stationarity assumption. Table 3 also facilitates comparison with ARIMA-SV results. In all these tables, we estimate a parsimonious subset ARFIMA model. The parameters $\phi_2, \phi_3, \dots, \phi_{10}$ are not significant in our applications and fixed at zero.

The first two columns of Table 2 show the full sample exact maximum likelihood estimates for the parameters of a subset ARFIMA model (1) for y_t . We impose $\phi_2 = \phi_3 = \dots = \phi_{10} = 0$ as these parameters are insignificant in our applications. In this case σ_ε denotes the standard deviation of ε . The order of integration, d , is estimated at 0.39 and significantly different from 0. This implies that U.S. core inflation is long memory. Note that we impose $-1 < d < 0.5$ but it is not clear that we can reject $d > 0.5$. The AR component of the ARFIMA model mainly captures seasonal correlations via the parameter ϕ_{12} . The mean inflation μ is estimated at 0.32% per month, or 3.9% per year. The residual standard error σ_ε is large at 0.18% per month. The mean appears not to be significant. The inflation rate in 1980:7 is a negative additive outlier and very significant, and remains significant in all other estimation results. The diagnostic tests reported for the prediction errors are the normality test and the Box-Ljung Q statistic to test the white noise assumption. Both tests clearly reject the simple ARFIMA specification.

Columns 3 and 4 of Table 2 present the full sample estimates of the complete ARFIMA-SV model (1)-(2) with $\sigma_\varepsilon = \exp(\gamma/2)$, to facilitate comparison with the ARFIMA model without SV in Columns 1 and 2. In this case σ_ε is a measure of the average volatility. The SV extension increases the likelihood by 48 points. The distribution of the Likelihood Ratio test statistic for SV is nonstandard as $\sigma_\xi^2 = 0$ is on the boundary of the parameter space and ρ is not identified under the null. Despite this caveat we interpret this likelihood increase as a significant improvement. To further illustrate the significance, we plot the smoothed estimate of $\exp(h_t/2)$ in Figure 2, see the solid line. The volatility decrease in the early 1980s is spectacular and persistent. Moreover, as indicated in § 1.1, it has been documented in a range of models, for a range of macroeconomic variables and for a range of developed countries. It is often labelled The Great Moderation. The estimate for d decreases by 0.09, or 1.5 standard error by the SV extension and ϕ_{12} increases by a similar amount. The estimate of μ is also clearly affected. The SV component itself is nearly nonstationary as the autoregressive coefficient of volatility, ρ , is very close to one and the volatility of volatility, σ_ξ , is well identified at 0.145 with a standard error of 0.05. When volatility is stationary, the model specification of the volatility does not have a big impact on the estimates of AR parameters and d , but this large and persistent change in volatility has a marked effect and has to be taken into account when evaluating persistence measures.

The subsample periods in Table 2 are chosen to investigate potential structural breaks in the parameter estimates of the ARFIMA-SV model between the Great Inflation and the Post Great Moderation era. For the interpretation of the subsample estimates it is important to note that the exact maximum likelihood estimates are truly unconditional. Presample observations of y_t are not used to condition on. This is in contrast with the conditional least squares estimators for the long memory model as in Baillie, Chung, and Tieslau (1996). Columns 5 to 8 of Table 2 show the parameters of the ARFIMA and ARFIMA-SV model for data

centred around the Great Inflation period. The associated stochastic volatility estimate is shown in the left hand side of Figure 2 as a dashed line. The d parameter is now close to the nonstationary value of 0.5, and the parameter ρ is somewhat smaller than for the full sample. Naturally, the estimate of mean inflation is high, but so is its standard error. The diagnostics look better than in the full sample. There is still a sizeable difference in loglikelihood values between the ARFIMA model and ARFIMA-SV model, but this difference is much smaller than for the full sample estimates and the differences for d and the AR parameters across the ARFIMA and ARFIMA-SV specification are much smaller than in the full sample.

Finally, columns 9 to 12 of Table 2 display the results for the period after the Great Moderation. The volatility estimate is shown in the dotted line on the right hand side of Figure 2. The d parameter is apparently significantly smaller than in the first period and here the likelihood contribution of the SV-component is comparatively small. The volatility of volatility is slightly lower than in the Great Inflation period, and its significance is much lower than in the full sample. The ARMA parameters also show a significant change in comparison with the first period. The diagnostics are nearly satisfactory, even without any outlier correction.

Table 3 presents estimates of the same ARFIMA-SV model for the same periods as Table 2, but now for changes in inflation, Δy_t . Consequently, we estimate $d - 1$ instead of d and we impose $0 < d < 1.5$. The mean of y_t is no longer identified, so μ drops out of the model. The explanatory dummy variable could also be differenced, but we choose a specification with separate coefficients for the dummy variable and its lag. In the Great Inflation period, $d - 1$ equals -0.29, corresponding to a value of d of 0.71 and it seems significantly larger than 0.5. For the full sample $d - 1$ equals -0.71, nearly equivalent to the d of 0.29 in Table 2. The implicit estimate for d in the last subsample is 0.12, compared to 0.14 in Table 2. Table 3 therefore shows a significant drop in d by a margin of around 0.6, providing clear evidence of a structural break in d , but also in σ_ε and in the AR parameters. The volatility of volatility is similar across subsamples and again insignificant in comparison with the full sample estimate.

3.3 Rolling and Recursive estimation results

The estimation results of Section §3.2 become easier to interpret after a sensitivity analysis with rolling and recursive estimates. Each subsample needs to be sufficiently long to make sense in a long memory analysis. Therefore we start our sensitivity analysis with rolling estimates for subsample lengths of 200 months.

A selection of the corresponding rolling window estimates with one-standard-error bands is presented in Figure 3. The values for 1965:1 correspond to the estimation period 1965:1-1981:8, which truncates the first sample of Table 2 from the right. The values for 1991:2 show estimates for 1991:2-2007:9 truncating the last sample from the left. We present the rolling subsample estimates for four key parameters. The fractional integration order d captures the long memory behaviour, ϕ_{12} picks up the short memory behaviour including seasonality. The parameter σ_ε shows the overall volatility in the disturbances, while μ measures the unconditional mean. Figure 3 reveals that the d estimates do not change much if one shifts the start of the sample from 1965 through to 1974. We computed the same rolling estimates of the ARFIMA-SV model for Δy_t as in Table 3. The results were similar, except that the values for d in samples with starting date upto 1974 fluctuated around 0.6 instead of around 0.45.

The estimates for μ show the effect of the (near-)nonstationarity of y_t in the early samples. The standard error estimates for μ are very large, indicating serious empirical identification




Figure 2: Volatility estimates $\exp(h_t/2)$ in ARFIMA-SV model (1)-(2). Solid line: estimates based on full sample 1965:1-2007:9. Dotted lines: estimates based on samples 1965:1-1982:12 and 1985:1-2007:9.

problems for the unconditional mean. These problems do not occur in the second part of the sample, where the mean is well identified, although the rolling estimates indicate a drop in the mean when the 21st century data start getting included.

The other parameters begin to change when the subsamples start to include both Great Inflation and Post Great Moderation years. The constant parameter model is then seriously misspecified. The long memory parameter decreases, the parameter ϕ_{12} increases. The estimates do not change for starting dates between 1983 and 1986. The results for the second period in Table 2 do not vary a lot if one shifts the period backwards. The estimates for d and ϕ_{12} tend back to their Great Inflation values when the 2004 and 2005 values are included in the sample, but 2006 and 2007 indicate a new decrease in persistence. It is too early to tell whether this reversion to around $d = 0.2$ is significant.

Figure 4 displays the results of our last parameter stability analysis of the ARFIMA-sv model. This analysis is based on estimates for two consecutive subsamples where we shift the break date over time. The solid lines show standard recursive estimates for the first subsample where new observations are subsequently added at the end. The dashed lines presents the estimates of the remaining subsamples. The distance between the two lines in Figure 4 shows the difference in the estimates for two non-overlapping but consecutive subsamples. These distances form the basis of parameter instability tests and tests for structural change with unknown change point as in Quandt (1960) and Andrews (1993). Ling (2007) recently derived asymptotic theory for normalised Quandt-type Maximum Wald statistics in the context of QML (conditional least squares) estimation of ARFIMA models. He extended and improved related frequency domain results of Beran and Terrin (1996). Na and Lee (2007) show that tests can also be developed based on the moving estimates of time series models as in Figure 3 but this is beyond the scope of this paper. As indicated in §1 we do not present formal tests for parameter stability. The rejection of the null hypothesis of constant parameters for the full sample has already been firmly established in the literature and the Post Great Moderation period is too short to detect a structural break in the long memory properties of inflation.

For forecasting purposes it is interesting to focus on post-break estimates. Estimates of d based on subsamples with starting dates between 1985 and 1993 fluctuate around 0.15, in accordance with the rolling estimates of Figure 3. It is too early to tell whether the long memory characteristic of inflation has remained significant in the 21st century, but based on this evidence one can safely state that the current long run persistence is lower than in the Great Inflation period. The recursive estimates of ϕ_{12} follow a slowly drifting trend. The most significant changes occur in the first decade after the Great Inflation. The recursive estimates of σ_ε confirm the Great Moderation in the inflation variance. The recursive estimates of μ in the pre-break periods clearly show the start of the Great Inflation period. The postbreak estimates of μ are remarkably stable, even going back to 1988. Although the U.S. have not had an explicit inflation targeting policy, the estimation results for μ confirm the success in achieving long run inflation stability together with both a decrease in short run inflation uncertainty (see the rolling estimates of σ_ε in panel (c) of Figure 3) and a decrease of volatility (see the postbreak stochastic volatility estimates $\exp(h_t/2)$ in Figure 2).

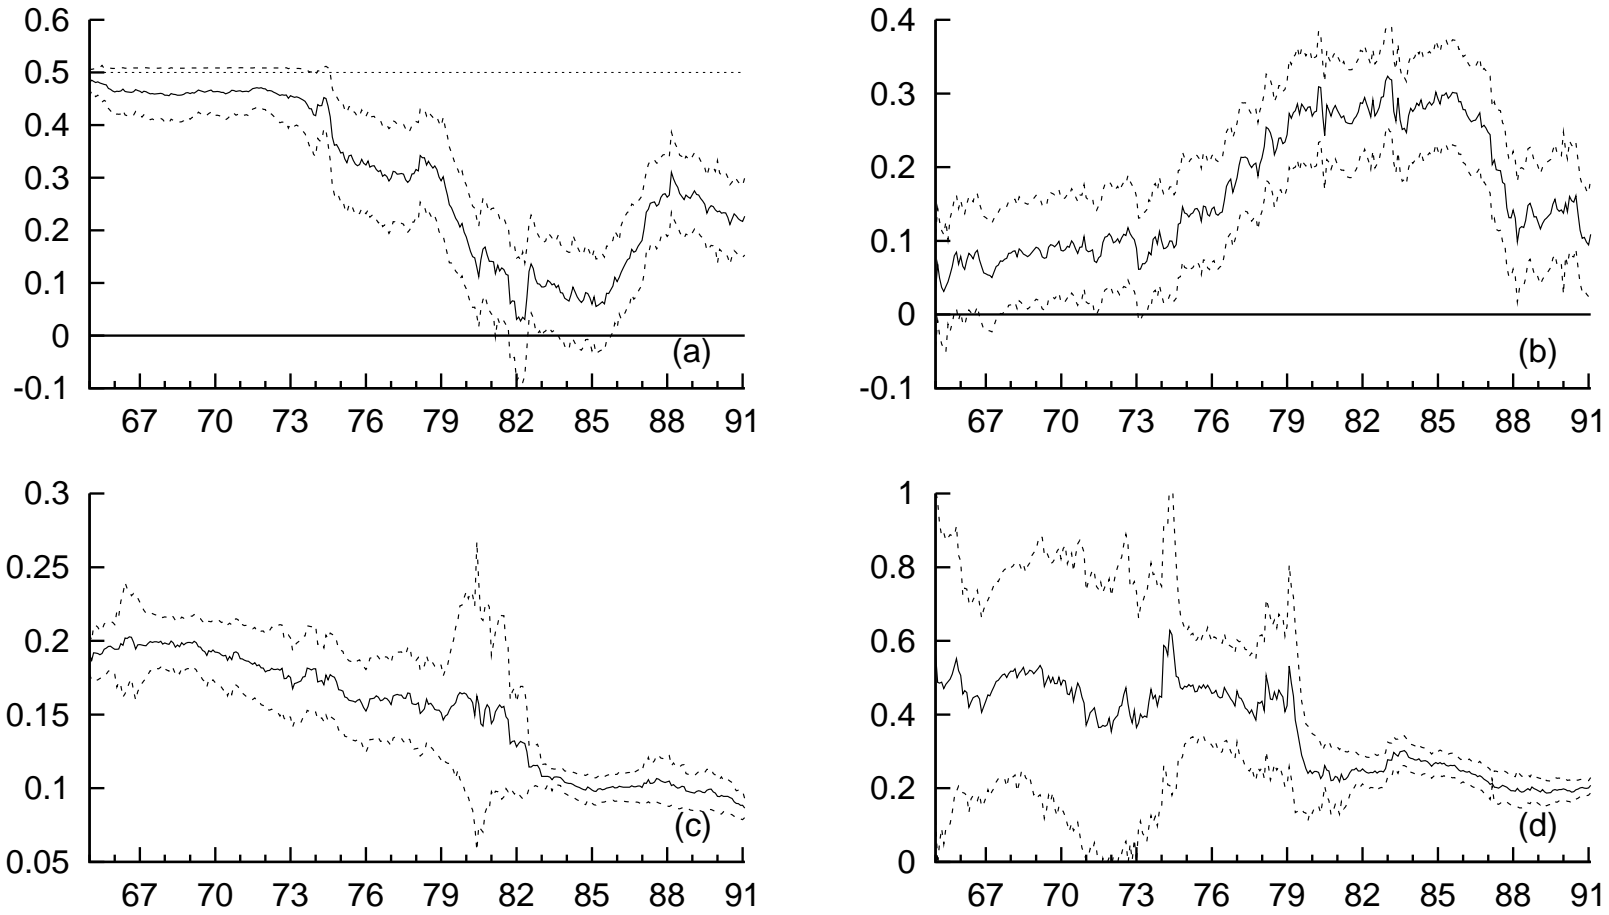


Figure 3: Selected Rolling Window parameter estimates for ARFIMA-SV model (1)-(2) for y_t (as specified in Table 2) with one-standard-deviation error bands. Panel (a): d , (b): ϕ_{12} , (c): σ_ε , (d): μ . Window width: 200 months. Sample periods: $t^* - (t^* + 199), t^* = 1, \dots, 314$. First period: 1965:1-1981:8, last period: 1991:2-2007:9. Horizontal line at 0.5 in panel (a) for fractional integration parameter d .

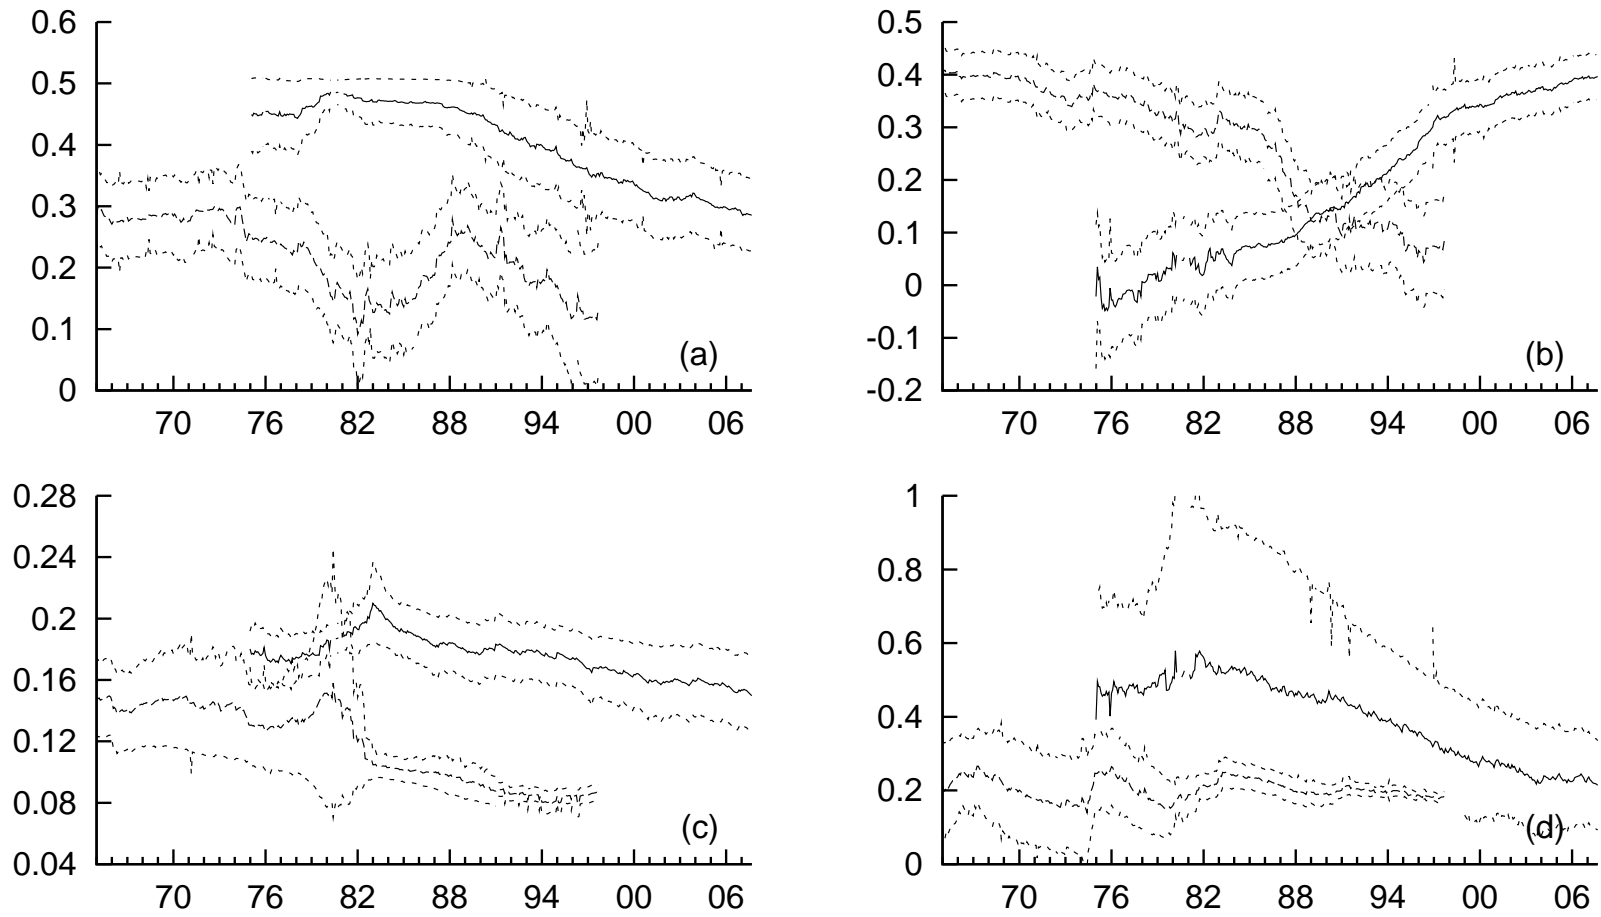


Figure 4: Selected recursive parameter estimates with one-standard deviation error bands for the ARFIMA-SV model (1)-(2) as specified in Table 2. Panel (a): d , (b): ϕ_{12} , (c): σ_ε , (d): μ . Solid line: estimates for first subsample, $1 - (t^* - 1), t^* = 120, \dots, 513$, shortest sample: 1965:1-1975:1. Dashed line: estimates for second subsample, $t^* - T, t^* = 1, \dots, 393$, shortest sample: 1997:9-2007:9.

3.4 Comparisons

This section compares our estimates with the results for an ARIMA-SV model, in particular model (1)-(2) with $d = 1$, $p = 14$ and $q = 1$. In this model the changes in the long run persistence are captured by changes in the MA(1) parameter θ_1 , see Stock and Watson (2007) and the references therein. Stock and Watson (2007) analysed these changes in the MA(1) parameter for different measures of U.S. inflation and interpreted these changes in terms of an unobserved components (UC) model. The equivalences between UC models and ARIMA models are discussed by Harvey (1989, §2.5). The ARIMA(0,1,1) model can be written as $(1 - L)y_t = \varepsilon_t + \theta_1\varepsilon_{t-1}$ where $\varepsilon_t \sim \mathcal{NID}(0, \sigma_\varepsilon^2)$ for $t = 1, \dots, T$ and with $-1 < \theta_1 < 1$. Values of θ_1 clearly above -1 correspond to a nonstationary $I(1)$ -process for y_t while values of θ_1 close to -1 correspond to a stationary $I(0)$ process.

Table 4 presents our results for the ARIMA-SV model and should be compared with Table 3. The significant change in persistence between subsamples is now captured by θ_1 which shifts from -0.66 to -0.96 , that is, from a clear $I(1)$ model to a nearly stationary specification. The ARIMA-SV likelihood in the first subsample is marginally lower than the ARFIMA-SV likelihood whereas it is somewhat higher in second subsample. The diagnostics for the ARIMA-SV model are slightly worse in the first subsample and slightly better in the second subsample. As the evidence of long memory is not strong in the second subperiod, one could have anticipated these results for the latter period.

We also compared our results with estimates for explicit UC-SV and UC-GARCH models as in Koopman and Bos (2004), but we omit the results as the resulting likelihoods were not higher than for the ARIMA-SV model presented here. In line with Table 4, those estimates show a drop in the importance of the trend component, an increase in the importance of seasonal changes and an overall decrease in the variance.

We have shown that our exact likelihood framework allows for straightforward comparisons between ARFIMA-SV and ARIMA-SV models using maximum likelihood estimation. We conclude from Tables 3 and 4 that we can reinterpret existing evidence on the shift in the long run dynamics based on a change in the MA(1) parameter θ_1 , as a shift in the long memory parameter d . This is an important finding because the long run characteristics of the ARFIMA-SV model differ in important respects from the long run behaviour of an ARIMA-SV model. For example, forecast intervals for the log of the price level based on ARFIMA-SV models and ARIMA-SV models are qualitatively different at long horizons. We agree with Stock and Watson (2007) that U.S. inflation may have become harder to forecast, but we provide a different interpretation of the evidence.

4 Summary and Conclusion

In this paper we have discussed and implemented exact maximum likelihood estimation for stationary ARFIMA models with stochastic volatility using Monte Carlo methods. The ARFIMA-SV model provides a good description of U.S. inflation, in particular for two subsamples: the Great Inflation period and the Post Great Moderation period. The fractional integration order d of U.S. inflation changed from around $d = 0.65$ during the Great Inflation period to about $d = 0.15$ after the Great Moderation. The constant parameter model is misspecified for the combined period, but the sv-part of the model clearly captures the Great Moderation in this case with a high volatility of volatility. Sensitivity analysis with rolling and recursive estimates confirms and extends these findings. We do not find evidence

of a change in the mean of inflation in the Post Great Moderation period. We compared our results with a related parameter stability analysis as in Stock and Watson (2007). For this purpose we used our exact likelihood framework to fit an ARIMA-SV model, where the MA(1)-parameter shifted from $\theta_1 = -0.66$ in the Great Inflation period, to $\theta_1 = -0.96$ in the Post Great Moderation period. Our analysis of parameter stability in the ARFIMA-SV model extends the existing applied econometric literature on this topic as it makes a different distinction between parameter changes in the long run dynamics and changes in the parameters for the short run dynamics while allowing for stochastic volatility.

Acknowledgements

We would like to thank Peter Boswijk and participants at the CREATES Symposium on long memory at the University of Aarhus, 2007, for valuable comments and suggestions on earlier versions of this paper.

A Approximating model

In this appendix we provide details for the construction of an approximating model for the ARFIMA-SV model (1)-(2), as discussed in §2.3. Consider $e_t = v_t^\sigma / \sqrt{f_t^\sigma}$ where v_t^σ and f_t^σ are obtained from the Durbin algorithm with $\sigma_t^2 = 1$. The approximating model is based on mean $E(e_t) = h_t + c_t$ and variance $\text{var}(e_t) = d_t$ where c_t and d_t , for $t = 1, \dots, T$, are constants across simulations. We have the approximating model given by

$$e_t = h_t + u_t, \quad u_t \sim \mathcal{NID}(c_t, d_t), \quad t = 1, \dots, T, \quad (9)$$

where h_t is modelled by (2). The values for c_t and d_t are determined by equating the first two derivatives with respect to h_t of the true density for $e_t \sim \mathcal{NID}(0, \exp(h_t))$ and the approximating model density for $e_t \sim \mathcal{NID}(h_t + c_t, d_t)$. The resulting set of equations need to be solved. Since the variable h_t is unobserved, optimal values for c_t and d_t need to be found iteratively based on estimates of h_t .

The conditional densities are given by

$$p(e|h; \psi) = \prod_{t=1}^T p_t, \quad g(e|h; \psi) = \prod_{t=1}^T g_t,$$

with $e = (e_1, \dots, e_T)$ and where $p(\cdot)$ refers to the true density, $g(\cdot)$ refers to the approximating density and

$$p_t = -0.5[\log 2\pi + h_t + \exp(-h_t)(v_t^\sigma)^2], \quad g_t = -0.5[\log 2\pi + \log d_t + (v_t^\sigma - c_t - h_t)^2 / d_t],$$

for $t = 1, \dots, n$. The first and second derivatives with respect to h_t are given by

$$\begin{aligned} \dot{p}_t &= -0.5[1 - \exp(-h_t)e_t^2], & \dot{g}_t &= (e_t - c_t - h_t) / d_t, \\ \ddot{p}_t &= -0.5 \exp(-h_t)e_t^2, & \ddot{g}_t &= -1 / d_t, \end{aligned}$$

for $t = 1, \dots, T$. Equating $\dot{p}_t = \dot{g}_t$ and $\ddot{p}_t = \ddot{g}_t$ leads to the solutions

$$c_t = e_t - h_t + 0.5d_t[1 - \exp(-h_t)e_t^2], \quad d_t = 2 \exp(h_t) / e_t^2,$$

which can be computed for a given value of h_t , for $t = 1, \dots, T$. We note that variance d_t is positive for any value of h_t . For given values of e_t , c_t and d_t , estimates of h_t can be based on the linear Gaussian state space model consisting of measurement equation (9) and transition equation (2) by using the Kalman filter and associated smoother, see Durbin and Koopman (2001). The resulting recursive procedure is repeated until convergence which usually takes a limited number of iterations, typically around five. The final model is the linear state space model with time-varying variances from which simulated series can be generated conditional on the observations e_1, \dots, e_T . This procedure is referred to as simulation smoothing; see Durbin and Koopman (2002) and the references therein. The resulting simulations can be regarded as realisations from $g(h|v^\sigma; \psi)$.

References

- Alogoskoufis, G. S. and R. Smith (1991). The Phillips curve, the persistence of inflation, and the Lucas critique: Evidence from exchange-rate regimes. *American Economic Review* 81, 1254–1275.
- Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica* 61, 821–856.
- Bai, J. and P. Perron (2003). Computation and analysis of multiple structural change models. *Journal of Applied Econometrics* 18, 1–22.
- Baillie, R. T. (1996). Long memory processes and fractional integration in econometrics. *Journal of Econometrics* 73, 5–59.
- 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.
- Ball, L. and N. G. Mankiw (1995). Relative-price changes as aggregate supply shocks. *Quarterly Journal of Economics* 110, 161–193.
- Benati, L. (2006). Investigating inflation persistence across monetary regimes. Report, Bank of England.
- Beran, J. (1994). *Statistics for Long-Memory Processes*. London: Chapman and Hall.
- Beran, J. and Y. Feng (2001). Local polynomial estimation with a FARIMA-GARCH error process. *Bernoulli* 7, 733–750.
- Beran, J. and N. Terrin (1996). Testing for a change of the long-memory parameter. *Biometrika* 83, 627–638.
- Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics* 51, 307–327.
- Bos, C. S., P. H. Franses, and M. Ooms (1999). Long memory and level shifts: Re-analyzing inflation rates. *Empirical Economics* 24, 427–449.
- Boswijk, H. P. and F. Klaassen (2004). Why frequency matters for unit root testing. Discussion paper 119/4, Tinbergen Institute, Amsterdam, The Netherlands.

- Bowdler, C. and L. Nunziata (2007). Trade union density and inflation performance: Evidence from OECD panel data. *Economica* 74, 135–159.
- Breidt, F. J., N. Crato, and P. De Lima (1998). On the detection and estimation of long memory in stochastic volatility. *Journal of Econometrics* 83, 325–348.
- Brockwell, A. E. (2007). Likelihood-based analysis of a class of generalized long-memory time series models. *Journal of Time Series Analysis* 28, 386–407.
- Brockwell, P. J. and R. A. Davis (1993). *Time Series: Theory and Methods (2nd ed.)*. USA: Springer-Verlag, New-York.
- Brodsky, J. and C. M. Hurvich (1999). Multi-step forecasting for long-memory processes. *Journal of Forecasting* 18, 59–75.
- Campa, J. M. and L. S. Goldberg (2005). Exchange-rate pass-through into import prices. *Review of Economics and Statistics* 87, 679–690.
- Cavaliere, T. and A. M. R. Taylor (2005). Stationarity tests under time-varying variances. *Econometric Theory* 21, 1112–1129.
- Cecchetti, S. G., P. Hooper, B. C. Kasman, K. L. Schoenholtz, and M. W. Watson (2007). Understanding the evolving inflation process. Report U.S. Monetary Policy Forum.
- Chan, N. H. and W. Palma (1998). State space modeling of long-memory processes. *Annals of Statistics* 26, 719–740.
- Cogley, T. and T. Sargent (2002). Evolving post-world war II U.S. inflation dynamics. In M. Gertler and K. Rogoff (Eds.), *NBER Macroeconomics Annual 2001*. Cambridge, U.S.: MIT Press.
- Cogley, T. and T. Sargent (2007). Inflation-gap persistence in the U.S. Working paper, University of California, Davis.
- Danielson, J. (1994). Stochastic volatility in asset prices: Estimation with simulated maximum likelihood. *Journal of Econometrics* 61, 375–400.
- Davidson, J. and N. Hashimzade (2008). Alternative frequency and time domain versions of Brownian motion. *Econometric Theory* 24, forthcoming, DOI: 10.1017/S0266466608080110.
- Dickey, D. A. and W. A. Fuller (1979). Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association* 74, 427–431.
- Diebold, F. X. and A. Inoue (2001). Long memory and regime switching. *Journal of Econometrics* 105, 131–159.
- Doornik, J. A. (2006). *An Object-oriented Matrix Programming Language - Ox*. London, U.K.: Timberlake Consultants Press.
- Doornik, J. A. and H. Hansen (1994). An omnibus test for univariate and multivariate normality. Working paper, Nuffield College, Oxford.
- Doornik, J. A. and M. Ooms (2003). Computational aspects of maximum likelihood estimation of autoregressive fractionally integrated moving average models. *Computational Statistics and Data Analysis* 42, 333–348.
- Doornik, J. A. and M. Ooms (2004). Inference and forecasting for ARFIMA models with an application to US and UK inflation. *Studies in Nonlinear Dynamics and Econometrics* 8(2), Article 14, <http://www.bepress.com/snnde/vol8/iss2/art14>.

- Durbin, J. (1960). The fitting of time series models. *International Statistical Review* 28, 233–244.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford, UK: Oxford University Press.
- Durbin, J. and S. J. Koopman (2002). A simple and efficient simulation smoother for state space time series analysis. *Biometrika* 89, 603–615.
- Durham, G. B. and A. Gallant (2002). Numerical techniques for maximum likelihood estimation of continuous-time diffusion. *Journal of Business and Economic Statistics* 20, 297–316.
- Emery, K. M. (1994). Inflation persistence and Fisher effects: Evidence of a regime change. *Journal of Economics and Business* 46, 141–152.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50, 987–1008.
- Engle, R. F. and A. D. Smith (1999). Stochastic permanent break. *Review of Economics and Statistics* 81, 553–574.
- Evans, M. and M. Wachtel (1993). Inflation regimes and the sources of inflation uncertainty. *Journal of Money, Credit, and Banking* 25, 475–510.
- Feng, Y., J. Beran, and K. Yu (2007). Modelling financial time series with SEMIFAR-GARCH model. *IMA Journal of Management Mathematics* 18, 395–412.
- Fuhrer, J. and G. Moore (1995). Inflation persistence. *Quarterly Journal of Economics* 110, 127–159.
- Galí, J. and M. Gertler (1999). Inflation dynamics: A structural econometric analysis. *Journal of Monetary Economics* 44, 195–222.
- Geweke, J. and S. Porter-Hudak (1983). The estimation and application of long-memory time series models. *Journal of Time Series Analysis* 4, 221–237.
- Golub, G. H. and C. F. Van Loan (1996). *Matrix Computations, 3rd Edition*. Baltimore, U.S.A.: Johns Hopkins University Press.
- Granger, C. W. J. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.
- Hamilton, J. D. (1990). Analysis of time series subject to changes in regime. *Journal of Econometrics* 45, 39–70.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman Filter*. Cambridge, UK: Cambridge University Press.
- Harvey, D. I., S. J. Leybourne, and A. M. R. Taylor (2006). Modified tests for a change in persistence. *Journal of Econometrics* 134, 441–469.
- Hassler, U. and J. Wolters (1995). Long memory in inflation rates: International evidence. *Journal of Business and Economic Statistics* 13, 37–46.
- Hooker, M. A. (2002). Are oil shocks inflationary? asymmetric and nonlinear specifications versus changes in regime. *Journal of Money, Credit and Banking* 34, 540–561.
- Hosking, J. R. M. (1981). Fractional differencing. *Biometrika* 68, 165–176.

- Kim, C.-J. (1993). Unobserved-Component time series models with Markov-Switching heteroscedasticity: Changes in regime and the link between inflation rates and inflation uncertainty. *Journal of Business and Economic Statistics* 11, 341–349.
- Kim, C.-J. and C. Nelson (1999). Has the U.S. economy become more stable? A Bayesian approach based on a Markov-Switching model of the business cycle. *Review of Economics and Statistics* 81, 608–616.
- Kim, C.-J., C. Nelson, and J. Piger (2004). The less volatile U.S. economy: A Bayesian investigation of timing, breadth, and potential explanations. *Journal of Business and Economic Statistics* 22, 80–93.
- Koopman, S. J. and C. S. Bos (2004). State space models with a common stochastic variance. *Journal of Business and Economic Statistics* 22, 346–357.
- Koopman, S. J., N. Shephard, and J. A. Doornik (1999). Statistical algorithms for models in state space using SsfPack 2.2. *Econometrics Journal* 2, 107–160.
- Kwiatkowski, D., P. C. B. Phillips, P. Schmidt, and Y. Shin (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics* 54, 159–178.
- Li, W. K. and A. E. McLeod (1986). Fractional time series modelling. *Biometrika* 73, 217–221.
- Ling, S. (2007). Testing for change points in time series models and limiting theorems for NED sequences. *Annals of Statistics* 35, 1213–1237.
- Ling, S. and W. K. Li (1997). On Fractionally Integrated Autoregressive Moving-Average models with Conditional Heteroskedasticity. *Journal of the American Statistical Association* 92, 1184–1194.
- Ljung, G. M. and G. E. P. Box (1978). On a measure of lack of fit in time series models. *Biometrika* 66, 67–72.
- Mandelbrot, B. B. (1969). Long-run linearity, locally Gaussian processes, H-spectra and infinite variances. *International Economic Review* 10, 82–111.
- Mandelbrot, B. B. and J. W. Van Ness (1968). Fractional Brownian motion, Fractional noises and applications. *SIAM Review* 10, 422–437.
- Marinucci, D. and P. M. Robinson (1999). Alternative forms of fractional Brownian motion. *Journal of Statistical Planning and Inference* 80, 111–122.
- McConnell, M. and G. Perez-Quiros (2000). Output fluctuation in the United States: What has changed since the early 1980's? *American Economic Review* 90, 1464–1476.
- Müller, U. K. (2005). Size and power of tests for stationarity in highly autocorrelated time series. *Journal of Econometrics* 128, 195–213.
- Müller, U. K. and M. W. Watson (2006). Testing models of low-frequency variability. Report W12671, National Bureau of Economic Research.
- Na, O. and S. Lee (2007). Moving estimates test with time varying bandwidth. *Journal of Multivariate Analysis* 98, 1356–1375.
- Nason, J. (2006). Instability in U.S. inflation 1967-2005. *Economic Review Q2*, 39–59.

- Nyblom, J. and A. C. Harvey (2001). Testing against smooth stochastic trends. *Journal of Applied Econometrics* 16, 415–429.
- Parke, W. R. (1999). What is fractional integration? *Review of Economics and Statistics* 81, 632–638.
- Perron, P. (1989). The great crash, the oil price shock, and the unit root hypothesis. *Econometrica* 57, 1361–1401.
- Pivetta, F. and R. Reis (2007). The persistence of inflation in the united states. *Journal of Economic Dynamics and Control* 31, 1326–1358.
- Primiceri, G. (2005). Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies* 72, 821–852.
- Quandt, R. E. (1960). Tests of the hypothesis that a linear regression system obeys two separate regimes. *Journal of the American Statistical Association* 55, 1184–1194.
- Robinson, P. M. (Ed.) (2003). *Time Series with Long Memory*, Oxford, UK. Oxford University Press.
- Sandmann, G. and S. J. Koopman (1998). Estimation of stochastic volatility models via Monte Carlo maximum likelihood. *Journal of Econometrics* 87, 271–301.
- Schweppe, F. (1965). Evaluation of likelihood functions for Gaussian signals. *IEEE transactions on Information Theory* 11, 61–70.
- Sensier, M. and D. J. C. van Dijk (2004). Testing for volatility changes in U.S. macroeconomic time series. *Review of Economics and Statistics* 86, 833–839.
- Shephard, N. (1996). Statistical aspects of ARCH and stochastic volatility. In D. Cox, D. Hinkley, and O. Barndorff-Nielsen (Eds.), *Time Series Models in Econometrics, Finance and Other Fields*, pp. 1–67. London, U.K.: Chapman and Hall.
- Shephard, N. (Ed.) (2005). *Stochastic Volatility Selected Readings*, Oxford, U.K. Oxford University Press.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84, 653–667.
- Sims, C. and T. Zha (2006). Were there regime switches in U.S. monetary policy? *American Economic Review* 96, 54–81.
- Sowell, F. (1992). Modeling long-run behavior with the fractional ARIMA model. *Journal of Monetary Economics* 29, 277–302.
- Stock, J. H. (1994). Unit roots, structural breaks and trends. In *Handbook of Econometrics, Volume IV*, pp. 2739–2841. Elsevier Science.
- Stock, J. H. and M. W. Watson (1999). Forecasting inflation. *Journal of Monetary Economics* 44, 293–335.
- 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.
- Taylor, S. J. (1982). Financial returns modelled by the product of two stochastic processes, a study of daily sugar prices, 1961-79. In O. D. Anderson (Ed.), *Time Series Analysis: Theory and Practice*, pp. 203–226. Amsterdam, The Netherlands: North Holland.
- Taylor, S. J. (1994). Modeling stochastic volatility. *Mathematical Finance* 4, 183–204.

Warne, A. and A. Vredin (2006). Unemployment and inflation regimes.
Studies in Nonlinear Dynamics and Econometrics 10(2), Article 2,
<http://www.bepress.com/snde/vol10/iss2/art2>.

Table 1: Descriptive statistics U.S. Core CPI inflation

Variable	min	max	mean	s.d.
P_t	32.6	211.6	113.92	58.09
100 $\Delta \log P_t$	-0.239	1.392	0.365	0.295
y_t	-0.246	1.386	0.365	0.274

NOTE: Source: BLS. Code P_t : CUUR0000SA0L1E. $y_t = \Delta \log P_t$ adjusted for fixed seasonal factors by OLS. Sample 1965:1–2007:9, $T = 513$.

Table 2: ARFIMA-SV Maximum Likelihood Estimation for y_t

Sample Column No. Parameter	1965:1–2007:9				1965:1–1982:12				1985:1–2007:9			
	1	2 s.e.	3	4 s.e.	5	6 s.e.	7	8 s.e.	9	10 s.e.	11	12 s.e.
d	0.386	(0.05)	0.285	(0.06)	0.474	(0.03)	0.474	(0.03)	0.131	(0.08)	0.136	(0.08)
ϕ_1	0.028	(0.07)	0.032	(0.07)	-0.038	(0.08)	-0.089	(0.08)	0.048	(0.10)	0.023	(0.10)
ϕ_{11}	0.108	(0.04)	0.131	(0.04)	0.086	(0.07)	0.058	(0.07)	0.175	(0.06)	0.160	(0.06)
ϕ_{12}	0.306	(0.04)	0.396	(0.05)	0.061	(0.07)	0.064	(0.07)	0.292	(0.06)	0.298	(0.06)
ϕ_{13}	0.029	(0.05)	0.080	(0.05)	-0.048	(0.07)	-0.056	(0.07)	0.129	(0.06)	0.125	(0.06)
σ_ε	0.175	(0.01)	0.150	(0.03)	0.208	(0.01)	0.207	(0.02)	0.103	(0.00)	0.101	(0.01)
σ_ξ			0.143	(0.04)			0.117	(0.07)			0.095	(0.07)
ρ			0.982	(0.01)			0.962	(0.04)			0.954	(0.06)
μ	0.320	(0.21)	0.214	(0.12)	0.443	(0.49)	0.521	(0.43)	0.252	(0.03)	0.236	(0.03)
$\beta(\text{i80:7})$	-1.240	(0.16)	-1.256	(0.25)	-1.105	(0.19)	-1.114	(0.21)				
LL	164.690		212.782		31.248		34.382		231.059		232.693	
T	513		513		216		216		273		273	
N	28.220	[0.00]	34.527	[0.00]	4.582	[0.10]	5.286	[0.07]	6.704	[0.04]	7.163	[0.03]
Q	38.569	[0.01]	57.829	[0.00]	13.833	[0.79]	15.218	[0.71]	28.373	[0.08]	28.829	[0.07]

NOTE: ML estimates for the parameters of model (1-2) with extra coefficients for dummy variables in the mean $\beta(\text{i80:7})$ and $\sigma_\varepsilon = \exp(\gamma/2)$. Full sample and two subsamples. Standard errors in parentheses. n: number of observations. LL: log-likelihood for y_t . N: Normality tests of Doornik and Hansen (1994). Q: Residual serial correlation test of Box-Ljung for 24 lags, see Ljung and Box (1978) and Li and McLeod (1986). The p -values are reported in brackets.

Table 3: ARFIMA-SV MLE results for Δy_t

Sample Parameter	1965:1–2007:9		1965:1–1982:12		1985:1–2007:9	
	s.e.		s.e.		s.e.	
$d - 1$	-0.707	(0.06)	-0.292	(0.08)	-0.882	(0.07)
ϕ_1	0.018	(0.07)	-0.298	(0.09)	0.017	(0.10)
ϕ_{11}	0.128	(0.04)	0.075	(0.07)	0.154	(0.06)
ϕ_{12}	0.392	(0.04)	0.061	(0.08)	0.291	(0.06)
ϕ_{13}	0.081	(0.05)	-0.049	(0.07)	0.110	(0.06)
σ_ε	0.149	(0.02)	0.204	(0.02)	0.101	(0.01)
σ_ξ	0.145	(0.04)	0.101	(0.07)	0.093	(0.08)
ρ	0.982	(0.01)	0.956	(0.05)	0.952	(0.06)
$\beta(\text{i80:7})$	-1.348	(0.25)	-1.287	(0.22)		
$\beta(\text{i80:8})$	1.139	(0.25)	0.999	(0.22)		
LL	214.524		37.034		232.542	
T	513		216		273	
N	32.202	[0.00]	3.643	[0.16]	2.595	[0.27]
Q	58.245	[0.00]	12.647	[0.86]	25.645	[0.14]

NOTE: ML estimates for the parameters of ARFIMA-SV model (1)-(2) with extra coefficients for dummy variables in the mean $\beta(\text{i80:7})$ and $\beta(\text{i80:8})$. $\sigma_\varepsilon = \exp(\gamma/2)$. Standard errors in parentheses. T : number of observations. LL: log-likelihood for Δy_t . N: Normality tests of Doornik and Hansen (1994). Q: Residual serial correlation test of Box-Ljung for 24 lags. The p -values are reported in brackets.

Table 4: ARIMA(14,1,1)-SV MLE results ($d=1$). Dependent variable: Δy_t

Sample Parameter	1965:1–2007:9		1965:1–1982:12		1985:1–2007:9	
	s.e.		s.e.		s.e.	
ϕ_1	0.230	(0.06)	0.059	(0.12)	0.080	(0.07)
ϕ_{11}	0.101	(0.04)	0.049	(0.07)	0.124	(0.06)
ϕ_{12}	0.347	(0.04)	0.029	(0.07)	0.258	(0.06)
ϕ_{13}	-0.021	(0.05)	-0.073	(0.07)	0.055	(0.06)
ϕ_{14}	-0.043	(0.04)	0.067	(0.08)	0.080	(0.06)
θ_1	-0.907	(0.03)	-0.660	(0.10)	-0.960	(0.03)
σ_ε	0.150	(0.02)	0.204	(0.02)	0.101	(0.01)
σ_ξ	0.156	(0.04)	0.133	(0.07)	0.080	(0.06)
ρ	0.979	(0.01)	0.953	(0.05)	0.967	(0.04)
$\beta(\text{i80:7})$	-1.329	(0.24)	-1.297	(0.22)		
$\beta(\text{i80:8})$	1.056	(0.24)	1.010	(0.22)		
LL	212.606		36.686		234.741	
T	513		216		273	
N	36.070	[0.00]	5.501	[0.06]	2.161	[0.34]
Q	64.671	[0.00]	18.987	[0.39]	21.949	[0.23]

NOTE: ML estimates for the parameters of model (1-2) with extra coefficients for dummy variables in the mean $\beta(\text{i80:7})$ and $\beta(\text{i80:8})$. $\sigma_\varepsilon = \exp(\gamma/2)$. Standard errors in parentheses. T : number of observations. LL: log-likelihood for Δy_t . N: Normality tests of Doornik and Hansen (1994). Q: Residual serial correlation test of Box-Ljung for 24 lags. The p -values are in brackets.