Inflation persistence in the euro-area, US, and new members of the EU: Evidence from time-varying coefficient models

Zsolt Darvas and Balázs Varga[•]

April 22, 2006

Abstract

This paper studies inflation persistence with time-varying-coefficient autoregressions in response to recently discovered structural breaks in historical inflation time series of the euro-area and the US. To this end, we compare the statistical properties of the well known ML estimation using the Kalman-filter and the less known Flexible Least Squares estimator by Monte Carlo simulation. We also suggest a procedure for selecting the weight for FLS based on an iterative Monte Carlo simulation technique calibrated to the time series in question. We apply the methods for the study of inflation persistence of the US, the euro-area and the new members of the EU.

Keywords: flexible least squares, inflation persistence, Kalman-filter, time-varying coefficient models

JEL classifications: C22, E31

[•] Darvas is Assistant Professor, Department of Mathematical Economics and Economic Analysis, Corvinus University of Budapest. E-mail: zsolt.darvas@uni-corvinus.hu. Varga is Ph.D. Candidate, Corvinus University of Budapest, E-mail: balazs.varga@stud.uni-corvinus.hu. A current version of the paper, the GAUSS codes used, and the data set are available at http://www.uni-corvinus.hu/darvas.

1 Introduction

This paper aims to contribute to the literature in three ways. First, we compare the statistical properties of two time-varying coefficient (TVC) methods: the maximum likelihood estimation of the state-space representation of TVC autoregression using the Kalman-filter, and the less known, distribution-free estimator of the same model via Flexible Least Squares (FLS), which was suggested by Kalaba and Tesfatsion (1988, 1989, 1990). Second, we propose a procedure to set the weight parameter needed for the FLS: to our knowledge, such a procedure is not yet available. The proposed procedure is based on an iterative Monte Carlo simulation technique calibrated to the time series in question. Third, we study inflation persistence (IP) in the euro-area and the US, for which IP has been extensively studied but not with TVC models. Fourth, we study IP in the new members of the EU using TVC models and compare persistence estimates for these countries to that of the euro-area, which also has an implication for the optimum currency area literature.

It is widely accepted that inflation is often exposed to numerous macroeconomic shocks which pull it away from its mean, which is generally identified by the central bank's inflation target. Shocks can be persistent or could have persistent effects on inflation due to, for example, nominal rigidities, leading to persistent deviations of inflation from its target. Therefore, knowing the persistence of these shocks and inflation deviations from target plays an essential role for the central bank whose primal aim is to achieve price stability. The adjustment of inflation towards its long run level after a shock can be characterized by the speed with which it converges back to its mean. The higher this speed is, the less complicated may be the central bank's price stability maintaining task. Inflation persistence is a measure of this convergence speed, based on different kinds of properties of the impulse response function within the model built to describe inflation.

Inflation persistence has been studied by various models, ranging from simple autoregressions to well-structured dynamic general equilibrium models.[1] In studying univariate autoregressive time series models, many authors found very high persistence or even could not reject the hypothesis of a unit root for a 50 year long sample stretching from the post world war II era, both in the US and in the euro-area. Latter studies found that inflation series have several structural breaks and most of these could be explained by corresponding historical events, for

[1] See, for example, Smets (2006) for a summary of the results of the ECB's Inflation Persistence Network.

example, the oil crises of the 70's. When studying the properties of the estimated autoregressive models for sub-periods identified by the break points, persistence turned out to be significantly smaller, particularly in the more recent periods. Hence, inflation persistence could be changing in time. Naturally, a change in inflation persistence could be the result of (a) change in the type of underlying shocks, (b) change on the persistence of the underlying shocks, (c) change in the monetary policy reaction function, (d) change in the way the economy responds to shocks or monetary policy actions, or (e) the fact that a linear approximation of an otherwise non-linear underlying structure is poor. A univariate autoregressive model estimated on different samples can not discriminate among these alternatives. Obviously, a time-varying coefficient autoregression also can not discriminate among these alternatives, but allows us to investigate changes in persistence more accurately and particularly, to highlight the dating and amplitude of breaks.

The rest of the paper is organized as follows. Section 2 describes Kalman-filtering and the less known FLS and presents are our iterative procedure to determine the weighting parameter of FLS. In Section 3 we perform a Monte Carlo study of the properties of Kalman-filtering and FLS for various data generating processes. Section 4 presents the empirical results for the US, euro-area, and new members of the EU. Finally, some concluding remarks are presented in Section 5.

2 Methodology

Time-variation in parameters of an economic model could be due to various reasons, for example, change in the behavior or agents, or could be the consequence that the functional form of the estimated model differs from the functional form of the underlying data generating process.[2] A standard approach to estimate time-varying coefficient models is the maximum likelihood estimation of a state-space representation of unobserved components models using the Kalman-filter to evaluate the likelihood function. However, this approach is built on the assumption of a certain distribution of the innovations, which is usually set to be the Gaussian distribution, which could not be valid.

[2] For example, most estimated economic models are linear, while the true data generating process could be non-linear.

Therefore, in this paper we also aim to compare Kalman-filtering to a less frequently used method, the so called Flexible Least Squares (FLS) introduced by Kalaba and Tesfatsion (1988, 1989, 1990). Since this estimation approach is less frequently used in economic applications we briefly describe it below.

2.1 Flexible Least Squares and its relation to Kalman-filtering

The FLS algorithm solves the time-varying linear regression problem with a minimal set of assumptions. Suppose y_t is the time t realization of a time series for which a time-varying coefficient model is fitted,

(1) $y_t = x_t'\beta_t + u_t$, $t = 1,...,T$,

where $x_t = (x_{0,t},...,x_{K-1,t})$ denotes a $K\times 1$ vector of known exogenous regressors (which can also contain the lagged values of y_t), $\beta_t = (\beta_{0,t},...,\beta_{K-1,t})$ denotes the $K\times 1$ vector of unknown coefficients to be estimated, which can change in time, and u_t is the approximation error.

Two main assumptions are needed for the formulation of a cost function to be minimized. First, the prior measurement specification states that the residual errors of the regression are small, that is,

(2) $y_t - x_t'b_t \approx 0$, $t = 1,...,T$.

Second, the prior dynamic specification declares that the vector of coefficients evolves slowly over time:

(3) $\beta_{t+1} - \beta_t \approx 0$, $t = 1,...,T$-1.

A basic problem for the investigator is to find a coefficient sequence estimate, $(\hat{\beta}_1,...,\hat{\beta}_T)$, which satisfies both of these prior assumptions in an acceptable manner. The idea of the FLS method is to assign two types of residual error to each possible coefficient sequence estimate. One of these consists of the sum of squared residual measurement errors:

(4) $r_M^2(\beta,T) = \sum_{t=1}^{T}(y_t - x_t'\beta_t)^2$,

matching the above mentioned measurement prior. The other – following the smoothness prior – is the sum of squared residual dynamic error, formally

(5) $$r_D^2(\beta,T)=\sum_{t=1}^{T-1}(\beta_{t+1}-\beta_t)'(\beta_{t+1}-\beta_t).$$

Taking the collection *P(T)* of all pairs of these two kinds of sum of squared residual errors as β ranges over the space of possible coefficient sequences defines the $\left(r_D^2(\beta,T),r_M^2(\beta,T)\right)$ residual possibility set on the positive quarter of the two-dimensional plane. The vector minimums of this set, analogously with the idea of a Pareto-efficient frontier, yield the $P_F(T)$ residual efficiency frontier. Each point of this frontier represents a corresponding sequence of estimates which can be found by minimizing the weighted sum of the two types of measurement errors. Thus, with a given μ weighting parameter, Kalaba and Tesfatison (1988) define the incompatibility cost assigned to any β coefficient sequence as

(6) $$C(\beta,\mu,T)=\mu\cdot r_D^2(\beta,T)+r_M^2(\beta,T).$$

Minimization of the incompatibility cost for β, given any $\mu>0$, leads to a unique estimate for β as $\hat{\beta}^{FLS}(\mu,T)=\left(\hat{\beta}_1^{FLS}(\mu,T),\ldots,\hat{\beta}_T^{FLS}(\mu,T)\right)$, which is the flexible least squares solution, conditional on μ and *N*. This conditional minimization is performed with a dynamic programming algorithm.

Consequently, there are a continuum number of FLS solutions for a given set of observations, depending on the weight parameter μ. The solutions lie between two extremes. First, if μ approaches zero, the incompatibility cost function places absolutely no weight on the smoothness prior. This means that while r_D^2 stays relatively large, r_M^2 will be brought down close to zero, resulting in a rather erratic sequence of estimates. Second, as μ becomes arbitrarily large, the cost function assigns all importance to the dynamic specification. This case yields the ordinary least squares (OLS) solution, i.e. r_M^2 is minimized subject to $r_D^2=0$.[3] Therefore, the selection of the weighing parameter is a highly critical part of the FLS procedure, as the appropriate coefficient sequence lies somewhere between the most variable and the fully stable – OLS – solution.

The major difference between FLS and Kalman-filtering is that the latter assumes a certain probability distribution for the measurement and dynamic errors and uses this probabilistic assumption to form a likelihood function, which, when maximized, allows a unique estimated sequence for the time-varying parameters. The FLS, on the contrary, does not require any probabilistic assumptions. The FLS only postulates a measurement prior and a dynamic prior

[3] Our estimations will also illustrate the convergence of the FLS estimate to the OLS estimate as μ increases.

(equations 2 and 3) and yields a continuum number of estimates along the residual efficiency frontier.[4]

2.2 A new procedure to set the weighting parameter of FLS

Kalaba and Tesfatsion (1988, 1989, 1990) suggested to study the properties of the FLS estimates along the residual efficiency frontier and to draw conclusions based on some general features of the estimates. Their approach is motivated by the fact that model estimation for economic processes is, intrinsically, a *multicriteria* optimization problem. The assumptions generally made for estimation typically fall into three categories: measurement, dynamic, and probabilistic.[5] However, an actual economic process will behave in a manner that is incompatible to some degree with each of these assumptions. Associated with any set of estimates there will be a set of discrepancy terms reflecting the incompatibility of the assumptions with the data and an econometrician undertaking the estimation would presumably want each type of discrepancy term to be small in some sense. Kalaba and Tesfatsion therefore suggest dropping probabilistic assumptions and studying the incompatibility of the data with assumptions by giving different weights to measurement and dynamic errors. Consequently, Kalaba and Tesfatsion do not suggest any particular way to set a given value for the weighting parameter, and we are not aware of such a suggestion by other authors as well.

Econometricians, however, are usually interested in a single estimate having certain properties, for example, unbiasedness, 'small' variance, and so on. To this end, both classical and Bayesisan econometricians generally postulate certain probabilistic assumptions in order to derive a point estimate and a confidence band reflecting the uncertainty of the point estimate. The probabilistic assumptions, though, could be dubious.

We aim to bridge the probability-free FLS approach yielding a continuum number of solutions with the standard probability-based techniques leading to a unique estimate by suggesting a procedure, which maintains the probability-free assumption of FLS but leads to a unique estimate. This estimate can be regarded, to a certain extent, as the 'most likely' estimate of the time-varying parameter sequence for the particular model and time series

---

[4] See Tesfatsion and Veitch (1990) for a money demand application.

[5] Assumptions on the cross-sectional dimension of the model add a further type.

under study. In practice, our procedure sets the weighting parameter of the FLS and regards the FLS estimate associated with this weighting parameter to be the 'most likely' estimate.

Our suggested procedure is based on the following iterative algorithm:

As a staring point, estimate an initial time-varying coefficient model (equation (1)) for the times series studied, that is, set an initial arbitrary μ and calculate the FLS solution.

Perform a stochastic simulation of the estimated model. This step can be view as a simulation of a model that fits to the times series under study (taking into account time variation in parameters). Note that while the FLS estimator is distribution free, we have to assume a certain distribution for the stochastic simulation.

For each of the simulated series select μ which minimizes, say, the mean quadratic difference between the FLS estimate and the parameter sequence of the data generating process used in Step 2. That is, for each of the simulated series perform many FLS estimates for various μ values with $\mu \in [\mu^L, \mu^U]$. The mean of all μ-estimates can be regarded as a first estimate of μ, which we denote as $\mu^{(1)}$.

Estimate the time-varying parameters of the model with FLS for the time series studied, using the chosen $\mu^{(1)}$ in Step 3.

Return to Step 2: perform a stochastic simulation of the estimated model of Step 4; then to Step 3: select μ for each of the simulated series and calculate their mean to arrive a new μ-estimate to be denoted as $\mu^{(2)}$; and to Step 4: estimate the model for the data using $\mu^{(2)}$ in the FLS algorithm.

Continue this iterative procedure till $\mu^{(i)}$ converges, that is, till $\mu^{(i)}- \mu^{(i-1)}<\varepsilon$.

The following selections should be made in order to make this procedure operative and to study its robustness.[6]

The initial μ in Step 1. We study the sensitivity of our procedure for a wide selection of initial μ.

[6] Some technical parameters are set to the following values: the number of Monte-Carlo draws in Step 2 is set to 10000; the possible range of μ is determined between $\mu^L = 0$ and $\mu^U = 10^6$; and the criteria for convergence, ε, is set to a 'small' number.

The probability distribution of the random numbers to be used for the stochastic simulation in Step 2. As a benchmark we use Gaussian random numbers but will study the sensitivity of the results to the choice of the distribution.

The loss function to be used in the selection of μ for each simulated series. In Step 3 above we indicated the mean quadratic difference, but we also study the sensitivity of the results other loss functions, such as the mean absolute difference.

NOTE: we have not yet completed all of these tasks and therefore do not report our preliminary results in this version of the paper. In both the simulation exercise of Section 3 and in the empirical estimation of Section 4 we report FLS results for various μ but not for our 'optimal' μ.

3 The ability of FLS and Kalman-filtering to capture time-varying parameters: a simulation study

In this section we compare the properties of Kalman-filtering and the FLS for the estimation of linear autoregressive models with time-varying coefficients by Monte Carlo simulation. To this end, we set up and calibrate various data generating processes (DGP), stochastically simulate them, and apply the FLS estimation technique to simulated time series with the aim of comparing the estimation results to the known properties of the DGP. Calibration of the DGPs is based on inflation time series of the euro-area and the US. Our study includes four types of time-varying coefficient DGPs:

Discrete shifts in parameters;

Linear deterministic change in parameters;

Sinusoidal deterministic change in parameters,

Unit root process of the parameters (stochastic).

The simulation study helps us to determine the accuracy of the Kalman-filtering and FLS method.

In the current version of the paper we study simple AR(1) processes, that is,

(7) $y_t = \beta_{0,t} + \beta_{1,t} \cdot y_{t-1} + u_t, \qquad t = 1, \ldots, T,$

where y_t is the time series under study (i.e., simulated series in this section and the inflation rate in the next empirical section); $\beta_{0,t}$ and $\beta_{1,t}$ are the time-varying coefficients, for which we assume the four alternative processes described above for the simulation study; and u_t is the error term, which will be set to a series of Gaussian random numbers. The initial condition, y_0, is always set to be equal to the initial mean of the process, $y_0 \equiv \beta_{0,0}/(1-\beta_{1,0})$. All simulations were performed using a sample size T = 200, which is typically used in the literature on one hand, and is close to our inflation sample size on the other.

As a starting point we apply OLS, FLS (with different μ-*s*) and Kalman-filtering/smoothing to a model with fixed parameters. The upper block of Table 1 shows the RMSE (Root Mean Squared Error) of the various estimation techniques when the parameters of the data generating process are constant. That is, assuming that $\beta_{0,t} = \beta_0 = 0.2$, $\beta_{1,t} = \beta_1 = 0.9$, and var($u_t$) = 0.25^2 in equation (7), we simulated equation (7) and applied the various estimation techniques. Repeating the simulation and estimation exercise 1000 times, we arrive at 1000 estimated parameter sequences for each estimator (the OLS estimator, obviously, yields a parameter, not a time-varying parameter sequence). We computed the quadratic difference between the estimated and true sequence, which – in the form of RMSE – are averaged across all repetitions and reported in the table. Among the various estimators the OLS captures the most adequately the constant nature of the parameters. All time-varying coefficient methods attribute some of the error variance to parameter changes. This result highlights the importance of applying tests for structural breaks before estimating any time-varying coefficient methods.

The rest of Table 1 shows RMSE of the estimators when the parameters change in time; the exact specifications will be detailed below. The plot of the parameter sequences are shown in the top panel of Figures 1 to 4. In Figure 4, which shows the case with unit root in parameters, obviously only a certain realization is shown for the parameters. The top panels of these figures also show a certain realization for y_t. The middle panels of the figures show, in addition to the true autoregressive parameter, $\beta_{1,t}$, the OLS estimate based on a fixed coefficient model, and four different FLS estimates corresponding to weights: μ = 10, 10^2,

10^3, and 10^4, and the Kalman-filter and Kalman-smoother estimates.[7] The bottom panels show the true constant, $\beta_{0,t}$ with its corresponding OLS, FLS and Kalman-filter estimates.

Figures 1-4 help us to get a visual impact of the nature of the FLS-estimator and Kalman-filter and smoother.

3.1 Discrete shifts in parameters

In Figure 1 we study the properties of our TVC model estimators when there are single discrete shifts in parameters of the data generating process. Specifically, the intercept of our simulated process shifts at observation $t = 60$ from 0.15 to 0.45, and the autoregressive parameter changes at observation $t = 140$ from 0.9 to 0.7. Notice that these parameters imply the same mean for the series until $t = 60$ and from $t = 140$. The simulated series, shown on the top panel of the figure, resembles somewhat to US inflation (see Figure 5).

The OLS does an awful job. The estimated parameters indicate a random walk without drift, as the estimated autoregressive parameter is almost 1 and the intercept is close to zero. This is not surprising and the biases of the OLS estimates in the face additive outlier have already been reported in the literature.

The FLS, although better than the OLS, but can not capture well the sudden changes in parameters. Depending on the value of μ, the estimated time-varying parameters are either rather smooth, or change at both shifts. For example, when $\mu = 10$, the autoregressive parameter rises from about 0.4 to 0.7 at observation $t = 60$, when only the intercept changed, and falls from 0.75 to 0.4 at $t = 140$. The estimated intercept also changes at both dates, although more smoothly, and has values much larger than implied by the data generating process. The Kalman-filter and smoother also can not capture well sudden discrete changes. The middle and bottom panels of Figure 1 indicates that for the certain draw shown both the Kalman-filtered and smoothed parameter estimate move along quite close to the FLS $\mu = 10^2$ sequence.

[7] For Kalman-filtering we assumed random walk processes for the parameters in the state-space representation in all cases. Probably Kalman-filter estimates could be improved if we assumed the true DGP for the time-varying parameters, but we motivate our choice for the following two reasons. First, the random walk assumption is the counterpart of the FLS assumptions. Second, when applying the Kalman-filter to a real-world series, the DGP of time variation is also unknown.

According to second panel of Table 1, the FLS estimator with $\mu = 10^2$ and $\mu = 10^3$ are the best and better than the Kalman-Filter and smoother.

3.2 Continuous deterministic change in parameters

In the above example we ignored the assumption of smooth changes in the underlying parameter sequences. Our next simulations thus produce more continuous variations: now we allow our parameters to change linearly in time: while $\beta_{0,t}$ rises from 0.1 to 0.9, the autoregressive parameter does just the opposite: it decreases from 0.9 to 0.1, again causing the mean to stay the same throughout the whole sample.

Figure 2 shows the results. The OLS crosses the true parameter values at about one third of the sample. The FLS does quite well now, especially with $\mu = 10^2$. In this case, after some initial misdirection up to about where the OLS line crosses the true sequence, it follows closely the parameter paths of the data generating process. The Kalman-filter and smoother perform only a little worse than the FLS estimator with $\mu = 10^2$, this can be especially seen on the third panel of Table 1.

3.3 Sinusoid change in $\beta_{1,t}$

The next data generating process assumes that the intercept remains constant and the autoregressive parameter follows a sinusoid process, namely, $\beta_{1,t} = 0.6 + 0.3 \cdot \sin(\pi / 100 \cdot t)$. Sinusoid change is a good example of smooth nonlinear variations, as we will see, the capabilities of the FLS estimator can be illustrated well through this case.

Results of a certain Monte-Carlo draw can be seen in Figure 3. The OLS estimates indicate a random walk process for y_t with the autoregressive parameter very close to one and even higher than the highest value of the whole data generating parameter sequence (0.9). The FLS estimator, on the other hand, does again very well when $\mu = 10^2$. It can nicely capture both the sinusoid path of $\beta_{1,t}$ and the constant value of $\beta_{0,t}$, although with a little remaining sinusoid variation in the latter. The Kalman-filter and smoother's estimates are similar to those of the FLS at $\mu = 10^2$ but again, the Monte Carlo simulation shows an advantage for the former in the fourth panel of Table 1.

3.4 Parameters follow random walk

Figure 4 shows estimation results of a certain Monte-Carlo draw when both parameters follow independent random walks without drift. Standard error of the innovations was set to 0.05 and the initial value of both parameters was set to 0.5. The OLS estimator indicates a value of 0.7 for the autoregressive parameter, which is larger than even the maximum of $\beta_{1,t}$ in this certain outcome shown on the figure. The FLS estimator, on the other hand, does quite well again when $\mu = 10^2$, and captures well the time-variation of $\beta_{1,t}$, and a little bit better again than the Kalman-smoother For $\beta_{0,t}$, on the other hand, the Kalman-smoother slightly outperforms the FLS, as can be seen on the bottom panel of Table 1.

3.5 Summing up

Let us summarize the main findings of the simulation study.

1. The OLS is the best when parameters are constant, but is generally upward biased even compared to the mean (or time-average) of the autoregressive parameter when there is time variation.
2. Neither FLS, nor the Kalman-filter and Kalman-smoother can capture sudden changes in parameters
3. The FLS with $\mu = 10^2$ can reasonably well capture more gradual changes in parameters, like continuous deterministic change, deterministic sinusoid path and random walk behavior, especially for the autoregressive parameter, which determines persistence and thus carries more importance for us.
4. The Kalman-filter and smoother is similarly good for capturing gradual changes in parameters, although a slightly worse than the FLS. The Kalman-smoother dominates the Kalman-filter in all cases we studied

4 Estimates of inflation persistence for the US, Euro-area, and Hungary[8]

The empirical part of the paper has two main goals. First, we study the euro-area and the US and would like to compare our time-varying coefficient results to some of the results of exuberant literature, most notably, to the results of the Inflation Persistence Network of the ECB (Smets, 2006). Second, we also study the time series of the newly joined EU members of Central Europe waiting for accession to the euro-area, for which much less research has been conducted even with fixed coefficient models, and we know about only one paper applying time-varying coefficient methods (Dossche – Everaert [2005]).

The case of euro-area candidate countries raises the issue of convergence. These countries must fulfill, among others, the Maastricht criterion related to the level of inflation. However, similar persistence to that of the euro-area will be crucial for the optimality of the common monetary policy. Consequently, our paper also has an implication for the optimum currency area literature.

The sample period of our data is the following:

US CPI	1957Q1 – 2005Q2
EA HICP	1970Q1 – 2003Q4 (taken from the AWM database)
Hungarian CPI	1976Q2 – 2005Q4

Inflation is defined as *Δln(price level) x100.*

Time-varying coefficient analysis is especially inevitable in the case of the euro-area and Hungary. The euro-area did not exist before 1999 and its data were constructed by aggregating country time series. It is rather likely that euro-area data include structural breaks. Hungary had a socialist economic system till 1990, when transition to a market economy had speeded up, accompanied by and a substantial economic downturn and pick up in inflation (partly due to price liberalization) in the early nineties. Therefore, it would be a complete nonsense to model Hungarian inflation with a fixed coefficient model incorporating both the planned and the market systems.[9] However, with time-varying coefficient models we can study changes in the parameters of the inflation process using a longer sample.

[8] In the current version of the paper we report results for only Hungary among the new members of the EU, but will augment the analysis for all new members.

[9] In applied research the starting year is usually set to 1995, which renders the number of observations short.

Estimation results are shown on Figures 5-7. We show two results for the FLS: with μ equal to 10 and 10^2, because these values proved to be the best in the simulation study presented in Section 3. For the Kalman-filter, we show both the filtered and smoothed estimate.

It is evident for all three countries that OLS estimates are likely upward biased: OLS parameter estimates are almost the highest considering all time-varying coefficient estimations and are very far from the time-average of them. Recall from the previous section that the OLS estimator proved to be upward biased, compared to the time average of the parameters, when there were changes in the parameters of the data generating process.

Inflation persistence tends to be higher in times of high inflation in the cases of all three countries. Thus, inflation persistence changes. For the Kalman-filter we also calculated the error bands which indicated that changes in inflation persistence are significant.

Considering the magnitudes, in the US, the high inflation persistence episode of the period of the oil crises showed an autoregressive parameter value of around 0.7-0.8, which declined to around 0.1-0.2 in recent years. In the euro-area, the autoregressive parameter declined from 0.4-0.6 to zero. In fact, three of the four estimates showed in the middle panel of Figures 6 even indicates a negative parameter, while our preferred FLS estimate with $\mu = 10^2$ indicates a value of 0.2. Hence, considering our preferred estimator inflation persistence in the US and the euro-area are similar. The picture is quite different for Hungary. Our preferred estimation (FLS with $\mu = 10^2$) indicates a parameter value of 0.3-0.4 in the socialist era, which, however, jumps at the times of the transition and the latest estimate is around 0.7-0.8. On the other hand, the four estimates at the end of the sample show a wide range, with the lowest value close to 0.2.

We have considered only an AR(1) process by now. However, autocorrelation tests for the innovation of the Kalman-filter indicates the US and euro-area innovations are highly autocorrelated, clearly indicating the simple AR(1) model is not adequate (Hungarian innovations, on the other hand, are not autocorrelated). Thus we will continue with the study of higher order models.

5 Summary

This paper studied inflation persistence with time-varying-coefficient autoregressions for the US, euro-area, and Hungary. We first compared the statistical properties of the well known ML estimation using the Kalman-filter and the less known Flexible Least Squares estimator by Monte Carlo simulation. We showed that the FLS estimator does capture many types of smooth parameter changes and works definitely not worse than the more conventional Kalman-filter and smoother, and even somewhat better. We have also suggested a procedure to set the weighting parameter of the FLS based on an iterative Monte Carlo simulation technique calibrated to the time series in question.

The FLS estimator with $\mu = 10$ and 10^2 seem to yield plausible results for empirical inflation series: parameters of the autoregressive model of inflation in the cases of the US, the Euro-area, and Hungary are time-varying, and the parameter sequences are interpretable in economic terms. Inflation persistence tends to be higher in times of high inflation in the cases of all three countries. In the US and euro-area inflation persistence declined to historically low levels, while persistence in Hungary is still higher. As we argued, similar persistence would be an important structural similarity in a currency union; hence, there is still room for improvement for Hungary in the run up to the euro-area.

6 References

Dossche, M. – Everaert, G (2005): Measuring Inflation Persistence – A Structural Time Series Approach. ECB Working Paper No. 495 (Jun 2005) ECB, Frankfurt am Main

Hansen, Bruce E. (2001): The New Econometrics of Structural Change: Dating Changes in U.S. Labor Productivity, Journal of Economic Perspectives 15, 117-128.

Kalaba, Robert and Leigh Tesfatsion (1988): The Flexible Least Squares Approach to Time-varying Linear Regression. Journal of Economic Dynamics and Control, Vol. 12, 43-48.

Kalaba, Robert and Leigh Tesfatsion (1989): Time-Varying Linear Regression Via Flexible Least Squares, Computers and Mathematics with Applications, Vol. 17, 1215-1245.

Kalaba, Robert and Leigh Tesfatsion (1990): Flexible Least Squares for Approximately Linear Systems, IEEE Transactions on Systems, Man, and Cybernetics, Vol. 20, 978-989.

Kalaba, Robert and Leigh Tesfatsion (1996): A Multicriteria Approach to Model Specification and Estimation, Computational Statistics and Data Analysis, Vol. 21, 193-214.

Marques, C. R. (2004): Inflation Persistence: Facts Or Artifacts? ECB Working Paper No. 371 (Jun 2004) ECB, Frankfurt am Main.

O'Reilly, G. – Whelan, K. (2004): Has Euro-area Inflation Persistence Changed over Time? ECB Working Paper No. 335 (Apr 2004) ECB, Frankfurt am Main

Smets, Frank (2006): Inflation persistence and price setting behaviour in the euro area - A summary of the IPN Evidence, mimeo, ECB, Frankfurt am Main.

Tesfatsion, Leigh and John Veitch (1990): U.S. Money Demand Instability: A Flexible Least Squares Approach, Journal of Economic Dynamics and Control, Vol. 14, 151-173.

Constant Parameters

	OLS	FLS ($\mu = 10^1$)	FLS ($\mu = 10^2$)	FLS ($\mu = 10^3$)	FLS ($\mu = 10^4$)	Kalman-Filter	Kalman-Smoother
Intercept	*0.066*	*0.896*	*0.387*	*0.157*	*0.084*	*0.417*	*0.337*
AR(1)	*0.032*	*0.465*	*0.203*	*0.083*	*0.042*	*0.204*	*0.165*

Discrete Shift in Parameters

	OLS	FLS ($\mu = 10^1$)	FLS ($\mu = 10^2$)	FLS ($\mu = 10^3$)	FLS ($\mu = 10^4$)	Kalman-Filter	Kalman-Smoother
Intercept	*0.338*	*0.645*	*0.169*	*0.203*	*0.313*	*0.482*	*0.432*
AR(1)	*0.168*	*0.330*	*0.102*	*0.096*	*0.151*	*0.204*	*0.176*

Linear Deterministic Change in Parameters

	OLS	FLS ($\mu = 10^1$)	FLS ($\mu = 10^2$)	FLS ($\mu = 10^3$)	FLS ($\mu = 10^4$)	Kalman-Filter	Kalman-Smoother
Intercept	*0.254*	*0.298*	*0.157*	*0.216*	*0.248*	*0.192*	*0.187*
AR(1)	*0.252*	*0.301*	*0.158*	*0.210*	*0.243*	*0.186*	*0.185*

Sinusoid Deterministic Change in Autoregressive Parameter

	OLS	FLS ($\mu = 10^1$)	FLS ($\mu = 10^2$)	FLS ($\mu = 10^3$)	FLS ($\mu = 10^4$)	Kalman-Filter	Kalman-Smoother
Intercept	*0.457*	*0.539*	*0.113*	*0.251*	*0.410*	*0.179*	*0.118*
AR(1)	*0.432*	*0.286*	*0.074*	*0.248*	*0.390*	*0.136*	*0.093*

Random Walk Change in Parameters

	OLS	FLS ($\mu = 10^1$)	FLS ($\mu = 10^2$)	FLS ($\mu = 10^3$)	FLS ($\mu = 10^4$)	Kalman-Filter	Kalman-Smoother
Intercept	*0.416*	*0.206*	*0.197*	*0.329*	*0.398*	*0.217*	*0.188*
AR(1)	*0.346*	*0.195*	*0.150*	*0.265*	*0.329*	*0.182*	*0.156*

Table 1: Average RMSE of estimated parameter sequences in the Monte Carlo simulation

Notes: Number of Monte-Carlo draws is 1000 for each specification.

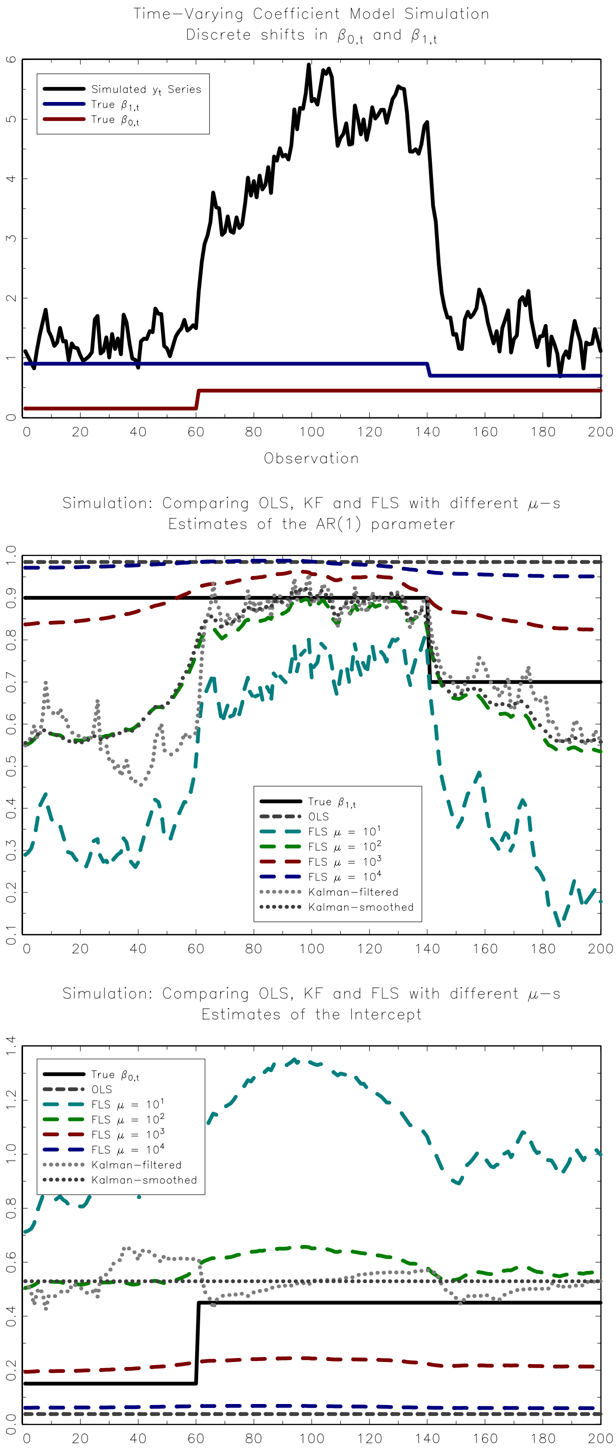

Figure 1: The ability of FLS and the Kalman-filter to capture discrete shift in parameters

<u>Top panel:</u> true processes for $\beta_{0,t}$ and $\beta_{1,t}$ and a certain realization for y_t.

<u>Middle panel:</u> the true process for $\beta_{1,t}$, its OLS estimate, and four different FLS estimates corresponding to weights: μ = 10, 10^2, 10^3, and 10^4, and the Kalman-filter and Kalman-smoother estimate.

<u>Bottom panel:</u> the true process for the constant, $\beta_{0,t}$, and its estimates.

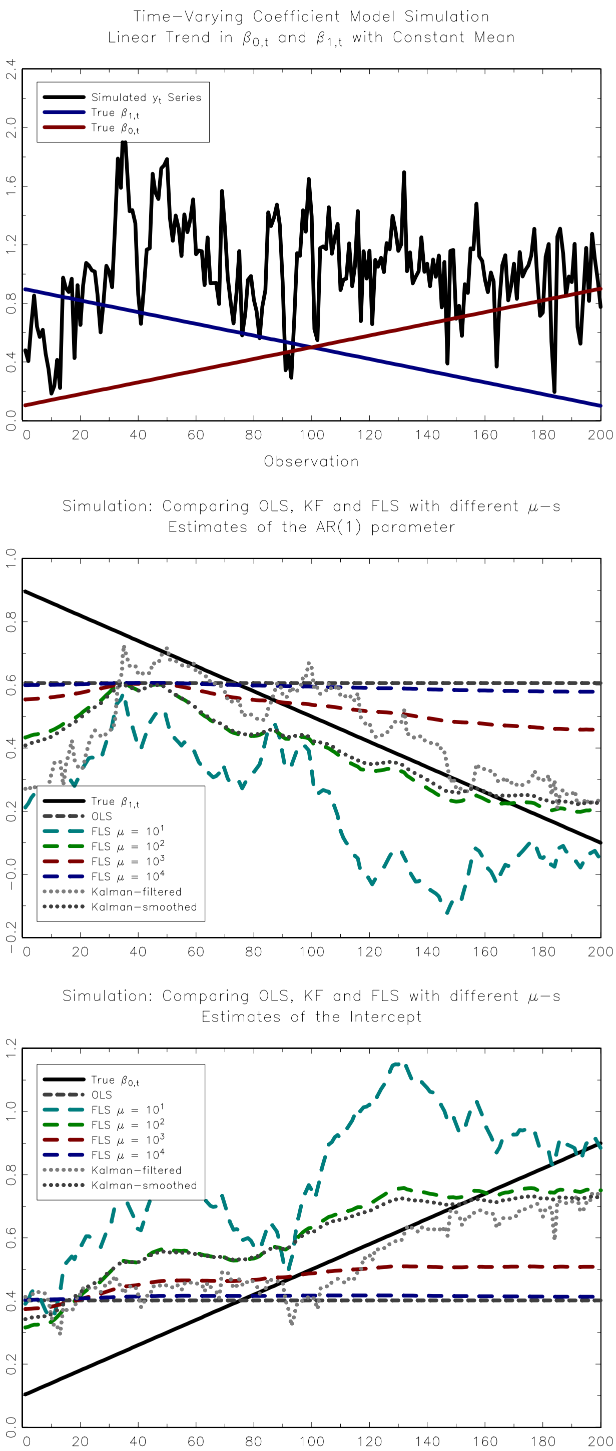

Figure 2: The ability of FLS and the Kalman-filter to capture continuous deterministic change in parameters

Top panel: true processes for $\beta_{0,t}$ and $\beta_{1,t}$ and a certain realization for y_t.

Middle panel: the true process for $\beta_{1,t}$, its OLS estimate, and four different FLS estimates corresponding to weights: μ = 10, 10^2, 10^3, and 10^4 and the Kalman-filter and Kalman-smoother estimate.

Bottom panel: the true process for the constant, $\beta_{0,t}$,and its corresponding estimates.

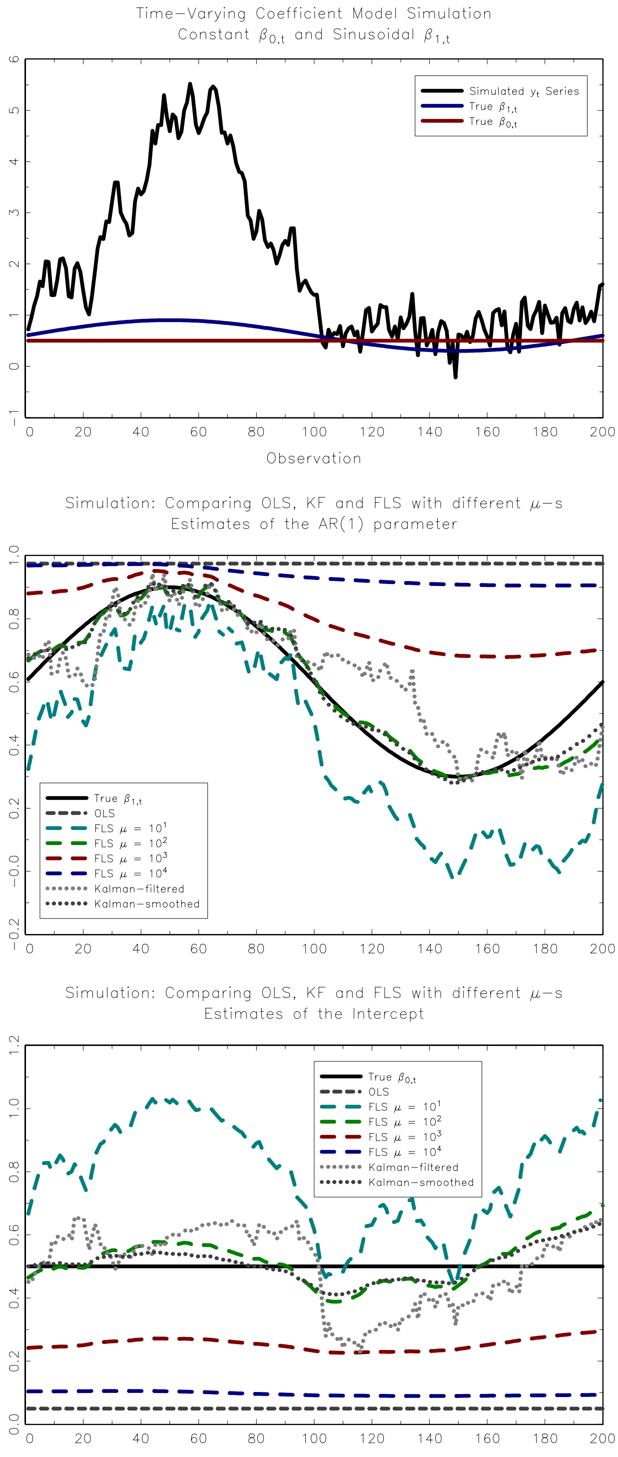

Figure 3: The ability of FLS and the Kalman-filter to capture sinusoid change in $\beta_{1,t}$

Top panel: true processes for $\beta_{0,t}$ and $\beta_{1,t}$ and a certain realization for y_t.

Middle panel: the true process for $\beta_{1,t}$, its OLS estimate, four different FLS estimates corresponding to weights: μ = 10, 10^2, 10^3, and 10^4, and the Kalman-filter and Kalman-smoother estimate.

Bottom panel: the true process for the constant, $\beta_{0,t}$, with its different estimates.

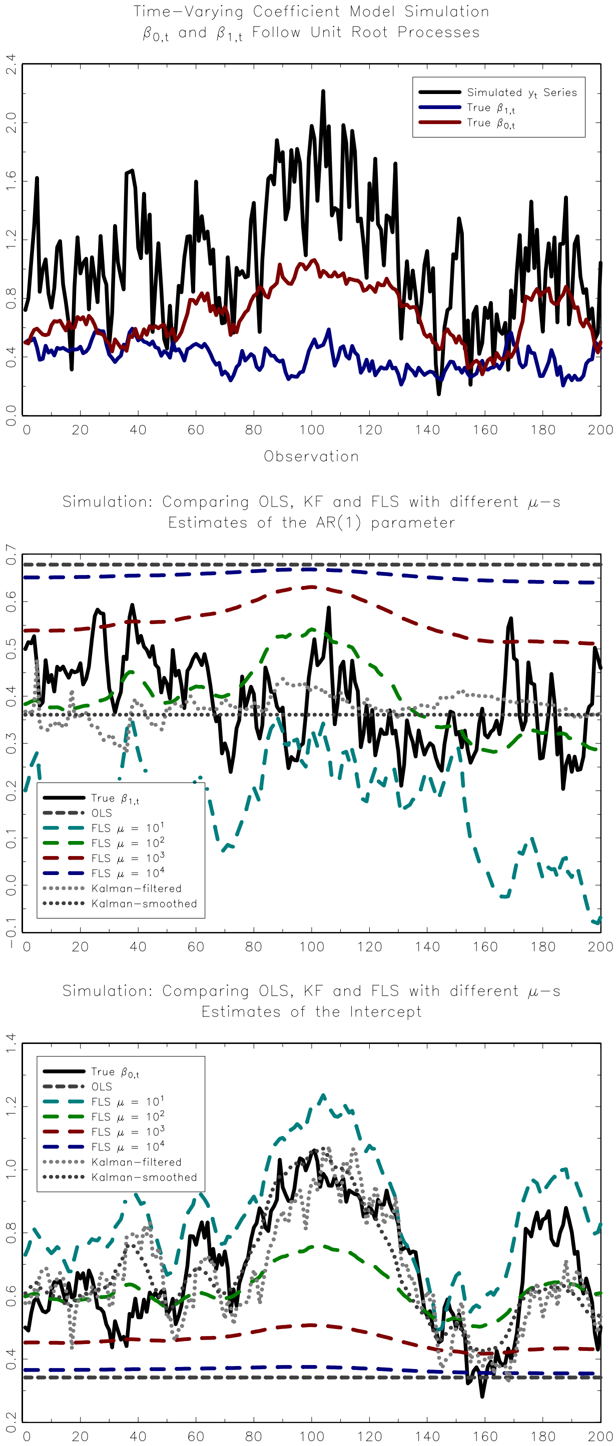

Figure 4: The ability of FLS and the Kalman-filter to capture random walk change in parameters

<u>*Top panel:*</u> true processes for $\beta_{0,t}$ and $\beta_{1,t}$ and a certain realization for y_t.

<u>*Middle panel:*</u> the true process for $\beta_{1,t}$, its OLS estimate, and four different FLS estimates corresponding to weights: μ = 10, 10^2, 10^3, and 10^4, and the Kalman-filter and Kalman-smoother estimate.

<u>*Bottom panel:*</u> the true process for the constant, $\beta_{0,t}$ and its estimates.

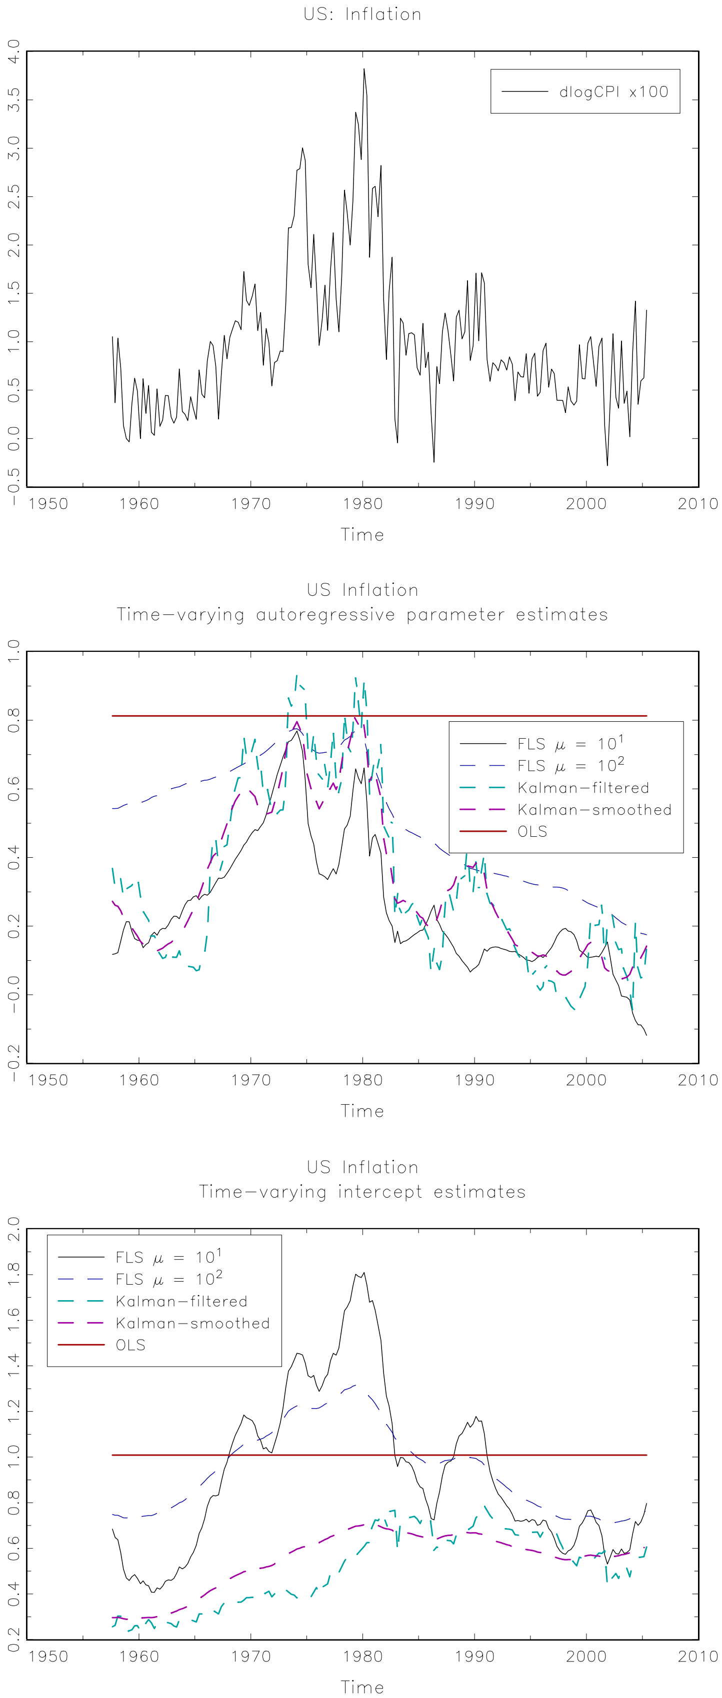

Figure 5: US inflation: Estimation results using OLS, FLS, and the Kalman-filter

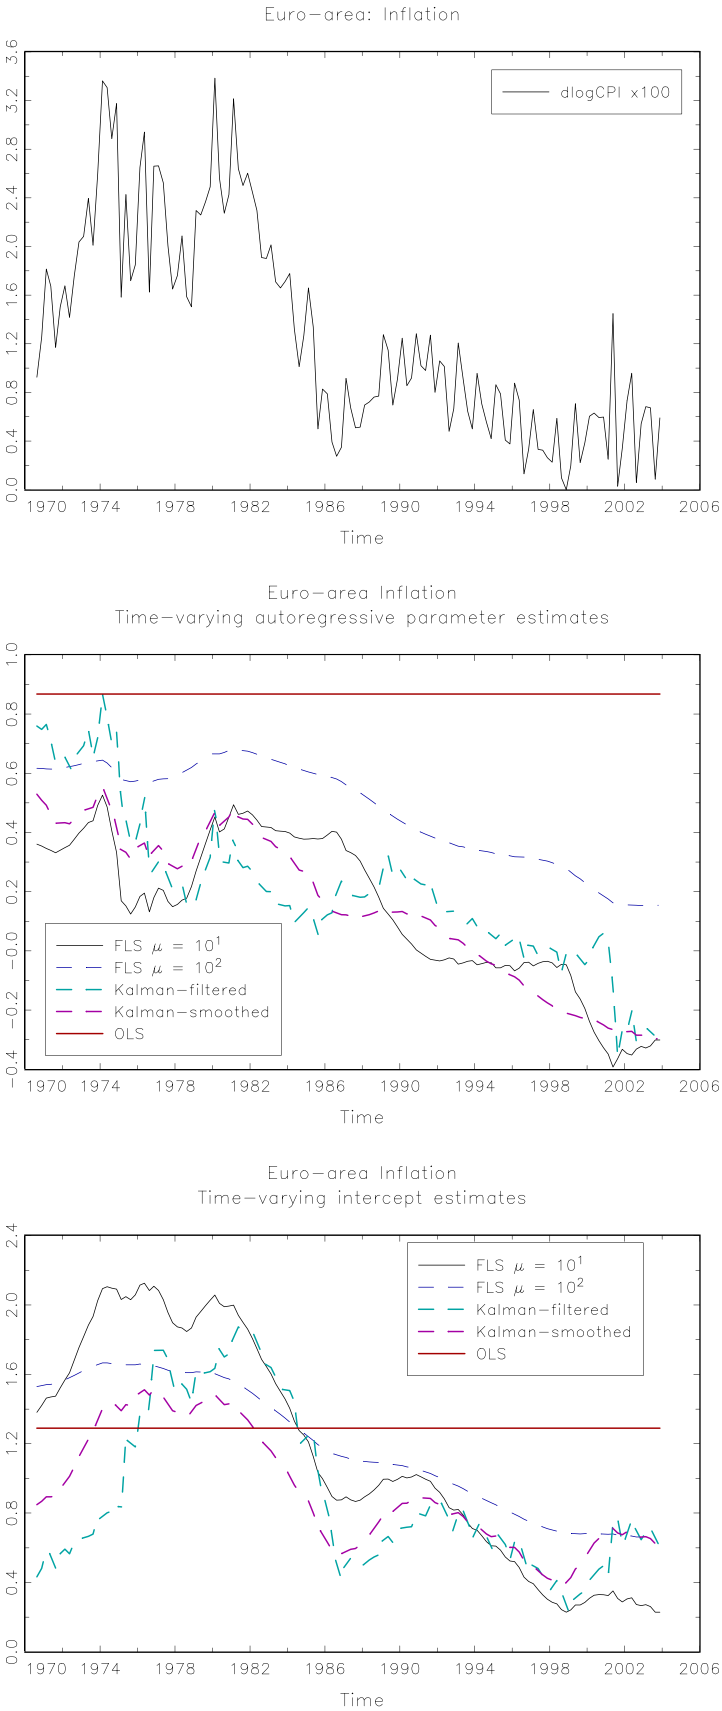

Figure 6: Euro-area inflation: Estimation results using OLS, FLS, and the Kalman-filter

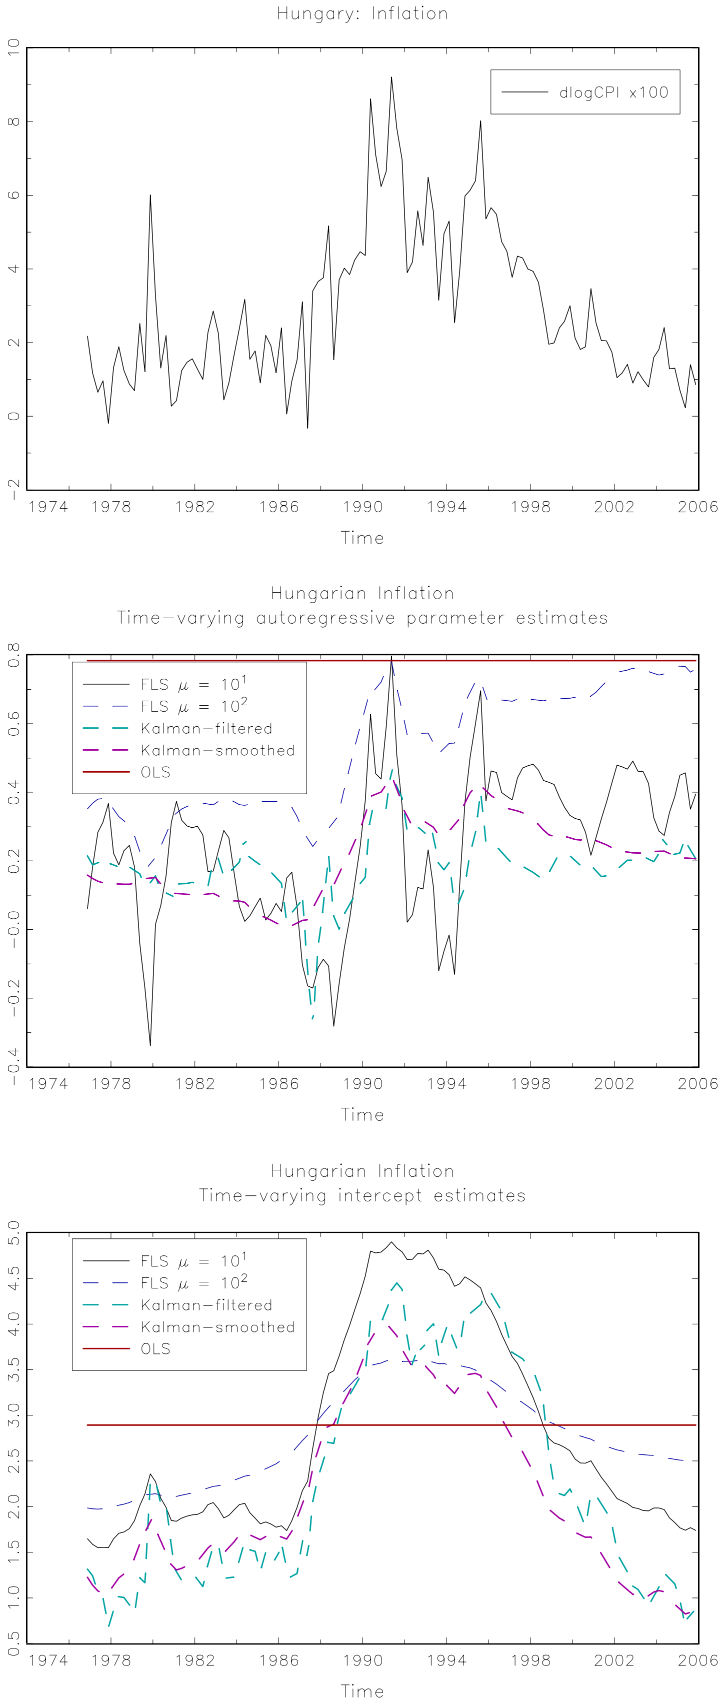

Figure 7: Hungarian inflation: Estimation results using OLS, FLS, and the Kalman-filter