

Time-Varying Persistence in US Inflation

Massimiliano Caporin · Rangan Gupta

Abstract The persistence property of inflation is an important issue not only for economists, but especially for central banks, given that the degree of inflation persistence determines the extent to which central banks can control inflation. Further, not only is it the level of inflation persistence that is important in economic analyses, but also the question of whether the persistence varies over time, for instance, across business cycle phases, is equally pertinent, since assuming constant persistence across states of the economy is sure to lead to misguided policy decisions. Against this backdrop, we extend the literature on long-memory models of inflation persistence for the US economy over the monthly period of 1920:1-2014:5, by developing an autoregressive fractionally integrated moving average-generalized autoregressive conditional heteroskedastic (ARFIMA-GARCH) model with a time-varying memory coefficient which varies across expansions and recessions. In sum, we find that inflation persistence does vary across recessions and expansions, with it being significantly higher in the former than in the latter. As an aside, we also show that persistence of inflation volatility is higher during expansions than in recessions. Understandably, our results have important policy implications. **Keywords:** Persistence; US Inflation Rate; Time-Varying Long Memory.

J.E.L. codes: C12, C13, C22, C51, E31, E52.

We would like to thank the Editor, Professor Robert M. Kunst, for giving us the opportunity to revise the paper, and two anonymous referees for their many helpful comments. However, any remaining errors are solely ours.

M. Caporin
University of Padova, Department of Statistical Sciences, Via Battisti 241, 35123 Padova, Italy, E-mail: massimiliano.caporin@unipd.it

R. Gupta
University of Pretoria, Department of Economics, Pretoria 0002, South Africa, E-mail: rangan.gupta@up.ac.za

1 Introduction

The persistence property of inflation is an important issue not only for economists, but especially for central banks, given that the degree of inflation persistence determines the extent to which central banks can control inflation. Understandably, the amount of research devoted (and still being carried out), given conflicting results¹ in analyzing the inflation persistence property for the US, as well as for other world economies, is voluminous to say the least.² Though various approaches³ have been used to analyze the degree of inflation persistence, autoregressive fractionally integrated moving average (ARFIMA) models is, perhaps, the most popular approach. This is simply because of the fact that the model nests the unit root and stationarity properties of the data, given its generalized form.

Not only is it the level of inflation persistence that is important in economic analyses, but also the question of whether the persistence varies over time, for instance, being contingent on the state of the economy, is equally pertinent. This is because, assuming constant persistence across business cycle phases is sure to lead to misguided policy decisions (as well as inaccurate forecasts). Against this backdrop, we extend the literature on long-memory models of inflation persistence for the US economy over the monthly period of 1920:1-2014:5, by developing an autoregressive fractionally integrated moving average-generalized autoregressive conditional heteroskedastic (ARFIMA-GARCH) model with a time-varying mem-

¹ In analyzing the issue of the degree of persistence of the shocks, a related controversy exists concerning the possible existence of a unit root in inflation. On one hand, Nelson and Schwert (1977), Barsky (1987), Ball and Cecchetti (1990), and Brunner and Hess (1993) provide evidence that U.S. inflation contains a unit root. On the other hand, Hassler and Wolters (1995), Baillie et al. (1996), Baum et al. (1999), Bos et al. (1999), Baillie et al. (2002), Hsu (2005), Lee (2005), Ajmi et al. (2008) and Hassler and Meller (2014) among others have found evidence that inflation is fractionally integrated, suggesting that the differencing parameter is significantly different from zero and unity.

² For a detailed survey in this regard, please refer to Balcilar, Gupta and Jooste (forthcoming) and Martins and Rodrigues (2014).

³ The econometric methods have covered various unit root tests, state-spaced-based time-varying) autoregressive models, and more recently, quantile regressions-based approaches. For a detailed literature review in this regard, refer to Tillmann and Wolters (2014) and Manzan and Zerom (forthcoming).

ory coefficient⁴ which varies across expansions and recessions.⁵ In sum, we find that inflation persistence does vary across recessions and expansions, with it being significantly higher in the former than in the latter. As an aside, we also show that persistence of inflation volatility, however, is higher during expansions than in recessions. The rest of the paper is organized as follows: section 2 discusses the data and the model, and section 3 is devoted to the results. Finally, section 4 concludes with some policy recommendations.

2 Model and data description

We compute US Inflation as a month-on-month percentage change in the US consumer price index (CPI) covering the monthly period of 1919:12-2014:5. Understandably, data on the inflation rate starts from 1920:1. The data on the CPI is sourced from the Global Financial Database. Note that, even though CPI data is available since 1876:1, the starting and end points of the sample coincide with the availability of reliable data at a monthly frequency at the time of writing this paper. To put it alternatively, we dropped the data before 1920, since the computed inflation rates over 1876:2 to 1919:12 had large amount of zeros and could possibly have made our results less reliable. The inflation rate data plotted over 1876:2 to 2014:5 vindicates our decision to start the analysis from 1920:1.

Given our objective of associating change in persistence with business cycle phases, we also recovered the timing of US recessions from the National Bureau of Economic Research (NBER). Table 1 reports the sequence of peaks and troughs on a quarterly basis. As our analyses consider monthly inflation time series, we set the end of recessions (and similarly of expansions) to the last month of the quarters

⁴ For a detailed survey on ARFIMA models with a time-varying long memory coefficient, the reader is referred to Boutahar et al. (2008) and Aloy et al. (2013).

⁵ Note that ever since the work of Granger and Hyung (2000), Diebold and Inoue (2001), and Mikosch and Starica (2004), it is well-known that spurious long-memory behavior can be detected in time series known to be theoretically short-memory, due to structural breaks (for a detailed literature review in this regard, refer to Tsay (2008), Tsay and Härdle (2008) and Hassler and Meller (2014). In our case, we deal with this issue by assuming that the break dates are known and that we have two regimes in expansions and recessions.




Fig. 1 Month-on-month US inflation from 1876:2 to 2014:5. Grey shades identify recession periods.

indicated in Table 1. Figure 1 plots the month-on-month inflation rate, along with the recession periods (highlighted by shaded areas).

To motivate the need for considering changes in inflation persistence across business cycle phases, we performed a preliminary analysis. In Table 2, we report a descriptive evaluation of the monthly inflation rate by comparing a few indicators, and conditioning on cycles. We note that the average inflation was lower during recessions, while the volatility was higher. We also found that inflation was skewed to the left during recessions, and to the right during expansions. This suggests that extreme values drove the mean and variance outcomes, given that during recessions we have instances of drops in month-on-month inflation, while during expansions we observe large inflation values. Notably, the inflation density is more leptokurtic during expansions, signalling a higher probability of observing extreme inflation movements.

| Peak | Trough |
|--------------------|--------------------|
| January 1920(I) | July 1921 (III) |
| May 1923(II) | July 1924 (III) |
| October 1926(III) | November 1927 (IV) |
| August 1929(III) | March 1933 (I) |
| May 1937(II) | June 1938 (II) |
| February 1945(I) | October 1945 (IV) |
| November 1948(IV) | October 1949 (IV) |
| July 1953(II) | May 1954 (II) |
| August 1957(III) | April 1958 (II) |
| April 1960(II) | February 1961 (I) |
| December 1969(IV) | November 1970 (IV) |
| November 1973(IV) | March 1975 (I) |
| January 1980(I) | July 1980 (III) |
| July 1981(III) | November 1982 (IV) |
| July 1990(III) | March 1991(I) |
| March 2001(I) | November 2001 (IV) |
| December 2007 (IV) | June 2009 (II) |

Table 1 Cyclical phases (quarter within parenthesis): source NBER - www.nber.org/cycles.html.

| | 1920:1-2014:5 | | |
|----------|---------------|--------|--------|
| | All | Exp | Rec |
| Mean | 0.002 | 0.003 | 0.000 |
| Median | 0.002 | 0.003 | 0.000 |
| St.dev | 0.006 | 0.005 | 0.008 |
| Min | -0.032 | -0.023 | -0.032 |
| Max | 0.057 | 0.057 | 0.030 |
| Q(1%) | -0.015 | -0.008 | -0.022 |
| Q(99%) | 0.018 | 0.019 | 0.015 |
| Skew | 0.308 | 1.931 | -0.604 |
| Kurt | 10.29 | 18.69 | 1.51 |
| IQ-range | 0.005 | 0.005 | 0.011 |
| N. zeros | 261 | 215 | 46 |
| N. obs. | 1133 | 244 | 891 |

Table 2 Descriptive analyses

If changes in persistence were present, one way of deducing this would be through the computation of the autocorrelation function conditional on the state of the economy (or state-dependent ACF). As in Caporin and Pres (2013), we considered the following estimators of the state-dependent autocorrelation function at lag k ,

$$\rho_R(k) = \frac{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t x_{t-k} S_t}{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t^2 S_t} \quad (1)$$

$$\rho_E(k) = \frac{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t x_{t-k} (1 - S_t)}{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t^2 (1 - S_t)} \quad (2)$$

where x_t is the month t inflation rate in deviation from the sample mean, R stands for recession and E for expansion, S_t is an indicator function assuming the value of 1 during recessions, T is the total sample size, $T_R(k)$ is the number of months associated with a recession ($\sum_{t=k+1}^T S_t = T_R(k)$), and $T_E(k)$ is the number of months where the economy is in expansion, $T_E(k) = T - T_R(k) - k$.

Note that the quantities in (1) and (2) represent only a preliminary check for the possible presence of changes in the persistence across states of the economy. Moreover, from a methodological point of view, we stress that, by construction, the presence of the indicator function S_t allows for computing the two serial dependence measures focusing on the time t business cycle phase. Therefore, if the lag k is not too large and the cycle phases are not too short, $\rho_R(k)$ and $\rho_E(k)$ will be based on different samples of data. On the contrary, for large values of k , it is possible that the two serial dependence measures are computed on data belonging to different economic states. Nevertheless, both $\rho_R(k)$ and $\rho_E(k)$ monitor the dependence on past data of the current observation and allow us to verify if a change in the state of the economy influences such a dependence. Therefore, the fact that past data belongs to different cycle phases might have had an effect on the evaluation of the correlations, but this did not prevent an interpretation of the state dependent autocorrelation measures. The two quantities in (1) and (2) will allow anyway to determine if past data become more or less important in affecting the current inflation level. The Appendix reports further details on the effect of overlapping regimes on the computation of the state dependent autocorrelation measures $\rho_R(k)$ and $\rho_E(k)$.

Figure 2 reports the autocorrelations computed with the standard estimator, and by conditioning on the state of the economy, with the sample period starting in






Fig. 2 Autocorrelation functions (ACF) for the sample from 1920:1 computed with the classic estimator and by conditioning on the cycle phase.

1920:1. Notably, we have evidence of changes in the persistence over the cyclical phases. The plots suggest that the inflation time series might be characterized by the presence of a long-memory behavior. However, a similar pattern could be generated

by a short-memory process with specific parameter values, such as those of an autoregressive model with coefficients close to the non-stationarity bound.

Further, the reported correlograms do not show clear evidence of seasonal patterns. However, seasonal variation is known to generally affect inflation as is shown in Ajmi et al. (2008) and references cited therein. Therefore, we decided to remove the seasonal pattern for our sample starting in 1920:1 with the modified X-13 method of Balcilar, Gupta and Jooste (forthcoming), and Balcilar, Gupta and Uwilingiye (forthcoming), which allowed us to deal with seasonality for data over 50 years of data.⁶

To shed some light on the possible presence of long-memory and, at the same time, allow for a change in the persistence across states of the economy, we modelled the CPI-based inflation rates as in Caporin and Pres (2013), with a Time Varying Auto Regressive Fractionally Integrated Moving Average (TV-ARFIMA) model. Notably, the model allowed for a change in the memory and, therefore, in the inflation persistence over time according to a time-function, such as the change in seasons, or conditionally upon a knowledge of changes in the business cycle phases. Let y_t be the inflation rate, which we modelled as follows:

$$(1 - L)^{d_t} \Phi(L) (y_t - \mu) = \Theta(L) \varepsilon_t \quad (3)$$

where d_t is a time varying memory coefficient, μ is the unconditional mean of the inflation, $\Phi(L)$ is an Auto Regressive (AR) polynomial, $\Theta(L)$ is a Moving Average (MA) polynomial and ε_t is the inflation shock which we assume to be identically and independently distributed with a zero mean and a standard deviation of σ_ε . The change in the memory coefficients has a step-wise evolution over time according to changes in the phases of the business cycle. If we denote R_t as the recession dummy, the memory coefficient assumed two possible values:

⁶ FORTRAN codes used here to conduct seasonal adjustment for long-span data is available upon request from the authors.

$$d_t = \begin{cases} d_R & R_t = 1 \\ d_E & R_t = 0 \end{cases} \quad (4)$$

where d_E and d_R denote the memory level during expansions and recessions, respectively.

In ARMA-type models, we measure the persistence of a series by means of the model parameters and through their role in the moving average infinite representation of the model. Larger persistence means that shocks do produce effects on the series realizations for longer time. In the TVARFIMA model, both long-memory and short-memory parameters affect persistence, as both play a role in the $MA(\infty)$ model representation. As a consequence, there will be two different degrees of persistence over recessions and expansions. However, when comparing persistence across the two states, given that the short-memory parameters are time-invariant, we can reason in a *ceteris paribus* way. Therefore, the occurrence of equal memory coefficients means equality of persistence across the two business cycle phases. Alternatively, a larger (smaller) memory coefficient in one cyclical phase means that this phase has larger (smaller) persistence.

Given that the change in persistence varies according to an exogenous variable, and is thus not exactly a function of time, the TVARFIMA model of Caporin and Pres (2013) becomes similar to the Threshold ARFIMA (TARFIMA) of Goldman et al. (2013). However, it differs from the TARFIMA model since the regime change is not associated with the level of the dependent variable (the CPI-based inflation rate), as happens in dynamic models with threshold-driven regimes. On the contrary, the change in the parameter structure depends on an exogenous variable, which in turn, is the dummy capturing recessions. The model can thus be compared with the approach of Haldrup and Nielsen (2006a,b), and could be considered to be equivalent to a Markov switching model where regimes are known and evolve in a fully independent way.

In order to capture the possible presence of heteroskedasticity, the model was extended by allowing the error variance to change over time according to the TVFI-GARCH model of Caporin and Pres (2009). Therefore, the variance of the inflation shock was denoted as σ_t^2 and evolved as:

$$\sigma_t^2 = \omega + \beta(L) \sigma_{t-1}^2 + \left[1 - \beta(L) - \Psi(L)(1-L)^{\lambda_t}\right] \varepsilon_t^2 \quad (5)$$

where $\beta(L)$ and $\Psi(L)$ are short-memory polynomials and λ_t is the variance memory parameter that changes over time according to the business cycle phases as follows:

$$\lambda_t = \begin{cases} \lambda_R & R_t = 1 \\ \lambda_E & R_t = 0 \end{cases} \quad (6)$$

where R_t is the recession dummy. Similarly to what was stated for the dynamics of the mean, the persistence of conditional variance depends on both short- and long-memory coefficients. However, short-memory parameters are time-invariant and thus, *ceteris paribus*, larger values of the long-memory coefficient are indicative of larger persistence.

Notably, in the dynamics of both the mean and the variance, the evolution of the persistence affects the current (time t) observation, irrespective of the fact that past data, on which the time t observation depends, belongs to different business cycle phases. What the model provides is a change in the dependence over past data, conditional on the fact that at time t we are in a given economic state. Therefore, depending on the state, the serial correlation might be stronger or weaker, and thus lead to a larger or smaller impact of past observations, irrespective of whether these are associated with an expansion or a recession.

Estimation of the models is performed by maximum likelihood methods. The full model can be estimated in a single step but, to simplify the computational complexity, a two-step approach is also feasible (at the cost of a loss in efficiency). In

the latter case, we first estimate the mean dynamics (the TV-ARFIMA part of the model) and, conditional on the first step estimation results, we recover mean residuals and estimate the variance dynamics (the TV-FIGARCH part of the model). Caporin and Pres (2013), by means of a Monte Carlo study, show that small sample performances of maximum likelihood estimators converge to the true parameter values with increasing sample sizes. Moreover, as expected, the estimator's density dispersion decreases with the sample size. The authors also provide evidence supporting using the model within a forecasting exercise conditional on a knowledge of the evolution of seasons. We stress that a theory for efficiency, consistency and asymptotic normality of the maximum likelihood estimator for the TV-ARFIMA and TV-FIGARCH models is still missing, as well as a theory for short-memory specifications with time-varying parameters. We point out that those models belong to the broad family of threshold models, among which we can distinguish between specifications where the coefficients change sign according to the modelled variable moving above and below a threshold (self-exciting specifications) and specifications where the change in regimes is due to an exogenous but observable variable, as in our case. While for the former the theory is available, for the latter it is not, as is pointed out by Tong (2011).

The model complexity also motivates focussing on the presence of time-variation only in the memory coefficient. From a theoretical point of view, all parameters could be allowed to change over seasons or business cycle phases. However, this increases the number of parameters to be estimated and the computational burden in the model's estimation. In this paper, we thus preferred to focus on specifications where only the long-memory coefficients change across business cycle phases.

3 Persistence in US inflation

We fitted the TVARFIMA-TVFIGARCH on the monthly inflation time series starting from 1920. In order to highlight the improvement of our proposed modelling strategy, we compared it with specifications where memory was either absent or

| | ARMA | | ARFIMA | | TVARFIMA | |
|---------------|--------|---------|--------|---------|----------|---------|
| | Coeff. | S.E. | Coeff. | S.E. | Coeff. | S.E. |
| d_R | | | 0.366 | 0.052 | 0.477 | 0.072 |
| d_E | | | | | 0.303 | 0.038 |
| ϕ_1 | 0.667 | 0.092 | | | | |
| ϕ_4 | 0.139 | 0.057 | 0.110 | 0.048 | 0.106 | 0.047 |
| ϕ_7 | 0.074 | 0.037 | | | | |
| θ_1 | 0.299 | 0.115 | | | | |
| θ_{12} | 0.147 | 0.064 | 0.190 | 0.061 | 0.189 | 0.062 |
| Llik | | 383.99 | | 389.42 | | 396.09 |
| AIC | | -755.97 | | -780.85 | | -782.18 |
| BIC | | -725.77 | | -750.72 | | -757.02 |
| Q(3) | | 0.63 | | 0.77 | | 0.44 |
| Q(12) | | 0.74 | | 0.24 | | 0.21 |
| Q(24) | | 0.00 | | 0.00 | | 0.01 |
| LR | | | | | | 0.00 |

Table 3 Estimation output for competing models on monthly inflation from 1920:1. S.E. denotes the standard error. R denotes the recession coefficients, and E , expansion coefficients. $Llik$ is the model's likelihood, while AIC and BIC are the Akaike and Bayesian information criteria, respectively. $Q(j)$ denotes the Ljung-Box test for the residual correlation at lag j , for which we report the p-value. LR corresponds to the p-value of the likelihood ratio test for the null of equal memory across the two business cycle phases.

was time invariant. Consequently, we estimated ARMA and ARFIMA specifications for the mean, while for the variance, we considered both GARCH and FIGARCH models. Given that we adopted a two-step estimation strategy, we report separate results for the mean and the variance dynamics. We estimated the competing variance models on the residuals of the most appropriate mean specification. Further, in order to confirm the need for time-variation in the memory coefficients, we report a likelihood ratio test between the time-varying and time-invariant long-memory specifications. Parameter symbols correspond with those adopted in the general model representations introduced in the previous section. Estimation tables report the parameters, the quasi maximum likelihood standard errors, as well as the model likelihood, the AIC and BIC information criteria, and the Ljung-Box test for the residual correlation of selected lags. To determine the appropriate specification of the short-memory orders, for both the mean and variance dynamic, we took into account both the significance of the estimated coefficients and the values of the information criteria.

| | GARCH | | FIGARCH | | TVFIGARCH | |
|-------------|--------|----------|---------|----------|-----------|----------|
| | Coeff. | S.E. | Coeff. | S.E. | Coeff. | S.E. |
| ω | 0.001 | 0.001 | 0.000 | 0.000 | 0.001 | 0.007 |
| λ_R | | | 0.410 | 0.049 | 0.299 | 0.037 |
| λ_E | | | | | 0.454 | 0.057 |
| α_1 | 0.149 | 0.055 | | | | |
| β_1 | 0.850 | 0.051 | 0.167 | 0.093 | 0.164 | 0.083 |
| Llik | | 681.59 | | 677.35 | | 682.32 |
| AIC | | -1357.19 | | -1348.70 | | -1358.64 |
| BIC | | -1342.09 | | -1333.60 | | -1343.54 |
| Q(3) | | 0.92 | | 0.90 | | 0.90 |
| Q(12) | | 0.99 | | 0.99 | | 0.99 |
| Q(24) | | 0.99 | | 0.99 | | 0.99 |
| LR | | | | | | 0.00 |

Table 4 Estimation output for competing models on the residuals of the TVARFIMA model for monthly inflation from 1920:1. S.E. denotes the standard error. R denotes the recession coefficients, and E , the expansion coefficients. $Llik$ is the model likelihood, while AIC and BIC are the Akaike and Bayesian information criteria, respectively. $Q(j)$ denotes the Ljung-Box test for residual correlation at lag j , for which we report the p-value. We performed the test on squared standardized residuals (mean residuals divided by the estimated conditional standard deviations) to test for the presence of residual heteroskedasticity. LR corresponds to the p-value of the likelihood ratio test for the null of equal memory across the two business cycle phases.

Before, we describe the details of the results regarding persistence, we observe the following characteristics of the models estimated: As can be seen from Table 3, residual correlation is found to be present for lags above 12, which could possibly be due to some element in the ACF being quite large. Generalizing the models with higher lags did not improve the results, since these additional coefficients turned out to be statistically insignificant, which in turn, was also vindicated by higher values of the AIC and the BIC. The ARMA specification was found to have a couple of economically significant lags, namely lag-lengths 4 and 7, while the ARFIMA and TVARFIMA were estimated with a lag-length of only 4. In addition, we note that the models required a lag 12 on the MA, which is possibly a by-product of the seasonal adjustment, or due to some stochastic seasonality. Importantly, the LR test was in favour of the TVARFIMA, which in turn, was also confirmed by both the AIC and the BIC. Turning now to the variance specifications in Table 4, the GARCH model is substantially equivalent to TVFIGARCH, both in terms of the standardized residual serial correlation, and in terms of the likelihood. Note that

the TVFIGARCH model is more flexible than the FIGARCH model, as it accounts for time-varying persistence, and a likelihood ratio test indicates a preference for the former. A similar test is not available for the comparison of TVFIGARCH against GARCH, with them being non-nested models. For the FIGARCH model, whose likelihood is inferior to both the GARCH and the TVFIGARCH, we observe that the parameters are significant, as in the other variance specifications.

Tables 3 and 4 report the estimated coefficients, and we stress that the memory parameters are always statistically significant. Our empirical results suggests that inflation persistence is changing across business cycle phases, and is higher during recessions. In fact, while the memory coefficient in the mean is equal to 0.477 for recessions, it decreases to 0.303 in periods of expansion. Conversely, the variance persistence moves from 0.299 in recessions up to 0.454 during expansions.⁷ We believe this is a relevant finding, and is supported by the overwhelming rejection of the Likelihood ratio test for equality of the memory coefficients across the business cycle phases. The null hypothesis is $d_E = d_R$, and the test statistic, assumed to be asymptotically distributed as a Chi-square with 1 degree of freedom, provides a p -value less than 0.001.

The observation that the persistence is higher in recessions than in expansions might be motivated by the fact that recessions are shorter and more clearly identifiable than expansionary phases. Supporting evidence is provided by the average duration of the business cycle phases: while for recessions we have an average duration of 14.3 months, for expansions the value increases to 52.3 months. Further, during expansions we might have oscillations in the growth cycle that could affect the persistence. In addition, another line of reasoning could be that while in our model we have two states, in reality we might also have stagnation, which in our analysis, is associated with expansion, and hence, could affect the persistence during expansion. As an aside, we also show that the persistence of inflation volatility is higher during expansions than during recessions. This is somewhat expected, as

⁷ These results were also qualitatively the same when we considered the longer sample starting in 1876:2. Complete details of these results are available upon request from the authors.

it implies that volatility persists on low regimes during expansions, while we have greater uncertainty and stronger reactions to shocks during recessions (less persistence induces a bigger reaction of volatility to innovations).

4 Concluding remarks

Given that the degree of inflation persistence determines the extent to which central banks can control inflation, the persistence property of inflation is an important issue. Also, not only is it the level of inflation persistence that is important in economic analyses, but also the question of whether the persistence varies over time, for instance, across business cycle phases, is equally pertinent. This is understandable since assuming constant persistence across expansionary and contractionary states of the economy is sure to lead to misguided policy decisions. Against this backdrop, we extend the literature on long-memory models of inflation persistence for the US economy over the monthly period of 1920:1-2014:5, by developing an autoregressive fractionally integrated moving-average-generalized autoregressive conditional heteroskedastic (ARFIMA-GARCH) model, with a time-varying memory coefficient that varies across expansions and recessions. In sum, we find that inflation persistence does vary across recessions and expansions, with it being significantly higher in the former (with the mean of inflation being relatively lower) than in the latter. As an aside, we also show that the persistence of inflation volatility, however, is higher during expansions (with the mean of volatility being relatively lower) than in recessions. Based on our findings, we can deduce the following in terms of monetary policy behaviour, conditioned on expansions or recessions, in the US historically: since higher persistence of inflation and its volatility associates itself with lower mean values of inflation and its volatility, monetary policy seems to be more proactive when the authorities perceive inflation and its volatility to be 'too high' on average. When inflation and its volatility are perceived to be low on average, monetary policy is aimed at keeping the levels where they are, implying higher persistence. In other words, the Federal Reserve seems to reduce the level of

inflation during expansions and the volatility of the same during recessions,⁸ while it tries to stabilize inflation during recessions and its volatility during expansions. Understandably, our results imply that the policy stance of the Federal Reserve to movements in inflation seems to be asymmetric, with it depending on not only its concern with the level or volatility of inflation, but also with whether the economy is in a recession or an expansion. It must be highlighted here that, besides price stability, the Federal Reserve is also concerned with the output gap. Given this, our discussion on its response to inflation is not foolproof, but with the analysis conditioned on business cycle phases, we accommodate for the mandate of maximum sustainable output and employment simultaneously in an implicit fashion.

References

Ajmi, A. N., Nasr, A. B., and Boutahar, M., 2008, Seasonal nonlinear long memory model for the US inflation rates, *Computational Economics*, 31(3), 243-254.

Aloy M., Tong C-L., Péguin-Feissolle A., and Dufrénot, G., 2013, A smooth transition long-memory model, *Studies in Nonlinear Dynamics & Econometrics*, 17(3), 281-296.

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

Baillie, R. T., Han, Y. W., and Kwon, T. G., 2002, Further long memory properties of inflationary shocks, *Southern Economic Journal*, 496-510.

Balcilar M., Gupta R., and Uwilingiye, J., Forthcoming. Dynamic Relationship between Oil Price and Inflation in South Africa, *Journal of Developing Areas*.

⁸ Given the so-called Cukierman and Meltzer (1986) effect, which implies that inflation volatility (uncertainty) leads to higher inflation, this result is perhaps an indication that the monetary authority does not intend to have higher inflation levels in the future when the economy recovers, to get into an expansion mode.

Balcilar M., Gupta R., and Jooste, C., Forthcoming. Analysing South Africa's inflation persistence using an ARFIMA model with Markov-switching fractional differencing parameter, *Journal of Developing Areas*.

Ball, L., Cecchetti, S. G., and Gordon, R. J., 1990, Inflation and uncertainty at short and long horizons, *Brookings Papers on Economic Activity*, 215-254.

Barsky, R. B., 1987, The Fisher hypothesis and the forecastability and persistence of inflation, *Journal of Monetary Economics*, 19(1), 3-24.

Baum, C. F., Barkoulas, J. T., and Caglayan, M., 1999, Persistence in international inflation rates, *Southern Economic Journal*, 900-913.

Bos, C., Franses, P. H., and Ooms, M., 1999, Re-analyzing inflation rates: Evidence of long memory and level shifts, *Empirical Economics*, 24, 427-449.

Boutahar, M., Dufrénot, G., and Péguin-Feissolle, A., 2008, A simple fractionally integrated model with a time-varying long memory parameter d_t , *Computational Economics*, 31(3), 225-241.

Brunner, A. D., and Hess, G. D., 1993, Are higher levels of inflation less predictable? A state-dependent conditional heteroscedasticity approach, *Journal of Business and Economic Statistics*, 11(2), 187-197.

Caporin M., and Pres, J., 2009. Forecasting temperature indices with time-varying long-memory models. Working paper 88-2009. Department of Economics and Management, University of Padua. Available: <http://www.decon.unipd.it/assets/pdf/wp/20090088.pdf>.

Caporin, M., and Pres, J., 2013, Forecasting Temperature Indices Density with Time-Varying Long-Memory Models, *Journal of Forecasting*, 32, 339-352.

Cukierman, A. and Meltzer, A., 1986. A theory of ambiguity, credibility, and inflation under discretion and asymmetric information, *Econometrica*, 54, 1099-1128.

Diebold, F.X and Inoue, A., 2001, Long memory and regime switching, *Journal of Econometrics*, 105, 131-159.

Goldman, E., Nam, J., Tsurumi, H., and Wang, J., 2013, Regimes and long memory in realized volatility, *Studies in Nonlinear Dynamic and Econometrics*, 17(5), 521-549.

Granger, C.W.J., and Hyung, N., 2000, Occasional Structural Breaks and Long Memory, revised version of UCSD Working Paper 99-14.

Haldrup, N., and Nielsen, M.O., 2006a, A regime switching long memory model for electricity prices, *Journal of Econometrics*, 135, 349-376.

Haldrup, N., and Nielsen, M.O., 2006b, Directional congestion and regime switching in a long memory model for electricity prices, *Studies in Nonlinear Dynamics and Econometrics*, 10, 1367.

Hassler, U., and Meller, B., 2014, Detecting multiple breaks in long memory the case of US inflation, *Empirical Economics*, 46(2), 653-680.

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

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

Lee, J., 2005, Estimating memory parameter in the US inflation rate, *Economics Letters*, 87, 207-210.

Manzan, S., and Zerom, D., forthcoming, Asymmetric quantile persistence and predictability: the case of U.S. inflation, *Oxford Bulletin of Economics and Statistics*.

Martins, L. F., and Rodrigues, P. M. M., 2014, Testing for persistence change in fractionally integrated models: An application to world inflation rates, *Computational Statistics and Data Analysis*, vol. 76, 502-522.

Mikosch, T., and Starica, C., 2004, Change of structure in financial time series, long range dependence and the GARCH models, *Econometrics*, EconWPA 0412004.

Nelson, C. R., and Schwert, G. W., 1977, Short-term interest rates as predictors of inflation: On testing the hypothesis that the real rate of interest is constant, *The American Economic Review*, 67(3), 478-486.

Tillmann, P., and Wolters, M. H., 2014, The changing dynamics of US inflation persistence: a quantile regression approach, Kiel Working Paper No. 1951.

Tong, H., 2011, Threshold models in time series analysis - 30 years on, *Statistics and its interface*, 4, 107-118.

Tsay, W. J., and Härdle, W. K., 2009, A generalized ARFIMA process with Markov-switching fractional differencing parameter, *Journal of Statistical Computation and Simulation*, 79(5), 731-745.

Tsay, W.J. 2008. Analysing inflation by the ARFIMA model with Markov-switching fractional differencing parameter, Mimeo, The Institute of Economics, Academia Sinica, Taiwan. Available at: <http://idv.sinica.edu.tw/wjtsay/pdf/example.pdf>.

A State-dependent Auto Correlations

We focus here on the computation of the two auto correlation measures $\rho_R(k)$ and $\rho_E(k)$, whose expression we report below, to clarify the following discussion:

$$\rho_0(k) = \frac{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t x_{t-k} S_t}{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t^2 S_t},$$

$$\rho_1(k) = \frac{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t x_{t-k} (1 - S_t)}{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t^2 (1 - S_t)}.$$

Note that we re-labelled states as 0 and 1 to be more general. We postulated that the zero-mean series x_t 's dynamic behaviour changes over time and moves between two states or regimes. The variable S_t monitors the states evolution, and takes value 0 in the first state and value 1 in the second state. We refer to the previous expressions as the estimators of the State-Dependent Auto Correlation Functions.

This Appendix discusses two aspects. Firstly, the appropriateness of the state-dependent Auto Correlation Functions to highlight differences in the serial dependence across states. Secondly, we note that the estimation of the state-dependent Auto Correlation Functions following the above reported estimators might suffer for distortions due to the fact that observations x_t and x_{t-k} might belong to two different states. We evaluated the impact of those distortions.

We tackled the two points by means of simulations, computing the average, across simulations, or the empirical ACF obtained from simulated series. We believe that the derivation of the theoretical expressions of the state-dependent ACF is beyond the scopes of the present paper and postpone that to future contributions.

We simulated data from the following TVARFIMA stochastic process

$$(1 - L)^{d_t} x_t = \varepsilon_t$$

with ε_t distributed as a standardized Normal, and

$$d_t = \begin{cases} d_0 & S_t = 0 \\ d_1 & S_t = 1 \end{cases}.$$

Moreover, we assumed that the variable S_t moves from one state to the other every M observations. Further, we noted that if $d_0 = d_1$ the model collapsed to an ARFIMA(0,d,0).

We first show that the ACF of the TVARFIMA computed on the full sample is a mixture of the ACF of the TVARFIMA computed on the two states. We thus estimated and averaged the ACFs of the TVARFIMA process over $N = 500$ simulated series, characterized by a change of state every $M = 50$ observations. We assumed that the timing of the state change is known. For those series




Fig. A.1 Simulated empirical ACF from TVARFIMA with two states, with memory coefficients $d = 0.45$ and $d = 0.1$, respectively.




Fig. A.2 Simulated empirical ACF from TVARFIMA with two states, with memory coefficients $d = 0.4$ and $d = 0.2$, respectively.

we considered a sample size of $T = 1000$ observations. Note that in the simulation we adopted a pre-sample of T observations to control for the dependence on initial values. We considered two cases: $d_0 = 0.45$ and $d_1 = 0.1$; $d_0 = 0.4$ and $d_1 = 0.2$. We provided two plots, Figure A.1 and A.2, where we report the empirical ACFs (averaged across the N simulations) for the full sample as well as those computed on each state according to the above-reported estimators.

From the two figures, it emerged clearly how the full sample ACF lies in the middle between the ACFs computed on the two states. The ACFs get closer when the persistence of the two states is closer, as we might have expected. Consequently, identifying long-range dependence by means




Fig. A.3 Simulated state-dependent empirical ACF from TVARFIMA with two states, with memory coefficients $d = 0.45$ and $d = 0.1$, respectively.

of a full-sample ACF might not have allowed us to detect the presence of two states characterized by very different memory degrees, as Figure A.1 shows.

We now move to the evaluation of the effect associated with the introduction in the computation of state-dependent ACF of observations belonging to different regimes. For that purpose we compared the estimators we proposed with the alternative ones reported below:

$$\delta_0(k) = \frac{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t x_{t-k} S_t S_{t-k}}{\frac{1}{T_R(k)-k} \sum_{t=k+1}^T x_t^2 S_t S_{t-k}},$$

$$\delta_1(k) = \frac{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t x_{t-k} (1 - S_t) (1 - S_{t-k})}{\frac{1}{T_E(k)-k} \sum_{t=k+1}^T x_t^2 (1 - S_t) (1 - S_{t-k})}.$$

These two competing estimators differ from the one we previously adopted in the sense that the pairs of observations $x_t x_{t-k}$ contribute to the evaluation of the ACF only if they both belong to the same state. We provide then simulated empirical ACF (averaging across simulations) using the same data generating process previously adopted. The following two figures, Figure A.3 and A.4, focus on the two different generators.

By analyzing the figures, we note that the state-dependent ACFs are contaminated by the presence of observations belonging to different states. In fact, the ACF for the higher memory states is downward biased, being contaminated by observations coming from a state with lower memory. We observe the opposite effect on the ACF of the lower memory state. This effect is, obviously, more evident if we have states with closer memory degree, or if one of the two states has very low memory.




Fig. A.4 Simulated state-dependent empirical ACF from TVARFIMA with two states, with memory coefficients $d = 0.4$ and $d = 0.2$, respectively.

While on the one side this evidence could suggest the use of δ_0 and δ_1 in place of ρ_0 and ρ_1 , we must stress that the evaluation of δ_0 and δ_1 is limited by the size of the observed states. In fact, if states have, in a limiting case, an equal size of M observations, the maximum lag for which we can compute δ_0 and δ_1 is exactly M , but at the cost of sensibly reducing the accuracy of the estimators for lags close to M , due to the limited number of pairs $x_t x_{t-k}$ contributing to their evaluation. In practice, we might compute the estimators up to a fraction of M .

Consequently, the use of ρ_0 and ρ_1 is advisable, as they could allow for evaluating the state dependent ACFs, even if the duration of states is limited. We must, however, advise the user on the presence of the distortions due to the inclusion, in the estimator evaluation, of observations belonging to different states. Nevertheless, as the state-dependent ACF should be adopted only within an identification step, and thus should be preliminary to any estimation, we do believe that the presence of distortions will have a limited impact. Those ACF would just suggest a possible presence of states with differing memory degrees, and only a model estimation would allow making a final decision.