Asymmetric quantile persistence and predictability: the case of U.S. inflation*

SEBASTIANO MANZAN[†] and DAWIT ZEROM[‡]

[†] *Department of Economics & Finance, Baruch College, CUNY*

e-mail: sebastiano.manzan@baruch.cuny.edu

[‡] *California State University-Fullerton*

e-mail: dzerom@fullerton.edu

Abstract

This paper investigates the evidence of time-variation and asymmetry in the persistence of U.S. inflation. We compare the out-of-sample performance of different forecasting models and find that quantile forecasts from an Auto-Regressive (AR) model with level-dependent volatility are at least as accurate as the forecasts of the Quantile Auto-Regressive model, in particular for the core inflation measures. Our results indicate that the persistence of core inflation has been relatively constant and high, but it declined for the headline inflation measures. We also find that the asymmetric persistence of inflation shocks can be mostly attributed to the positive relation between inflation level and its volatility.

JEL Classification: C22, C53, E31, E52

Keywords: Conditional quantiles; Distribution; Inflation persistence; Forecasting;

*We would like to thank the Editor and two referees for many insightful comments. Financial support by the PSC-CUNY Research Award Program and the Bert W. Wasserman fund is gratefully acknowledged. The usual disclaimers apply.

I. Introduction

There is a vast literature that investigates the, possibly changing, time series persistence of inflation with unsettled conclusions. Fuhrer (2010) is a recent overview that discusses the empirical and theoretical aspects of inflation persistence. On the empirical side, Fuhrer suggests that the evidence seems to point to the conclusion that U.S. inflation persistence has declined starting from the second half of the 1980s, in particular when considering the Consumer Price Index (CPI) and the Personal Consumption Expenditure (PCE) price index, although there is less convincing evidence when analyzing the GDP deflator and core measures of inflation. A standard approach to evaluate the persistence of inflation are unit root tests that typically conclude that inflation is non-stationarity when using samples up to 1984, but stationary when only the post-1984 period is considered. The decline in inflation persistence after the mid-1980s has been established also by Cogley and Sargent (2002) using a Bayesian VAR model with time-varying parameters. They show that persistence was high during the 1970s and the first half of the 1980s and low in the 1960s and since the mid-1980s, which also suggests the existence of a relationship between the persistence and the level of inflation. On the other hand, Stock (2002) and Pivetta and Reis (2007) argue that inflation persistence has been mostly high and constant over time and conjecture that findings of its decline might be the spurious result of neglecting to account for the fall in volatility that occurred after the mid-1980s. In addition, Tsong and Lee (2011) and Tillmann and Wolters (2012) apply the Quantile Auto-Regressive (QAR) model proposed by Koenker and Xiao (2004, 2006) to U.S. and international inflation rates at the quarterly frequency. They interpret the evidence of asymmetric persistence as suggesting that positive shocks (high conditional quantiles) to inflation have a permanent effect on inflation while negative ones (low conditional quantiles) have a quick mean-reverting effect. They also find that their results are, to a large extent, robust to considering breaks in the series which might spuriously lead to the conclusion that persistence varies over time.

In this paper we aim at re-examining the evidence on U.S. inflation persistence, in particular its variation over time and the notion that it might be asymmetric across the inflation distribution. A possible explanation for the contradictory results in the literature is that inflation persistence and volatility seem to be positively related to the level of inflation so that misspecification in the dynamics of one component leads to the conclusion of time variation in the other component. In order to distinguish between these alternative explanations we perform an out-of-sample test of the predictive accuracy of a semi-parametric specification, the QAR model, to that of an Auto-Regressive (AR) time series model with constant persistence and conditional standard deviation of the errors depending on the level of inflation (see Evans, 1991 and Brunner and Hess, 1996, among others). This approach can be interpreted as an out-of-sample specification test in which the performance of a parametric model is compared to that of a non-parametric one, with the performance being measured by their forecast accuracy. In particular, we compare forecasts of quantiles, rather than the conditional mean, given our interest in evaluating the persistence of inflation along its distribution, and in explaining its determinants. We thus interpret findings of equal or higher accuracy of the parametric forecasts (relative to the QAR forecasts) as supporting the heteroskedastic AR model for inflation dynamics. On the other hand, if the parametric forecasts are less accurate compared to the QAR ones, we conclude that the former is inadequate, either because of misspecification of the volatility equation or because persistence might be time-varying and/or asymmetric. In addition, we also consider a QAR specification in which we impose a unit root at all quantiles and we refer to this model as Quantile Unit Root (QUR). Comparing the accuracy of the QUR quantile forecasts to those of the QAR and heteroskedastic AR models provides an evaluation of the loss of (out-of-sample) predictability that derives from assuming non-stationarity at specific parts of the inflation distribution.

We forecast core and headline CPI and PCE inflation at the monthly frequency from January 1985 until June 2011. The evidence from the out-of-sample comparison indicates that the quantile forecasts of the heteroskedastic AR model are equally or more accurate relative to those from the QAR which confirms its validity as a model for inflation dynamics. The rolling estimates of inflation persistence in the heteroskedastic AR model are found to be fairly stable around 0.9 in the out-of-sample period for core CPI and PCE, which agrees with the findings of Stock (2002) and Pivetta and Reis (2007) that accounting for the volatility of inflation leads to persistence estimates that are high and relatively constant. However, for the headline measures we find that inflation persistence experiences a decline to 0.2-0.3 in the mid-2000s which can be attributed to the effect of excluding the high-inflation period of the 1970s from the estimation sample. For these series, we find that the persistence of the series seems to have declined in recent years in addition to the evidence of changing inflation volatility. Furthermore, the estimate of the parameter of lagged inflation in the volatility equation is stable and significantly positive which reinforces the view that inflation uncertainty increases with its level. If this is case, the asymmetric pattern of the quantile persistence coefficients can be attributed to the heteroskedasticity of inflation, as opposed to alternative explanations such as the asymmetric effect of shocks (as in Tsong and Lee, 2011, and Tillmann and Wolters, 2012). Another result is that forecasting the inflation quantiles using the QUR model that assumes a unit root at all quantiles delivers less accurate forecasts relative to the QAR and heteroskedastic AR models, in particular for quantiles below the median. This indicates that modeling the first differences of inflation, instead of its level, leads to a significant loss of predictive power when the interest is in forecasting the inflation distribution.

The paper is organized as follows. In Section II. we discuss the QAR model and in Section III. we provide the in-sample estimation results and the non-stationarity tests in a quantile framework. We then evaluate the out-of-sample performance of the quantile forecasts from the QAR, QUR, and heteroskedastic AR model in Section IV. and provide a detailed discussion of the results in Section VI.. Finally, Section VII. summarizes the findings of the paper.

II. A quantile model of inflation

We define the annualized inflation rate as $\pi_t = 1200 \log(P_t/P_{t-1})$, with P_t denoting the price index in month t. Several approaches have been considered in order to measure the persistence of inflation. For the purpose of this paper, we focus only on one of these measures which defines persistence, denoted by ρ, as $\rho \equiv \sum_{j=1}^{p} \alpha_j$, with the α_j's representing the coefficients of an AR(p) model for inflation, i.e.

$$\pi_t = \mu + \sum_{j=1}^{p} \alpha_j \pi_{t-j} + \varepsilon_t \equiv \mu + \rho \pi_{t-1} + \sum_{j=1}^{q} \beta_j \Delta \pi_{t-j} + \varepsilon_t \tag{1}$$

where ϵ_t is an *i.i.d.* random variable with mean zero and variance σ^2, and the β_j's are linear combinations of the α_j's and $q = p - 1$. Large values of ρ (with $|\rho| < 1$) correspond to high persistence of the series in the sense that the cumulative effect of a shock to π_t is given by $1/(1-\rho)$. In the limiting case that $\rho = 1$ the inflation process π_t contains a unit root and hence it is considered non-stationary. This measure of persistence focuses on the characteristics of the inflation dynamics at the center of the distribution, and thus neglects the possibility that the dynamics might be different in other parts of the distribution of π_t. To investigate this issue, Koenker and Xiao (2004, 2006) extend the constant (homogeneous) persistence model in (1) into a Quantile Auto-Regressive - QAR(q) - model given by

$$Q_{\pi_t}(\tau|\mathcal{F}_{t-1}) = \mu(\tau) + \rho(\tau)\pi_{t-1} + \sum_{j=1}^{q} \beta_j(\tau)\Delta\pi_{t-j}. \tag{2}$$

where $\tau \in (0,1)$, $Q_{\pi_t}(\tau|\mathcal{F}_{t-1})$ represents the quantile of π_t conditional on information available at time $t-1$, $\mathcal{F}_{t-1}$, and the parameters $\mu(\tau)$, $\rho(\tau)$ and $\beta_j(\tau)$ are allowed to vary across different quantiles τ. We refer to $\rho(\tau)$ as the quantile persistence (at level τ) of inflation since it generalizes the persistence parameter in the conditional mean model in Equation 1. Koenker and Xiao apply the QAR model to several economic variables and find an asymmetric pattern in the $\hat{\rho}(\tau)$, with large persistence estimates at high quantiles and significantly smaller at low quantiles. Similar results are also found for U.S. and international inflation measures by Tsong and Lee (2011) and Tillmann and Wolters (2012). These findings are interpreted as evidence that shocks to inflation have an asymmetric effect at different parts of its distribution, so that positive (negative) shocks contribute to increase (decrease) the persistence of inflation. Koenker and Xiao (2006) argue that the QAR can be considered as a general (approximate) specification for parametric models that allow for asymmetric persistence of shocks in economic variables (e.g., Beaudry and Koop, 1993).

The heterogeneity of the $\rho(\tau)$ coefficients could also be explained by a model in which the volatility of shocks to inflation increases with the level of inflation. This hypothesis has received considerable attention in the past by, among others, Evans (1991), Evans and Wachtel (1993), and Brunner and Hess (1996). We can relax the hypothesis of homoskedasticity of the error term in the AR specification in Equation 1 by assuming that it has a time-varying standard deviation, denoted by σ_t, that is,

$$\pi_t = \mu + \rho\pi_{t-1} + \sum_{j=1}^{q} \beta_j \Delta\pi_{t-j} + \sigma_t \varepsilon_t, \quad \varepsilon_t \sim N(0,1) \tag{3}$$

where $|\rho| < 1$ and the conditional standard deviation is given by $\sigma_t = c_0 + c_1\pi_{t-1} + \sum_{k=1}^{r} c_{k+1}\Delta\pi_{t-k}$. This assumption implies that the volatility of inflation is a linear function of the past levels of inflation and the existing evidence suggests that this relationship is likely to be positive, so that we expect the coefficient c_1 to be positive. The conditional quantiles of inflation implied by the heteroskedastic AR model in Equation 3 are given by

$$\begin{aligned} Q_{\pi_t}(\tau|\mathcal{F}_{t-1}) &= [\mu + c_0\Phi^{-1}(\tau)] + [\rho + c_1\Phi^{-1}(\tau)]\pi_{t-1} + \sum_{j=1}^{k}[\beta_j + c_{j+1}\Phi^{-1}(\tau)]\Delta\pi_{t-j} \\ &= \mu(\tau) + \rho(\tau)\pi_{t-1} + \sum_{j=1}^{k}\beta_j(\tau)\Delta\pi_{t-j}. \end{aligned} \tag{4}$$

where $k = max(q, r)$ and $\Phi^{-1}(\tau)$ denotes the τ-th quantile of ε_t. This model, which we denote by AR-HET, can thus be interpreted as a restricted version of the QAR specification with the quantile coefficients of π_{t-1} given by $\rho(\tau) = \rho + c_1\Phi^{-1}(\tau)$. For positive values of the parameter c_1 the model produces heterogeneous values for the persistence parameter across τ, similarly to the asymmetric persistence explanation. While at the median the quantile persistence parameter $\rho(\tau)$ equals ρ, it is larger than ρ for $\tau > 0.5$ since $\Phi^{-1}(\tau)$ takes positive values (and $c_1 > 0$). On the other hand, for quantiles below the median $\Phi^{-1}(\tau)$ is negative such that the quantile persistence is lower than ρ. Therefore, the AR-HET model implies conditional quantiles that display heterogeneous values for $\rho(\tau)$, although this heterogeneity is entirely driven by the relationship between inflation volatility and the level of inflation, rather than being the outcome of the asymmetric effect of positive and negative shocks.

We thus have two possible explanations for inflation dynamics that are consistent with the QAR specification and produce the type of heterogeneity in the persistence parameter that is observed in the data. While the AR-HET model implies a constant persistence ρ and a positive value for c_1, the hypothesis of asymmetric persistence of shocks can be interpreted, in terms of the AR-HET specification, as time-variation of the coefficient ρ with the level of inflation and c_1 equal to zero. Combining the analysis of the QAR and AR-HET models allows us to disentangle if the channel of asymmetric quantile persistence operates mostly through ρ or c_1, or quite possibly both. The controversy on the source of inflation persistence is a vivid one, with some papers arguing for its decrease starting from the mid-1980s (see Levin and Piger, 2006, Benati, 2008, Cogley and Sargent, 2002, and Cogley *et al.*, 2010 among others), while other papers contending that persistence has remained constant and high (see Stock, 2002, and Pivetta and Reis, 2007). In particular, the latter papers suggest that findings of a decrease in persistence might be the spurious outcome of neglecting to account for the fall in volatility, so that changes in volatility can be mistaken for changes in persistence.

It is important to be able to discriminate between these alternative explanations, in particular when the interest is to provide a framework for the conduct of monetary policy. In the next Section we review the in-sample results on the estimation of the QAR model on several headline and core inflation measures at the monthly frequency and provide results about their stationarity from a quantile perspective. We then try to distinguish between the asymmetric persistence and level-dependent volatility hypotheses by performing an out-of-sample comparison of the accuracy of quantile forecasts from the QAR model in Equation 2 and the heteroskedastic AR model in Equation 3.

III. The in-sample evidence

We estimate the QAR model in Equation 2 using (seasonally adjusted) monthly price indexes from January 1959 to June 2011 obtained from the Federal Reserve Bank of Saint Louis (FRED) database for the following four measures of inflation: the Consumer Price Index for all items (CPI), CPI excluding food and energy (CPILFE), Personal Consumption Expenditure deflator (PCE) and PCE excluding food and energy (PCELFE).

We evaluate the hypothesis of non-stationarity of the inflation rate using the tests proposed by Koenker and Xiao (2004), which assume that the quantile coefficients $\rho(\tau)$ measure the local persistence of inflation. The null hypothesis of unit-root behavior at any fixed τ for $\tau \in \mathcal{T}$, i.e. $H_0 : \rho(\tau) = 1$, can be tested using the following two statistics:

$$t_n(\tau) = \frac{f_n(F_n^{-1}(\tau))}{\sqrt{\tau(1-\tau)}}(Y_{-1}^T P_X Y_{-1})^{1/2}(\hat{\rho}(\tau) - 1) \quad \text{and} \quad U_n(\tau) = n(\hat{\rho}(\tau) - 1) \tag{5}$$

where Y_{-1} is the vector of lagged dependent variables (π_{t-1}), P_X is the projection matrix onto the space orthogonal to $X = (1, \Delta\pi_{t-1}, \cdots, \Delta\pi_{t-q})$. A consistent estimator of $f(F^{-1}(\tau))$ is given by $f_n(F_n^{-1}(\tau)) = 2h_n / \left[\bar{x}^T(\hat{\theta}(\tau + h_n) - \hat{\theta}(\tau - h_n))\right]$, where $\bar{x}$ is a vector of averages of $x_t = (1, \pi_{t-1}, \Delta\pi_{t-1}, \cdots, \Delta\pi_{t-q})^T$, the parameter vector $\hat{\theta}(\tau) = (\hat{\mu}(\tau), \hat{\rho}(\tau), \hat{\beta}_1(\tau), \cdots, \hat{\beta}_q(\tau))$ of the QAR(q) model is obtained from the usual quantile regression and h_n is a bandwidth which is set equal to $h_n = n^{-1/5}[4.5\phi^4\left(\Phi^{-1}(\tau)\right)/(2(\Phi^{-1}(\tau))^2 + 1)^2]^{1/5}$, where $\Phi(\cdot)$ represents the CDF of the standard normal distribution and $\phi(\cdot)$ indicates the standard normal density function.

Koenker and Xiao (2004) also propose two tests that aim at evaluating the overall non-stationarity of the series over a range of quantiles $\tau \in \mathcal{T}$, rather than at a specific quantile level. The tests for the null hypothesis that $\alpha(\tau) = 1 \; \forall \; \tau \in \mathcal{T}$ are Kolmogorov-Smirnov (KS) type statistics and are given by $\text{QKS}_t = \sup_{\tau \in \mathcal{T}} |t_n(\tau)|$ and $\text{QKS}_\alpha = \sup_{\tau \in \mathcal{T}} |U_n(\tau)|$, where $t_n(\tau)$ and $U_n(\tau)$ are given above. Rejections of the null hypothesis of unit-root for π_t means that π_t is not a constant (homogeneous) unit root process. This can happen even when the conventional Augmented Dickey-Fuller (ADF) test fails to reject unit root. When the KS tests reject the null hypothesis (implying non-constant unit root for π_t), the estimated $\rho(\tau)$ values together with their respective individual tests ($t_n(\tau)$ and $U_n(\tau)$) can be informative in uncovering possible asymmetry in the persistence behavior of inflation. To derive critical values for the above tests, we implement the bootstrap approach of Koenker and Xiao (2004) which consists of generating π_t^* under the null hypothesis of a unit root to ensure the non-stationarity of the bootstrap sample.

Estimates of the persistence term $\hat{\rho}(\tau)$ for several $\tau \in [0.1, 0.9]$ and the tests in (5) are provided at the top of Table 1. The results indicate that, for all inflation measures and using both QKS_t and QKS_α, the null hypothesis that $\rho(\tau) = 1 \; \forall \; \tau \in [0.1, 0.9]$ is rejected at the 1% level. This suggests that inflation does not exhibit a constant (or homogeneous) unit root behavior along its distribution, and the same conclusion can also be reached by examining the individual $\hat{\rho}(\tau)$ values at various quantiles τ. The estimates of $\rho(\tau)$ follow an asymmetric pattern with values increasing toward one with the quantile τ. For example, the coefficient estimate is as small as 0.643 for $\tau = 0.1$ (CPI) and as large as 1.13 for $\tau = 0.9$ (CPILFE). One interpretation of these results is that inflation is stationary at low quantiles since both $t_n(\tau)$ and $U_n(\tau)$ tests reject the null of unit root at those τ values in favor of the alternative $\rho(\tau) < 1$. On the other hand, the coefficient estimates at higher quantiles (above the median) are close to, or even above, unity thus implying that inflation displays unit root or even explosive behavior at higher quantiles. At those quantiles far above the median, $t_n(\tau)$ and $U_n(\tau)$ either fail to reject the unit-root hypothesis $\rho(\tau) = 1$ or if they reject the unit-root hypothesis, it is in favor of the alternative $\rho(\tau) > 1$ - suggesting explosive behavior.

The Table also reports the OLS or mean-based estimate of persistence $\hat{\rho}$ corresponding to the AR model in (1). Except for PCELFE, the null hypothesis of unit root is rejected in favor of $\rho < 1$, at 5% level for CPI and PCE and at 10% level for CPILFE using asymptotic critical values for the ADF_t statistics. A qualitatively similar conclusion is also suggested when using a confidence interval approach using the grid bootstrap of Hansen (1999). Hence, the mean-based persistence approach largely suggests global stationarity of inflation (with the exception of PCELFE). This conclusion implicitly means that inflation is overall stationary across its distribution.

Another interesting hypothesis to evaluate is whether the conditional distributions of the inflation measures have experienced structural breaks which might have occurred at the center of the distribution and/or on the tails. Oka and Qu (2011) recently proposed tests for structural parameter change for quantile regression models that allow to evaluate the hypothesis of parameter constancy at the individual quantile level or of the whole conditional distribution. These tests are thus more informative relative to break tests that focus on the conditional mean since they also provide information about possible breaks in the tails of the conditional distribution and their possible asymmetry. We implemented the $SQ_{\tau}(l+1|l)$ test proposed by Oka and Qu (2011) for $l = 0, 1,$ and 2 and for $\tau = 0.1, 0.3, 0.5, 0.7$ and 0.9 and Table 1 provides the break dates for the four inflation measures using a 5% significance level. The results indicate only one break at most quantiles which occurs at high quantiles for PCE in 1980 and for core CPI in 1980 and 1984. For core PCE the break occurs in 1975-1976 at the two extreme quantiles, while for CPI we find that the tails experienced a break at the beginning of the 2000s whilst the center of the distribution in 1980. The quantile break test thus indicates that the conditional distribution might have experienced significant changes which mostly happened at the outer quantiles, in particular on the right tail of the distribution, and can be associated with changes in the volatility of inflation, rather than its mean. The only inflation measures for which we find evidence of a break at the center of the distribution is CPI in April 1980.

Finally, we also estimate the AR-HET model in Equation 3 by ML under the assumption of normality of the errors on the full sample from 1959 to 2011. We select the lag orders in the mean and standard deviation equations by BIC criteria. Table 1 provides the coefficient estimates for π_t in the mean equation ($\hat{\rho}$) and in the volatility equation ($\hat{c}_1$). The estimation results indicate that persistence is high for all inflation measures with values ranging from 0.89 to 0.956, as well as the parameter of π_t in the volatility equation which are significant at the 5% level, except for CPI inflation. The fact that the coefficient estimates of c_1 are positive supports the hypothesis of a positive dependence between inflation volatility and its level.

IV. The out-of-sample evidence

The results of the previous Section provide evidence of the asymmetric pattern of the $\rho(\tau)$ estimates for both headline and core U.S. inflation measures at the monthly frequency. This pattern can arise when shocks to inflation have an asymmetric effect above and below the median but also when inflation volatility increases with its level. In addition, the presence of structural breaks might partly explain the findings of asymmetric persistence and level-dependence of inflation volatility. In order to distinguish between these different explanations, we evaluate and compare the (out-of-sample) quantile forecasts of the semi-parametric QAR model relative to those produced by the heteroskedastic AR model. We thus take the practical perspective of a forecaster who uses predictive models consistent with each hypothesis with the aim of forecasting the quantiles of the inflation distribution. The forecaster has also to decide on the sample period to use in estimation. We consider the cases of a long and short rolling window that allow to evaluate the potential effect of breaks on the forecasting accuracy since the short window is able to adjust more rapidly to changes of the inflation dynamics relative to the long window. In addition, we also consider a third specification which we refer to as Quantile Unit Root (QUR) model since we impose on the QAR specification the hypothesis that $\rho(\tau) = 1$ at all quantiles. The reason for including this model in the analysis is that it allows the evaluation of the loss of forecasting accuracy (compared to the QAR model) that derives from the assumption of a unit root at all quantile levels.

Models

In out-of-sample forecasting the horizons of interest are typically longer compared to the one-month horizon which was considered in the in-sample analysis of the previous Section. We thus consider two forecasting horizon, denoted by h, equal to 1 and 12 months ahead, respectively. For the multi-period ahead forecasts the inflation rate is defined as $\pi_{t+h}^h = (1200/h)\,[\ln(P_{t+h}) - \ln(P_t)]$, where P_t denotes a price index in month t. We begin the out-of-sample forecast exercise in January 1985 and end in June 2011 (318 monthly forecasts) and, as discussed above, estimate the model parameters on a long and a short rolling window of 300 and 120 months, respectively. The AR lags q and r are selected by the BIC criterion. To summarize, the models considered to produce quantile forecasts for π_{t+h}^h are:

- **QAR**: the Quantile Auto-Regressive (QAR) model is given by $Q_{\pi_{t+h}^h}(\tau|\mathcal{F}_t) = \mu(\tau) + \rho(\tau)\pi_t + \sum_{j=1}^{q} \beta_j(\tau)\Delta\pi_{t+1-j}$. We estimate the model on rolling windows of size 120 and 300 months which we denote by Short Rolling Window (SRW) and Long Rolling Window (LRW), respectively.

- **QUR**: the Quantile Unit Root (QUR) model is a restricted version of the QAR model in which we assume that $\rho(\tau) = 1\ \forall\tau$. We can then reformulate the QAR model in terms of changes of inflation and estimate the following quantile model: $Q_{\Delta\pi_{t+h}^h}(\tau|\mathcal{F}_t) = \mu(\tau)+\sum_{j=1}^{q} \beta_j(\tau)\Delta\pi_{t+1-j}$, where $\Delta\pi_{t+h}^h = \pi_{t+h} - \pi_t$. The quantile forecasts for the level of inflation are given by $Q_{\pi^h_{t+h}}(\tau|\mathcal{F}_t) = \pi_t + Q_{\Delta\pi^h_{t+h}}(\tau|\mathcal{F}_t)$. We estimate the model on a long rolling window and denote it by QUR LRW. From the comparison of the quantile forecasts for QUR LRW and QAR LRW we expect to reject the hypothesis of equal accuracy if the assumption of a unit root is unwarranted (at a specific quantile level) thus leading to inferior forecast accuracy. As discussed in the previous Section, we expect QUR and QAR to provide similar performance at high quantiles since the $\hat{\rho}(\tau)$ are close to one, but QAR is likely to outperform QUR at low quantiles given that the estimates of the quantile persistence is significantly lower than one.

- **AR-HET**: the Heteroskedastic AR model is given by $\pi^h_{t+h} = \mu + \rho\pi_t + \sum_{j=1}^{q}\beta_j\Delta\pi_{t+1-j} + \sigma_{t+h}\varepsilon_{t+h}$, where the volatility is given by $\sigma_{t+h} = c_0 + c_1\pi_t + \sum_{k=1}^{r} c_{k+1}\Delta\pi_{t+1-k}$. The model is estimated by gaussian ML with σ_{t+h} constrained to be positive. The conditional quantile of inflation is then constructed as $Q_{\pi^h_{t+h}}(\tau|\mathcal{F}_t) = \hat{\mu} + \hat{\rho}\pi_t + \sum_{j=1}^{q}\hat{\beta}_j\Delta\pi_{t+1-j} + \hat{\sigma}_{t+h}\hat{\Phi}^{-1}(\tau)$, where $\hat{\Phi}^{-1}(\tau)$ denotes the τ-th quantile of the EDF of the residuals $\hat{\epsilon}_{t+h}$ and $\hat{\sigma}_{t+h} = \hat{c}_0 + \hat{c}_1\pi_t + \sum_{k=1}^{r}\hat{c}_{k+1}\Delta\pi_{t+1-k}$. Also in this case we estimate the model both on a Long and Short Rolling Window and denote the model as AR-HET LRW and AR-HET SRW, respectively.

In addition, we aim to compare the quantile forecasts and we discuss next the statistical approach that we follow in this paper.

Evaluation

There are several methods proposed in the literature to evaluate density and distribution forecasts that mostly differ in the score (loss) function they assume to evaluate the relation between the forecast and the (future) realization of the variable. Given the scope of this paper, a natural choice of score function is the *Quantile Score* (QS) proposed by Gneiting and Raftery (2007) which is specifically aimed at the evaluation of quantile forecasts. Denote by $QS^i_{t+h}(\tau)$ the Quantile Score statistic that evaluates the τ-level quantile forecast of model i ($i =$QUR LRW, QAR LRW, QAR SRW, and AR-HET) in month t and for horizon h, which is given by:

$$QS^i_{t+h}(\tau) = \left[\pi^h_{t+h} - Q^i_{\pi^h_{t+h}}(\tau|\mathcal{F}_t)\right]\left[\mathcal{I}\left(\pi^h_{t+h} \leq Q^i_{\pi^h_{t+h}}(\tau|\mathcal{F}_t)\right) - \tau\right] \tag{6}$$

where $\mathcal{I}(\cdot)$ represents the indicator function, $Q^i_{\pi^h_{t+h}}(\tau|\mathcal{F}_t)$ is the τ-level h-period ahead quantile forecast of model i based on information available at time t, and π^h_{t+h} is the realization of h-period inflation at time $t+h$. This score function represents the check (tick) function that is employed for quantile regression estimation (see Koenker and Bassett, 1978). $QS^i_{t+h}(\tau)$ is interpreted as a loss and it can be used to compare the performance of two competing models. Given the quantile forecasts for inflation at time $t+h$ of model i and j, we conclude that model j is more accurate than model i if $QS^i_{t+h}(\tau) > QS^j_{t+h}(\tau)$, and vice-versa if i is more accurate than j.

To evaluate the statistical significance of any difference in performance between two models, as measured by the score function discussed above, we follow the approach of Giacomini and White (2006) and Amisano and Giacomini (2007). A test statistic for the null hypothesis of equal forecast accuracy at a given quantile τ, $QS^i_{t+h}(\tau) = QS^j_{t+h}(\tau)$ (for $t = 1, \cdots, P$, where P is the total number of forecasts), is given by

$$t = \left(\overline{QS}^j_h(\tau) - \overline{QS}^i_h(\tau)\right)/\sigma$$

where $\overline{QS}^i_h(\tau)$ and $\overline{QS}^j_h(\tau)$ denote the sample average of the quantile scores of model i and j, and σ represents the HAC standard error of the difference in scores. The test statistic t is asymptotically standard normal and rejections for negative values indicate that model j significantly outperforms model i (and vice-versa for positive values).

V. Results

The results of the QS test are reported in Table 2 for $h = 1$ and 3 for $h = 12$. The entries represent the test statistic for the null hypothesis of equal accuracy of the quantile forecasts produced by the null model (first column of the Table) relative to the alternative model (second column) which is standard normally distributed. We consider as benchmarks the QAR LRW and the AR-HET LRW models and we include as competing models the QUR LRW, QAR SRW, and AR-HET SRW. A negative value of the test statistic indicates that the forecasts of the alternative model are more accurate compared to those produced by the benchmark, and the opposite when the statistic is positive. The discussion is organized based on the pairwise model comparisons that address the issue of the non-stationarity of inflation, the possibility of a structural break, and the ability of the heteroskedastic AR model to explain the asymmetric persistence pattern.

QAR LRW *vs* QUR LRW

The first question relates to the cost, in terms of forecasting accuracy, of assuming a unit root at all quantiles of the QAR model. We investigate this issue by comparing the accuracy of the QAR LRW forecasts to those produced by the QUR LRW model. The results in Table 2 for $h = 1$ suggest that, for all inflation measures, the QS test statistics are large and positive for quantiles below the median and mostly not significantly different from zero at high quantiles. Hence, the evidence indicates that the QAR LRW benchmark forecasts provides more accurate out-of-sample forecasts at low quantiles relative to the competing QUR LRW forecasts, a finding which is consistent with the in-sample evidence. This suggests that assuming inflation is non-stationary, and thus differencing the series, might lead to a loss of forecast accuracy relative to the QAR model, although this mostly happens at low quantiles. These results hold also when the forecast horizon is equal to 12 months as reported in Table 3.

QAR LRW *vs* QAR SRW

Several papers have argued that inflation has experienced a structural break in terms of a mean shift (see Levin and Piger, 2006) or in terms of its persistence (see Cogley and Sargent, 2002) with the break typically dated in the mid to end of the 1980s. If a break indeed occurred, we would expect the QAR forecasts produced with a short rolling window to outperform the long rolling ones since the short estimation window adapts faster to the parameter change. The comparison of the performance of the QAR LRW and QAR SRW indicate that the shorter estimation window outperforms the long one for CPI (for $h = 12$) and core CPI (for both $h = 1$ and 12) for $\tau \geq 0.8$, but have similar accuracies in forecasting the central and lower quantiles. This suggests that a short rolling window delivers higher predictive accuracy when the interest is to forecast the CPI measures, in particular in the top part of the distribution which is consistent with the in-sample quantile break test results. However, for PCE and core PCE inflation we find that both estimation windows are, overall, equally accurate at both forecast horizons which indicates that the in-sample breaks found at high quantiles did not affect negatively the forecast performance.

QAR LRW *vs* AR-HET LRW and SRW

Comparing the QAR LRW and AR-HET quantile forecasts contributes to shed light on the ability of the heteroskedastic AR model to explain the observed pattern of the persistence coefficient $\rho(\tau)$. At $h = 1$ we find that the parametric model estimated on a short rolling window outperforms the quantile model at high quantiles for core CPI, while AR-HET LRW does significantly worse when forecasting CPI. For the PCE measures the two models provide equally accurate forecasts. Instead, when considering the annual horizon we find that the parametric model outperforms the quantile model forecasts at high quantiles, in particular when using a long estimation window (except for CPI). This suggests that the heteroskedastic AR model provides quantile forecasts that are equally and, in some part of the distribution, more accurate relative to QAR and can thus be considered a valid explanation for the quantile persistence pattern discussed in the previous Section.

AR-HET LRW *vs* QAR-SRW

The last comparison we are interested in evaluating is between the parametric model for inflation and the QAR model estimated on a short rolling window. For both models we find evidence that they outperformed the QAR LRW, at least for some of the inflation measures, and it is thus interesting to compare the (relative) accuracy of their forecasts. The results in Table 2 show that for $h = 1$ QAR SRW outperforms AR-HET at low quantiles for CPI and at high quantiles for core CPI, but they are equally accurate for the PCE measures. Instead, the results in Table 3 for

$h = 12$ indicate that the two quantile forecasting models perform equally, except at the center of the conditional distribution for core PCE in which AR-HET outperforms QAR SRW. Overall, this shows that the parametric model that assumes that inflation volatility depends on its level provides quantile forecasts that are as accurate as the QAR forecasts obtained on a short rolling window. Hence, at the one-year horizon it seems that even for the CPI measures for which we found the short window outperforms the long window, this better performance is explained by the parametric model thus indicating that the mechanism explaining inflation dynamics might be the level-dependent volatility rather than the existence of breaks in its distribution.

AR-HET LRW *vs* AR-HET SRW

The short estimation window provides more accurate forecasts relative to the LRW at the one-month horizon when forecasting Core CPI at high quantiles and at low quantiles for CPI. However, for $h = 12$ the SRW is outperformed by the LRW for several inflation measures and mostly at low quantiles.

VI. Discussion

The findings discussed in the previous Section indicate that assuming a positive relationship between inflation uncertainty and the past level of inflation delivers out-of-sample forecasts that perform similarly, for most measures, relative to the QAR model. This suggests that the hypothesis of level-dependent inflation volatility should be considered as a potential explanation for the heterogeneous persistence of inflation at different quantile levels, along with the possibility that inflation shocks have an asymmetric effect on its dynamics. In this Section we discuss in more detail some characteristics of the AR-HET and QAR models and evaluate whether there is evidence of time variation in the relative forecast performance of the two models.

Figure 1 represents a time series plot of the estimates for the persistence coefficient, ρ, and the coefficient of the inflation level in the volatility equation, c_1, in the AR-HET model for $h = 1$. The first observation in the graphs represents the parameter estimate produced at the end of January 1984 (to forecast January 1985) and the subsequent values have been estimated on a rolling window of 300 months (approximately 25 years). The plots for the core inflation measures show that the rolling estimates of the persistence parameter ρ (continuous line) oscillate closely around the full-sample estimates (dashed lines) and in both cases these estimates are high and close to 1. However, for the headline measures the evidence suggests that the persistence estimates are high and relatively constant up to the early 2000s and then tend to decrease to stabilize around a level between 0.2 and 0.4. Since we are using a 25 years rolling window, the estimates produced starting from the early 2000s are obtained on a rolling window that progressively excludes the high-inflation period of the 1970s. This contributes to lower the persistence of inflation for the headline measures, but remarkably does not have any effect on the persistence estimate of the core measures. A possible explanation for this difference is the role of the energy price shocks of the 1970s that affected more decisively the headline measures rather than the core ones. The evidence of a break in persistence across inflation measures is also consistent, at least partly, with results available in the literature which are based on various inflation measures (e.g., the GDP deflator or the CPI Index) and different frequencies (typically quarterly rather than monthly as in this study). In terms of the dependence of the past inflation level in the volatility equation, we find that it oscillates around the full-sample estimate with no clear tendency to decrease or increase during the sample. This suggests the dependence of inflation uncertainty on its level is an assumption which seems robust over time.

In order to evaluate the ability of the heteroskedastic AR model to account for the heterogeneous pattern of $\hat{\rho}(\tau)$, we estimate the following quantile regression model at each forecast date:

$$Q_{\tilde{\pi}^h_{t+h}(\tau)}(\tau|\mathcal{F}_t) = \tilde{\mu}(\tau) + \tilde{\rho}(\tau)\pi_t + \sum_{j=1}^{J} \tilde{\beta}_j \Delta\pi_{t+1-j} \tag{7}$$

where $\tilde{\pi}^h_{t+h}(\tau) = \pi^h_{t+h} - \hat{\rho}^{\text{AR-HET}}(\tau)\pi_t$ and $\hat{\rho}^{\text{AR-HET}}(\tau) = \hat{\rho} + \hat{c}_1\hat{\Phi}^{-1}(\tau)$ represents the quantile persistence of inflation implied by the AR-HET model (see Equation 3) based on the estimates of ρ and c_1. $\hat{\Phi}^{-1}(\tau)$ denotes the τ-level empirical quantiles of the residuals. The variable $\tilde{\pi}^h_{t+h}(\tau)$ can be interpreted as the h-month annualized inflation rate filtered by the persistence component that can be attributed to the AR-HET model and measured by $\hat{\rho}^{\text{AR-HET}}(\tau)\pi_t$. We thus expect the estimates of $\tilde{\rho}(\tau)$ in Equation 7 to be close to zero at all quantiles if the AR-HET model is the correct specification for the inflation dynamics. Alternatively, a positive value of $\tilde{\rho}(\tau)$ indicates that the AR-HET model underestimates the QAR coefficient of π_t at quantile τ, and the opposite for negative values. Figure 2 provides the box-plots of the 318 estimates of $\tilde{\rho}(\tau)$ as a function of τ for the four inflation measures and $h = 1$. The results show a similar pattern across measures in which the estimates fluctuate close to zero in the central part of the distribution (for τ between 0.3 and 0.7), but depart from zero at the tails with the deviations, in most cases, smaller than 0.3 in magnitude. The mostly negative values of the coefficient at low quantiles indicates that the parametric model predicts higher quantile persistence relative to the estimates from the quantile regression model. On the other hand, we find mostly positive values at high quantiles which indicate that AR-HET underestimates (local) inflation persistence. While these deviations do not affect significantly the relative forecast accuracy at low quantiles, they seem to have an impact on the forecast precision of AR-HET at the top quantiles.

Fluctuation Test

An additional test that can be conducted to evaluate and compare the forecasts is to assess whether the relative performance of two models has changed over time, as opposed to the evaluation discussed in Section 4 which is based on the full out-of-sample period (see Table 3). This could be useful in the context of this paper since, as suggested by several papers, inflation persistence might have declined after 1984 and thus could have changed the relative forecast accuracy of the models. The evidence from the quantile break test in Table 1 seems to support the occurrence of breaks, in particular at the top of the distribution. We thus implement a so-called fluctuation version of the QS test as proposed by Giacomini and Rossi (2010) which allows us to assess the evidence of changes in the relative performance of the QAR and AR-HET forecasts at a given quantile. The fluctuation QS test that compares the forecasts of model i and j at quantile level τ in month t and for window size m, denoted by $fQS_{t,m}^{i,j}(\tau)$, is given by

$$fQS_{t,m}^{i,j}(\tau) = \left(\frac{1}{m} \sum_{s=t-m}^{t} QS_{s|s-h}^{i}(\tau) - QS_{s|s-h}^{j}(\tau) \right) / \sigma_{t,m} \tag{8}$$

where $t = T + m, \cdots, T + P$ and T indicates the first forecast month and P the total number of forecasts. In practice, we use $m = 120$ which is equivalent to 10 years of monthly data and the 5 and 10% two-sided critical values from the non-standard asymptotic distribution are 3.012 and 2.766, respectively, while the one sided are 2.770 and 2.482. We index the fluctuation test $fQS_{t,m}^{i,j}(\tau)$ with the last observation of the estimation window so that the first value of the test refers to December 1994 and corresponds to the 120-month window from January 1985 to December 1994. The standard deviation of the statistic, $\sigma_{t,m}$, is obtained as the Newey-West standard errors for the difference in QS over the testing window.

Figure 3 shows the time series of the fluctuation QS test statistic for $\tau = 0.1$ (left column) and 0.9 (right column) for the four inflation measures considered when the QAR LRW benchmark is compared to QUR LRW, QAR SRW and AR-HET. The dotted lines represent the critical values for the null hypothesis of equal accuracy of the quantile forecasts against the one-sided alternative that QAR LRW outperforms the alternative forecasts (on the positive side) or the opposite (on the negative side). Hence, values of the fluctuation test outside the critical value on the positive side indicate that QAR LRW is more accurate in forecasting the quantile of interest, while for rejections that occur on the negative side we conclude that the alternative models outperform the benchmark (at a specific point in time). Considering first the performance of the QUR LRW forecasts relative to QAR LRW for $h = 12$, the full out-of-sample results in Table 3 indicate that assuming a unit root at all quantiles produces forecasts that are significantly less accurate at low quantiles, but equally performing at high quantiles, relative to the benchmark. This result is confirmed by the fluctuation test that shows that the test statistic for $\tau = 0.1$ is mostly large and positive, which suggests that the benchmark forecasts are superior to those of the QUR alternative and in the first part of the out-of-sample period for the two headline inflation measures. When considering the short rolling window QAR forecasts, the fluctuation test indicates the test statistics are mostly within the one-sided critical values for $\tau = 0.1$, but become significantly negative for $\tau = 0.9$ (except for core PCE) which suggests their higher accuracy relative to the long rolling QAR benchmark. This is also consistent with our earlier discussion of the higher performance of the short rolling window relative to the long one. Finally, the comparison of the QAR LRW and AR-HET forecasts shows that for $\tau = 0.10$ the forecast performance of the two models is, to a large extent, comparable, but at high quantiles the fluctuation test is significantly negative (except for headline CPI) in particular in the first half of the out-of-sample period. We also report the fluctuation test when the benchmark model is the AR-HET in Figure 4. One interesting comparison is between AR-HET and QAR SRW since, as discussed above, both were able to outperform, for most measures, the QAR LRW at the top quantiles. In the direct comparison of these two models we find that they have similar performance over time at both low and high quantiles for most inflation measures, but the short-rolling window seems to provide more accurate forecasts relative to the AR-HET forecasts for $\tau = 0.9$ for CPI and core CPI.

VII. Conclusion

Several studies have investigated the persistence of inflation with mixed conclusions about its level and variability over time. In this paper we consider this issue in the context of a heteroskedastic Auto-Regressive model for inflation where the conditional standard deviation of the error is a function of the past level and changes of inflation. The assumption that the volatility of inflation depends on its level is an established relationship in the literature which, if not properly accounted for, might lead to spurious evidence of time-variation in persistence. Our results indicate that persistence for the core CPI and PCE inflation measures has been high and fairly constant once the significant dependence of inflation volatility on its level is accounted for. In addition, we also find that this model provides equally or more accurate out-of-sample forecasts relative to a semi-parametric Quantile Auto-Regressive model, in particular at high quantiles. The fluctuation analysis shows that the better performance of the heteroskedastic AR model at high quantiles mostly occurs up to the early 2000s which can thus be characterized as a period of declining inflation volatility rather than persistence. We also compare the out-of-sample forecasts of the parametric model with a quantile model estimated on a 10-year rolling window which we consider short enough to capture parameter changes that might have occurred in the inflation process. We find no significant difference in accuracy between the forecasts of the two models and thus conclude that the decline of inflation and its volatility have been the main feature of the post-1984, rather than changes in persistence or breaks in the mean. These results refer to the core inflation measures and are particularly relevant since U.S. monetary policymakers seem to adopt the core measures, in particular PCE, as the main inflation indicator.

On the other hand, for the headline inflation measures we find that the oil price shocks had a permanent effect on the persistence of inflation which tends to decline once the 1970s are dropped from the estimation window. For these inflation measures we find that persistence declined even when we allow for a volatility component which depends on the inflation level. In terms of the out-of-sample performance, we find that the heteroskedastic AR model performs equally well relative to the short-rolling window quantile model at the annual horizon but it is outperformed a the one-month horizon. Overall, the findings point to the fact that the different measures show different persistence characteristics which, at least partly, can explain the differences in the existing literature that in some cases consider the quarterly GDP deflator (e.g., Pivetta and Reis, 2007) rather than the quarterly CPI index (e.g., Cogley and Sargent, 2002).

References

Amisano, G. and Giacomini, R. (2007). 'Comparing density forecasts via Weighted Likelihood Ratio Tests' *Journal of Business & Economic Statistics*, Vol. 25, pp. 177–190.

Beaudry, P. and Koop, G. (1993). 'Do recession change permanently output?' *Journal of Monetary Economics*, Vol. 31, pp. 149–163.

Benati, L. (2008). 'Investigating inflation persistence across monetary regimes' *Quarterly Journal of Economics*, Vol. 123 , pp. 1005–1060.

Brunner, A. and Hess (1996). 'Are higher levels of inflation less predictable? A state-dependent conditional heteroskedasticity approach' *Journal of Business & Economic Statistics*, Vol. 11, pp. 187–197.

Cogley, T., Primiceri and Sargent, T. (2010). 'Inflation-Gap persistence in the US' *American Economic Journal: Macroeconomics*, Vol. 2, pp. 43–69.

Cogley, T. and Sargent, T.J. (2002). 'Evolving Post-World War II U.S. inflation dynamics' In *NBER Macroeconomics Annual 2001* (eds B.S. Bernanke and K. Rogoff). MIT Press.

Evans, M. (1991). 'Discovering the link between inflation rates and inflation uncertainty' *Journal of Money, Credit, and Banking*, Vol. 23, pp. 169–184.

Evans, M. and Wachtel, P. (1993). 'Inflation regimes and the sources of inflation uncertainty' *Journal of Money, Credit, and Banking*, Vol. 25, pp. 475–511.

Fuhrer, J.C. (2010). 'Inflation persistence' In *Handbook of Monetary Economics, Volume 3* (eds Benjamin M. Friedman and Michael Woodford), Vol. 3, pp. 423 – 486. Elsevier.

Giacomini, R. and Rossi, B. (2010). 'Forecast comparisons in unstable environments' *Journal of Applied Econometrics*, Vol. 25, pp. 595–620.

Giacomini, R. and White, H. (2006). 'Tests of conditional predictive ability' *Econometrica*, Vol. 74, pp. 1545–1578.

Gneiting, T. and Raftery, A.E. (2007). 'Strictly proper scoring rules, prediction, and estimation' *Journal of the American Statistical Association*, Vol. 102, pp. 359–378.

Hansen, B. (1999). 'The grid bootstrap and the autoregressive model' *Review of Economics and Statistics*, Vol. 81, pp. 594–607.

Koenker, R. and Bassett, G. (1978). 'Regression quantiles' *Econometrica*, Vol. 46, pp. 33–50.

Koenker, R. and Xiao, Z. (2004). 'Unit root quantile autoregression inference' *Journal of the American Statistical Association*, Vol. 99, pp. 775–787.

Koenker, R. and Xiao, Z. (2006). 'Quantile autoregression' *Journal of the American Statistical Association*, Vol. 101, pp. 980–990.

Levin, A. and Piger, J. (2006). 'Is inflation persistence intrinsic in industrial economies?' *working paper*.

Oka, T. and Qu, Z. (2011). 'Estimating structural changes in regression quantiles' *Journal of Econometrics*, Vol. 162, pp. 248–267.

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

Stock, J.H. (2002). 'Comment on: Evolving post-world war ii U.S. inflation dynamics' In *NBER Macroeconomics Annual 2001* (eds B.S. Bernanke and K. Rogoff). MIT Press.

Tillmann, P. and Wolters, M.H. (2012). 'The changing dynamics of us inflation persistence: a quantile regression approach' *working paper*.

Tsong, C.-C. and Lee, C.-F. (2011). 'Asymmetric inflation dynamics: evidence from quantile regression analysis' *Journal of Macroeconomics*, Vol. 33, pp. 668–680.

Table 1

In-sample testing results

	CPI ($q=9$)			CPILFE ($q=7$)			PCE ($q=7$)			PCELFE ($q=5$)		
	$\hat{\rho}(\tau)$	$t_n(\tau)$	$U_n(\tau)$	$\hat{\rho}(\tau)$	$t_n(\tau)$	$U_n(\tau)$	$\hat{\rho}(\tau)$	$t_n(\tau)$	$U_n(\tau)$	$\hat{\rho}(\tau)$	$t_n(\tau)$	$U_n(\tau)$
τ=0.10	0.643	-3.77*	-214.77*	0.721	-6.06*	-167.94*	0.794	-3.29*	-124.34*	0.838	-3.838*	-97.90*
0.30	0.792	-4.94*	-125.08*	0.905	-2.42#	-57.21#	0.823	-4.97*	-106.98*	0.897	-3.53*	62.17*
0.50	0.900	-2.38#	-60.05#	0.998	-0.06	-1.11	0.916	-2.12$^{\odot}$	-50.81#	0.940	-2.20$^{\odot}$	-36.47$^{\odot}$
0.70	1.006	0.13	3.55	1.030	0.61#	18.29*	0.936	-1.52	-38.91#	0.982	-0.59	-10.87
0.90	0.969	-0.37	-18.07	1.131	2.49*	78.93*	1.009	0.16	5.45	1.069	1.25*	41.54*
QKS_t		5.22*			6.06*			5.69*			5.00*	
QKS_α			214.77*			167.94*			159.2*			101.00*
$\hat{\rho}$(OLS)	0.878	(0.829, 0.984)$^{(a)}$		0.913	(0.876, 0.998)$^{(a)}$		0.888	(0.846, 0.968)$^{(a)}$		0.940	(0.912, 1.004)$^{(a)}$	
ADF_t	-2.98#			-2.64$^{\odot}$			-3.29#			-2.40		
Break test date(s):												
τ=0.10		Feb 2001									Feb 1976	
0.30		Sept 1981/Feb 2001										
0.50		Apr 1980										
0.70		Mar 1980			Nov 1980			Mar 1980				
0.90		Mar 2004			Jan 1984			Oct 1980			Feb 1975	
AR-HET estimates:												
ρ		0.898*			0.956*			0.890*			0.928*	
c_1		0.015			0.068*			0.027*			0.068*	

The reported lag order q is selected by BIC. Significance at 1, 5, and 10% are denoted by *, # and $\odot$, respectively. The 90% confidence intervals for ρ denoted by $^{(a)}$ are computed using the grid bootstrap method of Hansen (1999) based on the non-studentized statistics, $(\hat{\rho}-\rho)$ based on a fine grid G of 200 and 1500 bootstrap replications. The significance for the ADF_t test is based on asymptotic critical values. The break dates for the Oka-Qu quantile break test are obtained for maximum two breaks and using a 5% significance level.

Table 2

QS Test ($h = 1$)

Null	Alternative	QS(τ) Test										
		0.1	0.15	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.85	0.9
							CPI					
QAR LRW	QUR LRW	*3.21*	*3.66*	*3.26*	*2.83*	*2.07*	*2.17*	*2.07*	*2.04*	1.46	1.36	1.50
	QAR SRW	-1.03	-0.36	-0.32	0.48	0.66	1.24	1.16	1.13	0.82	0.74	0.78
	AR-HET LRW	1.47	*1.67*	*1.82*	*2.38*	*2.56*	*2.66*	*1.86*	*2.02*	*2.22*	*2.81*	*2.69*
	AR-HET SRW	-0.15	-0.07	0.05	0.60	0.68	1.45	1.32	0.98	0.49	0.65	1.03
AR-HET LRW	QUR LRW	0.53	0.86	0.95	0.55	0.10	0.28	0.67	0.54	-0.25	-0.59	-0.53
	QAR SRW	**-2.45**	**-1.99**	**-1.72**	-1.28	-1.19	-0.55	-0.09	-0.17	-0.72	-1.34	-1.39
	AR-HET SRW	**-1.79**	-1.61	-1.24	-0.99	-0.91	-0.26	0.15	-0.20	-0.78	-0.96	-0.58
							Core CPI					
QAR LWR	QUR LRW	*5.89*	*3.67*	*2.89*	*2.59*	1.42	0.99	1.27	0.73	*2.00*	*2.00*	1.30
	QAR SRW	0.37	-0.31	-0.39	-0.01	0.23	0.44	0.19	-0.75	**-1.86**	**-2.38**	**-2.93**
	AR-HET LRW	-0.61	-0.89	**-1.81**	-1.18	0.06	1.22	*1.98*	0.98	-0.74	-1.17	-1.10
	AR-HET SRW	-0.44	-0.38	-0.45	0.17	0.26	0.08	0.01	-1.15	**-2.62**	**-2.68**	**-2.38**
AR-HET LRW	QUR LRW	*5.50*	*3.79*	*3.79*	*3.08*	0.97	-0.53	-1.22	-0.45	*1.80*	*2.17*	1.52
	QAR SRW	0.73	0.11	0.42	0.54	0.23	-0.19	-1.02	-1.36	-1.62	**-1.81**	**-2.06**
	AR-HET SRW	-0.14	0.04	0.30	0.62	0.24	-0.56	-1.22	**-1.86**	**-2.51**	**-2.28**	**-1.84**
							PCE					
QAR LRW	QUR LRW	*3.70*	*3.13*	*2.76*	*2.44*	*2.37*	*2.28*	*1.92*	1.35	1.38	1.53	1.66
	QAR SRW	0.60	0.05	0.22	0.48	0.58	1.12	1.18	0.74	*1.99*	*2.11*	*2.04*
	AR-HET LRW	0.86	-0.15	-0.80	-0.48	-0.57	-0.28	0.43	0.26	0.39	0.62	0.21
	AR-HET SRW	0.12	-0.45	-0.63	-0.14	0.58	1.27	1.59	*1.86*	*2.29*	*2.11*	1.34
AR-HET LWR	QUR LRW	*1.72*	*2.02*	*2.39*	*1.98*	*2.09*	*1.78*	1.13	0.86	0.89	0.92	1.24
	QAR SRW	-0.11	0.18	0.75	0.75	0.85	1.24	0.88	0.52	1.43	1.36	1.60
	AR-HET SRW	-0.57	-0.32	-0.17	0.09	0.81	1.38	1.38	1.62	*1.83*	1.56	1.11
							Core PCE					
QAR LRW	QUR LRW	*3.41*	*2.88*	*2.90*	*2.37*	1.30	0.53	0.65	0.58	0.38	0.85	1.48
	QAR SRW	1.14	0.75	1.17	1.36	0.96	0.77	0.88	0.44	0.12	-0.11	-0.48
	AR-HET LRW	-0.90	-0.48	-0.30	-0.89	-1.23	-1.56	-1.24	-0.53	1.10	*1.76*	1.17
	AR-HET SRW	0.76	0.82	1.16	1.22	0.52	0.03	0.33	0.22	0.18	0.07	-0.03
AR-HET LRW	QUR LRW	*3.97*	*3.09*	*2.84*	*2.50*	*1.81*	1.58	1.47	0.79	-0.81	-1.30	-0.34
	QAR SRW	1.58	0.94	1.20	1.54	1.35	1.33	1.31	0.67	-0.33	-0.84	-1.03
	AR-HET SRW	1.27	1.04	1.23	1.45	0.91	0.61	0.75	0.43	-0.28	-0.74	-0.67

The first column shows the benchmark model, the second column the alternative model, and columns 3 to 13 report the $QS(\tau)$ test statistic which is standard normally distributed (for $\tau = 0.10, \cdots, 0.9$). Negative values of the test statistic indicate that the alternative model outperforms the benchmark and the values in bold denote significance at 5% against this one-sided hypothesis. Instead, positive values indicate that the benchmark outperforms the alternative model and the statistics in italics denote significance against this one-sided hypothesis. The forecasting horizon h is equal to 1 and the out-of-sample forecast period starts in 1985.

Table 3

QS Test ($h = 12$)

Null	Alternative	QS(τ) Test										
		0.1	0.15	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.85	0.9
							CPI					
QAR LRW	QUR LRW	*4.15*	*4.34*	*4.41*	*3.44*	*2.31*	1.41	1.01	0.82	0.57	0.60	0.54
	QAR SRW	-0.55	0.18	0.51	0.97	0.94	0.51	-0.12	-0.74	-1.50	**-1.82**	**-2.03**
	AR-HET LRW	1.42	*2.49*	*1.86*	1.56	*1.86*	*2.13*	*1.82*	1.22	-0.02	-0.52	-0.50
	AR-HET SRW	1.43	*1.74*	*1.90*	*1.89*	*1.81*	1.50	0.89	0.18	-0.66	-0.93	-1.13
AR-HET LRW	QUR LRW	*3.94*	*3.66*	*3.61*	*2.58*	1.45	0.64	0.26	0.28	0.53	0.68	0.68
	QAR SRW	-0.99	-0.32	0.14	0.64	0.27	-0.43	-1.06	-1.15	-1.34	-1.50	**-1.69**
	AR-HET SRW	1.31	1.46	*1.67*	*1.70*	1.39	0.89	0.17	-0.26	-0.60	-0.76	-0.97
							Core CPI					
QAR LRW	QUR LRW	*4.70*	*3.50*	*2.48*	0.60	-0.44	-0.98	-1.36	**-1.64**	-0.56	0.40	0.65
	QAR SRW	0.90	1.50	1.23	1.01	0.75	0.60	0.28	-0.43	-1.24	**-1.79**	**-2.33**
	AR-HET LRW	-0.09	-0.94	-1.28	-0.59	-0.56	-0.23	-0.20	-1.29	**-2.22**	**-2.67**	**-2.66**
	AR-HET SRW	*2.20*	*2.01*	*1.74*	1.38	1.10	0.76	0.31	-0.42	-1.40	**-2.08**	**-2.55**
AR-HET LRW	QUR LRW	*4.50*	*3.62*	*2.90*	0.95	-0.23	-1.15	-1.63	-1.18	*1.72*	*3.33*	*3.46*
	QAR SRW	1.08	*1.82*	*1.67*	1.31	1.01	0.75	0.38	-0.04	-0.43	-0.63	-1.04
	AR-HET SRW	*2.03*	*2.02*	*1.89*	1.59	1.32	0.90	0.40	-0.11	-0.71	-1.13	-1.61
							PCE					
QAR LRW	QUR LRW	*4.09*	*3.91*	*3.71*	*2.80*	*2.03*	1.42	1.02	0.49	-0.27	-0.14	0.11
	QAR SRW	-0.76	-0.39	-0.19	0.14	0.32	0.28	0.26	0.06	-0.57	-0.76	-1.05
	AR-HET LRW	-0.36	-0.13	-0.09	-0.21	-0.06	0.00	0.01	-0.10	**-2.76**	**-3.65**	**-4.80**
	AR-HET SRW	-0.11	0.34	0.60	0.97	1.01	0.83	0.65	0.51	0.09	0.07	0.17
AR-HET LRW	QUR LRW	*3.83*	*3.38*	*2.87*	*2.08*	1.37	0.91	0.61	0.37	0.63	1.23	*2.22*
	QAR SRW	-0.73	-0.40	-0.19	0.25	0.43	0.36	0.32	0.12	0.29	0.40	0.64
	AR-HET SRW	-0.04	0.40	0.66	1.08	1.08	0.88	0.68	0.56	0.62	0.70	0.84
							Core PCE					
QAR LRW	QUR LRW	*6.11*	*4.45*	*2.87*	1.29	0.31	-0.24	-0.60	-0.86	-1.13	-1.08	-0.72
	QAR SRW	1.43	1.44	1.43	1.50	1.51	1.51	1.43	1.27	-0.13	-1.18	**-1.96**
	AR-HET LRW	0.81	1.06	0.76	0.59	0.04	-0.18	0.06	0.56	-0.40	**-2.21**	**-3.10**
	AR-HET SRW	*1.67*	*1.75*	*1.69*	1.37	0.88	1.09	1.14	1.09	0.90	0.69	0.35
AR-HET LRW	QUR LRW	*4.67*	*2.99*	*2.01*	0.87	0.23	-0.13	-0.48	-0.84	-0.82	-0.24	1.42
	QAR SRW	1.54	1.47	1.50	1.61	*1.77*	*1.86*	*1.78*	1.31	0.01	-0.40	-0.50
	AR-HET SRW	*1.74*	*1.76*	*1.74*	1.42	0.90	1.18	1.30	1.14	0.97	0.87	0.72

The first column shows the benchmark model, the second column the alternative model, and columns 3 to 13 report the $QS(\tau)$ test statistic which is standard normally distributed (for $\tau = 0.10, \cdots, 0.9$). Negative values of the test statistic indicate that the alternative model outperforms the benchmark and the values in bold denote significance at 5% against this one-sided hypothesis. Instead, positive values indicate that the benchmark outperforms the alternative model and the statistics in italics denote significance against this one-sided hypothesis. The forecasting horizon h is equal to 12 and the out-of-sample forecast period starts in 1985.

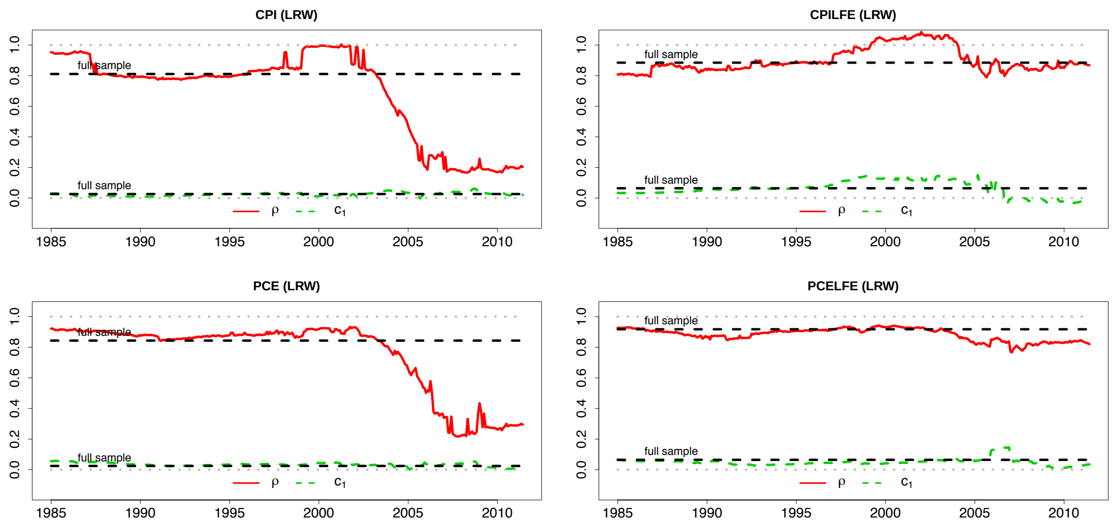

Figure 1. The continuous line represents the persistence parameter in Equation 3 and the dashed line the coefficient of π_t in the volatility equation for the heteroskedastic AR model. The time series represent the estimates of the two parameters on rolling window of 300 months and $h = 1$.

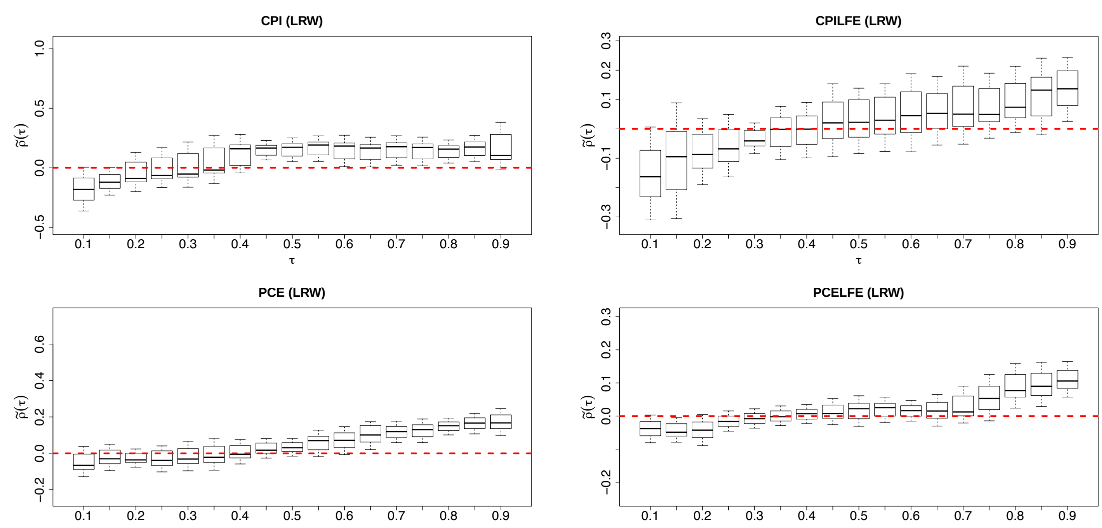

Figure 2. Boxplot of the time series of the estimates of $\tilde{\rho}(\tau)$ in Equation 7 for $h = 1$ as a function of the quantile level τ (for $\tau = 0.1, \cdots, 0.9$).

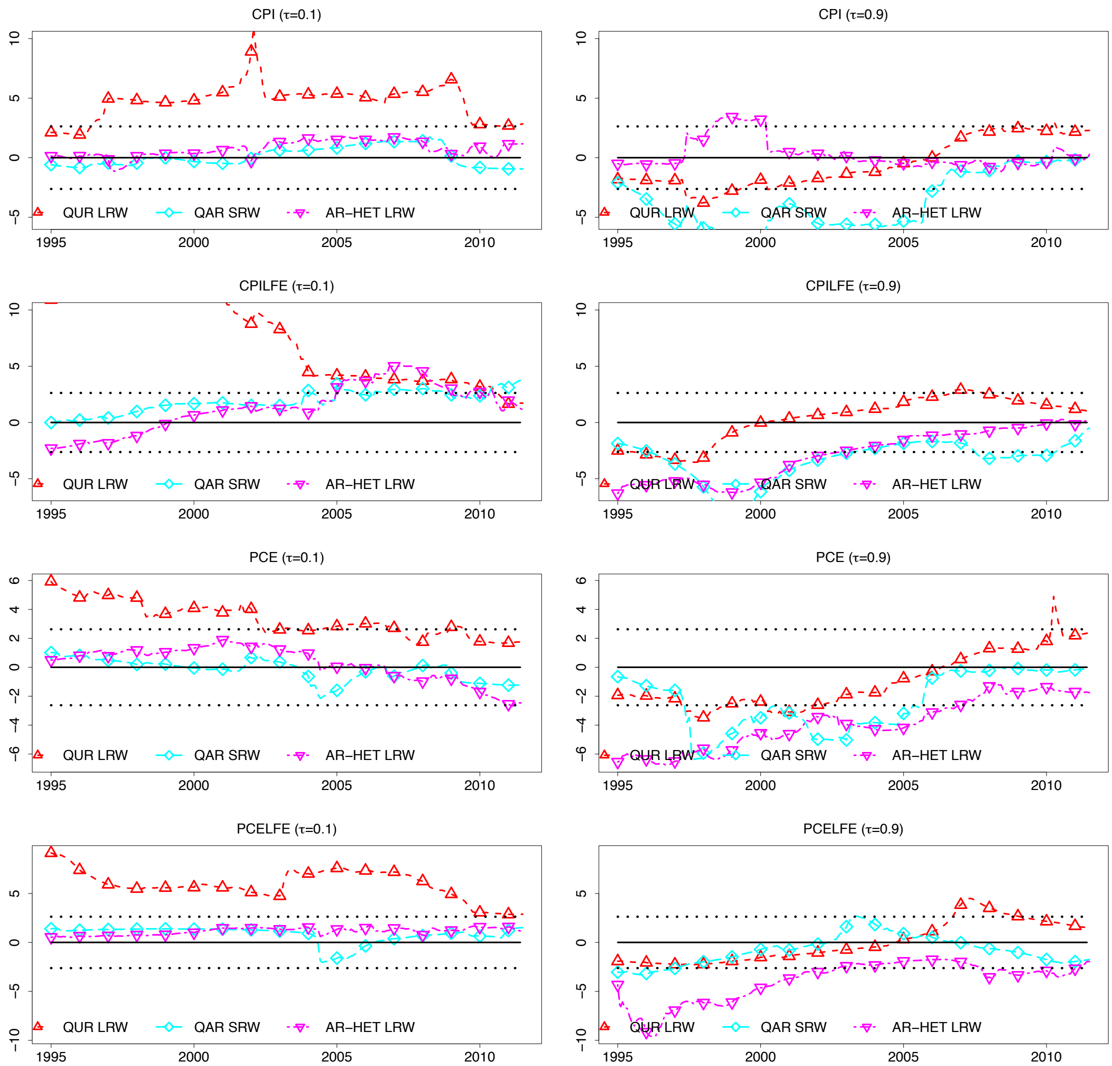

Figure 3. Fluctuation QS test for $h = 12$ and $\tau = 0.1, 0.9$ with benchmark model the long rolling QAR. The dashed lines represent the 5% one-sided critical values and rejections for negative values means that the alternative models outperform the benchmark, and the opposite for rejections in the positive side.

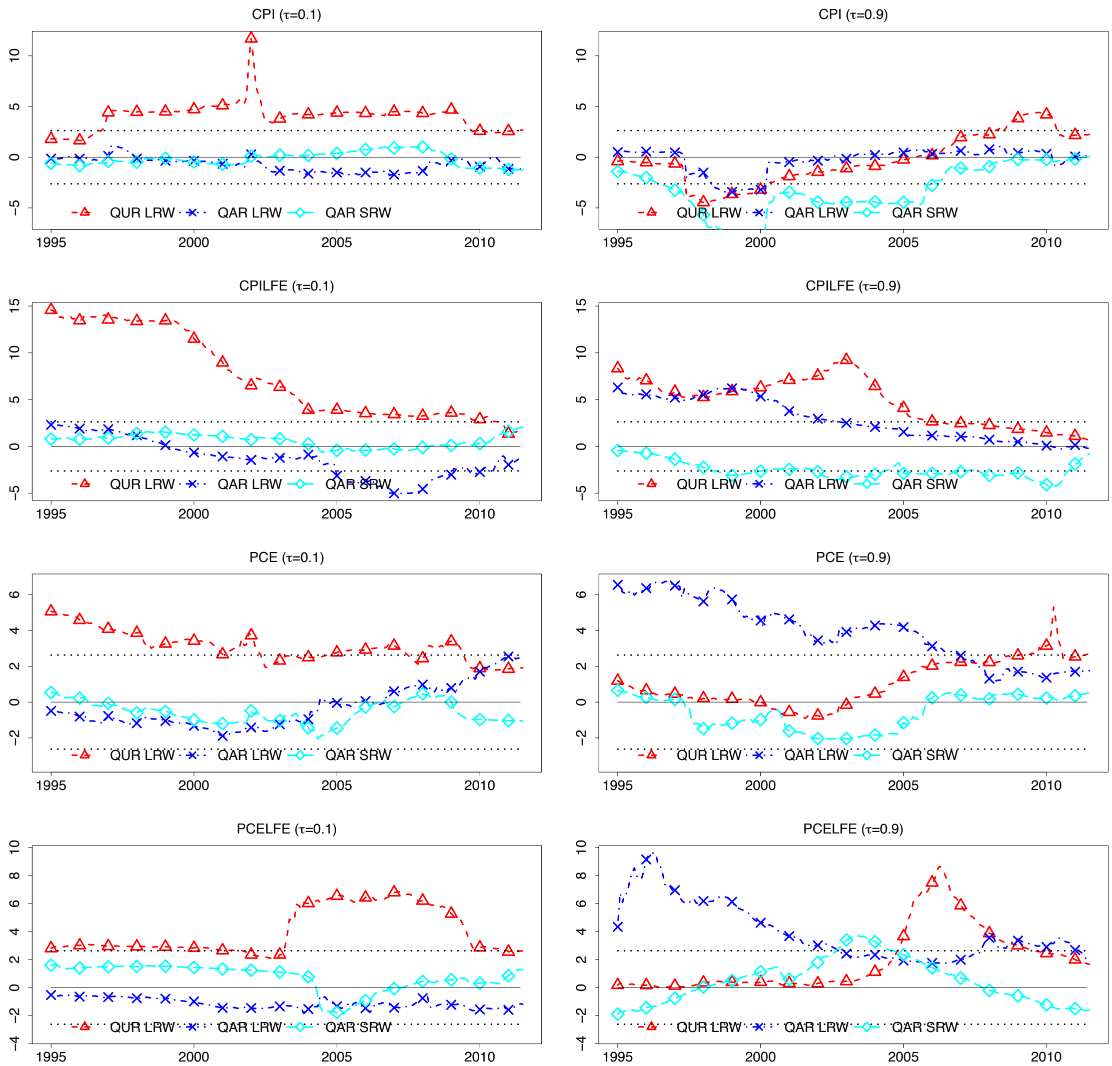

Figure 4. Fluctuation QS test for $h = 12$ and $\tau = 0.1, 0.9$ with benchmark model the long rolling HET. The dashed lines represent the 5% one-sided critical values and rejections for negative values means that the alternative models outperform the benchmark, and the opposite for rejections in the positive side.