

MPRA

Munich Personal RePEc Archive

A Time-Varying Expectations Formation Mechanism

Bovi, Maurizio

December 2019

Online at <https://mpra.ub.uni-muenchen.de/97624/>
MPRA Paper No. 97624, posted 18 Dec 2019 12:25 UTC

A Time-Varying Expectations Formation Mechanism

Maurizio Bovi[†]

Abstract

We propose an expectations formation mechanism (EFM) aimed to explain the median – hence lay – forecaster’s year-ahead inflation predictions. The EFM is a time-varying combination of long-run expectations, current inflation and uncertainty with weights naively calibrated according to inflation dynamics. Earning fixed income, in fact, the median forecaster has an aversion toward underestimation that increases with inflation. To allow for occasional – albeit unintentional – cost-minimizing calibrations, the EFM nests various forecasting rules. Data from the Michigan Survey of Consumers sustains the argued behavior and contributes to interpret some puzzling price dynamics such as the missing disinflation and reflation.

JEL codes: C53, C83, D84, E31, E37

Keywords: Survey expectations, Inflation, Time-Varying Parameters

[†] “Sapienza” University of Rome and ISTAT - Italian National Institute of Statistics. E-mail mbovi@istat.it; +39 06 46737465. The opinions expressed herein are those of the author and do not necessarily reflect the views of the mentioned institutions. I would like to thank two anonymous referees for their helpful comments.

1. Introduction

Despite inflation expectations are universally recognized to play a key role in economics, there is still relatively little knowledge about the way in which individuals form and update expectations about future inflation (Yellen, 2017; Armantier et al., 2013).

Our goal is to propose and test a time-varying expectations formation mechanism (EFM) aimed to explain the central tendency of lay forecasters' short-run inflation expectations across disparate price dynamics.

From the empirical standpoint, we use median expectations for U.S. inflation as gathered by the University of Michigan survey of Consumers (UMC).¹ We focus on median expectations (henceforth, the “consensus”) because the central tendency of the general public's expectations is paramount for both macroeconomic models and policies.

We focus on price expectations for their key role in economics as well as because U.S. inflation dynamics have recently been both potentially very informative and puzzling. For instance, it seems interesting to learn the possible change in the way people form expectations before and throughout the Great Moderation (GM) period and/or in the aftermath of the 2008 crisis.

As per the puzzling behavior of inflation dynamics, Friedrich (2016) has investigated the “twin puzzle” of higher-than-expected inflation despite economic slack from 2009 to 2011 and weakening inflation despite economic recovery starting in 2012. Inserting survey-based inflation expectations in a Phillips curve for 1995 to 2013, Friedrich finds that these measures of inflation expectations account for the twin puzzle. Coibion and Gorodnichenko (Coibion and Gorodnichenko, 2015) suggest that the 2009-2011 missing deflation can be explained by the rise of household inflation expectations induced by rising oil prices.² Though these works correctly stress the role of expectations in inflation dynamics (cf., also, Leduc and Sill, 2013), they left unexplored several aspects of the EFM. This said, the evidence that consumers' inflation expectations impinge on actual inflation offers another motivation for this paper.

From the theoretical standpoint, we make three main assumptions.

First, we claim that the consensus on one-step ahead expectations (π_{t+12}^e) is biased toward overestimation.³ This is because of the stylized fact that most agents earn fixed-incomes such as

¹ Curtin (1996) argues that the median must be preferred to the mean as measure of central tendency in the Michigan survey data (Section 3).

² The Phillips curve is built on firms' expectations. Nonetheless, there is some evidence that firms' inflation expectations track those of households (Coibion and Gorodnichenko, 2015).

³ Throughout the paper, the subscript indicates the number of months (see Section 3).

wages and pensions.⁴ The median forecaster, accordingly, forms expectations to reduce the probability of incurring in negative forecast errors ($\pi_{t+12}^e - \pi_{t+12}$). The idea behind this assumption is not new and it is well established but, as far as we know, never used as we do here. Shiller found evidence that, among the general public, the largest concern with inflation was that it lowers people's standard of living (Shiller, 1997). Similar findings are reported by other authors (Shafir, et al., 1997; Akerlof, et al., 2000). This view is also coherent with the stylized fact that wage/pension indexation systems are more widespread in higher inflation periods/countries (Hendricks and Kahn, 1985; Holland, 1995; Whitehouse, 2009). Though a minor issue, even overestimation may be costly so that forecasters must somewhat mediate between under and overestimation. This leads to the second assumption.

Second, we hypothesize that the EFM used for predicting π_{t+12}^e is a time-varying combination of current inflation (π_t), expected long-run inflation (π_{LR}^e), and uncertainty (σ_t^2). Assuming that inflation expectations are based on inflation is natural and common in the literature (see, e.g., Evans and Honkapohja, 2001). We insert the long-run expected inflation in the EFM because short-run expectations could be (come) anchored to some reference level of inflation (Bernanke, 2010). A flexible model should allow this (testable) case, especially in light of the let-the-data-speak approach that we follow in this paper. Uncertainty enters the EFM because of the assumed bias toward overestimation – higher uncertainty may induce higher forecasting errors. It is important to note that we expect uncertainty significantly enters the EFM only when it increases the probability of making negative forecast errors. The assumption of a time-varying weighting scheme is due to the sensitivity of fixed-income agents to inflation patterns. Since the 70s, indeed, the U.S. has been recording “tremendous changes in inflation dynamics” (p. 1, Stock and Watson, 2010).

Third, as per estimation methods we posit that lay forecasters do not use “optimization procedures”. In so doing we follow strands of the literature such as Simon's bounded rationality (Simon, 1957) and Shiller's irrational exuberance (Shiller, 2000). We are also in line with Kirman (Kirman 2006, p. 271), which argued that “individuals are not optimizers, they are adaptive and only behave “as if” they optimize”, and Fuster et al. (Fuster et al., 2010) which claimed that in forming expectations “agents unconsciously combine many insights”. Though possibly sub-optimal, predictions are neither random nor aimless. In the present paper, specifically, lay forecasters follow the informal logic described in Section 2, naively combining the parameters of the EFM to avoid underestimation.

In their informal calibration efforts, agents may occasionally end up with predictions congruent with optimal forecasting rules. Some parametrizations of the proposed EFM, in fact, are compatible with

⁴ According to the US Census Bureau, the number of total nonfarm payroll plus the number of retirees in the US amounts to more than 70% of the adult population.

both cost-minimizing and sub-optimal forecasting rules (Section 2.1). The contextual presence of rational (cost-minimizing) and sub-optimal behavior in the EFM is a common trait in the existing literature. If Muth-rationality is still the workhorse of macroeconomic models, evidence of departures from it is now well-established (Evans and Honkapohja, 2001). In this sense, our approach is also in line with - though more general than - that of “natural expectations”, whereas forecasts are modeled as lying in between rational and naïve expectations (Fuster et al., 2010).

This paper complements other strands of the research on expectations, too. The potential emergence of disparate EFMs brings our approach close to the body of work emphasizing the possibility that agents switch across different forecasting models (Brock and Hommes, 1997). A difference is that in the present paper the switch is not triggered by models’ predictive ability⁵ but, rather, by the will to reduce the probability of underestimations that we posit to be dependent on the inflation context. Capistran (2008) discusses the possibility of optimally biased expectations and studies the change in the shape parameter of an asymmetric loss function over some sub-samples. However, he imposes optimally biased expectations and deals with professional forecasters. Trehan (2015) compares the predictive ability of different forecasting models with the Michigan survey’s inflation expectations over time, but the author uses symmetric metrics. To the extent that consumers have motives to distort their forecasts, biased expectations are not an indication of poor predictive skill. In our setting, rationally cautious expectations may emerge. And, indeed, they show up (Section 4).

The proposed EFM fits well with the UMC data. In addition, the evidence confirms that the EFM changes over time and, in particular, that it varies according to the theorized *modus operandi*. Moreover, the proposed behavioral model can explain some puzzling events recently witnessed by the U.S. inflation. The missing disinflation, e.g., seems to be linked to a calibration mainly based on π_{LR}^e (that was greater than current inflation) and on uncertainty. In the following months, short run expectations were still determined by long term expectations, but with a smaller weight. Moreover, the effect of uncertainty shrunk. This may have hampered the reflation predicted by standard models. The paper proceeds as follows. Section 2 describes the proposed expectations formation mechanism, Section 3 the data, and Section 4 collects the empirical results. Concluding remarks and two Appendices close the paper.

⁵ Moreover, this literature is based on MSE statistics to select optimal models. As we shall see, instead, asymmetric losses permit optimally biased expectations (Section 2).

2. The Expectations Formation Mechanism

2.1 Parametrizations leading to Optimal and Sub-Optimal Consensus

We believe that lay forecasters do not perform optimal predictions but also that, albeit unintentionally, optimal predictions may nonetheless materialize. The proposed EFM tries to square this backdrop. Equation 2.1.1 formalizes the EFM for the next year inflation (π_{t+12}^e) as a time-varying combination of current inflation (π_t) long-run expected inflation (π_{LR}^e), and uncertainty about the increase in future prices (σ_t^2):

$$\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \partial_t \sigma_t^2 \quad (2.1.1)$$

where α_t , β_t and ∂_t are the time-varying parameters.

This EFM is flexible enough to embed some possible optimal and sub-optimal EFM proposed by the literature without imposing any particular mechanism other than the mentioned calibration (described in Section 2.2).

To be sure, we claim that lay forecasters minimize no loss function. Rather, we argue that common people informally calibrate the EFM to avoid underestimation and that, in so doing, the consensus might turn out to be compatible, ex post, with cost-minimizing predictions.

A peculiar parametrization of equation 2.1.1 generates optimal predictions under Linex loss. The Linex loss function may be defined as $L(a; \Delta) = b(e^{a\Delta} - a\Delta - 1)$, where $\Delta \equiv (\pi_{t+12}^e - \pi_{t+12})$, $b > 0$ is a scale parameter, $a \neq 0$ is the asymmetry parameter that controls the aversion of the forecaster towards either positive ($a > 0$) or negative ($a < 0$) forecast errors (Patton and Timmermann, 2007).

Combined with a (conditional to the information set at time t) Normal random variable with mean $\mu_{t+h,t}$ and variance $\sigma_{t+h,t}^2$, it can be shown that the optimal h -step ahead prediction that minimizes the expected loss satisfies ($h=12$):

$$\pi_{t+12,t}^e = \mu_{t+12,t} - a\sigma_{t+12,t}^2/2 \quad (2.1.2)$$

where the conditional moments are here estimated via an AR(12)-GARCH(1,1) model (Appendix B). In this GARCH setting, the conditional variance ($\sigma_{t+12,t}^2$, henceforth CV) plays the role of uncertainty (σ_t^2). Visual inspection of equations 2.1.1 and 2.1.2 reveals that expectations are based on Linex loss if $-a/2 = \partial_t$ and $\mu_{t+12,t} = (\alpha_t \pi_{LR}^e + \beta_t \pi_t)$. As per this latter, specifically, we assume that expectations are based on Linex loss if $(\alpha_t + \beta_t) = 1$. This optimality constraint emerges hypothesizing that π_{LR}^e is the

long-run mean of actual inflation.⁶ As per uncertainty, if our interpretation is correct underestimation is costlier than overestimation of the same magnitude. That is, we expect $a < 0$ and, accordingly, $\partial_t > 0$. Note that in the literature, usually, the shape parameter is considered to capture a time-invariant behavior (see, e.g., Zellner, 1986; Patton and Timmermann, 2007). In view of our flexible approach to agents' calibration efforts, instead, we allow ∂_t both to change and to be zero. All in all, equation 2.1.1 allows considering Linex loss without imposing it.

The case whereas the conditional mean is a weighted average of current and reference inflation and, in addition, $\partial_t = -a/2 = 0$ is important. For small values of $|a|$, in fact, the Linex is virtually symmetric and not far from a squared error loss function (Patton and Timmermann, 2007). In equation 2.1.2 MSE-compatible expectations would emerge if $a=0$, which makes clear that the optimal forecast is the conditional mean. The point is that under squared loss the optimal point forecast is unbiased: $E(\Delta)=0$. Thus, the proposed EFM is flexible enough to allow the emergence of both asymmetric loss - implying optimally biased expectations - and symmetric loss - implying optimally unbiased expectations.^{7,8}

While equation 2.1.1 permits the appearance of optimal predictions, other parametrizations make it compatible with "classical", possibly sub-optimal, EFM.⁹ The case $\alpha_t = \partial_t = 0$, $\beta_t = 1$ would lead to naïve expectations ($\pi_{t+12}^e = \pi_t$). If $\alpha_t = \partial_t = 0$ and if inflation follows an exponential smoothing pattern or, equivalently, an ARIMA(0,1,1), then the EFM may be thought of as an adaptive expectations scheme ($\pi_{t+12}^e = \pi_t = \lambda \sum_{i=0}^{\infty} (1 - \lambda)^i \pi_{t-i}$).¹⁰

2.2. Lay Forecasters' Way to Calibrate the Weighting Scheme

As said, in parametrizing the proposed EFM consumers might obtain both optimal and sub-optimal predictions. The next step is to explain how and why ordinary people tune the weights.

⁶ Appendix B shows the details about the optimality constraint. It also reports evidence that in the long-run, when both inflation and uncertainty are low, conditional mean and reference inflation tend to be equal. Sections 2 and 3 give more details on the theoretical links between the weights α_t , β_t and ∂_t .

⁷ Section 4.3 presents evidence on these restrictions.

⁸ Although the reference to the Linex loss function helps to discuss the proposed EFM, one can obtain essentially the same conclusions examining any generic asymmetric loss function, provided only weak restrictions are put on the form of that loss function and the data generating process (Patton and Timmermann, 2007). The requirement of the data-generating process is that the variable of interest is conditionally location scale distributed, and the requirement of the loss function is that it is homogeneous in the forecast error.

⁹ In fact, typically, classical EFMs deal with price levels rather than price dynamics. Note, also, that adaptive expectations (AE) are not necessarily sub-optimal (Evans and Ramey, 2006).

¹⁰ Recursive substitutions of the adaptive expectation (AE) formula $\pi_{t+1}^e = \pi_t^e + \lambda(\pi_t - \pi_t^e)$ show that AE are equivalent to an exponentially weighted moving average ($0 < \lambda < 1$), i.e., $\lambda \sum_{i=0}^{\infty} (1 - \lambda)^i \pi_{t-i}$.

The median consumer is a fixed-income agent and aims to avoid underestimation. This is the main scope of her forecasts. However, even overestimation bears some cost. E.g., it could lead to overconsumption. A balanced solution seems necessary. We posit that lay forecasters do not optimize but, rather, just informally combine the weights to avoid underestimation (Kirman, 2006; Fuster et al., 2010). Due to their limited skill and information set consumers may obtain sub-optimal outcomes, but they are happy even with *satisfactory* outcomes (Simon, 1957). To this end, we then claim that their mental tuning takes into account of the inflationary environment and goes as follows.

The variable π_{LR}^e may be thought of as the reference inflation, i.e., the price dynamic expected in “normal” circumstances (i.e., low inflation level and uncertainty) and/or as the trend inflation. In our setting there are reasons to think that lay forecasters naively predict π_{LR}^e via a sort of multiyear moving average of the actual inflation in normal inflationary stances (see Section 3 and Appendix B). This implies two things. First, π_{LR}^e is relatively exogenous with respect to π_{t+12}^e – trend inflation changes slowly and, if any, significant effects triggered by π_{t+12}^e can only emerge slowly. Second, when inflation is unusually high(low), then π_{LR}^e is smaller(larger) than current inflation. In inflationary environments, hence, π_{LR}^e should affect short-run expectations only weakly: the fear of underestimations leads the consensus to consider more π_t than π_{LR}^e . By the same token, when inflation is tendentially low and close to - or below - the reference rate, π_{t+12}^e should be mostly driven by π_{LR}^e because this reduces the probability of making negative forecasting errors ($\pi_{t+12}^e - \pi_{t+12}$). This behavior induces a strong negative correlation between the weights α_t and β_t .

The aversion toward underestimation leads to focus on a peculiar kind of uncertainty. In the present framework, specifically, only the uncertainty about the size of the increment in future prices affects short-run expectations. The uncertainty raising the probability of overestimation is a minor issue for fixed-income agents. Otherwise stated, shocks amplifying the uncertainty about the magnitude of the increase in next-year inflation have a larger effect on π_{t+12}^e than shocks inducing same-size uncertainty about the dimension of the decrease in next-year inflation.

Separating out the four possible combinations of inflation and uncertainty helps to make progress. On the one side, it clarifies the behavioral model that we think to be behind agents’ one-year-ahead predictions. On the other side, it allows the emergence of a number of testable propositions. Despite equation 2.1.1 includes a number of forecasting rules as special cases, thus, it nonetheless imposes sufficient discipline to obtain reliable evidence.

Case A: Low inflation, Low uncertainty. In this ideal environment the reference level of inflation is likely to be not smaller than current inflation and, hence, it is a suitable benchmark for lay forecasters with a propensity to overestimate. Thus, agents have less need to consider current inflation and

uncertainty to form their short-term expectations. In terms of equation 2.1.1 agents put the largest weight on the trend inflation: $\alpha_t \rightarrow 1$; $\beta_t, \partial_t \rightarrow 0$. Henceforth, we will refer to this parametrization as EFM-A. Should agents set exactly $\partial_t=0$ and $(\alpha_t+\beta_t)=1$, then, their predictions would turn out to be optimal in the MSE sense (henceforth EFM-M).

Case B: High inflation, Low uncertainty. When inflation is high, the EFM-A would lead to sizeable underestimations and π_t becomes more influential than π_{LR}^e in predicting next-year prices. In the agents' EFM, thus, α_t gets smaller and β_t escalates. Since inflationary environments are particularly dangerous for fixed-income agents, even relatively low uncertainty might affect their predictions. With respect to EFM-A, we conjecture an increase in the size of ∂_t . In fact, we expect a positive coefficient (cf. Section 2.1) and the selected forecasting model evolves to EFM-B: $\alpha_t \rightarrow 0$, $\beta_t \rightarrow 1$; $\partial_t > 0$. Unlike Case A, therefore, EFM-B with $(\alpha_t+\beta_t)=1$ is congruent with optimally biased expectations (henceforth EFM-L).

Case C: High inflation, High uncertainty. In this dramatically dynamic stance, the occurrence of costly negative forecast errors is more likely than in any other environment. Cautious forecasters react increasing the weight they put on σ_t^2 and on π_t . As a consequence of this re-balance, the reference level will have the smallest weight. EFM-C, i.e. EFM-B with a smaller α_t and both larger β_t and ∂_t materializes. If data show $(\alpha_t+\beta_t)=1$ then, again, the chosen combination of the weights would resemble to an asymmetric loss minimization exercise.

Case D: Low inflation, High uncertainty. As in the Case A, low inflation shifts the emphasis toward π_{LR}^e . Consequently, the role of π_t in the EFM narrows. Despite the probability to incur in underestimations is small, large uncertainty leads the weight ∂_t to be significantly positive. After all the adjustments are made, the new parametrization of equation 2.1.1 is EFM-D: $\alpha_t \rightarrow 1$, $\beta_t \rightarrow 0$; $\partial_t > 0$. If $(\alpha_t+\beta_t)=1$, then lay forecasters' weighting scheme makes the EFM coincident with asymmetric loss minimization.

To recap, amid “tremendous changes in inflation dynamics” (p. 1, Stock and Watson, 2010) tuning the weights of equation 2.1.1 is an as necessary as predictable job for fixed-income agents. Agents modulate the EFM via bounded rationality but, ex post, their predictions might also occasionally turn out to be optimal (cost-minimizing). Table 2.2.1 synthetizes the assumed testable framework.

Table 2.2.1 ABOUT HERE

2.3 Updating Rules

The EFM formalized in equation 2.1.1 posits which variables are considered by lay forecasters. Section 2.2 addresses the behavioral model behind their naïve modulation of the weighting scheme. The natural next step is to indicate how agents update the weights α_t , β_t , and ∂_t over time. Following the adaptive learning literature (Evans and Honkapohja, 2001) we hypothesize that the consensus is updated whenever new data are available. Unlike the adaptive learning literature, yet, we do not believe that lay forecasters act as econometricians (Bovi, 2014). We hypothesize that lay forecasters adopt some informal approach in both calibrating the weighting scheme (as described in Section 2.2) and in updating it. Otherwise the coherence of our setting would be compromised.

This said we now discuss whether, ex post, the consumers' informal updating rule turns out to look like recursive least square (RLS), rolling windows regressions or Kalman filter iterations. A formal digression permits to underline further details on why we believe that rolling iterations should be the better way to approximate the actual, intuitive and informal, reasoning behind the lay forecasters' updating.

Consider the following ordinary least squares regression: $\pi_{t+12}^e = \theta_t' x_t + \varepsilon_t$, where $\theta_t = (\alpha_t, \beta_t, \partial_t)'$, $x_t = (\pi_{LR}^e, \pi_t, \sigma_t^2)'$ and ε_t is a white noise sequence $(0, H_t)$. According to the RLS formula, parameters are updated as follows (Sargent, 1999):

$$\hat{\theta}_t = \hat{\theta}_{t-1} + t^{-1} R_t^{-1} x_t v_t \quad (2.3.1a)$$

$$R_t = R_{t-1} + t^{-1} [x_t x_t' - R_{t-1}] \quad (2.3.1b)$$

where R_t is the matrix of second moments of x_t , t^{-1} is the gain and v_t is the forecast error $(\pi_{t+12}^e - \hat{\theta}_{t-1}' x_t)$.

The RLS is not suitable for tracking time varying systems because the gain $t^{-1} \rightarrow 0$ as $t \rightarrow \infty$. In other words, the impact of successive prediction errors upon the estimate of $\hat{\theta}_t$ will diminish as the amount of information already incorporated in the estimate increases. A gain decreasing in this way may be dangerous for fixed-income agents predicting in inflationary environments - unsatisfactory outcomes are more likely to materialize. Despite their bounded rationality, the fear to underestimate suggests that lay forecasters' updating rule should produce outcomes far from this kind of procedure.

A possible step forward is to modify the RLS algorithm to produce a rolling regression. In fact, rolling window OLS estimations may be thought of as a constant gain algorithm (Sargent, 1999). In the presence of quickly evolving environments as those in the U.S., hence, rolling window estimations results may be closer to those stemming from the actual updating procedure behind the consensus.

Our flexible approach suggests involving in the discussion the Kalman filter, too. This latter offers a theoretical basis from which to develop an optimal updating algorithm. In our setting lay forecasters' predictions may be, albeit unintentionally, cost-minimizing ex post so that, for a matter of coherence, we must admit the possibility that their informal updating rule could be occasionally similar to optimal updating schemes.

The Kalman filter generalizes the previous procedures explicitly modeling the time varying process of the parameters (Hamilton, 1994; Sargent, 1999).

Parameters are modeled as a random walk, $\theta_t = \theta_{t-1} + n_t$ where n_t is a (vector valued) zero-mean process with covariance matrix Ω_t .

The Kalman recursion may then be formalized by the following three equations

$$\begin{aligned}\hat{\theta}_t &= \hat{\theta}_{t-1} + k_t v_t \\ k_t &= \frac{(P_{t-1} + \Omega_t)x_t}{H_t + x_t'(P_{t-1} + \Omega_t)x_t} \\ P_t &= P_{t-1} - \frac{(P_{t-1} + \Omega_t)x_t x_t'(P_{t-1} + \Omega_t)}{H_t + x_t'(P_{t-1} + \Omega_t)x_t} + \Omega_t\end{aligned}\tag{2.3.2}$$

where $P_t = E(\theta_t - \hat{\theta}_t)(\theta_t - \hat{\theta}_t)'$.

Setting $H_t=1$ and $\Omega_t=0$ gives back the RLS. Otherwise stated, the presence of the matrix Ω_t is one of the differences, and advantages, of the Kalman procedure – it lets some parameters to change significantly faster than others. In other terms, the Kalman filter is appropriate for representing continuous and optimal revisions.

It is perhaps worth stressing that we claim that these econometric exercises are not - and cannot be - performed by the respondents of consumers surveys. As we shall see in Section 4, rather, they are the econometrician's (i.e., the author's) attempt to approximate the time-varying parameters employed by the survey respondents. It should not be surprising that we expect that the median consumer's updating method should look more like a rolling procedure than an RLS or a Kalman filter. As already noticed, in fact, the former is too dangerous for fixed-income agents, and the latter approach is too demanding for lay forecasters.

3. Data

3.1. Inflation and Expectations

The data for the US inflation rate is the annual rate of change of the monthly Personal Consumption Expenditures Price Index (P), $100[\ln(P_t/P_{t-12})]$. Data are taken from the FRED database and cover the period January 1960 - October 2017 for the annual rate of change.

We use headline inflation because measures leaving out, e.g., food and energy, are congruent with neither our framework nor with the survey data (cf. Bullard, 2013, and below).¹¹

The Consumer Price Index (CPI) is another suitable inflation rate in our setting. The evidence reported in Appendix A shows that Personal Consumption Expenditures (PCE) inflation rather than CPI inflation seems to be the price concept behind UMC survey responses. Appendix A also indicates that results are robust regardless of the inflation metric used.

We use final revision data for two main reasons.¹² First, it is due to data availability. Long run real-time data for PCE inflation are only available on quarterly basis (vintages of monthly real-time data start from 2000). For CPI inflation, the first vintage for quarterly data is available since 1994q3 that for monthly inflation starts from 1998m11. In view of the empirical analyses of Section 4 and of monthly survey data, thus, we prefer to use final revision data.

Second, and possibly more important, real-time data could lead to underestimate the amount of information agents actually had at the time of the forecast. In practice, the average individual knows her own consumption, income, the prices she pays for consumption goods, and the like. At the time of their forecasting exercise, that is, it is likely that consumers know the situation that will be mirrored by revised data rather than that tell by real-time data (that are based on a smaller information set). In passing, it sounds incorrect that some authors equip lay forecasters with non-trivial econometric tools and at the same time subtract from the agents' information set this kind of knowledge. This knowledge, indeed, is one of the main reasons why households' surveys are so widespread and closely examined.

As per households' expectations, a unique data set can be obtained from the Surveys of Consumers, administered on a monthly basis by the Survey Research Center of the University of Michigan since January 1978. The goal of the Surveys of Consumers is to estimate economic expectations on a representative sample of all households located in the United States. The survey elicits consumers'

¹¹ Opposite reasons lead to not consider the GDP deflator. Using a deflator including expenditures such as investments or exports is too inclusive a measure of inflation from the consumers' standpoint.

¹² We have nonetheless performed some of the empirical exercises of Section 4 with real-time data. Unreported (but available upon request) results give further support to our evidence.

expectations of inflation rate for the next 12 months that are quantified using answers to the following question:

How many cents on the dollar do you expect prices to go up or down on the average, during the next 12 months?

The wording of this question should induce respondents to contemplate all prices with no exception. Also, the query does not refer to any particular price index, in that suggesting the use of all-item metrics.

As a proxy for the long-run expected inflation we consider five-ten year ahead inflation forecasts as quantified by using answers to the following question:¹³

By about what percent per year do you expect prices to go up or down, on the average, during the next 5 to 10 years?

A review of the estimates of inflation expectations indicates that for comparisons over time, the median should be the preferred measure of the central tendency of the response distribution. Curtin (Curtin, 1996) suggests using the median because it is not affected by truncated outliers that, instead, affect the estimate of the mean. Though the review is limited to the first part of our sample, we follow Curtin's hint and we focus on the median.

Fig. 3.1.1 and 3.1.2 HERE

Figure 3.1.1 highlights a number of "Great" episodes (Stock and Watson, 2010), as well as disparate dynamics in the US inflation. In the late '70s and early '80s the U.S. witnessed a double-digit inflation (the Great Inflation). Starting from 1981, price dynamics quickly declined until 1986 (the Great Disinflation), reaching values lower than 2 percent. In the following years, inflation returned to increase with a peak of about 5% at the end of 1990. Since then there was a new and protracted slowdown in price dynamics (2% on average), with a semester of negative rates in 2010 (Lower panel).

Figure 3.1.2 displays a tendency toward overestimation - the average of the median forecaster's prediction error ($\pi_{t+12}^e - \pi_{t+12}$) is 0.5%. This is even more evident for the mean forecaster.¹⁴ Overestimation seems to be seamlessly since the early 90s. The average of the median forecaster's prediction error in the period 1990m1-2017m10 is 0.9%. According to the behavioral model set up

¹³ This time series is available since 1979m02 with some missing value until 90m04 that we have filled using the cardinal spline method with tension parameter 0.1. In fact, there is some previous value. Yet, the query was different, so the responses are not comparable.

¹⁴ Mean expectations are closely correlated with, but systematically higher than, the median expectations. On average, the mean is 100 basis point higher than the median and their correlation coefficient is 0.98.

in Section 2 the reason is the propensity toward overestimation that, when inflation is low, leads lay forecasters' expectations to depend more heavily on the reference inflation (π_{LR}^e). This latter, in fact, has been almost always higher than actual inflation throughout the years 1990-2017 for two main reasons. Firstly because, given the wording of the query relative to π_{LR}^e it, unsurprisingly, appears to behave like a multiyear moving average of inflation (Appendix B). Moreover, this is in line with the evidence that, in their naïve updating, consumers look at rolling windows (Section 4). Figure 3.1.1 (Upper panel) shows that whereas in the first - high inflation - part of the sample π_{LR}^e was tendentially lower than current inflation, the reduction in inflation has reversed the situation. Secondly, in spite of the low inflation π_{LR}^e maintains strictly positive possibly because of the effect of the FED's strategy on it. Since the early '90s, the FED has been following an implicit, but strongly maintained and subsequently formalized, nominal anchor in the long run (Goodfriend, 2003; Mishkin, 2004). Figure 3.1.2 displays some occasional episode of underestimation that we shall examine in Section 4.2.

The proposed logic may also be suggestive of why data unveil that the US median forecaster has never figured out deflation as a real possibility. When inflation is low, short-run expectations are strictly linked to the trend inflation and this latter remained positive for the above-mentioned reasons. There exist other possible explanations for that. For instance, information rigidities (Coibion and Gorodnichenko, 2012) may hamper quicker reactions to a new environment, and deflation was perhaps too short to produce negative expectations. Strictly positive expectations may then be due to the fact that deflation was a brand-new experience for the median forecaster, and the evidence suggests that inflation expectations are linked to inflation experience (Lombardelli and Saleheen, 2003; Malmendier and Nagel, 2016). Our interpretation seems to complement rather than to contradict the others.

3.2. Inflation Uncertainty and Expectations Uncertainty

Uncertainty is hard-to-measure. It is even truer when dealing with the uncertainty perceived by lay forecasters in their predictive exercises. We use three proxies for the uncertainty that we believe particularly appropriate for our research interest.

The first proxy for uncertainty is the conditional variance of the AR(12)-GARCH(1,1) model of inflation estimated in Appendix B. It turns out to be high when inflation is either high or low (cf. Figure 3.2.1). As mentioned in Section 2, the conditional setting may give useful indications to our goals. This said, we claim that the other two proxies that we use are more suitable in the present setting. The key point is that these two proxies come from the same surveys upon which is based our

variable of interest. Thus, they are naturally closer to the mental exercise elicited from survey respondents than proxies based on hard data such as the conditional variance (CV).

The second proxy is the uncertainty index elaborated by Binder (Binder, 2017), which is based on the UMC survey data. The basic assumption behind the Binder's index (BI), which can take values from 0 to 100, is that agents that are sufficiently uncertain about their forecast on inflation choose from a set of round numbers when responding to the survey. An index equal 100 implies that all respondents indicate round numbers and, hence, that all agents are uncertain about their forecasts. Similarly to the CV, the BI is found to be countercyclical and high when inflation is either high or low (Binder, 2017). As we shall see, this latter feature turns out to be important in the present setting.

The other measure of uncertainty from UMC surveys is the share of interviewed indicating that they expect prices to rise, but reply that they do not know how much prices will increase (henceforth, DK_U). A possible drawback of DK_U is that it may reflect ambiguity or unwillingness to express inflation expectations quantitatively for reasons different from uncertainty. Nonetheless, several authors have used it as proxy for uncertainty with some success (Curtin, 2007; Blanchflower and Kelly, 2008). On the other hand, even the BI may reflect unwillingness to express precise quantitative expectations for reasons different from uncertainty. Furthermore, we have specific a priori about the role of DK_U in the EFM, we then maintain that the sign of its coefficient is positive and, finally, we offer evidence about that (Section 4).

We believe that the information content of these three proxies is complementary, in that helping to obtain robust evidence. We also posit that DK_U is the most pertinent proxy for the kind of uncertainty that we need to include in the proposed EFM. As argued in Section 2, in fact, a suitable proxy for uncertainty in our context should detect higher uncertainty only when there is wider uncertainty surrounding the size of the increase in future inflation. Instead, the proxy should remain small when the uncertainty is about the magnitude of the decrease in future inflation. To this end, by construction, DK_U provides more useful information with respect to the other two proxies.

Figure 3.2.1 displays the three proxies.

Figure 3.2.1 ABOUT HERE

Full-sample (1979m2-2017m10) correlations coefficients are: $\rho_{DK_U,BI}=0.17$, $\rho_{DK_U,CV}=0.19$, $\rho_{BI,CV}=0.40$. These statistics confirm *i*) that DK_U may actually capture different aspects of uncertainty with respect to the other two metrics, and *ii*) that these latter are more strongly connected to each other than to DK_U. A relatively common trait between all the proxies emerges in the Great Disinflation (Upper panel), when all measures detected declining levels of uncertainty. Possibly

putting aside the early ‘90s,¹⁵ uncertainty shrunk until the mid-90s and then remained relatively low. Examining those years, unsurprisingly several authors have argued that inflation level and inflation uncertainty go hand in hand (Cukierman and Wachtel, 1979; Ball, 1992; Capistran and Timmermann, 2009).

The period of relatively low uncertainty, the Great Moderation (GM), came to an end about in 2000. Wars and terrorist attacks at the beginning of the new millennium and the 2008 financial crisis are among the causes of these developments (Bloom, 2014). It is worth emphasizing that, unlike previous episodes and in contrast with the mentioned view that uncertainty and inflation are positively correlated, in the last two decades the typical scenario is featured by low inflation and high uncertainty (Lower panel, and Bloom, 2014). To our aims, Figures 3.1.1 and 3.1.2 sustain the importance to offer and test a ductile EFM and recommend a careful choice of the proxy for uncertainty.

The Lower panel of Figure 3.2.1 emphasizes the different behavior of the proposed proxies for uncertainty when the system is hit by an adverse shock. Unlike the other two proxies, in fact, DK_U did not raise significantly despite the 09/11 shock. Rather, it remained as relatively constant as the inflation level.¹⁶ This is reflected in the virtually zero-correlation between them: $\rho_{\pi,DK_U}=0.02$. Instead, the other two measures were negatively correlated with inflation: $\rho_{\pi,CV}=-0.75$, $\rho_{\pi,BI}=-0.44$. A similar story can be told for the 2008 shock. Over the period 2008-2009, DK_U mirrored the decline in inflation (showing a correlation of $\rho_{\pi,DK_U}=0.61$), whereas the other two proxies remained negatively correlated with inflation: $\rho_{\pi,CV}=-0.65$, $\rho_{\pi,BI}=-0.35$. All considered, the behavior of DK_U is what one should expect from a proxy aimed to capture the uncertainty surrounding the extent of next-year price increases. Section 4 offers further insights on that.

4. Empirical Evidence

4.1. How do Lay Forecasters Update? Do Inflation Expectations Fit the proposed Taxonomy?

Section 3 shows a rich variety of inflation level and uncertainty, offering the occasion to shade some light on the above mentioned questions.

To answer to the first question, we perform rolling estimations and Kalman iterations to see which is closer to the lay forecasters’ naïve updating procedure and, then and more importantly, to use the horse race winner as proxy for the time-varying parameters employed by the survey respondents. That is,

¹⁵ Inflation reached 5% at the end of 1990 (cf. Figure 3.1.1).

¹⁶ In the period 2001m08-2002m03 inflation slowed down from 2% to 1%.

we exploit the time varying parameter deriving from one of our econometric exercises “as if” they were the coefficients used by U.S. consumers. As shown below, the choice among the econometric outcomes is established in terms of correlations and RMSE.

Taking stock of the digression of Section 2, the state-space representation of equation 2.1.1 that we use for the Kalman procedure is

$$\begin{aligned} \pi_{t+12}^e &= \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{DK_U,t} + \delta_{BI,t} + \delta_{CV,t} + e_t & e_t &\sim N(0, \sigma_e^2) \\ \gamma_{t+1} &= \gamma_t + e_{\gamma,t} & e_{\gamma,t} &\sim N(0, \sigma_{e_\gamma}^2) \end{aligned} \tag{4.1.1}$$

where $\gamma_t = (\alpha_t, \beta_t, \delta_{DK_U,t}, \delta_{BI,t}, \delta_{CV,t})$. The corresponding five transition equations imply that all the five coefficients of the state-space representation 4.1.1 are assumed to follow a random walk.

We have estimated the parameters via ML and we have initialized the recursion using the coefficients and the coefficient covariance matrix from an OLS estimation of the signal equation of the system 4.1.1.¹⁷

FIGURE 4.1.1 ABOUT HERE

The broad picture emerging from Figure 4.1.1¹⁸ sustains the testable propositions summarized in Table 2.2.1. We postpone further comments after having performed rolling windows estimations, i.e., the alternative way that we propose to approximate the naïve updating rule actually performed by the median consumer.

In estimating the time-varying parameters of the signal equation in 4.1.1 via rolling windows we have used windows sizes of 24 up to 60 months with a step-size of 6 months. The fit tends to worsen marginally as we use larger windows. The RMSE for the five-years window size the is 0.32 to be compared with an RMSE of 0.27 for the two-years window size. The evidence is substantially confirmed with any windows size. In the following Figure 4.1.2., we therefore only report the evidence based on 24-months rolling windows.

FIGURE 4.1.2 ABOUT HERE

¹⁷ We have also initialized the recursion using diffuse priors. Outcomes remain substantially the same. E.g. the RMSE for diffuse priors is 0.38 to be compared with 0.37 obtained using priors based on the OLS regression.

¹⁸ Figure 4.1.1 starts from 1981 because the rolling coefficients of Figure 4.1.2 start from 1981. In the case of the Kalman filter this should also reduce the issues linked to the initial estimates.

The correlation between fitted and actual expectations computed from the Kalman and the rolling approach is high, respectively, 0.92, and 0.96. Crucially, these correlations suggest that lay forecasters' updating exercises may be better proxied by rolling windows iterations rather than by Kalman updating. Estimated residuals are in line with this conclusion, too. The RMSE stemming from the Kalman filter is 0.37 (MAE=0.25), whereas the RMSE from rolling windows estimations is 0.27 (MAE=0.19). All considered the evidence confirms our reasoning and suggests limiting further comments to rolling coefficients, which we henceforth use as if they were those actually adopted by U.S. consumers.¹⁹

The leftmost point estimation in the panels of Figure 4.1.2 is the coefficient referring to the regression relative to the sample period 1979m2-1981m1, the second point refers to the regression relative to the sample 1979m3-1981m2 and so on. Likewise for the 95% confidence bands. The whole procedure consists of 441 regressions.

The overall picture emerging from the rolling estimations is congruent with our a priori. The long-term behavior of the weights, $\alpha_t \rightarrow 1$ and $\beta_t \rightarrow 0$, mirrors the historical decline in inflation. The correlation between these two parameters is highly negative: the full-sample coefficient is -0.75.

Figure 4.1.2 also underlines the different effect of the proxies for uncertainty on short-run expectations. This even more evident for DK_U which, e.g., unlike the other proxies never shows significantly negative coefficients. A positive ∂_t is in line with our a priori, and the detected difference between the three proxies is unsurprising in light of the digression made in Section 3.2.

In Section 3.2 we have also argued for the superiority of DK_U with respect to the other two measures. Table 4.1.1 collects the results of a horserace comparing the three proxies under scrutiny.

Table 4.1.1 ABOUT HERE

Table 4.1.1 sustains the superiority of DK_U: the regression with DK_U decisively displays the better fit. Full-sample regressions also permit to give further support to the choice of approximating the time-varying parameters with those stemming from rolling windows estimations. The RMSE obtained from full-sample regressions is much higher than that obtained via rolling windows (0.27).²⁰

¹⁹ To save space we do not report evidence showing that, in terms of RMSE, rolling windows is closer to the lay forecasters updating rule than the RLS. Recall that we are not testing whether lay forecasters use rolling windows or Kalman filter. We do not see lay forecasters as econometricians. We are testing whether lay forecasters' updating rule is better proxied by rolling windows than Kalman filter iterations.

²⁰ The RMSE for the full-sample estimation of the signal equation in the system 4.1.1 is 0.53.

All in all, we are now in a position to focus on the evidence stemming from rolling estimations limited to DK_U.

4.2 Can the Proposed EFM Explain Peculiar Episodes of U.S. Inflation?

Figure 4.2.1 aims to emphasize the instinctive reaction of lay forecasters' expectations in peculiar episodes that we can illustrate within our framework. It reports the rolling (two-years window size) coefficients estimated using the equation $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{DK_U,t}$.

FIGURE 4.2.1 ABOUT HERE

The effect of the Great Inflation/Disinflation on the quick and wide shift of emphasis between current and long-run expected inflation is easily recognized in Figure 4.2.1. The high inflation of the early '80s led lay forecasters to promptly switch the weight from π_{LR}^e to current inflation. The disinflation triggered a similarly rapid - but opposite - switch.²¹ These rapid and wide adjustments requested by those chaotic years ended with negative forecasting errors (Section 3.1). It is perhaps worth noticing that *i*) in the early '80s even professional forecasters significantly underestimated inflation (Romer and Romer, 2002; Bordo and Orphanides, 2013), and that *ii*) we posit a time-varying EFM managed with bounded rationality, not that agents are able to avoid underestimation all the time. In Section 4.3 we offer further evidence on the sub-optimal behavior featuring the early '80s.

The parameters α_t and β_t display a less wide but similar evolution in 1991-1992, when the relative regressions include the bounce of inflation recorded in the late 1990. Again, fast inflation dynamics induced negative forecasting errors (Section 3.2).

Reflecting the GM, in the subsequent years the pattern of α_t and β_t was relatively smoother, with the former approaching one and the latter pointing to zero. The GM also affected the weight given to uncertainty, which was rather small and statistically zero at any significance level. Those years witnessed only positive prediction errors.

In the aftermath of the 09/11 attack, the coefficient α_t abruptly collapsed from a value of 0.68 to zero in few months. π_{LR}^e was immediately substituted by π_t in the EFM: β_t promptly raised from 0.26 to 1.²² Agents' reaction was again rather strong and fast. If our interpretation is correct, the adjustment was due to the lay forecasters' belief that in that brand-new and uncertain situation the reference

²¹ Note that the coefficients reported in Figure 4.2.1 in, e.g., 1983m1 derive from the 1981m2-1983m1 regression.

²² Point estimates are: $\alpha_{2001m8}=0.68$, $\alpha_{2001m9}=0.74$, $\alpha_{2001m10}=0.48$, $\alpha_{2001m11}=0.04$; $\beta_{2001m8}=0.26$, $\beta_{2001m9}=0.18$, $\beta_{2001m10}=0.53$, $\beta_{2001m11}=1.02$.

inflation was a poor guidance for their short-run predictions. The uncertainty generated by the terrorist attack, then, led to a significantly positive ∂_t until the mid-2004.

The response to the 2008 shock was different, but again congruent with our view. Though the shock led to a reduction in α_t and to an increase in β_t , the correction was much less pronounced than that recorded in the aftermath of the 2001 shock. This is so because after the 2008 shock current inflation immediately contracted, recording negative values: from 4% to -0.2% in the period 2008m08-2009m03. In 2001, instead, inflation was only marginally declining and did not go below zero (it slowed down from 2% in 2001m08 to 1% in 2002m03). Hence, despite the negative effect of the 2008 shock on α_t , the propensity toward overestimation hampered forecasters to give too much weight to negative inflation rates. The deflation, likewise, made the effect of the uncertainty surrounding the probability of underestimation neither as sustained nor as significant as in the aftermath of the 2001 shock. In both shocks negative forecasting errors showed up.

To the extent that consumers' inflation expectations impinge on actual inflation, the puzzling behavior of the 2009-2011 missing disinflation may then be solved by taking into account the evolution of the forecasters' EFM after the 2008 shock. Inflation did not fall as much as expected by standard models because, as said, the deflationary shock led prone-to-overestimation agents to operate a limited trade-off between α_t and β_t . Albeit narrowed as argued, then, uncertainty gave further support to higher expectations. As the effect of the 2008 shock on π_{LR}^e watered down, agents started increasing α_t , in that generating the missing reflation. The reduction in DK_U and, accordingly, in its coefficient offers further support to the lack of reflation. In terms of our taxonomy (Table 2.2.1), the missing reflation is explained by the adoption of the EFM-A: $\alpha_t \rightarrow 1$; $\beta_t \rightarrow 0$, $\partial_t = 0$. In the very last part of the sample three developments can be noticed: *i*) the sustained role of the reference inflation, *ii*) the re-emergence of a small, but significant, role of uncertainty, and *iii*) the zero-weight given to current inflation. This latter is as historically rare as intriguing aspect of the current EFM. Side-by-side with the mentioned effect of the FED's moves on α_t , another explanation of the zero-weight put on the current inflation is that agents may have learned that, in a low inflation era, a satisfactory (in the sense of H. Simon) way to avoid underestimation is to anchor their short-run expectations to their long-run expectations (without forgetting the role of the uncertainty).

4.3 Are Inflation Expectations Sub-Optimal, Optimally Biased or Optimally Unbiased?

Though lay forecasters aim to avoid underestimation without using optimizations procedures, they may unintentionally choose parametrizations occasionally leading to cost-minimizing predictions.

We stress again that the flexibility of the EFM allows this outcome. The restrictions that we want to test are the following (ε_t^i , $i=U, L, M$, are the estimated residuals):

$$\left. \begin{aligned}
 \text{Unrestricted Model (EFM-U)} \quad & \pi_{t+12}^e = \alpha_i \pi_{LR}^e + \beta_i \pi_t + \delta_{DK_U,t} DK_U + \varepsilon_t^U \\
 \text{Linex-restricted Model}^{23} \text{ (EFM-L)} \quad & \pi_{t+12}^e = \alpha_i \pi_{LR}^e + (1-\alpha_i) \pi_t + \delta_{DK_U,t} DK_U + \varepsilon_t^L \\
 \text{MSE-restricted Model (EFM-M)} \quad & \pi_{t+12}^e = \alpha_i \pi_{LR}^e + (1-\alpha_i) \pi_t + \varepsilon_t^M
 \end{aligned} \right\} \quad 4.3.1$$

Working with these three nested models we have computed F-statistics to see which model outperforms the others in which period.

FIGURE 4.3.1 ABOUT HERE

Figure 4.3.1 reports the p-values of the F-tests stemming from comparing the sums of squared residuals relative to the three nested EFM introduced in equation 4.3.1 estimated with rolling least squares (all with two-years window size).

In the Upper panel the first star quantifies, on the vertical axis, the p-value of the F-tests obtained in the horserace EFM-U vs EFM-L over the period 1979m2-1981m1; the second star refers to the period 1979m3-1981m2, and so on up to 2015m11-2017m10 for a total of 441 regressions/stars. Likewise for the other two panels. The horizontal solid line at the 0.05 level highlights the 95% confidence band.

When the stars are below the 0.05 line in both the Upper and the Middle panel, then the agents' preferred combination of weights is EFM-U. Instead, EFM-M turns out to be the chosen mechanism whereas the stars are above the 0.05 line in both the Middle and the Lower panel. Finally, the EFM-L is the selected scheme whereas the stars are above the 0.05 line in the Upper and below this line in the Lower panel. In fact, in this latter horserace EFM-L is the unrestricted model. It should be clear that in the cases whereas EFM-U beats the other two then the Lower panel, i.e., the last match (EFM-L vs. EFM-M), is redundant.

The shaded areas in the three panels tentatively highlight the periods in which the chosen set of weights is compatible with the EFM-U (Upper panel), the EFM-M (Middle panel), and the EFM-L (Lower panel). The cut-off is evidently not perfect. For instance, the EFM-U was the preferred model even in the return of inflation in the early '90s. Figure 4.3.1 just aims to underline the big picture as clearly as possible.

²³ The reference to the Linex cost function is not restrictive because the Linex can be generalized under weak requirements (Section 2).

Throughout the Great Inflation/Disinflation the vast majority of the stars in both the Upper and Middle panels are below the 95% confidence band, indicating that agents have selected parameters such a way that their forecasts are based on EFM-U. An explanation is that when inflation is volatile, then it is hard for lay forecasters to calibrate the weights that would lead to optimal predictions (as well as to avoid underestimations. Cf. Figure 3.1.2).

A possible confirmation of the foregoing comes from the Great Moderation, when data point out that lay forecasters' expectations were optimal. Specifically, the most used EFM in those years was that with MSE-restrictions (Middle panel). It may be so because in that ideal period with low inflation level²⁴ and uncertainty was easier to select, albeit unintentionally, a weighting scheme congruent with optimal predictions. Forecasters, moreover, had no need to consider uncertainty and the optimal predictions turned out to be unbiased. Otherwise stated, during the GM agents have been predicting using $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + (1 - \alpha_t) \pi_t$. This is somewhat in line with the findings of Atkeson and Ohanian (2001), stressing that a simple random walk (RW) outperformed a NAIRU model as predictor of inflation during the period 1985-1999.²⁵ More in general, our findings sustain the existing evidence showing that in periods whereas the slack variable does not enter the Phillips curve uncertainty does not affect expectations (Ball, Mankiw, and Romer 1988; Akerlof, Dickens, and Perry 1996; Stock and Watson, 2009; IMF, 2013).

The Lower panel informs that since the end of the GM there seems to be a prevalence of optimally biased predictions. According to our reading, the increase in uncertainty with respect to the ideal conditions of the low-inflation-GM period persuaded lay forecasters to enlarge $\delta_{DK,U,t}$, in that switching from unbiased to biased expectations. Again possibly thanks to low inflation rates, yet, predictions continued to be cost-minimizing.

Though our main aim is to try to decipher how median forecasts are formed, some consideration on monetary policy naturally arises. The beginning of the EFM-M period (mid-90s) coincides with the start of the FED's strategy involving an implicit, but strongly maintained, nominal anchor in the form of an overriding concern to control inflation in the long run (Goodfriend, 2003; Mishkin, 2004). In other words, it seems that the Greenspan's skill in communicating that the FED was pursuing a policy aimed to maintain a stable and moderate inflation pushed lay forecasters to consider only the inflation level. Calibrating efforts made the consensus compatible with optimally unbiased forecasts and small positive errors showed up.

²⁴ The average inflation rate for the period 1995-2000 is 1.7%.

²⁵ This result has been confirmed by several authors (e.g., Stock and Watson, 2009)

5. Concluding Remarks

This paper has introduced a novel Expectations Formation Mechanism (EFM) aimed to explain how the median, hence lay, forecaster predicts one-year ahead inflation rates in disparate environments.

The EFM is a combination of long-run expectations, current inflation and uncertainty.

Lay forecasters manage this EFM naively and informally, but their predictions are neither aimless nor casual. As a fixed-income earner, the median forecaster has an aversion toward underestimation and she predicts accordingly, being happy to achieve satisfactory outcomes in the sense of H. Simon. The EFM has time-varying weights because, due to the fixed-income, the higher the level of inflation, the higher the possibility of unsatisfactory forecasts, i.e. to suffer from underestimation. A tuning of the weighting scheme is needed.

Albeit unintentionally and occasionally, lay forecasters' predictions may turn out to be optimal. In fact, the EFM nests several existing expectations schemes (naïve, optimally unbiased and optimally biased). It also raises a number of testable predictions.

Working with the data gathered by the University of Michigan Survey of Consumers, the evidence supports the proposed behavioral model, shedding some light on several episodes of U.S. inflation and on consensus dynamics.

Amid particularly challenging periods – such as, e.g., the Great Inflation/Disinflation of the “80s - lay forecasters' expectations turn out to be often sub-optimal.

When inflation level and uncertainty shrink - as, e.g., in the Great Moderation - agents opt for an EFM taking into account only inflation levels. Possibly because of the favorable environment, agents' informal way to set the parameters leads to a weighting scheme compatible with optimally unbiased predictions.

When inflation is low and uncertainty is high - as, e.g., in the last two decades - lay forecasters' instinctive calibration leads to form optimally biased predictions.

The FED's implicit, but widely publicized, move to have a target long-run inflation rate seems to have induced agents to put a higher weight on their reference inflation rate and to give less weight to the other two variables included in their EFM.

The adoption of optimally unbiased expectations throughout the GM is congruent with the Atkeson-Ohanian (2001) random walk model: when inflation is low and stable the Phillips curve flattens and short-run expectations are only based on inflation levels.

Predictions grounded on trend inflation (higher than current inflation) and uncertainty (which was not small) may explain the 2009-2011 missing disinflation. By the same token, the reduction in

uncertainty side-by-side with the persistently strong reference to the long-run expected inflation may explain the lack of reflation expected after 2011.

Last but not least, it results that since mid-2016 forecasters have completely eliminated current inflation from their EFM. A choice that makes the current EFM as historically rare as intriguing from several standpoints. Virtually all the observed negative forecasting errors ($\pi_{t+12}^e - \pi_{t+12}$) seem to be induced by rapid changes in inflation and/or by deflationary shocks such as those of 2001 and 2008. In these two latter episodes a stronger weight on the trend inflation would have avoided the recorded underestimation. Side-by-side with the FED's inflation targeting strategy, this may have convinced lay forecasters to increase the link between short-run and long-run expectations. In other terms, fixed-income agents may have eventually learnt to disregard the current inflation when predicting inflation for the next year.

References

- Akerlof G.A., Dickens W.T. and G.L. Perry. 1996. “The Macroeconomics of Low Inflation” *Brookings Papers on Economic Activity*, 1, pp. 1-59.
- Armantier, O, W Bruine de Bruin, G Potter, G Topa, W. van der Klaauw, and B. Zafar. 2013. “Measuring Inflation Expectations”, *Annual Review of Economics* 5: 273–301.
- Atkeson, A. and L.E. Ohanian. 2001. “Are Phillips Curves Useful for Forecasting Inflation?” *Federal Reserve Bank of Minneapolis Quarterly Review* 25(1): 2–11. Available at <http://www.minneapolisfed.org/research/QR/QR2511.pdf>.
- Ball, L. 1992. Why does High Inflation Raise Inflation Uncertainty? *Journal of Monetary Economics* 29, 371-388.
- Ball, L., N. Mankiw, and D. Romer. 1988. “The New Keynesian Economics and the Output-Inflation Trade-off,” *Brookings Papers on Economic Activity*, 19, 1–82.
- Bernanke, B. 2010. “The Economic Outlook and Monetary Policy.” Speech, Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming, August 27. <http://www.federalreserve.gov/newsevents/speech/bernanke20100827a.htm>
- Binder, C.C. 2017. “Measuring uncertainty based on rounding: New method and application to inflation expectations”, *Journal of Monetary Economics*, 90, 1-12.
- Blanchflower, D. and R. Kelly. 2008. “Macroeconomic literacy, numeracy and the implications for monetary policy,” *Bank of England Working Paper*.
- Bloom, N. 2014. “Fluctuations in Uncertainty”, *Journal of Economic Perspectives* Vol. 28, Number 2, pp. 153–176.
- Bollerslev, T. 1986. “Generalized Autoregressive Conditional Heteroskedasticity.” *Journal of Econometrics*, 31, 307–27.
- Bordo, M. and A. Orphanides. 2013. *The great inflation: The rebirth of modern central banking*, Chicago: University of Chicago
- Bovi, M. (2014) “Are the representative agent’s beliefs based on efficient econometric models?”, *Journal of Economic Dynamics and Control*, 37, 3, pp. 633–648.
- Brock, W.A. and C.H. Hommes. 1997. A rational route to randomness. *Econometrica* 65, 1059–1160.
- Bullard, J.B. 2013. “CPI vs. PCE inflation: choosing a standard measure,” *The Regional Economist*, Federal Reserve Bank of St. Louis, July.
- Capistrán, C. 2008. Bias in federal reserve inflation forecasts: Is the federal reserve irrational or just cautious? *Journal of Monetary Economics*, 55(8):1415-27.
- Capistrán, C. and A. Timmermann. 2009. “Disagreement and Biases in Inflation Expectations,” *Journal of Money, Credit and Banking*, vol. 41(2-3), pp. 365-396.

- Coibion, O. and Y. Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116-159.
- Coibion, O. and Y. Gorodnichenko. 2015. "Is The Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7 (1), 197-232.
- Cukierman, A. and P. Wachtel. 1979. "Differential Inflationary Expectations and the Variability of the Rate of Inflation: Theory and Evidence" *American Economic Review* 69(4), 595-609.
- Curtin, R. 1996. "Procedure to Estimate Price Expectations", Manuscript, University of Michigan Survey Research Center .
- Curtin, R. 2007. "What U.S. Consumers Know About Economic Conditions", Technical report, University of Michigan.
- Evans, G. and G. Ramey (2006) "Adaptive expectations, underparameterization and the Lucas critique", *Journal of Monetary Economics*, 53, 2, pp. 249-264
- Evans, G.W. and S. Honkapohja. 2001. *Learning and Expectations in Macroeconomics*, Princeton University Press, Princeton
- Friedrich, C. 2016. Global inflation dynamics in the post-crisis period: What explains the puzzles? *Economics Letters*, 142(2):31–34
- Fuster, A., D. Laibson, and M. Brock. 2010. "Natural Expectations and Macroeconomic Fluctuations." *Journal of Economic Perspectives* 24 (4): 67-84.
- Goodfriend, M. 2003. "Inflation Targeting in the United States?" *NBER WP*, No. 9981
- Gorodnichenko, Y. and O. Coibion. 2012. "What can survey forecasts tell us about informational rigidities?", *Journal of Political Economy* 120, 116-159.
- Hamilton, J.D. 1994. *Time Series Analysis*. Princeton University Press, Princeton.
- Hendricks, W.E., and L.M. Kahn. 1985. *Wage Indexation in the United States: Cola or Uncola?* Ballinger Publishing Co.
- Holland, A. S. 1995. "Inflation and Wage Indexation in the Postwar United States." *Review of Economics and Statistics* 77(1): 172–76.
- International Monetary Fund (IMF). 2013. "The Dog That Didn't Bark: Has Inflation Been Muzzled or Was It Just Sleeping?" In *World Economic Outlook*, 79–96. Washington, DC, April.
- Kirman, A.P. (2006), "Demand Theory and General Equilibrium: From Explanation to Introspection, a Journey Down the Wrong Road," *History of Political Economy*, 38, 246-280.
- Leduc, S and K. Sill. 2013. "Expectations and Economic Fluctuations: An Analysis Using Survey Data" *Review of Economics and Statistics*, Vol. 95, No. 4:1352-1367
- Lombardelli, C. and J. Saleheen. 2003. Public Expectations of UK Inflation. *Bank of England Quarterly Bulletin* 43(3), 281-290.

- Malmendier, U. and S. Nagel. 2016. “Learning from Inflation Experiences”. *Quarterly Journal of Economics*; 131 (1): 53-87. doi: 10.1093/qje/qjv037
- Mishkin, F. 2004. “Why the Federal Reserve Should Adopt Inflation Targeting” *International Finance* 7:1, pp. 117–127
- Patton, A. and A. Timmermann. 2007. “Properties of Optimal Forecasts under Asymmetric Loss and Nonlinearity.” *Journal of Econometrics*, 140, 884–918.
- Romer, C.D. and D.H. Romer. 2002. “The Evolution of Economic Understanding and Postwar Stabilization Policy,” in *Rethinking Stabilization Policy*. Kansas City, MO: Federal Reserve Bank of Kansas City, pp.11-78
- Sargent, T. 1999. *The Conquest of American Inflation*. Princeton University Press. Princeton
- Shafir, E. P. Diamond, and A.Tversky. 1997. “Money Illusion”, *Quarterly Journal of Economics* 112(2): 341–74.
- Shiller, R.J. 1997. “Why Do People Dislike Inflation?” In *Reducing Inflation: Motivation and Strategy*, edited by Christina D. Romer and David H. Romer. University of Chicago Press.
- Shiller, R.J. 2000. *Irrational Exuberance*, Princeton Univ. Press
- Simon, Herbert (1957). “A Behavioral Model of Rational Choice”, in *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Social Setting*. New York: Wiley.
- Stock, J.H. and M.W. Watson. 2009. Phillips Curve Inflation Forecasts in *Understanding Inflation and the Implications for Monetary Policy, a Phillips Curve Retrospective*, Federal Reserve Bank of Boston
- Stock, J.H. and M.W. Watson. 2010. “Modeling Inflation After the Crisis” NBER Working Paper, n. 16488.
- Trehan, B. 2015. “Survey Measures of Expected Inflation and the Inflation Process,” *Journal of Money, Credit and Banking*, 47: 207–222. doi:10.1111/jmcb.12174
- Whitehouse, E. 2009. “Pensions, Purchasing-Power Risk, Inflation and Indexation”, *OECD Social, Employment and Migration Working Papers*, 77, OECD Publishing, Paris. <http://dx.doi.org/10.1787/227182142567>.
- Yellen, J.L. 2017. “Inflation, uncertainty, and monetary policy”, *Business Economics*, 52: 194. <https://doi.org/10.1057/s11369-017-0057-x>
- Zellner, A. 1986. “Bayesian estimation and prediction using asymmetric loss functions”, *Journal of the American Statistical Association*, 81, 446-451.

Appendix A. Robustness Check: Consumer Price Index (CPI) inflation

CPI inflation is defined as the annual rate of change $100[\ln(P_t/P_{t-12})]$, where P is the monthly Consumer price index. Data are taken from the FRED database and we limit the sample period to that used for PCE inflation (1960m1-2017m10).

As per PCE inflation, we use headline inflation and final revision data (cf. Section 3.1).

It is perhaps worth noticing that the CPI is used to Cost-Of-Living Adjustments (COLA), while the Federal Reserve states its goal for inflation in terms of the PCE. These indexes measure different concepts of inflation (Bullard, 2013). For instance, the CPI tends to show more inflation than the PCE. The 1979m02-2017m10 average for the former is 3.4%, the latter shows a mean of 2.8%. This said, their correlation coefficient is as high as 0.98.

We reiterate the same exercises of the main text with the obvious exception that in this Appendix π_t is the CPI inflation and that the conditional variance refers to the GARCH model of CPI inflation estimated in Appendix B. Tables and figures replicate the corresponding ones of the main text and, to ease comparisons, each of them has the same heading and number of its correspondent in the main text. Only a prefix “A” is added.

Figures A4.1.1 and A4.1.2 collect the evidence relative to, respectively, Kalman filter and rolling windows. The former is based on the system 4.1.1 of the main text (Section 4.1). It is performed estimating the parameters via ML and initializing the recursion using as prior mean and variance the coefficients and the coefficient covariance matrix deriving from the OLS estimation of equation 2.3.1.²⁶ Rolling regressions (with window size of 24-months) are based on the signal equation of the mentioned system 4.1.1.

FIGURE A4.1.1 and A4.1.2 ABOUT HERE

The big picture substantially mirrors that obtained for the PCE inflation.

Comparing the results obtained using the Kalman and the rolling windows approach, it results that the informal lay forecasters updating scheme resembles more to this latter than the former. The RMSE stemming from the Kalman filter is 0.34 (MAE=0.24), whereas the RMSE from rolling windows estimations is 0.28 (MAE=0.20). The correlation between fitted and actual expectations computed from the Kalman and the rolling approach is, respectively, 0.93, and 0.95, suggesting that the proposed EFM closely mimics the evolution of the median forecaster’s expectations. These statistics

²⁶ We have also initialized the recursions using diffuse priors. The fit marginally worsens, but results remain substantially the same.

sustain the findings obtained in the case of the PCE inflation. Also, the evidence points out that the fit slightly improves when we use the PCE inflation. Thus, this latter seems to be closer than CPI inflation to the target of agents' forecasts. The evidence indicates an admittedly minor refinement,²⁷ but it must be considered that, as mentioned, PCE and CPI inflation are strongly correlated. These findings are in a sense in line with the conclusions of an extensive analysis by the FOMC which, prior to 2000, focused on CPI inflation and then changed to PCE inflation (Bullard, 2013). Specifically, the decision was grounded on the fact *i*) that the expenditure weights in the PCE can change as people substitute away from some goods and services toward others, and *ii*) that the PCE includes more comprehensive coverage of goods and services. It can be finally observed that there is evidence that the lay forecasters' EFM is influenced by the FED and that this latter targets PCE inflation. Hence, there are reasons to believe that agents think about PCE inflation rather than CPI inflation when they predict price dynamics.

Table A4.1.1 collects the results of the long-run horserace comparing the three proxies for uncertainty under scrutiny.

Table A4.1.1 ABOUT HERE

Consistently with the choice of estimating time-varying parameters, the RMSE obtained from the full-sample regression of the signal equation in 4.1.1 is almost the double with respect to that obtained via rolling windows: 0.52. More importantly, Table A4.1.1 supports the superiority of DK_U as proxy for the uncertainty to be included in the EFM.

Thus, as in the main text, we offer further evidence focusing on rolling estimations with DK_U.

Figure 4.2.1 reports evidence on the reaction of agents to peculiar episodes.

FIGURE A4.2.1 HERE

Figure A4.2.1 quite closely mimics its correspondent in the main text (Section 4.2). Just to mention an important statistic that we claim to be strongly negative, the correlation between α_t and β_t is -0.70 to be compared with -0.77 of the PCE inflation case.

FIGURE A4.3.1 HERE

²⁷ In the case of PCE inflation the RMSE stemming from rolling windows estimations is 0.27 (MAE=0.19) and the correlation between fitted and actual expectations is 0.96 (Section 4).

Figure A4.3.1 shows that relatively more heterogeneous outcomes show up when dealing with the question of whether agents' expectations are sub-optimal or optimally (un)biased. For instance, whereas Figure 4.3.1 in the main text shows that the EFM-L is the almost always the preferred weighting scheme from the mid-2004, Figure A4.3.1 underlines that the persistent use of this EFM showed up one year later. In any case, the broad picture is substantially similar to Figure 4.3.1, so we have replicated the same cut-off proposed for the PCE inflation.

All considered, the evidence obtained using CPI inflation gives further support to our conclusions.

Appendix B. The Conditional Setting

In this Appendix we report the estimate of the AR-GARCH model of inflation behind the conditional variance used as proxy for uncertainty in the main text and in Appendix A.²⁸

Following Bollerslev (1986), we have modeled inflation as an AR(12)-GARCH(1,1):²⁹

$$\begin{aligned}\pi_{t+1,t} &= \phi_0 + \sum_{p=1}^{12} \phi_p \pi_{t-p,t-p} + \varepsilon_{t+1} \\ \sigma_{t+1,t}^2 &= \omega + \alpha_1 \varepsilon_t^2 + \beta_1 \sigma_{t,t-1}^2\end{aligned}\tag{B.1}$$

Since the Jarque-Bera statistic have shown that residuals are leptokurtic, we have modeled the thick tail in the residuals assuming that the errors follow a Student's t-distribution.³⁰

Table B.1 ABOUT HERE

Section 2 introduced the optimality condition $\mu_{t+12,t} = \alpha_t \pi_{LR}^e + (1-\alpha_t) \pi_t$. This weighted average can be derived assuming that the trend inflation, π_{LR}^e , is equal³¹ to the long-run mean, $\phi_0 / (1 - \sum_{p=1}^{12} \phi_p)$, of inflation:³² $\phi_0 = (1 - \sum_{p=1}^{12} \phi_p) \pi_{LR}^e$.

Data sustain this hypothesis. Reiterating the one-step ahead prediction from the mean equation of the GARCH B.1 virtually replicates current inflation. Its full sample correlation with π_t is larger than 0.99 for both PCE and CPI inflation. As the forecast horizon enlarges, the forecast tends to π_{LR}^e . The equation for the conditional mean of PCE inflation detects a long-run value of 3.2%, which is exactly the long-run expected inflation emerging in persistently ideal environments such as the last three decades - i.e. not considering the Great Inflation/Disinflation - of the sample under scrutiny (Figure 3.2.1).³³ The value for CPI inflation is a bit larger: 3.4%.

²⁸ The AR-GARCH estimate seems to imply that we assume that lay forecasters know future inflation – the model is estimated using full sample data. However, in the present setting we think of the conditional variance as just a possible quantification – moreover not the best one - of the uncertainty surrounding the forecasting exercise (Section 2 and 3).

²⁹ With quarterly data, Bollerslev (1986) found that an AR(4) model for the mean of inflation and a GARCH(1,1) model for the conditional variance provide a good description of US inflation. Following a general-to-specific procedure, we have included twelve lags of inflation in the mean equation. Then, we have eliminated the insignificant coefficients in order to minimize usual information criteria (AIC, SBIC).

³⁰ Using the Gaussian distribution or Generalized Error Distribution (tail parameter 1.5) results are virtually the same.

³¹ Moreover, this is in line with the evidence that, in their naïve updating, consumers look at rolling windows (Section 4).

³² ADF-tests (only constant) strongly reject the null of unit root for CPI inflation and marginally for PCE inflation (P-value 0.1). ADF-tests reject this null for both metrics in the last 30 years (after the Great Disinflation). The sum the estimated coefficients of the lagged inflation in the mean equation is lower than one for both inflations. Wald-tests reject the restriction that they sum to one (P-value lower than 0.01).

³³ In the early ‘90s inflation was above 4% (Figure 3.1.1).

Confirming existing evidence, finally, Table B.1 shows that the conditional variance follows a quite persistent pattern for both metrics.

Table 2.2.1. Inflation Expectations Formation Mechanisms (EFM) and Inflation Environments. A Taxonomy

		Inflation Uncertainty	
		Low	High
Inflation Level	Low	<u>EFM-A</u> $\alpha_t \rightarrow 1; \beta_t, \partial_t \rightarrow 0$ If $(\alpha_t + \beta_t) = 1$ and $\partial_t = 0$, then, optimal unbiased expectations	<u>EFM-D</u> $\alpha_t \rightarrow 1; \beta_t \rightarrow 0; \partial_t > 0$ If $(\alpha_t + \beta_t) = 1$ and $\partial_t > 0$, then, optimal biased expectations
	High	<u>EFM-B</u> $\alpha_t \rightarrow 0; \beta_t \rightarrow 1; \partial_t > 0$ If $(\alpha_t + \beta_t) = 1$ and $\partial_t > 0$, then, optimal biased expectations	<u>EFM-C</u> EFM-B with smaller α_t and larger β_t and ∂_t If $(\alpha_t + \beta_t) = 1$ and $\partial_t > 0$, then, optimal biased expectations

Note: Time varying coefficients α_t, β_t , and ∂_t refers to Equation 2.1.1: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \partial_t \sigma_t^2$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectations, π_t =current inflation; σ_t^2 =uncertainty. $\partial_t > 0$ implies preference toward overestimation.

Table 4.1.1. Full-sample (1979m02-2017m10) EFM weighting scheme

	α	β	δ		Log-Lik	RMSE	
DK_U	0.41***	0.41***	0.11***		-382.3	0.55	
BI	0.58***	0.31***		0.01*	-483.2	0.68	
CV	0.62***	0.33***			0.04**	-478.5	0.68

Note: Coefficients refers to the time invariant version of Equation 2.1.1: $\pi_{t+12}^e = \alpha \pi_{LR}^e + \beta \pi_t + \delta \sigma_t^2$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =PCE inflation; σ_t^2 =uncertainty=DK_U, BI, CV. DK_U is the share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI is the Binder's index taking into account the number of respondents, defined uncertain, indicating round numbers. CV is the conditional variance of an AR(12)-GARCH(1,1) model of inflation (Appendix B). The last column reports the p-values of the Wald test on the restriction $\alpha + \beta = 1$. $\delta > 0$ implies preference toward overestimation (Section 2). *** significant at 1% (** 5%, * 10%). HAC-robust standard errors.

Table A4.1.1. Full-sample (1979m02-2017m10) EFM weighting scheme

	α	β	δ		Log-Lik	RMSE	
DK_U	0.47***	0.32***	0.09***		-372.2	0.54	
BI	0.52***	0.30***		0.01**	-456.9	0.65	
CV	0.60***	0.30***			0.02***	-450.3	0.64

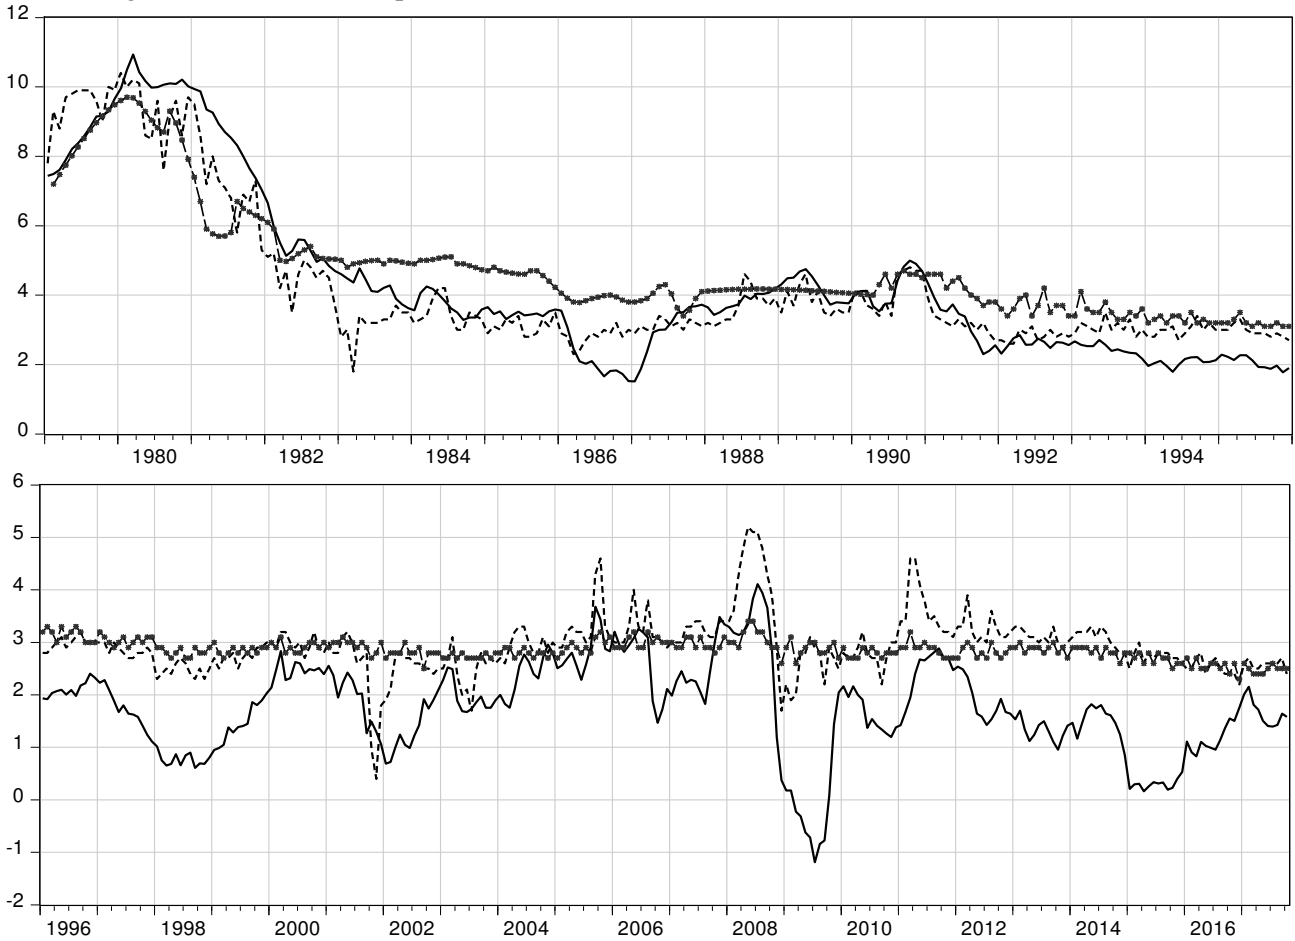
Note: Coefficients refers to the time invariant version of Equation 2.1.1: $\pi_{t+12}^e = \alpha \pi_{LR}^e + \beta \pi_t + \delta \sigma_t^2$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =CPI inflation; σ_t^2 =uncertainty=DK_U, BI, CV. DK_U is the share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI is the Binder's index taking into account the number of respondents, defined uncertain, indicating round numbers. CV is the conditional variance of an AR(12)-GARCH(1,1) model of inflation (Appendix B). $\delta > 0$ implies preference toward overestimation (Section 2). *** significant at 1% (** 5%, * 10%). HAC-robust standard errors.

Table B.1. AR-GARCH model of US Inflation. Sample 1960m11-2017m10

PCE Inflation			CPI Inflation		
	Mean	Variance		Mean	Variance
Const	0.03**		Const	0.05**	
ϕ_1	1.33***		ϕ_1	1.14**	
ϕ_2	-0.31***		ϕ_3	-0.24*	
ϕ_{10}	-0.03***		ϕ_4	0.13**	
			ϕ_{12}	-0.04**	
ω		0.00*	ω		0.00
α_1		0.07**	α_1		0.07**
β_1		0.89***	β_1		0.91**
Adj. R2 = 0.99			Adj. R2 = 0.98		
T-Dist (DoF) 7.1***			T-Dist (DoF) 6.0***		
Sum of Sq. Res = 36.2			Sum of Sq. Res = 79.8		
Log-Lik. = 87.5			Log-Lik. = -155.5		
T=684			T=684		


Note: Student's t distribution (BFGS / Marquardt steps), Coefficient covariance computed using QML sandwich with observed Hessian; Presample variance: backcast (parameter=0.7). *** significant at 1% (** 5%, * 10%). HAC-robust standard errors.

Fig. 3.1.1. Actual and Expected Inflation




Note: Solid line: actual inflation (annual rate of change of the monthly PCE Index). Dashed line: median expected price change next 12 months. Line with stars: long-term inflation expectations (5-10 years). Expectations are taken from the Michigan Survey of Consumers.

Fig. 3.1.2. Prediction Errors in Lay Forecasters' Expectations



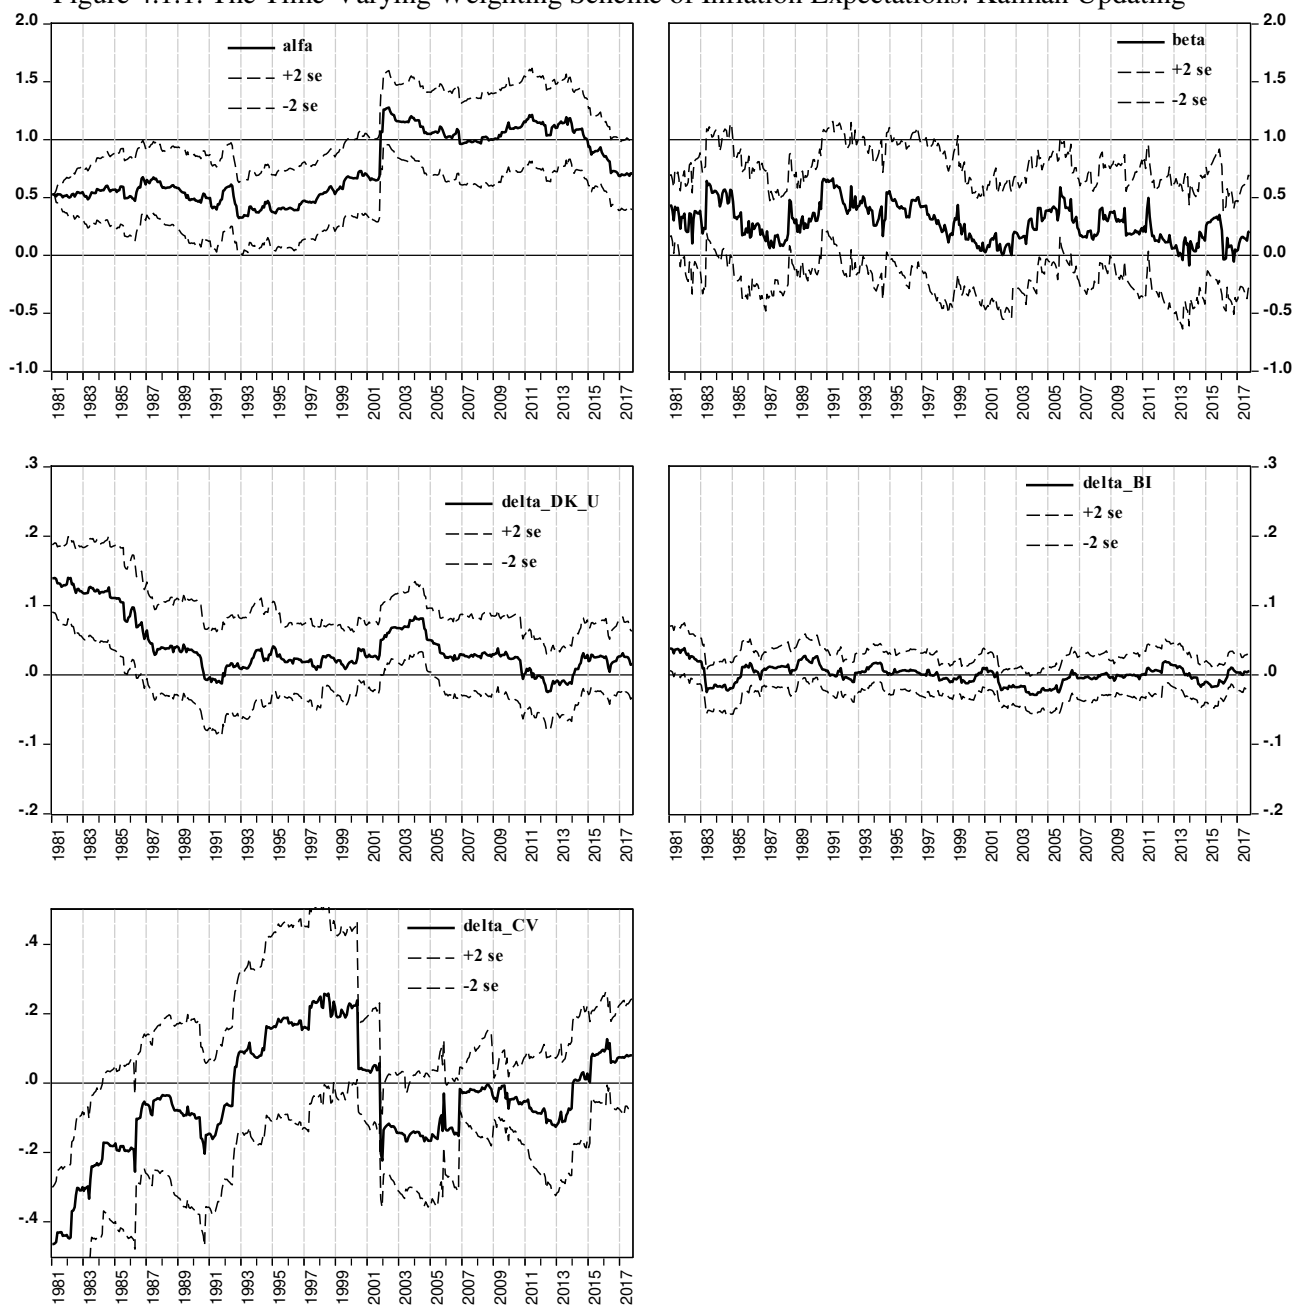
Note: Solid line = $(\pi_{t+12}^e - \pi_{t+12})$, where π_{t+12}^e =one-year ahead expectations (median); π_{t+12} =PCE inflation; Dashed line = $(\pi_{t+12}^{e,M} - \pi_{t+12})$, where $\pi_{t+12}^{e,M}$ =one-year ahead expectations (mean). Expectations are taken from the Michigan Survey of Consumers.

Figure 3.2.1. Three Proxies for Inflation Uncertainty




Note. The proxies for uncertainty are BI (solid line, rhs scale), DK_U (x-x, %, lhs scale) and CV (dashed line, %, lhs scale). BI=Binder's index. A BI equal 100 implies that all respondents indicate round numbers and, hence, that all agents are uncertain about their forecasts. DK_U=share of interviewed indicating that they expected prices to rise but replied that they didn't know how much prices would increase. CV=conditional variance of an AR(12)-GARCH(1,1) model of inflation (Appendix B).

Figure 4.1.1. The Time-Varying Weighting Scheme of Inflation Expectations. Kalman Updating




Note: The time-varying coefficients α_t , β_t , $\delta_{dk_u,t}$, $\delta_{BI,t}$, $\delta_{CV,t}$ are estimated via ML using the Kalman filter procedure and refer to the signal equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{dk_u,t} + \delta_{BI,t} + \delta_{CV,t}$, where π_{t+12}^e = one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =PCE inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI=Binder's index, which is the share of uncertain respondents, i.e. respondents indicating round numbers. CV=conditional variance of the AR(12)-GARCH(1,1) model of inflation estimated in Appendix B. $\partial_t > 0$ implies preference toward overestimation (Section 2). HAC-robust standard errors.

Figure 4.1.2. The Time-Varying Weighting Scheme of Inflation Expectations. Rolling Updating



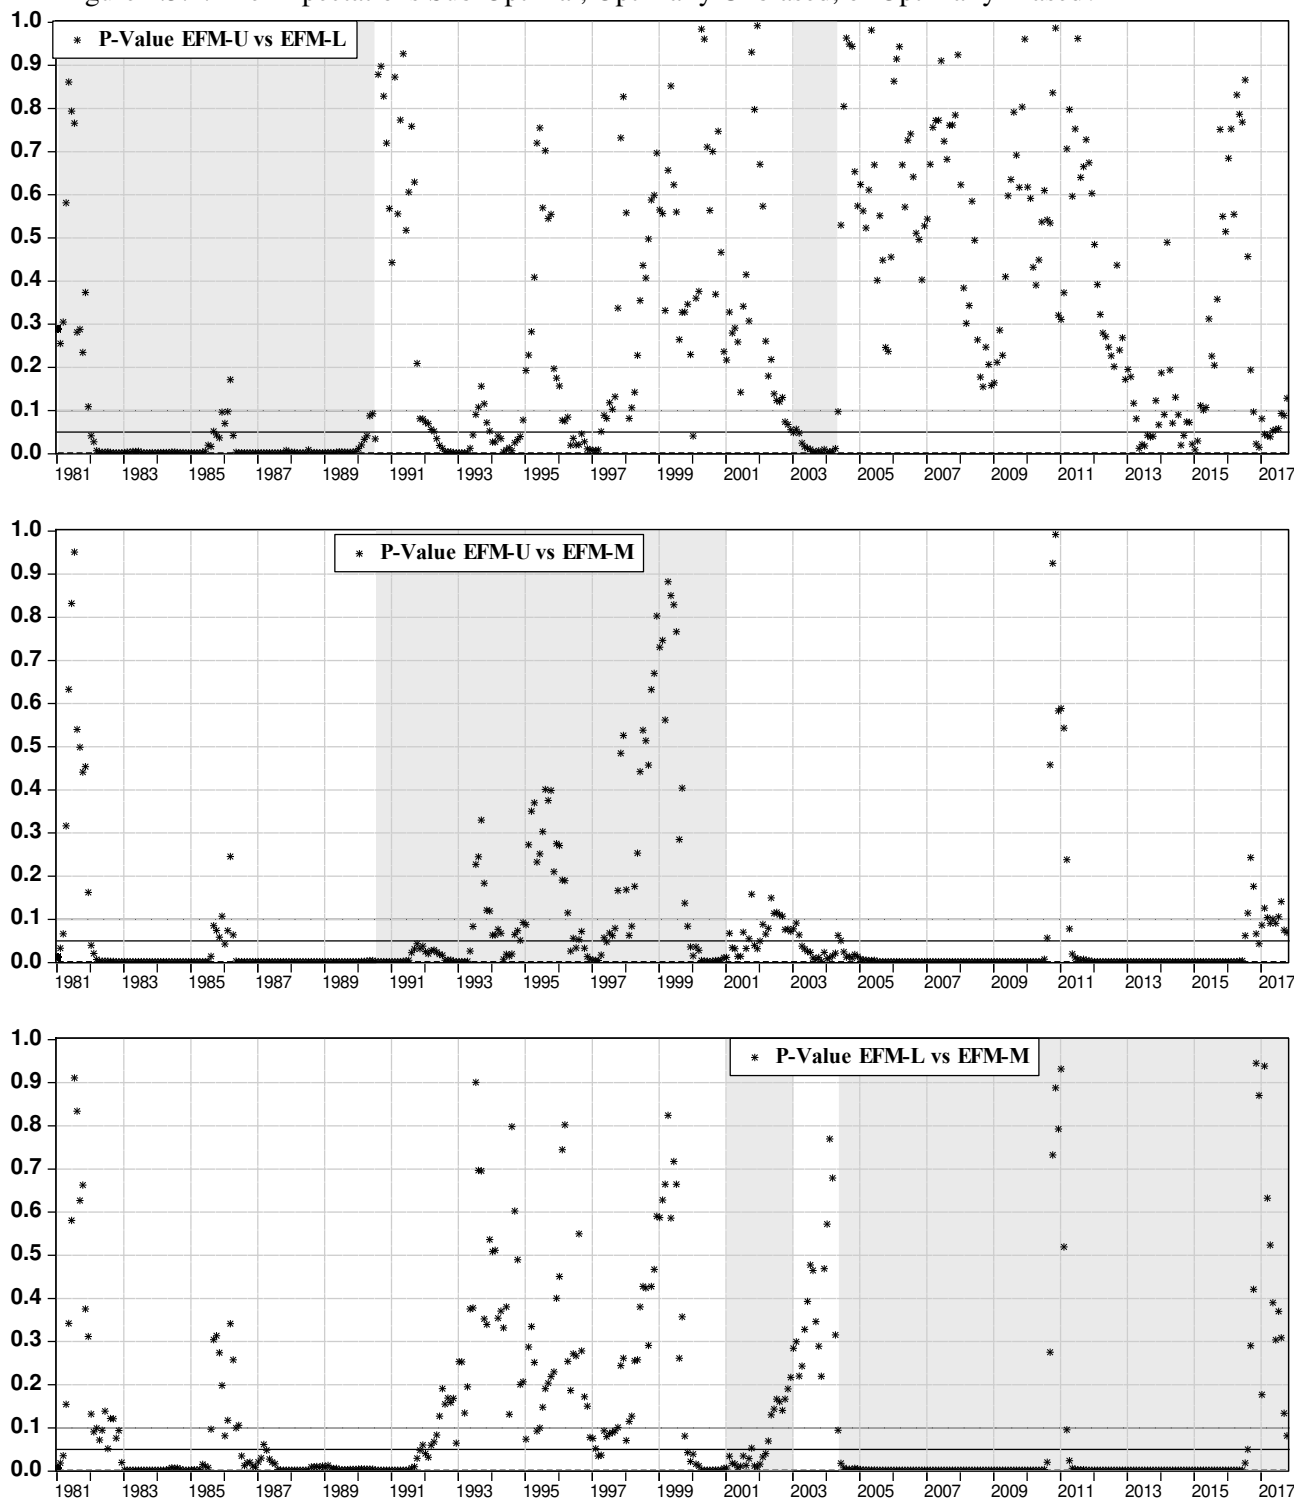
Note. Time varying coefficients refer to the equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{dk,u,t} + \delta_{BI,t} + \delta_{CV,t}$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =PCE inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI=Binder's index, which is the share of uncertain respondents, i.e. respondents indicating round numbers. CV=conditional variance of the AR(12)-GARCH(1,1) model of inflation estimated in Appendix B. $\partial_t > 0$ implies preference toward overestimation (Section 2). Coefficients are estimated via two-year rolling regressions. The sample of the first rolling regression is 1979m2-1981m1 and the relative coefficients are displayed in 1981m1. Likewise for the subsequent 441 regressions. Dashed lines are 95% confidence bands. HAC-robust standard errors.

Figure 4.2.1. Reactions of Lay Forecasters' EFM During Some Salient Episode



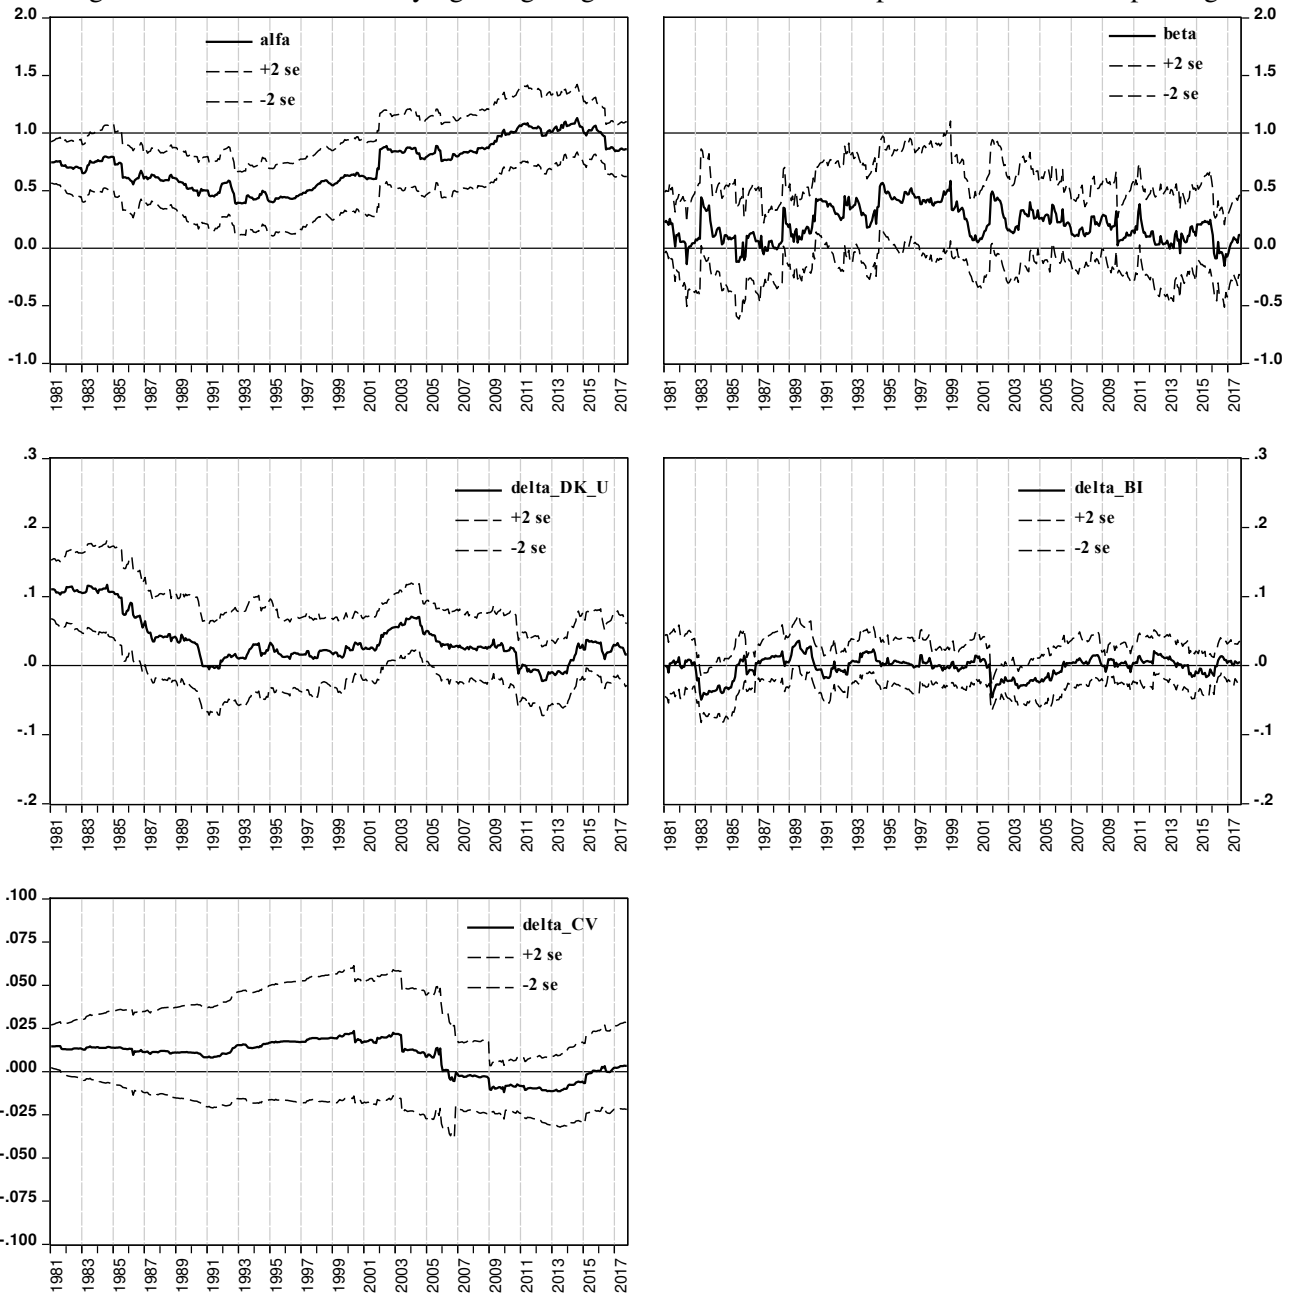
Note. Rolling coefficients α_t , β_t , and ∂_t (rhs scale) refer to the equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \partial_{DK_U,t}$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =PCE inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. $\partial_t > 0$ implies preference toward overestimation (Section 2). The sample of the first rolling regression is 1979m2-1981m1 and the relative coefficients are displayed in the corresponding panel in 1981m1. Likewise for the subsequent 441 regressions. Histograms quantify (lhs-scale) the size of the t-stat of the corresponding coefficient. Dashed lines are the 95% confidence bands. HAC-robust standard errors.

Figure 4.3.1. Are Expectations Sub-Optimal, Optimally Unbiased, or Optimally Biased?




Note. Models are: Unrestricted=EFM-U: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \partial_{DKU,t} + \varepsilon_t^U$; Linex-restricted=EFM-L: $\alpha_t = 1 - \beta_t$; MSE-restricted=EFM-M: $\alpha_t = 1 - \beta_t$, $\partial_{DKU,t} = 0$. All the model is estimated with rolling regressions (two-year window size). π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =PCE inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. Stars indicate the P-value of F-tests comparing the sums of squared residuals of nested models. In the Upper panel, e.g., the first star quantifies the P-value of F-test relative to the horseshoe EFM-U vs EFM-L over the period 1979m2 - 1981m1, the second star refers to the period 1979m3 - 1981m2 and so on. In the lower Panel the unrestricted model is the Linex and the only restriction is $\delta_i = 0$. Clearly, if EFM-U turns out to be better than both EFM-L and EFM-M, the last race (EFM-L vs EFM-M) is redundant. The shaded areas tentatively highlight periods in which the best model is EFM-U (Upper Panel), or EFM-M (Middle Panel), or EFM-L (Lower Panel). HAC-robust standard errors.

Figure A4.1.1. The Time-Varying Weighting Scheme of Inflation Expectations. Kalman Updating




Note: The time-varying coefficients α_t , β_t , $\delta_{DK_U,t}$, $\delta_{BI,t}$, $\delta_{CV,t}$ are estimated via ML using the Kalman filter procedure and refer to the signal equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{DK_U,t} + \delta_{BI,t} + \delta_{CV,t}$, where $\pi_{t+12}^e =$ one-year ahead expectations; $\pi_{LR}^e =$ five-ten years ahead expectation, $\pi_t =$ CPI inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI=Binder's index, which is the share of uncertain respondents, i.e. respondents indicating round numbers. CV=conditional variance of the AR(12)-GARCH(1,1) model of inflation estimated in Appendix B. $\partial_t > 0$ implies preference toward overestimation (Section 2). HAC-robust standard errors.

Figure A4.1.2. The Time-Varying Weighting Scheme of Inflation Expectations. Rolling Updating



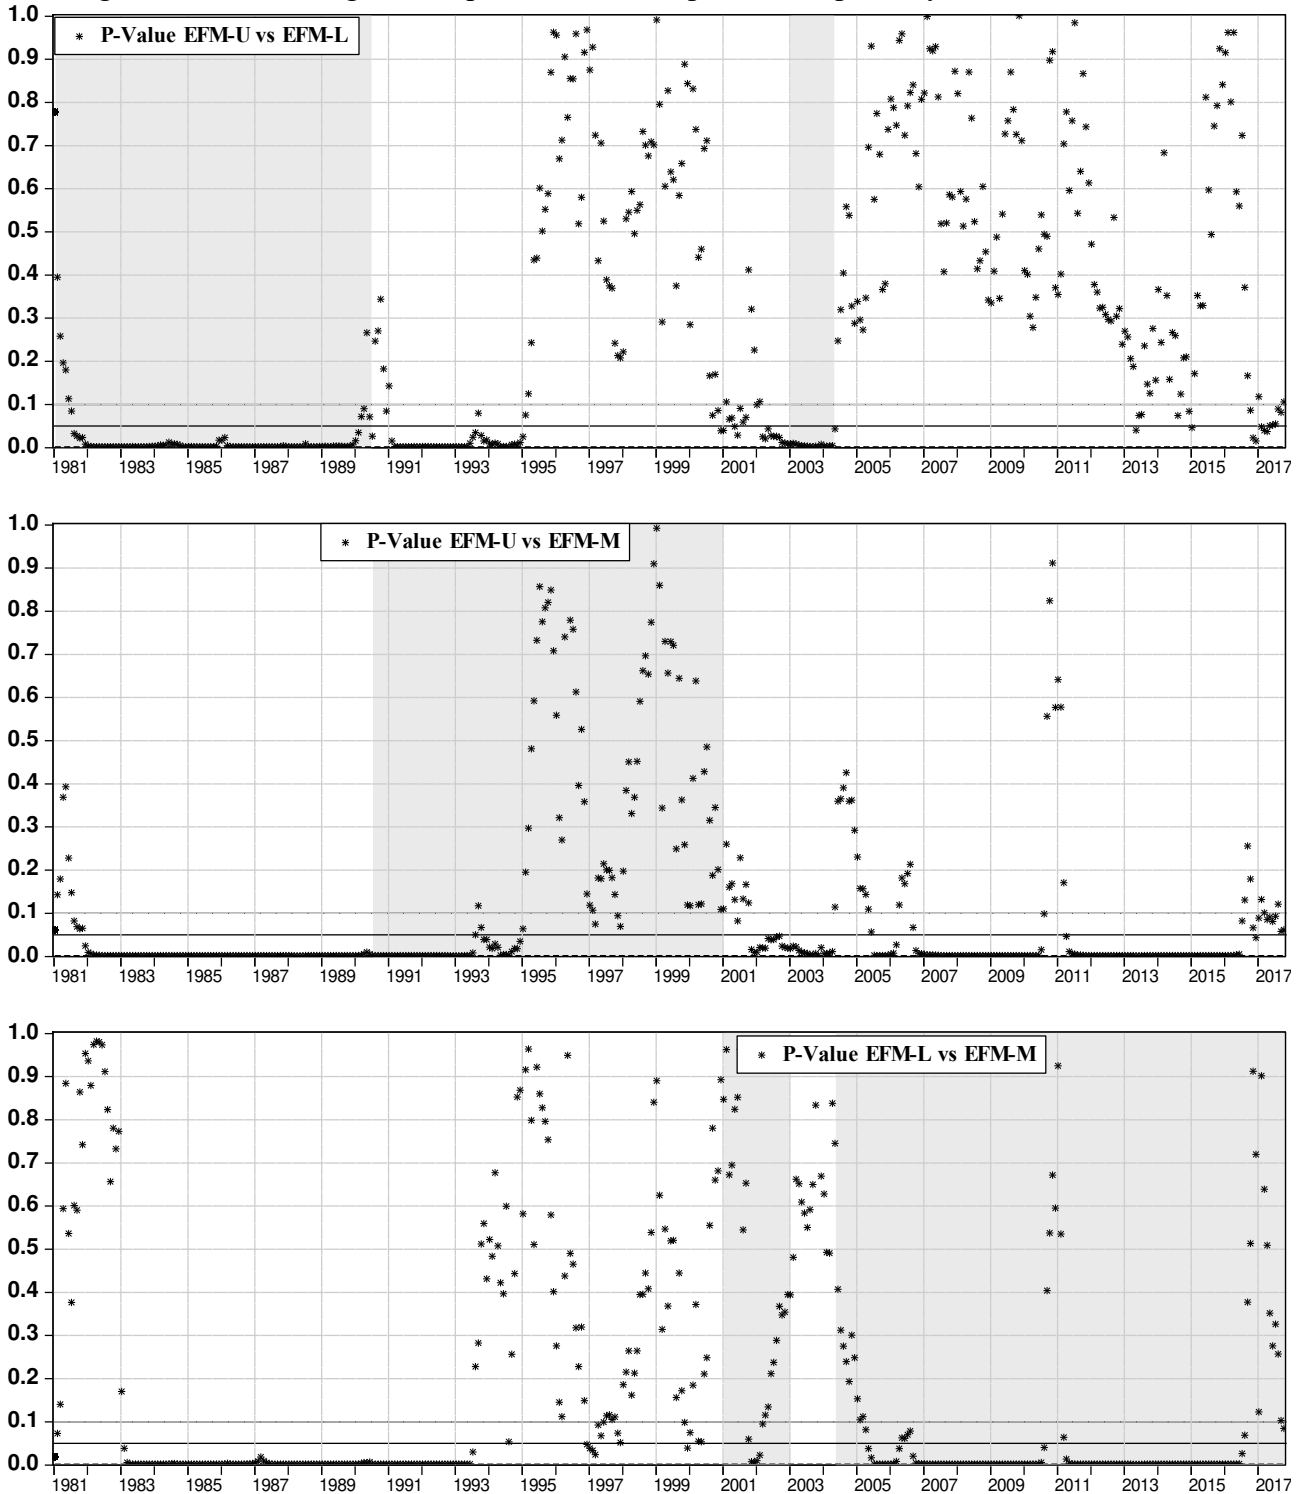
Note. Time varying coefficients refer to the equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{dk_u,t} + \delta_{BI,t} + \delta_{CV,t}$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =CPI inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. BI=Binder's index, which is the share of uncertain respondents, i.e. respondents indicating round numbers. CV=conditional variance of the AR(12)-GARCH(1,1) model of inflation estimated in Appendix B. $\partial_t > 0$ implies preference toward overestimation (Section 2). Coefficients are estimated via two-year rolling regressions. The sample of the first rolling regression is 1979m2-1981m1 and the relative coefficients are displayed in 1981m1. Likewise for the subsequent 441 regressions. Dashed lines are 95% confidence bands. HAC-robust standard errors.

Figure A4.2.1. Reactions of Lay Forecasters' EFM During Some Salient Episode



Note. Rolling coefficients α_t , β_t , and δ_t (rhs scale) refer to the equation: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta_t \pi_t + \delta_{DK_U,t}$, where π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =CPI inflation; DK_U=share of interviewed indicating that they expected prices to rise, but replied that they didn't know how much prices would increase. $\delta_t > 0$ implies preference toward overestimation (Section 2). The sample of the first rolling regression is 1979m2-1981m1 and the relative coefficients are displayed in the corresponding panel in 1981m1. Likewise for the subsequent 441 regressions. Histograms quantify (lhs-scale) the size of the t-stat of the corresponding coefficient. Dashed lines are the 95% confidence bands. HAC-robust standard errors.

Figure A4.3.1. Are Agents' Expectations Sub-Optimal or Optimally (un)biased?



Note. Models are: Unrestricted=EFM-U: $\pi_{t+12}^e = \alpha_t \pi_{LR}^e + \beta \pi_t + \partial_{DK_U,t} + \varepsilon_t^U$; Linex-restricted=EFM-L: $\alpha_t = 1 - \beta$; MSE-restricted=EFM-M: $\alpha_t = 1 - \beta$, $\partial_{DK_U,t} = 0$. All the model is estimated with rolling regressions (two-year window size). π_{t+12}^e =one-year ahead expectations; π_{LR}^e = five-ten years ahead expectation, π_t =CPI inflation; DK_U=share of interviewed indicating that they expected prices to rise but replied that they didn't know how much prices would increase. Stars indicate the P-value of F-tests comparing the sums of squared residuals of nested models. In the Upper panel, e.g., the first star quantifies the P-value of F-test relative to the horseshoe EFM-U vs EFM-L over the period 1979m2 - 1981m1, the second star refers to the period 1979m3 - 1981m2 and so on. In the lower Panel the unrestricted model is the Linex and the only restriction is $\delta=0$. Clearly, if EFM-U turns out to be better than both EFM-L and EFM-M, the last race (EFM-L vs EFM-M) is redundant. The shaded areas tentatively highlight periods in which the best model is EFM-U (Upper Panel), or EFM-M (Middle Panel), or EFM-L (Lower Panel). HAC-robust standard errors.