The Taming of the Skew: Asymmetric Inflation Risk and Monetary Policy*

Andrea De Polis† Leonardo Melosi‡ Ivan Petrella§

This draft: September 2024

Abstract

Inflation risk in U.S. data varies considerably over time and has often been asymmetric. A model incorporating time-varying asymmetric risk achieves better forecasting accuracy than a state-of-the-art symmetric model, providing results comparable to the performance of professional forecasters. The optimal monetary strategy calls for the Central Bank to counterbalance shifts in the direction of inflation risk by adjusting its inflation target –a strategy we term *risk-adjusted inflation targeting* (RAIT). By adjusting the modal inflation scenario, the Central Bank can realign average inflation with its desired long-run inflation objective and effectively anchor the private sector's expectations.

Keywords: *TBC*

JEL classification: *TBC*

*We would like to thank Gianluca Benigno, François Gourio and the seminar participants at the Federal Reserve Bank of Chicago. The views in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of Chicago or any other person associated with the Federal Reserve System.

†University of Strathclyde & ESCOE. andrea.de-polis@strath.ac.uk

‡University of Warwick and Federal Reserve Bank of Chicago & CEPR. leonardo.melosi@warwick.ac.uk

§University of Warwick & CEPR. ivan.petrella@wbs.ac.uk

1 Introduction

After nearly three decades of taking the backseat, inflation has once again become a significant concern for market participants and policymakers worldwide. This concern arises not only from recent increases in inflation observed in many countries but also from a growing perception that the risk of inflation spikes might have intensified. Data from the *Survey of Professional Forecasters* (SPF), as shown in Figure 1, supports this perception by indicating a distinct shift in the perceived balance of inflation risks in recent years. Before the COVID-19 pandemic, professional forecasters anticipated a downside tilt in inflation risks. However, the most recent period reveals a pronounced shift towards upside risks. Panel (b) illustrates the consensus subjective probabilities of professional forecasters for various inflation outcome ranges in the first quarter of 2009 and the first quarter of 2023. These distributions reveal that perceived risk can at times be asymmetric. Using the distributional data from SPF projections, panel (c) presents a measure of perceived skewness, calculated as the difference between the probabilities on either side of the modal predictions. When assessing inflation risk, professional forecasters rarely consider it to be symmetric, and shifts in the perceived balance of risks can be substantial.

Despite the crucial role of risk assessment in shaping professional forecasters' views on inflation, common modeling tools are unable to estimate real-time shifts in the balance of inflation risk from the data. In this paper, we devise a model that allows for time-varying asymmetric risk in the inflation process and show that it outperforms leading – symmetric – forecasting models routinely used by forecasters. Our estimates reveal that risks to the inflation outlook vary significantly over time in U.S. data, and exhibit strong and persistent asymmetry. Based on this evidence, we then demonstrate that the optimal monetary strategy must take into account these fluctuations in the direction of inflation risk.

We begin by showing that a simple subsample analysis of U.S. inflation data reveals significant variation in the properties of inflation risk during the postwar period. Formal testing procedures fail to reject the null hypothesis of asymmetry in the inflation process, while providing strong evidence in favor of its time variation. Building on this preliminary evidence, we estimate the time-varying mean, variance, and skewness of the predictive distribution of quarterly U.S. Personal Consumption

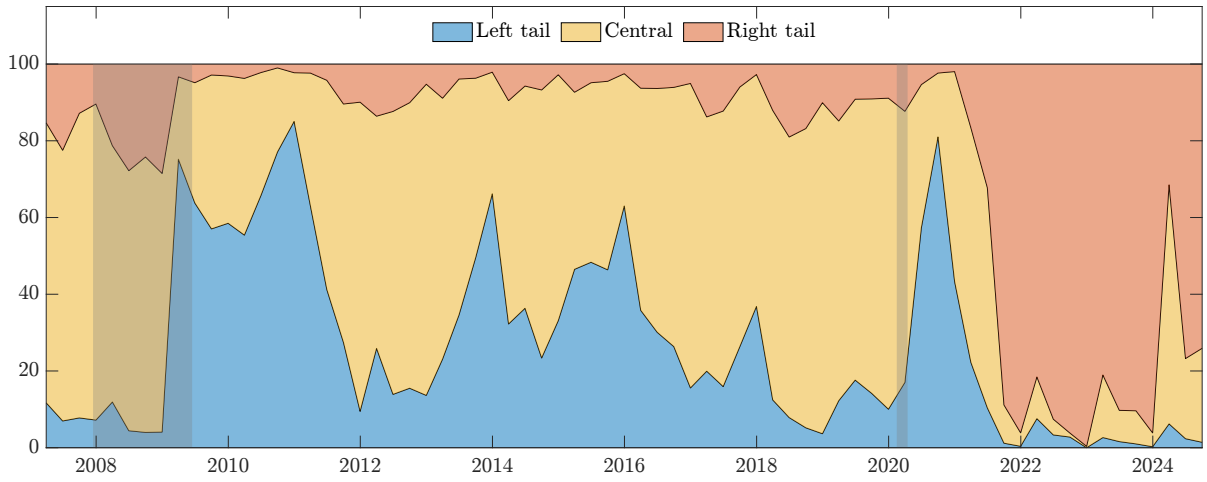

(a) SPF's predictive distribution

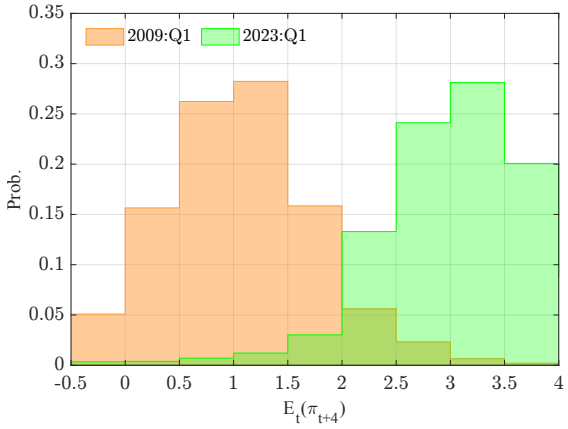

(b) Skewed forecasts

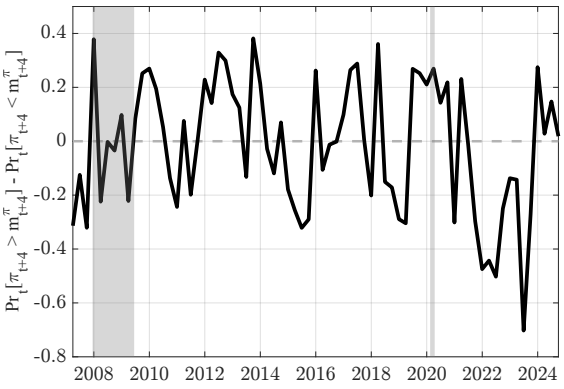

(c) SPF's implied skewness

Figure 1: Skewness in Inflation Expectations

Note: The top panel reports Survey of Professional Forecasters' interval forecasts for year end. We define left tail as the probability of inflation expectations below 1.5%, central corresponds to expectations in the $[1.5\%, 2.5\%)$ interval, whereas the right tail is defined as expectations above 2.5%. Panel (b) reports the predictive densities for 2009Q1 and 2023Q1. Panel (c) shows the implied skewness of SPF's predictive distributions. Gray shaded areas represent NBER recessions.

Expenditure excluding food and energy (core PCE) inflation. This enables us to provide a real-time assessment of the dynamics of the balance of risks to the inflation outlook. Our findings indicate that shifts in the balance of inflation risks have been substantial, persistent, and frequent throughout the postwar period.

We model the long- and short-run components of inflation's conditional moments, following a trend-cycle framework. Our estimates for the time-varying mean and volatility of core PCE inflation indicate persistent deviations of these moments from their trend components, with significant gaps emerging in the 1970s and early 1980s. In the aftermath of the Global Financial Crisis,

our model attributes the low inflation rates (the so-called *deflationary bias*) to persistent negative cyclical factors, while trend inflation remained stable around the Federal Reserve's 2% objective. This negative gap is largely driven by the persistently negative inflation skewness we observe starting in the mid-1990s. By contrast, the recent surge in inflation is marked by substantial positive skewness, with patterns and levels comparable to those seen in the early 1970s.

Accounting for time-varying asymmetry in the inflation process leads to significant improvements in out-of-sample forecasting accuracy. Specifically, we show that our model outperforms the unobserved component with stochastic volatility (UCSV) model of Stock and Watson (2002) in terms of point, density, and event predictions. We observe substantial gains in forecasting the tails of the predictive distributions, resulting in reduced forecast errors during periods of rapid and significant inflation changes. Additionally, we find that our model's event forecast accuracy is comparable to that of SPF projections, a challenging benchmark, with some improvements noted in predicting the recent surge in inflation.

We then move to examine the implications of time-varying asymmetry in the inflation process for monetary policy. We analyze optimal monetary policy in the presence of asymmetric shocks to the natural rate of interest within an otherwise standard New Keynesian model. If the Central Bank adopts a suboptimal symmetric strategy aimed at stabilizing inflation around an objective, the asymmetry in the distribution of the shocks will prevent inflation from averaging at the desired level. The asymmetry introduces skewness in inflation outcomes, distorting the long-term dynamics of inflation, the output gap, and the nominal interest rate. As a result, agents' inflation expectations will develop a persistent bias away from the Central Bank's objective.

The optimal monetary strategy must take into account time-varying asymmetry in inflation risk. Specifically, the Central Bank sets an inflation target that adjusts the modal inflation scenario to counteract the expected direction of inflation risk.[1] We refer to this optimal monetary policy strategy as Risk-Adjusted Inflation Targeting (RAIT). Our findings suggest that the Flexible Average Inflation Targeting (FAIT) approach – adopted by the Federal Reserve in 2020 – performs poorly during periods of rapid shifts in inflation risk, as seen in 2021, due to its reliance on adjusting

[1]We use the term *inflation objective* to refer to the desired level of inflation, and the term *inflation target* as a policy variable set by the central bank to minimize its quadratic loss function. The loss function is minimized by steering inflation expectations close to the inflation objective.

the target based on past inflation misses. The RAIT can be considered as a generalization of the forecasting inflation targeting proposed by Svensson and Woodford (2004) in a environment with time-varying asymmetric inflation risk.

We integrate the predictions of the New Keynesian model with our empirical findings to demonstrate that the persistent negative skew in inflation following the Great Financial Crisis would have required the Central Bank to adopt an interest rate path aligned with a modal inflation forecast exceeding the target by approximately 30 to 50 basis points. Adopting RAIT during that period could have eliminated the deflationary bias, which was a key factor behind the Federal Reserve's framework revision in 2020. In post COVID years, however, with upside risks dominating the inflation outlook, monetary policy can address the upward bias in inflation expectations by targeting average modal inflation levels slightly below the desired objective. Without this adjustment, the risk of consistently overshooting the desired inflation objective becomes a tangible concern, potentially undermining the credibility and effectiveness of monetary policy.

Literature Review The literature on inflation forecasting has primarily emphasized the importance of accounting for slow-moving trends in the data and time-varying uncertainty (see, e.g., Stock and Watson, 2002; Faust and Wright, 2013; Ascari and Sbordone, 2014). Far less research has focused on the risks of inflation or deflation.[2]

The paper most closely related to ours is Le Bihan et al. (2023), which introduces a new real-time measure of underlying inflation that accounts for time-varying changes in asymmetric risks to the inflation outlook. Their indicator is based on a multivariate regime-switching framework jointly estimated using disaggregated sub-components of the Euro Area's harmonized index of consumer prices (HICP). However, because they use disaggregated data in their estimation, their sample period is shorter than ours and does not include the previous episode of persistently elevated inflation in the 1970s. Additionally, they do not examine the implications for monetary policy.

Only recently the attention has moved to the modelling of the whole density of inflation outcomes (Manzan and Zerom, 2013, 2015; Lopez-Salido and Loria, 2020; Korobilis et al., 2021). In contrast to previous work, that mostly rely on quantile regression approaches, we follow Delle Monache

[2]Some recent exceptions are Andrade et al. (2014); Hilscher et al. (2022).

et al. (2024) and devise a parametric model for the whole density of US core PCE. The model allows for asymmetric innovations, drawn from a Skew-t distribution (see Arellano-Valle et al., 2005), and relies on the score-driven framework of Harvey (2013) and Creal et al. (2013) to set up laws of motion for the parameters, as in Delle Monache and Petrella (2017).[3] Following Stock and Watson (2007), we allow time-varying moments to feature trend components, mainly driven by structural policies, in line with Cogley and Sbordone (2008) and Ascari and Sbordone (2014), and cyclical variations, aimed at capturing transitory, short-lived factors that can temporarily affect price dynamics ("cost-push" and demand forces, as in Gordon, 1970).

We also contribute to the literature on the optimal inflation targeting (Svensson, 1997; Giannoni and Woodford, 2004). Our work is connected to the literature that investigates risk management approaches in monetary policy. Dolado et al. (2004); Surico (2007); Kilian and Manganelli (2007, 2008) study monetary policy under an asymmetric loss function. These models generally suggest that the optimal policy rule may involve nonlinear terms for the output gap and inflation if policymakers have asymmetric preferences for allowing output or inflation to deviate above or below their targets. Evans et al. (2020) argue that the zero lower bound necessitates a looser monetary policy under uncertainty, leading to an optimal delay in policy liftoff. Bianchi et al. (2021) show that the deflationary bias caused by the risk of recurrently hitting the Zero Lower Bound (ZLB) constraint can be eliminated if the Central Bank adopts an asymmetric monetary strategy. This strategy requires the Central Bank to respond more strongly to inflation deviations below target than to those above target. In this paper, we characterize a monetary policy rule that is optimal whether inflation risk is tilted to the downside or the upside, and whether the balance of risk happens to change abruptly. In addition, none of these papers estimate the time-varying asymmetric risks in the data using a frontier forecasting model.

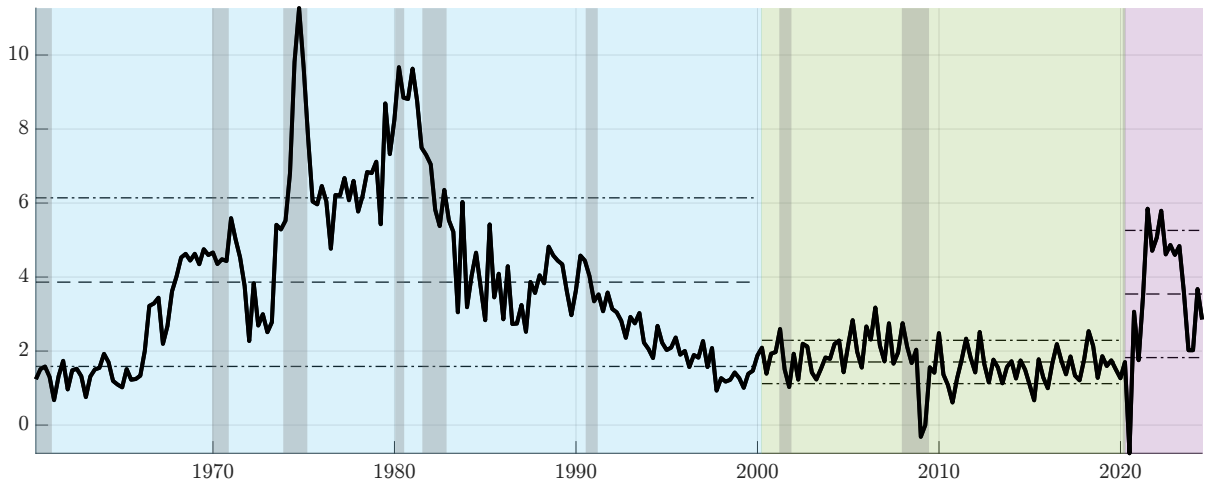

(a) Core PCE

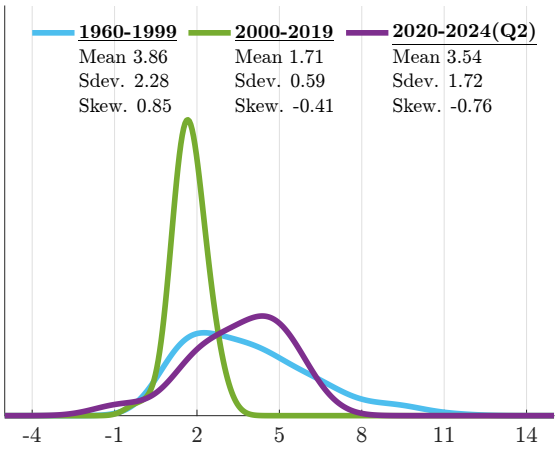

(b) Long-run distributions

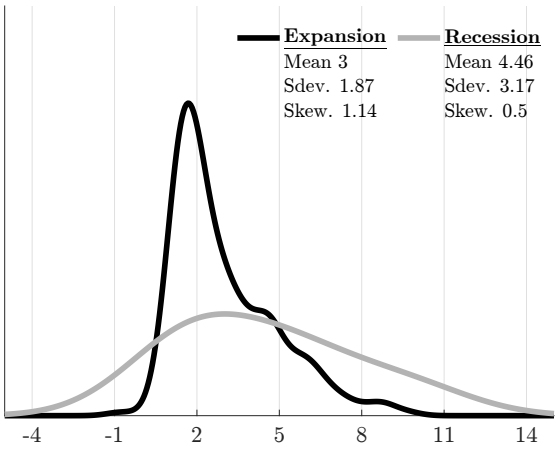

(c) Business cycle distributions

Figure 2: The Ever-Changing Faces of Inflation

Note: Panel (a) shows US core PCE from 1960Q1 to 2024Q2. Shaded areas divided the sample into three subsamples: 1960-1999, 2000-2019 and 2020-2024Q2. Dashed lines report the sample mean of each subperiod, along with the 68% confidence interval (dot-dashed). Panel (b) and (c) reports sample densities at the low and business cycle frequencies, respectively. Gray shaded areas represent NBER recessions.

2 Risks to price stability

Over the last 60 years, inflation has experienced substantial variability. The Great Inflation of the 1970s seems to be in stark contrast with the slow and steady price growth in the aftermath of the Great Financial Crisis (GFC). Recently, in the post-pandemic years, inflation has experienced a sustained surge which triggered the fastest hiking cycle devised in the history of the Federal Reserve. In panel (a) of Figure 2 we show the evolution of quarterly US core PCE inflation, from

[3]Score-driven dynamics provide, under some general conditions, optimal updates in the informational theoretic sense (Blasques et al., 2014).

1960Q1 to 2024Q2. In the same figure, we highlight three main subperiods: 1960 to 1999, 2000 to 2019, and 2020 to 2024Q2. Panel (b) displays the subsample distributions, derived from simple kernel fitting to the raw data, along with the corresponding sample moments.

In the first subperiod, inflation reaches its post-war peak and exhibits significant variability. This is reflected in a sample distribution with a mean of approximately 4%, nearly double the 2% inflation target formally adopted in 2012.[4] The volatility during this period is around 2%, and the distribution shows a significantly positive skewness, evident from the pronounced right tail in panel (b). However, by the turn of the century, the properties of inflation had changed markedly. The sample density in the lower-left panel shows a clear peak with low variability, and the sample average nearly halved to around 1.7%. This shift in the sample average was accompanied by similar changes in higher-order moments. Inflation volatility decreased from approximately 2.2% to about 0.6%, while the sample skewness reversed sharply, shifting from 0.85 to -0.4.[5] The fitted distributions over these subsamples show a common mode, approximately around 2%, suggesting that the stark differences in the means are largely attributable to different level of skewness.

The final subperiod, beginning with the outbreak of the COVID-19 pandemic, marks another turning point. The inflation distribution during this period largely resembles that of the first subsample, both in terms of mean and volatility. Notably, the distribution exhibits negative skewness, necessary to explain the rapid rise in inflation during 2021 and 2022, and the swift reversal observed over the last year.[6] Should inflation converge towards the target, the estimated properties of this subsample would increasingly resemble those of the 1960-1999 period, where a large surge in inflation is reflected in positive sample skewness. In panel (c), we further divide the sample into expansion and recession periods.[7] While both distributions tend to peak around the 2% target, recessions are characterized by greater volatility and symmetric risk, whereas inflation

[4]It is well documented that prior to 2012, the FOMC informally targeted inflation around 2% (see, e.g., Bullard, 2018).

[5]For interpretative clarity, we calibrate two Skew-Normal distributions (see, e.g., Mudholkar and Hutson, 2000) to match the sample moments in the two subperiods. A change in skewness from 0.8 to -0.95 corresponds to an increase in the probability of observing realizations below the mode from 0.15 to 0.75.

[6]The negative skewness we report highlight two facts: first, the local mean of inflation has abruptly shifted up, thereby assigning any realization below 4% to the left-side of the sample distribution. Second, using rolling estimates to estimates higher-order moments can lead to counterintuitive results, difficult to reconcile with the data. We will come back to the second point in Section 3.

[7]Recessions are defined as all the quarter falling between business cycle peaks and trough, as reported by the NBER US Business Cycle Dating Committee. Expansions define all the remaining quarters.

Table 1: Time variation in higher order moments

	Q	Q^*	N
	Homoskedastic		
Asymmetry	367.31***	371.60***	4.18***
	Heteroskedastic		
$Scale^2$	369.36***	373.67***	1.50***
Asymmetry	35.65***	36.07***	0.79***

Note: Q is the portmanteau test, Q* is the Ljung-Box extension (with automatic lag selection) and N corresponds to the Nyblom test. Q and Q* are distributed as a χ^2_1, while N is distributed as a Cramer von-Mises distribution with 1 degree of freedom. * p $<$ 10%, ** p $<$ 5%, *** p $<$ 1%.

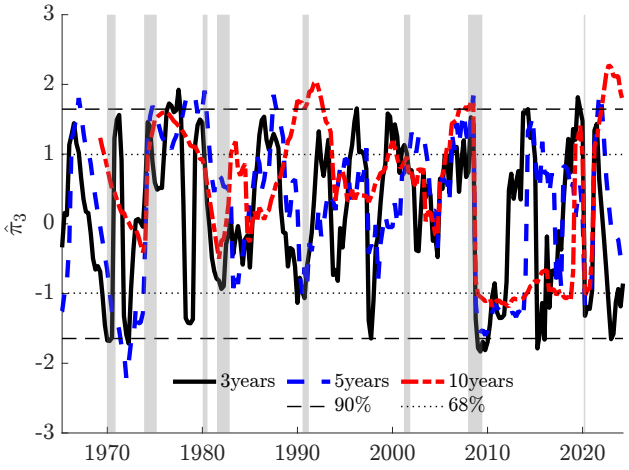

Figure 3: Rolling Bai and Ng (2005) tests

Note: The figure reports rolling Bai and Ng (2005) test statistics for US core PCE, using windows of 3, 5 and 10 years, and the the 68 and 90% critical values.

outcomes during expansions display distinct positive skewness.

Overall, the data suggest that skewness is a defining feature of inflation and, more importantly, that the asymmetry of inflation risk seem to significantly vary over time. In Table 1, we formally test the significance of these time-varying properties using three alternative parametric Lagrange Multiplier tests: a Q test, an adjusted Q* test, and the Nyblom test, following Delle Monache et al. (2024).[8][9] The upper panel presents the results without allowing for time-varying volatility, while the second panel incorporates time-varying volatility. In all cases, the stability of the asymmetry is decisively rejected at the 1% significance level.

Finally, in Figure 3, we present the Bai and Ng (2005) statistics for sample skewness, along with 68% and 90% critical values, calculated using rolling windows of 5 and 10 years. The results provide strong evidence of time variation in sample skewness, with test statistics showing significant fluctuations over the sample period. This includes periods of significantly positive skewness, such as in the mid-1970s, early 1980s, and at the end of the sample, as well as significantly negative skewness in the decade following the GFC.

These tests underscore the importance of tracking the evolution of asymmetry of inflation risk

[8]These tests consist of fitting the data to a *Skew-t* distribution, defined by parameters of location, scale and asymmetry. A score-based test can then be used to test for the stability of the fixed parameters. We also account for the possibility of fat tails in the distribution (see Harvey and Thiele, 2016; De Polis, 2023).

[9]In Appendix B we show that the test results hold for different definitions of inflation.

in real-time, and highlights a trade-off in estimating time-varying skewness with rolling measures. On the one hand, large look-back windows reduce the dominance of isolated outlying observations, improving the estimation accuracy of the third moment. On the other hand, longer sample periods reduce the sensitivity to time variation, leading to delayed detection of changes. The marked differences in the estimates based on a 5-year or a 10-year window reported in Figure 3 uncover significant variation in skewness, but also highlight substantial lags in its detection.

3 A statistical model of inflation risk

The evidence reported in the previous section highlights that inflation risk and uncertainty have dramatically shifted over the last 60 years. In what follows, we now introduce a statistical model which allows to capture these feature of the data.

Let $\pi_t = 400 \log(p_t/p_{t-1})$ be annualized, quarter-on-quarter (core) inflation, and assume that at each point in time the distribution of π_t can be characterized by Skew-t (Skt) distribution with time-varying location (μ_t), scale (σ_t), and shape (ϱ_t) parameters:

$$\pi_t \sim Skt_\nu(\mu_t, \sigma_t^2, \varrho_t), \tag{1}$$

where ν denotes the, time invariant, degrees of freedom. The distribution of inflation realizations is positively (negatively) skewed for $\varrho > 0$ ($\varrho < 0$). Therefore, the underlying right- and left-risk around the central scenario (mode), μ_t, can be retrieved as $\sigma_t(1-\rho)$ and $\sigma_t(1+\rho)$. It is worth noting that this specification allows as special cases the symmetric Student-t distribution, when $\varrho_t = 0$, the epsilon-Skew-Gaussian (Mudholkar and Hutson, 2000) for $\nu \to \infty$, and the Gaussian density when both conditions hold. Thus, we allow for, but do not impose, asymmetric innovation terms.

Following a long tradition in modelling the statistical properties of inflation (see e.g., Cogley, 2002; Stock and Watson, 2007; Faust and Wright, 2013), we treat the time-varying parameters as unobserved components that can be learned in real-time from the variation in the data. Unlike

Stock and Watson (2007), we opt for an observation-driven updating process.[10] Specifically, let $\delta_t = \log(\sigma_t)$ and $\gamma_t = \text{arctanh}(\varrho_t)$, we postulate that each element $f_{i,t}$ of $f_t = (\mu_{t|t-1}, \delta_{t|t-1}, \gamma_{t|t-1})'$ features a permanent and transitory component: $f_{i,t} = \bar{f}_{i,t} + \tilde{f}_{i,t}$, which evolve as:

$$\bar{f}_{i,t} = \bar{f}_{i,t-1} + a_i s_{i,t-1}, \tag{2}$$

$$\tilde{f}_{i,t} = \phi_i \tilde{f}_{i,t-1} + b_i s_{i,t-1}. \tag{3}$$

Updates of the time-varying parameters are proportional to $s_{i,t-1}$, which is the *scaled score* of the conditional distribution (as in Creal et al., 2013; Harvey, 2013). The scale score vector, $s_t = (s_{\mu,t}, s_{\sigma,t}, s_{\varrho,t})'$, is defined as $s_t = \mathcal{S}_t \nabla_t$, where ∇_t is the gradient of the likelihood function ℓ_t with respect to the dynamic parameters, and the scaling matrix $\mathcal{S}_t$ is proportional to the inverse of the diagonal of the Information matrix, $\mathcal{I}_t = \mathbb{E}[\nabla\nabla']$. Intuitively, the score vector translates the new information contained in the latest data release, summarized by the "prediction error", $\varepsilon_t = \pi_t - \mu_t$, into an update for the time-varying parameters characterizing the predictive distribution of inflation. In our setting:[11]

$$\begin{bmatrix} s_{\mu,t} \\ s_{\gamma,t} \\ s_{\delta,t} \end{bmatrix} = \begin{bmatrix} \sqrt{\frac{(1+3\eta)(1-\varrho_t^2)}{(1+\eta)}} w_t \zeta_t \\ \sqrt{\frac{(1+3\eta)}{2}} (w_t \zeta_t^2 - 1) \\ sgn(\varepsilon_t) \sqrt{\frac{(1+3\eta)(1-sgn(\varepsilon_t)\varrho_t)}{3(1+\eta)(1+sgn(\varepsilon_t)\varrho_t)}} w_t \zeta_t^2 \end{bmatrix}, \tag{4}$$

where $\zeta_t = \frac{\varepsilon_t}{\sigma_t}$, $\eta = \frac{1}{\nu}$, $w_t = \frac{(1+\eta)}{(1+sgn(\varepsilon_t)\varrho_t)^2+\eta\zeta_t^2}$, and $sgn(\cdot)$ is the sign function. In practice, updates driven by the scaled score are (generally) guaranteed to reduce the distance between the conditional and the true (unobserved) predictive distribution, easily allowing for non-Gaussian features.[12]

Consider a symmetric Gaussian environment, where $\eta = 0$, and $\varrho = 0$ and $w_t = 1$ $\forall t$; the location parameter, which now represents the mean of the distribution, would update according to $s_{\mu,t} = \frac{\varepsilon_t}{\sigma_t}$, and the variance following $s_{\sigma,t} = \frac{(\varepsilon_t^2 - \sigma_t^2)}{\sigma_t^2}$. That is, in line with standard Kalman

[10] In an observation-driven model, current parameters are deterministic functions of lagged dependent variables as well as contemporaneous and lagged exogenous variables. In parameter-driven models, parameters vary over time as dynamic processes with idiosyncratic innovations. See Cox (1981).

[11] Appendix A we provide detailed derivations for the score and updates of the time varying parameters of the model (see also Delle Monache et al., 2024).

[12] See Blasques et al. (2015, 2022)

filter learning (see, e.g., Cogley, 2002), score driven updates imply changes in the mean that are proportional to the prediction error, with strength inversely proportional to the variability of the data. Similarly, the variance is updated proportionally to the difference between the variability of the prediction error (ε_t^2) and the expected variability of the inflation process (σ_t^2). Allowing for fat tails, that is $w_t \neq 1$, makes the updating mechanism robust to large, unanticipated prediction errors, such that the central scenario and the associated uncertainty are less responsive to outlying observations (see. e.g., DelleMonache and Petrella, 2017; Antolín-Díaz et al., 2024). Finally, when asymmetry is introduced, the updating mechanism weights prediction errors differently, depending on their sing. For example, when the conditional distribution is left skewed, parameters react more to unexpected positive news, rather than to negative prediction errors, which are expected to be more likely to occur. Consistently with this mechanism, the asymmetry parameter always updates in the direction of the prediction error.

Alternative approaches to learn the entire distribution of inflation in real-time generally rely on quantile regressions (Manzan and Zerom, 2013, 2015). However, these methods fall short of accounting for several of the stylized facts about inflation. On the other hand, whereas the model we propose inherits many features of common unobserved components specifications (along the lines of Stock and Watson, 2007), it comes with the additional flexibility of allowing for skewness in the predictive distribution of inflation. Notably, this specification admits a broader definition of risk to inflation through the interaction of the first three moments of the distribution.

Basic Features of the Model. A defining feature of the Skew-t model in Equation (1) is the fact that asymmetry directly affects the first two moments of the distribution. Specifically, one can show that:

$$\mathbb{E}_{t-1}[\pi_t] = \mu_t + g(\eta)\sigma_t\varrho_t, \qquad g(\eta) = \frac{4\mathcal{C}(\eta)}{1-\eta}, \tag{5}$$

and

$$Var_{t-1}(\pi_t) = \sigma_t^2\left(\frac{1}{1-2\eta} + h(\eta)\varrho_t^2\right), \qquad h(\eta) = \frac{3}{1-2\eta} - g(\eta)^2, \tag{6}$$

where the notation $\mathbb{E}_{t-1}$ indicates the expectation conditional on information at time $t-1$. Equations (5) and (6) highlight two important features: first, when $\varrho = 0$, the mean and variance collapse to the first two moments of a t distribution with $\frac{1}{\eta}$ degrees of freedom, μ_t and $\frac{\sigma_t^2}{\eta-1}$ respectively; that is, the model allows for *symmetric* distributions. Second, asymmetry creates a wedge between the central scenario, e.g., the mode, and the expected value. This wedge has the same sign of the prevalent asymmetry, and it is quantitatively more relevant as the distribution becomes more disperse, as $\frac{\partial \mathbb{E}_{t-1}[\pi_t]}{\partial \varrho_t} > 0, \ \forall t$. On the other hand, asymmetry increase the variance when positive and decreases it when $\varrho_t < 0$.[13] Therefore, procyclical variations in inflation skewness are reflected into a time-varying correlation between the mean and volatility of the process.

Differently, we can compute skewness in closed form as (see De Polis, 2023),

$$Skew_{t-1}(\pi_t) = \frac{g(\eta)\varrho\left[1+\eta-\varrho^2\left(5-2g(\eta)^2+(10g(\eta)^2-19)\eta-12g(\eta)^2\eta^2\right)\right]}{(1-3\eta)(1-2\eta)\left(\frac{1}{1-2\eta}+h(\eta)\varrho^2\right)^{\frac{3}{2}}}, \tag{7}$$

noting that this moment only depends on the asymmetry parameter and on the estimated degrees of freedom.

Estimation. The parameters of the model and the associated conditional distribution of inflation are estimated using Bayesian methods as in Delle Monache et al. (2024). We use minnesota-type priors for the the persistence of the transitory components. Loadings on the score components are Inverse Gamma distributed, with mean and standard deviation equal to 0.01 and 0.001 for the permanent loadings, a, and 0.025 and 0.015 for the transitory loadings, b. This choice reflects the view that transitory parameters are slower to react to news compared to the transitory components. Furthermore, the prior ensures that the filter is invertible (Blasques et al., 2022), that is it reduces the possibility of overshooting the updates in the direction of the (local) optimum, and assumes conservative views on parameters time variation. Lastly, we assume an inverse gamma prior for η. We set up an adaptive Random-Walk Metropolis-Hastings algorithm (ARWMH, Haario et al., 1999). Credible sets for both static and time-varying parameters are obtained from the empirical

[13] $\frac{\partial Var(\pi_t|\Pi_{t-1})}{\partial \varrho_t} = 2h(\eta)\sigma_t^2\varrho_t$, and since $h(\eta) > 0$ for $\nu > 3$, the shift in the variance will be of the same sign as the level of the shape parameter (thus of the same sign to the level of the conditional skewness).

distribution functions arising from the resampling. See Delle Monache et al. (2024) for a detailed discussion.

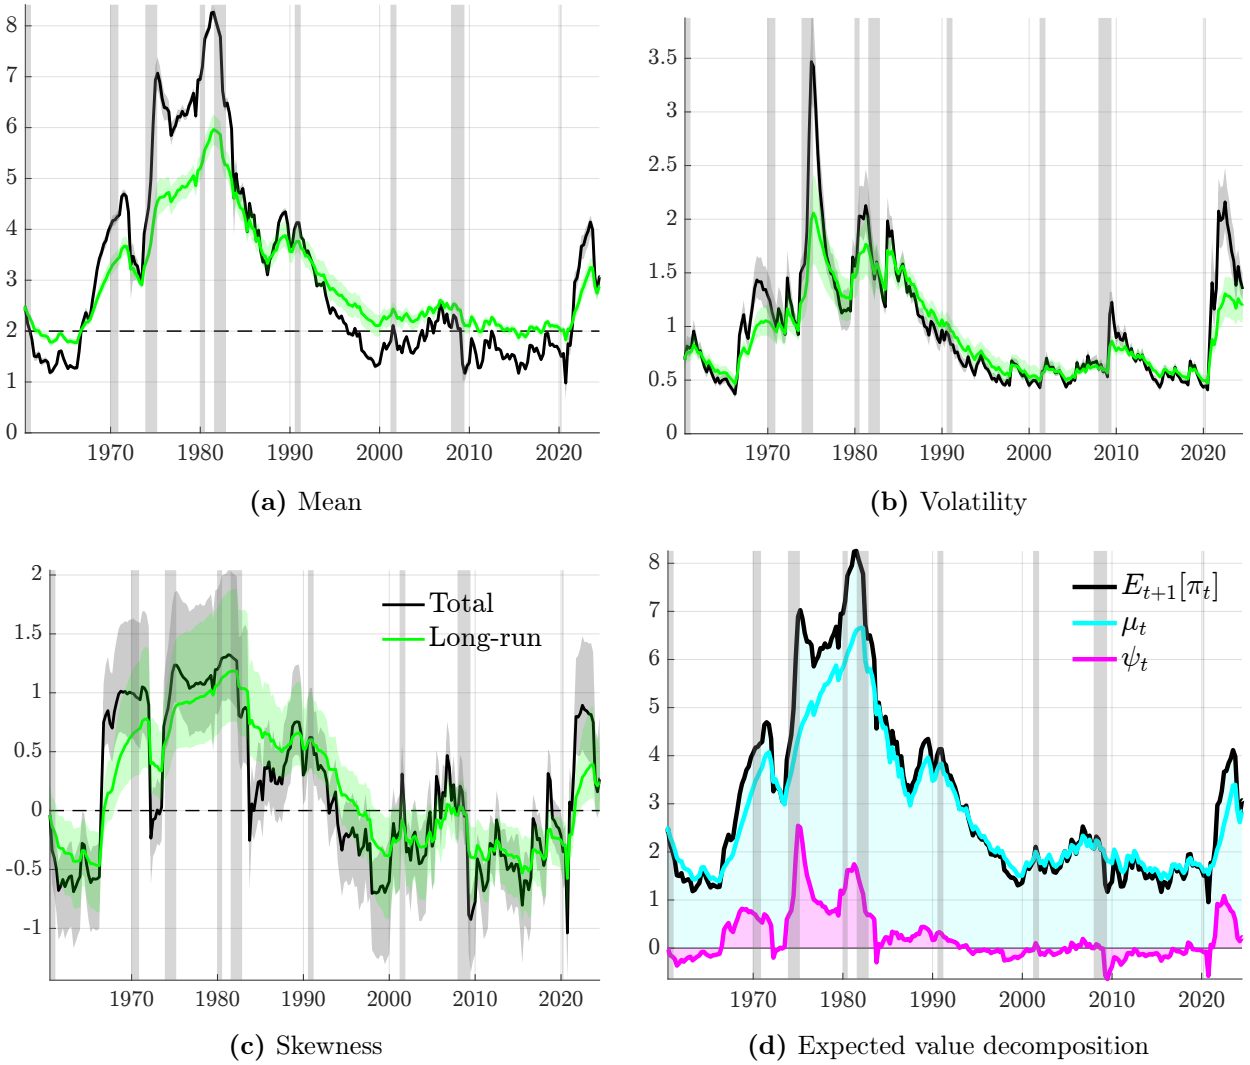

(a) Mean **(b)** Volatility

(c) Skewness **(d)** Expected value decomposition

Figure 4: Time-varying moments of inflation

Note: The panels report mean, volatility and skewness of US core PCE. Blue lines represent total moments, red lines correspond to long-run components only. Bands report 68 and 96% credible intervals. Gray shaded areas represent NBER recessions.

3.1 The asymmetric dynamics of inflation risks

Figure 4 displays the estimated time-varying moments. We report in black the total moment (e.g., computed using the total parameters, $f_t = \bar{f}_t + \tilde{f}_t$), whereas the persistent components (e.g., setting $f_t = \bar{f}_t$) are in green. The model reveals significant time variation across all moments. The time-varying mean (panel (a)) reflects the well-documented trend in inflation, which rises in the

mid-1960s, declines from the early 1980s, and stabilizes near a 2% target by the mid-1990s (see, e.g., Stock and Watson, 2016). The recent inflationary episode is marked by a sharp increase in both average expected inflation and its long-term component, with a noticeable reversal in the last few observations. Inflation volatility (panel (b)) peaks in the mid-1970s, remaining high until the late 1980s, and is well-contained until early 2020, when it sharply increases starting in Q2. Unlike the mean, inflation volatility exhibits clear cyclicality, rising significantly during recessions.

The skewness estimates (panel (c)) indicate moderate negative skewness in the 1960s, with increasing upside risks from the late 1960s, peaking in the late 1970s, and then declining from the early 1980s. Upside risks persist until the mid-1990s, when the skew shifts to negative. Downside risk dominates until the post-COVID inflationary episode, except for the period before the GFC, where risks are balanced. The model captures a marked increase in negative skewness during the pandemic, followed by a rapid rise in upside risk. By the end of 2020, substantial upside risks emerge, reaching levels comparable to those seen during the Great Inflation of the 1970s by mid-2021.Notably, the risk distribution during the latest inflationary episode closely resembles the environment of the mid-1970s, both qualitatively and quantitatively, while the low and stable inflation period before COVID mirrored the stable inflation era of the 1960s. It is worth noticing that, contrary to the mean, where the transitory components remains highly persistent, skewness shows far less transitory deviations. This is due to a lower estimate for the autocorrelation of the transitory component of asymmetry compared to that of the location, but also to smaller learning rates, which make the former less sensitive to noisy prediction errors.[14]

Following Equation (5), panel (d) presents the decomposition of expected inflation into the location (the most likely expected outcome) and the tilt induced by asymmetric risks around it, ψ. Inflation risk significantly influences inflation expectations, introducing a substantial upside bias during the high inflation periods of the 1970s and the post-COVID era. Moreover, negative skewness contributed to an downward bias between 30 and 50 basis points in expected inflation during the decade leading up to COVID.

Whereas more is known about inflation time-varying mean and variance (see, e.g., Stock and Watson, 2002), our model offers novel insights on the dynamics of inflation skewness. We compare

[14]See Table 7 in Appendix C.

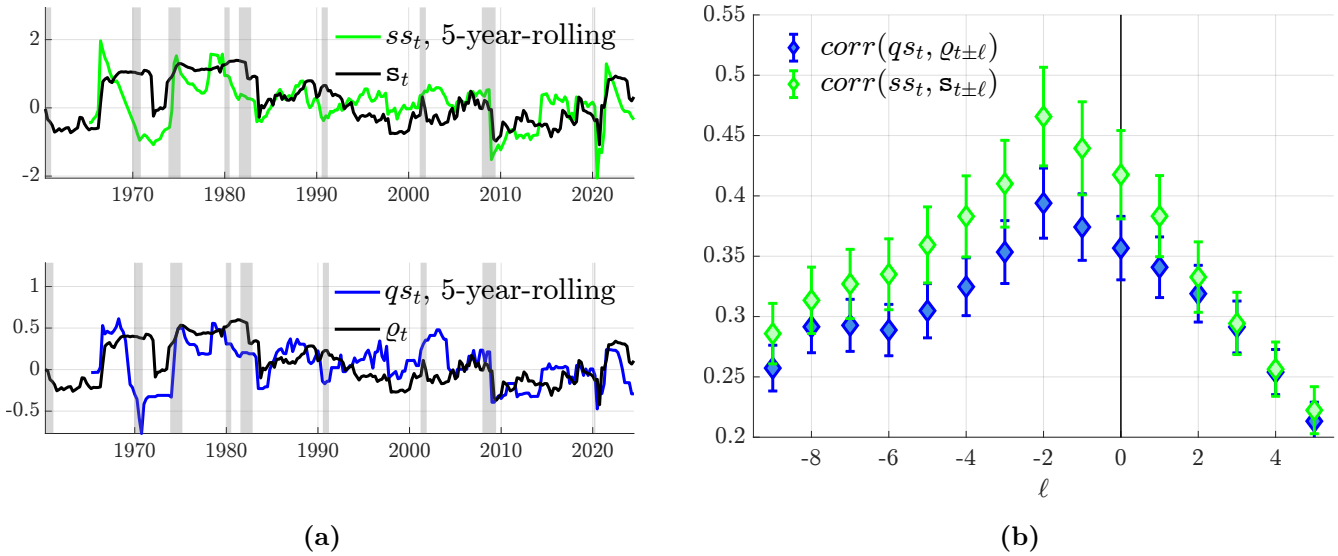

(a) **(b)**

Figure 5: Model-based vs rolling measures of skewness

Note: The figure reports cross-correlations between data-based and model-based measures of inflation skewness. We report in blue the correlation between rolling quantile skewness and the estimated coefficient of asymmetry, and in green the correlation between rolling sample skewness and the conditional skewness produced my the model. Gray shaded areas represent NBER recessions.

our model-based measures of skewness (e.g., sample skewness, $\mathtt{s}_t$ and the asymmetry parameter, ϱ_t) against rolling estimates solely based on data. Panel (a) in Figure 5 compares sample skewness against a a 5-year rolling Pearson's skewness, ss_t, reported in green. Despite patterns evolve similarly, ss_t is a noisier measure of skewness, due to its sensitivity to individual observations. In the bottom figure, we compare estimates of ϱ_t against a robust quantile-based measure of skewness, qs_t.[15] The impact of outliers is evident, especially with large negative data points from the second quarter of 2020, which continue to distort post-2021 qs_t estimates. In contrast, model-based skewness estimates are less affected by extreme values, due to the robust updating mechanism. Furthermore, rolling estimates assume constant skewness within the sample window, which makes them slow to adjust, particularly during major inflation shifts, presenting a significant challenge for real-time risk assessment. Differently, our model's skewness estimates are updated in real-time with each new inflation release. While all measures generally align in capturing underlying risk and its evolution, panel (b) shows that our skewness measures incorporates changes in inflation quicker, leading by an average of two quarters – a clear advantage for real-time monitoring of inflation risk.

[15]Notice that, incidentally, both measures are bounded between -1 and 1.

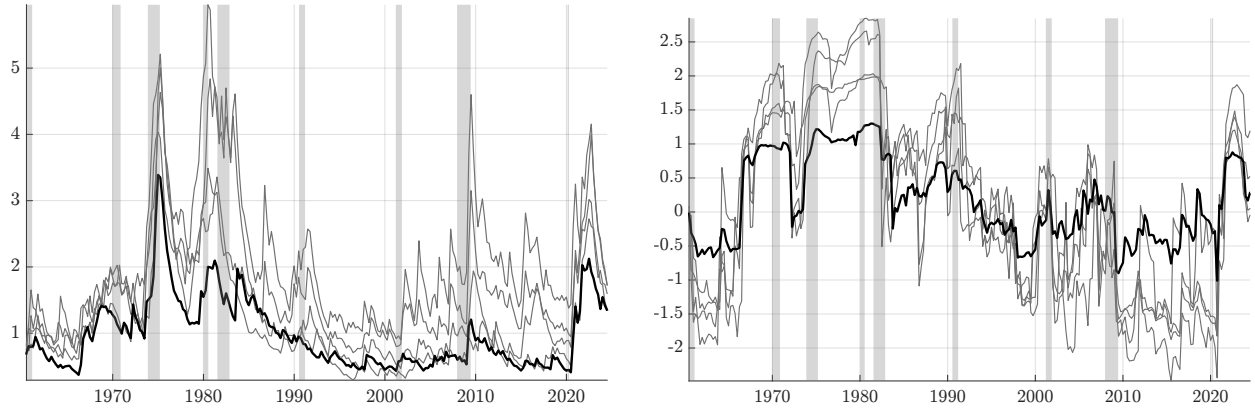

(a) Volatility **(b)** Skewness

Figure 6: Risk across different inflation measures

Note: The panels report the full moment median estimates volatilities (a) and skewness (b) for different measures of inflation. Black lines indicate estimates for core PCE. Other inflation measures we consider are: GDP deflator, headline PCE, headline CPI and core CPI. Gray shaded areas represent NBER recessions.

Other inflation measures The evidence reported so far is based on data for core PCE, which is the measure preferred by the FOMC to gauge price stability. Nevertheless, estimating the model on different inflation measures lends support to a generalization of our in-sample findings. We consider four additional inflation measures: the GDP deflator, headline PCE, and headline and core CPI. Figure 6 shows the estimated dynamics of inflation volatility and skewness across the different measures, highlighting in blacks that of core PCE. Two comments are in order. First, the dynamics of the two moments is extremely similar for all measures. With varying magnitudes, volatilities spike around recessions, and remain persistently high soon after. Skewness follow humped-shape patterns in the 1970s and 1980s, then moving downward since the 1990s, remaining negative until the pandemic period. Second, it's important to note that among all these measures, it emerges that core PCE is the least "risky", with both volatility and skewness exhibiting the least amount variation.

3.2 Forecasting

Here, we assess the out-of-sample forecasting performance of our model, with a focus gauging the added value of accounting for time-varying skewness. Specifically, we set up a real-time forecasting exercise where for each inflation vintage we produce up to twelve-step ahead forecasts

Table 2: Out-of-sample comparison

	h = 1	h = 2	h = 3	h = 4	h = 8
MSFE	**0.835** (0.016)	**0.854** (0.021)	**0.861** (0.031)	**0.859** (0.010)	**0.951** (0.004)
CRPS	**0.936** (0.012)	**0.939** (0.011)	**0.934** (0.002)	**0.927** (0.004)	**0.966** (0.002)
	CRPS decomposition				
Right	**0.926** (0.019)	**0.932** (0.018)	**0.931** (0.010)	**0.934** (0.013)	**0.965** (0.071)
Left	**0.949** (0.002)	**0.940** (0.006)	**0.937** (0.001)	**0.923** (0.001)	**0.962** (0.002)
Center	**0.933** (0.004)	**0.942** (0.011)	**0.935** (0.002)	**0.923** (0.005)	**0.971** (0.004)
	Event Forecasts				
$\pi_{t+h} < 1.5$	**0.945** (0.015)	**0.940** (0.005)	**0.931** (0.001)	**0.950** (0.026)	**0.960** (0.040)
$\pi_{t+h} > 2.5$	**0.910** (0.031)	**0.969** (0.089)	**0.960** (0.036)	**0.966** (0.032)	0.986 (0.132)
$1.5 \leq \pi_{t+h} \geq 2.5$	**0.939** (0.013)	**0.947** (0.009)	**0.947** (0.016)	0.981 (0.176)	**0.940** (0.001)

Note: The table report the relative performance of Stock and Watson (2002) UCSV model against our *Skt* model. Results are reported in ratios, with our model being at the numerator; values smaller than 1 imply superior predictive accuracy of the *SKt* model. The out-of-sample period runs from 2000Q1 to 2024Q2. Values in **bold** are significant at the 10% level.

for the whole density of core PCE inflation, starting from 2000Q1. We evaluate the forecasting performance of the model against the UCSV model of Stock and Watson (2002) which represents a solid benchmark model, widely employed by policy institutions to predict inflation outcomes.[16] We compare the two models in their ability to produce accurate point, density and event forecasts. That is, we evaluate the mean squared forecast error (MSFE) for point accuracy; we measure density forecasting ability by means of Gneiting and Ranjan (2011) weighted quantile scores. These scoring rules allow us to evaluate the overall density forecast accuracy using the Continuously Ranked Probability Score (CRPS), a loss function that compares the predictive distribution function to a step function which moves from 0 to 1 on the realization point ("perfect forecast"),[17] but also to focus on the prediction of specific parts of the distribution, where we focus on the right and left tails, and the central body of the predictive densities.[18]

[16] We implement the Bayesian version following Chan (2013).

[17] The CRPS can be interpreted as a generalization of the mean absolute error for density forecasts.

[18] Left tail forecasts are defined up to the 25^{th} quantile. Similarly, the right tail considers above the 95^{th} quantile. The remaining quantiles characterize to the center of the distribution.

Table 3: Event forecast comparison against SPF

$\pi_t^{Q4} \leq 1.5$				$\pi_t^{Q4} \geq 2.5$				$1.5 \leq \pi_t^{Q4} \leq 2.5$			
h = 1	h = 2	h = 3	h = 4	h = 1	h = 2	h = 3	h = 4	h = 1	h = 2	h = 3	h = 4
1.196	0.892	1.159	0.893	0.853	0.164	0.335	0.599	1.177	0.895	1.327	1.125

Note: The table reports the ratio of the Brier score of our *Skt* model over the SPF's for event predictions. The target variable is Q4-over-Q4 core PCE. The evaluation sample runs from 2007Q1 due to SPF data availability.

Results for the comparison are reported in Table 2, where for each loss function and forecast horizon we report the ratio of the score achieved by our model over that of the UCSV; values smaller than unity point at a superior accuracy of our preferred model. We report p-values for Diebold and Mariano (1995) test in parentheses.

The results strongly support the superiority of our model against the benchmark, across all horizon and forecast exercises. Gains in point forecast range from 25% for short-horizons to 7$ over the medium run. Smaller, yet sizable, gains are observed in CRPS scores. Notably, our model delivers improvements up to 8% in upside risks forecast. Overall, we uphold the relevance of accounting for inflation skewness as a way to improve forecasting accuracy. Nonetheless, the UCSV model not only does not feature skewness, but it also overlooks the presence of fat tails. To avoid mistaking gains coming from modelling inflation skewness from those related to fat tails, we reproduce the same results comparing our model against a specification that does not allow for any asymmetry, similarly to Delle Monache and Petrella (2017). Results reported in Appendix C validate our claim that modelling skewness is a key step to improve model fit and forecasting accuracy for inflation.

3.3 Event forecast and comparison with SPF

In the bottom part of Table 2 we also compare the ability of the models to produce events forecasts. Specifically, we use the Briers score to evaluate predictions of the events that π_{t+h} is lower than 1.5%, greater than 2.5% or that it falls within these thresholds. Again, our model provide improvements over the benchmark UCSV model, especially over the short-horizon.

We also compare our event forecast predictions against consensus. For this exercise we target

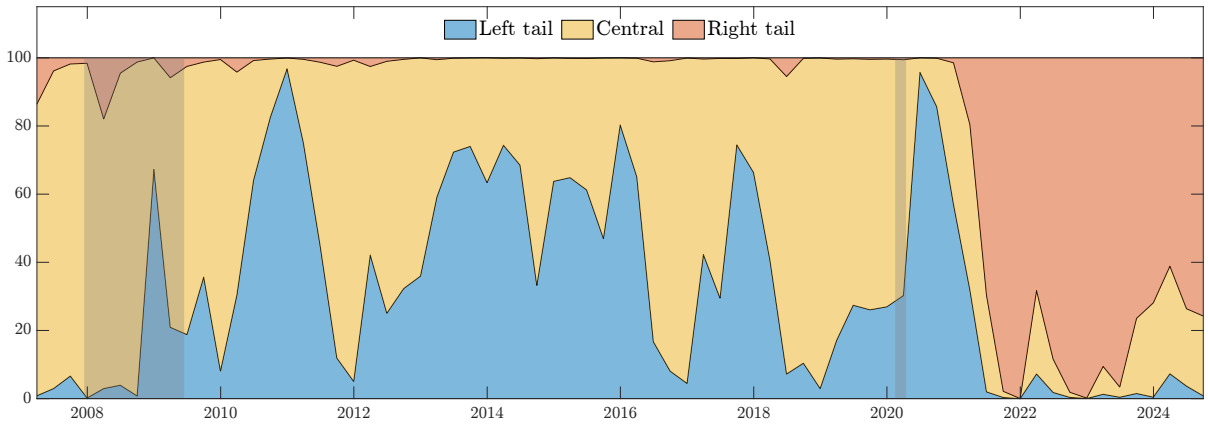

(a) *Skt* predictive distribution

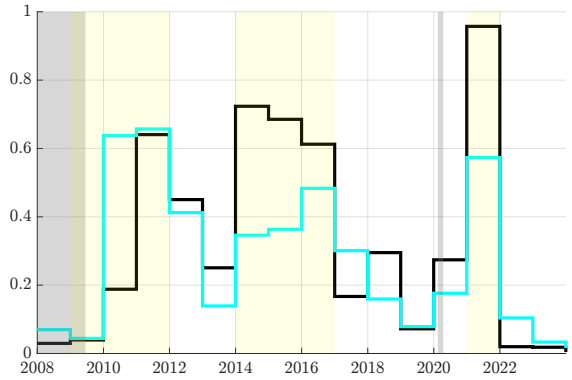

(b) $P\left(\pi_{t+2}^{Q4} < 1.5\%\right)$

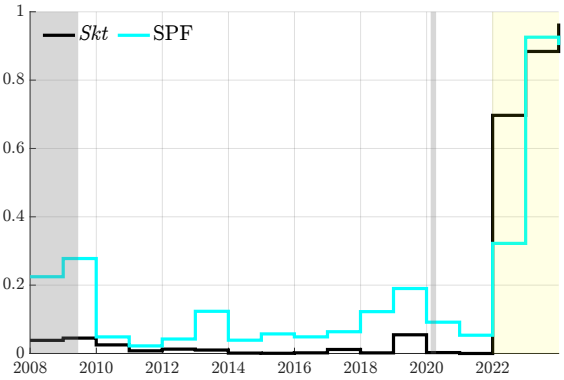

(c) $P\left(\pi_{t+2}^{Q4} > 2.5\%\right)$

Figure 7: Event forecasts

Note: The top panel report the interval forecasts produced by the *Skt* model. We define left tail as the probability of inflation expectations below 1.5%, central corresponds to expectations in the [1.5%, 2.5%) interval, whereas the right tail is defined as expectations above 2.5%. Panels (b) and (c) report a comparison of the two-step-ahead predicted probability of Q4-over-Q4 inflation being below 1.5% and above 2.5%, repsectively; yellow shaded region represent the events. The sample runs from 2007Q1 to 2023Q4. Gray shaded areas represent NBER recessions.

Q4-over-Q4 core PCE, π_t^{Q4}, starting in 2007Q1, to match SPF data. First, in panel (a) of Figure 7 we show the interval predictions produced by our model, as reported in the top panel of Figure 1. Upon visual inspection, it seems that our model is able to produce interval predictions that are in line with consensus. In the bottom panels, we also plot the evolution of the predicted probability of π_t^{Q4} being below 1.5% (panel (b)) or above 2.5% (panel (c)). These figure show that our model produces, on average, timelier assessments of the event probabilities. Strikingly, the *Skt* model detects little to no probability of overshooting the 1.5-2.5% interval for the whole period between the GFC and the post-pandemic bout of inflation, which is captured with greater precision.

We summarize a formal evaluation of these predictions in Table 3, where we report ratios of

Brier scores for the *Skt* model over the SPF for the three events and for up to four-step-ahead predictions.[19] Overall, the table show that our model and consensus deliver comparable accuracy, with the exception of $P\left(\pi_{t+1}^{Q4} > 2.5\%\right)$, where our model appears to produce timelier assessments to these probabilities (see Figure 7 panel (c)). Notice, however, that the ratios reported in the Table are computed with a small number of observations – due to SPF data availability.

4 Implications for Monetary Policy

The deflationary bias observed since the Great Financial Crisis has prompted policymakers to renew their commitment to maintaining price stability. Over the last two years, two major central banks, the Federal Reserve and the European Central Bank, have reviewed their approaches to keeping inflation close to a target level of about 2% over the medium term. The major update from these *strategy reviews* was a shift from a symmetric inflation target –where central banks react equally to inflation above or below the target (the so-called "bygones-be-bygones")– to an *asymmetric* target. Under this approach, for example, the policymaker allows inflation to run higher than the target after periods of below-target levels (see, e.g., FED, 2021; Reichlin et al., 2021). Within this new policy framework, the optimal monetary policy response to inflation fluctuations needs to account for the evidence of asymmetric risks.

In this section, we formally examine the implications of asymmetric macroeconomic risk for monetary policy. We begin by analyzing optimal policy within the foundational three-equation New Keynesian (NK) DSGE model, accounting for asymmetric risk. Following this theoretical exploration, we transition to an empirically focused framework, where real-time tracking of time-varying skewness in risk is translated into a corresponding asymmetric adjustment in monetary policy strategy.

[19] Notice that for $h = 1$, the prediction uses only out-of-sample values. As h approaches 4, up to 3 observed data points are used in the computation of Q4-over-Q4 inflation.

4.1 Skewed risk in a simplified NK model

The textbook NK model (e.g. Clarida et al., 1999; Woodford, 2003; Galí, 2008), log-linearized around its unique steady-state equilibrium, results in the following equations:

$$
\begin{aligned}
y_t &= E_t y_{t+1} - \varsigma^{-1}\left(\hat{i}_t - \mathbb{E}_t\,\hat{\pi}_{t+1}\right), \\
\hat{\pi}_t &= \kappa(y_t - y_t^{\star}) + \beta\,\mathbb{E}_t\,\hat{\pi}_{t+1}, \\
\hat{i}_t &= \hat{\pi}_t + \phi_\pi\left(\hat{\pi}_t - \hat{\pi}_t^{\star}\right), \\
y_t^{\star} &= \omega\hat{\alpha}_t,
\end{aligned}
$$

where ς is the inverse of the elasticity of intertemporal substitution, β the discount factor, and ω and κ are convolutions of the deep parameters.[20] Output and the fully-flexible price output are considered in deviation from steady state, whereas the inflation rate is defined $\hat{\pi}_t = \pi_t - \overline{\pi}$, where $\overline{\pi} = 2\%$ (annualized) – consistent with the Federal Reserve objective for the annualized inflation rate, and the nominal rate as $\hat{i}_t = i_t - (\overline{r} + \overline{\pi})$. The Central Bank decides upon a target central scenario for inflation, $\hat{\pi}_t^{\star}$, and reacts to any deviation of inflation from this target, where $\phi_\pi > 0$ measure the strength of the interest rate response.

To simplify the derivations, we assume that the TFP shock is *iid*, and we allow its distribution to belong to general class of (potentially) asymmetric distributions, such that $\hat{\alpha}_t \sim F\left(\mu_\alpha, \sigma_\alpha, \varrho_{\alpha,t}\right)$, where μ is the mode, σ is the scale parameter and $\varrho_{\alpha,t}$ measures the asymmetry of the mass about the mode, such that the distribution features positive (negative) skewness when $\varrho > 0$ ($\varrho < 0$).[21] Consistent with our empirical findings, we assume that the asymmetry is time varying.[22] In what follows we will assume $\mu_\alpha = 0$, i.e. the most likely outcome for the shock, ex-ante, is zero. Yet, for $\varrho_{\alpha,t} \neq 0$, we have $\mathbb{E}_t\left[\hat{\alpha}_{t+1}\right] = \psi_{\alpha,t}$, which is positive (negative) for $\varrho_{\alpha,t} > 0$ ($\varrho_{\alpha,t} < 0$), that is the presence of asymmetry gives rise to a wedge between the central scenario and the expected

[20] In a standard textbook NK model $\omega = (1+\eta)/(\eta+\varsigma)$ and $\kappa = (1-\phi)(1-\phi\beta)(\varsigma+\eta)/\phi$, η is the inverse of Frisch elasticity of labor supply and ϕ measures nominal price stickiness.

[21] The notation $F(\mu, \sigma, \varrho)$ indicates a general class of location-scale skew distributions (Arellano-Valle et al., 2005). Densities belonging to this class are *closed under affine transformations* due to the location-scale invariance of the skewness (for positive slopes), and any symmetric distribution is allowed as a special case when $\varrho = 0$. In this latter case, μ (also) becomes the mean, but σ does not necessarily indicates the standard deviation, as fat tails might be allowed. The Skew-t distribution used in the previous section belongs to this general class.

[22] We simplify the model assuming that the scale of the distribution, σ_α, is fixed.

value. Specifically, $\psi(\sigma_\alpha, \varrho_{\alpha,t})$ is a monotonically increasing function of $\varrho_{\alpha,t}$ which tilts the mean forecast away from the modal forecast in the direction of the skewness; the wedge between the two becomes larges as the distribution becomes more dispersed.

After some manipulations, we can rewrite the model as:[23]

$$x_t = \mathbb{E}_t[x_{t+1}] - \varsigma^{-1}\left(\hat{\pi}_t + \phi_\pi\left(\hat{\pi}_t - \hat{\pi}_t^\star\right) - \mathbb{E}_t[\hat{\pi}_{t+1}] - \hat{r}_t^\star\right), \tag{8}$$

$$\hat{\pi}_t = \kappa x_t + \beta\, \mathbb{E}_t[\hat{\pi}_{t+1}], \tag{9}$$

where $\hat{r}_t^\star = \varsigma\omega\, \mathbb{E}_t[\Delta\hat{\alpha}_{t+1}]$ denotes the natural rate of interest, and $x_t = y_t - y_t^\star$ is the output gap. It follows that movements of the natural rate of interest are entirely driven by changes to the moments of the the TFP shock, i.e. $r_t^\star \sim iidF\left(\mu^\star, \sigma^\star, \rho_t^\star\right)$, where $\mu^\star = \varsigma\omega\, \mathbb{E}_t\left[\psi_{\alpha,t+1}\right]$, $\sigma^\star = \varsigma\omega\sigma_\alpha$, $\rho^\star = -\rho_{\alpha,t}$. Note that $\mathbb{E}\left[r_t^\star\right] = \mu^\star + \psi_t^\star = \varsigma\omega\left(\mathbb{E}_t\left[\psi_{\alpha,t+1}\right] - \psi_{\alpha,t}\right)$, such that $\mathbb{E}_t\left[r_t^\star\right] \neq 0$ if $\mathbb{E}_t\left[\psi_{\alpha,t+1}\right] \neq \psi_{\alpha,t}$.

To recover the unique equilibrium allocations, we guess that a solution for output gap and inflation can be written as a linear function of $r_t^\star$ and $\pi_t^\star$, where the latter is determined by the optimal policy of the Central Bank. We use the method of undetermined coefficients to solve for the unique solution of the system. In this model, macroeconomic outcomes inherit the properties of the stochastic shock, in that $\mathbb{E}_t[x_{t+1}]$ and $\mathbb{E}_t[\hat{\pi}_{t+1}]$ are not necessarily equal to zero, since the technology shock may be asymmetric. Therefore, any solution of the system is characterized by the following conditions:

$$\begin{bmatrix} \mathbb{E}_t\left[x_{t+1}\right] \\ \mathbb{E}_t\left[\pi_{t+1}\right] \end{bmatrix} = \frac{1}{\varsigma + \kappa\left(1+\phi_\pi\right)} \begin{bmatrix} \left(\frac{\varsigma+\kappa}{\phi_\pi\kappa} + 1\right)(1-\beta) \\ \frac{\varsigma+\kappa}{\phi_\pi} + \kappa \end{bmatrix} \mathbb{E}_t\left[r_{t+1}^\star - \phi_\pi\pi_{t+1}^\star\right]. \tag{10}$$

Assume now that the monetary authority chooses the forward-looking, time-varying inflation trend, $\left\{\hat{\pi}_{t+h}^\star\right\}_{h=0}^{H}$, to minimize the *symmetric* loss function: $\mathcal{L} = \mathbb{E}[\hat{\pi}_{t+1}]^2$.[24] The optimal monetary policy commands that the Central Bank set $\hat{\pi}_{t+1}^\star$ so that $\mathbb{E}[\pi_{t+1}] = \overline{\pi}$. Equation (10) highlight that, in order to meet the symmetric target, monetary policy needs to offset the asymmetry of

[23] A full derivation of the solution to the model is available in Appendix D.

[24] The effects of announcing a new monetary policy strategy is studied by Hoffmann et al. (2022); Coibion et al. (2023)

macroeconomic risks.[25] Hence, the optimal policy requires that

$$\hat{\pi}^{\star}_{t+1} = -\phi_{\pi}\left(\mu^{\star} + \psi^{\star}_{t}\right).$$

If no shocks occur, $\hat{r}^{\star}_{t} = 0$ and by announcing $\hat{\pi}^{\star}_{t} \neq 0$, the Central Bank drives inflation away from the objective $\bar{\pi}$. On the other hand, when shocks are non-zero, average inflation will be on target. Therefore, even if the Central Bank has a symmetric mandate, optimal policy requires an asymmetric strategy, i.e. a *risk-adjusted inflation targeting* (RAIT), which adjusts the central scenario by the amount necessary to offset the effect of expected risk asymmetry.

It is also worth noting that the RAIT strategy holds some substantial differences compared to other available strategies, such as the *flexible average inflation targeting* (FAIT, see Mertens and Williams, 2019). For example, FAIT only requires corrections for past mistakes,

$$\hat{\pi}^{\star}_{t} = -(1-\rho^{\star})\hat{\pi}_{t-1} + \rho^{\star}\hat{\pi}^{\star}_{t-1}, \tag{11}$$

where the parameter ρ determines the look-back period. Differently, RAIT is purely forward-looking, requiring to produce a prediction of the balance of risks to future inflation.

4.2 Operationalizing the framework

The framework laid out above considers the optimal solution for a classical one-period problem. In practice, however, a one-period target might lead the Central Bank to overreact to single, extreme predictions, or require infeasible policy actions. Therefore, here we present an operationalizable version of the RAIT, where the monetary authority targets a forward-looking inflation average, computed over multiple periods.[26] The strategy still assumes a symmetric average inflation target with a loss function defined by the squared deviation of the forward-looking average inflation, π^{A}_{t}, from the pre-specified objective $\overline{\pi}$: $E_t\left(\pi^{A}_{t} - \overline{\pi}\right)^2$, where $\pi^{A}_{t} = \sum_{j=1}^{H_{Lead}} \pi_{t+j}$. The

[25] Note that the output gap as well is zero in expectations.

[26] The loss function can be easily extended to allow for past mistakes as well. This requires defining average inflation over $H_{Lag}+H_{Lead}$ periods: $\pi^{A}_{t} = \sum_{j=-H_{Lag}+1}^{H_{Lead}} \pi_{t+j} = \frac{H_{Lag}}{H_{Lag}+H_{Lead}}\pi^{L}_{t} + \frac{H_{Lead}}{H_{Lag}+H_{Lead}}\pi^{F}_{t}$, where π^{L}_{t} and π^{F}_{t} denote the past and future components of average inflation, defined as $\pi^{L}_{t} = \frac{1}{H_{Lag}}\sum_{h=1}^{H_{Lag}} \pi_{t+1-j}$, $\pi^{F}_{t} = \frac{1}{H_{Lead}}\sum_{h=1}^{H_{Lead}} \pi_{t+j}$.

optimal policy requires $E_t\left(\pi_t^A\right) = \overline{\pi}$.

If $\pi_{t+h} \sim N\left(\mu_{t+h}, \sigma_{t+h}^2\right)$, the risk around the central scenario is symmetric, and the optimal policy implies that the Central Bank sets interest rates consistent with the central scenario, so that

$$\frac{1}{H_{Lead}} \sum_{h=1}^{H_{Lead}} \mu_{t+h|t} = \overline{\pi}. \tag{12}$$

However, when inflation risks are not symmetric, that is $\pi_{t+h} \sim F\left(\mu_{t+h}, \sigma_{t+h}, \varrho_{t+h}\right)$,[27] the optimal policy action need to account for such risks. Hence, for $\mathbb{E}_t[\pi_{t+h}] = \mu_{t+h|t} + \psi\left(\sigma_{t+h|t}, \varrho_{t+h|t}\right)$,[28] when the modal forecast deviates from the expected forecast the Central Bank is expected to target:

$$\underbrace{\frac{1}{H_{Lead}} \sum_{h=1}^{H_{Lead}} \mu_{t+h|t} - \overline{\pi}}_{\text{Optimal MP bias } (\pi_t^\star)} = \underbrace{-\frac{1}{H_{Lead}} \sum_{h=1}^{H_{Lead}} \psi\left(\sigma_{t+h|t}, \varrho_{t+h|t}\right)}_{\text{Expected risk offset}}. \tag{13}$$

Optimal policy requires the Central Bank to set the interest rate so as achieve a central scenario that deviates from the target ("optimal monetary policy bias") by the amount necessary to make up for the presence of asymmetric risks over the forecast horizon. When risks around the central scenario are unbalanced (i.e., $\varrho_{t+h|t} \neq 0$ for some h), the Central Bank should target modal forecasts that overshoot (undershoot) the target when there is negative (positive) skewness in inflation risk. Failing to reflect the negative (positive) skew in inflation risk makes overshooting (undershooting) the target more likely leading to persistent deviations for the target.

Notice that, by setting $H_{Lead} = 1$, we return to the setting described above. In that setting, the optimal bias, $\pi_{t+1}^\star$, captures the asymmetry of macroeconomic risks inherited by the risk associated with the stochastic shock. In reality, however, macroeconomic risks arise from multiple shocks, with varying properties over time. Equation (13) highlights that the optimal policy under an AIT objective only requires the Central Bank to produce real-time forecasts of inflation risk.

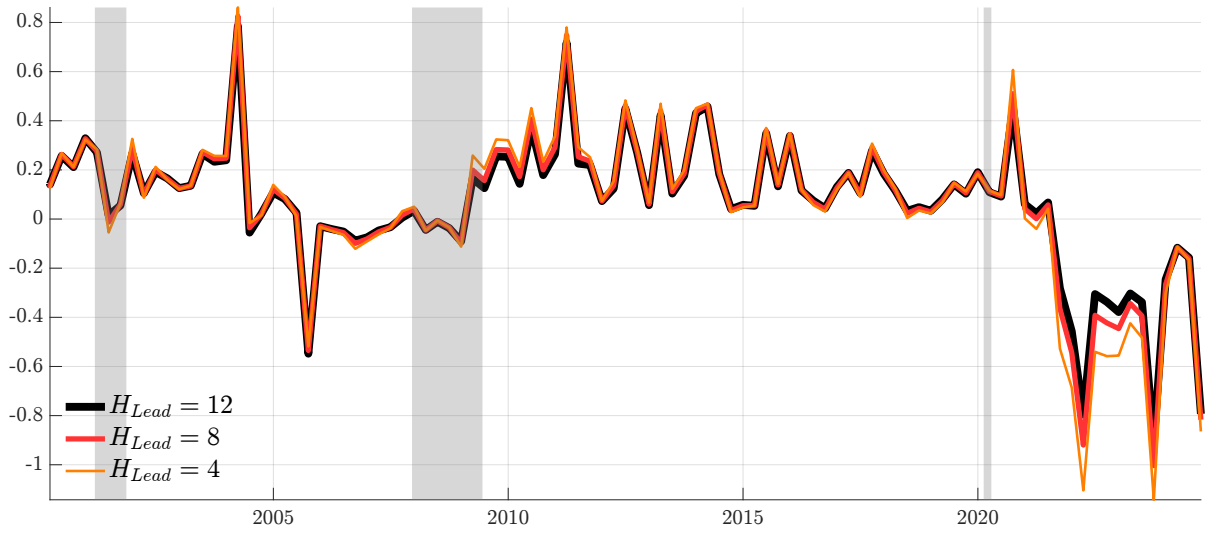

(a)

Figure 8: Real-rime expected risk offset

Note: The figure reports the real-time expected risk offset, as reported in Equation (13). We consider three H_{Lead} values of 4, 8 and 12 quarters. The sample goes from 2000 Q1 to 2023 Q4. Gray shaded areas represent NBER recessions.

5 Inflation risk and the design of AIT strategies

In this section we show how estimates of inflation risks produced by our model can be exploited by the monetary authority to design an appropriate approach to AIT.

A well-designed AIT framework requires two key elements: first, the desired inflation objective, $\pi^\star$, that is the inflation level that the Central Bank wants to achieve, and second, the definition of *"average inflation"*. Specifically, the latter entails determining a *look-back* and a *look-forward* period, over which the Central Bank assesses development in inflation (i.e. H_{Lag} and H_{Lead}, respectively). While there is a wide consensus among Central Banks to set the inflation objective to (or close to) 2% (see Bernanke, 2003), less attention has been devoted to the careful choice and the implications of different periods over which the average target is evaluated. Intuitively, however, a short look-back period is undesirable, as it would require the Central Bank to implement policies that can be unfeasible or too aggressive, whereas it is expected that the look-forward period account for the lags in the transmission of monetary policy. Based on evidence in (**?** and **?**), a look-forward period of 2 to 3 years seems an appropriate choice to capture the full effects of policy

[27] Asymmetric densities have been widely used to communicate future inflation outlooks since the mid-1990s (see, e.g., Wallis, 1999). See (Wallis, 2014) for an overview of the history of asymmetric (two-piece) distributions.

[28] For the distribution considered in this paper, $\psi = g(\eta)\sigma_t\varrho_t$, see equation Equation (5).

adjustments.

The FAIT approach in Equation (11), for example, rests on the assumption that past mistakes are sole guidance to inform the policy action, therefore setting $H_{Lead} = 0$. The strength of recent mistakes are regulated by the coefficient $\rho^{\star}$, which implicitly reflect the choice of H_{Lag},[29] On the other hand, the RAIT framework we introduce is based on the idea that what really matters for the conduit of monetary policy are the expected risks that can hinder a smooth approach of inflation to its target value. That is, $H_{Lag} = 0$, but H_{Lead} needs to be set in line with a realistic forecast horizon. In Figure 8 we show the evolution of the expected risk offset (see Equation (13)) estimated in real-time with our *Skt* model. We consider look-ahead periods of four, eight and twelve quarters. Regardless of H_{Lead}, all series suggest similar patterns, with minimal deviations required by shorter look-ahead periods emerging only during periods of substantial inflationary pressures. This results is mainly due to the fact that transitory deviations of the skewness are weakly autocorrelated, such that the predicted level of risk quickly reverts to its permanent component. The model suggests that a positive make up offset of 20 to 40 basis points was necessary to keep inflation at target in the aftermath of the GFC. The failure to target a modal scenario slightly above the 2% inflation objective could explain the persistent deflationary bias observed in realized inflation prints. In the post-COVID years, the model advises that large negative offsets might be required going forward.

The implicit inflation targets required by the FAIT and RAIT approaches are compared in Figure 9. Specifically, we compare the requirement form the FAIT approach, which is by definition a smooth weighted average of historical deviations from the target and the current inflation miss. For the RAIT, we produce a similar smoothed target, but using the real-time average expected levels of risk, with a look-ahead periods of one year. We opt for a longer look-back period of two years for the smoothed averages. The policy recommendation provided by the two approaches turn out to be strikingly different.

First, FAIT always require more aggressive policies. Under FAIT, the Central Bank is required to target expected inflation levels around 2.4% in the decade following the GFC, due to the persistent undershooting of the target. The large misses in the post-pandemic era, on the other hand, call for a massive undershooting of the objective, which seem unrealistic to achieve. Differently,

[29]For example, a smoothing coefficient of 0.95 corresponds to a look-back period of about 2 years.

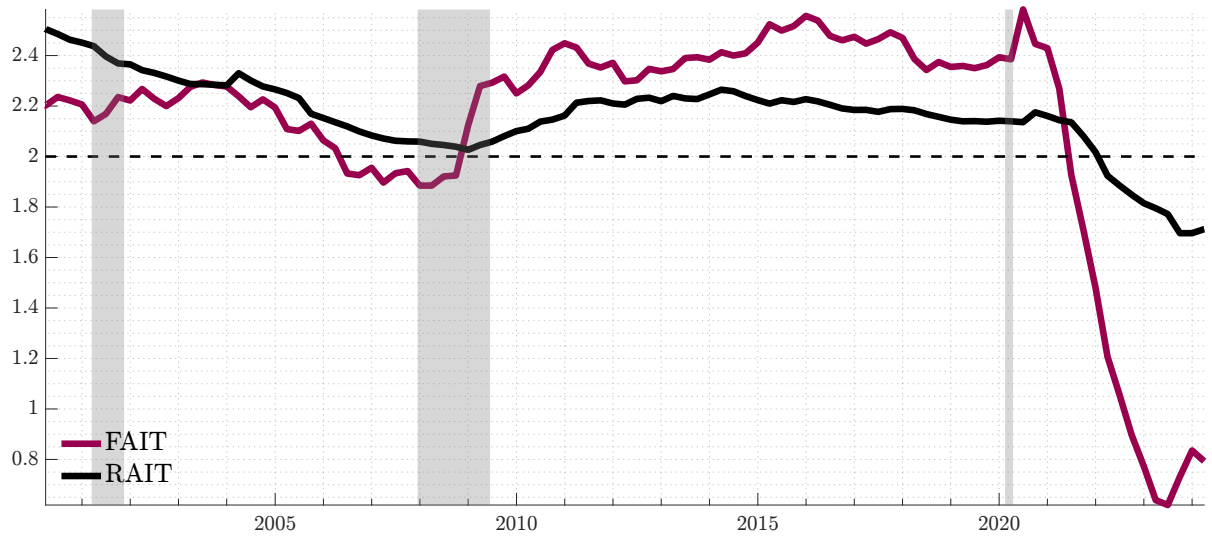

Figure 9: Inflation target under FAIT and RAIT

Note: The figure reports smoothed, real-time estimates of the implicit inflation target required under the FAIR and RAIT approaches. The sample goes from 2000 Q1 to 2023 Q4. Gray shaded areas represent NBER recessions.

the RAIT, being a purely forward-looking rule based on risk, prescribes inflation targets that are more stable and realistic in size. In the post-GCF, RAIT would have required less aggressive monetary policy actions, consistent with a target expected inflation level of 2.2%. Intuitively, if inflation projections are skewed towards the downside, the monetary authority which incorporates this expectation in the policy response is better off by targeting a central scenario that is above the desired target by a factor proportional to the expected risk in the forecasting period. In recent years, the forward-looking nature of the RAIT requires an undershooting of the objective by just 0.4% (against the 1.2% commanded by the FAIT), as expected risks are expected to recede.

Are mistakes predictable? Information about next period's miss can be extracted from our real-time density forecasts. Figure 10, panel (a) reports the predicted probability that next inflation print will be lower than an year before (yellow-shaded areas). The figure compares this probability with a counterfactual probability where the policymaker disregards the asymmetry in risk (e.g., setting $\varrho_t = 0$). Ignoring the asymmetry in risks tends to underestimate the probability of lower prints in periods of downside risk, such as after the Great Recession. When our model quickly updates the sign of the skewness, around 2022, the likelihood of lower prints becomes systematically smaller when asymmetry is accounted for.

Hence, monitoring the asymmetry of inflation risks in real-time is an essential tool for accurately

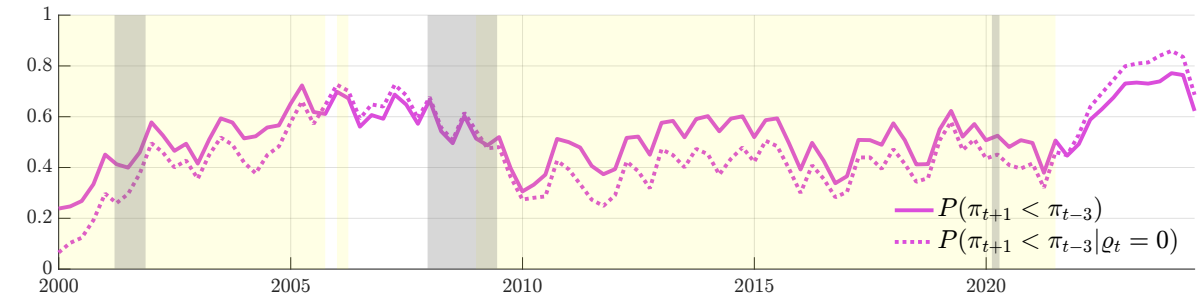

(a) Probability of higher past mistakes

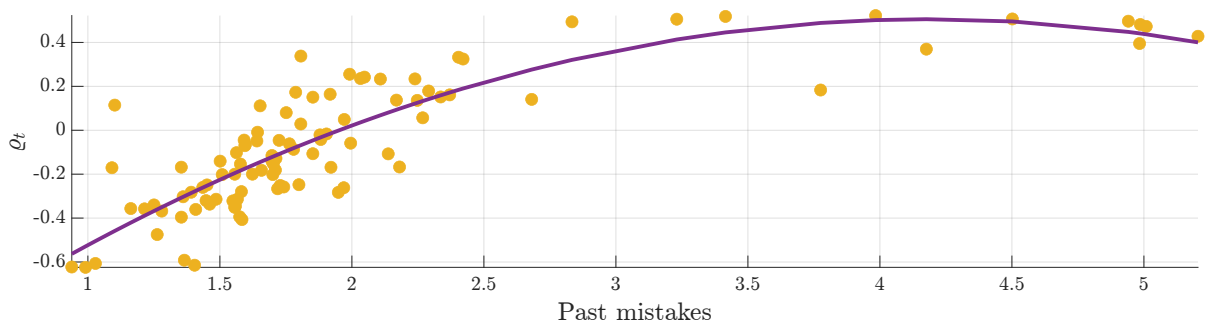

(b) Correlation between past mistakes and asymmetry

Figure 10: Expected asymmetry as predictor of target misses

Note: The top panel report the real-time assessment of the probability of increasing the past mistake components of the optimal monetary policy bias, defines as $P_t\left(\pi_{t+1} < \pi_{t-3}\right)$. The dotted lines reports a counterfactual probability computed assuming symmetry of the predictive distribution, that is $P_t\left(\pi_{t+1} < \pi_{t-3}|\varrho_t = 0\right)$. Yellow shaded bands represent period when realized inflation was higher than π_{t-4}. Panel (b) show the correlation between the past mistakes components and the forward-looking, filtered asymmetry estimated in real-time, ϱ_t. The sample goes from 2000 Q1 to 2023 Q4. Gray shaded areas represent NBER recessions.

gauging the costs of under/overshooting optimal policy prescriptions. In the bottom panel, we further show that tracking asymmetry in real time can further inform about the expected mistakes. When downside (upside) risk dominates, prolonged periods of negative (positive) inflation gaps are more likely. Even if the Central Bank were to ignore future risk asymmetry (i.e., $H_{Lead} = 0$), tracking inflation skewness can still provide insight into the sing an magnitude of next periods inflation misses.

6 Concluding remarks

Central Banks' mandates of price stability are strenuous and complex tasks, that require careful calibration of policy interventions and a consistent assessment of the risks associated to them. Among these risks stands the risk associated with future realizations of inflation, that implicitly

define the target of the price stability mandate.

We show that the process describing the dynamics of core PCE in the postwar period features frequent and persistent changes to the balance of risks. We propose a model to track the whole distribution of inflation in real-time which retains most of the desirable features of the trend-cycle model of Stock and Watson (2002), but further allows for time variation in the skewness of the predictive distributions of the inflation process. We show that inflation skewness has varied substantially over time, shifting from being positive in the 1970s and 1980s to negative from 1990s until the recent bout of inflation in the post-COVID period. Furthermore, we highlight that such model enhancement is crucial to improve point, density and event forecasts.

In the context of the recently adopted average inflation targeting, we argue that the optimal conduit of monetary policy ought to account for these persistent changes in the prevalent risks to the inflation outlook. Specifically, in an AIT framework with a forward-looking component, optimal policy is required to make-up for the perceived tilt in the balance of risks. Failure to accommodate the presence of these risks might lead to prolonged periods of missing the infaltion target.

References

Andrade, P., E. Ghysels, and J. Idier (2014): "Inflation risk measures and their informational content," *Available at SSRN 2439607*.

Antolín-Díaz, J., T. Drechsel, and I. Petrella (2024): "Advances in nowcasting economic activity: The role of heterogeneous dynamics and fat tails," *Journal of Econometrics*, 238.

Arellano-Valle, R. B., H. W. Gómez, and F. A. Quintana (2005): "Statistical inference for a general class of asymmetric distributions," *Journal of Statistical Planning and Inference*, 128, 427–443.

Ascari, G. and A. M. Sbordone (2014): "The macroeconomics of trend inflation," *Journal of Economic Literature*, 52, 679–739.

Bai, J. and S. Ng (2005): "Tests for skewness, kurtosis, and normality for time series data," *Journal of Business & Economic Statistics*, 23, 49–60.

Bernanke, B. S. (2003): "Remarks by Governor Ben S. Bernanke At the 28th Annual Policy Conference: Inflation Targeting: Prospects and Problems, Federal Reserve Bank of St. Louis, St. Louis, Missouri October 17, 2003," .

Bianchi, F., L. Melosi, and M. Rottner (2021): "Hitting the elusive inflation target," *Journal of Monetary Economics*, 124, 107–122.

Blasques, F., S. J. Koopman, and A. Lucas (2014): "Stationarity and ergodicity of univariate generalized autoregressive score processes," *Electronic Journal of Statistics*, 8, 1088–1112.

——— (2015): "Information-theoretic optimality of observation-driven time series models for continuous responses," *Biometrika*, 102, 325–343.

Blasques, F., J. van Brummelen, S. J. Koopman, and A. Lucas (2022): "Maximum likelihood estimation for score-driven models," *Journal of Econometrics*, 227, 325–346.

Bullard, J. B. (2018): "What Is the Best Strategy for Extending the U.S. Economy's Expansion?: a presentation at the CFA Society Chicago—Distinguished Speaker Series Breakfast, Chicago, Ill," Speech 320, Federal Reserve Bank of St. Louis.

Chan, J. C. (2013): "Moving average stochastic volatility models with application to inflation forecast," *Journal of Econometrics*, 176, 162–172.

Clarida, R., J. Gali, and M. Gertler (1999): "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 37, 1661–1707.

Cogley, T. (2002): "A Simple Adaptive Measure of Core Inflation," *Journal of Money, Credit and Banking*, 34, 94–113.

Cogley, T. and A. M. Sbordone (2008): "Trend inflation, indexation, and inflation persistence in the New Keynesian Phillips curve," *American Economic Review*, 98, 2101–2126.

Coibion, O., Y. Gorodnichenko, E. S. Knotek, and R. Schoenle (2023): "Average inflation targeting and household expectations," *Journal of Political Economy Macroeconomics*, 1, 403–446.

Cox, D. R. (1981): "Statistical analysis of time series: Some recent developments," *Scandinavian Journal of Statistics*, 93–115.

Creal, D., S. J. Koopman, and A. Lucas (2013): "Generalized autoregressive score models with applications," *Journal of Applied Econometrics*, 28, 777–795.

De Polis, A. (2023): "Conditional asymmetries and downside risks in macroeconomic and financial time series," Ph.D. thesis, University of Warwick.

Delle Monache, D., A. De Polis, and I. Petrella (2024): "Modeling and forecasting macroeconomic downside risk," *Journal of Business & Economic Statistics*, 42, 1010–1025.

Delle Monache, D. and I. Petrella (2017): "Adaptive models and heavy tails with an application to inflation forecasting," *International Journal of Forecasting*, 33, 482–501.

DelleMonache, D. and I. Petrella (2017): "Adaptive models and heavy tails with an application to inflation forecasting," *International Journal of Forecasting*, 33, 482–501.

Diebold, F. X. and R. S. Mariano (1995): "Comparing predictive accuracy," *Journal of Business & Economic Statistics*, 20, 134–144.

Dolado, J., P. R. María-Dolores, and R.-M. F. J. (2004): "Nonlinear Monetary Policy Rules: Some New Evidence for the U.S," *Studies in Nonlinear Dynamics & Econometrics*, 8, 1–34.

Evans, C., J. Fisher, F. Gourio, and S. Krane (2020): "Risk Management for Monetary Policy Near the Zero Lower Bound," Conference draft, Brooking Paper on Economic Activity.

Faust, J. and J. H. Wright (2013): "Forecasting inflation," in *Handbook of Economic Forecasting*, Elsevier, vol. 2, 2–56.

FED (2021): "Review of Monetary Policy Strategy, Tools, and Communications," .

Galí, J. (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

Giannoni, M. and M. Woodford (2004): "Optimal inflation-targeting rules," in *The inflation-targeting debate*, University of Chicago Press, 93–172.

Gneiting, T. and R. Ranjan (2011): "Comparing density forecasts using threshold-and quantile-weighted scoring rules," *Journal of Business & Economic Statistics*, 29, 411–422.

Gómez, H. W., F. J. Torres, and H. Bolfarine (2007): "Large-sample inference for the epsilon-skew-t distribution," *Communications in Statistics—Theory and Methods*, 36, 73–81.

Gordon, R. J. (1970): "The recent acceleration of inflation and its lessons for the future," *Brookings Papers on Economic Activity*, 1970, 8–47.

Haario, H., E. Saksman, and J. Tamminen (1999): "Adaptive proposal distribution for random walk Metropolis algorithm," *Computational Statistics*, 14, 375–396.

Harvey, A. and S. Thiele (2016): "Testing against changing correlation," *Journal of Empirical Finance*, 38, 575–589.

Harvey, A. C. (2013): *Dynamic models for volatility and heavy tails: with applications to financial and economic time series*, vol. 52, Cambridge University Press.

Hilscher, J., A. Raviv, and R. Reis (2022): "How likely is an inflation disaster?" .

Hoffmann, M., E. Moench, L. Pavlova, and G. Schultefrankenfeld (2022): "Would households understand average inflation targeting?" *Journal of Monetary Economics*, 129, 52–66.

Kilian, L. and S. Manganelli (2007): "Quantifying the risk of deflation," *Journal of Money, Credit and Banking*, 39, 561–590.

——— (2008): "The central banker as a risk manager: Estimating the Federal Reserve's preferences under Greenspan," *Journal of Money, Credit and Banking*, 40, 1103–1129.

Korobilis, D., B. Landau, A. Musso, and A. Phella (2021): "The time-varying evolution of inflation risks," .

Le Bihan, H., D. Leiva-Leon, and M. Pacce (2023): "Underlying inflation and asymmetric risks," Working Paper Series 2848, European Central Bank.

Lopez-Salido, D. and F. Loria (2020): "Inflation at risk," .

Manzan, S. and D. Zerom (2013): "Are macroeconomic variables useful for forecasting the distribution of US inflation?" *International Journal of Forecasting*, 29, 469–478.

——— (2015): "Asymmetric quantile persistence and predictability: the case of US inflation," *Oxford Bulletin of Economics and Statistics*, 77, 297–318.

Mertens, T. M. and J. C. Williams (2019): "Tying Down the Anchor: Monetary Policy Rules and the Lower Bound on Interest Rates," Federal Reserve Bank of San Francisco, Working Paper 2019-14, 2019.

Mudholkar, G. S. and A. D. Hutson (2000): "The epsilon–skew–normal distribution for analyzing near-normal data," *Journal of Statistical Planning and Inference*, 83, 291–309.

Reichlin, L., K. Adam, W. J. McKibbin, M. McMahon, R. Reis, G. Ricco, and W. di Mauro B (2021): "The ECB strategy: The 2021 review and its future," Tech. rep., CEPR Press.

Stock, J. H. and M. W. Watson (2002): "Has the business cycle changed and why?" *NBER Macroeconomics Annual Report*, 17, 159–218.

——— (2007): "Why has US inflation become harder to forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

——— (2016): "Core inflation and trend inflation," *Review of Economics and Statistics*, 98, 770–784.

Surico, P. (2007): "The Fed's monetary policy rule and U.S. inflation: The case of asymmetric preferences," *Journal of Economic Dynamics and Control*, 31, 305–324.

Svensson, L. E. (1997): "Inflation forecast targeting: Implementing and monitoring inflation targets," *European Economic Review*, 41, 1111–1146.

Svensson, L. E. and M. Woodford (2004): "Implementing optimal policy through inflation-forecast targeting," in *The inflation-targeting debate*, University of Chicago Press, 19–92.

Wallis, K. F. (1999): "Asymmetric density forecasts of inflation and the Bank of England's fan chart," *National Institute Economic Review*, 167, 106–112.

——— (2014): "The two-piece normal, binormal, or double Gaussian distribution: its origin and rediscoveries," *Statistical Science*, 106–112.

Woodford, M. (2003): "Interest and prices," .

A Score-driven framework

A.1 Score derivations

The scaled score s_t is a non-linear function of past observations and past parameters' values. For $\ell_t = \log \mathcal{D}(\theta, f_t)$ being the Skew-t of Gómez et al. (2007), $y_t|Y_{t-1} \sim Skt_\nu(\mu_t, \sigma_t^2, \rho_t)$, the log-likelihood takes the form

$$\ell_t(r_t|\theta, \mathcal{F}_{t-1}) = \log \mathcal{C}(\nu) - \frac{1}{2}\log \sigma_t^2 - \frac{1+\nu}{2}\log\left[1 + \frac{\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}\right], \tag{14}$$
$$\log \mathcal{C}(\nu) = \log\Gamma\left(\frac{\nu+1}{2}\right) - \log\Gamma\left(\frac{\nu}{2}\right) - \frac{1}{2}\log\nu - \frac{1}{2}\log\pi,$$

where $\Gamma(\cdot)$ is the Gamma function and $\nu > 3$ are the degrees of freedom. Differentiating (14) with respect to location, scale and asymmetry we obtain the gradient vector $\nabla_t = \left[\frac{\partial \ell_t}{\partial \mu}, \frac{\partial \ell_t}{\partial \sigma_t^2}, \frac{\partial \ell_t}{\partial \rho_t}\right]'$. Recall that $\varepsilon_t = y_t - \mu_t$, $\zeta_t = \frac{\varepsilon_t}{\sigma_t}$ and let

$$f(\mu_t, \sigma_t^2, \rho_t) = 1 + \frac{\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2} = \frac{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2 + \varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}$$

To avoid overburdening the notation, in what follows $\frac{\partial f(x)}{\partial x} = f'_x$ and $a = -\frac{1+\nu}{2}$. The score with respect to the location parameter reads

$$\frac{\partial \ell_t}{\partial \mu_t} = w_t \frac{\zeta_t}{\sigma_t}, \qquad \text{with} \qquad w_t = \frac{\nu+1}{\nu\left(1+\mathrm{s}\left(\varepsilon_t\right)\rho_t\right)^2 + \zeta_t^2}.$$

Proof. Define

$$g(\mu_t) = a \log f(\mu_t, \sigma_t^2, \rho_t),$$

such that $\frac{\partial \ell_t}{\partial \mu_t} = \frac{\partial g(\mu_t)}{\partial \mu_t} = a\frac{f'_{\mu_t}}{f(\mu_t, \sigma_t^2, \rho_t)}$. For

$$f'_{\mu_t} = -\frac{2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}\varepsilon_t,$$

it follows:

$$\begin{aligned}\frac{\partial \ell_t}{\partial \mu_t} &= \frac{1+\nu}{2}\frac{2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}\cdot\varepsilon_t\cdot\frac{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2+\varepsilon_t^2}\\ &= \frac{(1+\nu)}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2+\varepsilon_t^2}\varepsilon_t\\ &= \omega_t\frac{\zeta_t}{\sigma_t}\end{aligned}$$

. □

The score with respect to the squared scale parameter reads

$$\frac{\partial \ell_t}{\partial \sigma_t^2} = \frac{(w_t\zeta_t^2-1)}{2\sigma_t^2}.$$

Proof. Define

$$g(\sigma_t^2) = -\frac{\log\sigma_t^2}{2} + a\log f(\mu_t,\sigma_t^2,\rho_t),$$

such that $\frac{\partial \ell_t}{\partial \sigma_t^2} = \frac{\partial g(\sigma_t^2)}{\partial \sigma_t^2} = -\frac{1}{2\sigma_t^2} + a\frac{f'_{\sigma_t^2}}{f(\mu_t,\sigma_t^2,\rho_t)}$, with $f'_{\sigma_t^2} = -\frac{\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^4}$. It follows that:

$$\begin{aligned}\frac{\partial \ell_t}{\partial \sigma_t^2} &= -\frac{1}{2\sigma_t^2} - \frac{1+\nu}{2}\cdot\left[-\frac{\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^4}\cdot\frac{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2+\varepsilon_t^2}\right]\\ &= -\frac{1}{2\sigma_t^2} - \frac{1+\nu}{2}\cdot\left[-\frac{\varepsilon_t^2}{\sigma_t^2}\cdot\frac{1}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2+\varepsilon_t^2}\right]\\ &= -\frac{1}{2\sigma_t^2} + \frac{w_t\zeta_t^2}{2\sigma_t^2} = \frac{(w_t\zeta_t^2-1)}{2\sigma_t^2}\end{aligned}$$

. □

The score with respect to the shape parameter reads as

$$\frac{\partial \ell_t}{\partial \rho_t} = \frac{\mathrm{s}(\varepsilon_t)}{(1+\mathrm{s}(\varepsilon_t)\rho_t)}w_t\zeta_t^2.$$

Proof. Define

$$g(\rho_t) = a\log f(\mu_t,\sigma_t^2,\rho_t),$$

such that $\frac{\partial \ell_t}{\partial \rho_t} = \frac{\partial g(\rho_t)}{\partial \sigma_t^2} = a\frac{f'_{\rho_t}}{f(\mu_t,\sigma_t^2,\rho_t)}$, with $f'_{\rho_t} = -\frac{2(\mathrm{s}(\varepsilon_t)+\rho_t)\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^4\sigma_t^2}$. It follows that:

$$\begin{aligned}\frac{\partial \ell_t}{\partial \rho_t} &= \frac{1+\nu}{2}\cdot\frac{2(\mathrm{s}(\varepsilon_t)+\rho_t)\varepsilon_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^4\sigma_t^2}\cdot\frac{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2}{\nu(1+\mathrm{s}(\varepsilon_t)\rho_t)^2\sigma_t^2+\varepsilon_t^2} \\ &= \frac{(\mathrm{s}(\varepsilon_t)+\rho_t)\varepsilon_t^2}{(1+\mathrm{s}(\varepsilon_t)\rho_t)^2}\frac{w_t}{\sigma_t^2} = \frac{\mathrm{s}(\varepsilon_t)}{(1+\mathrm{s}(\varepsilon_t)\rho_t)}w_t\zeta_t^2\end{aligned}$$

. □

A.2 Scaled scores

Given we model $\gamma_t = \log\sigma_t$ and $\delta_t = atanh(\rho_t)$, for the chain rule we have:

$$\frac{\partial \ell_t}{\partial \gamma_t} = \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \gamma_t}, \quad \frac{\partial \ell_t}{\partial \delta_t} = \frac{\partial \ell_t}{\partial \rho_t}\frac{\partial \rho_t}{\partial \delta_t}, \tag{15}$$

where $\frac{\partial \sigma_t^2}{\partial \gamma_t} = 2\sigma_t^2$ and $\frac{\partial \rho_t}{\partial \delta_t} = (1-\rho_t^2)$. We can thus define the vector of interest as $f_t = (\mu_t, \gamma_t, \delta_t)'$ with the associated Jacobian matrix

$$J_t = \frac{\partial(\mu_t,\sigma_t^2,\rho_t)}{\partial f_t'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2\sigma_t^2 & 0 \\ 0 & 0 & 1-\rho_t^2 \end{bmatrix}. \tag{16}$$

The Fisher information matrix is computed as the expected value of outer product of the gradient vector. Given the degrees of freedom $\nu > 3$ this is computed as:

$$\mathcal{I}_t = \mathbb{E}_{t-1}[\nabla_t\nabla_t'] = \begin{bmatrix} \frac{(1+\nu)}{(\nu+3)(1-\rho_t^2)\sigma_t^2} & 0 & \frac{4(1+\nu)}{\sigma_t(1-\rho_t^2)(3+\nu)} \\ 0 & \frac{1}{2(3+\nu)\sigma_t^4} & 0 \\ \frac{4(1+\nu)}{\sigma_t(1-\rho_t^2)(3+\nu)} & 0 & \frac{3(1+\nu)}{(1-\rho_t^2)(3+\nu)} \end{bmatrix}. \tag{17}$$

As a result, the vector of scaled scores reads as:

$$\mathbf{s}_t = (J_t' \text{diag}(\mathcal{I}_t) J_t)^{-1} J_t' \nabla_t = \begin{bmatrix} s_{\mu t} \\ s_{\sigma t} \\ s_{\rho t} \end{bmatrix} = \chi \begin{bmatrix} (1-\rho_t^2) w_t \varepsilon_t \\ (\nu+1)(w_t \varepsilon_t^2 - \sigma_t^2) \\ \text{s}(\varepsilon_t)(1-\text{s}(\varepsilon_t)\rho_t) w_t \frac{\varepsilon_t^2}{3\sigma_t^2} \end{bmatrix}. \tag{18}$$

with $\chi = \frac{(\nu+3)}{(\nu+1)}$ and $w_t = \frac{\nu+1}{\nu(1+\text{s}(\varepsilon_t)\rho_t)^2+\zeta_t^2}$.

B Evidence for other inflation measures

In this appendix we report additional results about other policy relevant measures of inflation. Specifically, we consider, the GDP deflator, headline PCE and core and headline CPI. All samples go from 1960 Q1 to 2024 Q2.

Table 4 collect the test statistics for the detection of time variation in the asymmetry for all four measures of inflation. Overall, the null of restricted asymmetry is strongly rejected.

Table 4: Time variation in higher order moments

	Q	Q^*	N	Q	Q^*	N
	GDP Deflator			**Headline PCE**		
	Homoskedastic					
Shape	637.470***	644.910***	5.690***	303.820***	307.370***	6.460***
	Heteroskedastic					
$Scale^2$	597.120***	604.090***	4.050***	566.190***	572.800***	2.330***
Shape	154.150***	155.950***	2.780***	148.610***	150.350***	1.890***
	Core CPI			**Headline CPI**		
	Homoskedastic					
Shape	840.710***	850.480***	3.290***	407.600***	412.340***	4.220***
	Heteroskedastic					
$Scale^2$	556.980***	563.460***	3.810***	730.210***	738.700***	3.430***
Shape	185.210***	187.360***	3.260***	183.040***	185.160***	2.150***

Note: Q is the portmanteau test, Q* is the Ljung-Box extension (with automatic lag selection) and N corresponds to the Nyblom test. Q and Q* are distributed as a χ^2_1, while N is distributed as a Cramer von-Mises distribution with 1 degree of freedom. * p $<$ 10%, ** p $<$ 5%, *** p $<$ 1%.

Based on previous results, we also report the estimates for the time-varying moments of the four measures.

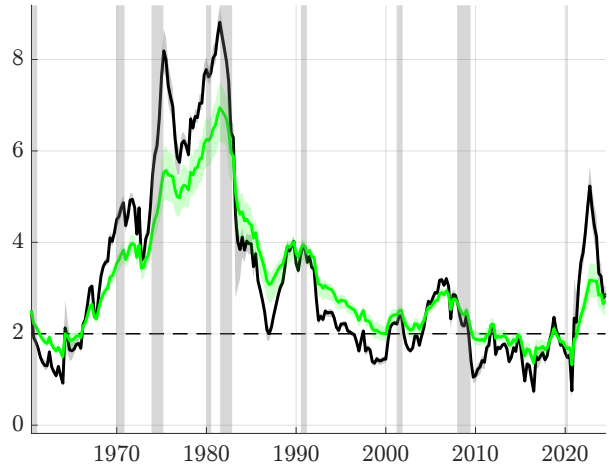

(a) GDP Deflator

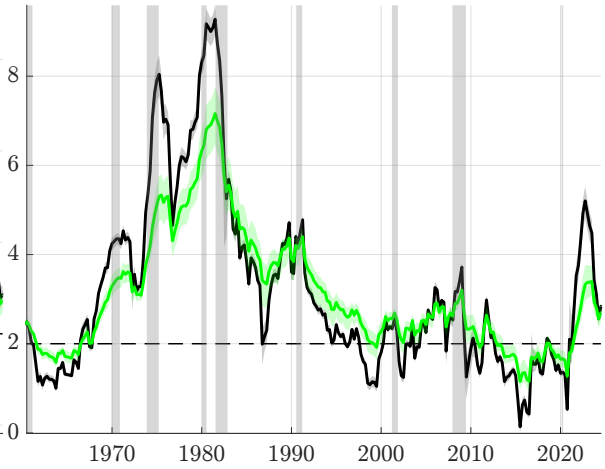

(b) Headline PCE

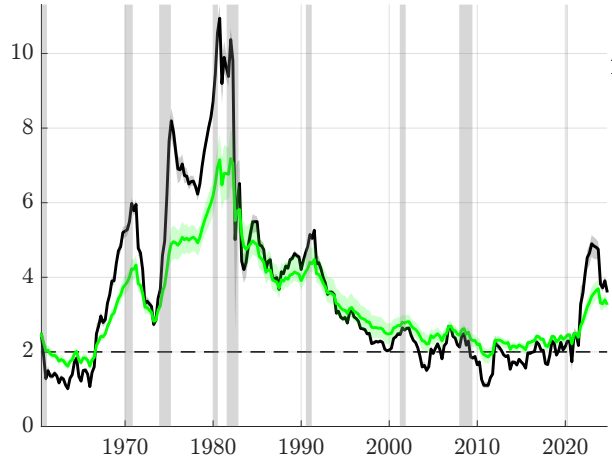

(c) Core CPI

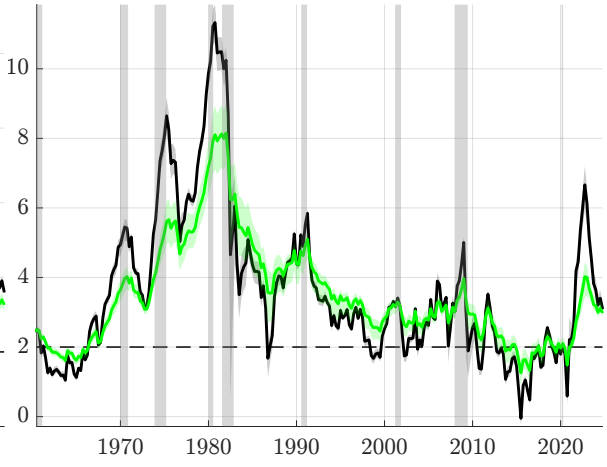

(d) Headline CPI

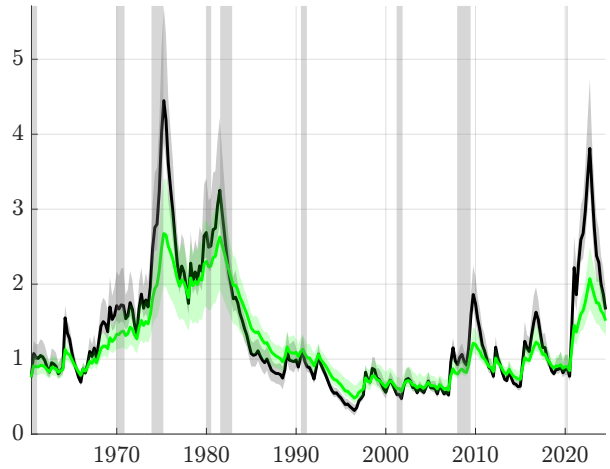

(a) GDP Deflator

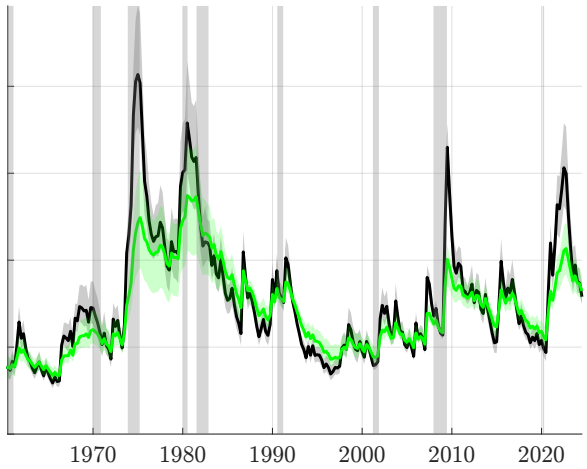

(b) Headline PCE

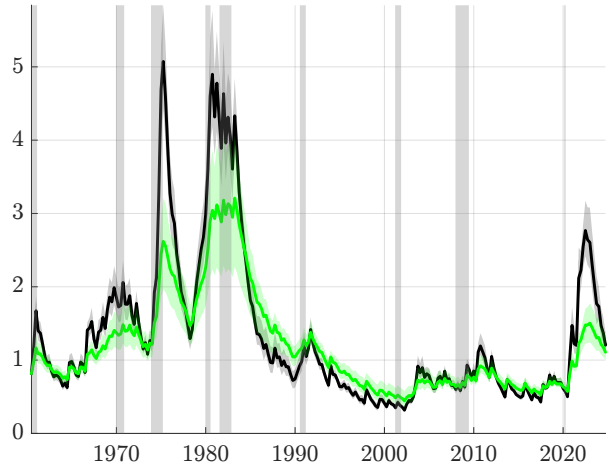

(c) Core CPI

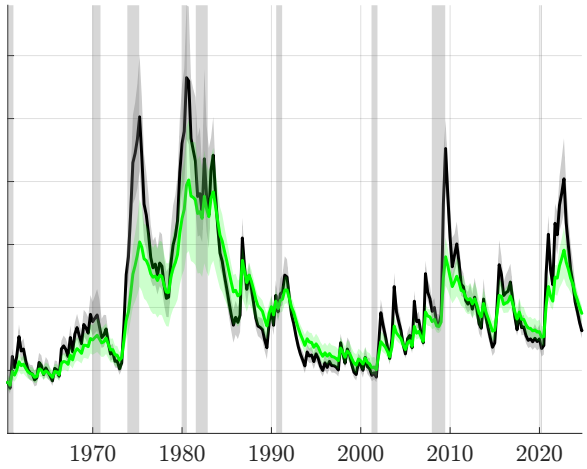

(d) Headline CPI

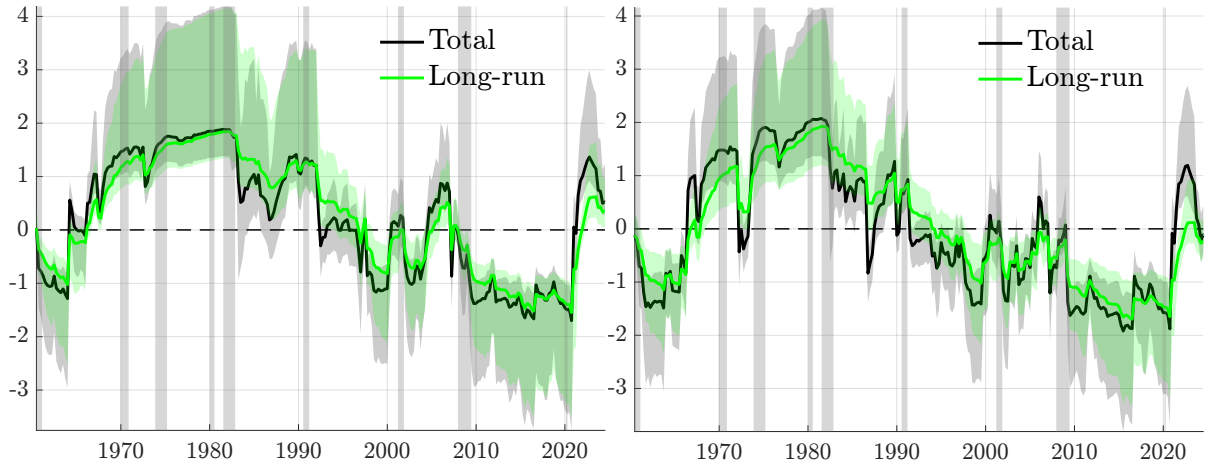

(a) GDP Deflator

(b) Headline PCE

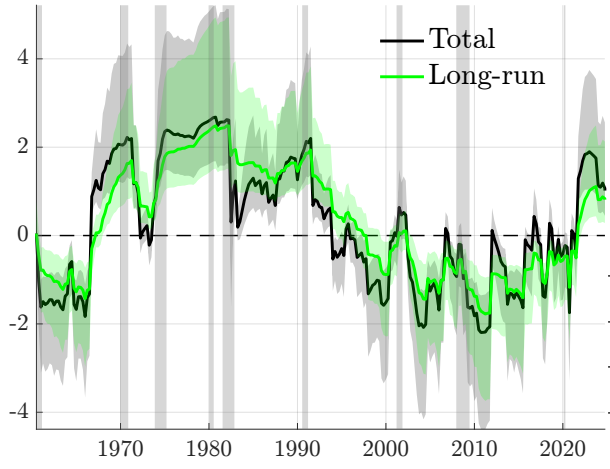

(c) Core CPI

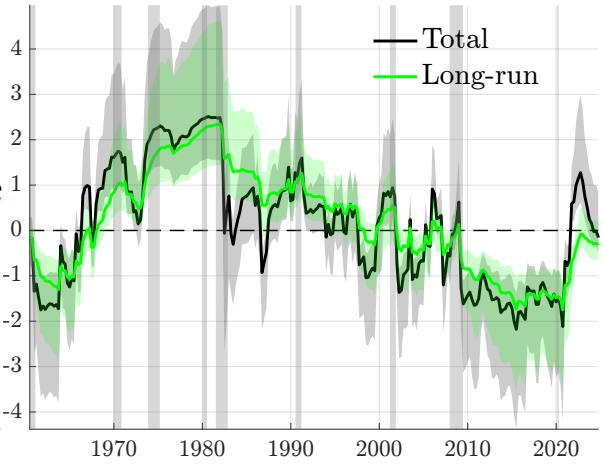

(d) Headline CPI

C Additional results

Table 5: Out-of-sample comparison - *Student t*

	h = 1	h = 2	h = 3	h = 4	h = 8
MSFE	**0.839** (0.085)	**0.854** (0.051)	**0.908** (0.081)	0.942 (0.129)	1.008 (0.782)
CRPS	**0.938** (0.074)	**0.950** (0.077)	**0.960** (0.081)	**0.959** (0.074)	0.996 (0.331)
	CRPS decomposition				
Right	0.927 (0.109)	**0.914** (0.075)	**0.935** (0.099)	**0.936** (0.074)	0.981 (0.116)
Left	**0.952** (0.088)	0.982 (0.258)	0.986 (0.197)	0.983 (0.200)	1.009 (0.808)
Center	**0.936** (0.058)	0.953 (0.117)	**0.960** (0.041)	**0.960** (0.060)	0.999 (0.476)

Note: The table report the relative performance of a t model against our *Skt* model. Results are reported in ratios, with our model being at the numerator; values smaller than 1 imply superior predictive accuracy of the *SKt* model. The out-of-sample period runs from 2000Q1 to 2024Q2. Values in **bold** are significant at the 10% level.

Table 6: Out-of-sample comparison - $\varrho = 0,\ \forall t$

	h = 1	h = 2	h = 3	h = 4	h = 8
MSFE	**0.832** (0.081)	0.884 (0.115)	**0.906** (0.057)	0.966 (0.190)	1.055 (0.997)
CRPS	0.947 (0.116)	0.961 (0.153)	**0.962** (0.075)	0.978 (0.166)	1.024 (0.983)
	CRPS decomposition				
Right	**0.928** (0.081)	**0.928** (0.053)	**0.947** (0.048)	0.962 (0.123)	1.012 (0.854)
Left	0.967 (0.242)	0.995 (0.440)	0.988 (0.339)	0.998 (0.461)	1.039 (0.994)
Center	0.946 (0.123)	0.961 (0.150)	**0.954** (0.053)	0.977 (0.155)	1.024 (0.970)

Note: The table report the relative performance of our *Skt* model when skewness is omitted ($\varrho_t = 0,\ \forall t$) against our *Skt* model. Results are reported in ratios, with our model being at the denominator; values smaller than 1 imply superior predictive accuracy of the *SKt* model. The out-of-sample period runs from 2000Q1 to 2024Q2. Values in **bold** are significant at the 10% level.

Table 7: Deep parameters estimates

Autocorrelations					
ϕ_μ	ϕ_γ	ϕ_δ			
0.990	0.853	0.803			
(0.006)	(0.068)	(0.055)			
Learning rates					
a_μ	b_μ	a_γ	b_γ	a_δ	b_δ
0.095	0.094	0.084	0.088	0.041	0.085
(0.005)	(0.005)	(0.013)	(0.011)	(0.011)	(0.013)
Degrees of freedom					
η					
0.130					
(0.035)					

Note: The table reports mean estimates of the deep parameters of the model. Parameters standard deviations are reported in parentheses.

D Solving the New Keynesian model with mark-up shock

Let us start from the baseline three-equation New Keynesian model for output, inflation and the interest rate,

$$
\begin{aligned}
x_t &= \mathbb{E}_t\left[x_{t+1}\right] - \varsigma^{-1}\left[i_t - \mathbb{E}_t\left[\pi_{t+1}\right] - r_t^*\right], \\
\pi_t &= \beta\,\mathbb{E}_t\left[\pi_{t+1}\right] + \kappa x_t, \\
i_t &= \pi_t + \phi_\pi\left(\pi_t - \pi_t^*\right),
\end{aligned}
$$

and let $r_t^* = \varsigma\left[\mathbb{E}_t\left[y_{t+1}^*\right] - y_t^*\right]$ and $y_t^* = \omega\alpha_t$ denote the natural rate of interest and the output under fully-flexible prices, respectively. $\alpha_t \sim iidF\left(0, \sigma_\alpha, \varrho_{\alpha,t}\right)$ is a TFP shock distributed according to a general class of unimodal, (possibly) asymmetric densities, such that $\mathbb{E}\left[\alpha_t\right] = \psi_{\alpha,t} \neq 0$ if $\varrho_{\alpha,t} \neq 0$. It follows that $y_t^* \sim iidF\left(0, \omega\sigma_\alpha, \varrho_{\alpha,t}\right)$ and $r_t^* = \varsigma\omega\left(\mathbb{E}_t\left[\alpha_{t+1}\right] - \alpha_t\right)$; therefore

$$
r_t^* \sim iidF\left(\mu_t^\star, \sigma_t^\star \varrho_t^\star\right)
$$

where $\mu_t^\star = \varsigma\omega\,\mathbb{E}_t\left[\psi_{\alpha,t+1}\right]$, $\sigma_t^\star = \varsigma\omega\sigma_\alpha$ and $\varrho_t^\star = -\rho_{\alpha,t}$. Note that $\mathbb{E}\left[r_t^*\right] = \varsigma\omega\left(\mathbb{E}_t\left[\psi_{\alpha,t+1}\right] - \psi_{\alpha,t}\right)$, and so $\mathbb{E}\left[r_t^*\right] \neq 0$ if $\mathbb{E}_t\left[\psi_{\alpha,t+1}\right] \neq \psi_{\alpha,t}$.

We can now rewrite the system in matrix form as:

$$\begin{bmatrix} \varsigma & (1+\phi_\pi) \\ -\kappa & 1 \end{bmatrix} \begin{bmatrix} y_t \\ \pi_t \end{bmatrix} = \begin{bmatrix} \varsigma & 1 \\ 0 & \beta \end{bmatrix} \begin{bmatrix} \mathbb{E}_t\left[y_{t+1}\right] \\ \mathbb{E}_t\left[\pi_{t+1}\right] \end{bmatrix} + \begin{bmatrix} 1 & \phi_\pi \\ 0 & 0 \end{bmatrix} \begin{bmatrix} r_t^* \\ \pi_t^* \end{bmatrix},$$

or

$$\Phi_0 Y_t = \Phi_1 \mathbb{E}_t\left[Y_{t+1}\right] + \Theta\varepsilon_t \tag{19}$$

in compact form. We solve the system by assuming a linear solution, $Y_t = c + B\varepsilon_t$, such that

$$\mathbb{E}_t\left[Y_{t+1}\right] = c + B\,\mathbb{E}_t\left[\varepsilon_{t+1}\right].$$

We now substitute the solution in Equation (19)

$$Y_t = \Phi_0^{-1}\Phi_1\left[c + B\,\mathbb{E}_t\left[\varepsilon_{t+1}\right]\right] + \Phi_0^{-1}\Theta\varepsilon_t \tag{20}$$

which requires that $B = \Phi_0^{-1}\Theta$; rearranging, we get

$$c = \left(\Phi_0 - \Phi_1\right)^{-1}\Phi_1\Phi_0^{-1}\Theta\,\mathbb{E}_t\left[\varepsilon_{t+1}\right].$$

To solve for c we need the following quantities

$$\Phi_0^{-1} = \begin{bmatrix} \varsigma & (1+\phi_\pi) \\ -\kappa & 1 \end{bmatrix}^{-1} = \frac{1}{\varsigma + \kappa(1+\phi_\pi)} \begin{bmatrix} 1 & -(\phi_\pi+1) \\ \kappa & \varsigma \end{bmatrix},$$

$$\begin{aligned} B &= \Phi_0^{-1}\Theta \\ &= \frac{1}{\varsigma + \kappa(1+\phi_\pi)} \begin{bmatrix} 1 & -(\phi_\pi+1) \\ \kappa & \varsigma \end{bmatrix} \begin{bmatrix} 1 & \phi_\pi \\ 0 & 0 \end{bmatrix} \\ &= \frac{1}{\varsigma + \kappa(1+\phi_\pi)} \begin{bmatrix} 1 & \phi_\pi \\ \kappa & \kappa\phi_\pi \end{bmatrix} \\ &= \frac{1}{\varsigma + \kappa(1+\phi_\pi)} \begin{bmatrix} 1 \\ \kappa \end{bmatrix} \begin{bmatrix} 1 & \phi_\pi \end{bmatrix}, \end{aligned}$$

$$\begin{aligned} (\Phi_0 - \Phi_1)^{-1} &= \left(\begin{bmatrix} \varsigma & (1+\phi_\pi) \\ -\kappa & 1 \end{bmatrix} - \begin{bmatrix} \varsigma & 1 \\ 0 & \beta \end{bmatrix} \right)^{-1} \\ &= \begin{bmatrix} \frac{1}{\kappa\phi_\pi}(1-\beta) & -\frac{1}{\kappa} \\ \frac{1}{\phi_\pi} & 0 \end{bmatrix}. \end{aligned}$$

We can now compute

$$\begin{aligned} (\Phi_0 - \Phi_1)^{-1}\Phi_1 &= \begin{bmatrix} \frac{1}{\kappa\phi_\pi}(1-\beta) & -\frac{1}{\kappa} \\ \frac{1}{\phi_\pi} & 0 \end{bmatrix} \begin{bmatrix} \varsigma & 1 \\ 0 & \beta \end{bmatrix} \\ &= \frac{1}{\phi_\pi} \begin{bmatrix} \frac{\varsigma}{\kappa}(1-\beta) & -\frac{\phi_\pi}{\kappa}\beta - \frac{1}{\kappa}(\beta-1) \\ \varsigma & 1 \end{bmatrix}, \end{aligned}$$

such that the solution for c reads

$$\begin{aligned}
c &= \left(\Phi_0 - \Phi_1\right)^{-1} \Phi_1 \Phi_0^{-1} \Theta \, \mathbb{E}_t \left[\varepsilon_{t+1}\right] \\
&= \frac{1}{\phi_\pi} \frac{1}{\varsigma + \kappa \left(1 + \phi_\pi\right)} \begin{bmatrix} \frac{\varsigma}{\kappa}\left(1-\beta\right) & -\frac{\phi_\pi}{\kappa}\beta - \frac{1}{\kappa}\left(\beta - 1\right) \\ \varsigma & 1 \end{bmatrix} \begin{bmatrix} 1 \\ \kappa \end{bmatrix} \begin{bmatrix} 1 & \phi_\pi \end{bmatrix} \mathbb{E}_t \left[\varepsilon_{t+1}\right] \\
&= \frac{1}{\phi_\pi \left[\varsigma + \kappa \left(1 + \phi_\pi\right)\right]} \begin{bmatrix} \left(\frac{\varsigma}{\kappa} + 1\right)\left(1-\beta\right) - \phi_\pi \beta \\ \varsigma + \kappa \end{bmatrix} \begin{bmatrix} 1 & \phi_\pi \end{bmatrix} \mathbb{E}_t \left[\varepsilon_{t+1}\right].
\end{aligned}$$

We now isolate $\mathbb{E}_t \left[\pi_{t+1}\right]$ to obtain:

$$\begin{aligned}
\mathbb{E}_t \left[\varepsilon_{t+1}\right] &= \begin{bmatrix} 0 & 1 \end{bmatrix} \mathbb{E}_t \left[Y_{t+1}\right] \\
&= \begin{bmatrix} 0 & 1 \end{bmatrix} \left[c + B \, \mathbb{E}_t \left[\varepsilon_{t+1}\right]\right] \\
&= \begin{bmatrix} 0 & 1 \end{bmatrix} \left[\left(\Phi_0 - \Phi_1\right)^{-1} \Phi_1 \Phi_0^{-1} \Theta + \Phi_0^{-1} \Theta\right] \mathbb{E}_t \left[\varepsilon_{t+1}\right] \\
&= \frac{1}{\varsigma + \kappa \left(1 + \phi_\pi\right)} \begin{bmatrix} 0 & 1 \end{bmatrix} \begin{bmatrix} \left[\frac{\varsigma + \kappa}{\phi_\pi \kappa} + 1\right]\left(1 - \beta\right) \\ \frac{\varsigma + \kappa}{\phi_\pi} + \kappa \end{bmatrix} \begin{bmatrix} 1 & \phi_\pi \end{bmatrix} \mathbb{E}_t \left[\varepsilon_{t+1}\right].
\end{aligned}$$

Optimal policy requires that $\mathbb{E}_t \left[\pi_{t+1}\right] = 0$, which implies

$$\mathbb{E}_t \left[r^*_{t+1} - \phi_\pi \pi^*_{t+1}\right] = 0$$

and so we can obtain

$$\begin{aligned}
\mathbb{E}_t \left[\pi^*_{t+1}\right] &= -\frac{1}{\phi_\pi} \mathbb{E}_t \left[r^*_{t+1}\right] \\
&= -\frac{\varsigma \omega}{\phi_\pi} \left(\mathbb{E}_t \left[\psi_{\alpha, t+1}\right] - \psi_{\alpha, t}\right),
\end{aligned} \tag{21}$$

which implies $\mathbb{E}_t \left[\pi^*_{t+1}\right] \neq 0$ if $\mathbb{E}_t \left[\psi_{\alpha, t+1}\right] \neq \psi_{\alpha, t}$.

We can finally recover the implied distribution of inflation in closed form. Under the optimal

policy $\mathbb{E}_t\left[r^*_{t+1} - \phi_\pi \pi^*_{t+1}\right] = 0$ such that $c = 0$. Hence,

$$\begin{aligned} Y_t &= B\varepsilon_t \\ &= \frac{1}{\varsigma + \kappa\left(1+\phi_\pi\right)} \begin{bmatrix} 1 \\ \\ \kappa \end{bmatrix} \left(r^*_t + \phi_\pi \pi^*_t\right); \end{aligned}$$

assuming π^*_t is constant, it is easy to show that

$$\begin{aligned} x_t &\sim iidF\left(\frac{\mu^\star + \phi_\pi \pi^*_t}{\varsigma + \kappa\left(1+\phi_\pi\right)}, \frac{\sigma^\star}{\varsigma + \kappa\left(1+\phi_\pi\right)}, \varrho^\star\right), \\ \pi_t &\sim iidF\left(\frac{\kappa\left(\mu^\star + \phi_\pi \pi^*_t\right)}{\varsigma + \kappa\left(1+\phi_\pi\right)}, \frac{\kappa\sigma^\star}{\varsigma + \kappa\left(1+\phi_\pi\right)}, \varrho^\star\right). \end{aligned}$$