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

Andrea De Polis† Leonardo Melosi‡ Ivan Petrella§

This draft: March 2025

Abstract

We develop a tractable methodology to investigate the macroeconomic implications of evolving asymmetric risks within a dynamic general equilibrium model. Optimal monetary policy requires the central bank to lean against the balance of inflation risks to mitigate the resulting bias on inflation expectations. Using U.S. data, we show that inflation risks vary significantly over time and are typically asymmetric. Using a quantitative general equilibrium model, we evaluate the counterfactual dynamics of a monetary policy strategy in which the central bank dynamically responds to real-time estimates of the balance of inflation risks.

Keywords: Asymmetric risks, optimal monetary policy, balance of inflation risks, risk-adjusted inflation targeting, flexible average inflation targeting.

JEL codes: E52, E31, C53.

*We would like to thank Guido Ascari, Gianluca Benigno, Domenico Giannone, Emmanuel Moench, Roberto Motto, and the seminar participants at the Federal Reserve Bank of Chicago, the Federal Reserve Bank of Cleveland, the De Nederlandsche Bank, the University of Lausanne, the Workshop on the Economics of Risk Econometric Tools and Policy Implications at Collegio Carlo Alberto, Turin, the workshop on Monitoring and Forecasting Macroeconomic and Financial Risk at the National Central Bank of Belgium, the 2nd Annual Non-Linearities in Macro Workshop at the Bank of England, and the Workshop on Empirical Monetary Economics at OFCE, Paris. Leonardo Melosi is grateful to the European Central Bank for its generous hospitality and support during the summer of 2024, where a significant portion of this work was conducted. Any views expressed in this paper are those of the authors and do not necessarily reflect the views of De Nederlandsche Bank or any other person associated with the Eurosystem.

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

‡University of Warwick, European University Institute, DNB & CEPR. leonardo.melosi@eui.eu

§Collegio Carlo Alberto, University of Turin & CEPR. ivan.petrella@carloalberto.org

The pandemic and war have underscored the need for the risk management framework to take full account of both upside and downside risks to inflation, as well as to the possibility that serious tensions may arise between the objectives of price stability and employment or growth.

Gita Gopinath, Jackson Hole Symposium, August 26, 2022

1 Introduction

Recent global events, such as the COVID-19 pandemic, supply chain disruptions and rising geopolitical tensions, have significantly amplified macroeconomic risks and heightened the likelihood of future inflation spikes and dire conseuqences for growth. This new regime marks a sharp departure from the pre-pandemic era, when macroeconomic risks were largely contained, and inflation frequently surprised to the downside. Concerns that these shifts in the balance of risks could give rise to significant tensions between the objectives of price stability and full employment —as emphasized by Gita Gopinath, Deputy Managing Director of the International Monetary Fund— have been widely discussed among policy makers and academics. In this paper, we investigate how shifts in the balance of macroeconomic risks influence the economy and how the central bank should respond.

We introduce a tractable methodology for analyzing how time-varying, asymmetric macroeconomic risks influence outcomes and optimal monetary policy in dynamic general equilibrium models. The model features inefficient shocks whose moments evolve stochastically over time. Agents are rational and fully aware of the distribution from which current shocks are drawn. They receive news about the evolving skewness of the distribution, shaping their perception of macroeconomic risks. These evolving perceptions create time-varying asymmetries in agents' beliefs, which, in turn, affect equilibrium outcomes and introduce an additional trade-off for the central bank.

In this model, we prove that optimal monetary policy requires central banks to actively counteract *inflation risks* arising from the skewed distribution of shocks. This approach reflects the need to counter the deanchoring of inflation expectations due to the asymmetric distribution of markup shocks. When risks are skewed to the upside, inflation expectations may exceed the central bank's target. In such cases, the central bank should optimally tighten monetary policy more aggressively than in the symmetric scenario to tame inflation expectations, which are elevated by

the increased likelihood of positive markup shocks in the coming periods. Conversely, a negative balance of risks may lead to deflationary pressures, leading to the deanchoring of expectations in the opposite direction. In this scenario, the central bank should ease monetary policy to stimulate aggregate demand and raise inflation expectations, thereby countering the deflationary pressure caused by the higher probability of negative markup shocks in the future.

For this normative prediction to be relevant in practice, there must be sufficient evidence that shifts in the balance of inflation risks are a significant feature of the data and that changes in risk asymmetry can be reliably tracked in real-time. Using data on U.S. core Personal Consumption Expenditure (PCE) inflation, we find robust evidence supporting time-varying asymmetry in the predictive distribution of inflation. Our analysis identifies statistically significant and frequent shifts in the balance of inflation risks throughout the postwar period, often following persistent, regime-like patterns. Incorporating time-varying skewness into inflation forecasts significantly enhances out-of-sample accuracy compared to standard benchmark models (Stock and Watson, 2007), achieving predictive performance on par with the Survey of Professional Forecasters (SPF).

Lastly, we define and evaluate a novel monetary policy strategy, the Risk-Adjusted Inflation Targeting (RAIT), within an estimated general equilibrium model (Smets and Wouters, 2007). The RAIT calls for the central bank to dynamically adjust its policy stance and communication to counter shifts in the balance of inflation risks. In the model, these adjustments are designed to completely shelter inflation expectations from distortions caused by asymmetric inflation risks. As a result, the RAIT ensures, by construction, that inflation expectations remain stable and anchored at the central bank's target in every period.

We compare the counterfactual dynamics of key macroeconomic variables under the RAIT to those observed in the data. During the recent inflation surge, real-time estimates indicate that inflation risks were heavily tilted to the upside. Under these conditions, the RAIT would have recommended an earlier liftoff in interest rates, reaching a peak similar to that observed in the data but slightly sooner. By early 2023, the RAIT would have called for a faster unwinding of the prior monetary tightening. The model predicts that, under the RAIT, inflation would have declined more quickly, reaching the Federal Reserve's target by early 2023 and stabilizing thereafter. By

initiating monetary easing in early 2023, the RAIT would have also supported the labor market, partially offsetting the effects of earlier tightening.

We compare the properties of the RAIT to those of Flexible Average Inflation Targeting (FAIT), which was adopted by the Federal Reserve in 2020 to address deflationary risks, which were dominant in the previous decade (Clarida, 2022). While FAIT relies on past inflation deviations from the 2% inflation objective to determine the extent of overshooting, the RAIT relies on the central bank's dynamic assessment of the evolving balance of inflation risks. Consequently, the RAIT offers greater flexibility in adapting to changes in the stochastic properties of inflation. This key feature makes the RAIT more responsive than FAIT to sudden shifts in inflation skewness, such as those observed at the onset of the COVID-19 pandemic and again in early 2023.

The decade preceding the COVID-19 pandemic was marked by persistently subdued inflation, reflected in our model as a sustained shift toward negative inflation asymmetry. Despite this, the remarkable stability of the macroeconomic environment resulted in fairly contained inflation volatility. Against this backdrop, the RAIT would have recommended looking through the deflationary bias. This finding underscores RAIT's priority on addressing the effects of unbalanced inflation risks on expectations. These effects are muted when inflation volatility is low. In contrast, the FAIT would have required the Federal Reserve to promise overshooting the target to compensate for past inflation misses.

Following Clarida et al. (1999); Woodford (2003); Galí (2008), we solve the optimal monetary policy problem in a linear-quadratic framework where only the first-order effects of time-varying asymmetric macroeconomic risks are relevant. While symmetric changes in risks are well-known to have no first-order effects on equilibrium outcomes, asymmetric shifts in risks do, as they influence agents' expectations. Positive (negative) skewness increases (decreases) expectations by pulling the mean toward the right (left) tail. Consequently, changes in the skewness of the shocks distribution affect agents' decisions at first order through their impact on expectations. In fact, expectations are the only moment of agents' beliefs that matters for optimal decisions in a linear model. Therefore, first-order effects of asymmetric risks can be studied through the log-linear approximation of the model around its steady-state equilibrium.

To conduct this analysis, we augment the model with zero-mean Gaussian anticipated markup shocks that serve to tilt agents' expectations of future shocks, consistent with the stochastic changes to the skewness of their distribution. By design, these dummy anticipated shocks do not necessarily materialize, resembling noise, sentiments, or shifts in pure beliefs. Hence, we call this framework the *beliefs representation of asymmetric risks*.

While focusing on the first-order effects of asymmetric macroeconomic risks limits the number of channels through which risks can influence equilibrium allocations and prices, we highlight two main advantages. First, the linear-quadratic beliefs representation allows the characterization of the solution to the optimal policy problem analytically and within a framework comparable to that employed by the literature on optimal monetary policy under symmetric shocks (Clarida et al., 1999; Woodford, 2003; Galí, 2008).

A second advantage is that the laws of motion in the beliefs representation are linear and all shocks – i.e. the structural ones and the dummy anticipated ones – are Gaussian. Therefore, this representation of the model can be solved with extremely fast and accurate off-the-shelf techniques. This tractability greatly facilitates the analysis of real-time changes in the balance of inflation risks and the study of policy counterfactuals within a quantitative dynamic general equilibrium framework – e.g. models like those introduced by Christiano, Eichenbaum, and Evans (2005); Smets and Wouters (2007).

Finally, it is important to emphasize that the proposed framework is generic and can be integrated with various methods for estimating skewness, not just the one adopted in this paper.

Literature Review Our work contributes to the relatively sparse literature on risk management approaches in monetary policy. Dolado et al. (2004) and Surico (2007) demonstrate that asymmetric loss functions can lead to optimal policy rules incorporating nonlinear terms for inflation and the output gap, reflecting policymakers' asymmetric preferences. Kilian and Manganelli (2008) highlight the importance of balancing upside and downside risks under an asymmetric loss function, while Kilian and Manganelli (2007) measure inflation risks under time-varying symmetric volatility. In these frameworks, central banks may target an average inflation rate different from their stated target when the costs of undershooting and overshooting are asymmetric. This bias

disappears with a symmetric loss function. In this paper, we always assume that the loss function is symmetric. Our key contribution is to show that when inflation risk is asymmetric, the optimal response must account for the balance of risks, even under a symmetric target. Asymmetric inflation risk creates a divergence between the modal forecast and the expected inflation path, and failure to address this divergence results in systematic deviations from the target on average.

Evans et al. (2020) prove that the ZLB requires a looser monetary policy under uncertainty, leading to an optimal delay in policy liftoff. Moreover, they provide both narrative and statistical evidence that the Federal Reserve has frequently employed risk management when setting the policy rate. Bianchi et al. (2021) show that the deflationary bias caused by the risk of recurrently hitting the ZLB constraint can be mitigated if the central bank adopts an asymmetric monetary strategy. Such a strategy entails responding more strongly to inflation deviations below target than to those above target.

Unlike these papers, our work examines the importance of the central bank responding to changes in the balance of macroeconomic risks beyond those exclusively associated with the ZLB constraint. Furthermore, none of these studies estimate time-varying asymmetric risks in the data using a new forecasting model or demonstrate how such estimates can be integrated with a quantitative structural model to explicitly analyze the implications of monetary policy strategies that target inflation risks.

This paper examines the empirical implications of inflation-targeting measures and is therefore connected to the longstanding literature on optimal inflation targeting (Giannoni and Woodford, 2004). By ensuring that the central bank's mean inflation forecasts align with its policy objectives, RAIT falls within the broader category of strategies studied by Svensson (1997), commonly known as forecast inflation targeting. A key distinction from the existing literature lies in our framing of the optimal monetary policy problem within a modeling environment where the balance of macroeconomic risks evolves stochastically.[1] Another feature that sets our approach apart from the academic literature on forecast inflation targeting is that we do not restrict the central bank to optimally respond only to current and lagged predetermined variables. Instead, in our framework,

[1] Svensson (2003) discusses how in practice central banks should implement its reaction function when macroeconomic risks are unbalanced, reviewing and critiquing certain practices—such as the Bank of England's decision to display the mode rather than the mean in fan charts for output and inflation.

the central bank can respond to endogenous, forward-looking variables, explicitly accounting for the two-way feedback between monetary policy reaction functions and rational agents' beliefs.

Our paper is also related to the literature on inflation forecasting. This literature has largely focused on the importance of accounting for slow-moving trends in the data and time-varying uncertainty (see, e.g., Cogley and Sargent, 2005; Stock and Watson, 2007; Faust and Wright, 2013; Ascari and Sbordone, 2014). Far less attention has been devoted to the risks of inflation or deflation. Some recent exceptions include Andrade et al. (2014); Hilscher et al. (2022). Differently from our approach, that only requires a single time series to extract predictive inflation distributions, Andrade et al. (2014) rely on survey data to extract measures of time-varying perceived inflation asymmetry and show that this has predictive power over future inflation. Hilscher et al. (2022) introduces a methodology to extract market-based tail probabilities from options data. Their approach relies on the availability and involved manipulations of inflation derivatives data. However, we find an high correlation between the predicted disaster probabilities with the same measures extracted in real-time from our model.[2]

Only recently the attention has moved to the modeling of the whole density of inflation outcomes (Manzan and Zerom, 2013, 2015; Korobilis et al., 2021; López-Salido and Loria, 2024). Most of the previous work rely on quantile regression approaches, popularized by Adrian et al. (2019) in the context of quantify the effect of macro-financial risks on the distribution of future growth. In contrast, we 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 et al. (2024). 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; Blanchard and Bernanke, 2023).

Two recent papers are directly related to our empirical contribution. First, Le Bihan et al.

[2]Disaster probabilities are defined as the probability of inflation being above 4% and 5% over the next ten years.

(2024) introduces a new real-time measure of underlying inflation that incorporates 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). Second, López-Salido and Loria (2024) document substantial variability in the tails of inflation, and relates their dynamics to deteriorating financial conditions and macroeconomic factors. Both papers employ sample periods that are shorter than ours and do not cover the earlier episode of persistently elevated inflation in the 1970s. Moreover, whereas López-Salido and Loria (2024) related their findings about tail dynamics to the model of Gertler et al. (2020), none of the two contributions explore the implications of asymmetric inflation risks for optimal monetary policy through the lens of DSGE models.

Structure The remainder of the paper is organized as follows. In Section 3, we characterize the optimal monetary policy within a New-Keynesian model featuring asymmetric macroeconomic risks. To solve this problem, we show how to obtain the beliefs representation of a model with asymmetric risks. We then estimate the skewness of US inflation in real-time using an econometric model Section 4. In Section 5 we calibrate the beliefs representation of a quantitative DSGE model to match the estimates of the balance of inflation risks coming for the econometric model and conduct a number of counterfactual exercises. In Section 6, we conclude.

2 Time-varying asymmetric risks in general equilibrium

In this section, we lay out key concepts used throughout the paper. We first examine how asymmetric shocks affect equilibrium outcomes in a general equilibrium setting. Then, we show how shifts in the asymmetry of future macroeconomic risks alter these outcomes through first-order effects. Finally, we demonstrate that these changes can be understood as revisions in agents' beliefs about the future distribution of shocks. Proofs of the following propositions are reported in Appendix A.

Let $\mathbf{Z}_t$ collects endogenous and (persistent) exogenous variables and $\mathbf{S}_t$ be a vector of iid exogenous shocks to the system. A competitive equilibrium can be defined as a system of expectational

difference equations of the form

$$E_t \left[f \left(\mathbf{Z}_{t+1}, \mathbf{Z}_t, \mathbf{Z}_{t-1}, \mathbf{S}_t \right) \right] = 0, \tag{1}$$

where f depends on agents' preferences, available technology and constraints, and E_t denotes the expectation operator conditional on the information set available at time t.

Let us now assume $\mathbf{S}_t = \left[\mathbf{S}_t^S \; S_t^a \right]'$, where $\mathbf{S}_t^s$ is a set of symmetric shocks, e.g., Normally distributed, whereas S_t^a represents a potentially *asymmetric* shock with distribution $\mathcal{F}_t$, where we assume, without loss of generality, the mode to be zero and the asymmetry to vary over time.[3] Agents information set includes the state of the economy up to period $t-1$, the history of the shocks, and past and current news about the asymmetry of the skewed shock distributions.

Proposition 1. *The first-order approximation of any model of the form* (1) *is*

$$\mathbf{A}_0 \mathbf{z}_t = \mathbf{A}_f E_t \mathbf{z}_{t+1} + \mathbf{A}_b \mathbf{z}_{t-1} + \mathbf{B}_s \mathbf{s}_t^s + b_a s_t^a, \tag{2}$$

independently of the distribution of the shocks, $\mathbf{s}_t$.

Proposition 1 highlights that only expectations matter for the first-order solution of the model, where lower case letters define log-deviations from steady-states. It should be noted that, when shocks are asymmetric, changes in the risk of the distribution induced by changes in the level of asymmetry by affecting the expected value of the shock, remain relevant even at the first-order.

The skewness –or asymmetry– of the distribution $\mathcal{F}_t$ is defined by the ratio of the mass to the right and to the left of the mode, $r_t = \frac{\int_0^\infty \mathcal{F}_t dx}{\int_{-\infty}^0 \mathcal{F}_t dx}$. Risks are skewed to the downside (upside) when $0 < r_t < 1$ ($r_t > 1$), while $r_t = 1$ indicates a symmetric distribution.[4]

Agents learn about changes in the asymmetry of the shock distribution by receiving news about $\varrho_t \in (-1, 1)$, a parameter that pins down the asymmetry of the distribution $\mathcal{F}_t$, defined as

[3] We focus on asymmetric risk in a unimodal setting. However, the framework can also accommodate multimodal risks. In such cases, skewness arises as a byproduct of multimodality and, up to *first-order*, its impact remains limited to its effect on expected realizations, with the sign captured by the skew in risk. Higher-order effects may be significant when the skewness is very large but these cases are beyond the scope of our analysis.

[4] This definition restricts the class of distributions we consider to the family of *two-piece* distribution, originally introduced by Fechner (1897). Nevertheless, this class of distribution is quite flexible, and encompasses most of the well-known univariate skewed distributions (see, e.g., Fernández and Steel, 1998; Arellano-Valle et al., 2005).

$\varrho_t = \frac{r_t - 1}{r_t + 1}$, such that for $\varrho_t > 0$ ($\varrho_t < 0$) the distribution is positively (negatively) skewed; $\varrho_t = 0$ implies symmetry (Arellano-Valle et al., 2005).[5]

An immediate consequence of asymmetry is that the mean and the mode of the distribution diverge. We can now characterize this wedge as a function of the asymmetry parameter, such that $E_t s^a_{t+j} = \kappa\sigma\varrho_t$. Hence, revisions about the expectation of future shocks' asymmetry induce a response of the endogenous variables that is proportional to the update in the degree of asymmetry of the distributions. Proposition 2 formalizes this concept and provides the exact representation of the response.

Proposition 2. *Updates in the expectations of the asymmetry parameter lead to revisions in the expectations about future shocks, i.e.* $E_t s^a_{t+j} - E_{t-1} s^a_{t+j} = k\sigma^2 \varrho^j_t$, *and thus the following impulse response function for the endogenous variables* $h-$*period ahead:*

$$E_t \mathbf{z}_{t+h} - E_{t-1}\mathbf{z}_{t+h} = \mathbf{M}_{(j,h)} k\sigma^2 \varrho^j_t, \; h \geq 0, \tag{3}$$

where $\varrho^j_t = E_t\varrho_{t+j} - E_{t-1}\varrho_{t+j}$ *denotes the revisions in the asymmetry of the shock distribution, and* $\mathbf{M}_{(j,h)}$ *is a convolution of the system matrices.*

Therefor, agents re-optimize their behavior when they receive news about changes in the asymmetry of future shocks, thereby affecting equilibrium allocations.

2.1 Belief representation

Let us now consider an alternative model with linear solution given by

$$\mathbf{A}_0\mathbf{z}_t = \mathbf{A}_f \widetilde{E}_t \mathbf{z}_{t+1} + \mathbf{A}_b \mathbf{z}_{t-1} + \mathbf{B}\widetilde{\mathbf{s}}_t, \tag{4}$$

where $\widetilde{E}_t$ is the *belief expectation* operator, such that

$$\widetilde{E}_t \widetilde{\mathbf{s}}_{t+j} = E_t \widetilde{\mathbf{s}}_{t+j} + \mathbf{c} E_t \psi_{t+j}, \tag{5}$$

[5]The asymmetry parameter, ϱ_t, is a convenient reparametrization of r_t, which is centered and symmetric around zero.

where $E_t\psi_{t+j}$ represent agents' beliefs about the mean of the distribution of some of the shocks hitting the economy a time $t+j$, with $E\psi_t = 0$, and $\mathbf{c}$ is a selection vector. We assume that the risk of the shocks $\widetilde{\mathbf{s}}_t$ is symmetric, i.e., $E_t\widetilde{\mathbf{s}}_{t+1} = \mathbf{0}, \forall t$. Proposition 3 establishes the first-order equivalence between the effect of asymmetric shocks and that of belief shocks.

Proposition 3. *Revisions in the asymmetry of future shocks' distribution can be equivalently represented as revisions in beliefs about the mean of future shocks,*

$$\mathbf{c}\psi_t^j \equiv \mathbf{\Xi} k\sigma^2 \varrho_t^j, \tag{6}$$

where $\psi_t^j = E_t\psi_{t+j} - E_{t-1}\psi_{t+j}$ denotes the revisions in beliefs about the mean of the shocks.

3 Optimal monetary policy with asymmetric risks

In this section, we analyze the implications of time-varying asymmetric macroeconomic risk for optimal monetary policy in a standard New Keynesian model with sticky prices (Galí, 2008, chapter 3). The model is shown in Appendix B. We begin by showing the benchmark (well-known) case of optimal monetary policy in a linear-quadratic model with symmetric macroeconomic risk. Then, we consider the case when risks are asymmetric within the same linear-quadratic framework. We derive the beliefs representation of the asymmetric risks, and then we characterize the optimal policy analytically. Finally, we consider the case where risks are asymmetric but the central bank looks through this feature when deciding its optimal policy. This last case offers valuable insights into the optimal monetary strategy under asymmetric macroeconomic risks, which will be further explored through numerical simulations at the end of the section.

3.1 The case of symmetric risks

Let us assume that the central bank can fully commit, with complete credibility, to a policy plan by selecting a state-contingent sequence of inflation deviations from its target and output gaps,

$\{\hat{\pi}_t, \hat{x}_t\}_{t=0}^{\infty}$, so as to maximize the quadratic approximation of the household's utility function:

$$\frac{1}{2} E_0 \sum_{t=0}^{\infty} \beta^t \left(\hat{\pi}_t^2 + \alpha_x \hat{x}_t^2 \right)$$

subject to the sequence of constraints given by the Phillips curve

$$\hat{\pi}_t = \beta E_t \pi_{t+1} + \kappa \hat{x}_t + u_t,$$

where price markup shocks follow an autoregressive process:

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u, \tag{7}$$

driven by i.i.d. shocks governed by a symmetric time-invariant distribution $\mathcal{F}_u$. The objective function is the quadratic approximation of the household's standard utility function – shown in Appendix B – and the weight on the output gap is $\alpha_x = \kappa/\varepsilon$, where ε denotes the elasticity of substitution between differentiated goods produced by monopolistically competitive firms and κ denotes the slope of the Phillips curve. All variables are expressed in log-deviations from their steady-state value.

It is easy to show that the following optimality condition must hold:

$$\hat{x}_t = -\frac{\kappa}{\alpha_x} \bar{p}_t, \tag{8}$$

where $\bar{p}_t = p_t - p_{-1}$ is the cumulative inflation rate over period 0 through period t, with p_t denoting the log of the price level at time t, and p_{-1} is the implicit target given by the price level prevailing one period before the optimal plan is chosen by the central bank. As it is well known, this optimality condition can be interpreted as a targeting rule that the central bank is required to follow in every period in order to implement the optimal policy. Under the optimal policy, the central bank sets the sign and the size of the output gap in proportion to the deviations of the price level from its implicit target (Galí, 2008, chapter 5).

Following Galí (2008), we show that the price level under the optimal policy is: $\bar{p}_t = \eta \bar{p}_{t-1} + \lambda u_t$

and the optimal monetary rule (under the assumption that $\rho_u = 0$) follows:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t, \tag{9}$$

where σ denotes the household's coefficient of risk aversion and the coefficients η and λ are function of other structural parameters of the New Keynesian model as shown in Appendix B. The appendix also includes detailed derivations of the equilibrium equations.

3.2 The case of asymmetric risks

Let us now introduce asymmetric macroeconomic risks. Specifically, we assume that the stochastic process driving markups u_t is the same process of the symmetric case in Equation (7), except that the iid shocks ε_t^u are now drawn from a two-pieces, zero-mode probability density function, $\varepsilon_t^u \sim \mathcal{F}_{u,t}$, whose skewness evolves stochastically over time.

The skewness of the distribution $\mathcal{F}_{u,t}$ is defined by the ratio $r_{u,t} = \frac{\int_0^\infty \mathcal{F}_{u,t} d\varepsilon_t^u}{\int_{-\infty}^0 \mathcal{F}_{u,t} d\varepsilon_t^u}$, where the numerator (denominator) is the probability that the realization of the shock ε_t^u is larger (smaller) than the mode, which, without loss of generality, we set equal to 0.[6] Risks are skewed to the downside (upside) when $0 < r_{u,t} < 1$ ($r_{u,t} > 1$), while $r_{u,t} = 1$ indicates a symmetric distribution. We define the asymmetry parameter as $\varrho_{t+j} \equiv \left(\frac{r_{u,t+j}-1}{r_{u,t+j}+1}\right)$, whose values range from -1 and 1.

In period t, agents receive a sequence of news, ϱ_t^j, leading them to revise their beliefs about the asymmetry parameter, ϱ_{t+j}, of the future distributions of markup shocks, $\mathcal{F}_{u,t+j}$ for $j = \{1, \ldots, J\}$. Formally, we define the news as follows: $\varrho_t^j \equiv E_t\varrho_{t+j} - E_{t-1}\varrho_{t+j}$. We assume that news shocks are drawn from a distribution with zero mean.[7] Zero mean implies that agents do not expect to revise their expectations regarding the asymmetry parameter ϱ_t^j next period.

[6] Assuming a mode of zero is consistent with the standard assumption that innovations to shocks have an expected value of zero when risk is symmetric. While we focus on asymmetric risk in a unimodal setting, the framework can also accommodate multimodal risks. In such cases, skewness arises as a byproduct of multimodality and, up to *first-order*, its impact remains limited to its effect on expected realizations, with the sign captured by the skew in risk. Higher-order effects may be significant when the skewness is very large but these cases are beyond the scope of our analysis.

[7] Note that the probability density function (pdf) of the news ϱ_t^j is defined between $-1 - E_{t-1}\varrho_{t+j}$ and $1 - E_{t-1}\varrho_{t+j}$. To have zero mean the pdf must be asymmetric, putting more weight on negative (positive) realizations if $E_{t-1}\varrho_{t-1} > 0$ ($E_{t-1}\varrho_{t-1} < 0$). Therefore, the pdf of the news ϱ_t^j is state dependent as its support depends on past news—specifically $E_{t-1}\varrho_{t+j} = \sum_{h=j}^{J} \varrho_{t-1+j-h}^{h+1}$.

Moreover, we define the forecast error of the asymmetry parameter realized at time t as $\varrho_t^0 \equiv \varrho_t - E_{t-1}\varrho_t$. It then follows that asymmetry parameter in each period can be expressed as a combination of past news and the current surprise, such that $\varrho_t = \sum_{j=0}^{J} \varrho_{t-j}^j$. Eventually, rational agents learn the shape of the distribution of current shocks, but they keep revising their beliefs about the distributions from which future shocks will be drawn in every period.

Given the assumption that news and forecast errors have zero mean, the unconditional expectation of the skewness of the distribution of shocks is zero; i.e. $E(\varrho_t) = 0$. Thus, agents expect the distribution of shocks to be symmetric in steady state and beyond the longest horizon, J, at which they receive news.

News-driven revisions to the asymmetry parameter prompt agents to update their expectations about future shock realizations, influencing their decisions in each period. Consequently, positive (negative) asymmetry increases (decreases) expectations by shifting the mean toward the right (left) tail. Formally, the expected realization of the markup shock j periods ahead is defined as (see, e.g., Arellano-Valle et al., 2005; Azzalini and Capitanio, 2014).

$$E_t(\varepsilon_{t+j}^u) = k\sigma_u E_t(\varrho_{t+j}),\ \forall j \in \{1, \ldots, J\}, \tag{10}$$

where k is a constant, σ_u represents the scale of the distribution of the shocks $\mathcal{F}_{u,t+j}$—capturing the dispersion of the distribution, which is time invariant, and agents' expectations of the asymmetry parameter, which depend on the news received up to period t; that is, $E_t\varrho_{t+j} = \sum_{h=j}^{J} \varrho_{t+j-h}^h$.

The Beliefs Representation. We now show this model with asymmetric risks admits a *beliefs representation*. This representation predicts the same equilibrium outcomes but features only symmetric, zero-mean shocks. One advantage of this representation is that it enables the analytical characterization of the optimal monetary policy strategy in the presence of stochastic and asymmetric macroeconomic risks.

The *beliefs representation* of the model preserves the linearized equilibrium laws of motion governing the dynamics of the model with asymmetric risks (henceforth, the *actual model*.) The main point of departure from the actual model is the presence of *dummy surprise and anticipated*

markup shocks.

The set of dummy surprise and anticipated markup shocks, φ_t^j for $j \in \{0, 1, \ldots, J\}$, such that price markup shocks in the beliefs representation are defined as follows:

$$\varepsilon_t^u = \sum_{j=0}^{J} \varphi_{t-j}^j. \tag{11}$$

The dummy anticipated shocks provide the beliefs representation with a precious tool to account for the effects of news on the asymmetric parameters on expectations at horizons 1 through J periods ahead in the actual model. Recall that these effects are defined in equation (10).

To replicate these effects in the beliefs representation of the model, we impose that the dummy anticipated shocks must satisfy the following condition:

$$\varphi_t^j = E_t \varepsilon_{t+j}^u - E_{t-1} \varepsilon_{t+j}^u = k \sigma_u \varrho_t^j, \tag{12}$$

for every period t and for every horizon $j \in \{1, \ldots, J\}$. Note that the first equality stems from the definition of the dummy anticipated shocks in equation (11) and the second equality is derived by taking the first difference of equation (10), recalling that $E_t \varrho_{t+j} = \sum_{h=j}^{J} \varrho_{t+j-h}^h$.

In the actual economy, news about future asymmetry parameters, ϱ_t^j, change the agents' beliefs about the wedge between the mean and the zero mode of future shock's distributions at every horizon $j = \{1, ..., J\}$. In the beliefs representation, the changes in these wedges are replicated through the dummy anticipated shocks, φ_t^j, which play the role of a location shifter for the distribution of markup shocks.

Finally, the dummy surprise shock, φ_t^0, is defined in every period to ensure that the realization of the price markup shocks in the actual model and its beliefs representation are identical in every period. This can be achieved by imposing the condition:

$$\varphi_t^0 = \varepsilon_t^u - \sum_{j=1}^{J} \varphi_{t-j}^j, \tag{13}$$

for every period t, where ε_t^u denotes the price markup shock in period t in the actual model. This

restriction uniquely determines the realization of the dummy surprise shock, φ_t^0, which combines the actual realization of the shock, ε_t^u, with the past dummy anticipated shocks, φ_{t-j}^j, for $j \in \{1, \ldots, J\}$. As explained, these anticipated shocks capture agents' evolving beliefs about the asymmetry of the distribution of future markup shocks, ϱ_{t+j}, for $j \in \{1, 2, \ldots, J\}$.

Equation (13) makes it clear that since the anticipated dummy shocks capture the effects of evolving asymmetric risks on agents' expectations, they do not necessarily materialize in realized shocks; rather, they play the role of noise, sentiments, or shifts in pure beliefs. To see this, consider the case in which the actual realization of the markup shock in period t is zero ($\varepsilon_t^u = 0$). Then the surprise shock, φ_t^0, will be solely pinned down by the need to offset all past news, implying that all the past news agents have received turn out to be just noise. Formally, $\varphi_t^0 = -\sum_{j=1}^{J} \varphi_{t-j}^j$.

To sum up, the first set of restrictions on the dummy anticipated shocks implies that agents in this representation share the same expectations of future realizations of the asymmetric shock in the actual economy, $E_t \varepsilon_{t+j}^u$, for $j \in \{1, \ldots, J\}$. These expectations are nonzero because agents recognize that the distributions of future shocks are skewed, with skewness evolving stochastically.

The second restriction ensures that the realizations of price markup shocks in the beliefs representation are identical to those in the actual economy in every period. Consequently, the beliefs representation of a model with stochastic skewness yields equilibrium dynamics identical to those implied by the first-order approximation of the original model.

To see this equivalence of the representation, recall that in a linear model, such as our economy with stochastic skewness, expectations are the only moments of beliefs that matter for agents' optimal decisions. Therefore, anticipated changes in skewness influence optimal decisions solely through their effects on expectations. The beliefs representation replicates these effects exactly through the dummy shocks.

Since all shocks in the beliefs representation are symmetric, this framework is highly tractable, allowing for an analytical characterization of the optimal monetary policy strategy under stochastic changes in the skewness of the shock distribution.

Optimal monetary policy under time-varying asymmetric macroeconomic risks. To simplify the characterization of the optimal monetary policy, we assume that agents receive news

about the skewness of the distribution only up to $J = 1$ periods ahead.[8] Under this assumption, we write the autoregressive process for the price markup shocks in the beliefs representation of the model as follows: $u_t = \rho_u u_{t-1} + \left(\varphi_t^0 + \varphi_{t-1}^1\right)$ where φ_t^j, for $j \in \{0, 1\}$, are Gaussian shocks known in period t that are expected to change price markup in period $t + j$. The actual price markup shock is unanticipated and normally distributed; that is, $\varepsilon_t^u \sim \mathcal{N}(0, \sigma_{u,t}^2)$. As in the symmetric case, agents know the future evolution of the standard deviation of this distribution.

In the beliefs representation, the optimal policy and the equilibrium dynamics of the price level and the output gap under this policy can be characterized analytically and read as follows:

$$\bar{p}_t = \eta \bar{p}_{t-1} + \lambda u_t + \zeta \varphi_t^1,$$
$$\hat{x}_t = \eta \hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\left[\lambda u_t + \zeta \varphi_t^1\right].$$

Detailed derivations are provided in Appendix B. Note that, in this linearized model, macroeconomic risks introduce a wedge in agents' expectations relative to the modal scenario (i.e., agents' expectations in the symmetric case). Consequently, the equilibrium dynamics of prices and the output gap are shifted upward or downward, depending on the direction of the risks (φ_t^1).

The optimal interest rate rule in the beliefs representation (assuming $\rho_u = 0$) is:[9]

$$\hat{i}_t = -(1 - \eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\varphi_t^1, \tag{14}$$

where the first term on the right-hand side represents the optimal rule under symmetry, and the second captures the policy rate adjustment required to offset the effects of the balance of risks on agents' expectations, φ_t^1. This rule suggests that the optimal policy requires the central bank to respond to the effects of the changing balance of macroeconomic risks on agents' expectations.

[8] This assumption will be relaxed in the quantitative sections of the paper. Generalizing the optimal monetary response to the case where $J > 1$ is straightforward but adds little to the core intuition behind the optimality of leaning against inflation risks.

[9] The derivation of the interest-rate rule satisfying the determinacy of one stable rational expectations equilibrium is also shown in Appendix B.

3.3 The case of asymmetric risks with an unwitting central bank

Assume now that the central bank chooses its optimal policy without internalizing that macroeconomic risks are asymmetric, whereas agents are aware of it. This case makes an useful counterfactual scenario to draw intuition regarding the key features of the optimal policy under asymmetric shocks.

When the central bank ignores that macroeconomic risks can be unbalanced, it will aim to achieve the price level that would arise if the distribution of shocks were symmetric (conditional on the previous period's price level). Specifically, the central bank chooses the output gap $\hat{x}_t$ such that

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^s, \tag{15}$$

where $\bar{p}_t^s = p_t^s - p_{-1}$ and p_t^s denotes the equilibrium price level the central wants to achieve based on its beliefs that shocks are symmetrically distributed. Specifically, this "wrong" price level is defined as: $\bar{p}_t^s = \eta\bar{p}_{t-1} + \lambda u_t$. Note that this price level is not exactly the same as that in the fully symmetric case as the central bank conditions its policy action today on the equilibrium price level from the previous period, $\bar{p}_{t-1}$. This assumption implies that "bygones are bygones," as the central bank does not attempt to correct past mistakes – i.e. $\bar{p}_{t-1} - \bar{p}_{t-1}^s$ in its current policy actions.[10]

The central bank – unaware of macroeconomic risks – targets the deviations of the wrong price level ($\bar{p}_t^s \neq \bar{p}_t$) from the implicit target, p_{-1}. Comparing the above optimality condition with Equation (8), which holds also in the case of asymmetric risk studied in the previous section, reveals that the central bank sets the sign and the size of the output gap in proportion to the deviations of the price level ($\bar{p}_t^s$) that is not achievable because of the presence of asymmetric macroeconomic risks. As a result, the output gap will be suboptimal and under an equilibrium price level that is different from that targeted by the central bank. The equilibrium price level in the case the central bank overlooks the presence of asymmetric macroeconomic risk is shown in the Appendix B.

[10]It is important to point out that this case is not intended to be realistic – at some point, the central bank realizes that the targeted price level cannot be achieved. This case is introduced to highlight key insights regarding the implications of optimal policy under asymmetric macroeconomic risks.

3.4 A illustrative example

We now present a simulation to illustrate how an optimizing central bank should respond to changes in the balance of macroeconomic risks. The simulation assumes that the distribution of price markup shocks in period 2 is negatively skewed, represented by a dummy anticipated shock $\varphi_t^1 = -0.8$ at $t = 1$. Agents are assumed to be fully aware of this asymmetry starting in period 1, when the simulation begins. In period 2, a negative price markup shock is drawn from the negatively skewed distribution; that is, $\varepsilon_t^u = -0.8$ with $t = 2$. In all other periods, shocks are zero, and their distribution is symmetric. The economy is assumed to be at the non-risky steady state in period 0. Furthermore, the markup shocks are modeled as independently and identically distributed (IID), such that $\rho_u = 0$.

Results of this simulation are shown in Figure 1. We set the values of the model parameters following Galí (2008), Chapter 3, which we report in Appendix B. The symmetric risks case is represented by the black solid line – where we set the skewness to zero or $\varphi_t^1 = 0$ in any period t. In the asymmetric risks case, where the central bank is aware of the balance of risks (blue dashed line with circle markers), the distribution of markup shocks exhibits negative skewness.

The period-by-period effects of skewness on agents' expectations about the next period's markup shocks are illustrated by the green solid line with square markers in the bottom-right panel. This line shows the dynamics of the anticipated shock, φ_t^1, in the beliefs representation. Since the distribution of shocks is negatively skewed only in period 2, the effects on expectations arise exclusively in period 1, as indicated by the spike in the green line during the first period. The red dashed-dotted line represents the case in which risks are asymmetric, but the central bank does not take this into account (the case of the unwitting central bank.)

To understand the implications of optimal policy in the case of asymmetric risks, one should compare the blue dashed line with circle markers to the red dashed-dotted line. The difference between these two lines captures the effects of the optimal response of monetary policy to asymmetric risks on the macroeconomic variables.

In the first period, the unwitting central bank chooses the same output as in the symmetric case (the black solid line) because it aims to achieve the price level that would arise under symmetry.

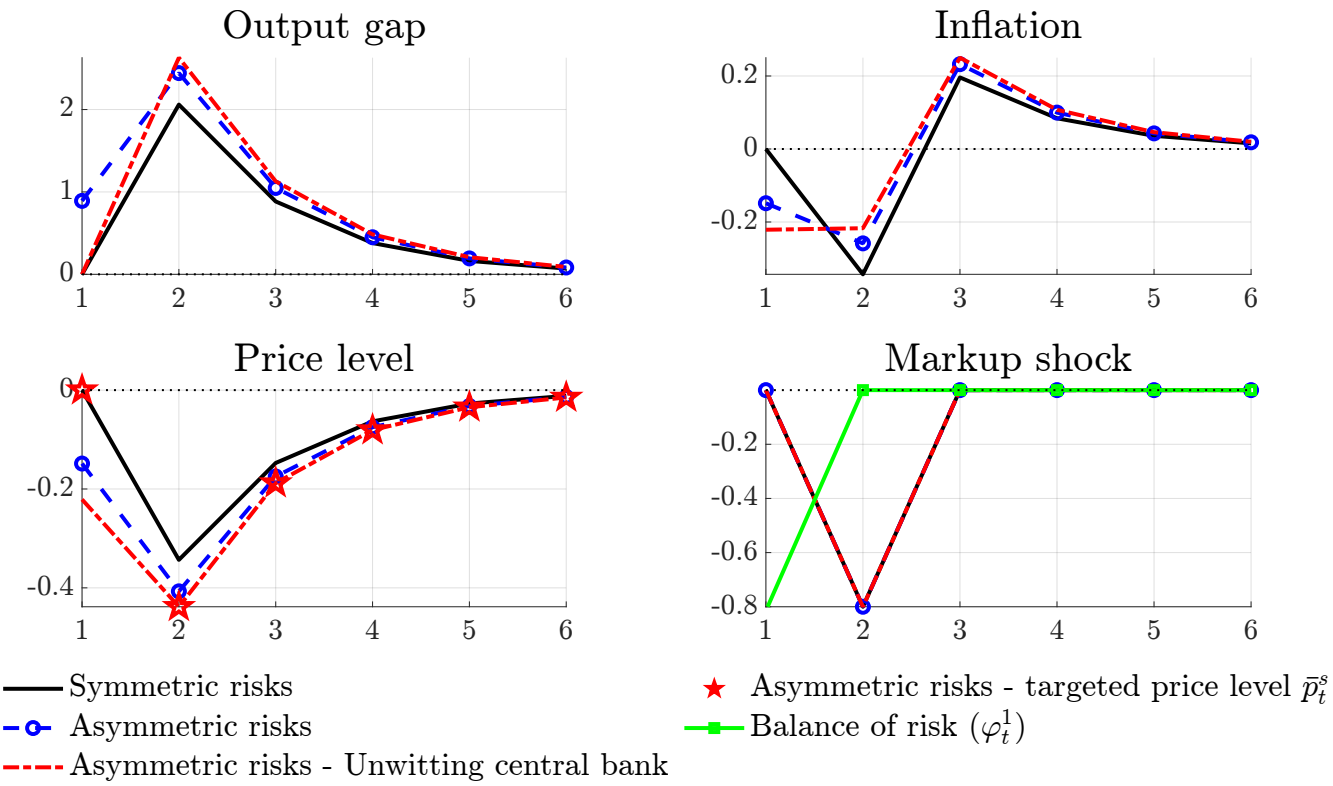


Figure 1: Optimal monetary policy under symmetric and asymmetric macroeconomic risks.

Note: The panels report the macroeconomic effects of a markup shocks drawn from a symmetric distribution (the black solid line) and from an asymmetric distribution (the blue dotted line with circle markers) under optimal monetary policy. The shock hits the economy in period 2, which is the only period when its distribution is assumed to be negatively skewed. The red dashed-dotted line denotes the counterfactual case where risks are asymmetric but the central bank does not take it into account when solving the optimal monetary policy problem. In the lower right panel, the price markup shock, u_t is shown as well as the skewness (φ_t^1) – green line with square markers. In the lower left panel, the line with red stars denotes the price targeted by the central bank that overlooks the unbalance of risks, $\bar{p}_t^s$.

However, agents' expectations are distorted by the negative skewness in the distribution of markup shocks for the next period. As a result, the desired price level turns out to be unattainable, with the equilibrium price level being significantly lower. This is reflected in the red dashed-dotted line falling below the black solid line at time 1 in the bottom-left panel.[11]

When the central bank accounts for negative inflation risks (the blue dashed line with circle markers), it optimally seeks to overheat the economy to counteract these deflationary risks. Consequently, the optimal response to negatively skewed markup shocks is to boost the output gap, $\hat{x}_t$, and the price level, $\hat{p}_t$. Compare the blue dashed line with circle markers (optimal policy) to the red dashed-dotted line (suboptimal policy failing to account for deflationary risks) in period

[11]The desired output gap under symmetric shocks is zero, as no shocks hit the economy in period 1.

1 in the left panels. The difference between this two line in period 1 is entirely explained by the central bank's optimal response to the negative balance of inflation risks.

In period 2, a negative markup shock materializes, as illustrated by the markup shock shown in the lower-right panel. At this point, agents understand that from now on, shocks will be drawn from a symmetric distribution. Consequently, in this period and in the subsequent ones, the central bank follows the same policy under the symmetric and the asymmetric risks cases. Despite adopting the same policy strategy and operating in the same macroeconomic environment, the shift in the balance of risks in period 1 leaves persistent effects on macroeconomic outcomes. This is due to the path dependence introduced by optimal commitment.[12]

In summary, in the presence of asymmetric risks (period 1), optimal policy requires the central bank to counteract the direction of the inflation risk. Specifically, an optimizing central bank must shift agents' inflation expectations in the opposite direction of the risk. In Figure 1, this shift in the central scenario is depicted by the difference between the blue dashed line with circle markers and the red dashed-dotted line in period 1 in the bottom-left panel. The central bank accomplishes this by accepting a larger output gap in period 1. Conversely, if the balance of risks is tilted to the upside, it is optimal for the central bank to cool down the economy, thereby pushing the modal inflation scenario downward and accepting a negative output gap.

Not surprisingly, in the symmetric case, optimal monetary policy achieves a higher level of welfare in every period, underscoring the fact that macroeconomic risks can exacerbate the trade-offs faced by the central bank.

4 An econometric framework for inflation risk

Having established that changes in the asymmetry of inflation risk warrant a response from the central bank, it is important to investigate whether this phenomenon is indeed a significant

[12]However, the unwitting central bank continues to target the wrong price level in period 2, resulting in a suboptimal output gap. The targeted price level is incorrect because the central bank failed to meet its price level target in the previous period. If the unwitting central bank were to target the correct price level, it would align with the price level shown in the symmetric risk case (the solid black line). It should be noted, however, that the unwitting central bank can achieve the (wrong) targeted price level in period 2. As shown in the bottom-left chart, the red star lies exactly on the red dashed-dotted line in period 2. This occurs because future shocks are symmetrically distributed, thereby eliminating the bias in agents' expectations.

feature of real-world data.

As a first step, we formally test for time variation in inflation's third moment, finding strong evidence to reject the null of no variation. Therefore, we propose a time series model that allows us to track the time variation of the full predictive density of inflation in real-time. The model we propose inherits many features of common unobserved components specifications (along the lines of Stock and Watson, 2007), and it comes with the additional flexibility of allowing for skewness in the predictive distribution of inflation.

Using this model, we provide fresh insights into the stochastic properties of inflation, emphasizing the dynamics of skewness. We show that shifts in risk asymmetry drive significant changes in the balance of risks, which have varied over time. For example, the balance of risks has been predominantly positive during periods such as the 1970s and the post-COVID recovery, while it turned negative following the Great Recession. Moreover, we demonstrate that explicitly accounting for time-varying skewness in the inflation process enhances out-of-sample forecasting performance. These findings highlight the asymmetric nature of inflation outcomes and underscore the value of incorporating the third moment to deepen our understanding of inflation dynamics.

4.1 Is there any evidence of time-varying inflation skewness?

As the initial step in our analysis, we test for evidence of time variation in the asymmetry of the predictive distributions of core PCE inflation.

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

$$\pi_t \sim Skt_\nu(\mu_t, \sigma_t^2, \varrho_t), \tag{16}$$

where ν denotes the, time invariant, degrees of freedom. The degree of the asymmetry is risk is captured by the parameter $\varrho_t \in (-1, 1)$. The distribution of inflation realizations is positively (negatively) skewed for $\varrho_t > 0$ ($\varrho_t < 0$) and the underlying right- and left-risk around the central scenario (mode), μ_t, can be retrieved as $\sigma_t(1-\varrho_t)$ and $\sigma_t(1+\varrho_t)$. This specification allows as special

Table 1: Time variation in higher order moments

	Q	Q^*	N	Q	Q^*	N
	Homoskedastic			*Heteroskedastic*		
$Scale^2$				369.36^{***}	373.67^{***}	1.50^{***}
Asymmetry	367.31^{***}	371.60^{***}	4.18^{***}	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 $< 1\%$.

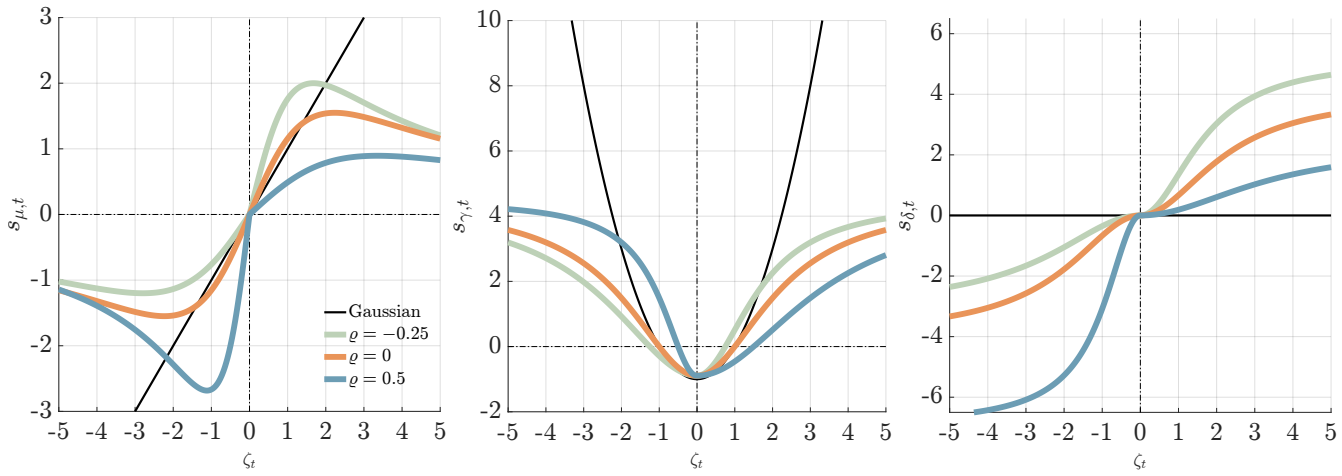
cases: (i) the symmetric Student-t distribution when $\varrho_t = 0$, (ii) the epsilon-Skew-Gaussian for $\nu \to \infty$, and (iii) the Gaussian density when both conditions hold jointly. Thus, we allow for, but do not impose, asymmetric innovation terms.

Within this framework, we formally test the null hypothesis that ϱ_t remains constant over the entire sample. Table 1 presents the results of three parametric Lagrange Multiplier tests: the Q test, the adjusted Q* test, and the Nyblom test as outlined by Delle Monache et al. (2024). These tests differ in their chosen alternative hypotheses, resulting in varying statistical power (see, e.g., Harvey, 2013). We consider two scenarios: one assuming constant volatility over time and another allowing for time-varying volatility. In both cases, the tests strongly reject the null hypothesis of symmetry at the 1% confidence level.

We also conduct additional tests using rolling estimates of inflation asymmetry. These tests further underscore the importance of accounting for time-varying skewness when tracking the conditional distribution of inflation (see Appendix F).

4.2 Model Specification

Having established the importance of accounting for the evolution of inflation asymmetry in assessing inflation risk, the next step is to evaluate its practical significance. Following a long tradition in modeling the stochastic 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



(a) Location **(b)** Scale **(c)** Shape

Figure 2: Parameter updating

Note: Note: The panels report the scaled scores for different values of the standardized prediction error $\zeta_t = \frac{\varepsilon_t}{\sigma_t}$. We consider the Gaussian case (black), the symmetric $Skt_5(\mu_t, \sigma_t, 0)$ (orange), and negatively (green) and positively (blue) $Skt_5(\mu_t, \sigma_t, \varrho_t)$.

Watson (2007), we opt for an observation-driven updating process.[13] Specifically, let $\delta_t = \log(\sigma_t)$ and $\gamma_t = \text{arctanh}(\varrho_t)$, we postulate that each element $f_{i,t}$ of $f_t = (\mu_t, \delta_t, \gamma_t)'$ features a permanent and transitory component: $f_{i,t} = \bar{f}_{i,t} + \tilde{f}_{i,t}$, which evolve as:

$$\bar{f}_{i,t} = \bar{f}_{i,t-1} + a_i s_{i,t-1}, \tag{17}$$

$$\tilde{f}_{i,t} = \phi_i \tilde{f}_{i,t-1} + b_i s_{i,t-1}. \tag{18}$$

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).[14]

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; learning rates, a_i and b_i, regulate the strength

[13] 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).

[14] The scaled 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 with respect to the dynamic parameters; the scaling matrix $\mathcal{S}_t$ is proportional to the inverse of the diagonal of the Information matrix, $\mathcal{I}_t = E[\nabla\nabla']$. 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. See Blasques et al. (2015, 2022)

of the updates.[15] We illustrate the updating mechanism in Figure 2.

Consider a symmetric Gaussian environment (black lines), where $\nu \to \infty$ and $\varrho_t = 0$ in every period. The location and scale parameters –which now represent the mean and standard deviation of the distribution– update in line with standard Kalman filter learning (see, e.g., Cogley, 2002). Updates of the mean are proportional to the prediction error, with strength inversely proportional to the variability of the data, and the volatility is updated proportional to the difference between the variability of the prediction error and the expected variability of the inflation process. Allowing for fat tails makes the updating mechanism robust to large, unanticipated prediction errors (orange lines; see. e.g., Delle Monache and Petrella, 2017; Antolín-Díaz et al., 2024) When asymmetry is introduced, the updating mechanism weights prediction errors differently, depending on their sign. For example, when the conditional distribution is left skewed (green lines), 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, large deviations of inflation from the expected central scenario imply updates of the asymmetry parameter in the direction of the prediction error.

The parameters of the model and the associated conditional distribution of inflation are estimated using Bayesian methods as in Delle Monache et al. (2024); refer to Appendix C for additional details. An extensive Monte Carlo exercise, reported in Appendix D, demonstrates that the procedure detects skewness only when it is present and remains robust to (changing) correlations between location and scale.

Expected value under asymmetry A defining feature of any skew-distribution, $p(\pi|\mu_t, \sigma_t, \varrho_t, \nu)$, is the fact that asymmetry directly affects the first moment of the distribution. Specifically, in the

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

case of the Skew-t distribution in Equation (16), one can show that $\forall h > 0$

$$\begin{aligned} E_t \pi_{t+h} &= \int_{\mathbb{R}} \pi p(\pi | \mu_{t+h}, \sigma_{t+h}, \varrho_{t+h}, \nu) d\pi, \\ &= \mu_{t+h} + \underbrace{g(\nu)\sigma_{t+h}\varrho_{t+h}}_{\psi_{t+h}} \end{aligned} \tag{19}$$

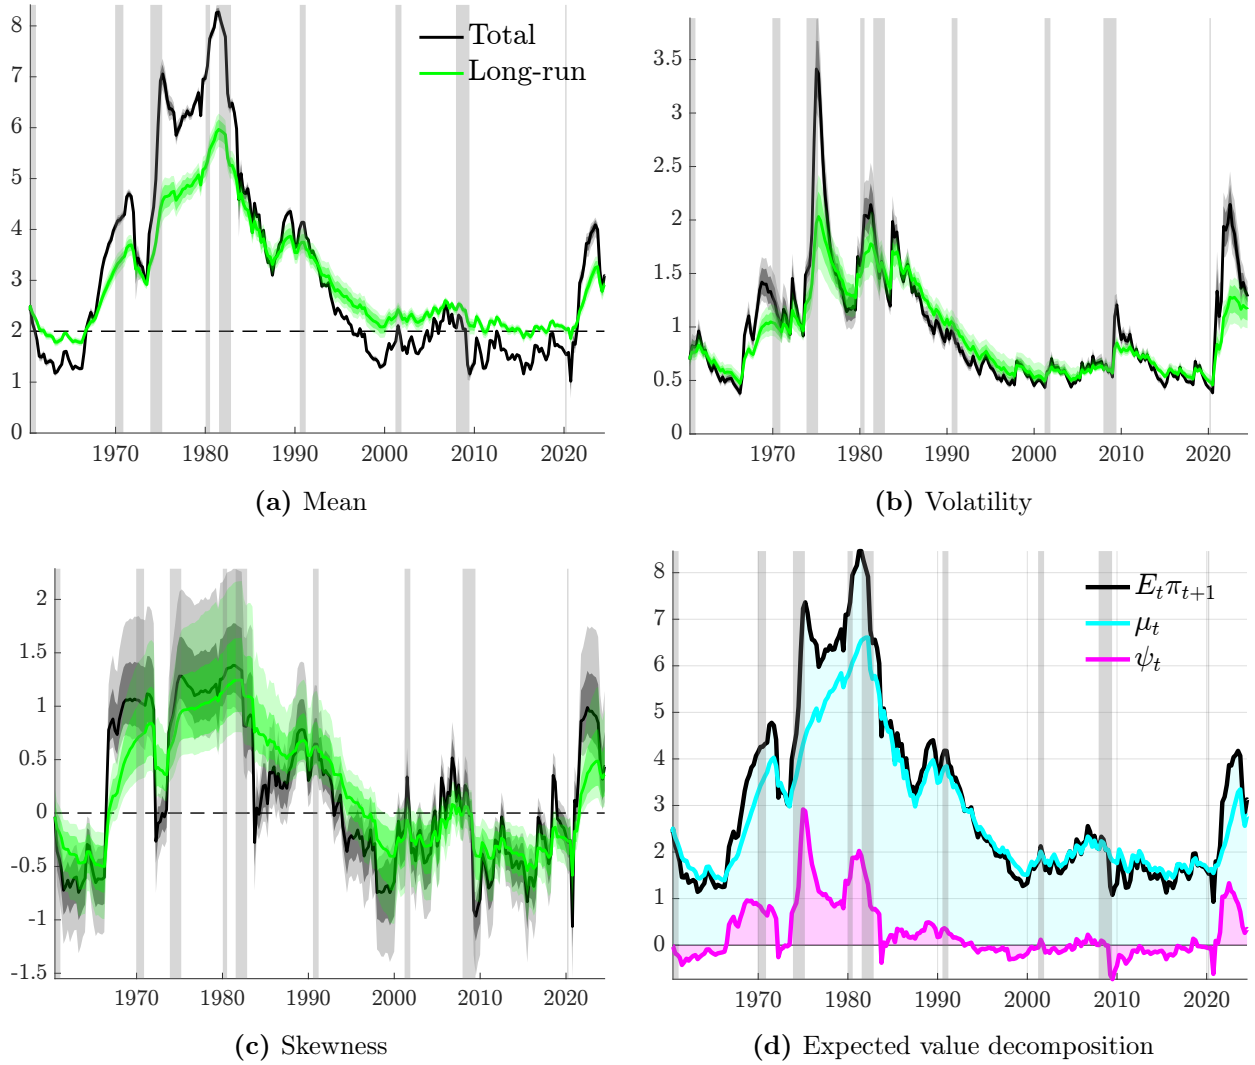
where $g(\nu) = \frac{4\nu\mathcal{C}(\nu)}{\nu-1}$. Therefore, the expected value can be represented as the sum of the mode and a component, ψ_{t+h}, that is a function of the asymmetry parameter. That is, asymmetry creates a wedge between the central scenario, e.g., the mode, and the expected value. This wedge has the same sign of the prevalent asymmetry and is quantitatively more relevant as the distribution becomes more dispersed, as $\frac{\partial E_{t-1}\pi_t}{\partial \varrho_t} > 0, \ \forall t$.

De Polis (2024) provides detailed derivations and closed forms for the variance and skewness of the Skew-t model. The variance depends on σ_t and ρ_t but is independent of the sign of risk asymmetry, whereas skewness depends solely on ρ_t. As a result, procyclical variations in inflation skewness lead to a time-varying correlation between the mean and volatility of the process. When risk is skewed to the upside, an increase in risk corresponds to rising inflation expectations. Conversely, when risk is negatively skewed, an increase in risk results in a decline in expected inflation.

In Section 5 we will exploit Equation (19) to introduce asymmetry in the shocks of a quantitative DSGE to understand the optimal monetary policy response when inflation risks is skewed.

4.3 In-sample inference

The estimated model offers novel insights into the time-varying stochastic properties of the inflation process. Figure 3 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 (labeled as "Long-run and computed setting $f_t = \bar{f}_t$) are in green. The model reveals significant time variation across all moments. The time-varying mean, reported in Figure 3a, 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).



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

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

Figure 3: 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.

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 (Figure 3b) peaks in the mid-1970s, remaining high until the late 1980s, and is well-contained until early 2020, when it sharply increases starting in the second quarter. Unlike the mean, inflation volatility exhibits clear cyclicality, rising significantly during recessions.

The skewness estimates in Figure 3c 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 estimated skewness during the latest inflationary episode closely resembles in magnitude the environment of the mid-1970s, 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.[16]

Following Equation (19), Figure 3d presents the decomposition of expected inflation into the location (the most likely expected outcome) and the tilt induced by the balance of risks around it, ψ. Inflation risk significantly influences inflation expectations, introducing substantial upside bias during the 1970s and the post-COVID-19 era. Negative skewness contributed to a downward bias in expected inflation during the decade leading up to COVID-19. However, low inflation volatility during this period mitigated the average effect of asymmetry on expectations, limiting its impact to around 20 basis points, despite the markedly negative skewness. The subdued volatility of inflation observed during this time has significant implications for how monetary policy should address the persistent negative skewness in the post-Great Recession era. This is an important point to which we will return in the next section.

Whereas more is known about inflation's time-varying mean and variance (see, e.g., Stock and Watson, 2007), our model provides novel insights into the dynamics of inflation skewness. We compare our model-based measures of skewness (e.g., sample skewness, $\mathtt{s_t}$, and the asymmetry parameter, ϱ_t) against rolling estimates derived solely from core PCE data. Panel (a) of Figure 4 compares the estimated skewness from our model to the sample skewness estimated on a 5-year rolling window, ss_t, shown in green. Although both measures exhibit similar patterns, ss_t is

[16] See Table F3 in Appendix F.

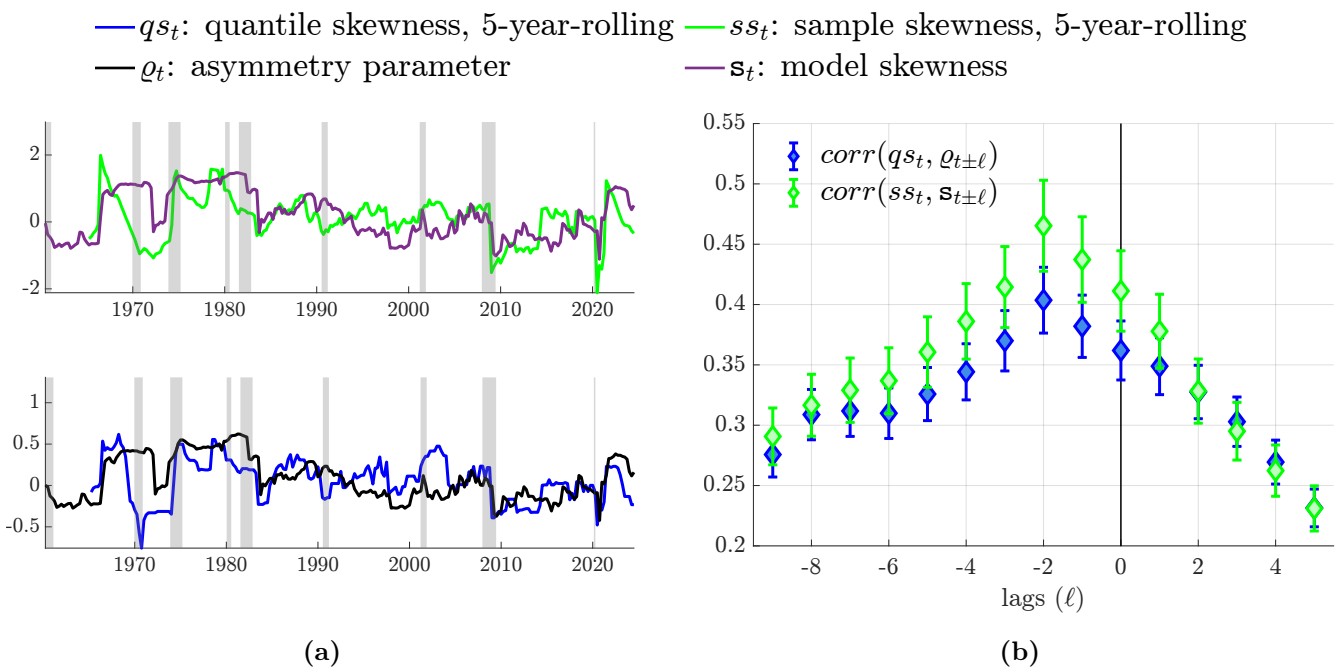


Figure 4: Model-based vs rolling measures of skewness

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.

noticeably noisier due to its sensitivity to individual observations. In the bottom figure, we compute (5-years rolling) estimates of robust quantile-based skewness measure, qs_t. Broadly speaking, this measures the degree of asymmetry in the probability mass on the right and left to the mode. As such this is directly comparable to our asymmetry parameter ρ_t in our model.[17]

Although rolling measures of sample skewness can be easily computed from raw data, rolling window estimators face a trade-off between estimation accuracy and sensitivity to time variation. Larger look-back windows enhance the precision of third-moment estimates by reducing the impact of isolated outliers but compromise the ability to detect and respond to changes in real time. Moreover, rolling estimates assume constant skewness within the sample window, causing them to adjust slowly, particularly during periods of significant inflation shifts—a critical drawback for real-time risk assessment. In contrast, our approach applies a one-sided discounting of past observations when estimating the model's time-varying parameters. This method delivers timely measures of

[17]The impact of outliers, while mitigated, continue to be visible especially the large negative data points from the second quarter of 2020, which continue to distort qs_t estimates post-2021. In contrast, model-based skewness estimates are less influenced by extreme values, due to a more robust updating mechanism.

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 (2007) 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.

inflation's evolving moments while maintaining robustness and minimizing noise in the estimates. While all measures generally align in capturing underlying risk and its evolution, Figure 4b shows the cross-correlation between data-based measures of skewness and lags of model-based estimates. Our skewness estimates respond to changes in inflation more quickly, leading the rolling measures by an average of two quarters –a clear advantage for real-time monitoring of inflation risk.

Appendix E extends the analysis to alternative measures of inflation, examining inflation rates derived from headline PCE, core and headline CPI, and the GDP deflator. Time-varying asymmetry in risk emerges as a robust feature of inflation, regardless of the specific price index used. Periods of upside or downside skewed risk tend to align across these indicators. Notably, core PCE inflation appears to exhibit a more attenuated balance of risk compared to other inflation measures.

4.4 Out-of-sample performance of the model

We assess the out-of-sample forecasting performance of our model, focusing on 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 for the whole density of core PCE inflation, starting from 2000Q1.[18] We evaluate the forecasting performance of the model against the UCSV model of Stock and Watson (2007) which represents a solid benchmark model, widely employed to predict inflation outcomes. We compare the two models in their ability to produce accurate point, density and event forecasts. Specifically, we evaluate the mean squared forecast error (MSFE) for point accuracy and we use Gneiting and Ranjan (2011) quantile-weighted CRPS to assess density forecasting accuracy. These scoring rules measure the squared difference between the forecast cumulative distribution function and the "perfect forecast", that is a step function which moves from 0 to 1 on the realization point. Furthermore, we evaluate the accuracy of predicting the right and left tails, and the central body of the predictive densities.[19]

Results of the comparison are presented in Table 2, where we report, for each loss function and forecast horizon, the ratio of the score achieved by our model to that of the UCSV benchmark. Values below unity indicate superior accuracy of our preferred model, p-values for the Diebold and Mariano (1995) test are provided in parentheses. The results strongly support the superiority of our model over the benchmark across all horizons and forecast exercises. Gains in point forecasts range from 15% at short horizons to 5% over the medium term.

Smaller yet significant gains are observed in CRPS scores, with our model notably delivering improvements of up to 8% in forecasting upside risks. These results underline the importance of accounting for inflation skewness as a means to enhance forecasting accuracy. However, the UCSV model not only lacks a mechanism for capturing skewness but also overlooks the presence of fat tails in the data. To ensure that the observed gains stem specifically from modeling inflation skewness rather than from addressing fat tails, we replicate the analysis using a specification that excludes any asymmetry, following a similar approach to Delle Monache and Petrella (2017). The

[18] We start the exercise in 2000Q1 due to the availability of real-time data vintages.

[19] Left tail forecasts are defined up to the 25^{th} quantile. Similarly, the right tail considers above the 75^{th} quantile. The remaining quantiles characterize to the center of the distribution.

Table 3: Event forecast comparison against SPF

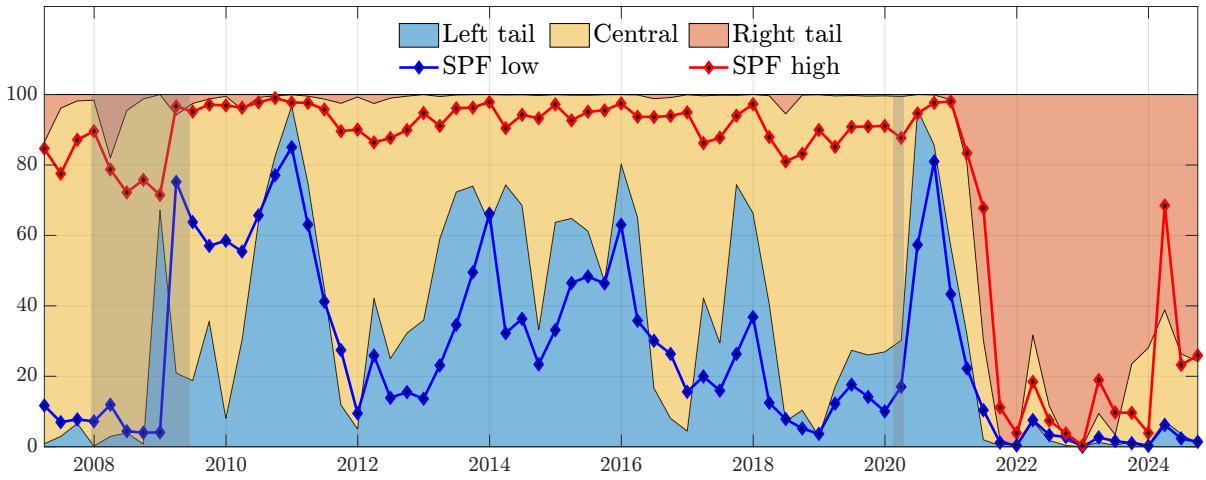
$\pi_t^{Q4} < 1.5$				$\pi_t^{Q4} > 2.5$				$1.5 \leq \pi_t^{Q4} \leq 2.5$			
h = 1	h = 2	h = 3	h = 4	h = 1	h = 2	h = 3	h = 4	h = 1	h = 2	h = 3	h = 4
1.196	0.892	1.159	0.893	0.853	0.164	0.335	0.599	1.177	0.895	1.327	1.125

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.

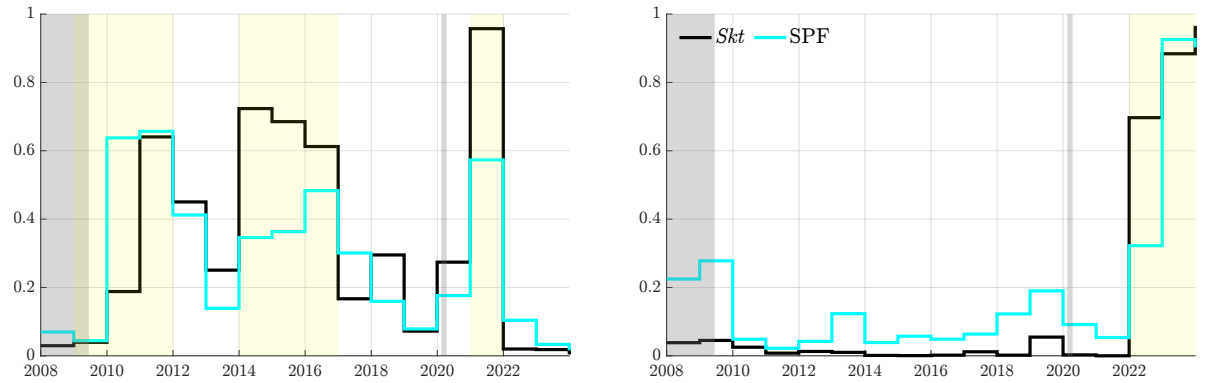
results, reported in Appendix F, confirm our assertion that incorporating skewness is a critical step for improving model fit and inflation forecasting accuracy.

Event forecast and comparison with SPF We now compare the ability of the competing models to produce event forecasts. Specifically, we use the Brier score, defined as the mean squared difference between the predicted probability of a binary outcome and its occurrence, to evaluate predictions of events where π_{t+h} is lower than 1.5%, greater than 2.5%, or falls within these thresholds. The comparison with the UCVS model is reported in the bottom section of Table 2. Again, our model provide improvements over the benchmark model, especially over the short-horizon.

We also compare our event forecast predictions against the SPF. Adams et al. (2021) show that both upside and downside risks to the median SPF forecast fluctuate over time, and that lower quantiles of the predictive distributions are generally more stable than the upper quantiles. For this exercise, we target Q4-over-Q4 core PCE, π_t^{Q4}, starting in 2007Q1 to match SPF data. First, panel (a) of Figure 5 shows the interval predictions generated by our model, overlaid with the SPF projections. Our model produces interval predictions that align closely with the SPF's. In the bottom panels, we plot the evolution of the predicted probability of π_t^{Q4} being below 1.5% (Figure 5b) or above 2.5% (Figure 5c) 6 months ahead of the realization; the predicted events are represented by the yellow-shaded areas. These figures demonstrate that our model provides timelier assessments of event probabilities, in that the black lines lie above the bright blue ones when predicted events realized (yellow shaded regions). Notably, our model detects little to no probability of overshooting the 1.5–2.5% interval throughout the period between the GFC and the



(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 5: 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.

post-pandemic inflation surge, while capturing the latter with greater precision.

A formal evaluation of these event predictions is summarized in Table 3, where we report the ratios of Brier scores for our model relative to the SPF for the three events and up to four-step-ahead predictions.[20] Overall, the table indicates that our model and the SPF achieve comparable accuracy, except for $P\left(\pi_{t+1}^{Q4} > 2.5\%\right)$, where our model provides more timely probability assessments (see Figure 5, panel (c)). These gains stem from the model's consistently lower event probability

[20]For $h = 1$, predictions are based solely on out-of-sample values. As h approaches 4, up to three observed data points are used in the computation of Q4-over-Q4 inflation. It is important to note that the ratios reported in the table are based on a small sample size due to the limited availability of SPF data.

assessments compared to the SPF during the decade following the Great Recession. Additionally, the model adjusts more quickly, reaching higher values in early 2022, accurately timing the inflation spike of the post-COVID period.

5 Quantitative structural analysis

In this section, we perform policy counterfactuals to highlight the importance of accounting for asymmetric inflation risk, using a prominent empirical DSGE model estimated with U.S. data (Smets and Wouters, 2007). Specifically, we consider a version of the model in which price markup shocks follow a Skew-t distribution with time-varying moments. As in the stylized model considered in Section 3, agents receive news about the evolving skewness in the distribution of these shocks up to $J = 20$ quarters ahead. We assume that, in each period, agents expect these moments to change in line with the forecasts provided by the model estimated in the previous section.[21]

To study the DSGE model with changing asymmetric inflation risks, we proceed as follows. First, we derive the beliefs representation of the asymmetric risks for this model. Second, we calibrate the beliefs representation to match the effects of the inflation asymmetry on inflation expectations, which we estimated in the previous section. Third, we analyze the macroeconomic implications of introducing the RAIT – a central bank communication strategy designed to anchor expectations to target by counteracting the effects of asymmetric risks – with particular focus on the post-pandemic period. Finally, we compare the implications of the RAIT with those of the FAIT, which is the current framework employed by the Federal Reserve.

Analogously to what shown in Section 3.2, to construct the beliefs representation of asymmetric inflation risks for the Smets and Wouters (2007) model, we augment the price Phillips Curve with

[21] It should be noted that whereas the model we introduced in the previous section delivers accurate and timely real-time estimates of U.S. inflation skewness, the beliefs representation can accommodate various methods (or combinations thereof) for estimating skewness, extending beyond our preferred model.

a set of dummy anticipated price markup shocks as follows:[22]

$$\hat{\pi}_t = \pi_1 \hat{\pi}_{t-1} + \pi_2 E_t \hat{\pi}_{t+1} - \pi_3 \mu_t^p + \varepsilon_t^p + \sum_{j=0}^{J} \varphi_{t-j}^j, \tag{20}$$

where the variable μ_t^p represents firms' price mark-ups, the shock ε_t^p stands for the actual price markup shock, which follows a Gaussian ARMA process, and π_1, π_2, and π_3 are the standard parameters in this equation, as defined in Smets and Wouters (2007). The last term on the right-hand side captures anticipated and unanticipated dummy mark-up shocks, with φ_t^j representing the mark-up shock revealed in period t but expected to materialize in period $t+j$. These anticipated shocks capture the effects of the skewness of the Skew-t distribution of future price markup shocks on agents's expectations about the realization of these shocks.

In the beliefs representation, the expected price markup shock h periods ahead can be decomposed as follows:

$$E_t\left[\varepsilon_{t+h}^p + \sum_{j=0}^{J} \varphi_{t-j}^j\right] = \underbrace{E_t(\varepsilon_{t+h}^p)}_{\text{mode } (\mu_{t+h|t}^p)} + \underbrace{\sum_{j=h}^{J} \varphi_{t+h-j}^j}_{\text{balance of risks } (\psi_{t+h|t}^p)}, \tag{21}$$

which mimic the structure of Equation (19), since we assumed that price markup shocks distribute as a Skew-t distribution. The balance-of-risks term on the right of the decomposition captures the effects of skewness in the distribution of the shocks on their expected value.

It then follows that the time-t revision to agents' expectations about the realization of the price markup shock h periods ahead is given by

$$E_t - E_{t-1}\left[\varepsilon_{t+h}^p + \sum_{j=0}^{J} \varphi_{t-j}^j\right] = \left(\mu_{t+h|t}^p - \mu_{t+h|t-1}^p\right) + \varphi_t^h. \tag{22}$$

Note that effects of the revision to the skewness from period $t-1$ to period t on the expected value of the shock in period $t+h$ are captured by the dummy anticipated shock, φ_t^h.

[22] This equation embeds an identification assumption: specifically, we assume that only expectations of future price markup shocks drive the changes in inflation risks estimated by our forecasting model. However, this approach can be generalized to accommodate alternative identification assumptions by incorporating a variety of dummy anticipated shocks that contribute to these risks.

How to make sure that the beliefs representation match the estimated effects of the balance inflation risks on expectations? First, realize that since all shocks (the seven structural shocks of the Smets and Wouters (2007) model plus the $J+1$ dummy shocks) are normally distributed, the linear approximation of the model in its beliefs representation can, therefore, be solved quickly and efficiently using off-the-shelf techniques. The solution of the model is standard and can be expressed as follows:

$$\mathbf{s}_t = \Gamma \mathbf{s}_{t-1} + \Omega \varepsilon_t, \tag{23}$$

where the vector $\mathbf{s}_t$ contains the model's state variables, and the vector ε_t includes all shocks of the beliefs representation of the model. The matrices Γ and Ω are the solution matrices, which depend on the structural parameters of the model.

Second, we have to align the beliefs representation to match the inflation asymmetry estimated in the previous section. To achieve this, we find the vector of $J+1$ dummy surprise and anticipated shocks that satisfy the following system of $J+1$ linear equations:

$$\begin{bmatrix} -\sum_{j=1}^{J} \varphi_{t-j}^{j} \\ \psi_{t+1|t} - \psi_{t+1|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{0}_{1\times J} \\ \Omega^S & \Omega^N \end{bmatrix} \begin{bmatrix} \varphi_t^0 \\ \varphi_t^1 \\ \vdots \\ \varphi_t^J \end{bmatrix}, \tag{24}$$

where Ω^S and Ω^N represent the contemporaneous effects of surprise and anticipated dummy markup shocks on inflation expectations in the model.

The first equation ensures that the effects of dummy anticipated shocks remain confined to the realm of beliefs and never materialize into an actual shock to the markup. As noted in Section 3, the function of these shocks is solely to capture the effects of asymmetric risks on agents' expectations. The first term on the left-hand side aggregates all past anticipated shocks expected to affect the economy in the current period, t. The equation imposes that the surprise dummy shock, φ_t^0, fully offsets these past anticipated shocks. In this sense, the dummy anticipated shocks in the beliefs

representation can be interpreted as pure beliefs or sentiments.

Moving to the last J equations and starting from left-hand side, $\psi_{t+j|t} - \psi_{t+j|t-1}$, with $j \in \{1, \ldots, J\}$, denote the revisions to the mean-mode inflation wedge estimated by the forecasting model – Equation (19) – in the data. On the right-hand side, we have the surprise and anticipated dummy markup shocks. The anticipated shocks $\left\{\varphi_t^j\right\}_{j=1}^{J}$ are multiplied by the matrix Ω^N which returns the effects of these shocks on inflation expectations at each horizons $(1, ..., J)$. So in the last J equations, we are imposing that the dummy anticipated shocks moves inflation expectations to match the revisions to the mean-mode inflation wedge estimated in the data at any horizon one through J quarters out. As explained earlier, the surprise shock that is meant to wipe out the effects of the past anticipated shocks has an effect on inflation expectations, which we have to take into account. These effects are captured by the vector $\Omega^S \varphi_t^0$.

We solve these system of linear equations in Equation (24) to obtain the realization of dummy surprise and anticipated price markup shocks, $\left\{\varphi_t^j\right\}_{j=0}^{J}$, that allows the beliefs representation to exactly match the effects of the estimated revisions to inflation skewness on agents' inflation expectations in every period of our estimation sample.

The next step is to introduce the RAIT, which is defined as a central bank communication strategy designed to anchor expectations to target by counteracting the effects of asymmetric risks we estimated in the data on inflation expectations. Under the RAIT, average inflation over the medium run aligns with the central bank's target, provided that the distribution of inflation outcomes is accurately estimated for every period and horizon.

Implementing the RAIT requires the central bank to tilt the expected central inflation scenario to counter the direction of the perceived risks. Specifically, if the balance of risks suggests upward inflation pressures, the central bank would temporarily aim to undershoot its target. Conversely, if the balance of risks points to downward inflation pressures, the central bank would temporarily try to overshoot its inflation target.

We take two simple steps to practically implement the RAIT in the Smets and Wouters (2007) model with time-varying balance of inflation risks. First, we enhance the monetary policy reaction

function as follows:

$$\hat{r}_t = \rho\hat{r}_{t-1} + (1-\rho)\left[r_x\hat{x}_t + r_{\Delta x}\Delta\hat{x}_t + r_\pi\left(\hat{\pi}_t - \underbrace{\sum_{j=1}^{J}\hat{\pi}^{\star}_{t|t-j}}_{\hat{\pi}^{\text{RAIT}}_t}\right)\right] + \varepsilon^r_t, \tag{25}$$

where the variable $\hat{\pi}^{\star}_{t|t-j}$ represents the forward guidance shocks that capture the central bank's communications in period $t-j$ regarding its intention to overshoot or undershoot the inflation target. Meanwhile, the variable $\hat{\pi}^{\text{RAIT}}_t$ denotes the magnitude of the overshoot or undershoot in period t as announced in the central bank's past communications. A negative (positive) value of $\hat{\pi}^{\text{RAIT}}_{t+h}$ indicates the central bank's intention to undershoot (overshoot) its inflation target h periods from now. In the model, the inflation target is set by a fixed parameter calibrated to align with a 2% annualized inflation rate. The variables $\hat{r}_t$, $\hat{x}_t$, and $\Delta\hat{x}_t$ denote the nominal interest rate, the output gap, and the first difference of the output gap, respectively, in log deviations from their steady-state value. The parameters in the monetary policy rule are standard and defined in Smets and Wouters (2007).

Second, under the RAIT, in each period t, the central bank communicates revisions to its overshooting or undershooting policy in response to changes in the balance of inflation risks and their impact on agents' expectations, as estimated from the data in Section 4. These impacts, denoted as $\psi_{t+j|t} - \psi_{t+j|t-1}$, could lead to a de-anchoring of inflation, which the central bank seeks to prevent by updating markets on its future intention to overshoot or undershoot its long-run inflation target. When implemented effectively, these communications fully offset the effects of the changing balance of inflation risks on expectations, thereby restoring the anchoring of inflation expectations.

To obtain this communication strategy in the beliefs representation of the Smets and Wouter's model, we require the sequence of forward guidance shocks, $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=0}^{J}$ to satisfy the following

system of J linear equations:

$$-\begin{bmatrix} \psi_{t+1|t} - \psi_{t+1|t-1} \\ \psi_{t+2|t} - \psi_{t+2|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = \Omega_{FG} \begin{bmatrix} \hat{\pi}^{\star}_{t+1|t} \\ \hat{\pi}^{\star}_{t+2|t} \\ \vdots \\ \hat{\pi}^{\star}_{t+J|t} \end{bmatrix}, \tag{26}$$

The $J \times J$ matrix Ω^{FG} captures the impact response of the sequence of forward guidance shocks issued at time t on current inflation expectations (from 1 quarter through J periods ahead). This matrix can be readily obtained by solving the beliefs representation of the Smets and Wouters (2007) model, augmented with the monetary policy rule that includes forward guidance shocks, $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=1}^{J}$ – see Equation (25). The solution to such a model can be obtained using standard solvers for linear Rational Expectations models, such as *Dynare*. The vector on the left-hand side represents the effects of the estimated revisions to the balance of risks on inflation expectations, as defined in Equation (19). The system of equation imposes that the shocks capturing the forward guidance issued in period t under the RAIT – $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=1}^{J}$ – are chosen by the central bank so as to offset these effects. Hence, the vector capturing these effects enters with a negative sign.

The sequence of forward guidance shocks solving the system of linear equations in Equation (26) ensures that the estimated balance of inflation risks has no effect on expectations. If the central bank shapes its forward guidance according to the RAIT, inflation expectations will remain anchored at all horizons, and average inflation will converge to the target in the medium term, as the effects of inflation skewness on inflation expectations are fully neutralized by the strategy.

Forward guidance puzzle? The RAIT is implemented through forward guidance. The literature has highlighted the forward guidance puzzle—i.e., the implausible predictions of macroeconomic models regarding the effects of forward guidance. However, this puzzle does not arise in our analysis. As shown by Del Negro, Giannone, Giannoni, and Tambalotti (2023); Bianchi, Melosi, and Nicolò (2024); Maliar and Taylor (2019), the forward guidance puzzle emerges when forward guidance shocks are used to match a specific path for the interest rate, as this approach implicitly

removes the response of active monetary policy to those shocks. In contrast, our framework sets forward guidance shocks to match the future path of inflation that restores inflation expectations—distorted by asymmetries in risk—back to target. In doing so, we account for the feedback effects of monetary policy rather than suppressing them.

An alternative way to interpret forward guidance under the RAIT. The implementation of the RAIT does not necessarily require communication about temporary overshooting or undershooting of the inflation target. Technically, the forward guidance shocks $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=0}^{J}$, appropriately rescaled by $-(1-\rho)r_{\pi}$, can be interpreted as stochastic upward or downward shifts in the intercept of the reaction function, affecting the expected future path of interest rates. In the case of the RAIT, these shifts are tied to the central bank's assessment of the balance of inflation risks. The changes in the future interest rate path would, therefore, be driven by forward guidance, designed to inform the public about the appropriate reaction function required to counterbalance variations in the balance of inflation risks.

5.1 The RAIT in the post-pandemic era

The first counterfactual exercise illustrates how the central bank would have communicated its future policy actions under the RAIT during the challenging period of the post-pandemic inflation surge.[23] In this exercise, we align the beliefs representation to match the forecasts of the mean-mode wedge estimated using the Skew-t model introduced in Section 4. This involves setting the dummy surprise and anticipated price markup shocks to satisfy Equation (24). We choose an horizon of five years when solving that system of equations; that is, $J = 20$ quarters.

In Figure 6, we illustrate how the beliefs representation translates predictions of the balance of risks for inflation into communications about the path of future interest rates, or forward guidance.[24]

[23]Estimates reported in Smets and Wouters (2007) cover the period 1966–2004, during which the authors estimate a Phillips curve slope of approximately 0.03. However, many studies suggest that the slope of the Phillips curve during the recent inflation surge was significantly steeper. We therefore calibrate the parameters feeding into the Phillips curve to achieve a slope of approximately 0.3.

[24]For this exercise, we applied a six-quarter moving average to the estimated balance of risks to smooth out the noise introduced by the forecasting model, which is particularly evident when one wants to estimate the skewness during such an unprecedented high-volatility period. Arguably, policymakers would assess the dynamics of inflation

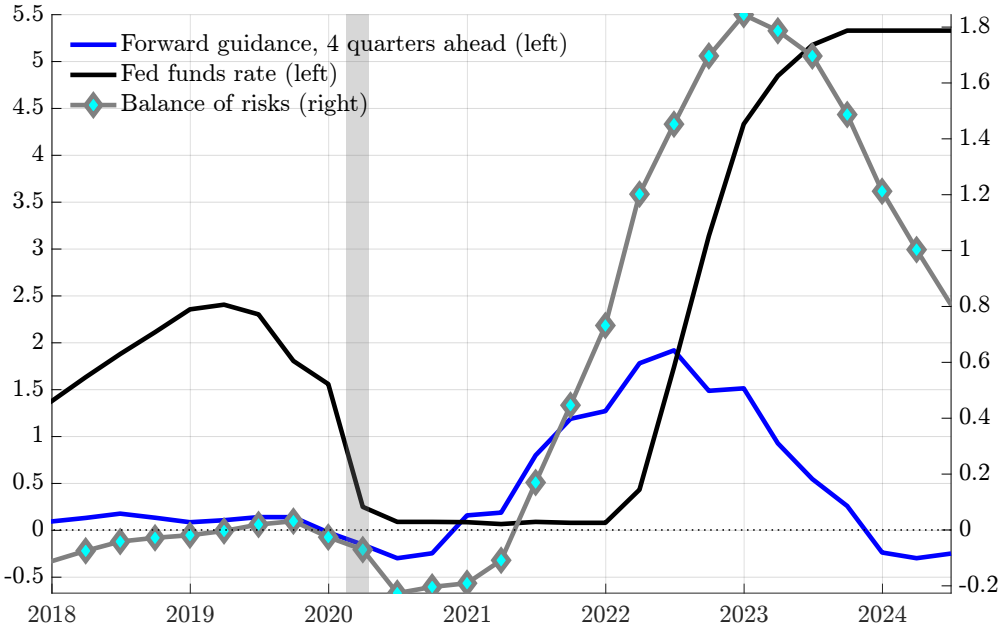


Figure 6: Forward guidance under the RAIT

Note: The figure reports the forward guidance on the policy rate recommended by the RAIT (left axis) from 2018 to 2024Q2, again the balance of inflation risk (right axis), estimated from the model introduced in Section 4. Gray shaded areas represent NBER recessions.

In Figure 6, the blue solid line represents the inflation overshoot communicated over the past four quarters to offset the effects of the balance of risks estimated during that period (the solid gray line with blue diamond markers.) Since forward guidance typically refers to the likely path of future interest rates, we express these revisions in interest rate units. Specifically, the solid blue line is computed as follows:

$$\text{Forward guidance}_t = -(1-\rho) r_\pi \sum_{j=1}^{4} \hat{\pi}^{\star}_{t+j|t},$$

where the coefficients ensure that the forward guidance shocks are appropriately mapped into the interest rate space.

The solid black line depicts the observed federal funds rate, the policy rate controlled by the Federal Reserve. The solid gray line with blue diamond markers represents the estimated balance of risks, $\psi_{t|t} - \psi_{t|t-1}$, obtained from the forecasting model introduced in Section 4. The balance of risks is expressed in inflation units and, hence, is mapped to the right axis of the chart.

risks using a combination of model averaging and judgment, especially in periods of heightened macroeconomic volatility, such as the post-pandemic era. The smoothing applied to the balance of risks reflects this consideration.

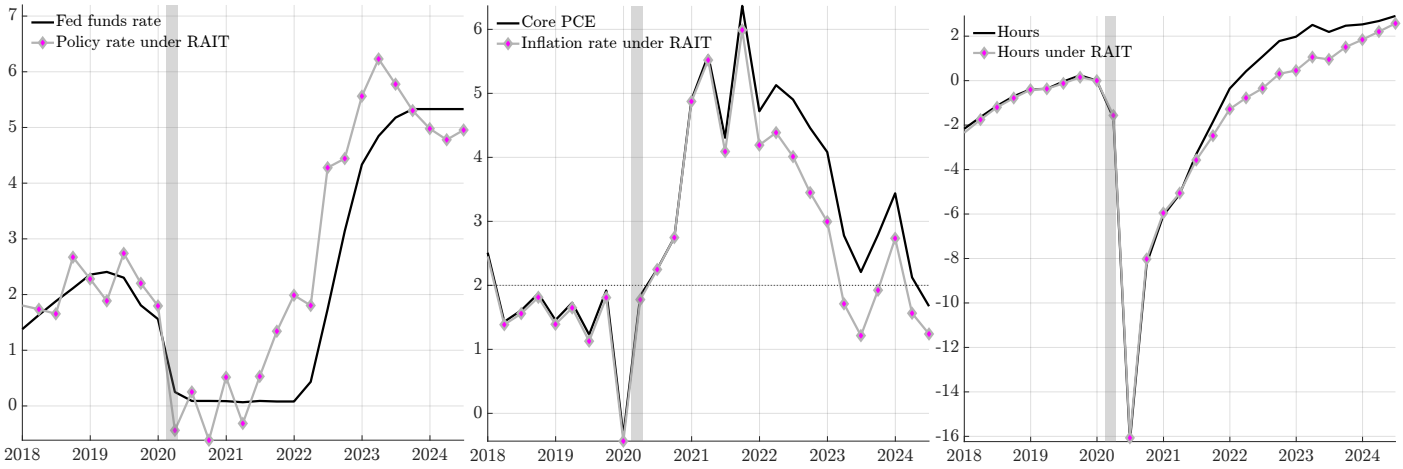
Our model's real-time assessment of the rapid changes in the balance of risks during the second half of 2020 leads the RAIT to recommend increasingly hawkish communication of the central bank's policy stance through early 2022. This hawkish forward guidance is somewhat mitigated in the following year, becoming essentially neutral by 2024, when risk assessments are considerably more balanced.

Comparing forward guidance with the actual path of the federal funds rate suggests that the central bank began raising the policy rate with some delays compared to what the RAIT would have suggested. Conversely, the RAIT would have recommended the central bank start lowering rates slightly earlier than observed.

As we will show, this delay in lowering the interest rate is broadly consistent with the recommendations of the Federal Reserve's framework. Specifically, this delay reflects the FAIT's limited adaptability to changes in the stochastic environment of inflation, a limitation that the RAIT seeks to address through real-time assessment of the skewness in the probability distribution of future inflation outcomes. We will return to this critical point in subsequent discussions.

In Figure 7, we show the counterfactual dynamics of the federal funds rate, core PCE inflation rate, and hours worked implied by the RAIT alongside the actual realizations. These counterfactual variables are computed by first simulating the interest rate, inflation, and hours from the beliefs representation of the model using only the forward guidance shocks, which solve the system of equations Equation (26), and then adding the observed federal funds rate, core PCE inflation rate, and hours worked (in logs) to the simulates series. This additive approach is appropriate because the beliefs representation of the model is linear.

Starting in mid-2021, the RAIT-consistent policy rate exceeds the actual policy rate. This positive gap reflects the prevailing upside risks to inflation, which necessitate tighter monetary policy to counter the impact of more likely future inflation spikes on agents' current inflation expectations. A substantial tightening is prescribed in 2022, following a strong resurgence of inflation after a brief quarter of slower price growth. It is important to note that the estimation of the skewness in the inflation distribution is conducted in real time, making abrupt revisions of risks more frequent – particularly in periods of heightened volatility, such as the post-pandemic



(a) Interest rate **(b)** Core PCE inflation **(c)** Hours worked

Figure 7: Macro dynamics under RAIT

Note: The plots report the counterfactual dynamics of the policy rate, inflation and hours worked consistent with the RAIT framework, against the actual data (black). Inflation is defined as annualized quarter-on-quarter core PCE. Hours worked are the logarithm of hours worked in the nonfarm business sector, normalized to 0 in 2019Q4. Gray shaded areas represent NBER recessions.

era.

The alternative policy path implied by the RAIT would have anchored inflation expectations and driven inflation to the 2% target one year earlier (see the middle chart of Figure 7). RAIT would have delivered a slower, but more balanced recovery of the labor market (see right chart of Figure 7). Specifically, the faster tightening under the RAIT would have significantly slowed the recovery of the labor market in 2022 and much of 2023. Yet, the strategy's early call for easing monetary policy in 2023 would have stimulated the labor market, helping it partially regain the momentum lost due to the earlier tighter policy.

At the peak of the monetary policy response, the RAIT prescribes a policy rate of approximately 5.5%. Notably, the peak of the tightening is remarkably similar to the level ultimately chosen by the Federal Open Market Committee (FOMC), even though the RAIT would have prescribed starting the tightening sooner and reaching that peak two quarters earlier. Moreover, consistent with the forward guidance analysis shown in Figure 6, the RAIT would have recommended a more rapid unwinding of the tightening. This pattern reflects the swift retrenchment of positive inflation risks, as estimated in real time by our econometric model.

5.2 FAIT vs. RAIT

We now compare the implications of adopting the FAIT or the RAIT for the central bank's communications. Both strategies aim to counterbalance asymmetric inflation risks by clarifying to the public that the central bank may tolerate an overshooting or undershooting of inflation relative to the 2% target in the short run.

One of the reasons behind the Federal Reserve's adoption of the FAIT was the heightened risk of encountering the ZLB constraint, driven by the low interest rate environment of the past decade (Clarida, 2022). This limited the central bank's ability to stabilize the economy during recessions, tilting the balance of inflation risks to the downside. In this context, the FAIT serves as a strategy whose objective is to counterbalance this balance of risks by enabling the central bank to convey its willingness to tolerate an overshoot of inflation for a period of time.

Similar to the RAIT, the FAIT aims to shift the distribution of inflation outcomes to counter the balance of inflation risks and reanchor inflation expectations to the desired central bank's target.[25] However, the two strategies differ in the conditionality underlying their respective communications regarding the opportunity to overshoot or undershoot the target. Under the FAIT, the central bank seeks to overshoot the target if the repeatedly binding ZLB constraint has lowered the past average inflation rate. In contrast, under the RAIT, the central bank communicates its intention to overshoot the target if the balance of inflation risks – estimated through models or policymakers' judgment – is tilted to the downside.

Under the FAIT, the central bank decides on its temporary inflation overshoot by accounting for past deviations of inflation from the 2% medium-run inflation target. Specifically, we operationalize the FAIT target ($\bar{\pi}_t^{\text{FAIT}}$) by evaluating its effects on the monetary policy stance:

$$\hat{\pi}_t^{\text{FAIT}} = \rho_F \hat{\pi}_{t-1}^{\text{FAIT}} - (1-\rho_F)\left(\pi_t^{\text{Data}} - \bar{\pi}\right), \tag{27}$$

where π_t^{Data} stands for the inflation rate observed in the data and $\bar{\pi}$ represents the 50 basis point

[25] The FAIT adopted by the Federal Reserve in August 2020 was asymmetric, as it did not account for the possibility of undershooting the FOMC's inflation objective. In this analysis, we consider a symmetric FAIT framework, as we examine an environment where inflation risks may be unbalanced either to the upside or the downside.

inflation target, corresponding to 2% annual inflation. Thus, the elements within the round brackets on the right-hand side capture the current miss of inflation relative to the 2% target. A positive value of $\hat{\pi}_t^{\text{FAIT}}$ indicates that, under the FAIT, the central bank would aim for an inflation overshoot to counteract the prevailing negative inflationary pressures observed in past years. The parameter $\rho_F \in (0,1)$ represents the Central Bank's memory when assessing deviations of inflation from the target.[26] The monetary policy rule under the FAIT is given by:

$$\hat{r}_t = \rho r_{t-1} + (1-\rho)\left[r_x \hat{x}_t + r_{\Delta x}\Delta\hat{x}_t + r_\pi\left(\hat{\pi}_t - \hat{\pi}_t^{\text{FAIT}}\right)\right] + \varepsilon_t^r. \tag{28}$$

In Figure 8, we compare the FAIT parameter, which captures the average of past inflation misses, $\hat{\pi}_t^{\text{FAIT}}$, with the RAIT parameter, $\bar{\pi}_t^{\text{RAIT}}$, which reflects the average revisions to the inflation target over a J-period horizon, induced by forward guidance provided at time t.

$$\begin{aligned}\bar{\pi}_t^{\text{RAIT}} &= \frac{1}{J}\sum_{j=1}^{J} E_t - E_{t-1}\left(\hat{\pi}_{t+j}^{\text{RAIT}}\right) \\ &= \frac{1}{J}\sum_{j=1}^{J} \hat{\pi}_{t+j|t}^{\star}\end{aligned}$$

The FAIT parameter in the plot was largely positive during the past decade when U.S. inflation consistently ran below target. Under the FAIT, the central bank would have been required to communicate its intention to create an inflation overshoot. In contrast, the predictions of the RAIT are very different. The RAIT would not have required an inflation overshoot to achieve this goal.

The reason lies in the RAIT's focus on the effects of the balance of inflation risks on expectations – captured by the wedge shown in Equation (19). As indicated in that equation, the effects remain small if the volatility, represented by the scale parameter σ, is low. When the predictive distribution of inflation is not highly diffuse (low σ_t), the impact of its shape – ϱ_t – on inflation expectations is muted. During the past decade, although the predictive distribution was negatively skewed due to

[26] A small ρ_F indicates short memory, leading the Central Bank to quickly attempt to correct recent deviations, resulting in high variability in FAIT. Conversely, a large ρ_F results in a smoother but slower-to-adjust target.

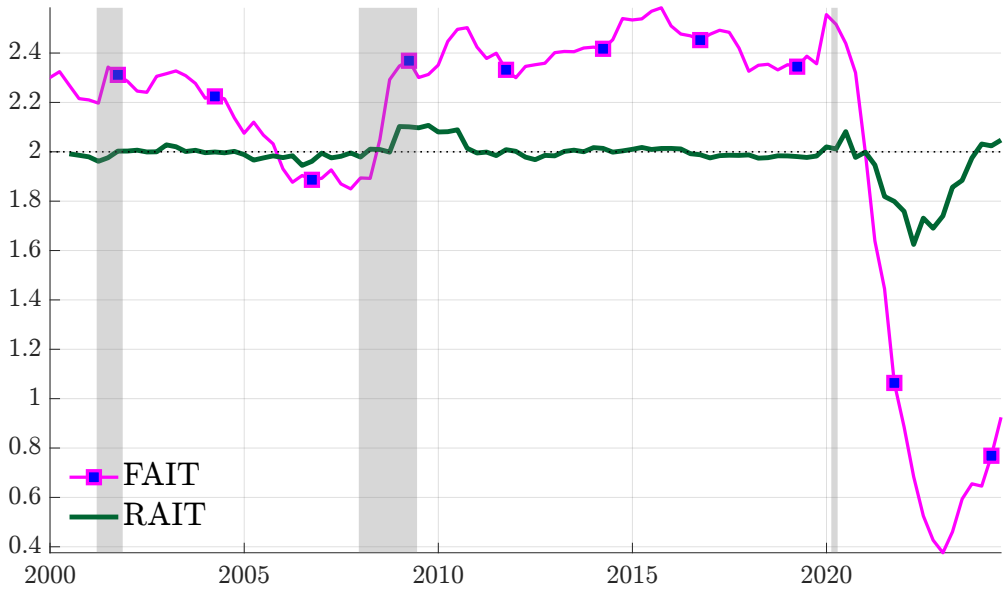


Figure 8: FAIT vs. RAIT

Note: The figure reports the (implicit) temporary inflation target implied by the FAIT (the pink solid line with blue circle markers, $2\% + \hat{\pi}_t^{\text{FAIT}}$) and the RAIT (the green line, $2\% + \hat{\pi}_t^{\text{RAIT}}$). The former is calculated assuming $\rho_F = 0.90$. Gray shaded areas represent NBER recessions.

continuous inflation misses, inflation volatility reached historical lows. As a result, the mean-mode wedge in Equation (19) – ψ_t – was small.

Consequently, the RAIT would have recommended maintaining the Federal Reserve's symmetric framework in the past decade, without requiring any adjustments to address negative skewness in inflation outcomes. This is always the case when inflation volatility is low.

The FAIT parameter lags behind the RAIT parameter during the inflationary surge of 2021, reducing the proactiveness of monetary policy. The FAIT is more backward-looking and less responsive to evolving stochastic conditions affecting future inflation dynamics than the RAIT. When the distribution of inflation outcomes shifts rapidly, the FAIT is slower than the RAIT in providing the central bank with the appropriate policy stance.

At the end of the sample period, the FAIT entirely misses the quick re-balancing of inflation risks. In contrast, the RAIT, with its greater adaptability, would have enabled the central bank to promptly communicate a shift in the balance of inflation risks, justifying monetary easing as early as the beginning of 2023. This more favorable balance of risks was detected in real time by the econometric model presented in Section 4.

The RAIT emphasizes communicating the evolution of the balance of inflation risks to provide

clarity on future policy decisions. Implementing the RAIT requires access to multiple models to reliably assess changes in the balance of inflation risks. Additionally, our econometric model can be enhanced with predictors such as fiscal policy measures, labor market prices, geopolitical indices, and commodity price indices to construct scenario analyses that further support forward-looking communications by a central bank adopting the RAIT.

De Polis et al. (2023) illustrate how incorporating macroeconomic and financial predictors into the forecasting model enhances its ability to predict changes in the balance of inflation risks accurately, thereby improving its practical utility for the RAIT implementation.

6 Concluding remarks

This paper explores the role of time-varying inflation skewness in shaping monetary policy. Using a time-series model to estimate the predictive distribution of U.S. core PCE inflation, we demonstrate that inflation skewness fluctuates significantly over time, affecting the balance of risks and enhancing forecasting accuracy during periods of heightened volatility.

Time-varying skewness introduces a wedge between expected and modal inflation, necessitating optimal monetary policy to "lean against" these risks by adjusting the modal scenario in the opposite direction of the balance of risks.

We assess Risk-Adjusted Inflation Targeting (RAIT), a strategy that incorporates real-time assessments of the balance of inflation risks into central bank communications, allowing for temporary overshoots or undershoots of the inflation target. Unlike FAIT, which advocates inflation overshooting in response to the past decade of persistently low price dynamics, the RAIT would have maintained a symmetric approach to inflation stabilization. Furthermore, during the recent inflation surge, the RAIT would have recommended earlier tightening and a more rapid initiation of policy normalization.

This study underscores the importance of integrating real-time assessments of the balance of inflation risks into policy frameworks, enabling central banks to more effectively achieve macroeconomic stability. Future research could refine RAIT's implementation and explore alternative communication strategies.

References

Adams, P. A., T. Adrian, N. Boyarchenko, and D. Giannone (2021): "Forecasting macroeconomic risks," *International Journal of Forecasting*, 37, 1173–1191.

Adrian, T., N. Boyarchenko, and D. Giannone (2019): "Vulnerable Growth," *American Economic Review*, 109, 1263–89.

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.

Azzalini, A. and A. Capitanio (2014): *The Skew-Normal and Related Families*, IMS monographs, Cambridge University Press.

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

Bianchi, F., L. Melosi, and G. Nicolò (2024): "Is There a Forward Guidance Puzzle?" BSE mimeo, Barcelona.

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

Blanchard, O. J. and B. S. Bernanke (2023): "What Caused the US Pandemic-Era Inflation?" NBER Working Papers 31417, National Bureau of Economic Research, Inc.

Blasques, F., S. J. Koopman, and A. Lucas (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.

Chahrour, R. and K. Jurado (2018): "News or Noise? The Missing Link," *American Economic Review*, 108, 1702–36.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005): "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113, 1–45.

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

Clarida, R. H. (2022): "The Federal Reserve's New Framework: Context and Consequences," Finance and Economics Discussion Series 2022-001, Board of Governors of the Federal Reserve System (U.S.).

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

Cogley, T. and T. J. Sargent (2005): "Drifts and volatilities: monetary policies and outcomes in the post WWII US," *Review of Economic Dynamics*, 8, 262–302.

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.

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. (2024): "Testing for Conditional Skewness with Epsilon-Skew-t Distributions." .

De Polis, A., L. Melosi, and I. Petrella (2023): "The ever-changing challenges to price stability," in *12th European Central Bank Conference on Forecasting Techniques*.

Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2023): "The Forward Guidance Puzzle," *Journal of Political Economy Macroeconomics*, 1, 43–79.

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.

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 F. J. Ruge-Murcia (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.

Fechner, G. T. (1897): *Kollektivmasslehre*, Engelmann.

Fernández, C. and M. F. Steel (1998): "On Bayesian modeling of fat tails and skewness," *Journal of the American Statistical Association*, 93, 359–371.

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

Gertler, M., N. Kiyotaki, and A. Prestipino (2020): "A macroeconomic model with financial panics," *The Review of Economic Studies*, 87, 240–288.

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. 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?" CEPR Discussion Papers 17224, C.E.P.R. Discussion Papers.

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," Working Paper Series 2600, European Central Bank.

Le Bihan, H., D. Leiva-León, and M. Pacce (2024): "Underlying inflation and asymmetric risks," *Review of Economics and Statistics*, 1–45.

López-Salido, D. and F. Loria (2024): "Inflation at risk," *Journal of Monetary Economics*, 145.

Maliar, L. and J. B. Taylor (2019): "Forward Guidance: Is It Useful Away from the Lower Bound?" NBER Working Papers 26053, National Bureau of Economic Research, Inc.

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.

Sims, C. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

Smets, F. and R. Wouters (2007): "Shocks and frictions in US business cycles: A Bayesian DSGE approach," *American economic review*, 97, 586–606.

Stock, J. H. and M. W. Watson (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. O. (2003): "What is wrong with Taylor rules? Using judgment in monetary policy through targeting rules," *Journal of Economic Literature*, 41, 426–477.

Woodford, M. (2003): *Interest and prices*, Princeton University Press Princeton.

A Proofs

A.1 Proof of Proposition 1

Proof. Let $\mathbf{S}^\star$ and $\mathbf{Z}^\star$ define the steady-state of the model. The first-order Taylor expansion of Equation (1) around the steady state reads

$$f(\mathbf{Z}_{t+1}, \mathbf{Z}_t, \mathbf{S}_t) \approx f_{Z_{t+1}}(\mathbf{S}^\star, \mathbf{Z}^\star)(\mathbf{Z}_{t+1} - \mathbf{Z}^\star) + f_{Z_t}(\mathbf{S}^\star, \mathbf{Z}^\star) + (\mathbf{Z}_t - \mathbf{Z}^\star) + f_{Z_{t-1}}(\mathbf{S}^\star, \mathbf{Z}^\star)(\mathbf{Z}_{t-1} - \mathbf{Z}^\star) + f_{S_t}(\mathbf{S}^\star, \mathbf{Z}^\star)(\mathbf{S}_t - \mathbf{S}^\star).$$

Defining log-deviations from steady-states with lower case letters and normalizing $\mathbf{S}^\star = \mathbf{1}$ delivers a linear solution of the form in (1) where $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}$ and $\mathbf{B}^a$ are independent of the skewness of the shocks in $\mathbf{s}_t^a$. □

A.2 Proof of Proposition 2

Proof. The solution of a linear rational model of the form in eq. Equation (2) can be written as

$$\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\mathbf{s}_t^s \; s_t^a\right]' + \mathbf{\Theta}_y \sum_{j=1}^{\infty} \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \left[E_t \mathbf{s}_{t+j}^s \; E_t s_{t+j}^a\right]'.$$

where the matrices $\mathbf{\Theta}_0$, $\mathbf{\Theta}_1$, $\mathbf{\Theta}_y$, and $\mathbf{\Theta}_z$ are functions of $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}_s$ and b_a (see, e.g., Sims, 2002, Sec. 4). In a fully symmetric setting, where $\varrho_t = 0, \forall t$, the above expression simplifies to $\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\mathbf{s}_t^s \; s_t^a\right]'$. When shock distributions are skewed, on the other hand, expectations are a function of the asymmetry of the distribution.

For the general class of skew-distributions we consider (Arellano-Valle et al., 2005), the expected value of a generic skew-variate $X \sim \mathcal{F}(\mu, \sigma, \varrho)$ takes the form $EX = \mu + \kappa\sigma\varrho$, where μ represents the mode of the distribution, σ captures the dispersion, ϱ summarizes the asymmetry of the distribution about the mode, κ is a constant that depends on the choice of $\mathcal{F}$. Therefor, for

$s_t^a \sim \mathcal{F}(0, \sigma, \varrho_t)$,

$$E_t s_{t+j}^a = k\sigma^2 E_t \varrho_{t+j},$$

such that, when risk evolves over time, shifts in expectations of future asymmetry lead agents to revise their decisions today. Therefore,

$$\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\mathbf{s}_t^s \; s_t^a\right]' + \mathbf{\Theta}_y \sum_{j=1}^{\infty} \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \mathbf{\Xi} k\sigma^2 E_t \varrho_{t+j}.$$

where $\mathbf{\Xi}$ is a section vector. Defining revisions in the asymmetry of risk j periods in advance as $\varrho_t^j = E_t \varrho_{t+j} - E_{t-1}\varrho_{t+j}$, such that

$$E_t \mathbf{z}_{t+h} - E_{t-1}\mathbf{z}_{t+h} = \mathbf{\Theta}_1^h \mathbf{\Theta}_y \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \mathbf{\Xi} k\sigma^2 \varrho_t^j;$$

setting $\mathbf{M}_{(j,h)} = \mathbf{\Theta}_1^h \mathbf{\Theta}_y \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \mathbf{\Xi}$ we recover Equation (3). □

A.3 Proof of Proposition 3

Proof. The solution of a linear rational model of the form in eq. Equation (2) can be written as

$$\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \widetilde{\mathbf{s}}_t + \mathbf{\Theta}_y \sum_{j=1}^{\infty} \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \widetilde{E}_t \widetilde{\mathbf{s}}_{t+j}$$

where the matrices $\mathbf{\Theta}_0$, $\mathbf{\Theta}_1$, $\mathbf{\Theta}_y$, and $\mathbf{\Theta}_z$ are functions of $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$ and $\mathbf{B}$ (see, e.g., Sims, 2002, Sec. 4). Therefore, given the definition of *beliefs expectations* in Eq. (5), the rational expectations solution of the system can be written as

$$\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \widetilde{\mathbf{s}}_t + \mathbf{\Theta}_y \sum_{j=1}^{\infty} \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \mathbf{c} E_t \psi_{t+j}$$

Following Chahrour and Jurado (2018), we define revisions in the asymmetry of risk j periods in advance ad $\psi_t^j = E_t\psi_{t+j} - E_{t-1}\psi_{t+j}$, such that

$$\mathbf{z}_t - E_{t-1}\mathbf{z}_t = \mathbf{\Theta}_y\mathbf{\Theta}_f^{j-1}\mathbf{\Theta}_z\mathbf{c}\psi_t^j.$$

By setting

$$\mathbf{c}\psi_t^j = \mathbf{\Xi}k\sigma^2\varrho_t^j,$$

the impact of revisions about the asymmetry of shocks' distributions is equivalent to shifts in beliefs about future realizations of the shocks. For $\widetilde{\mathbf{s}}_t = [\mathbf{s}_t^s \; s_t^a]'$ and $\mathbf{B} = [\mathbf{B}_s \; b_a]$, the responses of the two models are observationally equivalent up the first order. □

B Solving the New Keynesian model with mark-up shock

In this appendix, we provide detailed derivations of the solutions for the optimal monetary problem under asymmetric risks, outlined in Section 3. This problem is derived in the basic New Keynesian model presented in Chapter 3 of Galí (2008). We briefly outline the housholds and firms problems, and then derive the optimal monetary policy rules.

B.1 The baseline New Keynesian model

The model features a continuum of households, firms, and a central bank. Households work, consume, and save. Firms hire labor in a competitive market, produce goods, and set prices, with the market for goods being monopolistically competitive. Price setting is subject to the Calvo lottery. Monetary policy is optimal under the assumption that the central bank can commit to future actions. The optimal monetary problem faced by the central bank is discussed in the next section of this appendix.

Households. An infinitely-lived representative agent seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t U(C_t, N_t),$$

subject to the budget constraint

$$P_t C_t + Q_t B_t \leq B_{t-1} + W_t N_t + T_t, \text{ for } t = 1, 2, \ldots,$$

where C_t and P_t represent the quantity and price of the aggregate consumption index, Q_t is the price of a nominal bond purchased in quantity B_t, N_t denotes the hours worked, remunerated at the nominal wage W_t; finally, T_t is a lump-sum transfer. There is a CES technology aggregating differentiated goods produced by firms into the consumption bundle, C_t, consumed by the representative household. The parameter β is the deterministic discount factor of households.

Assuming the utility function to be

$$U(C_t, N_t) = \frac{C_t^{1-\sigma}}{1-\sigma} - \frac{N_t^{1-\varphi}}{1-\varphi},$$

where σ is the parameter of relative risk aversion and φ is the inverse Frisch elasticity of labor supply. Households' demand for a differentiated good is downward sloping with respect to the price of the good.

Firms. A continuum of monopolistic firms produce differentiated goods using the same technology

$$Y_t(i) = A_t N_t(i)^{1-\alpha},$$

where A_t denotes the technology level, and α is the production function scale parameter. All firms face the same downward-sloping demand schedule obtained from solving the household's problem of choosing goods of different variety.

We assume that, every period, a firm can re-optimize its prices with probability $1 - \theta$. This

friction gives rise to the following dynamics for the aggregate price level, expressed in log-terms,

$$\pi_t = (1-\theta)(p_t^* - p_{t-1}),$$

where p_t^* is the re-optimized price level and p_{t-1} is the average price level for the non-re-optimizing firms.

B.2 The optimal monetary policy problem

In this model, the output gap and inflation in deviations from steady state are governed by a standard IS and a Phillips curve

$$\sigma \hat{x}_t = \sigma E_t \hat{x}_{t+1} - \hat{i}_t + E_t \hat{\pi}_{t+1}, \tag{B1}$$

$$\hat{\pi}_t = \beta E_t \pi_{t+1} + \kappa \hat{x}_t + u_t, \tag{B2}$$

where $i_t = -\log Q_t$ is the short-term nominal rate and κ is the slope of the Phillips curve. The process driving price markups u_t can be expressed as follows:

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u \tag{B3}$$

where $u_t \sim iidF(0, \sigma_{u,t})$ is a shock to the price markup, with $F(0, \sigma_{u,t})$ being a general (unimodal) symmetric distribution, usually assumed to be the Normal distribution.

We assume that the central bank commits, with full credibility, to a policy plan consistent with a quadratic objective function in inflation deviations, $\hat{\pi}_t$ and the output gap, $\hat{x}_t$. Therefore, optimal monetary policy consists in choosing the state-contingent $\{\hat{\pi}_t, \hat{x}_t\}_{t=0}^{\infty}$ that minimizes

$$\frac{1}{2} E_0 \sum_{t=0}^{\infty} \beta^t \left(\hat{\pi}_t^2 + \alpha_x \hat{x}_t^2 \right),$$

subject to the sequence of constraints imposed by the Phillips curve above. Casting the problem

into its Lagrangian form

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2} \left(\hat{\pi}_t^2 + \alpha_x \hat{x}_t^2 \right) + \gamma_t \left(\hat{\pi}_t - \kappa \hat{x}_t - \beta \hat{\pi}_{t+1} \right) \right],$$

and differentiating with respect to $\hat{x}_t$ and $\hat{\pi}_t$ yields the optimality conditions

$$\alpha_x \hat{x}_t - \kappa \gamma_t = 0$$

$$\hat{\pi}_t + \gamma_t - \gamma_{t-1} = 0$$

that must hold for $t = 0, 1, 2, ...$; We set $\gamma_{-1} = 0$ in that Phillips curve constraint is not binding in period -1 for the central bank choosing the optimal plan in period 0.

Standard manipulations yield the following optimality conditions,

$$\hat{x}_0 = -\frac{\kappa}{\alpha_x} \hat{\pi}_0,$$
$$\hat{x}_t = \hat{x}_{t-1} - \frac{\kappa}{\alpha_x} \hat{\pi}_t, \ \forall t.$$

Define $\bar{p}_t = p_t - p_{-1}$ as the inflation rate over period 0 through period t, where p_t denotes the log of the price level at time t. We can now write the optimal targeting rule under commitment as

$$\hat{x}_t = -\frac{\kappa}{\alpha_x} \bar{p}_t, \tag{B4}$$

such that the optimizing central bank keeps output below or above the efficient level in proportion to the deviations of the price level from its implicit target. Plugging Equation (B4) into Equation (B2) we can recast the Phillips curve in as

$$\bar{p}_t = a \bar{p}_{t-1} + a\beta E_t \bar{p}_{t+1} + a u_t \tag{B5}$$

with $a \equiv \frac{\alpha_x}{\alpha_x(1+\beta)+\kappa^2}$.

B.2.1 The symmetric case

The stationary solution to Equation (B5) can be obtained by using the method of undetermined coefficients by conjecturing a solution of the form

$$\bar{p}_t = \eta \bar{p}_{t-1} + \lambda u_t, \tag{B6}$$

such that the expected value of the next period's price level is

$$E_t \bar{p}_{t+1} = \eta \bar{p}_t + \lambda \rho_u u_t. \tag{B7}$$

Substituting the expectations implied by the stationary solution yields

$$\begin{aligned} \bar{p}_t &= a\bar{p}_{t-1} + a\beta\left(\eta \bar{p}_t + \lambda \rho_u u_t\right) + a u_t \\ \bar{p}_t &= \underbrace{\frac{a}{1 - a\beta\eta}}_{\eta} \bar{p}_{t-1} + \underbrace{a\frac{1 + \beta\lambda\rho_u}{1 - a\beta\eta}}_{\lambda} u_t \end{aligned}$$

Solving for η and λ, we obtain

$$\eta = \frac{1 - \sqrt{1 - 4\beta a^2}}{2a\beta}, \tag{B8}$$

$$\lambda = \frac{a}{1 - a\beta(\eta + \rho_u)}, \tag{B9}$$

and we can express the equilibrium process for the output gap as

$$\hat{x}_t = \eta \hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\lambda u_t, \tag{B10}$$

for $t = 1, 2, ...$, with $\hat{x}_0 = -\frac{\kappa}{\alpha_x}\lambda u_0$.

Implementation. We assume that the price mark-up shocks are iid ($\rho_u = 0$) throughout this section.

The IS equation reads

$$\sigma \hat{x}_t = \sigma E_t \hat{x}_{t+1} - \hat{i}_t + E_t \hat{\pi}_{t+1},$$

and let us express this in terms of the price level

$$\sigma \hat{x}_t = \sigma E_t \hat{x}_{t+1} - \hat{i}_t + E_t \bar{p}_{t+1} - \bar{p}_t.$$

The optimality condition in Equation (B4) allows us to write the IS curve as

$$\left[1 - \sigma \frac{\kappa}{\alpha_x}\right] \bar{p}_t = \left[1 - \sigma \frac{\kappa}{\alpha_x}\right] E_t \bar{p}_{t+1} - \hat{i}_t.$$

Now, substituting the implied expectation derived in Equation (B7), and recalling that $\rho_u = 0$, we obtain the following *optimal monetary rule*:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma \frac{\kappa}{\alpha_x}\right] \bar{p}_t. \tag{B11}$$

To ensure determinacy, we need to replace $\bar{p}_t$ with its law of motion. Starting from Equation (B6),[27]

$$\bar{p}_t = \sum_{k=0}^{t} \eta^{k+1} u_{t-k};$$

substituting into the optimal rule and adding a subtracting $\phi_p \bar{p}_t$ yields

$$\hat{i}_t = -\left(\phi_p + (1-\eta)\left[1 - \sigma \frac{\kappa}{\alpha_x}\right]\right) \sum_{k=0}^{t} \eta^{k+1} u_{t-k} + \phi_p \bar{p}_t.$$

Provided that $\phi_p > 0$, the system of equations comprising the IS equation, the Phillips curve, and the optimal monetary rule in the above specification admits a unique stable rational expectations equilibrium (see, e.g., Galí, 2008).

[27] It can be shown that $\lambda = \eta$ when $\rho_u = 0$.

B.2.2 The asymmetric case

Let us now assume that the stochastic process driving the markup shock (Equation (B3)) is no longer symmetric. That is, let the shocks be $\tilde{\varepsilon}_t^u \sim iid\mathcal{F}(0, \sigma_{u,t}, \varrho_{u,t})$, where $\mathcal{F}(0, \sigma_{u,t}, \varrho_{u,t})$ represents a general (unimodal) distribution, which features asymmetry about 0 when $\varrho_{u,t} \neq 0$.

We now propose to represent the asymmetric shocks $\tilde{\varepsilon}_t^u$ as the linear combination of the symmetric shock ε_t^u and two news shocks,

$$u_t = \rho_u u_{t-1} + \varepsilon_t^u + \left(\psi_t^0 + \psi_{t-1}^1\right), \tag{B12}$$

where ψ_t^j is a shock known in period t and that will have affect in period $t+j$. For this representation to hold we impose the restriction $\psi_{t+1}^0 = -\psi_t^1$, where ψ_t^1 represents a surprise shock which agents are not aware of; we call this the *beliefs representation of asymmetric risk*. Here, we are considering the case where the distribution of the markup process is expected by economic agents to be skewed by just one period (i.e., in period t agents expect the distribution to be back to symmetric – $\varrho_{u,t+j} = 0$ for $j > 1$). However, this can be easily generalized to multiple periods.

Notice that now the expected value of the shock is potentially nonzero, for $E_t\varepsilon_{t+1}^u = \psi_t^1$, where $\psi_t^1 \sim \mathcal{N}(0, \sigma_1^2)$, and the expectation of next period's markup is given by $E_t u_{t+1} = \rho_u u_t + \psi_t^1$. Due to this asymmetry, next period's price level, $E_t\bar{p}_{t+1}$, are potentially distorted. We recompute the price process as above:

$$\bar{p}_t = \eta\bar{p}_{t-1} + \lambda u_t + \zeta\psi_t^1, \tag{B13}$$

and it follows that $E_t\bar{p}_{t+1} = \eta\bar{p}_t + \lambda\left(\rho_u u_t + \psi_t^1\right)$, since $Eu_{t+1} = \rho_u u_t + \psi_t^1$ and $E_t\psi_{t+1}^1 = 0$. Substituting the expectations implied by the stationary solution yields

$$\begin{aligned}
\bar{p}_t &= a\bar{p}_{t-1} + a\beta\left[\eta\bar{p}_t + \lambda\rho_u u_t + \lambda\psi_t^1\right] + au_t, \\
\bar{p}_t &= \underbrace{\frac{a}{1-a\beta\eta}}_{\eta}\bar{p}_{t-1} + \underbrace{a\frac{1+\beta\lambda\rho_u}{1-a\beta\eta}}_{\lambda}u_t + \underbrace{\frac{a\beta\lambda}{1-a\beta\eta}}_{\zeta}\psi_t^1.
\end{aligned}$$

Solving for η, λ and ζ we obtain

$$\eta = \frac{1 - \sqrt{1 - 4\beta a^2}}{2a\beta},$$
$$\lambda = \frac{a}{1 - a\beta(\eta + \rho_u)},$$
$$\zeta = \frac{a\beta\lambda}{1 - a\beta\eta}.$$

It follows that the equilibrium process for the output gap is

$$\hat{x}_t = \eta \hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\left[\lambda u_t + \zeta \psi_t^1\right], \tag{B14}$$

and $\hat{x}_0 = -\frac{\kappa}{\alpha_x}\left[\lambda u_0 + \zeta \psi_t^1\right]$.

Implementation. Again, let us assume $\rho_u = 0$ and let us rewrite the IS equation in terms of the price level under asymmetry and substitute the optimality condition such that

$$\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t = \left[1 - \sigma\frac{\kappa}{\alpha_x}\right] E_t \bar{p}_{t+1} - \hat{i}_t.$$

Substituting the expectation for the price level under asymmetry, and recalling that $\rho_u = 0$, we obtain the *optimal monetary rule under asymmetry*:

$$\hat{i}_t = -(1 - \eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\psi_t^1,$$

where the first term in the right-hand side is the same as in Equation (B11), and the last term captures how the the central bank needs to adjust the policy rate to take into the expectation bias due to asymmetric risks.

As before, we ensure determinacy by substituting into the rule the law of motion for $\bar{p}_t$,

$$\bar{p}_t = \sum_{k=0}^{t} \eta^{k+1}\left(u_{t-k} + \psi_{t-k}^1\right),$$

to obtain

$$\hat{i}_t = -\left(\phi_p + (1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\right)\sum_{k=0}^{t}\eta^{k+1}\left(u_{t-k} + \psi^1_{t-k}\right) + \phi_p\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\psi^1_t.$$

B.2.3 Asymmetric case with an unwitting central bank

We now consider the case in which the distribution of price markups are skewed but the central bank does not take this into account (or it does not know) and adopts the optimal policy under the incorrect assumption about shocks distributions. Markup shocks' asymmetry is captured by agents receiving surprises ψ^1_t, tilting their expectations about future realizations of the shocks away from the central scenario ($\rho_u u_{t-1}$) in every period t. We compare this scenario against the case of optimal monetary policy under symmetry.[28]

We start from the optimality condition between output gap and inflation

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}^s_t, \tag{B15}$$

where p^s_t is the price level under full symmetry ($\psi^0_t = 0$ in very period) defined

$$\bar{p}^s_t = \eta\bar{p}_{t-1} + \lambda u_t; \tag{B16}$$

note that this is not exactly the same price level as in fully symmetric case ($\bar{p}_t$), in that past suboptimal price levels, $\bar{p}_{t-1}$, are a given for the optimizing central bank.

We write the Phillips curve in terms of the price level,

$$(1+\beta)\bar{p}_t = \bar{p}_{t-1} + \kappa\hat{x}_t + \beta E_t\bar{p}_{t+1} + u_t,$$

and we conjecture that the difference between the two prices under symmetry that the central

[28] This exercise can be interpreted as the asymmetric case under the counterfactual assumption that the central bank does not try to lean against the asymmetry in the dynamics of the price level.

bank takes into account when solving its optimal problem is

$$\bar{p}_t - \bar{p}_t^s = \tau \psi_t^1.$$

By plugging this equation into the optimality condition in Equation (B15), and substituting into the Phillips curve expressed in terms of price level yields

$$\bar{p}_t = a\bar{p}_{t-1} + a\beta E_t \bar{p}_{t+1} + b\tau\psi_t^1 + au_t \tag{B17}$$

with $a \equiv \frac{\alpha_x}{\alpha_x(1+\beta)+\kappa^2}$ and $b \equiv \frac{\kappa^2}{\alpha_x(1+\beta)+\kappa^2}$.

As before, we conjecture the stationary solution, retrieve expectations of next period's price level and plug these into Equation (B17) to obtain

$$\bar{p}_t = \underbrace{\frac{a}{1-a\beta\eta}}_{\eta}\bar{p}_{t-1} + \underbrace{a\frac{1+\beta\lambda\rho_u}{1-a\beta\eta}}_{\lambda}u_t + \underbrace{\frac{a\beta\lambda+b\tau}{1-a\beta\eta}}_{\zeta}\psi_t^1.$$

Solving for the coefficients we obtain

$$\begin{aligned}
\eta &= \frac{1-\sqrt{1-4\beta a^2}}{2a\beta} \\
\lambda &= \frac{a}{1-a\beta(\eta+\rho_u)} \\
\zeta &= \frac{a\beta\lambda+b\tau}{1-a\beta\eta}
\end{aligned}$$

We recover τ by taking the difference between the price level, $\bar{p}$, and the price level under the symmetric policy – Equation (B16). Specifically,

$$\bar{p}_t - \bar{p}_t^s = \underbrace{\frac{a\beta\lambda+b\tau}{1-a\beta\eta}}_{\tau}\psi_t^1,$$

which leads to

$$\zeta = \tau = \frac{a\beta\lambda}{1 - a\beta\eta - b}. \tag{B18}$$

The equilibrium process for the output gap is

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^s, \tag{B19}$$

for $t = 1, 2, ...$ It should be noted that the central bank that overlooks the importance of the imbalance of risks will end up setting an output gap as a function of $\bar{p}_t^s$. However, this price level is not achievable by the central bank because of the balance of inflation risks, resulting in a suboptimal output gap and in a price level that is different from that targeted by the central bank, $\bar{p}_t^s$. Therefore,

$$\hat{x}_t \neq \frac{\kappa}{\alpha_x}\bar{p}_t. \tag{B20}$$

It should be noted that the output gap initially chosen by the central bank is the same as that chosen in the first case. This is because the price level targeted by the central bank, $\bar{p}_t^s$, is exactly the same as the price level target by the symmetric central bank in the symmetric case.[29] However, in the subsequent periods the output gap starts diverging in the two economies as the previous period's price level is different due to the skewness of the shocks distribution.

B.3 Calibration values for the numerical simulation

We follow Galí (2008) and set the following values for the parameters of the model: the elasticity of substitution, σ, to unity, the production function scale parameter, α, is set to one third, the elasticity of substitution among intermediate goods, ε, is equal to 6, and the slope of the Phillips curve, κ, is equal to 0.1275.

[29] We assume that in period $t = -1$, the economy is at the deterministic steady state equilibrium where risks are fully balanced, $\psi_{-i}^1 = 0$, $i = -1, -2, ...$

C Score-driven framework

C.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), $y_t|Y_{t-1} \sim Skt_\nu(\mu_t, \sigma_t^2, \varrho_t)$, the log-likelihood takes the form

$$\ell_t(r_t|\theta, \mathcal{F}_{t-1}) = \log \mathcal{C}(\nu) - \frac{1}{2}\log \sigma_t^2 - \frac{1+\nu}{2}\log\left[1 + \frac{\varepsilon_t^2}{\nu(1+sgn(\varepsilon_t)\varrho_t)^2\sigma_t^2}\right], \tag{C1}$$
$$\log \mathcal{C}(\nu) = \log\Gamma\left(\frac{\nu+1}{2}\right) - \log\Gamma\left(\frac{\nu}{2}\right) - \frac{1}{2}\log\nu - \frac{1}{2}\log\pi,$$

where $\Gamma(\cdot)$ is the Gamma function, $sgn(\cdot)$ is the sign function, and $\nu > 3$ are the degrees of freedom. Differentiating (C1) with respect to location, scale and asymmetry we obtain the gradient vector $\nabla_t = \left[\frac{\partial \ell_t}{\partial \mu}, \frac{\partial \ell_t}{\partial \sigma_t^2}, \frac{\partial \ell_t}{\partial \varrho_t}\right]'$. Recall that $\varepsilon_t = y_t - \mu_t$, $\zeta_t = \frac{\varepsilon_t}{\sigma_t}$ and let

$$f(\mu_t, \sigma_t^2, \varrho_t) = 1 + \frac{\varepsilon_t^2}{\nu(1+sgn(\varepsilon_t)\varrho_t)^2\sigma_t^2} = \frac{\nu(1+sgn(\varepsilon_t)\varrho_t)^2\sigma_t^2 + \varepsilon_t^2}{\nu(1+sgn(\varepsilon_t)\varrho_t)^2\sigma_t^2}$$

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

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

Proof. Define

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

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

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

it follows:

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

. □

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

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

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

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

. □

The score with respect to the shape parameter reads as

$$\frac{\partial \ell_t}{\partial \varrho_t} = \frac{sgn(\varepsilon_t)}{(1+sgn(\varepsilon_t)\varrho_t)}w_t\zeta_t^2.$$

Proof. Define

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

such that $\frac{\partial \ell_t}{\partial \varrho_t} = \frac{\partial g(\varrho_t)}{\partial \sigma_t^2} = a\frac{f'_{\varrho t}}{f(\mu_t, \sigma_t^2, \varrho_t)}$, with $f'_{\varrho_t} = -\frac{2(sgn(\varepsilon_t)+\varrho_t)\varepsilon_t^2}{\nu(1+sgn(\varepsilon_t)\varrho_t)^4\sigma_t^2}$. It follows that:

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

. □

C.2 Scaled scores

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

$$\frac{\partial \ell_t}{\partial \gamma_t} = \frac{\partial \ell_t}{\partial \sigma_t^2}\frac{\partial \sigma_t^2}{\partial \gamma_t}, \quad \frac{\partial \ell_t}{\partial \delta_t} = \frac{\partial \ell_t}{\partial \varrho_t}\frac{\partial \varrho_t}{\partial \delta_t}, \tag{C2}$$

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

$$J_t = \frac{\partial(\mu_t, \sigma_t^2, \varrho_t)}{\partial f_t'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2\sigma_t^2 & 0 \\ 0 & 0 & 1-\varrho_t^2 \end{bmatrix}. \tag{C3}$$

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:

$$\mathcal{I}_t = \mathbb{E}_{t-1}[\nabla_t \nabla_t'] = \begin{bmatrix} \frac{(1+\nu)}{(\nu+3)(1-\varrho_t^2)\sigma_t^2} & 0 & \frac{4(1+\nu)}{\sigma_t(1-\varrho_t^2)(3+\nu)} \\ 0 & \frac{1}{2(3+\nu)\sigma_t^4} & 0 \\ \frac{4(1+\nu)}{\sigma_t(1-\varrho_t^2)(3+\nu)} & 0 & \frac{3(1+\nu)}{(1-\varrho_t^2)(3+\nu)} \end{bmatrix}. \tag{C4}$$

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

$$\mathbf{s}_t = (J_t' \text{diag}(\mathcal{I}_t) J_t)^{-\frac{1}{2}} J_t' \nabla_t = \begin{bmatrix} s_{\mu,t} \\ s_{\gamma,t} \\ s_{\delta,t} \end{bmatrix} = \begin{bmatrix} \sqrt{\frac{(\nu+3)(1-\varrho_t^2)}{(\nu+1)}} w_t \zeta_t \\ \sqrt{\frac{(\nu+3)}{2\nu}} (w_t \zeta_t^2 - 1) \\ sgn(\varepsilon_t) \sqrt{\frac{(\nu+3)(1-sgn(\varepsilon_t)\varrho_t)}{3(\nu+1)(1+sgn(\varepsilon_t)\varrho_t)}} w_t \zeta_t^2 \end{bmatrix}. \tag{C5}$$

Full derivations for the Information matrix are provided in Delle Monache et al. (2024).

C.3 Bayesian estimation

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.

Posterior estimates of the parameters are obtained via simulation by means of an Adaptive Metropolis-Hastings algorithm (Haario et al., 1999). Given that estimated parameters lie in bounded regions of the parameter space, we augment the algorithm with a rejection step to prevent numerical instability due to invalid parameter draws. The algorithm is rather efficient, and a complete chain of 50000 draws can be obtained in less than 2 minutes. Section D in Delle Monache et al. (2024) provides detailed explanation of the algorithm.

D Monte Carlo analysis

We simulate T=250 observations from $Skt_\nu(\mu_t, \sigma_t, \varrho_t)$, for simulated values of the parameters of location, μ_t, scale, σ_t, and asymmetry, ϱ_t. Unless explicitly mentioned, we simulate the parameters independently, and we consider the following cases: no asymmetry, breaks in the asymmetry, fixed asymmetry with location-scale covariance, fixed asymmetry with location-scale covariance with breaks, time-varying asymmetry, and time-varying asymmetry with breaks in the location-scale covariance.

For all cases, we simulate the location and log-scale from first order Gaussian autoregressive processes, with autoregressive parameters equal to 0.9 and 0.99, repsectively, and variances set to 0.05 and 0.025. When we assume correlated innovations for the two parameters, we set this to 0.4. When we impose breaks in this correlation, we assume the relation abruptly shifts to 0.8 after 100 observations, and then falls to -0.4 after additional 50 observations. When time-varying, the asymmetry parameter is simulated from an AR(1) with persistence set to 0.9 and variance 0.025; when only breaks are considered, these occur on the 100th observation, moving from 0 to 0.25, and a sharp fall to -0.25 on the 150th observation.

Define $\delta_t = \log \sigma_t$, $\delta_t = \operatorname{arctanh} \varrho_t$, and $\varepsilon \sim \mathcal{N}(0,1)$, and let $chol()$ define the lower-triangular Choleski factor; here we report a summary of the six DGPs.

DGP1: no asymmetry

$$\begin{bmatrix} \mu_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.99 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \end{bmatrix} + chol\left(\begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix} \right) \varepsilon_t,$$

$$\gamma_t = 0, \ \forall t$$

DGP2: constant asymmetry with breaks

$$\begin{bmatrix} \mu_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.99 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \end{bmatrix} + chol\left(\begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix} \right) \varepsilon_t,$$

$$\gamma_t = \begin{cases} 0 & t \leq 100 \\ 0.25 & 100 < t \leq 150 \\ -0.25 & t < 150 \end{cases}$$

DGP3: no asymmetry and location-scale covariance

$$\begin{bmatrix} \mu_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.99 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \end{bmatrix} + chol\left(\begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \begin{bmatrix} 1 & .4 \\ .4 & 1 \end{bmatrix} \begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \right) \varepsilon_t,$$

$$\gamma_t = 0 \; \forall t$$

DGP4: no asymmetry and location-scale covariance with breaks

$$\begin{bmatrix} \mu_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.99 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \end{bmatrix} + chol\left(\begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \begin{bmatrix} 1 & \rho_t \\ \rho_t & 1 \end{bmatrix} \begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \right) \varepsilon_t,$$

$$\rho_t = \begin{cases} 0.4 & t \leq 100 \\ 0.8 & 100 < t \leq 150 \\ -0.4 & t < 150 \end{cases},$$

$$\gamma_t = 0 \; \forall t$$

DGP5: time-varying asymmetry

$$\begin{bmatrix} \mu_t \\ \delta_t \\ \gamma_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 & 0 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \\ \gamma_{t-1} \end{bmatrix} + chol \left(\begin{bmatrix} 0.05 & 0 & 0 \\ 0 & 0.025 & 0 \\ 0 & 0 & 0.025 \end{bmatrix} \right) \varepsilon_t$$

DGP5: time-varying asymmetry and correlated updates

$$\begin{bmatrix} \mu_t \\ \delta_t \\ \gamma_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 & 0 \\ 0 & 0.99 & 0 \\ 0 & 0 & 0.9 \end{bmatrix} \begin{bmatrix} \mu_{t-1} \\ \delta_{t-1} \\ \gamma_{t-1} \end{bmatrix} + chol \left(\begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \begin{bmatrix} 1 & \rho_t & 0.2 \\ \rho_t & 1 & 0.3 \\ 0.2 & 0.3 & 1 \end{bmatrix} \begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix}^{\frac{1}{2}} \right) \varepsilon_t,$$

$$\rho_t = \begin{cases} 0.4 & t \leq 100 \\ 0.8 & 100 < t \leq 150 \\ -0.4 & t < 150 \end{cases}$$

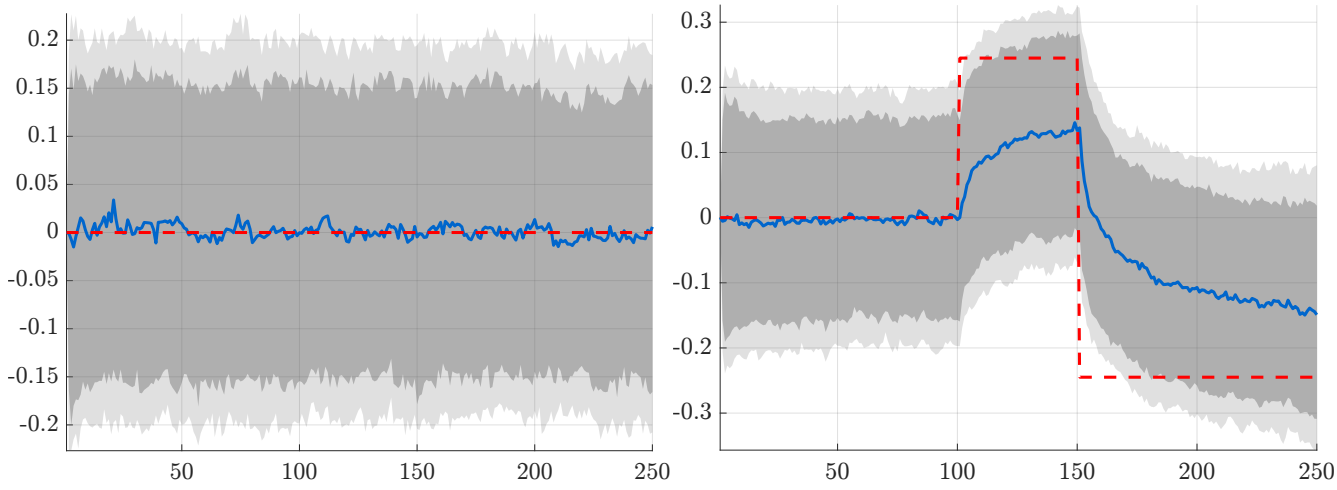
We report the results of this exercise in Figure D1. Specifically, for DGP1 to DGP4 we report in blue the estimated asymmetry, with 68% and 90% credible sets represented by shades of gray, against the simulated parameter, in red. For DGP5 and DGP6 we report the distribution of the difference between the estimated and the simulated asymmetry.

For the first DGP, data are simulated under the assumption of symmetry, with independent, time-varying location and volatility. We show that the model does not pick up any asymmetry when this is not a feature of the data. The second DGP considers the case in which the asymmetry parameter experiences a break from 0 to 0.25 after 100 observations, hence implying positively skewed distributions, and another jump to -0.25 after additional 50 observations; this second jump changes the sign of the skewness. Three comments are in order. First, as for DGP1, no asymmetry is detected when the true value is zero. Second, the parameter reacts promptly to the first jump, despite only 50 observations feature positive skewness. Third, the model quickly detects a turning point in the sing of the asymmetry, turning from positive to negative in less than 20 periods.

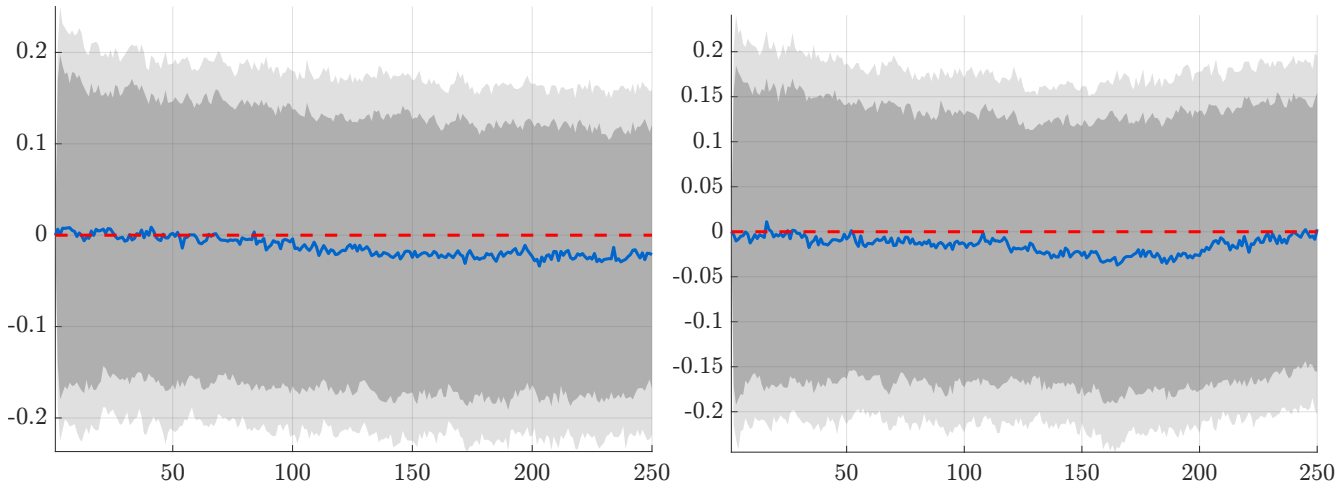
DGP3 and DGP4 are meant to provide reassurances that the model does not mistake correlations between the location and the scale for evidence of asymmetry. In DGP4 we further allow for the correlation to experience breaks, that flip the sing of the covariance between the two parameters. The reported results highlight that the model provides asymmetry estimates that are robust to such features of the data.

Finally, in DGP5 and DGP6 we simulate the asymmetry parameter to vary over time, as the other two parameters. The two DGPs differ in the covariance structure of the parameters: DGP5 assumes independent innovations to the processes, whereas DGP6 assumes a full covariance matrix, with the covariance between location and scale experiencing two breaks, as in DGP4. Once again, we document that our model is successful in detecting the correct sign and dynamics for the asymmetry parameter, even when all the parameters are correlated, and experience instability.

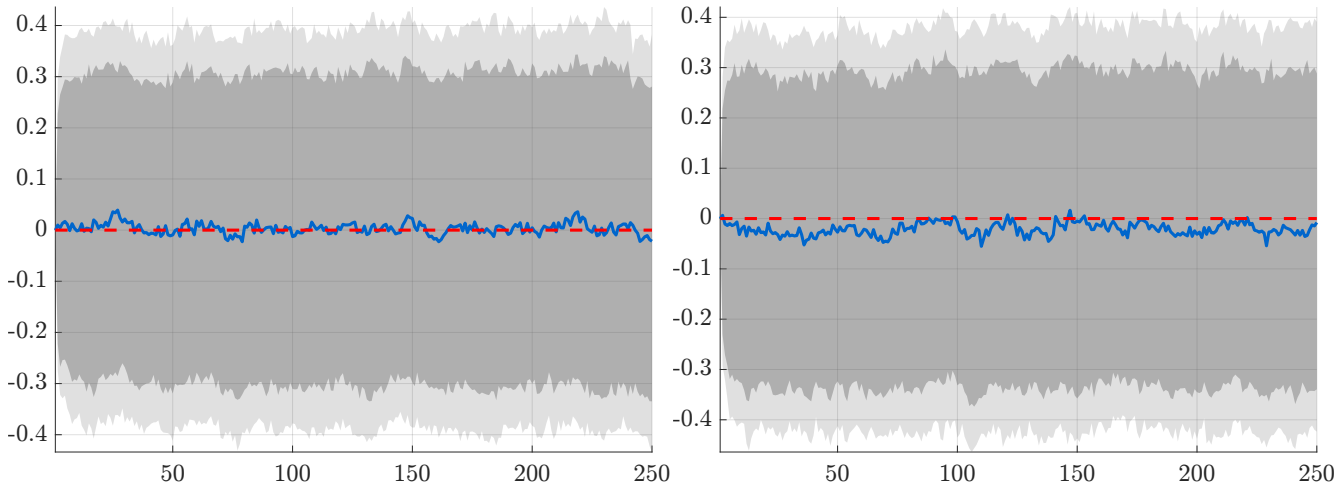
Returning to DGP2, we evaluate the ability of the model to distinguish permanent changes in the parameter against transitory moves. Figure D2 report the estimated long- and short-run components; notice that the two distributions add up to that reported in panel (b) of Figure D1. The model successfully discerns the persistence of the asymmetry in the data whereby it correctly picks up permanent changes. Interestingly, the short-run component shows short periods of increased volatility around the observations where the DGP jumps. This suggests that at first the model interprets new observations as transitory changes in the data, but as more evidence comes through, the long-run component quickly learns the new feature of the data, whereas the short-term component reverts to zero.



(a) DGP1 (b) DGP2



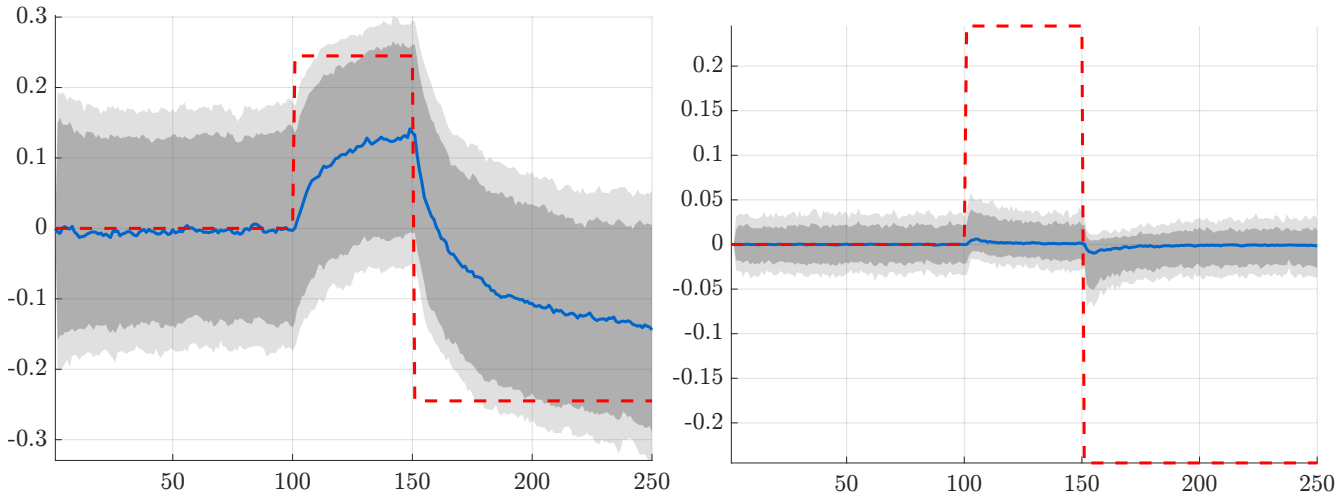
(c) DGP3 (d) DGP4



(e) DGP5 (f) DGP6

Figure D1: Estimated asymmetry

Note: The panels reports the estimated paths for the asymmetry parameters (blue) with the associated 68% and 90% credible sets. The asymmetry under the DGP is reported in red. For DGP 5 and 6 we report deviations of the estimated parameter from the simulated values. We consider T=250 observations for 1000 Monte Carlo replications.



(a) DGP1 **(b)** DGP2

Figure D2: Disentangling permanent changes

Note: The panel report the estimates long- (a) and short- (b) components of the asymmetry parameters estimated under DGP2. Median values are reported in blue, with the associated 68% and 90% credible sets in gray. The asymmetry under the DGP is reported in red. We consider T=250 observations for 1000 Monte Carlo replications.

Table E1: 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***

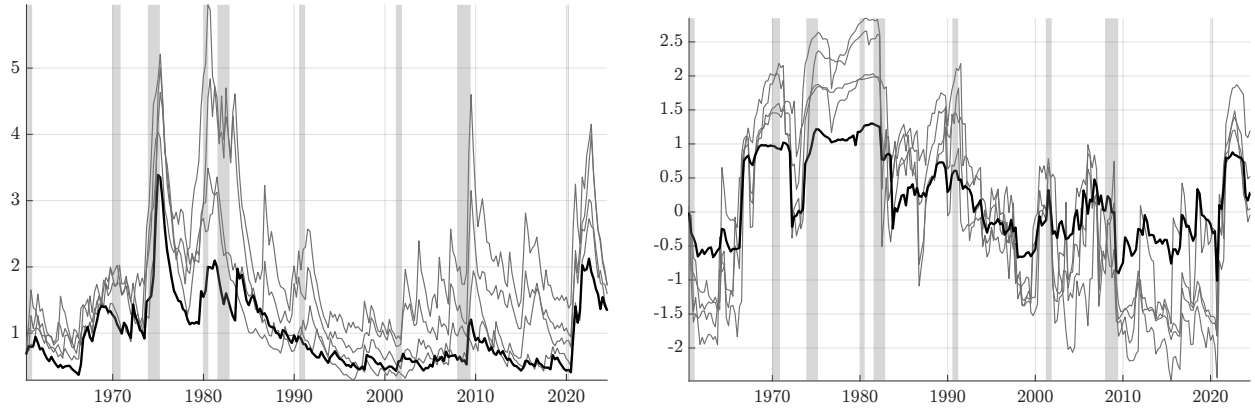
Note: Q is the portmanteau test, Q* is the Ljung-Box extension (with automatic lag selection) and N corresponds to the Nyblom test. Q and Q* are distributed as a χ^2_1, while N is distributed as a Cramer von-Mises distribution with 1 degree of freedom. * p $<$ 10%, ** p $<$ 5%, *** p $<$ 1%.

E Evidence for other inflation measures

In this appendix we report additional results about other policy relevant measures of inflation. The evidence reported in Section 4 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. Specifically, we consider, the GDP deflator, headline PCE and core and headline CPI. All samples go from 1960 Q1 to 2024 Q2.

Table E1 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.

Figure E1 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



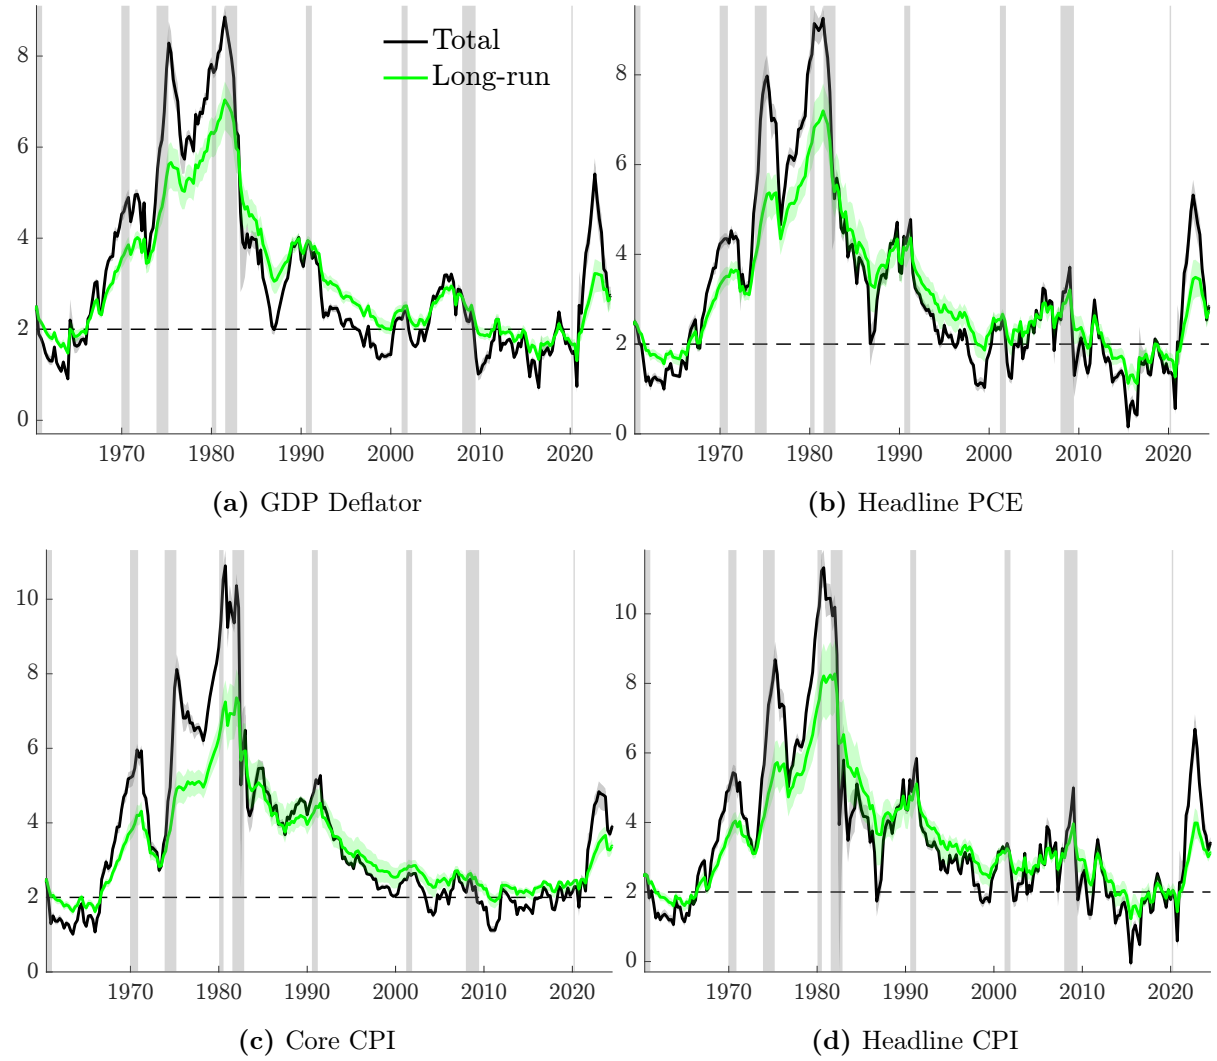
(a) Volatility **(b)** Skewness

Figure E1: 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.

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, core PCE shows the least variation in both volatility and skewness, appearing to be the more stable measure of price dynamics.

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



(a) GDP Deflator

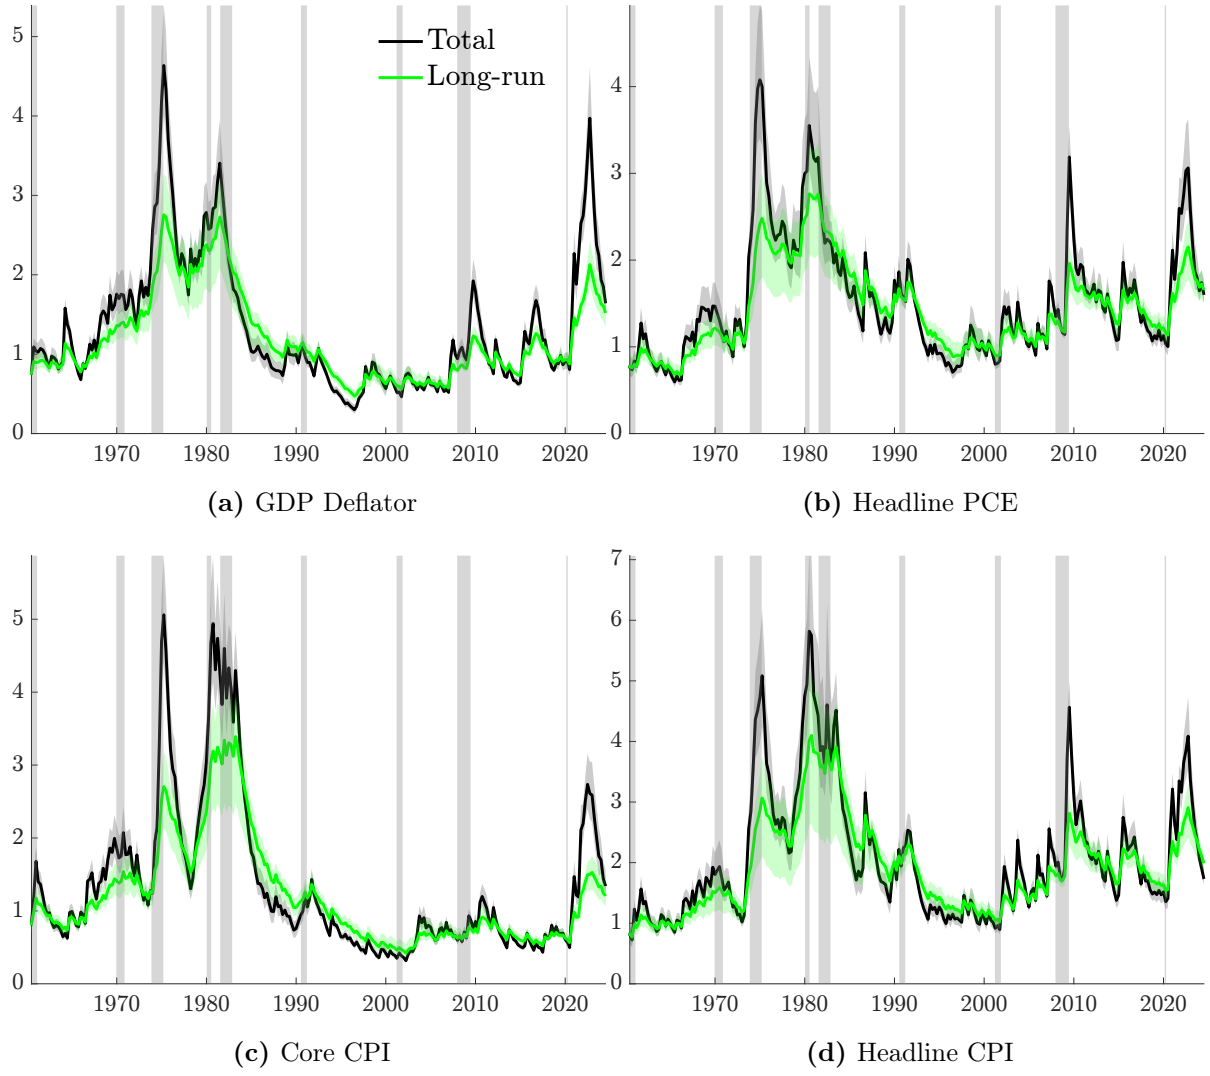
(b) Headline PCE

(c) Core CPI

(d) Headline CPI

Figure E2: Estimated mean for different inflation measures

Note: The panels report the estimated total (black) and long-run (green) mean for: (a) GDP deflator, (b) headline PCE, (c) core CPI, and (d) headline CPI. ray shaded areas represent NBER recessions.



(a) GDP Deflator

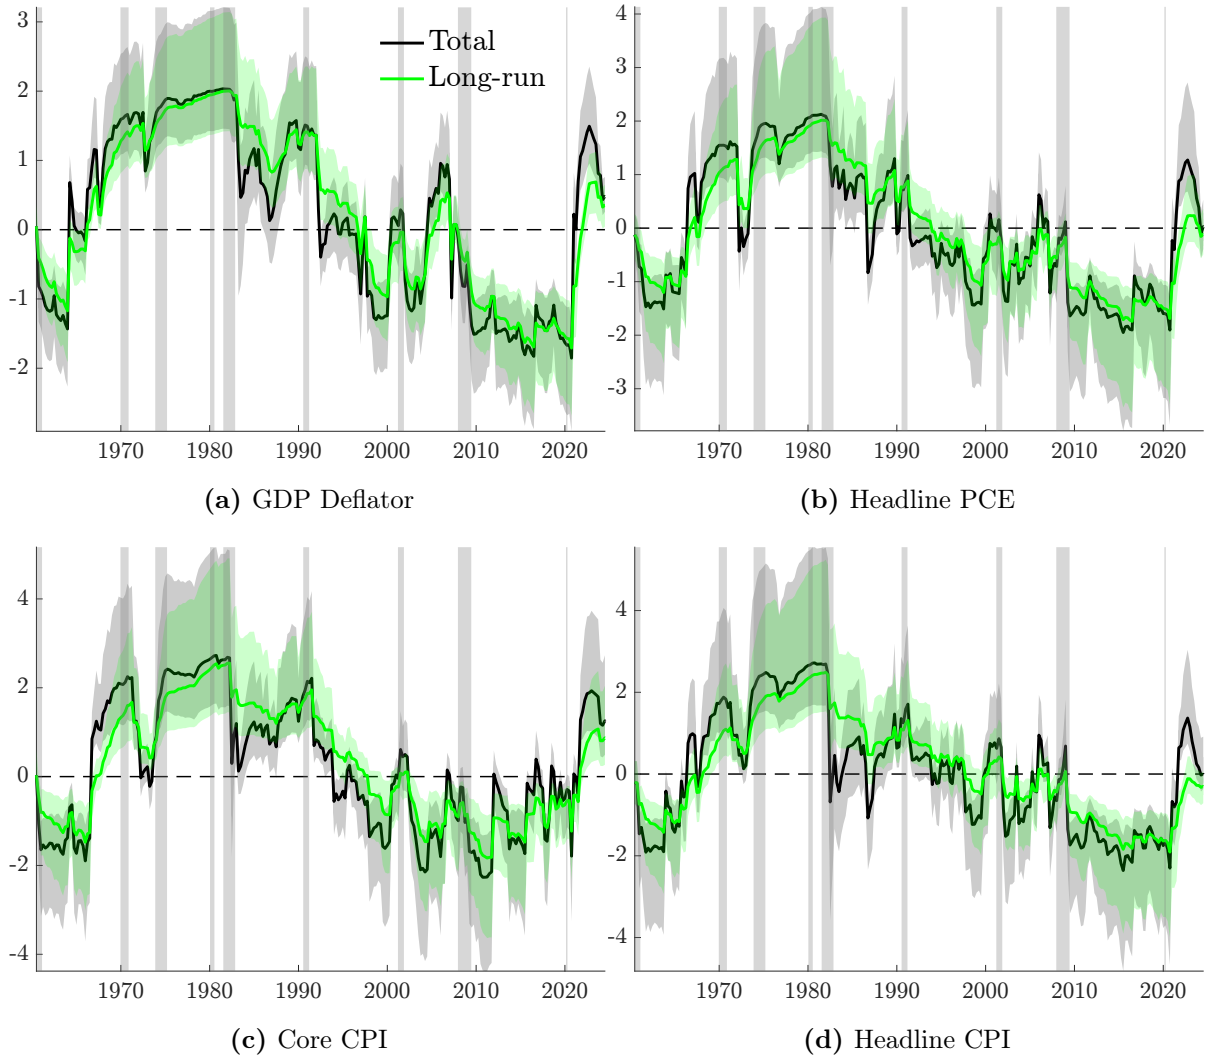
(b) Headline PCE

(c) Core CPI

(d) Headline CPI

Figure E3: Estimated volatility for different inflation measures

Note: The panels report the estimated total (black) and long-run (green) volatility for: (a) GDP deflator, (b) headline PCE, (c) core CPI, and (d) headline CPI. ray shaded areas represent NBER recessions.



(a) GDP Deflator **(b)** Headline PCE

(c) Core CPI **(d)** Headline CPI

Figure E4: Estimated skewness for different inflation measures

Note: The panels report the estimated total (black) and long-run (green) skewness for: (a) GDP deflator, (b) headline PCE, (c) core CPI, and (d) headline CPI. ray shaded areas represent NBER recessions.

F Additional results

Table F1: Out-of-sample comparison - *Student t*

	h = 1	h = 2	h = 3	h = 4	h = 8
MSFE	**0.839**	**0.854**	**0.908**	0.942	1.008
	(0.085)	(0.051)	(0.081)	(0.129)	(0.782)
CRPS	**0.938**	**0.950**	**0.960**	**0.959**	0.996
	(0.074)	(0.077)	(0.081)	(0.074)	(0.331)
	CRPS decomposition				
Right	0.927	**0.914**	**0.935**	**0.936**	0.981
	(0.109)	(0.075)	(0.099)	(0.074)	(0.116)
Left	**0.952**	0.982	0.986	0.983	1.009
	(0.088)	(0.258)	(0.197)	(0.200)	(0.808)
Center	**0.936**	0.953	**0.960**	**0.960**	0.999
	(0.058)	(0.117)	(0.041)	(0.060)	(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 F2: Out-of-sample comparison - $\varrho = 0,\ \forall t$

	h = 1	h = 2	h = 3	h = 4	h = 8
MSFE	**0.832**	0.884	**0.906**	0.966	1.055
	(0.081)	(0.115)	(0.057)	(0.190)	(0.997)
CRPS	0.947	0.961	**0.962**	0.978	1.024
	(0.116)	(0.153)	(0.075)	(0.166)	(0.983)
	CRPS decomposition				
Right	**0.928**	**0.928**	**0.947**	0.962	1.012
	(0.081)	(0.053)	(0.048)	(0.123)	(0.854)
Left	0.967	0.995	0.988	0.998	1.039
	(0.242)	(0.440)	(0.339)	(0.461)	(0.994)
Center	0.946	0.961	**0.954**	0.977	1.024
	(0.123)	(0.150)	(0.053)	(0.155)	(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 F3: Parameters estimates for the econometric model in Section 4

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.

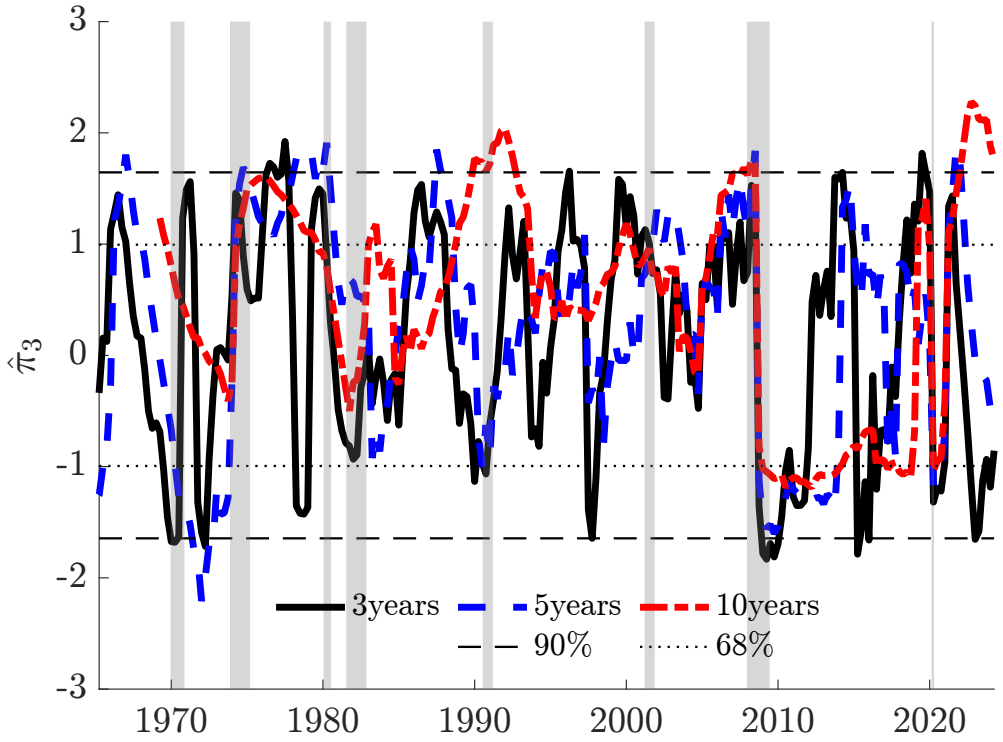


Figure F1: Bai and Ng (2005) rolling tests

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.