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

Andrea De Polis† Leonardo Melosi‡ Ivan Petrella§

This draft: July 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 consequences 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, First 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 general and can be integrated with various methods for estimating skewness, not only the one adopted in this paper.

Our method can also be applied to investigate the macroeconomic effects of different sources of risk skewness—such as trade risks, financial risks, climate risks, and technological disruptions. All that is required is that the methodology be applied to a model featuring those shocks and the relevant propagation channels.

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, 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

[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.

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

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. (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 Theoretical Framework

In this section, we highlight how shifts in the asymmetry of future macroeconomic risks influence equilibrium outcomes within a linear general equilibrium framework. We prove that changes in the perception of asymmetric risk can be interpreted as revisions in agents' beliefs about the future

distribution of shocks. The complete proof of this result, along with supporting claims, is provided in Appendix A. In this section, we prove two propositions showing that a linearized model in which agents receive signals about the likely evolution of the skewness of the shock distribution can be recast as a linear rational expectations model with news shocks. As a result, it can be solved using standard methods, such as *Gensys*.

2.1 Shifting asymmetries in linear general equilibrium models

Linearized rational expectations general equilibrium models, widely used in theoretical and empirical macroeconomic research, can be expressed as follows:[3]

$$\mathbf{A}_0\mathbf{z}_t = \mathbf{A}_f E_t\mathbf{z}_{t+1} + \mathbf{A}_b\mathbf{z}_{t-1} + \mathbf{B}_s\boldsymbol{\epsilon}_t^s + \mathbf{b}_a\epsilon_t^a. \tag{1}$$

where $\mathbf{z}_t$ collect endogenous and exogenous variables, the operator E_t denotes rational expectations conditional on the information set available at time t, which, as we will describe, contains the history of all the shocks. The vectors of i.i.d. disturbances are denoted by $\boldsymbol{\epsilon}_t^s$ and ϵ_t^a. The former includes all the symmetric shocks, e.g., normally distributed, whereas the latter is a possibly *asymmetric* shock with distribution $\mathsf{f}(\mu_a, \sigma_a, \varrho_a)$, where μ_a denotes the mode, σ_a the scale, and ϱ_a an asymmetry parameter that measures the degree of skewness.[4] For $\varrho_a > 0$ ($\varrho_a < 0$), the distribution f is positively (negatively) skewed, while $\varrho_a = 0$ implies symmetry. Asymmetric risk implies that the probability mass around the mode is unevenly distributed: risks are skewed to the downside (upside) when outcomes below (above) μ are more likely than those above (below) it.[5]

For ease of exposition, we focus on the family of two-piece distributions, introduced by Fechner (1897), which includes a large set of univariate skewed distributions (see, e.g., Fernández and Steel, 1998; Arellano-Valle et al., 2005). These distributions are flexible and convenient, making them widely used in finance, econometrics, and various other fields such as environmental and econology studies and engineering, for modeling asymmetric risks surrounding future events. In this class of

[3] In Appendix A.1, we derive the following log-linear system as a first-order Taylor expansion around the steady state of the full nonlinear model with skewed shocks.

[4] While we focus on the case with a single source of asymmetric risk, extending the analysis to multiple independent asymmetric shocks is straightforward and does not alter the underlying intuition.

[5] Let $r = \int_{\mu}^{\infty} \mathsf{f}\, dx / \int_{-\infty}^{\mu} \mathsf{f}\, dx$ be the ratio of probability mass right and left of the mode. Then, $\varrho = \frac{r-1}{r+1} \in (-1, 1)$.

distributions, the wedge between the expected value, $E\mathsf{f}$, and the mode μ_a (the *mean-mode wedge*, for short), depends on the linear mapping $E\mathsf{f} = \mu_a + \varkappa\sigma_a\varrho_a$, with $\varkappa > 0$. The exact definition of $\varkappa$ depends on the specific functional form of the skewed distribution.[6]

Time-varying asymmetric risk. A distinctive feature of our analysis is the assumption that the asymmetry parameter of the distribution f, denoted $\varrho_{a,t}$, varies over time. This means the economy can shift between periods of upside and downside risk, with the severity of each episode changing over time. We assume that the asymmetry parameter follows a stationary zero-mean process ($E(\varrho_{a,t}) = 0$), an assumption that we will relax in the empirical part of the paper. For simplicity in presenting the theoretical model, the other parameters—mode (μ_a) and scale (σ_a)—are held constant, with the mode to be set equal to zero. These restrictions can be relaxed, but doing so would complicate the analysis, which in this part of the paper is purely theoretical and intended to provide a tractable framework for studying the effects of asymmetry on optimal monetary policy within a linear-quadratic framework.[7] In the empirical analysis, we will jointly estimate the evolution of the three parameters of the predictive distribution of inflation in the data.

Agents beliefs. Agents are rational and perfectly informed. This means that they know the parameters in Equation (1), the shocks distributions, including the mode μ and the scale, σ, and the history of the realized shocks, including the past and current realizations of the asymmetry of the distribution f.

Agents receive news about the future evolution of the distribution's asymmetry, $\varrho_{a,t}$, as this is the stochastic parameter driving the mean-mode wedge. Specifically, in every period t, agents receive a set of noisy signals about the future evolution of the asymmetry parameter over the next $J > 0$ periods. At every horizon $j \in \{1, ..., J\}$ agents observe the signal $s_t^j = \varrho_{a,t+j} + \eta_t^j$, where η_t^j denote the zero-mean noise orthogonal to the future realizations of the asymmetric parameter.

[6]For the Skewed-Normal distribution, $\varkappa = \sqrt{2/\pi}$. In contrast, for the Skewed-t distribution—used to characterize inflation skewness in the following sections—$\varkappa$ depends on the degrees of freedom of the distribution (see, e.g., Arellano-Valle et al., 2005). Intuitively, when the distribution has fat (thin) tails, the effect of skewed risk on expected outcomes becomes larger (smaller).

[7]The mode may change stochastically. Provided that these changes do not offset the effects of skewness on the expected value of the shocks, the main conclusions of this section remain materially unaffected.

Rational agents solve the signal extraction problem, updating expectations about the asymmetry parameter j periods ahead. These updates are always iid, and are denoted as $\varrho^j_{a,t} \equiv E_t \varrho_{a,t+j} - E_{t-1}\varrho_{a,t+j}$.

Model dynamics, forecasts, and central scenarios. As a consequence of the certainty equivalence in a linear setting, only expectations matter for the first-order solution of the model.

Proposition 1. *The solution of the linear rational expectations model with time-varying skewness in Equation (1) reads:*

$$\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\boldsymbol{\epsilon}^s_t \ \epsilon^a_t\right]' + \mathbf{\Theta}_y \sum_{j=1}^{\infty} \mathbf{\Theta}_f^{j-1} \mathbf{\Theta}_z \mathbf{\Xi} E_t \epsilon^a_{t+j}, \tag{2}$$

where $\mathbf{\Xi}$ *is a selection vector, and 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* $\mathbf{b}_a$.

Proof. See Appendix A.2. □

Proposition 1 establishes that anticipated changes in the asymmetry parameter of future shock distributions, $\varrho_{a,t+j}$, affect equilibrium outcomes as they lead agents to revise their expectations about the future realizations of the shocks $E_t \epsilon^a_{t+j}$.

The third term on the right-hand side of Equation (2) captures the role of agents' beliefs about the evolution of asymmetry, which depend on the noisy signals agents receive. If the signals are not informative of the asymmetry of future shocks ϵ^a_{t+j}, or if the news leads to the belief that shocks will be drawn from symmetric distributions ($E_t \varrho_{t+j} = 0$ for all $j \in 1, \ldots, J$), then the expected value of future asymmetric shocks is zero, and the third term on the right-hand side of Equation (2) drops out. Thus, the matrices in that term capture how anticipated asymmetry in the distribution of future shocks tilts agents' forecast, $\mathbf{z}_{t+j}$, away from the central—or modal—forecast.

If agents did not receive informative signals about the evolution of the asymmetry in ε^a, or if they believed that future shocks would be drawn from symmetric distributions, Equation (2) simplifies to the familiar solution: $\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\boldsymbol{\epsilon}^s_t \ \epsilon^a_t\right]'$.[8]

[8]The matrices $\mathbf{\Theta}_0$, $\mathbf{\Theta}_1$, $\mathbf{\Theta}_y$, and $\mathbf{\Theta}_z$ can be easily computed using off-the-shelf linear rational expectations solvers, such as the *Gensys* routine introduced by Sims (2002).

Proposition 1 also highlights that the *central or modal forecast* of the state variables, $\mathbf{z}_{t+j}$, formed at time t, is given by $\mathbf{z}_{t+j} = \mathbf{\Theta}_1^j \mathbf{z}_t$. Yet, this expression coincides with the expected value only in the absence of anticipated skewed risk. When asymmetry is expected over the forecast horizon, the expected outcomes diverge from the modal forecast due to the influence of risk perceptions on agents' expectations.

First-order effects of expected changes in asymmetry. An important implication of modeling asymmetric shocks using a two-piece distribution is that expected changes in the asymmetry parameter of j-period-ahead shocks, $E_t \varrho_{a,t+j}$, affect equilibrium outcomes *only* through their impact on the expected value of future shocks, $E_t \epsilon^a_{t+j}$. Specifically, these effects on expectations are given by

$$E_t \epsilon^a_{t+j} = \varkappa \sigma_a E_t \varrho_{a,t+j}. \tag{3}$$

The model solution in Equation (2), together with Equation (3), makes clear how changes in the expected asymmetry of future risks influence macroeconomic dynamics.

Revisions in the asymmetry parameter lead to corresponding revisions in expectations about future shocks, i.e., $E_t \epsilon^a_{t+j} - E_{t-1} \epsilon^a_{t+j} = \varkappa \sigma_a \varrho^j_{a,t}$. Thus, even if the most likely outcome remains unchanged, a shift in the balance of risks prompts an immediate adjustment in current equilibrium allocations and expectations:

$$E_t \mathbf{z}_{t+h} - E_{t-1} \mathbf{z}_{t+h} = \mathsf{M}(j,h) \varkappa \sigma_a \varrho^j_{a,t}, \quad h \geq 0, \tag{4}$$

where $\mathsf{M}(j,h)$ is a convolution of the system matrices in Equation (2).[9]

2.2 Beliefs representation

We now show that, to a first-order approximation, a model with asymmetric risks can be equivalently represented as one with symmetric, zero-mean shocks, augmented by additional *beliefs* shocks. These belief shocks capture revisions to agents' expectations following the arrival of the signals. We refer to this formulation as the model's *beliefs representation*.

[9] See Appendix A.3 for detailed derivations.

This representation preserves the linearized equilibrium laws of motion that characterize the dynamics of the original model with asymmetric risks. The only point of departure is the introduction of *dummy surprise and anticipated shocks* that encode belief-driven distortions.

In the beliefs representation, define the asymmetric shock as the sum of past news shocks and current shock:

$$\epsilon_t^a \equiv \sum_{j=0}^{J} \varphi_{a,t-j}^j, \tag{5}$$

where φ_t^j denotes the dummy beliefs shocks.

The following proposition provides the restrictions on the stochastic process of the dummy beliefs shocks ensuring that the beliefs representation is equivalent to the actual model with time-varying asymmetry.

Proposition 2. *The equilibrium dynamics of the beliefs representation are identical to the ones of the actual economy—shown in Equation (2)— if the following two conditions hold: (i) The dummy anticipated shocks $\varphi_{a,t}^j$ match the revisions in expectations regarding the asymmetry of the distribution of future shocks at each horizon. That is, for all $j \in 1, \ldots, J$, the following condition must hold:*

$$\varphi_t^j = E_t \epsilon_{a,t+j} - E_{t-1}\epsilon_{a,t+j} = \varkappa \sigma_a \left(E_t \varrho_{a,t+j} - E_{t-1}\varrho_{a,t+j} \right) = \varkappa \sigma_a \varrho_{a,t}^j, \tag{6}$$

where $\varrho_{a,t}^j$ denotes agents' revision in expectations regarding the asymmetry parameter after observing the signal, s_t^j, associated with horizon j in the actual model. (ii) The surprise dummy shock must satisfy the following condition in every period t

$$\varphi_t^0 = \epsilon_{a,t} - \sum_{j=1}^{J} \varphi_{t-j}^j. \tag{7}$$

Proof. As shown in Equation (3), changes in the expected evolution of the asymmetry affect equilibrium outcomes only through their impact on expectations about the future realization of shocks. Condition (i) ensures that the dummy belief shocks replicate the same impact of anticipated asymmetry on the expectations of future shocks, $E_t \epsilon_{a,t+j}$. Condition (ii) ensures that the realization

of the shock $\epsilon_{a,t}$ exactly matches that in the actual economy in every period. Together, these two conditions guarantee that the beliefs representation replicates the actual economy exactly—both in terms of structure, realizations of shocks, and expectations about the future realizations of asymmetric shocks in every period. As a result, its equilibrium outcomes satisfy Equation (2) in Proposition 1. Detailed derivations are available in Appendix A.4 □

Establishing that the beliefs representation is indeed a representation of the economy with time-varying asymmetry of shocks is a key result. This is because the beliefs representation can be solved using off-the-shelf, highly efficient solution techniques such as *Gensys*. Consequently, once the beliefs representation is solved, it becomes straightforward to derive the matrices that determine the equilibrium laws of motion of the actual economy, as given by Equation (2).

It is worth emphasizing that the applicability of standard solution techniques for linear rational expectations models remains valid even when the underlying signal extraction problem is highly complex and may lead to non-Gaussian revisions to expectations about future shocks. In fact, these solvers typically do not require specifying the stochastic process of the shocks, beyond assuming that they are i.i.d. and mean zero—a condition that is not restrictive in most cases.

These shocks are i.i.d. because their stochastic process is tied to the process generating revisions in expectations about the asymmetry of future shock distributions, as shown in Proposition 2. In any rational expectations model, revisions to expectations about future shocks or variables are unpredictable and thus follow i.i.d. processes.

The beliefs representation can be used to study environments where agents receive information about future risk, conveyed more generally through changes in the scale and asymmetry of shock distributions. To a first-order approximation, anticipated changes in either parameter affect allocations by shifting the expected value of future shocks—thus preserving the relevance of the beliefs-based approach.

As shown in Equation (3), the effect of asymmetry depends on the scale of the distribution: a given shift in skewness has a larger impact on expectations when uncertainty is high (i.e., when σ_a is large). Likewise, changes in overall uncertainty influence average outcomes if the distribution is skewed, even when the asymmetry itself is unchanged. In such cases, the direction and magnitude

of the effect depend on the sign of the skew.

In the next sections, we apply this framework to analyze how shifts in perceived risk should inform the conduct of monetary policy.

3 Optimal monetary policy with asymmetric risks

In this section, we examine the implications of time-varying asymmetric macroeconomic risk for optimal monetary policy within a standard New Keynesian model with sticky prices. We begin by reviewing the well-known case of optimal policy in a linear-quadratic framework under symmetric risk, as presented in Galí (2008, Chapter 3 and 5). We then extend the analysis to account for shifts in asymmetric risk, using the tools developed in the previous section. Detailed derivations are provided in Appendix B.

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,[10] $\frac{1}{2}E_0 \sum_{t=0}^{\infty} \beta^t \left(\hat{\pi}_t^2 + \alpha_x \hat{x}_t^2\right)$, subject to the sequence of constraints given by the Phillips curve, $\hat{\pi}_t = \beta E_t \pi_{t+1} + \kappa \hat{x}_t + u_t$. We further assume that the cost-push shock follows $u_t \sim \text{i.i.d.}\, \mathcal{N}(0, \sigma_u)$. As a result, future uncertainty is governed by a symmetric, time-invariant distribution determined by the properties of the shock process.

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. p_t denotes the

[10] All variables are expressed in log-deviations from their steady-state value. 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.

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. 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.

Under the optimal policy, the price level evolves according to

$$\bar{p}_t = \eta \bar{p}_{t-1} + \lambda u_t, \tag{9}$$

and the corresponding optimal monetary policy rule is given by

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t, \tag{10}$$

where σ is the inverse of the household's intertemporal elasticity of substitution, and η and λ are functions of the structural parameters of the model.

3.2 The case of asymmetric risks

Let us now consider the case where the stochastic process driving the cost-push shock is no longer symmetric. Specifically, assume that the shocks follow $\tilde{u}_t \sim$ i.i.d. $\mathsf{f}(0, \sigma_u, \varrho_{u,t})$, where $\mathsf{f}(0, \sigma_u, \varrho_{u,t})$ denotes a general unimodal distribution centered at zero but asymmetric when $\varrho_{u,t} \neq 0$. In this setting, the perceived risk associated with the distribution evolves over time as agents update their views about the balance of upside and downside risks. As a result, the expected value of future shocks reflects these changes in perceived asymmetry. While the mode—the most likely outcome—remains at zero, the mean shifts with the asymmetry, so that $E_t\tilde{u}_{t+j} = \varkappa\sigma_u E_t\varrho_{u,t+j}$.

To simplify the characterization of optimal monetary policy, we assume that agents receive news about the skewness of the distribution only up to $J = 1$ period ahead. Beyond that, they assume all risk is symmetric. In this environment, at time t, agents observe the current realization and form expectations about future realizations of the shocks based on their current perception of the skewness of risk. This implies $E_t\tilde{u}_{t+1} = \varkappa\sigma_u E_t\varrho_{u,t+1}$, while $E_t\tilde{u}_{t+j} = 0$ for all $j > 1$. Moreover,

since $\varrho_{u,t+1} = \varrho^0_{u,t+1} + \varrho^1_{u,t}$, and assuming agents do not anticipate further revisions in perceived asymmetry (i.e., $E_t\varrho^0_{u,t+1} = 0$), it follows that $E_t\tilde{u}_{t+1} = \varkappa\sigma_u\varrho^1_{u,t}$, where $\varrho^1_{u,t} = E_t\varrho_{u,t+1} - E_{t-1}\varrho_{u,t+1}$. Expectations of future shocks therefore reflect the latest update in perceived risk asymmetry.

This information structure can be equivalently represented by introducing a tilt in agents' beliefs—that is, by assuming $E_t(\widetilde{u}_{t+1}) = \varphi_{t,t+1}$, where beliefs evolve over time according to $\varphi_{t,t+1} = \varphi^0_t + \varphi^1_{t-1}$. Setting $\varphi^1_t = \varkappa\sigma_u\varrho^1_{u,t}$, and imposing the restriction $\varphi^0_t = -\varphi^1_{t-1}$, ensures that belief revisions replicate the impact of changes in perceived asymmetry on the expected value of the shock, while leaving the actual realization unaffected. In other words, prior beliefs need not materialize and may turn out to be ex post noise. We then have that, up to a first-order approximation, the dynamics of a model with evolving perceptions of asymmetric risk can be equivalently represented by a model with symmetric shocks and belief distortions captured by dummy surprises.

Optimal monetary policy under time-varying asymmetric macroeconomic risks. 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:[11]

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

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 revision in the skewness of risks (as $\varphi^1_t \propto \varrho^1_{u,t}$).

Therefore, the optimal policy requires the central bank to respond to the effects of the changing balance of macroeconomic risks on agents' expectations. The optimal interest rate rule is:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\varphi^1_t, \tag{11}$$

[11] Detailed derivations are provided in Appendix B.

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.

3.3 The case of asymmetric risks with an unwitting central bank

Having established the optimal monetary policy response under full awareness of future risk asymmetries, we now explore the implications of a setting in which the central bank mistakenly assumes that macroeconomic risks are balanced. In this framework, the central bank fails to internalize the asymmetry in macroeconomic shocks, even though private agents correctly perceive and respond to such risks.

In this case, the central bank targets a price level consistent with the assumption that shocks are symmetrically distributed—conditional on the previous period's price level. It selects the output gap $\hat{x}_t$ according to the rule:

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^s, \tag{12}$$

where $\bar{p}_t^s = p_t^s - p_{-1}$ and p_t^s is the price level the central bank believes it should target under the mistaken belief of symmetry. This target evolves as: $\bar{p}_t^s = \eta\bar{p}_{t-1} + \lambda u_t$. Importantly, the central bank's target under asymmetric risk differs slightly from the fully symmetric benchmark. In this setting, policy is based solely on the previous period's equilibrium price level, $\bar{p}_{t-1}$, without attempting to revise or offset past errors. As a result, deviations between actual and perceived past price levels are not corrected—"bygones are bygones."

Ultimately, the central bank forms its policy under a misperception: it responds to deviations in a price level $\bar{p}_t^s$ that it believes is attainable, even though it is not, due to agents' responses to shifts in perceived risk asymmetry. Comparing this misinformed policy rule with the full-information benchmark (Equation (8)) reveals that the output gap is incorrectly determined. It is tied to a price level target that cannot be achieved, and attempts to reach it ultimately lead to larger deviations in both the output gap and the price level.

3.4 A illustrative example

We illustrate the role of optimal monetary policy in the presence of asymmetric risks using a simple example. We assume the economy is initially at the non-risky steady state. In period 1, agents learn that the distribution of cost-push shocks will be positively skewed in period 2. This upside risk is captured by a dummy anticipated shock $\varphi_t^1 = 0.8$. A positive cost-push shock then materializes in period 2, we assume $\varepsilon_t^u = 0.8$, while all other periods feature symmetric shocks with zero mean.

Figure 1 presents the simulation results.[12] The black solid line shows the benchmark case with symmetric risks ($\varphi_t^1 = 0$), while the blue line reflects optimal policy under asymmetric risk. The red line shows the outcome when the central bank ignores the asymmetry. The green line in the bottom-right panel shows the evolution of agents' expectations, captured by φ_t^1, which spikes only in period 1, consistent with the temporary asymmetry in period 2.

A comparison of the blue and red lines reveals the implications of accounting—or failing to account—for asymmetric risks. In period 1, the unwitting central bank follows the same policy as under symmetry, targeting a price level that is ultimately unattainable due to agents' expectations of increasing inflation. As a result, equilibrium inflation overshoots the target and is larger than what would be achieved with the optimal response to a shift in risk (bottom-left panel).

By contrast, a central bank that accounts for inflationary risk adjusts policy preemptively. It lowers the output gap and increases the price level in period 1 to offset the effects of a positive shock expected in period 2. The difference between the optimal and suboptimal policies in period 1 (blue vs. red line) reflects this forward-looking response. In period 2 and beyond, since risk returns to symmetry, both policies coincide, but the effects of the asymmetric shock persist due to the path dependency introduced by the central bank's commitment.

In summary, when risks are skewed, optimal policy must lean against the direction of asymmetry. Positive skewness in inflation risk calls for a policy stance that reduces inflation expectations and tolerates an output contraction. If risks were skewed to the downside, the central bank would instead overheat the economy. Finally, symmetric risk delivers higher welfare, highlighting the

[12] Model parameters follow Galí (2008, Chapter 3) and are reported in Appendix B.

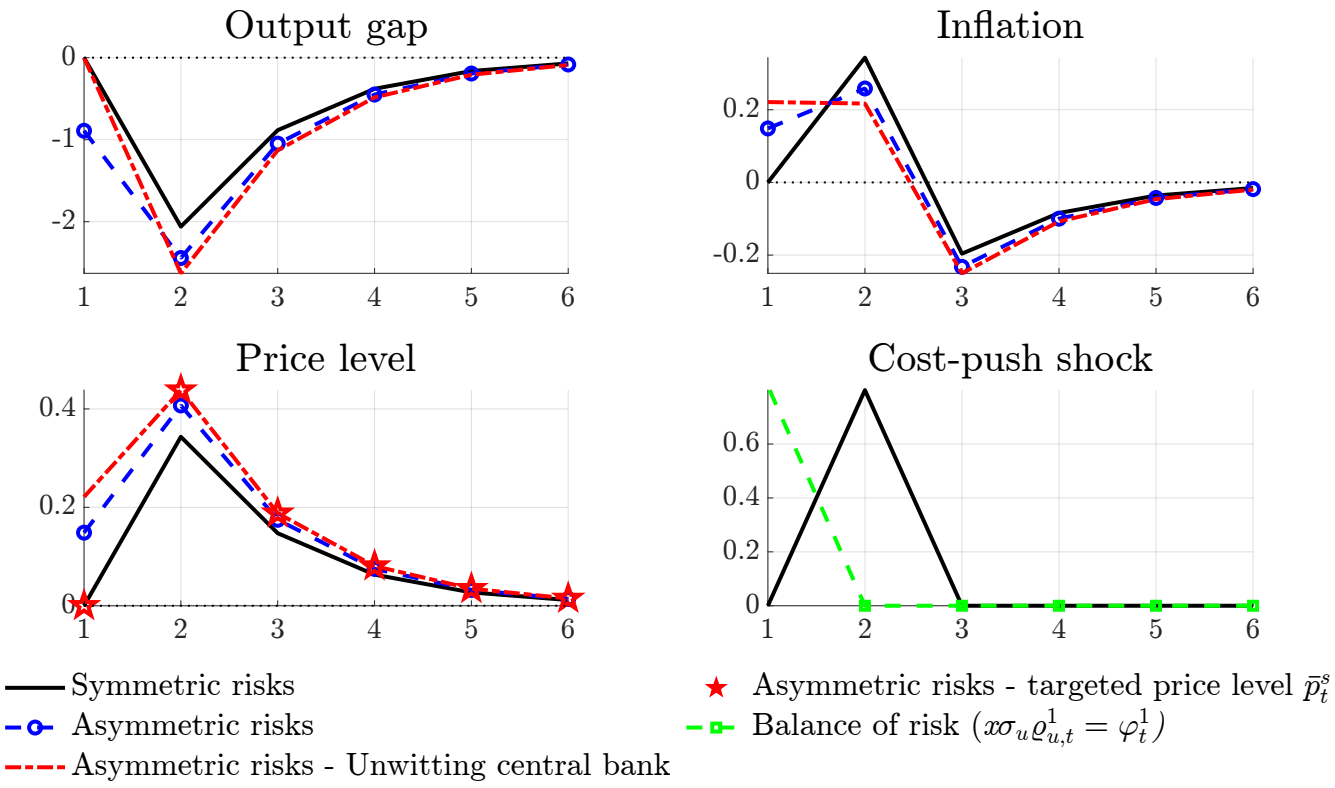


Figure 1: Optimal monetary policy under symmetric and asymmetric macroeconomic risks.

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

challenges that asymmetric risks pose for monetary policy.

4 An econometric framework for inflation risk

In the previous sections we established that changes in the asymmetry of inflation risk warrant a response from the central bank. We now propose a time series model that allows to estimate 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), but adds the flexibility of allowing for skewness in the predictive distributions.

4.1 Model specification

Let $\pi_t = 400 \log(p_t/p_{t-1})$ denote the annualized, quarter-on-quarter (core) PCE inflation rate, 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:

$$\pi_t \sim Skt_\nu(\mu_t, \sigma_t^2, \varrho_t), \tag{13}$$

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 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.

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 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[14]

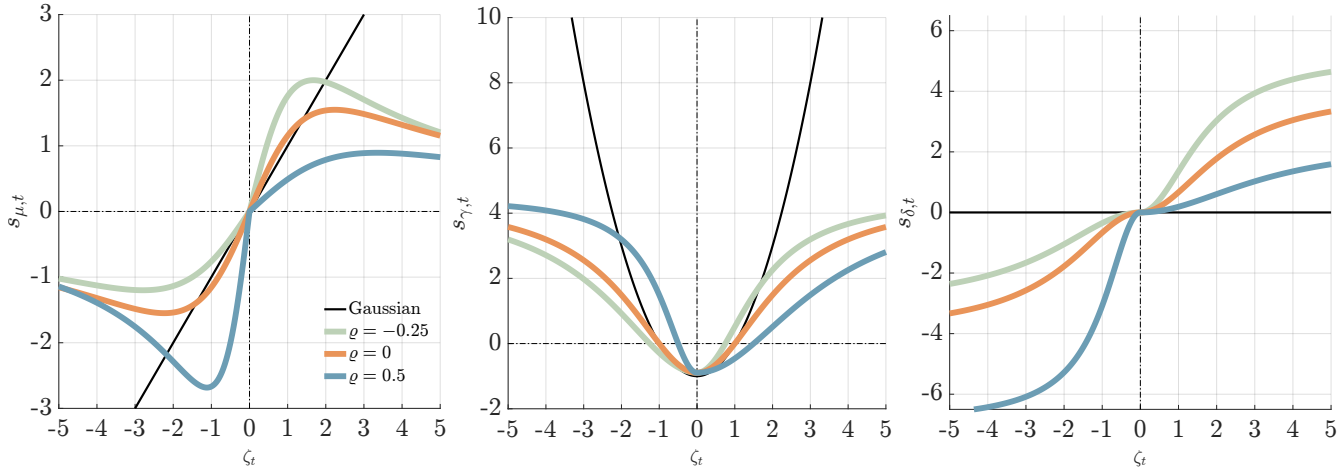
$$\bar{f}_{i,t} = \bar{f}_{i,t-1} + a_i s_{i,t-1}, \tag{14}$$

$$\tilde{f}_{i,t} = \phi_i \tilde{f}_{i,t-1} + b_i s_{i,t-1}. \tag{15}$$

Updates of the time-varying parameters are proportional to $s_{i,t-1}$, which is the *scaled score* of the

[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] Measures of model fit (see, e.g., Newton and Raftery, 1994; Spiegelhalter et al., 2002) favors the permanent-transitory decomposition over specifications which only features transitory components.



(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) $Skt_5(\mu_t, \sigma_t, \varrho_t)$.

conditional distribution (as in Creal et al., 2013; Harvey, 2013).[15]

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 of the updates.[16] 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

[15] 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)

[16] 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).

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 case of the Skew-t distribution in Equation (13), one can show that $\forall h > 0$

$$\begin{aligned} E_t\pi_{t+h} &= \int_{\mathbb{R}} \pi p(\pi|\mu_{t+h}, \sigma_{t+h}, \varrho_{t+h}, \nu)d\pi, \\ &= \mu_{t+h} + \underbrace{g(\nu)\sigma_{t+h}\varrho_{t+h}}_{\psi_{t+h}} \end{aligned} \tag{16}$$

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.

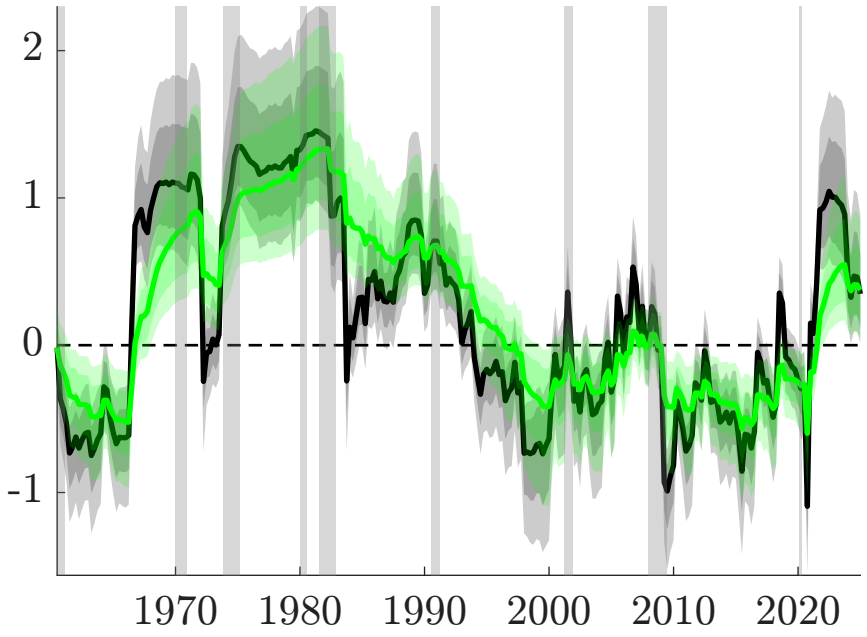


Figure 3: Time-varying skewness of Core PCE inflation

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.

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.

This decomposition in provides a useful tool to calibrate the asymmetry of economic shocks. In Section 5, we will exploit Equation (16) and the belief representation introduced in Section 2 to introduce asymmetry in the shocks of a quantitative DSGE to understand the optimal monetary policy response when inflation risks is skewed.

4.2 Asymmetric inflation risk

The estimated model offers novel insights into the time-varying stochastic properties of the inflation process. Estimates of the time-varying mean and volatility of US core PCE are broadly in line with measures obtained within the Stock and Watson (2007, 2016) framework.[17]

Figure 3 displays the estimated time-varying skewness. 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

[17] See Appendix E.

(labeled as "Long-run and computed setting $f_t = \bar{f}_t$) is in green. The estimates show 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.[18]

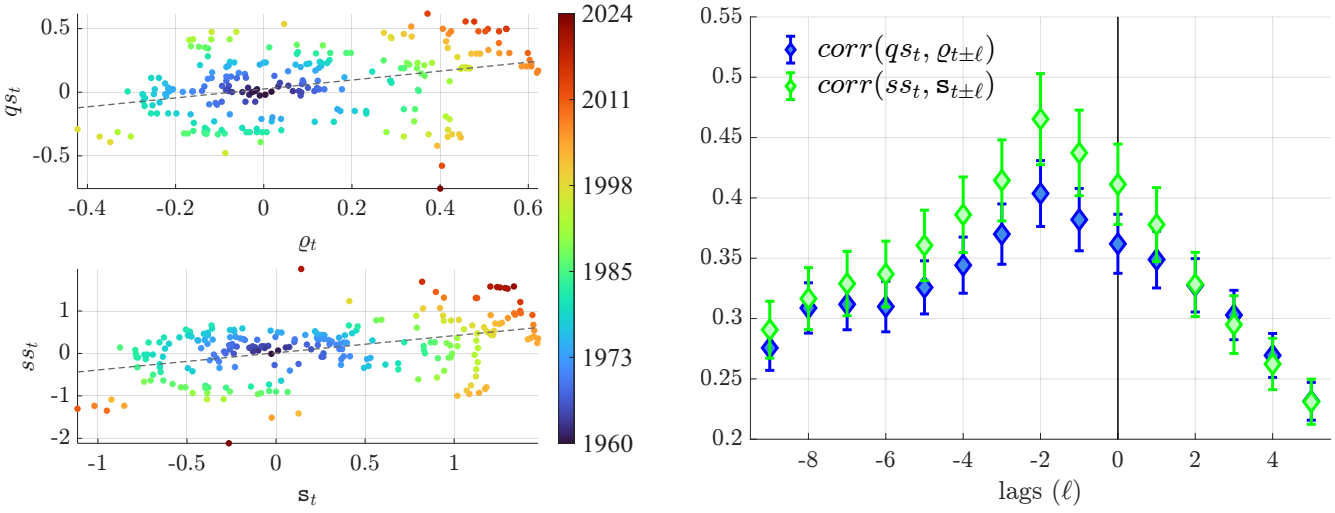
4.3 Skewness predictability and out-of-sample performance

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. The top panel of figure Figure 4 (a) reports the correlation over time between the data-based 5-year rolling quantile skewness (qs_t) and the asymmetry parameter, ϱ_t;[19] the bottom panel compares 5-year rolling sample skewness, ss_t, against the model-based conditional skewness, labeled $\mathtt{s_t}$. Although measures exhibit similar patterns, ss_t is noticeably noisier due to its sensitivity to individual observations. On the other hand, model-based skewness estimates are less influenced by extreme values, due to a more robust updating mechanism.

Although rolling measures of sample skewness can be easily computed from raw data, rolling

[18] See Table F4 in Appendix F.

[19] 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.



(a) Time series correlation

(b) Cross-correlations

Figure 4: Model-based vs rolling measures of skewness

Note: The figure reports (a) time series and (b) cross correlations between data-based and model-based measures of inflation skewness. Data-based measures are 5-year rolling quantile skewness (qs_t) and 5-year rolling sample skewness (ss_t). Model-based measures are the asymmetry parameter (ϱ_t) and the conditional skewness, $\mathbf{s}_t$.

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 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 4 (b) 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

Table 1: 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.

PCE inflation appears to exhibit a more attenuated balance of risk compared to other inflation measures.

Out-of-sample performance We set up a real-time forecasting exercise where for each inflation vintage we produce up to eight-step ahead forecasts for the whole density of core PCE inflation, starting from 2000Q1.[20] We compare our model against the widely used UCSV model of Stock and Watson (2007), in terms of point, density and event forecasts, measured by the mean squared forecast error (MSFE), the Continuously Ranked Probability Score (CRPS) of Gneiting and Ranjan (2011) and the Brier score, respectively.[21] Additional results are available in Appendix F.

The comparison strongly favors our model, across all metrics and forecast horizons, as presented in Table 1. The results, expressed as the ratio of the score achieved by our model to that of the

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

[21] The Continuously Ranked Probability Score (CRPS) scoring rule measure the squared difference between the predictive distribution function and the "perfect forecast". This score can be modified to highlight different regions of the predictive density (see, e.g. Gneiting and Ranjan, 2011).

UCSV benchmark, show significant gains in point forecasts range from 15% at short horizons to 5% over the medium term, where p-values for the Diebold and Mariano (1995) test are provided in parentheses. Improvements of up to 8% are observed in CRPS scores. These results underline the importance of accounting for inflation skewness as a means to enhance forecasting accuracy.

The UCSV model lacks a mechanism for capturing skewness but also overlooks the presence of fat tails in the data. Replicating the exercise using a specification that excludes any asymmetry (similarly to Delle Monache and Petrella, 2017) confirms the importance of incorporating skewness to improve model fit and inflation forecasting accuracy. Significant and substantial gains are also reported for the predictions of events where π_{t+h} is lower than 1.5%, greater than 2.5%, or falls within these thresholds.

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. We find out model to perform on par with the SPF.[22] Notably, for $P\left(\pi_{t+1}^{Q4} > 2.5\%\right)$ our model shows meaningful advantages, stemming from consistently lower event probability assessments compared to the SPF during the decade following the Great Recession, and a timely adjustment in early 2022, that accurately timed 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 cost-push shocks follow a Skew-t distribution with time-varying moments. As in the stylized model considered in Section 2, agents receive news about the time variation in the skewness of the distribution of these shocks, up to $J = 20$ quarters ahead. We assume that, in each period, agents expect moments

[22] We compare forecasts for the Q4-over-Q4 core PCE, which is the measure predicted by the SPF. The evaluation sample starts in 2007Q1 to match SPF data.

to change in line with the forecasts provided by the model estimated in the previous section.[23]

5.1 Belief representation of Smets and Wouters (2007)

We set up the belief representation of the Smets and Wouters (2007) model following Section 3.2. We augment the price Phillips Curve with a set of dummy anticipated cost-push shocks as follows:[24]

$$\hat{\pi}_t = \pi_1 \hat{\pi}_{t-1} + \pi_2 E_t \hat{\pi}_{t+1} - \pi_3 \mu_t^p + \varepsilon_t^p + \sum_{j=0}^{J} \varphi_{t-j}^j, \tag{17}$$

where the variable μ_t^p represents firms' price mark-ups, the shock ε_t^p stands for the actual cost-push 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 uticipated 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 cost-push shocks on agents's expectations about the realization of these shocks.

In the beliefs representation, the expected cost-push shock h periods ahead can be decomposed as follows:

$$E_t\left[\varepsilon_{t+h}^p + \sum_{j=0}^{J} \varphi_{t-j}^j\right] = \underbrace{E_t(\varepsilon_{t+h}^p)}_{\text{mode } (\mu_{t+h|t}^p)} + \underbrace{\sum_{j=h}^{J} \varphi_{t+h-j}^j}_{\text{balance of risks } (\psi_{t+h|t}^p)}, \tag{18}$$

which resembles the structure of Equation (16), since we assumed that cost-push 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

[23] 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.

[24] This equation embeds an identification assumption: specifically, we assume that only expectations of future cost-push 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.

cost-push shock h periods ahead is given by

$$E_t - E_{t-1}\left[\varepsilon^p_{t+h} + \sum_{j=0}^{J} \varphi^j_{t-j}\right] = \left(\mu^p_{t+h|t} - \mu^p_{t+h|t-1}\right) + \varphi^h_t. \tag{19}$$

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^h_t.

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 be solved quickly and efficiently using off-the-shelf techniques. The solution of the model is standard and can be expressed as follows:

$$\mathbf{s}_t = \Gamma \mathbf{s}_{t-1} + \Omega \boldsymbol{e}_t, \tag{20}$$

where the vector $\mathbf{s}_t$ contains the model's state variables, and the vector $\boldsymbol{e}_t$ includes all shocks as well as the belief dummies. The matrices Γ and Ω are the solution matrices, which depend on the structural parameters of the model.

5.2 Calibrating inflation risks

We now have to align the beliefs representation to match the inflation asymmetry estimated in the previous section. Equation (16) offers a clean way to decompose model based inflation expectations into a central scenario (the most likely expected outcome) and the tilt induced by the balance of risks around it, ψ; this decomposition is reported in Figure 5. 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.

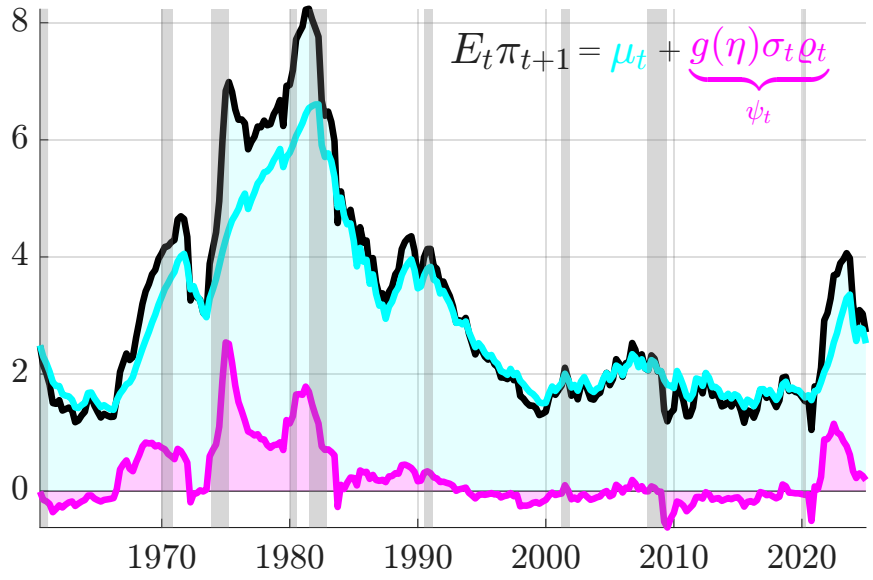


Figure 5: Expected inflation decomposition

Note: The figure report the decomposition of the model based measure of inflation expectations, reported in **??**. The total expectation (in black) is decomposed into a central scenario (μ_t, cyan) and the tilt induced by the balance of risks (ψ_t, magenta). Gray shaded areas represent NBER recessions.

Given $\psi_{t+J|t}$, the $J-$period-ahead forecast, conditional to time t, for the balance of inflation risks produced by Equation (16) (magenta line in Figure 5), we obtain the vector of $J+1$ dummy surprise and anticipated shocks as:

$$\begin{bmatrix} -\sum_{j=1}^{J}\varphi_{t-j}^{j} \\ \psi_{t+1|t} - \psi_{t+1|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{0}_{1\times J} \\ \Omega^S & \Omega^N \end{bmatrix} \begin{bmatrix} \varphi_t^0 \\ \varphi_t^1 \\ \vdots \\ \varphi_t^J \end{bmatrix}, \tag{21}$$

where Ω^S and Ω^N represent the contemporaneous effects of surprise and anticipated dummy cost-push 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 cost-push. As noted in Section 2, 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.

The last J equations map revisions to the mean-mode inflation wedge, $\psi_{t+j|t} - \psi_{t+j|t-1}$ estimated by the forecasting model, to the surprise and anticipated dummy cost-push shocks. The anticipated shocks $\left\{\varphi_t^j\right\}_{j=1}^{J}$ are multiplied by the matrix Ω^N which returns their effects on inflation expectations at each horizons $(1, ..., J)$. That is, the last J equations impose that the dummy anticipated shocks move inflation expectations to match the revisions to the mean-mode inflation wedge estimated in the data at each horizon.

5.3 Risk-Adjusted Inflation Targeting

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:

$$\hat{r}_t = \rho \hat{r}_{t-1} + (1-\rho)\left[r_x \hat{x}_t + r_{\Delta x} \Delta \hat{x}_t + r_\pi \left(\hat{\pi}_t - \underbrace{\sum_{j=1}^{J} \hat{\pi}^{\star}_{t|t-j}}_{\hat{\pi}_t^{\text{RAIT}}} \right) \right] + \varepsilon_t^r, \tag{22}$$

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. The variable $\hat{\pi}_t^{\text{RAIT}}$ 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 $\hat{\pi}_{t+h}^{\text{RAIT}}$ 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:

$$-\begin{bmatrix} \psi_{t+1|t} - \psi_{t+1|t-1} \\ \psi_{t+2|t} - \psi_{t+2|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = \Omega_{FG} \begin{bmatrix} \hat{\pi}^{\star}_{t+1|t} \\ \hat{\pi}^{\star}_{t+2|t} \\ \vdots \\ \hat{\pi}^{\star}_{t+J|t} \end{bmatrix}, \tag{23}$$

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. 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 (22). 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 (16). 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 (23) 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.

5.4 The RAIT in the post-pandemic era

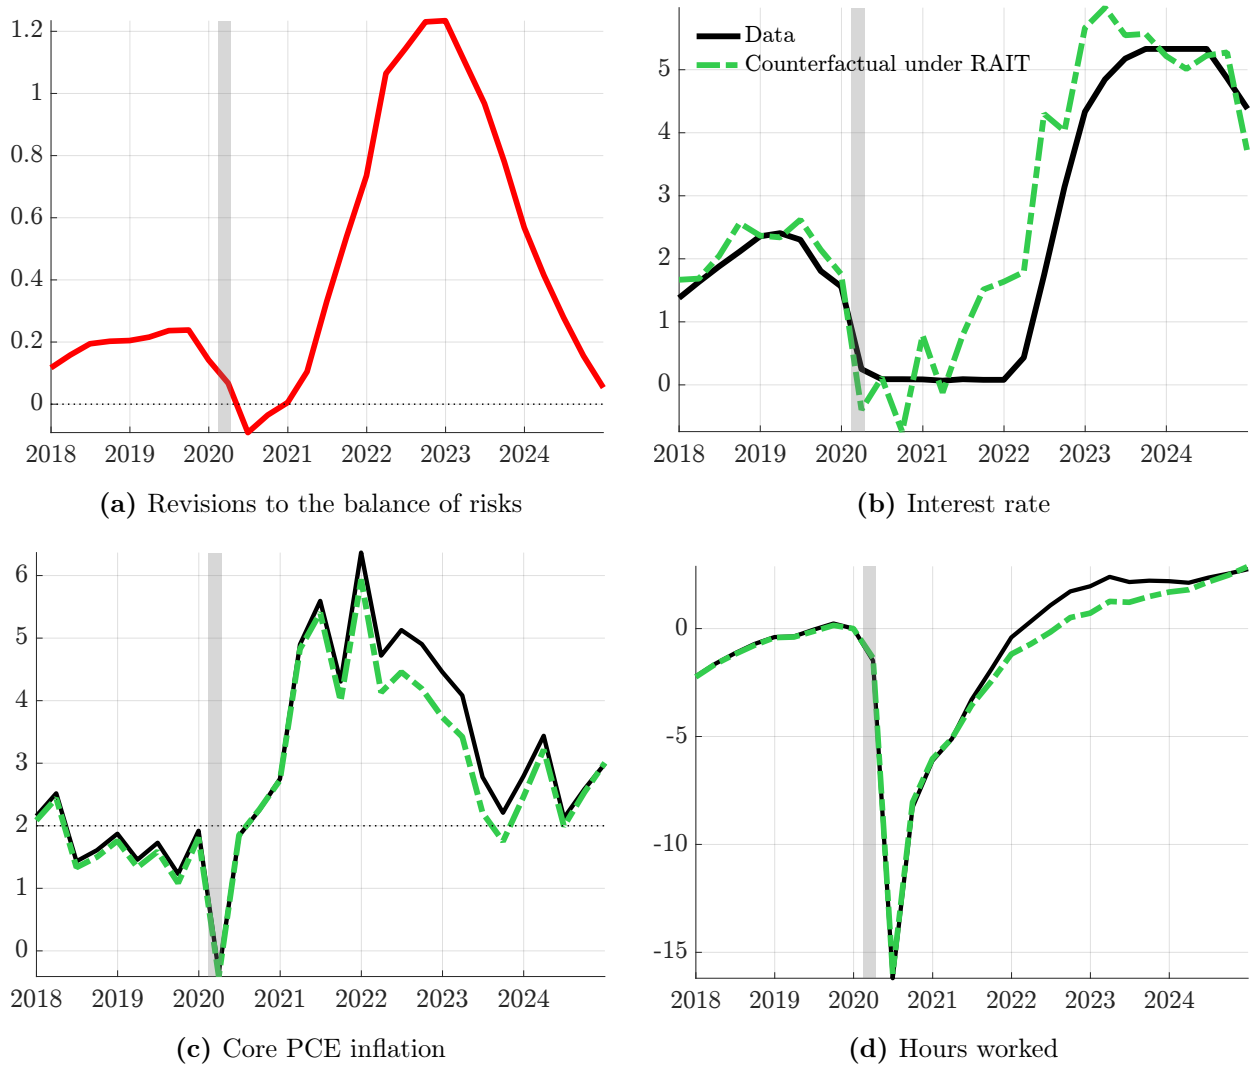
We illustrate the potential impact of the RAIT by evaluating counterfactual dynamics for key macroeconomic quantities.[25] Counterfactual paths are computed by first simulating variables from the beliefs representation of the model using only forward guidance shocks, which solve the system of equations Equation (23). Simulated paths are then added to observed data points. This additive approach is appropriate because the beliefs representation of the model is linear.

Counterfactuals for the federal funds rate, core PCE inflation rate, and hours worked implied by the RAIT are reported in Figure 6, alongside the actual realizations. The upper-left panel shows the estimated revisions to the balance of inflation risks, $\psi_{t|t} - \psi_{t|t-1}$, obtained from the forecasting model and used to simulate alternative dynamics under the RAIT. These estimates reveal that, as early as mid-2021, clear signs of rising upside inflation risk began to emerge. Upward revisions to inflation risk persist from mid-2021 through the end of 2023, after which they suddenly recede.

Starting in mid-2021, the RAIT-consistent policy rate began to exceed the actual policy rate. This positive gap reflects the prevailing upside risks to inflation, which call for a tighter monetary stance to offset the influence of more probable future inflation spikes on agents' current inflation expectations. A substantial tightening is prescribed in 2022, following a sudden increase in upside risks. It is important to note that the estimation of skewness in the inflation distribution—and its implications for the balance of risks—is conducted in real time. This makes abrupt revisions to risk assessments more frequent, especially during periods of heightened volatility such as the post-pandemic era.

The alternative policy path implied by the RAIT would have anchored inflation expectations and driven inflation to the 2% target one year earlier. RAIT would have delivered a slower, but

[25] 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.



(a) Revisions to the balance of risks

(b) Interest rate

(c) Core PCE inflation

(d) Hours worked

Figure 6: Macro dynamics under RAIT

Note: The plots report the estimated revisions to the balance of inflation risks and 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.

more balanced recovery of the labor market.

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

Table 2: Counterfactual Fed funds rate dynamics and forward guidance

	Fed funds rate		Counterfactual rate		Forward guidance	
	Q4 %	Cumulative bps	Q4 %	Cumulative bps	Q4 %	Cumulative bps
2021	0.080	−0.667	1.146	87.590	1.066	88.257
2022	4.333	425.333	5.350	420.433	1.017	−4.900
2023	5.330	99.667	5.244	−10.645	−0.086	−110.312
2024	4.380	−95.000	4.166	−107.715	−0.214	−12.715
Total	14.043	430.000	14.760	302.073	0.717	−127.927

Note: The table reports year-end and cumulative values of the realized Fed funds rate, their counterfactual path under the RAIT, as well as the forward guidance implied by the RAIT, from 2021 to 2024. Forward guidance is defined as $-(1-\rho)r_\pi \sum_{j=1}^{4} \hat{\pi}^\star_{t+j|t}$.

the Federal Open Market Committee (FOMC), even though the RAIT would have prescribed starting the tightening sooner and reaching that peak two quarters earlier.

We now illustrate how the prescriptions of the RAIT can be translated into communications about the path of future interest rates, or forward guidance. Table 2 compares the year-end values and the cumulative change of the realized Fed funds rate, the policy instrument controlled by the Fed, against the counterfactual interest rate under the RAIT. The last column isolates the contribution of central bank's communication via forward guidance (to interpreted as temporary deviations for the 2% target, in line with Equation (22)).

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 in 2021 up until through early 2023. 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. Consistent with the previous analysis, 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. That is, under the RAIT, the central bank could have started 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.

5.5 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.[26] 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'

[26]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.

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:

$$\hat{\pi}_t^{\text{FAIT}} = \rho_F \hat{\pi}_{t-1}^{\text{FAIT}} - (1-\rho_F)\left(\pi_t^{\text{Data}} - \bar{\pi}\right), \tag{24}$$

where π_t^{Data} stands for the inflation rate observed in the data and $\bar{\pi}$ represents the 50 basis point 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.[27] The monetary policy rule under the FAIT is given by:

$$\hat{r}_t = \rho r_{t-1} + (1-\rho)\left[r_x \hat{x}_t + r_{\Delta x}\Delta\hat{x}_t + r_\pi\left(\hat{\pi}_t - \hat{\pi}_t^{\text{FAIT}}\right)\right] + \varepsilon_t^r. \tag{25}$$

In Figure 7, 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

[27] 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.

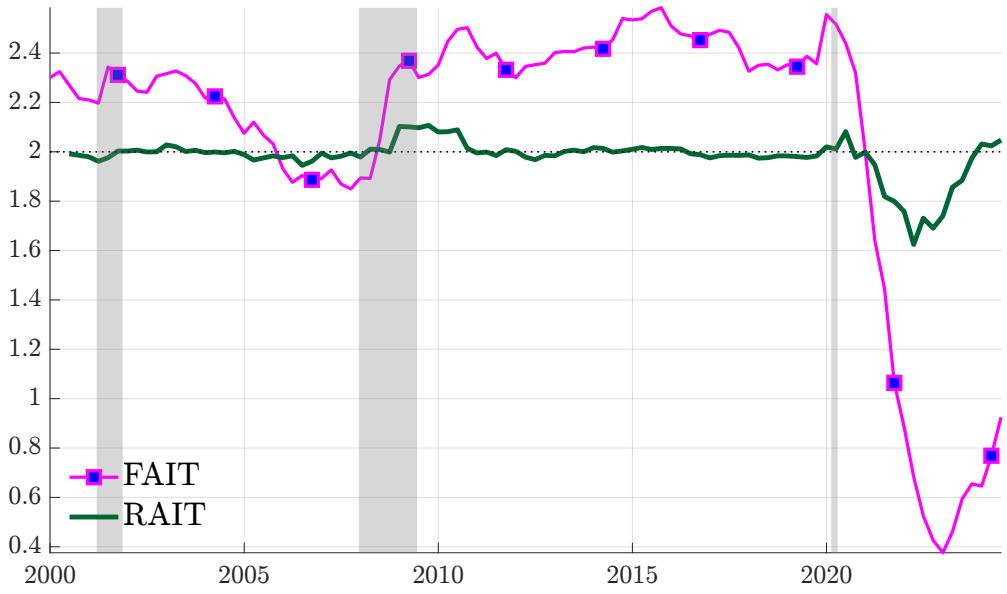


Figure 7: FAIT vs. RAIT

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.

communicate its intention to overshoot inflation. 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 (16). 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 continuous inflation misses, inflation volatility reached historical lows. As a result, the mean-mode wedge in Equation (16) – 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 ficial 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.

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.

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.

Newton, M. A. and A. E. Raftery (1994): "Approximate Bayesian inference with the weighted likelihood bootstrap," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 56, 3–26.

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.

Spiegelhalter, D. J., N. G. Best, B. P. Carlin, and A. Van Der Linde (2002): "Bayesian measures of model complexity and fit," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64, 583–639.

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 Log-linearized model with asymmetric risk

Let $\mathbf{Z}_t$ collect endogenous and (persistent) exogenous variables, and $\boldsymbol{\varepsilon}_t$ be a vector of i.i.d. exogenous shocks to the system. A stable Rational Expectations equilibrium is defined as a system of expectational difference equations of the form

$$E_t\left[f\left(\mathbf{Z}_{t+1}, \mathbf{Z}_t, \mathbf{Z}_{t-1}, \boldsymbol{\mathcal{E}}_t\right)\right] = 0, \tag{A1}$$

where f depends on agents' preferences, available technology, and constraints, and $\mathbb{E}_t$ denotes the expectation operator conditional on the information set available at time t.

Let us now assume $\boldsymbol{\mathcal{E}}_t = \exp(\boldsymbol{\varepsilon}_t)$ and $\boldsymbol{\varepsilon}_t = [\boldsymbol{\varepsilon}_t^s, \varepsilon_t^a]'$, where $\boldsymbol{\varepsilon}_t^s$ is a set of symmetric shocks—e.g., Normally distributed—while ε_t^a represents a potentially asymmetric shock with distribution. While we focus on the case with a single source of asymmetric risk, extending the analysis to multiple independent asymmetric shocks is straightforward and does not alter the underlying intuition.

Define $\boldsymbol{\mathcal{E}}^\star$ and $\mathbf{Z}^\star$ as the steady-state of the model. The first-order Taylor expansion of Equation (A1) around the (non-stochastic) steady state reads

$$\begin{aligned} f(\mathbf{Z}_{t+1}, \mathbf{Z}_t, \mathbf{Z}_{t-1}, \boldsymbol{\mathcal{E}}_t) \approx\ & f_{\mathbf{Z}_{t+1}}\left(\boldsymbol{\mathcal{E}}^\star, \mathbf{Z}^\star\right)\left(\mathbf{Z}_{t+1} - \mathbf{Z}^\star\right) + f_{\mathbf{Z}_t}\left(\boldsymbol{\mathcal{E}}^\star, \mathbf{Z}^\star\right) + \left(\mathbf{Z}_t - \mathbf{Z}^\star\right) \\ & + f_{\mathbf{Z}_{t-1}}\left(\boldsymbol{\mathcal{E}}^\star, \mathbf{Z}^\star\right)\left(\mathbf{Z}_{t-1} - \mathbf{Z}^\star\right) + f_{\boldsymbol{\mathcal{E}}_t}\left(\boldsymbol{\mathcal{E}}^\star, \mathbf{Z}^\star\right)\left(\boldsymbol{\mathcal{E}}_t - \boldsymbol{\mathcal{E}}^\star\right). \end{aligned}$$

where for any generic variable $\mathbf{X}$, $f_{\mathbf{X}}(\mathbf{X}) = \frac{\partial f(.)}{\partial \mathbf{X}}$ evaluated at $\mathbf{X}$. Defining log-deviations from steady-states with lower case letters and normalizing $\boldsymbol{\mathcal{E}}^\star = \mathbf{1}$ regardless of the specific shape of the asymmetry of risk, we obtain a linear rational model of the form in Equation (1).

A.2 Proof of Proposition 1

Proof. Starting from a linear solution of the form in Equation (1):

$$\mathbf{A}_0\mathbf{z}_t = \mathbf{A}_f E_t\mathbf{z}_{t+1} + \mathbf{A}_b\mathbf{z}_{t-1} + \mathbf{B}_s\boldsymbol{\epsilon}_t^s + \mathbf{b}_a\epsilon_t^a,$$

where $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}_s$ and $\mathbf{b}_a$ are independent of the skewness of the shocks in $\boldsymbol{\epsilon}_t^a$.

The solution of a linear rational model of the form in eq. Equation (1) can be written as

$$\begin{aligned}\mathbf{z}_t &= \boldsymbol{\Theta}_1 \mathbf{z}_{t-1} + \boldsymbol{\Theta}_0 \left[\boldsymbol{\epsilon}_t^s \ \epsilon_t^a\right]' + \boldsymbol{\Theta}_y \sum_{j=1}^{\infty} \boldsymbol{\Theta}_f^{j-1} \boldsymbol{\Theta}_z \left[E_t \boldsymbol{\epsilon}_{t+j}^s \ E_t \epsilon_{t+j}^a\right]'. \\ &= \boldsymbol{\Theta}_1 \mathbf{z}_{t-1} + \boldsymbol{\Theta}_0 \left[\boldsymbol{\epsilon}_t^s \ \epsilon_t^a\right]' + \boldsymbol{\Theta}_y \sum_{j=1}^{\infty} \boldsymbol{\Theta}_f^{j-1} \boldsymbol{\Theta}_z \boldsymbol{\Xi} E_t \epsilon_{t+j}^a.\end{aligned}$$

where the matrices $\boldsymbol{\Theta}_0$, $\boldsymbol{\Theta}_1$, $\boldsymbol{\Theta}_y$, and $\boldsymbol{\Theta}_z$ are functions of $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}_s$ and $\mathbf{b}_a$ (see, e.g., Sims, 2002, Sec. 4), whereas $\boldsymbol{\Xi}$ is a selection vector. The last row reflects the fact that $E_t \boldsymbol{\epsilon}_{t+j}^s = \mathbf{0}, \forall j$. In a fully symmetric setting, where $\varrho_{a,t} = 0$ and $E_t \epsilon_{t+j}^a = 0, \forall t, j$, the above expression simplifies to $\mathbf{z}_t = \boldsymbol{\Theta}_1 \mathbf{z}_{t-1} + \boldsymbol{\Theta}_0 \left[\boldsymbol{\epsilon}_t^s \ \epsilon_t^a\right]'$. □

A.3 Macroeconomic effects of revisions in the balance of risk

When shock distributions are skewed, 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 $X \sim \mathsf{f}(\mu, \sigma, \varrho)$ 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 f. Therefor, for $\epsilon_t^a \sim \mathcal{F}(0, \sigma, \varrho_{a,t})$,

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

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

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

Defining revisions in the asymmetry of risk j periods in advance as $\varrho_t^j = E_t \varrho_{a,t+j} - E_{t-1} \varrho_{a,t+j}$, implies that the h-step ahead impulse response of the system to a revision in the perceived asymmetry

of risk j-periods ahead can be defined as

$$E_t \mathbf{z}_{t+h} - E_{t-1}\mathbf{z}_{t+h} = \mathsf{M}(j,h) k\sigma^2 \varrho_t^j;$$

where $\mathsf{M}(j,h) = \boldsymbol{\Theta}_1^h \boldsymbol{\Theta}_y \boldsymbol{\Theta}_f^{j-1} \boldsymbol{\Theta}_z \boldsymbol{\Xi}$ is a convolution of the system matrices in Equation (2)

A.4 Proof of Proposition 2

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

$$\mathbf{A}_0 \mathbf{z}_t = \mathbf{A}_f \widetilde{E}_t \mathbf{z}_{t+1} + \mathbf{A}_b \mathbf{z}_{t-1} + \mathbf{B}\widetilde{\boldsymbol{\epsilon}}_t, \tag{A2}$$

where $\widetilde{E}_t$ denotes agents (potentially) distorted beliefs.

In this setting, agents' choices and equilibrium outcomes depend not only on current shocks but also on their beliefs about future shocks. Let $\widetilde{E}_t$ denote the *belief expectation* operator, defined as

$$\widetilde{E}_t \boldsymbol{\epsilon}_{t+j} = E_t \boldsymbol{\epsilon}_{t+j} + \boldsymbol{\varphi}_{t,t+j}, \tag{A3}$$

where $\boldsymbol{\varphi}_{t,t+j}$ captures agents' beliefs at time t about the mean of the distribution of certain shocks expected to hit the economy at time $t+j$. Under the maintained assumption that the shocks $\boldsymbol{\epsilon}_t$ are symmetric—i.e., $E_t \boldsymbol{\epsilon}_{t+1} = \mathbf{0}$ for all t—the term $\boldsymbol{\varphi}_{t,t+j}$ acts as a belief-based adjustment to the expected future shocks.

The solution of a linear rational model of the form in eq. Equation (A2) can be written as

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

where the matrices $\boldsymbol{\Theta}_0$, $\boldsymbol{\Theta}_1$, $\boldsymbol{\Theta}_y$, and $\boldsymbol{\Theta}_z$ are functions of $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}_s$ and $\mathbf{b}_a$ (see, e.g., Sims, 2002, Sec. 4). Therefore, given the definition of *beliefs expectations* in Eq. (A3), the rational

expectations solution of the system can be written as

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

Let $\boldsymbol{\varphi}_t^j = \boldsymbol{\varphi}_{t,t+j} - \boldsymbol{\varphi}_{t-1,t+j}$ denote a belief revision (or "dummy surprise"), capturing the change between periods $t-1$ and t in agents' beliefs about shocks occurring at time $t+j$. Belief distortions evolve as a reflection of agents' revised views, according to $\boldsymbol{\varphi}_{t,t+h} = \sum_{j=0}^{J} \boldsymbol{\varphi}_{t-j}^{h+j}$, where J is the maximum horizon over which agents project their distorted expectations. In line with Chahrour and Jurado (2018), we assume that agents do not anticipate future revisions to their beliefs—formally, $E_t \boldsymbol{\varphi}_{t+i}^j = 0$ for all $i \geq 1$.

It follows that revisions in beliefs lead to the following revision of current and expectations of future outcomes:

$$E_t \mathbf{z}_{t+j} - E_{t-1} \mathbf{z}_{t+j} = \boldsymbol{\Theta}_1^j \boldsymbol{\Theta}_y \boldsymbol{\Theta}_f^{j-1} \boldsymbol{\Theta}_z \boldsymbol{\varphi}_t^j, \quad h \geq 0. \tag{A4}$$

By setting

$$\varphi_t^j = \boldsymbol{\Xi} \boldsymbol{\varphi}_t^j = \varkappa \sigma^2 \varrho_{a,t}^j,$$

the impact of revisions about the asymmetry of shocks' distributions (see Equation (4)) is equivalent to shifts in beliefs about future realizations of the shocks (Equation (A4)) and the responses of the two models are observationally equivalent up the first order. □

B Solving the New Keynesian model with cost-push 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

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

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

where $i_t = -\log Q_t$ is the short-term nominal rate and κ is the slope of the Phillips curve. To simplify the exposition, we assume that cost-push shocks follow $u_t \sim \text{i.i.d.}\,\mathcal{N}(0,\sigma_u)$.[28]

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,

[28] Although we adopt the standard normality assumption, the results only require that the shock is univariate and symmetric. In this case, the mean and the central forecast coincide, implying $E_t u_{t+j} = 0$ for all $j \geq 1$, which is sufficient to derive the solution under symmetry.

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

$$\hat{x}_t = -\frac{\kappa}{\alpha_x} \bar{p}_t, \tag{B3}$$

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 (B3) into

Equation (B2) we can recast the Phillips curve in as

$$\bar{p}_t = a\bar{p}_{t-1} + a\beta E_t\bar{p}_{t+1} + au_t \tag{B4}$$

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

B.2.1 The symmetric case

The stationary solution to Equation (B4) can be obtained by using the method of undetermined coefficients by conjecturing a solution of the form

$$\bar{p}_t = \eta\bar{p}_{t-1} + \lambda u_t, \tag{B5}$$

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

$$E_t\bar{p}_{t+1} = \eta\bar{p}_t. \tag{B6}$$

Substituting the expectations implied by the stationary solution yields

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

Solving for η and λ, we obtain

$$\eta = \frac{1-\sqrt{1-4\beta a^2}}{2a\beta}, \tag{B7}$$

$$\lambda = \frac{a}{1-a\beta\eta}, \tag{B8}$$

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

$$\hat{x}_t = \eta\hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\lambda u_t, \tag{B9}$$

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

Implementation. The optimality condition in Equation (B3) allows us to write the IS curve in Equation (B1) 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 (B6), we obtain the following *optimal monetary rule*:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t. \tag{B10}$$

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).

B.2.2 The asymmetric case

Let us now consider the case where the stochastic process driving the cost-push shock is no longer symmetric. Specifically, assume that the shocks follow $\tilde{u}_t \sim$ i.i.d. $\mathsf{f}(0, \sigma_u, \varrho_{u,t})$, where $\mathsf{f}(0, \sigma_u, \varrho_{u,t})$ denotes a general unimodal distribution centered at zero but asymmetric when $\varrho_{u,t} \neq 0$. In this setting, the perceived risk associated with the distribution evolves over time as agents update their views about the balance of upside and downside risks. As a result, the expected value of future shocks reflects these changes in perceived asymmetry. While the mode—the most likely outcome—remains at zero, the mean shifts with the asymmetry, so that $E_t\tilde{u}_{t+j} = \varkappa\sigma_u E_t\varrho_{u,t+j}$.

To simplify the characterization of optimal monetary policy, we assume that agents receive news about the skewness of the distribution only up to $J = 1$ period ahead. Beyond that, they assume all risk is symmetric. In this environment, at time t, agents observe the current realization of the shock and form expectations about future realizations based on their current perception of asymmetry. This implies $E_t\tilde{u}_{t+1} = \varkappa\sigma_u E_t\varrho_{u,t+1}$, while $E_t\tilde{u}_{t+j} = 0$ for all $j > 1$. Moreover,

since $\varrho_{u,t+1} = \varrho^0_{u,t+1} + \varrho^1_{u,t}$, and assuming agents do not anticipate further revisions in perceived asymmetry (i.e., $E_t\varrho^0_{u,t+1} = 0$), it follows that $E_t\tilde{u}_{t+1} = \varkappa\sigma_u\varrho^1_{u,t}$, where $\varrho^1_{u,t} = E_t\varrho_{u,t+1} - E_{t-1}\varrho_{u,t+1}$. Expectations of future shocks therefore reflect the latest update in perceived risk asymmetry.

We now turn to the alternative *beliefs representation* of the information structure. In this framework, agents observe the realization of the shock at time t and update their beliefs about both the current shock, denoted by ψ^0_t, and the next period's shock, denoted by ψ^1_t. These belief distortions are mapped into expectations as follows:

$$E_t\tilde{u}_t = u_t + \psi^0_t + \psi^1_{t-1} = \tilde{u}_t, \tag{B11}$$

$$E_t\tilde{u}_{t+1} = E_t\psi^0_{t+1} + \psi^1_t = \psi^1_t. \tag{B12}$$

Equation (B11) imposes the restriction $\psi^0_t = -\psi^1_{t-1}$, requiring that the actual realization of the shock is not affected by prior beliefs—or equivalently, that earlier beliefs do not necessarily materialize and can turn out to be pure noise. Equation (B12) reflects the assumption that agents' expectations about future shocks are fully captured by the dummy belief ψ^1_t, and that they do not anticipate further updates in those beliefs, i.e., $E_t\psi^0_{t+1} = 0$. Furthermore, by setting $\psi^1_t = \varkappa\sigma_u\varrho^1_{u,t}$, we have that, up to a first-order approximation, the dynamics of a model with evolving perceptions of asymmetric risk can be equivalently represented by a model with symmetric shocks and belief distortions captured by dummy surprises.

In this setting, shifts in asymmetric risk affect current allocations and as a consequence next period's price level, $E_t\bar{p}_{t+1}$, are potentially distorted. We postulate that the price process under optimal policy evolves as:

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

and it follows that

$$E_t\bar{p}_{t+1} = \eta\bar{p}_t + \lambda E_t u_{t+1} + \zeta E_t\psi^1_{t+1}, \tag{B14}$$

$$= \eta\bar{p}_t + \lambda\psi^1_t \tag{B15}$$

since $E_t u_{t+1} = \psi_t^1$ and $E_t \psi_{t+1}^1 = 0$. Substituting the expectations implied by the stationary solution under optimal policy (Equation (B4)) yields

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

Solving for η, λ and ζ we obtain

$$\begin{aligned}
\eta &= \frac{1-\sqrt{1-4\beta a^2}}{2a\beta}, \\
\lambda &= \frac{a}{1-a\beta\eta}, \\
\zeta &= \frac{a\beta\lambda}{1-a\beta\eta}.
\end{aligned}$$

It follows that the equilibrium process for the output gap is

$$\hat{x}_t = \eta\hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\left[\lambda u_t + \zeta\psi_t^1\right], \tag{B16}$$

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

Implementation. 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, 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 (B10), and the last term captures how the central bank must adjust the policy rate in response to shifts in perceived risk asymmetry, specifically accounting for how these changes influence agents' expectations.[29]

B.2.3 Asymmetric case with an unwitting central bank

We now consider the case in which the distribution of cost-push shocks 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. Cost-push shocks' asymmetry is captured by agents receiving surprises ψ_t^1, tilting their expectations about future realizations of the shocks away from the central scenario in every period t. We compare this scenario against the case of optimal monetary policy under symmetry. This exercise can be also 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.

We start from the optimality condition between output gap and inflation

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^s, \tag{B17}$$

where p_t^s is the price level under full symmetry ($\psi_t^j = 0, \forall t, j$) defined

$$\bar{p}_t^s = \eta\bar{p}_{t-1} + \lambda u_t; \tag{B18}$$

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

[29] It can be shown that the conditions for determinacy are the same as in the model with symmetric risk.

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 (B17), and substituting into the Phillips curve expressed in terms of price level yields

$$\bar{p}_t = a\bar{p}_{t-1} + a\beta E_t \bar{p}_{t+1} + b\tau\psi_t^1 + au_t \tag{B19}$$

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 (B19) to obtain

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

Solving for the coefficients we obtain

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

We recover τ by taking the difference between the price level, $\bar{p}$, and the price level under the symmetric policy – Equation (B18). 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

$$\zeta = \tau = \frac{a\beta\lambda}{1 - a\beta\eta - b}. \tag{B20}$$

The equilibrium process for the output gap is

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^s, \tag{B21}$$

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,

$$\hat{x}_t \neq \frac{\kappa}{\alpha_x}\bar{p}_t. \tag{B22}$$

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.[30] 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.

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

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

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

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

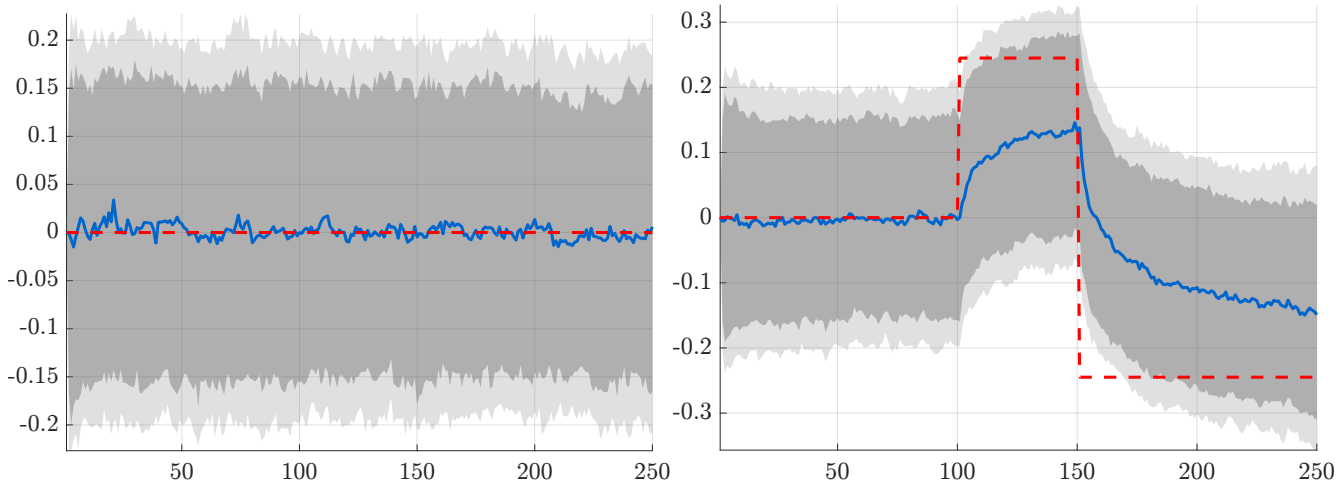
We report the results of this exercise in Figure D1. Specifically, for DGP1 to DGP4 we report in blue the estimated asymmetry, with 68% and 90% credible sets represented by shades of gray, against the simulated parameter, in red. For DGP5 and DGP6 we report the distribution of the difference between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

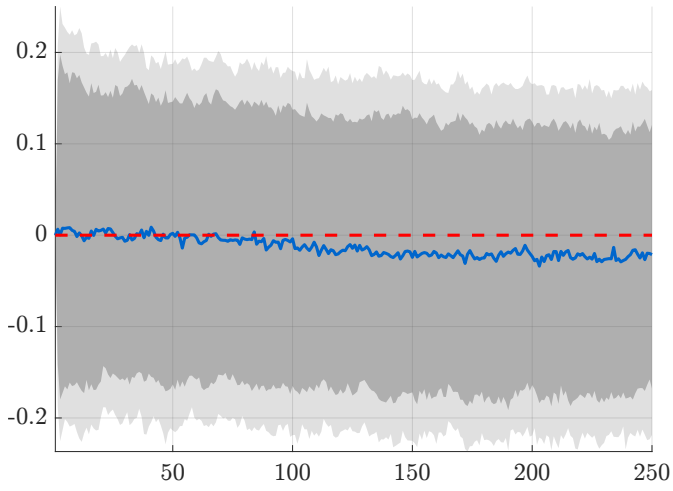
Finally, in DGP5 and DGP6 we simulate the asymmetry parameter to vary over time, as the other two parameters. The two DGPs differ in the covariance structure of the parameters: DGP5 assumes independent innovations to the processes, whereas DGP6 assumes a full covariance matrix, with the covariance between 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.

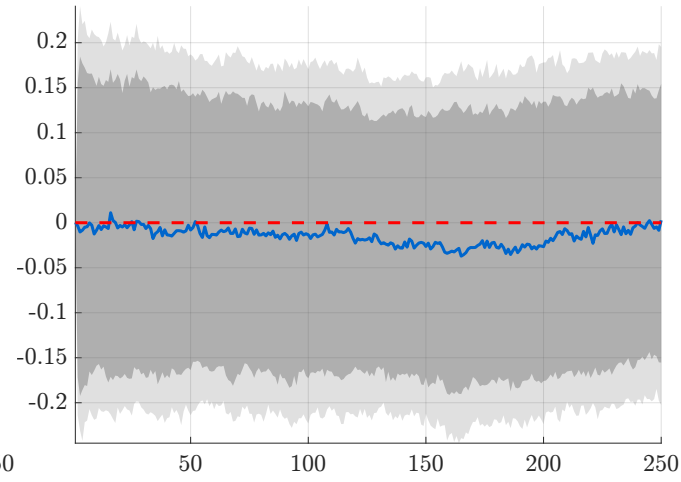


(a) DGP1

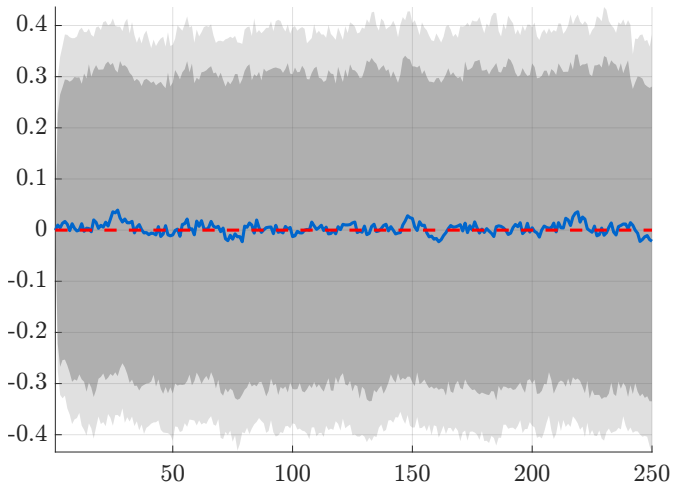
(b) DGP2



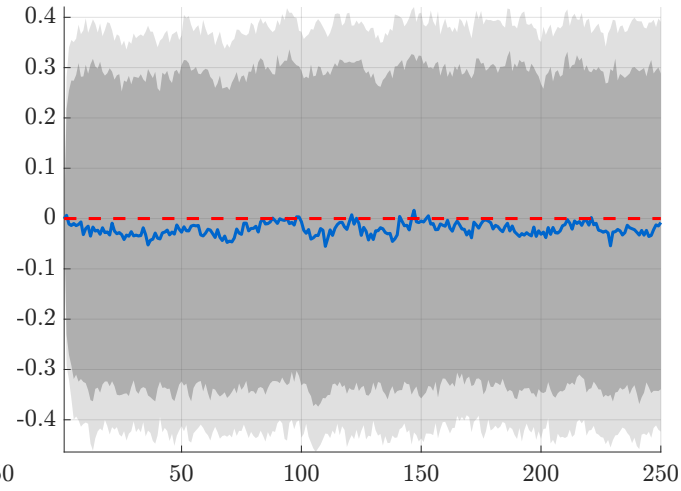
(c) DGP3



(d) DGP4



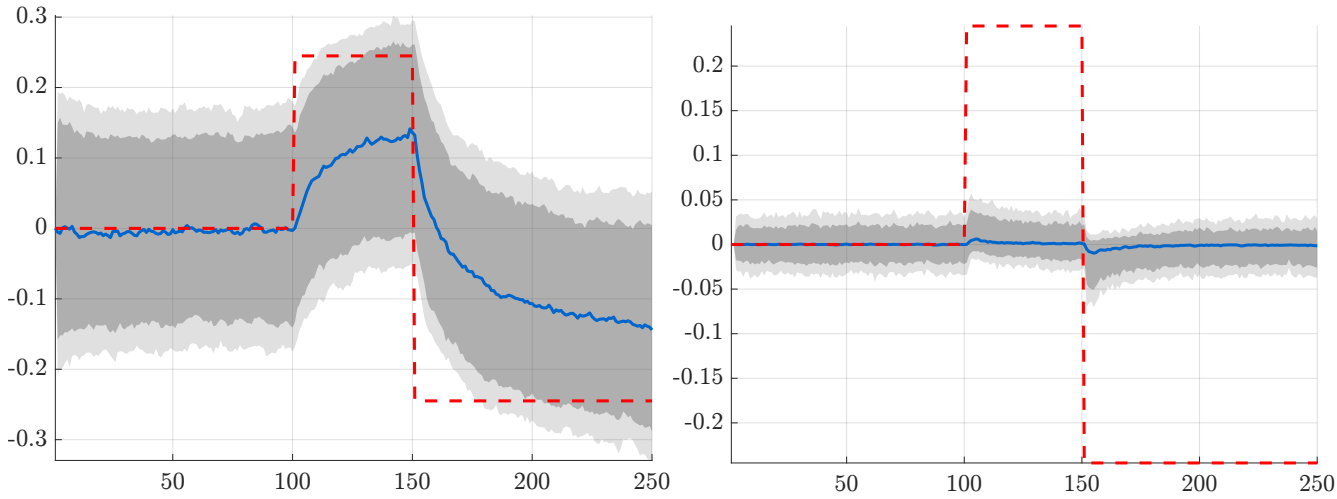
(e) DGP5



(f) DGP6

Figure D1: Estimated asymmetry

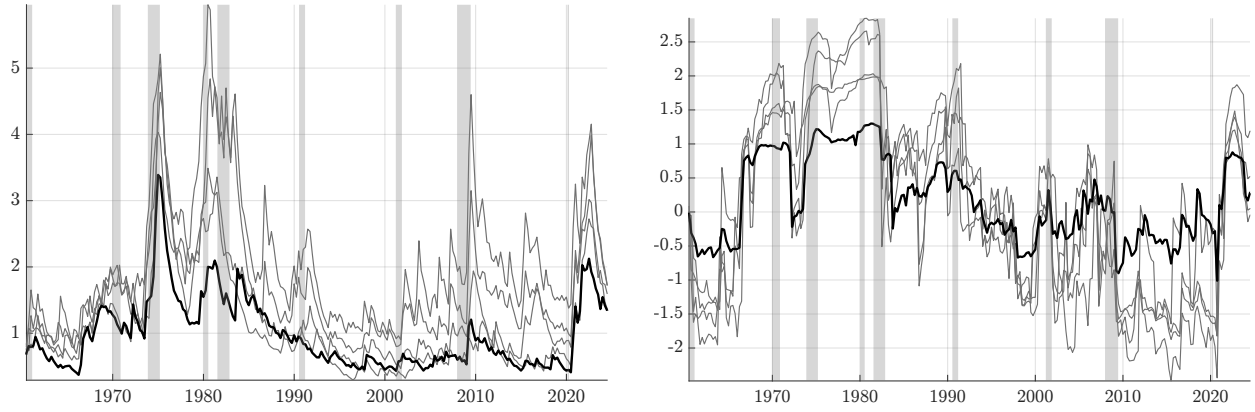
Note: The panels reports the estimated paths for the asymmetry parameters (blue) with the associated 68% and 90% credible sets. The asymmetry under the DGP is reported in red. For DGP 5 and 6 we report deviations of the estimated parameter from the simulated values. We consider T=250 observations for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



(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.

E Evidence of time variation in inflation risk for other inflation measures

In this appendix we report additional results about the time variation in the risks of policy relevant measures of inflation.

Within the framework introduced in Section 4, we cab formally test the null hypothesis that ϱ_t remains constant over the entire sample. Table E1 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.

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.

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).

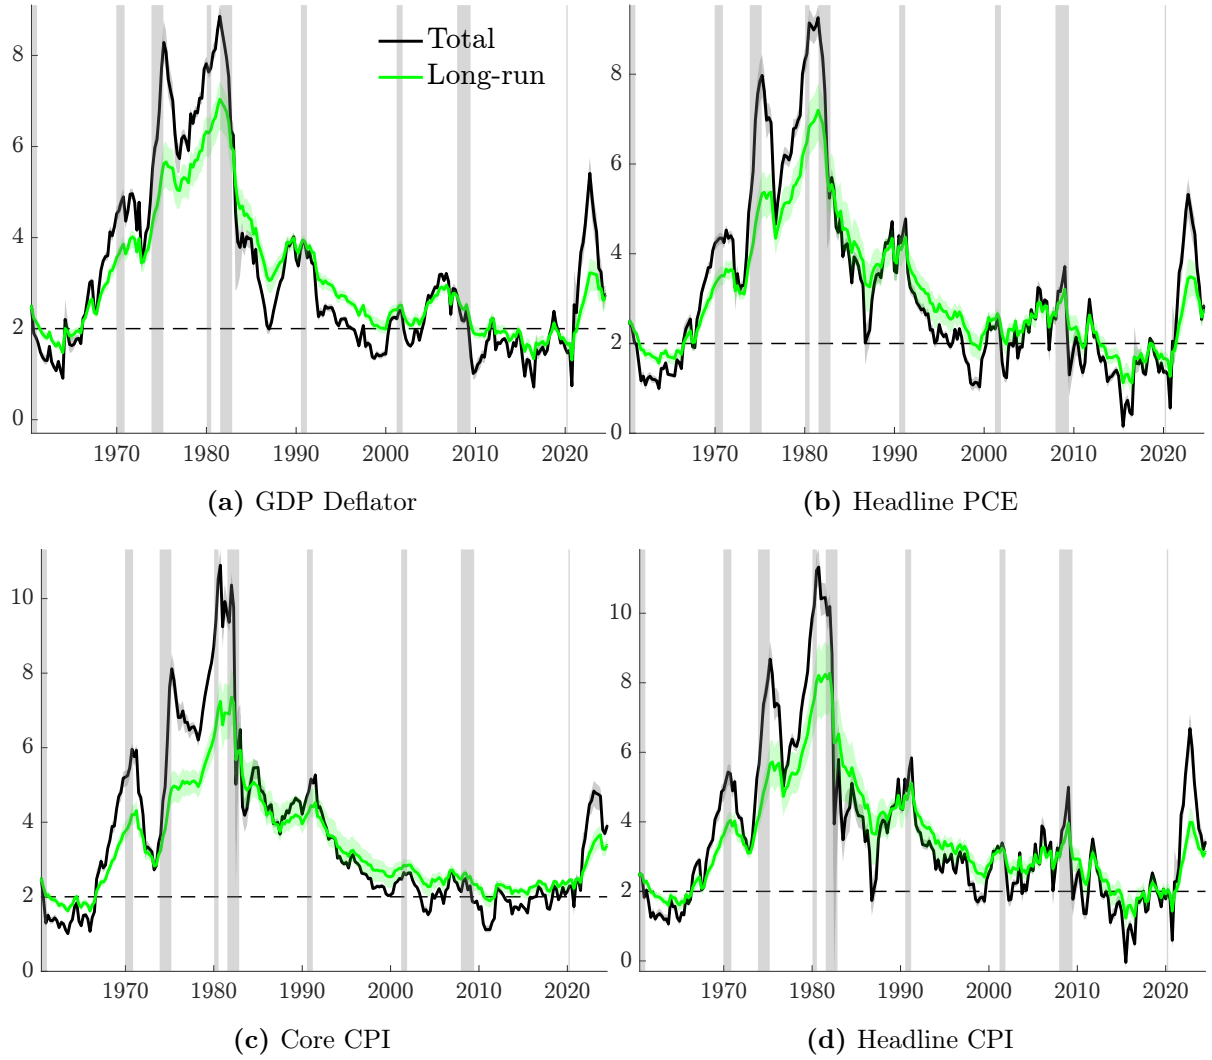
Figure E1 shows the estimated dynamics of inflation volatility and skewness across the different

Table E1: Time variation in higher order moments

	Q	Q^*	N	Q	Q^*	N
	Core PCE			**Headline PCE**		
	Homoskedastic					
Shape	367.31***	371.60***	4.18***	303.820***	307.370***	6.460***
	Heteroskedastic					
$Scale^2$	369.36***	373.67***	1.50***	566.190***	572.800***	2.330***
Shape	35.65***	36.07***	0.79***	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***
	GDP Deflator					
	Homoskedastic					
Shape	637.470***	644.910***	5.690***			
	Heteroskedastic					
$Scale^2$	597.120***	604.090***	4.050***			
Shape	154.150***	155.950***	2.780***			

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%.

measures, highlighting in blacks that of core PCE. Two comments are in order. First, the dynamics of the two moments is extremely similar for all measures. With varying magnitudes, volatilities spike around recessions, and remain persistently high soon after. Skewness follow humped-shape patterns in the 1970s and 1980s, then moving downward since the 1990s, remaining negative until the pandemic period. Second, it's important to note that, among all these measures, core PCE shows the least variation in both volatility and skewness, appearing to be the more stable measure of price dynamics.



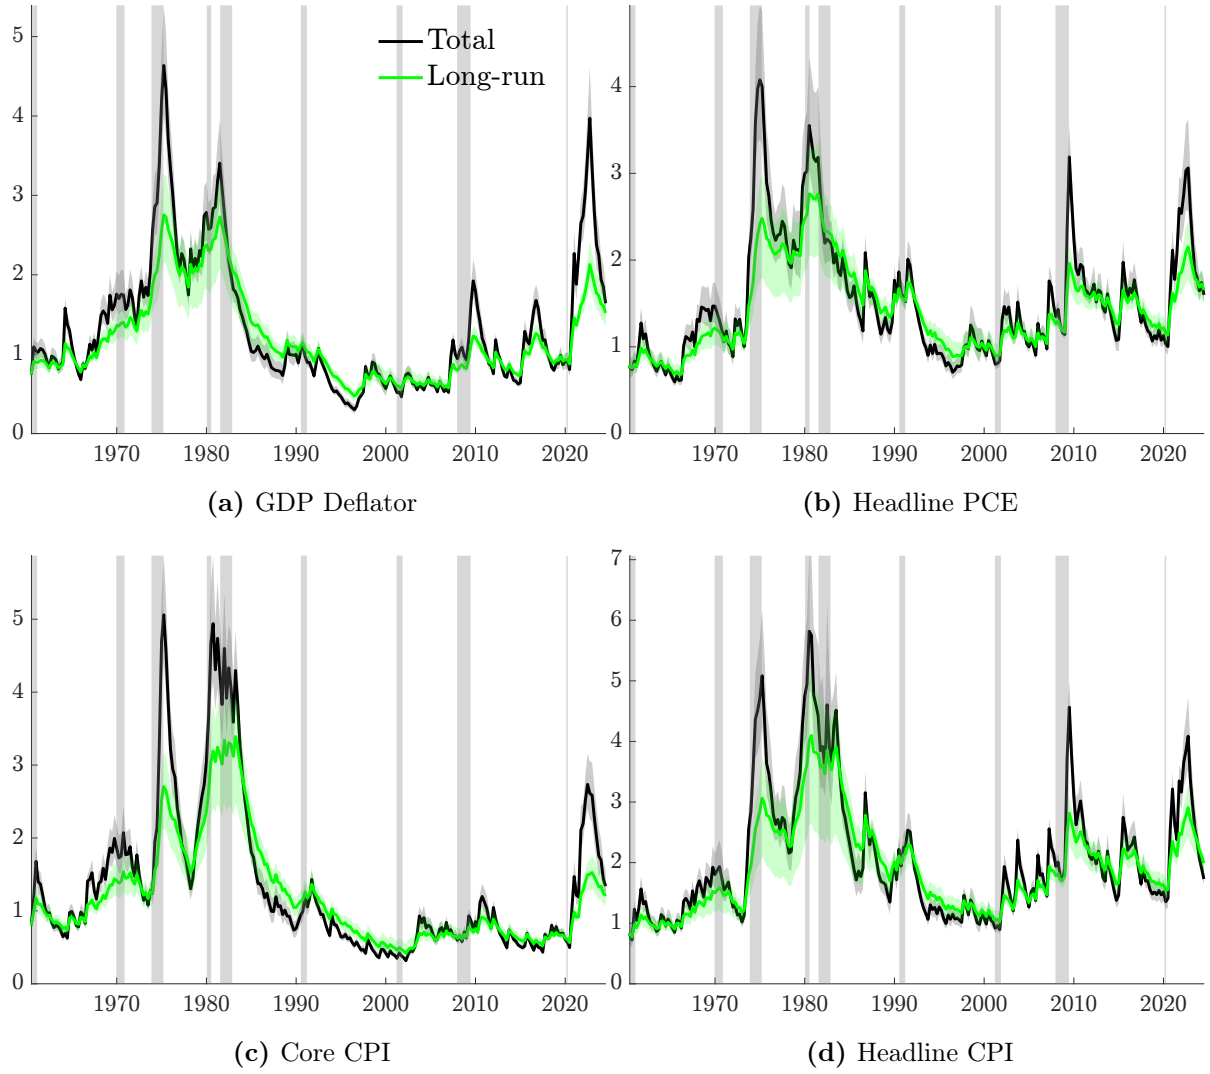
(a) GDP Deflator **(b)** Headline PCE

(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.

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

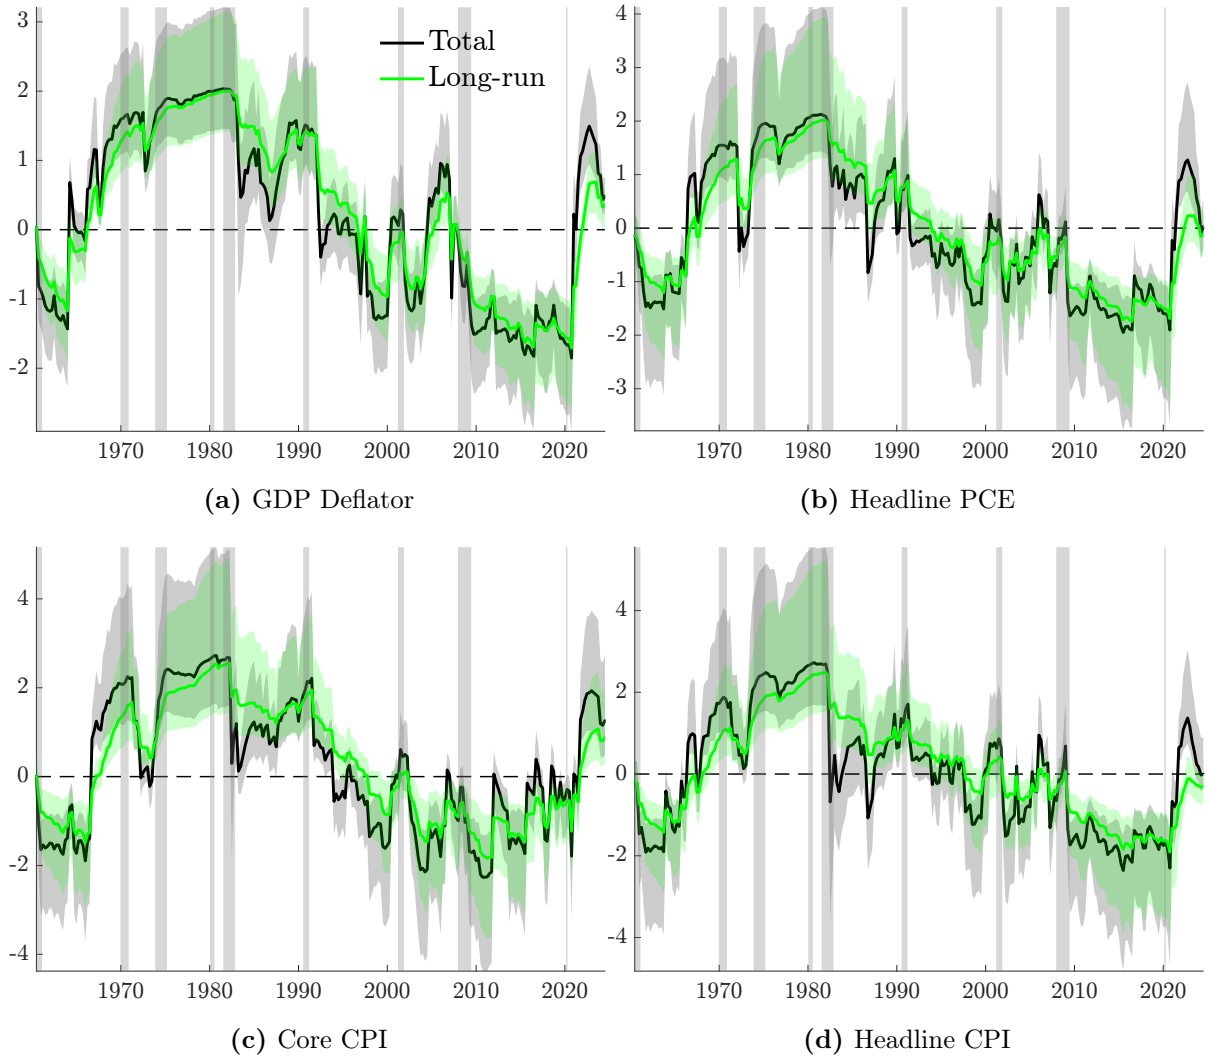


(a) GDP Deflator **(b)** Headline PCE

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

Figure 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.

Table F1: 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%.

F Additional results

Table F2: Out-of-sample comparison - *Student t*

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

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

Table F3: Out-of-sample comparison - $\varrho = 0,\ \forall t$

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

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

Table F4: Parameters estimates for the econometric model in Section 4

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

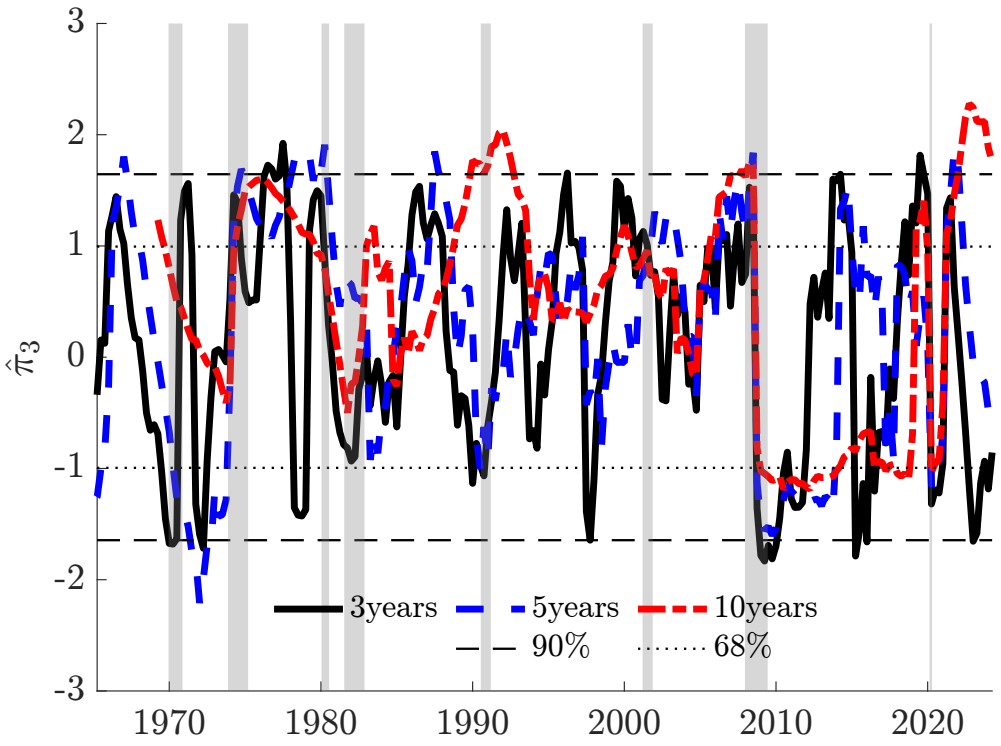


Figure F1: Bai and Ng (2005) rolling tests

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.