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

Andrea De Polis† Leonardo Melosi‡ Ivan Petrella§

This draft: July 2026
First draft: June 2024
Click here for the latest version

**Abstract**

Time-varying asymmetric inflation risks generate persistent stagflationary effects. A quantitative general equilibrium model with time-varying skewness in the distribution of cost-push shocks matches these effects. Central to the analysis is a representation theorem that provides a tractable characterization of a broad class of models with asymmetric shock distributions. The theorem enables a closed-form characterization of optimal monetary policy, prescribing that the central bank lean against the balance of inflation risks, while rendering quantitative general-equilibrium models with time-varying risks amenable to counterfactual and scenario analysis.

**Keywords:** Balance of risks, optimal monetary policy, asymmetric beliefs, policy trade-offs, news and noise, risk-adjusted inflation targeting, geopolitical risks.

**JEL codes:** E52, E31, C53.

*We would like to thank Guido Ascari, Gianluca Benigno, Gabriele Galati, Domenico Giannone, Elmar Mertens, Emmanuel Moench, Roberto Motto, Matthias Rottner, and the seminar participants at the DIW Berlin, University of Maryland, Johns Hopkins University, Federal Reserve Bank of St. Louis, Federal Reserve Bank of Chicago, the Federal Reserve Bank of Cleveland, the De Nederlandsche Bank, the University of Lausanne, the European Central Bank, 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, the Workshop on Empirical Monetary Economics at OFCE Paris, the Inflation: Drivers and Dynamics 2025 Conference at the ECB, and the ESCB Cluster on Monetary Economics at the Banque de France. 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 Banco de España or those of the Eurosystem.

†Banco de España. andrea.depolis@bde.es
‡European University Institute & 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, remarks delivered at the 2022 Jackson Hole Symposium

# 1 Introduction

The global economy has undergone a swift transition from a period of relative geopolitical stability to one characterized by heightened international tensions and recurrent supply-side disruptions. These have significantly changed the balance of macroeconomic risks, increasing the likelihood of both inflationary pressures and sharp economic slowdowns. Yet, despite their growing importance, the implications of these structural changes for macroeconomic dynamics and for the design of stabilization policies remain insufficiently understood.

This paper develops a tractable general equilibrium framework to examine how changes in the balance of risks propagate through the economy and to characterize the optimal monetary policy response. Central to our analysis is a new representation theorem that applies to a broad class of dynamic general equilibrium models in which the distribution of future structural shocks may be asymmetric and where this asymmetry evolves stochastically over time. In each period, agents receive news about the future distribution of shocks, including changes in skewness and tail risks. The theorem establishes that, to a first-order approximation, this type of models admit an equivalent *beliefs representation*.

In this representation, news about the evolving asymmetry of future shocks is translated into belief shocks. These shocks capture revisions in the mean–mode wedge of future shock distributions, thereby summarizing how changes in the balance of risks affect agents' expectations. The beliefs representation exactly characterizes the first-order macroeconomic effects of anticipated changes in risk asymmetry. We show that these effects account for most of the response in a nonlinear baseline New Keynesian model solved using global methods.

We exploit the tractability of the beliefs representation to augment a quantitative general equilibrium model in the tradition of Christiano et al. (2005) and Smets and Wouters (2007) with

news-driven revisions in the distribution of future cost-push shocks. We discipline these revisions using estimated moments of potentially skewed predictive densities for U.S. core PCE inflation. The model matches the macroeconomic effects of inflation skewness shocks estimated using local projections. We define these shocks as revisions in the estimated asymmetry of inflation predictive densities, orthogonal to current macroeconomic shocks. Upward revisions in inflation skewness generate a stagflationary pattern: inflation rises, aggregate demand weakens, and policy rates respond only modestly.

Quantitatively, the model reproduces the sizable stagflationary pressures observed in the data. Beyond validating the model with asymmetric cost-push risks, this result shows that first-order effects captured through the beliefs representation account for most macroeconomic effects of exogenous shifts in the balance of inflation risks observed in the data. Moreover, this finding supports the view that inflation skewness shocks largely stem from supply disruptions that generate cost-push pressures. Indeed, the inflation skewness estimated in the data closely tracks the fragmentation index proposed by Fernández-Villaverde et al. (2024).

The tractability of the beliefs representation also allows us to solve the Ramsey problem of optimal monetary policy under stochastic asymmetric risk in closed form and to compare it directly with the benchmark case of symmetric shocks studied by Clarida et al. (1999) and Woodford (2003). This analysis complements and extends the standard treatment of optimal monetary policy under Gaussian shocks by showing, in closed form, how fluctuations in the balance of inflation risks amplify policy trade-offs and alter the central bank's optimal policy response.

Optimal monetary policy actively *leans against asymmetric inflation risks*. When the distribution of future cost-push shocks tilts to the upside—so that positive shocks become relatively more likely—inflation expectations rise, calling for a monetary tightening that limits the increase in inflation by curbing aggregate demand. Conversely, when the distribution tilts to the downside—making negative shocks relatively more likely—policy should become more accommodative to support demand and mitigate disinflationary pressures. These predictions are consistent with the evidence in Cieslak et al. (2026), who show that FOMC members often favor a more forceful policy stance when they perceive the balance of risks to be tilted.

Motivated by these normative results, we study a novel monetary policy strategy—Risk-Adjusted Inflation Targeting (RAIT)—designed to offset the effects of shifts in the balance of inflation risks on expectations and preserve their anchoring. Under RAIT, the central bank adjusts both its current policy stance and its communication about the future path of policy rates to lean against fluctuations in inflation risk asymmetry.

We use the beliefs-based representation of the quantitative structural model to compare observed hours worked, core PCE inflation, and the federal funds rate with the counterfactual paths implied by a central bank implementing RAIT during the COVID and post-COVID periods. In this counterfactual exercise, the central bank systematically leans against *real-time* estimates of asymmetric inflation risks.

We find that the policy implied by the RAIT follows a three-phase pattern. In the first phase (2021Q2–2022Q3), the central bank tightens earlier than in the data, raising rates sooner and using forward guidance to signal a sustained tightening cycle ahead. This earlier and stronger response prevents inflation expectations from drifting upward in an environment characterized by heightened positive skewness in inflation risks. In the second phase (2022Q4–2023Q4), as tightening continues, forward guidance begins to signal that the hiking cycle is approaching its end. In the third phase (2024Q1–2024Q4), once inflation risks have receded, policy normalization begins and forward guidance becomes neutral, lowering the expected path of future policy rates.

**Literature Review** This paper contributes to the large literature studying the role of uncertainty and its time variation in driving business cycles (e.g., Bloom, 2009; Fernández-Villaverde et al., 2011; Jurado et al., 2015; Bloom et al., 2018).[1] By modeling uncertainty as the conditional variance of shocks, this literature has primarily focused on symmetric changes in the balance of risks. In these frameworks, uncertainty shocks are typically represented as mean-preserving increases in volatility that widen the distribution of future outcomes symmetrically around its center. We depart from this approach by studying changes in the asymmetry of shock distributions, and hence fluctuations in the relative importance of downside and upside risks for macroeconomic dynamics.

[1] See Fernández-Villaverde and Guerrón-Quintana (2020) for a comprehensive review.

A key contribution of this paper is to introduce a tractable approach to study time-varying asymmetric risks within quantitative general equilibrium models. The beliefs representation builds on the pioneering work of Chahrour and Jurado (2018, 2022), who show that a model with noisy signals about future realizations of exogenous variables is observationally equivalent to a model with news shocks. We depart from their framework by deriving a representation theorem that maps noisy signals about future shock distributions into revisions in the mean–mode wedge. This wedge captures, to a first-order approximation, changes in the skewness of shock distributions and summarizes the effects of asymmetric risks on equilibrium allocations. The resulting representation allows models with time-varying risk asymmetry to be solved using standard linear methods, making it possible to anchor belief shocks to real-time measures of the balance of inflation risks. Institutional evidence in Aruoba and Drechsel (2026) supports the relevance of this approach by showing that Federal Reserve staff forecasts combine modal projections with narrative assessments that track the evolving balance of risks around the outlook.

Our paper also contributes to the literature on monetary policy and macroeconomic risks. Existing work has emphasized risk management near the effective lower bound (Evans et al., 2016), the effects of monetary policy on the distribution of macroeconomic outcomes (Forni et al., 2024), and the policy challenges created by inflation scares and persistent cost-push disturbances (Erceg et al., 2024). We study a different but related problem: how policy should respond when the balance of inflation risks evolves stochastically and shifts expectations even before shocks materialize. By exploiting the beliefs representation, we show that this environment can be analyzed within a tractable linear-quadratic framework and derive the optimal policy in closed-form. The optimal response requires the central bank to lean against shifts in the balance of inflation risks.

The tractability of our approach also allows us to study monetary policy strategies that target inflation risks directly. In doing so, we extend the literature on optimal inflation targeting (Giannoni and Woodford, 2004) by formulating the policy problem in an environment where the balance of macroeconomic risks evolves stochastically. Our Risk-Adjusted Inflation Targeting strategy operationalizes the principle that policy should lean against the balance of inflation risks. While RAIT belongs to the class of forecast-based inflation-targeting strategies studied by Svensson (1997), it

extends that framework by allowing the central bank to respond to forward-looking endogenous variables that capture the interaction between policy, expectations, and evolving risk perceptions. In this sense, RAIT highlights the importance of responding preemptively to real-time shifts in the balance of inflation risks rather than reacting only after those risks materialize in realized inflation. Consistent with this mechanism, Cieslak et al. (2026) empirically link policymakers' perceptions of tail risks to tighter policy stances that cannot be justified by baseline forecasts alone.

On the empirical side, the paper relates to the growing literature on time-varying asymmetric risks in macroeconomic outcomes. A recent empirical literature initiated by Adrian et al. (2019) studies asymmetric risks in the predictive distribution of macroeconomic variables arising from macrofinancial conditions (see also Adrian et al., 2020, 2022). López-Salido and Loria (2024) show that inflation also displays asymmetric and time-varying risks. Related approaches measure inflation tail risks using option data (Dew-Becker, 2022; Hilscher et al., 2026), surveys of inflation expectations (Andrade et al., 2014; Mouabbi et al., 2025), or Markov-switching models (e.g., Le Bihan et al., 2024). These empirical approaches complement structural accounts in which changes in policy regimes can tilt the inflation outlook, including the possibility of switching to fiscal dominance (Bianchi, 2016; Bianchi and Melosi, 2017) or shifts in the monetary–fiscal policy mix (Kliem et al., 2016). Following Delle Monache et al. (2024), we estimate asymmetric inflation risks in real time within an unobserved-components framework that extends Stock and Watson (2007) to allow for time-varying skewness in the predictive distribution of inflation. Applying this procedure to a long postwar sample, spanning several low- and high-inflation regimes, allows us to identify revisions in the balance of inflation risks and trace their dynamic macroeconomic effects. We show that an increase in upside inflation risk generates stagflationary dynamics: inflation rises, aggregate demand weakens, and policy rates increase only modestly. This transmission differs from the evidence on business-cycle skewness in Salgado et al. (2019), where shifts in downside risk propagate primarily through demand-like contractions (see also Iseringhausen et al., 2026).

More broadly, our evidence extends work showing that expectations may be an independent source of aggregate fluctuations: while Coibion and Gorodnichenko (2026) emphasizes the macroeconomic effects of shifts in expectations, Ascari et al. (2023, 2024) show that changes in the distribu-

tion of inflation expectations—and in particular in expectation uncertainty—matter for aggregate dynamics. We focus on shifts in the asymmetry of the predictive distribution of inflation risk and embed them in a structural framework suited for policy analysis.

Finally, our focus on asymmetric risks relates to studies, starting with the seminal contribution of Barro (2006), that emphasize rare disasters as a source of large macroeconomic downturns (e.g., Gabaix, 2008, 2012; Gourio, 2008, 2012; Jordà et al., 2024). This literature typically relies on cross-country evidence to estimate the probability of severe, low-frequency disasters and studies their implications for aggregate fluctuations and asset prices. In contrast, our framework focuses on moderate but recurrent shifts in the balance of macroeconomic risks. Using real-time estimates, we document that these shifts are pervasive, economically meaningful, and systematically related to macroeconomic dynamics and monetary policy trade-offs.

**Structure** The remainder of the paper is organized as follows. Section 2 explains how shifts in risk affect macroeconomic allocations in a general DSGE framework and introduces the beliefs representation of a model with asymmetric risks. Section 3 provides empirical evidence on the relevance of inflation skewness in the U.S. In Section 4 we document the macroeconomic consequences of a change in the balance of inflation risks and compare these dynamics with those obtained from a general equilibrium model equipped with asymmetric cost-push shocks. Next, Section 5 characterizes optimal monetary policy within a New Keynesian model featuring asymmetric macroeconomic risks. In Section 6, we calibrate the beliefs representation of a quantitative DSGE model to match the estimated balance of risks from the econometric model and conduct policy counterfactuals. Section 7 concludes.

## 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. We provide 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. The complete proof of these results are provided in Appendix A.

The set of models to which our methodology can be applied is very broad. We make two additional assumptions. First, agents decisions are affected by their perceptions of the whole distribution of future shocks, which are assumed to be symmetric in the long-run.[2] Second, 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). A key feature of these unimodal distributions is that the mode is not affected by perturbations in the asymmetry of the density. These distributions are flexible and convenient, making them widely used in finance, econometrics, and various other fields such as environmental studies and engineering, for modeling asymmetric risks.

## 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:

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

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 $\epsilon_t^a$. The former includes all the symmetric shocks, e.g., normally distributed, whereas the latter is a possibly *asymmetric* shock with distribution $\mathsf{f}(m_a, \sigma_a, \varrho_a)$, where $m_a$ denotes the mode, $\sigma_a$ the scale, and $\varrho_a$ an asymmetry parameter that measures the degree of skewness.[3] Without loss of generality,

[2] As a consequence of the symmetry of shocks in the long-run, the only relevant steady state around which the model is log-linearized is the deterministic (or non-risky) one.

[3] 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.

the mode $m_a$ can be set to zero, as it can be absorbed into the definition of the steady state.[4]

For $\varrho_a > 0$ ($\varrho_a < 0$), $\mathsf{f}$ 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) $m$ are more likely than those above (below) it.[5] In Appendix A.1, we derive the log-linear system in Equation (1) as a first-order Taylor expansion around the steady state of the full nonlinear model with skewed shocks.

**Time-varying asymmetric risk.** In the class of two-piece distributions, $\varrho$ is proportional to the Pearson's coefficient of skewness. That is, the wedge between the expected value, $\mu$, and the mode $m_a$ (the *mean-mode wedge*, for short), implies the linear mapping $\mu = m_a + \varkappa\sigma_a\varrho_a$, with $\varkappa > 0$. The exact definition of $\varkappa$ depends on the specific functional form of the skewed distribution.[6]

A distinctive feature of our analysis is the assumption that the asymmetry parameter of the distribution $\mathsf{f}$ varies over time. We relax the i.i.d. assumption for $\epsilon_t^a$ such that serially independent errors allow the economy to shift between periods of upside and downside risk, with the severity of each episode changing over time, and we assume that the asymmetry parameter follows a stationary zero-mean process; i.e., $E(\varrho_{a,t}) = 0$. The other parameters, mode and scale are held constant, with the mode set to zero without loss of generality. These restrictions can be relaxed, but are maintained in the interest of analytical tractability.[7]

**Agents beliefs.** Agents are rational and perfectly informed. They know the parameters in Equation (1), the history of the realized shocks, and the shock distributions, including the mode $m$, the scale $\sigma$, the past and current realizations of the asymmetry of $\mathsf{f}$.

In every period $t$, agents receive news about the future evolution of the distribution's asymmetry, $\varrho_{a,t}$, which is the stochastic parameter governing the mean–mode wedge $(\mu - m_a)$. For each

[4]Recall that the distribution of shocks, $\epsilon_t^a$, is expected to become symmetric in the long run. As a result, the long-run mean and mode of $\epsilon_t^a$ coincide.

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

[6]For a Skewed-Normal distribution, $\varkappa = \sqrt{2/\pi}$. For the Skewed-$t$ distribution—used later in the paper—$\varkappa$ depends on the degrees of freedom (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]Introducing stochastic variation of the mode would not alter any results of this Section, provided that these changes do not completely offset the effects of skewness on the expected value of the shocks. In the empirical analysis, we will jointly estimate the evolution of the three parameters of the predictive distribution of inflation.

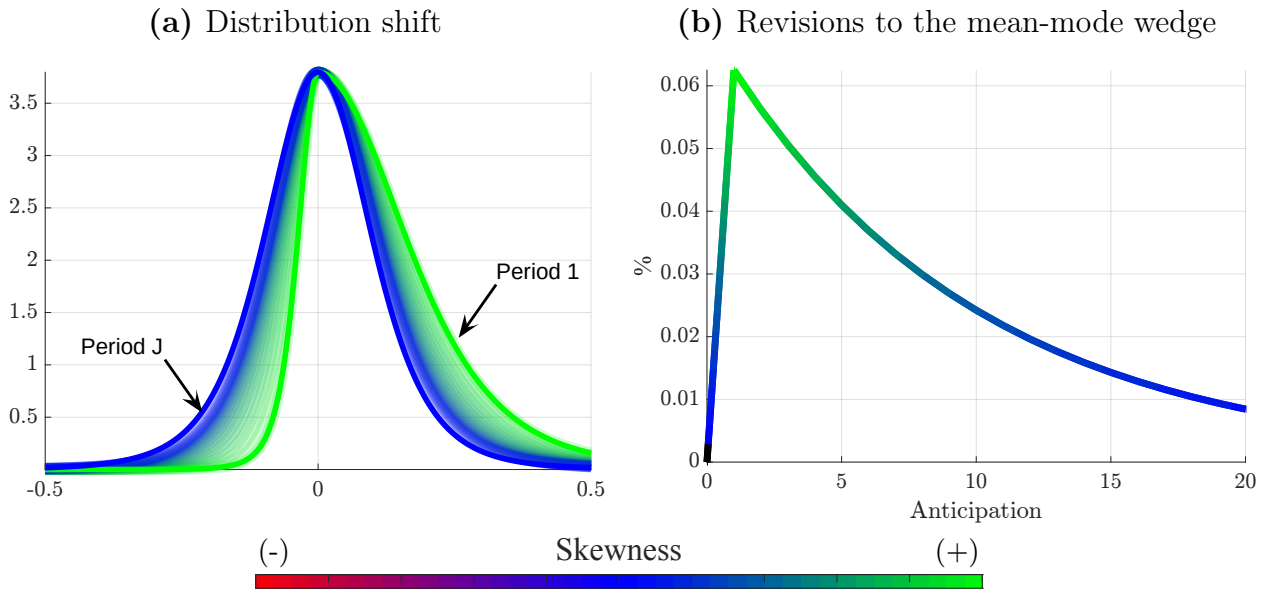



**Figure 1:** Agents' revision and shock distributions

*Note:* The figure illustrates the relation between agents' revision to shock asymmetry and the implied shock distributions. The left panel reports the change in the implied distributions for the asymmetry shocks. Positively (negatively) skewed distributions are colored in green (red); symmetric distributions are reported in blue. The right panel show the implied revisions to agents' beliefs about the asymmetry of macroeconomic shocks.

horizon $j \in \{1, \ldots, J\}$, agents observe the signal $s_t^j = \varrho_{a,t+j} + \eta_t^j$, where $\eta_t^j$ denotes zero-mean i.i.d. noise that is orthogonal to future realizations of the asymmetry parameter, and update their beliefs about future risks.

An example of such revisions is shown in the left panel of Figure 1. In this scenario, agents receive positive news about a temporary increase in the asymmetry of the future shock distribution, leading them to anticipate upside risks up to $J$ periods ahead. The left panel illustrates the sequence of shock distributions from the next period until the point at which risk returns to symmetry. As upside risks become more likely, the distributions become right-skewed (green curves); as the news fades, they gradually revert to symmetric distributions (blue curves). The right panel reports the corresponding shifts in the asymmetry parameter for the $J$ future shock distributions. The increase in upside skewness tilts the balance of risks upward, leading agents to expect positive future shocks, even though the most likely realization remains zero at all horizons.

Rational agents solve the signal extraction problem, updating expectations about the asymmetry parameter $j$ periods ahead. We denote these i.i.d. updates as $\varrho_{a,t}^j \equiv E_t \varrho_{a,t+j} - E_{t-1} \varrho_{a,t+j}$.

Within this setting, agents form expectations by agreeing on a central scenario for the variables of interest, and assessing the balance of risks around such modal forecast. This approach is consistent with how modern central banks generate and communicate predictions, as documented by Aruoba and Drechsel (2026).[8]

**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:*

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

*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 $E_t \epsilon_{t+j}^a$.

While the first two terms on the right-hand side of Equation (2) are standard, the third term 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 $\epsilon_{t+j}^a$, or if the news leads to the belief that shocks will be symmetric ($E_t \varrho_{t+j} = 0, \forall j$), the expected value of future asymmetric shocks is zero. As a consequence, the third term on the right-hand side of Equation (2) drops out, and the solution simplifies to the familiar form: $\mathbf{z}_t = \mathbf{\Theta}_1 \mathbf{z}_{t-1} + \mathbf{\Theta}_0 \left[\boldsymbol{\epsilon}_t^s \ \epsilon_t^a\right]'$.

Proposition 1 highlights that, when agents expect asymmetric shocks over the forecast horizon, the expected outcomes diverge from the modal forecast, $\mathbf{z}_{t+j} = \mathbf{\Theta}_1^j \mathbf{z}_t$, due to the influence of risk perceptions on agents' expectations.

[8] Note that when risks are balanced around the central scenario, modal forecasts coincide with 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 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 though the most likely outcome (the mode) 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).

**Alternative specifications for asymmetric risks.** It is important to highlight that our focus on the family of two-piece distribution is mainly dictated by convenience. Our framework holds more generally for a broader class of distributions that includes multimodal densities with one dominant mode; see Appendix B for a formal proof. These types of densities remain of interest when skewness emerges in the tails of a distribution, rather than around the mode. Within this class, the Pearson's coefficient of skewness, defined as the wedge between the mean and the mode, scaled by the dispersion of the distribution, remains a sufficient statistic to capture the first-order effects of changes in the asymmetry of the distribution.

## 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*. The key advantage of this representation is that it can be solved using off-the-shelf, standardized solution techniques.

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, we define the realizations of 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 $\varphi_t^j$ 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:*

$$\begin{aligned} \varphi_t^j &= E_t \epsilon_{a,t+j} - E_{t-1}\epsilon_{a,t+j} \\ &= \varkappa \sigma_a \left( E_t \varrho_{a,t+j} - E_{t-1}\varrho_{a,t+j} \right) \\ &= \varkappa \sigma_a \varrho_{a,t}^j, \qquad \forall j \in 1, \ldots, J, \end{aligned} \tag{6}$$

*where $\varrho_{a,t}^j$ denotes agents' revision about expected risk asymmetry after observing the signal, $s_t^j$, in*

*the actual model. (ii) The surprise dummy shock must satisfy the following condition*

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

*Proof.* See Appendix A.4. □

Condition $(i)$ ensures that the dummy belief shocks replicate the same impact of anticipated asymmetry on the expectations $E_t\epsilon_{a,t+j}$. Condition $(ii)$ ensures that, in every period, the realization of the shock $\epsilon_{a,t}$ exactly matches that in the actual economy. Together, these two conditions guarantee that the beliefs representation replicates the structure of the actual economy, the realizations of shocks, and agents' expectations about future realizations of asymmetric shocks.

Proposition 2 establishes the key result that the beliefs representation is isomorphic to an economy in which agents receive news regarding the evolving balance of risks. Hence, owing to its tractability, solving the belief representation allows to easily derive the equilibrium laws of motion of the actual economy in Equation (2). It is worth noting that standard solution techniques for linear rational expectations remains effective even when the signal extraction problem may lead to non-Gaussian revisions to expectations of future shocks. These solvers typically do not restrict the stochastic process of the shocks beyond the i.i.d. and mean zero assumptions, which are not restrictive in most cases (see Sims, 2002).

The beliefs representation can be used to study environments where agents receive news about future risks, conveyed through changes in the scale and asymmetry of future 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 the asymmetry depends on the scale of the distribution: a shift in skewness has a larger impact on expectations when uncertainty is high (i.e., when $\sigma_a$ is large). Likewise, when distributions are skewed, changes in uncertainty influence mean outcomes, even when asymmetry is unchanged. The direction and magnitude of these effect depend on the sign of the skew.

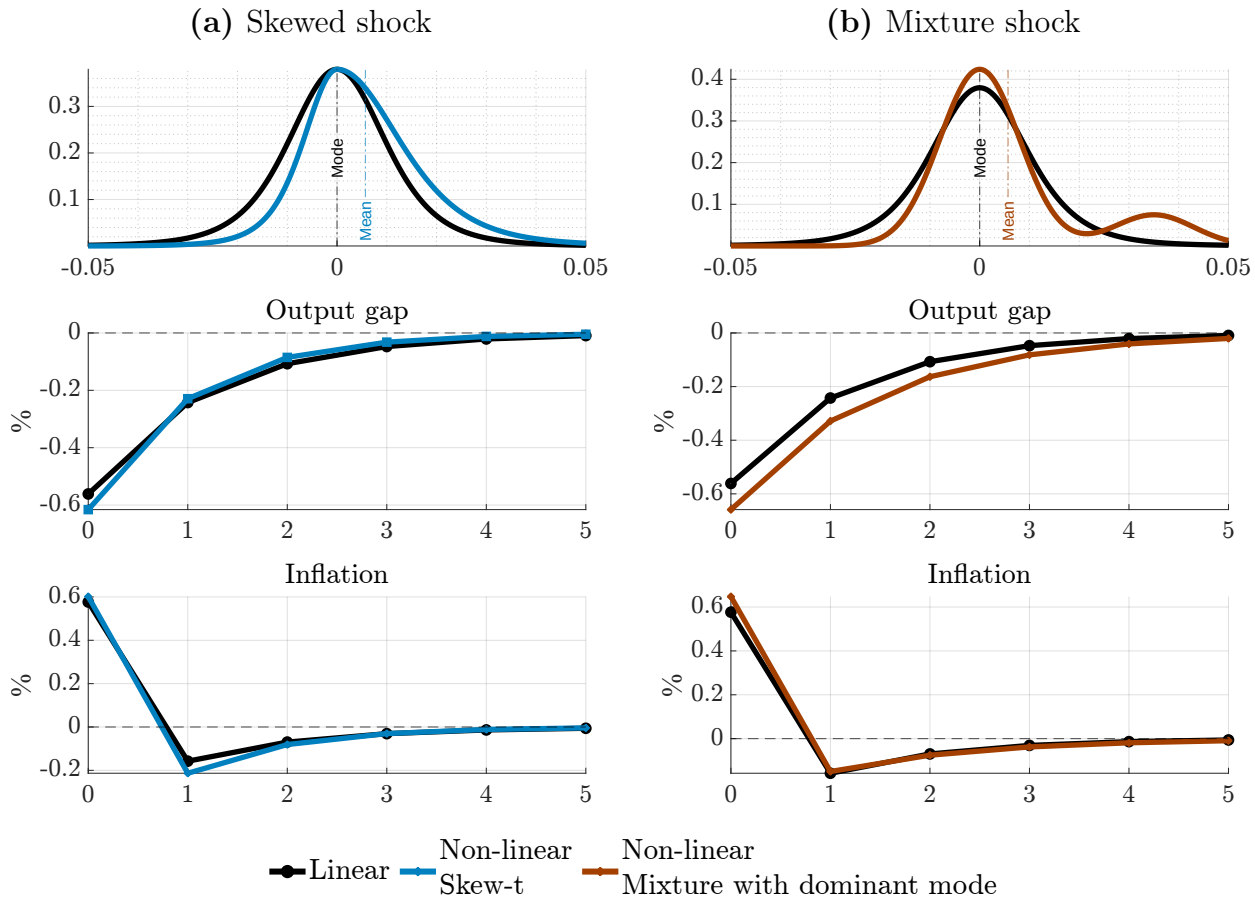



**Figure 2:** Non-linear propagation and the belief representation

*Note:* The panels compare the propagation of a positive anticipation shock about forthcoming upside risks to inflation in a fully non-linear model, solved with global methods, against the belief representation (black). The two columns consider a shocks drawn from a two-piece distribution (left, blue) and from a bimodal density with a dominant mode (right, red). The bottom panels report the response of the output gap and inflation. The shock size is adjusted within the two models to generate effects of the same size.

## 2.3 How good is the belief representation?

The beliefs representation offers a convenient way to characterize the first-order effects of shifts in risk asymmetry. A natural question is how accurate this approximation remains when higher moments of the distribution of future shocks change. To address this issue, we study a small-scale model with nominal frictions in the form of sticky prices, allowing agents to form expectations about future shocks with asymmetric risk. For illustration, we consider a scenario in which agents revise their beliefs about the shock occurring in the next period. An increase in perceived upside risk raises expected future shocks, while the most likely outcome remains zero—that is, skewness changes while the modal outcome is unchanged. Further details on the model, the calibration of

its parameters, and the nonlinear solution method used are provided in Appendix C.[9]

Figure 2 reports the responses of the output gap and inflation computed using a global solution method, against the first-order approximation implied by the beliefs representation of anticipated upside risk (black), constructed as in Proposition 2. We consider asymmetric shocks featuring skewness about the model (blue, column (a)) and skewness in the tails (red, column (b)). In both cases, price-setting firms partially incorporate the expected upside risks into current prices, raising inflation and compressing the output gap. The central bank, following a standard Taylor rule, responds by tightening monetary policy to contain inflationary pressures (more on this in Section 5). Importantly, the nonlinear responses closely overlap with those generated by the beliefs representation.

This finding suggests that nonlinearities in the consumption Euler equation and the New Keynesian Phillips curve are limited, so that the effects of asymmetric risks are largely captured by first-order terms. Consequently, incorporating higher-order terms in standard New Keynesian models appears to add little. Taken together, these results highlight the general validity and versatility of the beliefs representation for studying the macroeconomic implications of time-varying balances of macroeconomic risks.

# 3 Estimating the time-varying balance of inflation risks

The usefulness of the beliefs representation for studying time-varying macroeconomic risks rests on establishing that asymmetries in such risks vary meaningfully over time. The representation itself is agnostic about how the balance of risks is measured: what is required is a time series of revisions in the wedge between modal and expected inflation. We construct this object using a time-series model that estimates real-time predictive densities of inflation over a long postwar sample.[10] The model incorporates time-varying skewness in the inflation distribution and jointly

[9] Implementing this exercise requires specifying a time-varying distribution of future shocks. Allowing skewness to evolve continuously across multiple horizons would greatly expand the state space, rendering a fully nonlinear solution computationally infeasible.

[10] Alternative approaches could be used to construct this object, including Markov-switching models (e.g., Le Bihan et al., 2024), quantile regressions (e.g., López-Salido and Loria, 2024), survey data (e.g., Mouabbi et al., 2025), or option-implied measures (e.g., Hilscher et al., 2026). Measures of skewness can also be constructed from survey-based density forecasts (Coibion and Gorodnichenko, 2026).

estimates the location, scale, and asymmetry of inflation risks. We estimate the model using quarterly data on U.S. core PCE inflation from 1960Q1 to 2024Q4.

## 3.1 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 $\pi_t$ can be characterized by a Skew-t ($Skt$) distribution with time-varying location ($m_t$), scale ($\sigma_t$), and shape ($\varrho_t$) parameters:

$$\pi_t \sim Skt_\nu(m_t, \sigma_t^2, \varrho_t), \tag{8}$$

where $\nu$ denotes the, time invariant, degrees of freedom. The degree of the asymmetry of the distribution 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), $m_t$, can be retrieved as $\sigma_t(1 - \varrho_t)$ and $\sigma_t(1 + \varrho_t)$. Note that this specification allows as special case the symmetric Student-t distribution when $\varrho_t = 0$. 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), we treat the time-varying parameters as unobserved components that can be learned in real-time, allowing each parameter to feature a permanent and transitory component (as in Stock and Watson, 2007). Parameters time variation is accommodated as in Delle Monache et al. (2024);[11] see Appendix D for additional details.

## 3.2 Time-varying skewness of inflation risk

The estimated model provides new insights into the time-varying stochastic properties of the inflation process. While much is already known about the mean and variance of inflation (see, e.g.,

[11] Updates of the time-varying parameters are proportional to the scaled score of the conditional distribution (as in Creal et al., 2013; Harvey, 2013) The scaled score vector is defined as the gradient of the likelihood function with respect to the dynamic parameters, scaled by the inverse of the diagonal of the Information matrix. 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).

Stock and Watson, 2007), our approach sheds light on the dynamics of inflation skewness.[12]

Figure 3a reports the estimated time-varying skewness of inflation. Movements in skewed inflation risk are closely aligned with major shifts in the macroeconomic environment and are highly persistent. Episodes associated with binding supply constraints are typically accompanied by pronounced increases in upside inflation risk. The late 1960s and the post-COVID recovery—both characterized by tight capacity and supply bottlenecks—are marked by sharp increases in positive skewness. A similar rise in upside risk accompanies the oil price shocks of the mid-1970s, consistent with large and asymmetric cost-push disturbances. In contrast, declines in upside risk and increases in downside inflation risk coincide with episodes that constrained price adjustment or reflected excess supply, including the introduction of wage and price controls in the early 1970s and periods of oil market gluts, such as those observed in the early 1980s and in 2015–16.

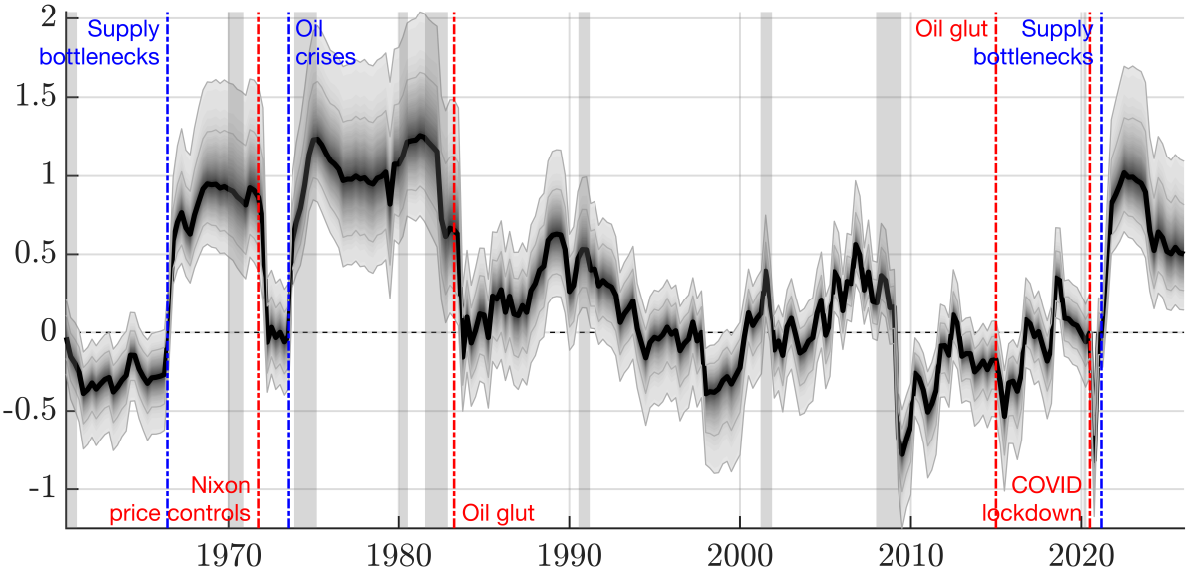
Over the full sample, inflation skewness is moderately negative during much of the early 1960s, turns positive in the late 1960s, and reaches its peak during the Great Inflation of the 1970s. Upside risk declines steadily beginning in the early 1980s and continues to recede through the mid-1990s, after which inflation risk is broadly symmetric. The Global Financial Crisis is associated with a marked deepening of negative skewness, followed by a prolonged period in which downside risk dominates. During the COVID-19 pandemic, skewness becomes sharply negative before reversing rapidly as economic activity recovers. By mid-2021, upside inflation risk reaches levels comparable to those observed in the late 1960s and during the Great Inflation.

These dynamics are also consistent with broader changes in the global economic environment. Our measure of inflation skewness closely comoves with the trade fragmentation index constructed by Fernández-Villaverde et al. (2024). In particular, the decline in skewness from the mid-1980s through the early 2000s coincides with the post-Cold War phase of deepening global integration and the expansion of global trade and production networks, while the sharp increase after 2020 aligns with renewed fragmentation following the pandemic and rising geopolitical tensions. The correlation between inflation skewness and the trade-related component of the fragmentation index is about 0.55.[13] These patterns suggests that changes in the global trade environment may be one

[12]Estimates of the time-varying mean and volatility of U.S. core PCE are broadly consistent with those obtained from the Stock and Watson (2007, 2016) framework (see Appendix D.7 for further analysis).

[13]The correlation is about 0.4 with the common component of the index and 0.5 with the financial subcomponent.

**(a)** Skewness




**(b)** Balance of risks

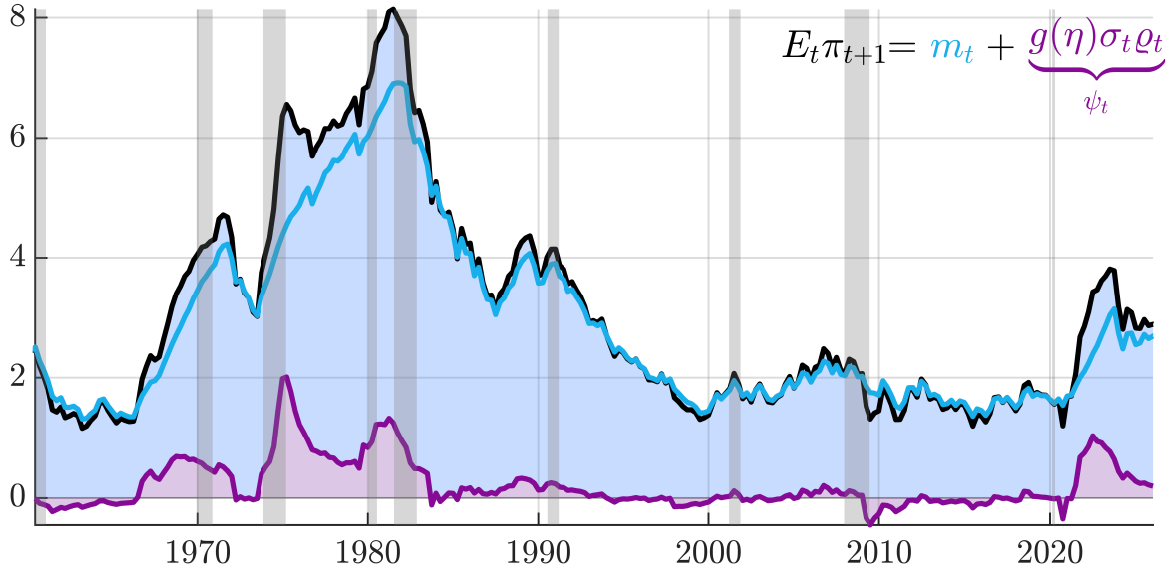



**Figure 3:** Time-varying asymmetric risks of Core PCE inflation

*Note:* The top panel plots the estimated time-varying skewness of U.S. core PCE inflation, along with 68% and 95% credible intervals. Vertical lines indicate key events changing inflation skewness. The bottom panel reports model based inflation expectations. The total expectation (in black) is decomposed into a central scenario ($m_t$, cyan) and the tilt induced by the balance of risks ($\psi_t$, purple). Gray shaded areas indicate NBER recessions.

of the drivers of inflation risk, consistent with the view that asymmetries in inflation outcomes partly reflect disruptions to production and to supply conditions.

### 3.3 Time-varying mean-mode wedge

A defining feature of any skew-distribution, $p(\pi|m_t, \sigma_t, \varrho_t, \nu)$, is the fact that asymmetry directly affects the first moment of the distribution. Specifically, in the case of the Skew-t distribution in Equation (8), one can show that $\forall h > 0$

$$E_t\pi_{t+h} = \int_{\mathbb{R}} \pi \times p(\pi|m_{t+h}, \sigma_{t+h}, \varrho_{t+h}, \nu)d\pi = m_{t+h} + \underbrace{g(\nu)\sigma_{t+h}\varrho_{t+h}}_{\psi_{t+h}} \tag{9}$$

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, $\psi_{t+h}$, that is a function of the asymmetry parameter. That is, asymmetry creates a wedge between the central scenario (i.e. the mode of the distribution, $m_{t+h}$), and the expected value. This wedge has the same sign of the prevalent asymmetry and is quantitatively relevant in more uncertain periods (i.e. when $\sigma_{t+h}$ is large). When risk is skewed to the upside (i.e. $\varrho_{t+h} > 0$), an increase in risk corresponds to rising inflation expectations. Conversely, when risk is negatively skewed (i.e. $\varrho_{t+h} < 0$), an increase in risk results in a decline in expected inflation.

Figure 3b illustrates the decompositions of model based inflation expectations into a central scenario (the most likely expected outcome) and the tilt induced by the balance of risks around it, $\psi$, which is itself a function of the skewness, based on Equation (9). Inflation risk significantly influences inflation expectations, especially in periods of high inflationary pressures and price volatility. The high labor costs and rising interest rates of the 1970s and 1980s are consistent with a prevalence of upside risks to inflation. Similarly, recent supply chain disruptions and geopolitical tensions could be associated with positive balance of risks estimated in the last 3 years of the sample.[14] Negative skewness contributed to a downward bias in expected inflation during the decade leading up to the COVID pandemic. Most of the post-Great Financial Crisis period has been characterized by the recurrence of zero lower bound spells, associated to strong deflationary concerns (see, e.g. Adam and Billi, 2007). However, remarkably low inflation volatility during these years 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

[14] De Polis et al. (2023) relate the dynamics of inflation moments to macroeconomic and financial predictors, suggesting that a mix of fiscal and monetary related factors account for the largest share of predictability.

during this time has significant implications for how monetary policy should address the persistent negative skewness in the post-Great Recession era.

## 3.4 Empirical validation

We now assess whether the estimated balance of inflation risks captures a robust and economically useful feature of the data. The evidence shows that time-varying skewness is statistically significant and reliably detected by the model. It also provides useful real-time information, improves forecast performance, and is present across alternative measures of inflation. Additional details and results are reported in Appendix D.

**Statistical evidence for time-varying skewness.** We provide additional evidence on time variation in the asymmetry of the predictive distribution of core PCE inflation using formal statistical tests. Parametric tests strongly reject the null of symmetry at the 1% level, even after allowing for time varying volatility. We also conduct targeted Monte Carlo exercises showing that our estimation procedure detects skewness only when it is present in the data. Additional experiments indicate that the procedure remains robust in the presence of both fixed and time-varying correlations between location and scale.

**Real-time information content.** Assessing how well our model-based measures of asymmetric risks captures changes in inflation risk is inherently challenging, as we only observe a single realization from a time-varying distribution. While ex-post skewness can be gauged using rolling measures derived from the data, these backward-looking estimates face important limitations. Our estimates consistently lead the rolling measures by about two quarters, providing a timely signal of changes in the balance of risks—a key advantage for monitoring the evolution of inflation risk.

**Out-of-sample forecasting performance.** As a further validation of our estimates of time-varying risk asymmetry, we conduct a real-time, out-of-sample forecasting exercise and show that incorporating time-varying skewness into inflation forecasts significantly improves predictive accuracy relative to standard benchmarks (Stock and Watson, 2007). Specifically, we compare the

**Table 1:** Real-time out-of-sample comparison

| | | | CRPS Decomposition | | | Event Forecasts | | |
|---|---|---|---|---|---|---|---|---|
| | MSFE | CRPS | Right | Left | Center | $\pi_{t+h} < 1.5$ | $1.5 \leq \pi_{t+h} \geq 2.5$ | $\pi_{t+h} > 2.5$ |
| h = 1 | **0.969** (0.011) | 0.995 (0.333) | 0.998 (0.424) | 0.992 (0.186) | 0.995 (0.364) | **0.956** (0.001) | **0.966** (0.004) | **0.967** (0.001) |
| h = 4 | **0.925** (0.000) | **0.958** (0.001) | **0.979** (0.052) | **0.940** (0.000) | **0.954** (0.000) | 0.981 (0.107) | 0.981 (0.115) | **0.987** (0.064) |
| h = 8 | **0.884** (0.000) | **0.927** (0.000) | **0.939** (0.000) | **0.921** (0.000) | **0.920** (0.000) | **0.975** (0.074) | **0.970** (0.014) | 1.007 (0.702) |

*Note:* The table report the relative performance of Stock and Watson (2007) UCSV model against our Skew-$t$ model. Results are reported in ratios, with our model being at the numerator; values smaller than 1 imply superior predictive accuracy of the Skew-$t$ model. The out-of-sample period runs from 2000Q1 to 2024Q4, and observations are collected from real-time vintages. Values in **bold** are significant at the 10% level.

models in terms of point, density, and event forecasts, evaluated using the mean squared forecast error (MSFE), the Continuously Ranked Probability Score (CRPS) of Gneiting and Ranjan (2011), and the Brier score, respectively. The comparison strongly favors our model across all metrics and forecast horizons, as reported in Table 1; p-values for the Diebold and Mariano (1995) test are shown in parentheses. Expressed as the ratio of the score achieved by our model to that of the UCSV benchmark, the results indicate significant gains in point forecasts ranging from 3% at short horizons to 10% at medium horizons. We also document improvements of up to 7% in CRPS scores and gains of 2 to 4% in event prediction.

The UCSV model lacks a mechanism for capturing skewness and overlooks the presence of fat tails in the data. Replicating the exercise using a specification that excludes asymmetry confirms the importance of incorporating skewness to improve model fit and inflation forecasting accuracy. Lastly, we find our real-time forecasts perform on par with the SPF. 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, anticipating the inflation spike of the post-COVID period.

**Robustness across inflation measures.** We extend 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 price index used. Periods of upside or downside skewed risk tend to align across these indicators. Notably, core PCE inflation appears to exhibit a more attenuated balance of risk compared to other inflation measures.

# 4 The macroeconomic effect of changes in the balance of risks

The previous section established that the balance of risks to U.S. inflation varies substantially over time and that tracking its evolution improves inflation forecasts. In this section, we evaluate to what extent a change in the balance of inflation risks has broader macroeconomic effects.

## 4.1 Empirical evidence

To assess the macroeconomic effects of a change in the balance of risks to U.S. inflation, we use a local projection framework (Jordà, 2005; Jordà and Taylor, 2025). Specifically, we estimate the responses of real GDP, the PCE headline index, and the federal funds rate to an increase in the asymmetry of the inflation process:

$$y_{t+h} - y_{t-1} = \alpha_h + \beta_h \Delta \varrho_t + \mathbf{X}_t \boldsymbol{\gamma}_h + \nu_{t+h}, \tag{10}$$

for $h = 0, ..., H$. The control vector includes a dummy capturing the effects of COVID, and a linear trend, a selection of contemporaneous identified shocks, and multiple lags of $y_t$ and $\varrho_t$ (as recommended by Montiel Olea and Plagborg-Møller, 2021).[15] The shocks include the oil supply news of Känzig (2021), the monetary policy shock of Jarociński and Karadi (2020), the tax shocks of Romer and Romer (2010), and the first difference of the TFP-adjusted series of Fernald (2012). Controlling for these shocks ensures that this specification identifies the dynamic effects of an exogenous upward shift in inflation risk asymmetry, referred to as inflation skewness shock.

The impulse response coefficients, $\{\beta_h\}_{h=0}^{H}$, measure the cumulative response of $y_t$ to a positive

[15] We include 8 lags, but results are robust to alternative specifications, including additional lags, excluding the linear trend, and adding further controls. Introducing the identified shocks contemporaneously effectively orthogonalizes the inflation skewness shock with respect to monetary policy, fiscal policy and oil supply shocks. Details on the data are reported in Appendix F.

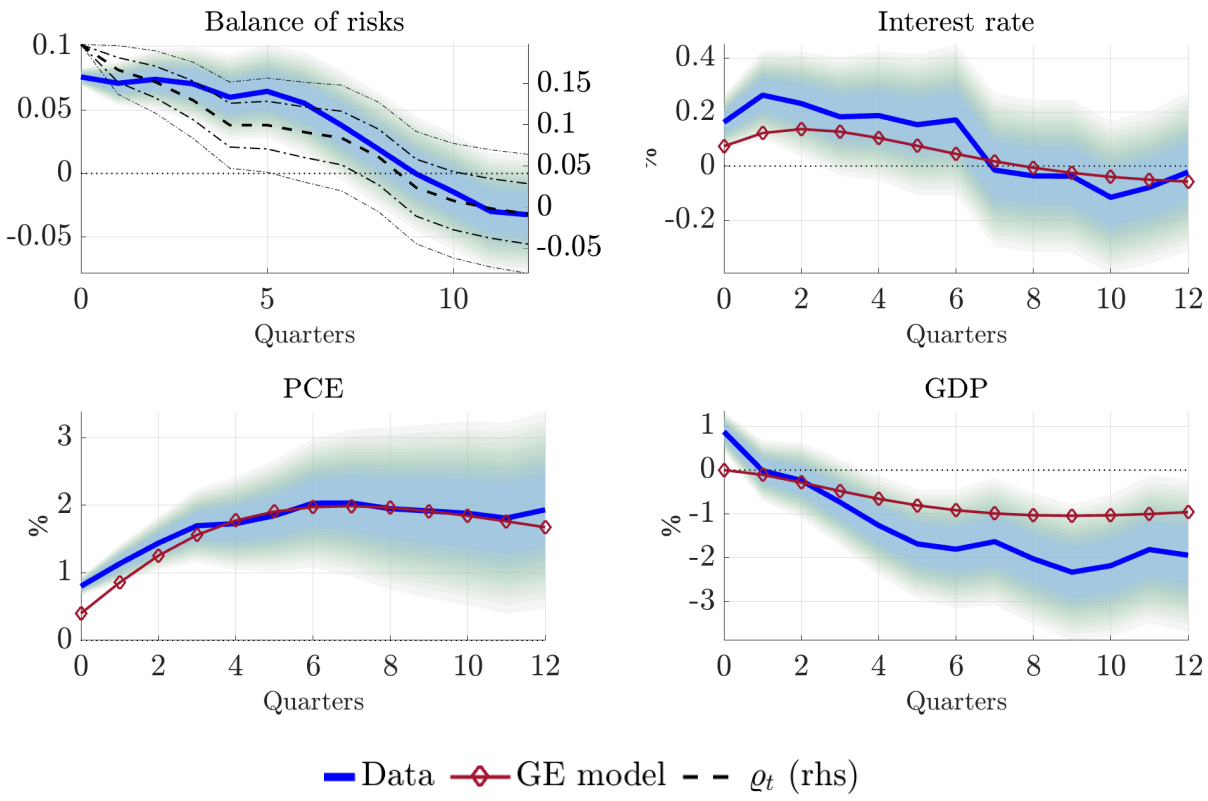



**Figure 4:** Impulse responses to an increase in upside inflation risks

*Note:* The panels report the responses of the balance of risks, the interest rate, headline and core PCE inflation, real GDP, and hours worked to a positive shock to upside inflation risks. Blue lines denote the responses estimated via local projections, while dark red lines correspond to the beliefs representation of the model estimated by Smets and Wouters (2007). The top-left panel also reports the response of the asymmetry parameter $\varrho_t$ (black line, left axis). Shaded areas and dot-dashed lines indicate 95% confidence intervals.

inflation skewness shock. In Figure 4, these responses are shown in blue, with bands denoting 95% confidence intervals. We normalize the IRFs so that, on impact, the shock generates a change in the balance of risks equal in magnitude to that observed at the onset of the Russian invasion of Ukraine in 2022Q1 (see Figure 3). This amounts to roughly 1.2 standard deviations increase in $\varrho_t$.

The upper-left panel reports the response of the balance of risks to an inflation skewness shock, together with the response of the asymmetry parameter, shown in black. The balance of risks rises on impact and then declines only gradually, with a half-life of about one year. A positive shift in the balance of risks is inflationary, with prices rising already within the same quarter. Inflation remains elevated in the near term and then turns negative, so that the cumulative effect on the price level is only partly sustained. GDP decline with a short delay and remain below baseline for roughly six to eight quarters. Inflationary pressures and weakening real activity warrant policy

rate to modestly increases on impact, before returning toward baseline levels within the first year.[16]

**Alternative identification schemes.** A potential concern is that residual revisions in inflation skewness may partly reflect omitted macroeconomic disturbances. To address this concern, we re-estimate the local projections using an IV design that instruments changes in skewness with the sign of large shifts in inflation asymmetry, while retaining the same contemporaneous controls for oil, monetary, fiscal, and productivity shocks as in the baseline specification.[17] Identification therefore comes from the direction of skewness changes around a handful of large events, rather than from all quarter-to-quarter movements in the estimated balance of risks. The IV responses, reported in Appendix F.1, remain close to the baseline estimates. As additional robustness checks, we perform two alternative exercises. First, we control for global fragmentation using the measure of Fernández-Villaverde et al. (2024), thereby isolating variation in inflation skewness unrelated to changes in the degree of global economic integration. Second, we add the estimated location and scale of inflation to the control set, so that identification comes from movements in asymmetry that are orthogonal to shifts in the central forecast and dispersion. Across all specifications, the results are consistent with those from the baseline and indicate that upward revisions in inflation risk behave like adverse supply shocks: they raise prices, depress aggregate demand, and lead to a cautious monetary policy response.

## 4.2 Skewness shock in general equilibrium

We now introduce an inflation skewness shock within an empirical DSGE model. Specifically, we augment the structural model estikmated by Smets and Wouters (2007) to allow for time variation in inflation risk arising from changes in expectations about future cost-push shocks.[18] These structural shocks are assumed to follow a Skew-$t$ distribution with time-varying moments,

[16] In Appendix F.1, we show that the response of headline inflation closely mirrors that of core inflation, and that the impulse responses of total hours worked are similarly aligned with those of GDP.

[17] Large shifts are selected mechanically as quarters in which the absolute change in inflation asymmetry exceeds a pre-specified threshold. We report results using thresholds of 1.7 and 2.5 standard deviations. The instrument is given by the sign of the change in asymmetry in those event quarters and is zero otherwise.

[18] While the framework can be generalized to accommodate alternative identification schemes by introducing additional dummy shocks that contribute to revisions in risk, we focus on a single mapping from inflation risk to cost-push shocks.

and agents learn about the evolution of the asymmetry parameter $\varrho_t$.

Based on these working assumptions, the beliefs representation of this popular structural model is obtained by augmenting the price Phillips curve with dummy surprise and news cost-push shocks. These shocks capture revisions in expectations associated with changes in inflation risk perceptions and are subject to the restrictions imposed by the representation theorem (Proposition 2):

$$\hat{\pi}_t = \pi_1 \hat{\pi}_{t-1} + \pi_2 E_t \hat{\pi}_{t+1} - \pi_3 \mu_t^p + \varepsilon_t^p + \sum_{j=0}^{J} \varphi_{t-j}^j, \tag{11}$$

where $\mu_t^p$ denotes firms' markups, $\varepsilon_t^p$ is the realized cost-push shock, and $\varphi_{t-j}^j$ denotes a dummy cost-push shock revealed at time $t-j$ and expected to materialize $j$ periods ahead.

The dummy shocks tilt the expected path of future cost-push shocks to capture the effect of changing skewness on expectations. Following the results in Section 2.2, we calibrate the shocks to match the revision in the balance of risks shown in the top-left panel of Figure 4. Parameters are set in line with the estimates of Smets and Wouters (2007) for U.S. data over 1966–2004.[19]

The resulting responses of prices, GDP, hours worked, and the interest rate are shown in Figure 4 as dark red lines. A positive skewness shock to future cost-push disturbances raises agents' inflation expectations. Forward-looking price setters respond by increasing prices, generating an immediate rise in inflation. The central bank reacts to the deterioration in inflation risks by tightening monetary policy over the first year, although the response remains quantitatively modest, in line with the empirical evidence. Higher prices and tighter policy weigh on real activity, producing a sustained decline in GDP and hours worked. Inflation rises sharply on impact and subsequently declines, as a result, the price level initially rises and then partially retraces over time.

This exercise shows that shifts in the balance of risks can generate sizable and persistent macroeconomic effects. It is worth highlighting that, despite the rich endogenous feedback of the quantitative general equilibrium model and the fact that the model parameters are chosen to match unconditional moments of the data, we document a striking correspondence between the simulated

[19] Under the original calibration, the Phillips curve is relatively flat and persistent, with a slope of about 0.03 over 1966–2004. More recent evidence points to a significantly steeper slope, especially during the latest inflation surge (see, e.g., Barnichon and Mesters, 2020). We therefore calibrate the relevant parameters to generate a steeper and less persistent Phillips curve, with slope 0.25. These are the only parameters that differ from the original modal estimates in Smets and Wouters (2007). Parameter values are reported in Appendix F.

conditional responses to an inflation skewness shock and the empirical estimates.

Taken together, these results suggest that a medium-scale New Keynesian model can successfully characterize the key transmission channels through which perceived shifts in the balance of future inflation risks affect the macroeconomy. This, in turn, makes the model a useful laboratory for conducting counterfactual exercises and evaluating alternative policy prescriptions explicitly designed to offset the effects of changes in the balance of inflation risks.

# 5 Optimal monetary policy with asymmetric risks

If changes in the balance of inflation risks have far-reaching macroeconomic effects, how should the central bank respond? In this section, we examine the implications of time-varying asymmetric macroeconomic risk for optimal monetary policy within the textbook New Keynesian model with sticky prices. We first review 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 E.

## 5.1 The case of symmetric risks

Let us assume that the central bank can fully commit, with credibility, to a policy plan by selecting a state-contingent sequence of inflation deviations from its target and output gaps, $\{\hat{\pi}_t, \hat{x}_t\}_{t=0}^{\infty}$, to minimize the loss function $-\frac{1}{2}E_0 \sum_{t=0}^{\infty} \beta^t \left(\hat{\pi}_t^2 + \alpha_x \hat{x}_t^2\right)$, subject to the sequence of constraints given by the Phillips curve, $\hat{\pi}_t = \beta E_t \pi_{t+1} + \kappa \hat{x}_t + u_t$. We further assume that the cost-push shock follows $u_t \sim \mathcal{N}(0, \sigma_u)$, implying symmetric risks.[20]

Under the optimal policy, the central bank sets the output gap proportional to the deviations of the price level from its implicit target: $\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t$, where $\bar{p}_t = p_t - p_{-1}$ denotes the cumulative inflation rate from the period preceding the implementation of optimal plan. This condition can

[20] All variables are expressed in log-deviations from their steady-state value. The objective function is the quadratic approximation of the household's utility function, with weight on the output gap $\alpha_x = \kappa/\varepsilon$, where $\varepsilon$ denotes the elasticity of substitution and $\kappa$ is the slope of the Phillips curve.

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 optimal policy, the price level evolves according to

$$\bar{p}_t = \eta \bar{p}_{t-1} + \lambda u_t, \tag{12}$$

and the corresponding optimal monetary policy rule is given by

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t, \tag{13}$$

where $\sigma$ is the inverse of the household's intertemporal elasticity of substitution, and $\eta$ and $\lambda$ are functions of the structural parameters of the model.

## 5.2 The case of asymmetric risks

Let us now consider the case where the stochastic process driving the cost-push shock is no longer symmetric. We assume shock are independent over time conditional on the realization of the asymmetry parameter $\tilde{u}_t|\varrho_{u,t} \sim \mathsf{f}(0, \sigma_u, \varrho_{u,t})$, where $\mathsf{f}(0, \sigma_u, \varrho_{u,t})$ denotes a general two-piece distribution centered at zero. The asymmetry of the distribution evolves over time, as agents update their views about the balance of risks, driven by $\varrho_{u,t} \neq 0$. As a result, the expected value of future shocks reflects changes in perceived asymmetry. While the mode—the most likely outcome—remains at zero, the mean shifts with asymmetry: $E_t\tilde{u}_{t+j} = \varkappa\sigma_u E_t\varrho_{u,t+j}$.

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

As discussed in Section 2.2, 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. Specifically, we introduce a tilt in agents' beliefs, $E_t(\widetilde{u}_{t+1}) = \varphi_{t,t+1}$, with beliefs evolving 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.

The optimal policy retains the form: $\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t$ but both the price level and the output gap are now influenced by shifts in risk. Specifically, their equilibrium dynamics under this policy can be characterized analytically as follows:

$$\bar{p}_t = \eta\bar{p}_{t-1} + \lambda u_t + \zeta\varphi^1_t. \tag{14}$$

Macroeconomic risks create a wedge between agents' expectations and the modal scenario—that is, the expectations they would hold under symmetric risk. As a result, positive (negative) shifts in the balance of risk lead to upward (downward) movements in prices, as firms incorporate upside (downside) risks to future costs into current pricing decisions (since $\varphi^1_t \propto \varrho^1_{u,t}$).[21]

The optimal policy requires the central bank to respond to the effects of the changing balance of macroeconomic risks on agents' expectations:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\varphi^1_t. \tag{15}$$

The first term represents the optimal rule under symmetry, while the second captures the policy rate adjustment required to offset the effects of the balance of risks on agents' expectations.

[21] The output gap moves in the opposite direction, reflecting the trade-off induced by cost-push shocks: $\hat{x}_t = \eta\hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\left[\lambda u_t + \zeta\varphi^1_t\right]$.

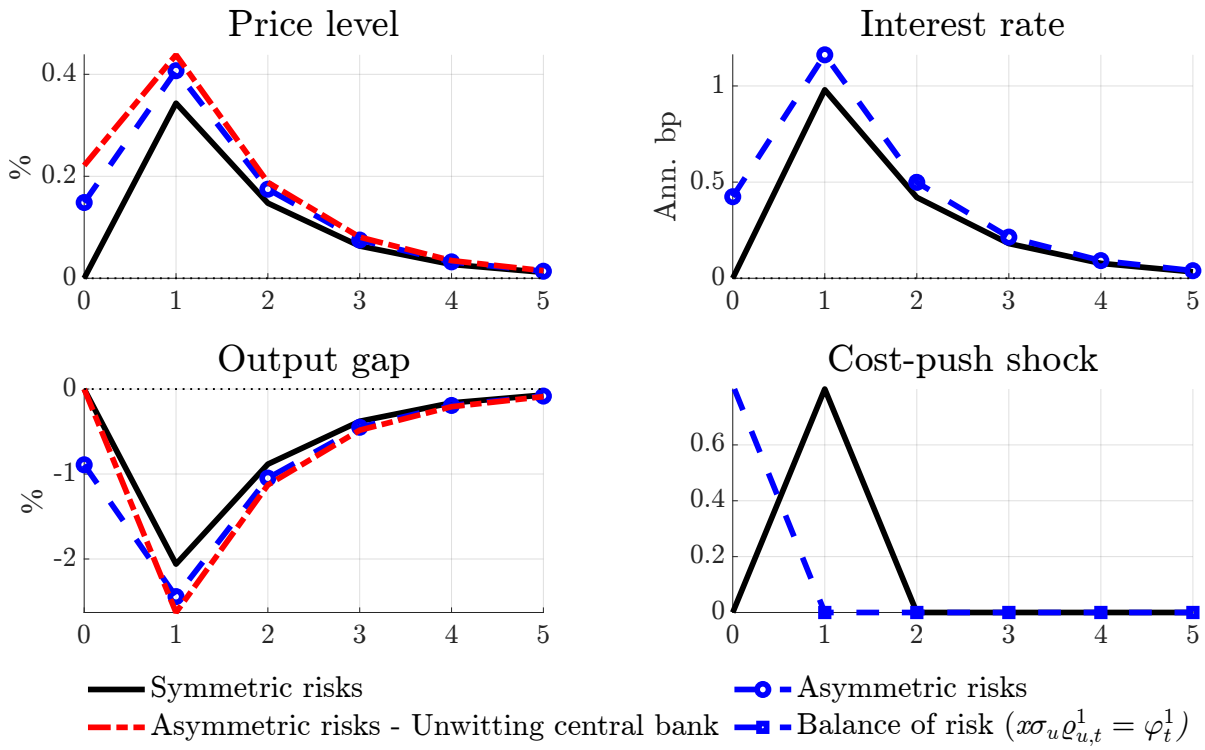



**Figure 5:** Optimal monetary policy under symmetric and asymmetric macroeconomic risks.

*Note:* The panels report the macroeconomic effects of a cost-push shocks drawn from a symmetric distribution (black) and from an asymmetric distribution (blue, dots) under optimal monetary policy. The red line denotes a counterfactual case where risks are asymmetric but the central bank does not take that into account and targets an unattainable price level $\bar{p}^s_t$. The lower right panel shows the cost-push shock, $u_t$ (black), which hits the economy in period 1, as well as the anticipated skewness $\varphi^1_t$, in blue (squares).

## 5.3 A illustrative example

We now illustrate the role of optimal monetary policy in the presence of asymmetric risks.[22] The response of the economy is shown in Figure 5. When the central bank is aware of a positive balance of risks (blue dashed line with circle markers), the distribution of markup shocks exhibits positive skewness. For comparison, we report the response to symmetric risks as black solid lines, where we set skewness to zero (i.e. $\phi^1_t = 0$ in every period $t$).

The economy is initially at its steady state. In period $t = 0$, agents learn that the distribution of cost-push shocks will be positively skewed in period $t = 1$. This upside risk is captured by a dummy anticipated shock $\varphi^1_0 = 0.8$. In period $t = 1$, a positive cost-push shock materializes, with $\varepsilon^u_1 = 0.8$, while all other periods feature symmetric shocks with zero mean. The bottom-left panel of Figure 5 reports the path of the cost-push shock and the evolution of agents' expectations,

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

represented by $\varphi_t^1$.

To ease exposition and build intuition, we consider an intermediate case in which risks are asymmetric but the central bank does not take this into account (an "unwitting" central bank), while the private sector does. This case, reported in red dashed–dotted lines, is defined and analyzed in Appendix E.

The effects of the optimal monetary policy response to asymmetric risks on macroeconomic variables are captured by the difference between the blue and the red lines in the left panels of Figure 5. In period 0, the unwitting central bank chooses the same output gap as in the symmetric case, aiming to achieve the price level that would prevail under symmetry. However, as agents' expectations are distorted by the positive skewness in the distribution of next period cost-push shocks, the targeted price level becomes unattainable. In turn, the equilibrium price level rises due to the upward shift of the New Keynesian Phillips curve, reflecting higher expected inflation in period 1 induced by skewed risks.

When the central bank accounts for expected upside risks to inflation, the optimal policy consists of leaning against the balance of risks, by preemptively contracting the output gap $\hat{x}_t$. Relative to a policy that ignores asymmetries, the price level increases less, reflecting the contractionary response to the shift in upside risks to inflation.

In period 1, a positive cost-push shock materializes in all scenarios, and, from this point on, agents learn that shocks are drawn from symmetric distributions. Persistent deviations between the blue and the black lines reflect path dependence induced by optimal commitment.

The upper-right panel shows the implementation of the optimal policy strategy: the central bank raises the interest rate to counteract upside risks in the inflation outlook by cooling down the economy.

In summary, when cost-push risks are asymmetric, the optimal policy responds by leaning against the direction of the inflation skewness. Positive skewness in inflation risks calls for a tighter policy stance that induces an output contraction to combat inflationary pressures. Conversely, downside inflation risks require the central bank to stimulate the economy beyond its natural level in order to generate inflationary pressures. More generally, symmetric risks are associated with

higher welfare, underscoring the challenges that asymmetric risks pose for monetary policy.

# 6 Risk-Adjusted Inflation Targeting (RAIT)

We now introduce a central bank communication strategy that anchors long-run inflation expectations by offsetting the effects of asymmetric risks: the Risk-Adjusted Inflation Targeting (RAIT). This strategy tilts monetary policy communication to lean against shifts in the balance of risks. Specifically, the central bank adjusts its projected policy rate path in response to revisions to inflation risks, signaling higher rates for upward pressure and lower rates for downward pressure. This forward guidance strategy reflects the key insight of our theoretical model in Section 5: optimal monetary policy should lean against the balance of inflation risks.

As in the previous section, we start from the empirical DSGE model of Smets and Wouters (2007) augmented with asymmetric shocks. To implement the RAIT, we add a time-varying inflation target to the monetary policy rule:

$$\hat{r}_t = \rho \hat{r}_{t-1} + (1-\rho)\left[ r_x \hat{x}_t + r_{\Delta x} \Delta \hat{x}_t + r_\pi \left( \hat{\pi}_t - \underbrace{\sum_{j=1}^{J} \hat{\pi}^{\star}_{t|t-j}}_{\hat{\pi}_t^{\text{RAIT}}} \right) \right] + \varepsilon_t^r, \tag{16}$$

where $\hat{r}_t$ and $\hat{x}_t$ are log deviations from steady state of the nominal interest rate and the output gap, respectively. The coefficients are specified as in Smets and Wouters (2007). The long-run inflation target is set at 2% per year.

The implementation of RAIT operates through forward guidance–that is, communication about the expected path of policy rates. To see this, the guidance shocks $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=1}^{J}$, rescaled by $-(1-\rho)r_\pi$, can be interpreted as stochastic shifts in the intercept of the policy rule, thereby influencing expectations about the future path of interest rates. Under RAIT, these shifts reflect the central bank's assessment of the balance of inflation risks and represent the adjustment in the reaction function required to offset risk asymmetries. Specifically, the term $\hat{\pi}^{\star}_{t|t-j}$ captures past forward-guidance communications, which are aggregated into $\hat{\pi}_t^{\text{RAIT}}$. A negative (positive) value of $\hat{\pi}_t^{\text{RAIT}}$

induces the central bank to tighten (loosen) the expected future path of policy rates.[23]

Under RAIT, the central bank adjusts its forward guidance in response to revisions in the balance of inflation risks. Formally, the central bank signals its intended adjustment to the future path of the policy rate by setting a sequence $\{\pi^{\star}_{t+j|t}\}_{j=1}^{J}$ such that the following condition holds:

$$E\left[\pi_{t+h} \mid \{\varphi_t^j\}_{j=1}^{J}\right] + E\left[\pi_{t+h} \mid \{\pi^{\star}_{t+j|t}\}_{j=1}^{J}\right] = 0, \tag{17}$$

for $h = 1, ..., J$, where $E\left[\pi_{t+h} \mid \{\mathtt{x}_{\mathtt{t}}^{\mathtt{j}}\}_{\mathtt{j}=1}^{\mathtt{J}}\right]$ denotes the response of expected inflation $h$ periods ahead to a revision in $\mathtt{x}$, $j$ periods into the future. Hence, with the RAIT-based guidance, average inflation converges to the level of inflation that would have prevailed in the absence of asymmetric risk over the forecast horizon: the effects of skewness on inflation expectations are fully neutralized by the strategy.

**Forward guidance puzzle?** The RAIT can be thought of as being implemented through forward guidance shocks—that is, anticipated monetary policy shocks, $\hat{\pi}^{\star}_{t|t-j}$. While the literature has identified the forward guidance puzzle—namely, the implausibly strong effects of forward guidance in standard models—this issue does not arise here. As shown by Maliar and Taylor (2019); Del Negro, Giannoni, and Patterson (2023); Bianchi, Melosi, and Nicolò (2024), the puzzle typically emerges when forward guidance shocks are used to match a pre-set interest rate path, thereby shutting down the endogenous policy response. In contrast, the RAIT relies on forward guidance shocks to restore inflation expectations—distorted by asymmetric risks—back to target, while preserving the endogenous feedback of monetary policy and thereby avoiding the puzzle.

## 6.1 Policy Counterfactual

We now present policy counterfactuals to evaluate the RAIT during the post-pandemic period. We use the Smets and Wouters (2007) model augmented with asymmetric cost-push shocks to

[23] This policy can equivalently be interpreted in the model as a time-varying inflation target that responds to shifts in the balance of inflation risks—hence the name RAIT. For instance, when the balance of risks signals upward (downward) inflation pressure, the central bank temporarily aims to undershoot (overshoot) its target as $\hat{\pi}^{\text{RAIT}}_{t+h}$ becomes negative (positive).

assess how the RAIT would have changed the post-pandemic macroeconomic dynamics in the U.S.; parameters are set as in Section 4.

The counterfactual analysis requires as input a revision in the balance of future risks. These revisions can be based on judgment–e.g. based—reflecting the policymaker's subjective view–or derived from model-based estimates. We adopt the latter approach and calibrate the dummy anticipated cost-push shocks to match the effects of revisions in the balance of inflation risks on inflation expectations, $\psi_{t+h|t} - \psi_{t+h|t-1}$, where $\psi_{t+h|t} = E_t\pi_{t+h} - m_{t+h|t}$, estimated in real time using the Skew-$t$ model of Section 3. This requires solving a system of $J+1$ linear equations in $J+1$ unknowns (i.e., the dummy surprise and anticipated cost-push shocks)—a computationally inexpensive task. The details of how to set up the system of equations are provided in Appendix F. We only consider transitory variations in the balance of risks, effectively removing the permanent components, to only account for deviations of risks from the steady state.

### 6.1.1 The RAIT in the post-pandemic inflation

The model explains U.S. data using all shocks, including the dummy shocks calibrated to match the first-order effects of real-time changes in the skewness of cost-push shocks, as estimated by the Skew-$t$ model. To quantify the effects of having followed the RAIT in the post-pandemic period, one simply needs to simulate the Smets and Wouters (2007) model using only the forward guidance shocks based on real-time revisions in the balance of inflation risks. This procedure is valid due to the linearity of the model and the orthogonality of the shocks involved.

Figure 6 shows the RAIT-implied federal funds rate, core PCE inflation, and hours worked, alongside actual data. Panel (a) displays the estimated revisions to inflation risk, $\{\psi_{t+h|t} - \psi_{t+h|t-1}\}_{h=1}^{20}$. These estimates indicate that signs of upside inflation risk began emerging since early 2021, with sustained revisions through the end of 2023. Forward guidance under the RAIT dynamics is designed to counteract these revisions.

Following a sharp increase in inflation risks, the RAIT calls for a rapid tightening as early as 2021. Compared to the observed policy path, the interest rate implied by the RAIT is consistently higher, reflecting the model's recommendation to adopt a tighter stance earlier—whereas the Fed's

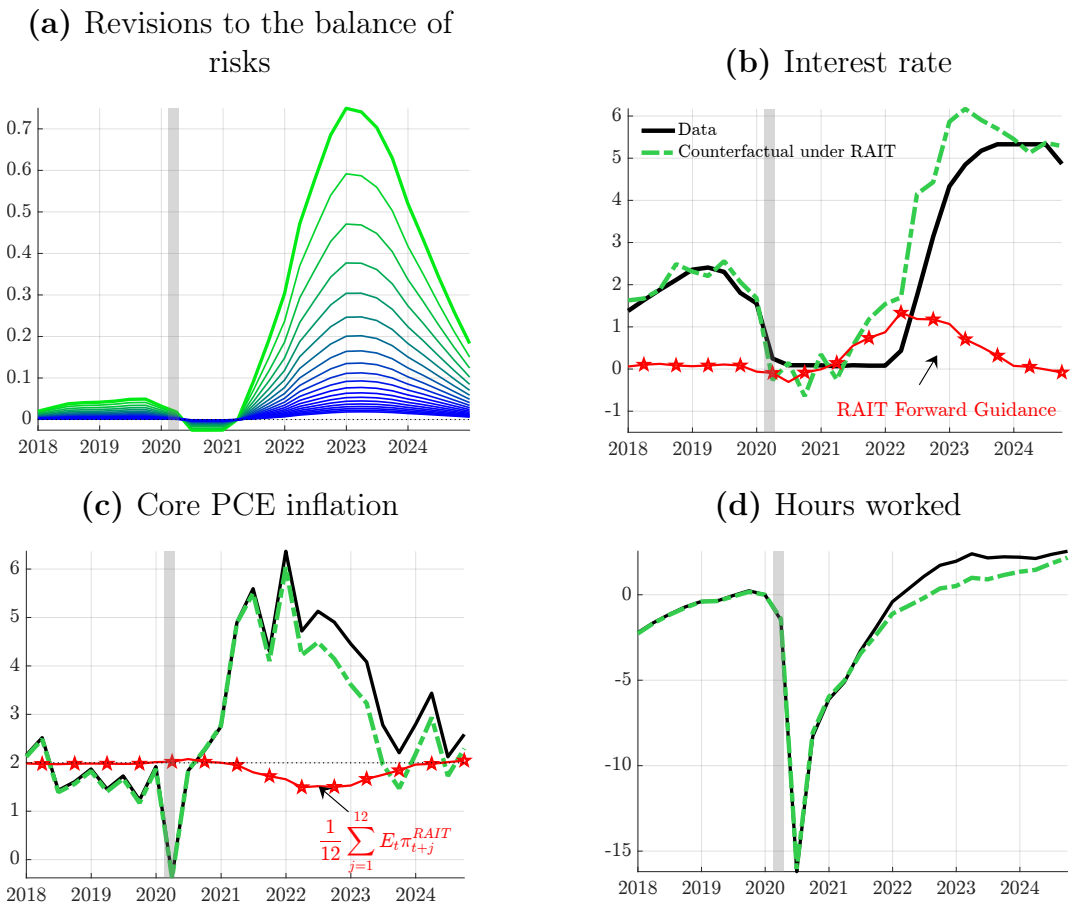



**Figure 6:** Macro dynamics under the RAIT

*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. Panel (b) also includes the forward guidance component of future monetary policy under RAIT, calculated as $-(1-\rho)r_\pi \frac{1}{12}\sum_{j=1}^{12}\hat{\pi}^{\star}_{t+j|t}$. Panel (c) also displays the expected average RAIT target over the next three years, $\frac{1}{12}\sum_{j=1}^{12} E_t\pi_{t+j}^{RAIT}$. Gray shaded areas represent NBER recessions.

actual hiking cycle began roughly three quarters later. Similarly, as the positive asymmetry of inflation began to decline in 2023, the RAIT would have favored an earlier easing.

Remarkably, at the peak of the response, the RAIT prescribes a policy rate broadly in line with the rate eventually set by the FOMC. Thus, the main difference lies in timing: the RAIT would have initiated tightening sooner and reached its peak two quarters earlier. Under the Risk-Adjusted Inflation Targeting (RAIT) framework, the counterfactual monetary policy response can be understood as unfolding in three distinct phases.

In the first phase (2021Q2 to 2022Q3), the central bank tightens monetary policy earlier than observed in the data, raising the federal funds rate sooner and more aggressively. This initial

adjustment is accompanied by strong forward guidance signaling a sustained tightening cycle, as represented by the red line in the top-right panel in Figure 6. In particular, the central bank communicates the likelihood of sustained rate increases over the following three years. Such an early, decisive response attempts to counteract the rising positive skewness in inflation risks, and prevent inflation expectations from drifting upward.

In the second phase (2022Q4 to 2023Q4), as inflationary pressures continues to materialize and the balance of risks remains tilted to the upside, the central bank keeps hiking rates. However, communication regarding the future path of interest rates becomes more measured. While policy still tightens rapidly, forward guidance begins to signal the end of the steeper hiking cycle. A moderation in the communication allows the central bank to maintain a restrictive stance while gradually preparing markets for the subsequent policy normalization, once inflation risks are tamed. This strategy is consistent with our real-time estimates, indicating that inflation skewness surged at short horizons but remained relatively anchored at longer horizons.

In the third phase (2024Q1 to 2024Q4), when inflation risks normalizes and short-term upside pressures recede, the central bank begints her transition toward full policy normalization. In this phase, the policy stance is eased, and a neutral forward guidance is aimed at lowering the expected path of future federal funds rates.

# 7 Conclusions

This paper shows that shifts in the balance of inflation risks are empirically important, macroeconomically consequential, and relevant for the conduct of monetary policy. Using U.S. core PCE inflation data, we document substantial time variation in inflation skewness and show that increases in upside inflation risk generate a stagflationary pattern: inflation rises, aggregate demand weakens, and monetary policy responds only modestly.

To study these effects, we develop a tractable framework in which changes in the asymmetry of future shocks can be represented in a structural model, to a first-order approximation, as belief shocks that shift expectations without necessarily altering realized fundamentals. Embedding this framework in a New Keynesian model, we show that inflation risks affect the optimal conduct of

monetary policy even when the central bank's objective is symmetric. In particular, policy should lean against the direction of inflation risk to preserve the anchoring of expectations.

Motivated by this insight, we propose a strategy, which we call RAIT, that adjusts policy communication in response to shifts in the balance of inflation risks. In post-COVID counterfactuals, RAIT implies earlier tightening and earlier announcements of policy normalization once inflation risks started to recede.

# References


Adam, K. and R. M. Billi (2007): "Discretionary monetary policy and the zero lower bound on nominal interest rates," *Journal of Monetary Economics*, 54, 728–752.

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

Adrian, T., F. Duarte, N. Liang, and P. Zabczyk (2020): "NKV: A New Keynesian Model with Vulnerability," *AEA Papers and Proceedings*, 110, 470–76.

Adrian, T., F. Grinberg, N. Liang, S. Malik, and J. Yu (2022): "The Term Structure of Growth-at-Risk," *American Economic Journal: Macroeconomics*, 14, 283–323.

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

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.

Aruoba, S. B. and T. Drechsel (2026): "Identifying Monetary Policy Shocks: A Natural Language Approach," *American Economic Journal: Macroeconomics*.

Ascari, G., P. Bonomolo, and A. Celani (2024): "The Macroeconomic Effects of Inflation Expectations: The Distribution Matters," Tech. Rep. 18937, Centre for Economic Policy Research.

Ascari, G., S. Fasani, J. Grazzini, and L. Rossi (2023): "Endogenous uncertainty and the macroeconomic impact of shocks to inflation expectations," *Journal of Monetary Economics*, 140, 48–63.

Barnichon, R. and G. Mesters (2020): "Identifying modern macro equations with old shocks," *The Quarterly Journal of Economics*, 135, 2255–2298.

Barro, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.

Bianchi, F. (2016): "Methods for measuring expectations and uncertainty in Markov-switching models," *Journal of Econometrics*, 190, 79–99.

Bianchi, F. and L. Melosi (2017): "Escaping the Great Recession," *American Economic Review*, 107, 1030–58.

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

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.

Bloom, N. (2009): "The impact of uncertainty shocks," *Econometrica*, 77, 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, 86, 1031–1065.

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

——— (2022): "Recoverability and Expectations-Driven Fluctuations," *The Review of Economic Studies*, 89, 214–239.

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.

Cieslak, A., S. Hansen, M. McMahon, and S. Xiao (2026): "Policymakers' uncertainty," Tech. rep., National Bureau of Economic Research.

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

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

Coibion, O. and Y. Gorodnichenko (2026): *Expectations matter: The new causal macroeconomics of surveys and experiments*, Princeton University Press.

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., 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., M. P. Giannoni, and C. Patterson (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.

Dew-Becker, I. (2022): "Real-Time Forward-Looking Skewness over the Business Cycle," Tech. rep., National Bureau of Economic Research.

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

Erceg, C. J., J. Lindé, and M. Trabandt (2024): *Monetary policy and inflation scares*, International Monetary Fund.

Evans, C., J. Fisher, F. Gourio, and S. Krane (2016): "Risk management for monetary policy near the zero lower bound," *Brookings papers on economic activity*, 2015, 141–219.

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

Fernald, J. G. (2012): "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

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

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramirez, and M. Uribe (2011): "Risk matters: The real effects of volatility shocks," *American Economic Review*, 101, 2530–2561.

Fernández-Villaverde, J. and P. A. Guerrón-Quintana (2020): "Uncertainty shocks and business cycle research," *Review of economic dynamics*, 37, S118–S146.

Fernández-Villaverde, J., T. Mineyama, and D. Song (2024): "Are We Fragmented Yet? Measuring Geopolitical Fragmentation and Its Causal Effect," NBER Working Papers 32638, National Bureau of Economic Research, Inc.

Forni, M., L. Gambetti, N. Maffei-Faccioli, and L. Sala (2024): "The Effects of Monetary Policy on Macroeconomic Risk," *European Economic Review*, 167, 104800.

Gabaix, X. (2008): "Variable Rare Disasters: A Tractable Theory of Ten Puzzles in Macro-Finance," *American Economic Review*, 98, 64–67.

——— (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 127, 645–700.

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

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

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

Gourio, F. (2008): "Disasters and Recoveries," *American Economic Review*, 98, 68–73.

——— (2012): "Disaster Risk and Business Cycles," *American Economic Review*, 102, 2734–2766.

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 (2026): "How likely is an inflation disaster?" *The Review of Financial Studies*, hhaf058.

Iseringhausen, M., I. Petrella, and K. Theodoridis (2026): "Aggregate Skewness and the Business Cycle," *The Review of Economics and Statistics*, forthcoming.

Jarociński, M. and P. Karadi (2020): "Deconstructing monetary policy surprises—the role of information shocks," *American Economic Journal: Macroeconomics*, 12, 1–43.

Jordà, Ò. (2005): "Estimation and inference of impulse responses by local projections," *American economic review*, 95, 161–182.

Jordà, Ò., M. Schularick, and A. M. Taylor (2024): "Disasters Everywhere: The Costs of Business Cycles Reconsidered," *IMF Economic Review*, 72, 116–151.

Jordà, Ò. and A. M. Taylor (2025): "Local projections," *Journal of Economic Literature*, 63, 59–110.

Jurado, K., S. C. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105, 1177–1216.

Känzig, D. R. (2021): "The macroeconomic effects of oil supply news: Evidence from OPEC announcements," *American Economic Review*, 111, 1092–1125.

Kliem, M., A. Kriwoluzky, and S. Sarferaz (2016): "On the Low-Frequency Relationship Between Public Deficits and Inflation," *Journal of Applied Econometrics*, 31, 566–583.

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, 103570.

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.

Montiel Olea, J. L. and M. Plagborg-Møller (2021): "Local projection inference is simpler and more robust than you think," *Econometrica*, 89, 1789–1823.

Mouabbi, S., J.-P. Renne, and A. Tschopp (2025): "Inflation and Growth Risk: Balancing the Scales with Surveys," *Available at SSRN 5274480*.

Romer, C. D. and D. H. Romer (2010): "The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks," *American economic review*, 100, 763–801.

Salgado, S., F. Guvenen, and N. Bloom (2019): "Skewed Business Cycles," Tech. rep., National Bureau of Economic Research.

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

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

Stock, J. H. and M. W. Watson (2007): "Why has US inflation become harder to forecast?" *Journal of Money, Credit and Banking*, 39, 3–33.

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

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

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

# The Taming of the Skew: Asymmetric Inflation Risk and Monetary Policy Online appendix
(Not for publication)

# 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 $\varepsilon_t^a$ 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

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

where the matrices $\mathbf{\Theta}_0$, $\mathbf{\Theta}_1$, $\mathbf{\Theta}_y$, and $\mathbf{\Theta}_z$ are functions of $\mathbf{A}_0$, $\mathbf{A}_b$, $\mathbf{A}_f$, $\mathbf{B}_s$ and $\mathbf{b}_a$ (see, e.g., Sims, 2002, Sec. 4), whereas $\mathbf{\Xi}$ 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 $\mathbf{z}_t = \mathbf{\Theta}_1\mathbf{z}_{t-1} + \mathbf{\Theta}_0\left[\boldsymbol{\epsilon}_t^s\ \epsilon_t^a\right]'$. □

## 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 $X \sim \mathsf{f}(m, \sigma, \varrho)$ takes the form $EX = m + \kappa\sigma\varrho$, where $m$ represents the mode of the distribution, $\sigma$ captures the dispersion, $\varrho$ summarizes the asymmetry of the distribution about the mode, $\kappa$ is a constant that depends on the choice of $\mathsf{f}$. By relaxing the i.i.d. assumption for $\epsilon_t^a \sim \mathcal{F}(0, \sigma, \varrho_{a,t})$, and assuming the shock to be serially independent, we can write

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

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

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

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 $\mathsf{M}(j,h) = \mathbf{\Theta}_1^h\mathbf{\Theta}_y\mathbf{\Theta}_f^{j-1}\mathbf{\Theta}_z\mathbf{\Xi}$ 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 Multimodal Risk and Mean–Mode Asymmetry

The analysis in the main text establishes a tight link between asymmetry and the mean–mode wedge for a two-piece distribution where, by design, the mode is not affected by perturbations in the asymmetry. This appendix shows that the same mechanism extends beyond that class of distributions. In particular, we show that the result continues to hold in more general environments in which the shock distribution may be multimodal, provided it admits a well-defined dominant

mode. Holding fixed the dominant modal realization, first-order movements in the mean are driven by shifts in asymmetry and are therefore linked to Pearson skewness.

## B.1 General setting

Let $Z$ be a real-valued random variable with distribution $F$ admitting a density $f$ with respect to Lebesgue measure. Assume that $\mathbb{E}[Z]$ and $\mathrm{Var}(Z)$ exist and that the standard deviation $\sigma \equiv \sqrt{\mathrm{Var}(Z)}$ is strictly positive.

**Definition.** The distribution $F$ is said to be *multimodal with a dominant mode* at $m \in \mathbb{R}$ if the following conditions hold:

(i) $m \in \arg\max_{x \in \mathbb{R}} f(x)$;

(ii) for any $x \neq m$ such that $x$ is a local maximizer of $f$, $f(x) < f(m)$;

(iii) there exists $\varepsilon_0 > 0$ such that for all sufficiently small $\varepsilon \in (0, \varepsilon_0)$,

$$\int_{m-\varepsilon}^{m+\varepsilon} f(t)\,dt \;>\; \sup_{x \neq m} \int_{x-\varepsilon}^{x+\varepsilon} f(t)\,dt.$$

Related notions of unimodality, weak unimodality, and modal dominance in densities are discussed in Dharmadhikari and Joag-Dev (1988) and Silverman (1986).

This general definition encompasses a wide array of distributional shapes, including multiple local maxima, while ensuring that a single outcome $m$ plays a distinguished role. Condition (i) identifies $m$ as a most likely realization in the pointwise sense. Condition (ii) rules out competing modes of equal likelihood, ensuring uniqueness of the dominant modal location. Condition (iii) imposes a weak non-degeneracy requirement: realizations in an arbitrarily small neighborhood of $m$ carry strictly more probability mass than neighborhoods of equal width around any other local mode. Importantly, this condition does not require that most probability mass be concentrated near $m$, but only that $m$ remain the locally most representative realization of the distribution.

Let $\mu \equiv \mathbb{E}[Z]$ denote the mean of the distribution. Given the existence of a dominant mode, the *Pearson first coefficient of skewness* is defined as

$$\psi \;\equiv\; \frac{\mu - m}{\sigma}.$$

Rewriting $\mu - m = \psi\sigma$ explicitly shows that the mean-mode wedge is sufficient to capture the degree of asymmetry in the class of distributions considered in this section.

Let us now consider changes in risk that alter the shape of the distribution $f$ through the reallocation of probability mass across outcomes, while leaving the dominant mode $m$ unchanged. Such changes may affect higher moments, tail behavior, or the configuration of secondary modes, but preserve the most likely realization of the shock. For any such change, the resulting variation in the mean satisfies

$$\Delta m = \sigma\Delta\psi + \psi\Delta\sigma.$$

Let $\psi_0 \equiv (\mu_0 - m)/\sigma_0$ denote the Pearson skewness coefficient of a reference distribution with dominant mode $m$. A local perturbation of the distribution induces

$$\Delta(\mu - m) = \sigma_0\Delta\psi + \psi_0\Delta\sigma + o(\|\Delta\|),$$

where $o(\|\Delta\|)$ collects higher-order terms. Assuming $|\psi_0| < 1$, that is the initial mean-mode wage is not larger than one standard deviation, the contribution of scale changes is small, unless variations in dispersion respond excessively to variations in asymmetry, that is[24]

$$\lim_{\psi \to 0} \frac{\psi \, d\sigma}{\sigma \, d\psi} = 0.$$

Under this condition, the second term in the decomposition above is asymptotically negligible relative to the first, and the mean–mode wedge remains locally proportional to the Pearson skewness coefficient even when the standard deviation is allowed to vary. Economically, this restriction requires that the mean–mode gap be small relative to the overall dispersion of the shock distribution.Under these conditions, asymmetric risk continues to admit a first-order characterization in terms of skewness alone.

## B.2 An illustrative example: mixture of normals

As a simple illustration, consider a mixture distribution with two Gaussian components. Let

$$X_1 \sim \mathcal{N}(\mu_1, \sigma^2), \qquad X_2 \sim \mathcal{N}(\mu_2, \sigma^2),$$

with $\mu_1 > \mu_2$, and let $p \in [0, 1]$ denote the mixture weight on $X_1$. The random variable $Y$ is drawn from the mixture distribution with density

$$g(y) = p\,\varphi(y; \mu_1, \sigma^2) + (1 - p)\,\varphi(y; \mu_2, \sigma^2),$$

where $\varphi(\cdot; \mu, \sigma^2)$ denotes the normal density with mean $\mu$ and variance $\sigma^2$.

Let the unique, dominant mode be defined as $m_y = \sup_{y \in \mathbb{R}} g(y)$. To ensure that the mixture admits a well-defined dominant mode, assume that (i) $|\mu_1 - \mu_2| > 2\sigma$ and (ii) $p \in (0.5, 1)$. The first condition guarantees bimodality of the mixture density, while the second ensures that the higher-mean component generates the dominant mode.[25] Under these conditions, the mode of $g$ is uniquely attained in a neighborhood of $\mu_1$.

The expected value of the mixture is given by

$$\begin{aligned} E[Y] &= p\mu_1 + (1 - p)\mu_2 \\ &= \mu_1 + (1 - p)\Delta, \end{aligned} \tag{B5}$$

where $\Delta \equiv \mu_2 - \mu_1 < 0$. The term $(1-p)\Delta$ captures the extent to which probability mass assigned to the lower-mean component shifts the mean away from the dominant outcome. As $p \to 1$, the mixture distribution converges weakly to $\mathcal{N}(\mu_1, \sigma^2)$.

Changes in the perceived balance of risks are naturally mapped into changes in the mixture probability $p$. News that increase the likelihood of favorable outcomes correspond to a reduction in $p$, reallocating probability mass toward the right tail of $g(y)$. Such shifts leave the dominant modal

[24] This condition holds in most economic environments in which the dominant outcome is representative of typical realizations. Cases with $|\psi_0| \geq 1$ correspond to extreme asymmetry, where most of the mass is placed far away from the single modal value.

in which higher-order terms and scale effects are generically first order and a local skewness-based characterization is no longer informative.

[25] See Silverman (1986) and Ray and Lindsay (2005) for rigorous analyses of modality and dominant modes in Gaussian mixtures.

outcome unchanged while increasing asymmetry and raising the mean relative to the mode. As a result, changes in $p$ primarily affect the Pearson skewness coefficient, which moves proportionally to $(1-p)$.

This mixture-of-normals example provides a concrete illustration of the general mechanism emphasized in the previous section: changes in asymmetric risk, implemented here through probability reweighting rather than changes in scale, translate into movements in expected outcomes even in a multimodal environment.

# C Asymmetric Risk in a Nonlinear New Keynesian Model

In this appendix, we provide additional details for the computational exercise presented in Section 2.3. In the first part, we derive the baseline New Keynesian model (see, e.g., Galí, 2008). We then describe how shifts in the asymmetry of inflation risk are incorporated into this framework and compare the resulting dynamics with the linear approximation obtained through the beliefs representation, as outlined in Section 2.2.

While the main text focuses on shifts in risk arising from changes in skewness within a two-piece distribution, the appendix concludes by replicating the exercise in an alternative environment where asymmetry emerges from the presence of a second mode in the distribution of the cost-push shock. In this case, the distribution features a dominant mode centered at zero alongside a secondary mode capturing the possibility of a large cost-push shock.

## C.1 A simplified 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 in a monopolistically competitive goods market. Price adjustment is subject to quadratic costs, following Rotemberg (1982). Monetary policy is conducted according to a standard Taylor rule with interest rate inertia. The economy is subject to a stochastic cost-push shock, which we model as an exogenous shift in the elasticity of substitution in the demand schedule.

**Households.** The representative household maximizes expected lifetime utility:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left( \frac{C_t^{1-\sigma}}{1-\sigma} - \chi \frac{N_t^{1-\varphi}}{1-\varphi} \right) \tag{C1}$$

where $\beta \in (0,1)$ is the discount factor, $\sigma > 0$ is the coefficient of relative risk aversion, $\varphi$ is the inverse Frisch elasticity of labor supply, and $\chi > 0$ is a scaling parameter for labor disutility. The maximization is subject to the following flow budget constraint in real terms:

$$C_t + b_t = w_t N_t + \frac{R_{t-1}}{\Pi_t} b_{t-1} + d_t \tag{C2}$$

where $C_t$ denotes the household's consumption of the final good and $N_t$ represents the total hours of labor supplied to firms. The variable $b_t \equiv B_t/P_t$ denotes the real value of one-period risk-free nominal bonds purchased at time $t$, paying the gross nominal interest rate $R_t$ in period $t+1$. The term $d_t$ represents real dividends distributed by the monopolistically competitive firms, and

$w_t$ denotes the real wage. $\Pi_t \equiv P_t/P_{t-1}$ represents the gross rate of inflation, where $P_t$ is the aggregate price level.

Optimality conditions yield the labor supply equation $w_t = \chi N_t^{\eta} C_t^{\sigma}$ and the consumption Euler equation:

$$1 = \beta R_t E_t \left[ \left( \frac{C_{t+1}}{C_t} \right)^{-\sigma} \frac{1}{\Pi_{t+1}} \right]. \tag{C3}$$

**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 $\alpha$ 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. Firms take the aggregate price level and aggregate consumption as given. A linear technology $Y_t = N_t$, implies that the real marginal cost equals the real wage ($mc_t = w_t$). Firms face quadratic price adjustment costs $\frac{\psi}{2}(\frac{\Pi_t}{\bar{\Pi}} - 1)^2 Y_t$, where $\psi \geq 0$ governs the degree of nominal rigidity and $\bar{\Pi}$ is the central bank's inflation target.

Profit maximization yields the nonlinear New Keynesian Phillips Curve (NKPC):

$$\psi \tilde{\Pi}_t \Pi_t = 1 - \epsilon_t + \epsilon_t w_t + \beta E_t \left[ \Lambda_{t,t+1} \psi \tilde{\Pi}_{t+1} \Pi_{t+1} \frac{Y_{t+1}}{Y_t} \right] \tag{C4}$$

where $\epsilon_t$ is the stochastic elasticity of substitution, $\tilde{\Pi}_t \equiv \Pi_t/\bar{\Pi} - 1$ represents the net inflation gap, and $\Lambda_{t,t+1} \equiv (C_{t+1}/C_t)^{-\sigma}$ is the relevant portion of the household's stochastic discount factor.

**Monetary Policy.** The central bank sets the nominal interest rate $R_t$ according to an inertial Taylor rule responding to deviations of inflation and output from their targets:

$$R_t = R_{t-1}^{\rho_r} \left[ \bar{R} \left( \frac{\Pi_t}{\bar{\Pi}} \right)^{\phi_\pi} \left( \frac{Y_t}{\bar{Y}} \right)^{\phi_y} \right]^{1-\rho_r} \tag{C5}$$

where $\rho_r \in [0, 1)$ determines the degree of interest rate smoothing, and $\phi_\pi > 1$ and $\phi_y \geq 0$ are the reaction coefficients to inflation and the output gap, respectively. $\bar{R}$ and $\bar{Y}$ denote the steady-state values of the nominal interest rate and output.

**Stochastic Process for the Cost-Push Shock.** We introduce an exogenous process $u_t$ that perturbs the elasticity of substitution between varieties, denoted by $\epsilon_t$:

$$\epsilon_t = \bar{\epsilon} \exp(-u_t) \tag{C6}$$

where $\bar{\epsilon} > 1$ is the steady-state elasticity of substitution. A positive realization of $u_t$ reduces the elasticity of substitution, thereby increasing the desired markup and generating inflationary pressure.

In each period, agents observe the realization of the cost-push shock, $u_t$, along with a signal characterizing the asymmetry of the distribution governing the realization of the cost-push shock in the subsequent period. The shock process is assumed to be i.i.d., with a distribution that

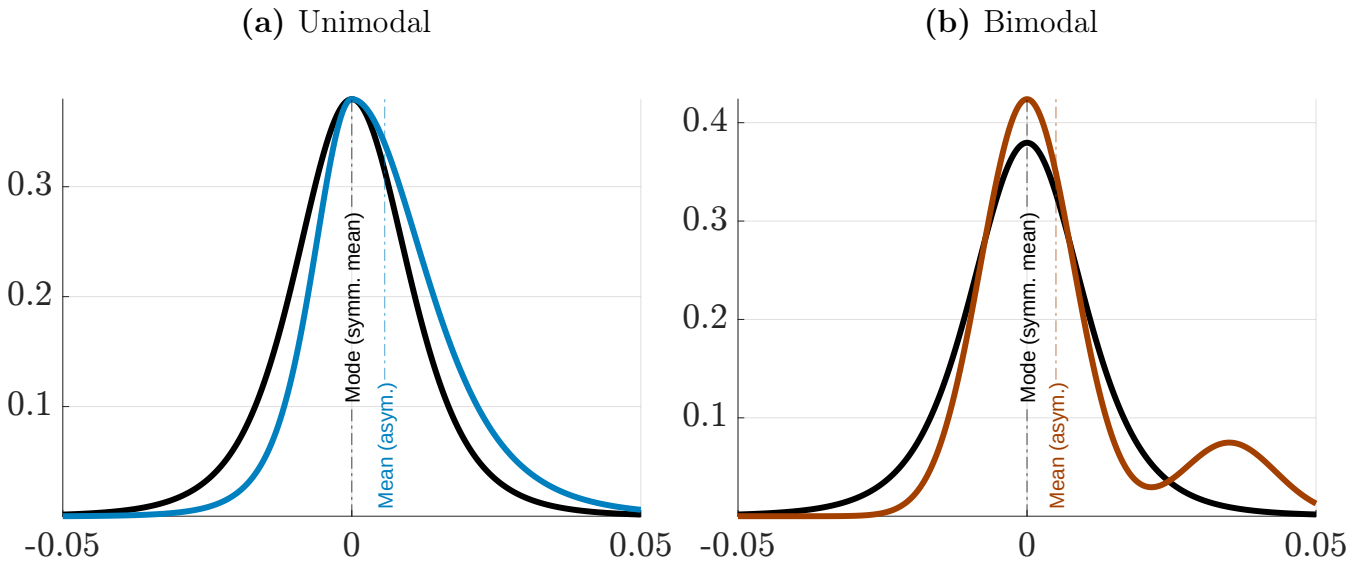



**Figure C1:** Conditional distributions of future cost-push shocks.

*Note:* The green line represents the Normal state. The red line represents the News state, characterized by positive skewness and a shifted conditional mean. In the left panel we report a unimodal skewed-distribution; we set the location to 0, unit scale and an asymmetry parameter of 0.3. In the right panel we report a gaussian mixture as the asymmetric risk case. We assume common variance for the two normals, and means of 0 and 0.035, respectively. The mixing probability assigns a probability of 0.8 to the first distribution.

may be symmetric or asymmetric. In the symmetric benchmark, the shock follows a standard Normal distribution. In the asymmetric case, we consider two alternative environments: a) the next-period cost-push shock follows a skewed Normal distribution with positive asymmetry; b) the shock distribution is a mixture of two Normal distributions, featuring a dominant mode at zero and a secondary mode corresponding to a "risk scenario" in which a large positive cost-push shock may occur.

Figure C1 illustrates the density functions associated with these two cases. Both specifications capture positively skewed risk. In each environment, while the modal realization of the future shock remains at zero, agents assign a higher probability to large inflationary outcomes in period $t+1$, either through increased probability mass in the right tail or through the emergence of a secondary mode.

Anticipated asymmetry in the distribution of future shocks generates a wedge between the most likely realization and the expected value of the shock. Although the modal realization remains unchanged, the shift in expectations affects current decisions through expectation terms in both the Euler equation and the New Keynesian Phillips Curve.

## C.2 Numerical Solution Method

We solve the nonlinear model using the global policy iteration algorithm implemented in the GDSGE framework (Cao et al., 2023). This approach allows us to accurately characterize the effects of shifts in perceived cost-push risk in a fully nonlinear environment, without relying on local perturbation methods.

**State space.** The state space consists of one endogenous state variable, the lagged nominal interest rate $R_{t-1}$, and an exogenous state $S_t$ governing the realization and perceived distribution of the cost-push shock. Formally, $S_t$ indexes both the current realization of the cost-push shock and the conditional distribution governing its realization in the subsequent period. The latter is represented by a transition matrix over a discrete grid of possible future states.

By appropriately calibrating the transition probabilities $\varpi$, we allow for situations in which the modal realization of the next-period cost-push shock remains at zero, while the conditional distribution becomes asymmetric. As discussed in the main text, we consider two alternative sources of asymmetry: (i) a unimodal environment in which the next-period shock follows a positively skewed distribution; and (ii) a bimodal environment featuring a dominant mode at zero and a secondary mode associated with a low-probability, high-inflation "risk scenario." In both cases, asymmetry operates through changes in the conditional distribution rather than through realized shocks.[26]

**Solution algorithm.** The solution procedure proceeds as follows.

1. **Discretization.** We discretize the continuous endogenous state $R_{t-1}$ on a dense grid of $N_R$ points. The exogenous state $S_t$ follows a finite-state Markov chain with $J$ nodes, corresponding to alternative perceived distributions of the next-period cost-push shock.

2. **Function approximation.** Let $x_t = \{C_t, \Pi_t, N_t, w_t, R_t, Y_t\}$ denote the vector of policy variables. We approximate the policy functions $\mathcal{P}(R_{t-1}, S_t)$ using multidimensional linear interpolation over the discretized state space.

3. **Policy iteration.** At iteration $n$, for each grid point $(R_i, S_k)$, we solve the system of nonlinear equilibrium conditions $F(\cdot) = 0$, defined by equations (C1)–(C6). Conditional expectations appearing in the Euler equation and the New Keynesian Phillips Curve are computed by integrating over future states using the transition matrix:

$$\mathbb{E}_t\left[f(x_{t+1})\right] = \sum_{m=1}^{J} \varpi_{km}\, f\left(\mathcal{P}^{(n-1)}(R_t, S_m)\right). \tag{C7}$$

   This step explicitly incorporates the anticipation channel associated with asymmetric risk. When agents perceive the next-period cost-push shock distribution to be skewed—either through increased right-tail probability mass or the emergence of a secondary mode—the transition probabilities shift weight toward higher future cost realizations. Importantly, these changes affect equilibrium outcomes even though the modal realization of the shock remains unchanged, isolating the role of perceived asymmetry from that of realized shocks.

4. **Convergence.** The policy functions are updated until convergence, defined as

$$\left\|\mathcal{P}^{(n)} - \mathcal{P}^{(n-1)}\right\|_\infty < \delta, \qquad \delta = 10^{-6}. \tag{C8}$$

[26] To reduce computational complexity, we assume that in the symmetric benchmark the cost-push shock always realizes at zero. With some probability, the economy transitions to a state in which agents perceive the next-period distribution of the shock to be asymmetric. In that case, the asymmetric distribution is discretized over $\tilde{J} = 7$ nodes, with probabilities reflecting the degree of skewness or the weight on the risk scenario. We verified that increasing the number of grid points for both $R_{t-1}$ and the shock discretization does not materially affect the results.

The resulting global solution allows us to study the macroeconomic effects of revisions in the perceived distribution of future cost-push shocks, fully accounting for household risk aversion and firms' nonlinear price-setting behavior.

## C.3 Simulating the response to shift in risk

**Table C1:** Model Parameter Calibration

| **Parameter** | | **Value** |
|---|---|---|
| *Steady State Values* | | |
| Inflation Target | $\bar{\pi}$ | 1.005 |
| Steady State Output | $\bar{y}$ | 1 |
| Steady State Nominal Interest Rate | $\bar{r}$ | 1.0101 |
| *Structural Parameters* | | |
| Discount Factor | $\beta$ | 0.995 $(\bar{\pi}/\bar{r})$ |
| Inverse EIS | $\sigma$ | 1 |
| Elasticity of Substitution | $\bar{\epsilon}$ | 6 |
| Price Adjustment Cost | $\psi$ | 25 |
| Inverse Frisch Elasticity | $\eta$ | 1.33 |
| Labor Disutility Weight | $\chi$ | 0.833 $((\bar{\epsilon}-1)/(\bar{\epsilon}\bar{y}^{\eta+\sigma}))$ |
| *Monetary Policy Rule* | | |
| Taylor Rule: Inflation | $\phi_\pi$ | 1.5 |
| Taylor Rule: Output | $\phi_y$ | 0.125 |
| Interest Rate Smoothing | $\rho_r$ | 0.8 |

*Note:* The table reports the parameters' calibration used in the simulation exercise.

Within this environment, we simulate a scenario in which agents perceive the distribution of the next-period cost-push shock to be asymmetric, while the realized shock equals zero in the subsequent period and thereafter. This design cleanly isolates the effects of changes in perceived risk from those associated with realized shocks.

By construction, the experiment generates a pure belief shock: agents revise expectations in response to perceived asymmetry, even though no cost-push shock ultimately materializes. Consequently, any movements in inflation, output, or the policy rate arise solely from changes in expectations rather than from contemporaneous or realized disturbances.

This setup allows us to directly assess the quantitative importance of the belief-driven channel emphasized by the beliefs representation. In particular, it enables a transparent comparison between the global nonlinear solution and its first-order approximation, highlighting the extent to which perceived risk asymmetry alone can generate meaningful macroeconomic dynamics.

## C.4 Linearized model

Log-linearizing the model around the deterministic steady state yields the following system, where variables with hats denote log-deviations from steady state (e.g., $\hat{x}_t \equiv \ln(Y_t/\bar{Y})$):

$$\hat{x}_t = E_t\hat{x}_{t+1} - \frac{1}{\sigma}(\hat{i}_t - E_t\hat{\pi}_{t+1}) \tag{C9}$$

$$\hat{\pi}_t = \beta E_t\hat{\pi}_{t+1} + \kappa\hat{x}_t + \gamma u_t \tag{C10}$$

$$\hat{i}_t = \rho_i\hat{i}_{t-1} + (1-\rho_i)(\phi_\pi\hat{\pi}_t + \phi_x\hat{x}_t) \tag{C11}$$

where $\kappa \equiv \frac{\bar{\epsilon}-1}{\psi}(\sigma+\eta)$ represents the slope of the Phillips curve and $\gamma$ is proportional to the Rothemberg parameter $\psi$.

Within this linearized environment, we study the effects of shifts in perceived risk asymmetry by exploiting the beliefs-representation equivalence established in Section 2.2. Specifically, changes in the perceived asymmetry of future cost-push shocks are introduced as belief (anticipated) shocks that affect expectations but do not alter the realized path of fundamentals. This linear benchmark provides a transparent point of comparison for the global nonlinear results discussed above.

# D A model for time-varying inflation risk: additional results

This appendix reports additional details and results to complement what reported in Section 3.

## D.1 Data

Table D1 report the data used throughout the paper, collected from FRED. The data sample ranges from 1960Q1 to 2024Q4.

**Table D1:** Data series

| Inflation measures | | Macro data | |
|---|---|---|---|
| Series | Fred ticker | Series | Fred ticker |
| core PCE | `PCEPILFE` | Real GDP | `GDPC1` |
| Headline PCE | `PCEPI` | Hours worked | `HOANBS` |
| core CPI | `CPILFESL` | Fed funds rate | `FEDFUNDS` |
| Headline CPI | `CPIAUCSL` | | |
| GDP deflator | `GDPDEF` | | |

*Note:* The table reports the data series and the respective Fred tickers.

We collect real-time data for core PCE inflation and SPF predictions from the Real-Time Data Research portal of the Federal Reserve Bank of Philadelphia.

## D.2 The model

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 $\pi_t$ can be characterized by a Skew-t

$(Skt)$ distribution with time-varying location $(m_t)$, scale $(\sigma_t)$, and shape $(\varrho_t)$ parameters:

$$\pi_t \sim Skt_\nu(m_t, \sigma_t^2, \varrho_t),$$

where $\nu$ denotes the, time invariant, degrees of freedom. The degree of the asymmetry is risk is captured by the parameter $\varrho_t \in (-1, 1)$. We postulate that each parameter features a permanent and transitory component, such that the location is

$$m_t = \bar{m}_t + \tilde{m}_t \tag{D1}$$
$$\bar{m}_t = \bar{m}_t + a_m s_{m,t} \tag{D2}$$
$$\tilde{m}_t = \phi_m \tilde{m}_t + b_m s_{m,t}. \tag{D3}$$

Let $\delta_t = \log(\sigma_t)$, then

$$\gamma_t = \bar{\gamma}_t + \tilde{\gamma}_t \tag{D4}$$
$$\bar{\gamma}_t = \bar{\gamma}_t + a_\gamma s_{\gamma,t} \tag{D5}$$
$$\tilde{\gamma}_t = \phi_\gamma \tilde{\gamma}_t + b_\gamma s_{\gamma,t}, \tag{D6}$$

such that the scale parameter is multiplicative in its two components. Finally, for $\gamma_t = \text{arctanh}(\varrho_t)$ we have

$$\delta_t = \bar{\delta}_t + \tilde{\delta}_t \tag{D7}$$
$$\bar{\delta}_t = \bar{\delta}_t + a_\delta s_{\delta,t} \tag{D8}$$
$$\tilde{\delta}_t = \phi_\delta \tilde{\delta}_t + b_\delta s_{\delta,t}. \tag{D9}$$

More compactly, for $f_t = (m_t, \delta_t, \gamma_t)'$,

$$\bar{f}_{i,t} = \bar{f}_{i,t-1} + a_i s_{i,t-1}, \tag{D10}$$
$$\tilde{f}_{i,t} = \phi_i \tilde{f}_{i,t-1} + b_i s_{i,t-1}. \tag{D11}$$

The time-varying parameters of the distribution are learned in real-time, with updates being proportional to $s_{i,t-1}$, which is the *scaled score* of the conditional distribution (as in Creal et al., 2013; Harvey, 2013).

**Scaled score** 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), $y_t|Y_{t-1} \sim Skt_\nu(m_t, \sigma_t^2, \varrho_t)$, the log-likelihood takes the form

$$\ell_t(r_t|\theta, \mathcal{F}_{t-1}) = \log \mathcal{C}(\nu) - \frac{1}{2}\log \sigma_t^2 - \frac{1+\nu}{2}\log\left[1 + \frac{\varepsilon_t^2}{\nu(1 + sgn(\varepsilon_t)\varrho_t)^2\sigma_t^2}\right], \tag{D12}$$
$$\log \mathcal{C}(\nu) = \log \Gamma\left(\frac{\nu+1}{2}\right) - \log \Gamma\left(\frac{\nu}{2}\right) - \frac{1}{2}\log \nu - \frac{1}{2}\log \pi,$$

where $\Gamma(\cdot)$ is the Gamma function, $sgn(\cdot)$ is the sign function, and $\nu > 3$ are the degrees of freedom. Differentiating Equation (D12) with respect to location, scale and asymmetry we obtain the gradient vector $\nabla_t = \left[\frac{\partial \ell_t}{\partial m}, \frac{\partial \ell_t}{\partial \sigma_t^2}, \frac{\partial \ell_t}{\partial \varrho_t}\right]'$, where $\varepsilon_t = y_t - m_t$, and the Fisher information matrix is obtained as $\mathcal{I}_t = \mathbb{E}_{t-1}[\nabla_t \nabla_t']$. We can thus define the vector of interest as $f_t = (m_t, \gamma_t, \delta_t)'$ with

the associated Jacobian matrix

$$J_t = \frac{\partial(m_t, \sigma_t^2, \varrho_t)}{\partial f_t'} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 2\sigma_t^2 & 0 \\ 0 & 0 & 1-\varrho_t^2 \end{bmatrix}. \tag{D13}$$

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

$$\mathbf{s}_t = (J_t' \text{diag}(\mathcal{I}_t) J_t)^{-\frac{1}{2}} J_t' \nabla_t = \begin{bmatrix} s_{m,t} \\ s_{\gamma,t} \\ s_{\delta,t} \end{bmatrix} = \begin{bmatrix} \sqrt{\frac{(\nu+3)(1-\varrho_t^2)}{(\nu+1)}} w_t \zeta_t \\ \sqrt{\frac{(\nu+3)}{2\nu}} (w_t \zeta_t^2 - 1) \\ sgn(\varepsilon_t) \sqrt{\frac{(\nu+3)(1-sgn(\varepsilon_t)\varrho_t)}{3(\nu+1)(1+sgn(\varepsilon_t)\varrho_t)}} w_t \zeta_t^2 \end{bmatrix}. \tag{D14}$$

Full derivations of the score vector, its properties and the Information matrix are provided in Delle Monache et al. (2024).

Intuitively, the score vector maps the new information contained in the latest data release — summarized by the prediction error, $\varepsilon_t = \pi_t - m_t$ — into updates of the time-varying parameters that characterize the predictive distribution of inflation. The learning-rate parameters, $a_i$ and $b_i$, govern the strength of these updates.

In a Gaussian environment, updates of the location resemble standard Kalman-filter learning, in which parameter adjustments are proportional to the prediction error. The scale is adjusted with respect to its unit expected value; it increases (decreases) when above (below) one. Allowing for fat tails further renders the updating mechanism robust to large, unanticipated errors (see, e.g., Delle Monache and Petrella, 2017; Antolín-Díaz et al., 2024).

When asymmetry is introduced, the updating mechanism weights prediction errors differently depending on their sign and magnitude. Small standardized prediction errors are interpreted as information about the central scenario and lead to only minor revisions of the shape parameter. In contrast, large deviations from the central forecast are taken as evidence of a shift in the balance of risks, prompting updates to the asymmetry parameter in the direction of the surprise. Consistent with this mechanism, large deviations of inflation from the expected central scenario induce adjustments in skewness that mirror the sign of the prediction error.

## D.3 Bayesian estimation

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.

( i ) Sample $\theta^j$ from a random walk candidate density;

( ii ) Reject draws that do not respect parameter bounds;

(iii) Compute $\{f_1, ..., f_T|\theta^j\}$, the log-likelihood $\ell(y|\theta^j)$, and the posterior $\pi(\theta^j|y)$;

(iv) Accept or reject $\theta^j$ according to the MH rule and update the sampling variance;

( v ) To improve the convergence of the algorithm, we allow for delayed acceptance in case a draw gets rejected (see, e.g., Sherlock et al., 2017).

(vi) Compute the statistics of interest as the percentiles of the empirical distribution function.

The algorithm is rather efficient, and a complete chain of 50000 draws can be obtained in about 2 minutes. Section D in Delle Monache et al. (2024) provides detailed explanation of the algorithm.

### D.3.1 Priors

The initial values for the permanent component of time-varying parameters are drawn from a multivariate Gaussian distribution, $\bar{f}_0 \sim \mathcal{N}(\mathbf{m}_0, \mathbf{s}_0)$, with mean vector $\mathbf{m}_0$ and covariance matrix $\mathbf{s}_0$. The mean vector is fixed to $[2.5 \;\; 0 \;\; 0]'$, which captures the properties of the first ten years of data. The covariance matrix is computed as $\mathbf{s}_0 = E[\nabla'_{\mathbf{g}_0} \nabla_{\mathbf{g}_0}]^{-1}$, where $\nabla_{\mathbf{m}_0} = \frac{\partial \mathbf{g}_0}{\partial \mathbf{m}_{0,j}}$ is the gradient of $\mathbf{g}_0$, a function matching sample moments to the three parameters of interest. We use Minnesota-type priors for the the persistence of the transitory components, $\phi \sim NID(m_\phi, \sigma_\phi) \cdot I_{(\phi \in \Phi)}$. We target high persistence, with $m_\phi \lesssim 1$, with a standard tightness of $\sigma_m = 0.2$, in line with Bayesian Vector Autoregressive models (see, e.g., Doan et al., 1984; Sims and Zha, 1998). We restrict the prior distribution to only span the stationary region, $\Phi$, by truncating the support, which gives rise to an improper prior distribution as in Cogley and Sargent (2005). Loadings on the score components are Gamma distributed with shape parameter equal to 2 and the scale parameters are set to obtain mean and standard deviation equal to 0.015 and 0.01 for the permanent loadings, $a$, and 0.07 and 0.05 for the transitory loadings, $b$. These hyperparameter values correspond to an a-priori half-life of 10 and 2 years for the two components, reflecting 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 a gamma prior for $\nu$, ensuring $\nu > 3$, $\nu \sim \mathcal{G}(2, 10) \cdot I_{(\eta \in \mathcal{H})}$.

### D.3.2 MCMC Convergence Diagnostics

We consider three types of diagnostics: the graphical inspection of cumulative means, the analysis of variance of several chains, and a test on the ergodic averages. In these exercises we check the convergence of the post burn-in sample, $\Theta^*$, for our preferred model specification. We draw $N = 30000$ replications and we keep the last $n^* = 3000$ replications, after thinning each two posterior draws to reduce the autocorrelation of the chains.

As a first step, we inspect the cumulative averages of the retained draws. A smooth and stable plot suggests that the chain has converged to the expected value of its invariant distribution. The cumulative averages plotted in Figure D1 appear all smooth and converging to the single value $\mathbb{E}[\Theta^*]$, suggesting that the chains have reached convergence.

Next, we turn to analysis of variance techniques, as advocated by Gelman and Rubin (1992). Due to the efficiency of our algorithm, we consider $m = 100$ parallel chains to compare the dispersion *between*, $B$, and *within*, $W$, chains, and test that the former is greater than the latter. Specifically,

$$B = \frac{n^*}{m-1} \sum_{i=1}^{m} (\overline{\theta}_i^* - \overline{\theta}^*)^2 \quad \text{and} \quad W = \frac{1}{m(n^*-1)} \sum_{i=1}^{m} \sum_{j=1}^{J} (\theta_{i,j}^* - \overline{\theta}_i^*)^2,$$

where $\overline{\theta}_i^* = \frac{1}{n^*} \sum_{n=1}^{n^*} \theta_n^*$ and $\overline{\theta}^* = \sum_{i=1}^{m} \overline{\theta}_i^*$. These statistics are used to consistently estimate the

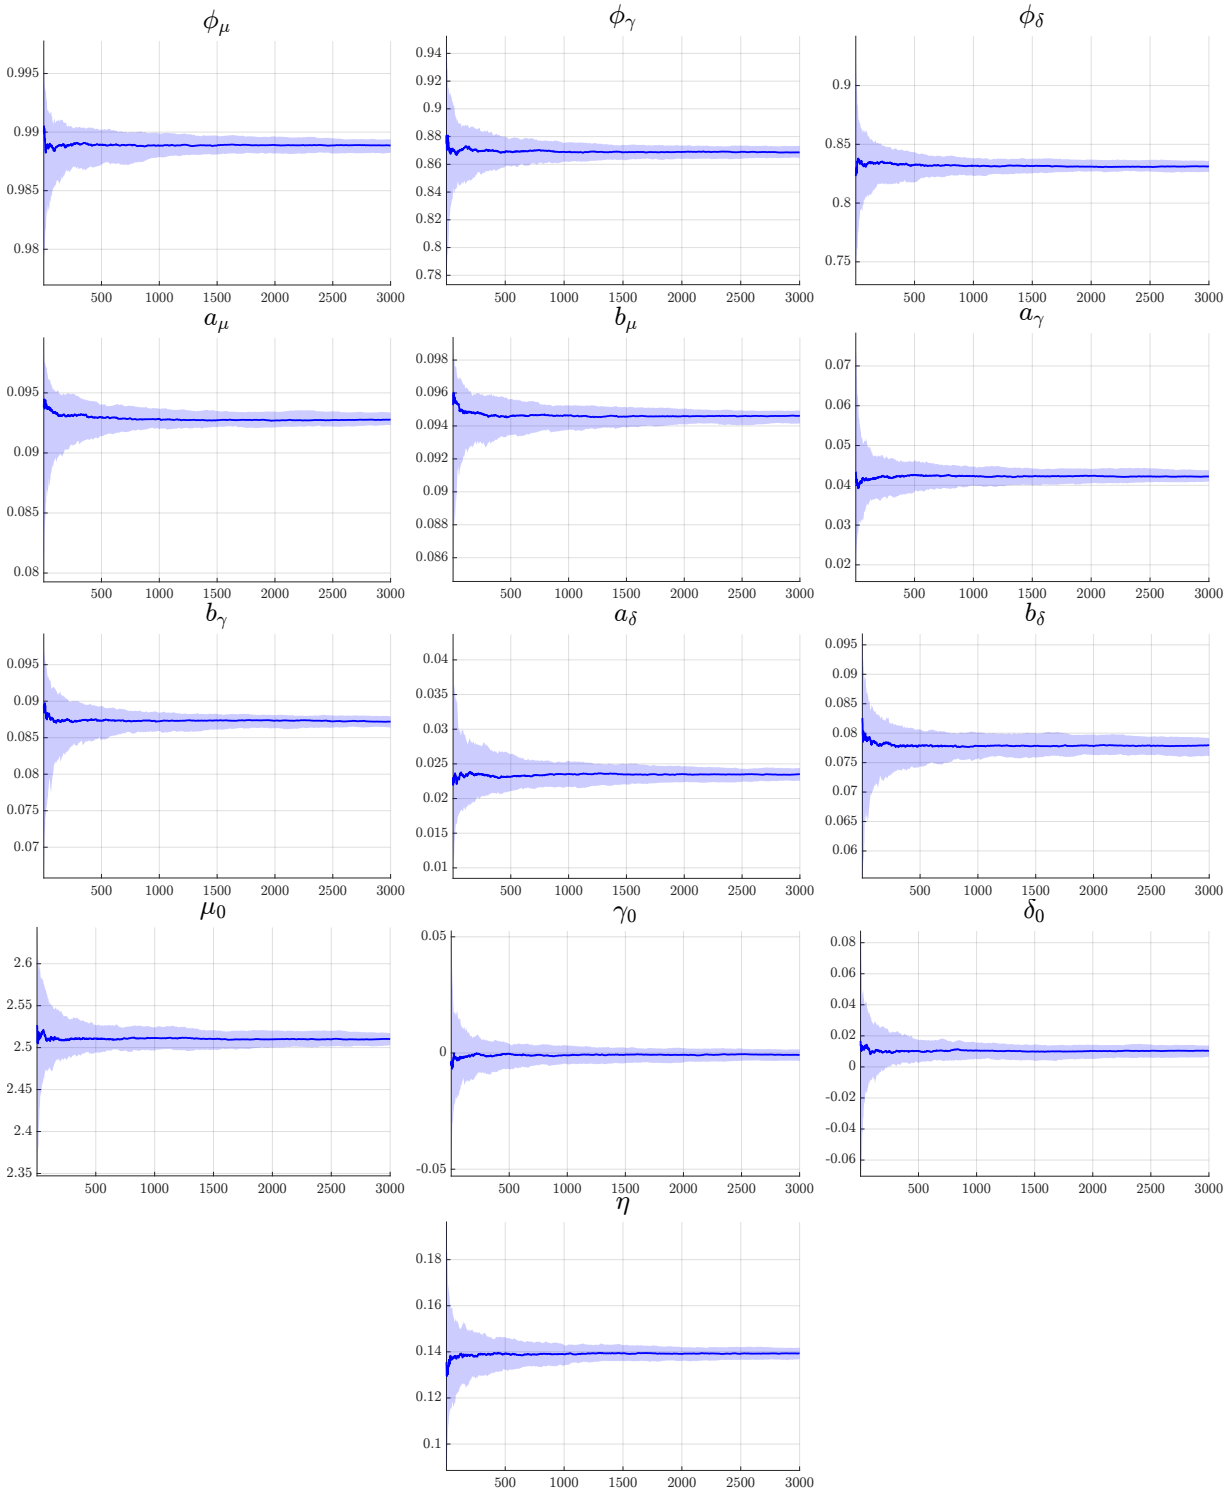



**Figure D1:** Cumulative averages

*Note:* The panels show the cumulative averages of the post burn-in sample. $\phi$s are the autoregressive parameters, $a$s and $b$s are the learning rate parameters for the transitory and permanent components, respectively, $m_0$, $\gamma_0$ and $\delta_0$ are the initial values of the permanent components and $\eta$ is the inverse of the degrees of freedom.

marginal posterior variance of $\theta^*$ as the weighted average of $W$ and $B$ as:

$$\hat{\sigma}^2_{\theta^*} = \frac{n-1}{n}W + \frac{1}{n}B.$$

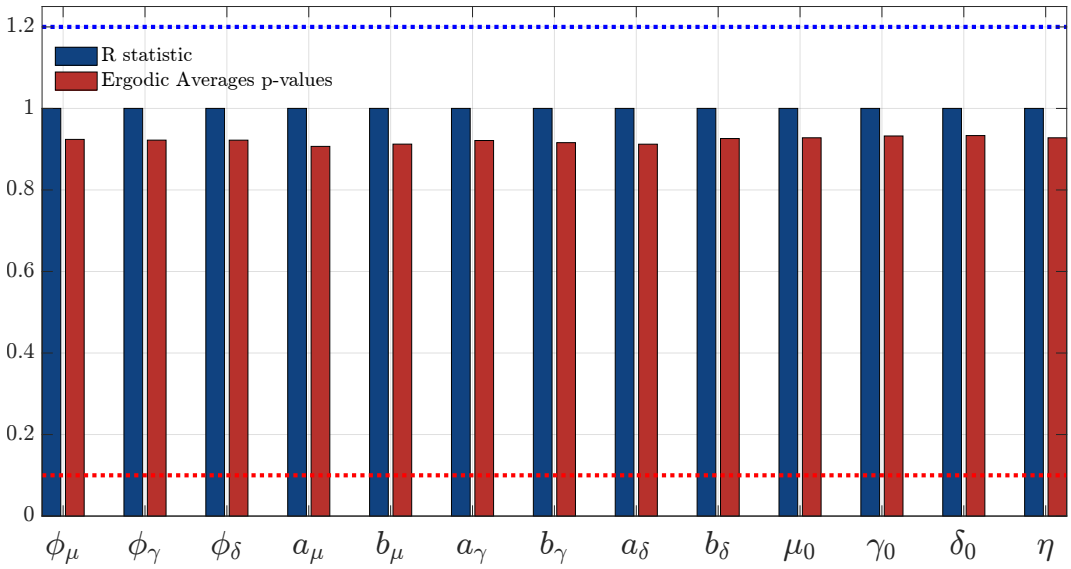



**Figure D2:** $R$ statistic and ergodic average p-values

*Note:* The bars represent Gelman and Rubin (1992) $R$ statistic (blue) and the p-values for Geweke (1992) test on ergodic averages (red) for each parameter of the model. The blue dotted line represents the critical value of 1.2 for the R statistics; the red dotted line represents the 10% critical value for the p-values. The statistics are computed for J=13 parameters and m=100 chains.

Convergence can thus be monitored by means of $\hat{R} = \sqrt{\hat{\sigma}^2_{\theta^*} W^{-1}} \geq 1$; Gelman (1996) suggest that $R < 1.2$ can be used as an indication to accept the convergence of the MCMC.

At last, we follow Geweke (1992) and test the ergodic averages of the time series of draws. We split the $n^*$ draws into a subsample $B$ of $n^*_b = 0.1n^*$ and another one $A$ of $n^*_a = 0.5n^*$, and compute

$$\overline{\theta}^*_b = \frac{1}{n^*_b} \sum_{j=1}^{n^*_b} \theta^*_j \quad \text{and} \quad \overline{\theta}^*_a = \frac{1}{n^*_a} \sum_{j=n^*_a}^{n^*} \theta^*_j.$$

As we only consider the post burn-in sample, $\overline{\theta}^*_b$ and $\overline{\theta}^*_a$ are the ergodic averages at the beginning and at the end of the convergence sample, and should therefore behave similarly. This can be tested by means of a simple *t-test* on the standardized difference between $\overline{\theta}^*_b$ and $\overline{\theta}^*_a$

$$z_G = \frac{\overline{\theta}^*_a - \overline{\theta}^*_b}{\sqrt{\hat{V}(\overline{\theta}^*_a) + \hat{V}(\overline{\theta}^*_b)}} \sim N(0,1).$$

We report the $R$ statistics and p-values for the ergodic averages in Figure D2. All the $R$ statistics (blue bars) are below the 1.2 threshold, identified by the dotted blue line. Similarly, the red bars show that we do not reject the null hypothesis for any of the parameters.

## D.4 Additional estimation results

The estimated time-varying moments of core PCE are reported in Figure D3, where we highlight the permanent components and the total moment in green and black, respectively. Specifically, moments are closed-form functions of the estimated time-varying parameters.

**Table D2:** Parameters estimates for the econometric model in Section 3

| Autocorrelations | | | | | |
|---|---|---|---|---|---|
| $\phi_m$ | $\phi_\gamma$ | $\phi_\delta$ | | | |
| 0.989 (0.007) | 0.866 (0.058) | 0.834 (0.062) | | | |
| **Learning rates** | | | | | |
| $a_m$ | $b_m$ | $a_\gamma$ | $b_\gamma$ | $a_\delta$ | $b_\delta$ |
| 0.092 (0.007) | 0.095 (0.005) | 0.043 (0.019) | 0.087 (0.010) | 0.024 (0.012) | 0.079 (0.016) |
| **Degrees of freedom** | | | | | |
| $\eta$ | | | | | |
| 0.141 (0.037) | | | | | |

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

$$E_{t-1}\pi_t = m_t + \underbrace{g(\eta)\sigma_t\varrho_t}_{\psi_t}, \qquad g(\eta) = \frac{4\mathcal{C}(\eta)}{1-\eta}, \tag{D15}$$

$$Var_{t-1}\pi_t = \sigma_t^2\left(\frac{1}{1-2\eta} + h(\eta)\varrho_t^2\right), \qquad h(\eta) = \frac{3}{1-2\eta} - g(\eta)^2, \tag{D16}$$

$$Skew_{t-1}\pi_t = \frac{g(\eta)\varrho_t\left[1+\eta-\varrho_t^2\left(5-2g(\eta)^2+(10g(\eta)^2-19)\eta-12g(\eta)^2\eta^2\right)\right]}{(1-3\eta)(1-2\eta)\left(\frac{1}{1-2\eta}+h(\eta)\varrho_t^2\right)^{\frac{3}{2}}}, \tag{D17}$$

noting that this moment only depends on the asymmetry parameter and on the estimated degrees of freedom (see De Polis, 2023).

We also report in Table D2 the estimated static parameters of the model, which consists of three autoregressive parameters for the transitory components, six learning rates for the pair of components of the time-varying parameters, and the fixed (inverse of the) degrees of freed on the distribution.

## D.5 Formally testing for inflation conditional skewness

We formally test for the evidence of time variation in the asymmetry of the predictive distributions of core PCE inflation. Assuming a flexible Skew-t specification for the likelihood function of inflation data, we test whether the asymmetry of the conditional distribution can be significantly assumed to be constant by examining the properties of the score function. Table D3 reports the results of three alternative parametric Lagrange Multiplier tests: a Q test, an adjusted Q* test,

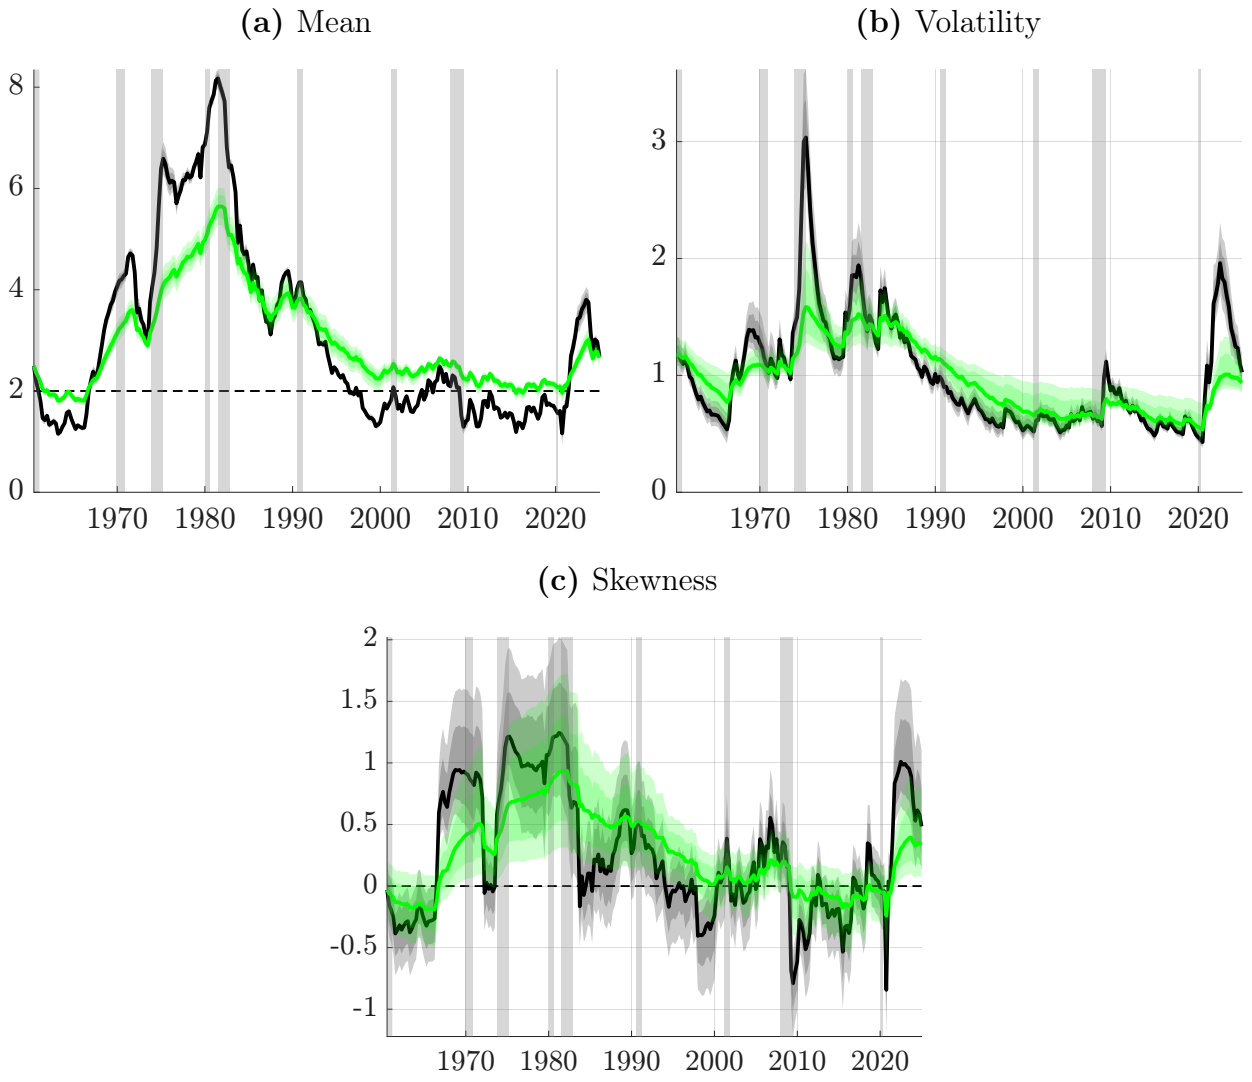



**Figure D3:** Time-varying moments of core PCE in the US

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

**Table D3:** 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 $\chi^2_1$, while N is distributed as a Cramer von-Mises distribution with 1 degree of freedom. * p $<$ 10%, ** p $<$ 5%, *** p $<$ 1%.

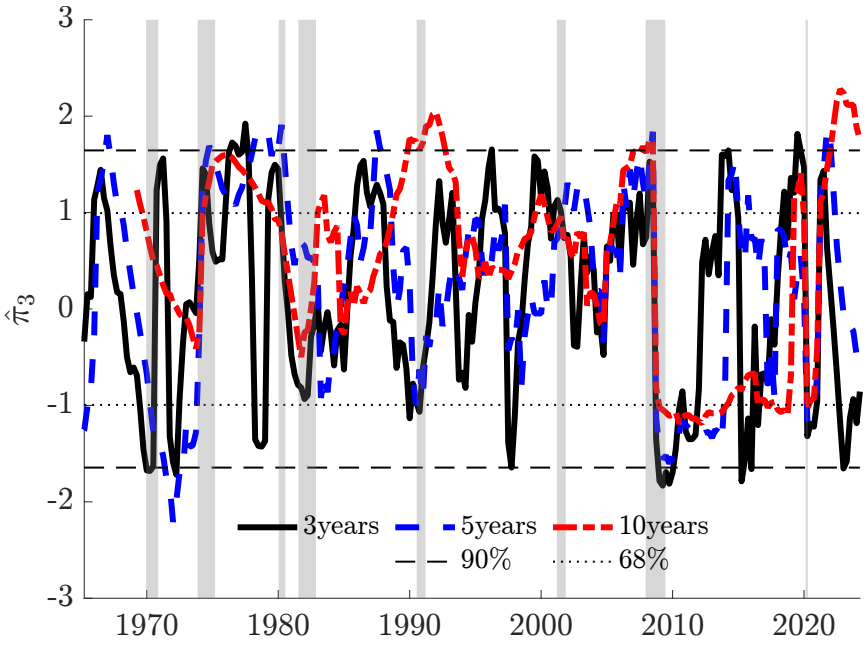



**Figure D4:** Bai and Ng (2005) rolling tests

*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.

and the Nyblom test, following Delle Monache et al. (2024).[27] All test strongly reject the null hypothesis of symmetry at the 1% confidence level; the right panel of the Table show an equally strong rejection after accounting for stochastic volatility. Similar results are obtained by applying Bai and Ng (2005) tests for unconditional skewness to an expanding window of data, as reported in Figure D4.

## D.6 Comparison with model-free measures of conditional skewness

Rolling window estimators of skewness strike a trade-off between precision and responsiveness. Longer windows improve statistical reliability by smoothing out outliers but reduce the ability to detect rapid changes in risk. They also assume constant skewness within the window, causing delayed adjustments during periods of sharp inflation movements, when timely assessment matters most.

To balance these trade-offs, we use 5-year rolling estimates and consider both sample skewness ($ss_t$) and quantile skewness ($qs_t$) as benchmarks. The former can be compared to our model's conditional skewness measure, $\mathtt{s}_t$, while the latter is conceptually closer to the asymmetry parameter, $\varrho_t$, as it captures the imbalance of probability mass around the mode. Unlike rolling statistics, our model applies one-sided discounting of past data to estimate time-varying moments, yielding more responsive and stable signals of inflation risk. Figure D5 shows that our model-based skewness estimates closely track rolling estimates, with maximum cross-correlations between 0.4 and 0.5.

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

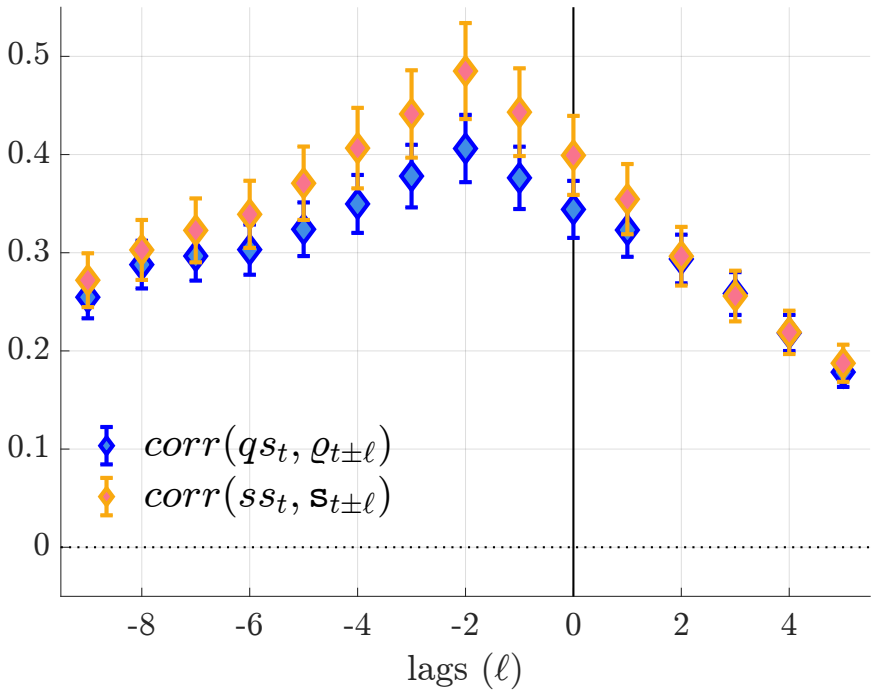



**Figure D5:** Cross-correlations

*Note:* The figure reports the cross-correlations between realized (ex-post) measures of inflation skewness—specifically, the 5-year rolling quantile skewness ($qs_t$) and the 5-year rolling sample skewness ($ss_t$)—and our model-based (ex-ante) measure of predictive asymmetry in risk, captured by the asymmetry parameter $\varrho_t$ and the conditional skewness $\mathtt{s}_t$, respectively.

## D.7 Time variation in inflation risk for other inflation measures

Table D4 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 D4 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.

Figure D6 shows the estimated dynamics of inflation volatility and skewness across the different measures, highlighting in blacks that of core PCE. Two comments are in order. First, the dynamics of the two moments is extremely similar for all measures. With varying magnitudes, volatilities spike around recessions, and remain persistently high soon after. Skewness follow humped-shape patterns in the 1970s and 1980s, then moving downward since the 1990s, remaining negative until the pandemic period. Second, it's important to note that, among all these measures, core PCE shows the least variation in both volatility and skewness, appearing to be the more stable measure of price dynamics. Based on previous results, we also report the estimates for the time-varying moments of the four measures.

**Table D4:** Time variation in higher order moments

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

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

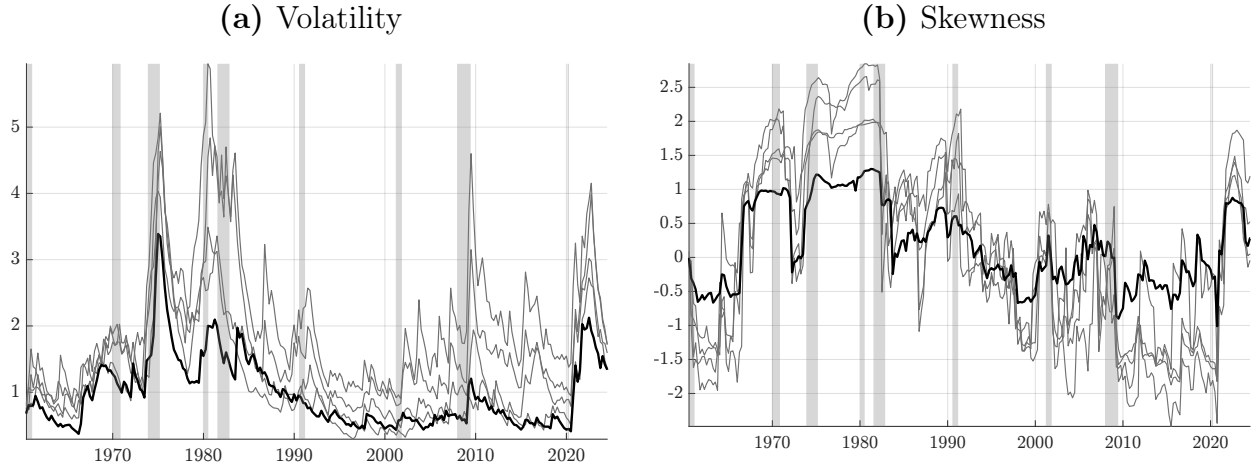



**Figure D6:** Risk across different inflation measures

*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.

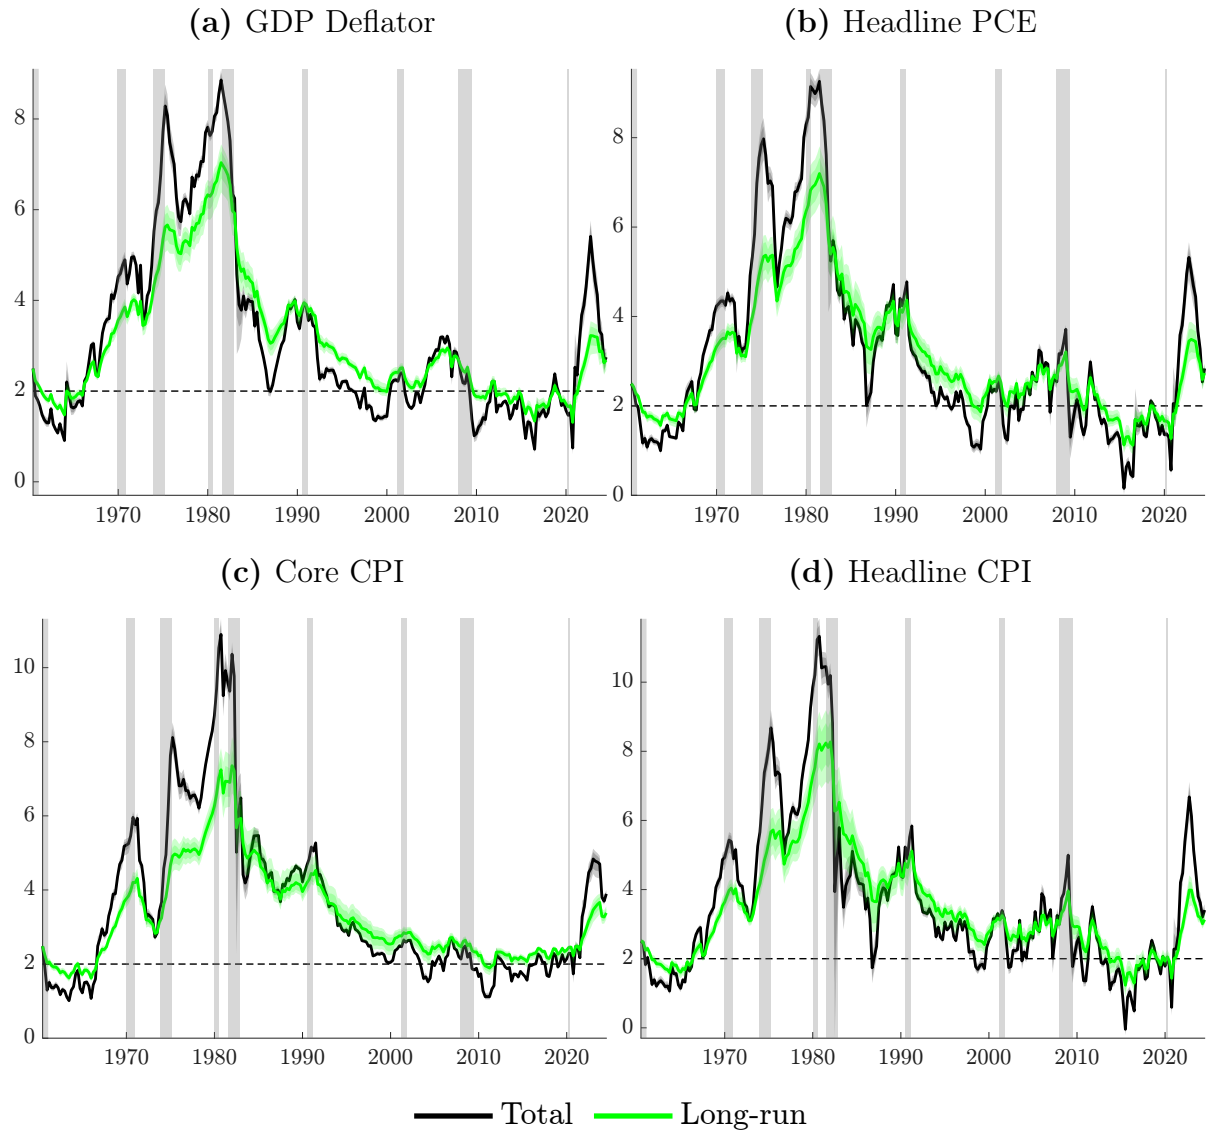



**Figure D7:** Estimated mean for different inflation measures

*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.

## D.8 Additional forecasting results

We conduct a full scale evaluation of the out-of-sample forecasting performance of our model. 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 due to the availability of real-time data vintages. 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. 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).

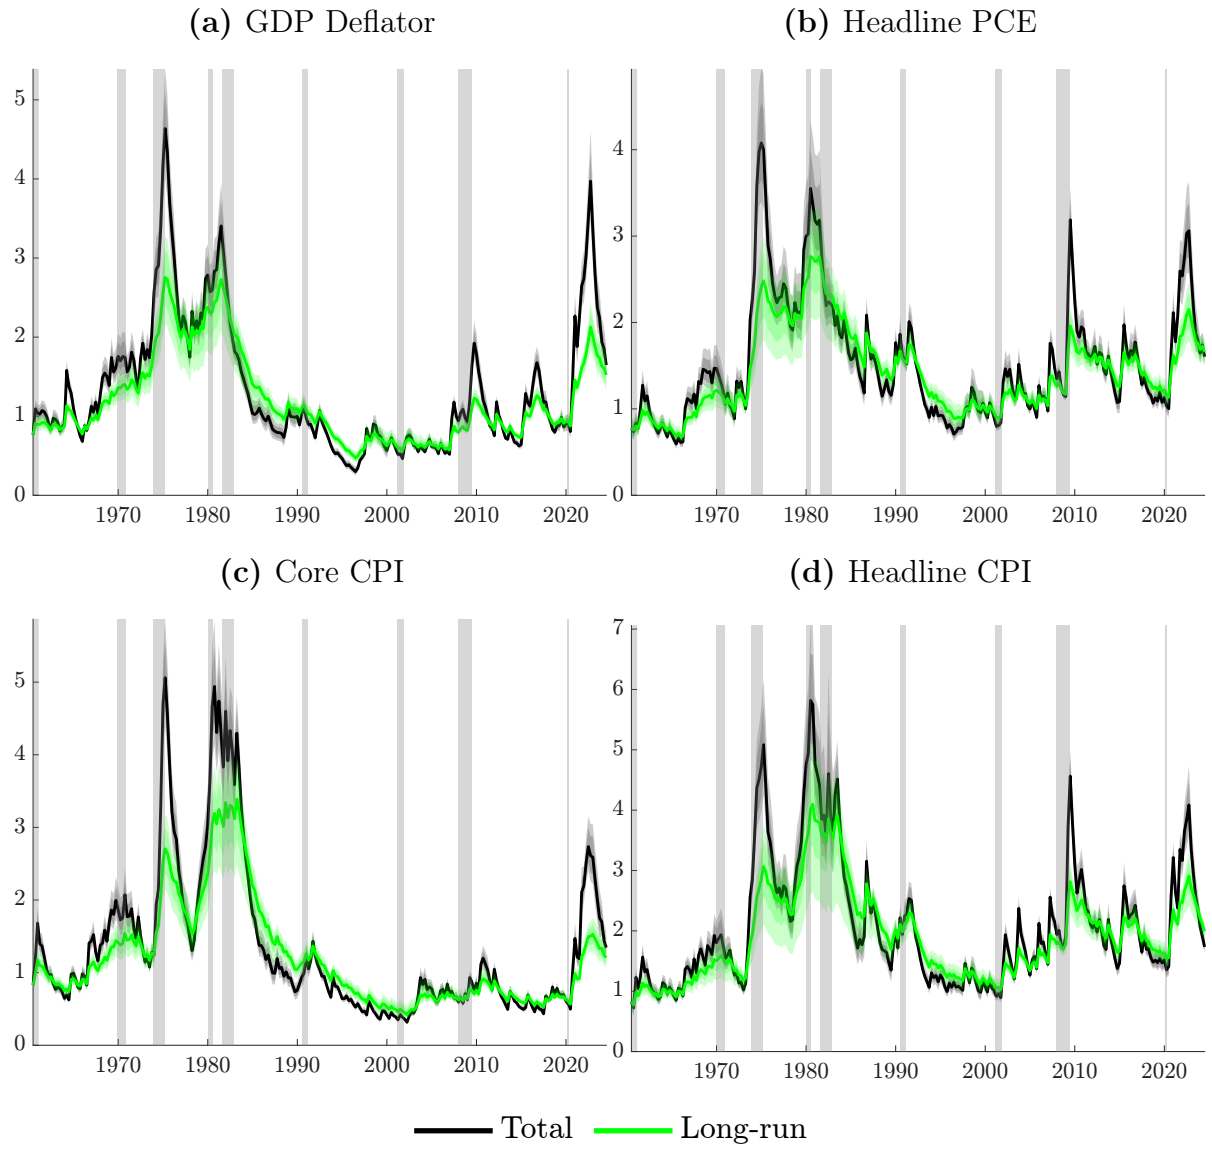



**Figure D8:** Estimated volatility for different inflation measures

*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.

As the UCSV model lacks a mechanism for capturing skewness and overlooks the presence of fat tails in the data, we replicate the exercise using a specification that excludes asymmetry (similarly to Delle Monache and Petrella, 2017). Results in Table D5 confirms the importance of incorporating skewness to improve model fit and inflation forecasting accuracy (see also Mouabbi et al., 2025).

Lastly, in Table D6 we compare our model to the event prediction of the SPF. 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.

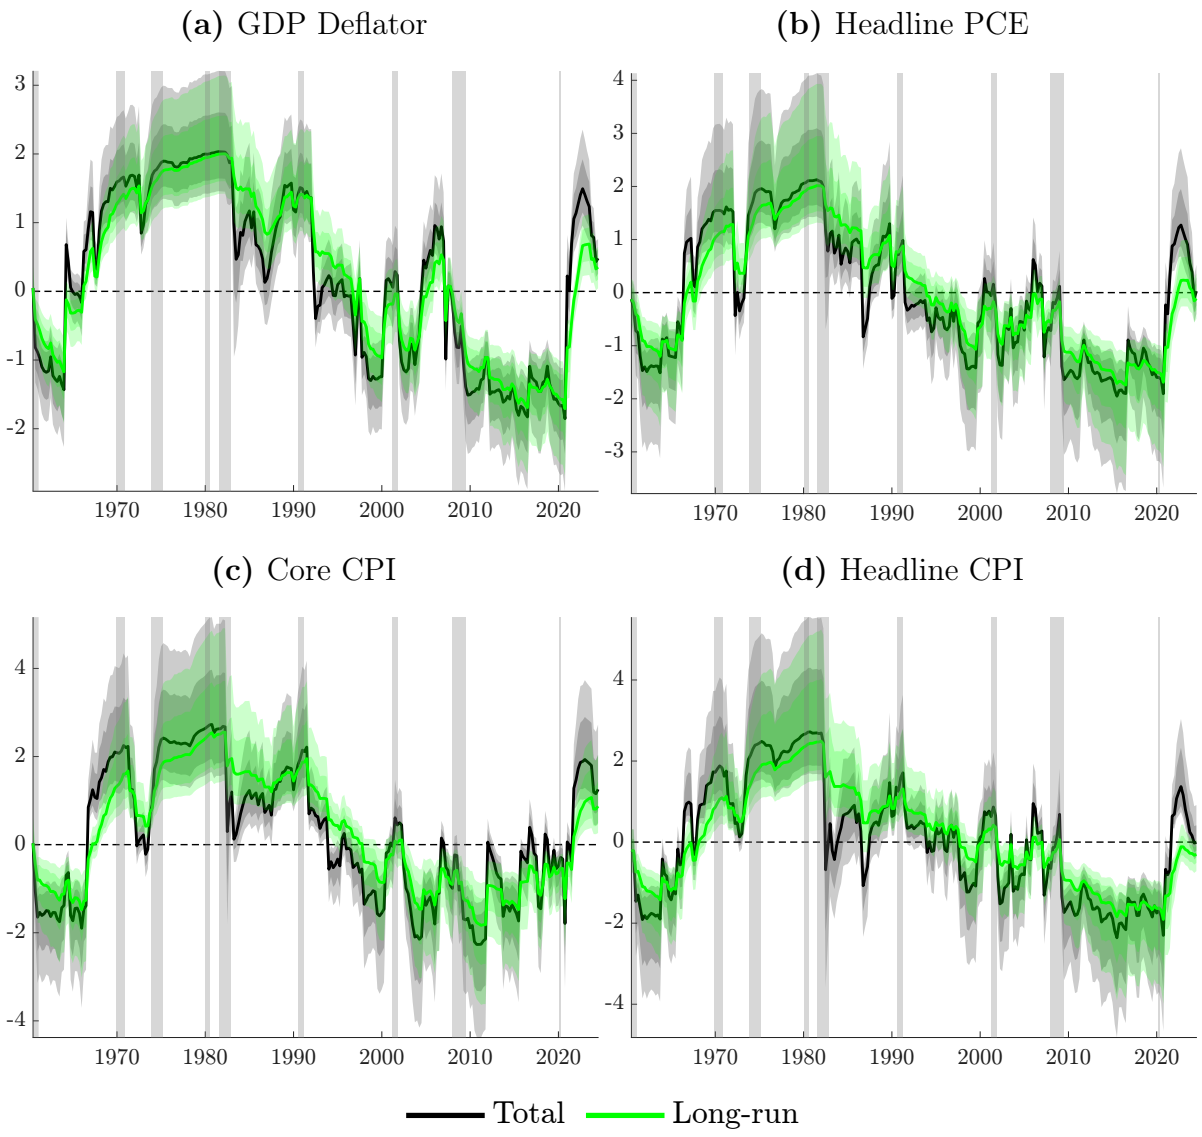



**Figure D9:** Estimated skewness for different inflation measures

*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 D5:** Out-of-sample comparison - *Student t*

| | h = 1 | h = 2 | h = 3 | h = 4 | h = 8 |
|---|---|---|---|---|---|
| MSFE | **0.865** (0.000) | **0.901** (0.000) | **0.926** (0.000) | **0.970** (0.000) | 1.006 (0.939) |
| CRPS | **0.957** (0.001) | **0.970** (0.000) | **0.969** (0.000) | **0.980** (0.006) | **0.995** (0.066) |
| | CRPS decomposition | | | | |
| Right | **0.945** (0.000) | **0.945** (0.000) | **0.952** (0.000) | **0.966** (0.000) | **0.985** (0.002) |
| Left | **0.966** (0.000) | 0.991 (0.142) | **0.983** (0.006) | 0.993 (0.146) | 1.004 (0.913) |
| Center | **0.959** (0.001) | **0.975** (0.008) | **0.972** (0.000) | **0.984** (0.007) | 0.997 (0.147) |

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

**Table D6:** Event forecast comparison against SPF

| | $\pi_t^{Q4} < 1.5\%$ | | | | $1.5\% \leq \pi_t^{Q4} \leq 2.5\%$ | | | | $\pi_t^{Q4} > 2.5\%$ | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | h = 1 | h = 2 | h = 3 | h = 4 | h = 1 | h = 2 | h = 3 | h = 4 | h = 1 | h = 2 | h = 3 | h = 4 |
| **SPF** | 1.216 | 0.929 | 1.217 | 0.890 | 1.118 | 0.928 | 1.376 | 1.120 | 0.910 | 0.235 | 0.530 | 0.847 |
| **UCSV** | 0.998 | 1.063 | 1.023 | 0.995 | 1.053 | 1.033 | 1.031 | 0.984 | 0.917 | 0.738 | 0.831 | 0.838 |

*Note:* The table reports the ratio of the Brier score of SPF and UCSV over our Skew-$t$ model event predictions. The target variable is Q4-over-Q4 core PCE. The evaluation sample runs from 2007Q1 due to SPF data availability.

## D.9 Monte Carlo analysis

We simulate T=250 observations from $Skt_\nu(m_t, \sigma_t, \varrho_t)$, for simulated values of the parameters of location, $m_t$, scale, $\sigma_t$, and asymmetry, $\varrho_t$. 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**

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

**DGP2: constant asymmetry with breaks**

$$\begin{bmatrix} m_t \\ \delta_t \end{bmatrix} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.99 \end{bmatrix} \begin{bmatrix} m_{t-1} \\ \delta_{t-1} \end{bmatrix} + chol\left( \begin{bmatrix} 0.05 & 0 \\ 0 & 0.025 \end{bmatrix} \right) \varepsilon_t,$$
$$\gamma_t = \begin{cases} 0 & t \leq 100 \\ 0.25 & 100 < t \leq 150 \\ -0.25 & t < 150 \end{cases}$$

**DGP3: no asymmetry and location-scale covariance**

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

**DGP4: no asymmetry and location-scale covariance with breaks**

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

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

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

**DGP5: time-varying asymmetry**

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

**DGP5: time-varying asymmetry and correlated updates**

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

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

We report the results of this exercise in Figure D10. 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 D11 report the estimated long- and short-run components; notice that the two distributions add up to that reported in panel (b) of Figure D10. 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.

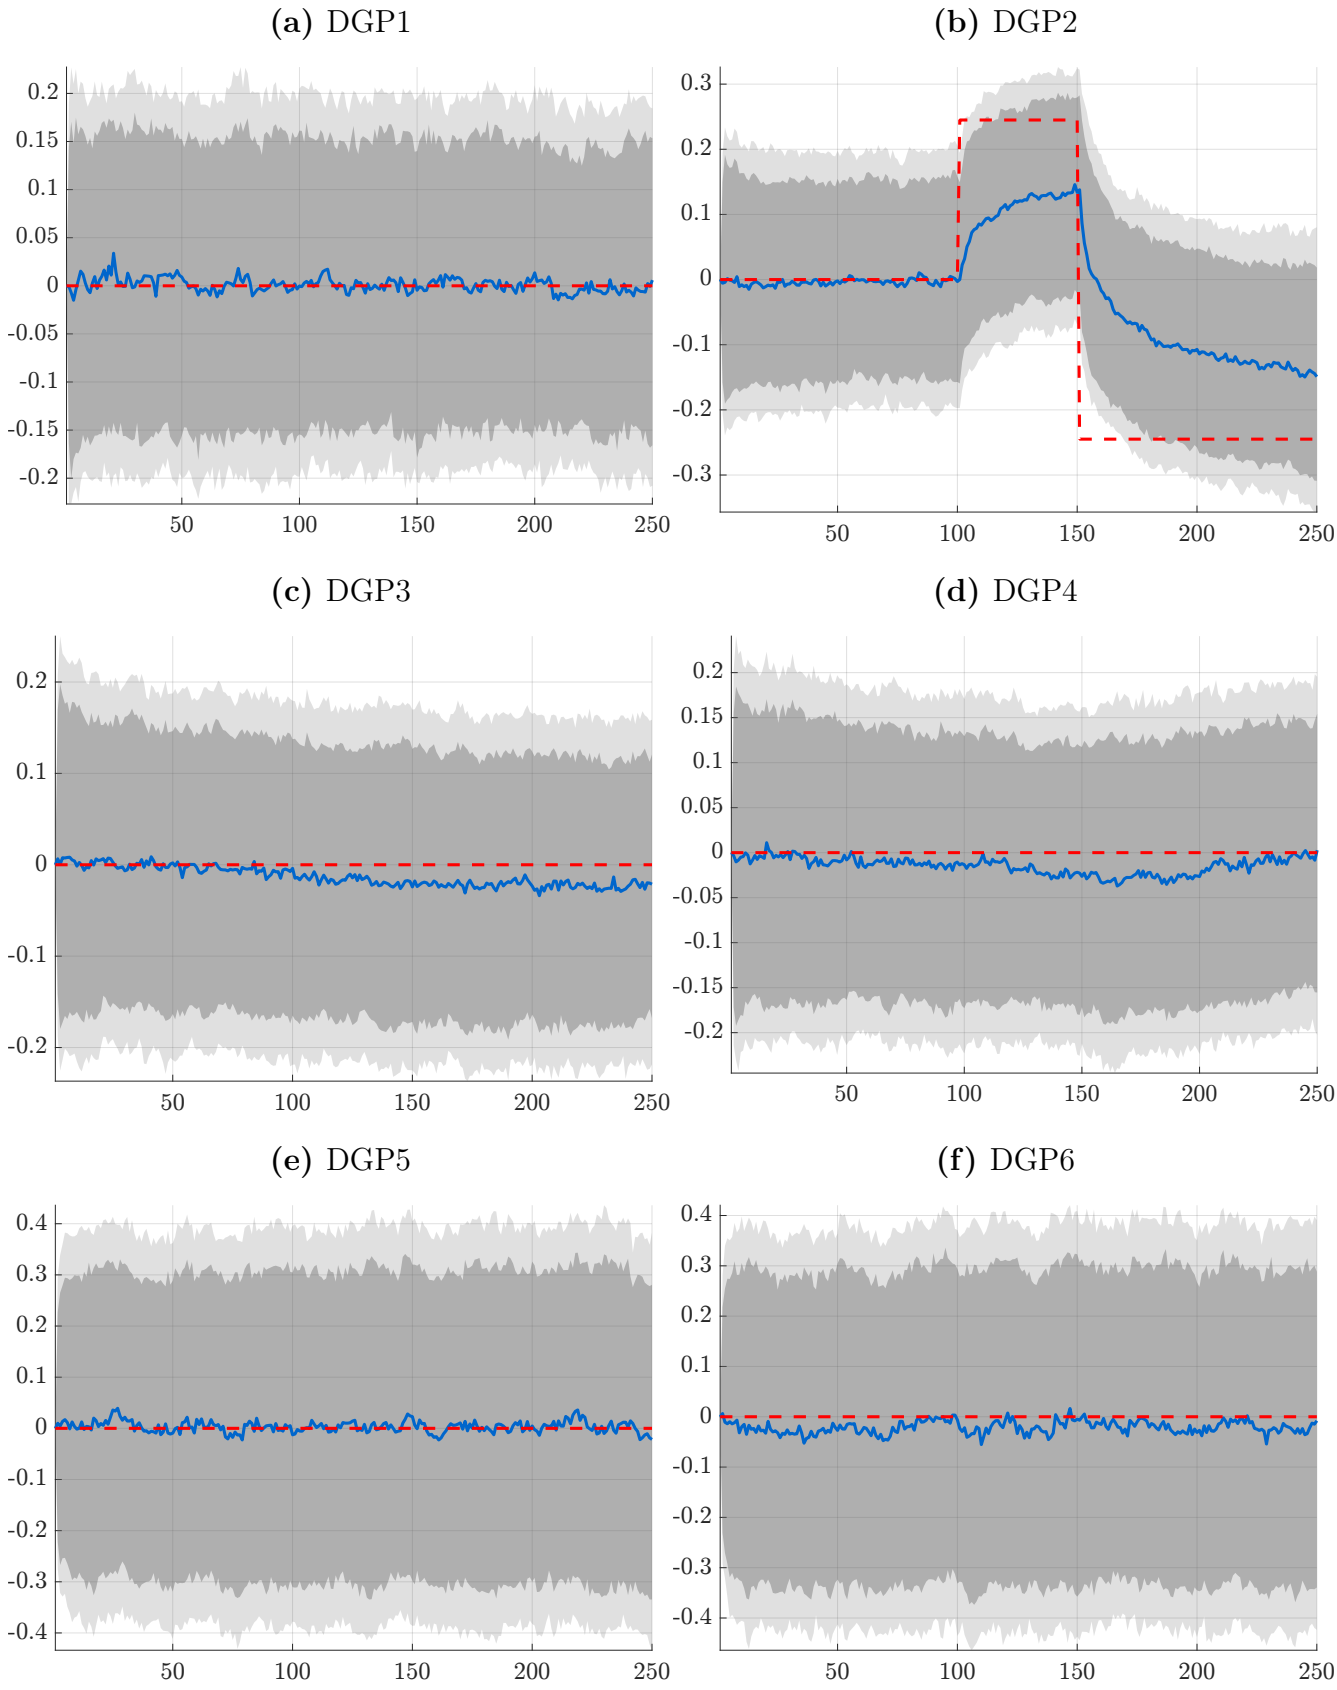



**Figure D10:** Estimated asymmetry

*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.

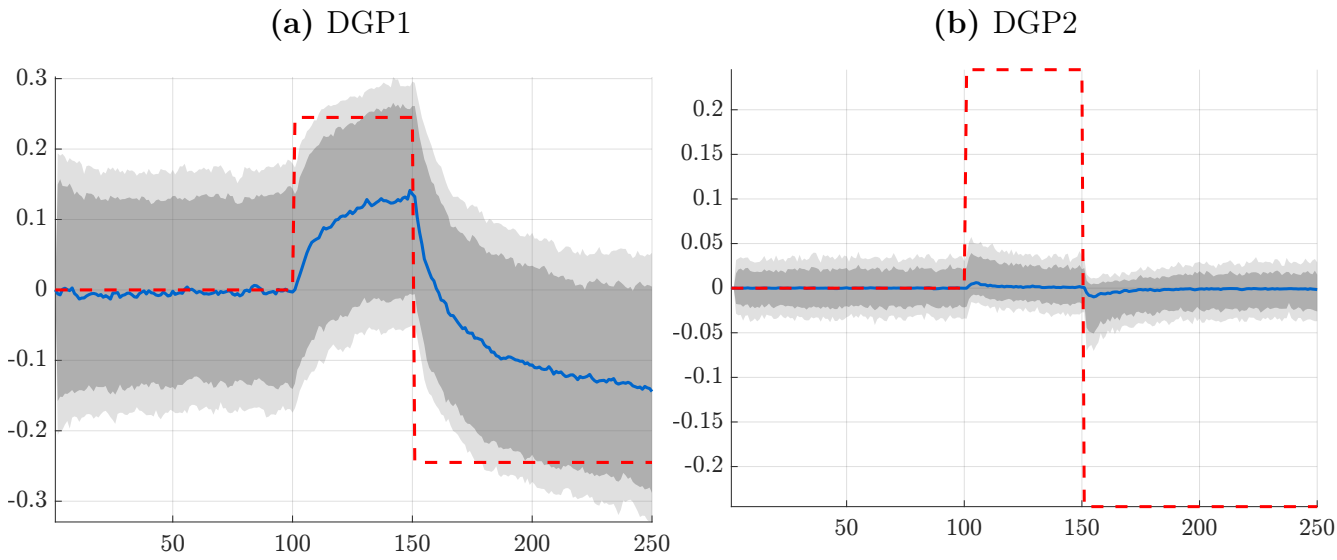



**Figure D11:** Disentangling permanent changes

*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.

# E A New Keynesian model with asymmetric cost-push shock

In this appendix, we provide detailed derivations of the solutions for the optimal monetary problem under asymmetric risks, outlined in Section 5. This problem is derived in the basic New Keynesian model presented in Appendix C and Chapter 3 of Galí (2008).

## E.1 The optimal monetary policy problem

To solve for the optimal monetary policy problem, we now astray from the presence of a Taylor rule (Equation (C11)). We assume that the central bank commits, with full credibility, to a policy plan consistent with a quadratic objective function in inflation deviations, $\hat{\pi}_t$ and the output gap, $\hat{x}_t$. Therefore, optimal monetary policy consists in choosing the state-contingent $\{\hat{\pi}_t, \hat{x}_t\}_{t=0}^{\infty}$ that minimizes

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

subject to the sequence of constraints imposed by the Phillips curve and the IS curve

$$\hat{x}_t = E_t \hat{x}_{t+1} - \frac{1}{\sigma}(\hat{i}_t - E_t \hat{\pi}_{t+1}) \tag{E1}$$

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t, \tag{E2}$$

derived in the previous Section.

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

$$\hat{x}_t = -\frac{\kappa}{\alpha_x} \bar{p}_t, \tag{E3}$$

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 (E3) into the Phillips

curve we obtain

$$\bar{p}_t = a\bar{p}_{t-1} + a\beta E_t\bar{p}_{t+1} + au_t \tag{E4}$$

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

### E.1.1 The symmetric case

The stationary solution to Equation (E4) can be obtained by using the method of undetermined coefficients by conjecturing a solution of the form

$$\bar{p}_t = \eta\bar{p}_{t-1} + \lambda u_t, \tag{E5}$$

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

$$E_t\bar{p}_{t+1} = \eta\bar{p}_t. \tag{E6}$$

Substituting the expectations implied by the stationary solution yields

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

Solving for $\eta$ and $\lambda$, we obtain

$$\eta = \frac{1-\sqrt{1-4\beta a^2}}{2a\beta}, \tag{E7}$$

$$\lambda = \frac{a}{1-a\beta\eta}, \tag{E8}$$

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

$$\hat{x}_t = \eta\hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\lambda u_t, \tag{E9}$$

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

**Implementation.** The optimality condition in Equation (E3) allows us to write the IS curve in Equation (E1) 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 (E6), we obtain the following *optimal monetary rule*:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t. \tag{E10}$$

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).

### E.1.2 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 $\psi^0_t$, and the next period's shock, denoted by $\psi^1_t$. These belief distortions are mapped into expectations as follows:

$$E_t\tilde{u}_t = u_t + \psi^0_t + \psi^1_{t-1} = \tilde{u}_t, \tag{E11}$$

$$E_t\tilde{u}_{t+1} = E_t\psi^0_{t+1} + \psi^1_t = \psi^1_t. \tag{E12}$$

Equation (E11) 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 (E12) reflects the assumption that agents' expectations about future shocks are fully captured by the dummy belief $\psi^1_t$, 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:

$$\bar{p}_t = \eta\bar{p}_{t-1} + \lambda u_t + \zeta\psi^1_t, \tag{E13}$$

and, since $E_t u_{t+1} = \psi^1_t$ and $E_t\psi^1_{t+1} = 0$, it follows that

$$E_t\bar{p}_{t+1} = \eta\bar{p}_t + \lambda\psi^1_t.$$

Substituting the expectations implied by the stationary solution under optimal policy (Equa-

tion (E4)) yields

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

Solving for $\eta$, $\lambda$ and $\zeta$ we obtain

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

Since under optimal policy $\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t$, the equilibrium process for the output gap reads as

$$\hat{x}_t = \eta\hat{x}_{t-1} - \frac{\kappa}{\alpha_x}\left[\lambda u_t + \zeta\psi_t^1\right], \tag{E14}$$

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*:

$$\hat{i}_t = -(1-\eta)\left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\bar{p}_t + \left[1 - \sigma\frac{\kappa}{\alpha_x}\right]\lambda\psi_t^1, \tag{E15}$$

where the first term in the right-hand side is the same as in Equation (E10), 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.[28]

### E.1.3 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. 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.

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

We start from the optimality condition between output gap and inflation

$$\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^m, \tag{E16}$$

where $\bar{p}_t^m$ is the (mispecified) price level the central bank is targeting, under the incorrect perception that of symmetric risk, i.e. under the assumption that $\psi_t^j = 0, \forall t, j$. This *mispecified* price level evolves according to

$$\bar{p}_t^m = \eta \bar{p}_{t-1} + \lambda u_t. \tag{E17}$$

Note that this is not exactly the same price level as in fully symmetric case, because prices $(\bar{p}_t)$ are affected by agents' beliefs of skewed risk. To see that, 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 conjecture that the difference between the actual prices and the prices that the central bank incorrectly targets when solving the optimal problem takes the form

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

By plugging this equation into the optimality condition in Equation (E16), and substituting into the Phillips curve expressed in terms of price level yields

$$\bar{p}_t = a\bar{p}_{t-1} + a\beta E_t \bar{p}_{t+1} + b\tau\psi_t^1 + a u_t \tag{E18}$$

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 (E18) to obtain

$$\bar{p}_t = \underbrace{\frac{a}{1-a\beta\eta}}_{\eta}\bar{p}_{t-1} + \underbrace{\frac{a}{1-a\beta\eta}}_{\lambda} u_t + \underbrace{\frac{a\beta\lambda + b\tau}{1-a\beta\eta}}_{\tilde{\zeta}} \psi_t^1. \tag{E19}$$

Solving for the coefficients we obtain

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

Comparing Equation (E19) with Equation (E13), note that the misspecified policy reaction implies $\tilde{\zeta} > \zeta$. As a result, by failing to respond to shifts in the asymmetry of risk, the central bank amplifies the sensitivity of inflation to those shifts.

We recover $\tau$ by taking the difference between the price level, $\bar{p}$, and the price level under the

symmetric policy – Equation (E17). Specifically,

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

which leads to

$$\tilde{\zeta} = \tau = \frac{a\beta\lambda}{1 - a\beta\eta - b}. \tag{E20}$$

It should be noted that a central bank that overlooks the role of imbalanced risks will end up setting the output gap as a function of $\bar{p}_t^m$. However, this price level is not attainable due to the presence of asymmetric inflation risks. As a result, the central bank achieves a suboptimal output gap and a realized price level that deviates from its target, $\bar{p}_t^m$.

Under the optimal policy, the equilibrium process for the output gap is given by $\hat{x}_t = -\frac{\kappa}{\alpha_x}\bar{p}_t^m$, for all $t$. Therefore, in the first period, $\hat{x}_0 = -\lambda\frac{\kappa}{\alpha_x}u_t$. That is, the central bank's response in the initial period—while ignoring the effects of skewed risk—results in an output gap equal to that under symmetric risk. This occurs because, in the first period, the price level targeted by the central bank, $\bar{p}_t^m$, coincides with the price level targeted by a central bank assuming symmetric risk.[29] However, in subsequent periods, the output gap begins to diverge across the two economies as the prior period's price level differs due to the skewness in the shock distribution. Specifically, with an unwitting central bank:

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

### E.2 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, $\sigma$, to unity, the production function scale parameter, $\alpha$, is set to one third, the elasticity of substitution among intermediate goods, $\varepsilon$, is equal to 6, and the slope of the Phillips curve, $\kappa$, is equal to 0.1275.

## F Inflation skewness shock: additional material

In this appendix we provide additional information about the data, methodologies and results for the exercises in Sections 4 and 6.

### F.1 Additional empirical specifications

In Figure 4, we report the responses of headline PCE, GDP, and the federal funds rate to an inflation skewness shock. In Figure F1, we extend the analysis to core PCE and hours worked in the nonfarm business sector. The responses are closely aligned with the baseline results. Core PCE rises on impact and remains persistently above baseline, indicating that the inflationary effects of an upward revision in inflation risk are not confined to the more volatile components of prices.

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

Likewise, the response of hours worked closely mirrors that of GDP, pointing to a broad-based weakening in real activity rather than to a composition effect specific to output measurement.

Our baseline specification assumes that, once we control for observed proxies for oil, monetary, fiscal, and productivity shocks, the residual component of revisions in inflation skewness can be interpreted as an exogenous shift in the asymmetry of inflation risk. In what follows, we assess the robustness of this interpretation under alternative identification schemes.

A first concern is that, despite the inclusion of these controls, the orthogonalized variation in skewness may still partly reflect endogenous responses to omitted shocks. To address this issue, we re-estimate the local projections using an instrumental-variables design. Specifically, we construct instruments for the skewness shock using the sign of changes in inflation asymmetry around a limited number of major events. We focus on episodes in which the absolute change in skewness exceeds 1.7 and 2.5 standard deviations, respectively. This procedure still allows other shocks to affect inflation skewness around those dates, but relies on the weaker assumption that no other disturbance displays a systematically similar sign pattern across this restricted set of large skewness episodes. The resulting impulse responses, reported in Figure F2, are close to the baseline and preserve its main qualitative implications.

As discussed in Section 3, our estimated skewness measure closely comoves with the Global Fragmentation Index (GFI) of Fernández-Villaverde et al. (2024), which captures the degree of global economic integration. As a second robustness check, we add the common component of the GFI as a contemporaneous control in the local projection specification.[30] Identification then relies on movements in $\Delta\varrho_t$ that are orthogonal to changes in global fragmentation, thereby removing potential contamination from developments in the international trade and financial environment. Figure F3 compares these responses with the baseline estimates in Section 4. While the response of the balance of risks differs somewhat over the first year, the responses of the interest rate, GDP, and PCE remain broadly similar and lie well within the confidence bands of the baseline specification.

Finally, we include the estimated location and scale parameters of the inflation distribution, $\mu_t$ and $\sigma_t$, as contemporaneous controls. This specification ensures that the identifying variation in asymmetry is orthogonal to shifts in the central forecast and in the dispersion of expected inflation. As shown in Figure F4, the estimated responses remain closely aligned with those of the baseline specification in Figure 4. This result suggests that the dynamics attributed to changes in the asymmetry parameter are largely independent of movements in the parameters governing the central location and dispersion of the conditional inflation distribution.

Taken together, these exercises confirm that the baseline findings are robust to alternative identification assumptions. Across specifications, upward revisions in inflation risk continue to behave like adverse supply shocks: they raise prices, depress aggregate demand, and call for a cautious monetary policy response.

[30] Using the disaggregated components of the GFI leaves the results unaltered.

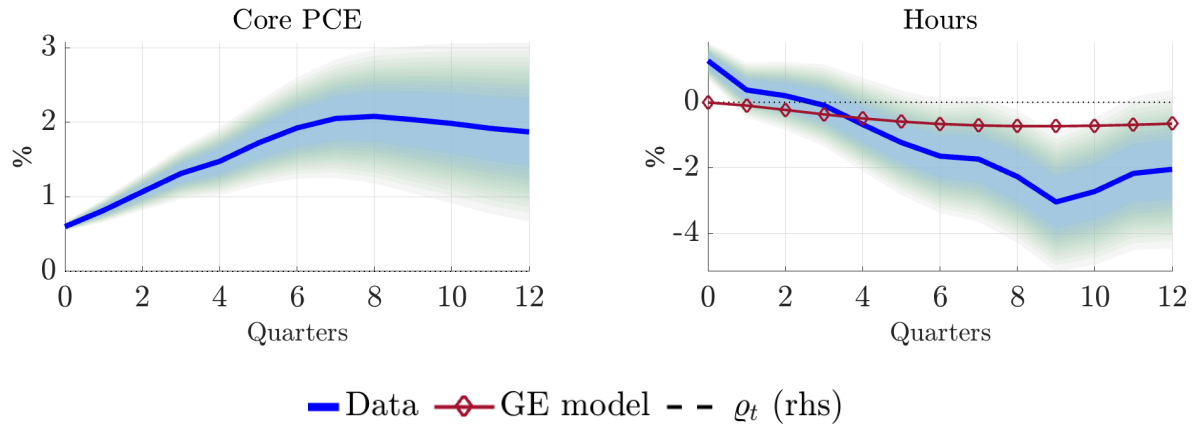



**Figure F1:** Impulse response to an increase in upside risks in inflation

*Note:* The panels report the response of core PCE index and hours worked to a positive increase in the balance of risks to inflation. Blue lines represent the responses estimated via local projections, whereas the dark red lines correspond to the beliefs representation of the Smets and Wouters (2007) model. Shaded areas report 95% confidence intervals.

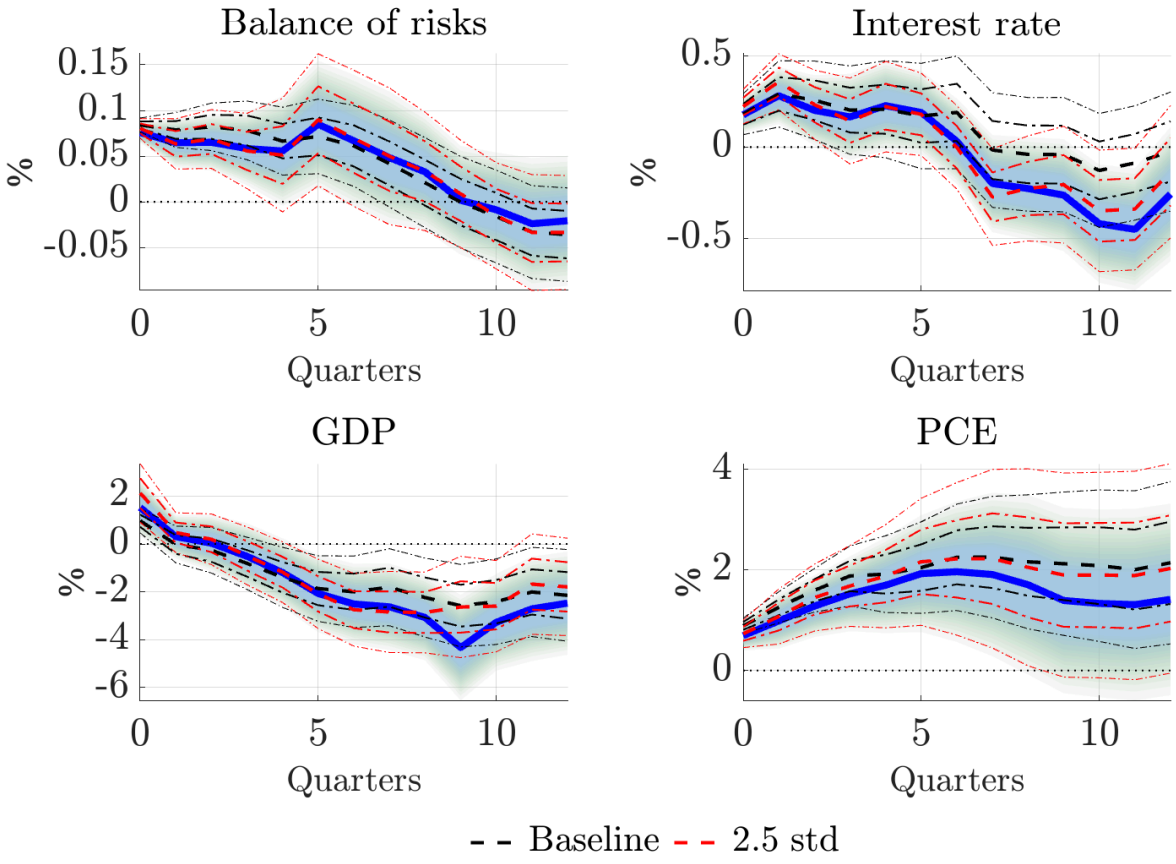



**Figure F2:** Impulse response to an increase in upside risks in inflation - LP-IV

*Note:* The panels report the response of the balance of risks, the interest rate, headline PCE index, and real GDP to a positive increase in the balance of risks to inflation. Blue lines represent the responses with the instrumented skewness shock using the sign of changes in inflation asymmetry with an absolute in exceess of 1.7 standard deviations; red lines indicate responses based on an instrument identified with a threshold level of 2.5 standard deviations. The black lines report the baseline response from Figure 4. Shaded areas and dash-dotted lines report 95% confidence intervals.

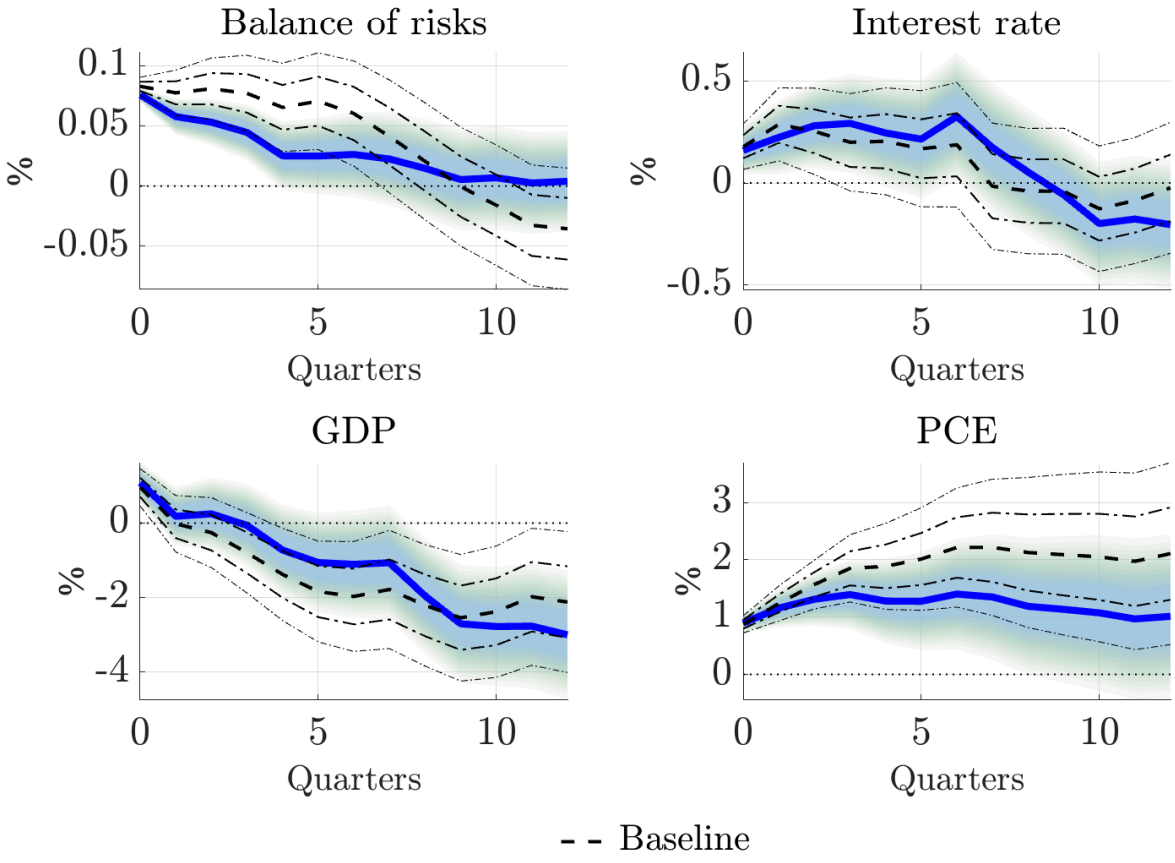



**Figure F3:** Impulse response to an increase in upside risks in inflation - GFI

*Note:* The panels report the response of the balance of risks, the interest rate, headline and core PCE indexes, real GDP and hours worked to a positive increase in the balance of risks to inflation. Blue lines represent the responses including the GFI of Fernández-Villaverde et al. (2024) as contemporaneous control, whereas the black lines report the baseline response from Figure 4. Shaded areas and dash-dotted lines report 95% confidence intervals.

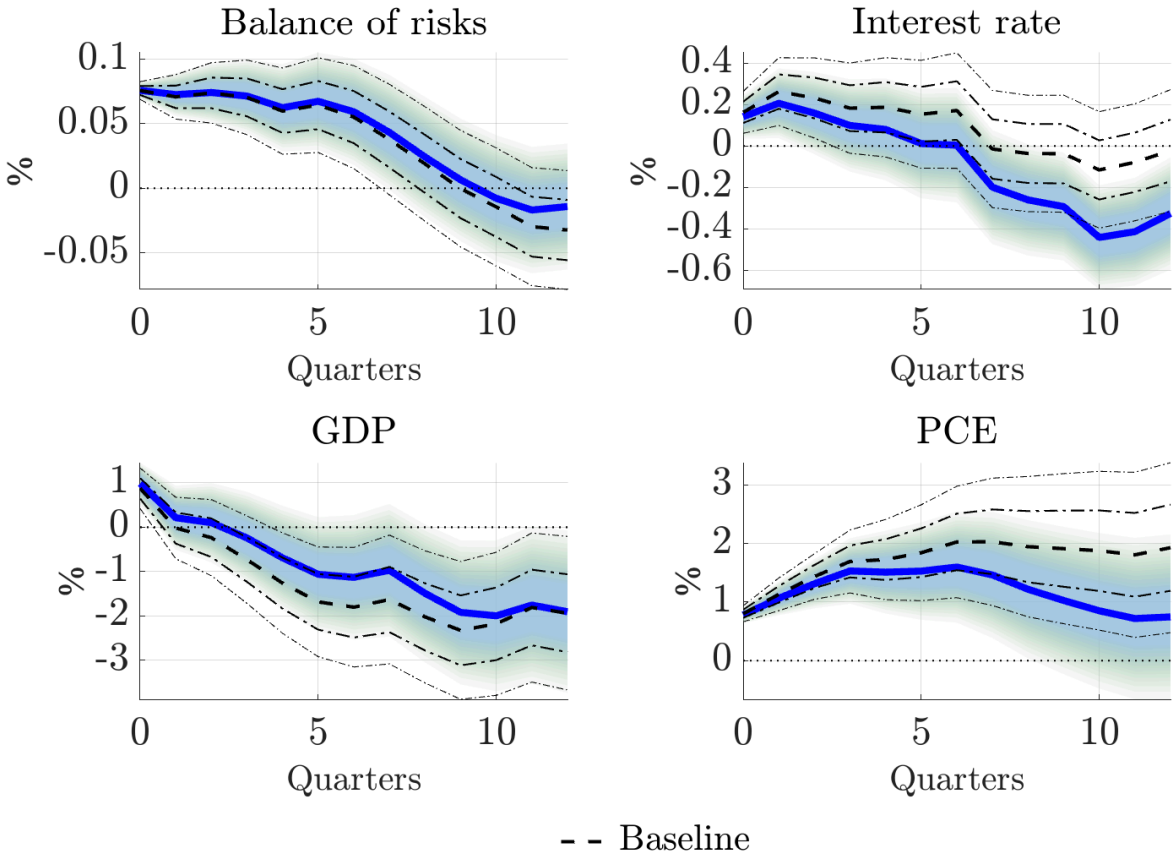



**Figure F4:** Impulse response to an increase in upside risks in inflation - parameters

*Note:* The panels report the response of the balance of risks, the interest rate, headline and core PCE indexes, real GDP and hours worked to a positive increase in the balance of risks to inflation. Blue lines represent the responses including the location and scale parameters as contemporaneous controls, whereas the black lines report the baseline response from Figure 4. Shaded areas and dash-dotted lines report 95% confidence intervals.

## F.2 Calibrating the inflation skewness shock

Under the belief representation, the model's solution follows a standard linear form:

$$\mathbf{s}_t = \mathbf{\Gamma}\mathbf{s}_{t-1} + \mathbf{\Omega}\mathbf{e}_t, \tag{F1}$$

where $\mathbf{s}_t$ includes all state variables and et captures structural shocks and dummy shocks. The solution matrices $\mathbf{\Gamma}$ and $\mathbf{\Omega}$ depend on the model's structural parameters. The model cast in this form can be solved with fast, off-the-shelf techniques.

For simplicity, we assumed that the only source of inflation risk asymmetry stems from cost-push shocks. We calibrate the dummy anticipated cost-push shocks to match the effects of the revisions in the balance of inflation risks on inflation expectations, $\psi_{t+h|t} - \psi_{t+h|t-1}$, where $\psi_{t+h|t} = E_t\pi_{t+h} - \mu_{t+h|t}$, estimated using the Skew-$t$ model of Section 3. See Equation (9). This requires solving the following system of $J+1$ linear equations:

$$\begin{bmatrix} -\sum_{j=1}^{J}\varphi_{t-j}^{j} \\ \psi_{t+1|t} - \psi_{t+1|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = \begin{bmatrix} 1 & \mathbf{0}_{1\times J} \\ \Omega^S & \Omega^N \end{bmatrix} \begin{bmatrix} \varphi_t^0 \\ \varphi_t^1 \\ \vdots \\ \varphi_t^J \end{bmatrix}, \tag{F2}$$

where $\Omega^S$ and $\Omega^N$ capture the effects of surprise and anticipated dummy cost-push shocks on inflation expectations in Equation (F1). These matrices are derived from the solution of the Smets and Wouters model with dummy cost-push shocks—Equation (F1). This solution is straightforward to obtain using widely available techniques. Inverting the matrix on the right-hand side yields the dummy shocks, $\{\varphi_t^j\}_{j=1}^J$ that ensure the structural model replicates the estimated effects of skewness on inflation expectations.

The first equation ensures dummy anticipated shocks do not materialize as actual cost-push shocks.[31] As explained in Proposition 2, these dummy shocks represent belief distortions and can be interpreted as pure beliefs or sentiments. The remaining equations map the revisions to inflation expectations—arising from the inflation asymmetry estimated in real time by the Skew-$t$ model—onto belief shocks. This final set of equations pins down the anticipated dummy shocks required to shift expectations in line with the estimated changes in perceived risk at each horizon.

### F.2.1 Steepening the NKPC in Smets and Wouters (2007)

Recent empirical evidence from the post-COVID inflation surge and US tariff shocks demonstrates a rapid, near-complete pass-through of supply costs to consumers, reflecting a markedly steeper New Keynesian Phillips curve. Consequently, we lower the Kimball curvature parameter to relax the assumption of strong strategic complementarities embedded in Smets and Wouters (2007), aligning the model's real rigidities with an environment where firms readily adjust prices in response to marginal cost shocks. Furthermore, slightly decreasing nominal Calvo rigidities helps achieving a NKPC slope around 0.25.

[31] The surprise shock needed to neutralize past anticipated shocks affects inflation expectations, which one has to take into account. These effects are captured by the vector $\Omega^S\varphi_t^0$.

## F.3 Forward guidance shock and the RAIT

Similarly, the sequence of forward guidance shocks $\{\hat{\pi}^{\star}_{t+j|t}\}_{j=1}^{J}$ in Equation (17):

$$\begin{bmatrix} \psi_{t+1|t} - \psi_{t+1|t-1} \\ \psi_{t+2|t} - \psi_{t+2|t-1} \\ \vdots \\ \psi_{t+J|t} - \psi_{t+J|t-1} \end{bmatrix} = -\Omega^{FG} \begin{bmatrix} \hat{\pi}^{\star}_{t+1|t} \\ \hat{\pi}^{\star}_{t+2|t} \\ \vdots \\ \hat{\pi}^{\star}_{t+J|t} \end{bmatrix}, \tag{F3}$$

where the $J \times J$ matrix $\Omega^{FG}$ captures the effect of forward guidance shocks on inflation expectations. This matrix, $\Omega_{FG}$, is obtained by solving the beliefs representation of the Smets and Wouters (2007) model, augmented with the RAIT-based monetary policy rule Equation (16).

The left-hand side of the system in Equation (F3) shows the impact of estimated risk revisions on inflation expectations, while the right-hand side represents the forward guidance shocks needed to neutralize that impact. These shocks can be easily obtained by inverting the matrix $\Omega_{FG}$. Solving this linear system ensures that forward guidance neutralizes the effects of the revised balance of inflation risks on expectations, thereby anchoring them as intended under the RAIT.

# References


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.

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

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.

Cao, D., W. Luo, and G. Nie (2023): "Global DSGE models," *Review of Economic Dynamics*, 51, 199–225.

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

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.

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

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

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

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

Dharmadhikari, S. W. and K. Joag-Dev (1988): *Unimodality, Convexity, and Applications*, Boston: Academic Press.

Doan, T., R. Litterman, and C. Sims (1984): "Forecasting and conditional projection using realistic prior distributions," *Econometric Reviews*, 3, 1–100.

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

Gelman, A. (1996): "Inference and Monitoring Convergence," in *Markov Chain Monte Carlo in Practice*, ed. by W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, London: Chapman and Hall, 131–143.

Gelman, A. and D. B. Rubin (1992): "Inference from iterative simulation using multiple sequences," *Statistical science*, 7, 457–472.

Geweke, J. (1992): "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments," in *Bayesian Statistics 4: Proceedings of the Fourth Valencia International Meeting, Dedicated to the memory of Morris H. DeGroot, 1931–1989*, Oxford University Press.

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.

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

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

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

Ray, S. and B. G. Lindsay (2005): "The topography of multivariate normal mixtures," *The Annals of Statistics*, 33, 2042–2065.

Rotemberg, J. J. (1982): "Sticky prices in the United States," *Journal of political economy*, 90, 1187–1211.

Sherlock, C., A. Golightly, and D. A. Henderson (2017): "Adaptive, delayed-acceptance MCMC for targets with expensive likelihoods," *Journal of Computational and Graphical Statistics*, 26, 434–444.

Silverman, B. W. (1986): *Density Estimation for Statistics and Data Analysis*, London: Chapman and Hall.

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

Sims, C. A. and T. Zha (1998): "Bayesian methods for dynamic multivariate models," *International Economic Review*, 949–968.

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