Reading Inflation Tails*

Sophie Guilloux-Nefussi[a], Magali Marx[a], Sarah Mouabbi[a], Jean-Paul Renne[b]

[a]*Banque de France, Paris, 75001, France*
[b]*University of Lausanne, CH-1015 Lausanne, Switzerland*

Abstract

We develop a framework to infer joint inflation-growth risks by combining survey density forecasts with inflation derivative prices. Survey histograms discipline marginal physical distributions of inflation and GDP growth, while inflation swaps, caps, and floors discipline the marginal nominal risk-neutral distribution of inflation. These marginal distributions leave the dependence between inflation and real activity only partially identified. Identification comes from the stochastic-discount-factor link between physical and risk-neutral distributions: since the market value of an inflation payoff depends on the real-activity state in which the inflation outcome occurs, differences between physical and risk-neutral inflation distributions are informative about joint outcomes for inflation and real activity. We implement this idea in a tractable Epstein-Zin model with a flexible joint distribution of inflation and real growth. Applied to the euro area, the framework uncovers substantial time variation in inflation-growth dependence and shows how this dependence shapes inflation risk premia and the pricing of inflation tails.

Keywords: Inflation risk, Inflation options, Survey expectations, Risk premia, Tail risk, Epstein-Zin preferences
JEL: E31, E44, G12, G13

*Email addresses: Sophie Guilloux-Nefussi (sophie.guilloux-nefussi@banque-france.fr); Magali Marx (magali.marx@banque-france.fr); Sarah Mouabbi (sarah.mouabbi@banque-france.fr); Jean-Paul Renne (jean-paul.renne@unil.ch; corresponding author). We thank participants at the NBB Conference on Structural Macro Modelling "The present, the past, and the future of DSGE models", the 13th Asset Pricing Workshop (BdF-BoE-University of York), and at the FGV EPGE seminar for useful comments. We also thank Michael Bauer, Mariano Massimiliano Croce, Burçin Kısacıkoğlu for useful comments. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem.

1. Introduction

Macroeconomic tail risks depend not only on the probability of extreme outcomes, but also on the states in which those outcomes occur. A high-inflation tail may reflect strong demand, with inflation accompanied by robust real activity, or an adverse supply shock, with inflation coupled with weak activity. These scenarios have different implications for policy and asset-pricing. The market value of an inflation payoff depends not only on the likelihood of the inflation outcome, but also on whether it materializes when marginal utility is high. Measuring inflation tail risk therefore requires information about the joint distribution of inflation and real activity, rather than the marginal distribution of inflation alone.

Forward-looking information about these risks comes from several sources, two of which are central to our analysis. Survey density forecasts describe beliefs under the physical measure. In the euro area, the ECB Survey of Professional Forecasters (SPF) is a central reference for policy analysis and is particularly useful because it reports probabilistic histogram forecasts for both inflation and GDP growth. Inflation derivative prices, by contrast, reveal Arrow-Debreu prices for future inflation outcomes–that is, how markets value payoffs contingent on different inflation states–and therefore discipline the nominal risk-neutral distribution of inflation. The wedge between the physical and risk-neutral measures is economically informative rather than incidental: it reflects which inflation outcomes provide insurance value and which correspond to costly states of the world. The empirical challenge is to exploit this wedge without treating it as a statistical adjustment, because the change of measure is governed by the stochastic discount factor linking beliefs, real activity, and asset prices.

Despite the richness of these data, they have not been used jointly to study inflation and GDP-growth risks. Inflation options are typically used to recover marginal risk-neutral distributions of inflation, while survey histograms are used to estimate marginal physical distributions. However, identical inflation and growth histograms can correspond to very different macroeconomic environments, depending on the dependence structure between the two variables. To our knowledge, inflation option-implied probabilities and SPF histogram

forecasts have not been jointly exploited to infer the joint distribution of inflation and GDP growth.[1]

Existing approaches typically exploit only one side of this information set. Some studies use inflation options to recover risk-neutral distributions and, in some cases, infer physical probabilities through a change of measure. Others rely on survey histograms to measure subjective distributions under the physical probability measure. Although both approaches are informative, neither on its own identifies how inflation tail risks are linked to real activity or priced across states of the world. This paper addresses this question.

This paper brings these two perspectives together. We estimate a tractable structural model using three empirical inputs: inflation derivatives, SPF inflation histograms, and SPF GDP-growth histograms. The survey histograms discipline the physical marginal distributions of inflation and real activity, while inflation swaps, caps, and floors identify the nominal risk-neutral distribution of inflation. Preferences pin down the stochastic discount factor linking the two measures, so that the wedge between them becomes informative about the joint distribution of inflation and growth. The model delivers closed-form expressions for both survey-bin probabilities and inflation-derivative prices, enabling all three data sources to be matched jointly at each point in time.

The key mechanism is straightforward. When low inflation occurs in weak-growth states, it is associated with high marginal utility and therefore receives greater weight under the risk-neutral measure. As a result, the risk-neutral distribution of inflation shifts toward lower outcomes relative to the physical distribution. Conversely, when high inflation occurs in weak-growth states, as in an inflationary recession, the risk-neutral distribution shifts toward higher inflation outcomes. The sign and magnitude of the inflation risk premium therefore depend on the joint distribution of inflation and real activity, rather than on the marginal distribution of inflation alone. Our framework is designed to allow this relationship to vary over time, rather than imposing a fixed disaster channel or a constant adjustment between risk-neutral and physical probabilities.

[1]Comparisons between inflation option-implied probabilities and survey-based inflation probabilities remain relatively rare; see, for example, the ECB boxes on inflation options (European Central Bank, 2013; Boninghausen et al., 2018) and Gimeno and Ibáñez (2018). These contributions focus on inflation distributions in isolation.

Our empirical analysis focuses on the euro area and combines data from euro-area inflation derivatives with the ECB SPF over the period from October 2009 to December 2025. The model fits both market prices and survey histograms for inflation and GDP growth, and uncovers substantial variation in the wedge between physical and risk-neutral distributions. In periods when weak activity is associated with low inflation, risk-neutral inflation distributions place greater weight on disinflationary outcomes. By contrast, during the inflation surge of 2021–2023, when high inflation becomes increasingly associated with adverse real outcomes, risk-neutral tails put more weight on high-inflation states. The results therefore show that the pricing of inflation tail risk varies systematically with the real-activity states in which those inflation outcomes are expected to occur.

The paper makes two main contributions. First, it develops a framework that uses the stochastic-discount-factor link between survey beliefs and derivative prices to infer joint inflation-growth risks. The framework is analytically tractable: it delivers closed-form expressions for both forecast-bin probabilities and inflation derivative prices, and it makes explicit which features of the joint distribution are disciplined by each data source. Second, the paper provides new evidence on the time variation and asymmetry of inflation and real-activity risk premia in the euro area. By integrating asset prices and survey information, the framework shows how macroeconomic tail risks are shaped by the interaction of beliefs, market prices of state-contingent payoffs, and the joint inflation-output structure.

The closest strand of the literature extracts inflation risk-neutral distributions from option prices. Kitsul and Wright (2013) use U.S. inflation caps and floors to recover option-implied inflation densities, while Gimeno and Ibáñez (2018) estimate euro-area inflation risk-neutral densities from inflation swaps, caps, and floors. Related measures are also employed in central-bank analyses of option-implied inflation probabilities (European Central Bank, 2013; Böninghausen et al., 2018). More recently, Hilscher et al. (2025) develop a structural framework that infers the probability of inflation disasters from inflation options and maps risk-neutral probabilities into physical probabilities through preferences. Their paper is an important benchmark because it places inflation options at the center of tail-risk measurement. Our contribution differs in an important respect. Rather than relying solely on option prices to infer physical inflation tails, we combine option-implied information with survey distributions for both inflation and GDP growth in order to discipline the joint inflation-

activity structure. This distinction matters because, as emphasized by Renne et al. (2026), the change of measure generally depends on states, horizons, and joint macroeconomic dynamics, rather than on a single scalar risk adjustment.

The paper also relates to work that extracts physical distributions directly from survey data, including studies of survey density forecasts and subjective uncertainty (Tay and Wallis, 2000; Engelberg et al., 2009; Boero et al., 2014), recent work that constructs, evaluates, or combines SPF predictive densities (Galvao et al., 2021; Clark et al., 2022, 2024), and analyses of euro-area SPF inflation risks (Andrade et al., 2014; Grishchenko et al., 2019; Mouabbi et al., 2025; Bassetti et al., 2024). More broadly, the paper complements the disaster-risk and tail-risk pricing literature, which shows how small probabilities of severe macroeconomic events can generate large asset-pricing effects (Barro, 2006; Gabaix, 2012; Wachter, 2013) and uses option prices to infer market compensation for tail events (Backus et al., 2011; Kelly and Jiang, 2014; Bollerslev et al., 2015; Barro and Liao, 2021). Relative to these strands of the literature, our focus is on how inflation tails interact with real activity, and on how inflation tail risks interact with real activity, and on how this interaction shapes the wedge between physical and risk-neutral macroeconomic distributions.

The remainder of the paper is organized as follows. Section 2 develops the identification intuition, and Section 3 introduces the model. Section 4 describes the data, while Section 5 presents and discusses the empirical results. Section 6 concludes. Technical details and additional results are deferred to the appendices.

2. Identification intuition

This section explains why the three data sources used in the paper are jointly informative. The objective is to identify the joint distribution of next-period inflation (π_1) and GDP growth (Δc_1). Survey histograms for inflation and GDP growth provide direct information about the corresponding physical marginal distributions, while inflation-derivative prices reveal how markets value payoffs contingent on inflation outcomes. Taken separately, neither source identifies an unrestricted joint distribution of inflation and real activity. However, their combination, disciplined by a stochastic discount factor, is informative about the dependence structure that governs risk pricing and helps distinguish demand-driven from supply-driven tail risks.

2.1. Why marginals are not enough

Suppose that each variable can take only two values, labeled low (L) and high (H). That is,

$$\pi_1 \in \{L_\pi, H_\pi\}, \qquad \Delta c_1 \in \{L_c, H_c\}.$$

The joint distribution of $(\pi_1, \Delta c_1)$ is then summarized by four probabilities:

$$p_{LL} = \mathbb{P}_0(\pi_1 = L_\pi, \Delta c_1 = L_c), \qquad p_{LH} = \mathbb{P}_0(\pi_1 = L_\pi, \Delta c_1 = H_c),$$

$$p_{HL} = \mathbb{P}_0(\pi_1 = H_\pi, \Delta c_1 = L_c), \qquad p_{HH} = \mathbb{P}_0(\pi_1 = H_\pi, \Delta c_1 = H_c),$$

with

$$p_{LL} + p_{LH} + p_{HL} + p_{HH} = 1.$$

Hence, the unrestricted joint distribution has three free parameters.

Assume first that we only observe the marginal distributions of inflation and GDP growth, as would be the case if we relied solely on survey density forecasts (e.g., using the SPF). These marginals imply

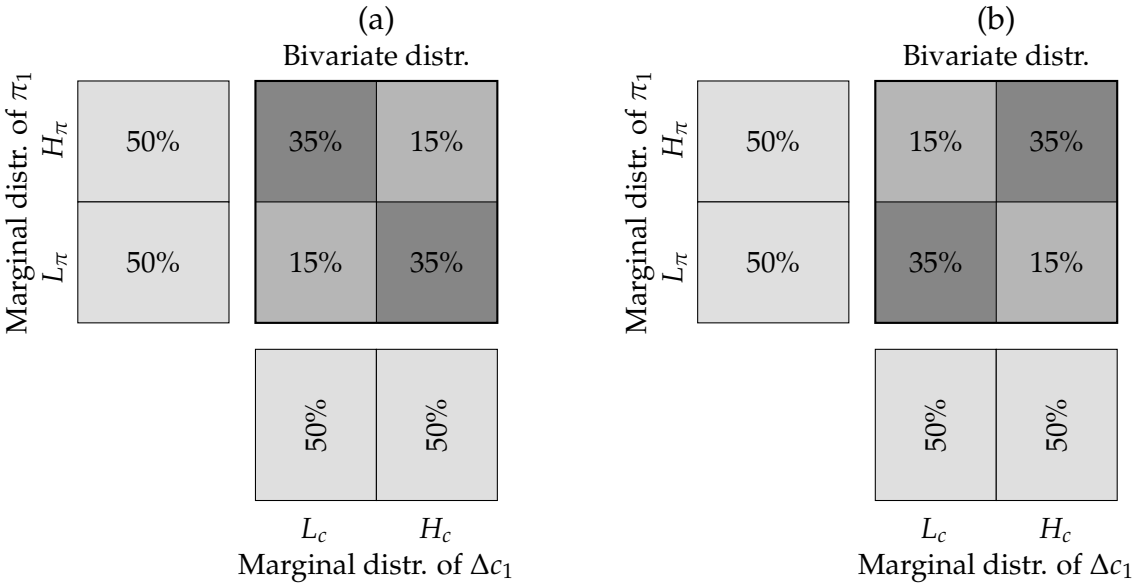
$$\mathbb{P}_0(\pi_1 = L_\pi) = p_{LL} + p_{LH}, \qquad \mathbb{P}_0(\Delta c_1 = L_c) = p_{LL} + p_{HL}.$$

Since the probabilities sum to one, these two marginal restrictions are not sufficient to identify the dependence structure between inflation and GDP growth. In particular, distinct joint distributions may share exactly the same marginals. Figure 1 illustrates this point: Panels (a) and (b) display two different joint distributions with identical marginals. In Panel (a), inflation and GDP growth are negatively associated, whereas in Panel (b) they are positively associated. If one considered inflation or growth in isolation, the two economies would appear identical. From a policy perspective, however, they differ markedly, because the dependence structure determines whether inflation tail risks coincide with weak or strong economic activity.

2.2. How asset prices help

Asset prices provide additional information because they weight states of the world through the stochastic discount factor. To see this, suppose that the stochastic discount factor between

Figure 1: **Discretized bivariate distributions of inflation π_1 and GDP growth Δc_1 with identical marginal distributions and different dependence structures.** In both panels, the marginal distributions of inflation and GDP growth assign a probability of 50% to each state. Panel (a) depicts a situation with negative correlation between π and Δc, while Panel (b) displays a positive correlation. Darker cells correspond to larger joint probabilities.



dates 0 and 1 is given by $\mathcal{M}_1$, and depends on consumption growth Δc_1.[2] Consider now a digital inflation claim that pays one unit at date 1 if inflation turns out to be low, that is, if $\pi_1 = L_\pi$. Its date-0 price is

$$\mathbb{E}_0\Big(\mathcal{M}_1 \mathbb{1}_{\{\pi_1 = L_\pi\}}\Big) = p_{LL}\, \mathcal{M}(L_c) + p_{LH}\, \mathcal{M}(H_c),$$

where $\mathcal{M}(L_c)$ and $\mathcal{M}(H_c)$ denote the realizations of the stochastic discount factor in the low- and high-growth states, respectively. If the stochastic discount factor were the same in the low- and high-growth states, this price would be redundant with the marginal physical probability of low inflation. It is not redundant when $\mathcal{M}(L_c) \neq \mathcal{M}(H_c)$. The pricing equation then delivers an additional restriction that depends on the joint distribution of inflation and GDP growth, not only on the marginal distribution of inflation.

The key point is that the price of this claim depends on how the low-inflation outcome is distributed across low- and high-growth states. If low inflation tends to occur in bad states, that is, when $\Delta c_1 = L_c$ and the stochastic discount factor is high, then the claim is more valu-

[2] A specific functional form for $\mathcal{M}_1$ will be introduced in the model section.

able than if the same low-inflation outcome mostly occurs in good states. Hence, comparing inflation tails under the physical and risk-neutral measures reveals information about how inflation outcomes are linked to real activity. This is the central idea of the paper: use the wedge between $\mathbb{P}$ and $\mathbb{Q}$, together with preferences, to infer the joint distribution of inflation and GDP growth.

This simple example is intended only to build intuition. In principle, one could attempt to recover an unrestricted discrete joint distribution from survey marginals and a sufficiently rich cross-section of prices for state-contingent payoffs. That is not the approach pursued in this paper. Instead, our empirical exercise is model-restricted: we impose a parsimonious parametric specification for the joint distribution of inflation and consumption growth and use preferences to discipline the change of measure. The specification is sufficiently flexible to capture rich dependence patterns, while remaining tractable and adequately disciplined by the joint information embedded in survey histograms and inflation-derivative prices.

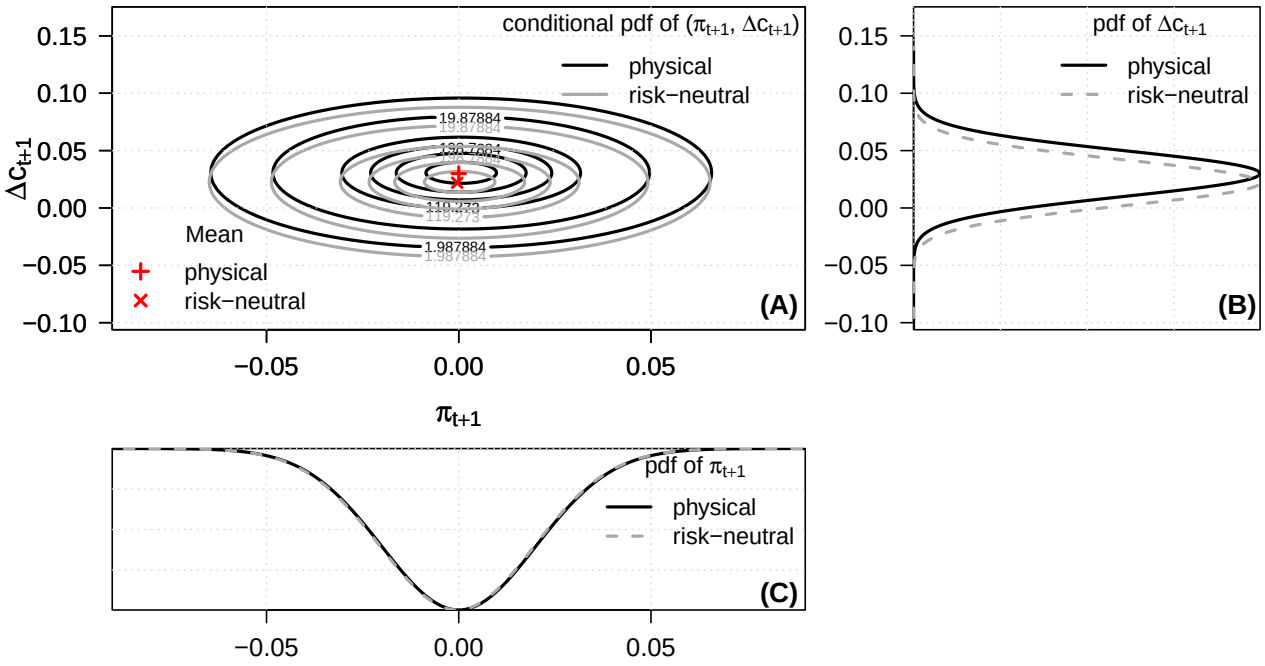
2.3. A tractable illustration

Before turning to the baseline Epstein-Zin model, we illustrate the same mechanism in a static CRRA setting with a two-component Gaussian mixture. This example is not the empirical model employed in the paper; its purpose is purely pedagogical. Because the risk-neutral distribution is available in closed form, the setup provides a transparent illustration of how the pricing kernel reweights different regions of the joint inflation-output distribution. The assumptions and derivations underlying the illustration are presented in Online Appendix C.

The example starts from a benchmark in which inflation and consumption growth are Gaussian and independent. In that case, risk-neutral reweighting of inflation does not contain a real-activity component: inflation outcomes are not associated with high- or low-consumption states. By contrast, the risk-neutral distribution of consumption growth shifts toward lower outcomes, because low-consumption states have high marginal utility and therefore receive greater weight under the pricing kernel.

Figure 3 then varies the dependence structure. The second row introduces positive comovement through a shared Gaussian shock. High-inflation states then tend to coincide with high consumption and low marginal utility, so this covariance component pushes the risk-neutral inflation distribution toward lower outcomes. The third row introduces a recession

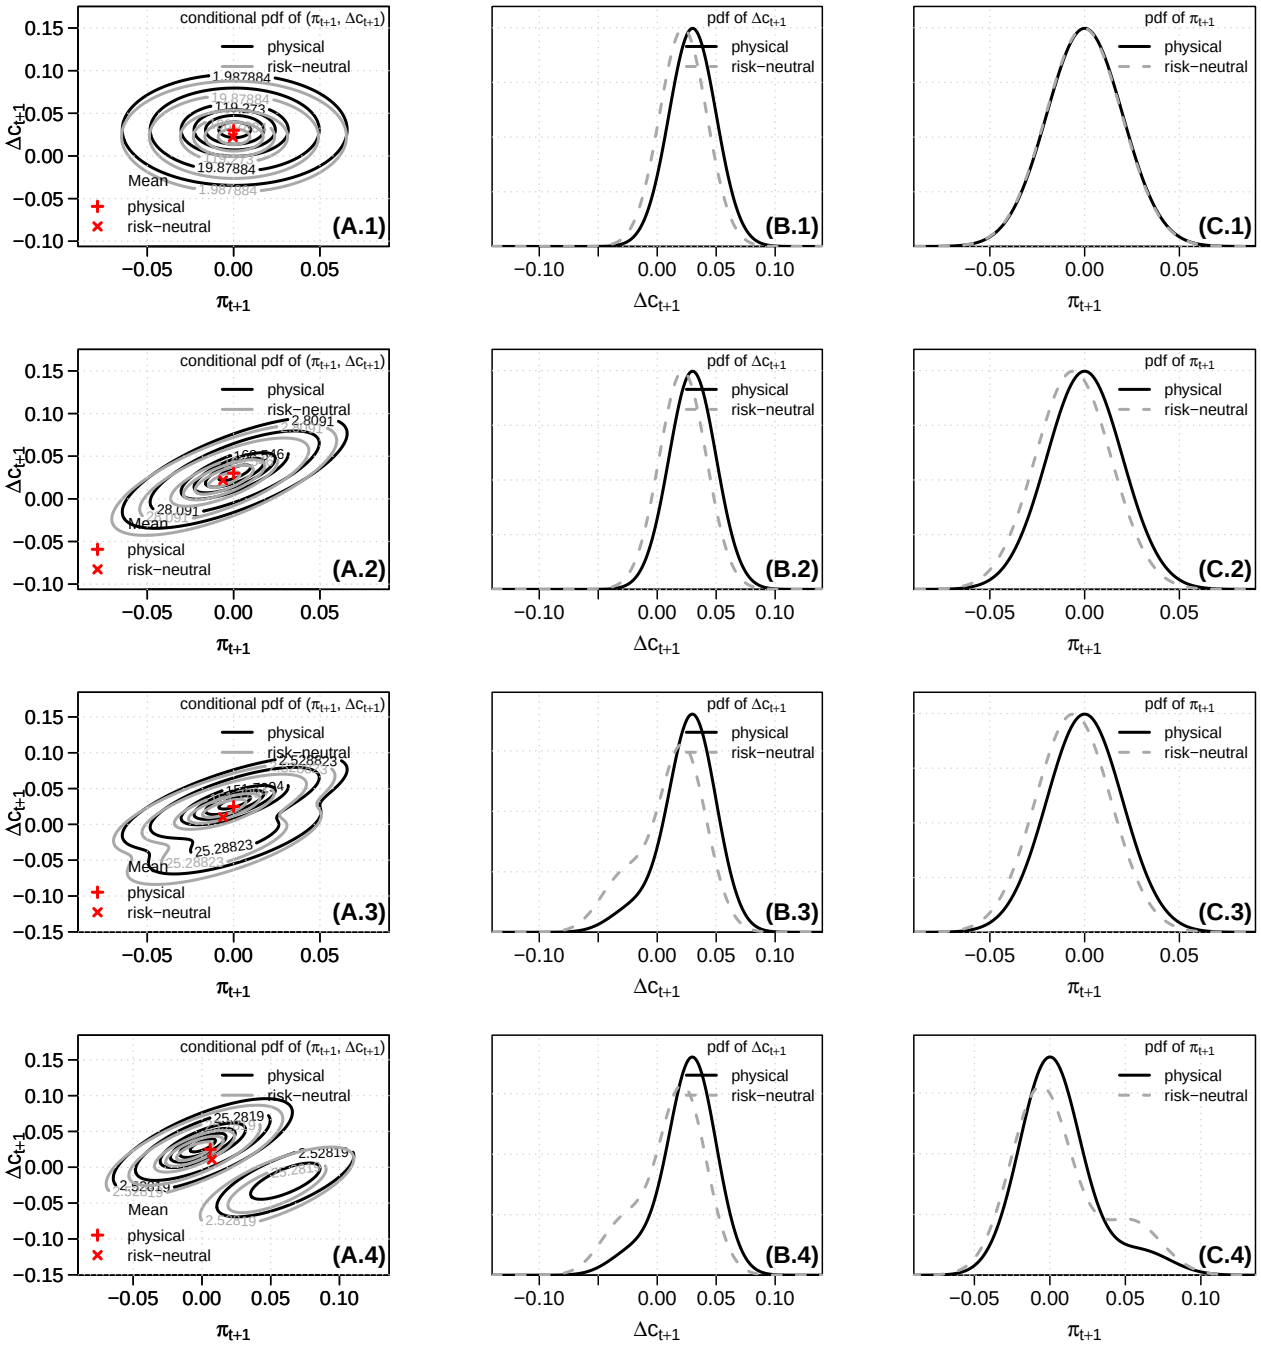
Figure 2: **Densities of Δc_1 and π_1 in the benchmark illustration.** Panel (A) shows the joint physical and risk-neutral densities of consumption growth Δc_1 and inflation π_1, together with their means; Panels (B) and (C) report their corresponding marginal densities. This benchmark Gaussian specification features independent inflation and consumption shocks, so there is no covariance-driven reweighting of inflation outcomes through real activity.



regime in which consumption falls but inflation is unchanged. The pricing kernel assigns more weight to this low-consumption regime, shifting the risk-neutral distribution of consumption without materially changing the inflation distribution. The fourth row considers an inflationary recession: low consumption and high inflation occur jointly. In that case, high-inflation outcomes are bad states of the world and therefore have high value in the market, pushing the risk-neutral inflation distribution to the right.

The key lesson is that the sign and magnitude of the inflation risk premium cannot be inferred from the marginal distribution of inflation alone, and depend on the joint distribution of inflation and real activity. This is why the empirical model combines SPF histograms for GDP growth with inflation-option prices and SPF inflation histograms: information on the distribution of GDP growth helps identify whether high- or low-inflation outcomes are associated with states of high marginal utility.

Figure 3: **Densities of Δc_1 and π_1 under four illustrative specifications.** Panels A.*i* show joint physical and risk-neutral densities of consumption growth and inflation, and Panels B.*i*–C.*i* their corresponding marginals. Specification 1 reproduces the benchmark Gaussian case. Specification 2 introduces positive comovement via a shared shock. Specification 3 adds a recession regime with lower consumption growth. Specification 4 features an inflationary recession with both lower output and higher inflation in the bad state.



3. Model

This section presents the baseline model used in the empirical analysis. Its role is to provide the disciplined change of measure needed to combine physical survey distributions with risk-neutral derivative prices. The framework extends the intuition developed in Section 2.3 along two dimensions. First, inflation and consumption growth are driven by a finite-state Markov process, allowing the joint distribution of nominal and real outcomes to vary across *persistent* macroeconomic regimes. Second, the representative agent has Epstein-Zin preferences, so the stochastic discount factor depends not only on contemporaneous consumption growth but also on continuation values. The resulting structure remains tractable while allowing the pricing of inflation tail risks to depend both on the macroeconomic regime prevailing at the payoff date and on the future prospects associated with that regime.

3.1. Macroeconomic dynamics

We begin by specifying the physical law of motion. The Markov state $\mathbf{z}_t$ summarizes the macroeconomic regime prevailing at date t. Conditional on the next-period regime, inflation and consumption growth are drawn from a regime-specific distribution.

Assumption 1 (Macroeconomic dynamics)**.** *The regime vector $\mathbf{z}_t$ is a J-dimensional regime indicator vector taking values in $\{\mathbf{e}_1, \ldots, \mathbf{e}_J\}$, with $\mathbf{e}_j$ denoting the jth column of the identity matrix $\mathbf{Id}_J$. It follows a time-homogeneous Markov chain with physical transition matrix $\mathbf{P} = \{p_{i,j}\}_{i,j\in\{1,\ldots,J\}^2}$, whose rows sum to one. Transition probabilities are given by*

$$p_{i,j} = \mathbb{P}_t\big(\mathbf{z}_{t+1} = \mathbf{e}_j \mid \mathbf{z}_t = \mathbf{e}_i\big)\,.$$

The horizon is normalized to one period. The notations Δc_{t+1} and π_{t+1} generalize the one-period Δc_1 and π_1 used in Section 2, and are specified as consumption growth and inflation between dates t and $t+1$ are specified as

$$\Delta c_{t+1} = \mathbf{x}'_{c,t+1}\mathbf{z}_{t+1}, \quad \textit{and} \quad \pi_{t+1} = \mathbf{x}'_{\pi,t+1}\mathbf{z}_{t+1}. \tag{1}$$

The vectors $\mathbf{x}_{c,t} = (x_{c,1,t}, \ldots, x_{c,J,t})'$ and $\mathbf{x}_{\pi,t} = (x_{\pi,1,t}, \ldots, x_{\pi,J,t})'$ are serially independent. For a given regime i, the distribution of $\mathbf{x}_{i,t} := (x_{c,i,t}, x_{\pi,i,t})'$ is characterized by its log-Laplace transform φ_i:

$$\varphi_i(\mathbf{v}) := \log \mathbb{E}\left[\exp\left(\mathbf{v}'\mathbf{x}_{i,t}\right)\right], \tag{2}$$

for any admissible vector $\mathbf{v} \in \mathbb{R}^2$*. We denote by*

$$\boldsymbol{\varphi}(\mathbf{v}) := \begin{pmatrix} \varphi_1(\mathbf{v}) \\ \vdots \\ \varphi_J(\mathbf{v}) \end{pmatrix}$$

the vector of regime-specific log-Laplace transforms.

Under Assumption 1, the physical dynamics of $(\mathbf{w}'_{t+1}, \mathbf{z}'_{t+1})'$, with $\mathbf{w}_{t+1} = (\Delta c_{t+1}, \pi_{t+1})'$, are characterized by

$$\psi_t(\mathbf{v}, \mathbf{s}) = \mathbb{E}_t \left[\exp(\mathbf{v}' \mathbf{w}_{t+1} + \mathbf{s}' \mathbf{z}_{t+1}) \right] = \exp \left[\boldsymbol{\varphi}(\mathbf{v}) + \mathbf{s} \right]' \mathbf{P}' \mathbf{z}_t. \tag{3}$$

This expression is useful because it summarizes both the conditional distribution of macroeconomic outcomes and the transition probabilities across regimes. It also makes the change of measure induced by the stochastic discount factor transparent.

3.2. Preferences and continuation values

We next specify preferences, whose role is to determine the stochastic discount factor that maps physical probabilities into risk-neutral probabilities. Here C_t denotes aggregate consumption at date t and $\mathbb{E}_t$ is the expectation conditional on date-t information.

Assumption 2 (Preferences)**.** *The representative agent has Epstein-Zin recursive preferences,*

$$U_t = \left[(1-\beta) C_t^{1-\frac{1}{\rho}} + \beta \left(\mathbb{E}_t \left[U_{t+1}^{1-\gamma} \right] \right)^{\frac{1-\frac{1}{\rho}}{1-\gamma}} \right]^{\frac{1}{1-\frac{1}{\rho}}}, \tag{4}$$

where $\beta \in (0,1)$ *is the subjective discount factor,* $\rho > 0$ *the intertemporal elasticity of substitution, and* $\gamma > 0$ *the coefficient of relative risk aversion.*

The recursive structure implies that the stochastic discount factor depends not only on contemporaneous consumption growth, but also on how the state reached next period changes expected utility from future consumption. In our finite-state economy, this continuation-value channel is therefore state dependent.

This state dependence admits a finite-dimensional fixed-point characterization. The vector $\mathbf{u}$ below collects the continuation-value component associated with each regime.

Proposition 1. *Under Assumptions 1 and 2, a value function of the form*

$$U_t = (\mathbf{u}'\mathbf{z}_t)\, C_t$$

solves the Epstein-Zin recursion (4) if and only if the vector $\mathbf{u} \in \mathbb{R}^J_+$ *satisfies*

$$\mathbf{u} = \left[(1-\beta)\mathbf{1}_J + \beta\left(\mathbf{P}\left(\mathbf{g}\odot\mathbf{u}^{1-\gamma}\right)\right)^{\frac{1-\frac{1}{\rho}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\rho}}}, \tag{5}$$

where $\mathbf{1}_J$ *denotes a* $J\times 1$ *vector of ones,* $\odot$ *denotes element-wise multiplication, all power operators are applied element-wise, and*

$$\mathbf{g} = \exp[\boldsymbol{\varphi}(\boldsymbol{\kappa})], \quad \textit{with} \quad \boldsymbol{\kappa} = \begin{pmatrix} 1-\gamma \\ 0 \end{pmatrix}. \tag{6}$$

Proof. See OA.B.1. □

3.3. Nominal risk-neutral dynamics

A risk-neutral measure is a change of probability induced by the stochastic discount factor and a choice of numeraire. It is not a description of agents' subjective beliefs. Rather, it is the probability measure under which prices can be written as expectations once payoffs are expressed in units of the chosen numeraire. Since the payoffs priced in our empirical application are nominal, we work with the nominal risk-neutral measure, denoted by $\mathbb{Q}^{(n)}$, which uses the one-period nominal bond as numeraire. For any one-period nominal payoff Y_{t+1},

$$\mathbb{E}_t^{\mathbb{Q}^{(n)}}(Y_{t+1}) = \frac{\mathbb{E}_t^{\mathbb{P}}[\mathcal{M}_{t,t+1}^{(n)} Y_{t+1}]}{\mathbb{E}_t^{\mathbb{P}}[\mathcal{M}_{t,t+1}^{(n)}]}. \tag{7}$$

Equivalently, $\mathcal{M}_{t,t+1}^{(n)}/\mathbb{E}_t^{\mathbb{P}}[\mathcal{M}_{t,t+1}^{(n)}]$ is the Radon-Nikodym derivative that maps conditional probabilities under $\mathbb{P}$ into conditional probabilities under $\mathbb{Q}^{(n)}$.

Other choices of numeraire would lead to other risk-neutral measures; for instance, a real risk-neutral measure would use the one-period real bond. The nominal measure is the relevant one here because inflation swaps, caps, and floors are nominal claims.

This change of measure is economically informative because differences between expectations under $\mathbb{P}$ and $\mathbb{Q}^{(n)}$ summarize risk premia, as in Section 5.2. The tractable feature of

the model is that the change of measure preserves the structure of the dynamics: under $\mathbb{Q}^{(n)}$, the process remains conditionally specified by the same Markov state, with tilted conditional distributions and reweighted transition probabilities. The vector $\mathbf{u}$ summarizes continuation values across regimes; through the Epstein-Zin stochastic discount factor, these continuation values affect the tilting. The following proposition characterizes the nominal risk-neutral law of motion for next-period inflation, consumption growth, and the Markov state.

Proposition 2. *Under Assumptions 1 and 2, with* $\mathbf{w}_{t+1} = (\Delta c_{t+1}, \pi_{t+1})'$, *define*

$$\boldsymbol{\alpha} := \begin{pmatrix} -\gamma \\ -1 \end{pmatrix}, \qquad \mathbf{m}^{(n)} := \exp\left[\boldsymbol{\varphi}(\boldsymbol{\alpha}) + \theta \log(\mathbf{u})\right], \quad \theta = \frac{1}{\rho} - \gamma, \tag{8}$$

where the exponential and logarithm are applied element-wise. Under the nominal risk-neutral measure ($\mathbb{Q}^{(n)}$), *the conditional law of* $(\mathbf{w}_{t+1}, \mathbf{z}_{t+1})$ *given* $\mathbf{z}_t$ *is characterized by*

$$\psi_t^{\mathbb{Q}^{(n)}}(\mathbf{v}, \mathbf{s}) := \mathbb{E}_t^{\mathbb{Q}^{(n)}}\left[\exp(\mathbf{v}'\mathbf{w}_{t+1} + \mathbf{s}'\mathbf{z}_{t+1})\right] = \exp\left[\boldsymbol{\varphi}^{\mathbb{Q}^{(n)}}(\mathbf{v}) + \mathbf{s}\right]' \mathbf{Q}^{(n)\prime} \mathbf{z}_t, \tag{9}$$

where

$$\mathbf{Q}^{(n)} := \operatorname{diag}\left(\frac{1}{\mathbf{P}\mathbf{m}^{(n)}}\right) \mathbf{P} \operatorname{diag}\left(\mathbf{m}^{(n)}\right), \tag{10}$$

and

$$\varphi_i^{\mathbb{Q}^{(n)}}(\mathbf{v}) := \varphi_i(\mathbf{v} + \boldsymbol{\alpha}) - \varphi_i(\boldsymbol{\alpha}), \quad i = 1, \ldots, J. \tag{11}$$

Proof. See OA.B.2. □

The proposition shows how the nominal risk-neutral dynamics are obtained from the physical dynamics. The conditional transform remains in the same form as under $\mathbb{P}$, but with two adjustments. First, the conditional distribution within each destination regime is exponentially tilted through $\boldsymbol{\varphi}^{\mathbb{Q}^{(n)}}$. Second, the transition matrix is changed from $\mathbf{P}$ to $\mathbf{Q}^{(n)}$. Equation (10) shows that this transition matrix is obtained by reweighting each destination regime by $m_j^{(n)}$ and then normalizing each row to sum to one. Hence nominal risk-neutral probabilities differ from physical probabilities both because outcomes within a regime are tilted and because future regimes receive different weights. This is the formal counterpart of the intuition in Section 2: inflation outcomes receive different market weights depending on the real-activity states and continuation values with which they are associated.

For the empirical implementation, we use a conditionally Gaussian specification.

Assumption 3 (Conditionally Gaussian case)**.** *Conditionally on the regime, the distribution of* $\mathbf{w}_{t+1}$ *is bivariate Gaussian:*

$$\mathbf{w}_{t+1} \mid \{\mathbf{z}_{t+1} = \mathbf{e}_i\} \sim \mathcal{N}(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i)\,.$$

Under Assumptions 1 and 3, the physical log-Laplace transform is

$$\varphi_i(\mathbf{v}) = \mathbf{v}'\boldsymbol{\mu}_i + \frac{1}{2}\mathbf{v}'\boldsymbol{\Sigma}_i\mathbf{v}, \quad i = 1, \ldots, J. \tag{12}$$

The nominal risk-neutral distribution remains conditionally Gaussian, with shifted conditional means.

Corollary 1 (Risk-neutral conditional distribution of $\mathbf{w}_{t+1}$ in the Gaussian case)**.** *Under Assumptions 1, 2, and 3,*

$$\mathbf{w}_{t+1} \mid \{\mathbf{z}_{t+1} = \mathbf{e}_i\} \sim \mathcal{N}\left(\boldsymbol{\mu}_i^{\mathbb{Q}^{(n)}}, \boldsymbol{\Sigma}_i\right), \quad \textit{with} \quad \boldsymbol{\mu}_i^{\mathbb{Q}^{(n)}} = \boldsymbol{\mu}_i + \boldsymbol{\Sigma}_i\boldsymbol{\alpha}.$$

The nominal risk-neutral transition matrix of the Markov state is $\mathbf{Q}^{(n)}$ *in* (10)*.*

Proof. See OA.B.3. □

3.4. Survey probabilities and pricing

The model delivers the objects matched in estimation: physical probabilities for SPF histogram bins and risk-neutral prices for inflation derivatives. These objects are generated from the same underlying state dynamics, but under different measures. SPF probabilities are computed under the physical measure, using $\mathbf{P}$. Inflation derivative prices are computed under the nominal risk-neutral measure, using $\mathbf{Q}^{(n)}$.

Consider first the survey probabilities. Conditional on the current regime $\mathbf{z}_t = \mathbf{e}_i$, the one-period-ahead physical mixture weights are given by the ith row of $\mathbf{P}$, denoted $\boldsymbol{\lambda}_t^{\mathbb{P}} = \mathbf{P}'\mathbf{z}_t$. In the conditionally Gaussian case, the model-implied c.d.f.s of inflation and consumption growth are

$$F_{\pi,t}^{\mathbb{P}}(a) = \sum_{j=1}^{J} \lambda_{j,t}^{\mathbb{P}} \Phi\left(\frac{a - \mu_{\pi,j}}{\sigma_{\pi,j}}\right), \tag{13}$$

$$F_{c,t}^{\mathbb{P}}(a) = \sum_{j=1}^{J} \lambda_{j,t}^{\mathbb{P}} \Phi\left(\frac{a - \mu_{c,j}}{\sigma_{c,j}}\right), \tag{14}$$

where $\mu_{\pi,j}$ and $\mu_{c,j}$ are the inflation and consumption-growth entries of $\boldsymbol{\mu}_j$, and $\sigma^2_{\pi,j}$ and $\sigma^2_{c,j}$ are the corresponding diagonal entries of $\boldsymbol{\Sigma}_j$. Hence, for any SPF bin $(a, b]$, the model-implied physical probability is $F^{\mathbb{P}}_{\pi,t}(b) - F^{\mathbb{P}}_{\pi,t}(a)$ for inflation and $F^{\mathbb{P}}_{c,t}(b) - F^{\mathbb{P}}_{c,t}(a)$ for real activity. In the empirical implementation, GDP-growth histogram probabilities are matched using the model-implied consumption growth distribution as the counterpart for real activity.

We now turn to inflation derivatives. Since these payoffs are nominal, they are priced under the nominal risk-neutral measure characterized in Proposition 2.

Proposition 3. *Under Assumptions 1, 2, and 3, the annualized inflation swap rate between dates t and $t+1$ is*

$$\pi^S_{t,1} = \log \psi^{\mathbb{Q}^{(n)}}_t \left(\begin{pmatrix} 0 \\ 1 \end{pmatrix}, \mathbf{0} \right). \tag{15}$$

Equivalently, denoting by $\mathbf{e}_\pi = (0,1)'$ and by i the current regime, $\mathbf{z}_t = \mathbf{e}_i$,

$$\exp(\pi^S_{t,1}) = \frac{\sum_{j=1}^J p_{ij} u_j^\theta \exp\left(\varphi_j(\boldsymbol{\alpha} + \mathbf{e}_\pi)\right)}{\sum_{j=1}^J p_{ij} u_j^\theta \exp\left(\varphi_j(\boldsymbol{\alpha})\right)}. \tag{16}$$

Proof. See OA.B.4. □

In estimation, the observed one-year inflation swap and the SPF point forecast are used as two exact restrictions on the two regime-specific inflation means. The first restriction matches the nominal risk-neutral exponential mean in (16); the second matches the physical mean of inflation. This joint normalization anchors the observed inflation-compensation wedge, while the survey histograms and option prices discipline the shapes of the physical and risk-neutral distributions. Appendix Appendix C provides the construction.

Proposition 4. *Under Assumptions 1, 2, and 3, the date-t prices of an inflation floor and cap with strike K and maturity $t+1$ are*

$$\begin{aligned} Floor_{\pi,t}(K) &= e^{-r_t+K} h(\boldsymbol{\theta}^{\mathbb{Q}^{(n)}}_{\pi,t}, K) - e^{-r_t} H(\boldsymbol{\theta}^{\mathbb{Q}^{(n)}}_{\pi,t}, 1, K), \\ Cap_{\pi,t}(K) &= e^{-r_t}[\exp(\pi^S_{t,1}) - \exp(K)] + Floor_{\pi,t}(K), \end{aligned}$$

where r_t is the continuously compounded one-period nominal risk-free rate, h and H are given in (A.1) and (A.5), respectively, and

$$\boldsymbol{\theta}^{\mathbb{Q}^{(n)}}_{\pi,t} = (\boldsymbol{\mu}^{\mathbb{Q}^{(n)}}_\pi, \boldsymbol{\sigma}_\pi, \boldsymbol{\lambda}^{\mathbb{Q}^{(n)}}_t), \quad \textit{with} \quad \boldsymbol{\lambda}^{\mathbb{Q}^{(n)}}_t = \mathbf{Q}^{(n)\prime} \mathbf{z}_t. \tag{17}$$

Proof. See OA.B.5. □

4. Data

Our empirical implementation combines three euro-area data sources at the one-year horizon over the period from October 2009 to December 2025. The first source consists of the ECB SPF histograms for one-year-ahead inflation and GDP growth. The second source is the cross section of inflation-linked derivative prices. These market prices summarize how financial markets value future inflation outcomes and therefore discipline the risk-neutral distribution of inflation, while the survey histograms reflect forecasters' assessments of the corresponding physical distributions of nominal and real outcomes. Bringing these sources together allows the estimation to jointly discipline inflation tail risks and the real-activity states with which they are associated.

On the survey side, we use the SPF histograms for one-year-ahead inflation and GDP growth. Because the published bins vary over time, and their endpoints can shift substantially during episodes of unusually large inflation or GDP fluctuations, we reconstruct the relevant support separately for each survey round and for each variable. The estimation then matches the cumulative probabilities implied by the harmonized histograms at the corresponding bin thresholds. For inflation, we additionally use the cross-sectional average of the SPF one-year point forecasts. This series anchors the physical mean through the joint normalization described in Appendix Appendix C. For GDP, the model counterpart is consumption growth, which we interpret as the representative-agent aggregate measure of real activity.

On the market side, we use one-year euro-area inflation swaps together with one-year zero-coupon floors and caps. The swap rate anchors the level of market-implied expected inflation, while floor and cap prices provide information about the shape of the risk-neutral distribution, especially its downside and upside tails. Although all market instruments have a one-year maturity, their payoffs are delivered in the regime reached at date $t+1$. With persistent regimes, this destination regime contains information about future consumption prospects, so the continuation-value vector $\mathbf{u}$ affects the one-period stochastic discount factor

and hence the risk-neutral distribution.[3] In the baseline estimation, floor strikes are set at -1, 0, 1, and 2 percent, while cap strikes are set at 2, 3, 4, and 5 percent. To account for time-varying market liquidity, we smoothen the daily market series using a backward-looking 22-trading-day moving average before aligning them with the survey dates.

The estimation sample begins with the first survey-round for which both inflation-option data and SPF distributions are jointly available. Additional details on data construction and recursive estimation are provided in Appendix B.

5. Estimation and Results

5.1. Model fit and estimation approach

We first describe how the model is brought to the data, and then use the fitted distributions to study the joint behavior of inflation and real activity under the physical and risk-neutral measures. At each survey round in the common sample, the estimation targets four sets of objects: the cumulative probabilities from the SPF inflation histogram, the cumulative probabilities from the SPF GDP-growth histogram, the cross section of inflation floor prices, and the cross section of inflation cap prices. The one-year SPF point forecast and the one-year inflation swap are imposed as exact restrictions on the physical and nominal risk-neutral inflation means, respectively. Section 3 describes the normalization. Since the model delivers closed-form expressions for all targeted objects, the estimation can be implemented date by date while preserving the structural link between physical and risk-neutral distributions through the stochastic discount factor.

The estimation is recursive. At each date, the loss function measures the discrepancies between observed and model-implied survey probabilities and option prices, together with regularization terms that discourage abrupt changes in parameters across adjacent dates. This smoothing is useful in a nonlinear environment with several cross-sectional objects observed at each date: it reduces the influence of local noise without imposing a constant parameter vector over the whole sample. The objective also penalizes implausibly high maxi-

[3] In a time-additive CRRA benchmark, this channel is absent because the one-period stochastic discount factor depends only on contemporaneous consumption growth. Under Epstein-Zin preferences, the stochastic discount factor also prices the continuation value associated with the future macroeconomic environment.

mum Sharpe ratios, and the preference parameters are restricted to economically meaningful ranges. Additional implementation details are collected in Appendix B.

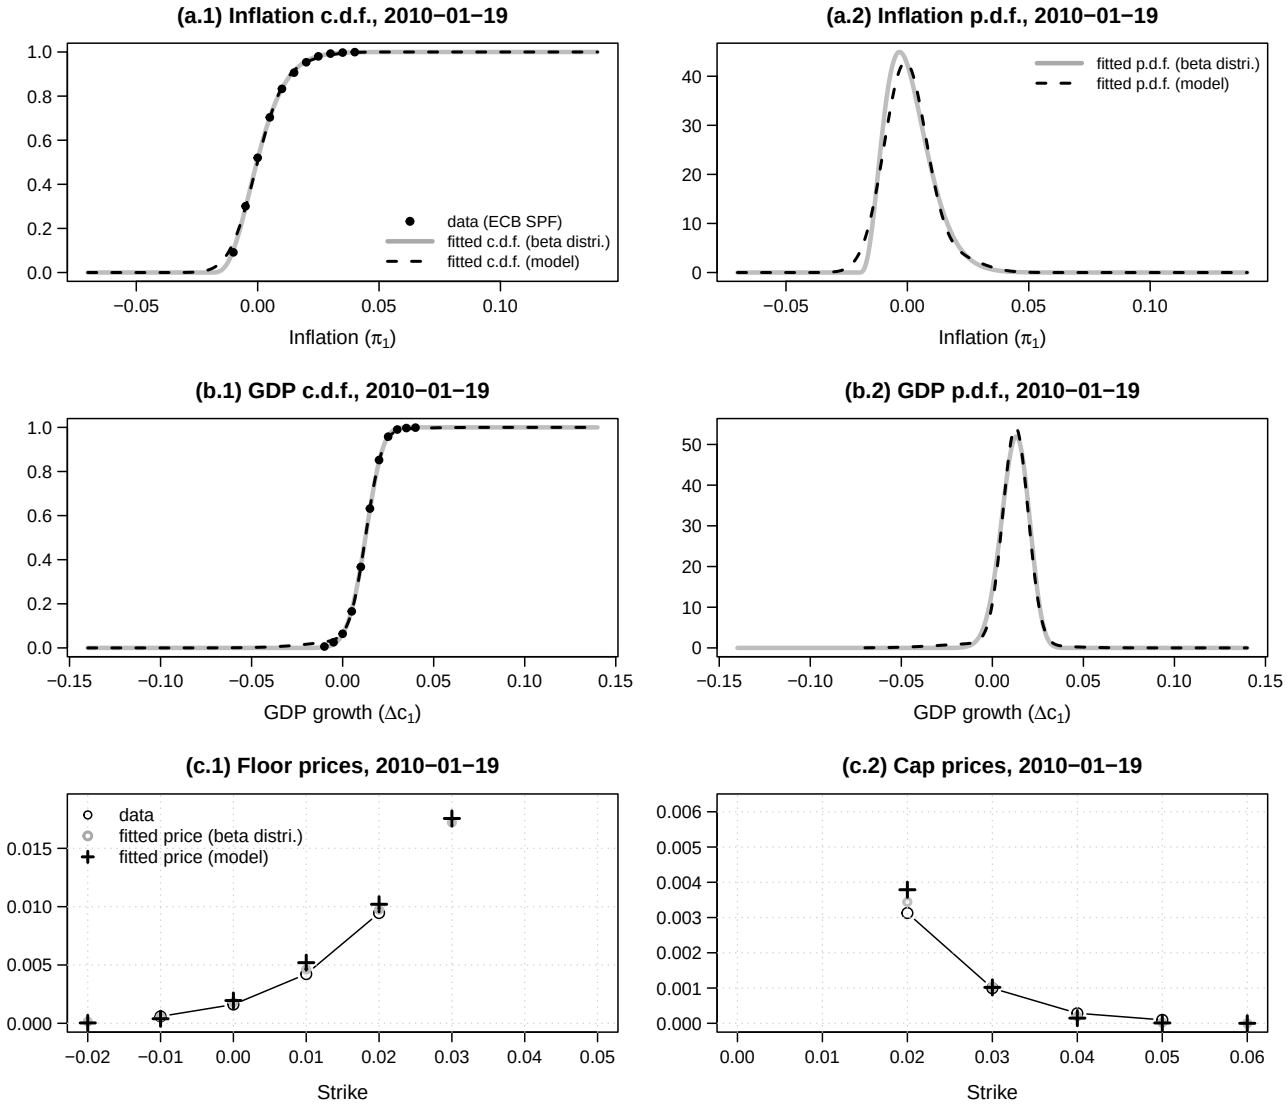
Figures 4 and 5 illustrate the fit for two representative dates, January 19, 2010 and January 9, 2025. In each case, the model is compared with the inflation and GDP survey c.d.f.s and with the observed cap and floor price schedules. For visual comparison, the figures also report beta-smoothed distributions, constructed in the spirit of the procedure described in Appendix 6.5 of Grishchenko et al. (2017). These smooth curves are not used in the structural estimation; they are included only to show how the model-implied distributions compare with flexible reduced-form summaries of the same survey data. Time-series fit diagnostics for all estimation dates are reported in Appendix D. These diagnostics are important for interpretation: the joint distributions studied below are only credible if the model first matches the marginal survey information and the inflation-option cross sections that discipline them.

Figure 6 then shows what the fitted model implies for the joint distribution on the same two dates. The figure displays the physical and nominal risk-neutral distributions of one-year-ahead inflation and GDP growth, together with the associated marginal distributions. It therefore moves from fit diagnostics to the central object of the paper: which regions of the inflation-growth distribution are reweighted by the pricing kernel. In January 2010, the nominal risk-neutral distribution assigns substantially more weight to states combining weaker real activity with different inflation outcomes than the physical distribution, implying a sizable wedge between survey beliefs and market prices. By contrast, in January 2025 the physical and risk-neutral distributions are much closer, consistent with a smaller reweighting of the joint distribution at that date. The comparison illustrates that the model is not only fitting marginal inflation and GDP distributions; it also recovers how markets price particular regions of the joint inflation-growth space.

5.2. Dependence and risk premia

We next examine how the dependence between inflation and real activity is priced over time. Figure 7 reports the model-implied covariance between one-year-ahead inflation and GDP growth under the physical measure, together with the corresponding mean wedges between the physical and nominal risk-neutral distributions. The covariance summarizes the

Figure 4: **Fit of the model on January 19, 2010.** The figure reports the fit of the model to the one-year-ahead SPF inflation and GDP distributions and to the cross section of one-year inflation floor and cap prices for January 19, 2010.

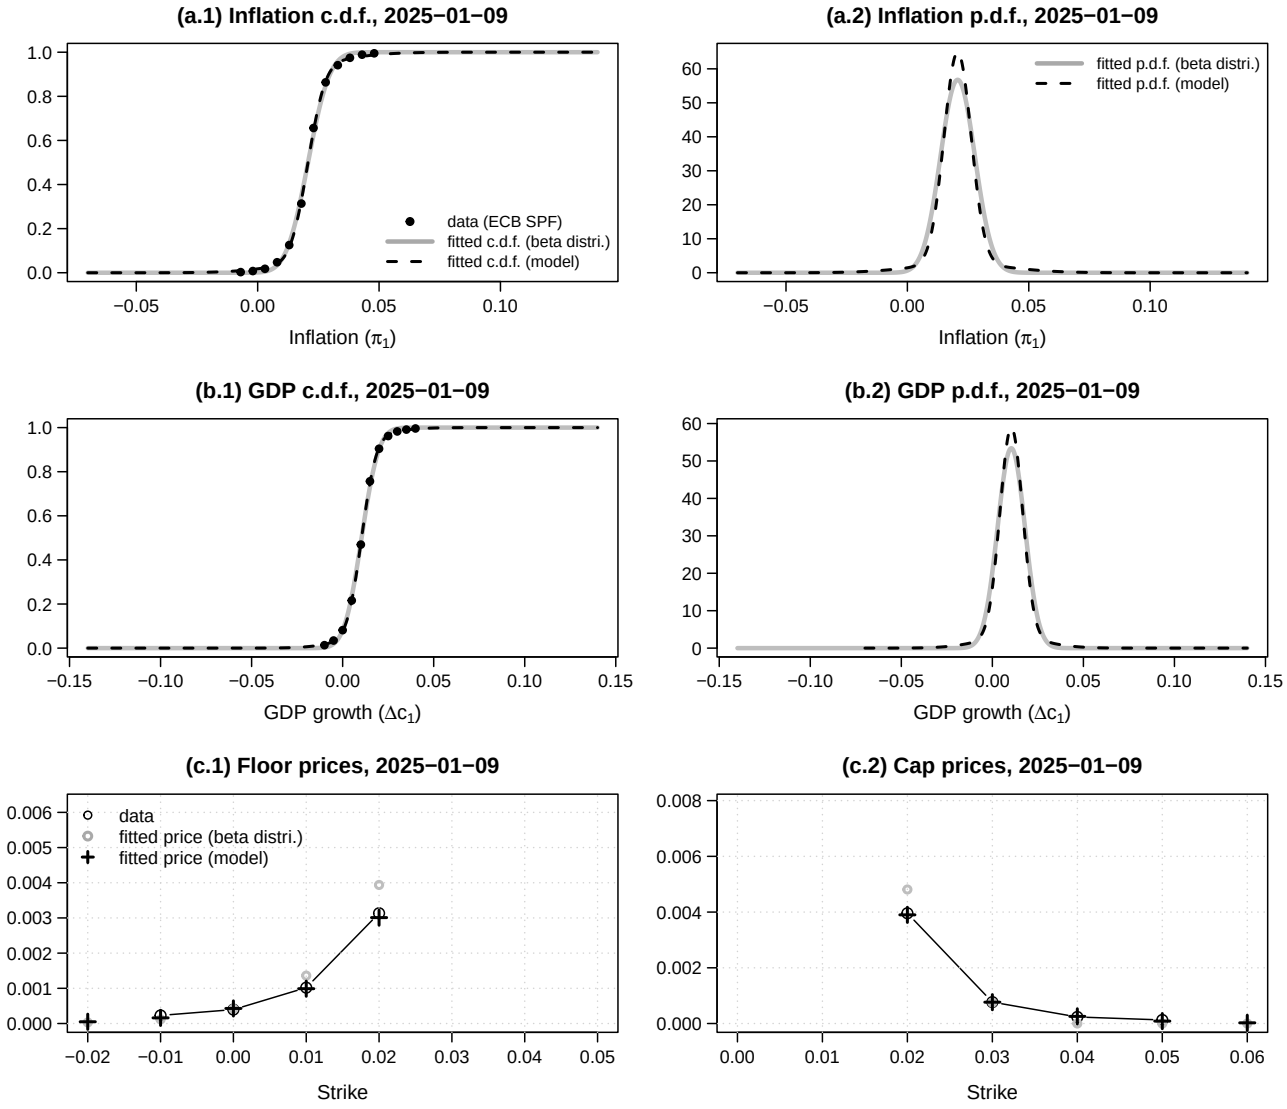


nature of macroeconomic risk; the mean wedges summarize how this risk is valued by the estimated nominal pricing kernel.

The interpretation follows from a simple pricing identity. Let $\overline{\mathcal{M}}^{(n)}_{t,t+1} = \mathcal{M}^{(n)}_{t,t+1} / \mathbb{E}^{\mathbb{P}}_t[\mathcal{M}^{(n)}_{t,t+1}]$ denote the normalized nominal pricing kernel. For a generic one-period payoff X_{t+1}, the fixed leg of a swap that exchanges this payoff against X^S_t satisfies

$$X^S_t = \mathbb{E}^{\mathbb{Q}^{(n)}}_t(X_{t+1}) = \mathbb{E}^{\mathbb{P}}_t(X_{t+1}) + \mathrm{Cov}^{\mathbb{P}}_t\left(\overline{\mathcal{M}}^{(n)}_{t,t+1}, X_{t+1}\right). \tag{18}$$

Figure 5: **Fit of the model on January 9, 2025.** The figure reports the fit of the model to the one-year-ahead SPF inflation and GDP distributions and to the cross section of one-year inflation floor and cap prices for January 9, 2025.

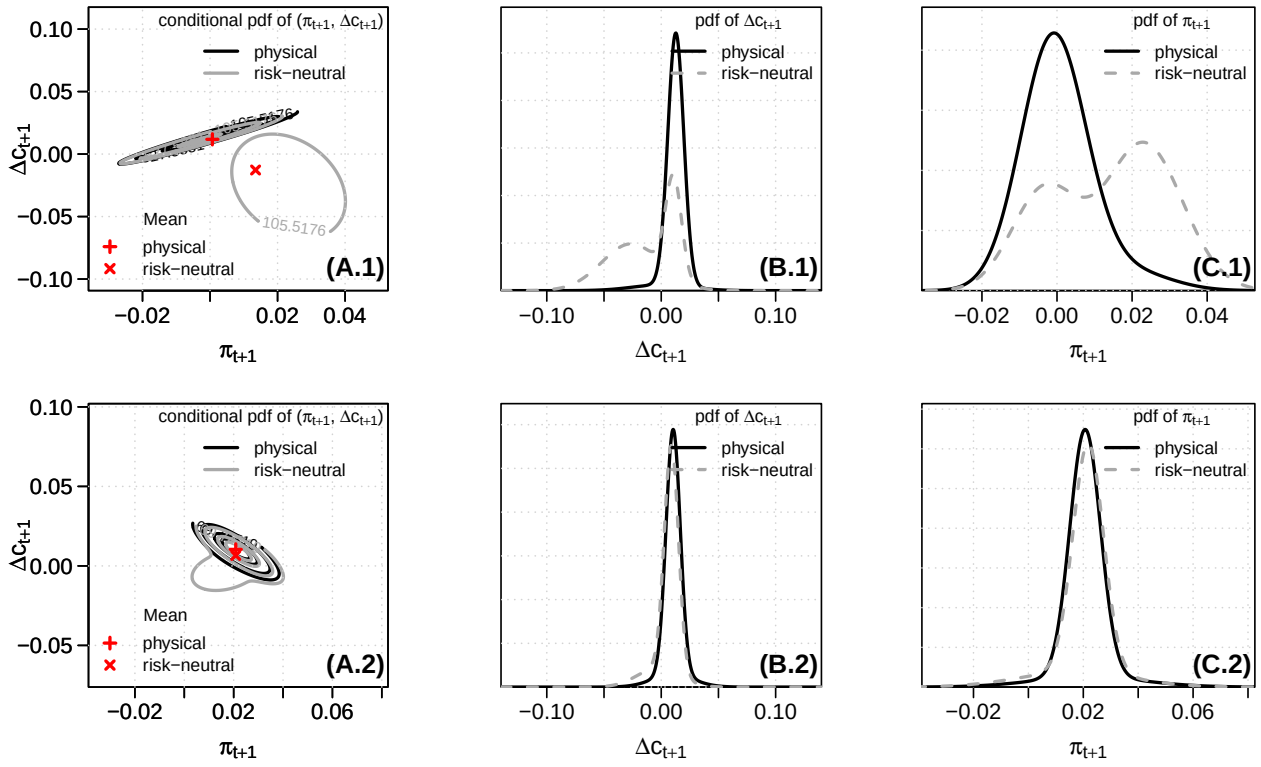


Thus, the mean wedge for any payoff is its covariance with the normalized nominal pricing kernel. A payoff has a high risk-neutral mean relative to its physical mean when it pays in high-marginal-utility states. This identity is the asset-pricing content of the comparison between survey-implied physical distributions and derivative-implied nominal risk-neutral distributions.

We report two such wedges. For inflation,

$$\mathrm{IRP}_t = \mathbb{E}_t^{\mathbb{Q}^{(n)}}(\pi_{t+1}) - \mathbb{E}_t^{\mathbb{P}}(\pi_{t+1}), \tag{19}$$

Figure 6: **Fitted physical and risk-neutral bivariate distributions on selected dates.** The figure reports the fitted one-year-ahead distributions for January 19, 2010 (row 1) and January 9, 2025 (row 2). Panels A.*i* show the joint density of GDP growth Δc_{t+1} and inflation π_{t+1} under the physical measure (black contours) and the nominal risk-neutral measure (grey contours), together with their means. Panels B.*i* report the corresponding marginal densities of GDP growth, and Panels C.*i* report the corresponding marginal densities of inflation.



which is the mean risk-neutral inflation wedge implied by the model.[4] For real activity, we report the wedge with the opposite sign,

$$\mathrm{GRP}_t = \mathbb{E}_t^{\mathbb{P}}(\Delta c_{t+1}) - \mathbb{E}_t^{\mathbb{Q}^{(n)}}(\Delta c_{t+1}). \tag{20}$$

With this convention, a positive GRP_t means that the nominal risk-neutral distribution puts more weight on weak-growth states than the physical distribution. This GDP-growth premium is model-implied rather than directly traded: it is the premium on a cash-settled pay-

[4]This object is closely related to, but not identical to, the observed inflation swap wedge $\pi_{t,1}^S - \mathbb{E}_t^{\mathbb{P}}(\pi_{t+1})$. The traded inflation swap in our data satisfies $\pi_{t,1}^S = \log \mathbb{E}_t^{\mathbb{Q}^{(n)}}[\exp(\pi_{t+1})]$ and therefore also contains convexity terms associated with the risk-neutral distribution of inflation. The general logic is standard in the literature on inflation risk premia, where inflation compensation reflects both expected inflation and compensation for bearing inflation risk (Ang et al., 2008; Haubrich et al., 2012; Grishchenko and Huang, 2013; Hördahl and Tristani, 2010).

off indexed to real GDP growth, using consumption growth as the model counterpart of real activity.[5]

Panel (a) is informative because the inflation-growth covariance determines how inflation loads on marginal utility. This link is transparent in the static CRRA-Gaussian benchmark, where the nominal SDF is proportional to $\exp(-\gamma \Delta c_{t+1} - \pi_{t+1})$. Exponential tilting of a Gaussian distribution then gives

$$\mathbb{E}_t^{\mathbb{Q}^{(n)}}(\pi_{t+1}) - \mathbb{E}_t^{\mathbb{P}}(\pi_{t+1}) = -\gamma \operatorname{Cov}_t^{\mathbb{P}}(\pi_{t+1}, \Delta c_{t+1}) - \operatorname{Var}_t^{\mathbb{P}}(\pi_{t+1}). \tag{21}$$

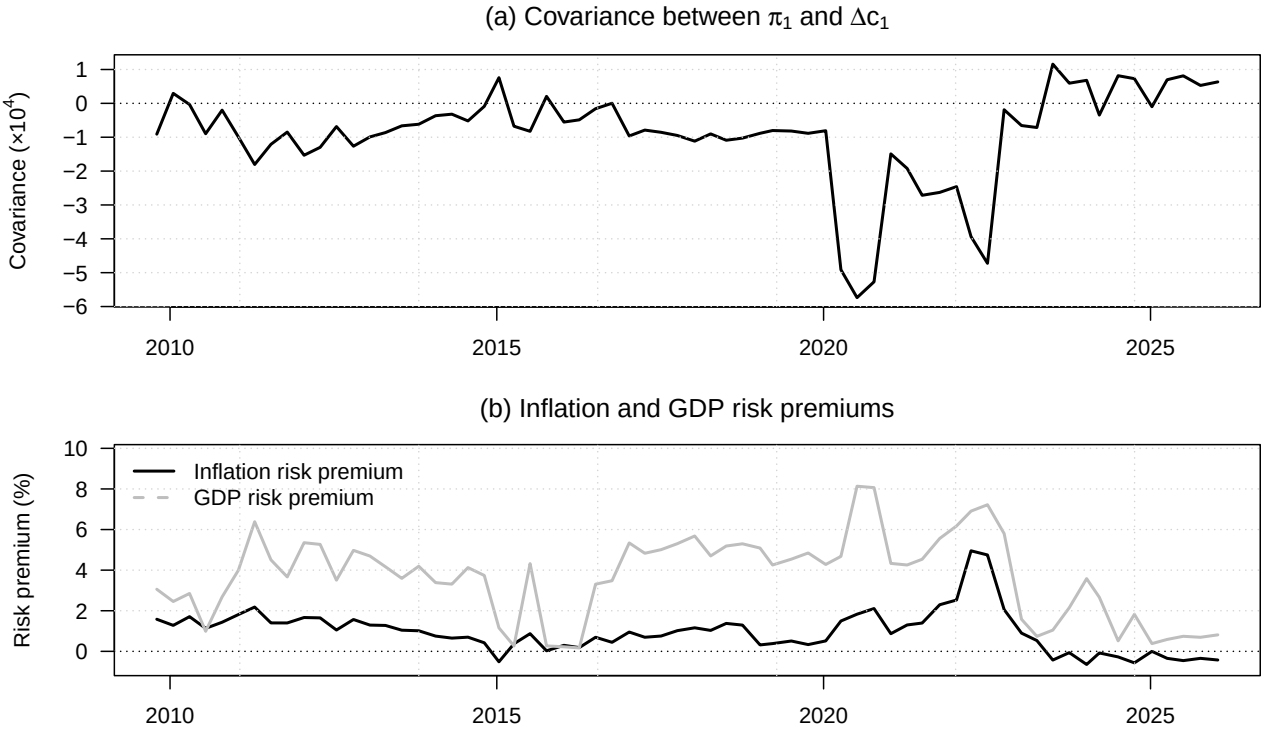
Thus, the inflation-growth covariance is not itself the primitive pricing covariance in (18); it enters through the exposure of the nominal SDF to real activity. A negative covariance means that high-inflation outcomes occur when consumption growth is low and marginal utility is high, which raises the real-activity component of the inflation premium. The variance term reflects the nominal nature of the pricing measure: inflation also enters the nominal SDF because it erodes nominal payoffs. In the Epstein-Zin baseline, this mechanism remains present, but the mapping from the covariance to the premium is no longer linear because the SDF also reflects continuation values, regime probabilities, and state-dependent risks.

Figure 7 shows that the inflation-growth covariance changes sign over the sample. This time variation is economically important. Negative covariance corresponds to episodes in which high inflation is associated with weak real activity, so inflation risk is concentrated in high-marginal-utility states and the inflation risk premium tends to be positive. When the covariance is positive, or when low-inflation outcomes are the ones associated with weak activity, the same pricing logic lowers the inflation premium and may turn it negative. The GDP-growth premium provides the complementary view from the real-activity margin: it

[5] It is not a premium on nominal GDP growth, which would involve the sum of real growth and inflation. The object is related to the macro-finance literature on GDP-linked securities, where growth-indexed claims embed risk premia because their payoffs are procyclical and therefore low in bad states (Athanasoulis et al., 1999; Borensztein and Mauro, 2004; Blanchard et al., 2016; Consiglio and Zenios, 2018; Eguren-Martin et al., 2021; Mouabbi et al., 2024). Since liquid GDP-linked bonds have not been issued by large advanced economies, this literature typically studies hypothetical or model-implied premia. Reported magnitudes range from a few tens of basis points to several hundred basis points, and signs depend on the convention. In Eguren-Martin et al. (2021), the premium is a component of the conventional-minus-GDP-linked yield breakeven for nominal GDP-linked bonds; because compensation for GDP risk raises GDP-linked yields, that component is negative and reaches about 700 basis points in absolute value at short maturities. Our GRP_t uses the opposite convention.

summarizes how much the estimated pricing kernel tilts the distribution of activity toward low-growth states.

Figure 7: **Time variation in dependence and risk premia.** Panel (a) reports the model-implied covariance between one-year-ahead inflation and GDP growth under the physical measure. Panel (b) reports the inflation risk premium, defined as $\mathbb{E}_t^{\mathbb{Q}^{(n)}}(\pi_{t+1}) - \mathbb{E}_t^{\mathbb{P}}(\pi_{t+1})$, and the GDP-growth risk premium, defined as $\mathbb{E}_t^{\mathbb{P}}(\Delta c_{t+1}) - \mathbb{E}_t^{\mathbb{Q}^{(n)}}(\Delta c_{t+1})$. Both premia are expressed in percentage points.



5.3. Tail-environment probabilities

The covariance and risk premia are useful summaries, but they compress the joint distribution into a small number of moments. Figures 8 and 9 provide more localized views by translating the estimated distributions into probabilities of selected regions of the one-year-ahead inflation-growth space. These probabilities should not be interpreted as probabilities of new structural shocks. They describe the likelihood that one-year-ahead inflation and growth fall within particular regions of the conditional joint distribution, given information available at the survey date.

The two figures are complementary. Figure 8 uses fixed regions of the $(\Delta c_{t+1}, \pi_{t+1})$ space, allowing probabilities to be compared over time in levels. Figure 9 instead uses thresholds

that move with the model-implied physical expectations at each date, which focuses on adverse deviations relative to the current expected macroeconomic environment. In both cases, each region is evaluated under the physical measure and under the nominal risk-neutral measure. The gap between the two measures is informative about which parts of the joint distribution are especially valuable to insure against.

When the risk-neutral probability of the low-growth/low-inflation region exceeds its physical probability, financial markets attach extra value to insurance against disinflationary recession states. Conversely, when the low-growth/high-inflation region is more heavily weighted under the risk-neutral measure, markets are particularly concerned about stagflationary or adverse supply-side outcomes. The time variation in these wedges provides a tail-based complement to the covariance and risk-premium evidence in Figure 7: rather than summarizing the pricing distortion with a single premium, the figures show which parts of the joint distribution are being reweighted by the pricing kernel.

Figure 9 complements the fixed-threshold evidence by measuring tail environments relative to the macroeconomic outlook prevailing at each survey date. The figure therefore asks whether markets place extra value on outcomes that are bad relative to current expectations, rather than on outcomes that are bad only in absolute terms. A widening gap between the risk-neutral and physical probabilities indicates that the corresponding relative tail–disinflationary recession or inflationary recession–has become particularly valuable to hedge.

Figure 8: **Physical and risk-neutral probabilities of selected macroeconomic tail environments.** Panels (a.1)–(b.1) define fixed regions in the $(\Delta c_{t+1}, \pi_{t+1})$ space. Panel (a) corresponds to $\Delta c_{t+1} < 0.5\%$ and $\pi_{t+1} < 1\%$. Panel (b) corresponds to $\Delta c_{t+1} < 0.5\%$ and $\pi_{t+1} > 3.5\%$. The shaded areas show the regions whose probabilities are computed. Panels (a.2)–(b.2) report the corresponding time series of probabilities under the physical measure (solid line) and under the nominal risk-neutral measure (dashed line).

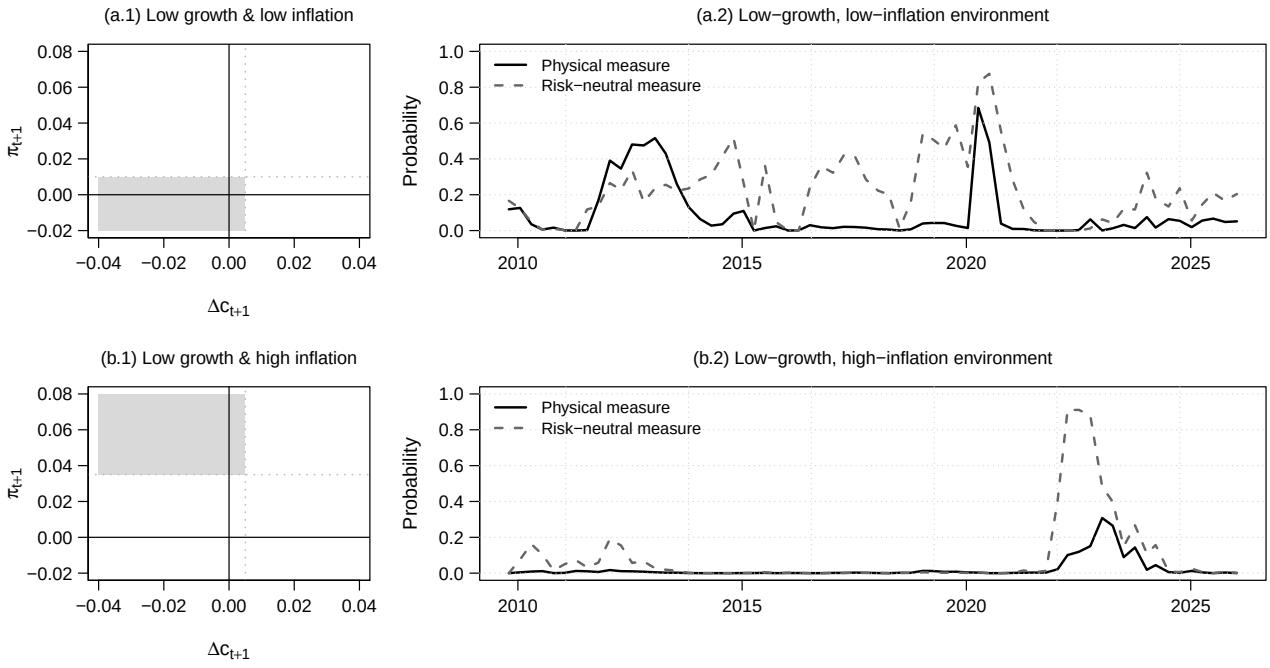
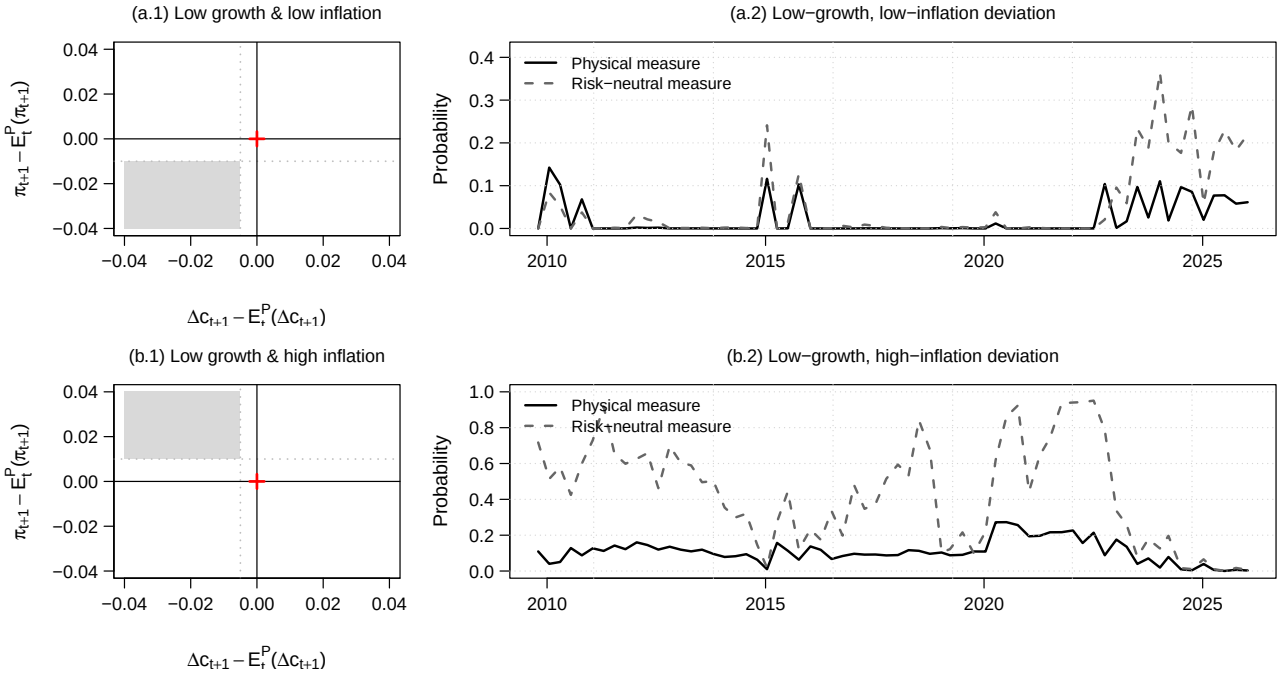


Figure 9: **Physical and risk-neutral probabilities of relative tail deviations.** Panels (a.1)–(b.1) define regions in deviation from the model-implied physical means. At each date, the growth threshold is $\mathbb{E}_t^{\mathbb{P}}(\Delta c_{t+1}) - 0.5\%$. Panel (a) uses the inflation threshold $\mathbb{E}_t^{\mathbb{P}}(\pi_{t+1}) - 1\%$, while Panel (b) uses the inflation threshold $\mathbb{E}_t^{\mathbb{P}}(\pi_{t+1}) + 1\%$. The shaded areas show the regions whose probabilities are computed after centering both variables at their physical means; the red cross marks the reference point $(0, 0)$, corresponding to the physical means. Panels (a.2)–(b.2) report the corresponding time series of probabilities under the physical measure (solid line) and under the nominal risk-neutral measure (dashed line).



6. Concluding remarks

This paper develops a framework for recovering joint inflation-growth risks from two forward-looking sources of information: survey density forecasts and inflation derivative prices. Survey histograms discipline physical distributions of inflation and real activity, while inflation swaps, caps, and floors discipline the nominal risk-neutral distribution of inflation. The identifying mechanism is the stochastic discount factor. Because it maps physical probabilities into risk-neutral probabilities, the wedge between survey-implied and market-implied inflation distributions is informative about the real-activity states in which inflation outcomes occur. This is the central contribution of the paper: the framework does not only ask how likely high- or low-inflation outcomes are, but also whether these outcomes are associated with weak or strong real activity.

The euro-area application shows that this joint perspective is empirically important. The dependence between inflation and growth varies substantially over time, and these movements are closely related to the sign and magnitude of inflation risk premia. Periods in which high inflation is associated with weak real activity generate a very different pricing of inflation tails from periods in which low inflation is associated with weak activity. The fitted model also delivers physical and risk-neutral probabilities of specific macroeconomic tail environments, such as disinflationary recessions and inflationary recessions. These probabilities show not only how likely such environments are perceived to be, but also which regions of the joint distribution are especially valuable to insure against. The broader message is that distributional surveys and derivative prices are most informative when used jointly.

References

Andrade, P., V. Fourel, E. Ghysels, and J. Idier (2014). The financial content of inflation risks in the euro area. *International Journal of Forecasting 30*(3), 648–659.

Ang, A., G. Bekaert, and M. Wei (2008). The term structure of real rates and expected inflation. *Journal of Finance 63*(2), 797–849.

Athanasoulis, S., R. J. Shiller, and E. V. Wincoop (1999). Macro markets and financial security. *Economic Policy Review 5*(1), 21–39.

Backus, D., M. Chernov, and I. Martin (2011). Disasters implied by equity index options. *The Journal of Finance 66*(6), 1969–2012.

Barro, R. (2006). Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics 121*(3), 823–866.

Barro, R. J. and G. Y. Liao (2021). Rare disaster probability and options pricing. *Journal of Financial Economics 139*(3), 750–769.

Bassetti, F., R. Casarin, and M. D. Negro (2024). A Bayesian approach for inference on probabilistic surveys. Staff Reports 1025, Federal Reserve Bank of New York. Revised August 2024.

Blanchard, O. J., P. Mauro, and J. Acalin (2016). The case for growth-indexed bonds in advanced economies today. Policy Briefs PB16-2, Peterson Institute for International Economics.

Boero, G., J. Smith, and K. F. Wallis (2014). The measurement and characteristics of professional forecasters' uncertainty. *Journal of Applied Econometrics doi: 10.1002/jae.2400*.

Bollerslev, T., V. Todorov, and L. Xu (2015). Tail risk premia and return predictability. *Journal of Financial Economics 118*(1), 113–134.

Böninghausen, B., G. Kidd, and R. de Vincent-Humphreys (2018). Interpreting option-implied probabilities. European Central Bank Economic Bulletin, Issue 6/2018, Box 1.

Borensztein, E. and P. Mauro (2004). The case for GDP-indexed bonds. *Economic Policy 19*(38), 165–216.

Clark, T. E., G. Ganics, and E. Mertens (2022). What is the predictive value of SPF point and density forecasts? Working Papers 22-37, Federal Reserve Bank of Cleveland.

Clark, T. E., G. Ganics, and E. Mertens (2024). Constructing fan charts from the ragged edge of SPF forecasts. Working Papers 2429, Banco de Espana.

Consiglio, A. and S. A. Zenios (2018). Pricing and hedging GDP-linked bonds in incomplete markets. *Journal of Economic Dynamics and Control 88*, 137–155.

Eguren-Martin, F., A. Meldrum, and W. Yan (2021). No-arbitrage pricing of GDP-linked bonds. *Journal of Banking and Finance 126*, 106075.

Engelberg, J., C. F. Manski, and J. Williams (2009). Comparing the point predictions and subjective probability distributions of professional forecasters. *Journal of Business & Economic Statistics 27*, 30–41.

European Central Bank (2013). The information content of inflation options. Monthly Bulletin, July 2013, Box 5.

Gabaix, X. (2012). Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics 127*(2), 645–700.

Galvao, A. B., A. Garratt, and J. Mitchell (2021). Does judgment improve macroeconomic density forecasts? *International Journal of Forecasting 37*(3), 1247–1260.

Gimeno, R. and A. Ibáñez (2018). The Eurozone (expected) inflation: An option's eyes view. *Journal of International Money and Finance 86*, 70–92.

Grishchenko, O., S. Mouabbi, and J.-P. Renne (2019). Measuring inflation anchoring and uncertainty: A U.S. and euro area comparison. *Journal of Money, Credit and Banking 51*(5), 1053–1096.

Grishchenko, O. V. and J.-Z. Huang (2013). Inflation risk premium: Evidence from the TIPS market. *Journal of Fixed Income 22*(4), 5–30.

Grishchenko, O. V., S. Mouabbi, and J.-P. Renne (2017). Measuring inflation anchoring and uncertainty: A U.S. and euro area comparison. Finance and Economics Discussion Series 2017-102, Board of Governors of the Federal Reserve System (U.S.).

Haubrich, J., G. Pennacchi, and P. Ritchken (2012). Inflation expectations, real rates, and risk premia: Evidence from inflation swaps. *Review of Financial Studies 25*(5), 1588–1629.

Hilscher, J., A. Raviv, and R. Reis (2025). How likely is an inflation disaster? *The Review of Financial Studies*, 744–782.

Hördahl, P. and O. Tristani (2010). Inflation risk premia in the U.S. and the euro area. ECB Working Paper Series No. 1270.

Kelly, B. and H. Jiang (2014). Editor's choice tail risk and asset prices. *The Review of Financial Studies 27*(10), 2841–2871.

Kitsul, Y. and J. H. Wright (2013). The economics of options-implied inflation probability density functions. *Journal of Financial Economics 110*(3), 696–711.

Mouabbi, S., J.-P. Renne, and J.-G. Sahuc (2024). Debt-stabilizing properties of GDP-linked securities: A macro-finance perspective. *Journal of Banking and Finance 162*, 107131.

Mouabbi, S., J.-P. Renne, and A. Tschopp (2025). Inflation and growth risk: Balancing the scales with surveys. Working paper, Banque de France.

Renne, J.-P., S. Guilloux-Nefussi, M. Marx, and S. Mouabbi (2026). How likely is an inflation disaster? a comment. Manuscript.

Tay, A. and K. F. Wallis (2000). Density forecasting: A survey. Econometric Society World Congress 2000 Contributed Papers 0370, Econometric Society.

Wachter, J. A. (2013). Can time-varying risk of rare disasters explain aggregate stock market volatility? *Journal of Finance 68*(3), 987–1035.

Appendix A. Properties of mixtures of Gaussian distributions

The proofs of the results in this appendix are reported in Online Appendix A.

Definition 1 (Gaussian mixture distribution)**.** *We say that the random variable X follows a Gaussian mixture distribution, which we denote $X \sim \mathcal{G}(\boldsymbol{\theta})$, with $\boldsymbol{\theta} = (\boldsymbol{\mu}_X, \boldsymbol{\sigma}_X, \boldsymbol{\lambda}_X)'$, where $\boldsymbol{\mu}_X$, $\boldsymbol{\sigma}_X$, and $\boldsymbol{\lambda}_X$ are J-dimensional vectors (and the components of $\boldsymbol{\lambda}_X$ are non-negative and sum to one), if*

$$X = \mathbf{B}'\boldsymbol{Y},$$

where $\boldsymbol{B}$ and $\boldsymbol{Y}$ are independent, with

$$\boldsymbol{Y} \sim \mathcal{N}\big(\boldsymbol{\mu}_X, \operatorname{diag}(\boldsymbol{\sigma}_X^2)\big), \quad \textit{and} \quad \mathbb{P}(\mathbf{B} = \mathbf{e}_j) = \lambda_{X,j}, \quad j = 1, \ldots, J,$$

where $\{\mathbf{e}_1, \ldots, \mathbf{e}_J\}$ is the canonical basis of $\mathbb{R}^J$.

Lemma 1 (C.d.f. of a mixture of Gaussian distributions)**.** *Let $X \sim \mathcal{G}(\boldsymbol{\theta})$, with $\boldsymbol{\theta} = (\boldsymbol{\mu}_X, \boldsymbol{\sigma}_X, \boldsymbol{\lambda}_X)'$, where $\boldsymbol{\mu}_X = (\mu_1, \ldots, \mu_J)'$, $\boldsymbol{\sigma}_X = (\sigma_1, \ldots, \sigma_J)'$, and $\boldsymbol{\lambda}_X = (\lambda_1, \ldots, \lambda_J)'$. Then, for any $a \in \mathbb{R}$,*

$$\mathbb{P}(X < a) = h(\boldsymbol{\theta}, a),$$

where

$$h(\boldsymbol{\theta}, a) = \sum_{j=1}^{J} \lambda_j \Phi\left(\frac{a - \mu_j}{\sigma_j}\right). \tag{A.1}$$

Lemma 2 (Laplace transform of a mixture of Gaussian distributions)**.** *Let $X \sim \mathcal{G}(\boldsymbol{\theta})$, with $\boldsymbol{\theta} = (\boldsymbol{\mu}_X, \boldsymbol{\sigma}_X, \boldsymbol{\lambda}_X)'$, where $\boldsymbol{\mu}_X = (\mu_1, \ldots, \mu_J)'$, $\boldsymbol{\sigma}_X = (\sigma_1, \ldots, \sigma_J)'$, and $\boldsymbol{\lambda}_X = (\lambda_1, \ldots, \lambda_J)'$. Then, for any $u \in \mathbb{R}$,*

$$\mathbb{E}[\exp(uX)] = \psi(u, \boldsymbol{\theta}),$$

where

$$\psi(u, \boldsymbol{\theta}) = \sum_{j=1}^{J} \lambda_j \exp\left(u\mu_j + \frac{u^2\sigma_j^2}{2}\right). \tag{A.2}$$

The next lemma is used to compute the truncated Laplace transform of a Gaussian mixture distribution:

Lemma 3. *If $X \sim \mathcal{N}(\mu, \sigma^2)$, then*

$$\mathbb{E}\left[\exp(X)\mathbb{1}_{\{X<a\}}\right] = g(\mu, \sigma, a), \tag{A.3}$$

where

$$g(\mu, \sigma, a) = \exp\left(\mu + \frac{\sigma^2}{2}\right) \Phi\left(\frac{a - \mu - \sigma^2}{\sigma}\right), \tag{A.4}$$

where Φ is the c.d.f. of $\mathcal{N}(0, 1)$.

Lemma 4 (Truncated Laplace transform of a mixture of Gaussian distributions)**.** *Let $u > 0$ and $X \sim \mathcal{G}(\boldsymbol{\theta})$, with $\boldsymbol{\theta} = (\boldsymbol{\mu}_X, \boldsymbol{\sigma}_X, \boldsymbol{\lambda}_X)'$, where $\boldsymbol{\mu}_X = (\mu_1, \ldots, \mu_J)'$, $\boldsymbol{\sigma}_X = (\sigma_1, \ldots, \sigma_J)'$, and $\boldsymbol{\lambda}_X = (\lambda_1, \ldots, \lambda_J)'$. Then*

$$\mathbb{E}\left[\exp(uX)\mathbb{1}_{\{X<a\}}\right] = H(\boldsymbol{\theta}, u, a),$$

where

$$H(\boldsymbol{\theta}, u, a) = \sum_{j=1}^{J} \lambda_j \, g(u\mu_j, u\sigma_j, ua), \tag{A.5}$$

and the function g is defined in (A.4)*.*

Lemma 5 (Bivariate c.d.f. of a mixture of Gaussian distributions)**.** *Let $(X, Y)'$ follow a bivariate Gaussian mixture distribution, that is*

$$(X, Y)' = \mathbf{B}'\mathbf{Z},$$

where $\mathbf{B}$ takes values in $\{\mathbf{e}_1, \ldots, \mathbf{e}_J\}$, is independent of $\boldsymbol{Z} = \left((X_1, Y_1)', \ldots, (X_J, Y_J)'\right)'$, and

$$(X_j, Y_j)' \sim \mathcal{N}(\boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j), \quad j = 1, \ldots, J,$$

with $\mathbb{P}(\mathbf{B} = \mathbf{e}_j) = \lambda_j$. Let $\boldsymbol{\theta} = \{\boldsymbol{\mu}_X, \boldsymbol{\sigma}_X, \boldsymbol{\lambda}_X\}$, where $\boldsymbol{\mu}_X = \{\boldsymbol{\mu}_1, \ldots, \boldsymbol{\mu}_J\}$, $\boldsymbol{\sigma}_X = \{\boldsymbol{\Sigma}_1, \ldots, \boldsymbol{\Sigma}_J\}$, and $\boldsymbol{\lambda}_X = \{\lambda_1, \ldots, \lambda_J\}$. Then, for any $(u, v) \in \mathbb{R}^2$,

$$\mathbb{P}(X < u, Y < v) = \sum_{j=1}^{J} \lambda_j \, \Phi_2\big(u, v; \boldsymbol{\mu}_j, \boldsymbol{\Sigma}_j\big),$$

where $\Phi_2(\cdot, \cdot; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the c.d.f. of a bivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$.

Appendix B. Data Construction and Estimation Details

This appendix provides additional details on the data construction and estimation procedure discussed in Section 4.

Appendix B.1. SPF histograms

The inflation and GDP data are extracted from the ECB SPF individual-forecast files. For each survey round, we retain the one-year-ahead histogram forecast and aggregate probabilities across respondents to obtain a cross-sectional average histogram. The target period associated with the one-year-ahead horizon depends on the survey round and on the variable considered, so the preprocessing code reconstructs the relevant target period separately for inflation and GDP growth.

For inflation, we also extract the individual one-year point forecasts from the same ECB files and average them across respondents. This series is read directly by the estimation code from the original SPF files, rather than being passed through the MATLAB preprocessing.

A practical difficulty is that the SPF bin structure changes over time, especially around periods of very low inflation or large GDP contractions. We therefore rebuild the support date by date, using the relevant lower and upper tail bins and the appropriate interior bin widths for each vintage. After harmonizing the support, each histogram is transformed into observed cumulative probabilities evaluated at the published bin thresholds. These cumulative probabilities are the survey probabilities targeted by the estimation.

Appendix B.2. Inflation derivative data

The market data are read from Bloomberg-based daily files containing euro-area inflation swap rates, zero-coupon floor prices, and zero-coupon cap prices. The empirical analysis focuses on one-year maturity instruments. The floor strike grid is $\{-1, 0, 1, 2\}$ percent and the cap strike grid is $\{2, 3, 4, 5\}$ percent. All rates and prices are converted to decimal units.

To reduce high-frequency noise, raw daily observations are replaced by a backward-looking 22-trading-day moving average. The resulting smoothed series contain the inflation swap, the nominal swap rate used for discounting, and the floor and cap price schedules. Before using the option prices, the option cross section is adjusted when needed to respect the monotonicity restrictions implied by static no-arbitrage across strikes.

Appendix B.3. Export to the estimation code

The preprocessing produces a market-data panel, an inflation SPF panel, and a GDP SPF panel. Dates are harmonized across sources, while the SPF histograms are stored in long format because the number of bin thresholds varies across survey rounds. The inflation SPF panel also contains the cross-sectional average point forecast. The resulting estimation dataset contains, for each survey date, the observed cumulative probabilities, the corresponding bin thresholds, the inflation swap rate, the nominal rate used for discounting, and the floor and cap price schedules.

Appendix B.4. Recursive estimation

The baseline Epstein-Zin model is estimated recursively over the common sample. At each date, the objective combines discrepancies between model-implied and observed inflation SPF cumulative probabilities, GDP SPF cumulative probabilities, floor prices, and cap prices. The objective also includes a penalty on excessive maximum Sharpe ratios. The two regime-specific inflation means are jointly backed out so that the model exactly matches both the observed inflation swap and the SPF point forecast.

To stabilize the nonlinear estimation, the criterion includes smoothing penalties on deviations of the current parameter vector from the previous date's estimate. A separate penalty discourages abrupt movements in the baseline-regime inflation mean implied by the joint normalization. The Epstein-Zin calibration fixes β and ρ, and estimates the remaining parameters within pre-specified bounds, including the risk-aversion parameter γ. In the baseline calibration, the maximum Sharpe-ratio penalty is activated above a value of 4, and the

upper bound on γ is set to 40. The optimization uses recursive starting values and retains the solution associated with the lowest objective value. The estimated parameters are then used to compute the model-implied survey c.d.f.s, option prices, joint physical and risk-neutral distributions, risk premia, and tail-environment probabilities reported in the main text.

Appendix C. Joint normalization of inflation means

This appendix describes the temporary implementation that exactly matches, at each date, the one-year SPF inflation point forecast and the one-year inflation swap rate. Let $\bar{\pi}_t^{\text{SPF}}$ denote the cross-sectional average SPF point forecast and suppose that the current regime is $i = 1$, as in the empirical implementation. The physical mean implied by the two-regime mixture is

$$\mathbb{E}_t^{\mathbb{P}}(\pi_{t+1}) = p_{11}\mu_{\pi,1} + p_{12}\mu_{\pi,2}. \tag{C.1}$$

We impose the physical-mean restriction exactly. Hence, for any candidate value of the baseline-regime mean, the second-regime mean is given by

$$\mu_{\pi,2} = \frac{\bar{\pi}_t^{\text{SPF}} - p_{11}\mu_{\pi,1}}{p_{12}}. \tag{C.2}$$

The swap restriction then reduces to a scalar nonlinear equation. We substitute (C.2) into the swap equation (16). The numerator of that equation does not depend on either inflation mean because $\boldsymbol{\alpha} + \mathbf{e}_\pi = (-\gamma, 0)'$. The continuation-value vector is also unaffected by the inflation means, since the fixed point in Proposition 1 depends on the log-Laplace transform evaluated at $(1-\gamma, 0)'$. Consequently, conditional on all other parameters, the swap restriction is a scalar nonlinear equation in $\mu_{\pi,1}$. We solve this equation numerically and recover $\mu_{\pi,2}$ from (C.2).

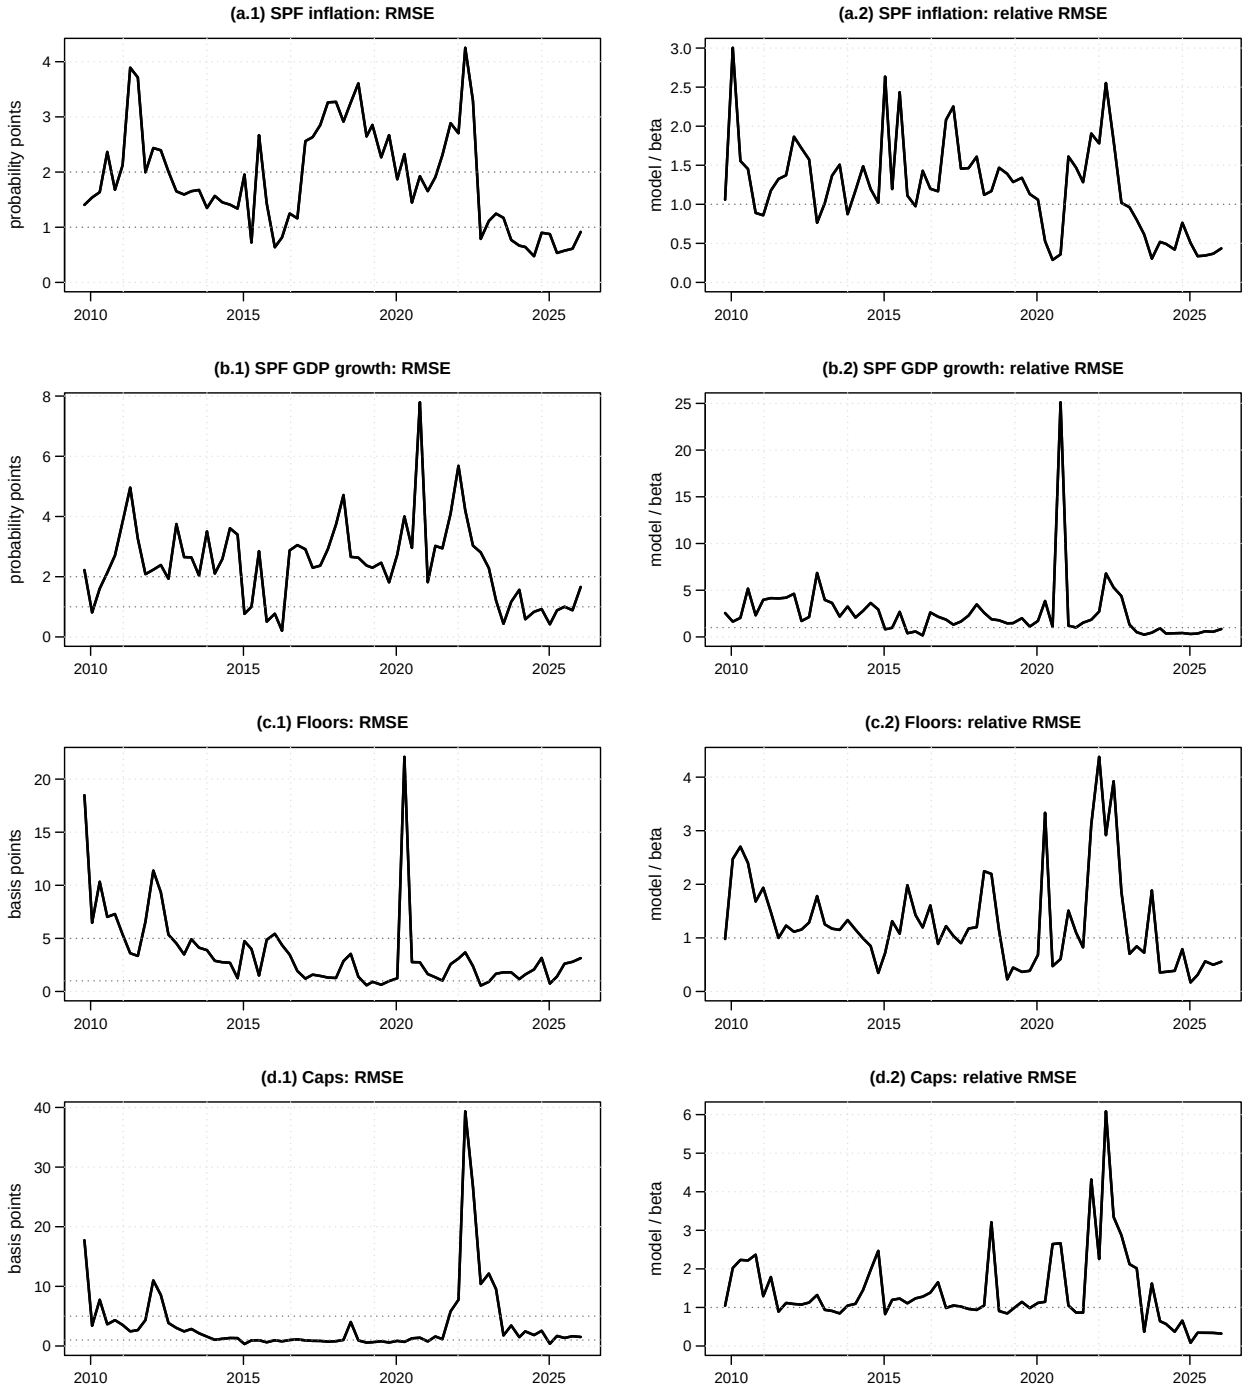
The two inflation means are therefore removed from the date-by-date optimizer. By construction, the fitted physical mean equals $\bar{\pi}_t^{\text{SPF}}$ and the fitted swap rate equals the observed swap rate. If the scalar equation admits more than one root, we retain the solution whose pair of regime-specific inflation means is closest to the SPF point forecast in squared Euclidean distance.

Appendix D. Model Fit Diagnostics

This appendix reports time-series diagnostics for the fit of the baseline Epstein-Zin specification. For each survey date, we compute root mean squared errors (RMSEs) separately for the four blocks of observations entering the estimation: inflation SPF cumulative probabilities, GDP-growth SPF cumulative probabilities, inflation floor prices, and inflation cap prices. The SPF errors are reported in probability points, while option-pricing errors are reported in basis points. The inflation point forecast and swap rate are not included because the two inflation means are jointly backed out to match them exactly.

To provide a benchmark, the figure also compares the model RMSE with the RMSE obtained from the beta-smoothed distributions used in the fit figures. For survey data, the benchmark is the beta c.d.f. evaluated at the SPF bin thresholds. For option prices, it is the price schedule implied by the beta-smoothed inflation distribution. The relative RMSE is the ratio of the model RMSE to the beta-benchmark RMSE, computed within each block and date; values below one indicate that the structural model fits better than the beta benchmark for that block and date.

Figure D.10: **Fit diagnostics for the baseline Epstein-Zin model.** The left column reports RMSEs for inflation SPF cumulative probabilities, GDP-growth SPF cumulative probabilities, inflation floor prices, and inflation cap prices. SPF errors are expressed in probability points; option-pricing errors are expressed in basis points. The right column reports the ratio of the model RMSE to the RMSE obtained from the beta-smoothed benchmark used for comparison in the fit figures. A value below one means that the structural model fits the corresponding block better than the beta benchmark on that date. The beta benchmark is not used in the structural estimation.



— Supplementary Material —

Reading Inflation Tails

Sophie GUILLOUX-NEFUSSI, Magali MARX, Sarah MOUABBI, and Jean-Paul RENNE

Online Appendix A. Proofs of Gaussian-mixture properties

OA.A.1. Proof of Lemma 1

Proof. Recall that $X = \mathbf{B}'\boldsymbol{Y}$, where $\mathbf{B}$ takes values in $\{\mathbf{e}_1, \ldots, \mathbf{e}_J\}$ and is independent of $\boldsymbol{Y} = (Y_1, \ldots, Y_J)'$, with $Y_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$.

We write

$$\mathbb{P}(X < a) = \mathbb{E}\big[\mathbf{1}_{\{X<a\}}\big] = \mathbb{E}\big[\mathbf{1}_{\{\mathbf{B}'\boldsymbol{Y}<a\}}\big].$$

Conditioning on $\mathbf{B}$ yields

$$\mathbb{P}(X < a) = \sum_{j=1}^{J} \mathbb{E}\Big[\mathbf{1}_{\{\mathbf{B}'\boldsymbol{Y}<a\}} \,\Big|\, \mathbf{B} = \mathbf{e}_j\Big]\, \mathbb{P}(\mathbf{B} = \mathbf{e}_j).$$

Since $\mathbf{B} = \mathbf{e}_j$ implies $\mathbf{B}'\boldsymbol{Y} = Y_j$, and using independence, we obtain

$$\mathbb{P}(X < a) = \sum_{j=1}^{J} \mathbb{P}(Y_j < a)\, \lambda_j.$$

Finally, since $Y_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$,

$$\mathbb{P}(Y_j < a) = \Phi\left(\frac{a - \mu_j}{\sigma_j}\right),$$

which gives the result. □

OA.A.2. Proof of Lemma 2

Proof. We have

$$\mathbb{E}[\exp(uX)] = \mathbb{E}\big[\exp(u\,\mathbf{B}'\boldsymbol{Y})\big].$$

Conditioning on $\mathbf{B}$ yields

$$\mathbb{E}[\exp(uX)] = \sum_{j=1}^{J} \mathbb{E}\Big[\exp(u\,\mathbf{B}'\boldsymbol{Y}) \mid \mathbf{B} = \mathbf{e}_j\Big]\,\mathbb{P}(\mathbf{B} = \mathbf{e}_j).$$

Since $\mathbf{B} = \mathbf{e}_j$ implies $\mathbf{B}'\boldsymbol{Y} = Y_j$, and using independence, we obtain

$$\mathbb{E}[\exp(uX)] = \sum_{j=1}^{J} \mathbb{E}[\exp(uY_j)]\,\lambda_j.$$

Finally, since $Y_j \sim \mathcal{N}(\mu_j, \sigma_j^2)$,

$$\mathbb{E}[\exp(uY_j)] = \exp\left(u\mu_j + \frac{u^2\sigma_j^2}{2}\right),$$

which gives the result. □

OA.A.3. Proof of Lemma 4

Proof. We have

$$\mathbb{E}\left[\exp(uX)\mathbb{1}_{\{X<a\}}\right] = \mathbb{E}\left[\exp\left(u\,\mathbf{B}'\boldsymbol{Y}\right)\mathbb{1}_{\{\mathbf{B}'\boldsymbol{Y}<a\}}\right].$$

Conditioning on $\mathbf{B}$ yields

$$\mathbb{E}\left[\exp(uX)\mathbb{1}_{\{X<a\}}\right] = \sum_{j=1}^{J} \mathbb{E}\Big[\exp\left(u\,\mathbf{B}'\boldsymbol{Y}\right)\mathbb{1}_{\{\mathbf{B}'\boldsymbol{Y}<a\}} \mid \mathbf{B} = \mathbf{e}_j\Big]\,\mathbb{P}(\mathbf{B} = \mathbf{e}_j).$$

Since $\mathbf{B} = \mathbf{e}_j$ implies $\mathbf{B}'\boldsymbol{Y} = Y_j$, and using independence, we obtain

$$\mathbb{E}\left[\exp(uX)\mathbb{1}_{\{X<a\}}\right] = \sum_{j=1}^{J} \lambda_j\,\mathbb{E}\left[\exp(uY_j)\mathbb{1}_{\{Y_j<a\}}\right].$$

Since $u > 0$, we can equivalently write $\{Y_j < a\} = \{uY_j < ua\}$, hence

$$\mathbb{E}\left[\exp(uX)\mathbb{1}_{\{X<a\}}\right] = \sum_{j=1}^{J} \lambda_j\,\mathbb{E}\left[\exp(uY_j)\mathbb{1}_{\{uY_j<ua\}}\right].$$

Applying Lemma 3 to each Gaussian component yields

$$\mathbb{E}\left[\exp(uY_j)\mathbb{1}_{\{uY_j<ua\}}\right] = g(u\mu_j, u\sigma_j, ua),$$

which gives the result. □

OA.A.4. Proof of Lemma 5

Proof. $\mathbf{B} = \mathbf{e}_j$ implies $(X, Y)' = (X_j, Y_j)'$. Hence

$$\mathbb{P}(X < u, Y < v) = \mathbb{E}\left[\mathbf{1}_{\{X<u, Y<v\}}\right] = \mathbb{E}\left[\mathbf{1}_{\{\mathbf{B}'\mathbf{Z}<(u,v)'\}}\right].$$

Conditioning on $\mathbf{B}$ yields

$$\mathbb{P}(X < u, Y < v) = \sum_{j=1}^{J} \mathbb{E}\left[\mathbf{1}_{\{\mathbf{B}'\mathbf{Z}<(u,v)'\}} \mid \mathbf{B} = \mathbf{e}_j\right] \mathbb{P}(\mathbf{B} = \mathbf{e}_j).$$

Using independence, we obtain

$$\mathbb{P}(X < u, Y < v) = \sum_{j=1}^{J} \lambda_j \, \mathbb{P}(X_j < u, Y_j < v),$$

which gives the result. □

Online Appendix B. Proofs

OA.B.1. Proof of Proposition 1

Posit a value function of the form $U_t = (\mathbf{u}'\mathbf{z}_t)C_t$, and substitute it into the right-hand side of (4). Using

$$C_{t+1} = C_t \exp(\mathbf{x}'_{c,t+1}\mathbf{z}_{t+1}),$$

we obtain

$$U_t = \left[(1-\beta)\, C_t^{1-\frac{1}{\rho}} + \beta \left(\mathbb{E}_t\left[\left\{C_t \exp(\mathbf{x}'_{c,t+1}\mathbf{z}_{t+1})(\mathbf{u}'\mathbf{z}_{t+1})\right\}^{1-\gamma}\right]\right)^{\frac{1-\frac{1}{\rho}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\rho}}}.$$

Factoring out $C_t^{1-1/\rho}$ and simplifying yields

$$\mathbf{u}'\mathbf{z}_t = \left[(1-\beta) + \beta \left(\mathbb{E}_t\left[\exp\big((1-\gamma)\mathbf{x}'_{c,t+1}\mathbf{z}_{t+1}\big)(\mathbf{u}'\mathbf{z}_{t+1})^{1-\gamma}\right]\right)^{\frac{1-\frac{1}{\rho}}{1-\gamma}}\right]^{\frac{1}{1-\frac{1}{\rho}}}.$$

Using the Markov property of $\mathbf{z}_t$ and the definition of $\mathbf{g}$,

$$\mathbb{E}_t\left[\exp\big\{(1-\gamma)\mathbf{x}'_{c,t+1}\mathbf{z}_{t+1}\big\}(\mathbf{u}'\mathbf{z}_{t+1})^{1-\gamma}\right] = \left(\mathbf{g} \odot \mathbf{u}^{1-\gamma}\right)'\mathbf{P}'\mathbf{z}_t.$$

Since the equality must hold for all realizations of $\mathbf{z}_t$, (5) follows.

OA.B.2. Proof of Proposition 2

Given the value function, the real SDF implied by Epstein-Zin preferences is

$$\mathcal{M}_{t,t+1}^{(r)} = \beta \left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\rho}} \left(\frac{U_{t+1}}{\left(\mathbb{E}_t[U_{t+1}^{1-\gamma}]\right)^{\frac{1}{1-\gamma}}}\right)^{\frac{1}{\rho}-\gamma}.$$

Substituting the consumption-growth specification and the value function gives

$$\mathcal{M}_{t,t+1}^{(r)} = \beta \exp\left(-\gamma\, \mathbf{x}_{c,t+1}'\mathbf{z}_{t+1}\right) \left(\frac{\mathbf{u}'\mathbf{z}_{t+1}}{\left(\mathbb{E}_t\left[\exp\big((1-\gamma)\mathbf{x}_{c,t+1}'\mathbf{z}_{t+1}\big)(\mathbf{u}'\mathbf{z}_{t+1})^{1-\gamma}\right]\right)^{\frac{1}{1-\gamma}}}\right)^{\frac{1}{\rho}-\gamma}.$$

The nominal SDF is $\mathcal{M}_{t,t+1}^{(n)} = \mathcal{M}_{t,t+1}^{(r)} \exp(-\pi_{t+1})$. Hence, up to a normalization known at date t, the nominal change of measure tilts next-period outcomes by

$$\exp\left(\boldsymbol{\alpha}'\mathbf{w}_{t+1} + \theta \log(\mathbf{u})'\mathbf{z}_{t+1}\right), \qquad \boldsymbol{\alpha} = (-\gamma, -1)', \quad \theta = 1/\rho - \gamma.$$

It follows that

$$\begin{aligned}\psi_t^{\mathbb{Q}^{(n)}}(\mathbf{v}, \mathbf{s}) &= \mathbb{E}_t\left[\exp(\mathbf{v}'\mathbf{w}_{t+1} + \mathbf{s}'\mathbf{z}_{t+1}) \frac{\mathcal{M}_{t,t+1}^{(n)}}{\mathbb{E}_t(\mathcal{M}_{t,t+1}^{(n)})}\right] \\ &= \frac{\exp\left[\boldsymbol{\varphi}(\mathbf{v}+\boldsymbol{\alpha}) + \theta \log(\mathbf{u}) + \mathbf{s}\right]' \mathbf{P}'\mathbf{z}_t}{\exp\left[\boldsymbol{\varphi}(\boldsymbol{\alpha}) + \theta \log(\mathbf{u})\right]' \mathbf{P}'\mathbf{z}_t}.\end{aligned}$$

Using the definitions of $\mathbf{m}^{(n)}$ and $\boldsymbol{\varphi}^{\mathbb{Q}^{(n)}}$ gives

$$\psi_t^{\mathbb{Q}^{(n)}}(\mathbf{v}, \mathbf{s}) = \exp\left[\boldsymbol{\varphi}^{\mathbb{Q}^{(n)}}(\mathbf{v}) + \mathbf{s}\right]' \operatorname{diag}\left(\mathbf{m}^{(n)}\right) \mathbf{P}' \operatorname{diag}\left(\frac{1}{\mathbf{P}\mathbf{m}^{(n)}}\right) \mathbf{z}_t,$$

which is (9) with $\mathbf{Q}^{(n)}$ defined in (10).

OA.B.3. Proof of Corollary 1

Under Assumption 3,

$$\varphi_i(\mathbf{v}) = \mathbf{v}'\boldsymbol{\mu}_i + \frac{1}{2}\mathbf{v}'\boldsymbol{\Sigma}_i\mathbf{v}.$$

Using (11),

$$\varphi_i^{\mathbb{Q}^{(n)}}(\mathbf{v}) = \varphi_i(\mathbf{v}+\boldsymbol{\alpha}) - \varphi_i(\boldsymbol{\alpha}) = \mathbf{v}'(\boldsymbol{\mu}_i + \boldsymbol{\Sigma}_i\boldsymbol{\alpha}) + \frac{1}{2}\mathbf{v}'\boldsymbol{\Sigma}_i\mathbf{v}.$$

Thus the covariance matrix is unchanged and the conditional mean is shifted to $\boldsymbol{\mu}_i + \boldsymbol{\Sigma}_i\boldsymbol{\alpha}$.

OA.B.4. Proof of Proposition 3

At date t, the fixed and floating legs of the inflation swap have the same value:

$$\mathbb{E}_t\left(\mathcal{M}_{t,t+1}^{(n)} \exp[\pi_{t+1}]\right) = \mathbb{E}_t\left(\mathcal{M}_{t,t+1}^{(n)}\right) \exp(\pi_{t,1}^S).$$

Therefore $\pi_{t,1}^S = \log \mathbb{E}_t^{\mathbb{Q}^{(n)}}[\exp(\pi_{t+1})]$. Applying Proposition 2 with $\mathbf{v} = \mathbf{e}_\pi$ and $\mathbf{s} = \mathbf{0}$ gives (15). Expanding the same expression conditional on $\mathbf{z}_t = \mathbf{e}_i$ gives (16).

OA.B.5. Proof of Proposition 4

The floor price is

$$\begin{aligned}
\text{Floor}_{\pi,t}(K) &= e^{-r_t}\mathbb{E}_t^{\mathbb{Q}^{(n)}}\left([\exp(K) - \exp(\pi_{t+1})]\mathbb{1}_{\{\pi_{t+1}<K\}}\right) \\
&= e^{-r_t+K}\mathbb{E}_t^{\mathbb{Q}^{(n)}}\left(\mathbb{1}_{\{\pi_{t+1}<K\}}\right) - e^{-r_t}\mathbb{E}_t^{\mathbb{Q}^{(n)}}\left(\exp(\pi_{t+1})\mathbb{1}_{\{\pi_{t+1}<K\}}\right).
\end{aligned}$$

The first term is the mixture c.d.f. $h(\boldsymbol{\theta}_{\pi,t}^{\mathbb{Q}^{(n)}}, K)$ and the second is the truncated mixture Laplace transform $H(\boldsymbol{\theta}_{\pi,t}^{\mathbb{Q}^{(n)}}, 1, K)$. The cap formula follows from put-call parity:

$$\text{Cap}_{\pi,t}(K) - \text{Floor}_{\pi,t}(K) = e^{-r_t}\left(\mathbb{E}_t^{\mathbb{Q}^{(n)}}[e^{\pi_{t+1}}] - e^K\right) = e^{-r_t}\left(e^{\pi_{t,1}^S} - e^K\right).$$

Online Appendix C. CRRA Mixture Illustration

This appendix collects the assumptions and formulas underlying the static CRRA illustration in Section 2.3 of the paper. The purpose of this benchmark is pedagogical: it isolates the mechanism through which the pricing kernel reweights the joint distribution of inflation and consumption growth. The empirical baseline in the paper is the Epstein-Zin Markov model of Section 3.

Because CRRA preferences are time-additive and non-recursive, the one-period stochastic discount factor depends only on consumption growth between dates 0 and 1. The illustration can therefore be written as a two-date economy, with no continuation-value state. This is precisely the feature that makes the benchmark analytically transparent, and also the reason why the Epstein-Zin baseline is needed in the main empirical analysis.

OA.C.1. Preferences and risk-neutral dynamics

Consider a two-period economy populated by identical agents with constant relative risk aversion (CRRA) preferences. The agents' date-0 utility is

$$U_0 = \frac{C_0^{1-\gamma}}{1-\gamma} + \mathbb{E}_0\left(\frac{C_1^{1-\gamma}}{1-\gamma}\right),$$

where C_i denotes consumption at date i, $\mathbb{E}_0$ is the expectation conditional on date-0 information, and γ is the coefficient of relative risk aversion.

Without loss of generality in this two-period environment, we omit the pure rate of preference for the present. The real stochastic discount factor between dates 0 and 1 is

$$\mathcal{M}_1^{(r)} = \exp(-\gamma\delta\Delta c_1),$$

where δ denotes the length of the period in years, and $\Delta c_1 = (1/\delta)\log(C_1/C_0)$ is annualized log consumption growth. The nominal SDF is

$$\mathcal{M}_1^{(n)} = \mathcal{M}_1^{(r)} \exp(-\delta\pi_1), \tag{OA.C.1}$$

where π_1 is inflation between dates 0 and 1.

OA.C.2. Consumption and inflation distributions

There are three sources of uncertainty: ε_1^c, ε_1^π, and B. The variable B is Bernoulli with parameter λ, and ε_1^c and ε_1^π are independent standard normal shocks, independent of B. Consumption growth and inflation are

$$\begin{aligned} \Delta c_1 &= (1-B)(\mu_c + \sigma_c\varepsilon_1^c) + B(\tilde{\mu}_c + \tilde{\sigma}_c\varepsilon_1^c), & \text{(OA.C.2)} \\ \pi_1 &= (1-B)(\mu_\pi + \sigma_{\pi,c}\varepsilon_1^c + \sigma_{\pi,\pi}\varepsilon_1^\pi) + B(\tilde{\mu}_\pi + \tilde{\sigma}_{\pi,c}\varepsilon_1^c + \tilde{\sigma}_{\pi,\pi}\varepsilon_1^\pi). & \text{(OA.C.3)} \end{aligned}$$

Let $\mathbf{w}_1 = (\Delta c_1, \pi_1)'$ and $\boldsymbol{\varepsilon}_1 = (\varepsilon_1^c, \varepsilon_1^\pi)'$. Then

$$\mathbf{w}_1 = (1-B)\left(\boldsymbol{\mu} + \boldsymbol{\Sigma}^{1/2}\boldsymbol{\varepsilon}_1\right) + B\left(\tilde{\boldsymbol{\mu}} + \tilde{\boldsymbol{\Sigma}}^{1/2}\boldsymbol{\varepsilon}_1\right), \tag{OA.C.4}$$

with

$$\boldsymbol{\mu} = \begin{pmatrix}\mu_c \\ \mu_\pi\end{pmatrix}, \quad \tilde{\boldsymbol{\mu}} = \begin{pmatrix}\tilde{\mu}_c \\ \tilde{\mu}_\pi\end{pmatrix}, \quad \boldsymbol{\Sigma}^{1/2} = \begin{pmatrix}\sigma_c & 0 \\ \sigma_{\pi,c} & \sigma_{\pi,\pi}\end{pmatrix}, \quad \tilde{\boldsymbol{\Sigma}}^{1/2} = \begin{pmatrix}\tilde{\sigma}_c & 0 \\ \tilde{\sigma}_{\pi,c} & \tilde{\sigma}_{\pi,\pi}\end{pmatrix}. \tag{OA.C.5}$$

OA.C.3. Risk-neutral transformation

Let

$$\boldsymbol{\alpha} = (-\gamma\delta, -\delta)'.$$

Under the nominal risk-neutral measure, the distribution of $\mathbf{w}_1$ remains a two-component Gaussian mixture:

$$\mathbf{w}_1 = (1-B^{\mathbb{Q}})\left(\boldsymbol{\mu}^{\mathbb{Q}} + \boldsymbol{\Sigma}^{1/2}\boldsymbol{\varepsilon}_1^{\mathbb{Q}}\right) + B^{\mathbb{Q}}\left(\tilde{\boldsymbol{\mu}}^{\mathbb{Q}} + \tilde{\boldsymbol{\Sigma}}^{1/2}\boldsymbol{\varepsilon}_1^{\mathbb{Q}}\right),$$

where $B^{\mathbb{Q}} \sim \mathcal{B}(\lambda^{\mathbb{Q}})$, $\boldsymbol{\varepsilon}_1^{\mathbb{Q}} \sim \mathcal{N}(\mathbf{0}, \mathbf{Id})$, and

$$\lambda^{\mathbb{Q}} = \frac{\lambda}{(1-\lambda)\exp\left[\boldsymbol{\alpha}'(\boldsymbol{\mu} - \tilde{\boldsymbol{\mu}}) + \frac{1}{2}\boldsymbol{\alpha}'(\boldsymbol{\Sigma} - \tilde{\boldsymbol{\Sigma}})\boldsymbol{\alpha}\right] + \lambda}. \tag{OA.C.6}$$

The conditional means are tilted according to

$$\boldsymbol{\mu}^{\mathbb{Q}} = \boldsymbol{\mu} + \boldsymbol{\Sigma}\boldsymbol{\alpha}, \qquad \tilde{\boldsymbol{\mu}}^{\mathbb{Q}} = \tilde{\boldsymbol{\mu}} + \tilde{\boldsymbol{\Sigma}}\boldsymbol{\alpha}. \tag{OA.C.7}$$

OA.C.4. Pricing formulas

For a univariate two-component Gaussian mixture with parameter vector $\boldsymbol{\theta} = (\mu_1, \sigma_1, \mu_2, \sigma_2, \lambda)'$, define

$$h(\boldsymbol{\theta}, a) = (1-\lambda)\Phi\left(\frac{a-\mu_1}{\sigma_1}\right) + \lambda\Phi\left(\frac{a-\mu_2}{\sigma_2}\right), \tag{OA.C.8}$$

and

$$\psi(u, \boldsymbol{\theta}) = (1-\lambda)\exp\left(u\mu_1 + \frac{u^2\sigma_1^2}{2}\right) + \lambda\exp\left(u\mu_2 + \frac{u^2\sigma_2^2}{2}\right). \tag{OA.C.9}$$

Also define

$$H(\boldsymbol{\theta}, u, a) = (1-\lambda)g(u\mu_1, u\sigma_1, ua) + \lambda g(u\mu_2, u\sigma_2, ua), \tag{OA.C.10}$$

where

$$g(\mu, \sigma, a) = \exp\left(\mu + \frac{\sigma^2}{2}\right)\Phi\left(\frac{a-\mu-\sigma^2}{\sigma}\right).$$

Let $\boldsymbol{\theta}_\pi^{\mathbb{Q}}$ denote the risk-neutral mixture parameters of inflation. The inflation swap rate is

$$\pi_0^S = \frac{1}{\delta}\log\psi(\delta, \boldsymbol{\theta}_\pi^{\mathbb{Q}}). \tag{OA.C.11}$$

The prices of an inflation floor and cap with strike K are

$$\begin{aligned} \text{Floor}_\pi(K) &= \exp(-\delta r + \delta K)h(\boldsymbol{\theta}_\pi^{\mathbb{Q}}, K) - \exp(-\delta r)H(\boldsymbol{\theta}_\pi^{\mathbb{Q}}, \delta, K), & \text{(OA.C.12)} \\ \text{Cap}_\pi(K) &= \exp(-\delta r)\left[\psi(\delta, \boldsymbol{\theta}_\pi^{\mathbb{Q}}) - \exp(\delta K)\right] + \text{Floor}_\pi(K). & \text{(OA.C.13)} \end{aligned}$$

OA.C.5. Maximum Sharpe ratio

Let $\psi_{\mathbf{w}}$ denote the Laplace transform of $\mathbf{w}_1$ under the physical measure. Since the nominal SDF is proportional to $\exp(\boldsymbol{\alpha}'\mathbf{w}_1)$, the Hansen-Jagannathan maximum Sharpe ratio is

$$\max SR = \sqrt{\exp\left[\log\psi_{\mathbf{w}}(2\boldsymbol{\alpha}) - 2\log\psi_{\mathbf{w}}(\boldsymbol{\alpha})\right] - 1}. \tag{OA.C.14}$$

OA.C.6. Swap-implied inflation mean

$$\pi_0^S = \frac{1}{\delta}\log\left((1-\lambda^{\mathbb{Q}})\exp\left(\delta\mu_\pi^{\mathbb{Q}} + \frac{\delta^2(\sigma_{\pi,c}^2 + \sigma_{\pi,\pi}^2)}{2}\right) + \lambda^{\mathbb{Q}}\exp\left(\delta\tilde{\mu}_\pi^{\mathbb{Q}} + \frac{\delta^2(\tilde{\sigma}_{\pi,c}^2 + \tilde{\sigma}_{\pi,\pi}^2)}{2}\right)\right).$$

Using that $\mu_\pi^{\mathbb{Q}} = \mu_\pi - \gamma\delta\sigma_c\sigma_{\pi,c}$, this gives:

$$\begin{aligned}\mu_\pi &= \gamma\delta\sigma_c\sigma_{\pi,c} - \frac{\delta(\sigma_{\pi,c}^2 + \sigma_{\pi,\pi}^2)}{2} \\ &\quad + \frac{1}{\delta}\log\left\{\frac{1}{1-\lambda^{\mathbb{Q}}}\left[\exp(\delta\pi_0^S) - \lambda^{\mathbb{Q}}\exp\left(\delta\tilde{\mu}_\pi^{\mathbb{Q}} + \frac{\delta^2(\tilde{\sigma}_{\pi,c}^2 + \tilde{\sigma}_{\pi,\pi}^2)}{2}\right)\right]\right\}.\end{aligned}$$

A first-order approximation gives

$$\mu_\pi = \gamma\delta\sigma_c\sigma_{\pi,c} + \frac{1}{1-\lambda^{\mathbb{Q}}}\pi_0^S - \frac{\delta(\sigma_{\pi,c}^2 + \sigma_{\pi,\pi}^2)}{2} - \frac{\lambda^{\mathbb{Q}}}{1-\lambda^{\mathbb{Q}}}\left(\tilde{\mu}_\pi^{\mathbb{Q}} + \frac{\delta(\tilde{\sigma}_{\pi,c}^2 + \tilde{\sigma}_{\pi,\pi}^2)}{2}\right)$$

Online Appendix D. Estimated Parameters

This appendix reports the time series of parameters estimated in the baseline Epstein-Zin specification. These parameters are obtained from the recursive estimation described in Appendix B. They are reported for transparency and are not interpreted individually in the main text, where the focus is on the model-implied distributions, risk premia, and tail-environment probabilities.

Figure D.11: **Estimated parameters of the baseline Epstein-Zin model.** The figure reports the date-by-date estimates of the parameters entering the two-regime conditional Gaussian specification and the calibrated risk-aversion parameter. The baseline-regime inflation mean is backed out from the inflation swap rate, as described in Appendix B. The intertemporal elasticity of substitution and subjective discount factor are fixed in the baseline calibration.

