

Measuring market-based core inflation expectations

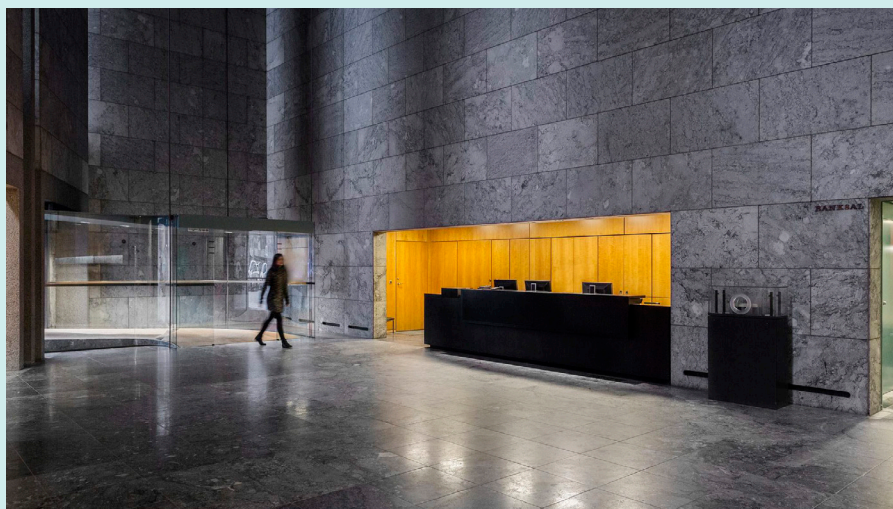
We build a novel term-structure model for pricing synthetic euro area core inflation-linked swaps (ILS). Our approach relies on a term-structure model of traded headline ILS rates, assumed to span core inflation. From these synthetic swaps, we can extract genuine euro area core inflation expectations of market participants.

Written by

Asger Munch Grønlund
Danmarks Nationalbank
amg@nationalbanken.dk

Kasper Jørgensen
European Central Bank
Kasper.Jorgensen@ecb.europa.eu

Fabian Schupp
European Central Bank
Fabian.Schupp@ecb.europa.eu



Keywords

Financial markets

Inflation and price development

Models

Monetary policy

Measuring Market-Based Core Inflation Expectations*

Asger Munch Grønlund[†]

Kasper Jørgensen[‡]

Fabian Schupp[§]

February 23, 2024

Abstract

We build a novel term structure model for pricing synthetic euro area core inflation-linked swaps, a hypothetical swap contract indexed to core inflation. Our approach relies on a term structure model of traded headline inflation-linked swap rates, which we assume span core inflation. The model provides estimates of market-based expectations for core inflation, as well as core inflation risk premia, at daily frequency, whereas core inflation expectations from surveys or macroeconomic projections are typically only available monthly or quarterly. We find that core inflation-linked swap rates are generally less volatile than headline inflation-linked swap rates and that market participants expected core inflation to be substantially more persistent than headline inflation following the 2022 energy price spike. Using an event-study methodology, we also find that monetary policy shocks significantly lower core inflation expectations.

Keywords: Inflation-linked swaps, affine term structure model, inflation expectations

JEL Codes: E31, E44, E52

*The views and opinions in this paper are those of the authors and do not necessarily reflect those of the European Central Bank, the Danish Central Bank or any of the institutions' staff.

[†]amg@nationalbanken.dk, Danmarks Nationalbank.

[‡]kasper.jorgensen@ecb.europa.eu, European Central Bank.

[§]fabian.schupp@ecb.europa.eu, European Central Bank.

Non-technical summary

Central banks are inherently interested in measuring inflation expectations. There are at least two reasons for the focus on inflation expectations: (i) these expectations play a crucial role in determining future actual inflation and (ii) long-term inflation expectations serve as an important indicator of how much trust people have in the central banks' ability to achieve its inflation target. Not all inflation is the same, however. Given the time it takes for monetary policy to filter through the economic system and have macroeconomic effects, policymakers prefer to track *underlying* inflation, which excludes the inflation contributions from the most volatile and typically temporary components of the basket of goods and services. For instance, the European Central Bank's Governing Council has recently emphasized that they will consider the "dynamics of underlying inflation" as a separate factor in making future decisions regarding interest rates.

In this paper, we aim to infer market-based *core* inflation expectations, which are otherwise not directly observable because no financial asset directly tied to core inflation exists. We do so by estimating a model of traded headline inflation-linked swap (ILS) rates, which we use for pricing synthetic euro area core ILS contracts, hypothetical contracts tied to core inflation. Our approach relies on a set of assumptions. First, we assume that HICP inflation is a linear combination of a small set of pricing factors that follow a simple vector-autoregressive process. Second, we assume that core inflation is a linear combination of the same set of pricing factors. We deem this second assumption reasonable because HICP inflation itself is a linear combination of core as well as energy and food inflation. Third, we assume that there are no arbitrage opportunities in the ILS market, such that headline and core ILS rates of any maturity are all linear combinations of the aforementioned pricing factors. Based on these assumptions, we fit the model to a monthly data set of inflation-linked swap rates linked to HICP inflation, realized HICP and core inflation, as well as survey information on expectations of headline and core inflation from the ECB's Survey of Professional Forecasters (SPF). This data is available from June 2005 until July 2023. The exception is survey expectations for core inflation, which are available as of fall 2016 only.

Based on those assumptions and data, our model produces plausible estimates for both headline and core inflation expectations. The expectations for headline as well as core inflation are more volatile at near-term maturities than at long-term maturities. Long-maturity inflation

expectations hover around 2 percent. The level of 2 percent and relatively low volatility of long-term inflation expectations suggests that inflation expectations are firmly anchored at the ECB's 2 percent inflation target.

One key advantage of our core inflation expectations measure is that it is based on financial market data, which means that it can be computed at, essentially, whatever frequency we observe headline ILS rates. Other measures of core inflation expectations, such as from surveys or projections based on macroeconomic models, are available at much lower frequency, usually quarterly or monthly. Therefore, the high-frequency feature of our estimates makes them potentially useful for at least two reasons: (i) it provides policymakers with a real-time read on core inflation expectations, and (ii) it facilitates event-studies that can give valuable information on the causal effects of, for example, monetary policy shocks on core inflation expectations.

To illustrate the usefulness of the high-frequency nature of our measure, we compute daily core inflation expectations and estimate the causal impact of monetary policy shocks by analyzing their changes around specific monetary policy events. We find that a surprise monetary policy tightening significantly lowers core inflation expectations at near-term horizons, whereas the response at longer-term horizons is more muted. This effect does not die out after just a few days, but is relatively persistent, which suggests that the impact on core inflation expectations indeed may have an impact also on wider macroeconomic outcomes.

1 Introduction

In pursuit of fulfilling their price stability mandates, central bank policymakers are acutely interested in inflation expectations. This is the case for at least two reasons: (i) these expectations are viewed as a key determinant of future realized inflation, and (ii) long-run inflation expectations are considered a useful gauge of central bank credibility. But all inflation is not created equal. Because monetary policy works with considerable lags, policymakers like to follow *underlying* inflation, stripping out the inflation contributions from the most volatile and typically transitory components in the consumption basket. The ECB's Governing Council has, for example, in several recent policy statements said that "the dynamics of underlying inflation" will be a separate factor in its future interest rate decisions.

In this paper, we measure market-based euro area *core* inflation expectations. These core expectations are not directly observable, but have to be inferred, because there is no financial instrument directly tied to core inflation. Our approach relies on a novel Gaussian affine term structure model of traded headline inflation-linked swap (ILS) rates, which we use for pricing synthetic euro area core ILS contracts, hypothetical contracts tied to core inflation. The key assumption underlying our approach is that the factors driving observable headline ILS rates span the factors driving core inflation. This assumption appears reasonably uncontroversial, as core inflation is a sub-component of headline inflation, which the observable headline ILS rates are tied to. Our estimates of core ILS rates reflect both market participants' genuine core inflation expectations and a core inflation risk premium, but our model explicitly allows for this decomposition. This means that our results give insights on both market participants' core inflation expectations as well as the perceived risks around the core inflation outlook.

We fit our model to three types of data: euro area ILS rates tied to the Harmonized Index of Consumer Prices (HICP) excluding tobacco with maturities of 1, 2, 3, 5, 7, and 10 years; realized month-on-month headline and core euro area HICP inflation; survey expectations from the Survey of Professional Forecasters (SPF) for both headline and core inflation at horizons 1-, 2-, and 5-years ahead. We use ILS data from June 2005 to July 2023, restricting the remaining data to this sample. The exception is the SPF, which did not ask respondent for their core inflation expectations before the SPF 2016Q4. While we use the other data going back to June 2005, the model does not have enough information to learn about how core inflation is related to the factors driving headline ILS data before the 2016Q4 SPF. For this reason, our core ILS estimates cover the period September 2016 through July 2023.¹ We estimate the model

¹As respondents are asked about their expectations around the start of the quarter, we attribute it to the end-month of the preceding quarter.

parameters by maximum likelihood using the Kalman filter, which can easily handle the fact that the SPF data is quarterly whereas the rest of our data for the estimation is monthly.

The model fits these data well. In particular, the fitted headline ILS rates track the moves in the observed series very closely, with an average root mean squared error (RMSE) of just a couple of basis points. This also means that the model captures well the fact that near-term headline ILS rates are more volatile than longer-term headline ILS rates and that ILS rates across maturities are highly persistent. The RMSEs for the included survey-based expectations are somewhat larger, but within a reasonable range of 5 to 30 bps, while the model also successfully captures the broad fluctuations in observed headline and core inflation.

The model-implied estimates of core ILS rates appear reasonable along several dimensions: (i) like realized core inflation is less volatile than headline inflation, the core ILS rates are less volatile than headline ILS rates, (ii) core ILS rates comove less with oil prices than headline ILS rates, (iii) the core inflation expectations, as reflected in core ILS rates, typically evolve similarly as the core inflation projections by Eurosystem staff, and (iv) consistent with market commentary at the time, core ILS rates suggest that market participants expected core inflation to be substantially more persistent than headline inflation following the 2022 energy price spike. Our model-implied decompositions of both headline and core ILS rates indicate a greater role for fluctuations of genuine inflation expectations in near-term ILS rates, while inflation risk premia dominate fluctuations in longer-dated ILS rates, such as the 5-year-in-5-years (5y5y) forward rate. It also appears that headline inflation risk premia are more volatile than core inflation risk premia, which may indicate that a substantial part of the headline inflation risk premia is earned due to energy and food inflation risks. Genuine headline and core inflation expectations are about equally volatile across different horizons. At the 5y5y horizon, both headline and core inflation expectations are very stable, hovering close to the ECB's 2 percent target, in line with the expectations reported in surveys.

One key advantage of our core inflation expectations measure is that it is based on financial market data, which means that it can be computed at, essentially, whatever frequency we observe headline ILS rates. Other measures of core inflation expectations, such as from surveys or projections based on macroeconomic models, are available at much lower frequency, usually quarterly or monthly, at best. This relative infrequency means that these measures may at times be of limited use for policymakers, as their signals can reflect stale information by the time of a policy meeting. The relative infrequency also means that core inflation expectations from surveys or macroeconomic models cannot be used to measure the immediate response to events like monetary policy decisions. To illustrate the usefulness of the high-frequency

nature of our measure, we compute daily core inflation expectations and estimate the causal impact of monetary policy shocks using the typical event-study methodology (see, for example, Kuttner (2001) and Gurkaynak et al. (2005)). We find that a monetary policy tightening shock significantly lowers core inflation expectations at near-term horizons, whereas the response at longer-term horizons is more muted. The effect does not die out after just a few days, but is relatively persistent, which suggests that the impact on core inflation expectations indeed has the potential to impact wider macroeconomic outcomes. Based on lower frequency surveys or macroeconomic models, similar inference would require different identifying restrictions, such as sign-, short- or long-run restrictions in VARs, which may be harder to defend.

In addition to their relative infrequency, these other sources of core inflation expectations typically focus on particular horizons, like the next one or two years. In contrast, we can compute core inflation expectations at any horizon, including long-run horizons like the average expectations over the 5y5y horizons, which is often used to assess the perceived credibility of a central banks' inflation objective.

To the best of our knowledge, we are the first to price core ILS rates and decompose them into market-based expectations for and risks around the core inflation outlook. But our approach is related to several strands of literature. The term structure model that we use belongs to the class of affine term structure models developed in Duffie and Kan (1996), Dai and Singleton (2002) and Duffee (2002), which are widely used for modelling the term structure of interest rates (see Kim and Wright (2005), Christensen et al. (2011), Joslin et al. (2011) and Adrian et al. (2013), among many others). Camba-Mendez and Werner (2017) also use this class of models for decomposing headline ILS rates into genuine inflation expectations and risk premia, while Christensen et al. (2010), Abrahams et al. (2016), and D'Amico et al. (2018) use similar models to extract headline inflation expectations and risk premia from nominal and real bonds, but these studies do not explore core expectations and premia. Our approach to inferring core ILS rates from headline ILS rates, realized headline and core inflation as well as survey expectations for headline and core inflation is also related to Ang et al. (2008), Martin et al. (2021) and Mouabbi et al. (2021) who use a similar methods to infer real interest rates from nominal bonds and realized inflation and GDP-linked bond yields from dividend swaps and realized and survey expectations for GDP growth, respectively. Ajello et al. (2020) is perhaps the study most closely related to ours. They estimate a term structure model for nominal yields that allow for different persistence of core, energy, and food inflation and use several other real activity data. Relative to their study, we separately measure core inflation expectations and risk premia, we provide core inflation expectations at a higher-frequency, and we provide evidence on the causal effects

of monetary policy shocks on core inflation expectations and risk premia.

The rest of the paper is structured as follows: Section 2 describes the synthetic core ILS contracts in detail and provide intuition behind the information embedded in the swaps, Section 3 presents term structure model for headline and core ILS rates, Section 4 describes our data and estimation procedure, Section 5 presents our empirical results, section 6 shows that a monetary policy tightening shock leads to a decline in market-based core inflation expectations, and Section 7 briefly summarizes our conclusions.

2 Core Inflation-Linked Swaps

In this section, we describe the synthetic core ILS contracts and provide some intuition behind movements in associated swap rates. Specifically, we decompose the synthetic core ILS rates into average expected core inflation over the lifetime of the swap contract and a core inflation risk premium that compensates investors for core inflation risk. We then compare this swap rate decomposition to that of the usual headline ILS contracts.

A zero-coupon core ILS contract, which we consider here, is a derivative product, where two counterparties agree to exchange cash flows at a pre-determined date. The cash flows consist of a fixed and a floating leg; the fixed leg is set when the swap contract is agreed, while the floating leg is tied to a core inflation index that is only realized at settlement. In particular, the protection buyer pays a fixed amount $e^{h \cdot i_{t,h}^{(c)}}$ to the protection seller (assuming a notional amount of 1), where $i_{t,h}^{(c)}$ is the agreed continuously compounded swap rate at time t on the contract that matures h years later. The protection seller pays a floating amount $\mathcal{I}_{t+h}^{(c)}/\mathcal{I}_t^{(c)}$ to the protection buyer (again, assuming a notional amount of 1), where $\mathcal{I}_t^{(c)}$ denotes the underlying core inflation index at time t .

The swap rate $i_{t,h}^{(c)}$ is negotiated between the two counterparts such that the price of the contract is zero at time t or, equivalently, that the prices of the fixed and floating legs are equal. Under the assumption that arbitrage opportunities do not exist, the price of the fixed and floating legs are $\mathbb{E}_t \left[\mathcal{M}_{t,t+h} e^{h \cdot i_{t,h}^{(c)}} \right] = \mathbb{E}_t \left[\mathcal{M}_{t,t+h} \right] e^{h \cdot i_{t,h}^{(c)}}$ and $\mathbb{E}_t \left[\mathcal{M}_{t,t+h} \frac{\mathcal{I}_{t+h}^{(c)}}{\mathcal{I}_t^{(c)}} \right]$, respectively, where $\mathcal{M}_{t,t+h} = \mathcal{M}_{t,t+1} \cdot \mathcal{M}_{t+1,t+2} \cdot \dots \cdot \mathcal{M}_{t+h-1,t+h}$ denotes the stochastic discount factor, which discounts nominal cash flows from time $t+h$ back to time t , and $\mathbb{E}_t[\cdot]$ denotes the expectation conditional on time t information. It then follows that the swap rate is given by

$$i_{t,h}^{(c)} = \frac{1}{h} \log \mathbb{E}_t \left[e^{\xi_{t,t+h} + \pi_{t,t+h}^{(c)}} \right], \quad (1)$$

where $\xi_{t,t+h} = \log \mathcal{M}_{t,t+h} - \log \mathbb{E}_t \left[\mathcal{M}_{t,t+h} \right]$ and $\pi_{t,t+h}^{(c)} = \log \mathcal{I}_{t+h}^{(c)} - \log \mathcal{I}_t^{(c)}$. Assuming that $\xi_{t,t+h}$

and $\pi_{t,t+h}^{(c)}$ are both normally distributed, as will be the case in our model below, we have that

$$i_{t,h}^{(c)} = \mathbb{E}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(c)} \right] + \mathbb{C}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(c)}, \xi_{t,t+h} \right] + \frac{h}{2} \mathbb{V}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(c)} \right], \quad (2)$$

where $\mathbb{C}_t[\cdot, \cdot]$ and $\mathbb{V}_t[\cdot]$ denotes covariance and variance conditional on time t information.²

The first term on the right-hand side of Equation (2) is the average expected core inflation rate over the lifetime of the swap contract. Intuitively, higher core inflation expectations increase the swap rate (*ceteris paribus*). Higher expected core inflation over the lifetime of the swap contract implies a higher expected cash flow on the floating leg of the contract. Thus, for higher core inflation expectations, the protection seller requires (and the protection buyer is willing to pay) a higher fixed leg rate to enter into the swap contract.

The second term on the right-hand side of Equation (2) is a core inflation risk premium. Suppose core inflation tends to be high when the stochastic discount factor is high (that is, in "bad times"). In that case, the floating leg cash flows will be high when they are valued the most. The protection seller requires, and the protection buyer is willing to pay, compensation for that risk. This would result in the covariance term being positive and thus the swap rate being above the average core inflation expectations over the lifetime of the swap contract. If core inflation instead tends to be high when the stochastic discount factor is low, the intuition is similar, but the risk premium is negative. The third and final term is a convexity term. In our model below, this term is constant over time and relatively small, so we will simply refer to the core inflation risk premium as the difference between the core ILS rate and the average expected core inflation over the lifetime of the swap contract.

It is clear that the decomposition in Equation (2) would equally apply to a swap contract

²To arrive at the expression in Equation (2), we use that $\mathbb{E}[e^X] = e^{\mu + \frac{1}{2}\sigma}$ when $X \sim \mathcal{N}(\mu, \sigma)$, which implies that $\mathbb{E} \left[e^{\xi_{t,t+h} + \pi_{t,t+h}^{(c)}} \right] = e^{\mathbb{E}_t[\xi_{t,t+h} + \pi_{t,t+h}^{(c)}] + \frac{1}{2} \mathbb{V}_t[\xi_{t,t+h} + \pi_{t,t+h}^{(c)}]} = e^{\mathbb{E}_t[\xi_{t,t+h}] + \mathbb{E}_t[\pi_{t,t+h}^{(c)}] + \frac{1}{2} \mathbb{V}_t[\xi_{t,t+h} + \pi_{t,t+h}^{(c)}] + \mathbb{C}_t[\xi_{t,t+h}, \pi_{t,t+h}^{(c)}]}$. Since $\mathbb{E}_t[\xi_{t,t+h}] = \mathbb{E}_t[\log \mathcal{M}_{t,t+h}] - \log \mathbb{E}_t[\mathcal{M}_{t,t+h}] = \mathbb{E}_t[\log \mathcal{M}_{t,t+h}] - \mathbb{E}_t[\log \mathcal{M}_{t,t+h}] - \frac{1}{2} \mathbb{V}_t[\log \mathcal{M}_{t,t+h}] = -\frac{1}{2} \mathbb{V}_t[\log \mathcal{M}_{t,t+h}]$, and because $-\frac{1}{2} \mathbb{V}_t[\log \mathcal{M}_{t,t+h}] = -\frac{1}{2} \mathbb{V}_t[\xi_{t,t+h}]$ as $\log \mathcal{M}_{t,t+h}$ is the only random component of $\xi_{t,t+h}$, it follows that $\mathbb{E} \left[e^{\xi_{t,t+h} + \pi_{t,t+h}^{(c)}} \right] = e^{\mathbb{E}_t[\pi_{t,t+h}^{(c)}] + \frac{1}{2} \mathbb{V}_t[\pi_{t,t+h}^{(c)}] + \mathbb{C}_t[\xi_{t,t+h}, \pi_{t,t+h}^{(c)}]}$. Equation (2) then also uses that $\pi_{t,t+h}^{(c)} = \log \mathcal{I}_{t+h}^{(c)} - \log \mathcal{I}_t^{(c)} = \log \mathcal{I}_{t+h}^{(c)} - \log \mathcal{I}_{t+h-1}^{(c)} + \log \mathcal{I}_{t+h-1}^{(c)} - \dots + \log \mathcal{I}_{t+1}^{(c)} - \log \mathcal{I}_t^{(c)} = \pi_{t+h-1,t+h}^{(c)} + \dots + \pi_{t,t+1}^{(c)}$.

written on headline inflation. In that case, the decomposition would read

$$\begin{aligned}
i_{t,h} &= \mathbb{E}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i} \right] + \mathbb{C}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}, \xi_{t,t+h} \right] + \text{convexity} \\
&= \omega \mathbb{E}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(c)} \right] + (1-\omega) \mathbb{E}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(fe)} \right] + \\
&\quad \omega \mathbb{C}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(c)}, \xi_{t,t+h} \right] + (1-\omega) \mathbb{C}_t \left[\frac{1}{h} \sum_{i=0}^{h-1} \pi_{t+i,t+1+i}^{(fe)}, \xi_{t,t+h} \right] + \text{convexity},
\end{aligned} \tag{3}$$

where $i_{t,h}$ denotes the swap rate on the contract that has headline inflation as the underlying index, $\pi_{t,t+h} = \omega \pi_{t,t+h}^{(c)} + (1-\omega) \pi_{t,t+h}^{(fe)}$ denotes headline inflation between time t and $t+h$, where ω is the weight on core inflation in the overall index and $\pi_{t,t+h}^{(fe)}$ is food and energy inflation. Thus, from Equation (3), it follows that an appropriately weighted difference between headline and core ILS rates would recover food and energy ILS rates. It then also follows that estimates of core ILS rates, as well as an expectation and risk premium decomposition of those rates, would help shed light on the drivers of headline ILS rates.

Readers may be more familiar with the difference between a nominal and inflation-linked bond as a measure market-based inflation compensation. Indeed, the decomposition of the ILS rates into inflation expectations and an inflation risk premium is, in theory, identical to that of the differential yield on a nominal and inflation-linked bond (see, among many other examples, D'Amico et al. (2018)). In practice, however, the two may differ due to factors such as liquidity or deflation protection options embedded in some inflation-linked bonds. In what follows, we focus on swap rates as the universe of outstanding inflation-linked bonds tied to euro area inflation is relatively small.

3 A Model of Headline and Core Inflation-Linked Swaps

In this section, we present a joint model for headline and core ILS rates. Our model belongs to the class of Gaussian Affine Term Structure Models (ATSMs). We present the headline ILS and core ILS parts of our model in sections 3.1 and 3.2, respectively.

3.1 Headline ILS rates

We assume that headline inflation is a linear combination of n_x factors X_t . That is,

$$\pi_{t,t+1} = \gamma_0 + \gamma_x X_{t+1} + \gamma_u u_{t+1}, \tag{4}$$

where γ_0 and γ_u are scalars, γ_x is a $1 \times n_x$ vector, and $u_{t+1} \sim \mathcal{N}(0, 1)$. The factors follow a first-order Gaussian vector autoregression (VAR), such that

$$X_{t+1} = \mu + \Phi X_t + \Sigma \varepsilon_{t+1}, \quad (5)$$

where μ is an $n_x \times 1$ vector, Φ is an $n_x \times n_x$ matrix, Σ is an $n_x \times n_x$ matrix, and $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$ is an $n_x \times 1$ vector of i.i.d. normal shocks. Finally, we assume that the stochastic discount factor takes the form

$$\mathcal{M}_{t,t+1}/\mathbb{E}_t[\mathcal{M}_{t,t+1}] = e^{-\frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1}}, \quad (6)$$

where λ_t is the $n_x \times 1$ vector with market prices of risks. These market prices of risks are affine in the factors. That is,

$$\lambda_t = \Sigma^{-1}(\lambda_0 + \lambda_x X_t), \quad (7)$$

where λ_0 is an $n_x \times 1$ vector and λ_x is an $n_x \times n_x$ matrix.

The model assumptions in equations (4) through (7), together with the no-arbitrage restriction in equation (1), implies that we can write headline ILS rates recursively as

$$e^{h \cdot i_{t,h}} = \mathbb{E}_t \left[e^{\xi_{t,t+h} + \pi_{t,t+h}} \right] = \mathbb{E}_t \left[e^{\pi_{t,t+1} - \frac{1}{2}\lambda_t'\lambda_t - \lambda_t'\varepsilon_{t+1} + (h-1) \cdot i_{t+1,h-1}} \right]. \quad (8)$$

It then follows, by guess and verification, that headline ILS rates of all maturities are affine in the factors. That is,

$$i_{t,h} = \frac{1}{h}(a_h + b_h X_t), \quad (9)$$

where the intercepts a_h and factor loadings b_h are given by the recursions

$$b_h = (b_{h-1} + \gamma_x)(\Phi - \lambda_x) \quad (10)$$

$$a_h = a_{h-1} + (b_{h-1} + \gamma_x)(\mu - \lambda_0) + \frac{1}{2}(b_{h-1} + \gamma_x)\Sigma\Sigma'(b_{h-1} + \gamma_x)' + \frac{1}{2}\gamma_u^2 + \gamma_0, \quad (11)$$

with the initial conditions $a_0 = 0$ and $b_0 = 0_{(1 \times n_x)}$. Appendix A provides details on the necessary derivations.

This part of our model is fairly standard, in the sense that it has the structure of the popular Gaussian ATSM developed in Dai and Singleton (2002) and Duffee (2002). These models are often used to model the term structure of nominal interest rates, or the term structures of nominal and real interest rates jointly. In joint nominal/real models, the inflation break-even rates have the same structure as the model we present here (see, among others, Ang et al. (2008), Christensen et al. (2010), Abrahams et al. (2016), Ajello et al. (2020) and D'Amico

et al. (2018)).

3.2 Core ILS rates

To have a joint model for headline and core ILS rates, we need one further assumption on the dynamics of realized core inflation. We assume, analogously to headline inflation, that core inflation is a linear combination of the n_x factors. That is,

$$\pi_{t,t+1}^{(c)} = \gamma_0^{(c)} + \gamma_x^{(c)} X_t + \gamma_v v_{t+1}, \quad (12)$$

where $\gamma_0^{(c)}$ and γ_v are scalars, $\gamma_x^{(c)}$ is a $1 \times n_x$ vector, and $v_{t+1} \sim \mathcal{N}(0, 1)$.

The assumption that core inflation is driven by the same set of factors as headline inflation should be relatively uncontroversial: since headline inflation is a weighted average of core and food and energy inflation, it should reflect any factors driving core inflation. If there are factors driving food and energy inflation, which do not show up in core inflation, then those factors should still show up in headline inflation. Equation (12) is consistent with this scenario when the relevant entry in $\gamma_x^{(c)}$ linking core inflation to those factors are set to zero, while the relevant entries in Equation (4) linking headline inflation to those factors would be non-zero.

Combining equation (12) with the rest of the model structure from above, it follows that core ILS rates must follow the recursions

$$e^{h \cdot i_{t,h}^{(c)}} = \mathbb{E}_t \left[e^{\pi_{t,t+1}^{(c)} - \frac{1}{2} \lambda_t' \lambda_t - \lambda_t' \varepsilon_{t+1} + (h-1) \cdot i_{t+1,h-1}^{(c)}} \right]. \quad (13)$$

Using all the same arguments as before, it follows that core ILS rates of all maturities are affine in the factors. That is,

$$i_{t,h}^{(c)} = \frac{1}{h} \left(a_h^{(c)} + b_h^{(c)} X_t \right), \quad (14)$$

where the intercepts $a_h^{(c)}$ and factor loadings $b_h^{(c)}$ are given by the recursions

$$b_h^{(c)} = \left(b_{h-1}^{(c)} + \gamma_x^{(c)} \right) (\Phi - \lambda_x) \quad (15)$$

$$a_h^{(c)} = a_{h-1}^{(c)} + \left(b_{h-1}^{(c)} + \gamma_x^{(c)} \right) (\mu - \lambda_0) + \frac{1}{2} \left(b_{h-1}^{(c)} + \gamma_x^{(c)} \right) \Sigma \Sigma' \left(b_{h-1}^{(c)} + \gamma_x^{(c)} \right)' + \frac{1}{2} \gamma_v^2 + \gamma_0^{(c)}, \quad (16)$$

with the initial conditions $a_0^{(c)} = 0$ and $b_0^{(c)} = 0_{(1 \times n_x)}$.

4 Identification, Data, and Estimation

In this section, we describe how we take our joint model for headline and core ILS rates to the data, when core ILS rate data do not exist. Section 4.1 discusses how we identify unobservable core ILS rates using only data on headline ILS rates and realized core inflation. Section 4.2 provides details and a preliminary look at the data that we use. In section 4.3, we outline our estimation procedure.

4.1 Identification

If we observed core ILS rates over time, then we could simply estimate the model parameters of the single-asset portion of our joint model in section 3.2 using standard approaches for estimating ATSMs. This is the approach taken by many previous studies, such as, for example, Kim and Wright (2005), Christensen et al. (2009), Joslin et al. (2011), Adrian et al. (2013) for the term structure of nominal bonds or Camba-Mendez and Werner (2017) for headline ILS rates. Alternatively, we could have estimated the entire joint model using a vector of both headline and core ILS rates over time. This would be similar to, for example, Christensen et al. (2010), Abrahams et al. (2016), D’Amico et al. (2018) for nominal and real bonds or Anderson et al. (2010), Meldrum et al. (2023) for multi-country nominal bonds.

But the fact that we do not observe core ILS rates means that we need to consider a different approach. In theory, we could identify all the parameters of our joint model using data on headline ILS rates and realized core inflation. To see this, consider the following two-step procedure. In step one, we would identify the latent factors and the parameters of the single-asset model in section 3.1. This would essentially be the approach and model estimated for HICP ILS in Camba-Mendez and Werner (2017). In step two, to be able to infer the factor loadings of core inflation, we would regress realized core inflation onto the estimated latent factors to identify the additional parameters in equation (12). In practice, we estimate all the model parameters in a single step, which should be more efficient. We also consider survey expectations for both headline and core inflation. Adding these data has two advantages. First, it helps identify the time series parameters for the factor dynamics because expected headline and core inflation are linear combinations of the expected factors in the model (as explained in Kim and Orphanides (2012) for nominal interest rates). Second, it gives additional observations from which we can infer the parameters $\gamma_0^{(c)}, \gamma_x^{(c)}$. Our approach to infer core ILS rates is similar to that used in, for example, Ang et al. (2008) to infer real yields from nominal yields and inflation or Martin et al. (2021) and Mouabbi et al. (2021) to infer the prices of GDP-linked bonds from dividend swaps and GDP growth.

The illustrative two-step procedure above makes it evident that the key assumption that we are making is that headline ILS rates span the relevant information for pricing core ILS rates. That is, headline ILS rates reveal all the relevant factors X_t . This seems like a reasonable assumption, because core inflation is a component of headline inflation. It is possible, however, that there are unspanned (or "hidden") factors. This could arise, for example, if there is a factor that moves core ILS rates up, but energy ILS rates down, such that headline ILS rates are exactly unchanged. In that case, we would not be able to infer this factor from headline ILS rates alone and we would then not be able to infer how core ILS rates load on that factor. In theory, these are "knife-edge" restrictions (Duffee (2011)): even a small non-zero loading would allow us to infer relevant the factors. In practice, however, small non-zero loadings may make identification difficult or imprecise.

Unfortunately, it is difficult to test the spanning assumption. Some previous tests of the spanning assumption (such as Joslin et al. (2014)) have relied on a regression of the variable that is assumed to be spanned on the variables that are supposed to span that variable. In our case, a regression of core inflation onto headline ILS rates. The idea behind this test is that the resulting R^2 will likely be very high, if the variable is indeed spanned. Bauer and Rudebusch (2017) show, however, that even if the variable of interest is fully spanned, we may get low R^2 's in these regressions if the spanning variables are observed with even tiny measurement errors. It may also be that some short-run fluctuations in the variable that is assumed spanned is unimportant for pricing the spanning variables. This would likewise lead to low R^2 's in the regression (Kim (2009)).

That said, there are some potential red flags to look for that may indicate issues related to the spanning assumption. One such red flag is if the model fits core inflation using a factor that does not explain much variation in headline ILS rates. That may indicate that a good portion of the variation in core inflation is unrelated to that priced into headline ILS rates. Another red flag is if the model provides a poor fit to either core inflation or headline ILS rates altogether. This may indicate that the model struggles to fit both sets of data using the same underlying factors and "gives up" along one of the dimensions. It appears, however, that we are not in any of these situations. We will elaborate on this claim in the empirical section below.

4.2 Data

We use three sets of data to estimate the model. The first set of data we consider is end-month zero-coupon headline ILS rates tied to the euro area Harmonized Index of Consumer Prices

Table 1: Sample moments of observable ILS rates

	1y	2y	3y	5y	7y	10y
Mean	1.6	1.5	1.6	1.6	1.7	1.8
Standard deviation	1.0	0.7	0.7	0.6	0.5	0.5
Autocorrelation	0.93	0.94	0.95	0.96	0.97	0.97

Note: ILS rates tied to euro area HICP excluding tobacco. The third row reports first order autocorrelation. Data are monthly from June 2005 to July 2023.

(excluding Tobacco) with maturities of 1-, 2-, 3-, 5-, 7-, and 10-years.³ These data are available from June 2005 and our sample end in July 2023. Table 1 provides some summary statistics for these variables: In our sample, the average headline ILS rates have been somewhat below 2 percent, but increasing with maturity. Longer maturity headline ILS rates are less volatile than nearer-term ones, and all headline ILS rates are highly persistent with auto-correlations close to one.⁴ Figure 1 shows time series of the raw data for ILS rates. Headline ILS rates of different maturities comove strongly and suggest a fairly low-dimensional factor structure. We confirm this by principal component analysis, which suggest that three principal components explain 99.96 percent of the variation in the cross section of headline ILS rates. We therefore, as is typical in the literature on the nominal term structure of interest rates, consider a three factor model (that is, $n_x = 3$).

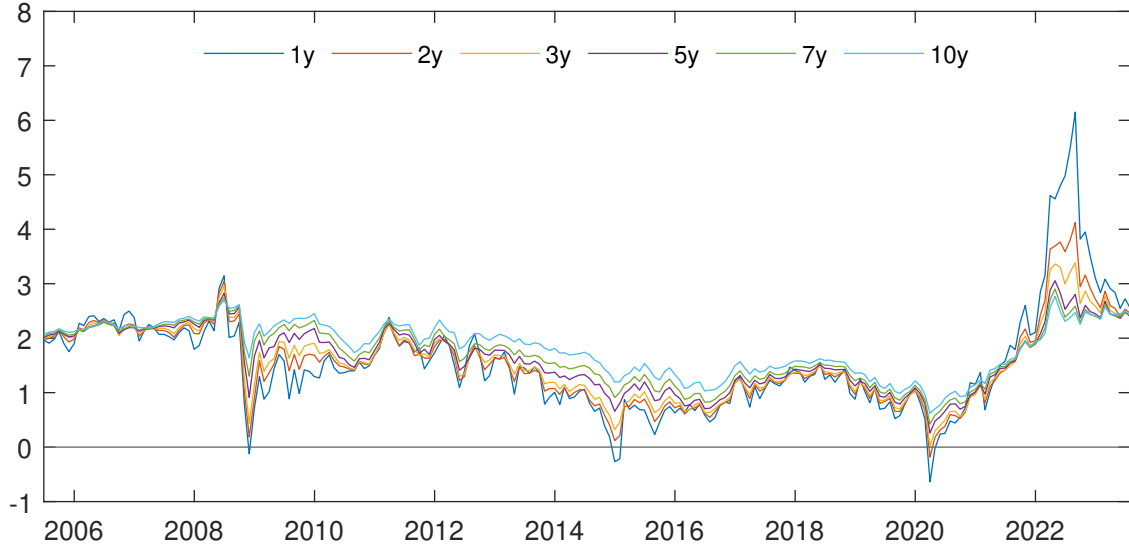
The second set of data we include are realized inflation data. For headline, we use month-to-month changes in euro area HICP, while for core we use euro area HICP excluding food, energy and tobacco. These data are also available monthly and cover the full June 2005 through July 2023 sample. While it is not strictly necessary to include the realized headline inflation measure in the estimation, we nevertheless choose to do so to treat headline and core inflation symmetrically.

Finally, the third set of data we include are survey-based inflation forecasts. We use expected headline and core inflation reported in the ECB’s Survey of Professional Forecasters (SPF). The SPF has been conducted since 1999 and asks a number of experts employed at financial and non-financial institutions (for instance research institutes) about - inter alia - their point forecast of euro area headline and core inflation over a range of horizons. We use the mean of survey-

³Note that ILS rates are subject to an indexation lag of three months. Following Camba-Mendez and Werner (2017), we adjust ILS rates for this indexation lag by interpolating the ILS-implied projection of the overall price index and computing purely forward-looking implied ILS rates from this price index projection.

⁴Given the relatively short sample, it is well-known that auto-correlations may have a downward bias in highly persistent data. In our case, we incorporate information from surveys, following Kim and Orphanides (2012), to alleviate this issue. Alternatively, one could apply a statistical bias correction as in Bauer et al. (2012)

Figure 1: Time series of observable ILS rates



Note: ILS rates tied to euro area HICP excluding tobacco. Data are monthly from June 2005 to July 2023.

respondents' expectations to headline and core inflation at the 1-, 2-, and 5-year horizons. The survey has quarterly frequency and is typically conducted at the beginning of each quarter, so we align it with the last end-month observation in the previous quarter. Before the fourth quarter of 2016, the SPF did not ask respondents for their core inflation expectations, so we are not able to use survey-based information about core inflation before then.

4.3 Estimation

Our model can be written as a linear-Gaussian state space system. The measurement equation reads

$$\begin{bmatrix} i_t \\ \pi_t \\ \pi_t^{(c)} \\ s_t \end{bmatrix} = \begin{bmatrix} a \\ \gamma_0 \\ \gamma_0^{(c)} \\ \psi_0 \end{bmatrix} + \begin{bmatrix} b \\ \gamma_x \\ \gamma_x^{(c)} \\ \psi_x \end{bmatrix} X_t + \begin{bmatrix} \sigma_i w_{i,t} \\ \gamma_u u_t \\ \gamma_v v_t \\ \sigma_s w_{s,t} \end{bmatrix}, \quad (17)$$

where i_t is a vector collecting headline ILS rates, a and b are given from the equations (10) and (11), s_t is a vector collecting headline and core survey expectations, and ψ_0 and ψ_x are survey loadings which follow from equations (4), (5), and (12). The variables $w_{i,t} \sim \mathcal{N}(0, I)$ and $w_{s,t} \sim \mathcal{N}(0, I)$ are measurement errors on headline ILS rates and survey expectations, respectively, where σ_i and σ_s denotes their standard deviations. The transition equation of the state space system is given by equation (5). We can then estimate the model parameters by maximum likelihood, using the Kalman filter to evaluate the likelihood function. Using the estimated model parameters and latent factors, we can then infer core ILS rates using equation

(14).

As the factors are treated as latent, we must impose identifying restrictions on the parameters of the ATSM (see, for example, Dai and Singleton (2002), Joslin et al. (2011), and Hamilton and Wu (2012)). We impose the restrictions that $\mu = 0_{(n_x \times 1)}$, that Φ is lower triangular, and that Σ is the identity matrix. We choose to estimate the standard deviation of headline ILS measurement errors (assuming it is the same across maturities) within the estimation procedure, but, as in Kim and Orphanides (2012), calibrate the standard deviation of survey measurement errors. We set the standard deviations at 10, 15, and 25 basis points for the 1-, 2-, and 5-year ahead expectations, respectively. Because we are interested in measuring market-based measures of core inflation, we calibrate them to somewhat higher values than the standard deviation on headline ILS rate measurement errors will turn out to be. In this way, we ensure that the estimation does not overemphasize survey information but rather track its broad movements over time.

5 Empirical Results

In this section, we report our main empirical findings. We report parameter estimates, implied factor loadings, and in-sample fit to both headline ILS rates, headline and core inflation expectations from surveys as well as realized headline and core inflation. We attempt to evaluate the validity of our spanning assumption (that core inflation is spanned by the three latent factors summarizing headline inflation dynamics). Then we compute synthetic core ILS rates implied by our model and decompose both headline and core ILS rates into genuine inflation expectations and inflation risk premia. Finally, to verify that our synthetic core ILS rates behave reasonably, we conduct two "out-of-sample" exercises. First, we show that genuine core inflation expectations from our model are broadly similar to forecasts of core inflation from the Eurosystem staff projections, and that revisions to expectations for a given calendar year comove positively. Second, we show the rather intuitive results that, while spot oil prices explain part of the variation in headline ILS rates, it explains much less of the variation in core ILS rates.

5.1 Parameter Estimates

Table 2 shows the estimated and calibrated parameters characterizing the dynamics of our three pricing factors (X_t) under the physical (P) and risk-neutral (Q) measures. Under both measures, shocks to the first and third factor display very high auto-correlation, while shocks to the second factor die out more quickly. The high persistence of the first and third factor

under the physical measure captures the high auto-correlation of the observable headline ILS rates reported in Table 1.

Turning to the model-implied factor loadings, headline ILS rates across maturities load about equally on the third factor (see Figure 2). The model-implied synthetic core ILS rates also load about equally on the third factor across the term structure, but the loadings are somewhat smaller compared to the loadings for headline ILS rates. Overall, the third factor appears to have the classic "level"-factor interpretation, moving both headline and core ILS rates up or down about equally across the term structure.

Headline ILS rates load more heavily on the second factor compared to core ILS rates. This is particularly pronounced for nearer-term maturities, as the loadings for both headline and core ILS rates decrease as a function of the maturity of the swap contract. As the second factor is also estimated to be the least persistent factor, the factor - although it is a purely latent and statistical factor - could be capturing effects of temporary energy price shocks. A temporary energy price shock should be expected to primarily affect relatively short-horizon headline ILS rates, while eventually dying out as the horizon increases. The effect on core ILS rates should be muted relative to the effect on headline ILS rates, as the energy shock is only expected to partially spill over to underlying inflation.

The first factor allows for somewhat interesting dynamics between headline and core ILS rates. Core ILS rates load positively on the first factor, while headline ILS rates load negatively on the first factor. This pattern allows a shock to the first factor to move headline and core ILS rate in opposite directions on impact.

5.2 Model fit to headline ILS rates, survey data and realized inflation

Overall, the model fits all of the input data quite well. Model-implied headline ILS rates are, across the maturity spectrum, virtually indistinguishable from the observed swap rates, with RMSEs in the order of just a few basis points (see Figure 3). This also means that the model captures well the fact that near-term ILS rates are more volatile than longer-term ILS.

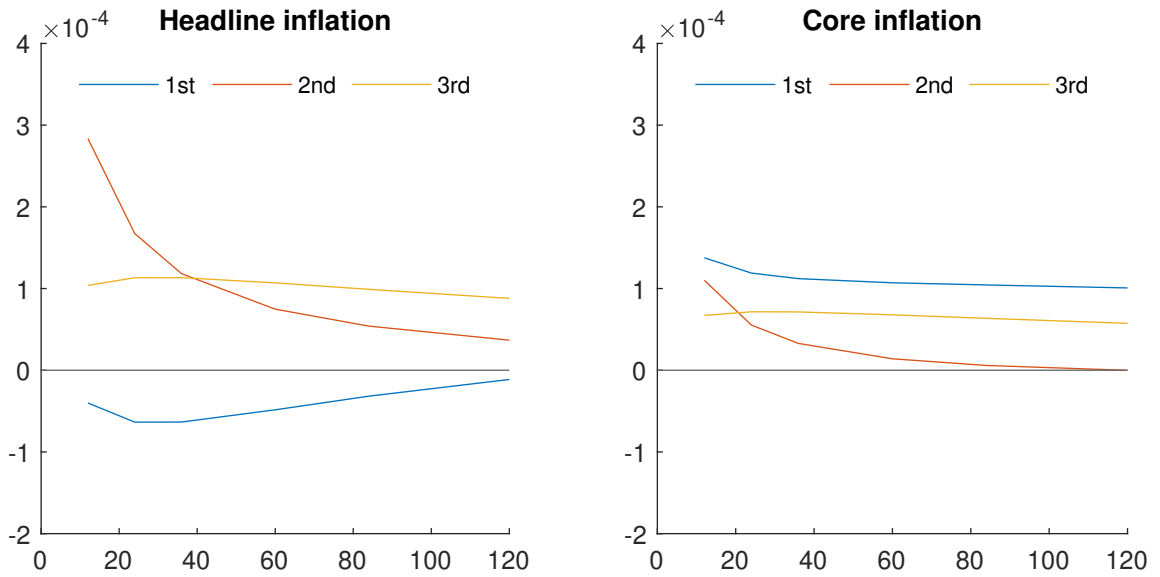
The model also fits headline and core inflation expectations as reported in surveys quite well, although the fit is not as tight as for ILS rates (see Figure 4). The RMSEs range between 5 and 30 basis points, with nearer-term forecast fitting errors generally larger than longer-horizon fitting errors. Again, the model successfully captures the pattern that more near-term survey forecasts show higher volatility than longer-term forecasts. In particular, the model captures well that 1-year ahead survey forecasts increased substantially in 2022, while 5-year ahead forecasts only edged higher.

Table 2: Estimated and calibrated parameters from the model

μ^P	0	0	0	μ^Q	0.06	0.15	-0.03
Φ^P	0.98			Φ^Q	0.98	-0.03	0.00
	-0.04	0.76			-0.05	0.84	0.02
	0.02	0.11	0.98		0.02	0.00	0.99
Σ	1	0	0				
	0	1	0				
	0	0	1				
γ_0	1.71			$\gamma_0^{(c)}$	1.67		
γ_x	0.10	0.72	0.05	$\gamma_x^{(c)}$	0.22	0.32	0.04

Note: μ^P is calibrated to zero, while Σ is calibrated to be an identity matrix. Φ^P is restricted to be lower triangular. The remaining parameters are estimated. γ_0 , γ_x , $\gamma_0^{(c)}$ and $\gamma_x^{(c)}$ are multiplied by a factor 1.000 to show granularity.

Figure 2: Factor loadings for headline and core inflation



Note: Estimated factor loadings for headline ILS rates and synthetic core ILS rates (equation (9) and (14)). The x-axis represents the maturity of ILS rates in months.

In terms of realized headline and core inflation, the model also shows a reasonable fit (see Figure 5). While the model does not fit all of the seasonal variation of the month-to-month inflation series, it captures the overall trend in headline and core inflation very well. Before 2016, the fitted core inflation series is somewhat above the realized one, potentially reflecting that the model has limited information about core inflation over this early period due to the lack of information about core inflation from surveys.

5.3 The spanning assumption

One key assumption in our model is that core inflation is a linear combination of the pricing factors spanning headline ILS rates. As discussed in section 4.1, validating this assumption is not straightforward, but there are some promising indications from the model output.

First, both headline and core inflation have non-zero loadings on all three pricing factors (see Figure 2). In fact, the fit of all information about core inflation – both the realized series and survey expectations – appear to be fitted about as well as the headline counterparts. This means that all of the factors driving headline ILS rates also appear important for core inflation, which means that we can indeed infer something about core inflation from headline ILS rates.

Second, none of the three factors appear unimportant for fitting headline ILS rates, as loadings are all reasonable away from zero (again, see Figure 2). That means that we are unlikely to be in the case where the model uses one factor to fit core inflation expectations and that factor is unimportant for headline ILS rates. This would be unfortunate because if there was a factor in core inflation that did not show up in headline ILS rates, we would not have been able to infer anything about its market price of risk parameters. This could have been the case if one of the factors moved core inflation and energy and food inflation in exactly offsetting direction, so the overall impact on headline inflation was exactly zero.

5.4 Implied core ILS rates

Based on our estimated model, we can calculate model-implied synthetic core ILS rates using the estimated factor time series and the model parameters. Figure 6 shows the time series of the model-implied core ILS rates along with the model-implied headline ILS rates. We report core ILS rates since 2016Q4, as the model does not have much information to estimate these prior to core survey expectations becoming available.

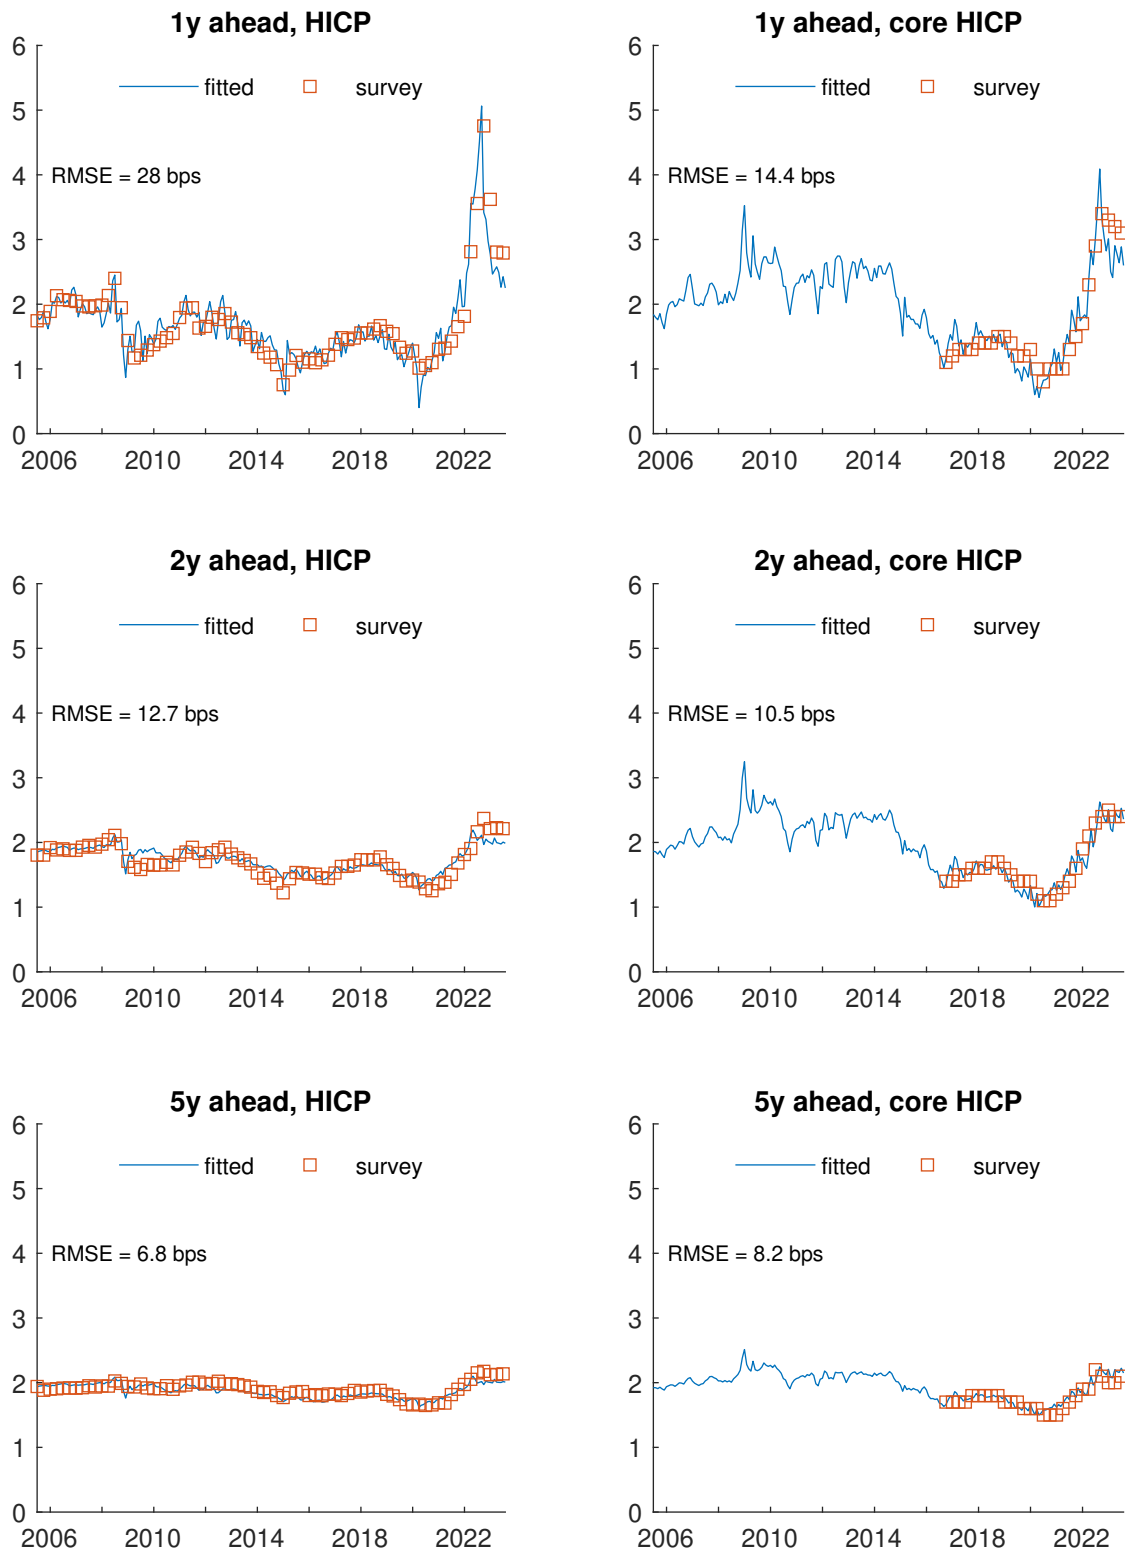
In the early parts of the overlapping sample, core ILS rates moved broadly in tandem with headline ILS rates, with both falling short of the ECB’s 2 percent inflation target across maturities. This reflects that realized inflation had been running relatively low for some time.

Figure 3: Model fit to ILS rates



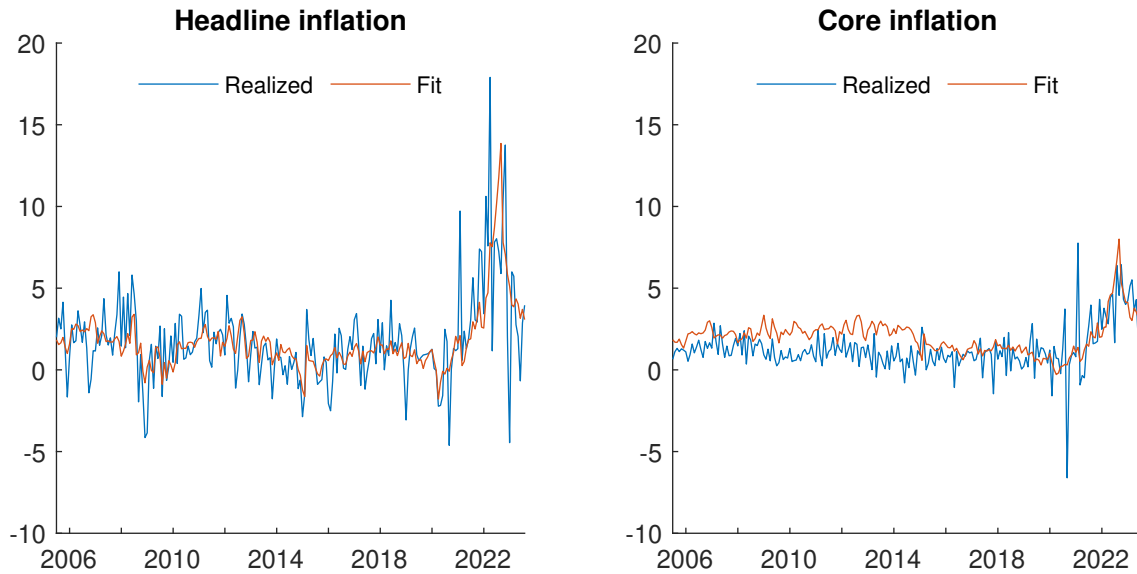
Note: Model-based headline ILS rates compared to raw data. Monthly data from June 2005 to July 2023.

Figure 4: Model fit to survey data



Note: Model-based genuine expectations to headline and core inflation under the physical probability measure. Monthly ILS data from June 2005 to July 2023. Headline and core inflation expectations from the Survey of Professional Forecasters at quarterly frequency.

Figure 5: Model fit realized headline and core inflation



Note: Month-on-month inflation for the euro area HICP (headline) and HICP less food and energy (core), June 2005 to July 2023.

Later in the sample, the two series differ more substantially. In particular, in early 2020, as the COVID-19 pandemic spread across Europe and the rest of the world, headline ILS rates fell markedly. Core ILS rates also declined, but less than headline ILS rates. The difference was more pronounced for the longest maturities, where the drop in core ILS rates was rather small. Over 2021 and 2022, both headline and core ILS rates trended up as inflationary pressures built. Headline ILS rates peaked substantially above core ILS rates in late 2022, as energy inflation shot up dramatically. As energy inflation eased, headline ILS rates came off their peak in early-to mid-2023 and even fell below core ILS rates for maturities of two years and above, consistent with the narrative that underlying inflation was proving stickier than headline inflation.

Tables 3 and 4 report selected moments for model-implied core and headline ILS rates, respectively. The unconditional moments are generally in line with standard intuition from above. Headline and core ILS rates roughly share unconditional means, which come out a touch below two percent, consistent with realized headline and core inflation trends over the sample. At near-term maturities, core ILS rates are less volatile than headline ILS rates, which seems intuitive as realized core inflation is also less volatile than headline. The difference in volatility narrows as the maturity of the contracts grows, however. As explained in detail in section 2 above, these unconditional moments may hide details about the volatility of genuine inflation expectations and inflation risk premia. Therefore, we next decompose both headline and core ILS rates into those components.

Figure 6: Model-implied core ILS rates



Note: Synthetic core ILS rates and fitted headline ILS rates from the model. Monthly data from June 2005 to July 2023. Synthetic core ILS only shown in the period, where SPF core expectations are available. The black line represents the ECB's current medium-term inflation target.

Table 3: Moments for model-implied core ILS rates

	1y	2y	3y	5y	7y	10y
Mean	1.6	1.5	1.6	1.6	1.7	1.9
Standard deviation	0.9	0.7	0.7	0.6	0.5	0.5
Autocorrelation (1 lag)	0.96	0.96	0.96	0.96	0.96	0.95

Note: Selected moments of the sample of core ILS rates from September 2016 to July 2023 (due to availability of SPF core inflation expectations).

Table 4: Moments for fitted headline ILS rates

	1y	2y	3y	5y	7y	10y
Mean	1.7	1.5	1.5	1.5	1.5	1.6
Standard deviation	1.3	0.9	0.8	0.7	0.6	0.5
Autocorrelation (1 lag)	0.95	0.96	0.96	0.97	0.97	0.98

Note: Selected moments of the sample of headline ILS rates from September 2016 to July 2023 (due to availability of SPF core inflation expectations).

5.5 Decomposing headline and core ILS rates

In order to decompose headline and ILS rates, we calculate model-implied average headline and core inflation over the same term structure as we have ILS rates. This is done using the expressions for headline and core inflation in equations (4) and (12) combined with the estimated factors and the factor dynamics in equation (5). The headline and core inflation risk premia are then the difference between the ILS rates and the genuine expectations.

Figures 7 and 8 show the model-implied time series of the decompositions into genuine expectations and risk premia for selected maturities for headline and core ILS rates, respectively. In both cases, we find that nearer-term ILS rates display more volatility in the genuine inflation expectations component relative to longer-term ILS rates. Further out the term structure, at the 5y5y forward horizon for example, the expectations component is rather stable, as most of the fluctuations in ILS rates are estimated to reflect risk premia. That said, genuine inflation expectations at the 5y5y horizon did increase from 1.9 percent in end 2020 and reached its peak of 2.1 percent in April 2022.⁵

The headline inflation risk premium generally fluctuates quite a bit more than the core

⁵The fact that long-term core inflation expectations were slightly below 2 percent before the pandemic could reflect that the ECB's quantitative inflation target was to keep inflation close to, but below 2 percent before its strategy review in September 2021.

inflation risk premium. The headline inflation risk premium was initially positive, but turned negative for a number of years in the early 2010's, before becoming positive again following the inflationary pressures in 2021. In contrast, the core inflation risk premium was rather stable around zero, until it increased amid the inflation surge in 2021 and 2022. As inflation risk premia are generally thought to be positive if inflation is expected to be high when marginal utility is high, for example in stagflationary environments, this suggests that core inflation risks had been reasonably balanced, although at a mean somewhat below the ECB's 2 percent target, before core inflation risks starting moving to the upside in 2021 following prolonged inflationary pressures.

Figure 9 compares the unconditional properties of these components more explicitly. Consistent with the sample considered, we find that genuine core inflation expectations fall a touch below genuine headline inflation expectations, on average, but the difference shrinks as the horizon considered increases. Likewise, the core inflation risk premium falls a touch below the headline inflation risk premium.

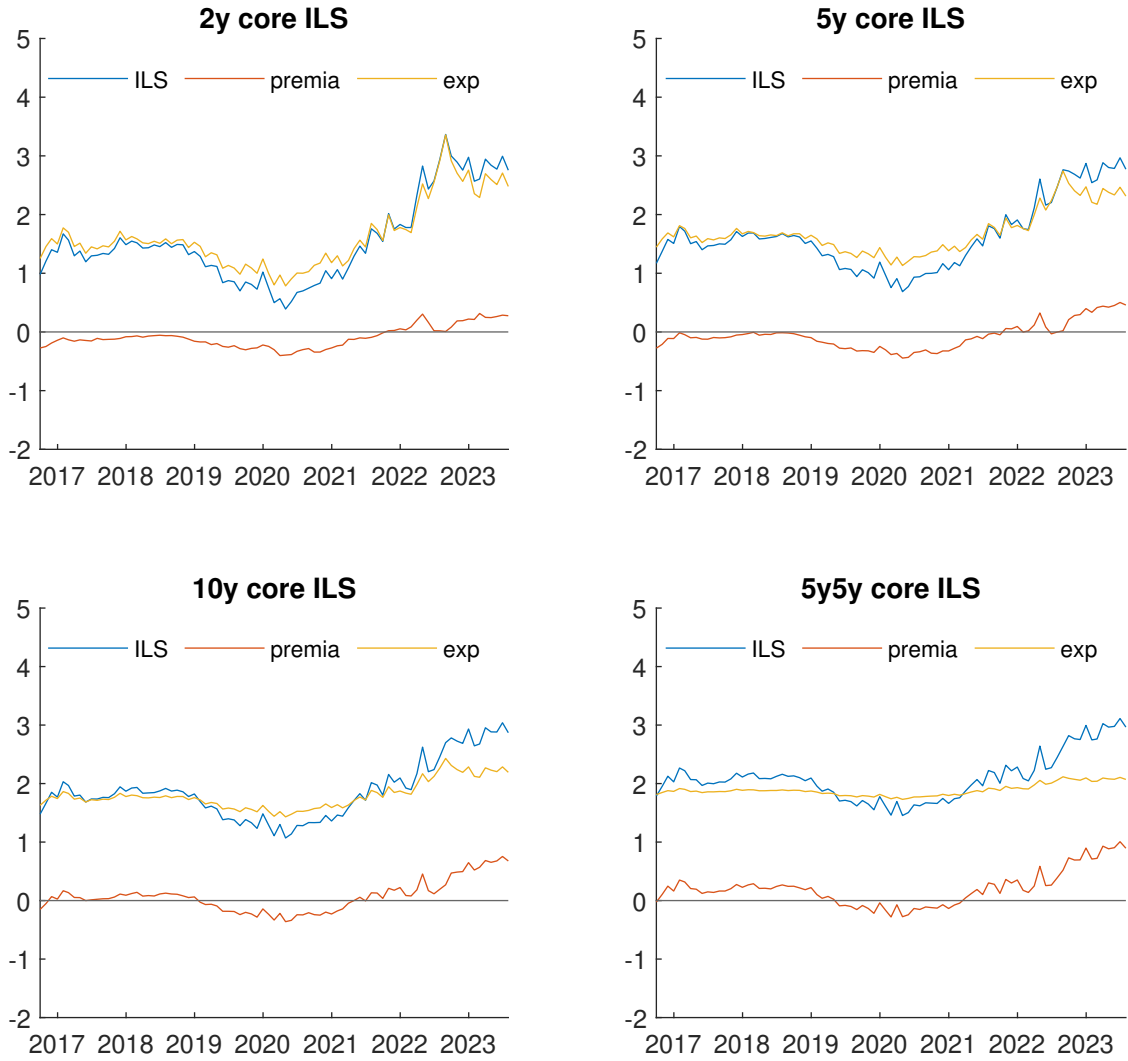
In terms of variation, headline and core inflation expectations are broadly equally variable and the variability declines with time to maturity, consistent with long-horizon inflation expectations being well anchored close to the ECB's 2 percent inflation target. In contrast, the headline inflation risk premium has been more volatile than the core inflation risk premium. This is consistent with the relatively big swings in the headline inflation risk premium, which was initially generally positive, then negative for several years, before turning positive again in recent years, while the core inflation risk premium was rather stable around zero, before trending up over 2021, 2022, and 2023. This pattern suggests that headline inflation risk premia are generally primarily fluctuating due to inflation risk associated with the food and energy components of the consumption basket.

5.6 Two "reasonability" exercises

As neither the core ILS rates nor genuine core inflation expectations from our model has an observable counterpart for comparison, we conduct two reasonability exercises to be comfortable about the reliability of our estimates. First, we compare our genuine core inflation expectations to the outcomes of the ECB and the Eurosystem's Macroeconomic Projection Exercise (MPE).⁶ As the MPE provides calendar-year based forecasts for core inflation, we calculate the same object from our model by constructing core inflation expectations for a given calendar year as the average expected change in the price level compared to the previous year. The genuine

⁶As we do not re-estimate the model for each forecasts, this constitutes a pseudo out-of-sample forecast exercise.

Figure 7: Decomposition of synthetic core ILS rates



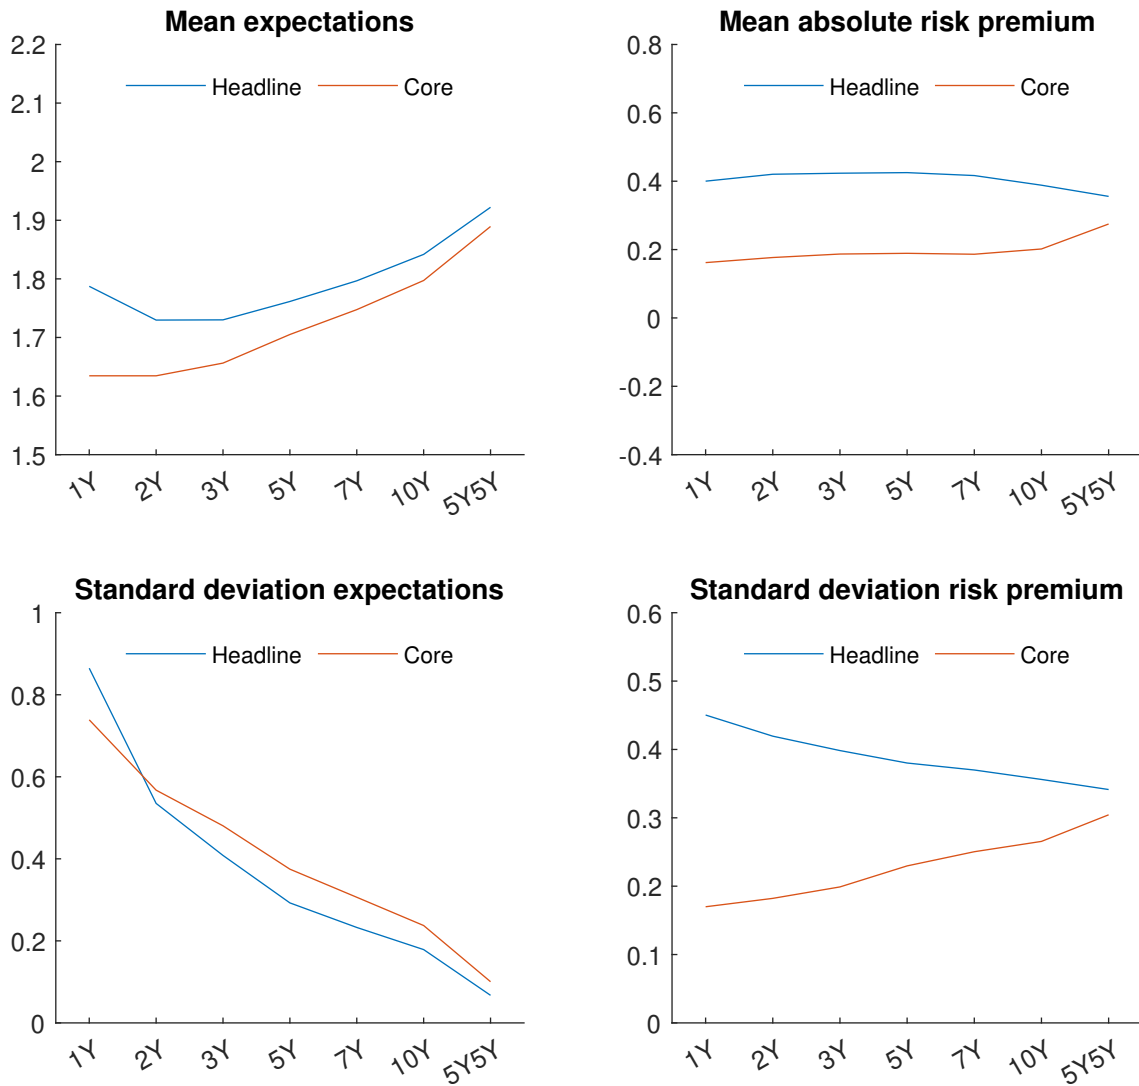
Note: Synthetic core ILS rates decomposed into genuine core inflation expectations and core inflation risk premia. Monthly data from September 2016 to July 2023.

Figure 8: Decomposition of ILS rates



Note: ILS rates decomposed into genuine core inflation expectations and core inflation risk premia. Monthly data from June 2005 to July 2023.

Figure 9: Moments for genuine inflation expectations and inflation risk premia



Note: Unconditional moments of (core) ILS rates decomposed into expectations and risk premia. Monthly data from September 2016 to July 2023.

Table 5: Percent of variation explained by spot oil prices

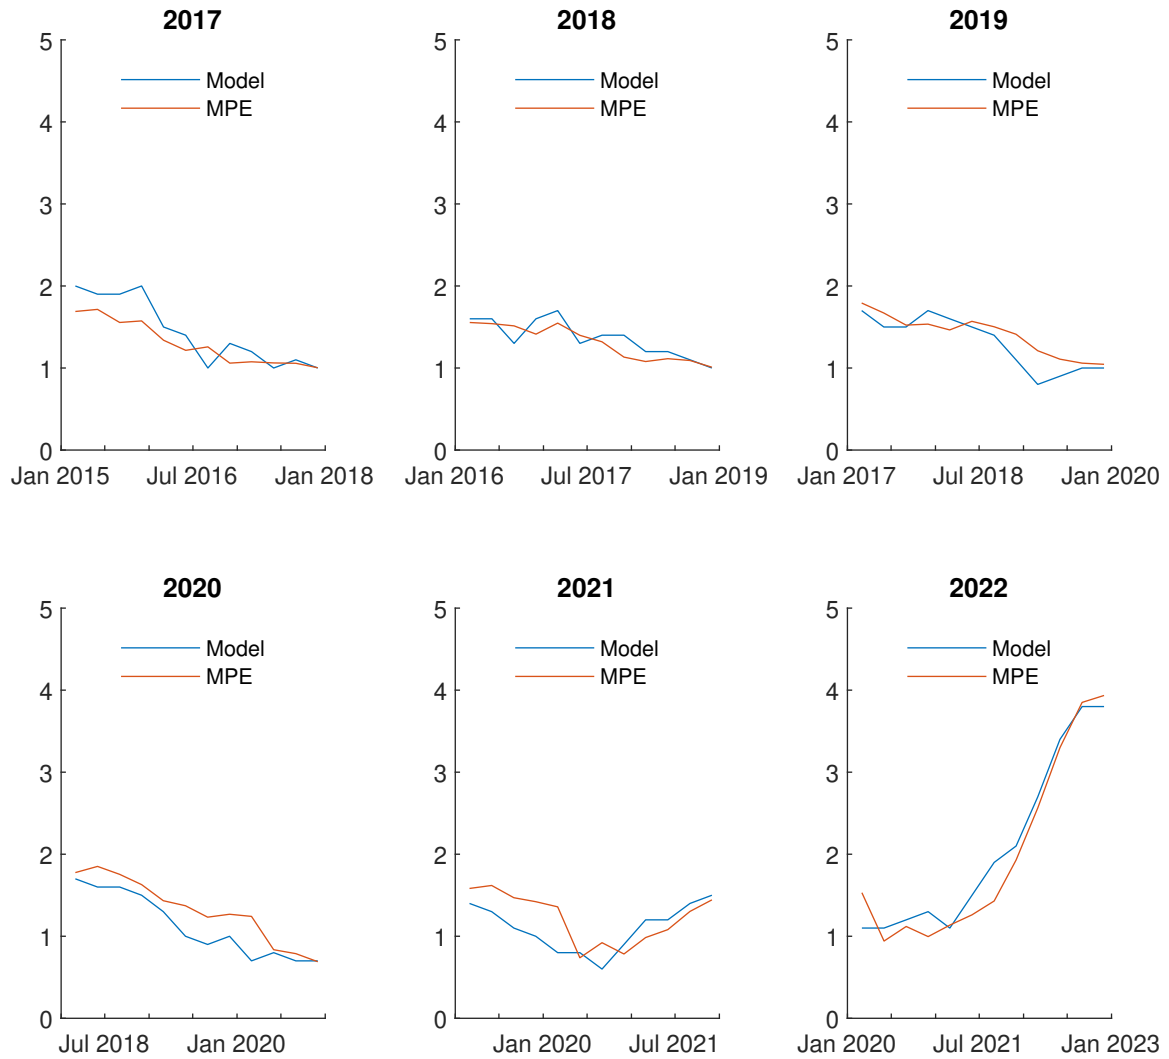
		1y	2y	3y	5y	7y	10y
Inflation-linked swap	Headline	17	21	21	20	17	13
	Core	0	0	0	0	0	1
Genuine expectations	Headline	9	10	10	10	10	11
	Core	0	0	0	0	0	0
Inflation risk premium	Headline	23	21	19	16	13	8
	Core	18	7	2	0	0	1

Note: Adjusted R^2 from regressing monthly changes in (core) ILS rates, genuine inflation expectations, and the inflation risk premium on monthly changes in the spot price of Brent oil. Monthly data from June 2005 to July 2023.

core inflation expectations from our model turns out to be broadly similar to core inflation expectations from the MPE and revisions to expectations implied by our model and the MPE tend to move in the same direction (see Figure 10). During 2021, for example, there were big upward revisions for 2022 inflation expectations from both our model and the MPE. In a similar vein, both our model and the MPE forecasts for the calendar year 2020 were revised down during 2019 and 2020. Hence, our model seems to pick up similar information about the core inflation outlook as the projections conducted by the ECB and the Eurosystem based on macroeconomic information.

As a second reasonability check, we regress monthly changes in our core ILS rates on monthly changes in the spot price of Brent oil. Table 5 reports R^2 's from such regressions for selected maturities. Across maturities, movements in the oil price explain very little or even none of the variation in our core ILS rates, while spot oil prices explain about 10-20 percent for headline ILS rates, depending on maturity. For the individual components of ILS rates, we also find that oil price fluctuations co-move more strongly with headline inflation expectations and risk premia compared to the counterparts for core inflation. These findings appear rather intuitively appealing. Unreported results suggest that the results are robust to using various maturities of futures prices on Brent oil rather than the spot price.

Figure 10: Evolution in fixed calendar-year expectations to core inflation



Note: "Model" is market participants genuine expectations to core inflation implied from our model. MPE refers to the Eurosystem's staff projections of core inflation. The frequency is quarterly to match that of the Eurosystem staff projections.

6 Monetary Policy Transmission to Headline and Core Inflation Expectations

One key advantage of our measure of core inflation expectations is that they can be computed at, essentially, whatever frequency that we observe headline ILS rates. In particular, we can compute our measure on a daily basis.⁷ In contrast, other measures are usually only available at much lower frequency. The SPF and the ECB staff projections, for example, are available only four times a year. Other surveys or macroeconomic projections may be available at a monthly frequency, but typically not more frequently than that. This high-frequency feature of our estimates make them useful for at least two reasons: (i) it provides policymakers with a real-time read on core inflation expectations, and (ii) it facilitates event-studies that can give valuable information on the causal effects of, for example, monetary policy shocks on core inflation expectations.

In this section, we make use of the high-frequency feature to measure the causal effects of monetary policy shocks on inflation expectations based on the event-study methodology. To measure monetary policy shocks mps_t , we use an updated version of the data from Altavilla et al. (2019). We identify monetary policy shocks as the first principal component of changes in the 1-, 3-, 6-month, and 1- through 10-year OIS rates tied to the €STR around ECB Governing Council monetary policy decisions. The changes in OIS rates are measured as the median quotes 10-20 minutes before the ECB press release and 10-20 minutes after the end of the press conference that follows the press release. We exclude events where the monetary policy shock and the Eurostoxx 50 index moves in opposite direction, because this likely signals that the event contained a central bank information shock (see Jarociński and Karadi (2020)). This could happen if the ECB Governing Council signalled that they judge the economy may be stronger/weaker than market participants expected. This may interfere with our inference for inflation expectations, as higher rates would be expected to lower inflation but a stronger economy would be expected to boost inflation. Our interest is not in these information shocks, but rather pure monetary policy shocks.

To measure the effect on core inflation expectations, we then run the standard event-study regressions

$$\Delta y_{t+h} = \alpha + \beta \cdot mps_t + u_{t+h}, \quad (18)$$

⁷Although we have estimated our model on monthly data above, we can re-formulate the model in daily terms by redefining the transition equation of our model. When implementing this, we fix the model parameters and add the daily observable headline ILS rates to our data set. We then re-run the Kalman filter and infer the state variables and core ILS rates day-by-day. The Kalman filter easily handles the fact that there are missing observations because the realized inflation series and SPF responses only are available at monthly and quarterly frequency, respectively.

where Δy_{t+h} measures the change in the outcome variable between time $t - 1$ (i.e. the closing price the day before the Governing Council meeting) and $t + h$. In our baseline specification, we consider the change in headline and core ILS rates for different maturities, but we also run the regressions for the genuine inflation expectations and inflation risk premia components separately. Our sample is limited by the availability of core ILS rates; while the sample is short, it is one with interesting ECB policy events, as it covers the pandemic period and the 2021-2023 bout of inflation and associated policy response. When implementing equation (18), we set h to be 3 days. That is, we measure the change in the relevant outcome variables as the difference in the value on the day before the ECB policy event and three days later. The choice of h involves a trade-off. On the one hand, picking h to be small runs the risk that the new information is not yet fully reflected in the ILS outcome variable. For example, Bahaj et al. (2023) report that fundamental shocks to UK ILS rates are incorporated in the pricing after 2-3 trading days. On the other hand, picking h to be too large runs the risk that inference is unnecessarily noisy due to other non-monetary news. We judge that 3 days is a reasonable compromise, but also report results for other choices below for robustness.

Table 6 reports the results for both headline and core ILS rate for maturities of 1-, 5-, and 10-year maturities. It also reports the results for the genuine expectations and risk premia components. We find that a one standard deviation monetary policy shock reduces 1-year headline ILS rates by 7 basis points, while it reduces 1-year core ILS rates by close to 3 basis points.⁸ Both effects are statistically different from zero, despite the relatively small sample. The fact that a monetary policy shock affects headline ILS rates more than core ILS rates suggest that monetary policy have larger effects on food and energy inflation expectations (or risk premia) than on core inflation expectations, for example because it is expected to affect commodity prices in the relatively near-term by dampening demand in the economy. That said, monetary policy shocks do also affect core ILS rates. At longer horizons, such as the 5- and 10-year maturities, the differences between the effect on headline and core ILS rates are smaller: we estimate that a monetary policy tightening shock of one standard deviation leads to a decline in 5-year headline and core ILS rates of about 3 and 1 basis points, respectively. For 10-year ILS rates, the same numbers are 2 and 1/2 basis points for headline and core. That the difference in the effect decreases with the horizon may indicate that the effect on food and energy inflation may die out more quickly than the effect on underlying inflation.

The effects of monetary policy shocks on headline and core ILS rates may reflect either genuine expectations or inflation risk premia. It could be that the response of headline ILS

⁸A one standard deviation shock correspond to OIS rates increasing by about 3 basis points across the term structure.

rates mostly reflects headline inflation risk premia, while the response of core ILS rates mostly reflects genuine core expectations. The two lower panels of table 6 separate out the effects into the expectation and risk premia components. We find that the headline inflation risk premium indeed does respond more strongly than the core inflation risk premium. This is particularly the case for the shortest maturities, but it is also true for the medium- and long horizons. That said, this does not fully explain the difference in the response of headline and core ILS rates. We still find that genuine headline inflation expectations responds more than genuine core inflation expectations for the 1-year horizon, while the difference is negligible at the medium- and long horizons.

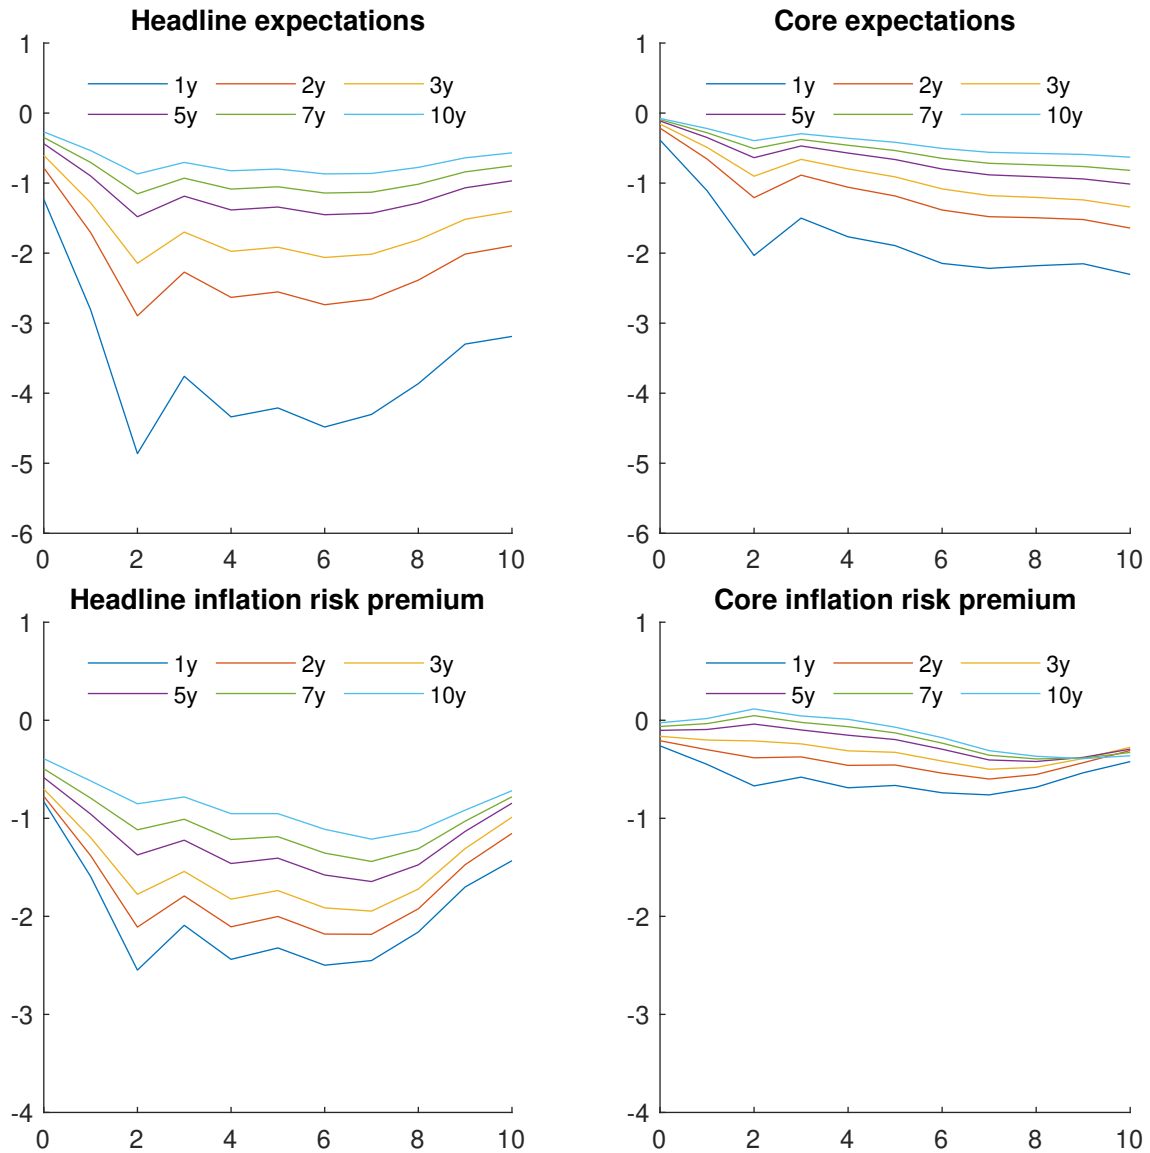
Figure 11 reports the results for varying choices of h . The main take away is that the results remain robust to different choices and that the baseline choice of $h = 3$ remains reasonable. The figure is also useful for judging the persistence of the effects of monetary policy shocks on inflation expectations. If the effects on inflation expectations died out quickly, over the course of a couple of days, say, it would be unlikely to ultimately have much significance for wider macroeconomic outcomes. But the figure suggests that the effects on inflation expectations, both headline and core, are rather persistent and still appear to last at least 10 business days after the initial shock. Although the signal gets noisier the wider the h -window becomes, this does provide suggestive evidence that monetary policy shocks have lasting effects on inflation expectations.

Table 6: Estimated effect of a one standard deviation monetary policy shock

		Headline			Core		
		1Y	5Y	10Y	1Y	5Y	10Y
ILS	$\hat{\beta}$	-7.41	-2.85	-1.72	-2.70	-0.68	-0.28
	t-stat	-4.13	-3.66	-3.58	-3.70	-1.79	-0.78
	R^2	0.35	0.30	0.29	0.30	0.09	0.02
Expectations	$\hat{\beta}$	-4.86	-1.48	-0.87	-2.03	-0.64	-0.39
	t-stat	-4.08	-4.22	-4.17	-2.83	-1.93	-1.99
	R^2	0.35	0.36	0.36	0.19	0.09	0.10
Risk premia	$\hat{\beta}$	-2.55	-1.37	-0.85	-0.67	-0.04	0.12
	t-stat	-3.75	-2.81	-2.70	-3.25	-0.21	0.53
	R^2	0.31	0.20	0.19	0.26	0.00	0.01

Note: Monetary policy shocks are the first principal component spanning changes in 1-, 3-, and 6-months as well as 1-10 year OIS rates during policy events. The series of shocks are normalised to have a unit standard deviation. The sample spans monetary policy events from September 2016 to June 2023, excluding information shocks (yields and stock prices moving in the same direction). The final number of events is 32. The responses are estimates of β in equation 18 with $h = 3$.

Figure 11: Estimated effects of a monetary policy shock on expectations and risk premia



Note: Monetary policy shocks are the first principal component spanning changes in 1-, 3-, and 6-months as well as 1-10 year OIS rates during policy events. The series of shocks are normalised to have a unit standard deviation. The sample spans monetary policy events from September 2016 to June 2023, excluding information shocks (yields and stock prices moving in the same direction). The final number of events is 32. The responses are estimates of β in equation 18 up to 10 trading days after the monetary policy event.

7 Conclusions

This paper proposes a no-arbitrage term structure model of observable headline ILS rates, which we use to estimate the swap rates on hypothetical euro area core ILS contracts. The key assumption underlying our approach is that traded headline ILS rates span core inflation, which

should be reasonably uncontroversial as core inflation is a sub-component of headline inflation.

We fit the model to euro area headline ILS rates, realized headline and core inflation, and both headline and core inflation expectations reported in the SPF. It generally achieves a good fit to all three types of input data. Likewise, the model produces plausible estimates for the level of both long-term forward headline and core inflation expectations. Both inflation measures hover around levels close to 2 percent, in line with the ECB's inflation target. Our model-implied core ILS rates have mostly evolved in the same direction as headline ILS rates, but exhibit lower volatility. Historically, our core ILS rates were below 2 percent before and during the early stages of the pandemic, before increasing to above 2 percent during 2021, when inflation pressures built. During 2023, core ILS rates have declined only slowly and less than headline ILS rates. Decomposing our core ILS rates into genuine core inflation expectations and core inflation risk premia shows that shorter maturities mainly reflect core inflation expectations, while the core inflation risk premium matters relatively more for longer maturities.

One key advantage of our core inflation expectations measure is that it is available at a daily frequency. In contrast, other measures of core inflation expectations, such as from surveys or macroeconomic projections, are only available monthly or even quarterly. The high-frequency nature of our measure is useful to, for example, policymakers that follow the economy in real-time. The high-frequency nature is also useful because it facilitates event-studies. We conduct an event-study, estimating the causal effect of monetary policy shocks on core inflation expectations. Our results suggest that a monetary policy tightening surprise significantly lowers near-term core inflation expectations, although less so than it lowers headline inflation expectations.

References

- Abrahams, M., T. Adrian, R. K. Crump, E. Moench, and R. Yu (2016). Decomposing real and nominal yield curves. *Journal of Monetary Economics* 84(C), 182–200.
- Adrian, T., R. K. Crump, and E. Moench (2013). Pricing the term structure with linear regressions. *Journal of Financial Economics* 110(1), 110–138.
- Ajello, A., L. Benzoni, and O. Chyruk (2020). Core and 'crust': Consumer prices and the term structure of interest rates. *Review of Financial Studies* 2019, 3719–3765.
- Altavilla, C., L. Brugnolini, R. S. Gürkaynak, R. Motto, and G. Ragusa (2019). Measuring euro area monetary policy. *Journal of Monetary Economics* 108, 162–179.
- Anderson, B., P. J. Hammond, and C. A. Ramezani (2010). Affine models of the joint dynamics of exchange rates and interest rates. *Journal of Financial and Quantitative Analysis* 45(5), 1341–1365.
- Ang, A., G. Bekart, and M. Wei (2008). The term structure of real rates and expected inflation. *Journal of Finance* 63(2), 797–849.
- Bahaj, S., R. Czech, S. Ding, and R. Reis (2023). The market for inflation risk. Bank of England Staff Working Paper 1028.
- Bauer, M. D. and G. D. Rudebusch (2017). Interest rates under falling stars. Working Paper Series 2017–16.
- Bauer, M. D., G. D. Rudebusch, and C. J. Wu (2012). Correcting estimation bias in dynamic term structure models. *Journal of Business and Economic Statistics* 30(3), 454–467.
- Camba-Mendez, G. and T. Werner (2017). The inflation risk premium in the post-Lehman period. ECB Working Paper 2033.
- Christensen, J. H. E., F. X. Diebold, and G. D. Rudebusch (2009). An arbitrage-free generalized Nelson-Siegel term structure model. *Econometrics Journal* 12(3), 33–64.
- Christensen, J. H. E., F. X. Diebold, and G. D. Rudebusch (2011). The affine arbitrage-free class of Nelson-Siegel term structure models. *Journal of Econometrics* 164(1), 4–20.
- Christensen, J. H. E., J. A. Lopez, and G. D. Rudebusch (2010). Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields. *Journal of Money, Credit and Banking* 42(s1), 143–178.
- Dai, Q. and K. J. Singleton (2002). Specification analysis of affine term structure models. *Journal of Finance* 55(5), 1943–1978.
- D'Amico, S., D. H. Kim, and M. Wei (2018). Tips from TIPS: the informational content of Treasury inflation-protected security prices. *Journal of Financial and Quantitative Analysis*.
- Duffee, G. R. (2002). Term premia and interest rate forecasts in affine models. *Journal of Finance* 57(1), 405–443.
- Duffee, G. R. (2011). Information in (and not in) the term structure. *Review of Financial Studies* 24(9), 2895–2934.
- Duffee, D. and R. Kan (1996). A yield-factor model of interest rates. *Mathematical Finance* 6(4), 379–406.
- Gürkaynak, R. S., B. Sack, and E. Swanson (2005). Do actions speak louder than words? the response of asset prices to monetary policy actions and statements. *International Journal of Central Banking* 1(1).
- Hamilton, J. D. and J. C. Wu (2012). Identification and estimation of gaussian affine term structure models. *Journal of Econometrics* 168(2), 315–331.
- Jarociński, M. and P. Karadi (2020). Deconstructing monetary policy surprises - the role of information shocks. *American Economic Journal: Macroeconomics* 12(2), 1–43.
- Joslin, S., M. Priebsch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *Journal of Finance* 69(3), 1197–1233.
- Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on Gaussian dynamic term structure models. *Review of Financial Studies* 24(3), 926–970.

- Kim, D. H. (2009). Challenges in macro-finance modeling. Federal reserve bank of st. louis review.
- Kim, D. H. and A. Orphanides (2012). Term structure estimation with survey data on interest rate forecasts. *Journal of Financial and Quantitative Analysis* 47(1), 241–272.
- Kim, D. H. and J. H. Wright (2005). An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates. Finance and Economic Discussion Series, Board of Governors of the Federal Reserve System 2005–33.
- Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics* 47(3), 523–544.
- Martin, F. E., A. Meldrum, and W. Yan (2021). No-arbitrage pricing of gdp-linked bonds. *Journal of Banking and Finance* 126.
- Meldrum, A., M. Raczko, and P. Spencer (2023). The information in joint term structures of bond yields. *Journal of International Money and Finance* 134.
- Mouabbi, S., J.-P. Renne, and J.-G. Sahuc (2021). Debt-stabilizing properties of gdp-linked securities: A Macro-Finance perspective. Banque de France Working Paper 844.

Appendix A. Deriving Pricing Recursions

Starting from equation (1), we have that

$$\begin{aligned} i_{t,h} &= \frac{1}{h} \log \mathbb{E}_t \left[e^{\xi_{t,t+h} + \pi_{t,t+h}} \right] \leftrightarrow \\ e^{h \cdot i_{t,h}} &= \mathbb{E}_t \left[e^{\xi_{t,t+h} + \pi_{t,t+h}} \right]. \end{aligned}$$

Because $\xi_{t,t+h}$ and $\pi_{t,t+h}$ are both "additive", in the sense that $\xi_{t,t+h} = \xi_{t,t+1} + \xi_{t+1,t+h}$ and $\pi_{t,t+h} = \pi_{t,t+1} + \pi_{t+1,t+h}$, we can then write the pricing equation as

$$e^{h \cdot i_{t,h}} = \mathbb{E}_t \left[e^{\xi_{t,t+1} + \pi_{t,t+1} + \xi_{t+1,t+h} + \pi_{t+1,t+h}} \right].$$

We can then use the law of iterated expectations to write

$$\begin{aligned} e^{h \cdot i_{t,h}} &= \mathbb{E}_t \left[\mathbb{E}_{t+1} \left[e^{\xi_{t,t+1} + \pi_{t,t+1} + \xi_{t+1,t+h} + \pi_{t+1,t+h}} \right] \right] \\ &= \mathbb{E}_t \left[e^{\xi_{t,t+1} + \pi_{t,t+1}} \mathbb{E}_{t+1} \left[e^{\xi_{t+1,t+h} + \pi_{t+1,t+h}} \right] \right] \\ &= \mathbb{E}_t \left[e^{\xi_{t,t+1} + \pi_{t,t+1}} e^{(h-1) \cdot i_{t+1,h-1}} \right] = \mathbb{E}_t \left[e^{\xi_{t,t+1} + \pi_{t,t+1} + (h-1) \cdot i_{t+1,h-1}} \right]. \end{aligned}$$

Then, we can plug in equations (4) and (6), to have

$$e^{h \cdot i_{t,h}} = \mathbb{E}_t \left[e^{\gamma_0 + \gamma_x X_{t+1} + \gamma_u u_{t+1} - \frac{1}{2} \lambda'_t \lambda_t - \lambda'_t \varepsilon_{t+1} + (h-1) \cdot i_{t+1,h-1}} \right].$$

Using equation (4), we have

$$\begin{aligned} e^{h \cdot i_{t,h}} &= \mathbb{E}_t \left[e^{\gamma_0 + \gamma_x \mu + \gamma_x \Phi X_t + \gamma_x \Sigma \varepsilon_{t+1} + \gamma_u u_{t+1} - \frac{1}{2} \lambda'_t \lambda_t - \lambda'_t \varepsilon_{t+1} + (h-1) \cdot i_{t+1,h-1}} \right] \\ &= e^{\gamma_0 + \gamma_x \mu + \gamma_x \Phi X_t - \frac{1}{2} \lambda'_t \lambda_t} \mathbb{E}_t \left[e^{\gamma_x \Sigma \varepsilon_{t+1} + \gamma_u u_{t+1} - \lambda'_t \varepsilon_{t+1} + (h-1) \cdot i_{t+1,h-1}} \right]. \end{aligned}$$

Now, assume that $(h-1) i_{t,h-1} = a_{h-1} + b_{h-1} X_t$. Plug this in to show that $h \cdot i_{t,h}$ is then also affine in the factors X_t . It follows that

$$\begin{aligned} e^{h \cdot i_{t,h}} &= e^{\gamma_0 + \gamma_x \mu + \gamma_x \Phi X_t - \frac{1}{2} \lambda'_t \lambda_t} \mathbb{E}_t \left[e^{\gamma_x \Sigma \varepsilon_{t+1} + \gamma_u u_{t+1} - \lambda'_t \varepsilon_{t+1} + a_{h-1} + b_{h-1} X_{t+1}} \right] \\ &= e^{\gamma_0 + \gamma_x \mu + \gamma_x \Phi X_t - \frac{1}{2} \lambda'_t \lambda_t} \mathbb{E}_t \left[e^{\gamma_x \Sigma \varepsilon_{t+1} + \gamma_u u_{t+1} - \lambda'_t \varepsilon_{t+1} + a_{h-1} + b_{h-1} \mu + b_{h-1} \Phi X_t + b_{h-1} \Sigma \varepsilon_{t+1}} \right] \\ &= e^{a_{h-1} + (\gamma_x + b_{h-1}) \mu + \gamma_0 + (\gamma_x + b_{h-1}) \Phi X_t - \frac{1}{2} \lambda'_t \lambda_t} \mathbb{E}_t \left[e^{\gamma_x \Sigma \varepsilon_{t+1} + \gamma_u u_{t+1} - \lambda'_t \varepsilon_{t+1} + b_{h-1} \Sigma \varepsilon_{t+1}} \right]. \end{aligned}$$

As ε_{t+1} and u_{t+1} are both standard normal and independent, we get that

$$\begin{aligned} e^{h \cdot i_{t,h}} &= e^{a_{h-1} + (\gamma_x + b_{h-1}) \mu + \gamma_0 + (\gamma_x + b_{h-1}) \Phi X_t - \frac{1}{2} \lambda'_t \lambda_t + \frac{1}{2} \gamma_u^2 + \frac{1}{2} ((\gamma_x + b_{h-1}) \Sigma - \lambda'_t) ((\gamma_x + b_{h-1}) \Sigma - \lambda'_t)'} \\ &= e^{a_{h-1} + (\gamma_x + b_{h-1}) \mu + \gamma_0 + (\gamma_x + b_{h-1}) \Phi X_t + \frac{1}{2} \gamma_u^2 + \frac{1}{2} (\gamma_x + b_{h-1}) \Sigma \Sigma' (\gamma_x + b_{h-1})' - (\gamma_x + b_{h-1}) \Sigma \lambda_t}. \end{aligned}$$

Then, using equation (7), we have

$$\begin{aligned} e^{h \cdot i_{t,h}} &= e^{a_{h-1} + (\gamma_x + b_{h-1}) \mu + \gamma_0 + (\gamma_x + b_{h-1}) \Phi X_t + \frac{1}{2} \gamma_u^2 + \frac{1}{2} (\gamma_x + b_{h-1}) \Sigma \Sigma' (\gamma_x + b_{h-1})' - (\gamma_x + b_{h-1}) (\lambda_0 + \lambda_x X_t)} \\ &= e^{a_{h-1} + (\gamma_x + b_{h-1}) (\mu - \lambda_0) + \gamma_0 + \frac{1}{2} \gamma_u^2 + \frac{1}{2} (\gamma_x + b_{h-1}) \Sigma \Sigma' (\gamma_x + b_{h-1})' + (\gamma_x + b_{h-1}) (\Phi - \lambda_x) X_t}. \end{aligned}$$

Hence, matching coefficient on the right- and left-hand sides of the equality, we have that $h \cdot i_{t,h} = a_h + b_h X_t$, where

$$b_h = (\gamma_x + b_{h-1}) (\Phi - \lambda_x)$$

$$a_h = a_{h-1} + (\gamma_x + b_{h-1}) (\mu - \lambda_0) + \gamma_0 + \frac{1}{2} \gamma_u^2 + \frac{1}{2} (\gamma_x + b_{h-1}) \Sigma \Sigma' (\gamma_x + b_{h-1})'.$$

That is, the swap rate for maturity h is affine in the factors if the swap rate for maturity $h-1$ is affine in the factors. This is the case, because the swap rate with $h=0$ is affine in the factors: it must have $a_0 = 0$ and $b_0 = 0_{(1 \times n_x)}$, since $h \cdot i_{t,h} = 0$ for all X_t when $h=0$.

Publications



NEWS

News may be a news article or an appetiser offering quick insight into one of Danmarks Nationalbank's more extensive publications. The series is targeted at people who need an easy overview and like a clear angle.



STATISTICAL NEWS

Statistical news focuses on the latest figures and trends in Danmarks Nationalbank's statistics. Statistical news is targeted at people who want quick insight into current financial data.



REPORT

Reports consist of recurring reports on Danmarks Nationalbank's areas of work and activities. Here you will find Danmarks Nationalbank's annual report, among other documents. Reports are targeted at people who need a status and update on the past period.



ANALYSIS

Analyses focus on current issues of particular relevance to Danmarks Nationalbank's objectives. The analyses may also contain Danmarks Nationalbank's recommendations. They include our projections for the Danish economy and our assessment of financial stability. Analyses are targeted at people with a broad interest in economic and financial matters.



ECONOMIC MEMO

Economic Memos provide insight into the analysis work being performed by Danmarks Nationalbank's employees. For example, Economic Memos contain background analyses and method descriptions. Economic Memos are primarily targeted at people who already have a knowledge of economic and financial analyses.



WORKING PAPER

Working Papers present research work by both Danmarks Nationalbank's employees and our partners. Working Papers are primarily targeted at professionals and people with an interest in central banking research as well as economics and finance in a broader sense.

Danmarks Nationalbank
Langelinie Allé 47
2100 Copenhagen Ø
+45 3363 6363



**DANMARKS
NATIONALBANK**