

Inflation (De-)Anchoring in the Euro Area

Valentin Burban¹, Bruno De Backer², Andreea Liliana Vladu³

September 2024, WP #965

ABSTRACT

This article measures the degree of potential de-anchoring of inflation expectations in the euro area vis-à-vis the inflation objective of the European Central Bank (ECB). A no-arbitrage term structure model that allows for a time-varying long-term mean of inflation expectations, π_t^* , is applied to inflation-linked swap (ILS) rates, while taking into account survey-based inflation forecasts. Estimates of π_t^* have been close to 2% since the mid-2000s, indicating that long-term inflation expectations have overall remained well anchored to the ECB's inflation objective. As this objective is however related to the "medium term", expectations components of various forward ILS rates are extracted: they appear to have been broadly anchored, with tentative signs of de-anchoring up to the two-year horizon. Using backcasted ILS rates, estimates of π_t^* are much above 2% in the early 1990s, but they converge to levels below 2% by the end of the decade when the ECB was established.

Keywords: Inflation-Linked Swap Rates, Surveys, No-Arbitrage, Shifting Endpoint, Inflation Expectations.

JEL classification: E31, E43, E47, E58.

¹ Banque de France and Aix-Marseille University, valentin.burban@banque-france.fr

² National Bank of Belgium, bruno.debacker@nbb.be

³ European Central Bank, andreea.liliana.vladu@ecb.europa.eu

We are grateful to Marie-Laure Barut-Etherington, Michael Bauer, Matthieu Bussière, Jens Christensen, Mátyás Farkas, Luís Fonseca, Olivier Garnier, Felix Geiger, Adriana Grasso, Guillaume Horny, Leonardo Iania, Kasper Jørgensen, Wolfgang Lemke, Magali Marx, Julien Matheron, Sarah Mouabbi, Alain Naef, Adrian Penalver, Luca Rossi, Glenn Rudebusch, Fabian Schupp, Aino Silvo, Pavel Tretiakov, Joris Wauters, Raf Wouters and to the participants in the IFABS 2022 Naples conference, the 2023 ECB workshop on extracting inflation expectations and risk premia from market measures, the 2023 SoFiE summer school in Brussels, the 2024 advanced workshop for central bankers, the first workshop of the research network on the "Challenges for Monetary Policy Transmission in a Changing World", and seminars at the European Central Bank, Banque de France and National Bank of Belgium for their comments and suggestions. We would also like to thank Mattia Alfero and Sofia Gori for research assistance. The views expressed do not necessarily reflect those of the European Central Bank, the Banque de France, the National Bank of Belgium or the Eurosystem.

NON-TECHNICAL SUMMARY

The anchoring of inflation expectations is fundamental to ensure price stability, and macroeconomic stability more in general. As a result of the lessons from history (such as the costly re-anchoring of inflation expectations at low levels in the 1980s) and modern institutional set-up (inflation-targeting, independent central banks), inflation expectations are closely monitored, in particular medium-term inflation expectations as they are a direct measure of the credibility of the monetary policy of central banks. The post-pandemic recovery triggered renewed discussions on de-anchoring inflation expectations in a context of inflation rates last seen in the 1970s, while at the same time the ECB 2021 strategy review highlighted "the need for a comprehensive framework for assessing (un)anchoring" of inflation expectations.


Speaking directly to the question of inflation de-anchoring in the euro area, this article proposes a no-arbitrage affine term structure model for inferring inflation expectations from market-based measures of inflation compensation. The novel feature of the model is that the infinite-horizon expected inflation rate is time varying (π_t^*). By contrast, standard stationary models feature a fixed endpoint (π^*), i.e. forecasts of the inflation rate are assumed to converge at any point in time to the same value.

The model proposed in this article also considers the ECB's survey of professional forecasters and the Consensus Economics survey when estimating inflation expectations. Jointly, these assumptions impact the decomposition of short-, medium- and longer-term euro area inflation-linked swap (ILS) rates into expectations and inflation risk premia. The baseline dataset of ILS rates used in the estimation spans the period going from June 2005 to December 2023. The question of whether there is evidence of a decrease in the volatility of the inflation trend with the ECB's establishment in 1998 is analysed with a longer dataset of backcasted ILS rates going back to 1992.

Overall, results show that the estimated long-term measure of inflation expectations has hardly ever been de-anchored in the euro area since 2005. Estimates go down to close to 1.80% in 2016 and reach again a similar trough during COVID-19. Expectations components of short-term ILS rates, such as the one-year rate, are not much affected by π_t^* , whereas those of longer-term (forward) rates are, including the one-year forward rate four years ahead (1y4y ILS rate) and the five-year forward rate five years ahead (5y5y). At these two horizons, inflation expectations have remained in line with the ECB objective. However, tentative signs of de-anchoring appear in the inflation expectations component of the one-year forward ILS rate two years ahead (1y2y), a relevant horizon as it broadly corresponds to the horizon of the Eurosystem/ECB staff projections.

Summing up, this paper addresses the challenge of reliably estimating medium-term inflation expectations, which is central for the determination of (optimal) monetary policy. With inflation expectations broadly well anchored in the medium term, i.e. with a credible medium-term inflation objective, monetary policy does not need to be particularly restrictive in order to bring inflation back to target. A soft landing of the economy can therefore be envisaged. But if medium-term inflation expectations were to de-anchor, monetary policy might have to be significantly more restrictive to bring inflation expectations back in line with the objective, which could happen at the expense of an economic recession.

Figure: Estimated π_t^* , inflation-linked swap (ILS) and surveys



(Dés-)ancrage de l'inflation dans la zone euro

RÉSUMÉ

Cet article mesure le degré de désancrage potentiel des anticipations d'inflation dans la zone euro par rapport à l'objectif d'inflation de la Banque centrale européenne (BCE). Un modèle de structure par terme des taux, faisant l'hypothèse d'absence d'opportunité d'arbitrage, et introduisant une mesure, variable dans le temps, des anticipations d'inflation à long terme, π_t^* , est appliqué aux taux de swaps indexés sur l'inflation (ILS), tout en tenant compte des prévisions d'inflation basées sur les enquêtes. Les estimations de π_t^* sont proches de 2% depuis le milieu des années 2000, indiquant que les anticipations d'inflation à long terme sont globalement restées bien ancrées à l'objectif d'inflation de la BCE. Comme cet objectif est cependant lié au "moyen terme", les composantes des anticipations de divers taux forward ILS sont extraites : elles semblent avoir été globalement ancrées, avec de faibles signes de désancrage jusqu'à l'horizon de deux ans. En utilisant des taux ILS rétropolés, les estimations de π_t^* sont bien supérieures à 2% au début des années 1990, mais convergent vers des niveaux inférieurs à 2% à la fin de la décennie lorsque la BCE a été établie.

Mots-clés : taux de swaps d'inflation, enquêtes, absence d'arbitrage, point final mouvant, anticipations d'inflation.

Les Documents de travail reflètent les idées personnelles de leurs auteurs et n'expriment pas nécessairement la position de la Banque de France. Ils sont disponibles sur publications.banque-france.fr

1 Introduction

"The firm anchoring of inflation expectations is critical under any circumstances, as it ensures that temporary movements in inflation do not feed into wages and prices and hence become permanent."

Mario Draghi, 21 November 2014

The anchoring of inflation expectations is fundamental to ensure price stability, and macroeconomic stability more in general, as theorised and shown by history. In the authoritative new Keynesian framework, inflation expectations affect the real economy through at least two channels: inflation expectations impact real interest rates (Fischer equation) and, to the extent that prices are sticky, firms form prices in a forward-looking way, implying that inflation expectations have a direct influence on current inflation (new Keynesian Phillips curve). Looking at history, the US experience of the 1970s-1980s is a prime example reminding that bringing inflation expectations down can imply significant monetary policy tightening, possibly at the cost of economic recessions. Since then, it is well accepted that a central bank can engineer a soft landing of the economy only if it is credible about its inflation objective, i.e. inflation expectations are anchored to the central bank's target. In the shorter history of the euro area, the experience of the 2010s, with inflation remaining persistently below the ECB's target, prompted increased attention to inflation expectations (hence the quote above by Mario Draghi). More recently, the post-pandemic recovery triggered renewed discussions on inflation expectations in a context of inflation rates last seen in the 1970s.

As a result of historical lessons, and in line with theory, inflation-targeting central banks monitor inflation expectations closely, in particular medium-term inflation expectations as they are a direct measure of the credibility of their monetary policy. In the euro area, the ECB used the discretion it enjoys to adjust the quantitative definition of its primary objective of price stability: inflation "below 2%" as of October 1998; "below, but close to 2%" as of May 2003, and "2%" as of July 2021, but the definition was always accompanied by the specification that the objective referred to the "medium term". This specification reflects the inevitability of short-term deviations of inflation from the target and the uncertainty in the transmission of monetary policy to the economy. It also takes into account that the appropriate monetary policy response

to a deviation of inflation from the target is context-specific.

The main aim of this article is to provide a model-based aggregate of several measures of medium-term inflation expectations. Feeding into the Governing Council's deliberations on the 2021 ECB strategy review, ECB (2021) highlighted "the need for a comprehensive framework for assessing (un)anchoring". One practical issue in monitoring medium-term inflation expectations is indeed that they can be measured in (many) different ways. Market-based measures might refer to inflation-linked swap (ILS) rates, break-even inflation rates, or fixings (short-term inflation swaps attached to specific dates). Expectations can also be measured via surveys sounding out professional forecasters, market or economic analysts, consumers or firms. While much work remains to be done, this article constitutes a first answer to ECB (2021) by providing a model encompassing both market-based and survey-based measures.

More precisely, the model estimates a time-varying measure of long-term inflation expectations, π_t^* (" π -star"), and estimates the inflation expectations components of ILS rates, while taking into account the ECB's survey of professional forecasters and the Consensus Economics survey. The expectations components entailed in different (forward) ILS rates speak to different views on what the "medium term" orientation of the ECB's inflation objective means, be it two, four or ten years ahead. The model relates ILS rates across maturities by no-arbitrage restrictions.

Overall, results show that the estimated long-term measure of inflation expectations has hardly ever been de-anchored in the euro area since 2005 (start of ILS data). Estimates go down to close to 1.80% in 2016 and reach again a similar trough during COVID-19. Expectations components of short-term ILS rates, such as the one-year rate, are not much affected by π_t^* , whereas those of longer-term (forward) rates are, including for the one-year forward rate four years ahead (1y4y ILS rate) and the five-year forward rate five years ahead (5y5y). At these two horizons, inflation expectations have always remained in line with the ECB objective. However, tentative signs of de-anchoring appear in the expectations component of the one-year forward ILS rate two years ahead (1y2y), a relevant horizon as it broadly corresponds to the horizon of the Eurosystem/ECB staff projections.

This article can be related to two strands of the literature: one on inflation anchoring/trends in the euro area; and the other on term structure models. The literature on inflation anchoring/trends looks at a range of data and resorts not only to a concept of anchoring in levels (e.g. close to 2% or not), but also in changes (responsiveness of inflation expectations to macroe-

conomic news or shocks), and in distributions (e.g. distributions of individuals responses in surveys or option-implied densities). It highlights various levels of concerns depending on the period analysed. Beechey et al. (2011) compare the evolution of long-run inflation expectations in the euro area and the United States, using evidence from financial markets and surveys of professional forecasters. They find that long-run inflation expectations are reasonably well anchored in both economies and that both sources of information suggests firmer anchoring in the euro area than in the United States (data from the 2000s). Similarly, Hördahl and Tristani (2014) find that, after correcting for liquidity and inflation risk premia, long-term inflation expectations extracted from French and US break-even inflation rates have remained remarkably stable at the peak of the financial crisis and throughout the Great Recession (data up to 2013). A few years later, Łyziak and Paloviita (2017) analyse the anchoring of inflation expectations of professional forecasters and consumers in the euro area to find that longer-term inflation expectations have become somewhat more sensitive to shorter-term ones and to actual inflation in the 2010s, suggesting some signs of de-anchoring. In line with these signs, Grishchenko et al. (2019) propose a measure of the anchoring of inflation expectations based on surveys that accounts for inflation uncertainty, and obtain results suggesting that, following the Great Recession, inflation anchoring improved in the United States, while mild de-anchoring occurred in the euro area. In the same vein, García and Werner (2021) find that the anchoring of euro area inflation expectations has weakened since 2013, and Corsello et al. (2021) conclude that long-term inflation expectations have de-anchored from the ECB's inflation aim as, among other things, the long-term expectations reported in the ECB's survey of professional forecasters have not regained the levels that prevailed before the 2013-14 disinflation. Brand et al. (2021) develop a semi-structural macroeconomic model with arbitrage-free yield-curve dynamics to estimate π_t^* , inter alia, which they find to be time-varying but overall relatively stable around 2% in the 2000s and 2010s.

Two recent papers take a similar approach to the present one to model euro area inflation expectations. Cecchetti et al. (2022) introduce a model that features time-varying long-term inflation with market-based risk-adjusted measures of expected inflation anchored to survey-based inflation forecasts (the ECB's survey of professional forecasters). Compared to the model presented below, their model features time-varying inflation volatility, is estimated based on options prices, and the identification of factors is carried out differently. Results also differ: they find a more volatile medium-term risk-adjusted expected inflation, falling to 0.8% in 2016

and being below 1% from mid-2019 to early 2021. In addition, Boeckx et al. (2024) develop a macro-finance model providing estimates of π_t^* that come close to those reported below. While they go further into the analysis of ILS rates, e.g. by carrying out structural decompositions, the present paper looks deeper into the implications of modelling choices for estimates of π_t^* and expectations components.

Regarding the literature on term structure models, the main challenge is to introduce a so-called shifting endpoint so that inflation expectations do not always converge to a constant long-run mean. Early work in this respect can be found in Kozicki and Tinsley (2001) and Dewachter and Lyrio (2006) who show the importance of long-run inflation expectations for (long-run) bond yields. Building on this work, Dewachter and Iania (2011) propose a macro-finance model of the US economy with shifting endpoints. However, the modelling approach of Bauer and Rudebusch (2020) is the one pursued below as it builds on the canonical Gaussian dynamic term structure model of Joslin et al. (2011) in which the pricing factors are observable portfolios of yields (ILS rates in our case), greatly facilitating model estimation. As shown in Camba-Méndez and Werner (2017), this canonical model can also be applied to ILS rates (see also Burbán et al. 2022). Bauer and Rudebusch (2020) model the shifting endpoint as an unspanned factor, as in Joslin et al. (2014). Besides, the model presented below relates to survey forecasts that have been shown to bring useful additional information to term structure models (Kim and Orphanides, 2012).

The article is organised as follows. Section 2 describes the ILS and survey data. Section 3 establishes stylised facts that are then used to guide modelling choices described in Section 4. Section 5 presents the main results, i.e. estimates of π_t^* and expectations components, while Section 6 breaks down these results along the building blocks of the model. Section 7 provides robustness tests. Section 8 concludes.

2 Data

This section introduces the market-based measures of inflation compensation and survey-based forecasts of inflation used to estimate the model. Unlike break-even inflation rates derived from inflation-linked and nominal sovereign bonds issued by specific euro area member states, ILS contracts refer to the euro area as a whole and are considered to be less prone to liquidity distortions (ECB, 2021). Compared to ILS rates, survey-based forecasts do not entail inflation

risk premia (allowing to consider them representative of "genuine" inflation expectations), but are observed less frequently, refer to a relatively small group of respondents and might suffer from other issues implying a misrepresentation of expectations. We focus on forecasts collected from Consensus Economics and the ECB's survey of professional forecasters.

2.1 Inflation-linked swap rates

ILS contracts exchange a fixed rate against a to-be-realised (floating) inflation rate, with both rates applied to a notional amount. Only one (net) payment is made at maturity, implying that ILS rates are zero-coupon rates. Standard no-arbitrage assumptions imply that the discounted expected pay-off received from the fixed-rate leg is equal to the discounted expected pay-off received from the floating-rate leg (so that the price of entering the swap is zero). Formally,

$$\tilde{\pi}_{t,M} = \sqrt[M]{\frac{E_t^{\mathbb{Q}} [I_{t+M \text{ years} - 3 \text{ months}}]}{I_{t-3 \text{ months}}}} - 1, \quad (1)$$

where $\tilde{\pi}_{t,M}$ is the annualised and discretely compounded end-of-month ILS rate of maturity M years observed at time (month) t , I_t is the reference price index, $I_{t+M \text{ years}}$ denotes (with somewhat loose notations) the reference price index in M years, and $E_t^{\mathbb{Q}}[\cdot]$ denotes the expectations operator under the risk-neutral probability measure \mathbb{Q} conditional on the information available up to time t . As per market convention, the reference price index is the euro area harmonised index of consumer prices excluding tobacco and it is lagged by three months (so that it is always observed when entering ILS contracts).¹ ILS rates $\tilde{\pi}_{t,M}$ are provided by the London Stock Exchange Group (LSEG) for maturities between one and ten years.²

In order to obtain forward-looking ILS rates, observed rates are adjusted for the indexation lag. In practice, this can be done by interpolating $E_t^{\mathbb{Q}} [I_t]$ (which is not observed at time t because of publication lags) from the observed $I_{t-3 \text{ months}}$ and $E_t^{\mathbb{Q}} [I_{t+1 \text{ year} - 3 \text{ months}}]$ implied by the 1y ILS rate (Equation 1). Further, $E_t^{\mathbb{Q}} [I_{t+M \text{ years}}]$ can be approximated from the implied $E_t^{\mathbb{Q}} [I_{t+M \text{ years} - 3 \text{ months}}]$ and $E_t^{\mathbb{Q}} [I_{t+(M+1) \text{ years} - 3 \text{ months}}]$, and so on and so forth for $E_t^{\mathbb{Q}} [I_{t+(M+2) \text{ years}}], \dots, E_t^{\mathbb{Q}} [I_{t+(M+10) \text{ years}}]$. This interpolation is done by log-linearisation on

¹The role played by tobacco in the dynamics of the harmonised index of consumer prices is marginal, so that the index excluding tobacco is virtually identical to the overall index.

²The mnemonics of the ILS data downloaded from LSEG are as follows: EUHCPT?Y_ICAP, where "?" must be replaced by the maturity of the rate (1, 2, 3, ...). End-of-month bid and ask annualised quotes are averaged to obtain time series of midpoint rates.

top of which a seasonal inflation pattern is added (the average pattern observed over the previous five years), following Camba-Méndez and Werner (2017, see Appendix A.1).

Computing ILS rates as continuously-compounded rates gives the following representation:

$$\pi_{t,M} = \frac{1}{M} \ln \left(\frac{E_t^{\mathbb{Q}} [I_{t+M \text{ years}}]}{E_t^{\mathbb{Q}} [I_t]} \right), \quad (2)$$

where $\pi_{t,M}$ is the continuously-compounded and indexation-lag-adjusted counterpart of $\tilde{\pi}_{t,M}$.

The left-side panel of Figure 1 displays the 1-year to 10-year continuously-compounded zero-coupon ILS rates adjusted for the indexation lag. The data provided by LSEG spans the period going from June 2005 to December 2023 (data to the right of the vertical dotted line in the figure), which constitutes the main sample used for the core analysis of this article. As of 2005, ILS rates are close to 2% but short-term rates drop significantly in late 2008 at the height of the global financial crisis (failure of Lehman Brothers). While short-term ILS rates had returned to close to 2% by the end of 2010, the cross-section of ILS rates then gradually declined to levels that would arguably be incompatible with the ECB's inflation target of, at the time, "below, but close to 2% over the medium term". For instance, the 5y5y ILS rate stood at about 1.5% by mid-2016. ILS rates recovered somewhat by mid-2018 but then slid again to low levels until the trough reached at the end of March 2020 in the midst of the COVID-19 crisis. ILS rates increased significantly over the next two years, with short-term ones reaching all-time highs in August 2022 before stabilizing to around 2.1% at the end of the sample.

A longer dataset of market-based inflation measures uses backcasted ILS rates from Burban and Schupp (2023) going back to 1992. Their backcasting procedure exploits the correlation between ILS rates and various economic variables.³ The longer backcasted time series allow to estimate more precisely, in principle, equilibrium quantities such as the long-term mean of ILS rates and to cover various policy regimes. Throughout the 1990s, ILS rates follow a downward path reflecting the trend of headline inflation across countries that would later be part of the euro area. ILS rates stabilised at levels close to 2% by the end of the decade. Acknowledging the possibility of inflation risk premia, these levels are compatible with the initial inflation objective of "below 2%" of the ECB.

³To ensure consistency along the cross-section, the 1-year, 5-year and 10-year ILS rates are backcasted and other maturity-specific ILS rates are interpolated.

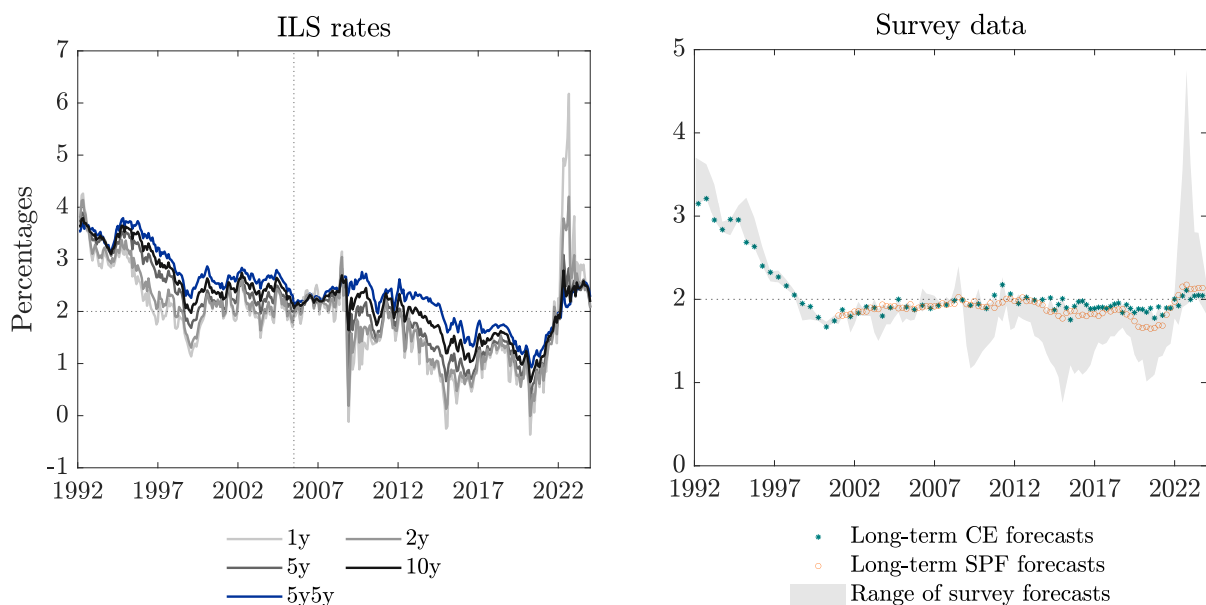


Figure 1: **ILS rates and inflation survey forecasts.** ILS rates are continuously compounded and adjusted for the indexation lag as in Camba-Méndez and Werner (2017), and backcasted before 2005 as in Burban and Schupp (2023) as highlighted by the dotted vertical line. The range of survey forecasts comprises forecasts from Consensus Economics (CE) and the ECB's SPF. "Long-term SPF forecasts" refer to average point forecasts for the annual inflation rate in the fourth or the fifth calendar year. "Long-term CE forecasts" refer to average point forecasts for the average annual inflation in 6 to 10 years time.

2.2 Inflation survey forecasts

We consider inflation survey forecasts from Consensus Economics (CE) and the ECB's Survey of Professional Forecasters (SPF). CE provides forecasts of the annual inflation rate for several calendar years, so-called "fixed-event" forecasts. Forecasts with horizons starting at the calendar year after next year and up to the fifth calendar year out are considered, amounting to four forecast horizons. Very short-term inflation forecasts (e.g. for the current and next calendar years) are discarded, so as to compare survey forecasts to purely model-based forecasts (avoiding keeping track of past real-time HICP realisations) and because these very short-term forecasts are less relevant to estimate the long-term level of inflation expectations. Besides, long-term forecasts for the average annual inflation rate in six to ten years time are also considered. These CE forecasts for the euro area as a whole are available since 2003, on a semiannual basis until 2014 and on a quarterly basis afterwards. Before 2003 and back until 1995, synthetic euro area inflation expectations are constructed based on GDP-weighted forecasts for the five largest euro area economies (Germany, France, Italy, Spain and the Netherlands), and for the three largest

economies back until 1992 (Germany, France, Italy).

From the ECB's SPF, average point forecasts for the fixed horizons of one and two years ahead are considered, as well as two fixed-event forecasts for the calendar year after the next and for the "long-run". This long-run horizon represents the annual inflation rate in the fourth calendar year for Q1 and Q2 vintages of the survey in a given year, and in the fifth calendar year for the Q3 and Q4 vintages. SPF forecasts are available at quarterly frequency since 2000.

The right-side panel of Figure 1 shows the time series of inflation survey forecasts we consider. Inflation survey forecasts show broadly similar dynamics to ILS rates, although fluctuations are less pronounced. Survey forecasts follow a downward trend up to the end of the 1990s. In subsequent years, long-term forecasts show only mild fluctuations at levels below but close to 2% until 2014, after which they gradually drop. Long-term survey forecasts start recovering as of mid-2020 to reach levels slightly above 2% by the end of 2023. Short-term inflation survey forecasts display more fluctuations than long-term forecasts, in particular during the COVID-19 pandemic when they reached historically high levels.

Overall, both ILS rates and survey-based forecasts provide tentative evidence of time variation in the long-run inflation expectations. Such variation can be modelled with a shifting endpoint for inflation expectations. Before introducing the model, the next section presents stylised facts of the euro area data.

3 Stylised facts on ILS rates and inflation trend proxies

Aiming to provide insights to guide the design of a dynamic term structure model in line with the data, this section analyses the statistical properties of ILS rates and of a range of observable proxies for the long-term inflation anchor, π_t^* , commonly used in the literature. A π_t^* proxy should reflect slow-moving dynamics of inflation expectations embedded in ILS contracts. ILS rates and observed π_t^* proxies are found to be cointegrated and, although the evidence is less clear-cut, π_t^* proxies appear to be unspanned by the cross-section of ILS rates.

3.1 Observed π_t^* proxies

The left-side panel of Figure 2 shows four observed π_t^* proxies based on survey forecasts and past inflation realisations used to derive a trend, along with the 5y5y ILS rate. Surveys include long-term forecasts from CE and the SPF. Since the long-term SPF horizon is either the fourth

or the fifth calendar year and the long-term CE horizon is six to ten years ahead, a third survey-based observed π_t^* proxy is considered by averaging the two long-term forecasts with the CE fifth calendar-year-ahead forecasts so as to entail all forecasts of at least five years (approximately). As survey forecasts are not available on a monthly basis, missing data points are linearly interpolated. Besides, an inflation trend is estimated following the Discounted Moving Average (DMA) approach of Cieslak and Povala (2015), based on monthly core inflation data (which is less volatile than headline inflation). The parametrization of the discounted moving average approach is presented in more details in Appendix A.1.




Figure 2: **Inflation trend proxies and cointegration residuals.** On the left panel, "DMA" stands for discounted moving average (Cieslak and Povala, 2015). The 5y5y ILS rate is observed as of June 2005 (right of vertical dashed line) and backcasted to 1992 based on Burban and Schupp (2023). On the right panel, cointegration residuals come from a regression of the 10y ILS rate on the observed π_t^* proxy taken as the average of long-term SPF and CE forecasts and the CE fifth calendar year forecasts.

3.2 Cointegration and common trend in ILS rates

ILS rates and observed π_t^* proxies are found to be non-stationary as the null hypothesis of a unit root fails to be rejected at conventional confidence levels by standard tests (unreported). Given the broadly similar tendencies displayed, ILS rates and observed π_t^* proxies appear as likely cointegrated. To illustrate this, the right-side panel of Figure 2 displays the residuals of the regression of the 10y ILS rate on one of the observed π_t^* proxies (the average of long-term SPF and CE forecasts and CE fifth calendar-year-ahead forecasts). The residuals remain relatively

| | Main sample, 2005-2023 | | | | Subsample, 2005-2019 | | | | Extended sample, 1992-2023 | | | |
|-----------------------------|------------------------|----------------------|----------------------|---------------------|----------------------|---------------------|----------------------|---------------------|----------------------------|----------------------|---------------------|--|
| | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | DMA | |
| Constant | -0.061*** (0.006) | -0.050*** (0.012) | -0.049*** (0.004) | 0.002 (0.002) | -0.060*** (0.011) | -0.025** (0.011) | -0.069*** (0.008) | -0.003* (0.002) | -0.008*** (0.003) | -0.008*** (0.003) | 0.010*** (0.001) | |
| π^* proxy | 4.135*** (0.326) | 3.517*** (0.590) | 3.557*** (0.237) | 1.192*** (0.155) | 4.078*** (0.578) | 2.250*** (0.564) | 4.628*** (0.438) | 1.554*** (0.112) | 1.499*** (0.113) | 1.450*** (0.110) | 0.655*** (0.038) | |
| R^2 | 60 | 27 | 73 | 66 | 44 | 11 | 69 | 80 | 59 | 53 | 79 | |
| $ADF_{OLS}Residual$ | -4.30*** | -3.51*** | -3.15*** | -3.20*** | -3.25*** | -2.48** | -2.46** | -4.87*** | -2.11** | -2.53** | -3.55*** | |
| $PP_{OLS}Residual$ | -3.72*** | -2.97*** | -3.91*** | -2.89*** | -2.78*** | -2.11** | -3.14*** | -4.39*** | -2.52** | -2.38** | -3.39*** | |
| $ADF_{DOLS}Residual$ | -2.01** | -1.89* | -1.64* | -3.08*** | -1.754* | -1.256 | -1.507 | -2.551** | -1.40 | -1.35 | -2.13** | |
| $PP_{DOLS}Residual$ | -2.00** | -1.95** | -1.65* | -3.11*** | -1.745* | -1.259 | -1.455 | -2.605*** | -1.40 | -1.39 | -2.14** | |
| Johansen trace $r = 0$ | 17.22** | 21.92*** | 12.94 | 28.86*** | 14.01* | 17.03** | 7.90 | 15.38* | 15.68** | 17.64** | 20.89*** | |
| Johansen trace $r = 1$ | 1.30 | 1.33 | 1.75 | 2.99* | 1.08 | 0.90 | 0.68 | 0.48 | 4.82** | 5.10** | 4.97** | |
| Johansen eigenvalue $r = 0$ | 15.92** | 20.59*** | 11.19 | 25.87*** | 12.93* | 16.13** | 7.22 | 14.89** | 10.86 | 12.55* | 15.92** | |
| Johansen eigenvalue $r = 1$ | 1.30 | 1.33 | 1.75 | 2.99* | 1.08 | 0.90 | 0.68 | 0.48 | 4.82** | 5.10** | 4.97** | |

*** p < 0.01, ** p < 0.05, * p < 0.1

Table 1: **Cointegration regressions and tests based on 10y ILS rates.** "DMA" stands for discounted moving average (Cieslak and Povala, 2015). Coefficients are based on regressions with Newey-West standard errors (using four lags) in parentheses. For the cointegration residuals from regressions estimated by OLS and dynamic OLS (DOLS), the second panel reports Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics. DOLS regressions include four leads and lags of the first difference of the 10y ILS rate and observed π_t^* proxy. The table also reports the Johansen trace and eigenvalue statistics, which tests whether the cointegration rank (r) among the ILS rates and π_t^* proxy is 0/1 against the alternative that it exceeds 0/1 (trace) or it equals 1/2 (eigenvalue). The number of lags for ADF and Johansen tests is selected based on the AIC. The data are observed at monthly frequency.

stable around zero, without any clear trend.

More formally, the results of the cointegration analysis for the 10y ILS rate are displayed in Table 1, while results for the 2y and 5y ILS rates are reported in Appendix A.2. Three sample periods are considered: (i) the main sample that spans the period going from 2005 to 2023; (ii) a subsample covering the years 2005-2019, i.e. stopping right before the COVID-19 outbreak in order to discard the volatility in short-term ILS rates observed during the pandemic; and (iii) an extended sample from 1992 to 2023 using backcasted ILS rates. As shown in the upper panel of the table, trend coefficient estimates are generally found to be (strongly significant and) greater than 1, as in Bauer and Rudebusch (2020). The lower panel of the table reports the results of cointegration tests. Augmented Dickey-Fuller and Phillips-Perron tests look at the stationarity of the residuals, and largely reject the null hypothesis of a unit root. This result is broadly supported by Johansen cointegration tests.

3.3 Spanning tests

If π_t^* is not fully explained by the term structure of ILS rates, then π_t^* is said to be unspanned by ILS rates. Two types of tests are carried out to verify the spanning hypothesis: regressions of observed π_t^* proxies on principal components of ILS rates; and regressions of excess returns on principal components and observed π_t^* proxies.

Following Joslin et al. (2014), observed π_t^* proxies are regressed on the first three principal components of ILS rates (which explain 99.9% of the total variation in ILS rates): coefficients of determination (R^2) of less than 100% would suggest that π_t^* is unspanned by ILS rates. Table 2 shows that R^2 are indeed generally found to be smaller than 100% across alternative observed π_t^* proxies and samples. Focusing on the main sample (2005-2023), for instance, the range of R^2 goes from 31 to 77%.

| | (1) | (2) | (3) | (4) |
|----------------------------|--------------|-----|-----|-----|
| | Avg. surveys | CE | SPF | DMA |
| Main sample, 2005-2023 | 62% | 31% | 77% | 68% |
| Subsample, 2005-2019 | 49% | 20% | 72% | 83% |
| Extended sample, 1992-2023 | 60% | 53% | | 81% |

Table 2: **Spanning tests as in Joslin et al. (2014)**. The table reports adjusted R^2 from regressions of observed π_t^* proxies on the first three principal components of ILS rates. The data are observed at monthly frequency. "Avg. surveys" refers to the average of long-term SPF and CE forecasts and the CE fifth calendar year forecasts.

As an additional, and probably more reliable test, observed π_t^* proxies are added as regressor to the first three principal components of ILS rates in regressions of excess returns, following Bauer and Hamilton (2018) and Bauer and Rudebusch (2017, 2020). Significant coefficients for observed π_t^* proxies would indicate that π_t^* helps to predict excess returns on top of the principal components, and that it is unspanned by ILS rates. The focus is on quarterly (overlapping) excess returns, as these returns suffer less from overlapping issues than annual returns, and as they show more pronounced links with the macroeconomy than monthly excess returns. Realised quarterly excess returns are computed as in Boeckx et al. (2024).⁴ The reverse regression approach of Wei and Wright (2013) is used, as in Cieslak and Povala (2015), in order to deal with the issue of overlapping excess returns.

Table 3 reports regressions estimates for excess returns averaged across maturities. Overall, results indicate that observed π_t^* proxies help to predict excess returns. Small-sample (one-sided) p-values computed as in Bauer and Hamilton (2018) are more often than not close to or below the 10% significance threshold. In particular, the coefficient of the observed π_t^* proxy is significant over the main sample when using an average across survey forecasts or only the long-term SPF forecasts. Besides, R^2 coefficients noticeably increase when an observed π_t^* proxy is added to the regression. Results for annual and monthly excess returns, reported in Appendix A.3, show broadly similar results although evidence in favour of an unspanned π_t^* is somewhat weaker. Besides, regressions of excess returns for maturity-specific ILS rates show that the coefficient of observed π_t^* proxies tend to be more significant for ILS rates of longer maturities (unreported).

⁴Specifically, realised quarterly excess returns are computed as follows: $rx_{t+3,m} = \frac{(m \times 4 - 1)}{4} \pi_{t+3,m-3} - m\pi_{t,m} - \frac{I_{t+3}}{I_t}$, where $\pi_{t,m}$ is the continuously-compounded ILS rate of maturity m months observed in month t adjusted for the indexation lag and $\frac{I_{t+3}}{I_t}$ represents the realised inflation over the quarter. In the first term on the right side, $\pi_{t+3,m-3}$ is approximated by $\pi_{t+3,m}$.

| | (1) | (2) | (3) | (4) | (5) |
|----------------------------|---------------------|----------------------|----------------------|----------------------|----------------------|
| | ILS rates only | Avg. surveys | CE | SPF | DMA |
| Main sample, 2005-2023 | | | | | |
| Constant | -0.003* (0.002) | 0.013 (0.009) | 0.006 (0.008) | 0.013 (0.009) | -0.003 (0.002) |
| <i>PC1</i> | 0.009 (0.017) | 0.041* (0.023) | 0.018 (0.017) | 0.062** (0.029) | 0.008 (0.039) |
| <i>PC2</i> | -0.156** (0.065) | -0.180** (0.071) | -0.164** (0.067) | -0.183*** (0.070) | -0.154** (0.077) |
| <i>PC3</i> | -0.009 (0.312) | 0.159 (0.304) | 0.075 (0.327) | 0.221 (0.280) | -0.012 (0.299) |
| π^* proxy | | -0.932* (0.525) | -0.493 (0.417) | -1.044* (0.536) | 0.008 (0.252) |
| | | [0.098] | [0.189] | [0.058] | [0.477] |
| R^2 | 3.98 | 6.09 | 4.37 | 6.70 | 3.53 |
| Subsample, 2005-2019 | | | | | |
| Constant | -0.002 (0.001) | 0.014 (0.010) | 0.009 (0.009) | 0.012 (0.011) | 0.001 (0.002) |
| <i>PC1</i> | 0.032 (0.025) | 0.057** (0.027) | 0.040 (0.024) | 0.065** (0.029) | 0.163*** (0.043) |
| <i>PC2</i> | -0.056 (0.075) | -0.097 (0.079) | -0.076 (0.078) | -0.105 (0.078) | -0.350*** (0.095) |
| <i>PC3</i> | 0.049 (0.380) | 0.260 (0.404) | 0.211 (0.425) | 0.221 (0.421) | 0.348 (0.311) |
| π^* proxy | | -0.949* (0.565) | -0.639 (0.496) | -0.897 (0.645) | -1.002*** (0.227) |
| | | [0.112] | [0.172] | [0.137] | [0.004] |
| R^2 | 1.65 | 4.27 | 3.14 | 2.75 | 16.04 |
| Extended sample, 1991-2023 | | | | | |
| Constant | -0.002 (0.001) | -0.000 (0.001) | -0.000 (0.002) | | -0.002* (0.001) |
| <i>PC1</i> | 0.022*** (0.008) | 0.036*** (0.014) | 0.032** (0.013) | | 0.039* (0.021) |
| <i>PC2</i> | -0.106** (0.041) | -0.121*** (0.044) | -0.118*** (0.044) | | -0.138*** (0.045) |
| <i>PC3</i> | -0.154 (0.288) | -0.161 (0.291) | -0.161 (0.292) | | -0.147 (0.287) |
| π^* proxy | | -0.122 (0.076) | -0.096 (0.072) | | -0.052 (0.054) |
| | | [0.135] | [0.187] | | [0.222] |
| R^2 | 6.52 | 7.26 | 6.95 | | 6.86 |

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 3: **Spanning tests based on one-period realised quarterly excess returns regressions.** Regressions of realised quarterly excess returns averaged across 2 to 10-year maturities. The three first principal components describing the term structure of the ILS rates are used as predictors. For each maturity, the ILS-only specification (column 1) is compared to a specification with an observed π_t^* proxy (columns 2 to 5). Regression coefficients are displayed along with Newey-West standard errors (four lags) in parentheses (with significance levels highlighted with stars), and small-sample (one-sided) p-values for the observed π_t^* proxy obtained with the bootstrap method of Bauer and Hamilton (2018) in brackets. The data are observed at monthly frequency.

4 Model

The model consists of three main building blocks. The first block is the affine term structure model of Joslin et al. (2011). A time-varying endpoint is added to the model following Bauer and Rudebusch (2020), which makes the second block. In that model, π_t^* is cointegrated with and unspanned by ILS rates, in line with the empirical evidence presented in the previous section. The third block consists in adding survey forecasts as in Kim and Orphanides (2012), i.e. by equalling these forecasts to model-implied expectations components up to an error term.

The model of Joslin et al. (2011) (first block) is summarised for convenience and completeness in Appendix B. This section proceeds with the introduction of a shifting endpoint and survey forecasts in the model.

4.1 A term structure model with a shifting endpoint

While the model of Joslin et al. (2011) features a stationary vector autoregressive (VAR) model for the dynamics of the pricing factors of ILS rates under the physical probability measure \mathbb{P} , Bauer and Rudebusch (2020) introduce non-stationarity as follows:

$$\mathcal{P}_t = \bar{\mathcal{P}} + \gamma\pi_t^* + \tilde{\mathcal{P}}_t, \quad \pi_t^* = \pi_{t-1}^* + \eta_t, \quad \tilde{\mathcal{P}}_t = \Phi\tilde{\mathcal{P}}_{t-1} + \tilde{\epsilon}_{t,\tilde{\mathcal{P}}}, \quad (3)$$

where \mathcal{P}_t is an $N \times 1$ vector of unobserved latent factors that come close to principal components once rotated as in Joslin et al. (2011) - henceforth, the "principal components" -, $\bar{\mathcal{P}}$ is an intercept vector, γ is a vector of loadings of the principal components on the trend π_t^* , and $\tilde{\mathcal{P}}_t$ represents the cyclical parts of the principal components. The trend π_t^* is assumed to follow a random walk with innovations $\eta_t \sim i.i.d. N(0, \sigma_\eta^2)$. The cyclical components, $\tilde{\mathcal{P}}_t$, follow a VAR model with one lag and no intercept (so that these cyclical components have a model-implied mean of zero) and innovations $\tilde{\epsilon}_{t,\tilde{\mathcal{P}}} \sim i.i.d. N(0, \tilde{\Omega})$.

The dynamics of the principal components under the risk-neutral measure \mathbb{Q} remain stationary as in Joslin et al. (2011). Specifically, these dynamics follow a VAR model with one lag.

To close the model, the short-term (one-month) ILS rate is assumed to be affine in the principal components:

$$\pi_t = \rho_{0,\mathcal{P}} + \rho_{1,\mathcal{P}}\mathcal{P}_t. \quad (4)$$

Under the assumption of stationarity of $\tilde{\mathcal{P}}_t$, i.e. given that the eigenvalues of Φ are all within the unit circle, the (time-varying) long-run mean of the principal components can be computed as:

$$\mathcal{P}_t^* \equiv \lim_{h \rightarrow \infty} E_t^{\mathbb{P}} [\mathcal{P}_{t+h}] = \bar{\mathcal{P}} + \gamma \pi_t^*. \quad (5)$$

This equation can be inserted in Equation 4 to derive two normalisation conditions in order to identify π_t^* as the long-run mean of the short-term ILS rate π_t , namely $\rho_{0,\mathcal{P}} + \rho_{1,\mathcal{P}} \bar{\mathcal{P}} = 0$ and $\rho_{1,\mathcal{P}} \cdot \gamma = 1$.

Equations 3 can be recast into a VAR by defining $Z_t = (\pi_t^*, \mathcal{P}_t)'$:

$$Z_t = \mu_Z + \Phi_Z Z_{t-1} + v_t, \quad v_t \sim i.i.d. N(0, \Omega_Z), \quad (6)$$

with μ_Z , Φ_Z and Ω_Z being functions of $\bar{\mathcal{P}}$, γ and σ_η as shown in Appendix C.

Given the above, it can be shown that ILS rates are affine functions of the principal components and that π_t^* is unspanned by ILS rates:

$$\Pi_t = A_{\mathcal{P}} + \begin{pmatrix} 0 \\ \vdots \\ B_{\mathcal{P}} \\ 0 \end{pmatrix} Z_t + \mu_t^Z, \quad (7)$$

where Π_t is a $J \times 1$ vector of ILS rates of different maturities ($J > N$), and $\mu_t^Z \sim i.i.d. N(0, I_J \otimes \sigma^2)$.

The model is estimated by maximising the likelihood of the pricing Equation 7, where Z_t is obtained via the Kalman filter. The filter is initialised so that the first observations of the principal components of observed ILS rates are matched, and with a π_t^* of 1.9% for the sample starting in 2005 (and 3.2% for the sample with backcasted ILS rates starting in 1992), in line with long-term survey forecasts. To start the maximum likelihood algorithm, starting values for parameters determining the cross-section of the term structure are taken from the estimation of the model without shifting endpoint, and starting values for the dynamics of Z_t (Equations 6) are obtained by ordinary least squares taking the average of long-term survey forecasts and CE fifth calendar year as proxy for π_t^* . The estimation is robust to considering different π_t^* proxies to obtain starting values, as long as σ_η is fixed (the importance of σ_η is highlighted in section 6).

4.2 Accounting for inflation survey forecasts

The approach of Kim and Orphanides (2012) is followed in order to add survey forecasts to the model. The goal is twofold: (1) add information to help the estimation of π_t^* , which seems warranted given the relatively short time series of observed ILS rates; and (2) inform the estimation of model-based expectations components. The inclusion of survey forecasts is also motivated by the results of Bańbura et al. (2021a, 2021b) and Diercks et al. (2023) suggesting that the ECB's SPF provides good point forecast performance, that it improves inflation forecasts for a wide range of time series models when taken into account, and that a combination of ILS rates and survey forecasts can lead to better forecast accuracy.

In practice, the Kalman filter is enhanced with nine additional measurement equations (five for the CE survey and four for the SPF) that equate the set of survey forecasts to the model-implied average inflation forecasts under the physical probability measure \mathbb{P} , while at the same time allowing for measurement errors. These survey measurement errors can be seen as reflecting the fact that models are subject to estimation and model uncertainty (imperfect model forecasts) or that surveys refer to a relatively small group of respondents (imperfect survey measurement), although they also accommodate the issue that aggregated (e.g. average) survey forecasts may differ from expectations of the marginal investor who determines asset prices, and the possibility that survey expectations suffer from misunderstandings or a misrepresentation of expectations induced by wrong incentives (like a reluctance to deviate too much from the consensus view).

Survey forecasts are connected to model-implied inflation expectations by exploiting the fact that both fixed-horizon and fixed-event inflation forecasts can be accurately approximated by a weighted sum of monthly inflation rates of the price index over the forecast horizon (Patton and Timmermann, 2011). For example, the SPF forecast for the $H =$ fifth calendar year ahead ($SPF5y$) represents approximately the weighted sum of 24 monthly growth rates of the price index covering the fourth and fifth calendar years:

$$f_{t,SPF5y}^{\pi} \approx \sum_{k=0}^{23} \omega_k E_t^{\mathbb{P}}[\pi_{t+l+36+k}], \quad (8)$$

where weights $\omega_k = 1 - |k - 11|/12$, $0 \leq k \leq 23$, and l represents the distance in number of months between the month of the survey and the end of the same calendar year.

Given the affine structure of the model according to Equation (7), each of these additional

measurement equations can be written as:

$$f_{t,H}^\pi = A_{H,Z}^{\mathbb{P}} + B_{H,Z}^{\mathbb{P}} Z_t + v_{t,z}, \quad (9)$$

where $f_{t,H}^\pi$ represents the average (across survey respondents) point inflation forecast for either fixed-horizon or fixed-event forecasts for horizon H , and $A_{H,Z}^{\mathbb{P}}$ and $B_{H,Z}^{\mathbb{P}}$ are computed as weighted loadings over the relevant time window spanning horizon H and are functions of μ_Z , Φ_Z , $\rho_{0,\mathcal{P}}$, and $\rho_{1,\mathcal{P}}$. The set-up of the Kalman filter accounts for the low frequency of survey forecasts (semiannually or quarterly) by having missing values in the additional measurement equations.

The survey measurement errors are modelled as autocorrelated to account for the possibility of momentum effects (i.e. the tendency of survey participants to keep their forecasts unchanged). Specifically, each of the nine measurement errors is assumed to follow an autoregressive model with one lag: $v_{t,z} = \rho_v v_{t-1,z} + \epsilon_{t,v}$, with $\epsilon_{t,v} \sim i.i.d. N(0, I_9 \otimes \sigma_f^2)$, and ρ_v and σ_f common across survey measurement error equations. In practice, the Kalman filter entails nine additional state equations (one for each of the survey measurement error term).

5 Main results

This section presents the main results, i.e. the results that come out of the model with all the features just introduced. The model is estimated with $N = 3$ principal components, which are the usual so-called level, slope, and curvature factors.⁵ The main results refer to the data sample going from 2005 to 2023.

5.1 Estimated π_t^*

The estimated π_t^* - henceforth $\hat{\pi}_t^*$ - displays historical variation but overall suggests little sign of de-anchoring in the history of the euro area, as shown in Figure 3. From 2005 to 2014, $\hat{\pi}_t^*$ hovers in a narrow range going from 1.89% to 1.99%, hence remaining consistent with the ECB objective of the time of an inflation rate below, but close to 2% over the medium term. Fluctuations were limited during the global financial crisis, which is in line with the relatively stable longer-term ILS rates and survey forecasts in spite of volatility in shorter-term gauges.

⁵As in Joslin et al. (2011), the rotated latent factors are found to be remarkably close to the principal components, with only the third (curvature) principal component being marginally more pronounced.

As of 2014, $\hat{\pi}_t^*$ starts following a downward trend with a trough at 1.83% reached in August 2016, while at the same time the ECB resorted to unconventional monetary policy tools, such as asset purchases, in addition to policy rate cuts to support inflation (expectations). Accordingly, the level of $\hat{\pi}_t^*$ remaining in line with the ECB objective of the time could partly be due to the supporting role played by accommodative monetary policy. At the end of 2016, $\hat{\pi}_t^*$ would somewhat recover, and the ECB would soon reduce the monetary policy impulse. However, it became gradually clearer in 2019 that measures of inflation expectations were tumbling again, dragging $\hat{\pi}_t^*$ along, and triggering renewed unconventional measures. A new trough was reached in the first quarter of 2020 in the wake of the COVID-19 outbreak in Europe, with $\hat{\pi}_t^*$ at 1.80%. But the prompt reaction of the ECB to boost unconventional measures helped to stabilise euro area economies. Eventually, $\hat{\pi}_t^*$ increased back to levels close to 2%, which became the new medium-term inflation objective following the strategy review of the ECB completed in July 2021.




Figure 3: $\hat{\pi}_t^*$ and observed proxies.

Long-term inflation survey forecasts are broadly in line with $\hat{\pi}_t^*$. While π_t^* refers in principle to an infinite horizon, long-term SPF forecasts have a horizon of about five years and long-term CE forecasts are an average of expected inflation over the six to ten-year horizon. Nevertheless, these horizons seem sufficiently far into the future for survey respondents to consider that economic shocks will mostly have faded away by that time, implying that their inflation forecasts should be close to the inflation equilibrium level they estimate (and close to the ECB objective

if inflation expectations are anchored).

The 5y5y ILS rate comoves strongly with $\hat{\pi}_t^*$, with a correlation of 97%, although it is substantially more volatile. According to the model, a spot ILS rate of maturity j , $\pi_{t,j}$, follows the dynamics $\pi_{t,j} = A_{\mathcal{P},j} + B_{\mathcal{P},j}\mathcal{P}_t = A_{\mathcal{P},j} + B_{\mathcal{P},j}\bar{\mathcal{P}} + B_{\mathcal{P},j}\gamma\pi_t^* + B_{\mathcal{P},j}\tilde{\mathcal{P}}_t$. Thus, the 5y5y ILS rate loads on π_t^* with weight $(2B_{\mathcal{P},10y} - B_{\mathcal{P},5y})\gamma$ and on the cyclical parts of the principal component with weights $2B_{\mathcal{P},10y} - B_{\mathcal{P},5y}$. The first two entries of the vector $2B_{\mathcal{P},10y} - B_{\mathcal{P},5y}$ are positive as the level factor loads more or less equally on ILS rates of different maturities and the slope factor loads more positively on longer-term rates. Since these two entries are associated with the first two entries of γ that tend to be large (as the level and slope factors load significantly on π_t^*), the 5y5y ILS rate dynamics are largely influenced by π_t^* , implying the high correlation. The estimated γ is (16.4, 9.1, 2.2)', with values indeed well above 1 as a result of the relatively stable $\hat{\pi}_t^*$ (the small fluctuations in $\hat{\pi}_t^*$ must be multiplied to match the fluctuations of the principal components). By contrast, the correlation between the 1y ILS rate and $\hat{\pi}_t^*$ is more limited, at 41%: the second entry in $B_{\mathcal{P},1y}$ is negative and multiplied by the large second entry of γ , which reduces the overall loading $B_{\mathcal{P},1y} \cdot \gamma$ on $\hat{\pi}_t^*$.

The fact that the 5y5y ILS rate is highly correlated with $\hat{\pi}_t^*$ is consistent with the idea that longer-term ILS rates depend more on structural forces, such as reflected in π_t^* , rather than cyclical components. As shown in Figure 4, the cyclical component of the 5y5y ILS rate shows only mild fluctuations. By contrast, the cyclical component becomes more prevalent for shorter-term ILS rates and explains most of the variation in the 1y rate.




Figure 4: **Trend and cyclical components in ILS rates.**

The importance of π_t^* in the dynamics of maturity-specific ILS rates is given by the loading $B_{\mathcal{P},j} \cdot \gamma$, which Figure 5 shows as a function of maturity j . The function starts at 1 for $j = 1$ month (i.e. the maturity of the implied short-term rate), following from the restriction $\rho_{1,\mathcal{P}} \cdot \gamma = 1$ imposed for identification purposes. As maturity j increases, the loadings of ILS rates on π_t^* increase. The function is concave as the model imposes that loadings converge to zero as the maturity increases (beyond the 10-year maturity) to infinity: stationarity under the risk-neutral probability measure \mathbb{Q} implies a constant short-term ILS rate in the long run. Empirical counterparts to the model-implied loadings are given by the cointegration analysis regressing ILS rates on an observed $\hat{\pi}_t^*$ proxy. Taking the average of SPF and CE survey forecasts for at least five years ahead as observed proxy, the empirical loadings are found to be broadly in line with the model-implied loadings of ILS rates of maturities of at least four years. There is a pronounced discrepancy between empirical and model-implied loadings for shorter-term ILS

rates, which comes from the identification restriction $\rho_{1,\mathcal{P}} \cdot \gamma = 1$. The shape of the $B_{\mathcal{P},j} \cdot \gamma$ function and correspondence to empirical estimates are similar to those in Bauer and Rudebusch (2020).




Figure 5: **ILS rates loadings on π_t^*** . The empirical coefficients come from regressions of ILS rates on the observed π_t^* proxy taken as the average of long-term SPF and CE forecasts and the CE fifth calendar year forecasts. The model parameter σ_η is calibrated in line with this observed π_t^* proxy, i.e. a standard deviation of 50 basis points per century.

5.2 Expectations components and inflation risk premia

While π_t^* is associated to an infinite horizon, the ECB's price stability objective refers to the "medium term". There is no consensus on what horizon is to be deemed best representative of the medium term, but there are several natural candidates. Financial market analysts often focus on the 5y5y horizon (which corresponds broadly to the longest horizon in the CE survey), probably as the horizon is distant enough in the future so as to consider that most of the effects of small temporary shocks will have faded away five years on, and also as it refers to a relatively liquid market segment. On the latter, liquidity on the ILS market is highest exactly for the 5-year and 10-year maturities (Boneva et al., 2019). As an alternative, the ECB's SPF probes into inflation expectations up to the four or fifth calendar year horizon, and refers to these forecasts as "long-term" inflation expectations. Finally, the Eurosystem/ECB staff regularly produces macroeconomic projections up to the second or third calendar year horizon. These projections are a major input to help the ECB's Governing Council to form a view on the inflation outlook.

Accordingly, the 5y5y, 1y4y and 1y2y ILS rates are considered relevant to refer to some

concept of "medium term". The 1y4y rate is the closest market-based counterpart of the long-term SPF forecast, and the 1y2y rate is the closest counterpart to the end-of-horizon staff projections. The 1y rate is kept in figures below as reference.

In order to assess the anchoring of inflation expectations at these horizons, one must extract the inflation expectations components embedded in ILS rates, i.e. strip ILS rates of inflation risk premia. The model presented above is well suited to perform such a decomposition in a consistent way for the cross-section of ILS rates. The expectations component of an ILS rate of maturity j years, $\pi_{t,j}^e$, is defined as the average expected (\mathbb{P} dynamics) short-term ILS rate over the next j years:

$$\pi_{t,j}^e = \frac{1}{12j} \sum_{h=1}^{12j} E_t^{\mathbb{P}}[\pi_{t+h}], \quad (10)$$

and the inflation risk premia, $\pi_{t,j}^{irp}$, is obtained by retrieving the expectations component from the ILS rate, such that $\pi_{t,j}^{irp} = \pi_{t,j} - \pi_{t,j}^e$.

Figure 6 shows that estimated expectations components tend to be relatively stable for longer-term ILS rates. For instance, the expectations component of the 5y5y ILS rate has always been close to 2%, with a minimum of 1.80% reached at the time of the COVID-19 outbreak. In fact, the expectations component of the 5y5y ILS rate shows similar dynamics to $\hat{\pi}_t^*$. This similarity comes from the relatively fast convergence of the short-term ILS rate to its long-run mean, which is precisely $\hat{\pi}_t^*$. As the 5y5y horizon is a quite distant one, forecasts of the short-term rate essentially hover around $\hat{\pi}_t^*$. The speed of convergence of forecasts depends on the eigenvalues of Φ , i.e. the feedback matrix in the VAR for the cyclical parts of principal components (Equations 3). As the highest eigenvalue of Φ is 0.95 (a number typically lower than in models with a fixed endpoint since it refers only to the persistence in the *cyclical* parts of principal components), the cyclical part in the forecasts of the short-term rate under the \mathbb{P} measure is close to zero at the 5y5y horizon (at the 5-year horizon, not much is left of the persistence implied by the eigenvalue: $0.95^{60} = 0.046$ and $0.95^{120} = 0.002$). By contrast, the expectations component of the 1y ILS rate has shown significant movements, ranging between 0.74% and 4.58% ($0.95^{12} = 0.540$).

The expectations components of the 1y4y and 1y2y ILS rates show somewhat more volatility than that of the 5y5y rate (or $\hat{\pi}_t^*$), which suggests somewhat more pronounced de-anchoring risks at these horizons. The expectations component of the 1y2y ILS rate, in particular, gets as low




Figure 6: **Inflation expectations components and inflation risk premia.**

as 1.60% during COVID-19 and 1.64% at the end of 2014. These relatively low values more clearly stretch the interpretation of the objective of "below, but close to 2%", and support the implementation of the APP and PEPP that shortly followed these tentative signs of de-anchored medium-term inflation expectations.

Looking into the sign of the estimated inflation risk premia, these have either been positive or negative, depending on the macroeconomic context. These premia followed a downward trend across the curve over most of the sample period and turned negative in the years preceding COVID-19. However, as the post-COVID-19 recovery took root in 2021 and gained momentum in 2022, inflation risk premia increased to become positive again. The return of inflation risk premia to positive values suggest that the macroeconomic environment is dominated by supply (rather than demand) shocks: inflation tends to increase when financial asset payoffs (in real

terms) are highly valued, i.e. when real activity declines and the marginal utility of consumption rises. Supply shocks emerged during the post-pandemic recovery from frozen global value chains, creating supply bottlenecks, and were exacerbated by Russia's war in Ukraine, which worsened the supply of energy, with particularly acute consequences for Europe as a net energy importer. Ongoing geopolitical tensions, global economic fragmentation risks, and the challenges posed by the energy transition likely continue to support positive inflation risk premia.

6 Breaking down the main results

This section sheds lights on the marginal impacts of the three main building blocks of the model: the standard model of Joslin et al. (2011) (with a fixed endpoint), the addition of a shifting endpoint as in Bauer and Rudebusch (2020), and the addition of survey forecasts as in Kim and Orphanides (2012).

One disadvantage of the model of Joslin et al. (2011, JSZ) is that it is not well suited to assess whether or not longer-term inflation expectations are de-anchored. As the model features a fixed endpoint (i.e. a constant π^*), forecasts of the short-term ILS rate converge at any point in time to the same value. Empirically, this value might or might not be estimated to be aligned with the ECB's inflation objective, depending on the history of ILS rates. Over the sample 2005-2023, the value of π^* is estimated at 1.7%, which is quite distant from today's ECB's objective of 2%. As a result of this relatively low fixed endpoint, the expectations component of the 5y5y ILS rate hovers around 1.7%, as shown in Figure 7. The fixed endpoint could even be estimated at a lower value, for instance if the model would be estimated on a sample ending before COVID-19. In that case, the model would only see the downward trend of ILS rates and the quite low last observations of these rates. Nevertheless, the JSZ model provides expectations components with reasonable variations for the shorter-term ILS rates, as these components are less influenced by the (fixed or shifting) endpoint. Alternatively, the fixed endpoint in the JSZ model could be calibrated at 2% for consistency with the ECB's objective (Burban et al., 2022), but the issue remains that the model is not well suited to assess potential de-anchoring of long-term inflation expectations.




Figure 7: **Expectations components.** "JSZ" refers to Joslin et al. (2011). "BR" refers to Bauer and Rudebusch (2020).

Adding a shifting endpoint to the JSZ model, as in Bauer and Rudebusch (2020, BR), allows in principle to discuss the potential de-anchoring of longer-term inflation expectations, but the estimation of the model is confronted with an identification issue. In practice, it turns out to be difficult to estimate π_t^* as the parameter σ_η (the standard error of the innovation term in the random walk model of π_t^*) is not well identified. Figure 8 shows that the average likelihood function of the BR model without surveys is flat for values of σ_η corresponding to a range going from 0 to 1% standard deviation per century for π_t^* .⁶ This identification issue is well-known and usually tackled with a tight prior in Bayesian econometrics (Bauer and Rudebusch, 2020;

⁶The translation of σ_η to a value corresponding to the standard deviation per century for π_t^* is given by the formula $\sqrt{(1200\sigma_\eta^2)}$.

Del Negro et al., 2017, 2019). As the equivalent of a tight prior is calibration in the frequentist maximum likelihood approach, the yellow lines in Figure 7 are obtained by calibrating σ_η at a value corresponding to a 50 basis point standard deviation per century, for illustrative purposes. These 50 basis points correspond to the standard deviation of the observed π_t^* proxy taken as the average across survey forecasts for a horizon of at least five years ahead. Following this choice, the BR model shows more variation for the expectations components of longer-term ILS rates than those derived from the JSZ model, as π_t^* is time-varying in the BR model. Yet, the JSZ and BR models provide similar estimates of expectations components in short-term ILS rates.




Figure 8: **Average log-likelihoods as a function of σ_η .** Values of σ_η are expressed in basis point standard deviation per century.

Adding surveys to the BR model offers two advantages: (1) it allows for a proper identification of σ_η ; and (2) it brings model-implied forecasts closer to survey forecasts. As shown by the likelihood function of the BR model with surveys, a global optimum is well-identified in Figure 8 at a standard deviation for $\hat{\pi}_t^*$ of 50 basis points standard deviation per century. With this optimum, yellow and magenta lines in Figure 7 can be compared directly as the only difference between the two is the inclusion of surveys in the BR model. The main impact of the inclusion of surveys is that the expectations components in the 1y4y and 5y5y ILS rates are estimated to be higher and flatter. This is in line with survey forecasts remaining closer to 2% than ILS rates throughout the sample period. As shown in Figure 9 (10), the fit of the CE (SPF) survey forecasts by the BR model with surveys clearly improves on the model without surveys, and is

overall quite accurate: the standard deviation of the survey measurement errors is 9 basis points across all survey forecasts.

Finally, an additional feature can be added to the model: the autocorrelation of survey measurement errors. This autocorrelation could account, for instance, for a momentum effect in survey respondents, i.e. the tendency to keep forecasts unchanged. As a further motivation, survey measurement errors of the BR model with surveys are found to be autocorrelated, prompting the need to model this autocorrelation. The autoregressive parameter of survey measurement errors is estimated to be quite high at 0.89. However, modelling this autocorrelation does not impact much model estimates, which is in line with Kim and Orphanides (2012). Survey forecasts are still tightly fit, with only a slightly larger standard deviation of survey measurement errors of 10 basis points, and estimates of $\hat{\pi}_t^*$ and expectations components are not materially changed (Figure 7). The small differences come almost entirely from calibrating σ_η at 30 basis points standard deviation per century in the BR model with surveys and autocorrelation. To be sure, the likelihood function displayed in Figure 8 is almost flat for values of σ_η below 30 basis points. Therefore, this value can be seen as a conservative choice, as it allows for more variation in $\hat{\pi}_t^*$, and hence greater probability of de-anchoring: even with a conservatively large σ_η , the model does not suggest that medium- to long-term inflation expectations have been de-anchored in the euro area.

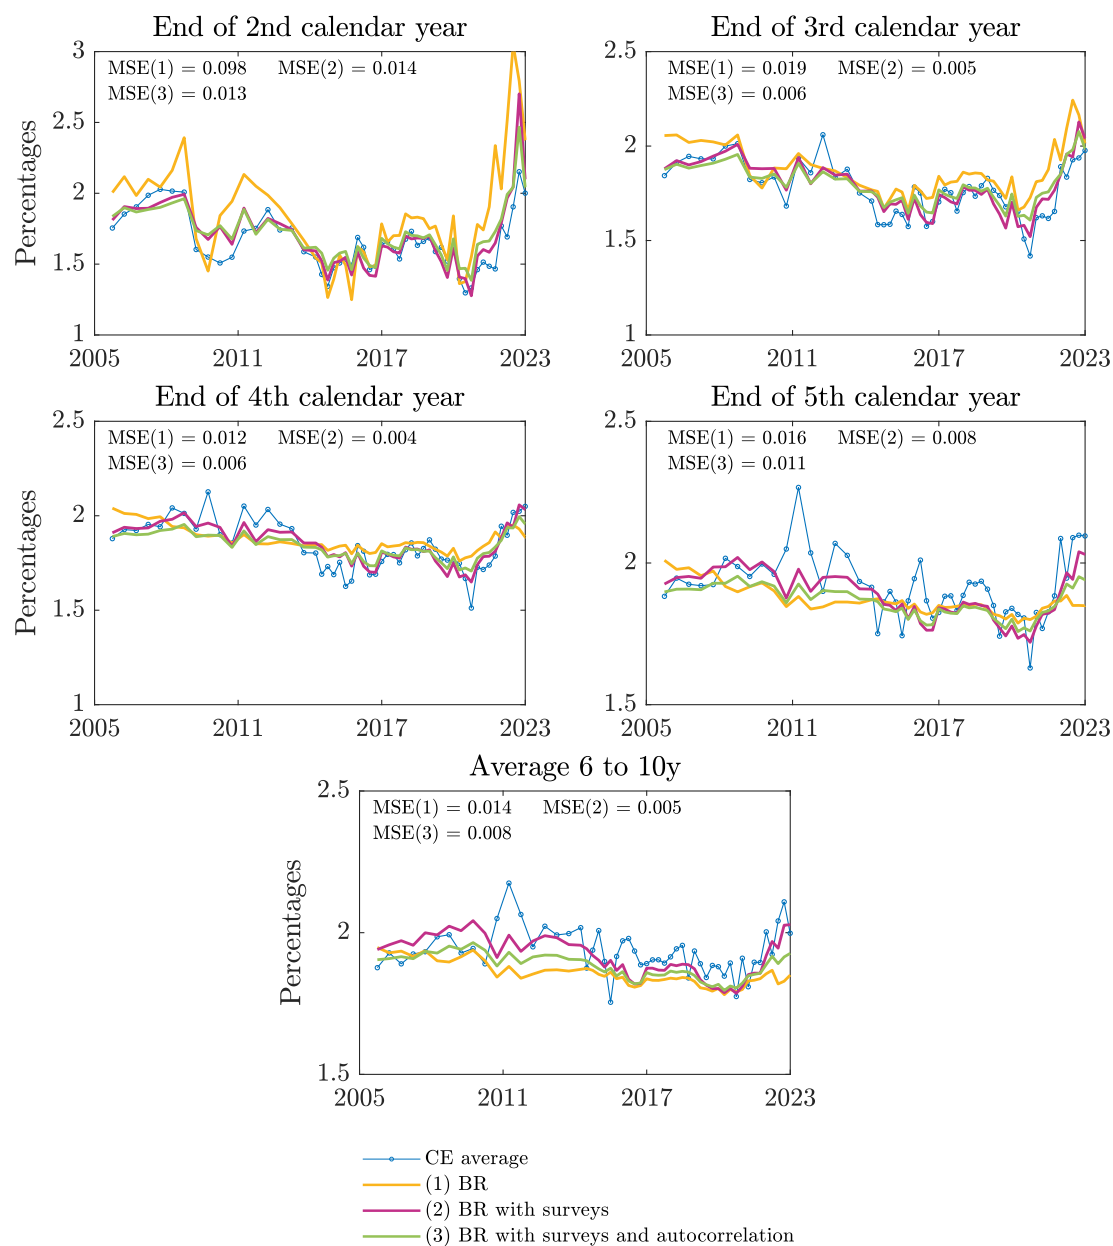


Figure 9: **Fit of Consensus Economics survey forecasts** "JSZ" refers to Joslin et al. (2011). "BR" refers to Bauer and Rudebusch (2020). The mean squared error of model $i \in \{1, 2, 3\}$ is computed as: $MSE_i = \frac{1}{T} \sum_{t=1}^T (f_t^i - \hat{f}_t^i)^2$.




Figure 10: **Fit of SPF forecasts.** "JSZ" refers to Joslin et al. (2011). "BR" refers to Bauer and Rudebusch (2020). The mean squared error of model $i \in \{1, 2, 3\}$ is computed as: $MSE_i = \frac{1}{T} \sum_{t=1}^T (f_t^i - \hat{f}_t^i)^2$.

7 Robustness tests

This section presents two types of robustness tests. The first type regards the crucial parameter σ_η that governs the volatility of $\hat{\pi}_t^*$. That parameter is allowed to change during COVID-19 and during the 1990s to allow for larger movements in $\hat{\pi}_t^*$. The question asked is thus: is there a regime shift during COVID-19 and/or during the 1990s? The second type of robustness tests is a series of exclusion restrictions, in an attempt to reduce the number of model parameters.

7.1 Regime shifts in the volatility of π_t^*

There is no sign of regime shift during the COVID-19 period. Allowing the parameter σ_η to increase as of February 2020 results in a lower likelihood than the model keeping σ_η constant at 30 basis points standard deviation per century for $\hat{\pi}_t^*$. It should be noted that the greater volatility in ILS rates as of February 2020 is concentrated in shorter-term rates, i.e. the rates that have less of a connection with π_t^* . Longer-term ILS rates, such as the 5y5y rate, did not show substantially more volatility, although they did converge relatively smoothly to just above 2%.

Figure 11 nevertheless plots the π_t^* that is estimated by doubling σ_η as of February 2020 (i.e. going from 30 to 60 basis points standard deviation per century for π_t^*). The estimated π_t^* is close to the baseline $\hat{\pi}_t^*$ of the main results, but increases more at the end of the sample to just reach 2%. While the latter seems desirable in the context of the ECB's medium-term objective, the model prefers the results with unchanged σ_η and a slightly lower π_t^* , probably because the other model parameters are constant through the sample period and adjusting them to accommodate a higher σ_η at the end of the sample causes a less good fit on the rest of the sample. The adjustment in other model parameters also explains why the estimated π_t^* is slightly different from the baseline $\hat{\pi}_t^*$ of the main results over the period 2005-2019.

In contrast to the lack of evidence in favour of a regime shift during COVID-19, there is ample evidence of a regime shift over the period 1992-1998. A regime shift is allowed for before October 1998, i.e. before the newly established ECB announced its inflation target of "below 2%". The regime shift is also allowed for as survey forecasts and backcasted ILS rates (including longer-term ones) show substantially higher volatility during the 1990s. The optimum σ_η before October 1998 is found at a value of 210 basis points standard deviation per century for $\hat{\pi}_t^*$, or seven times the optimum σ_η found for the main sample going from 2005 to 2023.

In spite of the clear evidence of a regime shift before October 1998, Figure 11 shows that the estimated π_t^* remains remarkably close to the $\hat{\pi}_t^*$ of the main results over the period 2005-2023. The main advantage of allowing for a regime shift in σ_η in the 1990s is to allow for $\hat{\pi}_t^*$ to converge fast to levels in line with what would soon become the ECB's objective. Allowing for a fast convergence is indeed necessary with π_t^* initialised at 3.2% in the Kalman filter, as suggested by the value of the long-term synthetic euro area CE surveys in January 1992. Starting from such a high level, $\hat{\pi}_t^*$ cannot get to less than 2% by the end of 1998 if σ_η is not large.




Figure 11: **Estimated π_t^* with regime shifts.**

7.2 Parameter restriction tests

Various restrictions on parameters of the model are tested as in Christensen et al. (2010, 2011), Christensen and Rudebusch (2012), and a series of subsequent papers.

Restricting $\tilde{\Omega}$ to be a diagonal matrix is not supported by a likelihood ratio test (p-value $\ll 1\%$). That is, the cyclical parts of principal components are found to be conditionally correlated. Nevertheless, this correlation is not significant from an economic point of view in the sense that the fit of ILS rates does not suffer much from imposing $\tilde{\Omega}$ to be diagonal. As shown in Table 4, the average root mean squared ILS measurement error is 1.58 basis points with $\tilde{\Omega}$ as full matrix, and 1.63 basis points with $\tilde{\Omega}$ diagonal. On this basis, the restriction of a diagonal $\tilde{\Omega}$ could be retained.

| | 1y | 2y | 3y | 4y | 5y | 6y | 7y | 8y | 9y | 10y | avg |
|-----------------------|------|------|------|------|------|------|------|------|------|------|------|
| $\tilde{\Omega}$ full | 1.15 | 3.38 | 1.43 | 1.30 | 1.77 | 1.77 | 1.27 | 0.68 | 1.06 | 2.02 | 1.58 |
| $\tilde{\Omega}$ diag | 1.20 | 3.38 | 1.41 | 1.29 | 1.76 | 1.75 | 1.27 | 0.80 | 1.25 | 2.15 | 1.63 |

Table 4: **ILS rates fitting errors.** The table reports root mean squared fitting errors. " $\tilde{\Omega}$ diag" refers to the restriction imposing that $\tilde{\Omega}$ is a diagonal matrix. " $\tilde{\Omega}$ full" refers to a full $\tilde{\Omega}$ matrix.

In addition, the various elements of Φ are tested (except the diagonal elements, i.e. the coefficients on own lags). Likelihood ratio tests clearly reject any exclusion restriction. The off-diagonal elements of Φ are tested one by one. First, all Φ matrices including only one zero

restriction are considered. As shown in Table 5 summarising the results, the most likely zero restriction would be imposed on the coefficient of the level factor in the slope equation, i.e. $\phi_{2,1} = 0$, but this restriction is highly rejected by a likelihood ratio test, as well as by the Akaike and Bayesian information criteria. Second, all the Φ matrices including two zero restrictions, including $\phi_{2,1} = 0$, are tested. The second most likely restriction would be imposed on the coefficient of the slope factor in the curvature equation ($\phi_{3,2} = 0$), but this restriction is also highly rejected. Idem for a third potential restriction on the coefficient of the slope factor in the level equation ($\phi_{1,2} = 0$). The exclusion restrictions taken together ($\phi_{2,1} = \phi_{3,2} = 0$, and $\phi_{2,1} = \phi_{3,2} = \phi_{1,2} = 0$), are also highly rejected. These tests were run while imposing $\tilde{\Omega}$ to be a diagonal matrix, but exclusion restrictions in Φ are also highly rejected with $\tilde{\Omega}$ being a full matrix.

| Restriction | Log-llk | p-value | AIC | BIC |
|--|----------|---------|-----------|-----------|
| Φ full | 17213.45 | - | -34384.91 | -34313.46 |
| $\phi_{2,1} = 0$ | 17202.92 | 0 | -34365.84 | -34297.79 |
| $\phi_{2,1} = \phi_{3,2} = 0$ | 17188.51 | 0 | -34339.03 | -34274.38 |
| $\phi_{2,1} = \phi_{3,2} = \phi_{1,2} = 0$ | 17183.18 | 0.1% | -34330.35 | -34269.11 |

Table 5: **Exclusion restriction tests.** Likelihood ratio tests are carried out for each additional exclusion restriction. ϕ refers to elements populating Φ with the first subscript referencing the line, and the second referencing the column.

8 Conclusion

From a policy point of view, the model presented above - which follows Bauer and Rudebusch (2020) for the modelling of a shifting inflation endpoint, π_t^* , and Kim and Orphanides (2012) for the inclusion of survey forecasts complementing the information entailed in ILS rates - provides evidence that long-term inflation expectations have been well-anchored with the ECB's inflation objective, as the estimated π_t^* has remained close to 2% since 2005. No evidence of regime shift is found during the COVID-19 pandemic, but a model estimated on ILS rates backcasted to 1992 requires a larger σ_η (the standard error of innovations in the random walk model of π_t^*) over the period 1992-1998, so that the estimated π_t^* converges to below 2% by the time of the establishment of the ECB.

From a technical point of view, the model provides a framework that allows the analysis of ILS rates in a way that is consistent with the sparser and more infrequent information provided by surveys. The model extracts the expectations components of ILS rates taking into account

survey forecasts, while at the same time allowing for imperfections in survey measurements. As a result, the model provides monthly estimates of not only π_t^* but also expectations components for ILS rates of virtually any maturity. In that sense, the framework is a first step in answering the call formulated in ECB (2021) on "the need for a comprehensive framework for assessing (un)anchoring". Next steps probably involve the integration of more data into the model, for instance related to consumer expectations, the dispersion of individual professional and consumer forecasts, and past inflation realisations.

Future research could pursue various other avenues. The finding that the identification of σ_η is solved when survey forecasts are taken into account could deserve more attention. There is scope for a more thorough investigation of possible regime shifts, possibly involving changes in more parameters than σ_η . Historical data (where regime shifts are more likely to be found) might require to look at the data of specific euro area Member States. Beyond the euro area, the model could be applied to various countries across the world. As ILS rates are observed every day, daily estimates of π_t^* could be derived.

Appendix A Additional results from the preliminary analysis

A.1 Calibration of Cieslak and Povala (2015)'s discounted moving average

The approach detailed in the online appendix of Cieslak and Povala (2015) is followed to calibrate the parameters driving the discounted moving average (DMA) of past realised core inflation.

The DMA is written as follow:

$$\hat{\pi}_t^* = (1 - \nu) \sum_{i=0}^{t-1} \nu^i \pi_{t-i}^{\text{core}}, \quad (\text{A.1})$$

where π_{t-i}^{core} is the year-on-year core inflation observed in month $t - i$. The sum is truncated at $N = 120$ months with a smoothing parameter, ν , of 0.975. The latter is determined by minimizing the fitting errors of an autoregressive model of order one - AR(1) - in gap form estimated using nonlinear least squares. The relation is the following :

$$\frac{HICP_{t+1}^{\text{core}}}{HICP_{t-11}^{\text{core}}} - \hat{\pi}_t^*(\nu, N) = \phi \left(\frac{HICP_t^{\text{core}}}{HICP_{t-12}^{\text{core}}} - \hat{\pi}_t^*(\nu, N) \right) + \epsilon_{t+1}, \quad (\text{A.2})$$

where $HICP_t^{\text{core}}$ is the core harmonised index of euro area consumer prices. Fitting errors are minimised by searching for parameter ν (keeping N fixed) that gives the most accurate core inflation forecasts and the best fit to inflation surveys, considering the CE forecasts introduced in Section 2. The left panel of Figure A.1 displays the root mean squared error (RMSE) for different values of ν , given $\phi = 0.996$. This value of ϕ is obtained by fitting an AR(1) process to core HICP inflation from January 1970 to December 2023. As shown on the panel, the RMSE is minimized for a value ν of 0.968. The right panel shows the results of a grid search over parameters ν and ϕ . The minimum MSE is located at $\nu = 0.975$ and $\phi = 0.918$.

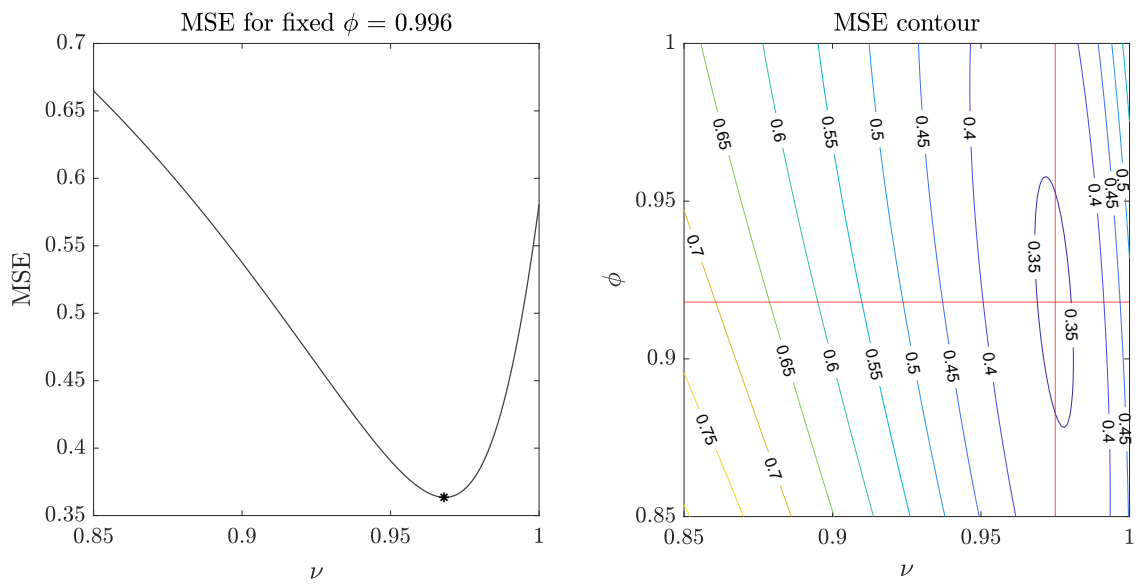


Figure A.1: **Core inflation mean squared errors for different gain parameters ν and ϕ .** Note : The mean squared error is computed as the average over CE survey forecasts at horizon j : $MSE_j = \frac{1}{T_j} \sum_{t=t_j}^{T_j} (f_{t+j} - \hat{f}_{t+j}^\nu)^2$. On the left panel, the minimum is reached at $\nu = 0.968$. On the right panel, the red lines point to the minimum MSE reached at $\nu = 0.975$ and $\phi = 0.918$.

A.2 Cointegration analysis

| | Main sample, 2005-2023 | | | | Subsample, 2005-2019 | | | | Extended sample, 1991-2023 | | | |
|-----------------------------|------------------------|----------------------|----------------------|----------------------|----------------------|---------------------|----------------------|----------------------|----------------------------|----------------------|----------------------|---------------------|
| | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | SPF | DMA |
| Constant | -0.093*** (0.013) | -0.073*** (0.020) | -0.080*** (0.009) | -0.004*** (0.003) | -0.070*** (0.012) | -0.026* (0.015) | -0.082*** (0.014) | -0.008*** (0.003) | -0.014*** (0.003) | -0.012*** (0.003) | -0.014*** (0.003) | 0.007*** (0.001) |
| π^* proxy | 5.639*** (0.676) | 4.522*** (1.017) | 5.037*** (0.474) | 1.387*** (0.185) | 4.385*** (0.628) | 2.057*** (0.770) | 5.082*** (0.762) | 1.671*** (0.231) | 1.593*** (0.112) | 1.517*** (0.112) | 1.593*** (0.112) | 0.648*** (0.046) |
| R^2 | 49 | 19 | 64 | 39 | 28 | 5 | 46 | 52 | 49 | 43 | 57 | 57 |
| ADF_{OLS} residual | -3.42*** | -3.67*** | -3.40*** | -3.59*** | -2.65*** | -2.30** | -2.36** | -2.99*** | -3.07*** | -2.98*** | -3.07*** | -3.58*** |
| PP_{OLS} residual | -4.24*** | -3.48*** | -4.66*** | -3.50*** | -3.42*** | -3.01*** | -3.45*** | -4.05*** | -3.82*** | -3.63*** | -3.82*** | -4.23*** |
| ADF_{DOLS} residual | -2.32** | -2.04** | -2.34** | -2.98*** | -1.96** | -1.69* | -1.39 | -2.53** | -1.82* | -1.63* | -1.82* | -2.71*** |
| PP_{DOLS} residual | -2.22** | -1.98** | -2.22** | -2.97*** | -1.81* | -1.58 | -1.16 | -2.47** | -1.78* | -1.60 | -1.78* | -2.65*** |
| Johansen trace $r = 0$ | 13.94* | 22.54*** | 18.38** | 37.08*** | 14.40* | 18.68** | 10.30 | 17.99** | 21.94*** | 22.07*** | 21.94*** | 42.18*** |
| Johansen trace $r = 1$ | 4.21** | 4.05** | 2.96* | 4.02** | 4.30** | 4.02** | 1.24 | 1.07 | 9.37*** | 6.38** | 9.37*** | 7.25*** |
| Johansen eigenvalue $r = 0$ | 9.73 | 18.49** | 15.42** | 33.06*** | 10.10 | 14.66** | 9.06 | 16.92** | 12.57* | 15.69** | 12.57* | 34.92*** |
| Johansen eigenvalue $r = 1$ | 4.21** | 4.05** | 2.96* | 4.02** | 4.30** | 4.02** | 1.24 | 1.07 | 9.37*** | 6.38** | 9.37*** | 7.25*** |

*** p < 0.01, ** p < 0.05, * p < 0.1

Table A.1: **Cointegration regressions and tests based on 2y ILS rates.** "DMA" stands for discounted moving average (Cieslak and Povala, 2015). Coefficients are based on regressions with Newey-West standard errors (using four lags) in parentheses. For the cointegration residuals from regressions estimated by OLS and dynamic OLS (DOLS), the second panel reports Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics. DOLS regressions include four leads and lags of the first difference of the 2y ILS rate and observed π_t^* proxy. The table also reports the Johansen trace and eigenvalue statistics, which tests whether the cointegration rank (r) among the ILS rates and π_t^* proxy is 0/1 against the alternative that it exceeds 0/1 (trace) or it equals 1/2 (eigenvalue). The number of lags for ADF and Johansen tests is selected based on the AIC. The data are observed at monthly frequency.

| | Main sample, 2005-2023 | | | | Subsample, 2005-2019 | | | | Extended sample, 1991-2023 | | | |
|-----------------------------|------------------------|----------------------|----------------------|----------------------|----------------------|---------------------|----------------------|----------------------|----------------------------|----------------------|---------------------|--|
| | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | SPF | DMA | Avg. surveys | CE | DMA | |
| Constant | -0.074*** (0.007) | -0.057*** (0.014) | -0.061*** (0.005) | -0.002*** (0.002) | -0.066*** (0.011) | -0.024** (0.013) | -0.077*** (0.012) | -0.007*** (0.002) | -0.013*** (0.003) | -0.012*** (0.003) | 0.007*** (0.001) | |
| π^* proxy | 4.662*** (0.381) | 3.748*** (0.696) | 4.099*** (0.253) | 1.304*** (0.162) | 4.248*** (0.593) | 2.071*** (0.666) | 4.916*** (0.615) | 1.687*** (0.161) | 1.631*** (0.112) | 1.572*** (0.109) | 0.698*** (0.038) | |
| R^2 | 54 | 22 | 69 | 56 | 35 | 7 | 57 | 70 | 59 | 52 | 75 | |
| $ADF_{OLS}Residual$ | -3.27*** | -3.39*** | -3.17*** | -3.45*** | -2.42** | -1.93* | -2.17** | -3.12*** | -2.34** | -2.29** | -3.14*** | |
| $PP_{OLS}Residual$ | -3.91*** | -3.10*** | -4.27*** | -3.29*** | -2.99*** | -2.47** | -3.12*** | -4.16*** | -3.00*** | -2.82*** | -3.85*** | |
| $ADF_{DOLS}Residual$ | -2.11** | -1.88* | -1.98** | -3.04*** | -1.82* | -1.44 | -1.34 | -2.55** | -1.38 | -1.31 | -2.30** | |
| $PP_{DOLS}Residual$ | -2.08** | -1.90* | -1.90* | -3.07*** | -1.75* | -1.38 | -1.19 | -2.53** | -1.40 | -1.35 | -2.30** | |
| Johansen trace $r = 0$ | 12.79 | 20.94*** | 14.45* | 32.68*** | 12.86 | 16.60** | 8.09 | 19.99*** | 22.55*** | 19.22** | 29.36*** | |
| Johansen trace $r = 1$ | 3.21* | 2.01 | 2.47 | 3.33* | 2.49 | 2.07 | 1.46 | 1.05 | 5.80** | 3.89** | 6.15** | |
| Johansen eigenvalue $r = 0$ | 9.58 | 18.93*** | 11.98 | 29.35*** | 10.37 | 14.53** | 6.63 | 18.94*** | 16.75** | 15.33** | 23.20*** | |
| Johansen eigenvalue $r = 1$ | 3.21* | 2.01 | 2.47 | 3.33* | 2.49 | 2.07 | 1.46 | 1.05 | 5.80** | 3.89** | 6.15** | |

*** p < 0.01, ** p < 0.05, * p < 0.1

Table A.2: **Cointegration regressions and tests based on 5y ILS rates.** "DMA" stands for discounted moving average (Cieslak and Povala, 2015). Coefficients are based on regressions with Newey-West standard errors (using four lags) in parentheses. For the cointegration residuals from regressions estimated by OLS and dynamic OLS (DOLS), the second panel reports Augmented Dickey-Fuller (ADF) and Phillips-Perron (PP) unit root test statistics. DOLS regressions include four leads and lags of the first difference of the 5y ILS rate and observed π_t^* proxy. The table also reports the Johansen trace and eigenvalue statistics, which tests whether the cointegration rank (r) among the ILS rates and π_t^* proxy is 0/1 against the alternative that it exceeds 0/1 (trace) or it equals 1/2 (eigenvalue). The number of lags for ADF and Johansen tests is selected based on the AIC. The data are observed at monthly frequency.

A.3 Spanning analysis

| | (1) ILS rates only | (2) Avg. surveys | (3) CE | (4) SPF | (5) DMA |
|----------------------------|-----------------------|-------------------------------|------------------------------|--------------------------------|---------------------------------|
| Main sample, 2005-2023 | | | | | |
| Constant | -0.001 (0.001) | 0.006* (0.003) | 0.002 (0.004) | 0.007** (0.003) | -0.001 (0.001) |
| <i>PC1</i> | 0.003 (0.008) | 0.017 (0.011) | 0.007 (0.009) | 0.030** (0.014) | 0.001 (0.015) |
| <i>PC2</i> | -0.037 (0.034) | -0.047 (0.034) | -0.039 (0.034) | -0.050 (0.034) | -0.032 (0.032) |
| <i>PC3</i> | 0.020 (0.165) | 0.092 (0.170) | 0.051 (0.180) | 0.138 (0.159) | 0.014 (0.167) |
| π^* proxy | | -0.397* (0.203) [0.072] | -0.178 (0.197) [0.248] | -0.527** (0.217) [0.018] | 0.017 (0.078) [0.549] |
| R^2 | -0.69 | 0.41 | -0.80 | 1.61 | -1.13 |
| Subsample, 2005-2019 | | | | | |
| Constant | -0.000 (0.001) | 0.004 (0.004) | 0.002 (0.004) | 0.003 (0.004) | 0.001 (0.001) |
| <i>PC1</i> | 0.011 (0.010) | 0.018* (0.010) | 0.013 (0.010) | 0.020* (0.011) | 0.049*** (0.017) |
| <i>PC2</i> | -0.017 (0.039) | -0.030 (0.041) | -0.021 (0.041) | -0.031 (0.041) | -0.103*** (0.037) |
| <i>PC3</i> | -0.060 (0.171) | 0.005 (0.181) | -0.024 (0.191) | -0.008 (0.184) | 0.024 (0.153) |
| π^* proxy | | -0.287 (0.215) [0.161] | -0.147 (0.203) [0.296] | -0.244 (0.266) [0.209] | -0.296*** (0.089) [0.012] |
| R^2 | -0.92 | -0.66 | -1.20 | -1.12 | 2.31 |
| Extended sample, 1991-2023 | | | | | |
| Constant | -0.001 (0.000) | -0.000 (0.001) | -0.000 (0.001) | | -0.001 (0.000) |
| <i>PC1</i> | 0.007** (0.003) | 0.012** (0.005) | 0.011** (0.005) | | 0.011 (0.008) |
| <i>PC2</i> | -0.027 (0.025) | -0.032 (0.025) | -0.031 (0.025) | | -0.034 (0.025) |
| <i>PC3</i> | -0.016 (0.156) | -0.019 (0.158) | -0.019 (0.158) | | -0.014 (0.157) |
| π^* proxy | | -0.043 (0.030) [0.148] | -0.032 (0.028) [0.231] | | -0.012 (0.021) [0.319] |
| R^2 | 1.02 | 1.18 | 1.00 | | 0.87 |

*** p < 0.01, ** p < 0.05, * p < 0.1

Table A.3: **Spanning tests based on one-period realised monthly excess returns regressions.** Regressions of realised monthly excess returns averaged across 2 to 10-year maturities. The three first principal components describing the term structure of the ILS rates are used as predictors. For each maturity, the ILS-only specification (column 1) is compared to a specification with an observed π_t^* proxy (columns 2 to 5). Regression coefficients are displayed along with Newey-West standard errors (four lags) in parentheses (with significance levels highlighted with stars), and small-sample (one-sided) p-values for the observed π_t^* proxy obtained with the bootstrap method of Bauer and Hamilton (2018) in brackets. The data are observed at monthly frequency.

| | (1) | (2) | (3) | (4) | (5) |
|----------------------------|----------------------|----------------------|----------------------|----------------------|----------------------|
| | ILS rates only | Avg. surveys | CE | SPF | DMA |
| Main sample, 2005-2023 | | | | | |
| Constant | -0.013*** (0.005) | -0.039 (0.028) | -0.041* (0.024) | -0.039 (0.033) | -0.018** (0.008) |
| <i>PC1</i> | 0.049 (0.041) | -0.001 (0.048) | 0.024 (0.036) | -0.031 (0.082) | -0.074 (0.113) |
| <i>PC2</i> | -0.559*** (0.147) | -0.519*** (0.138) | -0.536*** (0.136) | -0.514*** (0.140) | -0.205 (0.258) |
| <i>PC3</i> | 0.789 (0.753) | 0.469 (0.783) | 0.489 (0.742) | 0.352 (0.866) | 0.387 (0.662) |
| π^* proxy | | 1.547 (1.577) | 1.550 (1.186) | 1.709 (1.980) | 1.261 (1.092) |
| R^2 | 13.76 | 14.47 | 14.74 | 14.53 | 17.38 |
| Subsample, 2005-2019 | | | | | |
| Constant | -0.005* (0.003) | -0.005 (0.016) | -0.010 (0.016) | -0.016 (0.019) | 0.001 (0.003) |
| <i>PC1</i> | 0.068* (0.040) | 0.068 (0.047) | 0.065 (0.040) | 0.045 (0.054) | 0.304*** (0.052) |
| <i>PC2</i> | -0.288* (0.154) | -0.288* (0.162) | -0.281* (0.157) | -0.258 (0.168) | -0.821*** (0.147) |
| <i>PC3</i> | 0.076 (0.572) | 0.075 (0.583) | 0.007 (0.576) | -0.043 (0.633) | 0.605 (0.495) |
| π^* proxy | | 0.003 (0.921) | 0.277 (0.830) | 0.682 (1.162) | -1.786*** (0.311) |
| R^2 | 6.81 | 6.22 | 6.38 | 6.56 | 25.12 |
| Extended sample, 1991-2023 | | | | | |
| Constant | -0.010** (0.004) | -0.007* (0.004) | -0.007* (0.004) | | -0.011*** (0.004) |
| <i>PC1</i> | 0.086*** (0.024) | 0.111*** (0.036) | 0.106*** (0.035) | | 0.140*** (0.049) |
| <i>PC2</i> | -0.432*** (0.107) | -0.456*** (0.112) | -0.455*** (0.112) | | -0.531*** (0.110) |
| <i>PC3</i> | 0.499 (0.687) | 0.485 (0.687) | 0.482 (0.689) | | 0.514 (0.683) |
| π^* proxy | | -0.214 (0.151) | -0.192 (0.150) | | -0.162 (0.113) |
| R^2 | 20.64 | 20.98 | 20.92 | | 21.46 |

*** p < 0.01, ** p < 0.05, * p < 0.1

Table A.4: **Spanning tests based on one-period realised annual excess returns regressions.** Regressions of realised annual excess returns averaged across 2 to 10-year maturities. The three first principal components describing the term structure of the ILS rates are used as predictors. For each maturity, the ILS-only specification (column 1) is compared to a specification with an observed π_t^* proxy (columns 2 to 5). Regression coefficients are displayed along with Newey-West standard errors (four lags) in parentheses (with significance levels highlighted with stars), and small-sample (one-sided) p-values for the observed π_t^* proxy obtained with the bootstrap method of Bauer and Hamilton (2018) in brackets. The data are observed at monthly frequency.

Appendix B An affine term structure model with fixed endpoint

The specification of the model with a fixed endpoint follows Joslin et al. (2011). In this model, the latent factors, X , follow a vector autoregressive model of order one under both the physical and the risk-neutral probability measure, \mathbb{P} and \mathbb{Q} respectively, and the short-term inflation rate

is affine in the latent factors:

$$X_t = \mathcal{K}_{0,X}^{\mathbb{P}} + \mathcal{K}_{1,X}^{\mathbb{P}} X_{t-1} + \epsilon_{t,X}^{\mathbb{P}}, \quad (\text{B.3})$$

$$X_t = \mathcal{K}_{0,X}^{\mathbb{Q}} + \mathcal{K}_{1,X}^{\mathbb{Q}} X_{t-1} + \epsilon_{t,X}^{\mathbb{Q}}, \quad (\text{B.4})$$

$$\pi_{t,1} = \rho_{0,X} + \rho_{1,X} X_t, \quad (\text{B.5})$$

where $t = 1, \dots, T$ indexes time, X_t is a column vector of length N , and $\pi_{t,1}$ is the (implicit) short-term (one-month) ILS rate. Both $\epsilon_{t,X}^{\mathbb{P}}$ and $\epsilon_{t,X}^{\mathbb{Q}}$ are column vectors that are independently and identically distributed (*i.i.d.*) according to normal distributions with the same covariance matrix Ω_X , i.e. $\epsilon_{t,X}^{\mathbb{P}}, \epsilon_{t,X}^{\mathbb{Q}} \sim i.i.d. N(0, \Omega_X)$.

In order to ensure identification of the factors while keeping maximum flexibility (Dai and Singleton, 2000), $\mathcal{K}_{1,X}^{\mathbb{Q}}$ is diagonal (with entries assumed to be real, distinct and less than 1), $\rho_{1,X}$ is a vector of ones and, following the most general specification in Joslin et al. (2011), $\rho_{0,X} = 0$ and $\mathcal{K}_{0,X}^{\mathbb{Q}} = (k_{\infty}^{\mathbb{Q}}, 0, \dots, 0)'$.

Given the above, it can be shown that observed ILS rates are affine in X_t :

$$\pi_{t,m}^o = A_{m,X} + B_{m,X} X_t + \nu_{m,t}, \quad (\text{B.6})$$

where m indicates maturity in months, and $A_{m,X}$ and $B_{m,X}$ are determined by the usual recursions:

$$A_{m+1,X} = A_{m,X} + B'_{m,X} \mathcal{K}_{0,X}^{\mathbb{Q}} + \frac{1}{2} B'_{m,X} \Omega_X B_{m,X} - \rho_{0,X}, \quad (\text{B.7})$$

$$B_{m+1,X} = B_{m,X} \mathcal{K}_{1,X}^{\mathbb{Q}} - \rho_{1,X}, \quad (\text{B.8})$$

initiated with $A_{0,X} = 0$ and $B_{0,X} = 0$. To ease notations, the A 's and B 's of the rates that are matched by the model are often written in matrix form: $A_X = (\dots, A_{m,X}, \dots)'$ and $B_X = (\dots, B_{m,X}, \dots)'$. Measurement error terms are assumed to share the same variance across maturities $\nu_{m,t} \sim i.i.d. N(0, \sigma_{\nu}^2)$.

Joslin et al. (2011) show that considering the rates' principal components, \mathcal{P} , as perfectly observed rotated factors greatly facilitates the estimation of the model. In short, $\mathcal{P}_t^o = W \Pi_t^o = \mathcal{P}_t = W \Pi_t = W A_X + W B_X X_t$, where W is an $N \times J$ matrix of portfolio weights given by the eigenvectors associated with the N greatest eigenvalues of the J rates in Π_t ($N < J$), and

Π_t^o and Π_t collect all observed ILS rates and their model counterparts respectively. Given this rotation, the model can be rewritten as:

$$\mathcal{P}_t = \mathcal{K}_{0,\mathcal{P}}^{\mathbb{P}} + \mathcal{K}_{1,\mathcal{P}}^{\mathbb{P}}\mathcal{P}_{t-1} + \epsilon_{t,\mathcal{P}}^{\mathbb{P}}, \quad (\text{B.9})$$

$$\mathcal{P}_t = \mathcal{K}_{0,\mathcal{P}}^{\mathbb{Q}} + \mathcal{K}_{1,\mathcal{P}}^{\mathbb{Q}}\mathcal{P}_{t-1} + \epsilon_{t,\mathcal{P}}^{\mathbb{Q}}, \quad (\text{B.10})$$

$$\pi_t = \rho_{0,\mathcal{P}} + \rho_{1,\mathcal{P}}\mathcal{P}_t, \quad (\text{B.11})$$

$$\Pi_t = A_{\mathcal{P}} + B_{\mathcal{P}}\mathcal{P}_t + \nu_t, \quad (\text{B.12})$$

where $\epsilon_{t,\mathcal{P}}^{\mathbb{P}}, \epsilon_{t,\mathcal{P}}^{\mathbb{Q}} \sim i.i.d. N(0, \Omega_{\mathcal{P}})$, $\mathcal{K}_{0,\mathcal{P}}^{\mathbb{P}}, \mathcal{K}_{1,\mathcal{P}}^{\mathbb{P}}, \mathcal{K}_{0,\mathcal{P}}^{\mathbb{Q}}, \mathcal{K}_{1,\mathcal{P}}^{\mathbb{Q}}, \Omega_{\mathcal{P}}, \rho_{0,\mathcal{P}}, \rho_{1,\mathcal{P}}, A_{\mathcal{P}}$ and $B_{\mathcal{P}}$ are transformations of their counterparts in the model with X replacing \mathcal{P} , and $\nu_t \sim i.i.d. N(0, I_J \otimes \sigma_{\nu}^2)$.

In this setup, the model can easily be estimated in two steps. The first step simply consists in estimating $\mathcal{K}_{0,\mathcal{P}}^{\mathbb{P}}$ and $\mathcal{K}_{1,\mathcal{P}}^{\mathbb{P}}$ from Equation (B.9) by Ordinary Least Squares (OLS). Given these OLS estimates, the second step consists in maximising the joint likelihood of equations (B.9) and (B.12) over $\{k_{\infty}^{\mathbb{Q}}, \mathcal{K}_{1,X}^{\mathbb{Q}}, \Omega_{\mathcal{P}}, \sigma_{\nu}^2\}$, taking care that $A_{\mathcal{P}}$ and $B_{\mathcal{P}}$ are the transformations of A_X and B_X that obey the recursions in Equations (B.7)-(B.8). Conveniently, the set of parameters $\{k_{\infty}^{\mathbb{Q}}, \sigma_{\nu}^2\}$ can be concentrated out of the likelihood function, which implies that the likelihood (given $\mathcal{K}_{0,\mathcal{P}}^{\mathbb{P}}$ and $\mathcal{K}_{1,\mathcal{P}}^{\mathbb{P}}$) is maximised by searching for the optimal diagonal matrix $\mathcal{K}_{1,X}^{\mathbb{Q}}$ and the elements of the Cholesky decomposition of $\Omega_{\mathcal{P}}$, which in total represents $N * (N + 3)/2$ elements.⁷ If three principal components are considered ($N = 3$), only nine model parameters must be estimated.

Appendix C Details of the Bauer and Rudebusch (2020) model

As shown by Bauer and Rudebusch (2020), Equations 3 can be recast into a VAR by defining $Z_t = (\pi_t^*, \mathcal{P}_t)'$:

$$Z_t = \mu_Z + \Phi_Z Z_{t-1} + v_t, \quad v_t \sim N(0, \Omega_Z), \quad (\text{C.13})$$

where

⁷Concentrating σ_{ν}^2 out of the likelihood function is done by approximating it as the variance of the residuals in Equation B.12 given $A_{\mathcal{P}}$ and $B_{\mathcal{P}}$, see Joslin et al. (2011, footnote 24). Besides, $k_{\infty}^{\mathbb{Q}}$ can be concentrated out of the likelihood function as shown in Joslin, Le, and Singleton (2013).

$$\mu_Z = \begin{pmatrix} 0 \\ (I_N - \Phi)\bar{\mathcal{P}} \end{pmatrix}, \Phi_Z = \begin{pmatrix} 1 & 0_{1 \times N} \\ (I_N - \Phi)\gamma & \Phi \end{pmatrix}, \text{ and } \Omega_Z = \begin{pmatrix} \sigma_\eta^2 & \gamma'\sigma_\eta^2 \\ \gamma\sigma_\eta^2 & \Omega \end{pmatrix}.$$

The innovations to \mathcal{P}_t are $u_t = \gamma\eta_t + \tilde{\mu}_t$ and have covariance matrix $\Omega = \gamma\gamma'\sigma_\eta^2 + \tilde{\Omega}$.

References

- Bañbura, M., F. Brenna, J. Paredes, and F. Ravazzolo (2021a). Combining Bayesian VARs with survey density forecasts: does it pay off? *ECB Working Paper 2543*.
- Bañbura, M., D. Leiva-Leon, and J.-O. Menz (2021b). Do inflation expectations improve model-based inflation forecasts? *ECB Working Paper 2604*.
- Bauer, M. D. and J. D. Hamilton (2018). Robust bond risk premia. *The Review of Financial Studies* 31, 399–448.
- Bauer, M. D. and G. D. Rudebusch (2017). Resolving the spanning puzzle in macro-finance term structure models. *Review of Finance* 11, 511–553.
- Bauer, M. D. and G. D. Rudebusch (2020). Interest rates under falling stars. *American Economic Review* 5, 1316–1354.
- Beechey, M. J., B. K. Johansson, and A. T. Levin (2011). Are long-run inflation expectations anchored more firmly in the euro area than in the United States? *American Economic Journal: Macroeconomics* 3, 104–129.
- Boeckx, J., L. Iania, and J. Wauters (2024). Macroeconomic drivers of inflation expectations and inflation risk premia. *NBB Working Paper 446*.
- Boneva, L., B. Böninghausen, L. F. Rousová, and E. Letizia (2019). Derivatives transactions data and their use in central bank analysis. *ECB Economic Bulletin* 6/2019.
- Brand, C., G. Goy, and W. Lemke (2021). Natural rate chimera and bond pricing reality. *ECB Working Paper 2612*.
- Burban, V., B. D. Backer, F. Schupp, and A. Vladu (2022). Decomposing market-based measures of inflation compensation into inflation expectations and risk premia. *ECB Economic Bulletin* 8/2021.
- Burban, V. and F. Schupp (2023). Backcasting real rates and inflation expectations – combining market-based measures with historical data for related variables. *ECB Economic Bulletin* 2/2023.

- Camba-Méndez, G. and T. Werner (2017). The inflation risk premium in the post-Lehman period. *ECB Working Paper 2033*.
- Cecchetti, S., A. Grasso, and M. Pericoli (2022). An analysis of objective inflation expectations and inflation risk premia. *Banca d'Italia, Temi di Discussione 1380*.
- Christensen, J. H., F. X. Diebold, and G. D. Rudebusch (2011). The affine arbitrage-free class of Nelson–Siegel term structure models. *Journal of Econometrics* 164, 4–20.
- Christensen, J. H., 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, 143–178.
- Christensen, J. H. and G. D. Rudebusch (2012). The response of interest rates to US and UK quantitative easing. *The Economic Journal* 122, 385–414.
- Cieslak, A. and P. Povala (2015). Expected returns in treasury bonds. *The Review of Financial Studies* 28, 2859–2901.
- Corsello, F., S. Neri, and A. Tagliabracci (2021). Anchored or de-anchored? That is the question. *European Journal of Political Economy* 69, 1–19.
- Dai, Q. and K. Singleton (2000). Specification analysis of affine term structure models. *Journal of Finance* 55, 1943 – 1978.
- Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2017). Safety, liquidity, and the natural rate of interest. *Brookings Papers on Economic Activity Spring*, 235–316.
- Del Negro, M., D. Giannone, M. P. Giannoni, and A. Tambalotti (2019). Global trends in interest rates. *Journal of International Economics* 118, 248–262.
- Dewachter, H. and L. Iania (2011). An extended macro-finance model with financial factors. *Journal of Financial & Quantitative Analysis* 46(6), 1893–1916.
- Dewachter, H. and M. Lyrio (2006). Macro factors and the term structure of interest rates. *Journal of Money, Credit and Banking* 38(1), 119–140.
- Diercks, A. M., C. Campbell, S. Sharpe, and D. Soques (2023). The swaps strike back: Evaluating expectations of one-year inflation. *Finance and Economics Discussion Series (FEDS)*.

- ECB (2021). Inflation expectations and their role in Eurosystem forecasting. *ECB Occasional Paper 264*.
- García, J. A. and S. E. V. Werner (2021). Inflation news and euro-area inflation expectations. *International Journal of Central Banking* 17(3), 1–60.
- 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.
- Hördahl, P. and O. Tristani (2014). Inflation risk premia in the Euro Area and the United States. *International Journal of Central Banking* 10, 1–47.
- Joslin, S., A. Le, and K. J. Singleton (2013). Why Gaussian macro-finance term structure models are (nearly) unconstrained factor-VARs. *Journal of Financial Economics* 109(3), 604–622.
- Joslin, S., M. Priebisch, and K. J. Singleton (2014). Risk premiums in dynamic term structure models with unspanned macro risks. *Journal of Finance* 69(3), 1197–1232.
- 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. and A. Orphanides (2012). Term structure estimation with survey data on interest rate forecasts. *Journal of Financial and Quantitative Analysis* 47(1), 241–272.
- Kozicki, S. and P. A. Tinsley (2001). Shifting endpoints in the term structure of interest rates. *Journal of Monetary Economics* 47(3), 613–652.
- Łyziak, T. and M. Paloviita (2017). Anchoring of inflation expectations in the euro area: Recent evidence based on survey data. *European Journal of Political Economy* 46, 52–73.
- Patton, A. J. and A. Timmermann (2011). Predictability of output growth and inflation: A multi-horizon survey approach. *Journal of Business & Economic Statistics* 29(3), 397–410.
- Wei, M. and J. Wright (2013). Reverse regressions and long-horizon forecasting. *Journal of Applied Econometrics* 28(3), 353–371.