BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Survey-based daily estimates of inflation expectations and risk premia in the euro area

by Francesca Lilla and Gabriele Zinna

Number 991

December 2025

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Survey-based daily estimates of inflation expectations and risk premia in the euro area

by Francesca Lilla and Gabriele Zinna

Number 991 – December 2025

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of Banca d'Italia and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within Banca d'Italia, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

SURVEY-BASED DAILY ESTIMATES OF INFLATION EXPECTATIONS AND RISK PREMIA IN THE EURO AREA

by Francesca Lilla* and Gabriele Zinna*

Abstract

In this research note, we develop real-time, daily estimates of survey-based inflation expectations and inflation risk premia for the euro area using a simple, nearly model-free approach. We first estimate daily inflation expectations by projecting low-frequency survey measures onto a small set of daily financial variables. We then compute daily inflation risk premia as the difference between inflation-linked swap (ILS) forward rates and the estimated survey-based expectations. The resulting daily estimates provide timely insights for conjunctural analysis and monetary policy, offering a useful complement to measures derived from dynamic term structure models.

JEL Classification: F31, G12, G15.

Keywords: inflation expectations, inflation risk premia, monetary policy.

DOI: 10.32057/0.QEF.2025.991

* Bank of Italy, Directorate General for Economics, Statistics and Research – Economic Outlook and Monetary Policy Directorate.

1. Introduction[1]

Market-based measures of inflation compensation, such as inflation-linked swaps (ILS), are useful inputs to the conduct of monetary policy, especially because they are available daily and for many horizons. However, they embed both inflation expectations and inflation risk premia. While both components are highly informative, they convey substantially different signals, underscoring the importance of disentangling them. Inflation expectations reflect market participants' views on future inflation and on the degree of inflation anchoring, and are therefore key inputs for assessing the appropriate monetary policy stance. Inflation risk premia, instead, capture the compensation investors require for bearing inflation risk, providing valuable insights — also from a monetary policy perspective — into the expected correlation between output growth and inflation, and hence into the likely prevalence of demand- versus supply-driven shocks.[2] The decomposition of market-based measures of inflation compensation into expectations and risk premia has typically been achieved using dynamic term structure models, sometimes with the inclusion of survey information (e.g., Burban, De Backer, and Vladu 2024; Boeckx, Iania, and Wauters 2025).

Recently, a growing literature has estimated asset risk premia in a model-free manner, relying only on analyst surveys to capture investors' inflation expectations (e.g., Cieslak 2018; Schmeling, Schrimpf, and Steffensen 2022; Bocola et al. 2024). This is because surveys can be interpreted as a direct measure of investors' (subjective) expectations.[3]

[1]We are indebted for constructive comments and suggestions to Martina Cecioni, Stefano Neri, Alessandro Secchi, and Fabrizio Venditti. All errors are our responsibility. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Bank of Italy.

[2]From a theoretical perspective, the inflation risk premium should be positive if inflation tends to rise during economic downturns – when the stochastic discount factor (SDF), i.e. the marginal utility of consumption, is elevated – and negative otherwise. In a standard New Keynesian framework, this implies that the inflation risk premium will be positive when supply-driven shocks primarily shape the economy, as these shocks push inflation higher while depressing growth, thereby reducing the real value of nominal bond payments precisely when they are most valuable. Conversely, if demand-driven shocks predominate, the inflation risk premium should be negative.

[3]Under the full information rational expectations (FIRE) hypothesis, the subjective probability measure of agents coincides with the objective (or real-world, statistical) measure. This implies that the subjective inflation expectations and risk premia, which are obtained from surveys, are equal to the objective ones obtained from term structure models. If, instead, the FIRE hypothesis does not hold an important implication is that the perceived risk premium of market agents is not measured by the objective risk premium but instead by the subjective risk premium (see, for example, Coroneo and Golinski,

Therefore, by analysing asset prices alongside surveys, one can quantify the risk premium without relying on complex asset pricing models, which tend to suffer from identification problems and require the specification of a functional form for the SDF. A clear shortcoming in the use of surveys is that they provide inflation forecasts for only a few horizons and at low frequency, which makes them inherently less attractive for conjunctural analysis.

This note proposes a real-time, daily decomposition of market-based measures of inflation compensation into inflation expectations and risk premia using survey data in an (almost) model-free manner. To derive daily inflation expectations from quarterly surveys, we project survey inflation expectations onto a set of high-frequency financial variables in real time, following the approach of Aronovich and Meldrum (2021), recently refined by Bocola et al. (2024).[4] In particular, we use inflation forecasts from the ECB's Survey of Professional Forecasters (SPF) and Consensus Economics (CE), focusing on one-year expectations one- and four-year ahead and on those at the five-to-ten-year horizon (1y1y, 1y4y and 5y5y, respectively; 5y5y expectations are only available from CE; see Appendix I for a detailed description of the survey data). We then retrieve daily estimates of inflation risk premia by subtracting from the ILS rates the inflation expectations estimated in the projection step.

In doing so, we obtain estimates of inflation expectations and risk premia that are broadly consistent with those in the literature. For example, we find that inflation premia are negative in the period 2015-2021, while they turn positive during the following period of high inflation and they have been fluctuating around zero until the end of our sample period. Inflation expetations at shorter maturities are more volatile than longer-term ones, which remained close to the 2% ECB inflation target.

Our estimates are particularly useful for interpreting high-frequency movements in ILS rates in response to recent events. For example, they show that the rise in 1y4y and 5y5y ILS rates around the fiscal announcement on March 4 by Germany's chancellor-designate

2023). This is another reason why it is important to look at both estimates from surveys and term structure models to monitor inflation expectations and risk premia.

[4]In brief, whenever new data from inflation surveys become available, a regression – estimated using either parametric or non-parametric methods – is used to relate this information to a set of financial variables observed at daily frequency. This relationship is then used to infer, in real time, the evolution of inflation expectations based on the observed daily financial variables until the next survey release.

Merz was driven almost entirely by higher inflation risk premia, which increased by around 17 basis points. The decline in 1y1y and 1y4y ILS rates following Trump's tariff announcement on 2 April instead reflected a decline in both inflation risk premia and inflation expectations. As of 11 April 2025 (the last day covered by the analysis), longer-term inflation expectations are anchored at 2% and risk premia are slightly positive; for the 1y1y and 1y4y horizons inflation expectations stand at 1.8 and 1.9% and risk premia are slightly negative.

Overall, these novel high-frequency estimates of inflation expectations and risk premia should be seen as complementary to those obtained using dynamic term structure models.

2. Fitting surveys with financial factors

The main insight of the projection step is that inflation expectations over a given horizon are driven by a small set of observable factors. In particular, survey expectations are projected onto inflation compensation measures, consisting of 2-, 5- and 10-year Inflation-Linked Swap (ILS) rates, and the 2-year Overnight Index Swap (OIS) rate, which could be seen as proxy for monetary policy, as done in Bocola et al. (2024)[5]:

$$S_s^{(n)} = \beta_{n,0} + \beta_{n,1}\bar{X}_{s-d} + u_s^{(n)} \quad \text{for } s = [s07Q1, s07Q2, \ldots, s25Q1] \tag{1}$$

where n is the horizon of the survey (i.e., 1y1y, 1y4y, 5y5y), s is the day in which the survey is released, $s-d$ is the deadline for the respondents to submit their forecasts, which is typically few days before the release date, and $X_{s-d} = [ILS_{s-d}^{2y}, ILS_{s-d}^{5y}, ILS_{s-d}^{10y}, OIS_{s-d}^{2y}]$

[5] It is standard in the macro-finance literature to express the term structure of bond yields as a linear function of three linearly independent combinations of yields, namely the level, slope, and curvature. In fact, these three factors taken together explain almost entirely the information in the cross-section of bond yields. Using the level, slope, and curvature factors is essentially tantamount to using a short, medium, and long-term bond yield (e.g., Bauer and Rudebusch 2020). Three factors extracted from the cross-section of ILSs also seem to explain the bulk of the information in the cross-section of ILSs (e.g., Burban et al. 2024; Boeckx et al. 2025). Consistently, Aronovich and Meldrum (2021) use a short, medium, and long-term bond yield as conditioning information for fitting interest rate and inflation surveys. In modelling inflation surveys, we follow Bocola et al. (2024) in that we include not only the three ILS factors but also a short-term bond yield, which could be seen as a proxy for monetary policy (it is irrelevant if one were to use the 2-year real rate, as the two models would be observationally equivalent, given that also 2-year ILS are included). Later on, we will show that including this bond yield helps fit inflation surveys (see Appendix III)

collects the financial variables. We use the financial variables as w-day averages, $\bar{X}_{s-d} = \frac{1}{W}\sum_{i=0}^{W-1} X_{s-d-i}$. This averaging, in addition to reducing measurement error, captures the fact that respondents may use timely information available in the days prior to the deadline, and not just the information on the deadline date. Similar to Bocola et al. (2024), we use $w = 5$.[6] In sum, the regression is run at quarterly frequency, and the regressors are computed as averages of ILS and OIS rates over the five-day period ending on the deadline date.

We use both parametric and non-parametric approaches to capture time-variability in the relation between survey inflation expectations and regressors. We find that estimating the model over the full sample delivers a poor fitting of inflation expectations (see Appendix II). We therefore explore different ways to capture the time-variability in the relation between expectations and regressors; in particular, we consider three following three approaches.

I. *Recursive linear model (ReL):* We estimate Eq. (1) over expanding windows, using an initial window of six years, and then progressively add one quarter at a time as new surveys are released, to get a series of time-varying loadings, $\beta_{s,n}$.

II. *Rolling linear model (RoL):* We estimate Eq. (1) over rolling windows of six years, by adding one quarter at a time while removing the most distant quarter, to get a series of time-varying loadings, $\beta_{s,n}$.

III. *Local linear regression model (LLR):* We posit that the survey over a given horizon n is a non-parametric function of observable factors, $S_s^{(n)} = f^{(n)}(\bar{X}_{s-d}) + \varepsilon_s^{(n)}$, whereby $\bar{X}_{s-d}$ is the same set of ILS and OIS regressors as in Eq. (1), but $f^{(n)}(\cdot)$ is a non-parametric function of the regressors. Specifically, the fit $\hat{f}^{(n)}(x)$ at a target point x_0 is obtained by estimating a weighted linear regression model, whereby observations are weighted by their proximity to the target point. This implies that more distant observations from the target point take negligible weights, making the relation

[6] For example, suppose that the release date for the 2020:Q4 survey is s=30/10/2020, then the deadline to reply is $s - d$=09/10/2020. This implies that the $\bar{X}_{s-d}$, i.e. ILS and OIS variables, are computed as the averages over the five-day window that goes from 05/10/2020 to 09/10/2020.

between the $S^{(n)}$ and $\bar{X}$ locally linear. This model is estimated recursively, i.e. over expanding windows.[7]

Aronovich and Meldrum (2021) use the local linear regression method (LLR), while Bocola et al. (2024) use the linear model estimated over rolling windows (RoL), arguing that this simpler model performs reasonably well. In what follows, we assess the performance of each model and regressor in explaining survey data.

The fit improves considerably when allowing for time-variation in the loadings. Consistent with Bocola et al. (2024), we find that the RoL estimator produces results similar with those of the LLR estimator, although the latter seems to better capture more abrupt movements in survey data (see Figure A2 in Appendix II for SPF; evidence for CE is available upon request).

A more formal assessment of the relative goodness of fit of the models indicates that the LLR approach has the best performance, with fitting errors decreasing at longer horizons (Table 1). The LLR model outperforms the competing models, especially the ReL. The RoL model, is the second best; it improves upon the ReL at all maturities, likely as it better accounts for abrupt shifts in the relation between survey expectations and the conditioning information. The superior performance of the LLR model over the RoL model diminishes at longer horizons and essentially disappears at the 5y5y horizon. This suggests that nonlinearities are important, but also that linear models that account for shifts in the relationships between variables in a timely manner, like the linear rolling model, perform relatively well, especially for longer-horizon inflation expectations.

A "leave-one-out" exercise, which drops one of the regressors at the time, suggests that the full model, i.e. the model which includes the four variables, achieves the highest performance across all horizons (1y1y, 1y4y, and 5y5y) and surveys (either SPF or CE). This is generally true for both LLR and RoL.[8] Specifically, in this exercise each of the

[7]We refer to Aronovich and Meldrum (2021) for a detailed description of the local linear estimation method. Here we only note that we choose the bandwidth parameters using a standard leave-one-out cross-validation technique. These parameters control how much smoothing the model applies when estimating the regression function at each target point and therefore determine which nearby observations matter and how strongly they are weighted. Put differently, the bandwidth determines how fast the kernel weight drops as distance from the target point increases.

[8]Using LLR, the only exceptions concern the fit of the 1y1y and 1y4y inflation expectations of the

Table 1: Model fit

The first row shows the mean absolute percentage pricing errors (MAPPE) for the local linear regression model (LLR) when inflation expectations are measured according to the SPF and CE. The second and the third rows show the MAPPE of the alternative models relative to the LLR model. A ratio greater (smaller) than 1 indicates an underperformance (outperformance) of the LLR model. *ReL* means that the linear regression estimates start with a 6-year window and progressively increase in size until the end of the sample. *RoL* indicates that the linear regressions are estimated using a 6-year rolling window. *LLR* denotes the recourse to a pure non-linear model. The MAPPE is computed as $100 \times \frac{1}{N}\sum_{s=1}^{N}\left|\frac{S_s-\hat{S}_s}{S_s}\right|$, where S_s is the observed survey data, while $\hat{S}_s$ is the corresponding fitted value; N is the number of observations or fitted values.

	SPF		**CE**		
	1y1y	1y4y	1y1y	1y4y	5y5y
LLR	3.80	1.46	3.94	1.90	1.50
ReL/LLR	1.44	1.61	1.62	1.50	1.41
RoL/LLR	1.07	1.04	1.31	1.15	1.00

competing models includes three regressors and we compare the resulting MAPPEs from each of the restricted models with the full four-factor model, to evaluate each regressor contribution to the model performance (see Appendix III, Table A1). Of particular interest is that the model performance deteriorates substantially when the OIS^{2y} is excluded from the set of regressors. This evidence suggests that one should include all four regressors: a proxy for monetary policy, namely the short-term nominal rate, along with the inflation compensation variables of different tenors.[9]

3. Decomposing ILS rates into expectations and risk premia

Having identified the best-performing models, we then turn to present the daily estimates of inflation expectations and risk premia. Before doing so, we explain how we align survey data with the timing of high-frequency financial market data and briefly review the

CE. Specifically, the performance improves when ILS^{10y} (ILS^{2y}) is excluded from the regression for 1y1y (1y4y).

[9] To alleviate concerns about overfitting, we perform and discuss a Cross-Validation exercise for both LLR and RoL (we discard ReL given its poor performance) in Appendix IV. In essence, overfitting manifests when the model fits the training data well but struggles to generalize to "unseen" data. This is potentially a particularly relevant issue in our context, as we combine the estimated loadings with unseen data. However, Table A2 shows that overfitting is not an issue in any of the tested models, surveys, or forecast horizons.

decomposition of measures of inflation compensation into expectations and risk premia. We also highlight the main differences in obtaining such a decomposition using surveys versus term structure models.

From standard no-arbitrage asset pricing, we have:

$$IC_t^{(n)} = E_t^{\mathbb{Q}}[\bar{\pi}^{(n)}] = E_t^{j}[\bar{\pi}^{(n)}] + RP_t^{(n),j}, \quad j = \mathbb{S}, \mathbb{P} \tag{2}$$

where $IC_t^{(n)}$ is the inflation compensation measure (i.e., the ILS rate in this case) measured at time t with maturitiy n, $\bar{\pi}^{(n)}$ is the average inflation rate between time t and $t+n$, and $E_t^{\mathbb{Q}}[\bar{\pi}^{(n)}]$ denotes the conditional expectation under the $\mathbb{Q}$ risk-neutral pricing measure. This expectation can be written as the sum of the conditional expectation of the average inflation, either under the subjective measure from surveys ($j = \mathbb{S}$) or the objective measure from statistical models ($j = \mathbb{P}$), and the inflation risk premium, $RP_t^{(n)}$.[10] Under the full informational rational expectations (FIRE) hypothesis, the agents' subjective probability measure coincides with the objective measure, i.e. $\mathbb{S} = \mathbb{P}$.[11]

Thus, the inflation risk premium can be estimated using survey data as follows:

$$RP_t^{\mathbb{S},(n)} = ILS_t^{(n)} - E_t^{\mathbb{S}}[\bar{\pi}^{(n)}] = ILS_t^{(n)} - \hat{S}_t^{(n)} \tag{3}$$

However, the drawback is that surveys are available at very low frequency, being released for example quarterly, so that we could estimate the risk premium only for $t \in s = [s_{07Q1}, s_{07Q2}, \ldots, s_{25Q1}]$. Therefore, to obtain daily estimates of the inflation risk premium, we need first to project the surveys on a set of daily factors. Specifically, as

[10] Both subjective $E_t^{\mathbb{S}}[\cdot]$ and objective $E_t^{\mathbb{P}}[\cdot]$ expectations are expectations which do not include risk premia. The main difference is that the former are measured from surveys, while the latter are obtained from statistical models, as in dynamic term structure models (DTSMs). The distinction between the two becomes empirically blurred when the estimation of DTSMs uses surveys to anchor the objective expectation from the statistical model, by adding measurement equations which link the surveys with the objective expectations (allowing for a measurement error). This is for example the case in the term structure models of Cecchetti et al. (2022) and Burban et al. (2024).

[11] Cieslak (2018) emphasizes the possibility of a failure of the FIRE assumption, allowing for biases in beliefs. Nevertheless, the interpretation of the subjective risk premium remains unchanged regardless of whether FIRE holds or not, as it can be interpreted as the compensation required by the investors for holding risky assets, i.e., assets that are exposed to inflation risks. By contrast, the objective risk premium would be difficult to interpret, as it merely represents the measured ex-post risk premium based on the history of ILS (see Coroneo and Golinski 2023, for a similar point on the bond risk premium).

illustrated in the previous section, we need first to estimate,

$$S_s^{(n)} = g_s(\bar{X}_{s-d}) + \varepsilon_s^{(n)} \quad \text{for } s = [s_{07Q1}, s_{07Q2}, \ldots, s_{25Q1}], \tag{4}$$

where $\hat{g}_s(\cdot)$ can be either the *ReL*, *RoL*, or *LLR* approach. Then, using $\hat{g}_s(\cdot)$, we can retrieve the daily fitted surveys using the daily $\bar{X}_t$, for $s - d < t < s_+ - d$, where s_+ is the survey released after s, as:

$$\hat{S}_t^{(n)} = \hat{g}_s(\bar{X}_t). \tag{5}$$

To give a practical example, let s_{12Q4} denote the latest survey released for a given horizon n; the objective is to "update" this survey with the new information available until the release of the next quarter survey. We estimate the quarterly regression model of Eq. (4) from s_{07Q1} to s_{12Q4}, for example. Using Eq. (5), we then compute the fitted daily inflation expectations for all days that go from the day of the release for the survey s_{12Q4} to the day before the release of the next survey s_{13Q1}. So, using the estimated exposures $\hat{g}_s(\cdot)$ and the daily regressors $\bar{X}_t$, we obtain the inflation expectations, i.e. $\hat{S}_{t_1}, \hat{S}_{t_2}, \ldots$, at daily frequency until the next survey data become available. When a new survey is released, we re-estimate the model by using the surveys until s_{13Q1}, either with a rolling or recursive scheme, and we re-run the nowcasting exercise, following the same steps as before.

Finally, we compute the daily inflation risk premia from the ILS and the fitted daily surveys as:

$$RP_t^{\mathbb{S},(n)} = ILS_t^{(n)} - \hat{S}_t^{(n)} \approx ILS_t^{(n)} - E_t^{\mathbb{S}}[\bar{\pi}^{(n)}]. \tag{6}$$

These estimates of the risk premium differ from those obtained from dynamic term structure models:

$$RP_t^{\mathbb{P},(n)} = E_t^{\mathbb{Q}}[\bar{\pi}^{(n)}] - E_t^{\mathbb{P}}[\bar{\pi}^{(n)}], \tag{7}$$

for two main reasons. First, we assume that $ILS_t^{(n)}$ is a pure measure of risk-neutral inflation expectations, while term structure models allow for measurement error that should capture microstructure noise, illiquidity, for example. Nevertheless, in term structure models the difference between the observed $ILS_t^{(n)}$ and the model-implied $E_t^{\mathbb{Q}}[\bar{\pi}^{(n)}]$ is typically very small, being in the order of few basis points. Second, and more important, we use subjective expectations based on surveys rather than objective ones computed using autoregressive forecasting models. What is common between the two approaches is that both $E_t^{\mathbb{P}}[\bar{\pi}^{(n)}]$ and $\hat{S}_t^{(n)}$ are estimated using similar conditioning information consisting of a small set of financial factors (e.g., level, slope, and curvature).[12]

4. Empirical results

Figure 1 illustrates the daily estimates of the inflation expectations (top panel) and risk premia (bottom panel) for the 1y1y, 1y4y, and 5y5y horizons, covering the period from 16 November 2012 to 11 April 2025. Daily observations are computed as averages across the estimates obtained from the two best-performing models (LLR and RoL) and the two surveys under analysis (SPF, CE), when are both available. We present the results as averages for brevity, and also because model averaging typically helps achieve more accurate estimates.

Over the last decade, our measures of inflation expectations and risk premia are broadly consistent with those in the literature; for example, they are similar to those provided by the ECB (Burban et al. 2022).[13] Overall, we find that 1y1y expectations are more volatile than the longer-term expectations (Figure 1), despite using the same conditioning information.[14]

[12]Inflation expectations and risk premia obtained from term structure models have the advantage of being constructed imposing no-arbitrage across inflation rates of different maturities. For this reason, one can compute inflation expectations and risk premia for any horizon, not being confined to the maturity of the inflation compensation measures used in the estimation. By contrast, expectations and risk premia from surveys can only be constructed for the maturity for which surveys are available.

[13]Our estimates instead are somewhat different from those of Cecchetti et al. (2022). Indeed, they identify a positive risk premium at the 5y5y horizon throughout the entire sample period (2012Q4-2021Q1), while the component attributed to inflation expectations consistently remains below 2%. Their estimates also differ from the ones of the ECB, which show a negative inflation risk premium during the periods 2016-2017 and 2019-2021, alongside inflation expectations that persistently fall below 1.6%.

[14]An important point to note is that, despite daily estimates of inflation expectations display more variability than quarterly survey data due to the day-to-day fluctuations in financial data, they remain much less volatile than ILS forward rates, as shown in the Appendix V.

Figure 1: Estimates of daily inflation expectations and inflation risk premia

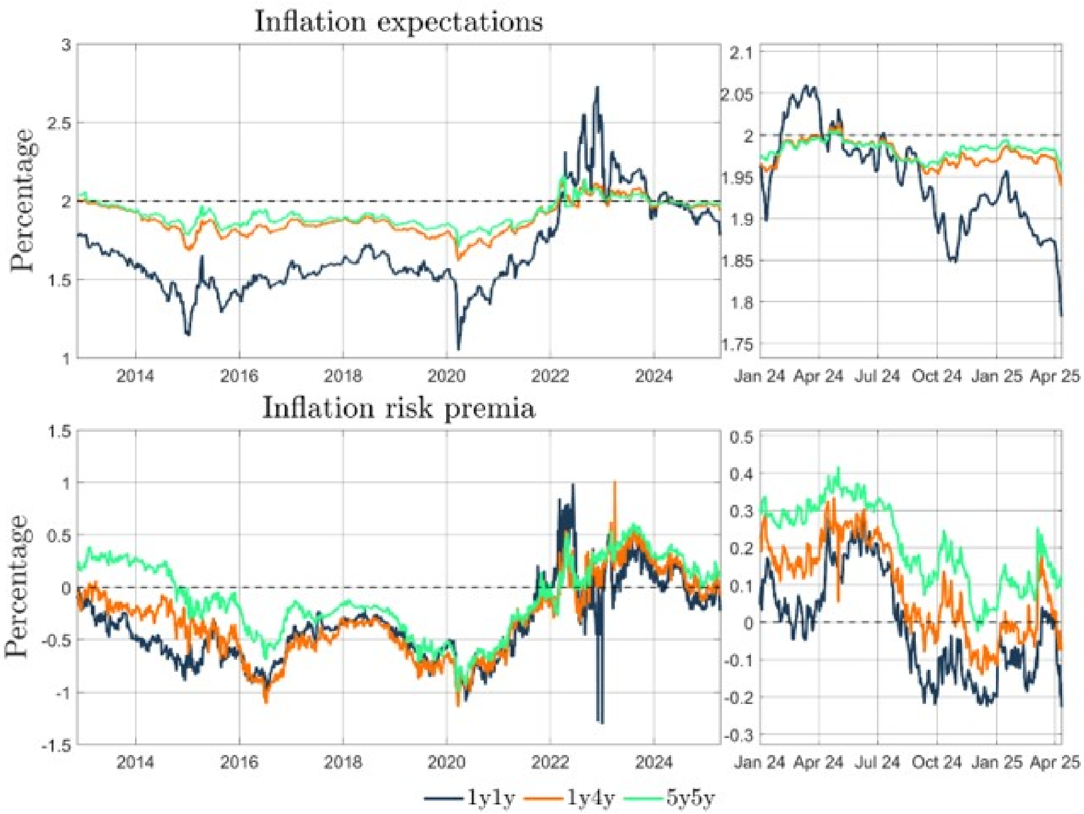

Note: The upper panel reports the estimates of expected inflation, calculated as the average of estimates from RoL and LLR approaches, based on data from SPF and CE both 1y1y and 1y4y inflation expectations. Estimates for 5y5y inflation expectations are computed as the average of RoL and LLR estimates from CE. The lower panel shows the corresponding inflation risk premia.

Moreover, from 2014 until the outbreak of Covid-19, short-term inflation expectations were significantly lower than their longer-term counterparts. With the inflation surge, however, expectations increased significantly at all horizons, and all exceeded the 2% ECB's price stability objective around the time of the Russian invasion of Ukraine, albeit to a very limited extent at longer horizons. Similarly, after years of negative values, inflation risk premia turned positive in 2022, consistent with expectations of a future prevalence of supply shocks. Indeed, positive inflation risk premia are consistent with investors expecting that the economy will be dominated by supply shocks over the considered horizon. In such scenarios, bonds act as poor hedges to equities (a proxy for the inverse of the SDF), as in a supply-driven economy recessions feature high inflation rates.[15] For this reason, in a supply-driven economy, investors are willing to pay a premium to hold financial instruments that offer protection to inflation risks.

Figure 2: Inflation risk premia and the balance of risk

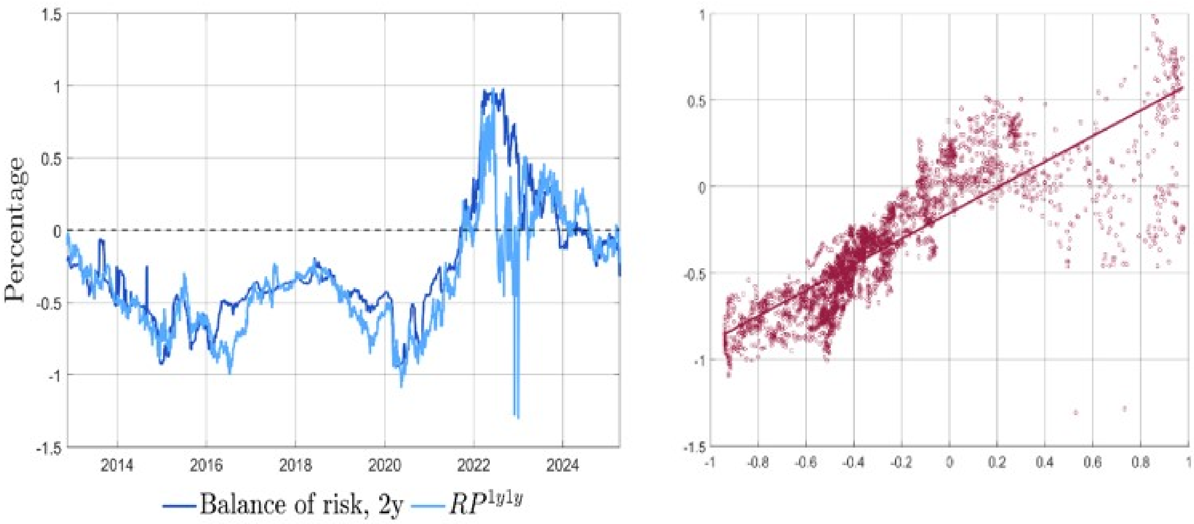

Note: The left panel reports the balance of risks (Cecchetti et al. 2021) which is constructed as the difference between the probability of 2-year inflation being higher than 3 per cent and lower than 1 per cent. The right panel shows the scatter plot between the balance of risks at the 2-year horizon and the estimated 1y1y inflation risk premium.

Interestingly, our estimates of inflation risk premia appear to be closely related to other measures of inflation risks, such as the balance of risk computed from ILS options (Figure 2). Indeed, our estimates of inflation risk premia and the balance of risks tend to move closely together; that is, inflation risk premia tend to be positive (negative) when the balance of risk is also positive (negative). Albeit to a lower extent, inflation risk premia also tend to co-move with estimates of the stock and bond correlation (Figure 3). The latter is also informative about the prevalence of supply versus demand shocks, as theory predicts that the stock and bond correlation should also switch sign with the correlation between output and inflation.[16]

Focusing on the recent disinflationary period, in 2024 short-term inflation expectations

[15]Put differently, in a supply-driven scenario the variable leg of the ILS contract offers an important hedge against the state of the economy, ILS returns are high (low) when GDP growth is low (high), and investors are willing to pay an inflation risk premium (on top of the expected inflation) to the holder of the fixed leg of the ILS contract.

[16]Nevertheless, the correlation between stock and bond could also reflect other financial factors, especially at higher frequencies. Moreover, a recent paper by Duffee (2022) documents that, although macro news accounts for a significant proportion of the stock-bond correlation, its effects are not always consistent with existing theoretical frameworks.

gradually declined, while medium to long-term ones remained close to the 2% target; during the year inflation risk premia also decreased towards zero. In the first months of 2024, inflation risk premia were positive, but they declined substantially from the spring onwards, turning negative at short and medium-term horizons in August 2024 as the turmoil in international financial markets unfolded, partly fuelled by fears of recession risks in the U.S. amid weaker demand. At that time euro area economic activity was slowing down while disinflation was proceeding, leading the ECB's Governing Council to start reducing the degree of monetary policy restrictiveness. Compared to the previous years, in 2024 our estimates show contained risk premia (Figure 1) pointing to balanced inflation risks as evidenced also from other indicators.

Figure 3: Inflation risk premia and stock-bond correlations

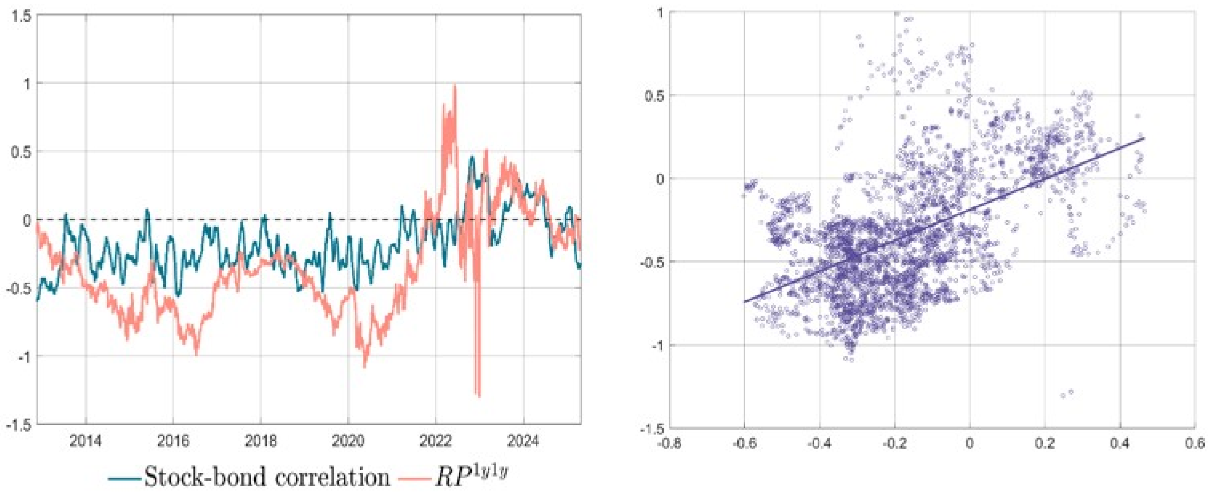

Note: The left panel reports the stock-bond correlation and the estimated 1y1y inflation risk premium. The correlation is the 20-day moving average of daily estimates of the stock-bond correlation, computed using infra-day changes (5 minutes) in the prices of futures on the Euro Stoxx 50 and the 10-year German Bund. The right panel shows the scatter plot between the stock-bond correlation and the estimated 1y1y inflation risk premium.

According to our estimates, Merz's announcement of German future fiscal spending on defence and infrastructure on 4 March 2025 resulted in higher inflation risk premia and unchanged inflation expectations (Figure 4). Following the announcement, ILSs rose sharply at the 1y4y and 5y5y horizons, driven by higher risk premia (from 3 to 5 March they increase by 18 and 16 basis points, respectively), on the back of largely unchanged inflation expectations. In that day ILS signalled a rapid repricing of medium to longer-term inflation risks by investors which however retraced in the following days as market participants' reassessed the economic outlook ahead of Trump's announcement of new tariffs.

Figure 4: Merz fiscal spending and Trump tariff announcements

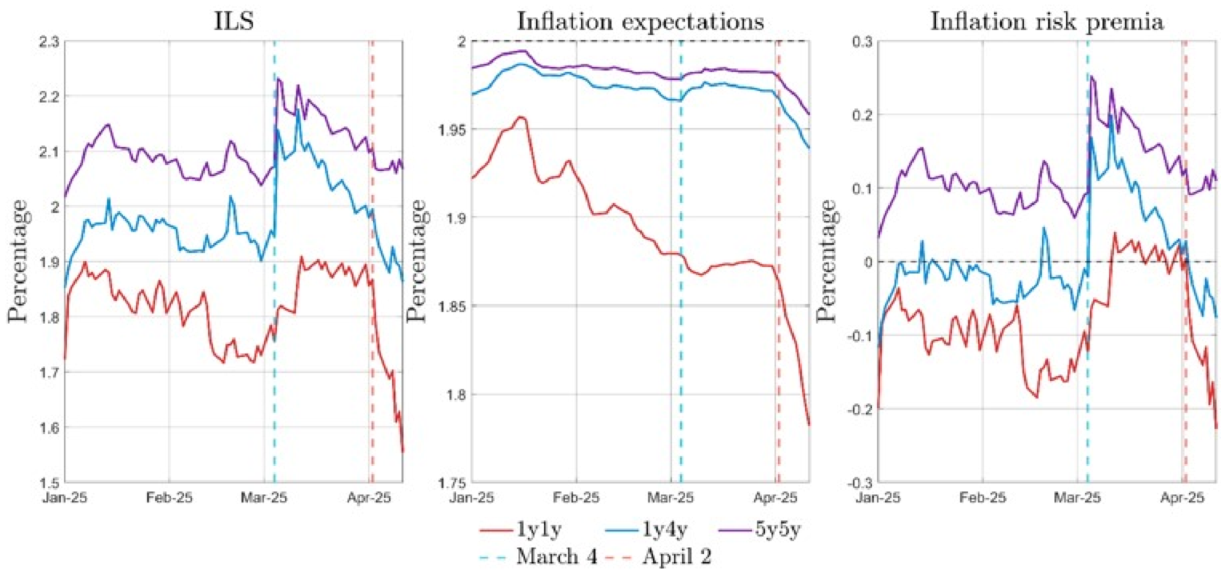

Note: The plots show the evolution of inflation-linked swap (ILS) rates, inflation expectations, and inflation risk premia for the 1y1y, 1y4y and 5y5y horizons from 1 January 2025 to 8 April 2025. Two key dates are highlighted: March 4, when German parlament approved a constitutional amendment to allow the government to take out substantial new loans for defense and infrastructure projects, and April 2, when U.S. President Trump announced new global tariffs.

Trump's tariff announcement on 2 April led to a decrease in both inflation expectations and risk premia. As the tariffs were larger and broader than expected, not only risk premia fell further, but also inflation expectations reduced, especially at the shorter horizon. For example, at the 1y1y horizon risk premia and expectations declined by 22 and 9 basis points, respectively, from their levels on April 1; they reduced further in the following days (Figure 4). This supports the view that tariffs could be disinflationary in the euro area.

References

Aronovich, A. and A. Meldrum (2021). High-frequency estimates of the natural real rate and inflation expectations. Technical Report 2021-034, Board of Governors of the Federal Reserve System.

Bauer, M. D. and G. D. Rudebusch (2020). Interest rates under falling stars. *American Economic Review 110*(5), 1316–1354.

Bocola, L., A. Dovis, K. Jørgensen, and R. Kirpalani (2024). Bond market views of the fed. Technical report, National Bureau of Economic Research.

Boeckx, J., L. Iania, and J. Wauters (2025). Macroeconomic drivers of inflation expectations and inflation risk premia. *Journal of Financial Econometrics 23*(1), 733–743. Oxford University Press.

Burban, V., B. De Backer, F. Schupp, and A. L. Vladu (2022). Decomposing market-based measures of inflation compensation into inflation expectations and risk premia. Technical Report Issue 8, European Central Bank.

Burban, V., B. De Backer, and A. L. Vladu (2024). Inflation (de-)anchoring in the euro area. Technical Report No. 457, National Bank of Belgium.

Cecchetti, S., D. Fantino, M. Riggi, A. Notarpietro, A. Tagliabracci, A. Tiseno, and R. Zizza (2021). Inflation expectations in the euro area: indicators, analyses and models used at Banca d'Italia. Technical Report No. 612, Banca d'Italia.

Cecchetti, S., A. Grasso, and M. Pericoli (2022). An analysis of objective inflation expectations and inflation risk premia. Technical Report No. 1380, Banca d'Italia.

Cieslak, A. (2018). Short-rate expectations and unexpected returns in Treasury bonds. *The Review of Financial Studies 31*(9).

Coroneo, L. and A. Golinski (2023). Information in (and not in) interest rates surveys. Technical report, University of York. Mimeo, available at SSRN 4339327.

Duffee, G. (2022). Macroeconomic news and stock–bond comovement. *Review of Finance*, 1–24.

Schmeling, M., A. Schrimpf, and S. Steffensen (2022). Monetary policy expectation errors. *Journal of Financial Economics 146*, 841–858.

I. Inflation surveys

In what follows, we briefly present the survey forecasts of euro area inflation used in the note. We consider the ECB's Survey of Professional Forecasters (SPF) and Consensus Economics (CE).

The SPF collects point forecasts of the annual inflation rate both at fixed and rolling horizons. We use the rolling horizon two years ahead of the latest available data and the longer-term horizon to measure 1y1y and 1y4y inflation expectations.[17] They are available at quarterly frequency starting from 1999.

CE provides forecasts for the annual inflation rate with horizons starting at the current calendar year and extending up to five years ahead, as well as a longer-term forecast capturing average yearly inflation from six to ten years ahead (5y5y hereafter). To retrieve the rolling or fixed horizon surveys from the calendar surveys, we use a simple linear interpolation.[18] CE survey forecasts for the euro area have been available since 2003 and are provided on a semi-annual basis until 2014. Before 2014, we obtain quarterly inflation expectations simply by linearly interpolating the semi-annual surveys.[19]

In this note, we do not consider very short-term surveys as the focus of this note is on estimating longer-term inflation expectations and risk premia. Thus, we centre the analysis on surveys which forecast yearly inflation one and four years ahead (1y1y and 1y4y, respectively); for CE, we also use the six-to-ten year ahead forecasts (5y5y).[20]

By looking at SPF and CE inflation expectations from 2007Q1 to 2025Q1 (unreported), it is apparent that shorter-term inflation forecasts exhibited greater fluctuations than longer-term forecasts and reached much lower values in the period from 2009 to 2021. Indeed, long-term inflation forecasts were on average closer to the 2% ECB's target, consistent with inflation level anchoring at longer horizons. At the beginning of 2025, both SPF and CE longer-term inflation expectations were very close to 2%, while the 1y1y expectations were only slightly below such level. CE and SPF forecasts tend to move similarly, but at times can deviate notably. For example, during the recent high inflation period, CE expectations reached somewhat higher values than the SPF ones at the 1y1y horizon, 3.1% and 2.4%, respectively. Possible explanations for the divergences are that (i) the pool of survey participants is different, and (ii) the exact time at which the two surveys are conducted also differs, especially in the early part of the sample.

[17]The longer-term forecast represents the annual inflation rate (year on year change) in the fourth calendar year for Q1 and Q2 vintages of the survey in each year and in the fifth calendar year for the Q3 and Q4 vintages. Although this survey cannot be considered a true rolling horizon, we treat it as a proxy for the 1y4y horizon, given the shifting reference window across survey rounds. The interpolation approach cannot be applied, as the calendar year forecasts four and five years ahead are not available.

[18]For example, the rolling 1y1y survey for the 2024Q2 vintage ($s_{24Q2}^{roll:1y1y}$) is computed as a weighted average of the calendar ones for the years 2025 ($s_{24Q2}^{cal:25}$) and 2026 ($s_{24Q2}^{cal:26}$): $s_{24Q2}^{roll:1y1y} = 0.75 \times s_{24Q2}^{cal:25} + 0.25 \times s_{24Q2}^{cal:26}$. For the 1y4y rolling horizon, the same formula is applied using the calendar surveys for the years 2028 and 2029: $s_{24Q2}^{roll:1y4y} = 0.75 \times s_{24Q2}^{cal:28} + 0.25 \times s_{24Q2}^{cal:29}$.

[19]Using CE surveys, one could recover expectations of average inflation from two to five years ahead from the published yearly figures, and the 5y5y forecasts as the average inflation rate from six to ten years ahead. We therefore extract both, given its relevance in the ECB definition of long-term expectations.

[20]Specifically, the 1y1y, 1y4y, and 5y5y horizons are defined in line with ECB conventions.

II. Regression fit for quarterly SPF and CE

We find that the fit of the linear model of Eq. (1) is rather poor, as the model seems not sufficiently accurate in capturing some of the broad movements in inflation expectations, particularly over longer-time horizons. For example, the model is not able to fully account for the increase in the 5y5y inflation expectations observed in 2022 (unreported). The linear model also exhibits a rather poor fit at the 1y1y and 1y4y horizons, both for SPF (Figure A1) and CE (unreported).

Figure A1: Regression fit for quarterly surveys: linear model, SPF

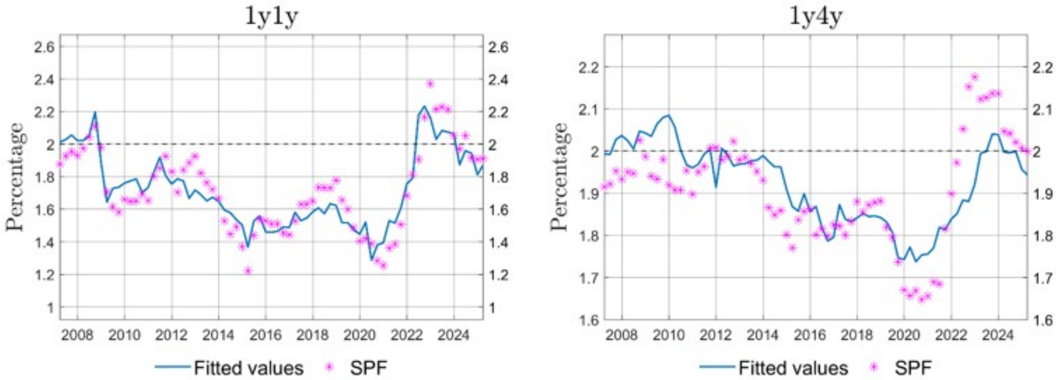

Note: The plots show the observed surveys (asterisk) and the quarterly fitted values, in the case of SPF. The fitted values refer to the linear regression in Eq. (1) estimated over the entire sample period, from 2007-Q1 to 2025-Q1.

Figure A2 reports the fitted 1y1y and 1y4y SPF inflation surveys obtained using the time-varying (ReL or RoL) and nonparametric (LLR) estimators (the CE evidence is available upon request).

Figure A2: Regression fit for quarterly surveys: alternative models, SPF

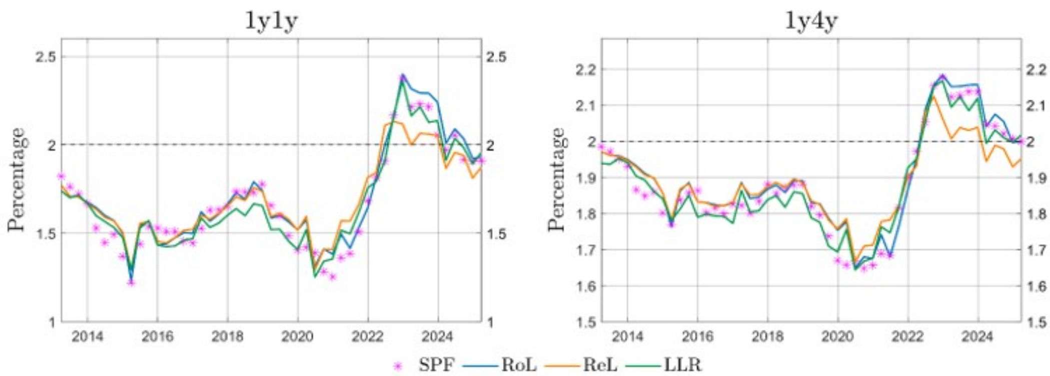

Note: The plots show the observed SPF surveys (asterisk) and the quarterly fitted values for the alternative estimation models. "RoL" indicates that the linear regressions are estimated using a 6-year rolling window. "ReL" means that the linear regression estimation starts with a 6-year window and progressively increases in size until the end of the sample. "LLR" presents the results for the pure non-linear model.

III. Regressors relevance

Table 1.D reports the results in terms of MAPPE from the "leave-one-out" exercise, where one regressor is excluded at a time (leaving 3 regressors per model) to assess the importance of each X-independent variable in predicting inflation expectations.

Table A1: Leave-one-out performance based on MAPPE

The table shows the MAPPE for the "leave-one-out" exercise for all alternative estimation models in the case of SPF and CE. The "X" columns indicate the regression performance when all independent variables are considered, namely $X = [ILS^{2y}, ILS^{5y}, ILS^{10y}, OIS^{2y}]$. The remaining columns show the results for the regression that excludes the variable written in the curly brackets. For example, $X_{(-ILS^{2y})}$ means that the ILS^{2y} variable is not considered as a regressor. In bold are reported the smallest MAPPEs for each survey and estimation method (LLR, RoL).

Panel A: SPF					
			LLR		
	X	$X_{(-ILS^{2y})}$	$X_{(-ILS^{5y})}$	$X_{(-ILS^{10y})}$	$X_{(-OIS^{2y})}$
1y1y	**3.80**	4.18	3.90	4.02	5.57
1y4y	**1.46**	1.81	1.59	1.50	2.24
			RoL		
	X	$X_{(-ILS^{2y})}$	$X_{(-ILS^{5y})}$	$X_{(-ILS^{10y})}$	$X_{(-OIS^{2y})}$
1y1y	**4.07**	4.61	4.34	4.36	4.53
1y4y	**1.52**	1.88	1.73	1.75	1.60
Panel B: CE					
			LLR		
	X	$X_{(-ILS^{2y})}$	$X_{(-ILS^{5y})}$	$X_{(-ILS^{10y})}$	$X_{(-OIS^{2y})}$
1y1y	3.94	5.74	4.18	**3.85**	5.78
1y4y	1.90	**1.65**	1.83	1.86	2.23
5y5y	**1.50**	1.54	1.63	1.73	1.93
			RoL		
	X	$X_{(-ILS^{2y})}$	$X_{(-ILS^{5y})}$	$X_{(-ILS^{10y})}$	$X_{(-OIS^{2y})}$
1y1y	**5.17**	5.80	5.57	5.45	5.56
1y4y	**2.19**	2.44	2.36	2.33	2.47
5y5y	**1.50**	1.76	1.71	1.55	1.69

IV. Overfitting

Overfitting is a common challenge in time-series modelling, occurring when a model fits the historical data too closely, capturing both true patterns and random noise. To mitigate this issue, it is essential to balance model complexity with generalization. In this note, we aim to reduce overfitting by using an appropriately sized rolling window which helps avoiding a too precise fitting of short-term fluctuations while enhancing the ability of the model to adapt to evolving data patterns. In LLR, instead, we use cross-validation techniques to select optimal bandwidths, along with a Gaussian kernel, as both choices help prevent overfitting by promoting a smoother fit and ensuring better generalization.

Despite these considerations regarding the two estimation approaches that are most suitable for our analysis, we perform a Leave-One-Out Cross-Validation (LOO-CV) exercise to detect any potential overfitting. Specifically, LOO-CV plays a critical role in detecting overfitting by providing an unbiased and accurate evaluation of model performance. Indeed, it ensures that each data point is used for both training and testing, making it a powerful tool for identifying when a model might be overfitting. This method helps ensure the model generalizes well to new, unseen data, rather than simply memorizing the training data.

In the LLR model, where the model is highly sensitive to the local data distribution, LOO-CV guarantees that each observation is used for both training and testing, minimizing bias and offering a robust assessment of the model's ability to generalize. Similarly, in the RoL model, where the model is updated with new data over time, LOO-CV enables out-of-sample testing by evaluating the model's predictions at each time point while preserving the temporal order and avoiding data leakage. This cross-validation method maximizes the use of available data, making it especially suitable for time series or small datasets, such as in our case.

In this analysis, we compare the performance of the LLR and RoL models across all dimensions (surveys and horizons) by evaluating in-sample and out-of-sample errors. For each model and survey at a specific horizon, a substantial gap between these metrics would be indicative of overfitting, as it suggests that the model fits the training data well but struggles to generalize to unseen data. To formally test for overfitting, we apply either the Paired t-test or the Wilcoxon signed-rank test, depending on the distribution of the error differences. The Paired t-test is typically used when the differences between in-sample and out-of-sample errors are approximately normally distributed, as it compares the means of two related groups. However, if the differences are not normally distributed or if there are outliers, the Wilcoxon signed-rank test, a non-parametric alternative, should be preferred. This test evaluates whether the median of the differences is significantly different from zero. In sum, by examining these error metrics and choosing the appropriate test, we can assess the models' ability to generalize and determine whether overfitting affects their predictive power on new, unseen data.

Table A2 presents the results of overfitting tests for both RoL and LLR models, assessing their in-sample (IS) and out-of-sample (OOS) performance. The IS RMSE measures how well the model fits the training data, while the OOS RMSE evaluates its ability to generalize to unseen data. For all models, surveys, and forecast horizons, the Wilcoxon test p-values exceed 0.05, indicating no statistically significant difference between in-sample and out-of-sample performance. This suggests that neither model suffers from overfitting, as their predictive accuracy remains consistent across training and validation sets. Overall, the results confirm that overfitting is not an issue in any of the tested models, surveys, or forecast horizons in this analysis.

Table A2: Leave-One-Out Cross-Validation for Overfitting

This table presents the results of overfitting tests for Rolling Estimation (RoL) and Local Linear Regression (LLR) in the case of SPF and CE at all horizons considered. The evaluation is based on the comparison between in-sample and out-of-sample Root Mean Squared Error (RMSE). A Wilcoxon signed-rank test is conducted to assess whether the difference between in-sample and out-of-sample errors is statistically significant. A p-value greater than 0.05 indicates that no significant overfitting is detected.

	Estimation approach	**Average IS RMSE**	**Average OOS RMSE**	**Wilcoxon p-value**	**Conclusion**
			SPF		
1y1y	RoL	0.0786	0.1198	0.1478	No significant overfitting
	LLR	0.0759	0.0924	0.1880	No significant overfitting
1y4y	RoL	0.0340	0.0458	0.7240	No significant overfitting
	LLR	0.0348	0.0391	0.4018	No significant overfitting
			CE		
1y1y	RoL	0.0980	0.1365	0.3627	No significant overfitting
	LLR	0.0761	0.1210	0.0560	No significant overfitting
1y4y	RoL	0.0492	0.0678	0.4062	No significant overfitting
	LLR	0.0465	0.0496	0.7188	No significant overfitting
5y5y	RoL	0.0495	0.0465	0.1809	No significant overfitting
	LLR	0.0434	0.0418	0.1634	No significant overfitting

V. Daily inflation expectation variability

A key advantage of our estimates of inflation expectations is that they are available at any frequency corresponding to market data, making these estimates particularly valuable for policymakers and practitioners who are interested in interpreting high-frequency movements in ILS rates. As is typical, we can view ILS as reflecting the sum of the average expected inflation and the inflation risk premium, which compensates investors for risk. Given that financial market prices fluctuate daily, our daily estimates also exhibit day-to-day variation. Given the tendency for inflation expectations and inflation risk premia to move together, we would expect daily changes in our estimates to be less volatile than those in ILS rates. If inflation expectations are well anchored, we would also expect them to remain relatively stable from one day to the next. By comparing the empirical distribution of the daily changes in forward ILS rates with that of our daily estimates of the survey-based inflation expectations, we can verify this.

Figures A3 and A4 show the empirical distributions for both types of surveys and for all the horizons considered in this note. Indeed, the resulting evidence confirms our prior in that the variability of daily changes in the estimated inflation expectations (represented by the red dashed lines) is significantly smaller than that of the raw forward ILS rates. For longer-term inflation expectations (1y4y and 5y5y), most daily changes fall between -1 basis point and +1 basis point. The distribution of daily changes in 1y1y inflation expectations is somewhat more spread out, though the majority still falls between -2 basis points and +2 basis points. These observations support our prior conclusion that long-horizon expectations are not subject to large daily fluctuations.

Figure A3: Daily changes of forward ILS and inflation expectations, SPF

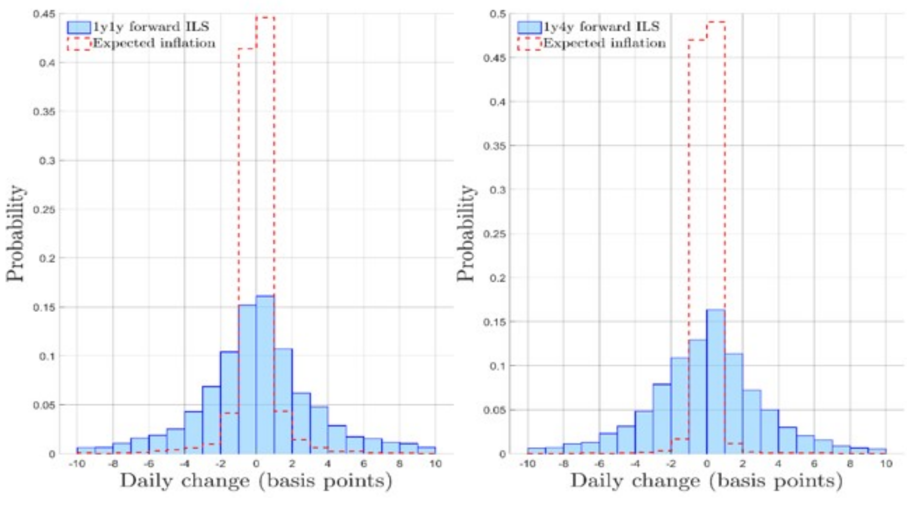

Note: The plots show the observed surveys (asterisk) and the quarterly fitted values for the alternative estimation models, in the case of SPF. The plots show histograms of daily changes in the forward ILS (blue bars) alongside histograms of daily changes in our estimates of inflation expectations (red dashed lines) from SPF. The probability for each bin on the x-axis refers to the proportion of daily changes falling in the 1-basis-point interval above the axis label. For instance, the bin labelled "0" refers to the range from 0 basis points to 1 basis point.

Figure A4: Daily changes of forward ILS and inflation expectations, Consensus Economics

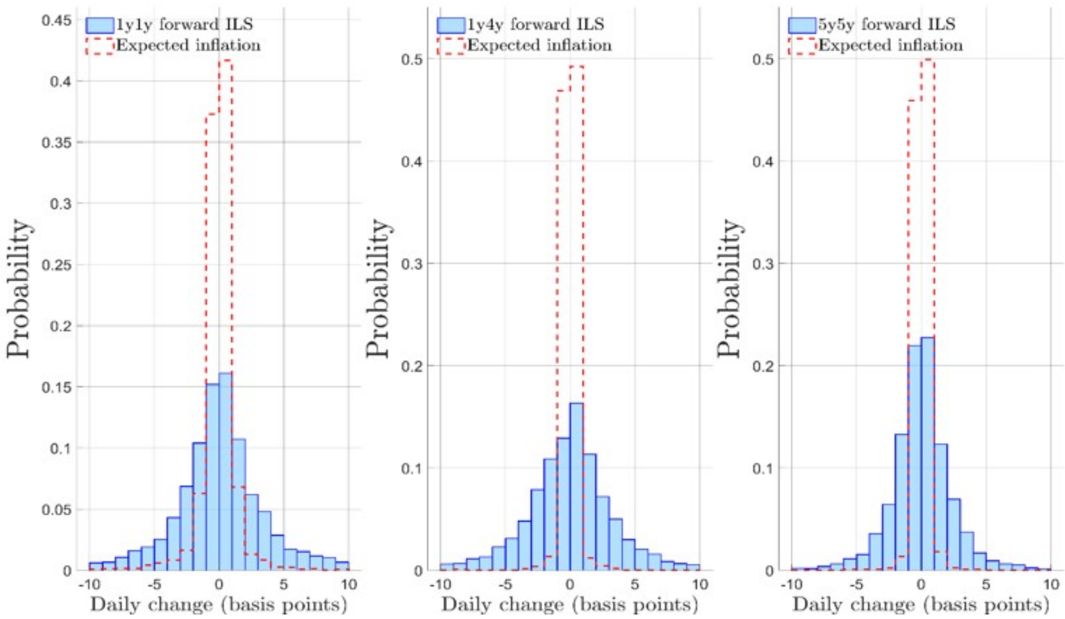

Note: The plots show the observed surveys (asterisk) and the quarterly fitted values for the alternative estimation models, in the case of CE. The plots show histograms of daily changes in the forward ILS (blue bars) alongside histograms of daily changes in our estimates of inflation expectations (red dashed lines) from CE. The probability for each bin on the x-axis refers to the proportion of daily changes falling in the 1-basis-point interval above the axis label. For instance, the bin labelled "0" refers to the range from 0 basis points to 1 basis point.