

Camba-Mendez, Gonzalo; Werner, Thomas

Working Paper

The inflation risk premium in the post-Lehman period

ECB Working Paper, No. 2033

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Camba-Mendez, Gonzalo; Werner, Thomas (2017) : The inflation risk premium in the post-Lehman period, ECB Working Paper, No. 2033, ISBN 978-92-899-2755-0, European Central Bank (ECB), Frankfurt a. M.,
<https://doi.org/10.2866/556561>

This Version is available at:

<https://hdl.handle.net/10419/162684>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK
EUROSYSTEM

Working Paper Series

Gonzalo Camba-Mendez, Thomas Werner

The inflation risk premium
in the post-Lehman period

No 2033 / March 2017

Disclaimer: This paper should not be reported as representing the views of the European Central Bank (ECB). The views expressed are those of the authors and do not necessarily reflect those of the ECB.

Abstract

In this paper we construct model-free and model-based indicators for the inflation risk premium in the US and the euro area. We study the impact of market liquidity, surprises from inflation data releases, inflation volatility and deflation fears on the inflation risk premium. For our analysis, we construct a special dataset with a broad range of indicators. The dataset is carefully constructed to ensure that at every point in time the series are aligned with the information set available to traders. Furthermore, we adopt a Bayesian variable selection procedure to deal with the strong multicollinearity in the variables that potentially can explain the movements in the inflation risk premium. We find that the inflation risk premium turned negative, on both sides of the Atlantic, during the post-Lehman period. This confirms the recent finding by Campbell et al. (2016) that nominal bonds are no longer “inflation bet” but have turned into “deflation hedges”. We also find, and contrary to common beliefs, that indicators of inflation uncertainty alone cannot explain the movements in the inflation risk premium in the post-Lehman period. The decline in the inflation risk premium seems mostly related to increased deflation fears and the belief that inflation will stay far away from the monetary policy target rather than declining inflation uncertainty. This in turn would suggest that central banks should not be complacent with low or even negative inflation risk premia.

JEL classification: E44, G17.

Keywords: Inflation linked swaps, inflation risk premium, inflation expectations.

NON-TECHNICAL SUMMARY

In this paper we construct model-free and model-based indicators for the inflation risk premium in the US and the euro area. We also study the impact of market liquidity, surprises from inflation data releases, inflation volatility and deflation fears on the inflation risk premium. The literature (see Fleckenstein et al. (2016)) is currently focusing on a possible recent change in the sign of the inflation risk premium. In particular Campbell et al. (2016) claim that nominal bonds changed from “inflation bet” to “deflation hedge”. Theoretically the sign of the inflation risk premium depends on the correlation of inflation with the marginal inter-temporal rate of substitution of consumption of the representative investor. As it is very difficult to measure empirically this correlation we apply two more direct approaches to estimate the inflation risk premium for the US and the euro area. Our first approach is “model-free” and similar to Soderlind (2011) and is based on comparing inflation forecasts from Consensus Economics with ILS rates. Our second “model-based” approach applies a standard affine term structure model to the term structure of ILS rates. We contribute to the literature by deriving a discrete time equivalent of the approach proposed by Beechey (2008) and Finlay and Wendel (2012) that can be easily applied to ILS data using standard programme code.

One methodological advantage of our analysis compared to the literature is the construction of a consistent data-set for ILS rates. Lags implicit in ILS contracts and settlement date conventions are taken into consideration when computing the “implicit” one-year-ahead inflation forecast embedded in the ILS curve. As we use Consensus Economics data for the construction of the model-free estimate of the inflation risk premium, we take as key dates for our regression analysis the dates on which the Consensus Economics data are collected. The financial market data will also be retrieved for these same dates. In this manner, the quoted prices for the ILS contracts and the survey data from Consensus are associated with the same information set. The inflation rates reported in Consensus Economics are not, however, aligned with the “year-on-year” (e.g. March-on-March) rates naturally embedded in ILS quotes. In this paper, we thus also propose a novel method to render inflation readings from Consensus Economics consistent with ILS quotes.

Another methodological advantage of our analysis is the application of a sophisticated Bayesian variable selection procedure for which we assess performance using a Monte Carlo study tailored towards our estimation strategy.

Our results provide strong evidence, robust across business areas and measures for the inflation risk premium, that the inflation risk premium turned negative, on both sides of the Atlantic, during the post-Lehman period. This confirms the recent finding by Campbell et al. (2016) that nominal bonds are no longer “inflation bet” but have turned into “deflation hedges”. We also find, and contrary to common beliefs, that indicators of inflation uncertainty alone cannot explain the movements in the inflation risk premium in the post-Lehman period. The decline in the inflation risk premium seems mostly related to increased deflation fears and the belief that inflation will stay far away from the monetary policy target rather than declining inflation uncertainty. This in turn would suggest that central banks should not be complacent with low or even negative inflation risk premia.

1 Introduction

From an investor's perspective, a bond that offers a nominal coupon payment of say 5% is more attractive when annual inflation turns out to be 1% rather than when it turns out to be 5%. Inflation risks are inherent in most investment strategies, and not surprisingly, investors commonly demand compensation for bearing those risks. The inflation risk premium is usually defined as the compensation demanded by investors to hold financial assets which are subject to inflation risks. Market-based measures of inflation expectations like the so-called break-even inflation rate (BEIR) and the inflation-linked swap (ILS) rate do not only contain information on inflation expectations but also on the inflation risk premium.¹

The literature (see Fleckenstein et al. (2016)) is currently focusing on a possible recent change in the sign of the inflation risk premium. In particular Campbell et al. (2016) claim that nominal bonds changed from “inflation bet” to “deflation hedge”. Theoretically the sign of the inflation risk premium depends on the correlation of inflation with the marginal inter-temporal rate of substitution of consumption of the representative investor. As it is very difficult to measure empirically this correlation we apply two more direct approaches to estimate the inflation risk premium for the US and the euro area. Our first approach is “model-free” and similar to Soderlind (2011) and is based on comparing inflation forecasts from Consensus Economics with ILS rates. Our second “model-based” approach applies a standard affine term structure model to the term structure of ILS rates. We contribute to the literature by deriving a discrete time equivalent of the approach proposed by Beechey (2008) and Finlay and Wendel (2012) that can be easily applied to ILS data using standard programme code. Our results provide strong evidence, robust across business areas and measures for the inflation risk premium, that the inflation risk premium turned negative during the great recession. The aim of the paper is also to shed more light on the main drivers of the inflation risk premium.

One methodological advantage of our analysis compared to the literature is the construction of a consistent data-set for ILS rates. Lags implicit in ILS contracts and settlement date conventions are taken into consideration when computing the “implicit” one-year-ahead inflation forecast embedded in the ILS curve. As we use Consensus Economics data for the construction of the model-free estimate of the inflation risk premium, we take as key dates for our regression analysis the dates on which the Consensus Economics data are collected. The financial market data will also be retrieved for these same dates. In this manner, the quoted prices for the ILS contracts and the survey data from Consensus are associated with

¹The BEIR is defined as the difference between the sovereign inflation-linked bond (a real bond) yield and an equivalent sovereign nominal bond yield.

the same information set. The inflation rates reported in Consensus Economics are not, however, aligned with the “year-on-year” (e.g. March-on-March) rates naturally embedded in ILS quotes. In this paper, we thus also propose a novel method to render inflation readings from Consensus Economics consistent with ILS quotes.

Another methodological advantage of our analysis is the application of a sophisticated Bayesian variable selection procedure for which we assess performance using a Monte Carlo study tailored towards our estimation strategy.

The paper is organised as follows. Section 2 explains the concept of the inflation risk premium and presents a model-free and a model-based strategy to estimate it. Section 3 provides a regression analysis to assess the main drivers of the inflation risk premium. This section also discusses a set of potential drivers of the inflation risk premium and elaborates on the empirical method adopted for our regression analysis. Section 4 provides some concluding remarks. Some detailed information on the construction of the data and the affine term structure model used for the model-based estimation of the inflation risk premium is presented in the appendix of the paper.

2 The inflation risk premium in ILS quotes

2.1 A conceptual discussion

Zero-coupon inflation-linked swaps (ILS) are the most commonly traded inflation derivatives in the euro area and the United States. One of the counterparts in the swap agreement will pay inflation and the other will pay an agreed fixed rate. The cash flow of the contract (for the net amount) will be paid at maturity. The agreed fixed rate is not, however, a perfect measure of inflation expectations as it includes an embedded inflation risk premium. To see this, let us think of the simple one-year-ahead zero-coupon ILS where one counterpart pays an agreed fixed rate, say π^{ILS} , and the other counterpart pays the inflation rate π . No money is exchanged upfront, and thus standard finance theory suggests that the price of the two future payoffs is the same. This means that the expected value of the product of the (stochastic) discount factor M and future inflation should be the same as the expected value of the product of M and π^{ILS} . That is:

$$E(M\pi^{ILS}) = E(M\pi) \quad (1)$$

Noting that $E(M) = (1+r)^{-1}$ and that r is the risk-free rate, simple standard algebra makes it possible to derive the following equation:²

$$\pi^{ILS} - E\pi = (1+r)\text{Cov}(M, \pi) \equiv IRP \quad (2)$$

The spread between the inflation rate quoted in the ILS and future expected inflation is called the inflation risk premium (IRP).

Note that according to equation (2) the inflation risk premium is proportional to the covariance between the stochastic discount factor and future inflation. From this perspective, the sign of the inflation risk premium is not set a priori and depends on the sign of this covariance term.

In macroeconomic theory the stochastic discount factor is equivalent to the marginal rate of substitution of consumption, or more precisely, the rate at which the investor is willing to substitute consumption today for consumption tomorrow. This marginal rate of substitution of consumption is dependent on the future marginal utility of consumption. If we assume that the marginal utility of consumption is higher when the level of consumption is low, then an asset that provides low (or even negative) returns when wealth is most needed (i.e. when the marginal utility of consumption is high or equivalently when consumption itself is low) should be held by the investor only if it offers a positive premium. Following this rationale, when the future movements of inflation are expected to be positively (negatively) correlated with the future marginal utility of consumption, there is a positive (negative) premium for holding financial assets that offer nominal payoffs. This is so because the “real” return on these assets deteriorates with inflation. To be more specific, in case market participants see the risk of “stagflation”, the inflation risk premium would be positive as investors like to hedge for high inflation that will “eat up” real returns in the economic downturn. In contrast, the expectation of a deflationary recession could lead to a negative inflation risk premium as nominal bonds perform well in case of deflation. The literature on the inflation risk premium, as surveyed by Fleckenstein et al. (2016), emphasises the possibility of a recent change in the sign of the inflation risk premium. In this regard the work by Campbell et al. (2016) is especially worth mentioning. Campbell et al. (2016) argue that nominal bonds changed from “inflation bets” to “deflation hedges” following the financial crisis.

As a direct measure of the correlation of inflation and marginal utility is not available the

²In the absence of arbitrage, for an investment offering a risk-free return of r_t , it should hold that $1 = E(M(1+r)) = (1+r)E(M)$. In the derivation of the result, use is made of the fact that for two random variables X and Y it follows that $E(XY) = E(X)E(Y) + \text{Cov}(X, Y)$.

literature uses several proxies for measuring potential changes in the correlation of inflation and the pricing kernel. Campbell et al. (2016) for example, investigate the correlation between equity prices and inflation in a semi-structural model framework. In this study we propose two alternative ways to directly measure the inflation risk premium.

Note that instead of pricing the two legs of the ILS contract using the pricing kernel as in (1), we could have chosen to price the two legs using the risk-neutral measure. The “discounted” payoffs of the two legs of the ILS contract under the risk neutral measure should be equivalent. This we could write as:

$$E^Q \left(\frac{\pi^{ILS}}{1+r} \right) = E^Q \left(\frac{\pi}{1+r} \right)$$

where E^Q denotes expectations under the risk-neutral measure. Noting that we take r , the risk-free rate as deterministic and the inflation swap rate π^{ILS} is agreed in advance and will not change after the realisation of π , it then easily follows that $\pi^{ILS} = E^Q \pi$, that is, the ILS quote is simply the expected value of future inflation under the risk-neutral measure. We can then use this result to adjust (2) and thus derive an alternative definition for the inflation risk premium, namely, that the inflation risk premium is the difference between the expected value of future inflation under the risk-neutral measure and the expected value of future inflation under the physical measure:

$$IRP \equiv E^Q(\pi) - E\pi. \quad (3)$$

This suggests two alternative approaches to estimate the inflation risk premium. First, a *model-free approach* whereby we use quotes of the inflation-linked swap, and we replace the expectation of inflation in (2) with available inflation forecasts from surveys like Consensus Economics. The inflation risk premium is thus the difference between the ILS quote and the survey inflation forecast.

Second, a *model-based approach*, whereby we formulate a model for the joint dynamics of the stochastic discount factor and the inflation process. The model will take the form of an affine term structure model for which the formulation of the risk neutral and the physical distributions are tractable, and we can proceed as in equation (3) to derive the inflation risk premium. We now turn to the detailed formulation of these two approaches.

2.2 Model-free estimate of the inflation risk premium

As a model-free indicator for the inflation risk premium we use the difference between the one-year-ahead ILS rate and the mean of the Consensus Economics one-year-ahead inflation

forecasts. We take as key dates t for our analysis those dates on which Consensus Economics data are collected. The time of collection is usually, but not always, the second Monday of the month. As for the computation of the one-year-ahead inflation from the ILS, this is computed using end-of-day quotes on the same dates as Consensus Economics data were collected and using the special procedure described in appendix A to construct the implied future price index from the ILS curve. This special procedure accounts for lags and settlement date conventions embedded in the ILS contracts, and thus ensures that we truly retrieve from the data a genuine one-year-ahead inflation forecast.

Two important considerations need to be borne in mind when adopting the method described in the appendix to construct the implied future price index from ILS quotes. First, we must ensure that we only use the information available to traders at the time of the quotes. At the time Consensus Economics data are collected every month, the t in our sample, data for the previous month's euro area overall HICP inflation is already available in the form of a "flash" release by Eurostat, which has been provided since November 2001. This release is neither the HICP excluding tobacco, nor the "final" release of the overall HICP. However, this index is revised only to a very small degree subsequently. We therefore take the view that information on the previous month's inflation is already available to traders when pricing euro area ILS contracts. For the United States, and for our sample that spans the period January 2008 to September 2016, the CPI for the previous month had never been released by the BLS at the time of the collection of the Consensus Economics data. This means, for example, that for our observation for February 2015 (which corresponds precisely to 9 February 2015), the latest CPI data released by the Bureau of Labor Statistics related to December 2014. Second, in the computation of future inflation from ILS quotes, we use the seasonal factors computed from the ratio of the seasonally adjusted and non-seasonally adjusted official price indexes the year before. This choice is motivated by the fact that in the United States the seasonally adjusted CPI index published by the Bureau of Labor Statistics is revised with the release of January data each year for the preceding five years. In the euro area the ECB publishes estimates of the seasonally adjusted HICP index excluding tobacco that are widely used by inflation traders.³ Previous data are revised with the release of the latest seasonally adjusted data. However, the seasonality pattern over the period under study has been relatively stable, and hence for the sake of simplicity and similarity when handling US data, we equally take the seasonal factors of the previous year as those employed by inflation traders, when adjusting the data.

³Eurostat, the Statistical Office of the European Union, does not publish a seasonally adjusted series of the HICP or of the HICP excluding Tobacco index.

The Consensus Economics data provides us with a forecast of an “annual” rate of change of inflation for the “current” and “next” year. These “annual” rates are not the simple “year-on-year” (e.g. December-on-December) rates naturally embedded in ILS quotes. In order to retrieve information on year-on-year inflation expectations from the information provided by available data and information from Consensus Economics, we also need to apply a special procedure. The details of this special procedure are left for appendix B.

It is also important to note that the survey data on inflation expectations from Consensus Economics for the euro area does not refer to the same price index employed as reference for the inflation-linked contract. Data from Consensus Economics relate to the overall HICP index, while the ILS contract is associated with the HICP excluding tobacco. The weight attached to the tobacco component in the overall index is, however, small (around 2.4%), and the historical average of the overall HICP is around 0.1% higher than that of the HICP index excluding tobacco. Therefore, the difference between the inflation expectations from consensus and the embedded inflation expectations from the ILS prices cannot be solely attributed to risk premia, but also to the disparity in the gap. However, we take the view that the disparity between the overall HICP index and that of the HICP excluding tobacco is small in size, and, while acknowledging that they potentially bias our results to some degree, we presume the effect to be small.⁴

2.3 Model-based estimate of the inflation risk premium

Alternatively, we construct a model-based estimate of the inflation risk premium using the approach proposed by Beechey (2008) and Finlay and Wendel (2012). We have adapted their continuous time framework to the discrete time case. The idea is to model the full inflation swap curve using an affine term structure framework. In appendix A it is explained that the inflation swap rate is fixed in a way that the cash flows of the inflation payer and the inflation receiver are equalised:

$$N \times [(1 + \pi_t^{ILS}(n))^n - 1] = N \times \left[\frac{P_{t+n}}{P_t} - 1 \right]$$

If we assume that the inflation swap rates are corrected for the indexation lag, the inflation swap rate for maturity n , defined in the equation above as $\pi_t^{ILS}(n)$, is exactly the inflation

⁴For the United Kingdom, data from Consensus Economics relate to the Retail Price Index excluding mortgage costs, usually denoted as RPI-X which thus adjusts the price index for the effects of changes in interest rates, while the ILS contract most commonly traded in the UK is linked to the overall Retail Price Index, standard RPI. The difference between RPI inflation and RPI-X inflation is, however, large and unstable, and this prevents us from replicating our analysis using UK data.

compensation paid in the market for expected changes in the price level, defined in the equation by P_{t+n}/P_t . So far we use π^{ILS} to be the discretely compounded inflation swap rate as quoted in the market. For the model analysis it is more convenient to use instead the continuously compounded inflation swap rate $y_{t,n}^\pi$, and also assume that inflation, π_t , is continuously compounded; we then get the following valuation equation for inflation swap rates:⁵

$$\begin{aligned} N \times [(1 + \pi_t^{ILS}(n))^n - 1] &= N \times E_t^Q \left[\frac{P_{t+n}}{P_t} - 1 \right] \\ N \times [e^{ny_{t,n}^\pi} - 1] &= N \times E_t^Q \left[\exp \left(\sum_{j=1}^n \pi_{t+j} \right) - 1 \right] \\ e^{ny_{t,n}^\pi} &= E_t^Q \left[\exp \left(\sum_{j=1}^n \pi_{t+j} \right) \right] \\ y_{t,n}^\pi &= \frac{1}{n} \log \left(E_t^Q \left[\exp \left(\sum_{j=1}^n \pi_{t+j} \right) \right] \right) \end{aligned}$$

Here E_t^Q is the expectation under the risk-neutral probability measure Q . The inflation swap rate for maturity n is, up to a Jensen inequality term, equal to the expected future inflation (over the period t to $t+n$). The expectation, however, has to be computed under the risk-neutral measure as markets normally require a compensation for risk. To model the expectations under the risk-neutral measure Q we further assume that inflation is driven by a vector of factors X_t as follows:

$$\pi_{t+1} = \pi_0 + \pi_1' X_t$$

where π_0 is a parameter and π_1 is a parameter vector. The vector of factors X_t is further assumed to be driven by the following dynamic under Q :

$$X_t = \mu^Q + \Phi^Q X_{t-1} + \Sigma \epsilon_t^Q$$

Then we show in the appendix that the inflation swap rates are linear functions of the underlying factors X_t :

$$y_{t,n}^\pi = A_n + B_n' X_t$$

where (see appendix) the coefficients A_n and B_n are determined by standard Riccati equations known from the affine term structure literature. The factor dynamic under the physical measure P is:

$$X_t = \mu^P + \Phi^P X_{t-1} + \Sigma \epsilon_t^P$$

Conditional on the dynamics of X_t the inflation risk premium for maturity n ($IRP_{t,n}$) can easily be computed as the difference of the expectation under the risk-neutral measure (the

⁵That is, $y_{t,n}^\pi = \ln(1 + \pi_t^{ILS}(n))$ and $\pi_t = \ln(P_t/P_{t-1})$.

inflation swap rate) and the physical measure⁶:

$$IRP_{t,n} = \frac{1}{n} \left(E_t^Q \left[\sum_{j=1}^n \pi_{t+j} \right] - E_t^P \left[\sum_{j=1}^n \pi_{t+j} \right] \right)$$

To estimate the model we apply the JSZ approach proposed by Joslin et al. (2011) and assume that the pricing factors X_t are the first three principle components of the inflation swap curve. This makes it possible to estimate the model via maximum likelihood without using the Kalman filter. As in the JSZ approach we assume that the pricing factors are observed without measurement errors. The observed individual inflation swap rates, however, can deviate from their theoretical values by a gaussian measurement error.

2.4 Estimates of the inflation risk premium

Chart 1, panel a, displays the one-year-ahead model-free indicator of the inflation risk premium in the euro area and the United States for the period January 2008 to September 2016. From 2011 onwards, the resemblance of the inflation risk premium across countries is quite striking, particularly in view of the quite different trends displayed by the Consensus inflation forecasts shown in panel b of the same Chart. Prior to 2011, when allegedly the US inflation-linked swap market was less liquid, the two ILS spreads were less closely correlated.

Chart 2 shows that both model-free and model-based measures of the inflation risk premium turned negative during the great recession. Indeed, Table 1 shows that the inflation risk premium was positive in the period before the great recession and turned strongly negative following the outbreak of the financial crisis in summer 2007. The dramatic downturn of the inflation risk premium, following the collapse of Lehman in 2008, is also striking. This extraordinary movement is explained by a period of extreme liquidity tensions and portfolio rebalancing triggered by the Lehman bankruptcy. Interestingly, this period is detected by our outlier adjustment method described later.

The overall movements of model-free and model-based measures are very similar. However, the model-free measures are considerably more volatile than the model-based measures. Model-free measures of the inflation risk premium are volatile as noisy movements of inflation swap rates at a monthly frequency are compared to relatively smooth survey-based measures of inflation expectations. In contrast, measurement errors in model-based esti-

⁶To see how this equation can be derived from the previous one please note that the variance under risk-neutral and physical measure is the same. The Jensen inequality term coming from interchanging the expectation and logarithm operators is therefore cancelled out.

mates of the inflation risk premium can absorb noise moments in the inflation swap rates to some extent.

It is also important to note that the fact that inflation projections from ILS are contaminated by risk premia does not necessarily make them worse for forecasting in practice. Chart 3 shows the rolling (over a 12-month window) root mean square error (RMSE) of various inflation projections. Interestingly, throughout the financial crisis period, the one-year-ahead projections from ILS quotes in the euro area were consistently better than Consensus Economics forecasts. In the United States the forecasting performance of Consensus Economics forecasts appears better than the ILS between 2008 and 2013. However, since 2013, the forecasting performance of both the ILS projections and those of Consensus are roughly equivalent. In terms of the actual quality of the forecast, Chart 3 appears to suggest that forecasting inflation for a one-year horizon is a difficult task. In both the euro area and the United States, a constant forecast was best for most of the sample, although over recent years it has provided the worse forecast. In any event, for the full sample the difference between the best performing model and the worse is marginal in absolute values.

Table 1: Inflation risk premium (in basis points)

Period	US			Euro area		
	survey-based	model-based	average	survey-based	model-based	average
2005-2007	43	9	26	19	3	11
2008-2016	-50	-30	-40	-23	-17	-20

3 Drivers of the inflation risk premium

For our empirical analysis the number of possible explanatory variables k to explain the inflation risk premium is not excessively large ($k = 10$). Given the sample size of around 100 observations, the use of ordinary least square methods might a priori seem feasible. However, strong collinearity among some of the regressors suggests that the variance of ordinary least square estimates might be abnormally high, leading to small t-statistics that may prompt us to treat many of the potential regressors as irrelevant. We argue that the use of Bayesian selection methods is preferable under these circumstances, and we justify our strategy with a Monte Carlo exercise that mimics the nature of our empirical study.

3.1 Potential drivers

3.1.1 Inflation uncertainty

We use four indicators of inflation uncertainty. Our first proxy for inflation uncertainty is the one also employed in Soderlind (2011), namely the implied volatility of long-term sovereign bond options, which we refer to as *bond volatility*. Strictly speaking this is not truly a direct measure of inflation uncertainty, but it is sensible to expect inflation uncertainty to be priced in nominal long-term bond yields, and so, it should also be reflected in the price of bond options.

Our second proxy is the monthly standard deviation of the daily changes in the one-year-ahead ILS computed with a one-month window, which we refer to as *realised volatility*. Additionally, and in view of the fact that inflation options have become more liquid, and more information is now available for empirical research, we compute, as our third proxy for inflation uncertainty, an estimate of the *implied volatility* measured from inflation options. To compute this measure we estimate a volatility smile from seven zero-coupon inflation floors with strike prices ranging from -2% to 1% and one year of maturity. We then take as our measure of implied volatility the fitted volatility for a strike price which is at the money, i.e. a strike price equal to the one-year-ahead inflation-linked swap quote. Implied volatilities are computed by means of the Black formula commonly employed for the valuation of zero-coupon inflation caps and floors.

Further to these measures we also include the difference between the mean and the median of the Consensus responses as a potential indicator of uncertainty; we refer to this indicator as the *Mean-Median Consensus*.⁷

3.1.2 Market liquidity

We will employ the difference between the ILS rate and the BEIR rate at the five-year maturity as a measure of market liquidity. We will denote the measure of market liquidity as *liquidity* in the tables below. In a non-arbitrage environment without frictions, the difference between the ILS rate and the BEIR rate should be zero. In a world with frictions, the spread should be related to the sum of the liquidity premium in the inflation-linked bond market and the ILS market, see e.g. Christensen and Gillan (2015). It is important to note that our

⁷Some studies favour the use of the median as a better measure of inflation expectations from survey responses than the mean, see e.g. Mankiw et al. (2003). To the extent that this is a reasonable assumption, this indicator may capture the possible bias of the mean forecast as a measure of inflation expectations.

model-free and model-based indicators for the inflation risk premium and our market liquidity proxy are both constructed using ILS data. This might imply a spurious correlation. We have chosen to compute our market liquidity measure for the five-year maturity rather than for the one-year maturity employed for our dependent variable, to partly address this issue. First, the changes in the five-year rate will not be fully driven by changes in expectations to the one-year rate. The market liquidity proxy computed for the five-year maturity may still be a relatively good proxy for liquidity conditions in the market. The BEIR for the euro area is computed as the average of the BEIR rates embedded in the French and German inflation-linked bonds. Details of how to estimate the BEIR are given in Ejlsing et al. (2007). For the United States we employ the BEIR series from the Gurkaynak et al. (2008) database.

3.1.3 Inflation disagreement

We have also constructed a measure of inflation disagreement from the individual responses of the Consensus Economics inflation forecasts based on the standard deviation of Consensus Economics expected inflation one year ahead, denoted below by *Consensus disagreement*. Our measure of Consensus disagreement is similar to the measure of inflation disagreement employed in published research, e.g. Soderlind (2011).

Bloomberg reports expectations of the outcome of the inflation data release a few days before the data are publicly disclosed by the statistical offices. The difference between the inflation data released by the statistical office and the expectations collected a few days before provides a measure of the surprise of the data release. We have collected data from Bloomberg surveys on the “normalised” surprises caused by the data release, and we denote this data series as *data surprise* in our tables below. Further to this we have also collected the reported standard deviation of the surprise of the inflation data release among respondents to the Bloomberg survey. We will refer to this data series as *data surprise uncertainty* when commenting on our results below. In interpreting the impact from the data releases in our analysis it is important to understand their timing, and how they relate to the dates on which Consensus data are collected (which serve as reference dates for our empirical analysis). In the euro area, the date of the inflation data surprise relates to the release by Eurostat of the “overall” HICP index at the end of the previous month, and it is thus information that by the time the Consensus Economics surveys are conducted is already a little dated, usually 8 to 10 days old. The surprise is a bit more dated for the United States, where the release of the CPI by the Bureau of Labor Statistics lags behind our reference date by two to three weeks. Data release surprises may nonetheless have an impact on the inflation risk premium.

3.1.4 Inflation and deflation fears

The addition of proxies for inflation and deflation fears to our set of indicators is especially important for the analysis of the inflation risk premium during the financial crisis period. In particular, the inflation risk premium may be as much a reflection of a deflation or inflationary bias settling in the mind of investors, as a reflection of levels of inflation volatility per se. Declining inflation rates are to the detriment of investors holding long inflation positions, and thus pessimism on short-term inflation dynamics should potentially trigger a decline in the inflation risk premium. We thus add two proxies for inflation and deflation fears. These measures will be defined with reference to “psychological” thresholds. In particular, the inflation fear may be associated with the price of the one-year-ahead inflation cap with a strike price equal to 4% which is comfortably above the inflation target of the central bank. In the same vein, the deflation fear may be defined by the price of the one-year-ahead inflation floor with a strike price of 0%, that is the threshold for true deflation. These will be referred to as *Cap at 4* and *Floor at 0* respectively.

3.2 A simulation study of our variable selection strategy

A simple and common method for model selection is to estimate the full model and discard variables which are not statistically significant. As stated above, this may be problematic when there is multicollinearity among the regressors. Thus we set ourselves to study whether selecting variables on the basis of whether the t-statistics of the ordinary least square regressions are large, leads indeed to worse results than selecting a model by means of Bayesian variable selection techniques in the presence of multicollinearity. For this purpose we will adjust and expand the Monte Carlo study presented in Ghosh and Ghattas (2015, Sec 3.1). In their study, Ghosh and Ghattas (2015) did not assess the potential use of standard t-statistics for model selection (something we would like to address here), but rather focused on the impact of the adoption of various prior distributions for Bayesian variable selection.

Variable selection methods. The variable selection methods that will be put to the test are the following. First, ordinary least squares, OLS^{**} and OLS^* , will refer to the selection of variables as pertaining to the true model whenever the t-statistics reject the hypothesis that the parameter associated with that variable is zero with a level of significance of 5% and 10% respectively. Second, the highest probability model (HPM) will be the selection

of that model which is assigned the largest posterior probability among all possible models. Third, the median probability model (MPM) of Barbieri and Berger (2004), i.e. the model which includes all variables whose marginal posterior probability of inclusion is larger than 50%. Fourth, selection using the Bayesian model averaging estimates, BMA^{**} and BMA^* , the same as for OLS, but now using the BMA estimated coefficients.

We will adopt a fairly standard Bayesian variable selection framework. In particular, we take the prior probability for a certain variable to be present in the true model as being independent Bernoulli with probability 0.5. This is equivalent to assuming that any single possible combination of variables is equally likely as a true model, that is any combination of variables has prior probability 2^{-10} of being the true model. We further adopt a standard conjugate non-informative prior structure for the parameters of a given model. We adopt an inverse gamma prior for σ^2 with parameters $\nu = 3$ and $\lambda = \hat{\sigma}^2$, where $\hat{\sigma}^2$ is the ordinary least square estimate of the error term computed from the full model, i.e. the model containing all the explanatory variables.

Bayesian selection methods in the presence of multicollinearity are particularly sensitive to the choice of the prior distribution of the parameters, as shown in Ghosh and Ghattas (2015). As in Ghosh and Ghattas (2015) we also put to the test the two most commonly encountered priors employed in the literature. The first are independently normal priors with mean zero and variance equal to $2.85^2 \times \sigma^2$. This prior was advocated by Raftery et al. (1997) and, as we will see below, leads to better results than alternative equally standard specifications. The second is the standard g-prior specification of Zellner (1986), where the parameter g is fixed using the strategy recommended by Fernandez et al. (2001). When commenting on our results below, we will use the letter i when referring to the adoption of independent priors, e.g. MPMi, and the letter g when referring to the adoption of the g-prior, e.g. MPMg. Both these standard conjugate priors have the convenient property that the marginal posterior probabilities of inclusion of a variable are analytically tractable. In our analysis the number of possible variables to include is not excessively large and we can thus compute analytically the posterior probabilities of inclusion for all possible combinations of the explanatory variables.⁸

⁸See Chipman et al. (2001, Sec 3) for details on the implementation of this type of priors, and full details on the exhaustive calculation of the posterior model probabilities. Note also that for the computation of the estimated variance of the model averaged parameters reported in the tables, say $\hat{\beta}_M$, we proceed as in Leamer (1978, p. 118), namely:

$$\hat{Var}(\beta_M) = \sum_i \hat{Var}(\beta_i) P(i) + \sum_i \left(\hat{\beta}_i - \hat{\beta}_M \right)^2 P(i)$$

Monte Carlo design. We take the number of candidate variable, k , as equal to 10. The last six variables are independent standard normally distributed random variables. The first candidate variables are generated from $x_{it} = z_t + u_{it} * (1 - \rho) \rho^{-1}$ for $i = 1$ to 4, and where z_t and u_{it} are standard independent normally distributed variables. This generates regressors which are mutually correlated, and with correlations coefficient equal to ρ . We then conduct two different Monte Carlo studies. In *Monte Carlos study A* we assume that the first four variables are assigned a coefficient equal to one, and the remaining six a coefficient equal to zero. In contrast, in our *Monte Carlos study B* the opposite is the case, i.e. the first four variables are assigned a coefficient of zero and the last six a coefficient of one. We will further conduct the simulation studies for values of ρ either fixed to 0.9 or to 0.5. These values are aligned with the correlation observed among the indicators listed in section 3.1. They are, of course, well below the value used in Ghosh and Ghattas (2015) which was 0.997, a value we judge is not commonly encountered in applied econometric studies that use monthly time series. We will employ two samples sizes, $N = 90$ and $N = 200$, the former equally aligned with our empirical analysis. Finally, the standard deviation of the error term σ_ε is set to either 4 or 6; these lead to R-squared coefficients of around 0.85 and 0.45 respectively.

Performance of the different methods is evaluated over 10,000 simulations. Tables 2 and 3 present various metrics to evaluate performance. Table 2 shows i) the average number of *false inclusions* (FI), that is the average number of variates that do not belong to the true model that are selected as valid regressors; ii) the average number of *false exclusions* (FE), that is the average number of variates that belong to the true model that fail to be selected; and iii) the percentage of times that the null model is selected (Null), that is, no single variate is selected. Table 3 shows instead the average *bias* and average root mean square error (RMSE) of the estimated parameters. The bias is presented for the estimated coefficient of the first and fifth parameter; recall that in the Monte Carlo study A, the first parameter is 1 and the fifth is zero, and vice-versa for the Monte Carlo Study B. The RMSE refers to the root mean square error of the 10 estimated parameters.⁹

Several lessons can be taken from these exercises. First, Bayesian selection techniques such as MPM and HPM perform better than OLS in identifying the true model. Second, this is particularly so when multicollinearity is large. Somehow worryingly, the OLS methods lead

where $\hat{Var}(\beta_i)$ is the estimated variance of the parameter β in model i , and $P(i)$ is the marginal probability of model i .

⁹Of course, when a variable is not selected the assigned value is zero.

to a large percentage of cases where no single variable is actually selected, e.g. sometimes the null model is selected 65% of the time. Third, Bayesian techniques such as BMA do not improve upon the OLS method for identifying the true model. Fourth, MPM and HPM with independent priors have a slight advantage compared to when the method is implemented with g-priors. This is indeed aligned with the result reported in Ghosh and Ghattas (2015). However, our results also reveal that for levels of multicollinearity more aligned with econometric applications, the disadvantage of using g-priors is not very large. It is also important to note that the OLS method does not even appear to have a major advantage over the MPM and HPM when the level of multicollinearity is low.

Moving to the results of Table 3, it transpires that, as expected, the reliability of the estimates in terms of bias and RMSE deteriorates in the presence of multicollinearity and when the standard deviation of the error term deteriorates. The RMSE of the posterior parameter estimates of the MPM and HPM cannot be said to be better or worse than the RMSE of the OLS method. However, the BMA methods appear clearly more reliable in terms of RMSE than both the OLS or the MPM and HPM.

All in all, this simulation study suggests that for our empirical analysis Bayesian methods are more suitable than the adoption of simple OLS techniques. For our empirical analysis we will choose to report the estimation results for the MPMi and BMAi methods which provide the best method for model selection and the best method to retrieve estimators with lower RMSE.

3.3 Regression results

We conduct our regression analysis using series in levels and in first-differences for robustness, and focus on the sample period January 2008 to September 2016. The series do appear to be stationary. However, serial correlation is non-negligible and may jeopardise the implementation of the Bayesian selection procedure that relies on normality assumptions. On the other hand, regressions using first differences might be subject to the standard problems of over-differencing. We thus also choose to add the lagged value of the inflation risk premium as a potential explanatory variable in our Bayesian selection strategy.¹⁰ In our regression analysis the explanatory variables are all normalised, that is their mean is adjusted to zero and the standard deviation to one. This renders the magnitude of the coefficients across

¹⁰The only variable that does not appear sensible to transform in first differences is the *data surprise*. It is thus left in levels in the reported results below.

variables comparable. Finally, estimation results removing outliers are also reported in the tables. Outliers are identified as those observations which deviate from their sample mean by more than three times their sample standard deviation. All candidate explanatory variables are shown in figures 4 and 5.

The results for the Bayesian model averaging estimation, where candidate models are weighted by their posterior probability of inclusion, are shown in Tables 4 and 5. The majority of the explanatory variables turn out not to be statistically significant. The *liquidity* indicator turns out to have a significant effect on the model-free indicator for the inflation risk premium for the US and the euro area (in the regression in levels) but is not significant when outliers are removed. For the model-based indicator, liquidity is significant in the level regressions for the US but again no longer significant when outliers are removed. The fact that the outlier removal changes the significance of the liquidity indicators consistently, suggests that liquidity was particularly a problem in the immediate months following the collapse of Lehman, as this is the period identified as an outlier. For the vast majority of specifications, indicators of realised inflation volatility have a negative, although mainly not significant, coefficient. In contrast, for a few US specifications the implied inflation volatility has a significant positive coefficient.

The regression results for our indicators of inflation and deflation fears suggest that the inflation risk premium is mainly a measure of deflation fears. The price of the inflation floor at zero (price for deflation protection) consistently shows a negative coefficient that is statistically significant for most specifications. The price of protecting against inflation above the target shows consistently a positive coefficient that is statistically significant for some specifications. The opposite sign of the coefficient associated with the indicators for inflation and deflation fears strongly indicates that the inflation risk premium is dominantly influenced by *the balance of inflation risk* rather than inflation volatility. If the risk of deflation increases, or the risk of inflation being higher than the target declines, the inflation risk premium declines and can even become negative. This interpretation is also supported by the regression results for the *inflation surprise* indicator. This indicator has a positive and statistically significant coefficient for the model-based inflation risk premium for the US. At the same time, the indicator of *surprise uncertainty* is not statistically significant. This again suggests that it is the *direction* of data surprises, rather than the uncertainty of inflation data releases, that is the main driver for the inflation risk premium. The sequence of negative surprises in the inflation releases we have seen since the great recession have brought down the inflation risk premium.

4 Conclusions

Our analysis can be condensed into two main results. First, the inflation risk premium that was, both for model-free and model-based measures, positive in the years before the collapse of Lehman has turned considerably negative since then. This is striking as most other premia, like credit risk premia, increased strongly amid the financial crisis. Second, indicators of inflation uncertainty, like disagreement on inflation forecasts and realised or historical inflation volatility, seem not to be the main driver of the inflation risk premium. What actually drives the inflation risk premium are indicators of inflation above the monetary policy target and indicators for deflation fears. The opposite sign of the indicators of inflation and deflation fears strongly indicates that the inflation risk premium is dominantly influenced by *the balance of inflation risk* rather than inflation volatility. This implies that a low, or even negative, inflation risk premium is primarily associated with strong deflation risk rather than low inflation uncertainty. Our results have considerable implication for monetary policy as central banks should not be complacent with low or even negative inflation risk premia.

References

- BARBIERI, M. AND J. BERGER (2004): “Optimal predictive model selection,” *Annals of Statistics*, 32, pp. 870–897.
- BEECHEY, M. J. (2008): “Lowering the Anchor: How the Bank of England’s Inflation-Targeting Policy have Shaped Inflation Expectations and Perception of Inflation Risk,” Working Paper, Federal Reserve Board.
- CAMPBELL, J., A. SUNDERAM AND L. VICEIRA (2016): “Inflation Bets or Deflation Hedges? The Changing Risks of Nominal Bonds,” *Critical Finance Review*.
- CHIPMAN, H., E. I. GEORGE AND R. E. MCCULLOCH (2001): *The Practical Implementation of Bayesian Model Selection*, Beachwood, OH: Institute of Mathematical Statistics, vol. 38 of *Lecture Notes–Monograph Series*, 65–116.
- CHRISTENSEN, J. H. E. AND J. M. GILLAN (2015): “Does Quantitative Easing Affect Market Liquidity?” Federal Reserve Bank of San Francisco Working Paper 2013-26.
- EJSING, J., J. GARCIA AND T. WERNER (2007): “The term structure of euro area break-even inflation rates: The impact of seasonality,” ECB Working Paper Series, No 830.
- FERNANDEZ, C., E. LEY AND M. STEEL (2001): “Benchmark priors for Bayesian model averaging,” *Journal of Econometrics*, 100, pp. 381–427.
- FINLAY, R. AND S. WENDEL (2012): “Estimating Inflation Expectations with a Limited Number of Inflation-Indexed Bonds,” *International Journal of Central Banking*, 8, pp. 111–142.
- FLECKENSTEIN, M., F. LONGSTAFF AND H. LUSTIG (2016): “Inflation-Adjusted Bonds and the Inflation Risk Premium,” in *Handbook of Fixed-Income Securities*, ed. by P. Veronesi, Wiley, 41–52.
- GHOSH, J. AND A. E. GHATTAS (2015): “Bayesian variable selection under collinearity,” *The American Statistician*, 69(3), pp. 165–173.
- GURKAYNAK, R. S., B. SACK AND J. H. WRIGHT (2008): “The TIPS Yield Curve and Inflation Compensation,” Federal Reserve Board, Finance and Economics Discussion Series 2008/05.
- JOSLIN, S., K. SINGLETON AND H. ZHU (2011): “A New Perspective on Gaussian Dynamic Term Structure Models,” *Review of Financial Studies*, 24(3), pp. 927–970.

- KNUPPEL, M. AND A. L. VLADU (2016): “Approximating fixed-horizon forecasts using fixed-event forecasts,” Deutsche Bundesbank Discussion Paper, No 28/2016.
- LEAMER, E. E. (1978): *Specification Searches*, New York: John Wiley & Sons.
- LIU, Z., E. VANGELISTA, I. KAMINSKA AND J. RELLEEN (2015): “The informational content of market-based measures of inflation expectations derives from government bonds and inflation swaps in the United Kingdom,” Staff Working Paper No. 551, Bank of England.
- MANKIW, N. G., R. REIS AND J. WOLFERS (2003): “Disagreement about inflation expectations,” *NBER Macroeconomics Annual*, 18, pp. 209–248.
- PATTON, A. J. AND A. TIMMERMANN (2011): “Predictability of output growth and inflation: a multiple-horizon survey approach,” *Journal of Business and Economic Statistics*, 29(3), pp. 397–410.
- RAFTERY, A. E., D. M. MADIGAN AND J. A. HOETING (1997): “Bayesian model averaging for linear regression models,” *Journal of the American Statistical Association*, 92, pp. 179–191.
- SODERLIND, P. (2011): “Inflation Risk Premia and survey evidence on macroeconomic uncertainty,” *International Journal of Central Banking*, 7(2), pp. 113–133.
- ZELLNER, A. (1986): “On assessing prior distributions and bayesian regression analysis with g-prior distributions,” in *Bayesian Inference and Decision Techniques: Essays in Honor of Bruno Finetti*, ed. by P. Goel and A. Zellner, North Holland/Elsevier, 233–243.

Technical Appendix

A Projected price index embedded in the ILS curve

One of the counterparts in the Zero-coupon inflation-linked swap agreement will pay inflation and the other will pay an agreed fixed rate. The cash flow of the contract (for the net amount) will be paid at maturity. If counterpart **A** pays the fixed rate, say $\pi_t^{ILS}(k)$, and counterpart **B** pays inflation, their respective cash-flows in a k -year maturity swap for a notional of N will be:

$$\begin{aligned}\text{Cash Flow A} &= N \times \left[\left(1 + \pi_t^{ILS}(k) \right)^k - 1 \right] \\ \text{Cash Flow B} &= N \times \left[\frac{I_{t+k}}{I_t} - 1 \right]\end{aligned}$$

where t is the date when the contract is settled and I_t the price index used as reference in the swap. Of course, the index k serves to denote those years for which there are quotes in the ILS market. In particular, zero-coupon swaps are actively traded for maturities of 1 to 10 years.

The zero coupon swap provides a picture of expected (under the risk neutral measure) path for the “reference” price index. If t is the settlement date used to compute the reference date, the price of an ILS contract of k -years of maturity provides us with the expected value of the price index in the future. Noting that the discounted value of the cash flow A and B should be equal at the time of settlement, it follows that:

$$I_{t+k} = I_t \times \left(1 + \pi_t^{ILS}(k) \right)^k \quad (\text{A-1})$$

This clearly illustrates how expectations on future inflation are embedded in ILS quotes. However, the computation of future inflation from ILS quotes is not as simple as suggested by equation (A-1). We now turn to this issue.

The ILS quotes can be used to retrieve expectations about the future path of inflation as recorded by the “official” price index. In this paper we will adopt the notation P_t to denote the “official” price index underlying the construction of the “reference” price index, I_t , actually employed in the payments of the ILS contract. With a slight abuse of notation, where before we generally adopted the index t to denote a date, we now also adopt the convention of splitting that index into the subindex y, m, d to denote the year, month and day of a given date t .¹¹ The reference price I_t adopted in ILS contracts is computed from

¹¹We further adopt the convention that subtracting one month from a date in January will also result in a subtraction of the year, even when this is not implicitly defined. That is $y, m - 1, d$ for the date 2015/*Jan*/25 would result into 2014/*Dec*/25.

lag values of P_t following market conventions. The rationale for adopting an “indexation” lag results from the fact that consumer price indexes are usually reported with a lag.

A.1 Euro area

In the most commonly zero coupon inflation-linked swap traded in the euro area, the price index is usually the non-seasonally adjusted Harmonised Index of Consumer Prices (HICP) excluding tobacco, and the reference date is the HICP three months before the time of settlement (three-month lag indexation).¹² The time of settlement is currently two working days. In sum, the index is defined as follows:

$$I_t = I_{y,m,d} = P_{y,m-3} \quad (\text{A-2})$$

When using the lag convention adopted in the euro area and the formulation of the reference price index, shown in equation (A-2), it follows that the price index projected in the swap contract in equation (A-1) results in:

$$P_{y+k,m-lag} = P_{y,m-lag} \times (1 + \pi_t^{ILS}(k))^k \quad (\text{A-3})$$

It would be unwise to compute the embedded price index in between say, $P_{y+k,m-lag}$ and $P_{y+k+1,m-lag}$, by linear interpolation. The presence of seasonality in price indexes implies that inter-month growth in the index over certain months is larger than on average. It is thus sensible to adjust the monthly interpolation to account for this regular seasonal pattern, by means of using seasonal adjustment factors, S_t .¹³ For computing the value of the index at a date τ , where $t_0 < \tau < t_1$ and values for the price index at P_{t_0} and P_{t_1} are available (either from the projections in equation (A-3) above, or because the index at date t_0 may have already been published by the Statistical Office), we would proceed as follows:

$$P_\tau = P_{t_0} \frac{S_\tau}{S_{t_0}} \left(\frac{P_{t_1} S_{t_0}}{P_{t_0} S_{t_1}} \right)^{\frac{\#(\mathcal{I})}{\#(\mathcal{K})}} \quad (\text{A-4})$$

where \mathcal{I} is the set of all months from t_0 to τ , \mathcal{K} the set of months from t_0 to t_1 (without counting t_0 on both accounts), and $\#(\mathcal{I})$ and $\#(\mathcal{K})$ denote respectively the number of months in \mathcal{I} and \mathcal{K} .¹⁴

¹²The HICP is the official measure of consumer price inflation in the euro area. The HICP is calculated according to a harmonised approach and a single set of definitions, thus providing a comparable measure of inflation across euro area countries. It is published by Eurostat, the Statistical Office of the European Commission.

¹³That is, if P_i^{SA} is the seasonally adjusted index on month i and P_i the non-seasonally adjusted index, the seasonal factor is defined by $S_i = P_i / P_i^{SA}$.

¹⁴The computed value for P_τ is such that P_τ^{SA} , $P_{t_0}^{SA}$ and $P_{t_1}^{SA}$ lie on the same straight line.

A.2 United States

The “official” price index underlying the ILS contract in the United States is the Consumer Price Index for All Urban Consumers published by the Bureau of Labor Statistics. The indexation lag is three months, and contracts are settled in two days. Furthermore, and in contrast with the euro area, in order to compute a daily reference price index from the monthly official data published by the Statistical Office, an interpolation procedure is adopted.¹⁵ The reference price index in the US ILS contract is thus computed from the official monthly data $P_{y,m}$ as:

$$I_t = I_{y,m,d} = P_{y,m-3} + \frac{d-1}{\# \text{ of days in } m} (P_{y,m-2} - P_{y,m-3}) \quad (\text{A-5})$$

(and where it is further assumed that $m-3$ gives the official index published by the Statistical office for the month lagging behind the current month by three months). The index I_t is commonly rounded to five decimal digits.

The derivation of the future path of the US CPI from ILS contracts is slightly more cumbersome. When using the lag convention adopted in the United States, the formulation of the reference price index, shown in equation (A-5), it follows that the price index projected in the swap contract in equation (A-1) results in:

$$(1 - \delta_{d,m}) P_{y+k,m-lag} + \delta_{d,m} P_{y+k,m-lag+1} = [(1 - \delta_{d,m}) P_{y,m-lag} + \delta_{d,m} P_{y,m-lag+1}] \times (1 + ilst(k))^k \quad (\text{A-6})$$

and where:

$$\delta_{d,m} = \frac{d-1}{\# \text{ of days in } m}$$

The right-hand side of (A-6), and including $\delta_{d,m}$, are known, but this does not allow us to obtain a projection for the future price index k -years ahead, because this is split into a weighted average of $P_{y+k,m-lag}$, and $P_{y+k,m-lag+1}$. Of course, we could assume that the projected future values of the index for both $P_{y+k,m-lag}$ and $P_{y+k,m-lag+1}$, are both aligned with the growth rate implied by the ILS rate. That is:

$$P_{y+k,m-lag} = P_{y,m-lag} \times (1 + h_t(k))^k$$

$$P_{y+k,m-lag+1} = P_{y,m-lag+1} \times (1 + h_t(k))^k$$

For computing the value of the index at a date τ , where $t_0 < \tau < t_1$ and values for the price index at P_{t_0} and P_{t_1} are available, we would simply employ the interpolating formula (A-4).

¹⁵There are, however, other zero-coupon inflation-linked swaps traded in the euro area that adopt interpolation of the price index as a convention, e.g. the French CPI excluding tobacco. These ILS contracts will not be used in the context of this paper

B Projected price index embedded in Consensus data

In order to retrieve information on year-on-year inflation expectations from the information provided by available data and information from Consensus Economics, we need to proceed as follows. First, we make use of the approximating formula of Patton and Timmermann (2011). This formula provides a representation of the yearly rates provided by Consensus Economics as a weighted sum of the monthly rates of growth of the price index. Namely,

$$z_t = \frac{\bar{P}_t}{\bar{P}_{t-12}} - 1 \approx \sum_{k=0}^{23} w_k y_{t-k} \quad (\text{B-7})$$

and where

$$\begin{aligned} \bar{P}_t &= \sum_{k=0}^{11} P_{t-k} \\ y_{t-k} &= \log \left(\frac{P_{t-k}}{P_{t-k-1}} \right) \\ w_k &= 1 - |k - 11| / 12 \end{aligned}$$

The data from Consensus correspond to z_t for the “current” and for the “next” year, while we would like to infer the year-on-year rates defined by $P_t/P_{t-12} - 1$. Second, when reporting the Consensus “current” estimate of annual inflation, the price indexes for a number of months, say k^* , in the “current” year are not available. By further using the standard identity linking seasonally adjusted and non-seasonally adjusted data, i.e. $P_t = P_t^{SA} S_t$, we could write the approximating formula in (B-7) as:

$$z_t^{current} \approx \sum_{k=k^*}^{23} w_k \log \left(\frac{P_{t-k}}{P_{t-k-1}} \right) + \sum_{k=0}^{k^*-1} w_k \log \left(\frac{P_{t-k}^{SA} S_{t-k}}{P_{t-k-1}^{SA} S_{t-k-1}} \right)$$

Of course, the price indexes for the last k^* are not available. Third, we assume that the monthly seasonally adjusted growth rate of the price index for the last k^* observations remains constant at a rate g .¹⁶ This allows to re-write $z_t^{current}$ as:

$$z_t^{current} \approx \mathcal{Y}_1 + g\mathcal{W} + \mathcal{Y}_2$$

with:

$$\begin{aligned} \mathcal{Y}_1 &= \sum_{k=k^*}^{23} w_k \log \left(\frac{P_{t-k}}{P_{t-k-1}} \right) \\ \mathcal{Y}_2 &= \sum_{k=0}^{k^*-1} w_k \log \left(\frac{S_{t-k}}{S_{t-k-1}} \right) \\ \mathcal{W} &= \sum_{k=0}^{k^*-1} w_k \end{aligned}$$

¹⁶While this appears at first a strong assumption, the seasonally adjusted monthly rates of inflation in both the euro area and the United States do not display strong serial correlation.

From where it follows that:

$$g \approx (z_t^{current} - \mathcal{Y}_1 - \mathcal{Y}_2) / \mathcal{W}$$

Finally, if τ is the last available price index in the “current” year, the extrapolated price index for next month can now be simply constructed as:

$$P_{\tau+1} = P_{\tau} \frac{S_{\tau+1}}{S_{\tau}} \exp \{g\}$$

This would provide the extrapolated price index for the remaining months of the year. This procedure could also be easily implemented for the extrapolation of the price index over the “next” year by simply using the available price indexes as well as the extrapolated price indexes of the “current” year. Year-on-year rates can then be constructed using the extrapolated path for the price index over the “current” and “next” year.¹⁷

C Affine term structure of inflation swap rates

In this appendix we derive Riccati equations for computing the coefficients A_n and B_n to price inflation swap rates:

$$y_{t,n}^{\pi} = A_n + B_n' X_t$$

We start with the valuation equation for inflation swap rates derived in the main text

$$e^{ny_{t,n}^{\pi}} = E_t^Q \left[\exp \left(\sum_{j=1}^n \pi_{t+j} \right) \right]$$

In order to derive the same Riccati equations known from nominal term structure models we invert the equation above¹⁸

$$e^{-ny_{t,n}^{\pi}} = E_t^Q \left[\exp \left(- \sum_{j=1}^n \pi_{t+j} \right) \right]$$

For an $n - 1$ maturity inflation swap rate in period $t + 1$ we get

$$e^{-(n-1)y_{t+1,n-1}^{\pi}} = E_{t+1}^Q \left[\exp \left(- \sum_{j=1}^{n-1} \pi_{t+1+j} \right) \right]$$

¹⁷Knuppel and Vladu (2016) provided an alternative method to compute year-on-year forecasts from Consensus data. Both their method and the method we propose in this paper rely on the approximating formula of Patton and Timmermann (2011) and implicitly assume constant monthly growth rates for the unobserved price indexes of the current year. However, and in contrast with the approach pursued in Knuppel and Vladu (2016), in our method i) we explicitly model seasonality, (and indeed the assumption of constant growth over the remaining months of the current year is only imposed on the seasonally adjusted growth rate) and ii) the projected path of the price index is fully consistent with the reported “current” and “next” inflation projections of Consensus. More importantly, our method is fully aligned with the method employed in the previous section to extrapolate the future price index from ILS quotes.

¹⁸We are using properties of the log-normal distribution. When Z is a normal distributed random variable with mean μ and standard deviation σ then $X = \exp(Z)$ is log-normal distributed with $E(X) = \exp(\mu + \frac{1}{2}\sigma^2)$. Therefore $E(X)^{-1} = \exp(-(\mu + \frac{1}{2}\sigma^2))$.

We can rearrange the equation for $e^{-ny_{t,n}^\pi}$ by using the law of iterated expectations,

$$\begin{aligned} e^{-ny_{t,n}^\pi} &= E_t^Q \left[\exp \left(- \sum_{j=1}^n \pi_{t+j} \right) \right] \\ &= E_t^Q \left[\exp(-\pi_{t+1}) \exp \left(- \sum_{j=2}^n \pi_{t+j} \right) \right] \\ &= E_t^Q \left[E_{t+1}^Q \left[\exp(-\pi_{t+1}) \exp \left(- \sum_{j=1}^{n-1} \pi_{t+1+j} \right) \right] \right] \end{aligned}$$

If we substitute the equation for $e^{-(n-1)y_{t+1,n-1}^\pi}$ and take logs on both sides we get a recursive equation for inflation swap rates; and if we further assume that inflation in period $t+1$ is known already in period t ,¹⁹ then it follows that:

$$-ny_{t,n}^\pi = \log \left(\exp(-\pi_{t+1}) E_t^Q \left[\exp(-(n-1)y_{t+1,n-1}^\pi) \right] \right)$$

Inflation swap rates are assumed to be linear affine functions for the state variables X_t

$$\begin{aligned} y_{t,n}^\pi &= -\frac{a_n}{n} - \frac{b'_n}{n} X_t = A_n + B'_n X_t \\ -ny_{t,n}^\pi &= a_n + b'_n X_t \end{aligned}$$

Substituting this definition into the recursive equation for swap rates we get

$$\begin{aligned} a_n + b'_n X_t &= -\pi_{t+1} + \log \left(E_t^Q \left[\exp(a_{n-1} + b'_{n-1} X_{t+1}) \right] \right) \\ &= -\pi_{t+1} + a_{n-1} + b'_{n-1} E_t^Q(X_{t+1}) + \frac{1}{2} b'_{n-1} V_t^Q(X_{t+1}) b_{n-1} \end{aligned}$$

Using the dynamics of the state variables under the Q measure and the fact that inflation is a linear affine function of the state variables

$$\begin{aligned} X_t &= \mu^Q + \Phi^Q X_{t-1} + \Sigma \epsilon_t^Q \\ \pi_{t+1} &= \pi_0 + \pi'_1 X_t \end{aligned}$$

we get

$$a_n + b'_n X_t = -\pi_0 - \pi'_1 X_t + a_{n-1} + b'_{n-1} \mu^Q + b'_{n-1} \Phi^Q X_t + \frac{1}{2} (b'_{n-1} \Sigma' \Sigma b_{n-1})$$

¹⁹In a model with monthly frequency this assumption just implies that inflation for the current month is already known at the beginning of the month. In the continuous time framework of Beechey (2008) and Finlay and Wendel (2012) this assumption is not needed to derive Riccati equations (Riccati ODEs) that are identical to the case of nominal bond yields. In our discrete time framework the Riccati equations derived below are approximately valid also without this assumption as shown by Liu et al. (2015).

and by equating constant terms and terms including X_t the standard Riccati equations are derived:

$$\begin{aligned} b_n &= -\pi_1 + (\Phi^Q)'b_{n-1} \\ a_n &= -\pi_0 + a_{n-1} + (\mu^Q)'b_{n-1} + \frac{1}{2}(b'_{n-1}\Sigma'\Sigma b_{n-1}) \end{aligned}$$

with $b_0 = 0$ and $a_0 = 0$.

Table 2: Monte Carlo Simulation for variable selection methods. Inclusion probabilities.

		$\sigma_\varepsilon = 4$						$\sigma_\varepsilon = 6$					
		$\rho = 0.9$			$\rho = 0.5$			$\rho = 0.9$			$\rho = 0.5$		
N	mod	FI	FE	Null	FI	FE	Null	FI	FE	Null	FI	FE	Null
Monte Carlo Study A													
90	OLS**	0.305	3.452	50.7%	0.306	1.217	0.0%	0.307	3.644	67.7%	0.302	2.472	6.9%
	OLS*	0.606	3.111	27.9%	0.602	0.814	0.0%	0.595	3.382	48.3%	0.600	1.966	1.1%
	BMAg**	0.020	3.562	56.2%	0.027	2.111	1.8%	0.024	3.738	73.9%	0.023	3.031	21.9%
	BMAg*	0.038	3.445	44.6%	0.043	1.834	0.3%	0.039	3.635	63.6%	0.040	2.802	11.2%
	MPMg	0.222	2.767	1.1%	0.235	0.932	0.0%	0.223	3.050	10.1%	0.225	1.836	0.0%
	HPMg	0.234	2.837	0.0%	0.245	1.037	0.0%	0.230	2.992	0.0%	0.240	1.922	0.0%
	BMAi**	0.020	3.562	56.4%	0.028	2.122	1.8%	0.023	3.742	74.2%	0.023	3.050	22.5%
	BMAi*	0.038	3.423	42.9%	0.045	1.858	0.4%	0.038	3.631	63.2%	0.039	2.835	12.3%
	MPMi	0.211	2.464	0.1%	0.237	0.982	0.0%	0.201	2.865	4.2%	0.214	1.937	0.0%
	HPMi	0.199	2.608	0.0%	0.245	1.097	0.0%	0.183	2.965	0.0%	0.218	2.024	0.0%
200	OLS**	0.298	2.951	17.5%	0.303	0.124	0.0%	0.295	3.426	47.7%	0.293	1.099	0.0%
	OLS*	0.594	2.509	3.9%	0.600	0.063	0.0%	0.597	3.085	25.1%	0.582	0.728	0.0%
	BMAg**	0.016	3.294	32.5%	0.015	0.740	0.0%	0.015	3.545	54.5%	0.016	2.145	1.4%
	BMAg*	0.027	3.107	18.7%	0.027	0.578	0.0%	0.026	3.433	43.3%	0.026	1.890	0.2%
	MPMg	0.137	2.251	0.0%	0.145	0.195	0.0%	0.136	2.862	1.8%	0.138	1.039	0.0%
	HPMg	0.141	2.204	0.0%	0.148	0.213	0.0%	0.140	2.938	0.0%	0.143	1.145	0.0%
	BMAi**	0.015	3.247	29.7%	0.015	0.781	0.0%	0.013	3.556	55.6%	0.015	2.194	1.6%
	BMAi*	0.024	3.035	15.7%	0.026	0.615	0.0%	0.022	3.430	43.0%	0.024	1.950	0.4%
	MPMi	0.120	1.991	0.0%	0.136	0.225	0.0%	0.111	2.695	0.5%	0.124	1.132	0.0%
	HPMi	0.118	2.103	0.0%	0.135	0.248	0.0%	0.103	2.819	0.0%	0.128	1.237	0.0%
Monte Carlo Study B													
90	OLS**	0.205	2.393	0.6%	0.196	2.395	0.7%	0.196	4.110	11.7%	0.198	4.102	11.5%
	OLS*	0.409	1.694	0.1%	0.401	1.697	0.1%	0.397	3.391	3.7%	0.401	3.375	3.6%
	BMAg**	0.006	4.491	20.5%	0.011	4.498	20.4%	0.006	5.487	59.4%	0.013	5.502	60.6%
	BMAg*	0.011	4.140	13.0%	0.021	4.146	12.8%	0.010	5.308	48.5%	0.021	5.316	49.3%
	MPMg	0.085	2.681	0.7%	0.133	2.677	0.8%	0.082	4.266	11.0%	0.133	4.264	11.5%
	HPMg	0.146	2.681	0.2%	0.164	2.680	0.4%	0.148	4.188	2.5%	0.173	4.203	3.5%
	BMAi**	0.014	4.495	21.2%	0.009	4.510	21.3%	0.010	5.504	60.7%	0.010	5.516	61.5%
	BMAi*	0.023	4.149	13.6%	0.016	4.161	13.6%	0.018	5.332	50.3%	0.016	5.340	50.6%
	MPMi	0.191	2.711	0.9%	0.095	2.728	1.0%	0.169	4.344	13.1%	0.086	4.344	12.6%
	HPMi	0.202	2.754	0.2%	0.114	2.762	0.3%	0.176	4.304	2.9%	0.108	4.309	2.7%
200	OLS**	0.196	0.436	0.0%	0.199	0.433	0.0%	0.205	2.221	0.2%	0.195	2.238	0.3%
	OLS*	0.397	0.229	0.0%	0.393	0.227	0.0%	0.407	1.550	0.0%	0.396	1.577	0.1%
	BMAg**	0.002	2.075	0.3%	0.008	2.074	0.4%	0.003	4.605	21.1%	0.006	4.602	20.8%
	BMAg*	0.004	1.725	0.1%	0.014	1.725	0.1%	0.005	4.283	13.5%	0.014	4.297	13.9%
	MPMg	0.038	0.763	0.0%	0.082	0.758	0.0%	0.042	2.965	0.9%	0.078	2.969	1.1%
	HPMg	0.068	0.751	0.0%	0.098	0.748	0.0%	0.073	2.974	0.2%	0.097	2.973	0.2%
	BMAi*	0.005	2.110	0.3%	0.006	2.114	0.5%	0.006	4.672	23.7%	0.004	4.669	23.0%
	BMAi*	0.009	1.763	0.1%	0.010	1.762	0.2%	0.010	4.362	15.3%	0.009	4.379	15.7%
	MPMi	0.076	0.800	0.0%	0.056	0.797	0.0%	0.075	3.102	1.5%	0.049	3.111	1.5%
	HPMi	0.089	0.793	0.0%	0.066	0.785	0.0%	0.086	3.140	0.3%	0.058	3.138	0.2%

NOTE: FI shows the average number of false inclusions; FE the average number of false exclusions; and *Null* the percentage of times that the null model is selected. See text for further details.

Table 3: Monte Carlo Simulation for variable selection methods. Bias and estimation error.

		$\sigma_\varepsilon = 4$						$\sigma_\varepsilon = 6$					
		$\rho = 0.9$			$\rho = 0.5$			$\rho = 0.9$			$\rho = 0.5$		
N	mod	$bias_1$	$bias_5$	RMSE	$bias_1$	$bias_5$	RMSE	$bias_1$	$bias_5$	RMSE	$bias_1$	$bias_5$	RMSE
Monte Carlo Study A													
90	OLS	0.010	-0.004	6.736	-0.009	0.001	1.868	-0.003	0.010	15.095	0.012	-0.008	4.185
	BMAg	-0.031	-0.001	5.032	-0.069	-0.000	1.461	-0.041	0.003	5.936	-0.087	-0.001	2.735
	MPMg	-0.033	0.001	9.639	-0.046	-0.001	1.787	-0.119	-0.000	11.526	-0.067	-0.000	3.994
	HPMg	-0.027	0.000	10.213	-0.051	-0.000	1.921	-0.036	0.002	12.446	-0.070	0.000	4.156
	BMAi	-0.014	-0.001	4.679	-0.062	-0.000	1.529	-0.028	0.003	6.321	-0.087	-0.002	2.878
	MPMi	-0.011	0.001	7.961	-0.050	-0.001	1.863	-0.061	0.000	11.403	-0.077	0.001	4.202
	HPMi	-0.015	0.000	8.707	-0.055	-0.001	2.011	-0.035	0.000	12.520	-0.085	0.001	4.357
200	OLS	0.004	-0.000	2.839	0.003	0.001	0.779	-0.002	-0.001	6.380	0.002	0.002	1.741
	BMAg	-0.022	0.000	3.627	-0.014	0.001	0.533	-0.042	-0.000	5.347	-0.053	0.000	1.470
	MPMg	-0.016	0.000	6.144	-0.003	0.001	0.514	-0.064	-0.001	10.097	-0.033	-0.000	1.816
	HPMg	-0.008	0.000	5.538	-0.003	0.001	0.532	-0.049	-0.001	10.830	-0.040	-0.000	1.943
	BMAi	-0.013	0.000	3.196	-0.010	0.001	0.557	-0.031	-0.001	4.957	-0.053	0.000	1.549
	MPMi	-0.010	0.000	4.938	-0.004	0.001	0.540	-0.049	-0.001	9.112	-0.040	0.000	1.943
	HPMi	-0.006	0.000	4.957	-0.004	0.001	0.563	-0.046	-0.001	9.902	-0.047	-0.000	2.074
Monte Carlo Study B													
90	OLS	0.001	0.004	6.766	0.001	-0.001	1.864	-0.008	-0.003	15.027	-0.004	0.005	4.194
	BMAg	0.001	-0.305	2.857	0.001	-0.306	2.548	-0.002	-0.467	4.540	0.001	-0.467	3.900
	MPMg	0.000	-0.277	3.736	0.000	-0.274	3.416	-0.001	-0.495	6.370	0.003	-0.496	5.736
	HPMg	0.000	-0.270	4.012	0.000	-0.268	3.459	-0.002	-0.475	6.944	0.004	-0.481	5.789
	BMAi	0.003	-0.302	3.254	0.001	-0.305	2.572	-0.002	-0.475	5.274	0.001	-0.476	3.932
	MPMi	0.004	-0.280	4.274	0.001	-0.282	3.444	0.002	-0.509	7.385	0.003	-0.515	5.752
	HPMi	0.006	-0.281	4.493	-0.001	-0.281	3.508	-0.000	-0.498	7.588	0.003	-0.507	5.791
200	OLS	0.004	0.004	2.798	0.003	0.001	0.779	0.009	0.008	6.391	0.002	0.004	1.773
	BMAg	0.001	-0.105	1.135	0.000	-0.109	1.070	0.003	-0.332	2.831	0.000	-0.335	2.686
	MPMg	-0.000	-0.061	1.174	-0.000	-0.066	1.116	0.002	-0.314	3.780	-0.001	-0.320	3.626
	HPMg	-0.000	-0.059	1.204	-0.000	-0.065	1.111	0.000	-0.312	3.903	-0.001	-0.317	3.652
	BMAi	0.001	-0.105	1.235	0.000	-0.109	1.090	0.003	-0.347	3.081	-0.000	-0.350	2.768
	MPMi	0.002	-0.066	1.288	-0.000	-0.070	1.145	0.003	-0.337	4.103	-0.002	-0.344	3.750
	HPMi	0.002	-0.064	1.324	-0.000	-0.069	1.135	0.003	-0.342	4.236	-0.002	-0.347	3.791

NOTE: $bias_1$ is the average bias of the estimated coefficient associated with the first variable; $bias_5$ that associated with the fifth variable. RMSE denotes the average root mean square error of the 10 estimated parameters. See text for further details.

Table 4: Model-free inflation risk premium estimation results.

	Full sample						Without outliers					
	In levels			In differences			In levels			In differences		
	BMA	MPM	US	BMA	MPM	US	BMA	MPM	US	BMA	MPM	US
Lag IRP	0.181** (0.035)	0.180 (0.184)	0.455** (0.198)	-0.067 (0.040)	-0.079 (0.167)	0.003 (0.014)	0.184** (0.039)	0.170 (0.189)	0.305** (0.041)	-0.010 (0.023)	-	-0.006 (0.017)
bond volatility	-0.002 (0.012)	-	-	0.000 (0.008)	-	-0.000 (0.035)	-0.003 (0.011)	-	0.004 (0.013)	0.008 (0.020)	-	0.024 (0.033)
realised volatility	-0.000 (0.008)	-	-0.071 (0.202)	-0.006 (0.018)	-	-0.010 (0.027)	-0.002 (0.011)	-	-0.003 (0.013)	-0.005 (0.016)	-	-0.005 (0.016)
implied volatility	0.006 (0.020)	-	0.266 (0.220)	-0.010 (0.028)	-	0.181** (0.051)	0.051 (0.061)	0.081 (0.212)	0.081 (0.082)	0.005 (0.019)	-	0.117** (0.037)
Mean-median Consensus	-0.005 (0.015)	-	0.005 (0.017)	-0.005 (0.016)	-	0.001 (0.010)	-0.001 (0.010)	-	0.001 (0.009)	-0.003 (0.013)	-	0.001 (0.010)
Liquidity	-0.107** (0.034)	-0.109 (0.174)	-0.198 (0.220)	-0.071 (0.038)	-0.080 (0.165)	-0.081 (0.044)	-0.071 (0.042)	-0.086 (0.170)	-0.036 (0.048)	-0.003 (0.014)	-	-0.009 (0.021)
Consensus disagreement	0.001 (0.010)	-	-	-0.000 (0.008)	-	-0.000 (0.010)	0.002 (0.012)	-	0.064 (0.064)	0.023 (0.033)	-	0.061 (0.040)
data surprise	0.001 (0.009)	-	0.088 (0.172)	0.000 (0.008)	-	0.038 (0.042)	0.001 (0.010)	-	0.060 (0.042)	-0.003 (0.012)	-	0.008 (0.020)
data surprise uncertainty	0.003 (0.012)	-	-	-0.003 (0.012)	-	-0.004 (0.014)	0.002 (0.011)	-	0.001 (0.010)	0.000 (0.008)	-	0.002 (0.011)
Floor at 0	-0.096** (0.036)	-0.095 (0.167)	-0.302 (0.234)	-0.096** (0.041)	-0.109 (0.167)	-0.295** (0.045)	-0.145** (0.047)	-0.161 (0.196)	-0.189** (0.078)	-0.028 (0.037)	-	-0.204** (0.036)
Cap at 4	0.187** (0.038)	0.190 (0.187)	-	0.108** (0.031)	0.108 (0.167)	0.005 (0.022)	0.117 (0.064)	0.116 (0.208)	0.066 (0.064)	0.090** (0.034)	0.099 (0.167)	0.029 (0.037)
R^2	0.642 105	0.639 105	0.878 105	0.335 104	0.328 104	0.436 104	0.655 92	0.656 92	0.756 94	0.202 88	0.125 88	0.404 85
obs												


NOTE: Standard deviations of estimated coefficients are reported in between brackets. Significance levels higher than 1%, 5% and 10% are denoted respectively with ***, ** and *. Obs is used to denote the number of available observations.

Table 5: Model-based inflation risk premium estimation results.

	Full sample						Without outliers					
	In levels			In differences			In levels			In differences		
	BMA	MPM	EA	BMA	MPM	EA	BMA	MPM	EA	BMA	MPM	EA
Lag IRP	0.095** (0.015)	0.090 (0.115)		-0.019 (0.017)	-0.026 (0.108)		0.108** (0.014)	0.115 (0.107)		-0.003 (0.007)	-	
bond volatility	0.001 (0.005)	-		0.001 (0.005)	-		0.006 (0.012)	-		0.007 (0.012)	-	
realised volatility	-0.001 (0.004)	-		-0.076** (0.027)	-		0.001 (0.004)	-		-0.001 (0.005)	-	
implied volatility	-0.003 (0.010)	-		0.114** (0.034)	-		0.002 (0.007)	-		0.002 (0.008)	-	
Mean-median Consensus	0.000 (0.003)	-		0.001 (0.012)	-		-0.005 (0.007)	-		-0.003 (0.008)	-	
Liquidity	-0.036 (0.019)	-0.042 (0.112)		0.006 (0.016)	-0.024 (0.105)		0.011 (0.011)	-		0.007 (0.007)	-	
Consensus disagreement	-0.000 (0.004)	-		-0.000 (0.004)	-		-0.004 (0.005)	-		-0.006 (0.003)	-	
data surprise	-0.006 (0.011)	-		0.121** (0.015)	-0.024 (0.107)		-0.006 (0.011)	-		0.059** (0.015)	-0.028 (0.103)	
data surprise uncertainty	0.001 (0.005)	-		0.000 (0.004)	-		0.002 (0.006)	-		0.001 (0.005)	-	
Floor at 0	-0.036* (0.018)	-0.041 (0.113)		-0.039* (0.018)	-0.046 (0.107)		-0.010 (0.015)	-		-0.007 (0.012)	-	
Cap at 4	0.028 (0.020)	0.034 (0.111)		0.026 (0.017)	0.031 (0.106)		0.004 (0.010)	-		0.014 (0.015)	0.027 (0.103)	
R^2	0.610 101	0.604 101		0.290 100	0.298 100		0.580 88	0.538 88		0.183 82	0.132 82	
obs												


NOTE: Standard deviations of estimated coefficients are reported in between brackets. Significance levels higher than 1%, 5% and 10% are denoted respectively with ***, ** and *. Obs is used to denote the number of available observations.

Figure 1: Inflation Linked Swap deviations from Consensus forecasts.



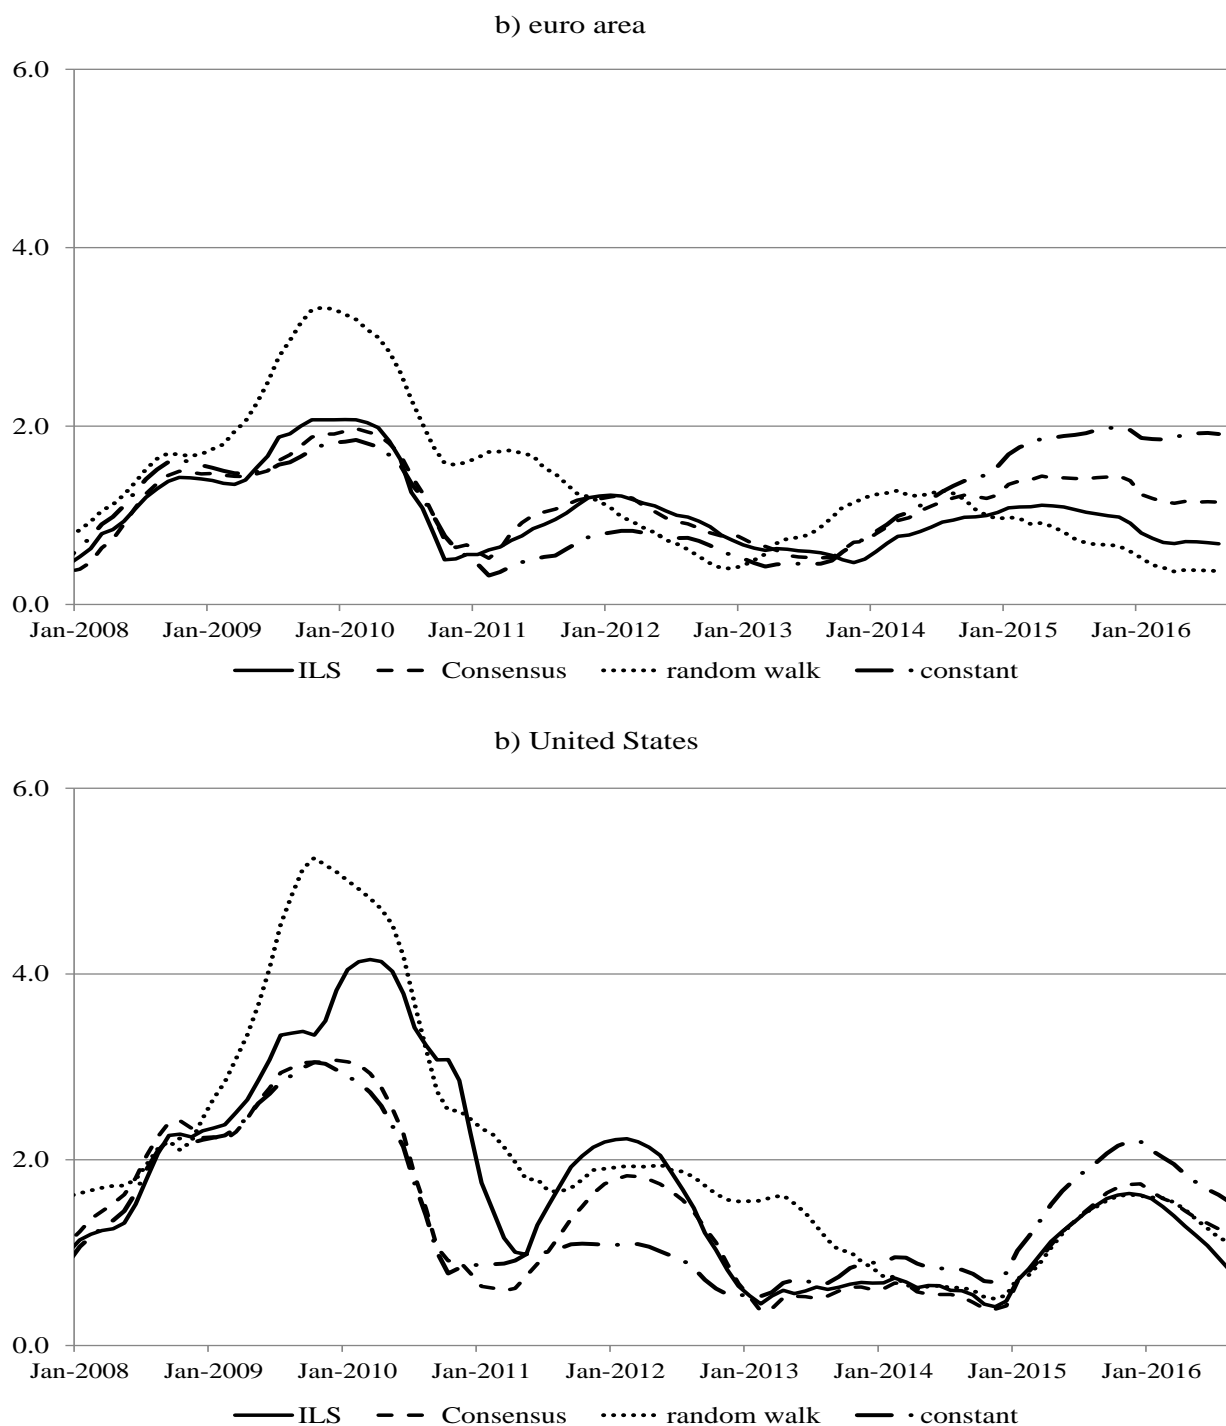
NOTE: Last observation is September 2016.

Figure 2: Model-free and model-based one-year ahead inflation risk premium.



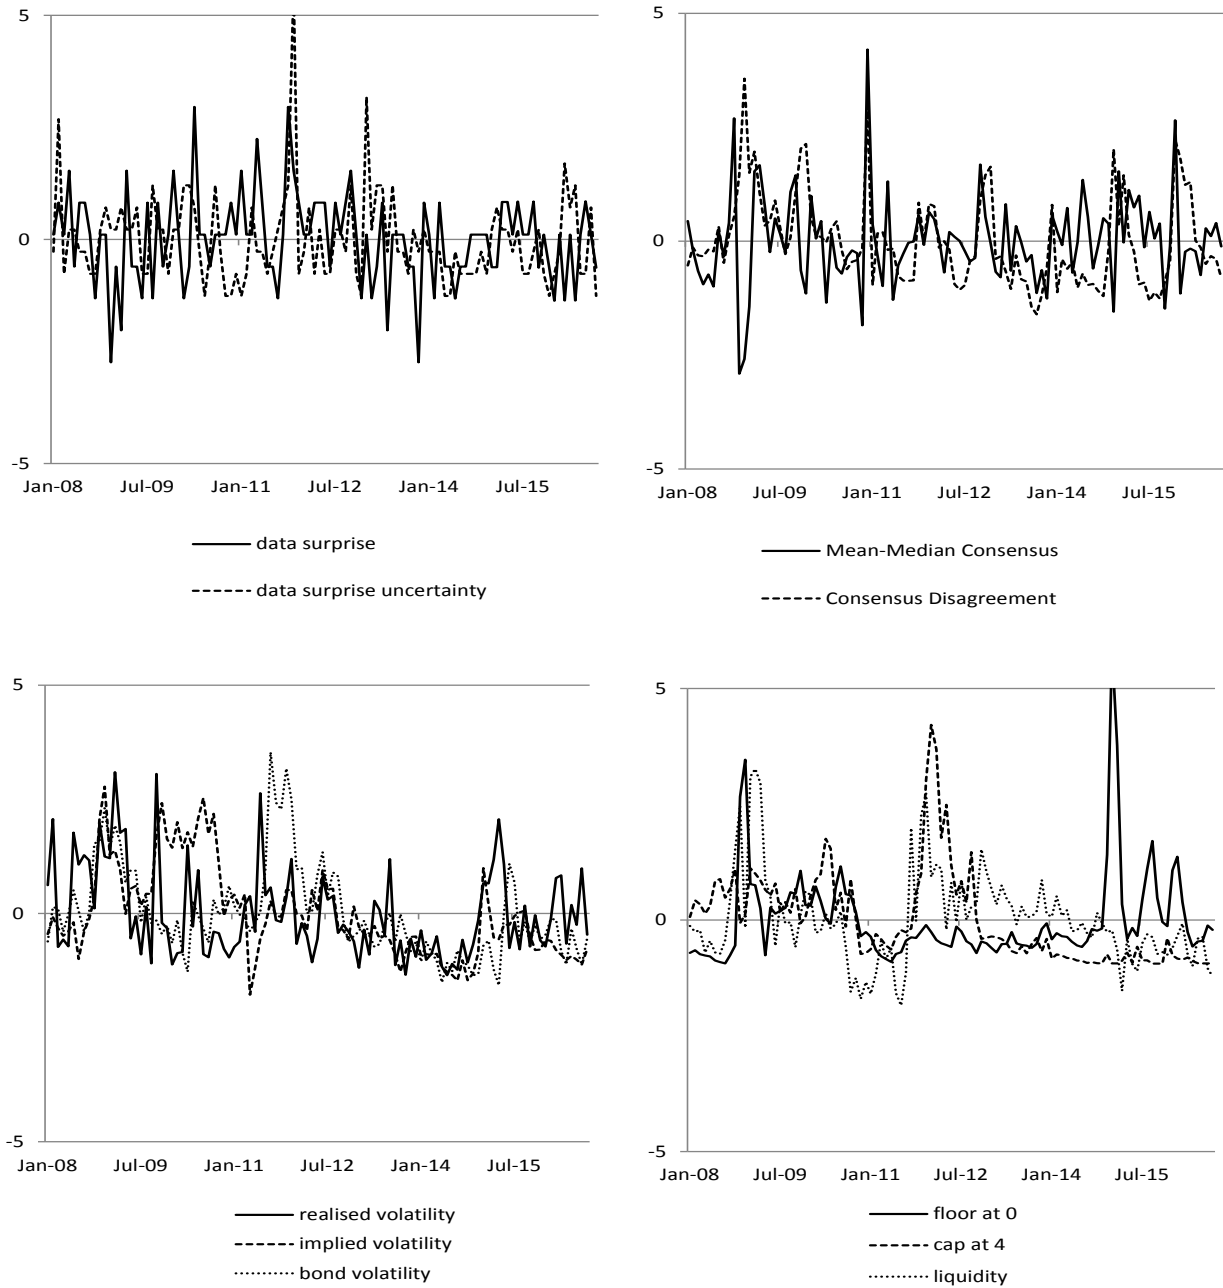
NOTE: Last observation is September 2016. See text for further details.

Figure 3: Rolling Root Mean Square errors of alternative CPI forecasts.



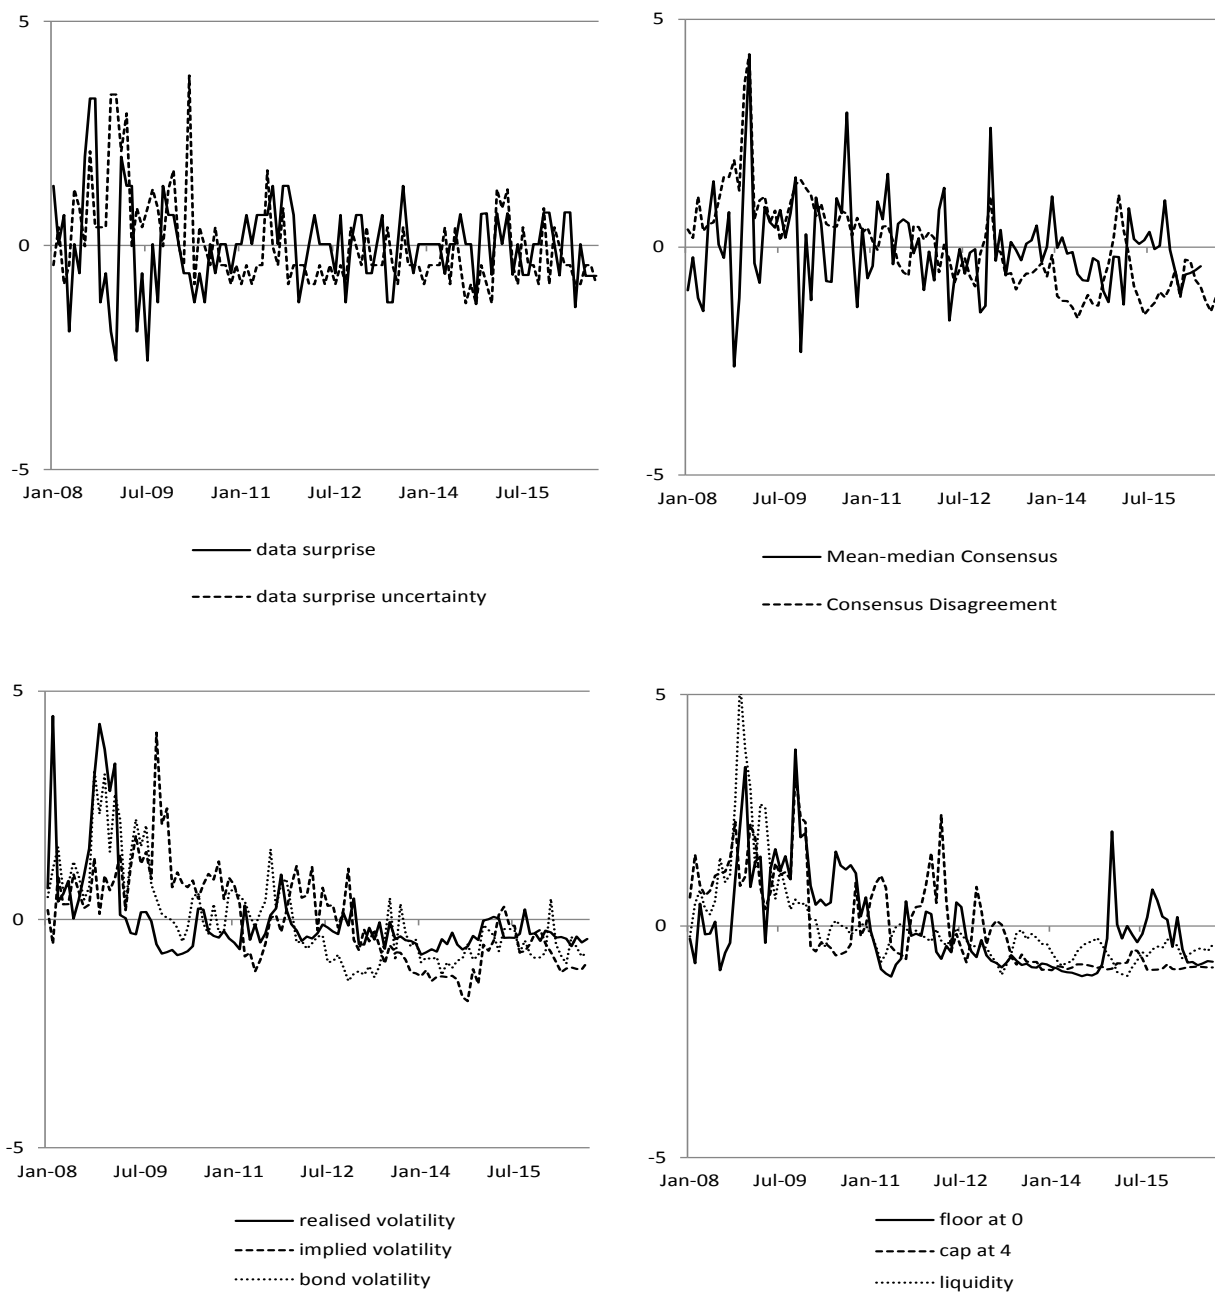
NOTE: Last observation is September 2016.

Figure 4: Explanatory variables for the euro area.



NOTE: Series shown have been normalised. Last observation is September 2016. See text for further details.

Figure 5: Explanatory variables for the United States.



NOTE: Series shown have been normalised. Last observation is September 2016. See text for further details.

Acknowledgements

Gonzalo Camba-Mendez

European Central Bank, Frankfurt, Germany; email: gonzalo.camba-mendez@ecb.europa.eu

Thomas Werner

European Central Bank, Frankfurt, Germany; email: thomas.werner@ecb.europa.eu

© European Central Bank, 2017

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

All rights reserved. Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the authors.

This paper can be downloaded without charge from www.ecb.europa.eu, from the [Social Science Research Network electronic library](#) or from [RePEc: Research Papers in Economics](#). Information on all of the papers published in the ECB Working Paper Series can be found on the [ECB's website](#).

ISSN	1725-2806 (pdf)	DOI	10.2866/556561 (pdf)
ISBN	978-92-899-2755-0 (pdf)	EU catalogue No	QB-AR-17-045-EN-N (pdf)