Is euro area lowflation here to stay ? Insights from a time-varying parameter model with survey data

Working Paper Research

by Arnoud Stevens and Joris Wauters

October 2018 No 355

NationalBank OF BELGIUM
Eurosystem

Is euro area lowflation here to stay? Insights from a time-varying parameter model with survey data.*

Arnoud Stevens†
National Bank of Belgium

Joris Wauters
National Bank of Belgium and Ghent University

15th October 2018

Preliminary draft prepared for the NBB 2018 International Conference

Abstract

Inflation has been persistently weak in the euro area despite the economic recovery since 2013. We investigate the sources behind this protracted low inflation by building a time-varying parameter model that jointly explains the dynamics of inflation and inflation expectations from the ECB's Survey of Professional Forecasters. We find that the inclusion of survey data strengthens the view that low inflation was mainly due to cyclical drivers. In particular, the model with survey expectations finds a more muted decline of trend inflation in recent years and a larger degree of economic slack. The impact of economic slack and import prices on inflation is found to have increased in recent years. We also find that survey expectations have become less persistent over the financial crisis period, and that including survey data improves the model's out-of-sample forecasting performance.

Keywords: inflation dynamics, trend inflation, survey-based inflation expectations, ECB Survey of Professional Forecasters, nonlinear state space model, Bayesian estimation, euro area
JEL Classification: E31, C11, C32

*This paper has been written for the National Bank of Belgium (NBB) 2018 International Conference on "Understanding inflation dynamics: the role of costs, mark-ups and expectations". Part of the research was undertaken while Stevens was on secondment at the European Central Bank (Prices and Costs Division). We would like to thank Jef Boeckx, Ferre De Graeve, Carlos Montes-Galdón, Rafael Wouters, members of the ECB's Prices and Costs Division and participants at the preparatory meetings of the conference for useful comments and suggestions. All remaining errors are our own. Any views expressed herein are those of the authors and do not necessarily reflect those of the National Bank of Belgium or the Eurosystem.

†National Bank of Belgium, Economics and Research Department, de Berlaimontlaan 14, B-1000 Brussels, Belgium. Email addresses: arnoud.stevens@nbb.be and joris.wauters@nbb.be

1 Introduction

The euro area enjoyed an economic recovery over the past five years, yet headline inflation remained persistently weak. Specifically, the year-on-year growth rate of the Harmonized Index of Consumer Prices (HICP) has averaged 0.8% since 2013, thus keeping a good distance from both its pre-crisis (1999-2007) average of 2.06%, as well as the European Central Bank (ECB)'s target of below, but close to, 2% in the medium term (Figure 1). Economists speak in this respect of a "missing inflation puzzle", because inflation was expected to be higher given the favourable economic recovery that followed the sovereign debt crisis (see, e.g., Ciccarelli and Osbat, 2017).

[INSERT FIGURE 1 HERE]

What explains this so-called *lowflation* period? Conceptually, we make a distinction between two causes. One the one hand, inflation could have been pushed down by *permanent* effects, which lowered its long-run trend. On the other hand, cyclical and thus *temporary* effects might have created a long-lasting downward drag on inflation. Even though the symptoms - low inflation - are similar in both cases, an accurate diagnosis remains important in order to identify the right monetary policy treatment. For instance, if cyclical effects are the main culprit, then inflation might still return to target within the medium term once the effects from shocks have faded out. However if the long-run trend has decreased, then additional stimulus might be necessary in order to re-align the trend with the target.

A recent stream of literature studies euro area inflation dynamics from different angles. A first strand uses constant parameter structural vector autoregressive (SVAR) models to distinguish between the different types of shocks affecting inflation (see, e.g., Neri et al., 2017, Jarociński and Bobeica, 2017, and Ciccarelli and Osbat, 2017). One important finding of this line of research is that shocks from both domestic and foreign origin have driven down inflation. A second strand of the literature uses time-varying parameter models in order to investigate the role of structural changes to inflation dynamics. The evidence from this literature points to a recent increasing sensitivity of inflation to cyclical conditions (see, e.g., Riggi and Venditti, 2015, Dany-Knedlik and Höltemoller, 2017, and Cordemans and Wauters, 2018), but also to a decline in inflation's long-run trend (Dany-Knedlik and Höltemoller, 2017).

This paper explores whether survey inflation expectations are helpful for measuring the contribution of trend and cyclical factors to low inflation. To this end, we estimate a time-varying parameter Phillips curve model that jointly explains macroeconomic data and inflation expectations from the ECB's Survey of Professional Forecasters, and contrast its dynamics with that of a model variant that abstracts from survey information. Forward-looking information variables might help to uncover shifts in inflation dynamics not necessarily captured in historical data. In fact, such forecasts typically incorporate new information regarding structural changes; including changes to the conduct of monetary policy, the formation of inflation expectations and price setting behaviour (see, e.g., Kozicki and Tinsley, 2012). Inflation expectations of professional experts have been considered as particularly attractive in this respect because of their proven success in forecasting inflation relative to statistical forecasts derived from past data (Ang et al., 2007; Faust and Wright, 2013).

We add to the empirical literature that models euro area inflation dynamics with the use of survey inflation expectations. Garcia and Poon (2018) and Banbura and van Vlodrop (2018) estimate models where inflation's long-run trend rate is linked to data from long-term survey forecasts or market-based expectations, and Grishchenko et al. (2017) incorporate the survey expected inflation distribution in a factor model with constant parameters and stochastic volatility. Yet, the combination of information from the expectations term structure (of short-, medium- and long-term expectations) has not yet been exploited for estimating the time variation in the dynamics of both the trend and cycle of euro area inflation. Our first contribution is to fill this gap and to show that the short and long end of the inflation expectations curve are complementary information sources for identifying structural changes.

Our proposed framework follows the spirit of Kozicki and Tinsley (2012) and other following work, including studies by Crump et al. (2016) and Winkelried (2017), in that the survey expectations for inflation are explained by the model-consistent forecast using the equations that describe the macroeconomic series. However, the literature typically adopts a constant parameter set-up for the transmission of cyclical factors to inflation. To the best of our knowledge, Mertens and Nason (2018) is the only other paper to provide a model-consistent treatment of inflation dynamics and the term structure of expectations in a time-varying parameter framework. They estimate the joint dynamics of US inflation and its expectations in an unobserved components model featuring sticky-expectations formation and time-varying parameters. Relative to their setup, our second contribution is that we further decompose the cyclical component of inflation into parts that relate to economic slack and import price developments, and infer their relative contributions to cyclical movements.

The baseline empirical model that we estimate is an open-economy unobserved-component Phillips curve model with time-varying parameters. The Phillips curve model provides a structural framework for explaining inflation dynamics. In its simplest variant, the Phillips curve prescribes a negative relationship between inflation and the amount of slack in economic activity. Modelling this Phillips curve relationship in an unobserved component set-up, in accordance with Stella and Stock (2013) and Chan et al. (2016), serves two purposes. First, it permits decomposing inflation into a permanent component called 'trend inflation' and short-run cyclical fluctuations. Second, it allows estimating a model-consistent slack measure and, therefore, covers concerns about official output gap measures providing a distorted picture of actual inflationary pressures.[1] To account for potential foreign price pressures, we add an international dimension to the traditional Phillips curve in the form of a variable measuring relative import prices. Finally, the presence of time variation in the model's parameters allows for measuring structural changes in the dynamics of the various determinants of inflation.

We consider two model extensions to the baseline model where the term structure of the SPF mean inflation predictions is linked to the model-consistent inflation expectations. One of these models allows for "forecast smoothing", in the sense that survey expectations can gradually respond to changes in the model forecast. However, we remain agnostic about the source of such smoothing behaviour. It could, for instance, be related to informational frictions, in accordance with the Mankiw and Reis (2002) sticky information paradigm, but it could also be grounded in the epidemiological interpretation of expectations' formation by Carroll (2003), emphasizing the tendency of well-informed professionals to keep their forecasts unchanged unless they are fully convinced about the usefulness of new incoming information.[2]

[1] Riggi and Venditti (2015) discuss the possibility that conventional output gap measures understate the amount of spare capacity in the euro area economy, thereby misguiding their relevance in pushing down inflation in recent years. In fact, alternative measures of slack estimates that match inflation dynamics after the sovereign debt crisis are found to perform better in predicting inflation in real time (see, e.g., Jarociński and Lenza, 2018).

[2] In its original interpretation the epidemiological model of expectations formation applies to agents not experienced in macroeconomic forecasting; such agents typically adjust their predictions with delay to information provided by the media. Applied to the forecasting behaviour of professional experts, Łyziak and Paloviita (2018) re-interpret this argument in terms of

The inclusion of survey data as endogenous variables to the model creates measurement equations which are nonlinear functions of the time-varying parameters. This precludes the use of standard Bayesian MCMC techniques for sampling from the posterior distribution. To estimate our nonlinear state space model, we follow Cogley (2005) and Koop and Potter (2011) and apply the single-move sampler of Carlin et al. (1992) in a Metropolis-Hastings within-Gibbs sampler. Broadly speaking, this implies some steps where we sample the time-varying parameters on a date-by-date basis, each time conditioning on the time-varying parameters from all other periods.

By comparing the estimates from the models with and without survey expectations data, we find that the inclusion of survey data strengthens the view that low inflation was mainly due to cyclical drivers. The results differ mainly in the estimated trends, in the sense that the models with survey expectations find a more muted decline of trend inflation in recent years and a lower natural rate of unemployment. As a result, the gap between the unemployment rate and its natural rate - a measure of economic slack - is larger and remains positive at the end of the sample for the models with expectations data. Concerning the time-varying coefficients that affect the transmission of the cyclical drivers to inflation, we find broadly similar evolutions. The Phillips curve slope flattens between 1990 and 2012, but then steepens sharply. The effect of import prices is found to have increased since 2000. The intrinsic inflation gap persistence is generally low and flat, but follows a short-lived spike in 2008-2009 in the models with survey data. In the model which allows for a time-varying degree of forecast smoothing by survey respondents, we find that the smoothing coefficient declines during the financial crisis period. Although this finding is consistent with more frequent updating by survey respondents, the persistence in survey expectations remains high in general. Overall, the models with survey data attribute the lowflation period since 2013 mainly to a downward drag from import prices and economic slack, rather than a strong reduction in trend inflation.

Our strategy of modelling the expectations term structure in a time-varying parameter model invokes complications relative to the typical approach in the literature, which uses either only long-term survey expectations or constant parameters for the transmission coefficients in the inflation gap. We therefore evaluate the usefulness of our approach in two ways. First, we perform an out-of-sample forecasting comparison between the different models. We find that the models with survey data outperform the baseline model both in terms of root mean squared errors and log predictive density scores. From the perspective of a practitioner who is concerned with inflation predictability, this finding supports the inclusion of survey data in a time-varying parameter inflation model. Second, we perform robustness checks to evaluate if the results remain the same when we simplify the model to *i)* using only long-term expectations data, or *ii)* featuring constant transmission coefficients. We find that these perturbations bring economically meaningful changes to the results, which leads us to conclude that short- and medium-term expectations bring useful information for detecting changes in the dynamics of the inflation gap.

The paper proceeds as follows. Section 2 presents the different empirical models. Section 3 contains details on the data, estimation method and the choice of priors. In Section 4, we present the full sample estimates of our model, including an assessment of the factors behind the lowflation and of the information content of SPF inflation expectations. In section 5, we undertake a forecasting exercise. This exercise serves to assess the value of SPF data in forecasting inflation, but also to evaluate the usefulness of survey data in understanding euro area inflation dynamics over the past two decades. Section 6 discusses two robustness

forecasters fearing reputational risks if their projections would be revised too frequently.

exercises. Finally, Section 7 concludes the paper.

2 Empirical models

Our goal is to determine whether survey data helps to diagnose whether lowflation in the euro area is driven by either *permanent* effects, i.e. shifts in inflation's long-term trend, or by *cyclical* and thus temporary effects. To this end, we estimate three different model specifications. The first specification, our *baseline model,* is an empirical Phillips curve model with time-varying parameters and stochastic volatility, which we estimate on macroeconomic time series alone.

The next two models extend the baseline model with information from surveys. These specifications add equations where inflation expectations series are fitted in a model-consistent way. By doing so, we can infer how the forecasters' expectations affect the parameter estimates and, therefore, the economic interpretation of the recent lowflation. The first model with survey data allows for a gradual adjustment of survey expectations to changes in the model forecast; the other assumes no type of rigidity. We discuss the three models in turn.

2.1 Baseline model

Our modelling framework builds on the bounded bivariate unobserved components model of Chan et al. (2016). Their model jointly estimates the degree of slack in the economy and the Phillips curve relationship between slack and inflation. The trends and coefficients are allowed to be time-varying in order to capture structural changes to the economy. Our baseline model extends their framework by also including the effects from import price inflation. The measurement equation for inflation takes the following form:

$$\pi_t - \tau_t^{\pi} = \rho_t^{\pi}\left(\pi_{t-1} - \tau_{t-1}^{\pi}\right) + \lambda_t\left(u_t - \tau_t^{u}\right) + \gamma_t\left(\pi_t^{m} - \tau_t^{m}\right) + \epsilon_t^{\pi}, \tag{1}$$

where π_t stands for inflation, u_t for the unemployment rate, and π_t^m for inflation in the relative price of imports. The trends for these variables are, respectively, τ_t^π, τ_t^u and τ_t^m. Hence, equation (1) states that the deviation of inflation from its long-term trend $(\pi_t - \tau_t^\pi)$, henceforth defined as the "inflation gap", is a stationary process that mean-reverts to zero. This gap depends on its own lag, the "unemployment gap" between the unemployment rate and its own trend $(u_t - \tau_t^u)$, and the "import price inflation gap" between relative import price inflation and its own trend $(\pi_t^m - \tau_t^m)$. The impact of these three variables on the inflation gap is measured by the time-varying coefficients ρ_t^π, λ_t and γ_t. They denote, respectively, the degree of intrinsic inflation gap persistence, the Phillips curve slope, and the impact effect from the import price inflation gap.

As in Chan et al. (2016), the unemployment gap is assumed to evolve as an autoregressive process of order 2:

$$u_t - \tau_t^{u} = \rho_1^{u}\left(u_{t-1} - \tau_{t-1}^{u}\right) + \rho_2^{u}\left(u_{t-2} - \tau_{t-2}^{u}\right) + \epsilon_t^{u}, \tag{2}$$

where the autoregressive parameters ρ_1^u and ρ_2^u are restricted to ensure stationarity.

The novelty relative to Chan et al. (2016), is the addition of the term $\gamma_t\left(\pi_t^m - \tau_t^m\right)$ in equation (1) - a time-varying effect from the import price inflation gap. This extension is motivated by the strong fluctuations of commodity prices during the financial crisis, and the hypothesis that globalization has made imported inflation more important in the recent period (IMF, 2013).[3] We also add a measurement equation for this gap which, for simplicity, is assumed to evolve as as an i.i.d. process:

$$\pi_t^m - \tau_t^m = \epsilon_t^m. \tag{3}$$

To account for changing volatility in the shocks to inflation, the measurement error equation (1) follows a stochastic volatility process as $\epsilon_t^\pi \sim N(0, e^{\psi_t})$, where $N\left(\mu, \sigma^2\right)$ denotes a Normal distribution with mean μ and variance σ^2. The remaining measurement errors are assumed to have constant volatility: $\epsilon_t^u \sim N(0, \sigma_u^2)$, and $\epsilon_t^m \sim N(0, \sigma_m^2)$.

The state equations for the time-varying trends $(\tau_t^\pi, \tau_t^u, \tau_t^m)$, the time-varying coefficients $(\rho_t^\pi, \lambda_t, \gamma_t)$, and the log of the volatility of the error term in equation (1), ψ_t, are assumed to follow independent random walks:

$$\tau_t^\pi = \tau_{t-1}^\pi + \eta_t^{\tau\pi} \tag{4}$$

$$\tau_t^u = \tau_{t-1}^u + \eta_t^{\tau u} \tag{5}$$

$$\tau_t^m = \tau_{t-1}^m + \eta_t^{\tau m} \tag{6}$$

$$\rho_t^\pi = \rho_{t-1}^\pi + \eta_t^\rho \tag{7}$$

$$\lambda_t = \lambda_{t-1} + \eta_t^\lambda \tag{8}$$

$$\gamma_t = \gamma_{t-1} + \eta_t^\gamma \tag{9}$$

$$\psi_t = \psi_{t-1} + \eta_t^\psi. \tag{10}$$

As in Chan et al. (2016), the parameters λ_t and ρ_t^π are bounded to lie in the $(-1, 0)$ and $(0, 1)$ intervals, respectively. These restrictions imply that the error terms for these state variables are drawn from truncated normal distributions as

$$\eta_t^\lambda \sim TN(-1 - \lambda_{t-1}^\pi, 0 - \lambda_{t-1}^\pi, 0, \sigma_\lambda^2) \tag{11}$$

$$\eta_t^\rho \sim TN(0 - \rho_{t-1}^\pi, 1 - \rho_{t-1}^\pi, 0, \sigma_\rho^2), \tag{12}$$

where $TN(a, b; \mu, \sigma^2)$ denotes the normal distribution with mean μ and variance σ^2 that is truncated to the interval (a, b).[4] The remaining error terms for the unbounded states follow independent normal distributions:

$$\eta_t^i \sim N(0, \sigma_i^2), \tag{13}$$

with $i \in (\tau\pi, \tau u, \tau m, \gamma, \psi)$.

In this framework, shifts in trend inflation τ_t^π have permanent effects on inflation. If the recent lowflation is mainly driven by a decline in trend inflation, it follows that inflation will eventually converge to the new lower level of τ_t^π in the long run. On the other hand, lowflation can also be driven by a persistent decline in the inflation gap $\pi_t - \tau_t^\pi$. Such a decline can be due to three factors. First, the direct effect from shocks to the inflation gap (ϵ_t^π). Second, indirect effects from shocks to the unemployment gap (ϵ_t^u) or the import price gap (ϵ_t^m). And third, changes in the impact coefficients $(\rho_t^\pi,\ \lambda_t,\ \gamma_t)$ which affect the transmission of these shocks and, ultimately, the speed at which inflation returns to its long-term trend τ_t^π.

[3] Other studies which measure a time-varying impact of import price inflation are Matheson and Stavrev (2013), IMF (2013), and Dany-Knedlik and Höltemoller (2017). In contrast to these studies, which de-mean the import price variable π_t^m prior to estimation, we also estimate a trend τ_t^m in order to account for the possible presence of slow-moving shifts in relative import price inflation which leave the cyclical nature of headline inflation unaffected.

[4] Chan et al. (2016) also impose bounds to trend inflation τ_t^π and the natural rate of unemployment τ_t^u. In our application, we leave these trends unbounded as we have no strong prior on their behaviour.

2.2 Extensions with survey data

To assess the role of survey data in uncovering the factors behind euro area lowflation, we also estimate two models which include survey expectations data on inflation. These two models build on the baseline model by adding measurement equations which connect the survey expected inflation series with the model-consistent inflation forecast, as derived using the equations described above.

SPF expectations model The first specification assumes that the reported survey expectation at time t for future inflation at time $t+h$ is a weighted average between the previous period's expectation, and the model-consistent inflation forecast. Specifically, we add a set of measurement equations

$$\pi^e_{t+h_1|t} = (1-\xi_{t-1})\, f_{h_1}\left(\theta_{t-1},\, Y^{t-1}\right) + \xi_{t-1}\pi^e_{t-1+h_1|t-1} + \epsilon^{h_1}_t \tag{14}$$

$$\vdots$$

$$\pi^e_{t+h_n|t} = (1-\xi_{t-1})\, f_{h_n}\left(\theta_{t-1},\, Y^{t-1}\right) + \xi_{t-1}\pi^e_{t-1+h_n|t-1} + \epsilon^{h_n}_t, \tag{15}$$

where $\pi^e_{t+h|t}$ represents the survey expectation for year-on-year inflation rate in period $t+h$ given the information at the start of period t. Each survey observation in quarter t is linked with the h quarter ahead model forecast in period t, $f_h\left(\theta_{t-1}, Y^{t-1}\right)$, the previous period's survey expectation for the same horizon h, and a measurement error. The weighting between the survey data and the model forecast is determined by the smoothing coefficient ξ_t, which we assume to evolve as an independent random walk that is bounded to lie within the $(0,1)$ interval:

$$\xi_t = \xi_{t-1} + \eta^{\xi}_t,$$

where

$$\eta^{\xi}_t \sim TN(0-\xi_{t-1}, 1-\xi_{t-1}, 0, \sigma^2_{\xi}).$$

Finally, the residuals are independently distributed as $\epsilon^{h_1}_t \sim N(0, \sigma^2_{h_1})$,..., and $\epsilon^{h_n}_t \sim N(0, \sigma^2_{h_n})$.

This specification allows for survey expectations to gradually adjust to changes in the underlying model forecast. To see this, subtract the previous period's expectation from both sides to deliver, using simplified notation, $\Delta\pi^e_t = (1-\xi)(f_t - \pi^e_{t-1}) + \epsilon_t$. That is, the change in reported survey expectation is a fraction of the gap between the current period's model forecast and the previous period's survey expectation. A larger ξ implies more sluggish adjustment of survey expectations to changes in the underlying model forecast. Similar specifications where survey expectations adjust gradually to either the model forecast or to other data can be found in Baele et al. (2015) and Łyziak and Paloviita (2018). Partial adjustment of survey expectations is typically motivated by informational rigidities, such as the sticky information framework of Mankiw and Reis (2002) and the epidemiological model of Carroll (2003), or by the strategic behaviour of forecasters who, on the one hand, aim to minimize forecast errors, but on the other hand want to avoid making large forecast revisions for reputational concerns (Łyziak and Paloviita, 2018). Our empirical setup is not tied to a specific theoretical model. Instead, it assumes that the reduced form of the forecasters' structural model can be approximated with the reduced form baseline model described above.

We define the information set at the beginning of period time t to include the ECB's SPF data, which are collected at the start of quarter t, and the macroeconomic data up to the previous quarter. Hence, the model forecast function is based on data stacked in Y^{t-1}, where $Y^{t-1} = (Y_1', ..., Y_{t-1}')'$ and $Y_t = (\pi_t, u_t, \pi^m_t)'$, as well as the relevant (constant and time-varying) model parameters in the previous quarter, collected in

θ_{t-1}. To generate the model forecast of inflation, we rewrite equations (1) to (3) as a reduced form VAR and iterate the model forward (appendix A.2). One complication at this stage, is how to treat the path of the time-varying parameters in the future. Following Cogley (2005) and Mertens and Nason (2018), we invoke the anticipated utility model and keep the future values of the time-varying parameters such as, e.g., trend inflation τ_t^π constant at their current values. This assumption can be interpreted as approximations to more complex systems, or a form of bounded rationality by the agents (Cogley, 2005).

SPF no smoothing model The third model is a nested case of the above SPF expectations model where the smoothing parameter $\xi_t = 0$ in all periods. This assumption changes the measurement equations (14) to (15) into:

$$\pi^e_{t+h_1|t} = f_{h_1}\left(\theta_{t-1}, Y^{t-1}\right) + \epsilon_t^{h_1} \tag{16}$$

$$\vdots$$

$$\pi^e_{t+h_n|t} = f_{h_n}\left(\theta_{t-1}, Y^{t-1}\right) + \epsilon_t^{h_n}. \tag{17}$$

We label this specification as our "SPF no smoothing model", because it assumes that model forecasts align with the survey data.

2.3 Relation with the literature

Following the work of Kozicki and Tinsley (2012), a burgeoning literature has emerged in which models with latent states, such as trend inflation, are estimated by explaining survey expectations and macroeconomic data jointly.[5] Relative to these studies, we innovate by jointly explaining macroeconomic data and expectations series in a nonlinear state space model with time-varying parameters. To the best of our knowledge, the only other paper to do so is Mertens and Nason (2018). They apply the sticky information framework to model US inflation and inflation expectations jointly in a time-varying parameter model which resembles our "SPF expectations model".[6] Relative to their work, we further decompose the inflation gap into components related to economic slack and import price inflation, such that we can zoom in on the drivers of the inflation gap.

3 Data and estimation method

This section explains some details related to the estimation. We first describe the data, followed by an outline of the Bayesian estimation method and, finally, a description of the priors.

[5] See, inter alia, Kozicki and Tinsley (2012), Crump et al. (2016), Nason and Smith (2016), and Winkelried (2017) for a state space model with constant parameters, and Grishchenko et al. (2017) for a state space model with constant parameters and stochastic volatility. For a model-consistent treatment of inflation expectations in a DSGE model, see e.g. Del Negro and Eusepi (2011), Smets et al. (2014) and Cui et al. (2015).

[6] In the sticky information framework, period t's expected inflation for period $t+h$ is a weighted average between the rational expectations forecast at time t and $\pi^e_{t+h_n|t-1}$, the previous period's inflation expectation *for the same end date.* Given that we use rolling event forecasts from the ECB's SPF, we cannot link $\pi^e_{t+h|t}$ to $\pi^e_{t+h|t-1}$ as in the sticky information framework, and use $\pi^e_{t-1+h_n|t-1}$ instead.

3.1 Data

Macroeconomic data Headline price inflation π_t is derived from the euro area Harmonised Index of Consumer Prices (HICP). Throughout the paper, the rate of inflation is defined as the annualised quarter-on-quarter growth rate of the price index: $\pi_t = 400ln(P_t/P_{t-1})$, where P_t is the price index and $ln(.)$ is the natural logarithm. Following Matheson and Stavrev (2013), we define the relative price of imports as the import-price deflator relative to the gross domestic product (GDP) deflator. The unemployment rate u_t is the euro area civilian unemployment rate. These series are obtained from the ECB's Statistical Data Warehouse (SDW). We transform the monthly unemployment rate and price indices to a quarterly frequency by taking the three months average, and backdate all series using the historical Area Wide Model (AWM) database (see appendix A.1 for details).

Inflation expectations data Our inflation expectations data is taken from the ECB's Survey of Professional Forecasters (SPF). The SPF is a quarterly survey that is conducted in January, April, July and October since 1999. Each quarter, around 60 EU-based professional forecasters deliver their point estimates and density forecasts for HICP inflation (annual percent change), GDP growth (annual percent change) and the unemployment rate (percentage of the labour force).

The survey contains individual level and aggregate data for two types of forecasts: rolling horizon and calendar horizon forecasts. The rolling horizon forecasts comprise one-year and two-year ahead forecasts starting from the latest date for which data are available. As the survey participants are provided with the inflation rate for the previous month, this implies that, for example, the 2018Q1 one-year and two-year ahead forecasts refer to year-on-year inflation in the month *December* 2018 and 2019, the 2018Q2 survey refers to March 2019 and 2020, etc. The calendar horizon forecasts refer to year-on-year inflation in the current calendar year, the next calendar year, and five calendar years ahead.[7]

Our empirical setup follows the "noise" interpretation from Smets et al. (2014). Specifically, we treat the rolling horizon forecasts and five-year ahead calendar year forecast as noisy indicators of the model forecasts for year-on-year inflation 3, 7 and 19 quarters ahead. This implies that in our implementation the systems of equations (14) to (15) and (16) to (17) each consist of three equations which refer to these horizons. Concerning the expectations data $\pi^e_{t+h|t}$, we use the (self-computed) mean of the professional forecasters' aggregate probability distribution for inflation. This implies that the model forecast - a conditional mean expectation - is matched with mean expectation from the survey distribution.[8] The evolution of the three survey expected inflation series is shown in Figure 2.

[INSERT FIGURE 2 HERE]

Estimation sample Our sample starts in 1990Q1, about one decade before the start of the SPF data in 1999Q1 in order to let the model 'learn' the parameter values before the start of the SPF data, and ends in 2017Q4. Our aim is to compare how an econometrician's ex-post assessment of the lowflation period changes when survey data is incorporated in the information set. Hence, we use revised data instead of real-time data, because we are not primarily interested in maximizing the fit of the survey expectations series.

[7] Since 2010Q3 the 2 calendar years ahead forecast is also provided.

[8] The mean of the probability distribution is very close to the average of the reported point forecasts for the one-year and two-year ahead forecasts. For the five-year ahead forecast, however, the mean of the distribution is persistently below the average point forecast in the lowflation period by about 10 basis points.

3.2 Gibbs sampling algorithm for the baseline model with SPF data

We estimate the model using Bayesian methods. Since our baseline model builds on that of Chan et al. (2016), an MCMC procedure similar to that from their appendix can be applied. However, the inclusion of survey data into the extended models requires additional care. As shown in equations (14) to (15) and (16) to (17), the measurement equations become nonlinear functions of the parameters due to the model forecast function $f_h(\theta_{t-1}, Y^{t-1})$. Therefore, additional steps are required in order to sample from these models. We summarize the Gibbs sampling algorithm for the SPF expectations model below, and refer the reader to the appendix for further details. Starting with initial values for the state variables, the Gibbs sampling algorithm separates the unknown parameters into separate blocks and draws from their conditional posterior distributions as follows:

1. Sample the error variances $\sigma_u^2, ..., \sigma_\psi^2$ conditional on all other parameters.
   - Following Chan et al. (2016), the error variances σ_ρ^2 and σ_λ^2 of the bounded states ρ_t^π and λ_t are both drawn with an independent Metropolis Hastings sampler.
   - The remaining variances are drawn independently from inverse Gamma distributions.
2. Sample the persistence parameters ρ_1^u, ρ_2^u conditional on all other parameters.
   - In this step we have to take into account that i) the parameters obey stationarity constraints, and that ii) they affect the likelihood terms in measurement equations (14) to (15) relating to the survey expectations data. We draw the persistence parameters using an independent Metropolis-Hastings step where the proposal distribution is based on the posterior distribution from an approximated model where the nonlinear model forecast functions $f_{h_1}\left(\theta_{t-1}, Y^{t-1}\right), ...,$ $f_{h_n}\left(\theta_{t-1}, Y^{t-1}\right)$ are linearized.
3. Sample the time-varying trends $\tau_t^\pi, \tau_t^u, \tau_t^m$ for $t = 1, ..., T$ conditional on the time-varying coefficients and error variances.
   - Conditional on the time-varying coefficients $\rho_t^\pi, \lambda_t, \gamma_t, \xi_t$ and the persistence parameters ρ_1^u and ρ_2^u it turns out that the model forecast functions $f_{h_1}\left(\theta_t, Y^t\right), ..., f_{h_n}\left(\theta_t, Y^t\right)$ are linear functions of the time-varying trends $\tau_t^\pi, \tau_t^u, \tau_t^m$. Hence, the model can be cast in a linear state space form, and the trends can be drawn with the Carter and Kohn (1994) algorithm.
4. Sample the time-varying coefficients $\rho_t^\pi, \lambda_t, \gamma_t$ and ξ_t for $t = 1, ..., T$ conditional on the time-varying trends and other parameters
   - This block involves two complications: *i*) these parameters enter nonlinearly in the model forecast functions $f_{h_1}\left(\theta_{t-1}, Y^{t-1}\right), ..., f_{h_n}\left(\theta_{t-1}, Y^{t-1}\right)$ of the equations of the survey expectations, and *ii*) the parameters ρ_t^π and λ_t are bounded to lie within certain intervals. To accommodate both features, we implement a single-move sampler (see Cogley, 2005; Koop and Potter, 2011), where for each period $t = j$ the time-varying coefficients are drawn conditional on the values for these coefficients from periods $t \neq j$, in addition to all other model parameters, using an independent Metropolis-Hastings sampler.
5. Sample stochastic volatility ψ_t for $t = 1, ..., T$ conditional on the state variables and other parameters

- The stochastic volatility terms ψ_t of the error term in the measurement equation (1) for inflation are drawn using the single-move sampler of Jacquier et al. (1994).

6. Go back to step 1 until the required number of draws has been reached.

We executed 250,000 replications of the Gibbs sampler and discarded the first 50,000. Finally, we stored every 20th draw in order to break the autocorrelation and economize on storage size. This leaves us with 10,000 posterior draws. To assess convergence, we inspected the recursive means of the retained draws at every 20th draw. The fact that there is little evidence of large fluctuations in the posterior means is taken as evidence in favour of convergence (see appendix A.5.).

3.3 Priors

In their analysis on U.S. data, Chan et al. (2016) use relatively uninformative priors which favour smooth transitions for the time-varying parameters. We closely follow their choices in our analysis. The priors of the error variances are independently distributed as $\sigma_i^2 \sim IG\left(\underline{\nu}_i, \underline{S}_i\right)$ for $i \in (u, m, \tau\pi, \tau u, \tau m, \rho, \lambda, \gamma, \psi, h_1, ..., h_n)$, where $IG\left(.,.\right)$ denotes the inverse-Gamma distribution. As in Chan et al. (2016), we set the degrees of freedom to a small value in order to make them relatively uninformative: $\underline{\nu}_i$=10 for $i \in$ (u, m, $\tau\pi$, τu, τm, ρ, λ, γ, ψ). The scale parameters $\underline{S}_{\tau\pi}$, $\underline{S}_{\tau u}$, $\underline{S}_{\tau m}$ are set to 0.09, such that $E\left(\sigma_{\tau\pi}^2\right)$=$E\left(\sigma_{\tau u}^2\right)$=$E\left(\sigma_{\tau m}^2\right)$=0.01. We deviate here from Chan et al. (2016), who set $\underline{S}_{\tau\pi}$ to 0.18 instead. Our prior mean implies that difference between two adjacent state observations, e.g. $\tau_t^\pi - \tau_{t-1}^\pi$, lies with about 95% probability between -0.2 and 0.2. The scale parameters for the time-varying coefficients, $\underline{S}_\rho$, $\underline{S}_\lambda$, $\underline{S}_\gamma$, and $\underline{S}_\xi$ are set to 0.018, which imply prior means $E\left(\sigma_\rho^2\right)$=$E\left(\sigma_\lambda^2\right)$=$E\left(\sigma_\gamma^2\right)$=$E\left(\sigma_\xi^2\right)$=0.002. This prior favours smoothly transitioning states, as the change between two periods has about 95% probability of lying between -0.09 and 0.09. For the stochastic volatility process and the unemployment gap equation, the scale parameters are set to $\underline{S}_u$=$\underline{S}_\psi$= 0.9, which implies a prior mean of $E\left(\sigma_u^2\right) = E\left(\sigma_\psi^2\right) = 0.1$. Concerning the measurement equations for the survey data, we set smaller values for the degrees of freedom: $\underline{\nu}_{h_1}$=... =$\underline{\nu}_{h_n}$=1.5, and scale parameters $\underline{S}_{h_1}$=...=$\underline{S}_{h_n}$=0.15. These values also imply prior means of $E\left(\sigma_{h_1}^2\right) = ... = E\left(\sigma_{h_n}^2\right) = 0.1$, but the variances are larger. Given the high volatility of import price inflation, we raise the scale parameter to $\underline{S}_m = 9$, such that $E\left(\sigma_m^2\right) = 1$. The persistence parameters parameters of the unemployment gap are normally distributed: $(\rho_1^u, \rho_2^u)' \sim N\left((1.8, -0.8)', 5I_2\right)$.

The initial values of the state variables are normally distributed with large variances: $\tau_0 \sim N\left(3.5, 10\right)$, $(\tau_0^u, \tau_{-1}^u)' \sim N\left(9I_2, 10I_2\right)$, $\tau_0^m \sim N\left(0, 10\right)$, $\rho_0^\pi \sim N\left(0, 1\right)$, $\lambda_0 \sim N\left(0, 1\right)$, $\gamma_0 \sim N\left(0, 10\right)$, $\psi_0 \sim N\left(0, 5\right)$, and $\xi_0 \sim N\left(0.7, 1\right)$.

4 Full sample results

This section describes the posterior estimates from our three models. It compares the results from the baseline model, estimated without survey expectations, against those from the two other specifications that incorporate expectations data. We discuss in turn the estimated time-varying trends, the time-varying coefficients, the error variances, and a historical decomposition of inflation.

4.1 Time-varying trends

Trend inflation τ_t^π Figure 3 shows the posterior median and 68% credible sets of the trend inflation τ_t^π estimates from the three models. The results from the baseline model (blue shaded area) are compared against those from the SPF expectations model in panel (a), and against those from the SPF no smoothing model in panel (b) (red, with dashed lines). The plots show that in all three cases τ_t^π trends downward from levels around 3% or higher in 1990, to about 2% by 1999 - the starting date of the European Monetary Union. Trend inflation estimates are subsequently stable around 2% between 1999 and 2011, after which they decline. The end-of-sample decline is the strongest in the baseline model, which delivers a final trend inflation estimate around 1.4%, whereas in the models with survey data the decline is more muted, and trend inflation lies around 1.7% at the end of the sample. Note also that the uncertainty bands are remarkably narrower for the models with survey data, especially for the SPF no smoothing model (panel b). As a result, the decline between 2011Q4 and 2017Q4 is only a posteriori significant at the 95% level for the SPF no smoothing model, even though it is significant at the less restrictive 68% level for all three models.

[INSERT FIGURE 3 HERE]

Natural rate of unemployment τ_t^u Estimates of the natural rate of unemployment τ_t^u from the baseline model show a marginal decline from 9.2% to 8.8% over the sample period (Figure 4). As a result, the unemployment gap - a measure of economic slack - was closed at the end of the sample according to baseline model. By contrast, the decline is more pronounced for the models with survey data. These models find a larger unemployment gap over the lowflation period, and that slack remained in the economy at the end of the sample. This is especially the case in the model without forecast smoothing, which finds an end-of-sample unemployment gap of 0.7%. The difference in τ_t^u values between the first (1990Q1) and last observations (2017Q4) is only a posteriori significant at the 95% level for the SPF no smoothing model.

[INSERT FIGURE 4 HERE]

Trend relative import price inflation τ_t^m We include a trend τ_t^m in order to capture any remaining slowly moving trend effects to relative import price inflation π_t^m. Figure 5 shows that the trend estimates are very similar across models and that, with the exception of the pre-1999 period, the trends move broadly sideways at levels around zero.

In sum, the extended models with survey expectations data point toward temporary, rather than permanent factors, as the main drivers behind low inflation in the euro area. In particular, these two models find a more muted decline of trend inflation - inflation's permanent component - and a larger degree of economic slack - in the form of a positive unemployment gap - over the lowflation period. In the next subsection we further explore the evolution of the temporary factors that influence the inflation gap.

[INSERT FIGURE 5 HERE]

4.2 Time-varying coefficients

Figure 6 shows the posterior estimates of the time-varying inflation gap persistence ρ_t^π, the Phillips curve slope λ_t, and the import price coefficient γ_t, and compares the posterior estimates from the baseline model against those from the models with survey data.

[INSERT FIGURE 6 HERE]

Intrinsic inflation gap persistence ρ_t^π In the baseline model, intrinsic inflation gap persistence ρ_t^π is found to trend slightly downward from 0.17 in 1990 to about 0.1 in 2000, and stabilizes thereafter. However, both models with survey data find a short-lived spike in inflation gap persistence which peaks in the 2008Q4-2009Q1 quarters. As we discuss below, this timing coincides with shifts in the one-year and two-year ahead SPF forecasts, which are also reflected in ξ_t changes. The SPF model without forecast smoothing finds higher intrinsic persistence at the start of the sample, which is related to the fact that this model also reports the lowest initial level of trend inflation (around 3%) and, therefore, a larger and more persistent inflation gap.

Phillips curve slope λ_t To facilitate comparison among coefficients, Figure 6 shows the Phillips curve slope as $-1 \times \lambda_t$. Hence, a decrease toward zero indicates a weaker impact and, therefore, a Phillips curve flattening, and an increase represents a Phillips curve steepening. Across all models, the Phillips curve slope estimates follow the same broad tendencies. First, a general slope flattening between 1990 and 2012 despite several short-lived steepening interruptions. Second, a sharp steepening which starts around the end of 2012 and, finally, a stabilization during the final years at values close to 0.3. The slope estimates from the SPF model without smoothing show several short-lived fluctuations, but follow the same broad trends as the SPF model with forecast smoothing. In particular, the end-of-sample steepening is the strongest in these models with survey data. Combined with the finding of economic slack in the post-2013 period, this steeper slope partly attributes low inflation to a stronger effect of economic slack, which remained present despite the continuous decline of the unemployment rate. This steepening is a posteriori significant at the 95% level in the SPF no smoothing model, and at the 68% level in both models with survey data.[9]

A second effect of the Phillips curve slope steepening, is that the (inherited) persistence of the inflation gap increases. Indeed, the euro area unemployment gap is a highly persistent process, where the sum of the autoregressive persistence parameters is close to one (Table 1).

Import price coefficient γ_t Estimates of the import price coefficient are very similar across all models.[10] The import price coefficients decline during the period from 1990 to 2000, but then follow an upward trend which ultimately stabilizes at the end of the sample.[11] The increase between 2000Q1 and 2017Q4 is a posteriori significant at the 95% level in all three models. Despite the apparently small magnitude of γ_t, import prices still have a meaningful impact due to the relatively large volatility of shocks to the import price gap. As we will show below, the general weakness of commodity prices during the post-2013 period, combined with this stronger impact coefficient, resulted in a important downward effect on inflation in recent years.

[9] Using different model specifications and estimation methods, Riggi and Venditti (2015), Dany-Knedlik and Höltemoller (2017) and Cordemans and Wauters (2018) also report a slope steepening of the euro area Phillips curve in the post-crisis period. Bulligan and Viviano (2017) find a steeper *wage* Phillips curve slope for Italy, Spain and France after the global financial crisis.

[10] The strong similarity across model estimates in γ_t and τ_t^m is possibly due to our simplifying assumption that the import price inflation gap is an i.i.d. process. This assumption implies that a shock to today's import price gap affects future inflation via the impact on today's inflation gap, and not via spillovers to the future import price gap. As a result, including the survey data on inflation expectations delivers no additional information for estimating γ_t. A robustness check which models the import price gap as an AR(2) process is on our to-do list.

[11] The overall evolution of γ_t resembles that from its counterpart in Dany-Knedlik and Höltemoller (2017) and Cordemans and Wauters (2018).

Smoothing parameter ξ_t Panel (a) of Figure 7 shows the evolution of the smoothing parameter ξ_t from the SPF expectations model. The pre-1999 period, for which no SPF data is available, is marked in grey. From 1999 onward, the smoothing coefficient ξ_t remains stable above 0.8 and then declines swiftly in the 2007-2008 period. ξ_t subsequently follows an upward trend, which is interrupted only shortly by a drop at the end of 2014. Thus, the evolution points to a lower tendency for forecast smoothing during the financial crisis period, but overall ξ_t remains at elevated levels above 0.6. These results resonate with Łyziak and Paloviita (2018) who, in a different setup, also report high smoothing coefficients for the euro area that also decline during the crisis period.[12]

Intuitively, the role of ξ_t in our SPF model is to allow for a gradual adjustment of survey expectations to changes in the underlying model forecast. Drops in the level of ξ_t would thus suggest a higher degree of comovement between the survey expectations and the underlying model forecast. To support this idea, panel (b) of Figure 7 shows the recursively estimated values of ξ_t, along with the evolution of the SPF one-year and two-year ahead expectations.[13] Compared to the full-sample estimates from panel (a), the median estimate of ξ_t shows two periods of more pronounced drops: the first in 2008Q4, and the second in 2015Q1. The plot shows that these declines coincide with periods wherein the SPF one-year and two-year forecasts experienced sharp changes.[14]

[INSERT FIGURE 7 HERE]

4.3 Error variances

Figure 8 shows the evolution of the time-varying standard deviation of ϵ_t^π - the error in the measurement equation (1) for inflation. The volatility of residual shocks to inflation trended upward in the period from about 1995 to 2008, but has been in decline since. The estimates from the all three models, as shown by the medians, are very similar to one another.

Table 1 shows the summary statistics (median and 16th and 84th percentiles) from the posterior distributions of the error variances that remain constant over time. Overall, the results are mostly similar across the three models. The variance of the residual in the import price equation, σ_m^2, is by far the largest due to the volatile nature of import price inflation. The variance of the residuals of the survey expectations data, shown as σ_{1y}^2, σ_{2y}^2 and σ_{5y}^2 for one-year ahead, two-year ahead and five-year ahead expectations, are smaller than those from the other measurement equations.

The degree of time-variation in the trends and coefficients is limited and does not deviate much from the prior mean. Of all time-varying parameters, the variation is the largest in the stochastic volatility series (σ_ψ^2), followed by trend inflation ($\sigma_{\tau\pi}^2$) and trend import price inflation ($\sigma_{\tau m}^2$).

[INSERT FIGURE 8 HERE]

[12] Łyziak and Paloviita (2018) estimate how the ECB's projections (instead of an underlying model forecast) spill over into the SPF expectations.

[13] The results in panel (b) reflect the estimated ξ_t using data up to period t, whereas those from panel (a) use the full sample. The recursive estimates are a byproduct of the forecasting exercise described in the next section.

[14] In the plot, the SPF data is shifted back one quarter in order to align it with the timing in the empirical model. Given this timing, the recursively estimated ξ_t drops strongly in 2008Q4, which is the quarter where the one-year and two-year ahead forecasts record their largest period-by-period changes. ξ_t also drops remarkably in 2015Q1, a period where the same SPF series also experience large changes.

4.4 Historical decomposition

To summarize the overall impact of these estimates to the dynamics of inflation, Figure 9 shows a historical decomposition of year-on-year headline inflation over the 1999-2017 period from the SPF expectations model.[15] Since trend inflation is generally close to 2% and only declines weakly in the post-2013 period, it does not explain the lowflation period. Instead, import prices had an important role in the inflation cycles after 2007. In the post-2009 period, economic slack had a protracted downward drag on inflation, which reverts to zero by the end of the sample. According to a back-of-the-envelope calculation, the variation in import prices accounted for 55% of the variation in inflation since 2007, while the shares are 21% for slack, 16% for the residual and 8% for trend inflation.

[INSERT FIGURE 9 HERE]

5 Is the SPF data helpful for forecasting?

We now evaluate and contrast the forecast performance of our three model specifications for forecasting future inflation at different horizons. This exercise serves two purposes. First, it provides an assessment of the value of SPF data in improving the forecast accuracy of our baseline model. Second, it offers the applied econometrician with an interest in forecasting a model evaluation device to discriminate among our different measures of trend inflation and cyclical inflation drivers. In fact, as surveys send more accurate signals about future inflation, they can be regarded as providing more valuable information about current inflation dynamics as well.

We perform a pseudo-out-of sample forecast exercise for the time span 1990Q1 to 2017Q4. We use an expanding window, re-estimating the model every quarter.[16] We evaluate the accuracy of both point and density forecasts, using root mean squared errors (RMSE) and cumulative log predictive densities (CLPD), respectively. The forecast horizons include the one quarter ahead as well as the one, two and five year ahead horizons. The one through five years ahead forecasts refer to expected year-on-year inflation rates; matching the SPF's rolling forecast horizon. We use iterated forecasts calculated with predictive simulation. More specifically, when forecasting using information through time t, predictive simulation is done of future values of the latent states and of the dependent variables. In the following, we first discuss aggregate forecast statistics, evaluating the average forecast performance over the forecast sample. Subsequently, we investigate the evolution of the relative forecast performance of our different model specifications over time.

5.1 Average forecast accuracy

Figure 10 displays the accuracy of point and density forecasts, as ratios of RMSEs of the extended model variants including survey forecasts relative to the baseline specification, estimated without survey expectations, (panel a) and as differences in CLPDs relative to the baseline model (panel b). The figure also contains the forecast statistics of two robustness checks (presented by the grey striped and dotted bars); these are discussed in Section 6.

[15] This decomposition of year-on-year inflation takes into account the time variation in all parameters. It is calculated as a four quarter moving average of the decomposition of quarter-on-quarter inflation. This transformation, along with the effects of ρ_t^π, explains the apparent persistence of the residuals in the figure.

[16] A caveat is that the results are based on 40,000 draws from the posterior sampler, where we discard the first 20,000 and keep every 10th draw in order to obtain 2,000 retained draws.

[INSERT FIGURE 10 HERE]

Overall, we find that accounting for survey data improves on the accuracy of our baseline model. Using CLPDs as the forecast metric, our model variants that include survey expectations (i.e., the SPF and SPF no smoothing models) forecast better than the baseline model without surveys at all forecast horizons. Using RMSEs, the SPF model variants forecast best at the short to medium horizons, whereas they perform about equally well as the baseline model at the 5 year forecast horizon.

The relative performance of our two model variants that incorporate expectations data is more mixed. The CLPD measure favours the model variant with smoothing at short horizons, but the variant without smoothing at the long horizon; although some caution is required in drawing strong conclusions from the long-term forecast statistics in that the number of fully independent observations is limited in the available data sample. The point forecasts of both models are more or less equally accurate: the differences between their RMSE statistics are small to modest at all forecast horizons. As a result, differences in forecast performance mainly relate to differences in higher moments of the predictive distribution.

In sum, we find that including information from survey forecasts helps to improve inflation forecasts, especially at short horizons. For the practitioner interested in inflation predictability, this finding supports the inclusion of survey data in model-based assessments of inflation dynamics. Our forecast evaluation exercise therefore provides confidence that the lowflation is mainly driven by cyclical, and thus temporary forces, rather than by downward revisions in long-term inflation expectations.

5.2 Relative forecast performance over time

It is instructive to investigate whether and how the value of the SPF data in predicting future inflation has evolved over time. To this end, we look to the evolution of the cumulative sums of the log predictive densities over the forecast sample, see Figure 11. We subtract the cumulative sums for the baseline model from that of our two model variants including survey forecasts (i.e., the SPF and SPF no smoothing models). Hence, declining cumulative sums indicate better performance of the baseline model in that period, and vice versa for upward sloping lines.

Two observations stand out. First, it appears that the better forecast accuracy of the SPF models manifests itself mainly after the outbreak of the global financial crisis in 2007. This is particularly true at short forecast horizons, in which case the baseline model and the two SPF model variants tend to have comparable accuracy prior to the crisis. Second, we see that the relative accuracy of the baseline model temporarily improved during the two most recent bouts of falling inflation: first in 2008, in the immediate aftermath of the financial crisis, then during the first one to two years of the lowflation episode since 2013. This result applies in particular to the SPF model without forecast smoothing and suggests that these drops in inflation came more as a surprise to the SPF than to our baseline model forecasts. Interestingly, the points in time where the relative accuracy of the SPF models starts to prevail again seem to broadly correspond with those episodes during which the forecast smoothing coefficient ξ_t drops (see Figure 7). This coincident co-movement may indicate that forecasters started to update their beliefs more actively as it became clear that they were doing an increasingly worse job in predicting inflation. As such beliefs are typically not only based on past events but also take account of future expected development, it might be expected that a drop in ξ_t leads to forecasts conveying more valuable information about future inflation dynamics.

6 Robustness analysis

Our strategy of modelling the expectations term structure in a time-varying parameter model invokes complications relative to the typical approach in the literature, which uses either only long-term survey expectations or constant parameters for the transmission coefficients in the inflation gap. We therefore perform robustness checks to evaluate if the results remain the same when we simplify the SPF model to *i)* using long-term expectations data only, or *ii)* having constant transmission coefficients. Concretely, in the first case, we take out all measurements equation from the system (14) to (15), except for the one related to the five-year ahead SPF inflation forecast. In the second robustness check, we estimate a model variant with constant Phillips curve coefficients (ρ^{π}, λ and γ), a time invariant smoothing coefficient (ξ) and static variances of the inflation gap residual (i.e., $\epsilon_t^{\pi} \sim N(0, \sigma_{\pi}^2)$).

First, we investigate the impact of the proposed perturbations on the estimated dynamics of the smoothing coefficient ξ_t, trend inflation τ_t^{π} and the natural rate of unemployment τ_t^{u}. We then compare the forecast performance of the SPF model with smoothing to that of its alternative variants. For ease of exposition, we label the constant parameter version of the model as the "CST model" and refer to the model variant excluding information on short-term inflation expectations as the "LT-exp model"

Robustness smoothing coefficient ξ_t We find a median estimate of $\xi = 0.81$ in our constant parameter estimation (Table 2). This value accords well with the average degree of forecast smoothing recorded over the sample in the SPF model with time variation. Interestingly, in the model using only long-term SPF forecasts, ξ_t is found to be broadly stable at a level just above 0.9 (not shown). The observed decline in ξ_t in our baseline estimate during the financial crisis period is therefore fully driven by developments in short-term inflation expectations. This is not surprising, since the short-end of the expectations curve experienced sharp changes consistent with Phillips curve predictions in recent years, whereas the long-end remained relatively more stable (see Figure 2).

[INSERT TABLE 2 HERE]

Robustness trend inflation τ_t^{π} and natural rate of unemployment τ_t^{u} Figure 12 displays the posterior median estimates of trend inflation and the natural rate of unemployment of the full fledged SPF model and its two robustness variants. For ease of analysis, the figure also plots the respective state estimates of the baseline model (i.e., the model specification without survey data). Of the four model variants, the SPF model finds the lowest value for the natural rate of unemployment, specifically at the end of the sample. Hence, compared to the SPF model, the model variants with either constant parameters or abstracting from short-term forecast data, attribute less weight to domestic slack in explaining the low inflation in the euro area.

[INSERT FIGURE 12 HERE]

The two robustness variants differ, however, in how they compensate for this more muted slack channel to fitting the low inflation. The model variant without short-term forecast data points to a stronger weakening in trend inflation in recent years, when compared to the SPF model that includes both short and long-term forecast. At first sight, this result might seem counter-intuitive, given the higher and more stable course of long-term relative to shorter-term inflation expectations and the fact that -in isolation- long-term expectations mainly inform about the trend. However, compared to the SPF model, the relatively high ξ_t estimate in LT-exp model, implies a much slower adjustment of survey expectations to changes in underlying model forecasts. This mechanism allows long-term survey forecasts to persistently overshoot trend inflation. Finally, note that the LT-exp model, focusing on long-term forecast data only, and the baseline model, neglecting the information of survey data all together, deliver similar estimates for both trend inflation and the natural rate of unemployment. Their difference with the SPF model thus relates to both models abstracting from information on short-term inflation forecasts. This observation suggests that short-horizon survey expectations may play an important role in measuring inflation's trend and cyclical factors.

The median trend inflation in the CST model is comparable to that obtained in the SPF model during most of the protracted period of below target inflation since 2013. Historical decomposition exercises (not shown) learn that the CST model attributes the more moderate contribution of domestic slack on the lowflation, relative to the SPF model, to mainly a stronger contribution from negative inflation shocks, but also to lower upward price pressures coming from foreign factors towards the end of the sample. These results relate to the inability of the CST model to fully account for the decline in volatility of the inflation gap residual as well as the increase in the impact of import price inflation as observed in the SPF model in recent years (see Figures 8 and panel (e) of Figure 6, respectively). In the SPF model the import price coefficient and the standard deviation of residual shocks to inflation averaged around $\bar{\gamma}_{2013-17} = 0.19$ and $\bar{\sigma}_{\pi,2013-17} = 0.47$ respectively since 2013, whereas their respective median values note at $\gamma = 0.15$ and $\sigma_\pi = 0.76$ in the constant parameter version of the model (see Table 2).

Forecast performance robustness specifications The full sample forecast statistics of our two robustness model specifications for forecasting inflation are presented by the grey striped and dotted bars in Figure 10. Similar to the approach in Section 5, we choose the benchmark model to be the baseline model excluding survey data.

Looking at the accuracy of point forecasts first, the RMSE indicates that the SPF model with forecast smoothing performs at least as well as its robustness variants with either constant parameters or abstracting from short-term forecast data. At short and medium horizons, the SPF model exhibits the best forecast performance. At the 5 year forecast horizon, the two robustness variants are comparable in accuracy to the SPF model. The picture is mostly similar when accounting for uncertainty around point forecasts. In fact, the CLPD statistics favour the SPF model in forecasting inflation at short to medium horizons. However, at the longer horizon of 5 year ahead, the model variant with constant parameters offers the best density forecasts; which follows from the fact that the CST model's forecasts are not subject to parameter uncertainty.

Overall, these findings support the view that including information from short-term survey forecasts and adding time-variation in parameters is useful in predicting future inflation. The fact that the fully fledged SPF model tends to forecast well relative to its variant only incorporating LT inflation expectations also suggests that short-term survey forecasts carry important information in addition to long-term expectations data about inflation's trend and cyclical factors; a point made earlier by Kozicki and Tinsley (2012).

7 Conclusion

This paper explores the role of survey inflation expectations in uncovering the relative importance of permanent effects and cyclical forces in explaining the protracted disinflation period experienced in the euro area since late 2012. To this end, we estimate - employing Bayesian Gibbs sampling techniques - a time-varying parameter Phillips curve model that jointly explains macroeconomic data and inflation expectations from the ECB's Survey of Professional Forecasters. To take account of potential systematic forecast biases in the SPF, we allow for some type of "forecast smoothing" behaviour, in the sense that survey expectations can gradually respond to changes in the model forecast.

Our results support the view that the lowflation is mainly driven by cyclical, and thus temporary forces, rather than by downward revisions in long-term inflation expectations. More specifically, compared to the estimated dynamics of a model variant that abstracts from survey information, we find a more muted decline of trend inflation in recent years and a larger degree of economic slack when survey data are incorporated in the analysis. At the same time, inflation has become more sensitive to both cyclical fluctuations and foreign price pressures in recent years. Moreover, we find that the inclusion of survey data improves the forecast accuracy of our model in predicting future inflation. Hence, from the perspective of a practitioner who is concerned with inflation predictability, this finding supports the usefulness of survey data in understanding current inflation dynamics. As a byproduct, our analysis provides an estimate of the speed with which survey respondents adjust their forecasts to changing macro-economic developments. We find that forecasters updated their believes more frequently during the financial crisis period, compared to the years before, but that overall survey expectations remained very persistent.

In addition to inflation forecasts, the ECB's SPF also contains expectations for unemployment and GDP growth rates. An interesting topic for future research is to investigate whether this data conveys additional valuable information about inflation's underlying driving sources.

Tables and Figures

Table 1: Summary statistics of the posterior distribution of the error variances and persistence parameters

Parameter	Baseline model	SPF expectations model	SPF no smoothing model
σ^2_u	0.026	0.027	0.029
	(0.023, 0.030)	(0.023, 0.031)	(0.026, 0.034)
σ^2_m	1.614	1.613	1.634
	(1.424, 1.837)	(1.426, 1.837)	(1.440, 1.866)
$\sigma^2_{\tau\pi}$	0.012	0.009	0.004
	(0.008, 0.016)	(0.007, 0.012)	(0.004, 0.005)
$\sigma^2_{\tau u}$	0.005	0.005	0.005
	(0.004, 0.006)	(0.004 0.006)	(0.004, 0.007)
$\sigma^2_{\tau m}$	0.010	0.010	0.010
	(0.007, 0.014)	(0.007, 0.014)	(0.007, 0.014)
σ^2_ρ	0.002	0.002	0.002
	(0.001, 0.002)	(0.001, 0.002)	(0.001, 0.002)
σ^2_λ	0.002	0.002	0.002
	(0.001, 0.003)	(0.001, 0.003)	(0.001, 0.002)
σ^2_γ	0.002	0.002	0.002
	(0.001, 0.002)	(0.001, 0.002)	(0.001, 0.002)
σ^2_ψ	0.077	0.078	0.080
	(0.059, 0.103)	(0.059, 0.103)	(0.061, 0.107)
σ^2_ξ		0.002	
		(0.001, 0.003)	
σ^2_{1y}		0.014	0.010
		(0.011, 0.017)	(0.009, 0.013)
σ^2_{2y}		0.007	0.007
		(0.006, 0.008)	(0.006, 0.008)
σ^2_{5y}		0.006	0.006
		(0.005, 0.007)	(0.005, 0.008)
ρ^u_1	1.816	1.766	1.645
	(1.742, 1.888)	(1.713, 1.817)	(1.600, 1.690)
ρ^u_2	-0.836	-0.794	-0.685
	(-0.908, -0.762)	(-0.843, -0.742)	(-0.727, -0.642)

Note: The summary statistics shown are the posterior median and, between parentheses, the 16th and 84th percentiles of the posterior distribution. The parameters σ^2_{1y}, σ^2_{2y} and σ^2_{5y} denote, respectively, the error variances from the measurement equations for the one-year ahead, two-year ahead, and five-year ahead inflation expectations. The ρ^u_1 and ρ^u_2 parameters denote the autoregressive parameters in the unemployment gap equation.

Table 2: Summary statistics of the posterior distribution of selected parameters in the Robustness estimation

Parameter	CST model
ρ_1^u	1.771
	(1.710, 1.828)
ρ_2^u	-0.797
	(-0.852, -0.737)
λ	0.214
	(0.142, 0.311)
γ	0.152
	(0.135, 0.169)
ρ^π	0.151
	(0.067, 0.245)
ξ	0.806
	(0.723, 0.878)
σ_π^2	0.760
	(0.710, 0.816)

Note: The summary statistics shown are the posterior median and, between parentheses, the 16th and 84th percentiles of the posterior distribution. The table contains selected parameters from the CST model estimation from the robustness section.

Figure 1: Euro area headline HICP inflation

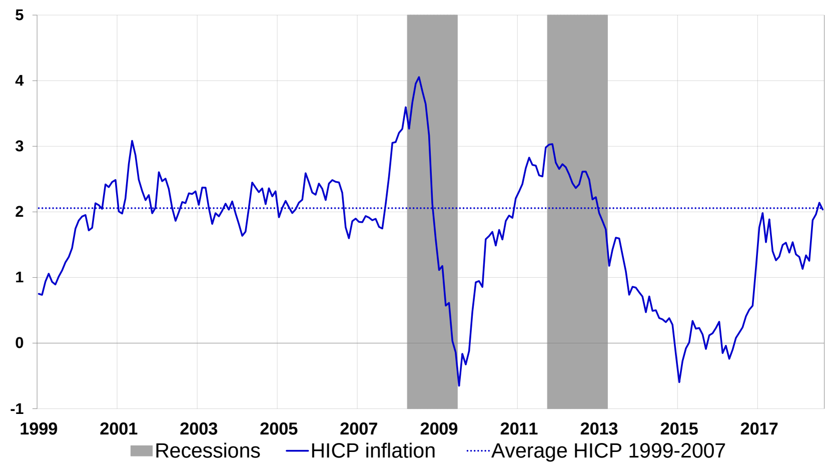

Note: Headline inflation is the year-on-year growth rate of the Harmonized Index of Consumer Prices (HICP; full line). The dotted line shows the pre-crisis (1999-2007) average of 2.06%. Sample: 1999m1 - 2018m8. Recession dates are taken from the CEPR Euro Area Business Cycle Dating Committee.

Figure 2: SPF inflation expectations data

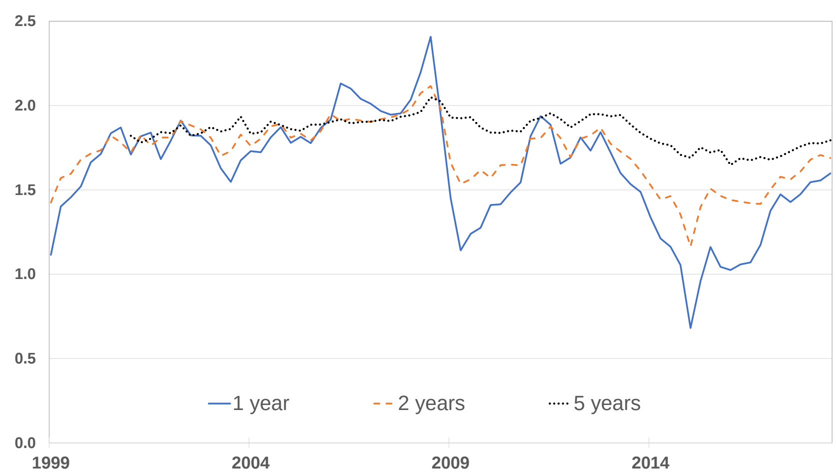

Note: The SPF expectations are the rolling horizon one-year and two-year ahead expectations and the five-year ahead calendar year inflation expectations. In all three cases, we report the computed mean from the aggregate probability distribution for year-on-year headline inflation. Sample: 1999Q1 - 2017Q4.

Figure 3: Inflation π_t and estimates of trend inflation τ_t^π

(a) Baseline and SPF models

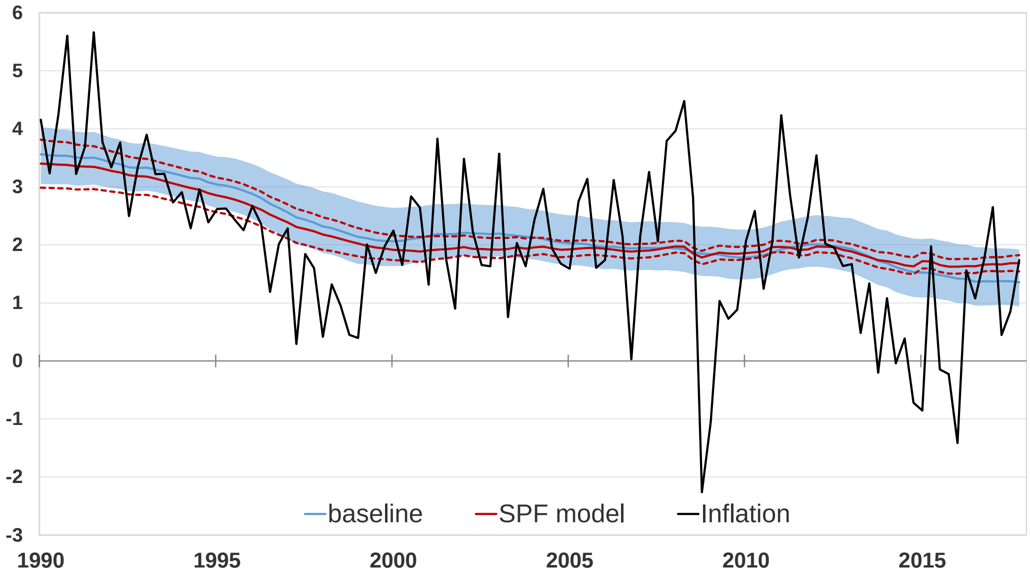

(b) Baseline and SPF no smoothing models

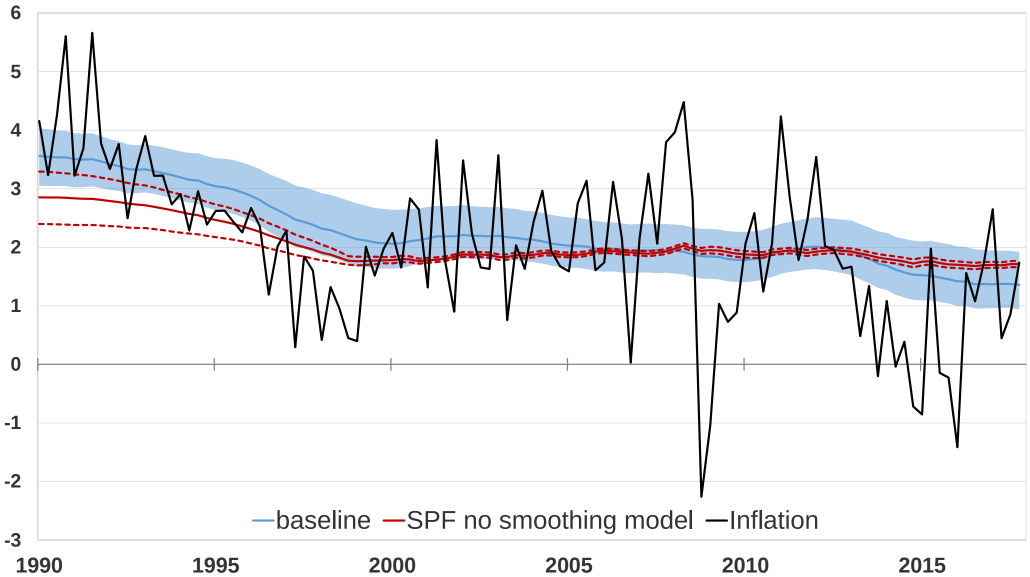

Note: The figures show headline inflation (black line) and estimates of trend inflation. Trend inflation estimates from the baseline model (blue shaded area) are compared against those from the SPF expectations model in panel (a), and against those from the SPF no smoothing model in panel (b) (both have bands with red dashed lines). The bands depict the median and 68% credible set.

Figure 4: Unemployment rate u_t and estimates of the natural rate of unemployment τ_t^u

(a) Baseline and SPF models

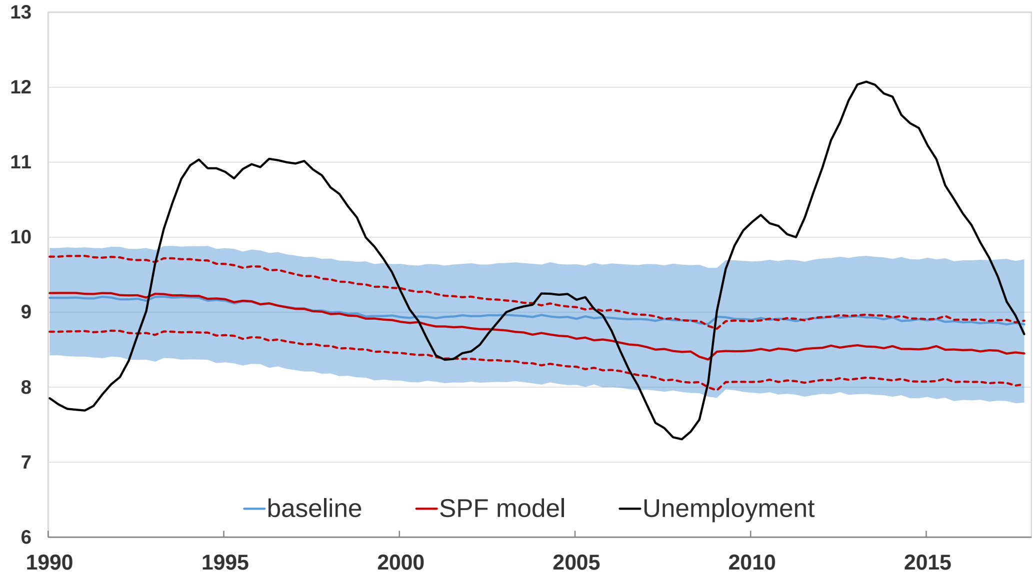

(b) Baseline and SPF no smoothing models

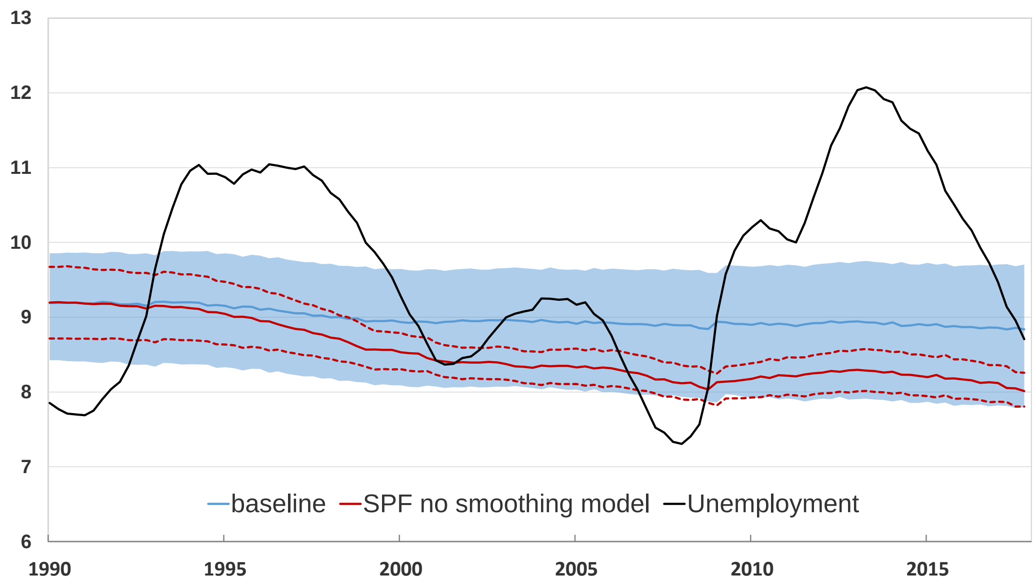

Note: The figures show the unemployment rate (black line) and estimates of the natural rate of unemployment. Natural rate estimates from the baseline model (blue shaded area) are compared against those from the SPF expectations model in panel (a), and against those from the SPF no smoothing model in panel (b) (both with red lines). The bands depict the median and 68% credible set.

Figure 5: Relative import price inflation π_t^m and its trend τ_t^m

(a) Baseline and SPF models

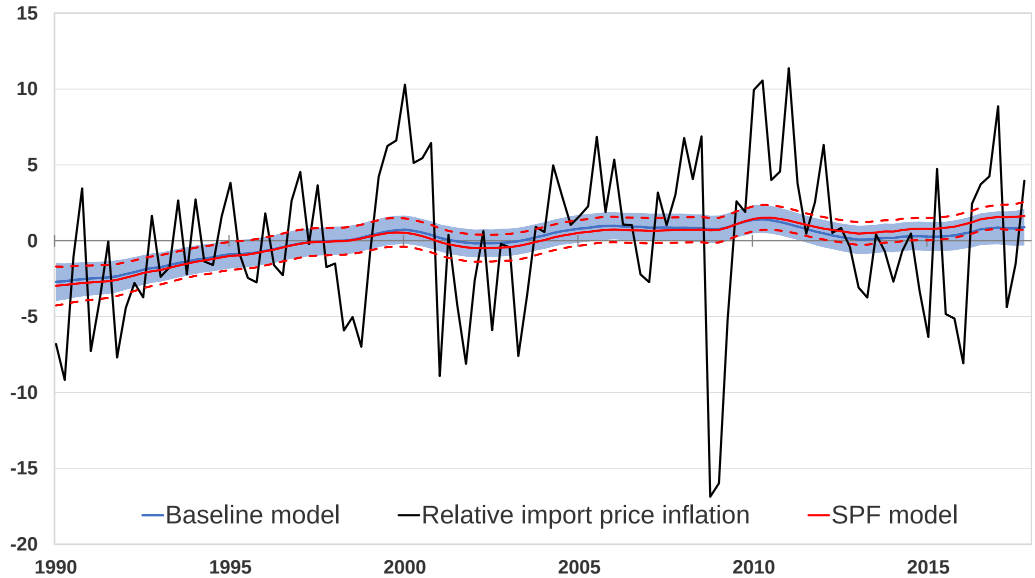

(b) Baseline and SPF no smoothing models

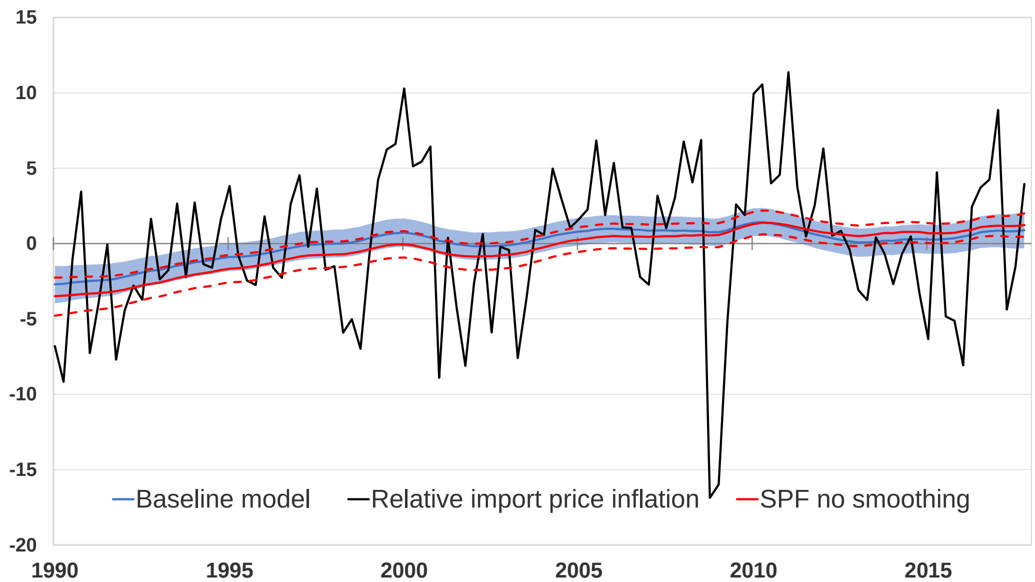

Note: The figures show the inflation rate of the relative price of imports (black line) and estimates of its trend. Trend estimates from the baseline model (blue shaded area) are compared against those from the SPF expectations model in panel (a), and against those from the SPF no smoothing model in panel (b) (both with red lines). The bands depict the median and 68% credible set.

Figure 6: Time-varying coefficients

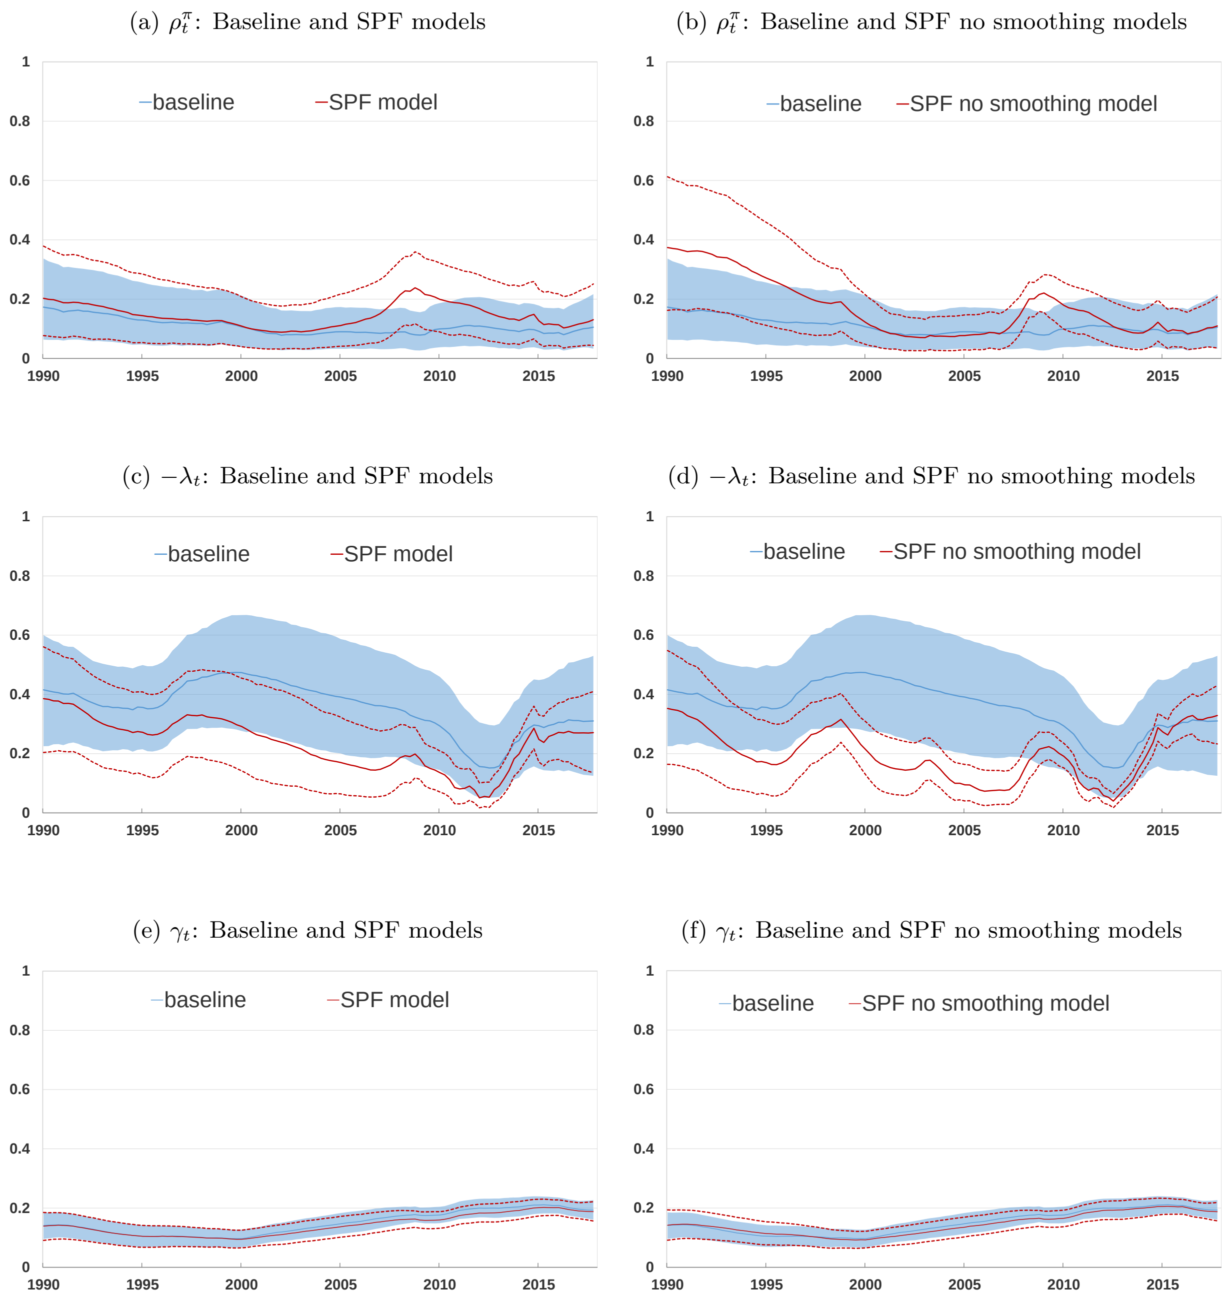

Note: The figures show the estimated time-varying inflation gap persistence ρ_t^π (panels a and b), Phillips curve slope λ_t (panels c and d), and the import price coefficient γ_t (panels e and f). The Phillips curve slope is shown as $-\lambda_t$ to facilitate comparison (an increase reflects a steepening of the slope). Estimates from the baseline model (blue shaded area) are compared against those from the SPF expectations model in the first column, and against those from the SPF no smoothing model in the second column (both with red lines). The bands depict the median and 68% credible set.

Figure 7: Estimated smoothing coefficient ξ_t

(a) Full sample estimates ξ_t (b) Recursively estimated ξ_t

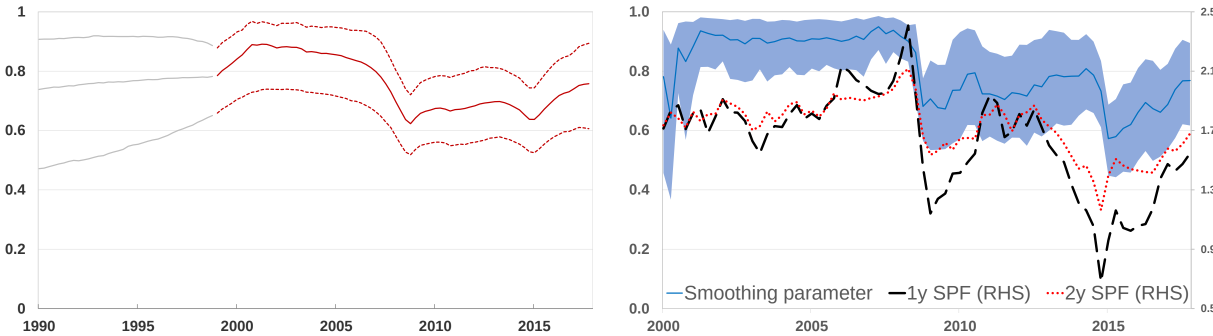

Note: Panel (a) shows the full sample estimates of the smoothing parameter ξ_t from the SPF expectations model. We mark the pre-1999 evolution of ξ_t in grey, because there is no SPF data available for this period. The bands depict the median and 68% credible set. Panel (b) shows the recursively estimated ξ_t values, along with the 1 and 2 year ahead SPF inflation expectations.

Figure 8: Time-varying stochastic volatility

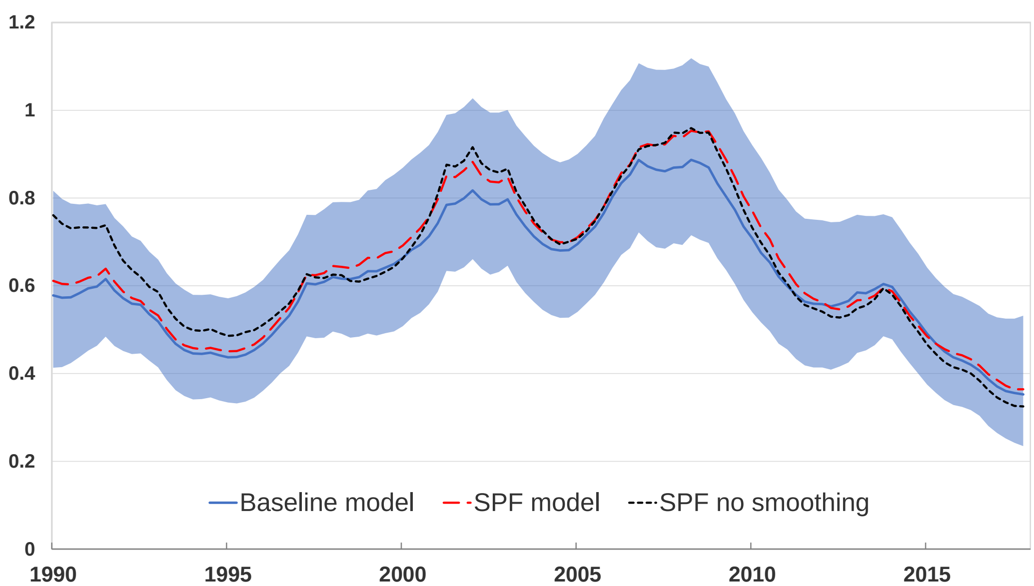

Note: Estimates for the time-varying standard deviation of ϵ_t^{π}, calculated as $e^{\psi_t/2}$, are shown for the baseline model (median and 68% credible sets), a well as the median estimates for the SPF expectations model and the SPF no smoothing model.

Figure 9: Historical decomposition of inflation with the SPF smoothing model

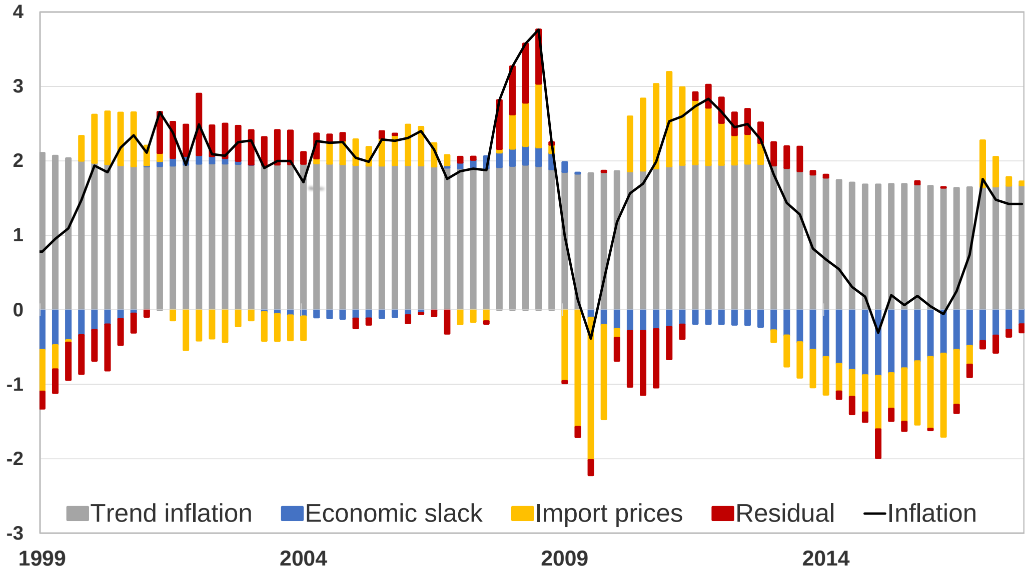

Note: The historical decomposition of headline inflation is calculated from the posterior mean of the time-varying parameters from the SPF expectations model.

Figure 10: Forecast performance statistics

(a) RMSE relative to baseline model

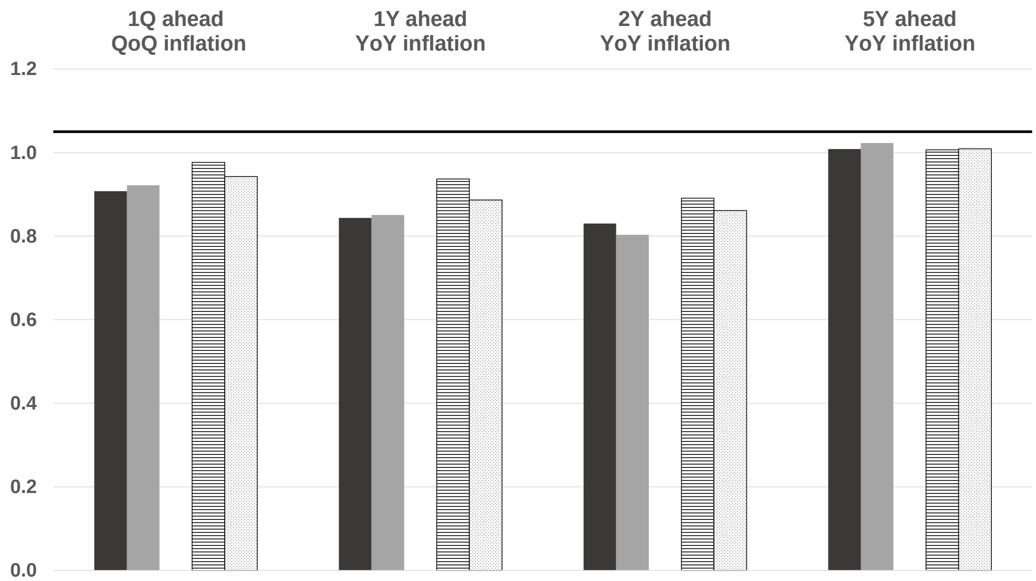

(b) CLPD relative to baseline model

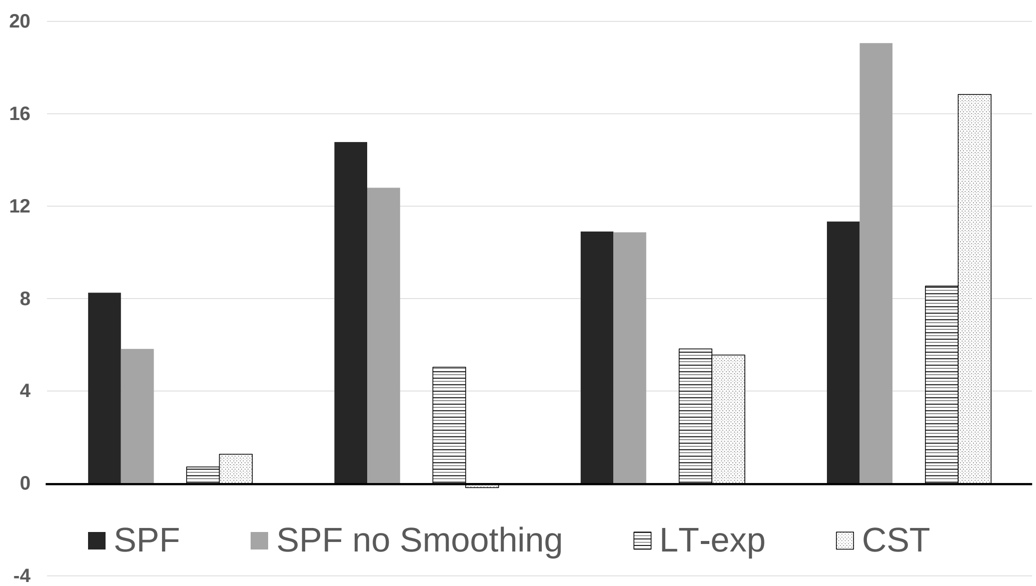

Note: The figures show the forecast performance relative to the baseline model. Panel (a) shows the root mean squared error (RMSE) results against the baseline model. Panel (b) shows the cumulative log-predictive density scores against the baseline model. A RMSE below 1 indicates better performance than the baseline model, and a CLPD above 0 indicates better density forecasting compared to the baseline model. The 'LT-exp' and 'CST' models refer to the robustness checks with *i)* only long-term expectations and *ii)* constant transmission parameters, respectively.

Figure 11: Forecasting comparison: CLPD over time

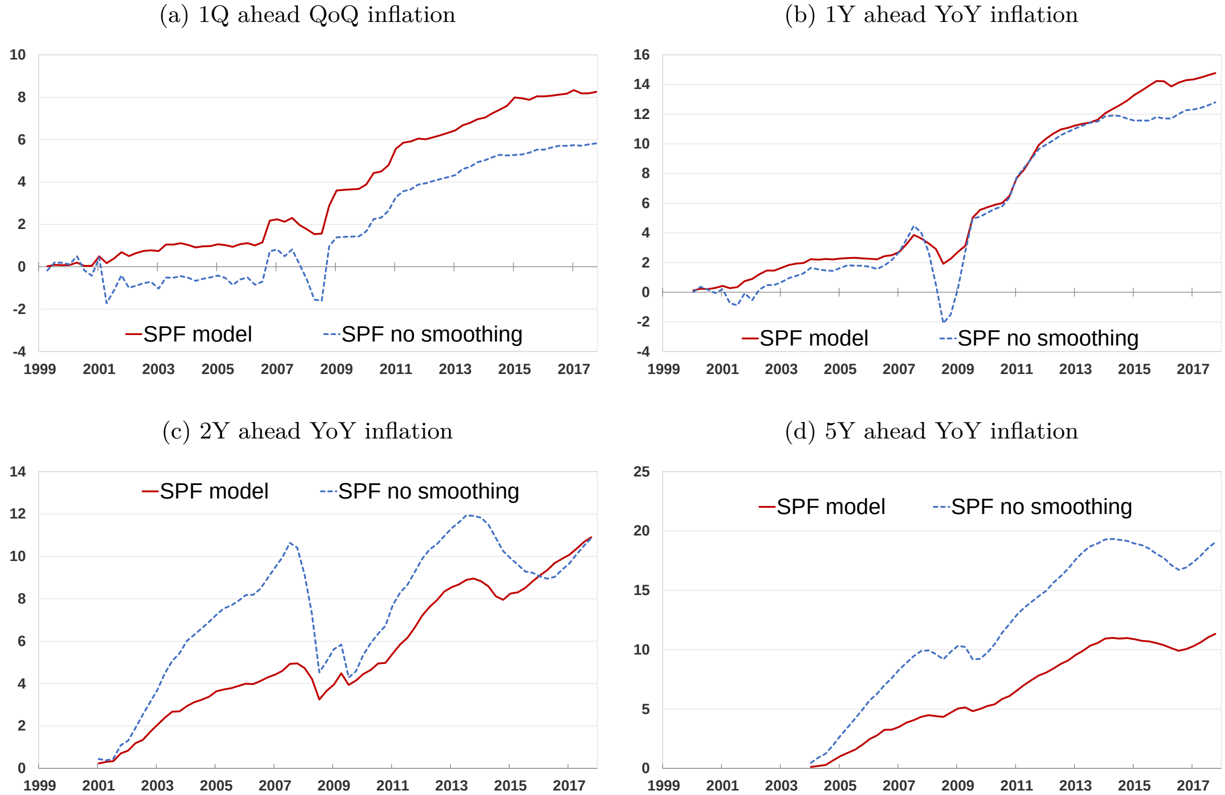

Note: The figures show the evolution of the cumulative log-predictive density score over time. The log-density score relative to the baseline model is shown in the full red line for the SPF model, and in the dashed blue line for the SPF no smoothing model.

Figure 12: Robustness checks: trend inflation τ_t^π and natural rate of unemployment τ_t^u

(a) Trend inflation τ_t^π estimates

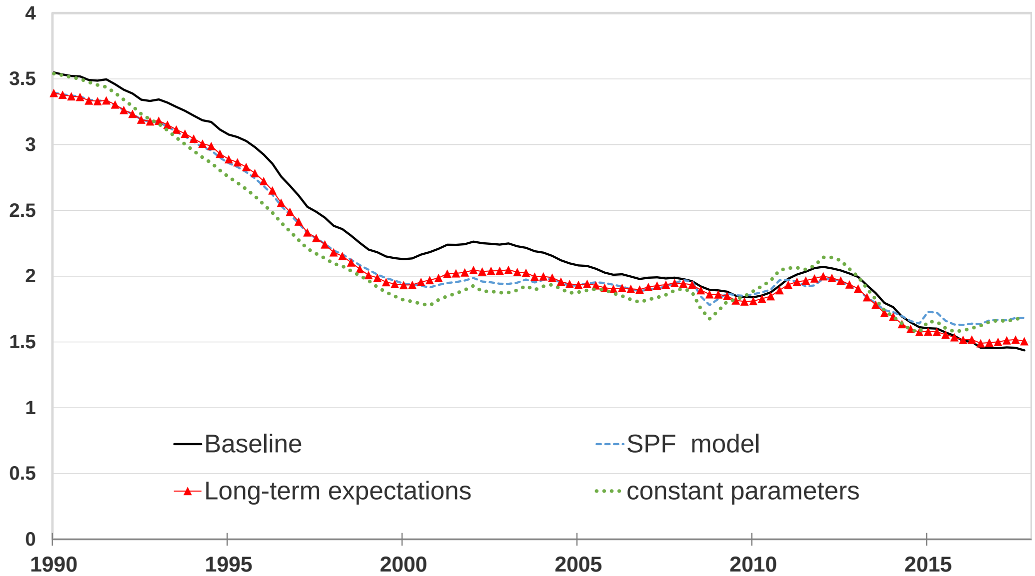

(b) Unemployment rate and estimated natural rates τ_t^u

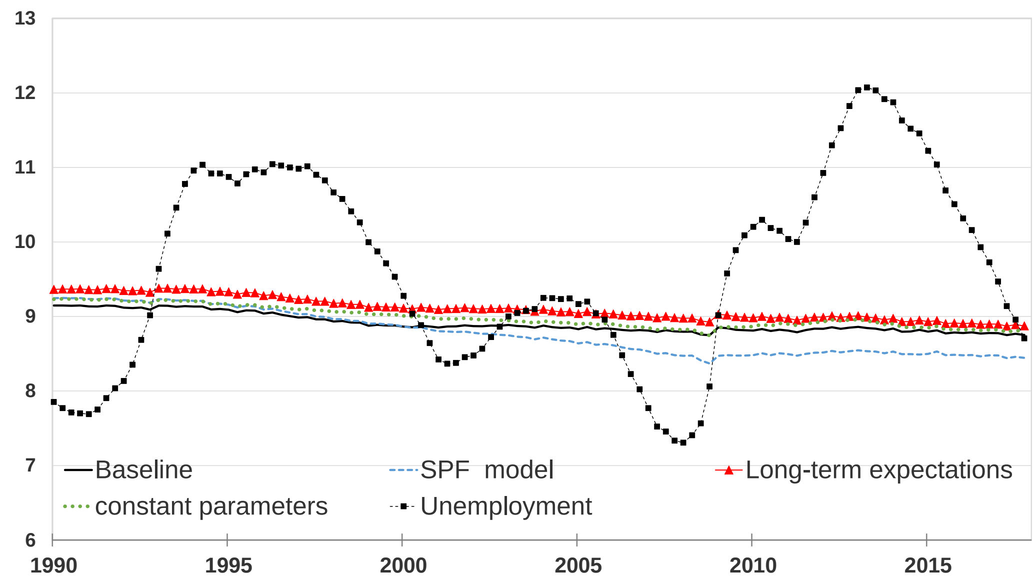

Note: The figures compare the median trend estimates from the different robustness checks. Panel (a) shows trend inflation τ_t^π estimates and panel (b) shows the natural rate of unemployment τ_t^u estimates. The 'LT-exp' and 'CST' models refer to the robustness checks with *i)* only long-term expectations and *ii)* constant transmission parameters, respectively.

Appendix

This appendix starts with a description of the data sources in Section A.1. and then provides more details concerning the estimation procedure. The inclusion of survey expectations data in the SPF and SPF no smoothing models implies nonlinear state space structures in the conditional posterior distributions, for which particular steps need to be taken. Section A.2. describes how we link the model forecast to the survey inflation forecasts. Subsequently, we explain the Gibbs sampler for the SPF model in Section A.3., and briefly discuss in Section A.4. how the algorithm changes for the baseline model and the forecast smoothing model.

A.1. Data

The macroeconomic series are sourced from the ECB's Statistical Data Warehouse database (SDW)[17] and correspond to the series published in the ECB's Economic Bulletin. We backdate these series using historical data from the Area Wide Model database (AWM)[18]. Specifically, we follow the AWM procedure and backdate price indexes and the unemployment rate using growth rates (Fagan et al., 2001, Annex 2). Prior to backdating, the AWM HICP price index was seasonally adjusted with the X13 procedure using JDEMETRA+ software.[19] Details are given in Table 3 below.

The inflation expectations series are sourced from the ECB's SPF webpage[20]. Specifically, we collect the aggregate probability distributions for inflation at the one-year, two-year and five-year ahead horizons, and compute the mean from these distributions at each point in time. Note that these are discrete probability distributions with bins such as [1.5%, 1.9%] , [2%, 2.4%], etc. To gauge the mean of the probability distribution at each point in time, we compute a weighted sum of the means of the bins. The weights are the probabilities assigned to the bins by the forecasters, and the mean value of each bin is the mean of the two outer points in the interval (e.g. for the [2%, 2.4%] interval it is (2%+2.4%)/2).

Table 3: Data

Variable	**Source (and *codes*)**
Headline inflation	SDW (*ICP.M.U2.Y.000000.3.INX*)
	AWM (*HICP*)
Unemployment rate	SDW (*STS.M.I8.S.UNEH.RTT000.4.000*)
	AWM (*URX*)
GDP deflator	SDW (*MNA.Q.Y.I8.W2.S1.S1.B.B1GQ._Z._Z._Z.IX.D.N*)
	AWM (*YED*)
Import price inflation	SDW (*MNA.Q.Y.I8.W1.S1.S1.C.P7._Z._Z._Z.IX.D.N*)
	AWM (*MTD*)
Inflation expectations	ECB SPF website

[17]See http://sdw.ecb.europa.eu/.
[18]See https://eabcn.org/page/area-wide-model.
[19]https://ec.europa.eu/eurostat/cros/content/software-jdemetra_en
[20]https://www.ecb.europa.eu/stats/ecb_surveys/survey_of_professional_forecasters/html/index.en.html.

A.2. Linking the survey expectations with the model forecast

▷ *Notation*

Our estimation sample ranges from 1990Q1 to 2017Q4, but the SPF inflation expectations data is only available from 1999Q1 onward. We denote the starting date of the SPF with $\bar{t}$, such that $1 < \bar{t} < T$, with T denoting the final observation. We explain below how the estimation algorithm differs for time periods $t < \bar{t}$ without the survey data compared to periods $t \geq \bar{t}$ which include it. Starting with some notation, denote $Y_t = (\pi_t, u_t, \pi_t^m)'$, and $Y^T = (Y_1', ..., Y_T')'$ as the vectors which stack the macro data, and let $Z_t = (\pi^e_{t+h_1|t}, ..., \pi^e_{t+h_n|t})'$, and $Z^T = (Z_{\bar{t}}', ..., Z_T')'$ denote the vectors which stack the survey expectations data collected in periods $\bar{t}, ..., T$. Similarly, the superscript T is used to indicate a vector of (stacked) time-varying parameters, e.g. as in $\lambda^T = (\lambda_1, ..., \lambda_T)'$. Also, let $\tau_t = (\tau_t^\pi, \tau_t^u, \tau_t^m)'$ and $\tau^T = (\tau_1', ..., \tau_T')'$ collect the time-varying trends, and define $\theta_t = \left(\tau_t', \tau_{t-1}', \rho_t^\pi, \lambda_t, \gamma_t, \rho_1^u, \rho_2^u\right)'$ as a vector that collects the relevant parameters for forecasting inflation using data up to period t. Finally, the detrended variables are given by $\tilde{\pi}_t = \pi_t - \tau_t^\pi$, $\tilde{u}_t = u_t - \tau_t^u$, and $\tilde{\pi}_t^m = \pi_t^m - \tau_t^m$, and they are collected in the vector $\tilde{Y}_t = (\tilde{\pi}_t, \tilde{u}_t, \tilde{\pi}_t^m)'$.

▷ *Rewriting the macro block in VAR / state space form*

Measurement equations (16) to (17) link survey expectations to the model forecast. To generate this model-consistent forecast, we rewrite equations (1) to (3) as a vector autoregressive model (VAR):

$$\underbrace{\begin{pmatrix} 1 & -\lambda_t & -\gamma_t \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{A_{0,t}} \begin{pmatrix} \tilde{\pi}_t \\ \tilde{u}_t \\ \tilde{\pi}_t^m \end{pmatrix} = \underbrace{\begin{pmatrix} \rho_t^\pi & 0 & 0 \\ 0 & \rho_1^u & 0 \\ 0 & 0 & 0 \end{pmatrix}}_{A_{1,t}} \begin{pmatrix} \tilde{\pi}_{t-1} \\ \tilde{u}_{t-1} \\ \tilde{\pi}_{t-1}^m \end{pmatrix} + \underbrace{\begin{pmatrix} 0 & 0 & 0 \\ 0 & \rho_2^u & 0 \\ 0 & 0 & 0 \end{pmatrix}}_{A_2} \begin{pmatrix} \tilde{\pi}_{t-2} \\ \tilde{u}_{t-2} \\ \tilde{\pi}_{t-2}^m \end{pmatrix}$$
$$+ \begin{pmatrix} \epsilon_t^\pi \\ \epsilon_t^u \\ \epsilon_t^m \end{pmatrix}.$$

This VAR, which describes our macroeconomic series, can be cast in state space form as

$$X_t = F_t X_{t-1} + e_t, \tag{18}$$

where the detrended variables are collected in the state vector $X_t = \left(\tilde{\pi}_t, \tilde{u}_t, \tilde{\pi}_t^m, \tilde{\pi}_{t-1}, \tilde{u}_{t-1}, \tilde{\pi}_{t-1}^m\right)'$, the transition dynamics are given by

$$F_t = \begin{pmatrix} A_{0,t}^{-1} A_{1,t} & A_{0,t}^{-1} A_2 \\ I_3 & 0_{3\times3} \end{pmatrix},$$

and the error terms are

$$e_t = \begin{pmatrix} A_{0,t}^{-1} \\ 0_{3\times3} \end{pmatrix} (\epsilon_t^\pi, \epsilon_t^u \epsilon_t^m)'.$$

▷ *Timing of the SPF data*

Two remarks are in order concerning the data. First, the SPF expectations data are collected at the start of each quarter and are based on data up to a certain month in the previous quarter. So, for example, the one-year ahead expected year-on-year inflation rate in the SPF from Q1 refers to year-on-year inflation in the month of *December* of the same year, etc. We follow the 'noise interpretation' of Smets et al. (2014), and consider these series as noisy indicators of year-on-year inflation of the quarter that contains the month to which is referred. Ergo, the Q1 SPF one-year ahead year-on-year expected inflation is taken as a measure for expected year-on-year inflation three quarters ahead, etc. In addition, although the five-year ahead forecast in the SPF is a calendar year forecast, we use it as a proxy for the rolling horizon forecast of year-on-year inflation five years ahead.

Second, given the timing of the SPF collection, we consider this series to indicate the forecasters' views on the trends and coefficients from the previous quarter. Concretely, each period t's SPF expectations series for h periods ahead inflation, $\pi^e_{t+h|t}$, is shifted one period back in time (to $t-1$) in our empirical application such that they inform on trends and transmission coefficients from quarter $t-1$.

▷ *Defining point forecasts for year-on-year inflation*

Inflation is defined as the annualised quarter-on-quarter growth rate of the price index: π_t=$400ln(P_t/P_{t-1})$, where P_t is the price index and $ln(.)$ is the natural logarithm. Denote π^a_t as the year-on-year inflation in quarter t, then $\pi^a_t = \frac{1}{4}\left(\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3}\right)$. We use equation (18) of the detrended model to generate forecasts of the inflation gap $E_t(\tilde{\pi}_{t+h})$. Rewrite this term as $E_t(\pi_{t+h} - \tau^\pi_{t+h})$=$E_t(\pi_{t+h})$−$E_t(\tau^\pi_{t+h})$, then it follows that the point forecast for inflation, $E_t(\pi_{t+h})$, is the sum of the forecasted inflation gap $E_t(\tilde{\pi}_t)$ and expected trend $E_t(\tau^\pi_{t+h})$.

▷ *Model forecast function* $f_h\left(\theta_{t-1}, Y^{t-1}\right)$

We generate model-consistent inflation forecasts by iterating the companion form (18) forward. We invoke the anticipated utility model (AUM) and keep the time-varying parameters fixed to their current states. Hence, when forecasting future inflation using data up to period t, we expect future trend inflation to equal τ^π_t. To match the survey expectations with the model forecast, we consider the forecasts for year-on-year inflation one, two and five years ahead. Recall that the SPF survey is conducted at the start of each quarter. Hence, in our analysis, we link the one-year ahead forecast $\pi^e_{t+3|t}$ with the model-implied three-quarter ahead forecast of year-on-year inflation, which is given by

$$\begin{aligned}
f_3\left(\theta_{t-1}, Y^{t-1}\right) &= \tau^\pi_{t-1} + e_1'\frac{1}{4}\left(\hat{X}_{t|t} + \hat{X}_{t+1|t} + \hat{X}_{t+2|t} + \hat{X}_{t+3|t}\right) \\
&= \tau^\pi_{t-1} + e_1'\frac{1}{4}\left(F_{t-1} + F^2_{t-1} + F^3_{t-1} + F^4_{t-1}\right)X_{t-1} \\
&= \tau^\pi_{t-1} + e_1'\frac{1}{4}F_{t-1}\left(I_6 - F_{t-1}\right)^{-1}\left(I_6 - F^4_{t-1}\right)X_{t-1},
\end{aligned}$$

where $\hat{X}_{i|j}$ stands for the forecasted state vector using the information known at the start of quarter j, i.e. the SPF data from quarter j and the macroeconomic data up to period $j-1$. The unit vector e_1, which has size 6×1 and contains 1 in row 1 and zero elsewhere, selects the element in the first row as this corresponds to the inflation forecast. Given that the state vector contains the detrended annualized quarter-on-quarter inflation rate, we construct the year-on-year detrended inflation rate as the 4 quarter moving average of the quarterly detrended inflation rates, and add trend inflation in order to generate the overall expected inflation rate. For the two-year and five-year ahead forecast of year-on-year inflation, we use

$$f_7\left(\theta_{t-1}, Y^{t-1}\right) = \tau^{\pi}_{t-1} + e_1' \frac{1}{4} F^5_{t-1} \left(I_6 - F_{t-1}\right)^{-1} \left(I_6 - F^4_{t-1}\right) X_{t-1}$$
$$f_{19}\left(\theta_{t-1}, Y^{t-1}\right) = \tau^{\pi}_{t-1} + e_1' \frac{1}{4} F^{17}_{t-1} \left(I_6 - F_{t-1}\right)^{-1} \left(I_6 - F^4_{t-1}\right) X_{t-1.}$$

It is important to note that $f_h\left(\theta_{t-1}, Y^{t-1}\right)$ is nonlinear in the parameters that are in θ_{t-1} - that is, all parameters except the time-varying trends. Since these parameters affect the likelihood function of the survey data, they will require additional attention in the Gibbs sampler, as explained below.

A.3. Gibbs sampler for the SPF no smoothing model

This section describes the Gibbs sampling algorithm for the SPF no smoothing model. The aim is to draw from the joint posterior of all unknown parameters by drawing iteratively from the conditional posterior distributions using a Gibbs sampling algorithm which goes through the following steps:

Initialize: We initialize all state variables and time-invariant parameters at their prior means.

Step 1: Draw the error variances from their conditional posterior distribution:
$p\left(\sigma^2_u, ..., \sigma^2_\psi | Y^T, Z^T, \tau^T, \rho^{\pi,T}, \lambda^T, \gamma^T, \psi^T\right)$. Observe that the error variances are conditionally independent given the data and the state variables. Therefore, we can draw them one by one from the appropriate distributions as described in the appendix of Chan et al. (2016). There are two cases to consider: error variances related to unbounded states and measurement equations, and error variances related to bounded states λ^T and $\rho^{\pi,T}$.

▷ *Error variances of measurement equations and unbounded states*

Using standard linear regression results, it follows that the conditional posteriors follow standard inverse gamma distributions. For example,

$$\left(\sigma^2_u | u^T, \tau^T_u, \rho^u_1, \rho^u_2\right) \sim IG\left(\underline{\nu}_u + \frac{T}{2}, \underline{\mathrm{S}}_u + \frac{1}{2}\sum_1^T \left(\epsilon^u_t\right)^2\right),$$

where the prior is defined as $\sigma^2_u \sim IG\left(\underline{\nu}_u, \underline{\mathrm{S}}_u\right)$. The error variances σ^2_m, $\sigma^2_{h_1}$,..., $\sigma^2_{h_n}$, $\sigma^2_{\tau\pi}$, $\sigma^2_{\tau u}$, $\sigma^2_{\tau m}$, σ^2_γ, and σ^2_ψ are drawn in a similar manner.

▷ *Error variances σ^2_ρ and σ^2_λ of the bounded states*

The error variances related to the bounded states $\rho^{\pi,T}$ and λ^T require a different approach because the error terms in those state equations are drawn from truncated normal distributions instead of normal distributions. Therefore,

$$\begin{aligned} logp\left(\sigma^2_\rho | \rho^{\pi,T}\right) &\propto logp\left(\rho^{\pi,T} | \sigma^2_\rho\right) + logp\left(\sigma^2_\rho\right) \\ &\propto -\frac{T-1}{2} log\sigma^2_\rho - \frac{1}{2\sigma^2_\rho}\sum_2^T \left(\rho^\pi_t - \rho^\pi_{t-1}\right)^2 \\ &\left\{-\sum_2^T log\left(\Phi\left(\frac{1-\rho^\pi_{t-1}}{\sigma_\rho}\right) - \Phi\left(\frac{-\rho^\pi_{t-1}}{\sigma_\rho}\right)\right)\right\} - \left(\underline{\nu}_\rho - 1\right) log\sigma^2_\rho - \frac{\underline{\mathrm{S}}_\rho}{\sigma^2_\rho}, \end{aligned}$$

which is a non-standard density due to the part that enters in brackets. Following Chan et al. (2016), we implement a Metropolis Hastings step with proposal density

$$IG\left(\underline{\nu}_{\rho}+\frac{T-1}{2},\underline{S}_{\rho}+\frac{1}{2}\sum_{2}^{T}\left(\rho_t^{\pi}-\rho_{t-1}^{\pi}\right)^2\right),$$

which is based on the above kernel, but discards the part in brackets. σ_{λ}^2 is drawn analogously.

Step 2: Sample the persistence parameters ρ_1^u, ρ_2^u from $p\left(\rho_1^u, \rho_2^u | u^T, Z^T, \lambda^T, \rho^{\pi,T}, \gamma^T, \tau^T, \sigma_u^2, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right)$. Notice that information for these parameters is given by the unemployment gap equation (2) and the measurement equations (16) to (17) for the survey data. If we disregard the latter set of equations, we obtain standard regression results (see Chan et al., 2016). However, the inclusion of survey data implies that these parameters enter nonlinearly in the function $f_h(\theta_{t-1}, Y^{t-1})$ for each period $t \geq \bar{t}$.

To draw from $\rho^u \equiv (\rho_1^u, \rho_2^u)'$, we apply an independent Metropolis Hastings step where the proposal distribution is based on a model where the nonlinear function $f_h(\theta_{t-1}, Y^{t-1})$ is linearized. Denote θ_{t-1/ρ^u} as the vector θ_{t-1} without the ρ^u elements. We can then rewrite $f_h(\theta_{t-1}, Y^{t-1})$ as $f_h(\theta_{t-1/\rho^u}, \rho^u, Y_{t-1})$, and take a first order approximation of the latter in the point ρ_0^u:

$$\begin{aligned} f_h(\theta_{t-1/\rho^u}, \rho^u, Y^{t-1}) &\approx f_h(\theta_{t-1/\rho^u}, \rho_0^u, Y^{t-1}) + \left(\frac{\partial f_h(\theta_{t-1/\rho^u}, \rho_0^u, Y^{t-1})}{\partial \rho_0^u}\right)'(\rho^u-\rho_0^u) \\ &\approx \underbrace{f_h(\theta_{t-1/\rho^u}, \rho_0^u, Y^{t-1}) - \left(\frac{\partial f_h(\theta_{t-1/\rho^u}, \rho_0^u, Y^{t-1})}{\partial \rho_0^u}\right)'\rho_0^u}_{c_t^h} + \underbrace{\left(\frac{\partial f_h(\theta_{t-1/\rho^u}, \rho_0^u, Y^{t-1})}{\partial \rho_0^u}\right)'}_{\tilde{x}_{h,t}^{u}{}'}\rho^u. \end{aligned}$$

In other words, we transform the measurement equations with survey data (16) to (17) into equations which are linear functions of ρ^u. Using standard regression results, this approximate model delivers closed form solutions for the conditional posterior of ρ^u, which is a normal distribution. Specifically, we rewrite the unemployment gap equation (2) as

$$\tilde{u}^T = X^u \rho^u + \epsilon^{u,T},$$

where $X^u = \begin{bmatrix} \tilde{u}_0 & \tilde{u}_{-1} \\ \vdots & \vdots \\ \tilde{u}_{T-1} & \tilde{u}_{T-2} \end{bmatrix}$, $var(\epsilon^{u,T}) = \sigma_u^2 I_T$, and we stack the linearized measurement equations for the survey data as

$$Z^T = C + \tilde{X}^u \rho^u + \epsilon^{z,T},$$

where $C = \left(c_{\bar{t}}', ..., c_T'\right)'$, $c_t = \left(c_t^{h_1}, ..., c_t^{h_n}\right)'$, $\tilde{X}^u = \begin{pmatrix} \tilde{x}_{\bar{t}}^u \\ \vdots \\ \tilde{x}_T^u \end{pmatrix}$, $\tilde{x}_t^u = \begin{pmatrix} \tilde{x}_{h_1,t}^{u}{}' \\ \vdots \\ \tilde{x}_{h_n,t}^{u}{}' \end{pmatrix}$, $\epsilon^{z,T} = \left(\epsilon_{\bar{t}}^{z}{}', ..., \epsilon_T^{z}{}'\right)'$,

$\epsilon_t^z = \left(\epsilon_t^{h_1}, ..., \epsilon_t^{h_n}\right)'$, and $var(\epsilon^{z,T}) \equiv \Omega_z = I_{T-\bar{t}+1} \otimes \begin{pmatrix} \sigma_{h_1}^2 & & 0 \\ & \ddots & \\ 0 & & \sigma_{h_n}^2 \end{pmatrix}$. Combining the likelihood functions with a normal prior $p\left(\rho^u\right) \sim N\left(\underline{\rho}^u, \underline{V}_u\right)$ leads to a normal conditional posterior $\left(\rho^u | Y^T, Z^T, ...\right) \sim N\left(\bar{\rho}^u, \bar{V}_u\right)$, where

$$\begin{aligned}
\bar{V}_u &= \left(\underline{V}_u^{-1} + X^{u\,\prime} X^u / \sigma_u^2 + \tilde{X}^{u\,\prime} \Omega_z^{-1} \tilde{X}^u\right)^{-1} \\
\bar{\rho}^u &= \bar{V}_u \left(\underline{V}_u^{-1} \underline{\rho}^u + X^{u\,\prime} \tilde{u}^T / \sigma_u^2 + \tilde{X}^{u\,\prime} \Omega_z^{-1} \left(Z^T - C\right)\right).
\end{aligned}$$

Using these results, we take a candidate draw from a t-distribution with degrees of freedom 10, mean $\bar{\rho}^u$ and variance $\bar{V}_u$, in order to endow the proposal distribution with fatter tails. If the candidate draw is non-stationary, we use the previous draw for ρ^u as the current draw. If it meets the stationarity conditions, we accept it with a certain probability according to the Metropolis Hastings procedure. We select ρ_0^u, the parameter values around which the function f_h is linearized, as the previously accepted draw.[21] In sum, we generate a candidate draw for the persistence parameters from an approximate model which linearizes the model forecast functions around the posterior mean from a model that discards the survey data.[22]

Step 3: Sample the time-varying trends $\tau_t^\pi, \tau_t^u, \tau_t^m$ for $t = 1, ..., T$. Conditional on the time-varying coefficients $\rho_t^\pi, \lambda_t, \gamma_t$, the model can be cast in a linear state space form, and the trends can be drawn with the Carter and Kohn (1994) algorithm. Building on the expressions from section A.1., we consider the following augmented state vector:

$$\begin{aligned}
\tilde{X}_t &= \left(\tau_t', X_t'\right)' \\
&= \left(\tau_t^\pi, \tau_t^u, \tau_t^m, \tilde{\pi}_t, \tilde{u}_t, \tilde{\pi}_t^m, \tilde{\pi}_{t-1}, \tilde{u}_{t-1}, \tilde{\pi}_{t-1}^m\right)'.
\end{aligned}$$

The measurement equations, spelled out for our implementation with three survey expectations series for inflation in periods $t \geq \bar{t}$, are built by stacking Y_t and Z_{t+1} in the left-hand side:

$$\begin{pmatrix} \pi_t \\ u_t \\ \pi_t^m \\ \pi_{t+4|t+1}^e \\ \pi_{t+8|t+1}^e \\ \pi_{t+20|t+1}^e \end{pmatrix} = \begin{pmatrix} I_3 & I_3 & 0_{3\times 3} \\ \hdashline 1 & 0_{1\times 2} & e_1' F_t / 4 \left(I_6 - F_t\right)^{-1} \left(I_6 - F_t^4\right) \\ 1 & 0_{1\times 2} & e_1' F_t^5 / 4 \left(I_6 - F_t\right)^{-1} \left(I_6 - F_t^4\right) \\ 1 & 0_{1\times 2} & e_1' F_t^{17} / 4 \left(I_6 - F_t\right)^{-1} \left(I_6 - F_t^4\right) \end{pmatrix} \tilde{X}_t + \begin{pmatrix} 0 \\ 0 \\ 0 \\ \epsilon_{t+1}^3 \\ \epsilon_{t+1}^7 \\ \epsilon_{t+1}^{19} \end{pmatrix},$$

which builds on the previously defined functions $f_3\left(\theta_{t-1}, Y^{t-1}\right)$, $f_7\left(\theta_{t-1}, Y^{t-1}\right)$ and $f_{19}\left(\theta_{t-1}, Y^{t-1}\right)$ that define the model forecasts. For the periods $t < \bar{t}$ without survey data, the left and right-hand side of the above expression are left-multiplied with the matrix $\left(I_3\, 0_{3\times 3}\right)$ in order to abstract from the survey data equations.

The state equations are given by

[21] We have also experimented with setting ρ_0^u to the conditional posterior mean that is obtained when the measurement equations for the survey data are ignored: $\rho_0^u = \left(\underline{V}_u^{-1} + X^{u\,\prime} X^u / \sigma_u^2\right)^{-1} \left(\underline{V}_u^{-1} \underline{\rho}^u + X^{u\,\prime} \tilde{u}^T / \sigma_u^2\right)$. However, results were found to be similar.

[22] To setup the C and $\tilde{X}^u$ matrices, we use the symbolic toolbox in Matlab to derive the Jacobian of the functions f_h, and use the `matlabFunction()` command to convert this symbolic expression into a vectorized Matlab function.

$$\tilde{X}_t = \left(\begin{array}{cc} I_3 & 0_{3\times6} \\ \hdashline 0_{6\times3} & F_t \end{array}\right) \tilde{X}_{t-1} + \left(\begin{array}{cc} I_3 & 0_{3\times3} \\ 0_{3\times3} & A_{0,t}^{-1} \\ \hdashline 0_{3\times6} & \end{array}\right) \begin{pmatrix} \eta_t^{\tau\pi} \\ \eta_t^{\tau u} \\ \eta_t^{\tau m} \\ \epsilon_t^{\pi} \\ \epsilon_t^{u} \\ \epsilon_t^{m} \end{pmatrix}.$$

Step 4: Sample the time-varying transmission coefficients $\rho^{\pi,T}, \lambda^T, \gamma^T$ for $t = 1, ..., T$. This block involves two complications. First, these time-varying parameters also enter nonlinearly in the model forecast functions $f_{h_1}\left(\theta_t, Y^{t-1}\right), ..., f_{h_n}\left(\theta_t, Y^{t-1}\right)$ for the survey expectations equations, which precludes setting up a linear state space model as in Step 3. Second, the states $\rho^{\pi,T}$ and λ^T are bounded to lie within certain intervals. To accommodate both features, we implement a single-move sampler based on Cogley (2005) and Koop and Potter (2011), where for each period $t = j$ the time-varying coefficients are drawn conditional on the values for these coefficients in periods $t \neq j$, in addition to the other model parameters, using an independent Metropolis-Hastings step.

Define $\delta_t = (\rho_t^{\pi}, \lambda_t, \gamma_t)'$ as the vector collecting the time-varying coefficients, covariance matrix $Q = diag\left(\sigma_\rho^2, \sigma_\lambda^2, \sigma_\gamma^2\right)$, and θ_{t/δ_t} as the vector θ_t excluding the δ_t elements. In each period $t = \bar{t}, ..., T-1$ the single move sampler draws from (see Koop and Potter, 2011, equation 15)[23]:

$$\begin{aligned} & p\left(\delta_t|\delta_{j\neq t}, Y^T, Z^T, \theta_{t/\delta_t}, \tau^T, \psi^T, \sigma_\lambda^2, \sigma_\rho^2, \sigma_\gamma^2, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) \propto \\ & \qquad p\left(Z_{t+1}|Y_t, Y_{t-1}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) p\left(Y_t|Y_{t-1}, Y_{t-2}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \psi_t, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) \\ & \qquad\qquad p\left(\delta_{t+1}|\delta_t, Q\right) p\left(\delta_t|\delta_{t-1}, Q\right) \frac{1\left(\delta_t \in A\right)}{R\left(\delta_t, Q\right)}. \end{aligned}$$

The two terms in the second line correspond to the likelihood function of the data, and the next three terms in the third line to the prior distribution. To draw from this conditional posterior distribution, we derive an analytical expression for a proposal density by extracting certain terms from the above expression. This is described in the following steps:

▷ *The likelihood function*

Concerning the joint likelihood of Y^T and Z^T, we only keep the parts which depend on δ_t, since the rest is absorbed by the integrating constant. We decompose the joint likelihood $p\left(Z^T, Y^T|\delta_t, \delta_{j\neq t}, ...\right)$ into the product of predictive likelihoods: $\prod_{t=1}^T p\left(Z_t, Y_t|Z^{t-1}, Y^{t-1}, \delta_t, \delta_{j\neq t}, ...\right)$, and then decompose the joint likelihood in each period t as

$$\begin{aligned} p\left(Z_t, Y_t|Z^{t-1}, Y^{t-1}, \delta_t, \delta_{j\neq t}, ...\right) &= p\left(Z_t|Z^{t-1}, Y^{t-1}, \delta_t, ...\right) p\left(Y_t|Z^{t-1}, Y^{t-1}, \delta_t, \delta_{j\neq t}, ...\right) \\ &= p\left(Z_t|Y_{t-1}, Y_{t-2}, \delta_{t-1}, ...\right) p\left(Y_t|Y_{t-1}, Y_{t-2}, \delta_t, ...\right). \end{aligned}$$

Including survey data into the model has an impact on the estimated trends, transmission coefficients, etc. But once we condition on past data and the model parameters, we consider Y_t and Z_t to be (conditionally)

[23] In the final period $t = T$, there is no δ_{T+1} to condition on, so the $p\left(\delta_{t+1}|\delta_t, Q\right)/R\left(\delta_t, Q\right)$ terms disappear from the conditional posterior density.

independently distributed. Of the whole joint likelihood function, only the vectors Z_{t+1} and Y_t depend on δ_t, and therefore these two likelihood terms remain in the above kernel of the posterior of δ_t.

By decomposing the likelihood function in this way, we can exploit the fact that in the $p\left(Y_t|Y_{t-1},Y_{t-2},\delta_t,...\right)$ component, Y_t is a linear function of δ_t. In particular, we have that

$$\begin{aligned}\tilde{\pi}_t &= \left(\tilde{\pi}_{t-1},\tilde{u}_t,\tilde{\pi}^m_t\right)\delta_t+\epsilon^\pi_t\\ &= X^\delta_t\delta_t+\epsilon^\pi_t,\end{aligned}$$

is the only part that depends on δ_t. For the periods $t<\bar{t}$ which do not contain survey data, $p\left(Y_t|Y_{t-1},Y_{t-2},\delta_t,...\right)$ is the only likelihood term. Our proposal distribution uses this expression in combination with the prior distribution. We now turn to the latter.

▷ *The prior distribution*

The elements of δ_t evolve as independent random walks which are subject to the constraints that ρ^π_t and λ_t lie within certain intervals. These restrictions are captured by the indicator function $1\left(\delta_t\in A\right)$, which equals 1 if δ_t satisfies the bounds and is zero otherwise. The term $R\left(\delta_t,Q\right)$ indicates the integrating constant from the restricted prior distribution $1\left(\delta_{t+1}\in 1\right)p\left(\delta_{t+1}|\delta_t,Q\right)$. Intuitively, the integrating constant measures the percentage of random draws from the normal distribution $p\left(\delta_{t+1}|\delta_t,Q\right)$ that would fall within the acceptance region (see Koop and Potter, 2011).

▷ *Calculating the integrating constant $R\left(\delta_t,Q\right)$*

Assuming that the elements of δ_t evolve as independent random walks allows us to derive analytical expressions for the integrating constant $R\left(\delta_t,Q\right)$. By decomposing the joint distribution as the product of three independent distributions, we obtain

$$p\left(\delta_{t+1}|\delta_t,Q\right)=p\left(\rho^\pi_{t+1}|\rho^\pi_t,\sigma^2_\pi\right)p\left(\lambda_{t+1}|\lambda_t,\sigma^2_\lambda\right)p\left(\gamma_{t+1}|\gamma_t,\sigma^2_\gamma\right).$$

Therefore, the integrating constant of the restricted prior is

$$\begin{aligned}R\left(\delta_t,Q\right) &= \int_0^1\int_{-1}^0\int_{-\infty}^\infty p\left(\delta_{t+1}|\delta_t,Q\right)d\rho^\pi_{t+1}d\lambda_{t+1}d\gamma_{t+1}\\ &= \int_0^1 p\left(\rho^\pi_{t+1}|\rho^\pi_t,\sigma^2_\pi\right)d\rho^\pi_{t+1}\int_{-1}^0 p\left(\lambda_{t+1}|\lambda_t,\sigma^2_\lambda\right)d\lambda_{t+1}\int_{-\infty}^\infty p\left(\gamma_{t+1}|\gamma_t,\sigma^2_\gamma\right)d\gamma_{t+1}\\ &= \left(\Phi\left(\frac{1-\rho^\pi_t}{\sigma_\pi}\right)-\Phi\left(\frac{0-\rho^\pi_t}{\sigma_\pi}\right)\right)\left(\Phi\left(\frac{-1-\lambda_t}{\sigma_\lambda}\right)-\Phi\left(\frac{0-\lambda_t}{\sigma_\lambda}\right)\right),\end{aligned}$$

where $\Phi\left(.\right)$ is the cumulative distribution of the standard normal density.

▷ *Combining terms into a proposal density*

Our proposal density combines the terms:

$$p\left(Y_t|Y_{t-1},Y_{t-2},\delta_t,...\right)p\left(\delta_{t+1}|\delta_t,Q\right)p\left(\delta_t|\delta_{t-1},Q\right),$$

because they lead to closed form solutions for the candidate draw δ^*_t (see Carlin et al., 1992, for the formulae). Given a prior $\delta_0\sim N\left(\underline{\delta}_0,\underline{Q}_0\right)$, we obtain

$$\left(\delta^*_t|\delta_{j\neq t},Y^T,Z^T...\right)\sim N\left(\bar{\delta}_t,\bar{\Sigma}_t\right),$$

where

$$
\begin{aligned}
\bar{\Sigma}_t &= \left(\underline{Q}_0^{-1} + Q^{-1}\right)^{-1} & t = 0 \\
&= \left(X_t^{\delta\prime} X_t^{\delta} / \sigma_{\pi,t}^2 + 2Q^{-1}\right)^{-1} & t = 1, ..., T-1 \\
&= \left(X_t^{\delta\prime} X_t^{\delta} / \sigma_{\pi,t}^2 + Q^{-1}\right)^{-1} & t = T,
\end{aligned}
$$

and

$$
\begin{aligned}
\bar{\delta}_t &= \bar{\Sigma}_t \left(\underline{Q}_0^{-1} \underline{\delta}_0 + Q^{-1} \delta_1\right) & t = 0 \\
&= \bar{\Sigma}_t \left(X_t^{\delta\prime} \tilde{\pi}_t / \sigma_{\pi,t}^2 + Q^{-1} \delta_{t-1} + Q^{-1} \delta_{t+1}\right) & t = 1, ..., T-1 \\
&= \bar{\Sigma}_t \left(X_t^{\delta\prime} \tilde{\pi}_T / \sigma_{\pi,t}^2 + Q^{-1} \delta_{T-1}\right) & t = T.
\end{aligned}
$$

With these ingredients, we generate a candidate draw δ_t^* in each period and evaluate the Metropolis Hastings acceptance probability. This acceptance probability uses the remaining likelihood term for Z_{t+1}, the indicator function $1\left(\delta_t^* \in A\right)$ and integrating constant $R\left(\delta_t^*, Q\right)$. Notice that the Z_{t+1} term will only be used in periods $\bar{t} - 1 \leq t \leq T - 1$ for which we have survey data. In the periods $t < \bar{t} - 1$, the acceptance probability only depends on $1\left(\delta_t^* \in A\right)$ and $R\left(\delta_t^*, Q\right)$.

Step 5: Sample the stochastic volatility ψ_t for $t = 1, ..., T$ conditional on all other parameters.
The stochastic volatility terms ψ_t of the error term in the measurement equation (1) for inflation are drawn using the single-move sampler of Jacquier et al. (1994).

Repeat: Go back to step 1 until the required number of draws has been reached.

A.4. Gibbs sampler for the baseline and SPF models

▷ *Baseline model*

In the baseline model specification without survey data the whole procedure becomes more simple. Steps 1 and 5 remains the same. Steps 2 is based on a normal conditional posterior, for which acceptance-rejection sampling can be used. Step 3 requires the correction for the absence of survey data (by left-multiplication) to be applied in each period. The expressions for Step 4 require the removal of the likelihood terms related to Z^T.

▷ *SPF model*

Estimation of the model which allows for forecast smoothing (see equations 14 to 15) is similar to that of the previous section. In this case, we also draw from the conditional posterior of ξ^T and restrict each ξ_t to lie in the interval $(0, 1)$. We adjust Step 4 and jointly draw $\rho_t^\pi, \lambda_t, \gamma_t$ and ξ_t by adjusting the expressions for the candidate draw δ_t^* and the Metropolis Hastings acceptance probability accordingly. For instance, $\delta_t = \left(\rho_t^\pi, \lambda_t, \gamma_t, \xi_t\right)'$, $Q = diag\left(\sigma_\rho^2, \sigma_\lambda^2, \sigma_\gamma^2, \sigma_\xi^2\right)$, and $X_t^\delta = \left(\tilde{\pi}_{t-1}, \tilde{u}_t, \tilde{\pi}_t^m, 0\right)$.

A.5. Convergence

We executed 250,000 replications of the Gibbs sampler and discarded the first 50,000. Finally, we stored every 20th draw in order to break the autocorrelation and economize on storage size. This leaves us with 10,000 posterior draws. To assess convergence, we inspected the recursive means or the retained draws at every 20th draw, as shown in Figure 13. The fact that there is little evidence of large fluctuations in the posterior means is taken as evidence in favour of convergence.

Figure 13: Recursive means of vectorized parameters from the SPF expectations model

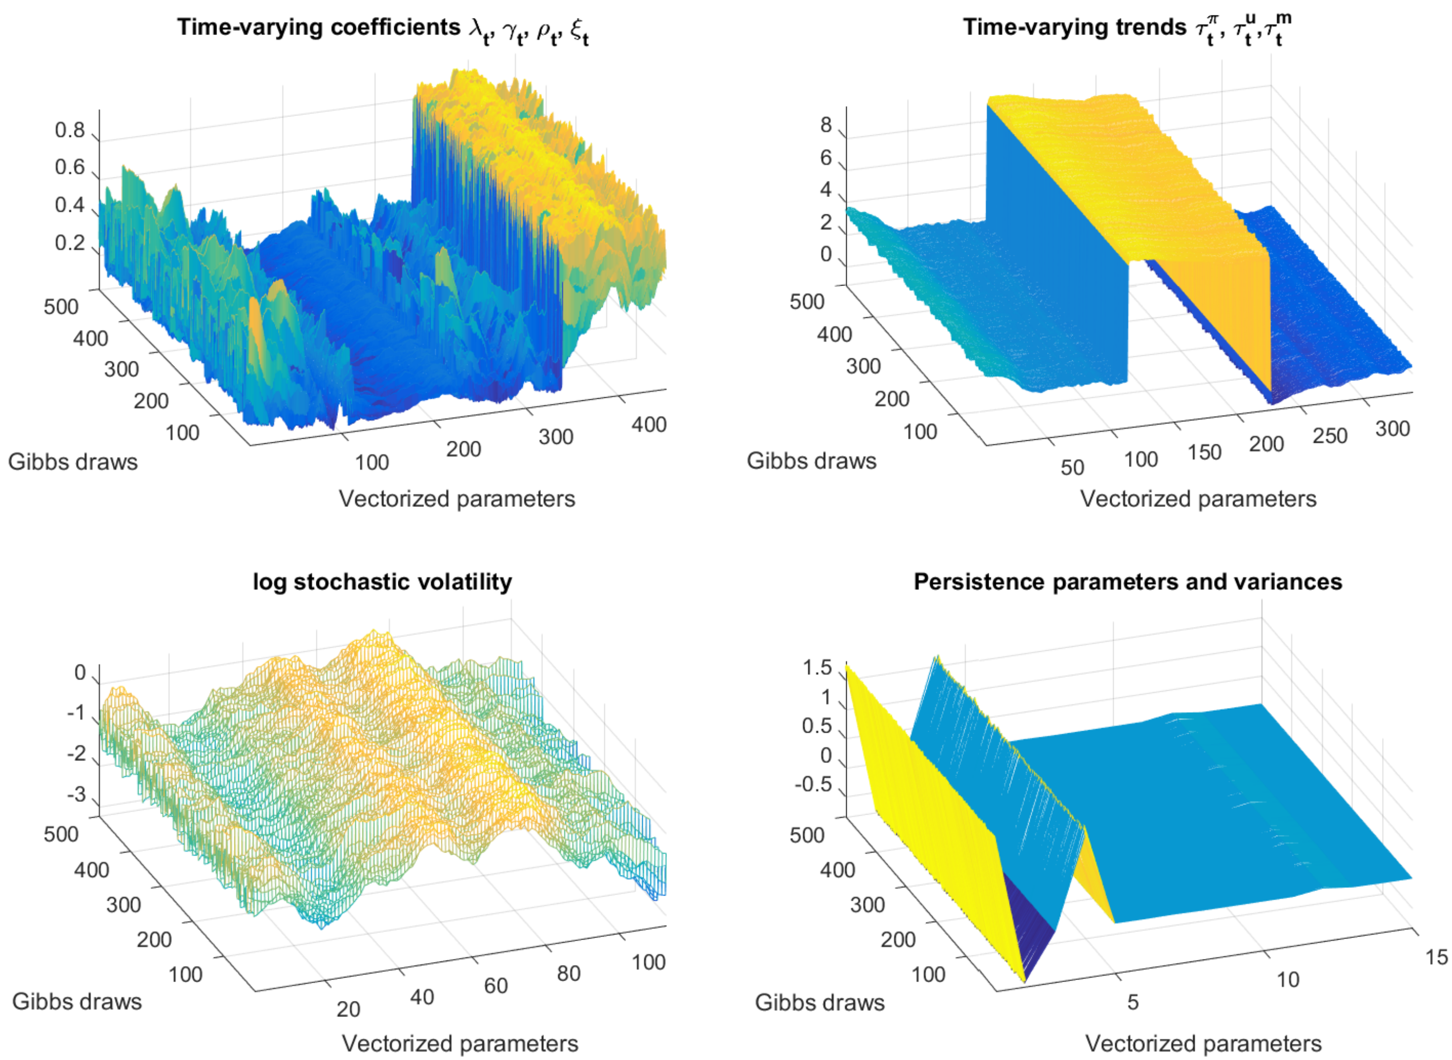

References

Ang, A., Bekaert, G., and Wei, M. (2007). Do macro variables, asset markets, or surveys forecast inflation better? *Journal of Monetary Economics*, 54(4):1163–1212.

Baele, L., Bekaert, G., Cho, S., Inghelbrecht, K., and Moreno, A. (2015). Macroeconomic regimes. *Journal of Monetary Economics*, 70:51 – 71.

Banbura, M. and van Vlodrop, A. (2018). Forecasting with Bayesian Vector Autoregressions with Time Variation in the Mean. Tinbergen Institute Discussion Papers 18-025/IV, Tinbergen Institute.

Bulligan, G. and Viviano, E. (2017). Has the wage Phillips curve changed in the euro area? *IZA Journal of Labor Policy*, 6(1):1–22.

Carlin, B. P., Polson, N. G., and Stoffer, D. S. (1992). A monte carlo approach to nonnormal and nonlinear state-space modeling. *Journal of the American Statistical Association*, 87(418):493–500.

Carroll, C. D. (2003). Macroeconomic Expectations of Households and Professional Forecasters. *The Quarterly Journal of Economics*, 118(1):269–298.

Carter, C. and Kohn, R. (1994). On gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Chan, J. C. C., Koop, G., and Potter, S. M. (2016). A Bounded Model of Time Variation in Trend Inflation, Nairu and the Phillips Curve. *Journal of Applied Econometrics*, 31(3):551–565.

Ciccarelli, M. and Osbat, C. (2017). Low inflation in the euro area: Causes and consequences. Occasional Paper Series 181, European Central Bank.

Cogley, T. (2005). Changing Beliefs and the Term Structure of Interest Rates: Cross-Equation Restrictions with Drifting Parameters. *Review of Economic Dynamics*, 8(2):420–451.

Cordemans, N. and Wauters, J. (2018). Are inflation and economic activity out of sync in the euro area? *Economic Review*, (i):79–96.

Crump, R. K., Eusepi, S., and Moench, E. (2016). The term structure of expectations and bond yields. Staff Reports 775, Federal Reserve Bank of New York, revised 01 Apr 2018.

Cui, W., Härdle, W. K., and Wang, W. (2015). Estimation of NAIRU with Inflation Expectation Data. SFB 649 Discussion Papers SFB649DP2015-010, Sonderforschungsbereich 649, Humboldt University, Berlin, Germany.

Dany-Knedlik, G. and Höltemoller, O. (2017). Inflation dynamics during the financial crisis in Europe: Cross-sectional identification of long-run inflation expectations. IWH Discussion Papers 10/2017, Halle Institute for Economic Research (IWH).

Del Negro, M. and Eusepi, S. (2011). Fitting observed inflation expectations. *Journal of Economic Dynamics and Control*, 35(12):2105–2131.

Fagan, G., Henry, J., and Mestre, R. (2001). An area-wide model (AWM) for the euro area. Working Paper Series 0042, European Central Bank.

Faust, J. and Wright, J. H. (2013). *Forecasting Inflation*, volume 2 of *Handbook of Economic Forecasting*, chapter 0, pages 2–56. Elsevier.

Garcia, J. A. and Poon, A. (2018). Trend Inflation and Inflation Compensation. IMF Working Papers 18/154, International Monetary Fund.

Grishchenko, O. V., Mouabbi, S., and Renne, J.-P. (2017). Measuring Inflation Anchoring and Uncertainty : A US and Euro Area Comparison. Finance and Economics Discussion Series 2017-102, Board of Governors of the Federal Reserve System (U.S.).

IMF (2013). The dog that didn't bark: Has inflation been muzzled or was it just sleeping? *World Economic Outlook, Chapter 3*, pages 79–95.

Jacquier, E., Polson, N. G., and Rossi, P. E. (1994). Bayesian analysis of stochastic volatility models. *Journal of Business & Economic Statistics*, 20(1):69–87.

Jarociński, M. and Bobeica, E. (2017). Missing disinflation and missing inflation: the puzzles that aren't. Working Paper Series 2000, European Central Bank.

Jarociński, M. and Lenza, M. (2018). An Inflation-Predicting Measure of the Output Gap in the Euro Area. *Journal of Money, Credit and Banking*, 50(6):1189–1224.

Koop, G. and Potter, S. M. (2011). Time varying VARs with inequality restrictions. *Journal of Economic Dynamics and Control*, 35(7):1126–1138.

Kozicki, S. and Tinsley, P. A. (2012). Effective Use of Survey Information in Estimating the Evolution of Expected Inflation. *Journal of Money, Credit and Banking*, 44(1):145–169.

Łyziak, T. and Paloviita, M. (2018). On the formation of inflation expectations in turbulent times: The case of the euro area. *Economic Modelling*, 72:132–139.

Mankiw, N. G. and Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Matheson, T. and Stavrev, E. (2013). The Great Recession and the inflation puzzle. *Economics Letters*, 120(3):468–472.

Mertens, E. and Nason, J. M. (2018). Inflation and professional forecast dynamics: an evaluation of stickiness, persistence, and volatility. BIS Working Papers 713, Bank for International Settlements.

Nason, J. M. and Smith, G. W. (2016). Sticky professional forecasts and the unobserved components model of us inflation. *Manuscript, Department of Economics, Queen's University.*

Neri, S., Nobili, A., and Conti, A. M. (2017). Low inflation and monetary policy in the euro area. Working Paper Series 2005, European Central Bank.

Riggi, M. and Venditti, F. (2015). Failing to Forecast Low Inflation and Phillips Curve Instability: A Euro-Area Perspective. *International Finance*, 18(1):47–68.

Smets, F., Warne, A., and Wouters, R. (2014). Professional forecasters and real-time forecasting with a DSGE model. *International Journal of Forecasting*, 30(4):981–995.

Stella, A. and Stock, J. H. (2013). A state-dependent model for inflation forecasting. International Finance Discussion Papers 1062, Board of Governors of the Federal Reserve System (U.S.).

Winkelried, D. (2017). Inferring Inflation Expectations from Fixed-Event Forecasts. *International Journal of Central Banking*, 13(2):1–31.

NATIONAL BANK OF BELGIUM - WORKING PAPERS SERIES

The Working Papers are available on the website of the Bank: http://www.nbb.be.

302. "The transmission mechanism of credit support policies in the Euro Area", by J. Boeckx, M. de Sola Perea and G. Peersman, *Research series*, October 2016.
303. "Bank capital (requirements) and credit supply: Evidence from pillar 2 decisions", by O. De Jonghe, H. Dewachter and S. Ongena, *Research series*, October 2016.
304. "Monetary and macroprudential policy games in a monetary union", by R. Dennis and P. Ilbas, *Research series*, October 2016.
305. "Forward guidance, quantitative easing, or both?, by F. De Graeve and K. Theodoridis, *Research series*, October 2016.
306. "The impact of sectoral macroprudential capital requirements on mortgage loan pricing: Evidence from the Belgian risk weight add-on", by S. Ferrari, M. Pirovano and P. Rovira Kaltwasser, *Research series*, October 2016.
307. "Assessing the role of interbank network structure in business and financial cycle analysis", by J-Y Gnabo and N.K. Scholtes, *Research series*, October 2016.
308. "The trade-off between monetary policy and bank stability", by M. Lamers, F. Mergaerts, E. Meuleman and R. Vander Vennet, *Research series,* October 2016.
309. "The response of euro area sovereign spreads to the ECB unconventional monetary policies", by H. Dewachter, L. Iania and J-C. Wijnandts, *Research series*, October 2016.
310. "The interdependence of monetary and macroprudential policy under the zero lower bound", by V. Lewis and S. Villa, *Research series*, October 2016.
311. "The impact of exporting on SME capital structure and debt maturity choices", by E. Maes, N. Dewaelheynes, C. Fuss and C. Van Hulle, *Research series*, October 2016.
312. "Heterogeneous firms and the micro origins of aggregate fluctuations", by G. Magerman, K. De Bruyne, E. Dhyne and J. Van Hove, *Research series*, October 2016.
313. "A dynamic factor model for forecasting house prices in Belgium", by M. Emiris, *Research series*, November 2016.
314. "La Belgique et l'Europe dans la tourmente monétaire des années 1970 – Entretiens avec Jacques van Ypersele", by I. Maes and S. Péters, *Research series*, December 2016.
315. "Creating associations to substitute banks' direct credit. Evidence from Belgium", by M. Bedayo, *Research series*, December 2016.
316. "The impact of export promotion on export market entry", by A. Schminke and J. Van Biesebroeck, *Research series*, December 2016.
317. "An estimated two-country EA-US model with limited exchange rate pass-through", by G. de Walque, Ph. Jeanfils, T. Lejeune, Y. Rychalovska and R. Wouters, *Research series,* March 2017.
318. Using bank loans as collateral in Europe: The role of liquidity and funding purposes", by F. Koulischer and P. Van Roy, *Research series*, April 2017.
319. "The impact of service and goods offshoring on employment: Firm-level evidence", by C. Ornaghi, I. Van Beveren and S. Vanormelingen, *Research series*, May 2017.
320. "On the estimation of panel fiscal reaction functions: Heterogeneity or fiscal fatigue?", by G. Everaert and S. Jansen, *Research series*, June 2017.
321. "Economic importance of the Belgian ports: Flemish maritime ports, Liège port complex and the port of Brussels - Report 2015", by C. Mathys, *Document series*, June 2017.
322. "Foreign banks as shock absorbers in the financial crisis?", by G. Barboni, *Research series*, June 2017.
323. "The IMF and precautionary lending: An empirical evaluation of the selectivity and effectiveness of the flexible credit line", by D. Essers and S. Ide, *Research series*, June 2017.
324. "Economic importance of air transport and airport activities in Belgium – Report 2015", by S. Vennix, *Document series*, July 2017.
325. "Economic importance of the logistics sector in Belgium", by H. De Doncker, *Document series*, July 2017.
326. "Identifying the provisioning policies of Belgian banks", by E. Arbak, *Research series*, July 2017.
327. "The impact of the mortgage interest and capital deduction scheme on the Belgian mortgage market", by A. Hoebeeck and K. Inghelbrecht, *Research series,* September 2017.
328. "Firm heterogeneity and aggregate business services exports: Micro evidence from Belgium, France, Germany and Spain", by A. Ariu, E. Biewen, S. Blank, G. Gaulier, M.J. González, Ph. Meinen, D. Mirza, C. Martín and P. Tello, *Research series*, September 2017.
329. "The interconnections between services and goods trade at the firm-level", by A. Ariu, H. Breinlichz, G. Corcosx, G. Mion, *Research series,* October 2017.

330. “Why do manufacturing firms produce services? Evidence for the servitization paradox in Belgium”, by P. Blanchard, C. Fuss and C. Mathieu, *Research series,* November 2017.

331. “Nowcasting real economic activity in the euro area: Assessing the impact of qualitative surveys”, by R. Basselier, D. de Antonio Liedo and G. Langenus, *Research series,* December 2017.

332. “Pockets of risk in the Belgian mortgage market: Evidence from the Household Finance and Consumption Survey (HFCS)”, by Ph. Du Caju, *Research series,* December 2017.

333. “The employment consequences of SMEs’ credit constraints in the wake of the great recession” by D. Cornille, F. Rycx and I. Tojerow, *Research series*, December 2017.

334. “Exchange rate movements, firm-level exports and heterogeneity”, by A. Berthou and E. Dhyne, *Research series*, January 2018.

335 “Nonparametric identification of unobserved technological heterogeneity in production”, by L. Cherchye, T. Demuynck, B. De Rock and M. Verschelde, *Research series*, February 2018.

336 “Compositional changes in aggregate productivity in an era of globalisation and financial crisis”, by C. Fuss and A. Theodorakopoulos, *Research series*, February 2018.

337. “Decomposing firm-product appeal: How important is consumer taste?”, by B. Y. Aw, Y. Lee and H. Vandenbussche, *Research series*, March 2018.

338 “Sensitivity of credit risk stress test results: Modelling issues with an application to Belgium”, by P. Van Roy, S. Ferrari and C. Vespro, *Research series*, March 2018.

339. “Paul van Zeeland and the first decade of the US Federal Reserve System: The analysis from a European central banker who was a student of Kemmerer”, by I. Maes and R. Gomez Betancourt, *Research series*, March 2018.

340. “One way to the top: How services boost the demand for goods”, by A. Ariu, F. Mayneris and M. Parenti, *Research series*, March 2018.

341 “Alexandre Lamfalussy and the monetary policy debates among central bankers during the Great Inflation”, by I. Maes and P. Clement, *Research series*, April 2018.

342. “The economic importance of the Belgian ports: Flemish maritime ports, Liège port complex and the port of Brussels – Report 2016”, by F. Coppens, C. Mathys, J.-P. Merckx, P. Ringoot and M. Van Kerckhoven, *Document series*, April 2018.

343. “The unemployment impact of product and labour market regulation: Evidence from European countries”, by C. Piton, *Research series,* June 2018.

344. “Trade and domestic production networks”, by F. Tintelnot, A. Ken Kikkawa, M. Mogstad, E. Dhyne, *Research series*, September 2018.

345. “Review essay: Central banking through the centuries”, by I. Maes, *Research series,* October 2018.

346 “IT and productivity: A firm level analysis”, by E. Dhyne, J. Konings, J. Van den Bosch, S. Vanormelingen, *Research series*, October 2018.

347 “Identifying credit supply shocks with bank-firm data: methods and applications”, by H. Degryse, O. De Jonghe, S. Jakovljević, Klaas Mulier, Glenn Schepens, *Research series*, October 2018.

348 “Can inflation expectations in business or consumer surveys improve inflation forecasts?”, by R. Basselier, D. de Antonio Liedo, J. Jonckheere and G. Langenus, *Research series*, October 2018.

349 “Quantile-based inflation risk models”, by E. Ghysels, L. Iania and J. Striaukas, *Research series*, October 2018.

350 “International food commodity prices and missing (dis)inflation in the euro area”, by G. Peersman, *Research series*, October 2018.

351 “Pipeline pressures and sectoral inflation dynamics”, by F. Smets, J. Tielens and J. Van Hove, *Research series*, October 2018.

352 “Price updating in production networks”, by C. Duprez and G. Magerman, *Research series*, October 2018.

353 “Dominant currencies. How firms choose currency invoicing and why it matters”, by M. Amiti, O. Itskhoki and J. Konings, *Research series*, October 2018.

354 “Endogenous forward guidance”, by B. Chafwehé and R. Oikonomou, R. Priftis and L. Vogel, *Research series*, October 2018.

355 “Is euro area lowflation here to stay? Insights from a time-varying parameter model with survey data”, by A. Stevens and J. Wauters, *Research series*, October 2018.

National Bank of Belgium
Limited liability company
RLP Brussels – Company's number: 0203.201.340
Registered office: boulevard de Berlaimont 14 – BE-1000 Brussels
www.nbb.be

Editor

Jan Smets
Governor of the National Bank of Belgium

Layout: Analysis and Research Group
Cover: NBB AG – Prepress & Image

Published in October 2018