Online Appendix to
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
Ghent University

13th March 2021

This appendix proceeds as follows. Section A.1 describes the data sources. Several subsequent sections provide more details on the estimation procedure. Section A.2 describes how we link the model forecast to the survey inflation forecasts; Section A.3 explains the Gibbs sampler for model M3 that discards survey data; Section A.4 briefly describes how the algorithm changes for the baseline model and model M2, and Section A.5 discusses the convergence results. In Section A.6, we describe additional estimation results. Finally, we provide more details concerning the model evaluation exercises in Section A.7 and discuss our sensitivity analyses to the priors and inflation expectations data in Section A.8.

A.1 Data

The macroeconomic series are sourced from the ECB's Statistical Data Warehouse database (SDW)[1] 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).[2] Specifically, we follow the AWM procedure and backdate price indexes and the unemployment rate using growth rates (Fagan et al., 2001, Annex 2). Before backdating, we seasonally adjust the AWM HICP price index with the X13 procedure using JDemetra+.[3] Details are given in Table A1 below.[4]

*E-mail addresses: arnoud.stevens@nbb.be and joris.wauters@nbb.be

[1] http://sdw.ecb.europa.eu/

[2] https://eabcn.org/page/area-wide-model

[3] https://ec.europa.eu/eurostat/cros/content/software-jdemetra_en

[4] All tables, figures, and several equations in this Online Appendix are labeled as A1, A2, etc. All reference labels without an 'A' thus refer to elements from the main text.

Table A1: 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

We retrieve the inflation expectations series from the ECB's SPF webpage.[5] More specifically, we collect the aggregate probability distributions for inflation at the one-year, two-years, and five-years 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). The resulting series are shown in Figure A1.

Figure A1: SPF inflation expectations data

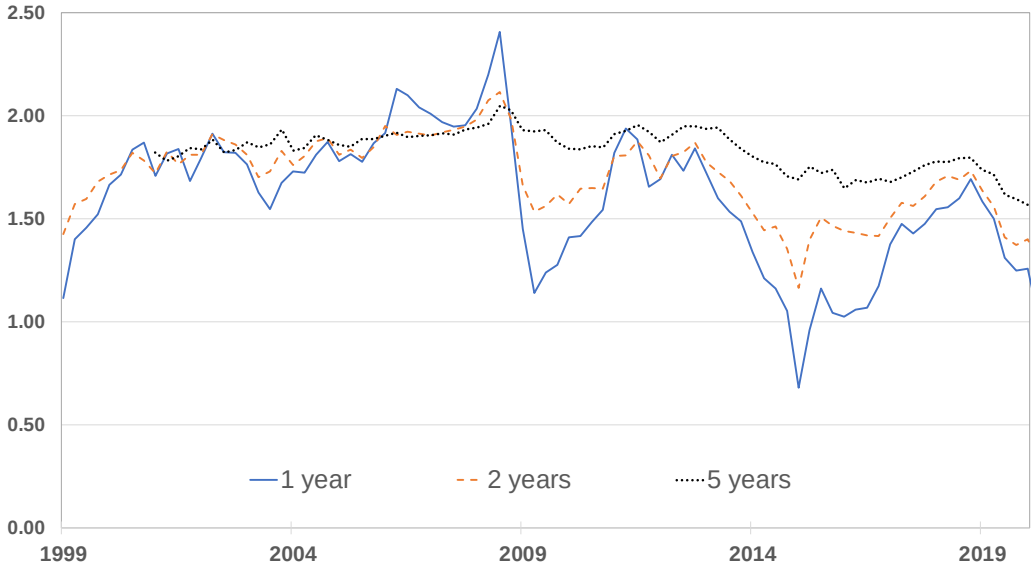

Note: The SPF expectations are the rolling horizon one-year and two-years ahead expectations and the five-years ahead calendar year inflation expectations. In all three cases, we report the computed mean from the aggregate probability distribution for year-on-year headline inflation. Sample: 1999Q1 - 2020Q1.

[5]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 2019Q1, but the SPF inflation expectations data is only available from 1999Q1 onwards. 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 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 time-varying parameters are stacked as, e.g., $\lambda^T = (\lambda_1, ..., \lambda_T)'$, the trends as $\tau^T = (\tau_1', ..., \tau_T')'$, where $\tau_t = (\tau_t^\pi, \tau_t^u, \tau_t^m)'$. We define $\theta_t = \left(\tau_t', \tau_{t-1}', \rho_t^\pi, \lambda_t, \gamma_t, \rho_1^u, \rho_2^u, \rho_1^m, \rho_2^m\right)'$ as a vector that collects the relevant parameters for forecasting inflation using data up to period t. Finally, the detrended (or 'gap') variables are $\tilde{\pi}_t = \pi_t - \tau_t^\pi$, $\tilde{u}_t = u_t - \tau_t^u$, and $\tilde{\pi}_t^m = \pi_t^m - \tau_t^m$; 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 (22) to (23) from the main text link survey expectations to the model forecast. To generate this model-consistent forecast, we rewrite equations (5) to (7) as a vector autoregressive model (VAR):

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

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

$$X_t = F_t X_{t-1} + e_t, \tag{A1}$$

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

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

and the error terms are

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

▷ *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. Since macroeconomic data is released with a lag, these expectations are based on data up to a certain month in the previous quarter. For example, the one-year-ahead expected year-on-year inflation rate in the SPF from Q1 refers to year-on-year inflation in December of the same calendar year. In turn, the Q2 survey refers to the year-on-year inflation in March of the next calendar year, and so on. We follow the 'noise interpretation' of Smets et al. (2014) and consider these series as noisy indicators of year-on-year inflation in the quarter that contains the month of reference. 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, and so on. 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 that SPF data is collected at the start of the quarter, we consider this series to indicate the forecasters' views on the trends and coefficients from the previous quarter. As a result, in our empirical application, we specify that each period t's SPF expectations series for h periods ahead inflation, $\pi_{t+h|t}^{e}$, is informative about the trends and transmission coefficients from quarter $t-1$ (see below).

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

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. Denote π_t^a as the year-on-year inflation in quarter t, then $\pi_t^a = \frac{1}{4}\left(\pi_t + \pi_{t-1} + \pi_{t-2} + \pi_{t-3}\right)$. We use equation (A1) of the detrended model to generate forecasts of the inflation gap $E_t(\tilde{\pi}_{t+h})$. After rewriting this term as $E_t(\pi_{t+h} - \tau_{t+h}^{\pi})$=$E_t(\pi_{t+h})$−$E_t(\tau_{t+h}^{\pi})$, it follows that the point forecast for inflation, $E_t(\pi_{t+h})$, is the sum of the forecasted inflation gap $E_t(\tilde{\pi}_{t+h})$ and expected trend $E_t(\tau_{t+h}^{\pi})$.

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

We generate model-consistent inflation forecasts by iterating equation (A1) 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 set expected future trend inflation equal to τ_t^{π}. 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 $\pi_{t+3|t}^{e}$ with the model-implied three-quarter-ahead forecast of year-on-year inflation using the previous period's data and states, as given by

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

where $\hat{X}_{i|j}$ stands for the forecast of the state vector in period i 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 annualised quarter-on-quarter inflation rate, we construct the year-on-year detrended inflation rate as the four-quarter moving average of the quarterly detrended inflation rates and add trend inflation to generate the overall expected inflation rate. For the two-years- and five-years-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 non-linear 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 require additional attention in the Gibbs sampler, as explained below.

A.3 Gibbs sampler for model M3

This section describes in detail the Gibbs sampling algorithm for model M3 that disables forecast smoothing. In the next section, we briefly explain the different steps for the baseline model, which allows for forecast smoothing, and model M2, which discards the survey data, relative to the algorithm for model M3.

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:

Initialise: We initialise all state variables and time-invariant parameters at their prior means.

Step 1: Draw the error variances from their conditional posterior distribution: $p\left(\sigma^2_u, ..., \sigma^2_{\psi} | Y^T, Z^T, \tau^T, \rho^{\pi,T}, \lambda^T, \gamma^T, \psi^T, \omega^T\right)$. 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 equation errors, 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}, σ^2_{ψ}, and σ^2_{ω} are drawn similarly.

▷ *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, the log-posterior density is

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

with $\Phi(.)$ denoting the cumulative distribution function (CDF) of the standard normal distribution. This is a non-standard density due to the part that goes in curly brackets. Following Chan et al. (2016), we implement a Metropolis-Hastings step with proposal density

$$IG\left(\underline{\nu}_\rho + \frac{T-1}{2}, \underline{\mathrm{S}}_\rho + \frac{1}{2}\sum_2^T \left(\rho_t^\pi - \rho_{t-1}^\pi\right)^2\right),$$

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 $\rho_1^u, \rho_2^u, \rho_1^m, \rho_2^m$ from $p\left(\rho_1^u, \rho_2^u, \rho_1^m, \rho_2^m|Y^T, Z^T, \lambda^T, \rho^{\pi,T}, \gamma^T, \tau^T, \omega^T, \sigma_u^2, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right)$. Note that information for these parameters is given by several measurement equations, i.e. the unemployment gap equation (6), the import price gap equation (7), and the survey data equations (22) to (23). If we disregard the equations related to survey data, we obtain standard regression results (see Chan et al., 2016). However, the inclusion of survey data implies that these parameters enter non-linearly in the function $f_h(\theta_{t-1}, Y^{t-1})$ for each period $t \geq \bar{t}$.

To draw from $\tilde{\rho} \equiv (\rho_1^u, \rho_2^u, \rho_1^m, \rho_2^m)'$, one approach would be to use maximisation routines to determine the mode and hessian of the posterior distribution, and use this information to draw from a proposal normal distribution. We have applied this approach at first, but found similar results at increased computational speed using the following method. We apply an independent Metropolis-Hastings step where the proposal distribution is based on an approximate model where the non-linear functions $f_h(\theta_{t-1}, Y^{t-1})$ are linearised.[6] Denote $\theta_{t-1/\tilde{\rho}}$ as the vector θ_{t-1} without the $\tilde{\rho}$ elements. We can then rewrite $f_h(\theta_{t-1}, Y^{t-1})$ as $f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}, Y_{t-1})$, and take a first order approximation of the latter in the point $\tilde{\rho}_0$:

$$\begin{aligned} f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}, Y^{t-1}) &\approx f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}_0, Y^{t-1}) + \left(\frac{\partial f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}_0, Y^{t-1})}{\partial \tilde{\rho}_0}\right)' (\tilde{\rho} - \tilde{\rho}_0) \\ &\approx \underbrace{f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}_0, Y^{t-1}) - \left(\frac{\partial f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}_0, Y^{t-1})}{\partial \tilde{\rho}_0}\right)' \tilde{\rho}_0}_{c_t^h} + \underbrace{\left(\frac{\partial f_h(\theta_{t-1/\tilde{\rho}}, \tilde{\rho}_0, Y^{t-1})}{\partial \tilde{\rho}_0}\right)'}_{\tilde{x}_{h,t}'} \tilde{\rho}. \end{aligned}$$

In other words, we transform the measurement equations with survey data (22) to (23) into equations which are linear functions of $\tilde{\rho}$. By decomposing the joint likelihood $p(Y^T, Z^T|...)$ as the product $p(Z^T|Y^T, ...) \times p(Y^T|...)$, and using standard regression results, this approximate model delivers closed form solutions for

[6] Canova and Forero (2015, Appendix E) discuss how a non-linear state-space model can be estimated by creating an approximate model through linearisation and then using this approximate model as proposal density.

the conditional posterior of $\tilde{\rho}$, which is a normal distribution. Specifically, we rewrite the unemployment gap equation (6) as

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

where $X^u = \begin{bmatrix} \tilde{u}_0 & \tilde{u}_{-1} & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \tilde{u}_{T-1} & \tilde{u}_{T-2} & 0 & 0 \end{bmatrix}$, with $var(\epsilon^{u,T}) = \sigma_u^2 I_T$, and we rewrite the import price gap equation (7) as

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

where $X^m = \begin{bmatrix} 0 & 0 & \tilde{\pi}_0^m & \tilde{\pi}_{-1}^m \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \tilde{\pi}_{T-1}^m & \tilde{\pi}_{T-2}^m \end{bmatrix}$, with $var(\epsilon^{m,T}) = \sigma_m^2 I_T$. Next, we stack the linearised measurement equations for the survey data as

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

where $C = \left(c_{\bar{t}}', ..., c_T'\right)'$, $c_t = \left(c_t^{h_1}, ..., c_t^{h_n}\right)'$, $\tilde{X} = \begin{pmatrix} \tilde{x}_{\bar{t}} \\ \vdots \\ \tilde{x}_T \end{pmatrix}$, $\tilde{x}_t = \begin{pmatrix} \tilde{x}_{h_1,t}{}' \\ \vdots \\ \tilde{x}_{h_n,t}{}' \end{pmatrix}$, $\epsilon^{z,T} = \left(\epsilon_{\bar{t}}^z{}', ..., \epsilon_T^z{}'\right)'$, $\epsilon_t^z = \left(\epsilon_t^{h_1}, ..., \epsilon_t^{h_n}\right)'$, 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 $p(\tilde{\rho}) \sim N(\underline{\rho}, \underline{V})$ leads to a normal conditional posterior $\left(\tilde{\rho}|Y^T, Z^T, ...\right) \sim N\left(\bar{\tilde{\rho}}, \bar{V}\right)$, where

$$\begin{aligned} \bar{V} &= \left(\underline{V}^{-1} + X^u{}'X^u/\sigma_u^2 + X^m{}'X^m/\sigma_m^2 + \tilde{X}'\Omega_z^{-1}\tilde{X}\right)^{-1} \\ \bar{\tilde{\rho}} &= \bar{V}\left(\underline{V}^{-1}\underline{\tilde{\rho}} + X^u{}'\tilde{u}^T/\sigma_u^2 + X^m{}'\tilde{\pi}^{m,T}/\sigma_m^2 + \tilde{X}'\Omega_z^{-1}\left(Z^T - C\right)\right). \end{aligned}$$

Using these results, we take a candidate draw from a t-distribution with degrees of freedom 10, mean $\bar{\tilde{\rho}}$ and variance $\bar{V}$, to endow the proposal distribution with fatter tails. If the candidate draw is non-stationary, we use the previous draw for $\tilde{\rho}$ as the current draw. However, if the candidate draw meets the stationarity conditions, we accept it with a certain probability according to the Metropolis-Hastings procedure. We select $\tilde{\rho}_0$, the parameter values around which the function f_h is linearised, as the previously accepted draw.[7] In sum, we generate a candidate draw for the persistence parameters from an approximate model which linearises the model forecast functions around the posterior mean from a model that discards the survey data.[8]

[7] We have also experimented with setting $\tilde{\rho}_0$ to the conditional posterior mean that is obtained when the measurement equations for the survey data are ignored: $\tilde{\rho}_0 = \left(\underline{V}^{-1} + X^u{}'X^u/\sigma_u^2 + X^m{}'X^m/\sigma_m^2\right)^{-1}\left(\underline{V}^{-1}\underline{\tilde{\rho}} + X^u{}'\tilde{u}^T/\sigma_u^2 + X^m{}'\tilde{\pi}^{m,T}/\sigma_m^2\right)$. However, results were found to be similar.

[8] To set up the C and $\tilde{X}$ 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 vectorised Matlab function.

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.2, we consider the following augmented state vector:

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

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:

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

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 sides of the above expression are left-multiplied with the matrix $(I_3\, 0_{3\times 3})$ in order to abstract from the survey data equations.

The state equations are given by

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

Step 4: Sample the time-varying transmission coefficients $\rho^{\pi,T}, \lambda^{T}, \gamma^{T}$ for $t = 1, ..., T$. This block has two complications. First, these time-varying parameters also enter non-linearly in the model forecast functions $f_{h_1}\left(\theta_{t-1}, Y^{t-1}\right), ..., f_{h_n}\left(\theta_{t-1}, Y^{t-1}\right)$ 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)[9]:

$$\begin{aligned} p\left(\delta_t|\delta_{j\neq t}, Y^T, Z^T, \theta_{t/\delta_t}, \tau^T, \psi^T, \sigma_\lambda^2, \sigma_\rho^2, \sigma_\gamma^2, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) \propto \\ p\left(Z_{t+1}|Y_t, Y_{t-1}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) \{p\left(Y_t|Y_{t-1}, Y_{t-2}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \psi_t, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) \\ p\left(\delta_{t+1}|\delta_t, Q\right) p\left(\delta_t|\delta_{t-1}, Q\right)\} \frac{1\left(\delta_t \in A\right)}{R\left(\delta_t, Q\right)}. \end{aligned} \tag{A2}$$

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 a proposal probability density function which is based on the parts between brackets in the above expression. This is described in the following steps:

▷ *The likelihood function*

We decompose the joint likelihood $p\left(Z^T, Y^T|\delta_t, \delta_{j\neq t}, ...\right)$ as $p\left(Z^T|Y^T, \delta_t, \delta_{j\neq t}, ...\right)p\left(Y^T|\delta_t, \delta_{j\neq t}, ...\right)$. This product can be further decomposed as $\prod_{i=1}^T p\left(Z_i|Z^{i-1}, Y^T, \delta_t, \delta_{j\neq t}, ...\right) \prod_{i=1}^T p\left(Y_i|Y^{i-1}, \delta_t, \delta_{j\neq t}, ...\right)$. From this joint likelihood we only keep the parts which depend on δ_t, since the rest is absorbed by the integrating constant. As a result, the only two remaining terms are

$$p\left(Z_{t+1}|Y_t, Y_{t-1}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right) p\left(Y_t|Y_{t-1}, Y_{t-2}, \delta_t, \theta_{t/\delta_t}, \tau_t, \tau_{t-1}, \psi_t, \sigma_{h_1}^2, ..., \sigma_{h_n}^2\right).$$

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

$$\begin{aligned} \tilde{\pi}_t &= \left(\tilde{\pi}_{t-1}, \tilde{u}_t, \tilde{\pi}_t^m\right)\delta_t + \epsilon_t^{\pi} \\ &= X_t^{\delta}\delta_t + \epsilon_t^{\pi}, \end{aligned}$$

as 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 under bound restrictions*

Building on the notation and reasoning from Koop and Potter (2011), the restricted prior distribution of δ_t consists of three terms: $p\left(\delta_t|\delta_{t-1}, Q\right)1\left(\delta_t \in A\right)/R\left(\delta_{t-1}, Q\right)$. The first term, $p\left(\delta_t|\delta_{t-1}, Q\right)$, is the unrestricted prior distribution of δ_t. Due to the random walk process, this is a normal distribution centred on δ_{t-1} with variance Q. Since we apply constraints that ρ_t^{π} and λ_t lie within certain bounds, an indicator function $1\left(\delta_t \in A\right)$ is added which equals 1 if δ_t meets the bounds and is zero otherwise. Finally, the term

[9] 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. Furthermore, in each period $t = 1, ..., \bar{t}$-1 for which we do not have survey data, the likelihood term $p\left(Z_{t+1}|...\right)$ from the second line drops out.

$R(\delta_{t-1}, Q)$ indicates the integrating constant that applies to the kernel of the restricted prior distribution $1(\delta_t \in 1)\, p(\delta_t|\delta_{t-1}, Q)$. Intuitively, the integrating constant measures the percentage of random draws from the normal distribution $p(\delta_t|\delta_{t-1}, Q)$ that would fall within the acceptance region (see Koop and Potter, 2011).

Note that in expression (A2) the terms $1(\delta_{t+1} \in A)/R(\delta_{t-1}, Q)$ are missing from the two restricted prior distributions. This is because, conditional on (previously accepted) δ_{t-1} and δ_{t+1} values, these two terms are absorbed by the integrating constant.

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

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

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

Therefore, the integrating constant of the restricted prior is

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

▷ *Combining terms into a proposal density*

Our proposal density combines the terms between brackets from expression (A2):

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

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

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

where

$$\begin{aligned}
\bar{\Sigma}_t &= \left(\underline{\Omega}_0^{-1} + Q^{-1}\right)^{-1} && t = 0 \\
&= \left(X^{\delta}_t{}' X^{\delta}_t/\sigma^2_{\pi,t} + 2Q^{-1}\right)^{-1} && t = 1, ..., T-1 \\
&= \left(X^{\delta}_t{}' X^{\delta}_t/\sigma^2_{\pi,t} + Q^{-1}\right)^{-1} && t = T,
\end{aligned}$$

and

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

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 survey data is applied in the estimation. 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 and ω_t for $t = 1, ..., T$ conditional on all other parameters. The stochastic volatility terms ψ_t and ω_t of the error terms in the measurement equations (5) and (7) for inflation and import price inflation are drawn separately using the multi-move sampler of Omori et al. (2007). Their method is similar to the commonly used approach of Kim et al. (1998) for drawing stochastic volatility states. However, it uses a 10-component mixture distribution to approximate a log χ_1^2 distribution that is more accurate than the 7-component mixture distribution of Kim et al. (1998).

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

A.4 Gibbs sampler for the baseline and M2 models

▷ *Baseline model (includes forecast smoothing)*

Estimation of the baseline model which allows for forecast smoothing (see equations 19 to 20) 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, $\delta_t = (\rho_t^\pi, \lambda_t, \gamma_t, \xi_t)'$, $Q = diag\left(\sigma_\rho^2, \sigma_\lambda^2, \sigma_\gamma^2, \sigma_\xi^2\right)$, and $X_t^\delta = (\tilde{\pi}_{t-1}, \tilde{u}_t, \tilde{\pi}_t^m, 0)$.

Note that this model requires two consecutive survey observations for estimation (the likelihood term for the survey data becomes $p\left(Z_{t+1}|Z_t, Y_t, Y_{t-1}, ...\right)$), which implies that survey data are used in periods $\bar{t} \leq t \leq T-1$.

▷ *Model M2 (discards survey data)*

In model M2 without survey data, the whole procedure becomes more straightforward. Steps 1 and 5 remain the same. Step 2 is now based on a normal conditional posterior distribution, for which acceptance-rejection sampling can be used. Step 3 requires that the correction for the absence of survey data (by left-multiplication) is applied in each period. The expressions for Step 4 require the removal of the likelihood terms related to Z^T.

A.5 Convergence

We executed 250,000 replications of the Gibbs sampler and discarded the first 50,000. We then stored every 20th draw to break the autocorrelation and economise on storage size, after which we have 10,000 posterior draws. We assessed convergence in several ways, first, through casual glances of the trace plots and by verifying that re-estimating the models delivered similar results. Second, we inspected the recursive means of the retained draws at every 20th draw, as shown in Figure A2. The fact that there is little evidence of large fluctuations in the posterior means is taken as evidence in favour of convergence.

Finally, Table A2 reports the inefficiency factors of the estimated parameters. The inefficiency factor is defined as

$$1 + 2\sum_{h=1}^{H} \rho(h),$$

where $\rho(h)$ is the sample autocorrelation at lag h, with H chosen large enough to let the autocorrelation taper off (Chan et al., 2013). As a rule-of-thumb, dividing the amount of Gibbs sampler replications S by the inefficiency factor indicates the amount of independent posterior draws or "effective sample size" (ESS). For instance, in case of independent draws $\sum_{h=1}^{H} \rho(h) \approx 0$ and $ESS = S$. In the table, we report the inefficiency factors for the time-varying parameters by sorting them and then taking the 25th, 50th, and 75th percentiles. The values were calculated from the post-burn-in sample of 200,000 draws (before selecting each 20th draw) and using $H = 100$.

As expected, the inefficiency factors are higher for the coefficients ρ^{π}, λ, ξ, as well as their error variances, due to imposing bounds and the use of a single-move sampler. Nevertheless, the values remain well below 200, which indicates that our large amount of posterior replications guarantees more than 1000 independent draws.

Table A2: Summary inefficiency factors: baseline model

Persistence parameters	ρ_1^u	34.7	ρ_2^u	35.0	ρ_1^m	24.0	ρ_2^m	13.3
Error variances	σ_u^2	2.6	σ_{1y}^2	22.3	σ_{2y}^2	9.3	σ_{5y}^2	2.1
	$\sigma_{\tau\pi}^2$	25.4	$\sigma_{\tau u}^2$	7.5	$\sigma_{\tau m}^2$	14.6		
	σ_{ρ}^2	50.6	σ_{λ}^2	60.1	σ_{γ}^2	27.1	σ_{ξ}^2	58.8
	σ_{ψ}^2	13.4	σ_{ω}^2	11.8				
Time-varying parameters	$\tau_{(25\%)}^{\pi}$	22.4	$\tau_{(50\%)}^{\pi}$	30.9	$\tau_{(75\%)}^{\pi}$	41.9		
	$\tau_{(25\%)}^{u}$	38.0	$\tau_{(50\%)}^{u}$	40.2	$\tau_{(75\%)}^{u}$	41.2		
	$\tau_{(25\%)}^{m}$	12.6	$\tau_{(50\%)}^{m}$	13.7	$\tau_{(75\%)}^{m}$	15.0		
	$\rho_{(25\%)}^{\pi}$	70.7	$\rho_{(50\%)}^{\pi}$	87.4	$\rho_{(75\%)}^{\pi}$	105.7		
	$\lambda_{(25\%)}$	96.2	$\lambda_{(50\%)}$	137.9	$\lambda_{(75\%)}$	152.8		
	$\gamma_{(25\%)}$	28.5	$\gamma_{(50\%)}$	39.8	$\gamma_{(75\%)}$	47.1		
	$\xi_{(25\%)}$	99.1	$\xi_{(50\%)}$	104.1	$\xi_{(75\%)}$	128.1		
	$\psi_{(25\%)}$	9.4	$\psi_{(50\%)}$	10.7	$\psi_{(75\%)}$	13.2		
	$\omega_{(25\%)}$	4.3	$\omega_{(50\%)}$	4.7	$\omega_{(75\%)}$	5.2		

Figure A2: Recursive means of vectorised parameters from the baseline model

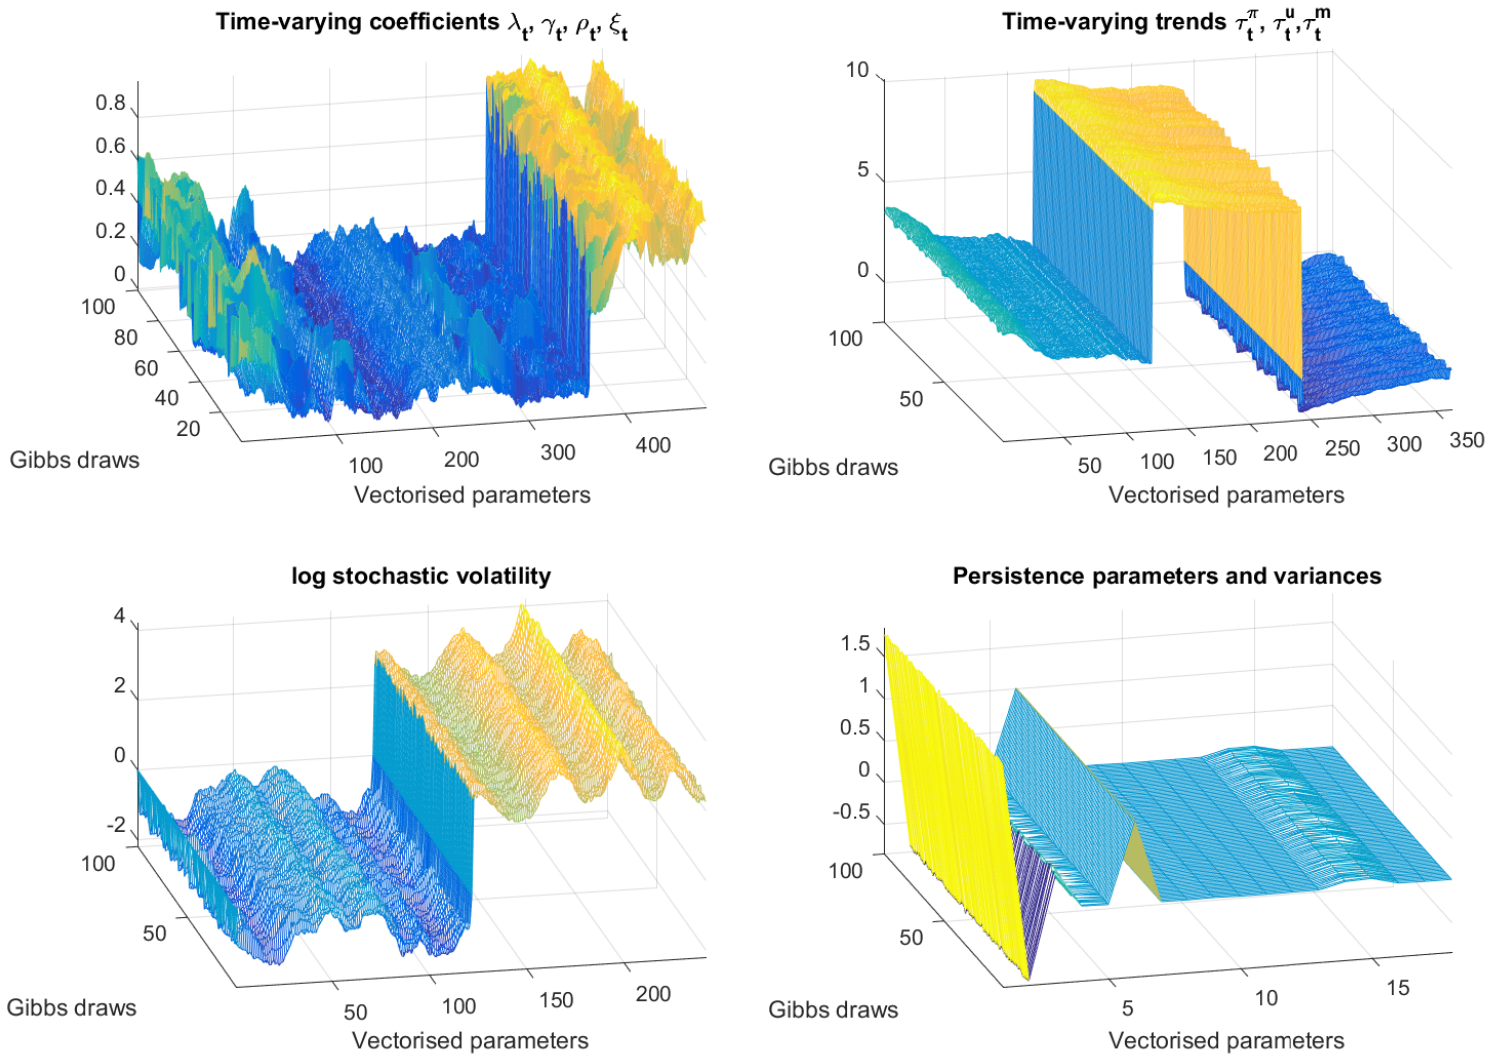

A.6 Additional estimation results

In the main text, we focus on the parameters of key interest and compare the baseline model with forecast smoothing with model M2 that discards survey data. In this section, we report the main estimation results of model M3, followed by the secondary results of the baseline, M2, and M3 models and, finally, a comparison with the estimation results from models M4 and M5.

A.6.1 Model M3: main results

Model M3 features no forecast smoothing but aligns the model forecasts contemporaneously with the survey expectations for inflation. This specification is in line with the approach in Kozicki and Tinsley (2012) and several follow-up studies. Figures A3 and A4 compare the baseline model with model M3 in terms of the estimated trend inflation and trend unemployment rates, as well as the time-varying transmission coefficients.

The main message is that model M3 puts an even stronger role for cyclical factors in explaining the lowflation in the euro area. Compared to the baseline model estimates, trend inflation τ_t^{π} is more stable and somewhat higher at the end of the sample, and the trend unemployment rate τ_t^{u} trends downward more strongly. As a result, this model points out that economic slack remained present during the entire lowflation period. In our view, the highly dynamic trend unemployment rate and very tight uncertainty bands around trend inflation are suspect results. Recall that our evidence against this model is that *i*) model M3 does not feature forecast smoothing ξ_t, while the estimates from the baseline model find a high degree of sluggishness, and *ii*) that model M3 is decidedly rejected against our baseline model in terms of a marginal likelihood comparison.

In terms of the time-varying transmission coefficients, Figure A4 shows that the same qualitative results tend to appear for λ_t, γ_t, and ρ_t^{π}. The only exception is that the Phillips curve slope λ_t tends to steepen again at the end of the sample for model M3, which was not found for the baseline and M2 models.

Figure A3: Estimates of trend inflation τ_t^{π} and trend unemployment τ_t^{u}

(a) Trend inflation τ_t^{π} estimates

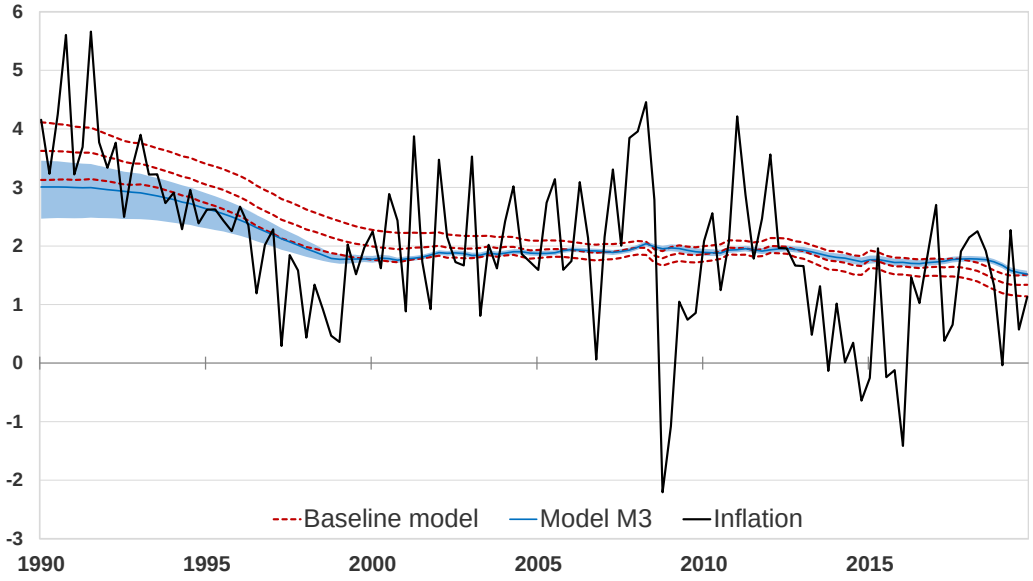

(b) Trend unemployment τ_t^{u} estimates

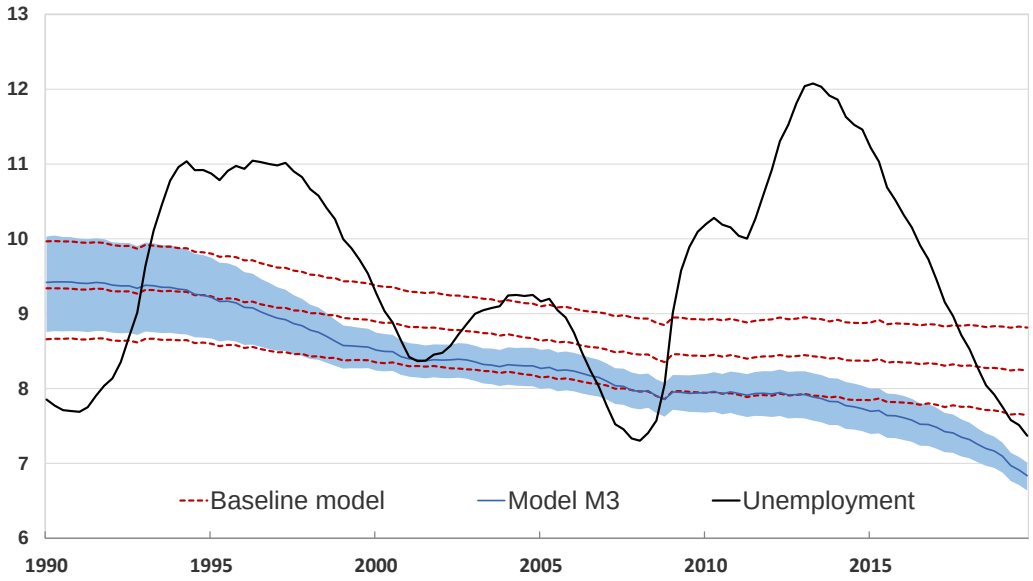

Note: The figures show estimates of trend inflation (panel a) and estimates trend unemployment (panel b). Trend estimates from the baseline model (red dashed lines) are compared with those from model M3 that disables forecast smoothing (blue shaded area). The bands depict the median and 68% credible set.

Figure A4: Time-varying transmission coefficients: Baseline model and model M3

(a) $-\lambda_t$: Phillips curve slope

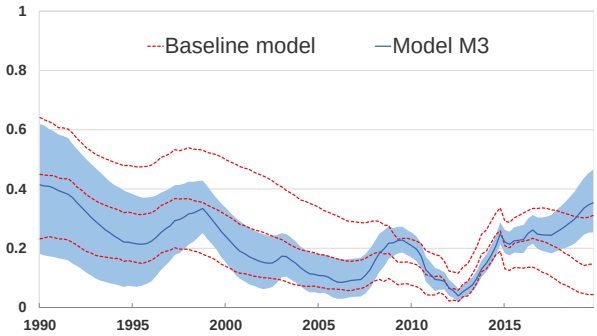

(b) γ_t: Import price coefficient

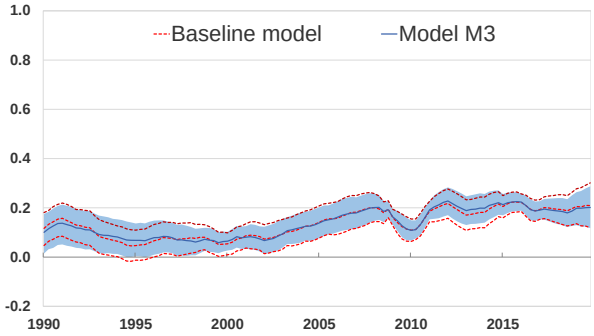

(c) ρ_t^{π}: Inflation gap persistence

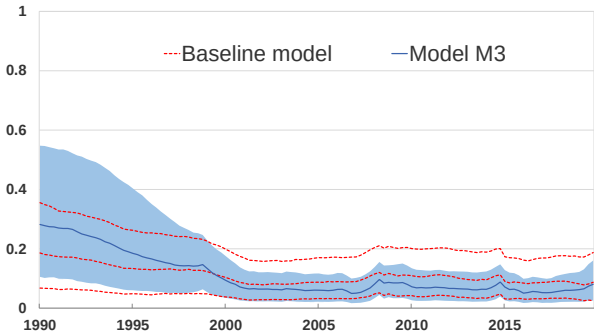

Note: The figures show the estimated time-varying Phillips curve slope λ_t (panel a), import price coefficient γ_t (panel b), and inflation gap persistence ρ_t^{π} (panel c). The Phillips curve slope is shown as $-\lambda_t$ to facilitate comparison (an increase reflects a steepening of the slope). Estimates from the baseline model (red dashed lines) are compared with those from model M3 that disables forecast smoothing (blue shaded area). The bands depict the median and 68% credible set.

A.6.2 Secondary results from baseline, M2 and M3 models

Trend relative import price inflation τ_t^m: Instead of demeaning relative import price inflation π_t^m, we also estimate a trend τ_t^m in order to capture any slow-moving trend changes. Figure A5 shows that the trend estimates evolve similarly across models; they creep upwards from negative figures in 1990 to values close to zero at the end of the sample.

Error variances: Figures A6 and A7 show the evolution of $e^{\psi_t/2}$ and $e^{\omega_t/2}$, which are, respectively, the time-varying standard deviations of the error terms ϵ_t^π and ϵ_t^m in the inflation gap equation (5) and import price equation (7). The volatility of residual shocks to headline inflation ($e^{\psi_t/2}$) trended broadly upward in the period from about 1995 to 2008, and then declined afterwards. For import price inflation, the volatility of shocks ($e^{\omega_t/2}$) undergoes stronger fluctuations and peaks in 2008. The estimates across the baseline model and models M2 and M3, as shown by the medians, are very similar.

Table A3 shows the summary statistics (median, 16th and 84th percentiles) from the posterior distributions of the error variances that remain constant over time. Overall, the estimates tend to be similar across the three models. 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-years ahead and five-years ahead expectations, are relatively small. Of all time-varying parameters, the variances are the largest in the stochastic volatility series (σ_ψ^2 and σ_ω^2), followed by trend inflation ($\sigma_{\tau\pi}^2$) and trend import price inflation ($\sigma_{\tau m}^2$).

Persistence parameters: Table A4 shows the estimated persistence parameters ρ_1^u, ρ_2^u, ρ_1^m, and ρ_2^m. The euro area unemployment gap is found to be highly persistent, with the sum of AR coefficients close to 1. This sum is only about half as large for the more volatile import price inflation gap.

Figure A5: Relative import price inflation π_t^m and its trend τ_t^m

(a) Baseline model and model M2

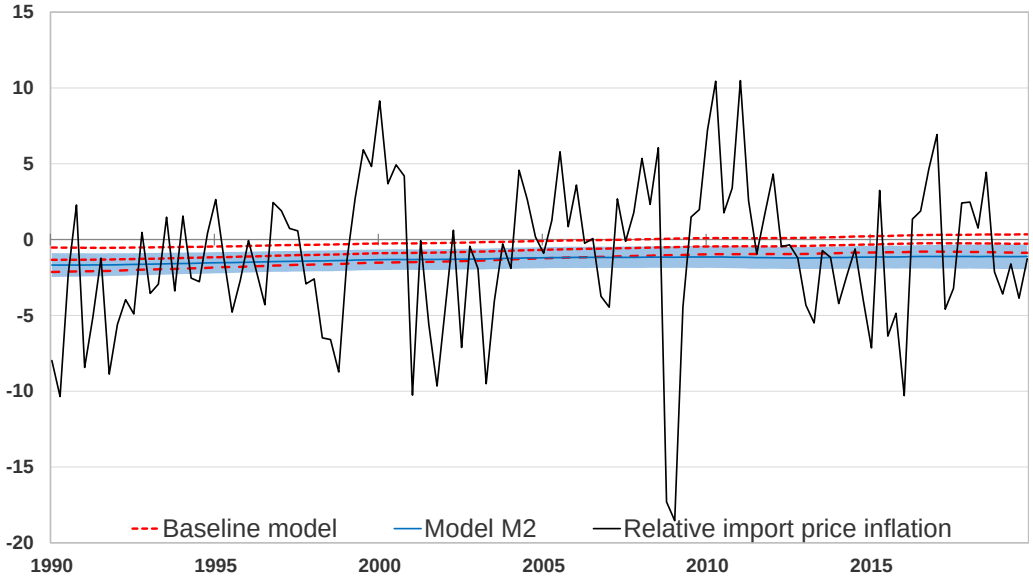

(b) Baseline model and model M3

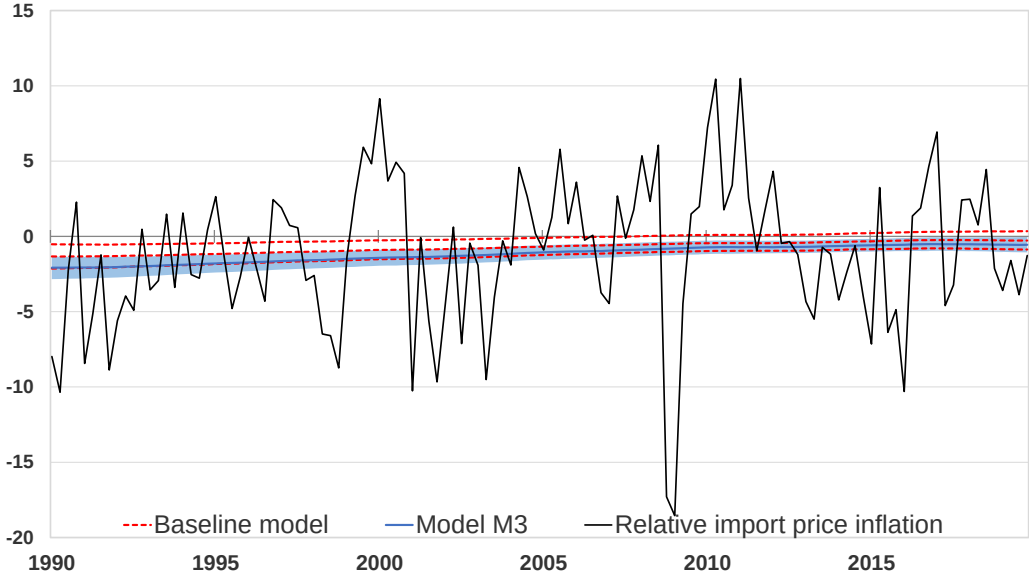

Note: The figures show the inflation rate of the relative price of imports (black line) and estimates of its trend. Trend estimates from the baseline model (red dashed lines) are compared with those from model M2 that excludes survey data in panel (a), and with those from model M3 that disables forecast smoothing in panel (b) (both as blue shaded areas). The bands depict the median and 68% credible set.

Figure A6: Time-varying stochastic volatility of inflation gap equation error

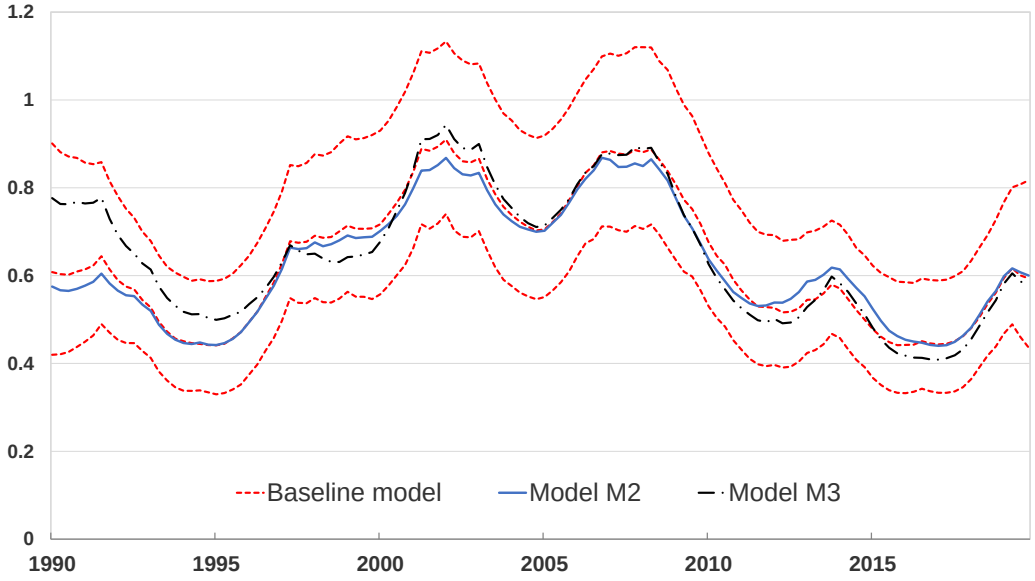

Note: Estimates for the time-varying standard deviation of ϵ_t^{π}, calculated as $e^{\psi_t/2}$, are shown for the baseline model (median and 68% credible sets; red dashed lines), a well as the median estimates for model M2 that discards survey data, and model M3 that disables forecast smoothing.

Figure A7: Time-varying stochastic volatility of import price inflation error

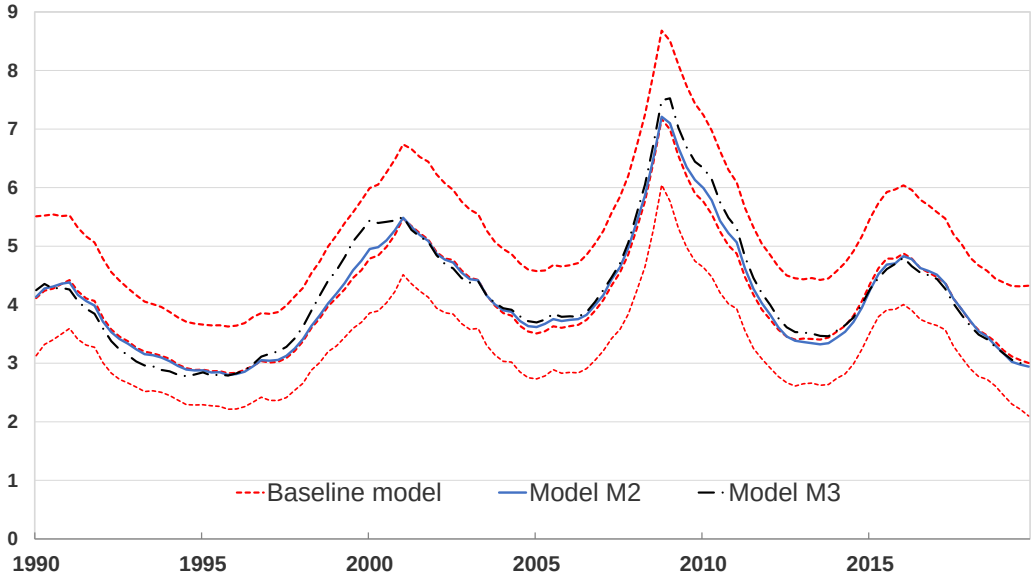

Note: Estimates for the time-varying standard deviation of ϵ_t^{m}, calculated as $e^{\omega_t/2}$, are shown for the baseline model (median and 68% credible sets; red dashed lines), a well as the median estimates for model M2 that discards survey data, and model M3 that disables forecast smoothing.

Table A3: Summary statistics of the error variances' posterior distribution

Parameter	Baseline model	Model M2	Model M3
σ_u^2	0.025	0.025	0.028
	(0.022, 0.029)	(0.022, 0.028)	(0.025, 0.033)
$\sigma_{\tau\pi}^2$	0.010	0.012	0.005
	(0.007, 0.013)	(0.009, 0.017)	(0.004, 0.006)
$\sigma_{\tau u}^2$	0.005	0.005	0.006
	(0.004 0.006)	(0.004, 0.006)	(0.005, 0.008)
$\sigma_{\tau m}^2$	0.009	0.009	0.010
	(0.007, 0.013)	(0.007, 0.012)	(0.007, 0.014)
σ_ρ^2	0.002	0.002	0.001
	(0.001, 0.002)	(0.001, 0.002)	(0.001, 0.002)
σ_λ^2	0.002	0.002	0.002
	(0.001, 0.003)	(0.001, 0.003)	(0.001, 0.002)
σ_γ^2	0.001	0.001	0.001
	(0.001, 0.002)	(0.001, 0.002)	(0.001, 0.002)
σ_ψ^2	0.079	0.079	0.082
	(0.060, 0.106)	(0.060, 0.106)	(0.062, 0.110)
σ_ω^2	0.082	0.083	0.083
	(0.063, 0.109)	(0.063, 0.110)	(0.064, 0.110)
σ_ξ^2	0.002		
	(0.001, 0.003)		
σ_{1y}^2	0.013		0.010
	(0.011, 0.015)		(0.008, 0.012)
σ_{2y}^2	0.007		0.006
	(0.006, 0.008)		(0.005, 0.007)
σ_{5y}^2	0.005		0.006
	(0.004, 0.006)		(0.005, 0.007)

Note: The summary statistics shown are the posterior median and, between round brackets, the 16th and 84th percentiles of the posterior distribution. The parameters σ_{1y}^2, σ_{2y}^2 and σ_{5y}^2 denote, respectively, the error variances from the measurement equations for the one-year-ahead, two-years-ahead, and five-years-ahead inflation expectations.

Table A4: Summary statistics of the persistence parameters

Parameter	Baseline model	Model M2	Model M3
ρ_1^u	1.782	1.823	1.659
	(1.727, 1.835)	(1.755, 1.891)	(1.613, 1.704)
ρ_2^u	-0.805	-0.842	-0.696
	(-0.856, -0.751)	(-0.909, -0.775)	(-0.739, -0.650)
ρ_1^m	0.428	0.381	0.142
	(0.340, 0.519)	(0.280, 0.479)	(0.080, 0.202)
ρ_2^m	0.030	0.026	0.061
	(-0.045, 0.105)	(-0.068, 0.119)	(0.016, 0.106)

Note: The summary statistics shown are the posterior median and, between round brackets, the 16th and 84th percentiles of the posterior distribution of the persistence parameters ρ_1^u and ρ_2^u related to the unemployment gap equation, and ρ_1^m and ρ_2^m related to the relative import price gap equation.

A.6.3 Comparison with models M4 and M5

In our estimations, we explained the term structure of expectations using models that allow for time-varying trends and transmission coefficients. The typical approach in the literature, however, is to use constant transmission coefficients in the inflation gap equation and to include only long-term survey expectations. Therefore, we perform two comparison checks to see whether the results change when we simplify our modelling approach. The first simplification, applied in model M4, is to use only long-term inflation expectations data at the five-years ahead horizon, and discard the equations related to the one- and two-years ahead forecasts. The second, applied in model M5, is to only allow for time variation in the trends. Model M5 thus features constant transmission coefficients (ρ^π, λ and γ), a constant smoothing coefficient (ξ), and static variances of the inflation gap residual ($\epsilon_t^\pi \sim N(0, \sigma_\pi^2)$, $\epsilon_t^m \sim N(0, \sigma_m^2)$).

In Section 6 of the main text, we assess the performance of the baseline model against these two simplifications in terms of inflation forecast accuracy and marginal likelihood evaluation. In this section, we investigate whether models M4 and M5 lead to qualitatively different results compared to the baseline model. More specifically, the next paragraphs discuss the implications for the estimated smoothing coefficient ξ_t, trend inflation τ_t^π, and the natural rate of unemployment τ_t^u.

Smoothing coefficient ξ_t: Model M5 with constant transmission coefficients delivers a median estimate of $\xi = 0.81$ (Table A5). This value fits well with the average degree of forecast smoothing recorded over the sample in the baseline model. Interestingly, for model M4 that uses 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. Indeed, the short end of the expectations curve experienced sharp changes in the later years of our sample, whereas the long end remained relatively more stable (Figure A1).

Table A5: Summary statistics of selected parameters of model M5

Model M5							
ρ_1^u	1.783	ρ_1^m	0.483	ρ^π	0.104	λ	0.201
	(1.721, 1.839)		(0.399, 0.564)		(0.041, 0.175)		(0.146, 0.273)
ρ_2^u	-0.801	ρ_2^m	0.012	ξ	0.806	γ	0.144
	(-0.857, -0.741)		(-0.061, 0.084)		(0.752, 0.859)		(0.128, 0.160)
σ_π^2	0.585	σ_m^2	17.207				
	(0.516, 0.670)		(15.294, 19.502)				

Note: The summary statistics shown are the posterior median and, between round brackets, the 16th and 84th percentiles of the posterior distribution. The table contains selected parameters from model M5, which uses constant transmission coefficients, smoothing parameter, and error variances.

Trend inflation τ_t^π: Panel (a) of Figure A8 shows the posterior median trend inflation estimates for the baseline model, model M2, and the two variants M4 and M5. The baseline and M5 models deliver similar trend inflation estimates during most of the protracted period of below-target inflation since 2013. At the end of the sample, these trend inflation rates are close to 30 basis points above the trend inflation from model M2 that excludes survey data. Model M4, which uses only long-term inflation expectations, points to a stronger weakening in trend inflation in recent years compared to the baseline model. This result might seem counter-intuitive. After all, long-term inflation expectations have been higher and more stable than shorter-term inflation expectations, and they should, in isolation, be the most informative about the trend. However, compared to the baseline model, the relatively higher ξ_t estimates in model M4 imply a slower adjustment of survey expectations to the underlying model forecasts. This mechanism allows for a more persistent deviation between long-term survey forecasts and trend inflation, which explains the lower trend inflation for model M4.

Natural rate of unemployment τ_t^u: Panel (b) of Figure A8 shows that, out of the four models, the baseline model tends to deliver the lowest natural rate of unemployment estimates, especially at the end of the sample. Again, model M5 delivers similar estimates. The natural rates of the models M2 and M4 form a second group with higher values over the sample period.

All told, a key message from this exercise is that including only long-term inflation expectations in the model can lead to different views regarding the trends and, therefore, the relative roles of permanent and cyclical factors in explaining the lowflation. In fact, Model M4 finds lower trend inflation and a higher natural rate, which points to less economic slack over the lowflation period and a more critical role for permanent effects. However, the more accurate forecasts of the baseline model than its M4 variant, as reported in Section 6 of the main text, suggests that short-term survey forecasts carry essential additional information about the inflation trend and cyclical factors —a point made previously by Kozicki and Tinsley (2012).

Figure A8: Model comparisons: trend inflation τ_t^{π} and natural rate of unemployment τ_t^{u}

(a) Trend inflation τ_t^{π} estimates

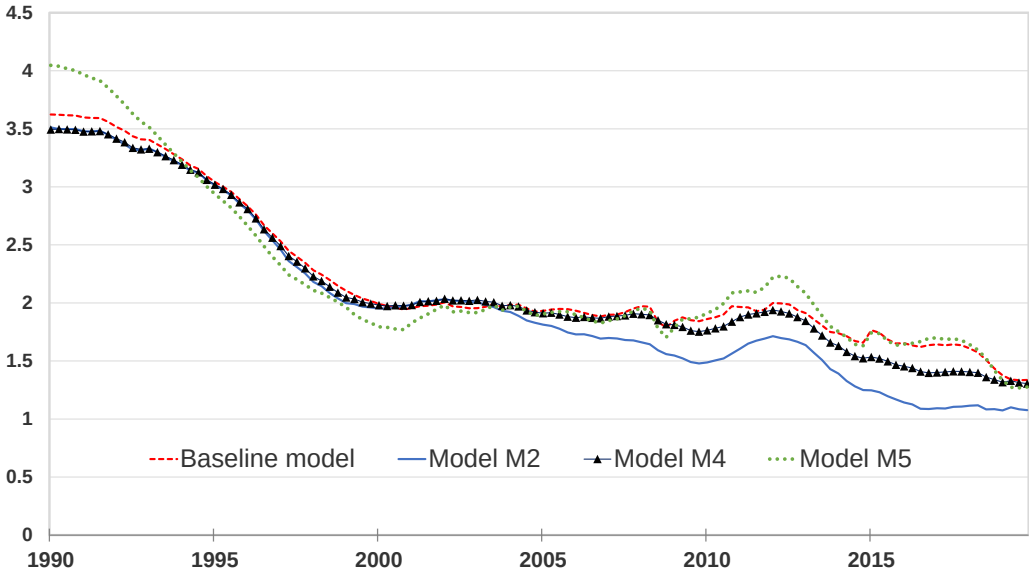

(b) Unemployment rate and estimated natural rates τ_t^{u}

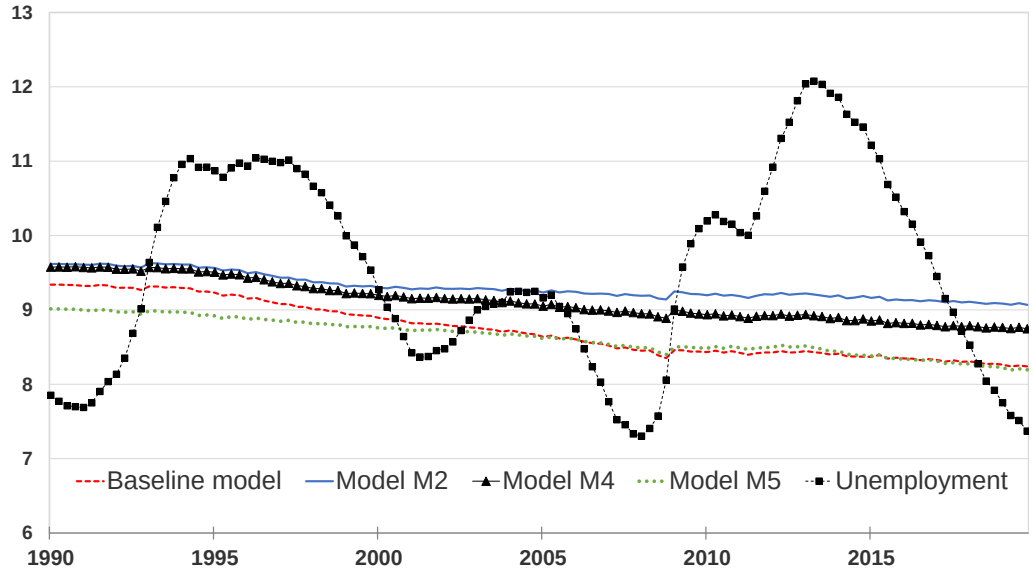

Note: The figures compare the median trend estimates from the different model variants. Panel (a) shows trend inflation τ_t^{π} estimates and panel (b) shows the natural rate of unemployment τ_t^{u} estimates. Models M4 and M5 refer to the model variants with (i) only long-term expectations and (ii) constant transmission parameters, smoothing coefficient, and error variances, respectively.

A.7 Model evaluation exercises: extra details

Forecasting details

Table 1 in the main text includes out-of-sample forecasting results from the Stock and Watson (2007) trend inflation model with stochastic volatility (UC-SV) and a random walk model. Here, we provide details concerning these models.

The UC-SV model is defined as

$$\pi_t = \tau_t^{\pi} + \eta_t$$
$$\tau_t^{\pi} = \tau_{t-1}^{\pi} + \epsilon_t,$$

where $\eta_t \sim N(0, \sigma_{\eta,t}^2)$, $\epsilon_t \sim N(0, \sigma_{\epsilon,t}^2)$. The error terms follow stochastic volatility processes:

$$\ln\sigma_{\eta,t}^2 = \ln\sigma_{\eta,t-1}^2 + \nu_{\eta,t}$$
$$\ln\sigma_{\epsilon,t}^2 = \ln\sigma_{\epsilon,t-1}^2 + \nu_{\epsilon,t},$$

and $\nu_{\eta,t} \sim N(0, \sigma_{\nu\eta}^2)$, $\epsilon_t \sim N(0, \sigma_{\nu\epsilon}^2)$. Although Stock and Watson calibrate $\sigma_{\nu\eta}^2 = \sigma_{\nu\epsilon}^2 = 0.2^2$, we set relatively uninformative priors for these variances and draw from their posterior distribution. More specifically, we set inverse-Gamma priors for $\sigma_{\nu\eta}^2$ and $\sigma_{\nu\eta}^2$ with 10 degrees of freedom and scale parameters such that their means are $E(\sigma_{\nu\eta}^2) = E(\sigma_{\nu\epsilon}^2) = 0.2^2$. We set diffuse initial values as $\tau_1^{\pi} \sim N(0, 9)$, $\ln\sigma_{\eta,1}^2 \sim N(0, 9)$, and $\ln\sigma_{\epsilon,1}^2 \sim N(0, 9)$. For each estimation, we discard the initial 10,000 burn-in draws from the MCMC sampler, and retain each 10th draw of the 50,000 post-burn-in sample. Our forecasts are generated by simulating future data using the remaining 5,000 MCMC sampler draws.

Our random walk forecast is generated by estimating the model

$$\pi_t = \pi_{t-1} + \epsilon_t,$$

where $\epsilon_t \sim N(0, \sigma^2)$. We set a flat (or uninformative) prior for σ^2, draw 100,000 draws from the posterior sampler for each estimation, and simulate future inflation data using the MCMC sampler output.

All models simulate future values of annualised quarter-on-quarter inflation. To create year-on-year inflation, we take the four-quarter moving average. Finally, we fit a normal distribution to the posterior draws from the MCMC sampler and evaluate the actual realisation of inflation at this probability density function.

Marginal likelihood estimation details

Let $y^{1:T} = (y_1', ..., y_T')'$ denote all data up to time T. Our aim is to estimate the marginal likelihood of the data given model M_k as $p(y^{1:T}|M_k)$. This computation is not trivial for models with many latent state variables, such as those estimated in our paper. One often applied approach is the modified harmonic mean estimator. However, Chan and Grant (2015) show that this approach leads to seriously biased estimates of the marginal likelihood in a model with many latent states. Moreover, the numerical standard errors can be so severely underestimated that the true marginal likelihood is not within the uncertainty bands. Therefore,

we opted for using a brute-force approach, which works as follows. We decompose the marginal likelihood as

$$p(y^{1:T}|M_k) = p(y_1|M_k)\prod_{t=1}^{T} p(y_{t+1}|y^{1:t}, M_k),$$

where $p(y_{t+1}|y^{1:t}, M_k)$ denotes the predictive likelihood. To evaluate each predictive likelihood element, a separate MCMC estimation is done and, based on the estimated parameters, a one-step-ahead forecast is evaluated at the realised values. More formally, collect model M_k's constant parameters in the vector β, and the time-varying parameters in $\theta^{1:T}$, and rewrite the predictive likelihood as (see the appendix of Lakdawala, 2015):

$$\begin{aligned}
p(y_{t+1}|y^{1:t}, M_k) &= \int\int p(y_{t+1}, \beta, \theta^{1:t+1}|y^{1:t}, M_k) d\beta d\theta^{1:t+1} \\
&= \int\int p(y_{t+1}|\beta, \theta^{1:t+1}, y^{1:t}, M_k) p(\beta, \theta^{1:t+1}|y^{1:t}, M_k) d\beta d\theta^{1:t+1} \\
&= \int\int p(y_{t+1}|\beta, \theta^{1:t+1}, y^{1:t}, M_k) p(\theta_{t+1}|\beta, \theta^{1:t}, y^{1:t}, M_k) p(\beta, \theta^{1:t}|y^{1:t}, M_k) d\beta d\theta^{1:t+1}.
\end{aligned}$$

In the last line, the term $p(\beta, \theta^{1:t}|y^{1:t}, M_k)$ refers to the posterior distribution using data up to time t. Given this distribution, a simulation is made for states in period $t+1$ using $p(\theta_{t+1}|\beta, \theta^{1:t}, y^{1:t}, M_k)$ (see also Cogley et al., 2005), such that the realized data from $t+1$ is evaluated using $p(y_{t+1}|\beta, \theta^{1:t+1}, y^{1:t}, M_k)$. Using the posterior draws from the MCMC sampler, we evaluate the predictive density for period $t+1$ as

$$\frac{1}{S}\sum_{i=1}^{S} p(y_{t+1}|\beta^{(i)}, \theta^{(i),1:t+1}, y^{1:t}, M_k),$$

where superscript (i) refers to draw i from the MCMC sampler's S total retained draws. This procedure is repeated by re-estimating the model $T-1$ times in order to build the marginal likelihood $p(y^{1:T}|M_k)$. We used the output from the pseudo-out-of-sample forecasting exercise from Section 6.1 for this purpose.[10] In particular, as the SPF data only start in 1999, the baseline, M2 and M3 models are the same from 1990 until 1998, which makes it unnecessary to repeat MCMC runs for each model separately for this subsample.

When compare the models against restricted versions with constant parameters, we use the following proper priors: $\lambda \sim TN(-1, 0; 0.5, 5)$, $\gamma \sim N(0.5, 5)$, $\rho^{\pi} \sim TN(0, 1; 0.5, 5)$, $\sigma_{\pi}^2 \sim IG(10, 9)$, and $\sigma_{\pi}^2 \sim IG(10, 9)$.

A.8 Prior and data sensitivity analyses

Prior sensitivity analysis

Since our baseline model is an extension of the model of Chan et al. (2016), we strongly follow their prior settings in the main text. They use weakly informative priors that favour smooth variation in the time-varying parameters. This strategy is necessary for flexible models such as those in this paper to avoid erratic estimates, but without being dogmatically imposing. In this section, we discuss the sensitivity of our state variable estimates to different prior settings.

[10] As in Chan et al. (2018), we discarded the initial four predictive likelihoods to reduce sensitivity to the priors. Moreover, to compare the baseline model with model M3, we discarded the initial periods for the SPF data because the baseline model uses a lagged SPF expectation in its measurement equation, while model M3 does not.

The initial values of our state variables have relatively large variances that make them uninformative (or diffuse). Therefore, we do not vary these settings. To verify the prior sensitivity of our estimated state variables (τ^{π}, τ^{u}, τ^{m}, ρ^{π}, λ, γ, ξ, ψ, ω), we re-estimated the baseline model four times under different settings for the error variances of the state variables. By adjusting the scale parameters of their inverse-Gamma priors, we considered the following four cases: *i*) Very loose: baseline prior means $\times$ 10; *ii*) Looser: baseline prior means $\times$ 4; *iii*) Tighter: baseline prior means / 4; *iv*) Very tight: baseline prior means / 10. Settings *i* and *ii* thus favour more time variation in the states, while *iii* and *iv* are more restrictive against time variation. To get a feel for how large these changes are, note first that the priors in our main text imply that, with 95% probability, the change in trend inflation $\tau_t^{\pi} - \tau_{t-1}^{\pi}$ lies in the (-0.2;0.2) interval. By contrast, for the four sensitivity scenarios this interval changes to, respectively, (-0.62;0.62), (-0.39,0.39), (-0.1,0.1), and (-0.06;0.06). The last setting particularly stacks the deck against time variation, as its 95% probability bands imply that it would take at least 16 years for trend inflation to decline by 50 basis points.

Overall, we find that the main results hold up well against these prior perturbations. Under case *i* and *ii*, more short-term fluctuations become apparent in the evolutions of the state variables. Under case *iii*, by contrast, the estimates become smoother but still show economically relevant time-variation. By contrast, the "Very tight" case *iv* kills the time variation in several parameters. Therefore, very conservative priors towards time variation are needed to end up with flat parameter evolutions. As an illustration, Figures A9 to A11 below show the estimated state variables of trend inflation τ^{π}, Phillips curve slope λ, and stochastic volatility ψ under the different prior settings. (Other charts are available upon request.)

Figure A9: Prior sensitivity analysis for trend inflation τ_t^π

(a) Very loose prior setting

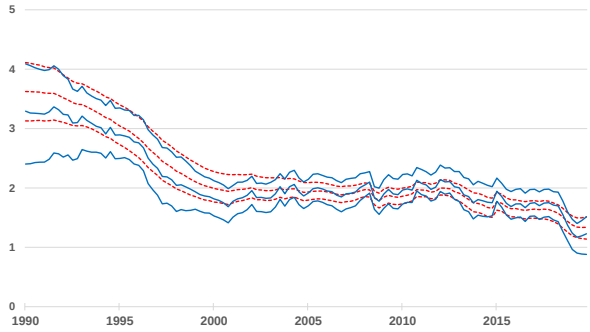

(b) Looser prior setting

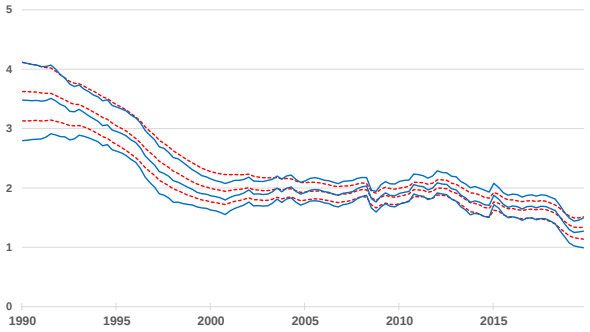

(c) Tighter prior setting

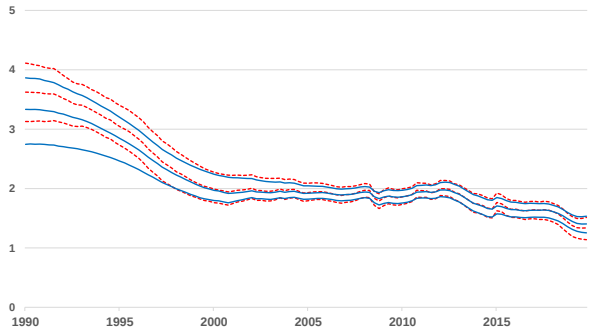

(d) Very tight prior setting

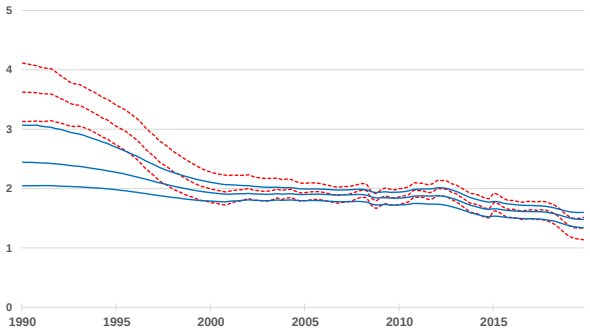

Note: The figure shows four panels where the baseline model estimates for trend inflation τ_t^π (red dashed lines) are compared against estimates under different prior settings for the error variances of the state variables (blue full lines). The bands depict the median and 68% credible set. The inverse-Gamma prior settings for the error variances are defined as follows: *i*) Very loose: baseline prior means $\times$ 10; *ii*) Looser: baseline prior means $\times$ 4; *iii*) Tighter: baseline prior means / 4; *iv*) Very tight: baseline prior means / 10.

Figure A10: Prior sensitivity analysis for the Phillips curve slope λ_t

(a) Very loose prior setting

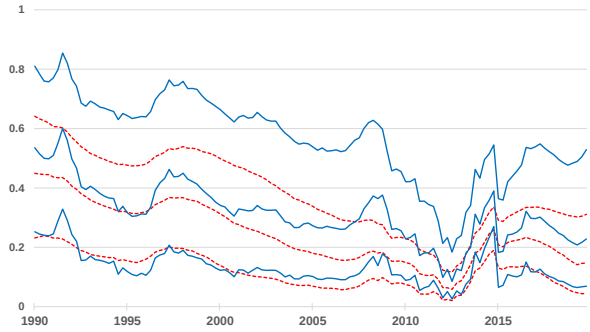

(b) Looser prior setting

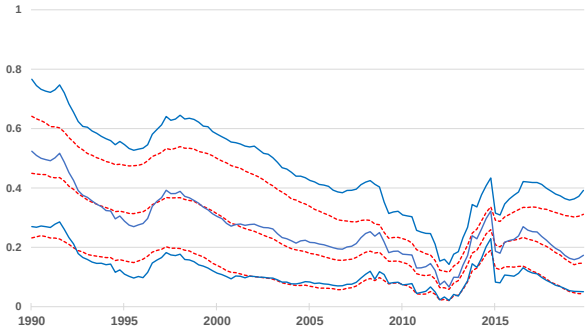

(c) Tighter prior setting

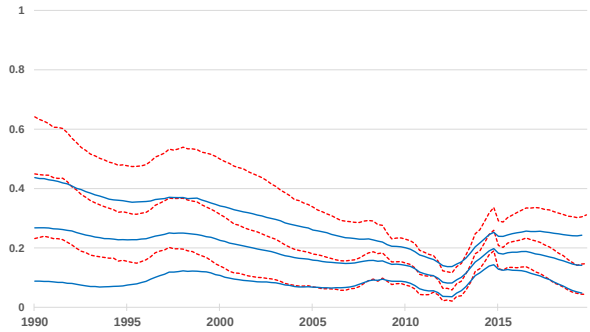

(d) Very tight prior setting

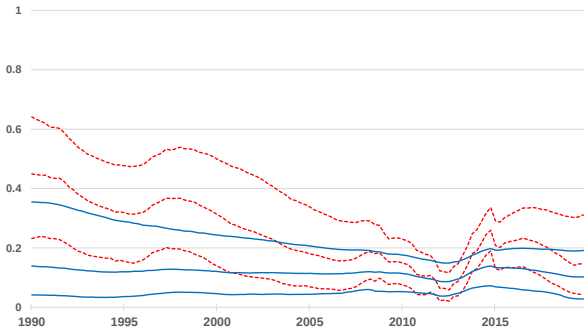

Note: The figure shows four panels where the baseline model estimates for Phillips curve slope λ_t (red dashed lines) are compared against estimates under different prior settings for the error variances of the state variables (blue full lines). The bands depict the median and 68% credible set. The inverse-Gamma prior settings for the error variances are defined as follows: *i*) Very loose: baseline prior means $\times$ 10; *ii*) Looser: baseline prior means $\times$ 4; *iii*) Tighter: baseline prior means / 4; *iv*) Very tight: baseline prior means / 10.

Figure A11: Prior sensitivity analysis for stochastic volatility ψ_t

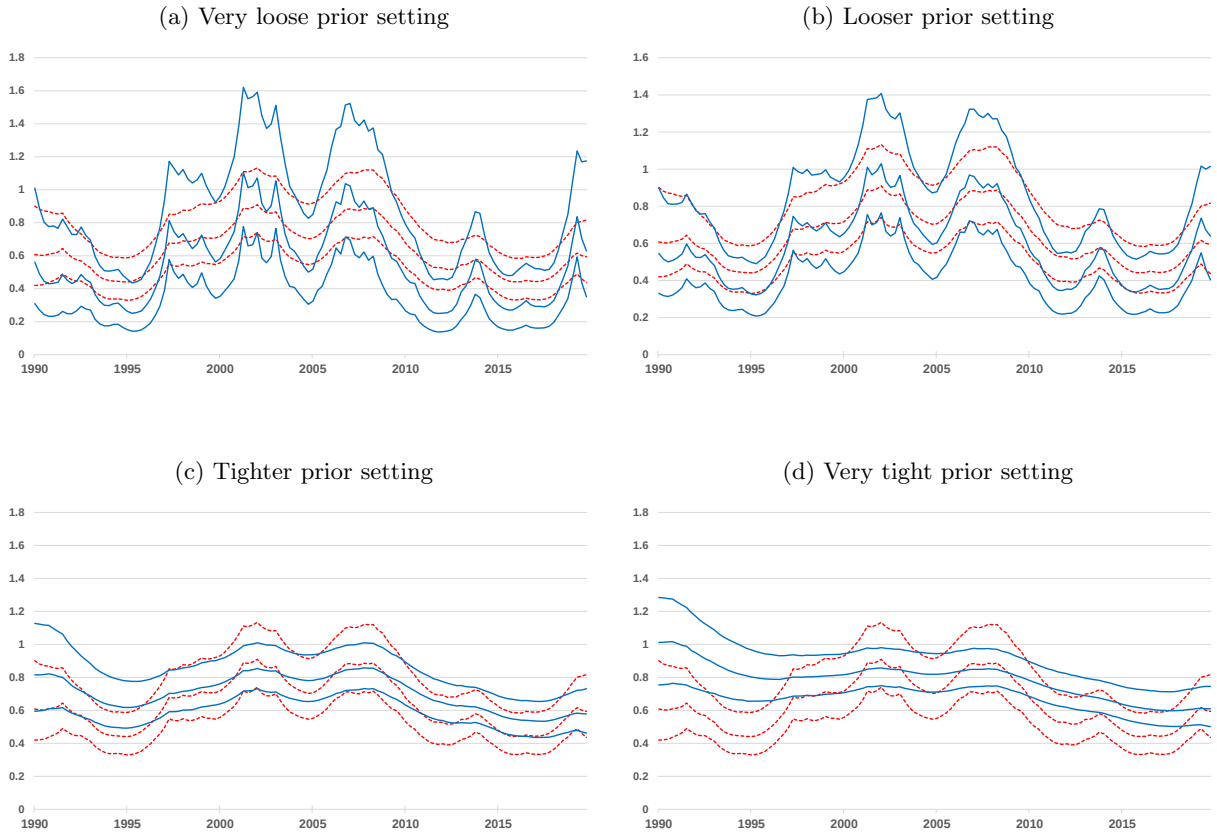

Note: The figure shows four panels where the baseline model estimates for stochastic volatility ψ_t (red dashed lines) are compared against estimates under different prior settings for the error variances of the state variables (blue full lines). The bands depict the median and 68% credible set. The inverse-Gamma prior settings for the error variances are defined as follows: *i*) Very loose: baseline prior means × 10; *ii*) Looser: baseline prior means × 4; *iii*) Tighter: baseline prior means / 4; *iv*) Very tight: baseline prior means / 10.

Sensitivity analysis to survey expectations data:

We have assessed the robustness of our baseline model estimates to using alternative inflation expectations series. We did not use market-based inflation expectations measures because they are strongly influenced by other factors such as risk-premia (see Section 1 of the main text). Moreover, the sample for such series only starts in 2005 for the euro area. We have, however, verified the sensitivity of our estimates to two other survey-based inflation expectations series, namely the SPF average point forecasts, and the Consensus Economics (CE) inflation expectations series. Concerning the former, recall that our main estimates use the (self-computed) mean of the SPF inflation expectations distributions at the one-, two-, and five-years ahead horizons. These series are very similar to the average point forecasts from SPF respondents at the first two horizons. However, at the long-term five-years ahead horizon, the average point forecast is, on average, ten basis points above the self-computed mean in the 2013-2019 period. Nevertheless, panel (a) of Figure A12 shows that the estimates of trend inflation are essentially the same when we use the average point forecasts in the estimation of the baseline model. Similarly, panel (b) of Figure A12 shows that trend inflation estimates are highly similar when we use inflation expectations from CE instead. The same follows for trend unemployment estimates τ_t^u (not shown). Hence, we conclude that our findings hold up well with

other inflation expectations measures.

Some remarks are in order concerning our implementation of the CE data. The CE dataset reports calendar year (or "fixed event") forecasts, which are different from the SPF rolling-horizon forecasts that we use. In the literature, the CE fixed-event forecasts have been used to approximate rolling-horizon forecasts. For instance, the one-year ahead rolling horizon expectation for inflation is constructed as a weighted average of the current and next calendar year inflation forecasts. (Similarly, the two-years ahead forecast is made as a weighted average of next two consecutive calendar years.) These weights will vary depending on the quarter (or month) of the year when the expectations are made. Dovern and Fritsche (2008) provide ad-hoc determined weights. However, Knüppel and Vladu (2016) determine optimal weights by assuming an underlying AR(1) data generating process for inflation and show that their weights can differ substantially from the ad-hoc determined weights.

Our approach for setting the weights is inspired by Knüppel and Vladu (2016) but contains some shortcuts. First, we performed an AR(1) regression on quarter-on-quarter annualised inflation to obtain the estimated mean, persistence parameter, and error variance. Second, we used these point estimates to simulate a long time series from the AR(1) model, together with the model-consistent rolling horizon and calendar year forecasts. Based on this simulated data, we ran an OLS regression of the rolling horizon forecasts on the calendar year forecasts in order to extract the optimal weights. These weights were applied to the CE fixed event forecasts to extract rolling-horizon forecasts.[11] Another caveat is that the CE data is available twice per year from 1990 to 2013, and quarterly afterwards. To enhance the comparability between the SPF and CE series, we used linear interpolation to fill the gaps and only used data starting in 1999.

Figure A12: Data sensitivity analysis for trend inflation τ_t^π

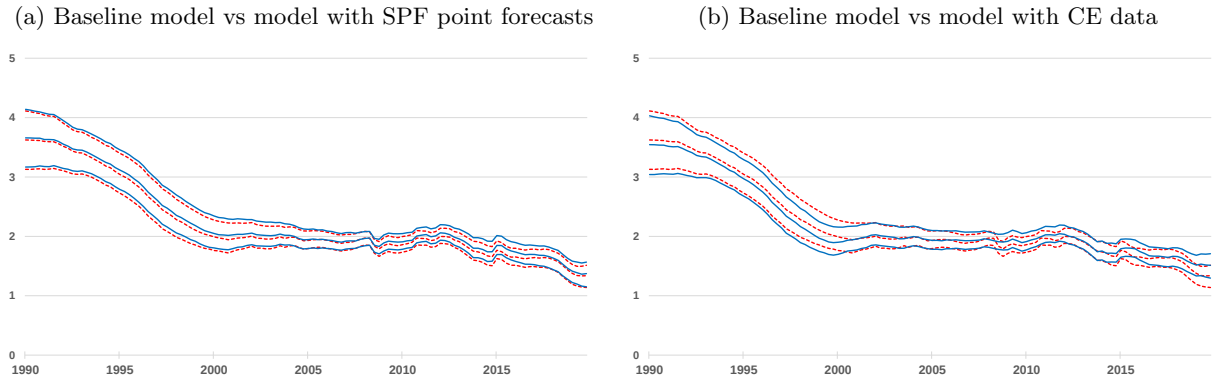

Note: The trend inflation τ_t^π estimates from the baseline model (red dashed lines) are compared against those from two models with different inflation expectations data (blue full lines). Panel (a) shows the comparison with a model using the SPF average point estimates, and panel (b) the comparison with a model using Consensus Economics (CE) expectations. The bands depict the median and 68% credible set.

References

Canova, F. and Forero, F. J. P. (2015). Estimating overidentified, nonrecursive, time-varying coefficients structural vector autoregressions. *Quantitative Economics*, 6(2):359–384.

[11] Since four-quarter ahead forecasts are available from CE starting in 2012, we found that our derived weights delivered substantially more accurate approximations compared to the ad-hoc weights, with a mean-square error of 0.01 vs 0.12.

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.

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

Chan, J. C., Clark, T. E., and Koop, G. (2018). A new model of inflation, trend inflation, and long-run inflation expectations. *Journal of Money, Credit and Banking*, 50(1):5–53.

Chan, J. C. and Grant, A. L. (2015). Pitfalls of estimating the marginal likelihood using the modified harmonic mean. *Economics Letters*, 131(C):29–33.

Chan, J. C. C., Koop, G., and Potter, S. M. (2013). A New Model of Trend Inflation. *Journal of Business & Economic Statistics*, 31(1):94–106.

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.

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.

Cogley, T., Morozov, S., and Sargent, T. J. (2005). Bayesian fan charts for U.K. inflation: Forecasting and sources of uncertainty in an evolving monetary system. *Journal of Economic Dynamics and Control*, 29(11):1893–1925.

Dovern, J. and Fritsche, U. (2008). Estimating Fundamental Cross-Section Dispersion from Fixed Event Forecasts. Discussion Papers of DIW Berlin 787, DIW Berlin, German Institute for Economic Research.

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

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models. *Review of Economic Studies*, 65(3):361–393.

Knüppel, M. and Vladu, A. L. (2016). Approximating fixed-horizon forecasts using fixed-event forecasts. Discussion Papers 28/2016, Deutsche Bundesbank.

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.

Lakdawala, A. (2015). Modeling Monetary Policy Dynamics: A Comparison of Regime Switching and Time Varying Parameter Approaches.

Omori, Y., Chib, S., Shephard, N., and Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140(2):425–449.

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.

Stock, J. H. and Watson, M. W. (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, 39(s1):3–33.