

Karlsson, Sune; Österholm, Pär

Working Paper

Is the US Phillips Curve Stable? Evidence from Bayesian VARs

Working Paper, No. 5/2018

Provided in Cooperation with:

Örebro University School of Business

Suggested Citation: Karlsson, Sune; Österholm, Pär (2018) : Is the US Phillips Curve Stable? Evidence from Bayesian VARs, Working Paper, No. 5/2018, Örebro University School of Business, Örebro

This Version is available at:

<https://hdl.handle.net/10419/244537>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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



WORKING PAPER

5/2018

Is the US Phillips Curve Stable? Evidence from Bayesian VARs

Sune Karlsson

School of Business, Örebro University

and

Pär Österholm

School of Business, Örebro University

Economics/Statistics

ISSN 1403-0586

<https://www.oru.se/institutioner/handelshogskolan/forskning/working-papers/>

Örebro University School of Business
701 82 Örebro
SWEDEN

Is the US Phillips Curve Stable? Evidence from Bayesian VARs*

Sune Karlsson[∇]
School of Business, Örebro University

and

Pär Österholm[◊]
School of Business, Örebro University

Abstract

Inflation did not fall as much as many economists expected as the Great Recession hit the US economy. One explanation suggested for this phenomenon is that the Phillips curve has become flatter. In this paper we investigate the stability of the US Phillips curve, employing Bayesian VARs to quarterly data from 1990Q1 to 2017Q3. We estimate bivariate models for PCE inflation and the unemployment rate under a number of different assumptions concerning the dynamics and covariance matrix. Specifically, we assess the importance of time-varying parameters and stochastic volatility. Using new tools for model selection, we find support for both time-varying parameters and stochastic volatility. Interpreting the Phillips curve as the inflation equation of our Bayesian VAR, we conclude that the US Phillips curve has been unstable. Our results also indicate that the Phillips curve may have been somewhat flatter between 2005 and 2013 than in the decade preceding that period. However, while the dynamic relations of the model appear to be subject to time variation, we note that the effect of a shock to the unemployment rate on inflation is not fundamentally different over time. Finally, a conditional forecasting exercise suggests that as far as the models are concerned, inflation may not have been unexpectedly high around the Great Recession.

JEL Classification: C11, C32, C52, E37

Keywords: Time-varying parameters, Stochastic volatility, Model selection, Inflation, Unemployment

* We are grateful to Todd Clark and participants at the workshop on "Central Bank Forecasting" at the Federal Reserve Bank of St Louis for valuable comments and Joshua Chan for sharing Matlab code.

[∇] Örebro University, School of Business, 701 82 Örebro, Sweden

e-mail: sune.karlsson@oru.se Phone: +46 19 301257

[◊] Örebro University, School of Business, 701 82 Örebro, Sweden

e-mail: par.osterholm@oru.se Phone: +46 19 301311

1. Introduction

In association with the Great Recession, the US unemployment rate rose from 4.8 percent in the fourth quarter of 2007 to 9.3 percent in the second quarter of 2009. Such a dramatic increase in unemployment could be expected to generate substantial downward pressure on inflation. And while the US inflation rate certainly fell during that period, the fall was smaller than expected by many economists. Similar developments were also found in many other countries, giving rise to a general discussion about the “missing disinflation”.¹ One explanation for this development is that the Phillips curve may have become flatter in recent years; see, for example, Bean (2006), Gaiotti (2010), Ihrig *et al.* (2010) and Kuttner and Robinson (2010).² The suggestion that the Phillips curve has become flatter is not undisputed though. For example, Blanchard *et al.* (2015) suggest that the Phillips curve largely has been stable since the early 1990s.³ While the question of the Phillips curve’s stability has generated an intense academic debate, it also matters to policy makers; for decades, different versions of the Phillips curve have been – and are still – widely used by central banks when forecasting inflation.

In this paper we add empirical evidence concerning the properties of the US Phillips curve. This is done in a Bayesian VAR framework where we estimate bivariate models using quarterly data of PCE inflation and the unemployment rate ranging from 1990Q1 to 2017Q3. The models are estimated under a number of different assumptions concerning the dynamics and covariance matrix. The different specifications cover two relevant aspects of the time series dynamics which have been discussed in the literature; the time-varying parameters is a way to incorporate gradual structural change – thereby allowing for the Phillips curve to evolve over time – whereas the stochastic volatilities allow the shocks that hit the economy to be heteroscedastic. This latter feature seems important given that we are studying a sample including the Great Recession. Having estimated the different models, we then employ recently developed methods for Bayesian model selection in order to establish the preferred specification.

In conducting this analysis, we contribute to the existing literature in several ways. First, by conducting formalized model selection in a Bayesian framework to distinguish between the models we provide a statistical measure of how important different time-varying features are. This can be contrasted to previous literature using time-varying parameters and/or stochastic volatilities where such modelling choices almost exclusively have been assumed rather than supported through statistical testing procedures; see, for example, Cogley and Sargent (2005) and Bianchi and Civelli (2015). Exceptions exist though, such as Koop *et al.* (2009). Thanks

¹ See, for example, the International Monetary Fund (2013).

² Globalisation is commonly suggested as an important reason for such a development. This explanation has, however, been questioned by Ball (2006).

³ Further criticism can be found in Coibion and Gorodnichenko (2015) and Blanchard (2016).

to recent developments in Bayesian estimation techniques by Chan and Eisenstat (2018), we can conduct these model-selection calculations despite the fairly complicated nature of this issue. Second, by analysing a low-dimensional system with observable variables, we offer information on this relevant topic in a fairly intuitive framework. In doing this, we move in the direction of Phillips’ (1958) original observation where simple correlations were in focus. At a more general level, using sparse specifications is of course not something new in the literature related to the Phillips curve; see, for example, parts of the analysis in Stock and Watson (1999) and Faust and Wright (2013), or the contributions by Clark and McCracken (2006), Dotsey *et al.* (2017) and Knotek and Zaman (2017) in which only small models are used.⁴ It does however stand in contrast to recent literature relying on larger models, such as Laséen and Taheri Sanjani (2016). Third, and from a policy perspective, we provide evidence concerning the environment in which monetary policy is acting today. For example, a flat Phillips curve means that the unemployment rate may have to become quite low in order to build up substantial inflationary pressure. Such information should prove useful to the Federal Reserve.

Briefly mentioning our results, we note that they indicate that both time-varying parameters and stochastic volatility appear to be relevant features. We accordingly conclude that the US Phillips curve has not been stable. In addition, there are indications that the Phillips curve may have been somewhat flatter between 2005 and 2013 than in the decade preceding that period. But while we find support for time variation in the relation between the unemployment rate and inflation, it is also the case that the effect of a shock to the unemployment rate is not fundamentally different over time. Taken together, our results suggest that a flatter Phillips curve may have contributed to the high inflation around the Great Recession together with a number of large positive shocks and high core inflation. Looking at conditional forecasts from the model though, it can be questioned whether inflation actually was surprisingly high in association with the Great Recession.

The remainder of this paper is organised as follows: In Section 2, we describe the methodological framework that we rely upon. Apart from presenting the detailed specifications of the Bayesian VARs, we pay special attention to the issue of model selection. We present data, conduct the empirical analysis and present our results in Section 3. Finally, Section 4 concludes.

2. Methodological framework

The Phillips curve is an analytical tool that has reached widespread fame since its introduction in 1958 and it is commonly used in both theoretical and empirical work. However, as well as being popular it also comes in

⁴ The model employed by Knotek and Zaman (2017) is a bivariate BVAR model with time-varying parameters and stochastic volatility and accordingly shares the most relevant features of the model used for the main analysis in this paper. Their study has a narrow focus on what such a model has to say about future inflation at the end of their sample though.

many different forms, where both the specification and variables included can differ.⁵ It is accordingly not completely transparent what one means when referring to “the Phillips curve”. In this paper, we will rely on a Bayesian VAR (BVAR) for our analysis. The reason for this is that the inflation equation of the BVAR can be seen as what King and Watson (1994, p. 172) referred to as the “*dynamic generalization of the Phillips curve*”.⁶ We will, however, in general focus on the system as a whole and not just the inflation equation. By doing this we reduce the risk that important dynamic effects between the two variables are omitted.

2.1 The Bayesian VAR model

We initially define $\mathbf{y}_t = (\mathbf{u}_t \ \pi_t)'$, where \mathbf{u}_t is the unemployment rate and π_t is PCE inflation, and then follow Chan and Eisenstat (2018) when specifying the most general model which will be used – the BVAR with time-varying parameters and stochastic volatility. The model is given as

$$\mathbf{B}_{0t}\mathbf{y}_t = \boldsymbol{\mu}_t + \mathbf{B}_{1t}\mathbf{y}_{t-1} + \dots + \mathbf{B}_{pt}\mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t \quad (1)$$

where \mathbf{B}_{0t} is a 2x2 lower triangular matrix with ones on the diagonal;⁷ $\boldsymbol{\mu}_t$ is a 2x1 vector of time-varying intercepts; $\mathbf{B}_{1t}, \dots, \mathbf{B}_{pt}$ are 2x2 matrices with the parameters describing the dynamics of the BVAR; and $\boldsymbol{\varepsilon}_t$ is a 2x1 vector of disturbances, $\boldsymbol{\varepsilon}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}_t)$, where $\boldsymbol{\Sigma}_t = \text{diag}(\exp(h_{1t}), \exp(h_{2t}))$.

The model can be re-written as

$$\mathbf{y}_t = \tilde{\mathbf{X}}_t\boldsymbol{\beta}_t + \mathbf{W}_t\boldsymbol{\gamma}_t + \boldsymbol{\varepsilon}_t \quad (2)$$

where $\tilde{\mathbf{X}}_t = \mathbf{I}_2 \otimes (\mathbf{1}, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p})$, $\boldsymbol{\beta}_t = \text{vec}((\boldsymbol{\mu}_t, \mathbf{B}_{1t} \dots, \mathbf{B}_{pt})')$, $\mathbf{W}_t = (0, -\mathbf{y}_{1t})'$ and $\boldsymbol{\gamma}_t$, which consists of the free elements of \mathbf{B}_{0t} stacked by rows, in this case becomes the scalar b_{21t}^0 ; $\boldsymbol{\varepsilon}_t$ is defined as above. Re-writing the model again, we can cast it as

$$\mathbf{y}_t = \mathbf{X}_t\boldsymbol{\theta}_t + \boldsymbol{\varepsilon}_t \quad (3)$$

where $\mathbf{X}_t = (\tilde{\mathbf{X}}_t, \mathbf{W}_t)$ and $\boldsymbol{\theta}_t = (\boldsymbol{\beta}'_t, \boldsymbol{\gamma}'_t)'$; $\boldsymbol{\varepsilon}_t$ is defined as above. Finally, the processes for the time-varying parameters and log-volatilities are specified as random walks:

$$\boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\eta}_t \quad (4)$$

⁵ Just to mention a few examples, it can be noted that Phillips' (1958) original observation largely consisted of that there was a negative relation between (wage) inflation and the unemployment rate. Phelps (1967) and Friedman (1968) augmented the Phillips curve with expectations. In the New-Keynesian literature, the Phillips curve is typically derived from micro foundations and often of hybrid type; see, for example, Galí and Gertler (1999).

⁶ Being non-structural and based on time-series models, our analysis also bears resemblance to analysis relying on simple single-equation regressions; see, for example, Svensson (2015).

⁷ In line with, for example, Cogley and Sargent (2005) and Primiceri (2005), we accordingly rely on a recursive structure in order to identify the orthogonal disturbances to the system.

$$\mathbf{h}_t = \mathbf{h}_{t-1} + \boldsymbol{\zeta}_t \quad (5)$$

where $\boldsymbol{\eta}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}_\theta)$ and $\boldsymbol{\zeta}_t \sim N(\mathbf{0}, \boldsymbol{\Sigma}_h)$. The model in equations (3) to (5) now has a generic state-space form. Due to this specification of $\boldsymbol{\theta}_t$, both $\boldsymbol{\beta}_t$ and $\boldsymbol{\gamma}_t$ can be integrated out analytically and numerical efficiency is improved (Eisenstat *et al.*, 2016).

The model described above – which has both time-varying parameters and stochastic volatility – is the most general model we will consider in our analysis. It seems highly relevant to allow for the possibility that both of these effects are present as they find support in earlier research. For example, King and Watson (1994, p. 209) found “*important evidence of econometric instability over subsamples*” for the Phillips curve. Also Stock and Watson (1999) found evidence of an unstable Phillips curve. Concerning the issue of shock volatility, Sims and Zha (2006) have pointed out the importance of allowing this to be time-varying and Stock and Watson (2012) found that relatively large shocks were the cause of the Great Recession. By imposing various restrictions on the most general model, we can assess how important different features are. For example, by making $\boldsymbol{\theta}_t$ and \mathbf{h}_t constant, we get a traditional BVAR with time-invariant parameters and covariance matrix. We will in what follows estimate and compare models under six different assumptions:

- i) Time-varying $\boldsymbol{\theta}_t$ and stochastic volatility.
- ii) Time-varying $\boldsymbol{\theta}_t$ and time-invariant covariance matrix.
- iii) Time-varying $\boldsymbol{\gamma}_t$, time-invariant $\boldsymbol{\beta}_t$ and stochastic volatility.
- iv) Time-varying $\boldsymbol{\beta}_t$, time-invariant $\boldsymbol{\gamma}_t$ and stochastic volatility.
- v) Time-invariant $\boldsymbol{\theta}_t$ and stochastic volatility.
- vi) Time-invariant $\boldsymbol{\theta}_t$ and covariance matrix.

To gain some additional insights into the model variations we consider, it is useful to also study the “reduced form” VAR

$$\mathbf{y}_t = \boldsymbol{\delta}_t + \mathbf{A}_{1t}\mathbf{y}_{t-1} + \dots + \mathbf{A}_{pt}\mathbf{y}_{t-p} + \mathbf{e}_t \quad (6)$$

with $\boldsymbol{\delta}_t = \mathbf{B}_{0t}^{-1}\boldsymbol{\mu}_t$, $\mathbf{A}_{it} = \mathbf{B}_{0t}^{-1}\mathbf{B}_{it}$ and $\mathbf{e}_t \sim N(\mathbf{0}, \boldsymbol{\Omega}_t)$ for $\boldsymbol{\Omega}_t = \mathbf{B}_{0t}^{-1}\boldsymbol{\Sigma}_t\mathbf{B}_{0t}^{-1'}$. Under model *i* all reduced form parameters, regression coefficients, error variances and covariances, are time varying. Model *ii* restricts $\boldsymbol{\Omega}_t$ to a limited variability mainly affecting the correlation structure through \mathbf{b}_{21t}^0 while model *iii* allows for full variability in $\boldsymbol{\Omega}_t$ with only limited variability in the regression parameters. In model *iv*, the regression parameters and variances have full variability while the correlation structure is relatively constant and in model *v* only the variances have full variability and the correlation structure only changes as a consequence of the time varying variances.

2.2 Priors and estimation

The choice of prior distributions and parameters of the priors can have a substantial effect on model comparisons based on marginal likelihoods and the priors must be set up with some care. Starting with the simplest model, the constant parameter VAR, a diffuse prior is specified with $\boldsymbol{\beta} \sim N(\mathbf{0}, 10\mathbf{I})$, $\boldsymbol{\gamma} \sim N(\mathbf{0}, 10\mathbf{I})$ and the diagonal elements of $\boldsymbol{\Sigma}$ are given independent inverse Gamma priors, $\sigma_i^2 \sim iG(v_{0i}, S_{0i})$, with $v_{0i} = 5$ and S_{0i} is chosen to match the prior mean, $E(\sigma_i^2) = S_{0i}/(v_{0i} - 1)$, to the residual variance of a univariate AR-model with the same lag length as the VAR.

For the models with time-varying parameters, $\boldsymbol{\beta}_t$ and/or $\boldsymbol{\gamma}_t$, the constant parameter prior is taken as the prior for the initial conditions, $\boldsymbol{\beta}_0 \sim N(\mathbf{0}, 10\mathbf{I})$, $\boldsymbol{\gamma}_0 \sim N(\mathbf{0}, 10\mathbf{I})$. The variance-covariance matrix of the state equation (4) is diagonal, $\boldsymbol{\Sigma}_\theta = \text{diag}(\sigma_{\theta i}^2)$, with inverse Gamma priors, $\sigma_{\theta i}^2 \sim iG(v_{\theta i}, S_{\theta i})$, where we set the shape parameters to $v_{\theta i} = 5$ and adjust the scale parameters to achieve a prior mean for $\sigma_{\theta i}^2$ of 0.01 for the intercepts and 0.0001 for the other regression parameters. With our sample size of $T = 111$, the implied prior variance of $\boldsymbol{\beta}_{iT}$ is about 11 for intercepts and 10.01 for the other regression parameters.

For the models with stochastic volatility, the initial state for the log-variance is given a normal prior, $\mathbf{h}_0 \sim N(\boldsymbol{\mu}_h, 0.25\mathbf{I})$. That is, the initial state for the variance is lognormally distributed and we select the elements of $\boldsymbol{\mu}_h$ to match the mean of the lognormal distribution with the residual variances of univariate AR-models as with the constant variance models. The variance-covariance of the state equation (5) is diagonal, $\boldsymbol{\Sigma}_h = \text{diag}(\sigma_{h1}^2, \sigma_{h2}^2)$, with inverse Gamma priors, $\sigma_{hi}^2 \sim iG(v_{hi}, S_{hi})$. The shape parameters are set to $v_{hi} = 5$ and the scale parameters to $S_{hi} = 0.04$ resulting in a prior mean for σ_{hi}^2 of 0.01.

We use the MCMC-sampler developed by Chan and Eisenstat (2018) for posterior inference. In its most general form it proceeds by sampling from the full conditional posteriors

$$p(\boldsymbol{\theta} | \mathbf{y}, \mathbf{h}, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_h, \boldsymbol{\theta}_0, \mathbf{h}_0) \quad (7)$$

$$p(\mathbf{h} | \mathbf{y}, \boldsymbol{\theta}, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_h, \boldsymbol{\theta}_0, \mathbf{h}_0) \quad (8)$$

$$p(\boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_h | \mathbf{y}, \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\theta}_0, \mathbf{h}_0) \quad (9)$$

$$p(\boldsymbol{\theta}_0, \mathbf{h}_0 | \mathbf{y}, \boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\Sigma}_\theta, \boldsymbol{\Sigma}_h) \quad (10)$$

Using the model formulation in (3), the full conditional posterior of $\boldsymbol{\theta} = (\boldsymbol{\theta}'_1, \dots, \boldsymbol{\theta}'_T)'$ is normal and can be sampled efficiently using the precision sampler of Chan and Jeliazkov (2009). For models when only $\boldsymbol{\beta}_t$ or $\boldsymbol{\gamma}_t$ is time-varying, $\boldsymbol{\beta}_t$ is sampled conditionally on $\boldsymbol{\gamma}_t$ and $\boldsymbol{\gamma}_t$ is sampled conditionally on $\boldsymbol{\beta}_t$, both with normal updates. In models with stochastic volatility, \mathbf{h} is sampled using the auxiliary mixture sampler of Kim *et al.* (1998). The diagonal elements of $\boldsymbol{\Sigma}_\theta$ and $\boldsymbol{\Sigma}_h$ are conditionally independent with inverse Gamma full

conditional posteriors. Finally $\boldsymbol{\theta}_0$ and \mathbf{h}_0 are conditionally independent with normal full conditional posteriors. For the models with constant variance, the last three steps are replaced by an update of the diagonal elements of $\boldsymbol{\Sigma}$ which are conditionally independent with inverse Gamma full conditional posteriors.

2.3 Model selection

In order to distinguish between the models and gain insight into the stability over time of the US Phillips curve, we assess the fit of the different models using the marginal likelihood. In a Bayesian setting the marginal likelihood is the appropriate measure of how well the model (and prior) agrees with the data. A particular advantage of the Bayesian framework is that it allows for statistically rigorous and straightforward comparison of non-nested models.⁸

The marginal likelihood for a model M is obtained by integrating out all unknown quantities from the joint distribution with the data

$$m(\mathbf{y}|M) = \int p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\xi}, M) p(\boldsymbol{\theta}|\boldsymbol{\xi}, M) p(\mathbf{h}|\boldsymbol{\xi}, M) p(\boldsymbol{\xi}, M) d\boldsymbol{\theta} d\mathbf{h} d\boldsymbol{\xi} \quad (11)$$

where $\boldsymbol{\xi}$ collects the parameters $\boldsymbol{\theta}_0$, \mathbf{h}_0 , $\boldsymbol{\Sigma}_{\boldsymbol{\theta}}$ and $\boldsymbol{\Sigma}_{\mathbf{h}}$ of the state equations (4) and (5). Model choice can then be made through direct comparison of marginal likelihoods or Bayes factors

$$B_{ij} = \frac{m(\mathbf{y}|M_i)}{m(\mathbf{y}|M_j)} = \frac{p(M_i|\mathbf{y})}{p(M_j|\mathbf{y})} \frac{p(M_i)}{p(M_j)} \quad (12)$$

measuring the relative fit of the models or – if we are willing to assign prior probabilities that the different models are correct – how our opinion about the models have changed after viewing the data. As we see no strong reason to a priori prefer one model over the other, we opt for a simple comparison of the marginal likelihoods.

Following Chan and Eisenstat (2018) we first obtain the integrated likelihood where $\boldsymbol{\theta}$ and \mathbf{h} are integrated out,

$$p(\mathbf{y}|\boldsymbol{\xi}, M) = \int p(\mathbf{y}|\boldsymbol{\theta}, \mathbf{h}, \boldsymbol{\xi}, M) p(\boldsymbol{\theta}|\boldsymbol{\xi}, M) p(\mathbf{h}|\boldsymbol{\xi}, M) d\boldsymbol{\theta} d\mathbf{h} = \int p(\mathbf{y}|\mathbf{h}, \boldsymbol{\xi}, M) p(\mathbf{h}|\boldsymbol{\xi}, M) d\mathbf{h} \quad (13)$$

where the time-varying or constant parameters in $\boldsymbol{\theta}$ can be integrated out analytically. An estimate of $p(\mathbf{y}|\boldsymbol{\xi}, M)$ is then obtained by importance sampling techniques as

⁸ Previous literature relying on classical methods have in some cases conducted some tests for structural stability but often not in a framework that is directly designed for the question at hand. Many times the analysis has been ad hoc and/or less formal; see, for example, Bernanke and Mihov (1998) who applied the test of Andrews (1993) or Abbate *et al.* (2016).

$$p(\widehat{\mathbf{y}}|\widehat{\boldsymbol{\xi}}, M) = \frac{1}{R} \sum_{i=1}^R \frac{p(\mathbf{y}|\mathbf{h}^i, \boldsymbol{\xi}, M)p(\mathbf{h}^i|\boldsymbol{\xi}, M)}{g(\mathbf{h}^i|\boldsymbol{\xi}, M)} \quad (14)$$

by sampling \mathbf{h}^i from a suitable importance density $g(\mathbf{h}^i|\boldsymbol{\xi}, M)$. The marginal likelihood is then estimated using an outer importance sampling loop for $\boldsymbol{\xi}$ as

$$m(\widehat{\mathbf{y}}|M) = \frac{1}{P} \sum_{i=1}^P \frac{p(\widehat{\mathbf{y}}|\widehat{\boldsymbol{\xi}}^i, M)p(\boldsymbol{\xi}^i, M)}{g(\boldsymbol{\xi}^i, M)} \quad (15)$$

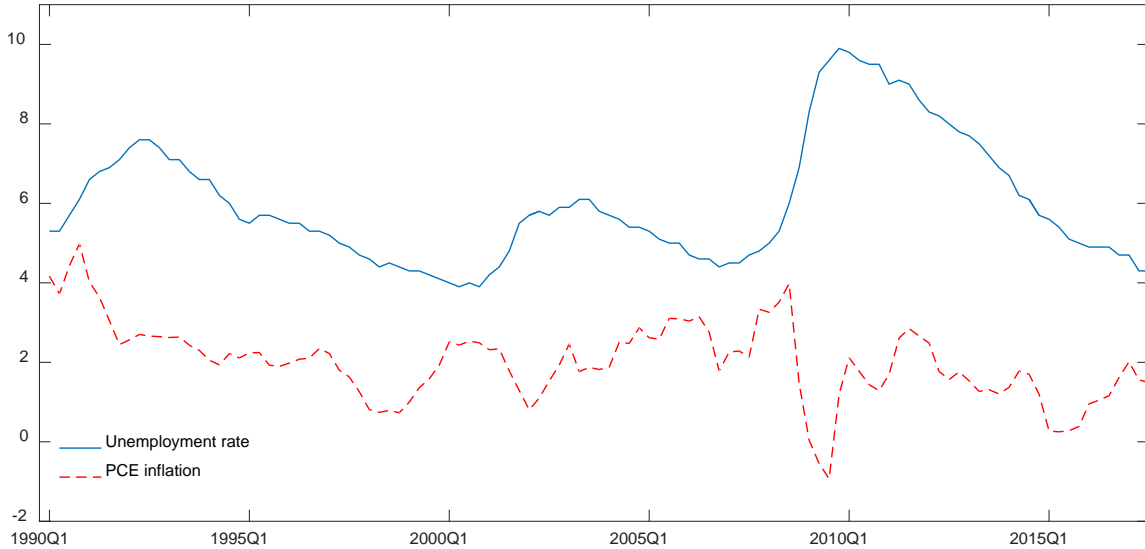
using draws of $\boldsymbol{\xi}^i$ from the importance density $g(\boldsymbol{\xi}^i, M)$ constructed using the cross-entropy method of Chan and Eisenstat (2015). The inner importance sampling step is omitted for the models with constant variances as the integrated likelihood is analytic. Details concerning the implementation can be found in Chan and Eisenstat (2018).

3. Empirical analysis

3.1 Data

Data on seasonally adjusted PCE deflator and unemployment rate – ranging from 1990Q1 to 2017Q3 – were sourced from the FRED database of the Federal Reserve Bank of St Louis. PCE inflation is calculated as $\pi_t = 100(P_t/P_{t-4} - 1)$ where P_t is the PCE deflator at time t . Data are shown in Figure 1.

Figure 1. Data



Note: Both variables are measured in percent.

We believe that 1990Q1 is a reasonable starting point for our analysis. Choosing a period which is excessively long – for example, if we were to begin in the 1950's – we argue that it is very unlikely that the Phillips curve could have been stable given the changes that the US economy has undergone over time. Using the sample

in this paper, we rely on reasonably recent data and hence estimate the model over a period where it would not be unrealistic to find stability; as pointed out above, it has been suggested by, for example, Blanchard *et al.* (2015) that the Phillips curve largely has been stable since the beginning of this sample.

3.2 Is there time variation?

We next set out to estimate and compare our models. The six different models had the following assumptions: *i*) time-varying θ_t and stochastic volatility, *ii*) time-varying θ_t and time-invariant covariance matrix, *iii*) time-varying γ_t , time-invariant β_t and stochastic volatility, *iv*) time-varying β_t , time-invariant γ_t and stochastic volatility, *v*) time-invariant θ_t and stochastic volatility and *vi*) time-invariant θ_t and covariance matrix. The lag length is set to $p = 4$ in all cases.

The log marginal likelihoods from the estimation of the models are shown in Table 1. As can be seen, the marginal likelihood is highest for the models with stochastic volatility (models *i*, *iii*, *iv* and *v*). While the difference is small between these models the Bayes factors favour models with time-varying parameters with the best model being the fully flexible model *i*. Turning to the constant variance models (*ii* and *vi*) it is clear that they have less support in the data and that allowing for time-varying parameters in model *ii* leads to a considerable improvement of the fit. There is thus support for time-varying parameters in the data. But it is also clear that the data prefer stochastic volatility over a time-invariant covariance matrix – a finding that also is in line with the criticism of Sims (2001) and Stock (2001) against Cogley and Sargent’s (2001) BVAR model which had time-varying parameters but a time-invariant covariance matrix.

Table 1. Marginal likelihood for the different BVAR specifications.

Model	$\ln(p(y M))$
i	-101.8
ii	-107.7
iii	-102.2
iv	-101.9
v	-102.4
vi	-112.3

Note: The models are the following: *i*) time-varying θ_t and stochastic volatility, *ii*) time-varying θ_t and time-invariant covariance matrix, *iii*) time-varying γ_t , time-invariant β_t and stochastic volatility, *iv*) time-varying β_t , time-invariant γ_t and stochastic volatility, *v*) time-invariant θ_t and stochastic volatility and *vi*) time-invariant θ_t and covariance matrix.

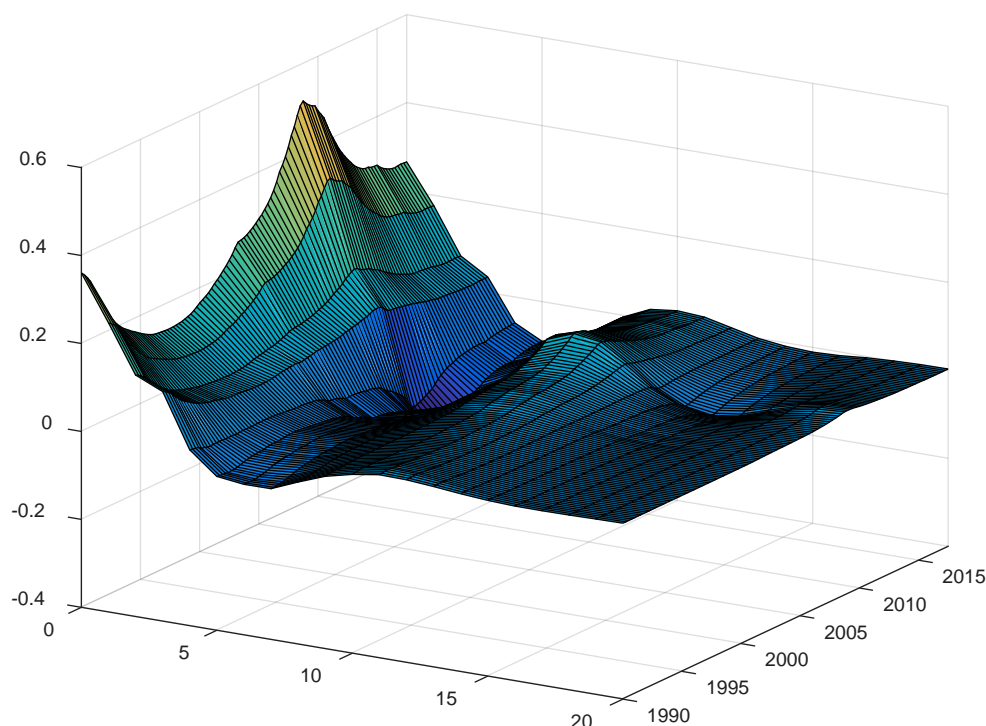
Summarising the above results, we have found support for both time-varying parameters and stochastic volatility. The fact that parameters appear to be time varying indicates that the Phillips curve is not stable over time. We next illustrate what this time variation looks like by illustrating the dynamic properties of model *i* in more detail.

3.3 Impulse-response functions

As was shown above, model i – that is, the model with time-varying θ_t and stochastic volatility – was the preferred one according to the marginal likelihood calculations. The impulse-response function from this model which describes the effect of a shock to inflation on inflation itself is given in Figure 2. We initially discuss this impulse-response function since it clearly shows why both time-varying parameters and stochastic volatility are favoured features.

As can be seen from the figure, the estimated standard deviation of the shocks to the inflation equation – that is, $[\exp(\hat{h}_{2t})]^{0.5}$ – varies a fair bit over time. In 1997, in what may be described as the peak of the Great Moderation, it was as low as 0.14. It then increased almost continuously until 2008Q4 when it reached a value of 0.51. Since then, the volatility of the shocks has decreased and it is by the end of the sample estimated to be approximately 0.27.

Figure 2. Impulse-response function for model i: The effect of shocks to inflation on inflation, 1990Q1-2017Q3.



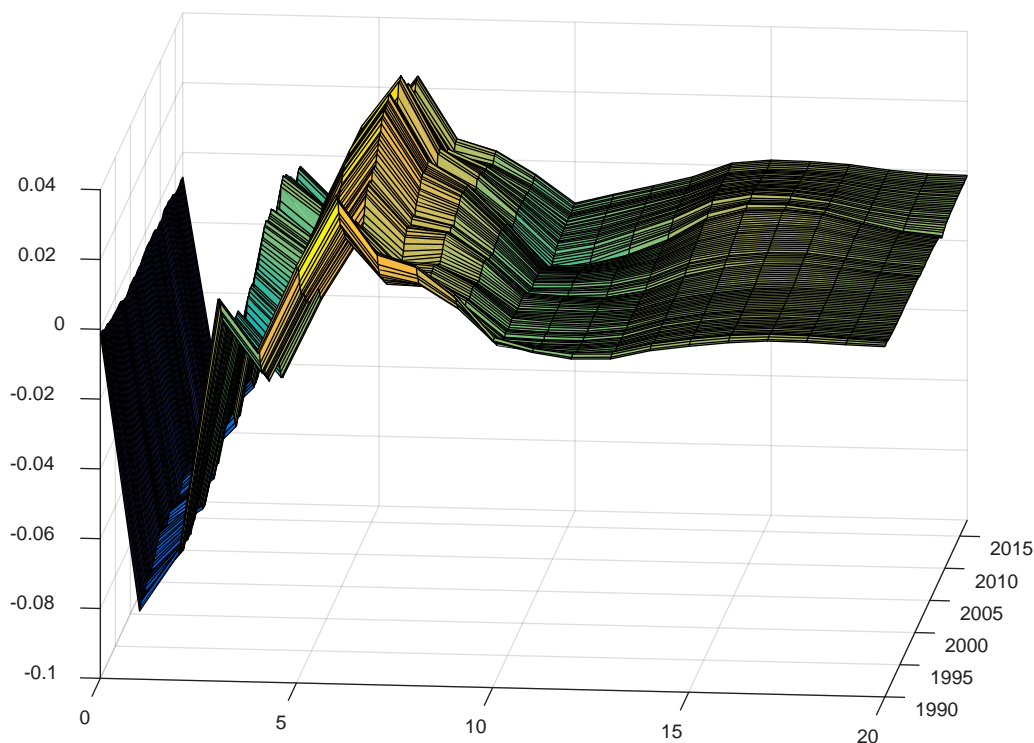
Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure 2 also illustrates that the dynamic properties of the system has changed. The fairly large variation in these impulse-response functions indicate that there is non-negligible time variation in the coefficients of the model, which leads us to question the stability of the US Phillips curve. This is also confirmed when looking

at the estimated coefficients of the inflation equation from the reduced form VAR in equation (6); see Figure A2 in the Appendix.

Turning to the impulse-response functions which perhaps are of primary interest to us, namely the effect that shocks to the unemployment rate have on inflation, these can be found in Figure 3. As is clearly shown in the figure, a shock to the unemployment rate has a negative effect on inflation; it lowers it by 0.06 to 0.08 percentage points at the one-quarter horizon and 0.03 percentage points at the two-quarter horizon. It can be noted that this effect is rather stable over time. The same thing can be said about the volatility of the shock to the unemployment rate; as can be seen from Figure A1 in the Appendix, the standard deviation of this shock has been between 0.11 and 0.14 during the entire sample.⁹ So while there is time variation also in the impulse-response functions shown in Figure 3, it is much less pronounced than what we saw in Figure 2.

Figure 3. Impulse-response function for model i: The effect of shocks to the unemployment rate on inflation, 1990Q1-2017Q3.



Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

The fact that shocks to the unemployment rate has had an effect on inflation that has not varied substantially over time seems to contradict the claim of a Phillips curve that has flattened markedly in recent years. However, the impulse-response functions describe the properties of the full bivariate system. A different way to assess the slope of the Phillips curve is to look at the sum of the coefficients on the lagged unemployment

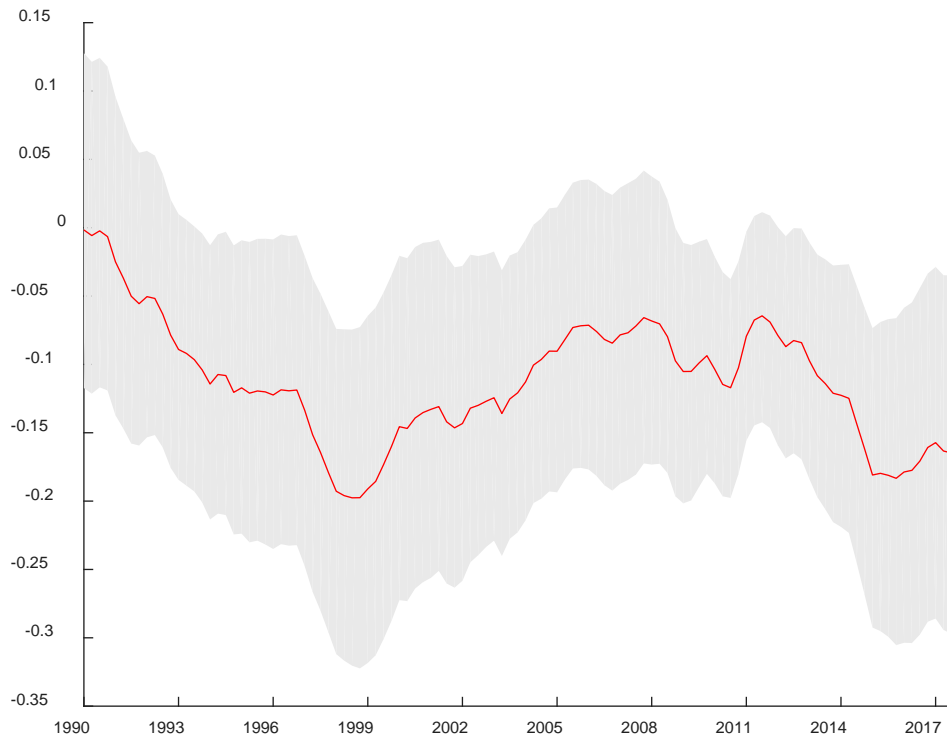
⁹ For completeness, the impulse-response functions showing the effect of shocks to inflation on the unemployment rate are given in Figure A3 in the Appendix.

rate in the inflation equation of the VAR in equation (6); see, for example, Knotek and Zaman (2017). We next turn to this issue.

3.4 The slope of the Phillips curve

The estimated parameters of the inflation equation from the reduced form VAR in equation (6) are given in Figure A2 in the Appendix. Being interested in the slope of the Phillips curve, we turn our focus to the sum of the coefficients on the unemployment rate which is shown in Figure 4. As can be seen, the point estimate indicates a non-negligible amount of variation within the sample. At the beginning of the sample, the slope is approximately zero but it then falls and reaches a low of -0.2 in 1998. Between 2005 and 2013, the slope is more moderate – hovering around -0.09 – only to once again fall and reach values of -0.18 towards the end of the sample. We should of course keep in mind that the point estimate is associated with a fair amount of uncertainty as illustrated by the 68 percent credible interval; changes should accordingly not be over-interpreted. It is nevertheless our best estimate and therefore interesting to note that the Phillips curve may have been flatter around the financial crisis than it typically has been during this sample. It hence seems that a flatter Phillips curve may have contributed to the surprisingly high inflation outcomes in the aftermath of the financial crisis.

Figure 4. Estimated slope of the Phillips curve for model i, 1990Q1-2017Q3.

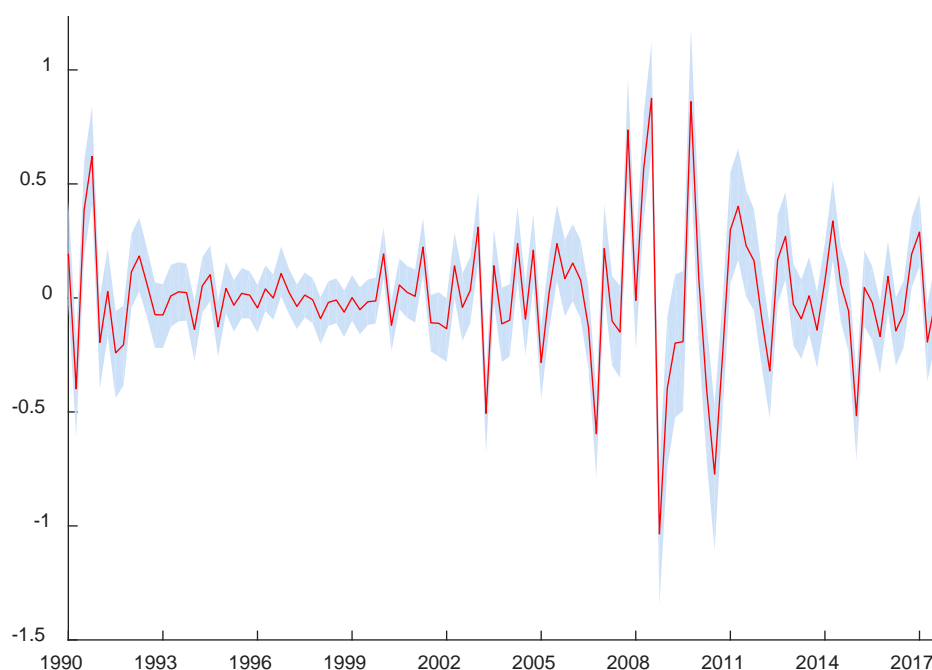


Note: Slope is measured as the sum of the coefficients on the lags of unemployment in the inflation equation in equation (6). Coloured band is 68% equal tail credible interval.

A somewhat flatter than usual Phillips curve was not the only thing that contributed to fairly high inflation around the financial crisis though. Within the model, a second explanation is to some extent offered by the

shocks that according to the model hit the economy. Figure 5 shows the estimated residuals to inflation. As can be seen from the figure, a number of large positive shocks hit the economy between 2007Q4 and 2008Q3. This contributed to a higher than expected inflation for the following quarters as shown by the impulse-response functions above. This is to some extent counteracted by the effect of the several large negative shocks that the model has identified in 2008Q4 and 2009Q1.¹⁰

Figure 5. Estimated shocks to inflation from model i, 1990Q1-2017Q3.



Note: Estimated shocks are measures in percentage points. Red line is based on median. Coloured band is 68% equal tail credible interval.

Before leaving the discussion of the slope of the Phillips curve behind, it also deserves to be pointed out that the Phillips curve has not been flatter than usual during the most recent years. Since 2015, the slope has been between -0.20 and -0.17 which is actually quite steep relative to the rest of the sample; a steeper curve can only be found in the late 1990s. This appears to at least to some extent contradict a claim that a flatter Phillips curve has been the cause of the weak inflation outcomes that the Federal Reserve have been struggling with during this period.

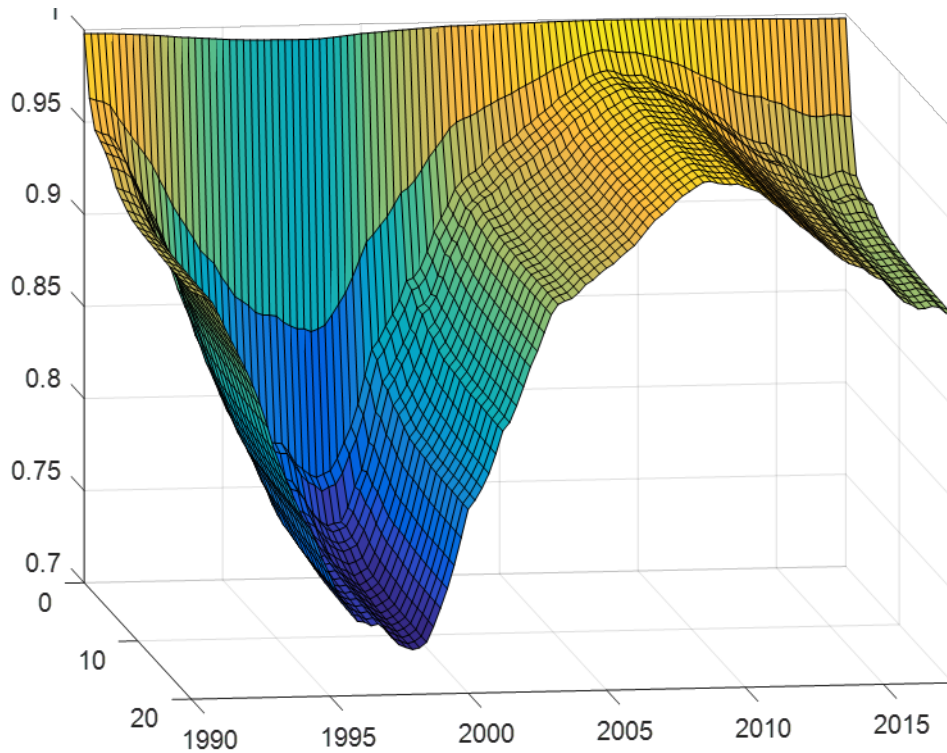
3.5 Variance decomposition

As has been shown above, there appears to have been changes over time to both the coefficients describing the dynamics of the system and the volatilities of the shocks. Another way of illustrating this fact is to conduct a variance decomposition.

¹⁰ Our finding that the transmission of shocks to the unemployment rate is not dramatically different around the financial crisis, and that large shocks to inflation appear to have contributed to the surprisingly high inflation, is to some extent also in line with previous literature. For example, Stock and Watson (2012) concluded – based on analysis from a high-dimensional dynamic factor model – that relatively large shocks, not changes in the transmission, caused the Great Recession.

Figure 6 shows the variance decomposition where model i has been used to calculate the share of the forecast error variance of inflation which is due to shocks to inflation. Note that since the model is bivariate, the share of the forecast error variance of inflation which is due to shocks to the unemployment rate is simply one minus the share in Figure 6. As is illustrated in the figure, the share of the forecast error of inflation which is due to shocks to inflation itself always is fairly large; regardless of time period or horizon, it never falls below 0.75. That said, there are substantial differences in this share. During the peak of the Great Moderation – that is, around 1997 – shocks to the unemployment rate explained approximately 27 percent of the forecast error variance of inflation at the twenty-quarter horizon. Around the Great Recession, on the other hand, almost all of the variance was explained by shocks to inflation itself. This result seems reasonable when looking at the shocks to inflation that the model has identified (in Figure 5) or their estimated standard deviation (which can be seen in Figure 2). It would take very large changes in the parameters in θ_t in order to offset the change in the volatility of inflation shocks given that the volatility of unemployment rate shocks has varied only moderately over time.

Figure 6. Variance decomposition for model i : Share of forecast error variance of inflation explained by shocks to inflation, 1990Q1-2017Q3.

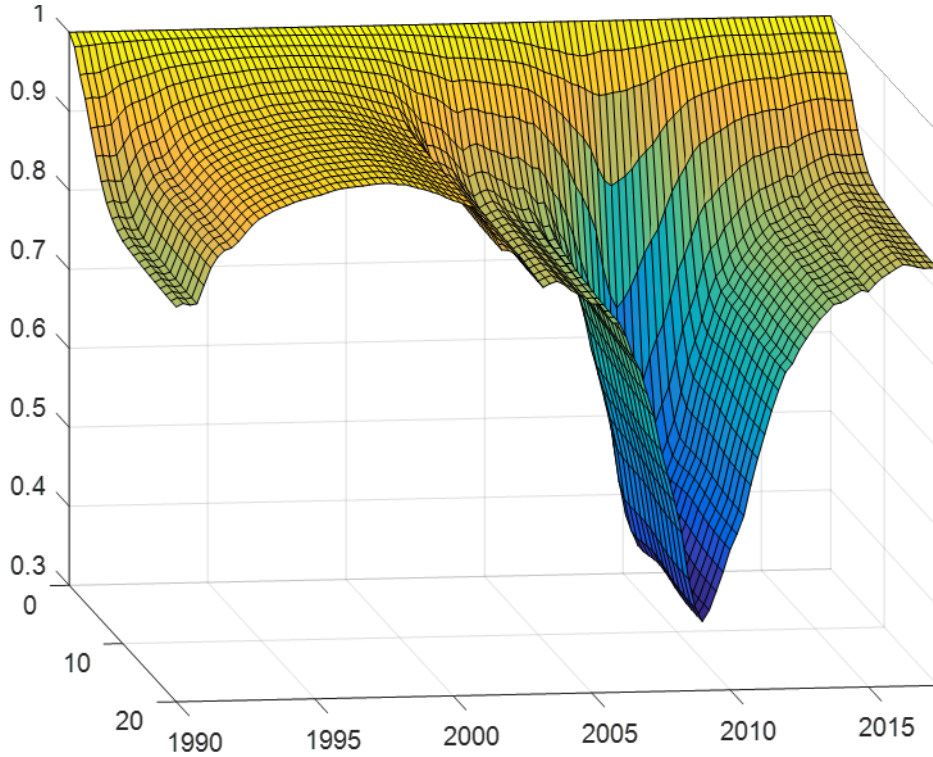


Note: Share of forecast error variance explained by shocks to inflation on vertical axis. Horizon in quarters and time on horizontal axes.

Turning to the forecast error variance decomposition of the unemployment rate, we can see from Figure 7 that this varies substantially more over time than that of inflation. For example, at the longest horizon – that is, 20 quarters – we find that the share explained by shocks to the unemployment rate varies from 0.95 in 1997 to only 0.39 in 2009Q1. At a general level the pattern is roughly the opposite of that found for inflation.

This is not surprising given that the model is bivariate and the previously noted behaviour of the volatilities of the shocks to inflation and the unemployment rate.

Figure 7. Variance decomposition for model i: Share of forecast error variance of unemployment rate explained by shocks to the unemployment rate, 1990Q1-2017Q3.



Note: Share of forecast error variance explained by shocks to the unemployment rate on vertical axis. Horizon in quarters and time on horizontal axes.

3.6 Time-varying parameters and estimated “core inflation”

The dynamics of the model have above been illustrated in a number of ways, including the traditional tools that impulse-response functions and variance decompositions constitute in the VAR literature. This has provided information concerning a number of features of the model, including the stability of the Phillips curve. However, in a model with time-varying parameters, there is information based on the change in parameters that occurs over time which is not communicated through these standard tools. Since this can be of importance when it comes to the analysis of the Phillips curve, we now address the issue of the estimated “core inflation” from the model.

Defining “core inflation”, we follow the terminology used by Cogley and Sargent (2001, 2005) and accordingly mean the value to which the inflation forecasts from the model will converge.¹¹ Employing the reduced form of the VAR in equation (6), we rewrite it in companion form as

¹¹ This concept is also closely related to what is denoted “trend inflation” in the literature; see, for example, Faust and Wright (2013) and Clark and Doh (2014) for discussions.

$$\tilde{\mathbf{y}}_t = \tilde{\boldsymbol{\delta}}_t + \mathbf{A}_t \tilde{\mathbf{y}}_{t-1} + \tilde{\mathbf{e}}_t \quad (16)$$

with

$$\tilde{\mathbf{y}}_t = (\mathbf{y}'_t, \mathbf{y}'_{t-1}, \dots, \mathbf{y}'_{t-p+1})', \quad (17)$$

$$\mathbf{A}_t = \begin{pmatrix} \mathbf{A}_{1t} & \mathbf{A}_{2t} & \cdots & \mathbf{A}_{p-1,t} & \mathbf{A}_{pt} \\ \mathbf{I} & \mathbf{0} & \cdots & & \mathbf{0} \\ \mathbf{0} & \ddots & & & \\ \vdots & & \ddots & & \vdots \\ \mathbf{0} & \cdots & \mathbf{0} & \mathbf{I} & \mathbf{0} \end{pmatrix} \quad (18)$$

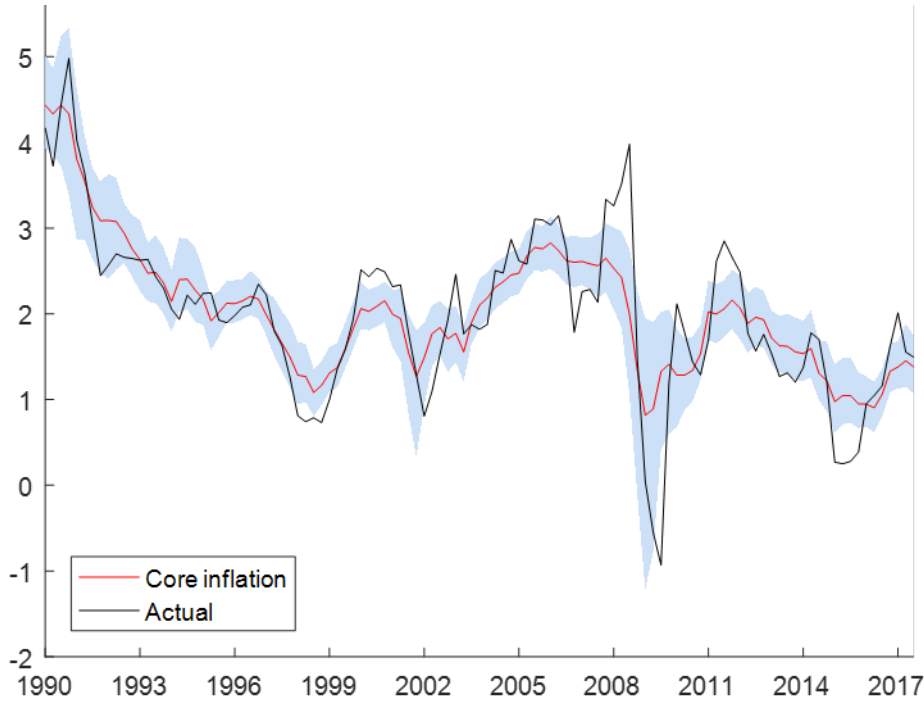
and

$$\tilde{\mathbf{e}}_t = (\mathbf{e}'_t, \mathbf{0}', \dots, \mathbf{0}')' \quad (19)$$

and solve for the unconditional mean of inflation as the second element of $\boldsymbol{\phi}_t = (\mathbf{I} - \mathbf{A}_t)^{-1} \tilde{\boldsymbol{\delta}}_t$.

The median estimated core inflation at each point in time from the model is shown in Figure 8. This shows that core inflation was fairly high when the Great Recession hit the US economy in the second quarter of 2007, namely 2.6 percent; it had reached this value after drifting up from 1.6 percent in 2003Q2. Core inflation also stayed reasonably high until 2008Q2 – when it was 2.4 percent – before falling sharply in 2008Q3. It hence seems that part of the high inflation during the Great Recession can be explained by high core inflation during its first part.

Figure 8. Estimated “core inflation” from model i, 1990Q1-2017Q3.



Note: “Core inflation” and inflation are measured in percent. Coloured band is 68% equal tail credible interval.

Figure 8 also illustrates the stubbornly low inflation during the last few years of our sample. As can be seen from the figure, core inflation has at the same time been low. The low inflationary pressure that we have seen over the last few years has hence been reflected in the estimated parameters of the model and caused estimated core inflation to decrease noticeably. For the Federal Reserve this means that while shocks to the unemployment rate might have roughly the same effect now as it has historically, there should be scope to allow for additional inflationary pressures to build up since such pressures could be needed in order to bring core inflation back up to the target level of two percent.¹²

Summing up the analysis in Sections 3.3 to 3.6 which describes the properties of the model, we note that both the time-varying parameters and the stochastic volatility leave highly visible prints. Both features also help explain the (perhaps) surprisingly high inflation around the financial crisis: There is an indication that the Phillips curve may have been somewhat flatter around the financial crisis and also that core inflation was high; in addition, a number of large positive shocks to inflation hit the system during 2007 and 2008.

3.7 Conditional forecasts

One common usage of econometric models of the type investigated above is forecasting. In many cases the models are simply employed to generate an endogenous forecast. It is, however, also common – particularly at policy institutions – to use the models for conditional forecasting. When generating conditional forecasts, one or more variables in the model are associated with paths (or distributions) that are assumed by the forecaster whereas the rest of the variables behave in accordance with the dynamics of the model. A common choice of variable to condition upon is the policy instrument of the central bank; see, for example, Sims (1982) and Leeper and Zha (2003). But applications range widely and practical policy questions include as diverse issues as effects of a US recession on Latin American GDP growth (International Monetary Fund, 2007) and effects on Swedish CPI inflation of a depreciating Swedish Krona (Sveriges Riksbank, 2016).¹³

Concerning the US Phillips curve as studied in this paper, we believe that a conditional forecasting exercise is of certain interest as it can provide us with information regarding the “missing disinflation” associated with Great Recession. In line with the analysis presented above, we conduct this exercise using model i . However, we also employ the BVAR with time-invariant θ_t and covariance matrix in order to provide a comparison with a more traditional specification.

¹² It can be noted that the formal inflation target that the Federal Reserve introduced in 2012 can be seen as being at conflict with models with time-varying parameters of the type used in this paper. Providing a clear focal point for future inflation, it can be argued that inflation forecasts should converge to the target level of two percent and that, for example, models with steady-state priors of the type suggested by Villani (2009) or Clark (2011) should be considered. However, we believe that the fact that the Federal Reserve's formal target has been active during only a reasonably small part of our sample is a good reason to employ our chosen framework.

¹³ Further contributions in the field of conditional forecasting – dealing with methodological issues and/or empirical applications – include Waggoner and Zha (1999), Hamilton and Herrera (2004), Österholm (2009), Baumeister and Kilian (2013), Clark and McCracken (2014) and Giannone *et al.* (2014).

We generate inflation forecasts for the period 2008Q1-2010Q4 conditional upon an assumed path for the unemployment rate; more specifically, this path is given by the actual unemployment rate. This will show us how the models would have predicted future inflation given perfect foresight regarding the unemployment rate.

Conditional forecasts can be produced in different ways; here we consider two methods that we believe are representative of current practice. The first method – employed in, for example, Österholm and Zettelmeyer (2008) – is recursive and myopic in the sense that it only looks one time period ahead at a time. It can be thought of as the result of a decision maker adjusting the future value of a policy variable to achieve a desired outcome in the next time period. Technically we consider restrictions of the form

$$\mathbf{R}_{t+1}\mathbf{y}_{t+1} = \mathbf{r}_{t+1} = \mathbf{R}_{t+1}(\hat{\mathbf{y}}_{t+1|t} + \mathbf{e}_{t+1}) \quad (20)$$

or

$$\mathbf{R}_{t+1}\mathbf{e}_{t+1} = \mathbf{r}_{t+1} - \mathbf{R}_{t+1}\hat{\mathbf{y}}_{t+1|t} = \mathbf{d}_{t+1} \quad (21)$$

where $\hat{\mathbf{y}}_{t+1|t}$ is the one step-ahead unconditional forecast. Given that \mathbf{e}_{t+1} is unconditionally normally distributed, $\mathbf{e}_{t+1} \sim N(\mathbf{0}, \mathbf{\Omega}_{t+1})$, the conditional distribution of \mathbf{e}_{t+1} given the restriction is

$$\begin{aligned} \mathbf{e}_{t+1} | \mathbf{R}_{t+1}\mathbf{e}_{t+1} = \mathbf{d}_{t+1} \\ \sim N(\mathbf{\Omega}_{t+1}\mathbf{R}_{t+1}(\mathbf{R}_{t+1}\mathbf{\Omega}_{t+1}\mathbf{R}_{t+1}')^{-1}\mathbf{d}_{t+1}, \mathbf{\Omega}_{t+1} - \mathbf{\Omega}_{t+1}\mathbf{R}_{t+1}(\mathbf{R}_{t+1}\mathbf{\Omega}_{t+1}\mathbf{R}_{t+1}')^{-1}\mathbf{\Omega}_{t+1}\mathbf{R}_{t+1}') \end{aligned} \quad (22)$$

and a draw from the conditional distribution, \mathbf{e}_{t+1}^c , is fed in to generate the conditional forecast, $\hat{\mathbf{y}}_{t+1}^c = \hat{\mathbf{y}}_{t+1|t} + \mathbf{e}_{t+1}^c$. For time $t + 2$ a new one step-ahead unconditional forecast is generated, taking $\hat{\mathbf{y}}_{t+1}^c$ as the actual outcome and the process is repeated recursively in this manner until the last forecast horizon.

The second approach – suggested by Waggoner and Zha (1999) – considers the restrictions jointly and is concerned with the path of the unconstrained variables over the forecast period. It addresses the question “what is the most likely path of the unconstrained variables given the restrictions”. Write the joint restrictions as

$$\mathbf{R}\mathbf{y}_{t+1:t+h} = \mathbf{r} = \begin{pmatrix} \mathbf{R}_{t+1} & & & \mathbf{0} \\ & \mathbf{R}_{t+2} & & \\ & & \ddots & \\ \mathbf{0} & & & \mathbf{R}_{t+h} \end{pmatrix} \begin{pmatrix} \mathbf{y}_{t+1} \\ \mathbf{y}_{t+2} \\ \vdots \\ \mathbf{y}_{t+h} \end{pmatrix} = \begin{pmatrix} \mathbf{r}_{t+1} \\ \mathbf{r}_{t+2} \\ \vdots \\ \mathbf{r}_{t+h} \end{pmatrix} \quad (23)$$

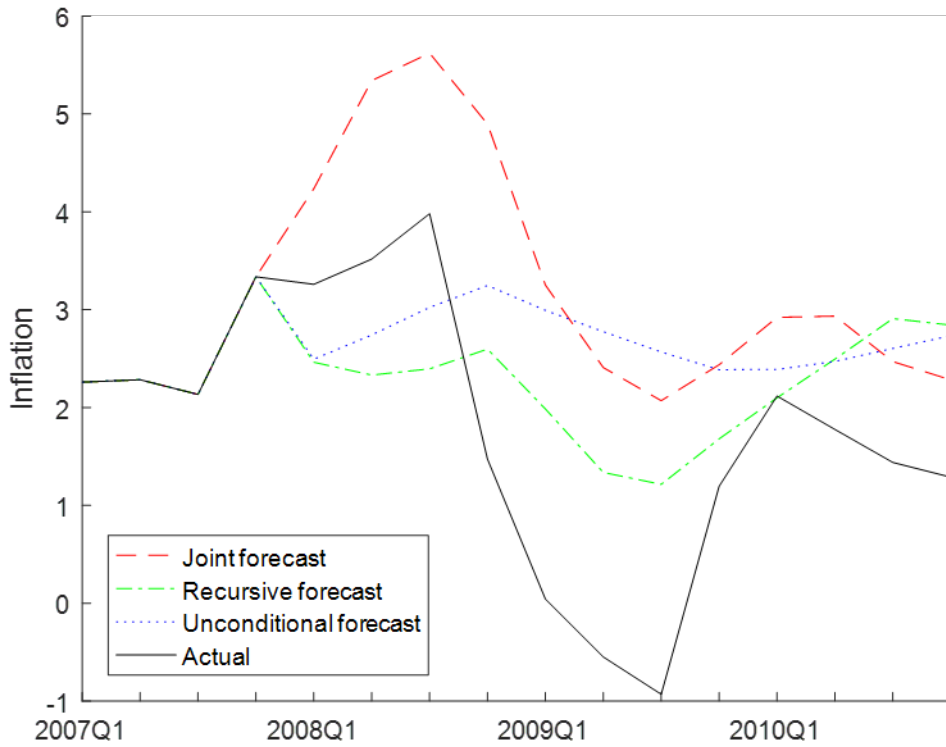
or

$$\mathbf{R}\tilde{\mathbf{e}}_{t+1:t+h} = \mathbf{r} - \mathbf{R}\hat{\mathbf{y}}_{t+1:t+h} = \mathbf{d} \quad (24)$$

for $\hat{\mathbf{y}}_{t+1:t+h} = (\hat{\mathbf{y}}'_{t+1|t}, \hat{\mathbf{y}}'_{t+2|t}, \dots, \hat{\mathbf{y}}'_{t+h|t})'$ the 1 to h step ahead unconditional forecasts and $\tilde{\mathbf{e}}_{t+1:t+h} = (\tilde{\mathbf{e}}'_{t+1|t}, \tilde{\mathbf{e}}'_{t+2|t}, \dots, \tilde{\mathbf{e}}'_{t+h|t})'$ the 1 to h step ahead forecast errors. The forecast errors are related to the disturbances $\mathbf{e}_{t+1:t+h} = (\mathbf{e}'_{t+1}, \mathbf{e}'_{t+2}, \dots, \mathbf{e}'_{t+h})'$ through the moving average coefficient matrices, e.g. $\tilde{\mathbf{e}}'_{t+i|t} = \sum_{j=0}^{i-1} \mathbf{C}_j \mathbf{e}'_{t+i-j}$ or $\tilde{\mathbf{e}}_{t+1:t+h} = \mathbf{C} \mathbf{e}_{t+1:t+h}$ and the restriction can be rewritten as $\tilde{\mathbf{R}} \mathbf{e}_{t+1:t+h} = \mathbf{R} \mathbf{C} \mathbf{e}_{t+1:t+h} = \mathbf{d}$. Finally let $\mathbf{\Omega}$ be the variance-covariance matrix of $\mathbf{e}_{t+1:t+h}$. A draw from the distribution of $\mathbf{e}_{t+1:t+h}$ conditional on $\tilde{\mathbf{R}} \mathbf{e}_{t+1:t+h} = \mathbf{d}$ can then be obtained in an analogous way to (22); see Waggoner and Zha (1999) and Jarocinski (2010) for details.

Figure 9 displays the conditional and unconditional forecasts of inflation based on model i . To obtain more precise estimates of the parameter values as of 2007Q4 we estimate the model on the full sample rather than data up to 2007Q4 where the estimates would suffer from “end of sample” uncertainty. This does accordingly not correspond to how a conditional forecasting exercise would have been conducted in real time; rather, it is based on today’s best estimate of the parameters of the model at the beginning of the financial crisis.

Figure 9. Conditional inflation forecasts using model i and estimation based on the full sample.



Note: Median forecasts. Parameter estimates based on full sample but dated 2007Q4. “Recursive forecast” refers to conditional forecasts having been generated based on the restrictions in equation (21). “Joint forecast” refers to conditional forecasts having been generated based on the restrictions in equation (24).

As can be seen from Figure 9, the joint conditional forecasts show larger swings than the recursive conditional forecasts. This is due to two factors: First, the joint forecasts take account for how shocks to inflation influences future unemployment whereas the shocks to inflation in the recursive forecast only depends on the

current shock to unemployment. Second, the volatility of inflation is much higher than the volatility of unemployment and shocks to inflation are much “cheaper” when considering the most likely path for the variables. Judging by the results in Figure 9, it is difficult to claim that there is much evidence of a missing disinflation. Admittedly the recursive conditional forecasts are somewhat lower than actual inflation during the first three quarters of 2008. However, the differences between the conditional forecasts and actual inflation are typically not very large and generally positive; they do not suggest that the model wanted substantially lower inflation than the one that materialized.

Since the forecasts in Figure 9 were based on parameters using the full sample (1990Q1-2017Q3), we also construct forecasts using data only up until 2007Q4 in order to see if this affects our conclusion. Using this shorter sample, we provide a better approximation to what the model would have suggested in real time.¹⁴ The results from this exercise are presented in Figure A4 in the Appendix. Some differences relative to the results presented in Figure 9 can be found, most importantly that the forecasts generated under joint restrictions are not associated with an initial upturn in inflation. Regardless of how the forecasts were made, it is clear that the model has no missing disinflation when it is estimated on this sample either.

Finally, we also generate conditional forecasts from the BVAR with time-invariant θ_t and covariance matrix based on the estimation sample 1990Q1-2007Q4. These results are presented in Figure A5 in the Appendix. As can be seen, there is still a discrepancy between the conditional forecasts generated using the recursive method and those based on the joint restrictions. It is utterly clear though that regardless of how the forecasts are made, the model does not suggest a sizable drop in inflation in response to the increasing unemployment rate. Once again, the model-based analysis suggests that there was no missing disinflation.

3.8 Sensitivity analysis

Having conducted the above analysis, we finally want to assess how sensitive our results are with respect to some key concepts, namely the inflation measure and our choice of priors. In this section we accordingly first conduct analysis using two alternative measures of inflation; we then vary some of our prior parameters (using the data employed in the main analysis).

3.8.1 CORE PCE INFLATION

When discussing monetary policy, it is common to relate central banks’ actions to “core” measures of inflation, that is, inflation measures that have had some volatile components removed. (Note that this is not related to “core inflation” which was discussed in Section 3.6; we can only agree that the terminology is

¹⁴ We do not use real-time data though, so this does not show exactly what the model would have predicted in real-time. However, given the purpose of the exercise, we do not believe that this is strictly necessary either.

unfortunate and could be confusing.) Since we have used PCE inflation in our analysis above, we here accordingly investigate if our results are robust to using core PCE inflation instead. In constructing this measure, food and energy prices have been removed from the PCE deflator. Inflation is calculated as $\pi_t^{core} = 100(P_t^c/P_{t-4}^c - 1)$ where P_t^c is the core PCE deflator at time t .

Figures A6 to A11 in the Appendix show selected results from this analysis based on model i . As can be seen from Figure A6, the impulse-response function of core PCE inflation with respect to a shock to the unemployment rate is qualitatively similar to what we saw in Figure 3 (when PCE inflation was used) in that there is an initial decline in inflation. It can be noted though that there appears to be somewhat more time variation in the effect when core PCE inflation is used. Turning to the slope of the Phillips curve in Figure A7, we find that using core PCE inflation does not change our findings from above; the results are very similar both qualitatively and quantitatively to those discussed in Section 3.4.

The variance decompositions in Figures A8 and A9, on the other hand, look distinctly different from those shown in Figures 6 and 7. We note that they have a much more even profile over time than when PCE inflation was used and that shocks to inflation also are found to be much less important. This latter fact is of course not particularly surprising as we have employed a less volatile measure of inflation.¹⁵ Regarding the estimated core inflation shown in Figure A10, we can see that a very similar story is being told as that in Figure 8. Core inflation was high leading up to the Great Recession but then fell and for the last few years of the sample it has been clearly below two percent.

Finally, turning to the conditional forecasts (based on estimation using the full sample) shown in Figure A11, they also confirm our conclusions from Section 3.7. The model does not suggest that there was any missing disinflation.

3.8.2 QUARTER-ON-QUARTER INFLATION

When modelling inflation it is also not obvious that it is the annual change in the PCE deflator that we want to use as our measure. Another relevant alternative is to use the (annualised) quarterly change, that is, we define PCE inflation as $\pi_t = 400(P_t/P_{t-1} - 1)$. Results based on this inflation measure and model i are shown in Figures A12 to A17 in the Appendix.

As can be seen from Figure A12, the impulse-response function of inflation with respect to the shocks to the unemployment rate looks qualitatively similar to that when annual changes were used (shown in Figure 3). Also the slope of the Phillips curve – given in Figure A13 – has roughly the same story attached to it as before; looking at the point estimate, the curve was flatter around the financial crisis than in the late 1990s or near the end of our sample.

¹⁵ The standard deviation of the shock to the inflation equation (not reported in detail but available upon request) ranges from approximately 0.06 to 0.1 over time.

Looking at the variance decompositions in Figures A14 and A15, we note that they also show the same patterns as in our main analysis. Shocks to inflation are less important for both inflation and the unemployment rate around the mid-1990s and peak in their contribution to the forecast error variance of both variables in 2010Q1.

As for core inflation, our previous picture is once again confirmed. Figure A16 shows that core inflation was above two percent going into the Great Recession and as before, it has been low near the end of the sample.

The last aspect that we consider is the conditional forecasts (once again based on estimation using the full sample). As can be seen from Figure A17, the conditional forecasts based on the recursive method are largely in line with the level of actual inflation, even though there is a tendency to under-predict inflation for the first three quarters of 2008 and over-predict it thereafter; in particular, the model does not do a good job in matching the substantial fall in inflation which took place in 2008Q4. Taken together though, we believe that while the model for the first three quarters wants slightly less inflation than what materialised, the overall picture is hardly one of missing disinflation. For the conditional forecasts based on the joint method, the picture is somewhat different. The model does a fairly good job in replicating the path of actual inflation, roughly catching its swings. However, the model tends to under-predict inflation during 2009 and 2010. Unlike the previous results, we here finally find something that could be interpreted as model support for the claim of missing disinflation. It should of course be kept in mind that this is based on the point forecast and while this – as pointed out above – is our best estimate, there is a fair amount of uncertainty associated with these forecasts.

3.8.3 PRIORS

The choice of prior parameters affects the marginal likelihood and can influence the ranking of the six models when comparing marginal likelihoods. As the difference between the models lie in the degree to which they allow for time-varying parameters or time-varying volatilities we focus on the prior parameters governing these features. More specifically we modify the prior means of the variance of the changes in the parameters and log-volatilities (η_t and ξ_t in the state equations (4) and (5)) by halving or doubling $S_{\theta i}$ and S_{hi} . The results of this exercise are displayed in Table 2. For ease of comparison we also repeat the results for our preferred specification (with $S_{\theta i} = 0.0004$ for coefficients on the lags of variables and 0.04 for intercepts and $S_{hi} = 0.04$) in column 2 of the Table.

Table 2. Marginal likelihood for different prior specifications.

Model	ln(p(y M))					
	$S_{\theta i}, S_{hi}$	$S_{\theta i}/2, S_{hi}/2$	$2 \times S_{\theta i}, S_{hi}/2$	$S_{\theta i}/2, 2 \times S_{hi}$	$2 \times S_{\theta i}, 2 \times S_{hi}$	$S_{\theta i}, 2 \times S_{hi}$
i	-101.8	-104.1	-106.0	-100.1	-103.4	-100.0
ii	-107.7	-109.1	-110.2	-109.1	-110.2	-107.7
iii	-102.2	-105.3	-104.0	-100.1	-98.7	-99.6
iv	-101.9	-104.0	-106.5	-100.3	-103.4	-100.1
v	-102.4	-105.0	-105.0	-100.0	-100.0	-100.0
vi	-112.3	-112.3	-112.3	-112.3	-112.3	-112.3

Note: The models are the following: i) time-varying θ_t and stochastic volatility, ii) time-varying θ_t and time-invariant covariance matrix, iii) time-varying γ_t , time-invariant β_t and stochastic volatility, iv) time-varying β_t , time-invariant γ_t and stochastic volatility, v) time-invariant θ_t and stochastic volatility and vi) time-invariant θ_t and covariance matrix.

The log marginal likelihood for the best model is shown in **bold and italics**, the second best model in **bold** and the third best in *italics*.

With the exception of model *vi*, which does not depend on $S_{\theta i}$ and S_{hi} , decreasing or increasing $S_{\theta i}$ tend to lead to a lower marginal likelihood, decreasing S_{hi} tend to lead to a lower marginal likelihood and increasing S_{hi} tend to lead to a higher marginal likelihood. While the ranking of the models change with the prior parameters the difference in marginal likelihood between the models with stochastic volatility (models *i*, *iii*, *iv* and *v*) is quite small and does not affect our overall conclusion that there is support for both time-varying parameters and time-varying volatility.

To illustrate the effect of the prior specification on the results reported above Figures A18 to A21 show selected results for the prior setting in column $S_{\theta i}/2, 2 \times S_{hi}$ of Table 2 that favors time-varying volatility and limits variability in the parameters. That is, $S_{\theta i} = 0.0002$ except for the intercepts where it is set to 0.02 and $S_{hi} = 0.08$. Overall the results are largely unaffected by the change in the prior specification. The shape of the impulse responses for inflation (Figures A18 and A19) are quite close although there is more variability in the volatilities as is evident from the lag zero response in inflation to a shock to inflation in Figure A18. The slope of the Phillips curve in Figure A20 shows less variability than Figure 4 but the overall shape is very similar, confirming the conclusions in section 3.4. The conditional forecasts in Figure A21 are also qualitatively very similar to Figure 9, confirming the conclusion of Section 3.7 that there was no missing disinflation.

4. Conclusions

The Phillips curve is a popular tool in macroeconomics which, among other things, is used to assess inflationary pressure in the economy. However, its usefulness for this purpose depends on having a good econometric specification. In this paper we have added empirical evidence concerning the properties of the US Phillips curve by analysing different specifications of BVAR models. Conducting Bayesian model selection using recently developed methods for this purpose, we find that that both time-varying parameters and stochastic volatility appear to be features that the BVAR should incorporate. Our results indicate that the Phillips curve may have been somewhat flatter between 2005 and 2013 than in the decade preceding that period. This fact – together with a number of large positive shocks and high core inflation – can help explain why inflation was high around the Great recession.

In addition, our findings from the conditional forecasting exercises conducted suggest that as far as the models are concerned, there may not have been that much of a missing disinflation associated with the Great Recession. Conditioning on the increase in the unemployment rate that actually took place, different model specifications and restrictions yield somewhat different results. However, the models did not in general predict that inflation should have been lower than what turned out to be the case. While we believe that all of the above information is useful, we also want to point out that it to some extent shows the downside of working with small non-structural models. It is after all difficult to give a deep economic interpretation of, for example, shocks to the inflation equation and the estimated core inflation.

From a policy perspective, our results imply two things. First, policy institutions and forecasters need to consider the fact that both dynamic relations and the volatility of shocks may be changing over time when building econometric models. Since methods for model selection have been improved, it is also the case that the econometric modelling choice can stand on firmer ground nowadays; where the econometrician previously largely had to make an assumption regarding the relevance of different specifications, they can now be formally compared. Second, the stubbornly low inflation in the United States the last few years appears, at least to some extent, to be due to low inflationary pressures which have affected the estimated core inflation. This means that there might be scope for the federal funds rate to be quite low even in light of a low unemployment rate. Such a policy response could be needed in order to let inflationary pressures build up in the economy and reliably bring inflation back up to the target level.

References

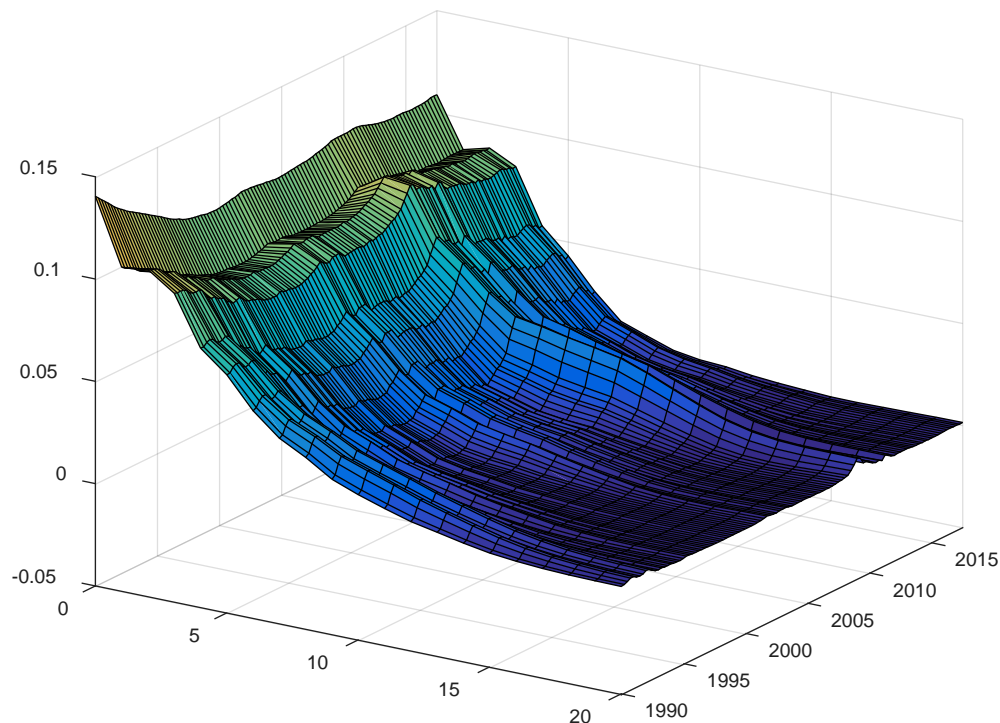
- Abbate, A., Eickmeier, S., Lemke, W. and Marcellino, M. (2016), “The Changing International Transmission of Financial Shocks: Evidence from a Classical Time-Varying FAVAR”, *Journal of Money, Credit and Banking* 48, 573-601.
- Andrews, D. K. (1993), “Tests for Parameter Instability and Structural Change with Unknown Change Point”, *Econometrica* 61, 821-856.
- Ball, L. M. (2006), “Has Globalization Changed Inflation?”, NBER Working Paper No. 12687.
- Baumeister, C. and Kilian, L. (2013), “Real-Time Analysis of Oil Price Risks Using Forecast Scenarios”, *IMF Economic Review* 62, 119-145.
- Bean, C. (2006), “Globalisation and Inflation”, *Bank of England Quarterly Bulletin* 46, 468-475.
- Bernanke, B. S. and Mihov, I. (1998), “Measuring Monetary Policy”, *Quarterly Journal of Economics* 113, 869-902.
- Bianchi, F. and Civelli, A. (2015), “Globalization and Inflation: Evidence from a Time Varying VAR”, *Review of Economic Dynamics* 18, 406-433,
- Blanchard, O., (2016), “The Phillips Curve: Back to the ’60s?”, *American Economic Review* 106, 31-34.
- Blanchard, O., Cerutti, E. and Summers, L. (2015), “Inflation and Activity – Two Explorations and Their Monetary Policy Implications”, NBER Working Paper No. 21726.
- Chan, J. C. C. and Eisenstat, E. (2015), “Marginal likelihood estimation with the cross-entropy method”, *Econometric Reviews* 34, 256-285.
- Chan, J. C. C. and Eisenstat, E. (2018), “Bayesian Model Comparison for Time-Varying Parameter VARs with Stochastic Volatility”, Forthcoming in *Journal of Applied Econometrics*.
- Chan, J. C. C. and Jeliazkov, I. (2009), “Efficient Simulation and Integrated Likelihood Estimation in State Space Models”, *International Journal of Mathematical Modelling and Numerical Optimisation* 1, 101-120.
- Clark, T. E. (2011), “Real-Time Density Forecasts from VARs with Stochastic Volatility”, *Journal of Business and Economic Statistics* 29, 327-341.
- Clark, T. E. and McCracken, M. W. (2006), “The Predictive Content of the Output Gap for Inflation: Resolving In-Sample and Out-of-Sample Evidence”, *Journal of Money, Credit and Banking* 38, 1127-1148.
- Clark, T. E. and McCracken, M. W. (2014), “Evaluating Conditional Forecasts from Vector Autoregressions”, Federal Reserve Bank of St. Louis Working Paper 2014-25.
- Clark, T. E. and Doh, T. (2014), “Evaluating Alternative Models of Trend Inflation”, *International Journal of Forecasting* 30, 426-448.
- Coibion, O. and Gorodnichenko, Y. (2015), “Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation”, *American Economic Journal: Macroeconomics* 7, 197-232
- Cogley, T., and Sargent, T. J. (2001). “Evolving Post-World War II U.S. Inflation Dynamics”, *NBER Macroeconomics Annual* 16, 331-373.
- Cogley, T., and Sargent, T. J. (2005). “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US”, *Review of Economic Dynamics* 8, 262-302.
- Dotsey, M., Fujita, S. and Stark, T. (2017), “Do Phillips Curves Conditionally Help to Forecast Inflation?”, Working Paper No. 17-26, Federal Reserve Bank of Philadelphia.

- Eisenstat, E., Chan, J. and Strachan, R. (2016), “Stochastic Model Specification Search for Time-Varying Parameter VARs”, *Econometric Reviews* 35, 1638-1665.
- Faust, J. and Wright, J. H. (2013), “Forecasting Inflation”, In: Elliott, G., Granger, C. and Timmermann, A. (eds), *Handbook of Economic Forecasting*, vol. 2A, North-Holland, Amsterdam.
- Friedman, M. (1968) “The Role of Monetary Policy”, *American Economic Review* 58, 1-17.
- Galí, J. and Gertler, M. (1999), “Inflation Dynamics: A Structural Econometric Analysis”, *Journal of Monetary Economics* 44, 195-222.
- Gaiotti, E. (2010), “Has Globalization Changed the Phillips Curve? Firm-Level Evidence on the Effect of Activity on Prices”, *International Journal of Central Banking* 6, 51-84.
- Giannone, D., Lenza, M., Momferatou, D. and Onorante, L. (2014), “Short-Term Inflation Projections: A Bayesian Vector Autoregressive Approach”, *International Journal of Forecasting* 30, 635-644.
- Hamilton, J. D. and Herrera, A. M. (2004), “Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy”, *Journal of Money, Credit and Banking* 36, 265-286.
- Ihrig, J., Kamin, S. B., Lindner, D. and Marquez, J. (2010), “Some Simple Tests of the Globalization and Inflation Hypothesis”, *International Finance* 13, 343-375.
- International Monetary Fund (2007), *Regional Economic Outlook, Western Hemisphere*, November 2007.
- International Monetary Fund (2013), *World Economic Outlook*, April 2013.
- Jarocinski, M. (2010), “Conditional Forecasts and Uncertainty About Forecast Revisions in Vector Autoregressions”, *Economics Letters* 108, 257-259.
- Kim, S., Shepherd, N. and Chib, S. (1998), “Stochastic volatility: Likelihood inference and comparison with ARCH models”, *Review of Economic Studies*, 65, 361–393.
- King, R. G. and Watson, M. W. (1994), “The post-War U.S. Phillips Curve: A Revisionist Econometric History”, *Carnegie-Rochester Conference Series on Public Policy* 41, 157-219.
- Knotek, E. S. and Zaman, S. (2017), “Have Inflation Dynamics Changed?”, *Federal Reserve Bank of Cleveland Economic Commentary* 2017-21.
- Koop, G., Leon-Gonzalez, R. and Strachan, R. W. (2009), “On the Evolution of the Monetary Policy Transmission Mechanism”, *Journal of Economic Dynamics and Control* 33, 997-1017.
- Kuttner, K. and Robinson, T. (2010), “Understanding the Flattening Phillips Curve”, *North American Journal of Economics and Finance* 21, 110-125.
- Laséen, S. and Taheri Sanjani, M. (2016), “Did the Global Financial Crisis Break the U.S. Phillips Curve?”, IMF Working Paper 16/126.
- Leeper, E. M. and Zha, T. (2003), “Modest Policy Interventions”, *Journal of Monetary Economics* 50, 1673-1700.
- Österholm, P. (2009), “Incorporating Judgement in Fan Charts”, *Scandinavian Journal of Economics* 111, 387-415.
- Österholm, P. and Zettelmeyer, J. (2008), “The Effect of External Conditions on Growth in Latin America”, *IMF Staff Papers* 55, 595-623.
- Phelps, E. S. (1967), “Phillips Curves, Expectations of Inflation, and Optimal Unemployment over Time”, *Economica* 34, 254-281.
- Phillips, A. W. (1958), “The Relationship between Unemployment and the Rate of Change of Money Wages in the United Kingdom 1861-1957”, *Economica* 25, 283-299.

- Primiceri, G. (2005). "Time Varying Structural Vector Autoregressions and Monetary Policy", *Review of Economic Studies* 72, 821-852.
- Sims, C. A. (1982), "Policy Analysis with Econometric Models", *Brookings Papers on Economic Activity* 1982:1, 107-152.
- Sims, C. A., (2001), "Comment on Sargent and Cogley's 'Evolving Post-World War II U.S. Inflation Dynamics'", *NBER Macroeconomics Annual* 16, 373-379.
- Sims, C. A. and Zha, T. (2006), "Were There Regime Switches in U.S. Monetary Policy?", *American Economic Review* 96, 54-81.
- Stock, J. H., (2001), "Comment on Sargent and Cogley's 'Evolving Post-World War II U.S. Inflation Dynamics'", *NBER Macroeconomics Annual* 16, 379-387.
- Stock, J. H. and Watson, M. W. (1999), "Forecasting Inflation", *Journal of Monetary Economics* 44, 293-335.
- Stock J. H. and Watson M. W. (2012), "Disentangling the Channels of the 2007-09 Recession", *Brookings Papers on Economic Activity* 44, 81-156.
- Svensson, L. E. O. (2015), "The Possible Unemployment Cost of Average Inflation below a Credible Target," *American Economic Journal: Macroeconomics* 7, 258-296.
- Sveriges Riksbank (2016), *Monetary Policy Report*, February.
- Villani, M. (2009), "Steady-State Priors for Vector Autoregressions", *Journal of Applied Econometrics* 24, 630-650.
- Waggoner, D. F. and Zha, T. (1999), "Conditional Forecasts in Dynamic Multivariate Models", *Review of Economics and Statistics* 81, 639-651.

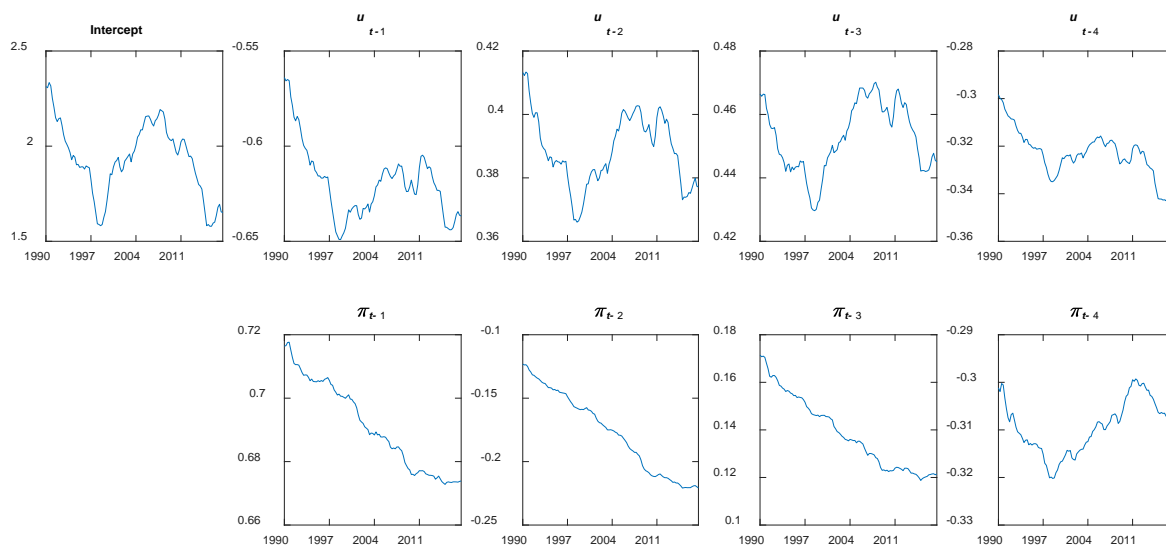
Appendix

Figure A1. Impulse-response functions for model i: The effect of shocks to the unemployment rate on the unemployment rate, 1990Q1-2017Q3.



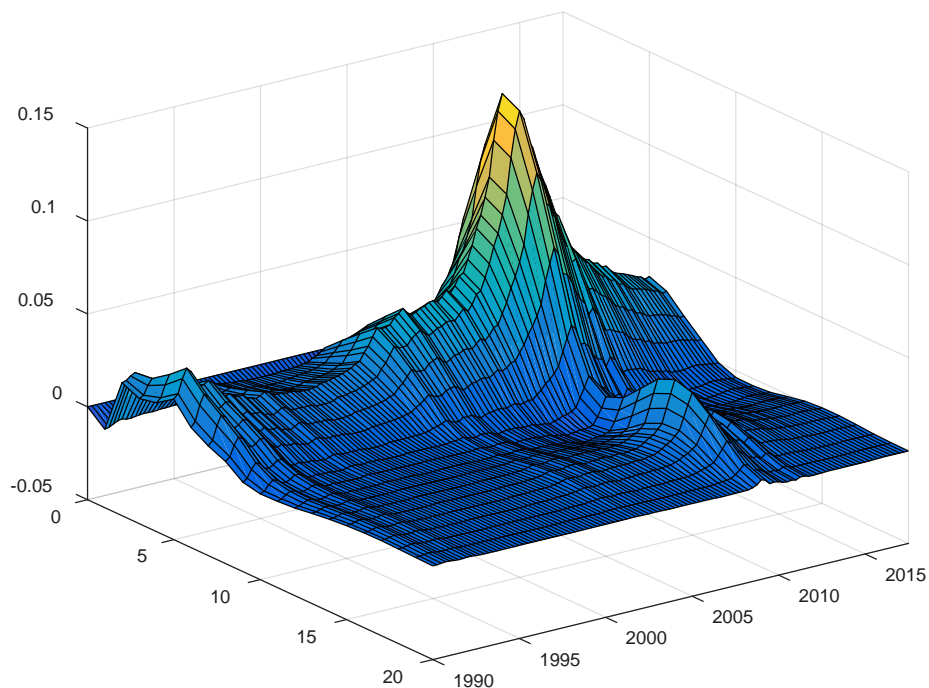
Note: Size of impulse is one standard deviation. Effect on the unemployment rate in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A2. Estimated coefficients of the inflation equation of model i, 1990Q1-2017Q3.



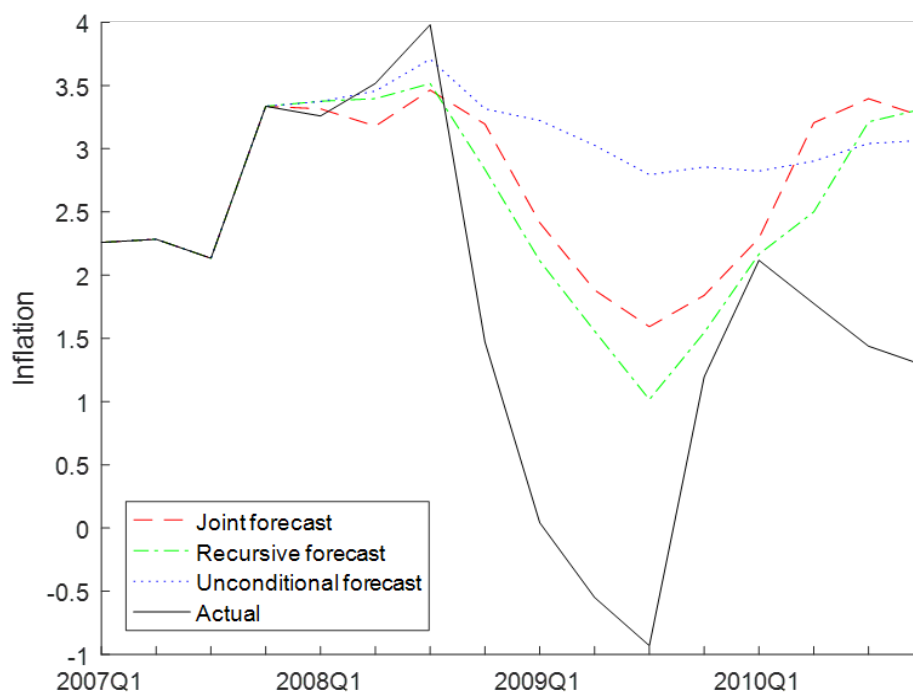
Note: Time on horizontal axes.

Figure A3. Impulse-response functions for model i: The effect of shocks to inflation on the unemployment rate, 1990Q1-2017Q3.



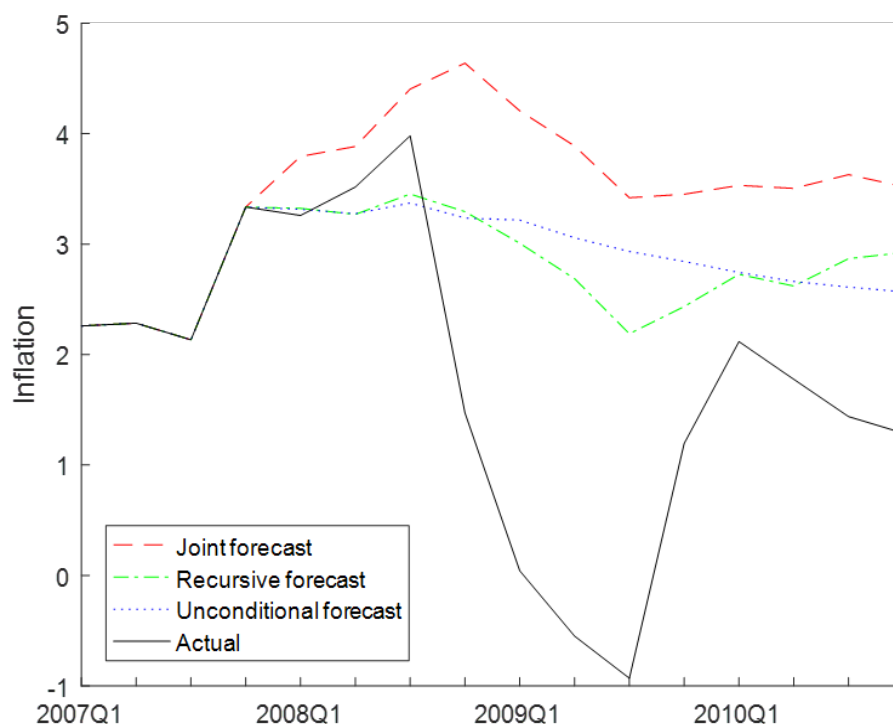
Note: Size of impulse is one standard deviation. Effect on the unemployment rate in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A4. Conditional inflation forecasts using model i and estimation based on the sample 1990Q1-2007Q4.



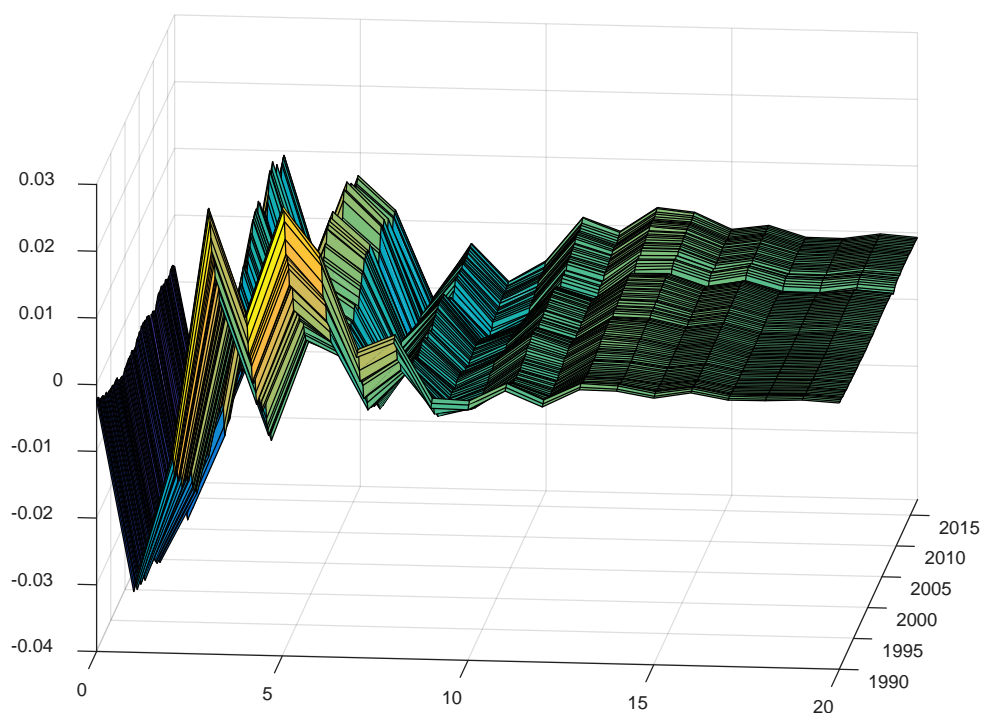
Note: Median forecasts. Parameter estimates are based on the sample 1990Q1-2007Q4. "Recursive forecast" refers to conditional forecasts having been generated based on the restrictions in equation (21). "Joint forecast" refers to conditional forecasts having been generated based on the restrictions in equation (24).

Figure A5. Conditional inflation forecasts using model vi and estimation based on the sample 1990Q1-2007Q4.



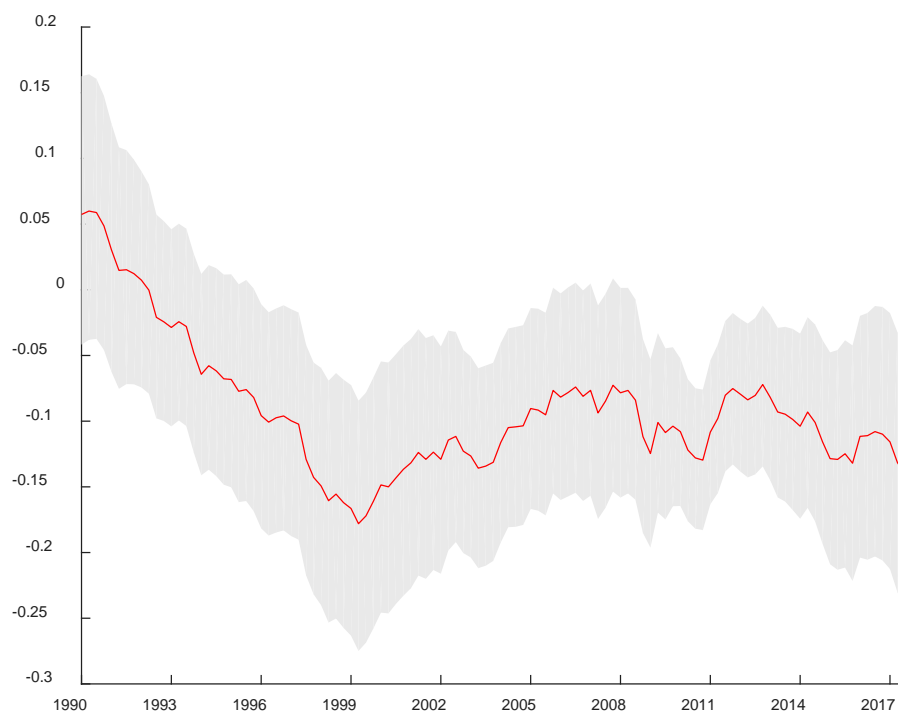
Note: Median forecasts. Parameter estimates are based on the sample 1990Q1-2007Q4. "Recursive forecast" refers to conditional forecasts having been generated based on the restrictions in equation (21). "Joint forecast" refers to conditional forecasts having been generated based on the restrictions in equation (24).

Figure A6. Impulse-response function for model i using core PCE inflation: The effect of shocks to the unemployment rate on core PCE inflation, 1990Q1-2017Q3.



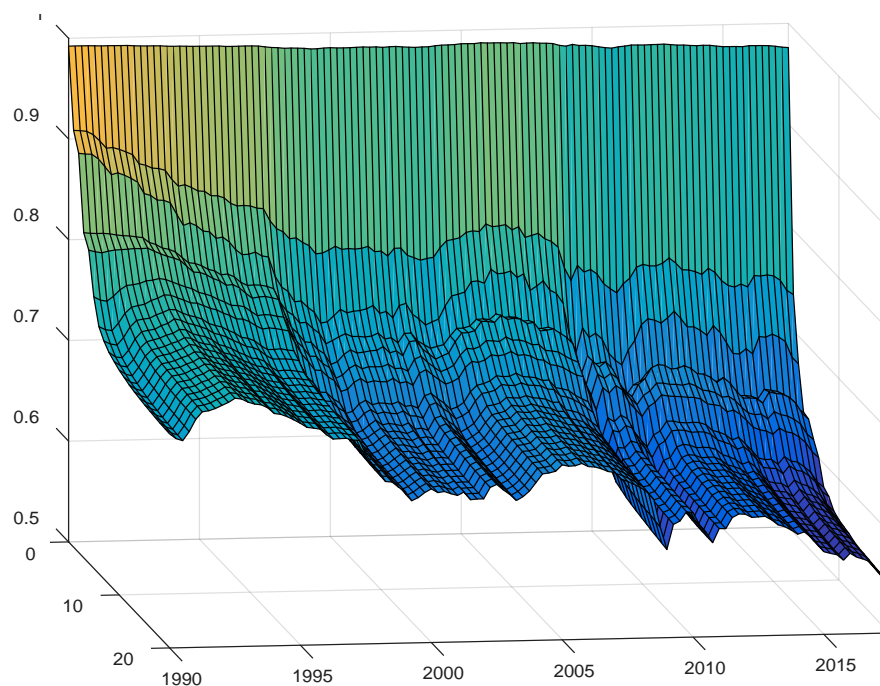
Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A7. Estimated slope of the Phillips curve for model i using core PCE inflation, 1990Q1-2017Q3.



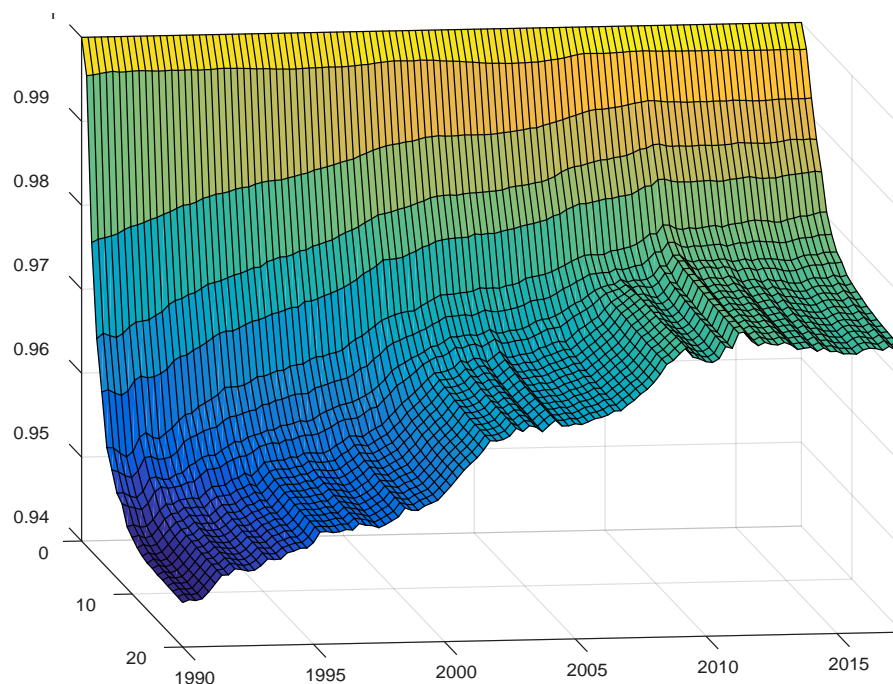
Note: Slope is measured as the sum of the coefficients on the lags of unemployment in the inflation equation in equation (6). Coloured band is 68% equal tail credible interval.

Figure A8. Variance decomposition for model i using core PCE inflation: Share of forecast error variance of PCE inflation explained by shocks to PCE inflation, 1990Q1-2017Q3.



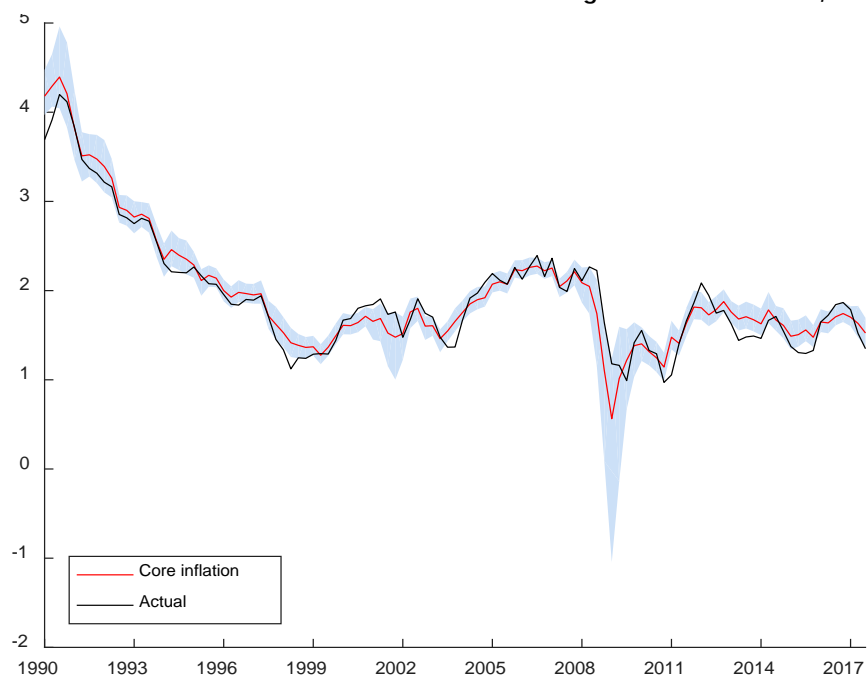
Note: Share of forecast error variance explained by shocks to inflation on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A9. Variance decomposition for model i using core PCE inflation: Share of forecast error variance of unemployment rate explained by shocks to the unemployment rate, 1990Q1-2017Q3.



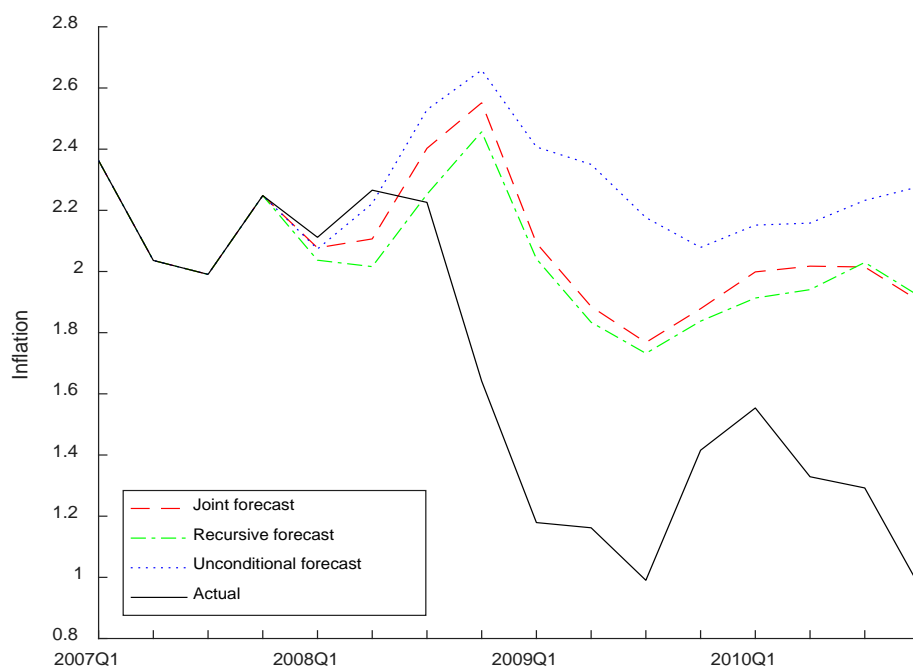
Note: Share of forecast error variance explained by shocks to the unemployment rate on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A10. Estimated "core inflation" from model i using core PCE inflation, 1990Q1-2017Q3.



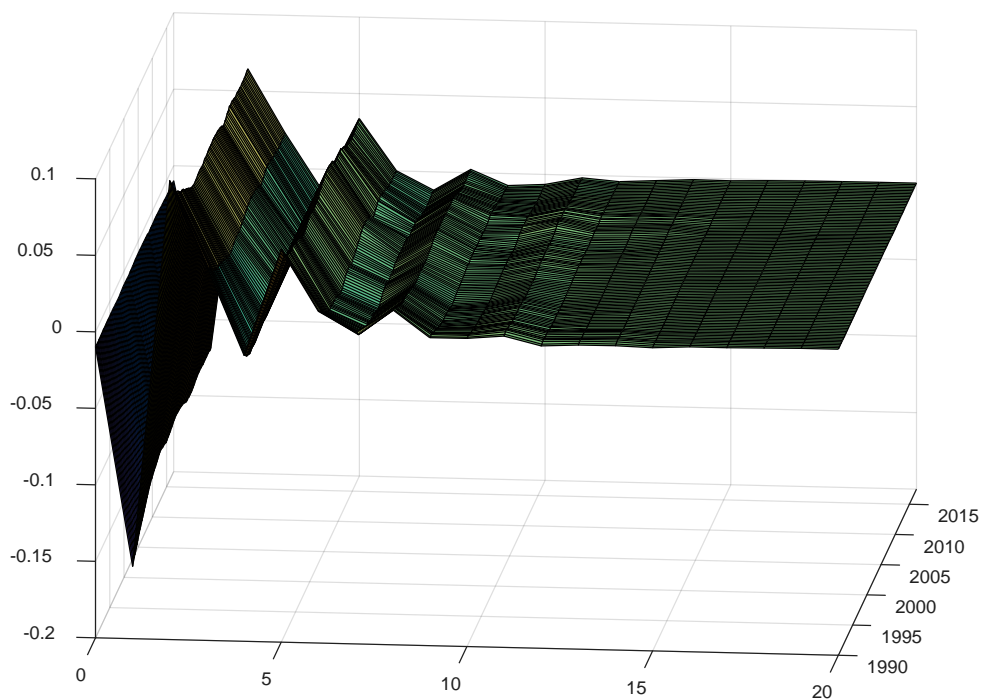
Note: "Core inflation" and inflation are measured in percent. Coloured band is 68% equal tail credible interval.

Figure A11. Conditional inflation forecasts using model i and core PCE inflation; estimation based on the full sample.



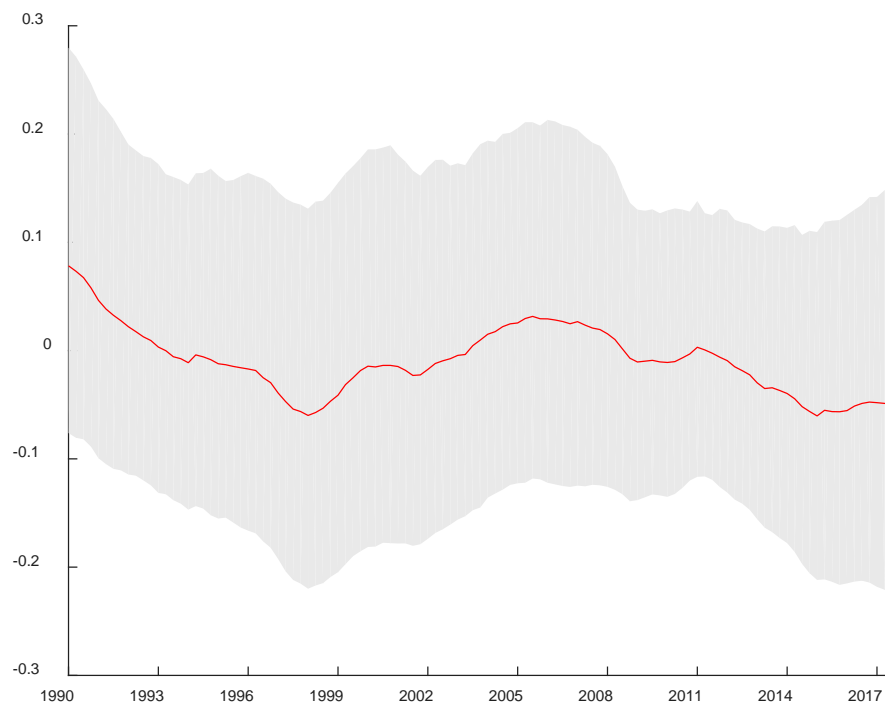
Note: Median forecasts. Parameter estimates based on full sample but dated 2007Q4. "Recursive forecast" refers to conditional forecasts having been generated based on the restrictions in equation (21). "Joint forecast" refers to conditional forecasts having been generated based on the restrictions in equation (24).

Figure A12. Impulse-response function for model i using quarter-on-quarter PCE inflation: The effect of shocks to the unemployment rate on PCE inflation, 1990Q1-2017Q3.



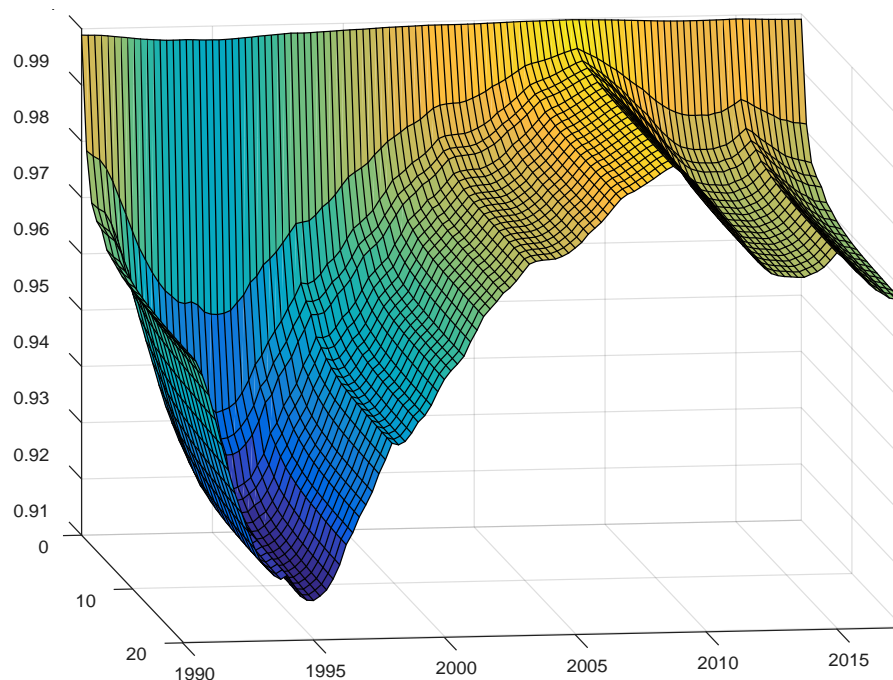
Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A13. Estimated slope of the Phillips curve for model i using quarter-on-quarter PCE inflation, 1990Q1-2017Q3.



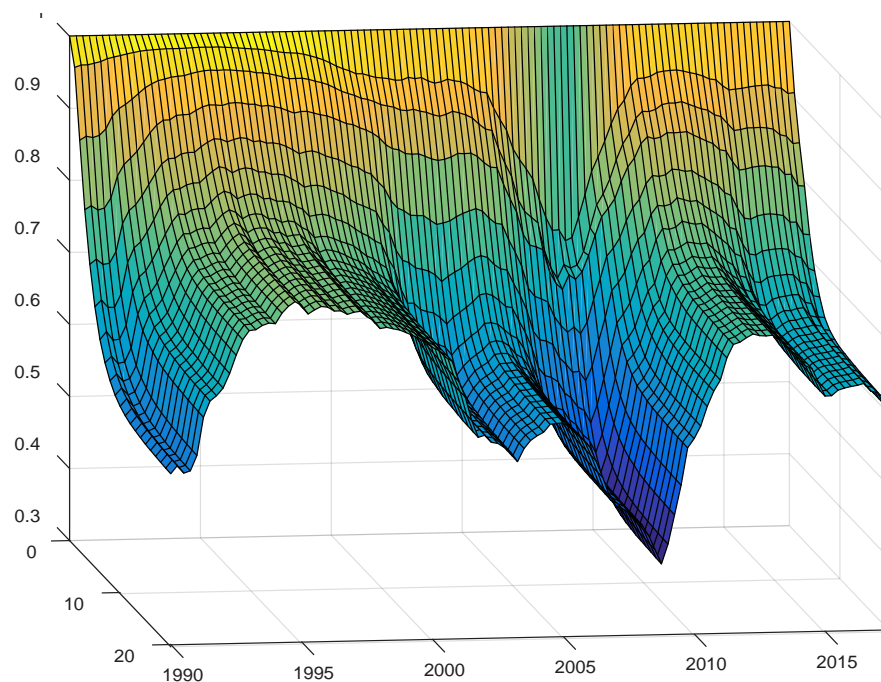
Note: Slope is measured as the sum of the coefficients on the lags of unemployment in the inflation equation in equation (6). Coloured band is 68% equal tail credible interval.

Figure A14. Variance decomposition for model i using quarter-on-quarter PCE inflation: Share of forecast error variance of PCE inflation explained by shocks to PCE inflation, 1990Q1-2017Q3.



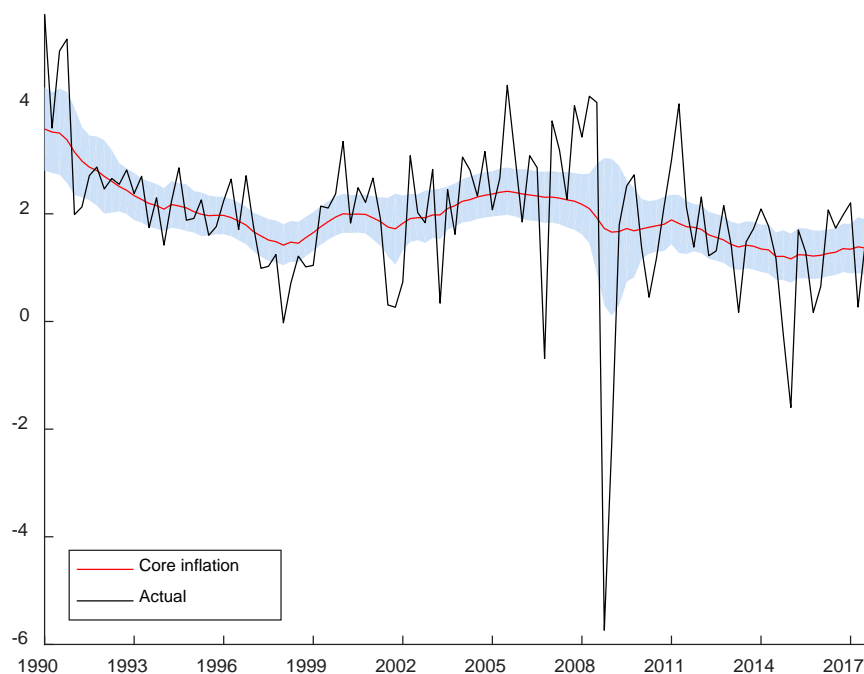
Note: Share of forecast error variance explained by shocks to inflation on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A15. Variance decomposition for model i using quarter-on-quarter PCE inflation: Share of forecast error variance of unemployment rate explained by shocks to the unemployment rate, 1990Q1-2017Q3.



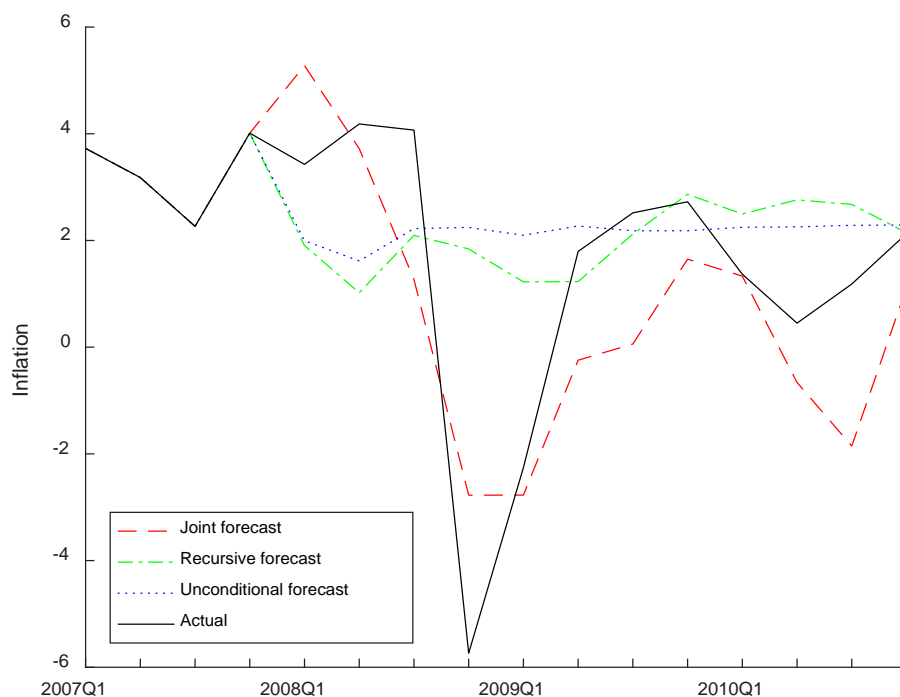
Note: Share of forecast error variance explained by shocks to the unemployment rate on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A16. Estimated “core inflation” from model i using quarter-on-quarter PCE inflation, 1990Q1-2017Q3.



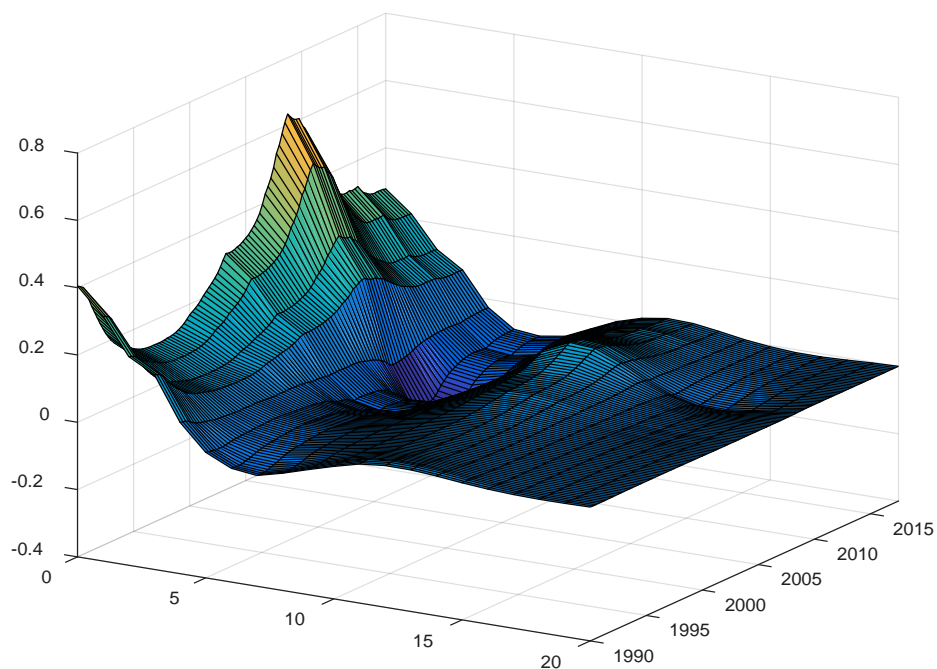
Note: “Core inflation” and inflation are measured in percent. Coloured band is 68% equal tail credible interval.

Figure A17. Conditional inflation forecasts using model i and quarter-on-quarter PCE inflation; estimation based on the full sample.



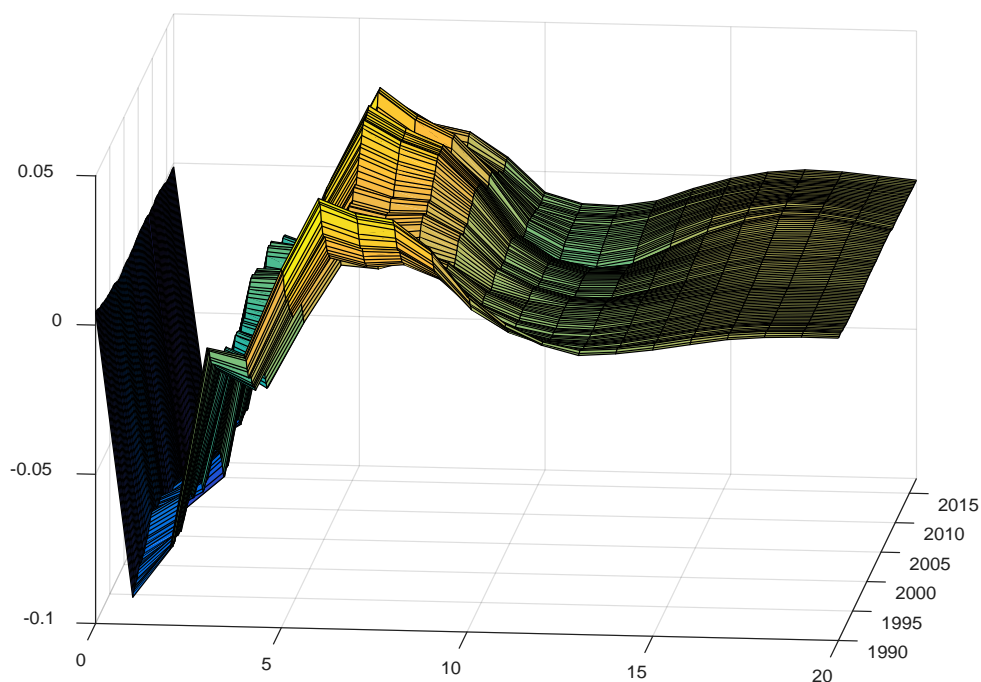
Note: Median forecasts. Parameter estimates based on full sample but dated 2007Q4. "Recursive forecast" refers to conditional forecasts having been generated based on the restrictions in equation (21). "Joint forecast" refers to conditional forecasts having been generated based on the restrictions in equation (24).

Figure A18. Impulse-response function for model i: The effect of shocks to inflation on inflation, 1990Q1-2017Q3. Alternative prior with $S_{\theta l} = 0.00005, 0.005$ and $S_{hl} = 0.02$.



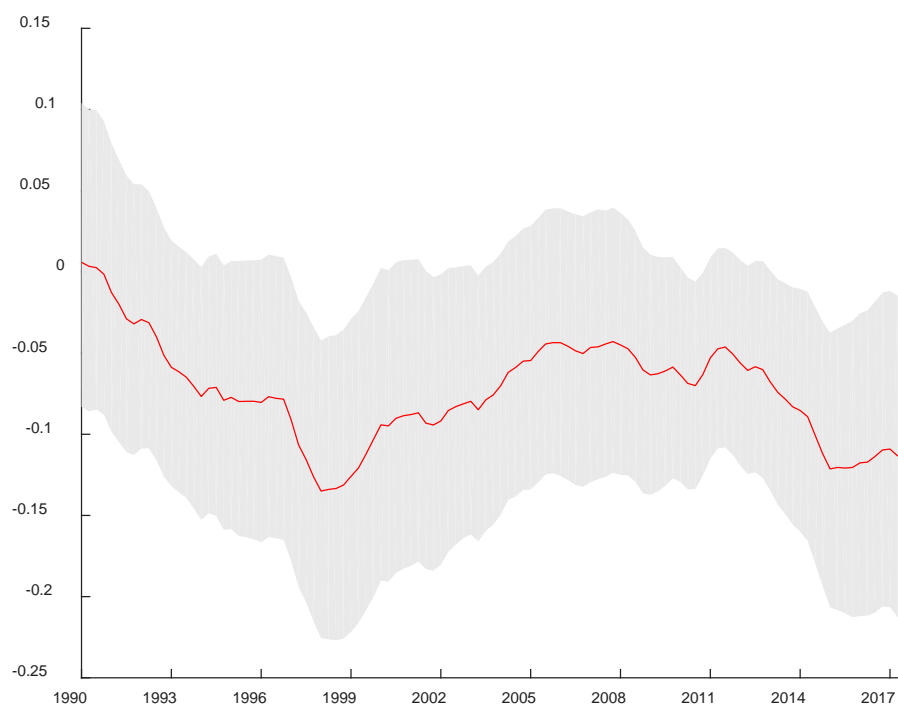
Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A19. Impulse-response function for model i: The effect of shocks to the unemployment rate on PCE inflation, 1990Q1-2017Q3. Alternative prior with $S_{\theta l} = 0.00005, 0.005$ and $S_{hl} = 0.02$.



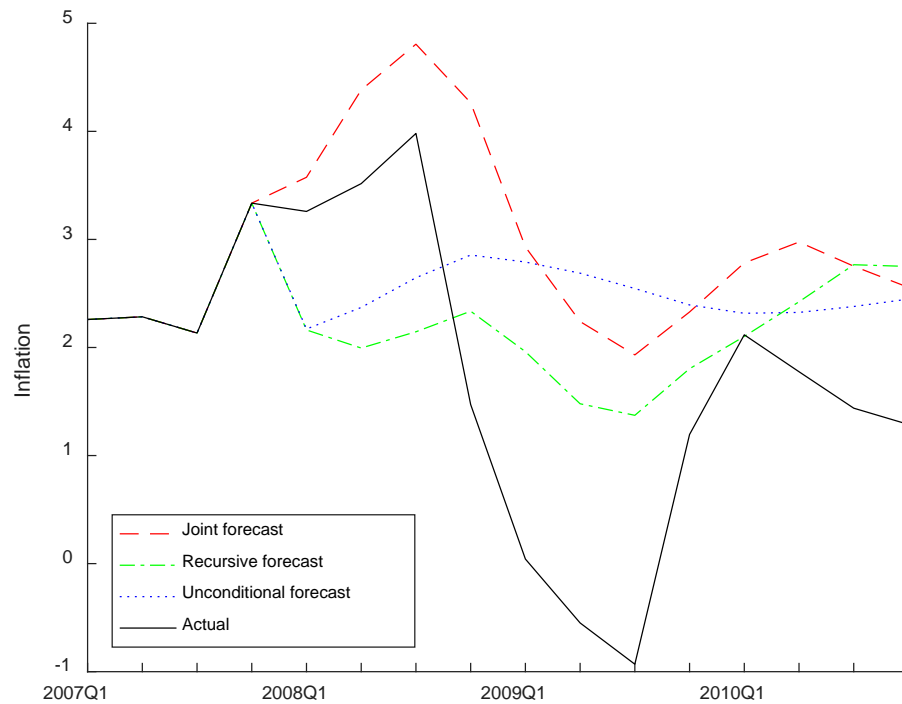
Note: Size of impulse is one standard deviation. Effect on inflation in percentage points on vertical axis. Horizon in quarters and time on horizontal axes.

Figure A20. Estimated slope of the Phillips curve for model i, 1990Q1-2017Q3. Alternative prior with $S_{\theta l} = 0.00005, 0.005$ and $S_{hl} = 0.02$.



Note: Slope is measured as the sum of the coefficients on the lags of unemployment in the inflation equation in equation (6). Coloured band is 68% equal tail credible interval.

Figure A21. Conditional inflation forecasts using model i; estimation based on the full sample. Alternative prior with $S_{\theta i} = 0.00005, 0.005$ and $S_{hi} = 0.02$.



Note: Median forecasts. Parameter estimates based on full sample but dated 2007Q4. "Recursive forecast" refers to conditional forecasts having been generated based on the restrictions in equation (21). "Joint forecast" refers to conditional forecasts having been generated based on the restrictions in equation (24).