Online Appendix for
A Model of the Fed's View on Inflation

Thomas Hasenzagl[1], Filippo Pellegrino[2], Lucrezia Reichlin[3], and Giovanni Ricco[4]

[1] *University of Minnesota*
[2] *London School of Economics and Political Science*
[3] *London Business School, Now-Casting Economics, and CEPR*
[4] *University of Warwick, OFCE-SciencesPo, and CEPR*

First draft: December 2017

This version: 25th August 2020

Contents

A Metropolis-Within-Gibbs **3**
A.1 Algorithm . . . 3

B Posteriors of all parameters **7**

C Robustness to Priors **11**

D Building up the Model **13**
D.1 Model A: small-scale model with AR(1) cycles and without dynamic heterogeneity . . . 13
D.2 Model B: Model A with ARMA(2,1) cycles . . . 18
D.3 Model C: Model B with EP cycle . . . 22
D.4 Model D: Model C with employment and core inflation . . . 26
D.5 Model E: Model D with University of Michigan inflation expectations . . . 30
D.6 Model F: Benchmark Model . . . 34

E Employment-to-Population Ratio **38**

F Global Activity **42**
F.1 Correlation of Cycles with Global Activity Indicators . . . 42
F.2 A Model with Global Indicators . . . 43

G Model Forecasting Performance **47**

H Priors for the Long-Run **52**

Appendix A Metropolis-Within-Gibbs

Starting with the work of Beveridge and Nelson (1981), Harvey (1985), and Clark (1987) several methodologies have been suggested in the literature to estimate tend-cycle models with unobserved components. As discussed in Harvey et al. (2007), frequentist techniques tend to deliver inaccurate estimates and – as a result – implausible cycles and trends, due to large estimation uncertainty. Conversely, Bayesian methods, which allow for the incorporation of a-priori knowledge into the model estimation, make it possible to consistently estimate both univariate and multivariate trend-cycle decompositions via efficient numerical methods.

In estimating our model we adopt a Metropolis-Within-Gibbs algorithm. However, since this approach tend to have slow performances in large dimensions, we run a simulation smoother only after the burn-in period to gain computational speed. In fact, during the burn-in period, we only employ a Kalman filter with exact diffuse initial conditions to estimate the likelihood function, as described in Koopman and Durbin (2000) and Durbin and Koopman (2012). This significantly increases the speed of the estimation which, given the large state-space of our model, is useful.

A.1 Algorithm

The algorithm is structured in two blocks: (1) a Partially Adaptive Metropolis (e.g., Herbst and Schorfheide, 2015) step for the estimation of the state-space parameters, (2) a Gibbs sampler to draw the unobserved states conditional on the model parameters. In a Partially Adaptive Metropolis the variance covariance matrix, Σ, of the candidate distribution is generated in an initialisation step.

Algorithm: Metropolis-Within-Gibbs

Initialisation

For $s = 1, \dots, n_s$ $(n_s = 40000)$

1. *Metropolis Algorithm*

   i. Draw a candidate vector for the unbounded parameters (θ_*), from a multivariate normal distribution with mean θ_{s-1} and variance $\omega\mathbb{I}$, where ω is a scaling constant used to get an acceptance rate between 25% and 35%

   ii. Set

$$\theta_s = \begin{cases} \theta_* & \text{with probability } \eta \\ \theta_{s-1} & \text{with probability } 1-\eta \end{cases} \tag{1}$$

   for

$$\eta = min\left(1, \frac{p(y \mid f(\theta_*)^{-1})\, p(f(\theta_*)^{-1})\, J(\theta_*)}{p(y \mid f(\theta_{s-1})^{-1})\, p(f(\theta_{s-1})^{-1})\, J(\theta_{s-1})}\right) \tag{2}$$

2. Discard the first $s = 1, \ldots, n_0$ $(n_0 = 20000)$ draws of θ_s.

Recursion

1. *Metropolis Algorithm*

   Set Σ to the sample covariance of the chain of θ_s, $(s = \{n_0, \ldots, n_s\})$, from the Initialisation step.

   For $q = 1, \ldots, n_q$ $(n_q = 20000)$

   i. Draw a candidate vector for the parameters (θ_*), from a multivariate normal distribution with mean θ_{q-1} and variance $\omega\Sigma$, where ω is set to have an acceptance rate between 25% and 35%

   ii. Set

$$\theta_q = \begin{cases} \theta_* & \text{with probability } \eta \\ \theta_{q-1} & \text{with probability } 1-\eta \end{cases} \tag{3}$$

   where η is defined as in the Initialisation step.

2. *Gibbs sampling*

   For $n_q > n_\emptyset$ for $n_\emptyset = 10000$ (burn-in period), apply the univariate approach for multivariate time series of Koopman and Durbin (2000) to the simulation smoother proposed in Durbin and Koopman (2002) to sample the

unobserved states, conditional on the parameters. In doing so, we follow the refinement proposed in Jarociński (2015).

3. Discard the first $q = 1, \ldots, n_\emptyset$ draws of θ_q.

Jacobian Most of these parameters are constrained (or bounded) in their support (e.g. the variances of the shocks are greater than zero). The standard approach used to tackle this problem is to transform the bounded parameters (Θ) so that the support of the transformed parameters (θ) is unbounded. Our Metropolis algorithm draws the model parameters in the unbounded space in order to avoid a-priori rejections and to obtain a more efficient estimation routine.[1] The following transformations have been applied to parameters with Normal, Inverse-Gamma and Uniform priors, respectively:

$$\begin{aligned}
\theta_j^N &= \Theta_j^N \\
\theta_j^{IG} &= \ln(\Theta_j^{IG} - a_j) \\
\theta_j^U &= \ln\left(\frac{\Theta_j^U - a_j}{b_j - \Theta_j^U}\right)
\end{aligned}$$

Where a_j and b_j are the lower and the upper bounds for the j-th parameter. These transformations are functions $f(\Theta) = \theta$, with inverses $f(\theta)^{-1} = \Theta$ given by:

$$\begin{aligned}
\Theta_j^N &= \theta_j^N \\
\Theta_j^{IG} &= \exp(\theta_j^{IG}) + a_j \\
\Theta_j^U &= \frac{a_j + b_j \exp(\theta_j^U)}{1 + \exp(\theta_j^U)}
\end{aligned}$$

[1] This description uses the same notation and a similar approach to the one described in Warne (2008)

These transformations must be taken into account when evaluating the natural logarithm of the prior densities in (2), by adding the Jacobians of the transformations of the variables:

$$\ln\left(\frac{d\Theta_j^N}{d\theta_j^N}\right) = 0$$

$$\ln\left(\frac{d\Theta_j^{IG}}{d\theta_j^{IG}}\right) = \theta_j^{IG}$$

$$\ln\left(\frac{d\Theta_j^U}{d\theta_j^U}\right) = \ln(b_j - a_j) + \theta_j^U - 2\ln(1 + \exp(\theta_j^U))$$

Code The code is written in Julia and it is available on GitHub.

Appendix B Posteriors of all parameters

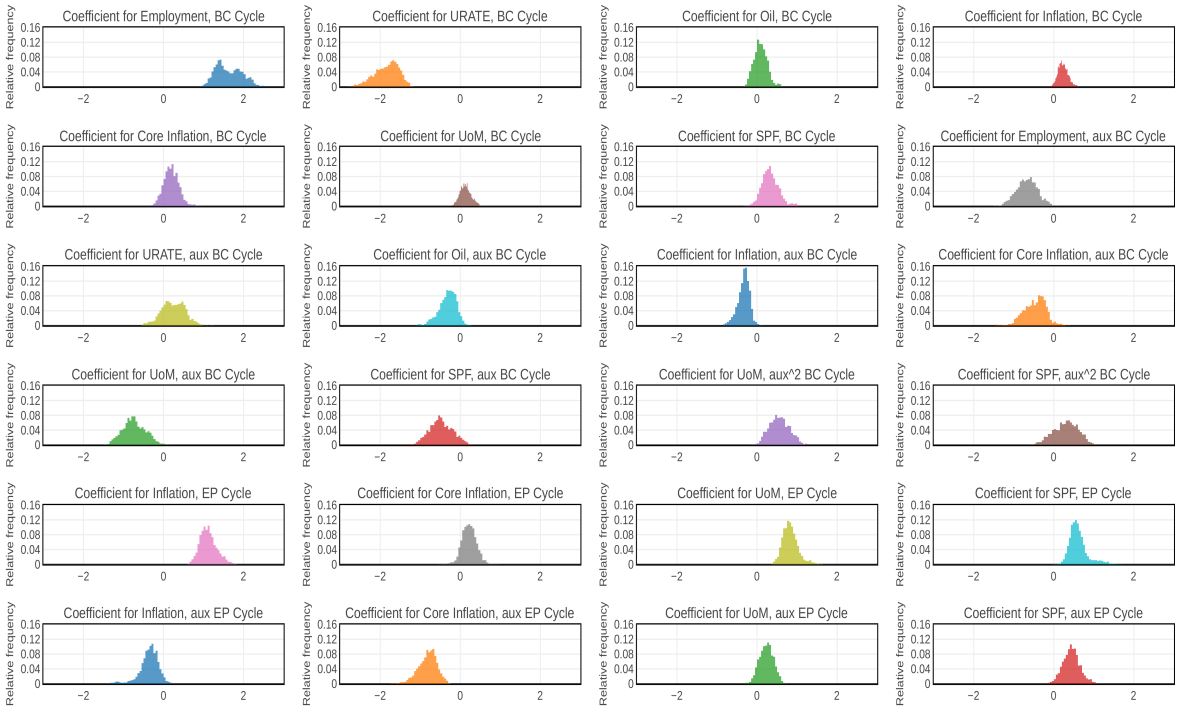

Figure 1: Coefficients

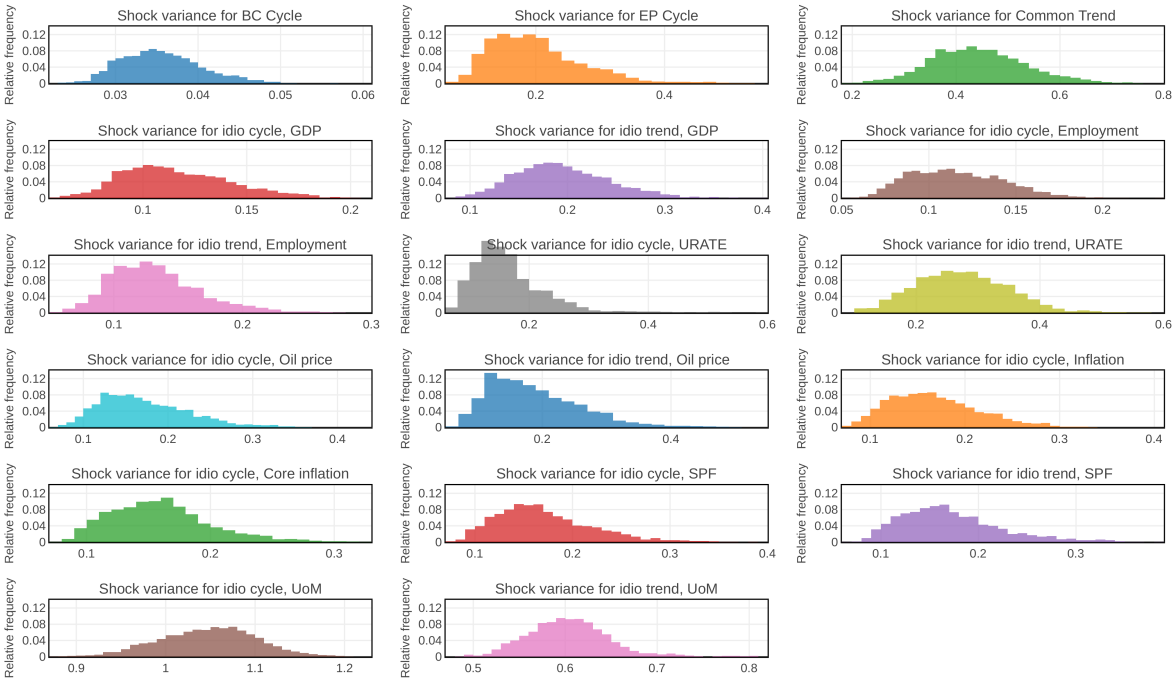

Figure 2: Variance of shocks

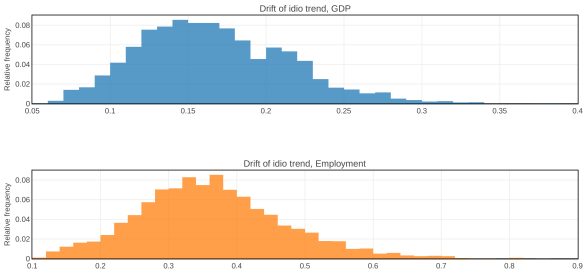

Figure 3: Drifts

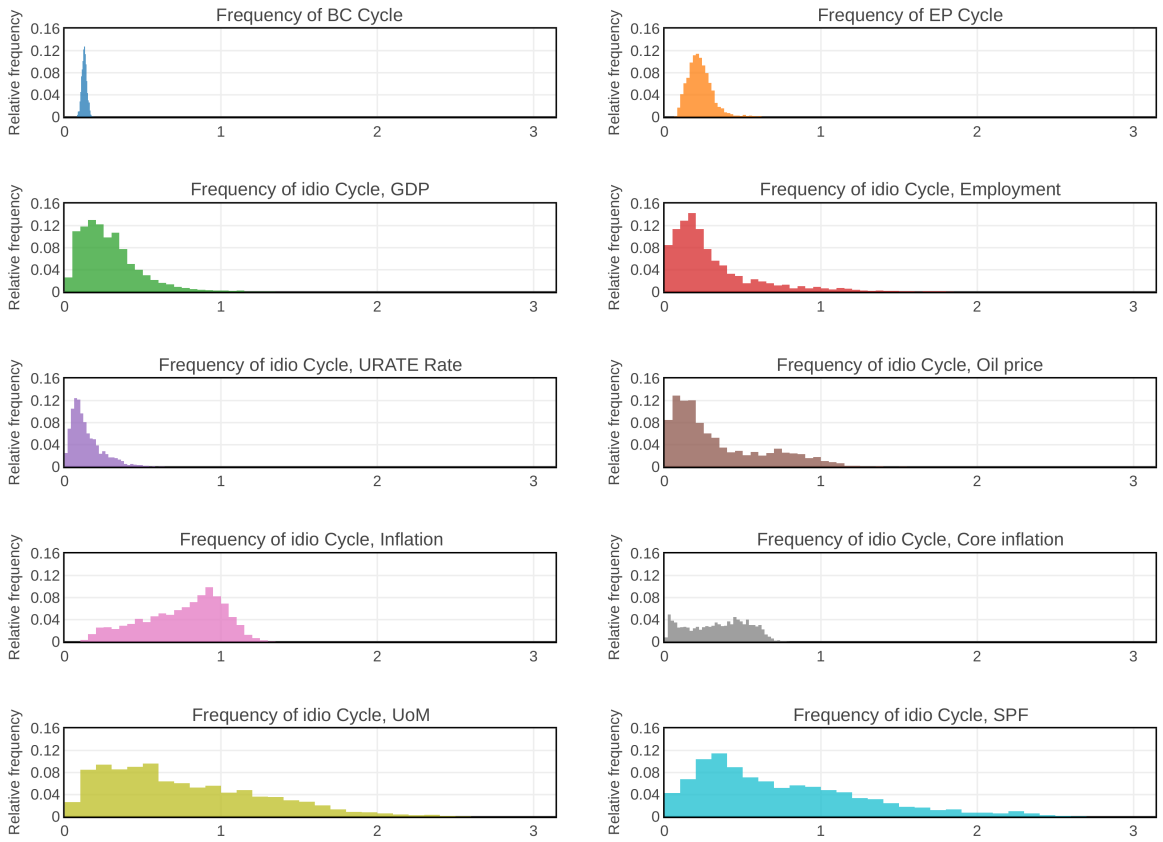

Figure 4: Frequency of cycles

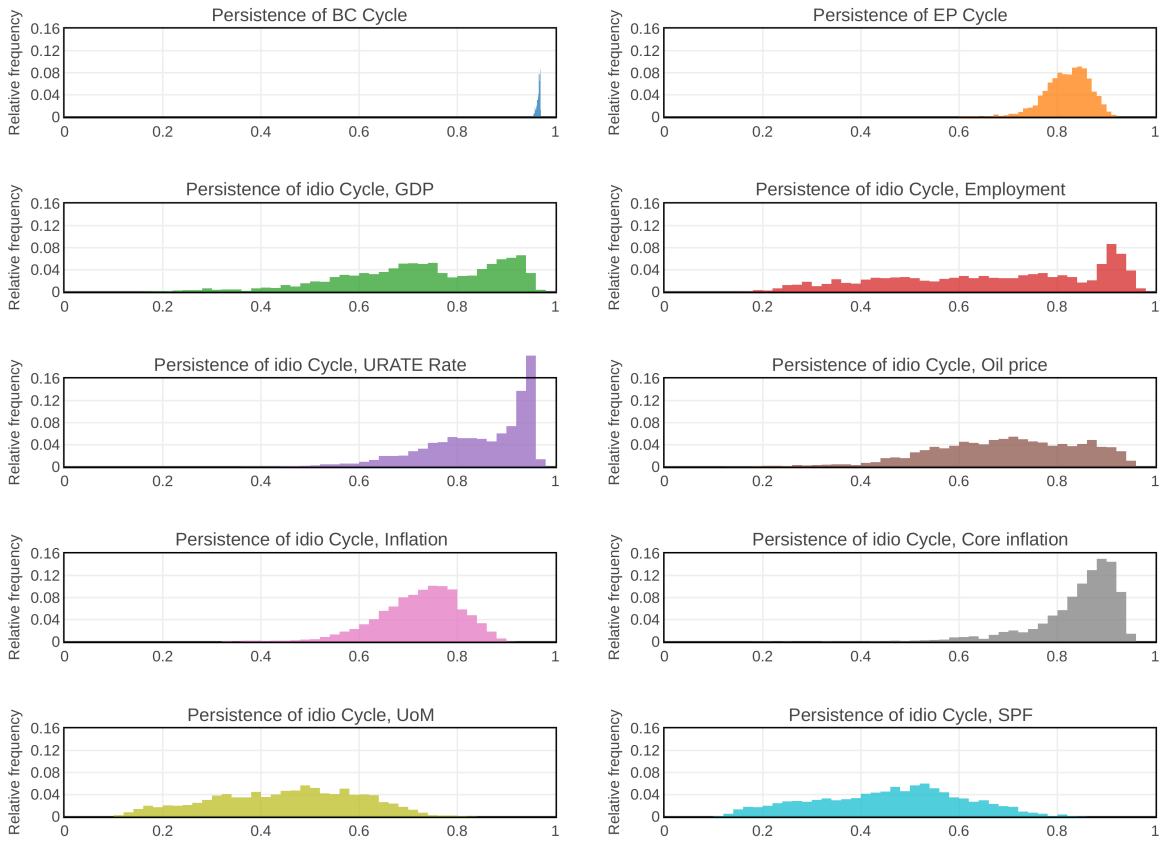

Figure 5: Persistence of cycles

Appendix C Robustness to Priors

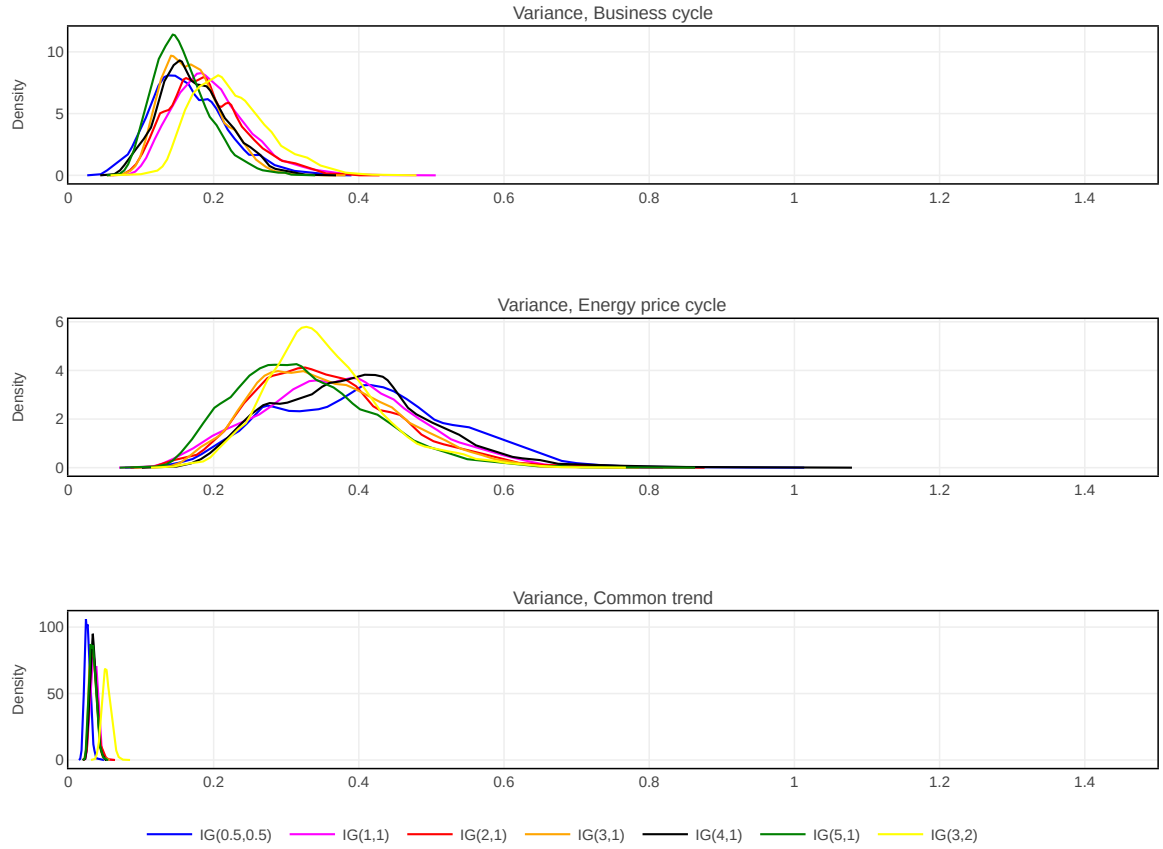

Figure 6: The chart reports the posterior distributions of the variance of the shocks to the business cycle (top), energy price cycle (middle), and common trend (bottom) when the inverse gamma priors of those variances have different shape and scale parameters.

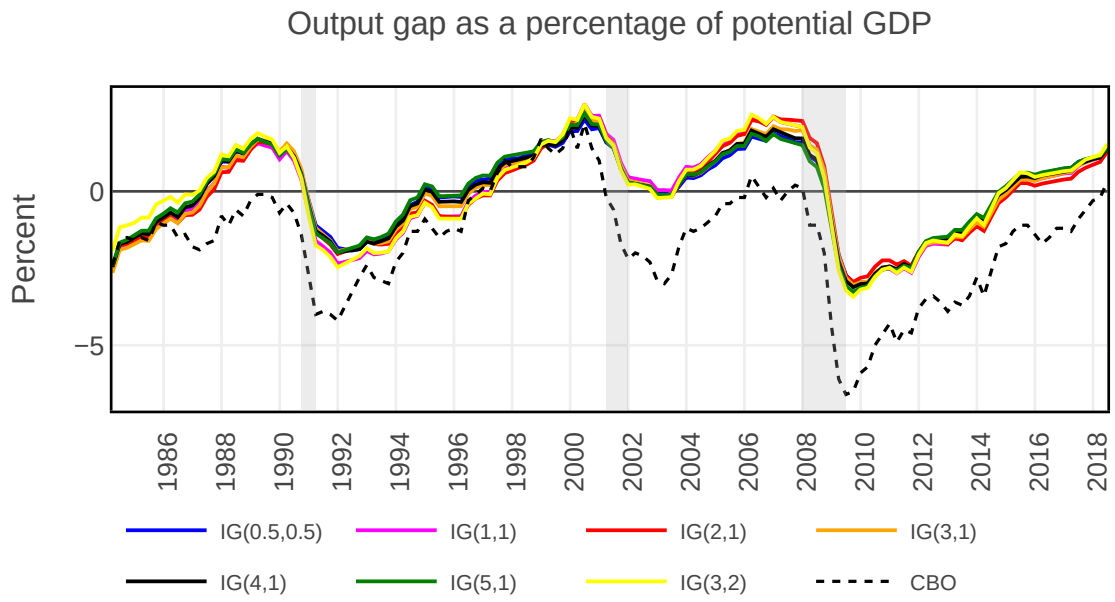

Figure 7: The chart reports the median of the output gap when the priors for the variance of the shocks to the common components are inverse gamma priors with different shape and scale parameters. The chart also reports the output gap from the CBO.

Appendix D Building up the Model

D.1 Model A: small-scale model with AR(1) cycles and without dynamic heterogeneity

For this model we used the data reported in the table below.

Table 1: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Unemployment rate	u_t	u	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price.

The observation equation for this model is:

$$\begin{pmatrix} y_t \\ u_t \\ oil_t \\ \pi_t \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \delta_{u,1} & 0 \\ \delta_{oil,1} & 0 \\ \delta_{\pi,1} & \phi_\pi \\ \delta_{spf,1}+\delta_{spf,2}L & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \mu_t^\pi \end{pmatrix} + \begin{pmatrix} \psi_t^y \\ \psi_t^u \\ \psi_t^{oil} \\ \psi_t^\pi \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^y \\ \mu_t^u \\ \mu_t^{oil} \\ 0 \\ \mu_t^{spf} \end{pmatrix} \tag{4}$$

where ϕ_π and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[2] The cycles are modelled as stationary AR(1) with a U(-0.97,

[2]In the empirical model, the series are standardised so that the standard deviations of their first differences are equal to one. For this reason, we normalise ϕ_π and ϕ_{spf} to the reciprocal of the standard deviation of the first difference of the respective variable.

0.97) prior on the autoregressive coefficients.

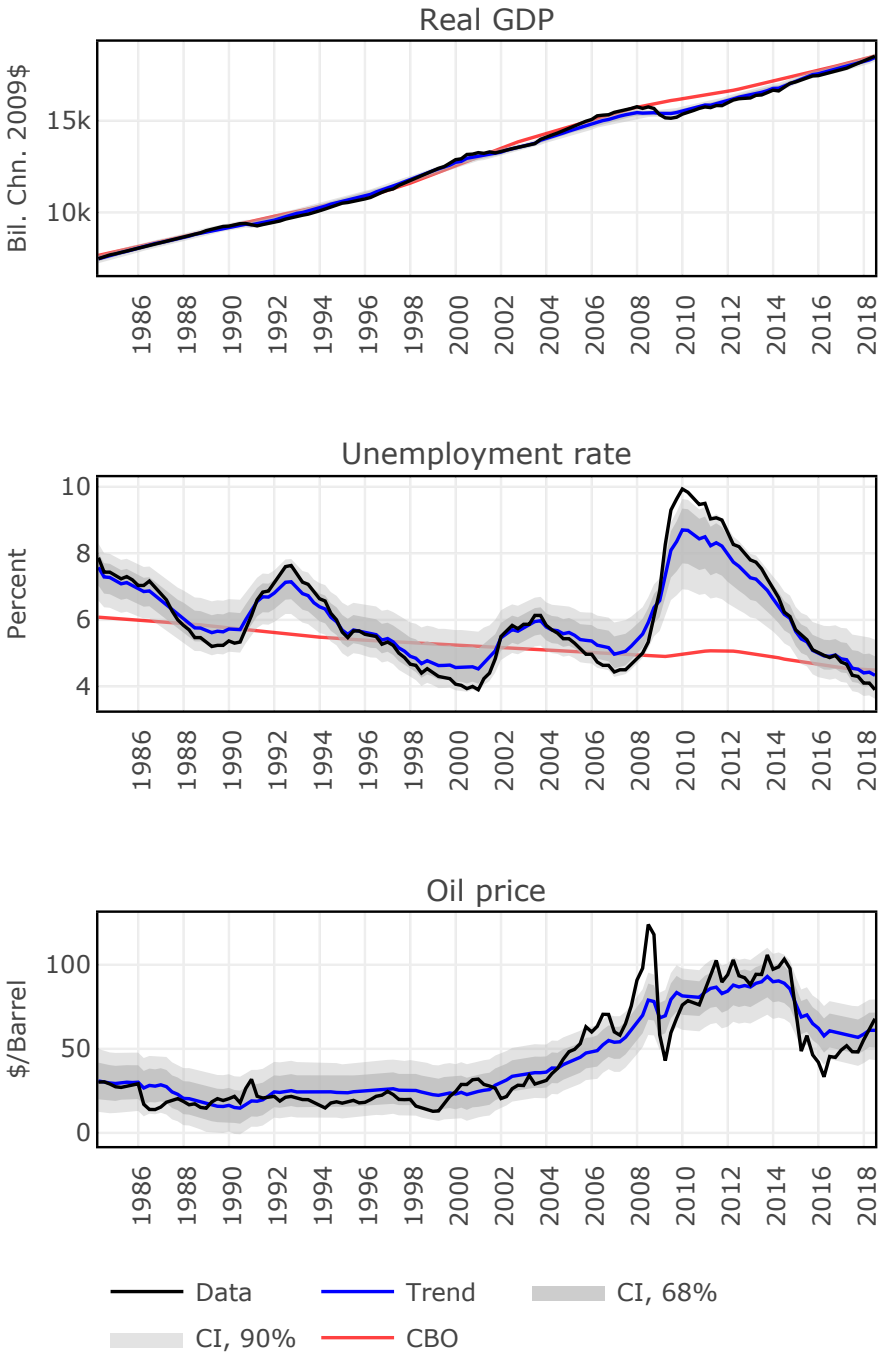

Figure 8: Independent trends of output, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

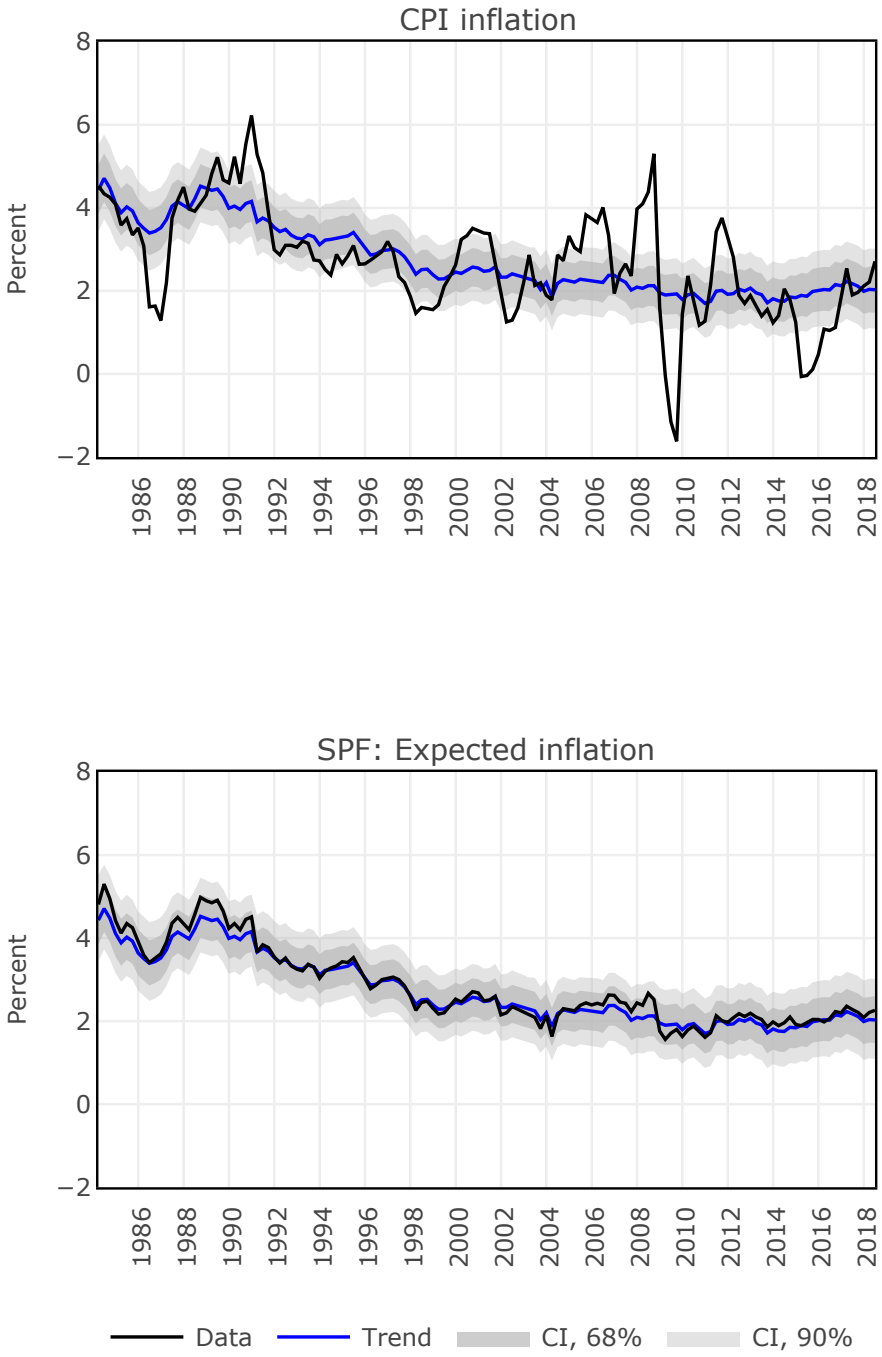

Figure 9: Trend common to CPI inflation and SPF inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

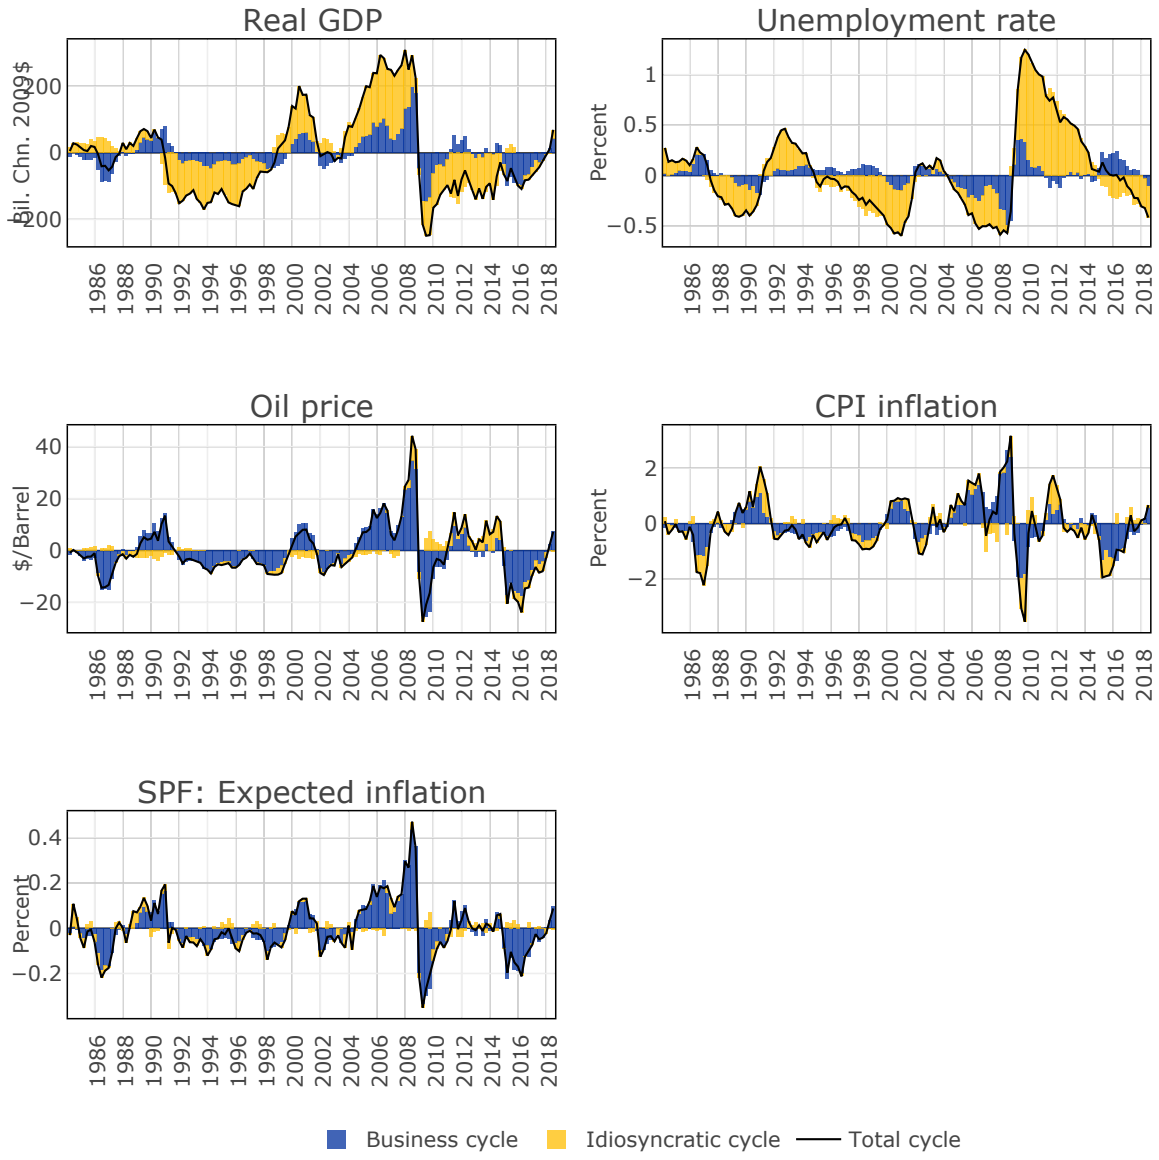

Figure 10: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), and idiosyncratic cycle (in yellow).

D.2 Model B: Model A with ARMA(2,1) cycles

Same data and observation equation used for the model A. Cycles are modelled as in the baseline trend-cycle model in the main text - i.e. as ARMA(2,1).

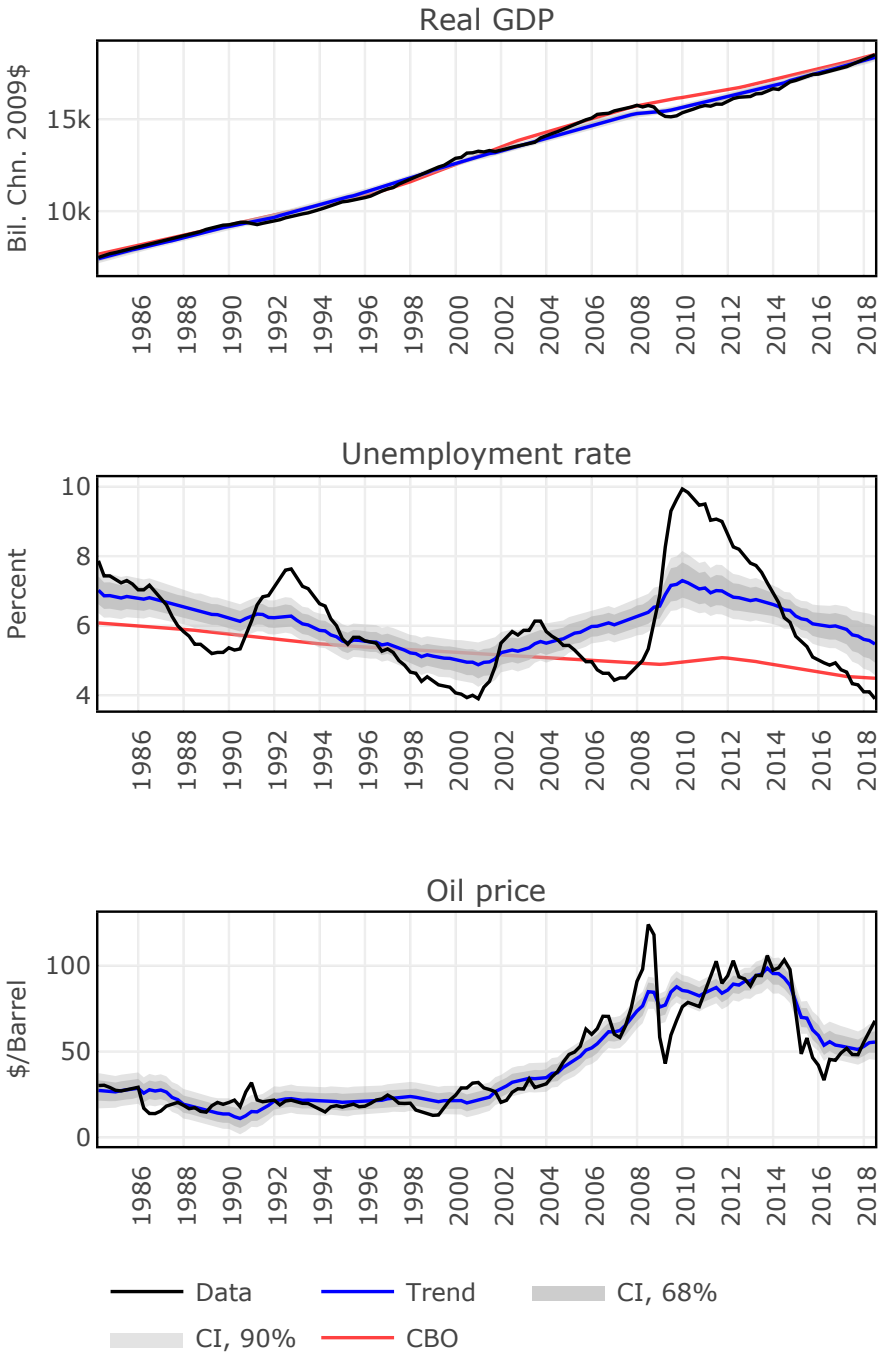

Figure 11: Independent trends of output, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

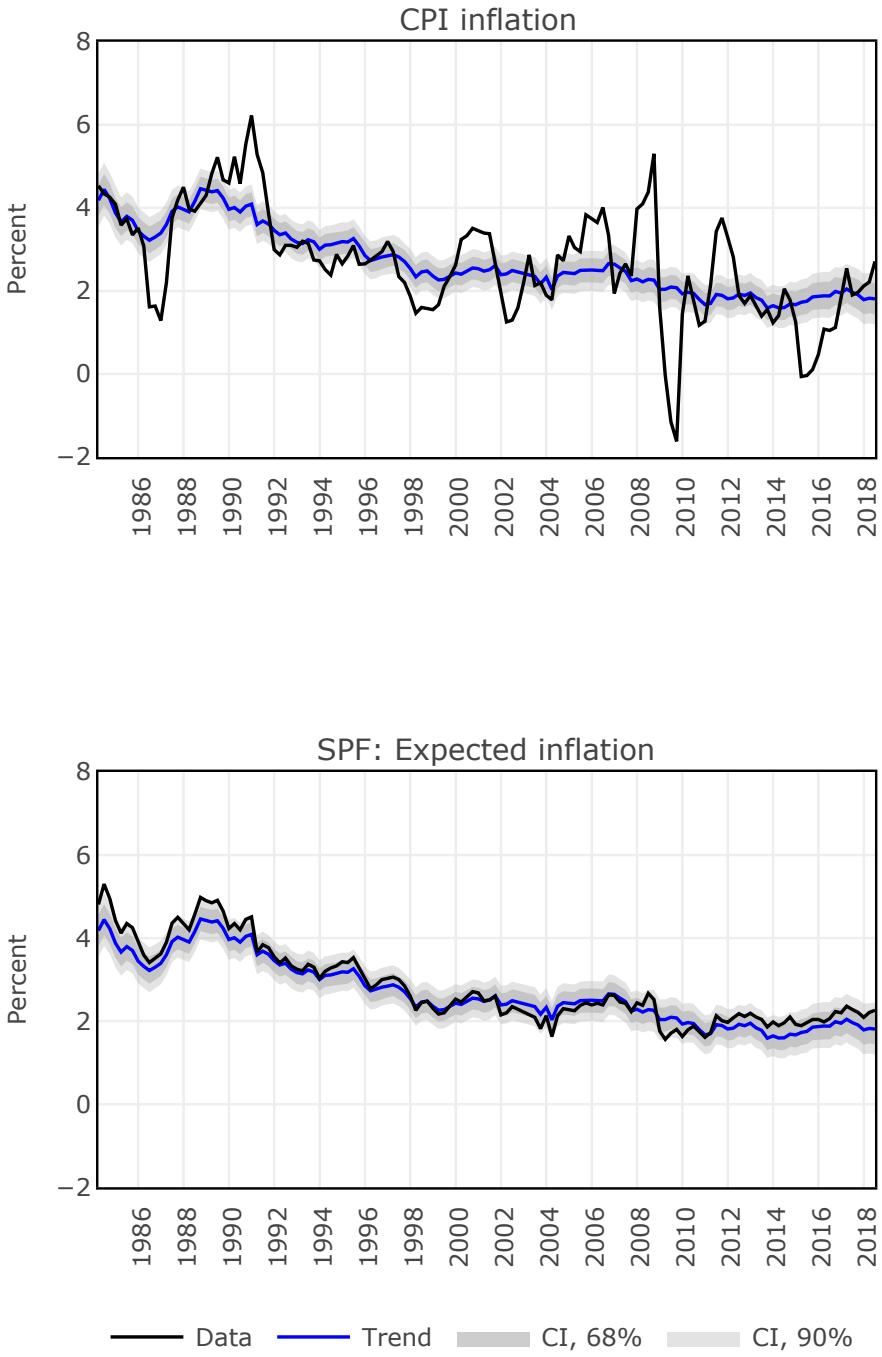

Figure 12: Trend common to CPI inflation and SPF inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

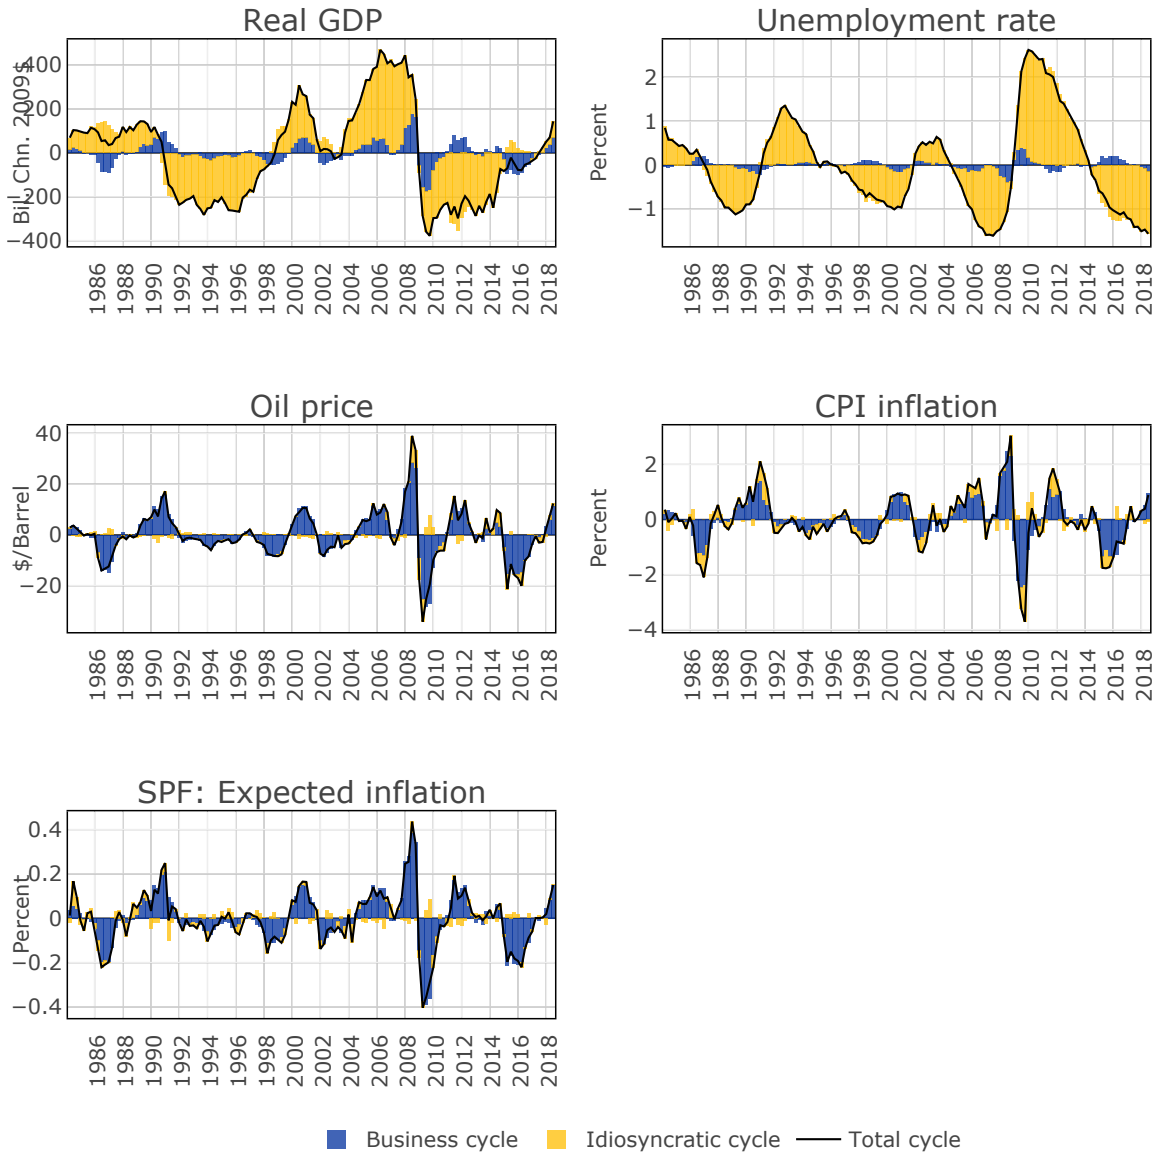

Figure 13: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), and idiosyncratic cycle (in yellow).

D.3 Model C: Model B with EP cycle

For this model we used the data reported in Table 1. The observation equation for this model is:

$$\begin{pmatrix} y_t \\ u_t \\ oil_t \\ \pi_t \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ \delta_{u,1} & 0 & 0 \\ \delta_{oil,1} & 1 & 0 \\ \delta_{\pi,1} & \gamma_{\pi,1} & \phi_\pi \\ \delta_{spf,1}+\delta_{spf,2}L & \gamma_{spf,1} & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \psi_t^{EP} \\ \mu_t^{\pi} \end{pmatrix} + \begin{pmatrix} \psi_t^{y} \\ \psi_t^{u} \\ \psi_t^{oil} \\ \psi_t^{\pi} \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^{y} \\ \mu_t^{u} \\ \mu_t^{oil} \\ 0 \\ \mu_t^{spf} \end{pmatrix} \tag{5}$$

where ϕ_π and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[3]

[3] In the empirical model, the series are standardised so that the standard deviations of their first differences are equal to one. For this reason, we normalise ϕ_π and ϕ_{spf} to the reciprocal of the standard deviation of the first difference of the respective variable.

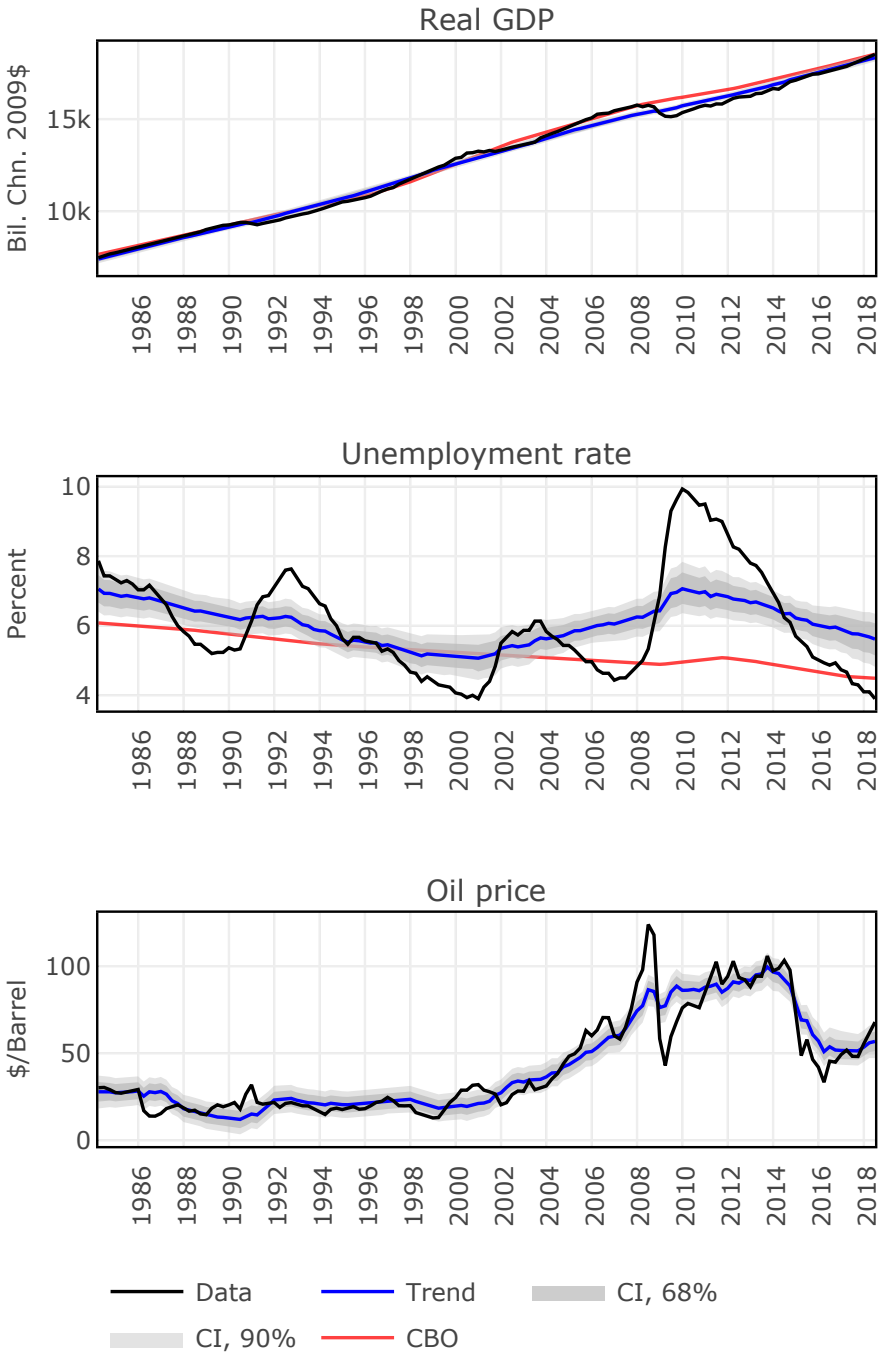

Figure 14: Independent trends of output, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

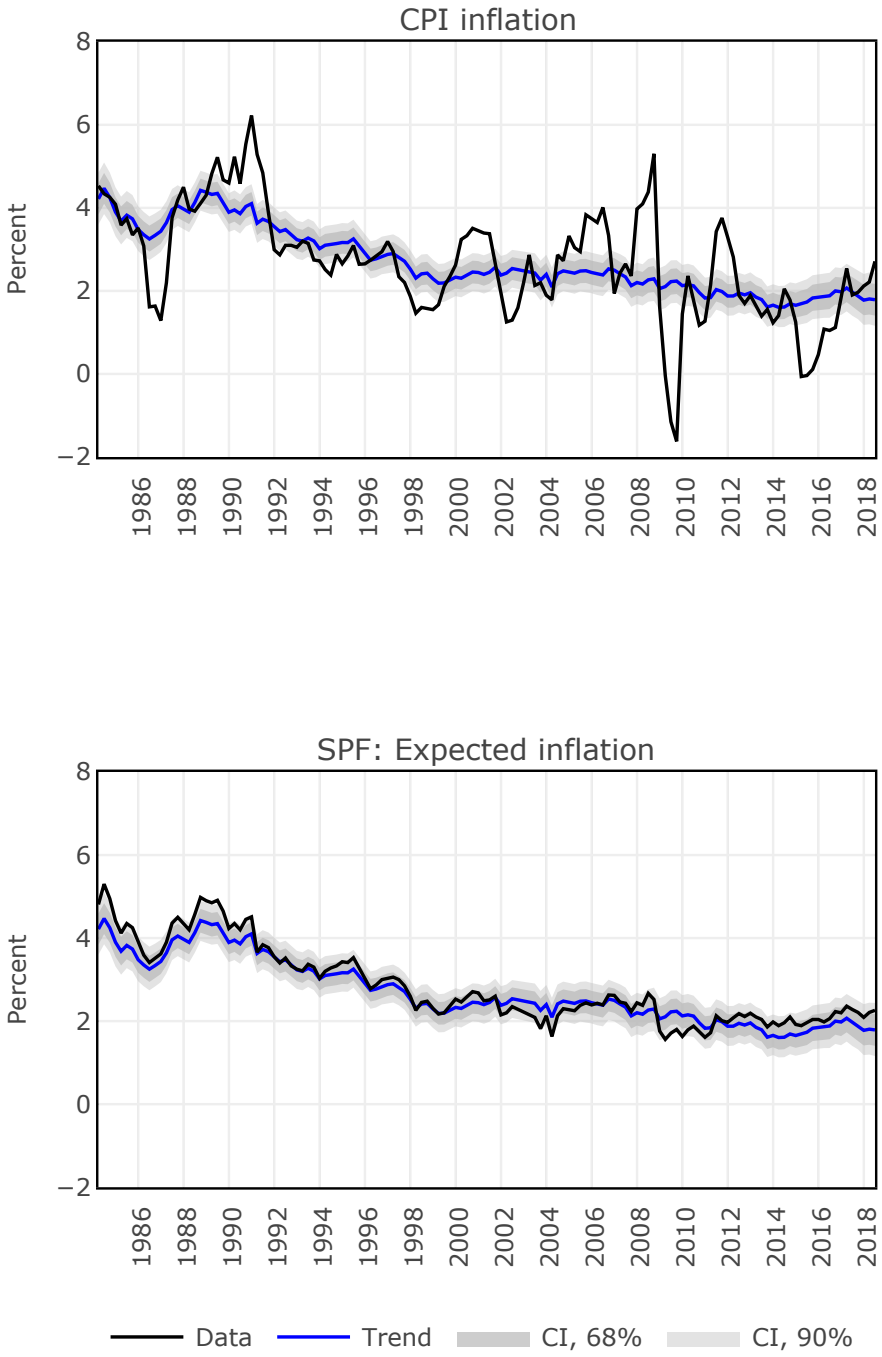

Figure 15: Trend common to CPI inflation and SPF inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

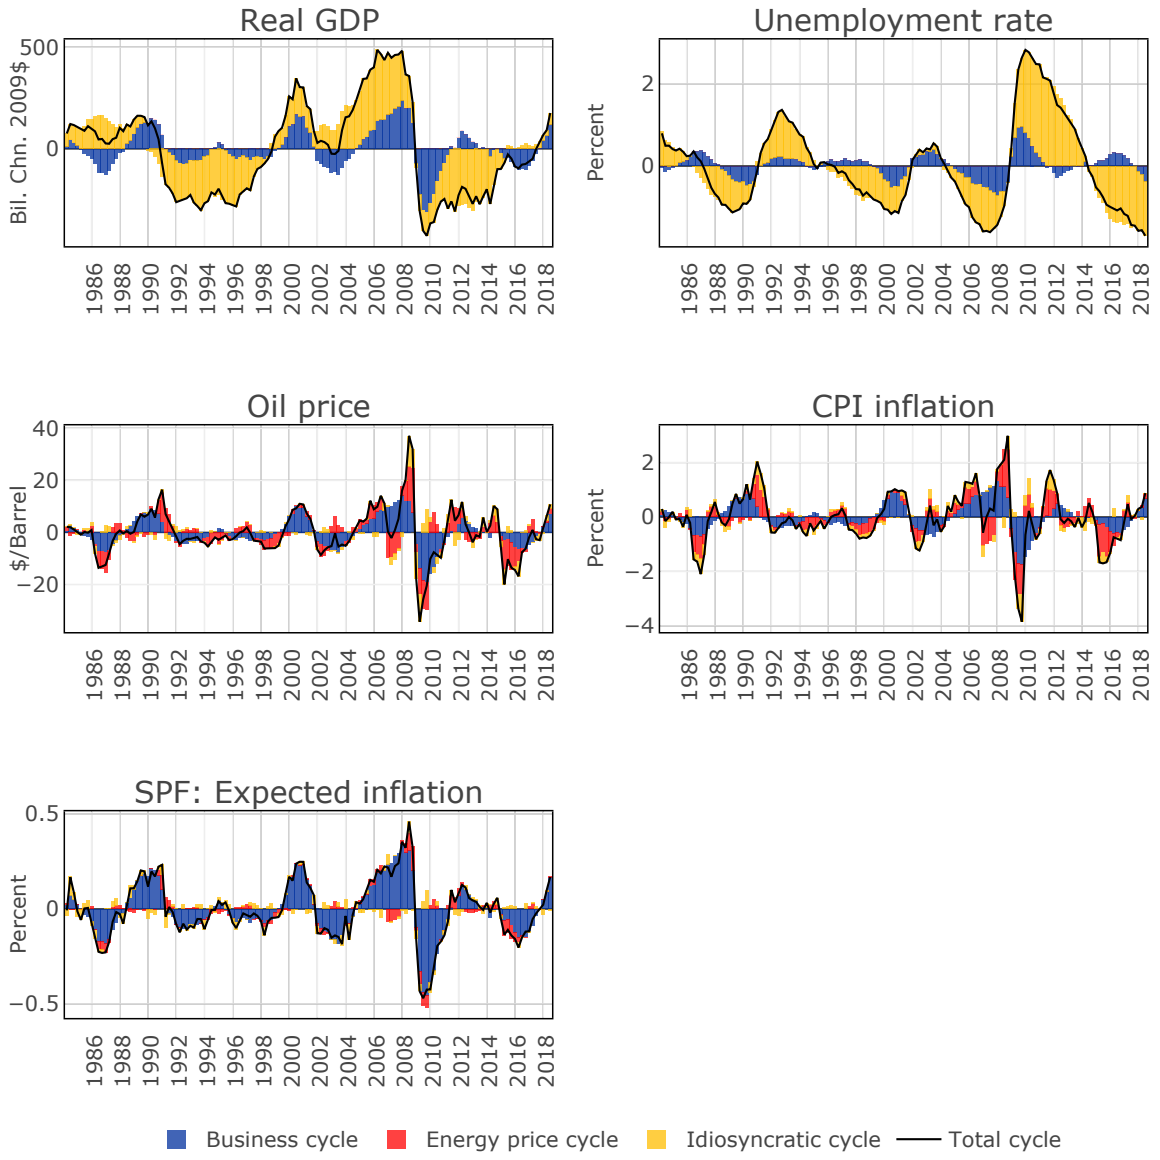

Figure 16: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

D.4 Model D: Model C with employment and core inflation

For this model we used the data reported in the table below.

Table 2: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Employment	e_t	e	Levels
Unemployment rate	u_t	u	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
Core CPI inflation	π_t^c	π^c	YoY
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price.

The observation equation for this model is:

$$\begin{pmatrix} y_t \\ e_t \\ u_t \\ oil_t \\ \pi_t \\ \pi_t^c \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ \delta_{e,1} & 0 & 0 \\ \delta_{u,1} & 0 & 0 \\ \delta_{oil,1} & 1 & 0 \\ \delta_{\pi,1} & \gamma_{\pi,1} & \phi_\pi \\ \delta_{\pi^c,1} & \gamma_{\pi^c,1} & \phi_{\pi^c} \\ \delta_{spf,1}+\delta_{spf,2}L & \gamma_{spf,1} & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \psi_t^{EP} \\ \mu_t^\pi \end{pmatrix} + \begin{pmatrix} \psi_t^y \\ \psi_t^e \\ \psi_t^u \\ \psi_t^{oil} \\ \psi_t^\pi \\ \psi_t^{\pi^c} \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^y \\ \mu_t^e \\ \mu_t^u \\ \mu_t^{oil} \\ 0 \\ 0 \\ \mu_t^{spf} \end{pmatrix} \tag{6}$$

where ϕ_π, ϕ_{π^c} and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[4]

[4]In the empirical model, the series are standardised so that the standard deviations of their first differences are equal to one. For this reason, we normalise ϕ_π, ϕ_{π^c} and ϕ_{spf} to the reciprocal of the standard deviation of the first difference of the respective variable.

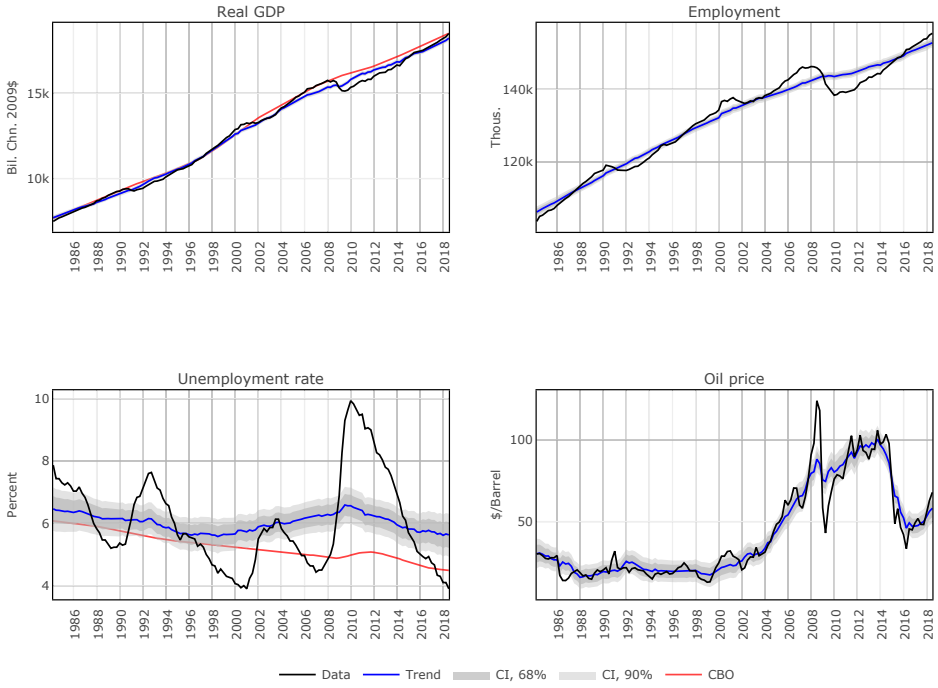

Figure 17: Independent trends of output, employment, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

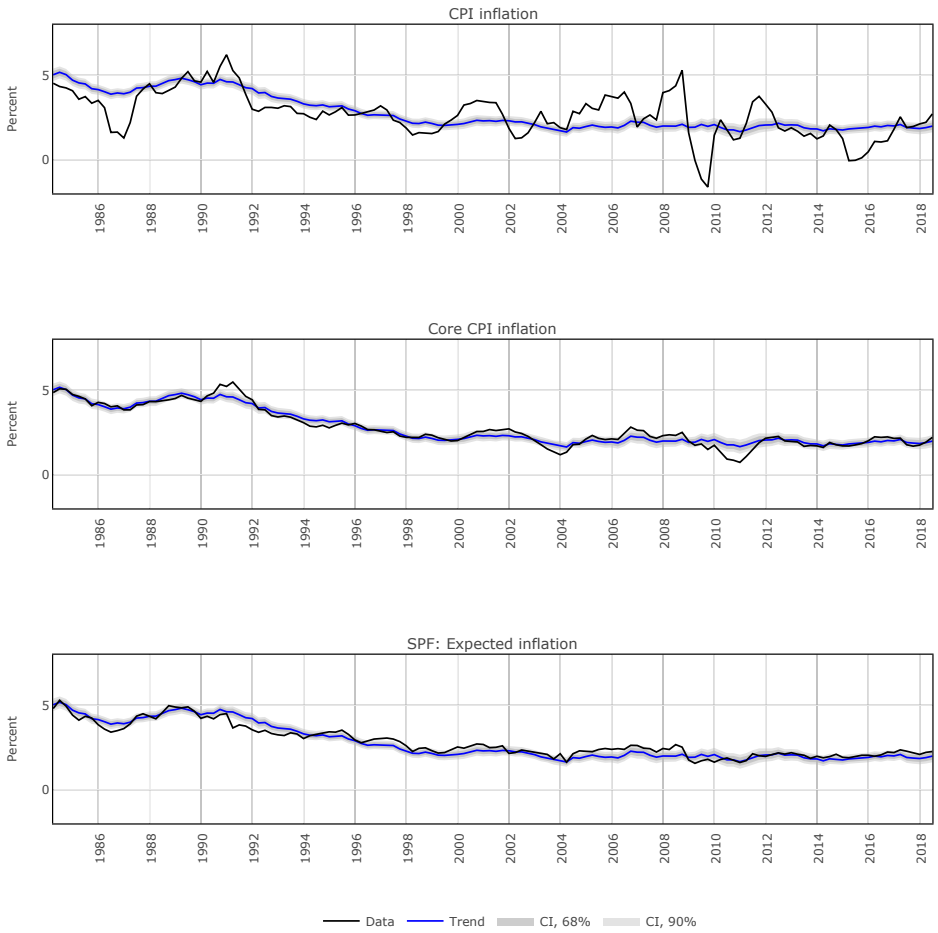

Figure 18: Trend common to CPI inflation, core CPI inflation, and inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

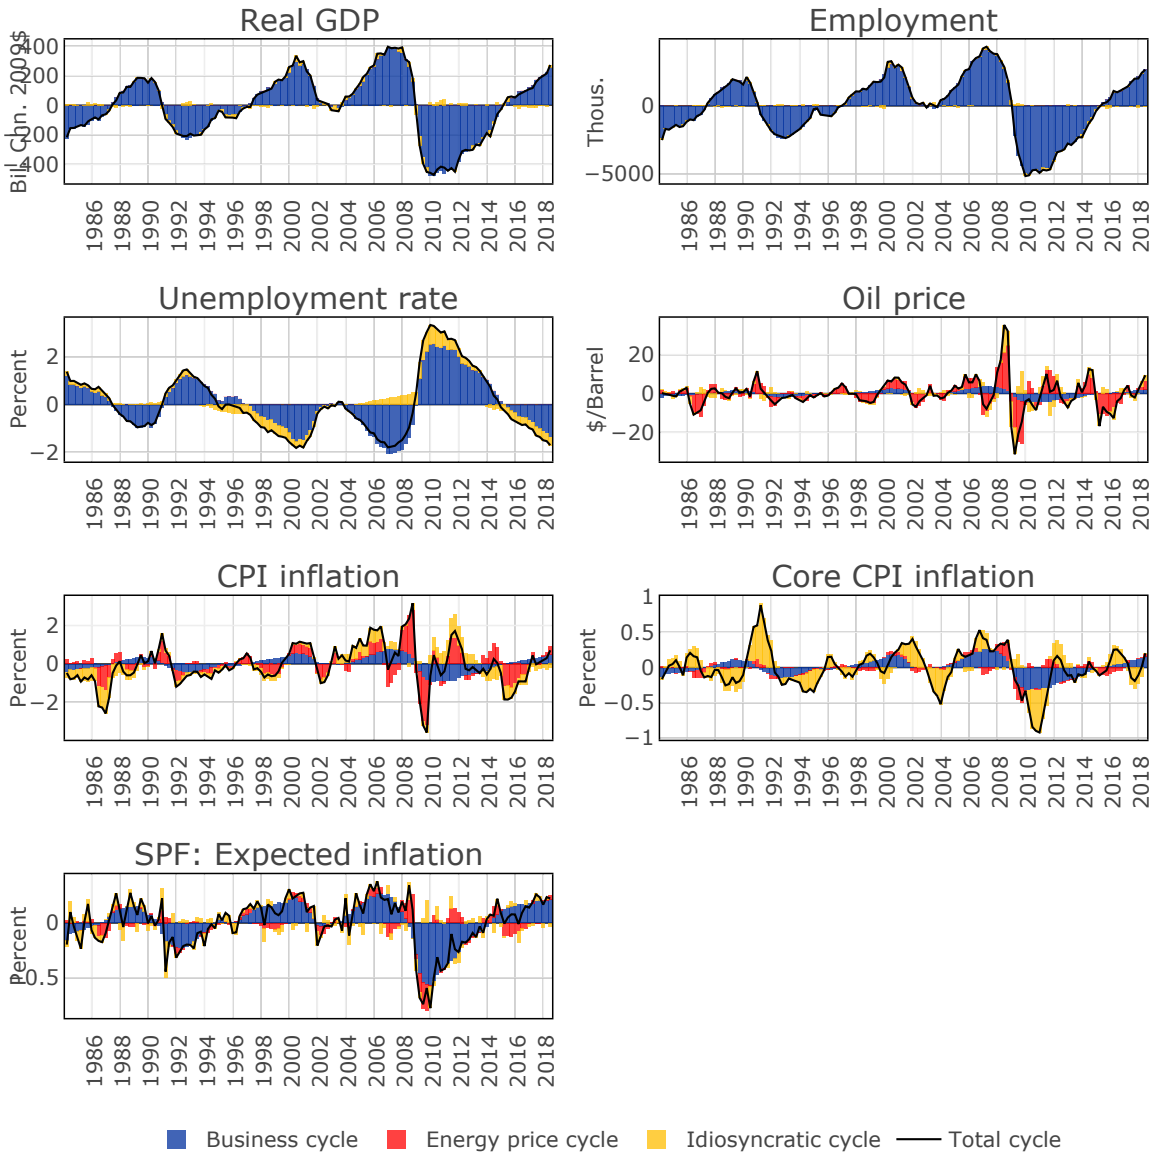

Figure 19: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

D.5 Model E: Model D with University of Michigan inflation expectations

For this model we used the data reported in the table below.

Table 3: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Employment	e_t	e	Levels
Unemployment rate	u_t	u	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
Core CPI inflation	π_t^c	π^c	YoY
UoM: Expected inflation	$F_t^{uom}\pi_{t+4}$	uom	Levels
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'UoM: Expected inflation' is the University of Michigan, 12-months ahead expected inflation rate. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price.

The observation equation for this model is:

$$\begin{pmatrix} y_t \\ e_t \\ u_t \\ oil_t \\ \pi_t \\ \pi_t^c \\ F_t^{uom}\pi_{t+4} \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ \delta_{e,1} & 0 & 0 \\ \delta_{u,1} & 0 & 0 \\ \delta_{oil,1} & 1 & 0 \\ \delta_{\pi,1} & \gamma_{\pi,1} & \phi_\pi \\ \delta_{\pi^c,1} & \gamma_{\pi^c,1} & \phi_{\pi^c} \\ \delta_{uom,1} + \delta_{uom,2}L & \gamma_{uom,1} & \phi_{uom} \\ \delta_{spf,1} + \delta_{spf,2}L & \gamma_{spf,1} & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \psi_t^{EP} \\ \mu_t^\pi \end{pmatrix} + \begin{pmatrix} \psi_t^y \\ \psi_t^e \\ \psi_t^u \\ \psi_t^{oil} \\ \psi_t^\pi \\ \psi_t^{\pi^c} \\ \psi_t^{uom} \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^y \\ \mu_t^e \\ \mu_t^u \\ \mu_t^{oil} \\ 0 \\ 0 \\ \mu_t^{uom} \\ \mu_t^{spf} \end{pmatrix} \tag{7}$$

where ϕ_π, ϕ_{π^c}, ϕ_{uom} and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[5]

[5] In the empirical model, the series are standardised so that the standard deviations of their first

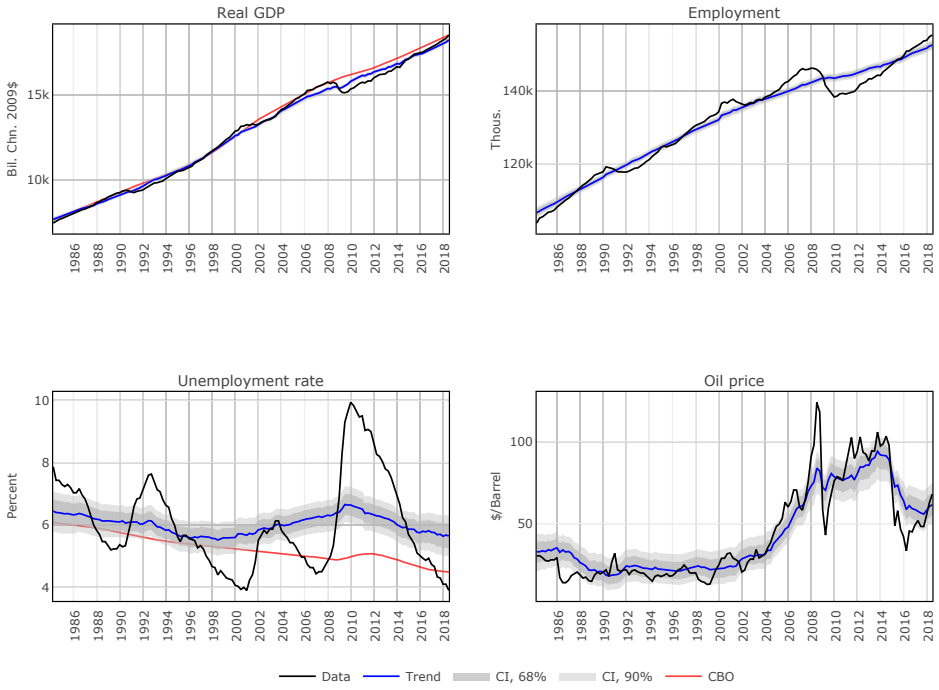

Figure 20: Independent trends of output, employment, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

differences are equal to one. For this reason, we normalise ϕ_π, ϕ_{π^c}, ϕ_{uom} and ϕ_{spf} to the reciprocal of the standard deviation of the first difference of the respective variable.

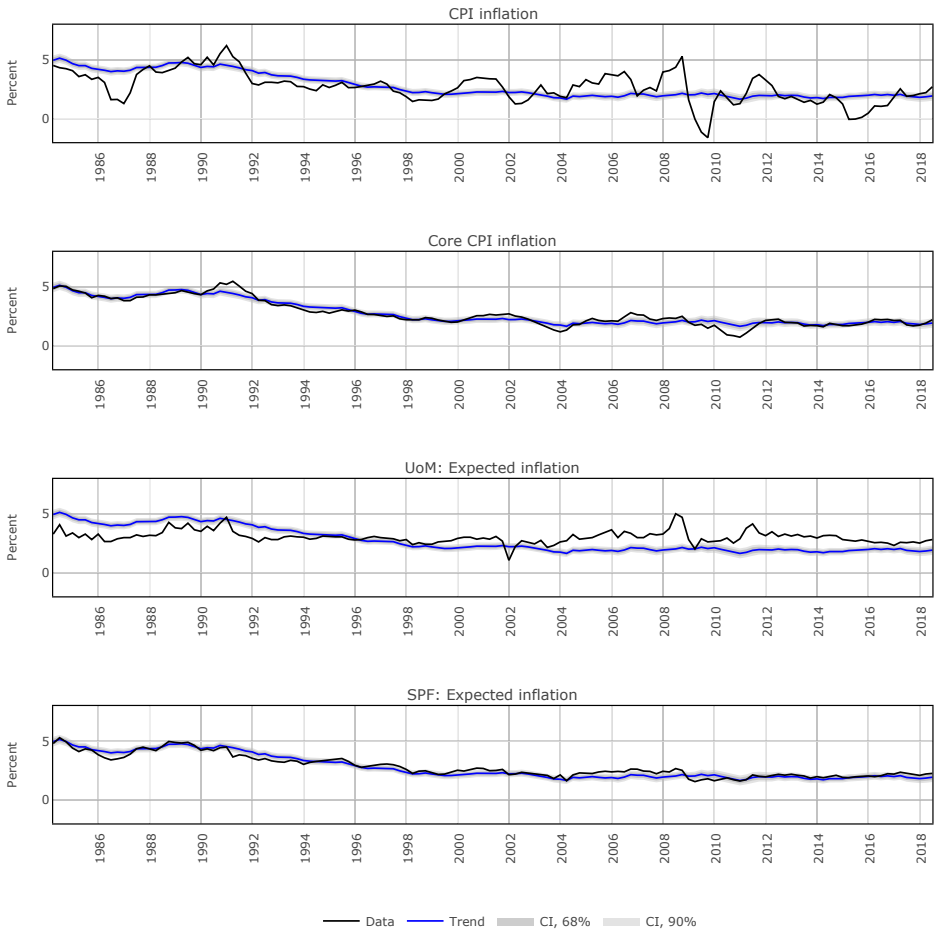

Figure 21: Trend common to CPI inflation, core CPI inflation, and inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

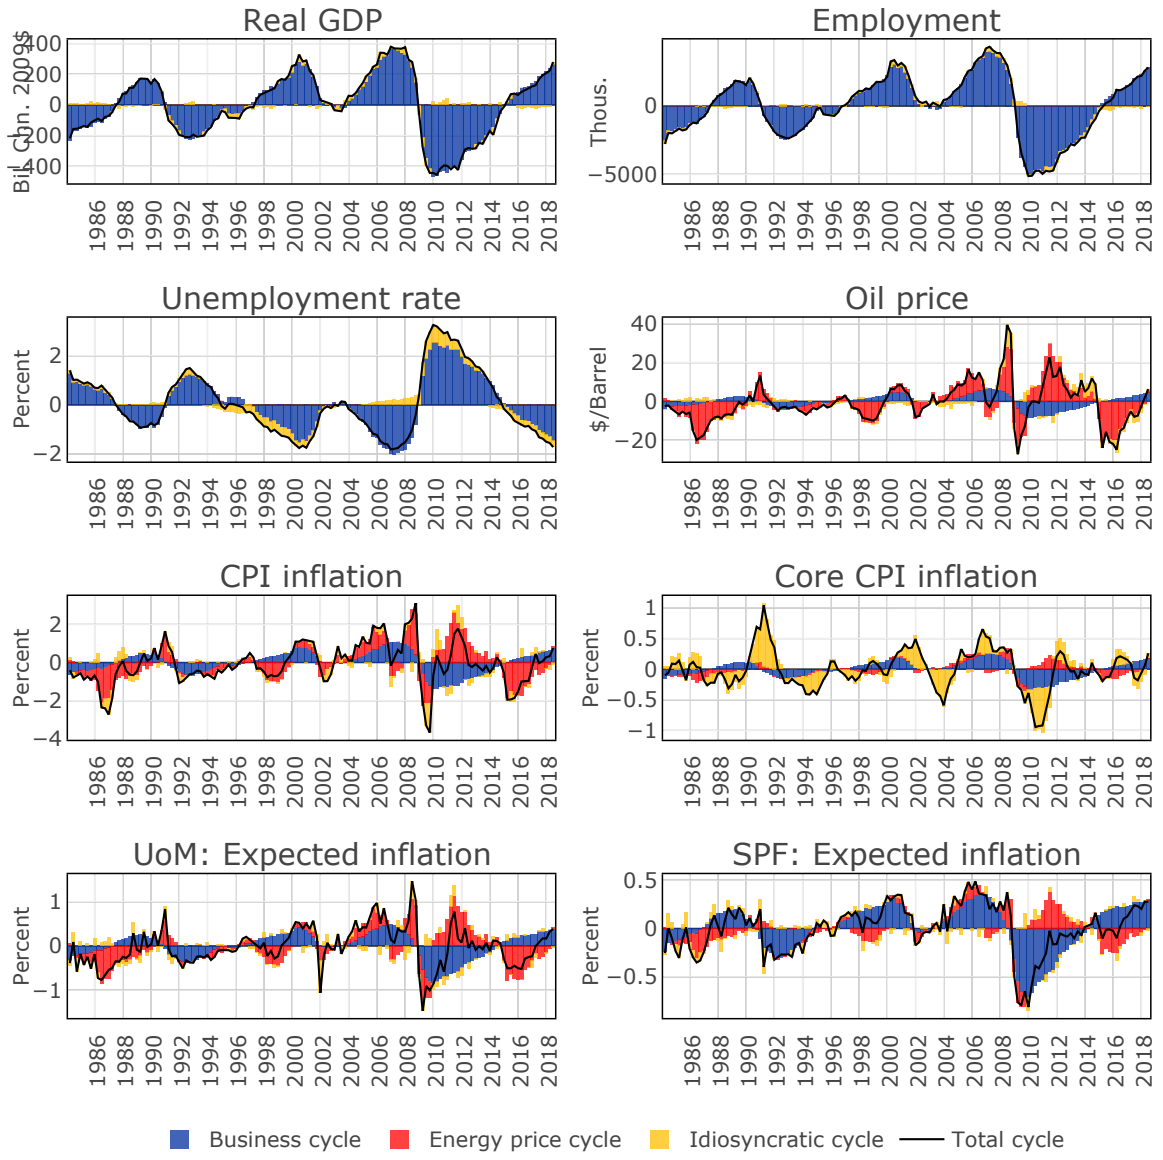

Figure 22: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

D.6 Model F: Benchmark Model

For the benchmark model[6] we used the data reported in the table below.

Table 4: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Employment	e_t	e	Levels
Unemployment rate	u_t	u	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
Core CPI inflation	π_t^c	π^c	YoY
UoM: Expected inflation	$F_t^{uom}\pi_{t+4}$	uom	Levels
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'UoM: Expected inflation' is the University of Michigan, 12-months ahead expected inflation rate. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price.

The observation equation for the benchmark model adopted in the paper is:

$$\begin{pmatrix} y_t \\ e_t \\ u_t \\ oil_t \\ \pi_t \\ \pi_t^c \\ F_t^{uom}\pi_{t+4} \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ \delta_{e,1}+\delta_{e,2}L & 0 & 0 \\ \delta_{u,1}+\delta_{u,2}L & 0 & 0 \\ \delta_{oil,1}+\delta_{oil,2}L & 1 & 0 \\ \delta_{\pi,1}+\delta_{\pi,2}L & \gamma_{\pi,1}+\gamma_{\pi,2}L & \phi_\pi \\ \delta_{\pi^c,1}+\delta_{\pi^c,2}L & \gamma_{\pi^c,1}+\gamma_{\pi^c,2}L & \phi_{\pi^c} \\ \delta_{uom,1}+\delta_{uom,2}L+\delta_{uom,3}L^2 & \gamma_{uom,1}+\gamma_{uom,2}L & \phi_{uom} \\ \delta_{spf,1}+\delta_{spf,2}L+\delta_{spf,3}L^2 & \gamma_{spf,1}+\gamma_{spf,2}L & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \psi_t^{EP} \\ \mu_t^{\pi} \end{pmatrix} + \begin{pmatrix} \psi_t^y \\ \psi_t^e \\ \psi_t^u \\ \psi_t^{oil} \\ \psi_t^{\pi} \\ \psi_t^{\pi^c} \\ \psi_t^{uom} \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^y \\ \mu_t^e \\ \mu_t^u \\ \mu_t^{oil} \\ 0 \\ 0 \\ \mu_t^{uom} \\ \mu_t^{spf} \end{pmatrix} \tag{8}$$

where ϕ_π, ϕ_{π^c}, ϕ_{uom}, and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[7]

[6] This is the baseline model described in the main text.

[7] In the empirical model, the series are standardised so that the standard deviations of their first differences are equal to one. For this reason, we normalise ϕ_π, ϕ_{π^c}, ϕ_{uom} and ϕ_{spf} to the reciprocal of the standard deviation of the first difference of the respective variable.

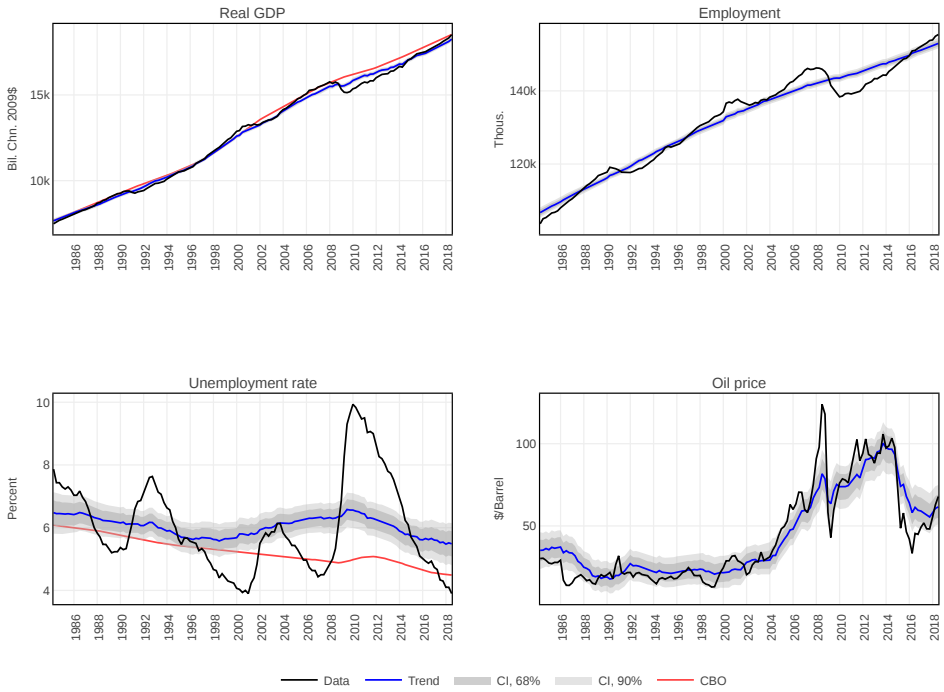

Figure 23: Independent trends of output, employment, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

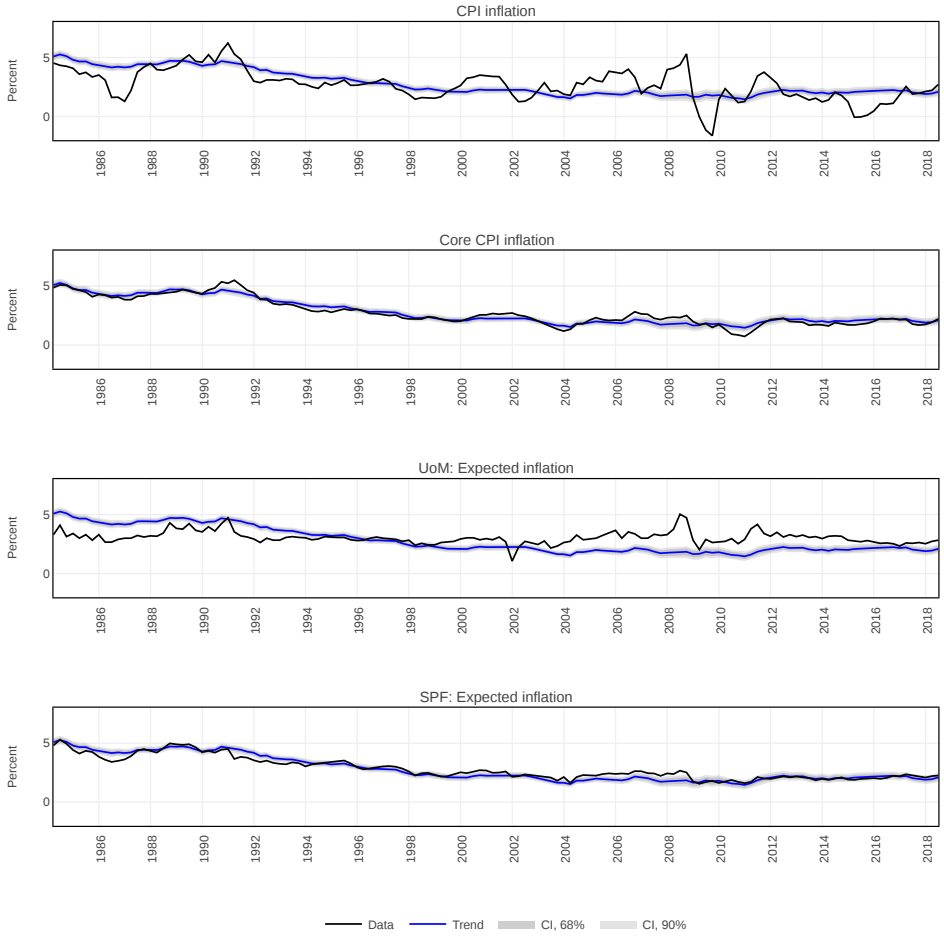

Figure 24: Trend common to CPI inflation, core CPI inflation, and inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

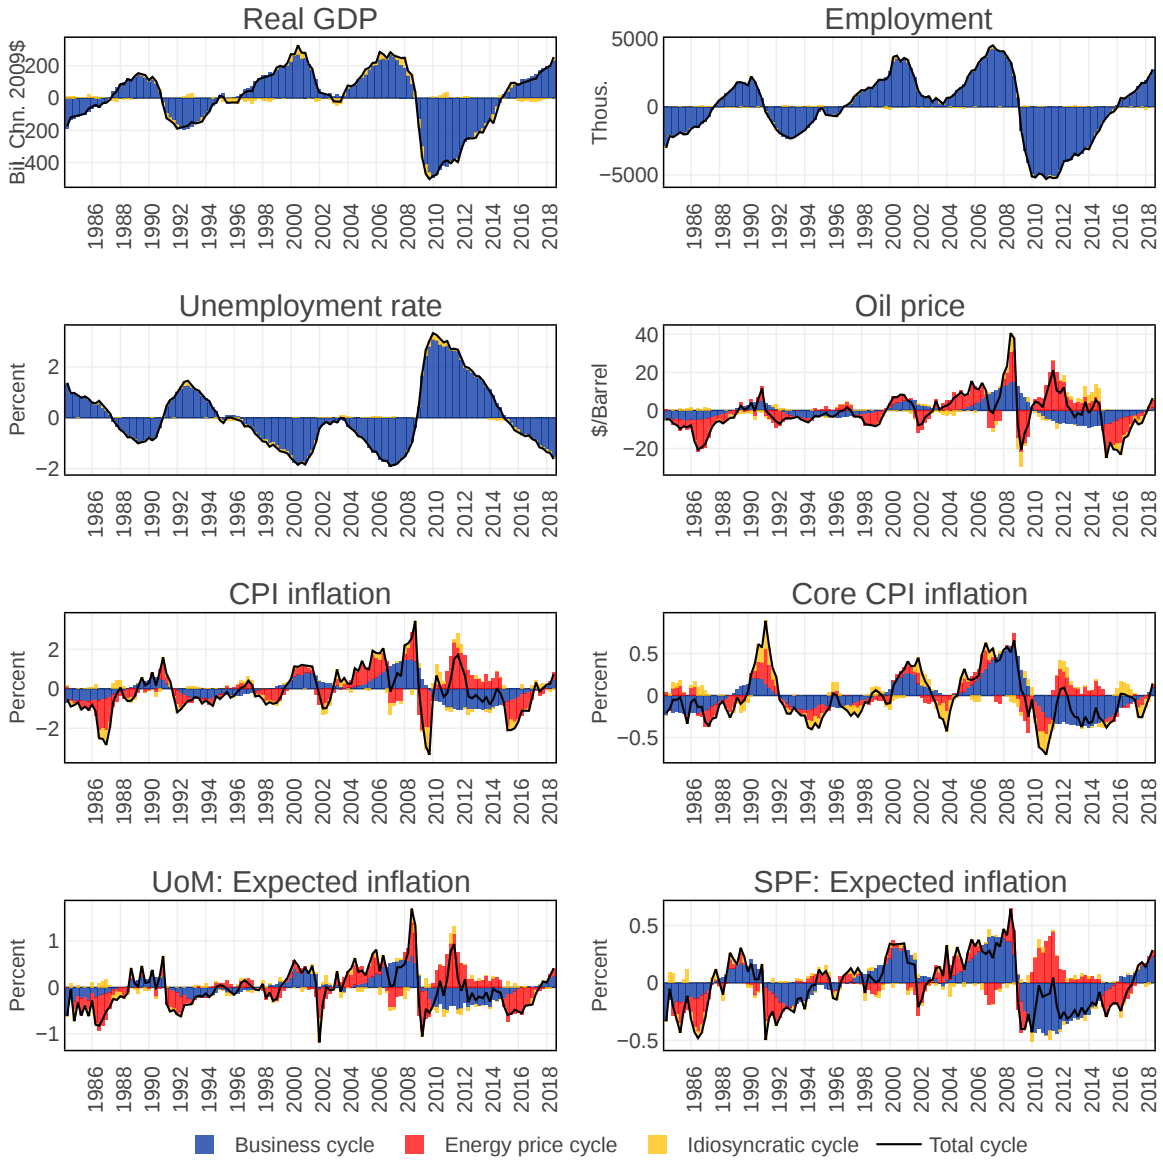

Figure 25: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

Appendix E Employment-to-Population Ratio

The state-space model is identical to the one reported in the main text.

Table 5: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Employment-to-population ratio	e_t	e	Levels
Unemployment rate	u_t	u	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
Core CPI inflation	π_t^c	π^c	YoY
UoM: Expected inflation	$F_t^{uom}\pi_{t+4}$	uom	Levels
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'UoM: Expected inflation' is the University of Michigan, 12-months ahead expected inflation rate. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price.

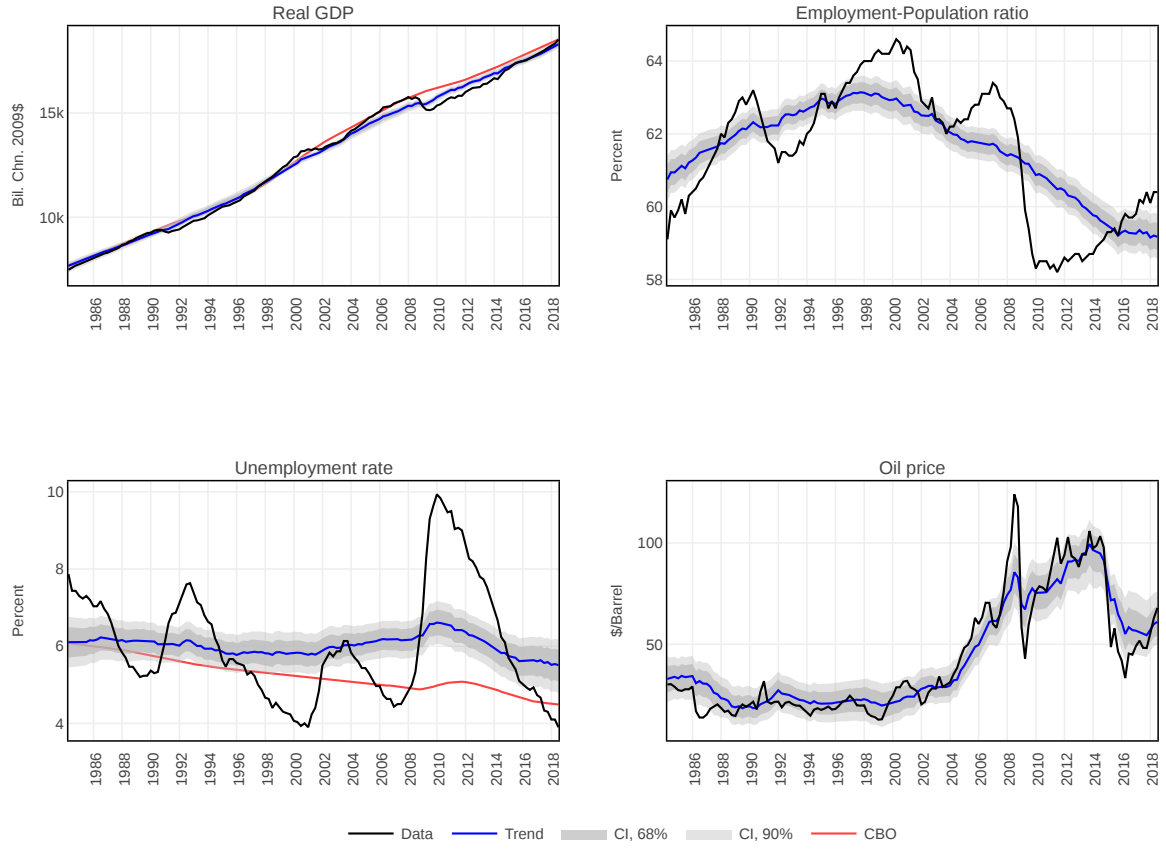

Figure 26: Independent trends of output, employment, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

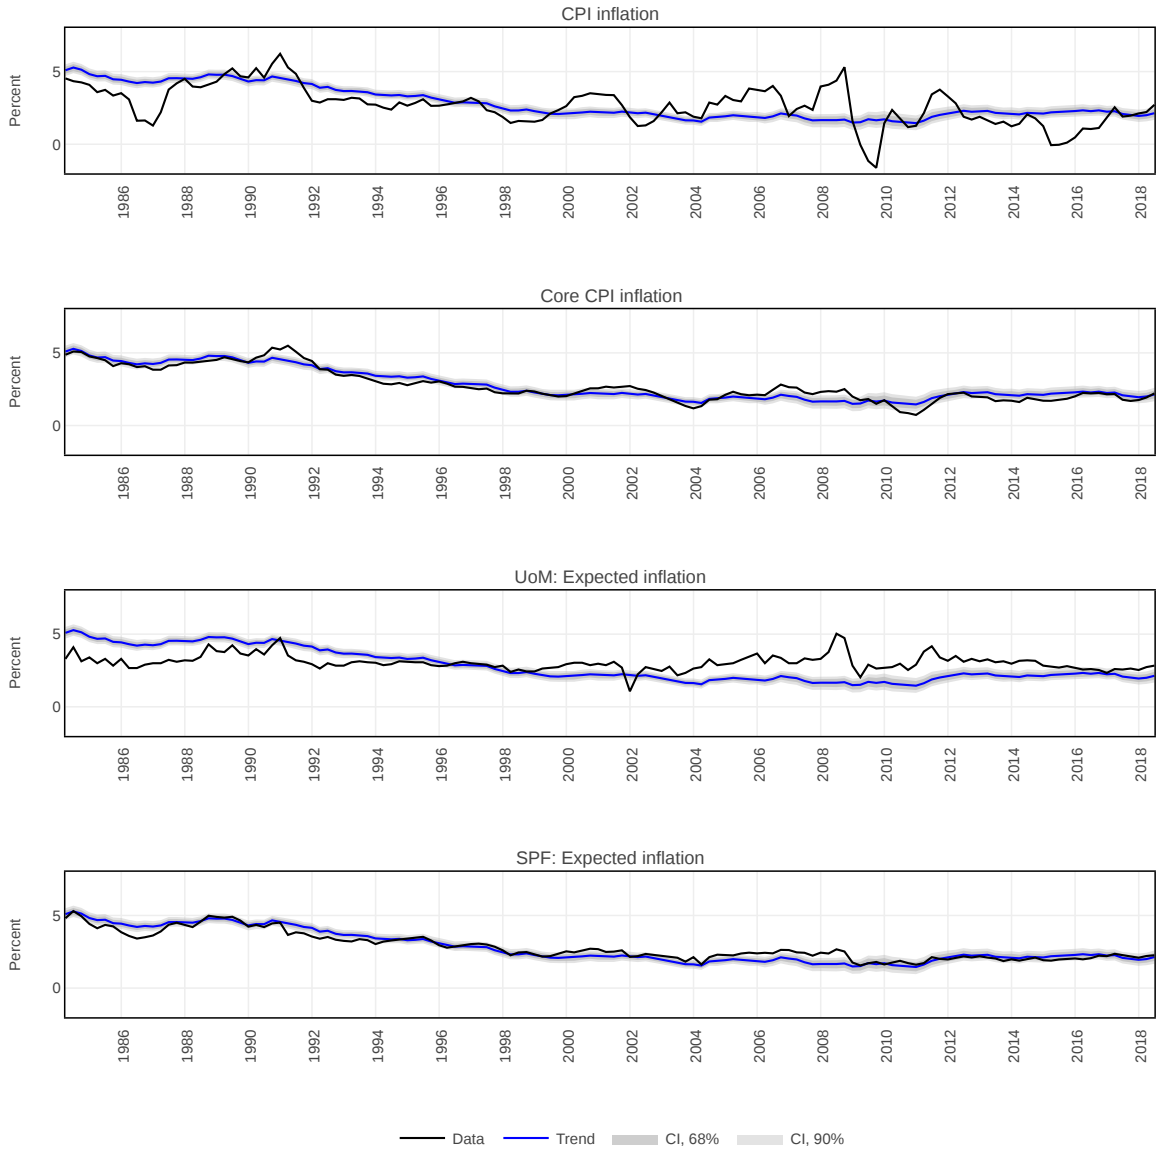

Figure 27: Trend common to CPI inflation, core CPI inflation, and inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

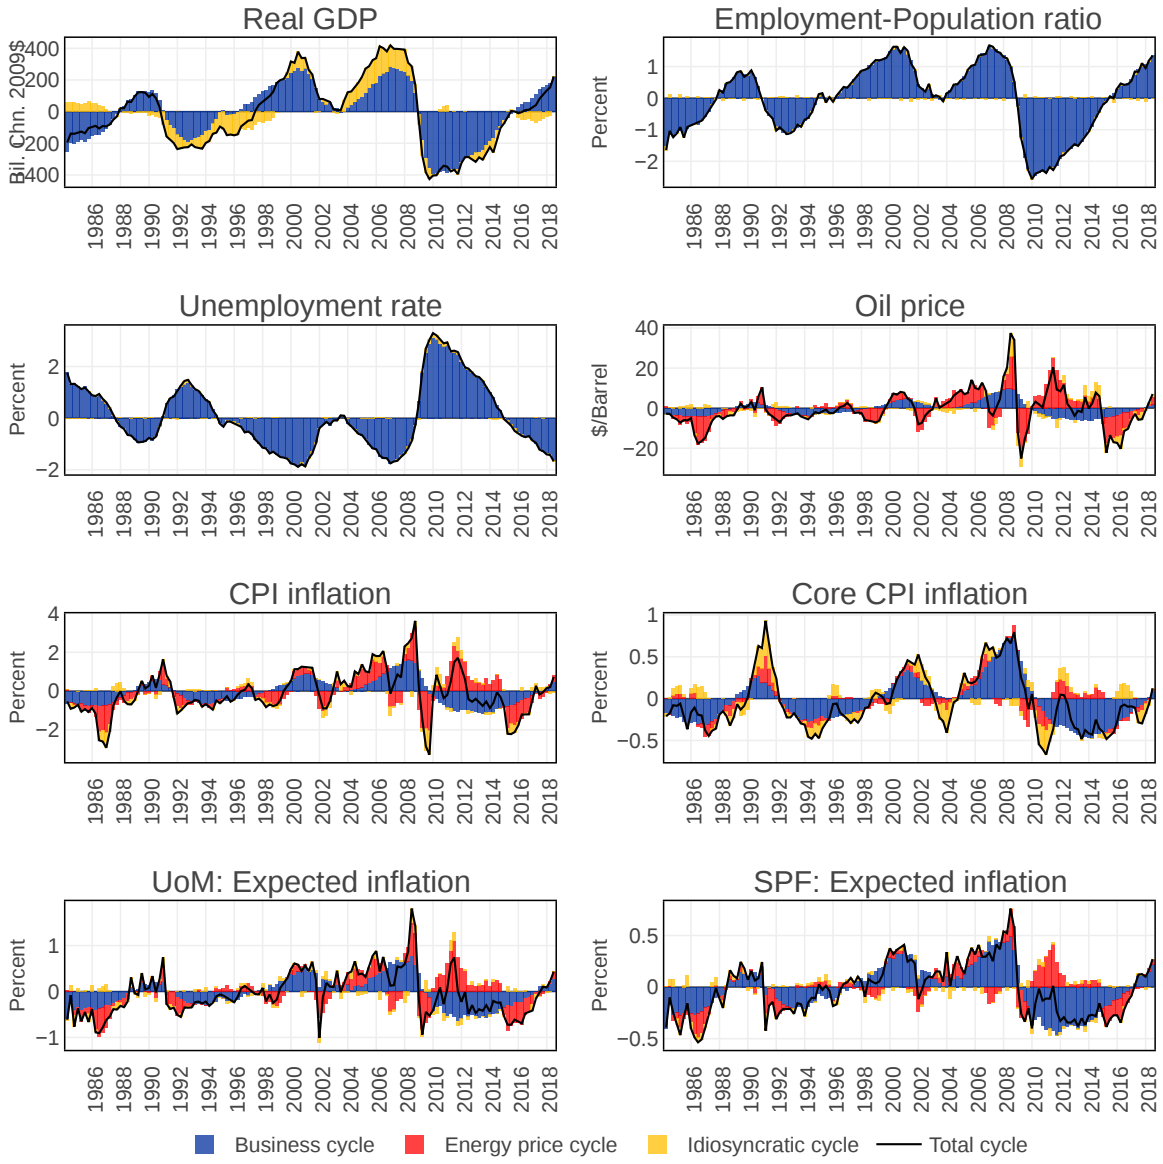

Figure 28: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

Appendix F Global Activity

F.1 Correlation of Cycles with Global Activity Indicators

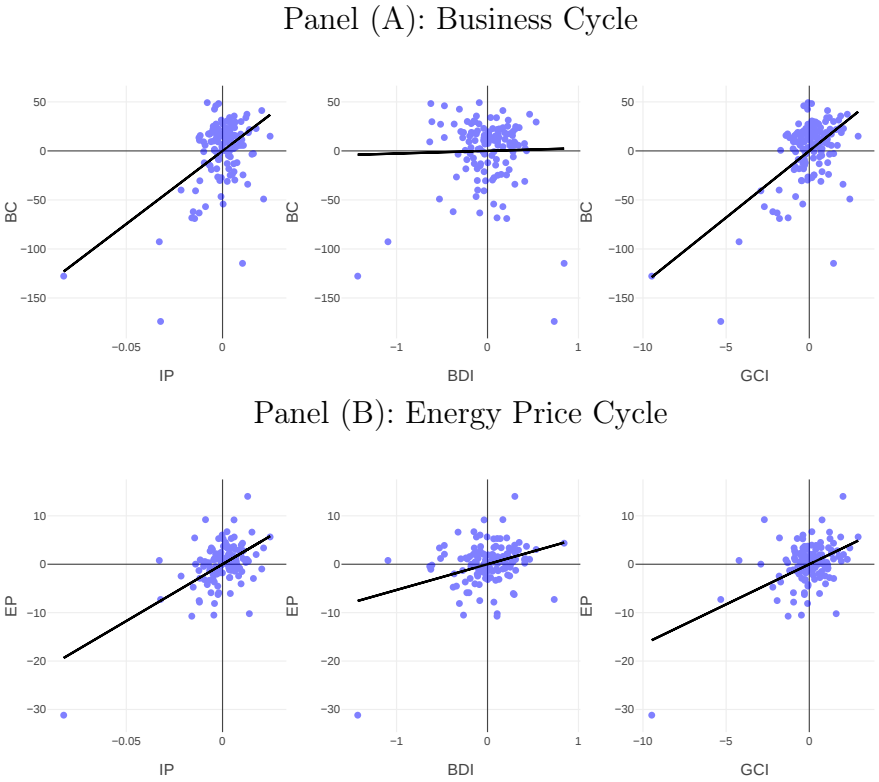

Figure 29: Scatter-plots and regression lines of the business cycle and the energy price cycle (differenced and standardised) on three global activity variables: (i) the Baltic Dry Index (BDI), initially proposed by Kilian (2009) but taken in level; (ii) the measure of global industrial production (GIP) proposed by Baumeister and Hamilton (2019) and based on the OECD methodology; and (iii) the Global Condition Index (GCI) of Cuba-Borda et al. (2018).

Table 6: Data and transformations

Variable	Symbol	Mnemonic	Transformation
Real GDP	y_t	y	Levels
Employment	e_t	e	Levels
Unemployment rate	u_t	u	Levels
Global industrial production	ip_t	ip	Levels
Baltic Dry Index	bdi_t	bdi	Levels
Oil price	oil_t	oil	Levels
CPI inflation	π_t	π	YoY
Core CPI inflation	π_t^c	π^c	YoY
UoM: Expected inflation	$F_t^{uom}\pi_{t+4}$	uom	Levels
SPF: Expected CPI	$F_t^{spf}\pi_{t+4}$	spf	Levels

Note: The table lists the macroeconomic variables used in the empirical model. 'UoM: Expected inflation' is the University of Michigan, 12-months ahead expected inflation rate. 'SPF: Expected CPI' is the Survey of Professional Forecasters, 4-quarters ahead expected CPI inflation rate. The oil price is the West Texas Intermediate Spot oil price. The Baltic Dry Index is issued daily by the London-based Baltic Exchange as a proxy for global activity. Global industrial production was downloaded from James Hamilton's webpage and it is part of the replication dataset of Baumeister and Hamilton (2019).

F.2 A Model with Global Indicators

Our model in $\tilde{x}_t := \{y_t, e_t, u_t, ip_t, bdi_t, oil_t, \pi_t, \pi_t^c, F_t^{uom}\pi_{t+4}, F_t^{spf}\pi_{t+4}\}$ can be written as

$$\begin{pmatrix} y_t \\ e_t \\ u_t \\ ip_t \\ bdi_t \\ oil_t \\ \pi_t \\ \pi_t^c \\ F_t^{uom}\pi_{t+4} \\ F_t^{spf}\pi_{t+4} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ \delta_{e,1}+\delta_{e,2}L & 0 & 0 \\ \delta_{u,1}+\delta_{u,2}L & 0 & 0 \\ \delta_{ip,1}+\delta_{ip,2}L & 1 & 0 \\ \delta_{bdi,1}+\delta_{bdi,2}L & \gamma_{bdi} & 0 \\ \delta_{oil,1}+\delta_{oil,2}L & \gamma_{oil,1}+\gamma_{oil,2}L & 0 \\ \delta_{\pi,1}+\delta_{\pi,2}L & \gamma_{\pi,1}+\gamma_{\pi,2}L & \phi_\pi \\ \delta_{\pi^c,1}+\delta_{\pi^c,2}L & \gamma_{\pi^c,1}+\gamma_{\pi^c,2}L & \phi_{\pi^c} \\ \delta_{uom,1}+\delta_{uom,2}L+\delta_{uom,3}L^2 & \gamma_{uom,1}+\gamma_{uom,2}L & \phi_{uom} \\ \delta_{spf,1}+\delta_{spf,2}L+\delta_{spf,3}L^2 & \gamma_{spf,1}+\gamma_{spf,2}L & \phi_{spf} \end{pmatrix} \begin{pmatrix} \widehat{\psi}_t \\ \psi_t^{GD} \\ \mu_t^\pi \end{pmatrix} + \begin{pmatrix} \psi_t^y \\ \psi_t^e \\ \psi_t^u \\ \psi_t^{ip} \\ \psi_t^{bdi} \\ \psi_t^{oil} \\ \psi_t^\pi \\ \psi_t^{\pi^c} \\ \psi_t^{uom} \\ \psi_t^{spf} \end{pmatrix} + \begin{pmatrix} \mu_t^y \\ \mu_t^e \\ \mu_t^u \\ \mu_t^{ip} \\ \mu_t^{bdi} \\ \mu_t^{oil} \\ 0 \\ 0 \\ \mu_t^{uom} \\ \mu_t^{spf} \end{pmatrix} \tag{9}$$

where ϕ_π, ϕ_{π^c}, ϕ_{uom}, and ϕ_{spf} are normalised to have unitary loading of inflation and inflation expectations on trend inflation.[8]

[8] In the empirical model, the series are standardised so that the standard deviations of their first differences are equal to one. For this reason, we normalise ϕ_π, ϕ_{π^c}, ϕ_{uom}, and ϕ_{spf} to the reciprocal of the

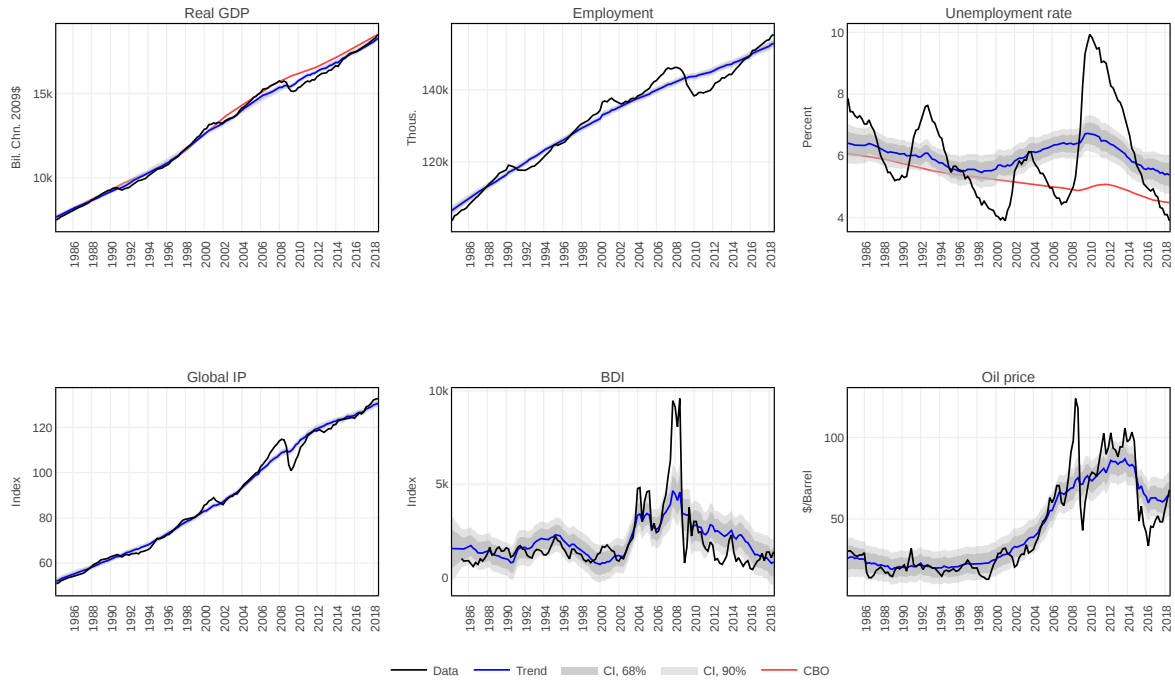

Figure 30: Independent trends of output, employment, unemployment, and oil prices (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model. The chart also reports the measures of potential outputs and NAIRU estimated by the CBO (in red).

standard deviation of the first difference of the respective variable.

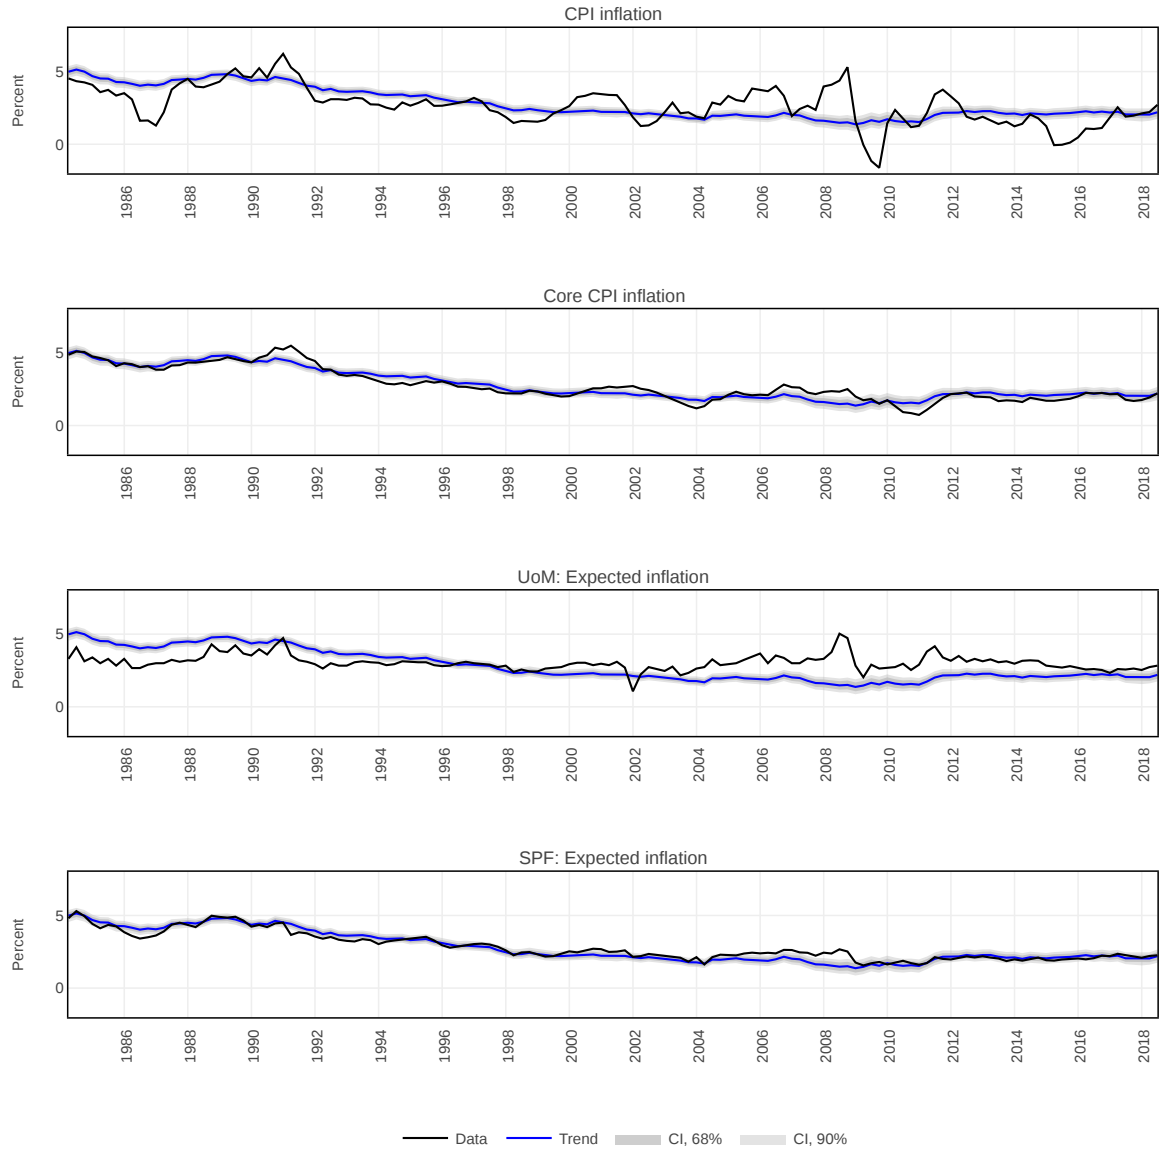

Figure 31: Trend common to CPI inflation, core CPI inflation, and inflation expectations (in blue), with coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model.

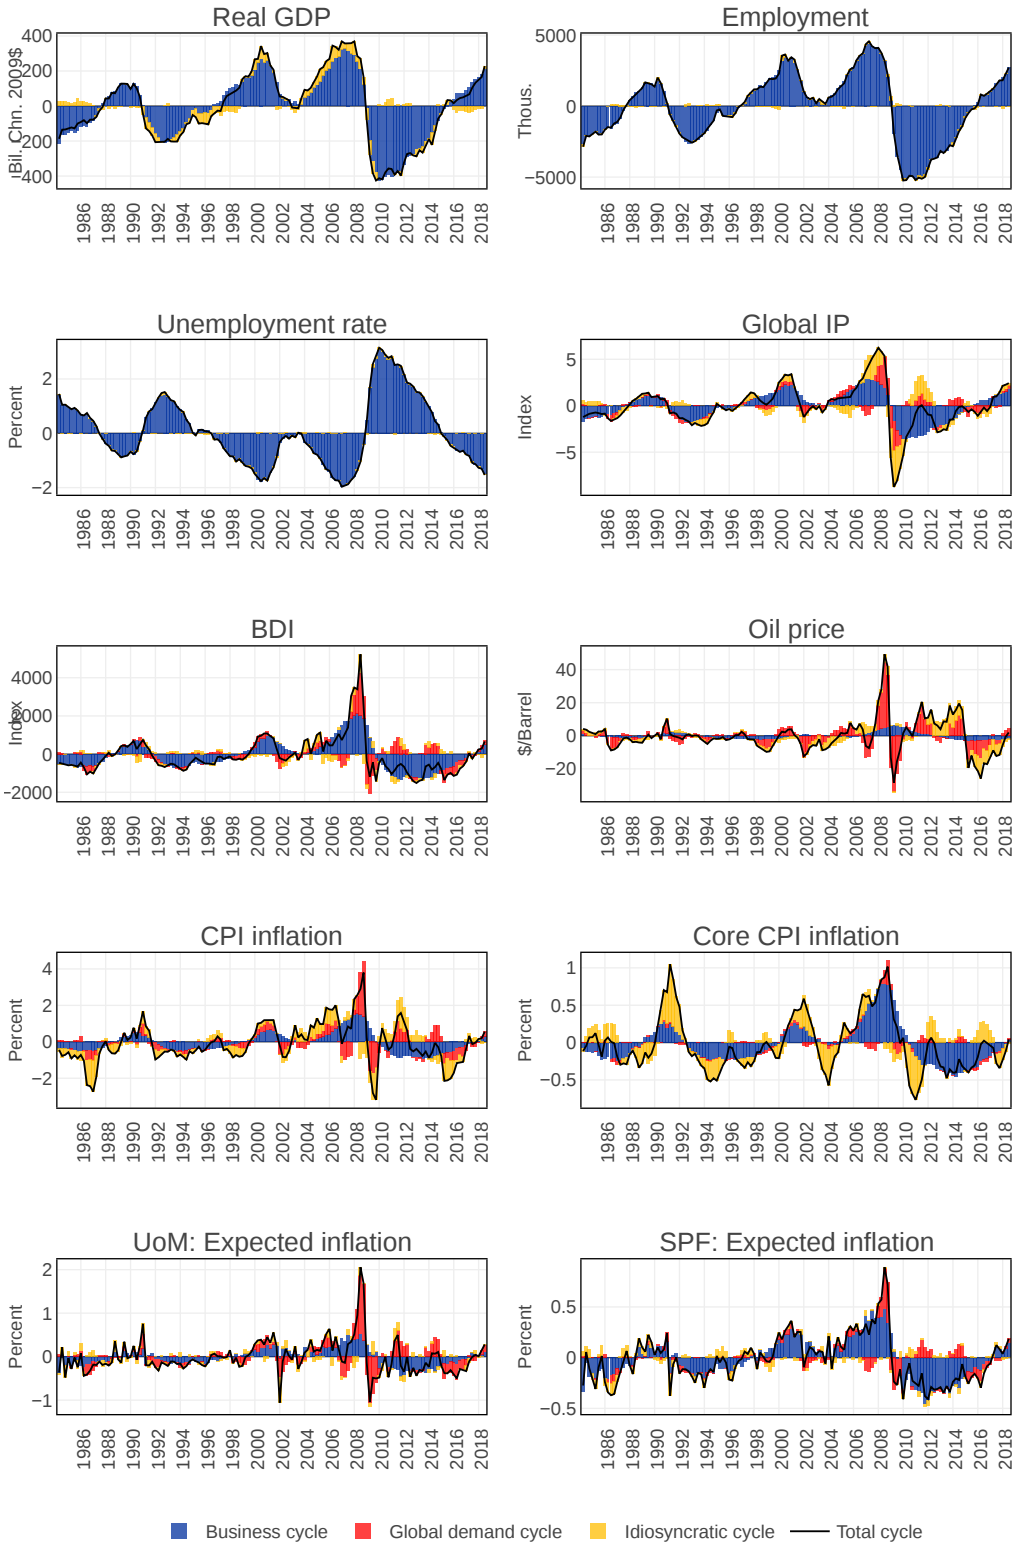

Figure 32: Historical decomposition of the cycles, as estimated by the model. The chart reports the Business cycle (in blue), Energy price cycle (in red), and idiosyncratic cycle (in yellow).

Appendix G Model Forecasting Performance

Our results show that a trend-cycle model, incorporating key economic relations and allowing for deviations of agents' forecasts from full information rational expectations, provides a coherent 'structural' interpretation of economic developments in the US from the 1980s onwards, based on fundamental and generally accepted economic relationships. While this is an important and desirable feature of the 'in-sample' behaviour of the model, an additional test of robustness and reliability of the model is provided by its out-of-sample behaviour.

In this section of the Online Appendix we provide an out-of-sample assessment of the model along two dimensions. First we look at trends and cycles extracted by the model in expanding samples, as it would happen in out-of-sample forecast, and check for their stability. This is important in assessing whether the historical decomposition provided by the model is reliable in a pseudo-real-time exercise. Second we test the out-of-sample forecasting performance of the model against two of the best performing models used for inflation forecasting. Forecasting inflation is notoriously difficult and good performance from such a complex model would provide indirect evidence of whether the model is able to capture important features of the data generating process.

Figure 33 shows the revisions of the two common cycles and of the inflation trend over time with an expanding data window. The model is re-estimated every quarter. The period from Q1 1984 to Q4 1998 is employed as the pre-sample, while the evaluation sample starts in Q1 1999 and ends in Q2 2018. Results show that trends and the common business cycle are fairly stable overall and provide an assessment of the development in the economy that is evenly consistent over the sample - including in the recessions. The energy price cycle provides a slightly less stable, albeit roughly coherent, reading of the contribution of energy fluctuations to prices.

The forecasting exercise is conducted in the same sample and again the period from Q1 1984 to Q4 1998 serves as the pre-sample. We use an expanding window and recursively

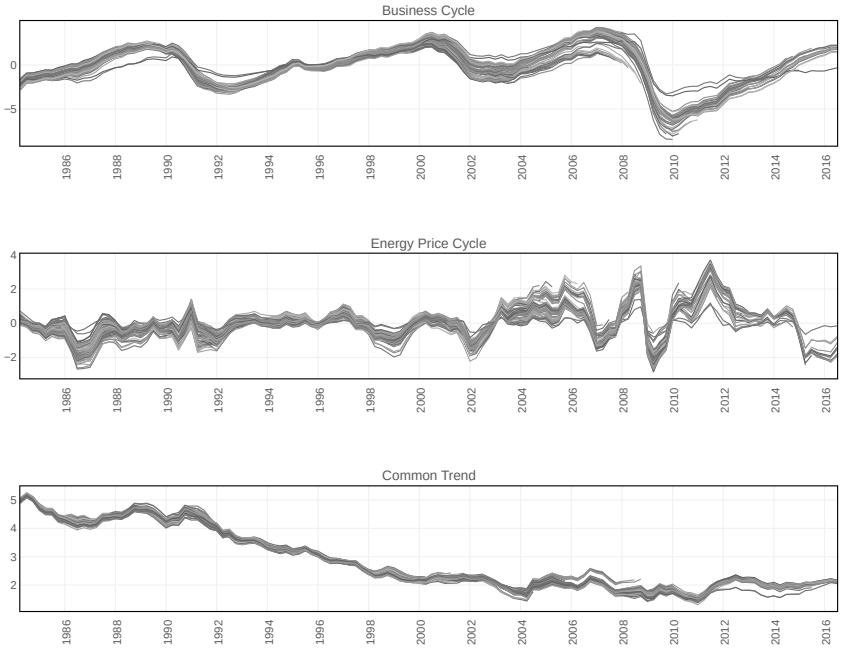

Figure 33: This chart shows the revisions of the business cycle (top), energy price cycle (middle), and common trend (bottom) as estimated during the OOS forecasting exercise.

forecast up to 8 quarters ahead. The final quarter that we condition the forecasts on is Q2 2016, so that the last 8-quarters-ahead forecast is for Q2 2018. In every quarter we reestimate the model, including the unobserved components and the coefficients. Apart from our model (TC), we consider (i) a BVAR where priors are set as in Giannone et al. (2015), (ii) a BVAR with "long-run" prior as in Giannone et al. (2019), and (iii) an univariate unobserved components IMA(1,1) with stochastic volatility model as suggested by Stock and Watson (2007) to be tough benchmarks for inflation forecasts. In setting the system with long-run priors we try to closely replicate the main assumptions on trends adopted in our trend-cycle model. In particular, we set long-run priors considering a common trend between CPI inflation, core CPI inflation, and inflation expectations. We allow for the difference between core CPI inflation and CPI inflation, and the difference between inflation expectations and CPI inflation to be stationary.[9]

For all models we report the root mean squared forecast errors relative to those of a

[9]In the Online Appendix H, we provide details on how the long-run priors are elicited, following an approach that is analogous to the one followed in designing the trend-cycle model.

random walk with drift for forecasting horizons of one, two, four, and eight quarters. We also report the test statistical significance of each benchmark model forecast against the trend-cycle model forecast using Diebold and Mariano (1995) with a quadratic loss function and the modification from Harvey et al. (1997). The Diebold-Mariano test assesses the null hypothesis that the benchmark model forecast and the trend-cycle model forecast are equally accurate. Hence, if the relative root mean squared forecast error of the benchmark model is larger than the relative root mean squared forecast error of the trend-cycle model and we can reject the null hypothesis of equal accuracy, we are allowed to conclude that the trend-cycle model forecast is statistically significantly more accurate than the benchmark model forecast.

Results are reported in Table 7. They show that the trend-cycle model outperforms all others for CPI inflation and core CPI inflation at the 4 and 8 quarters ahead horizons. Our conjecture is that our advantage with respect to the two BVARs is driven by the random walk trend which captures the slow-moving, low frequency component. This is consistent with the fact that the advantage of the trend-cycle model over the BVARs is statistically significant for core inflation at least at the 10% level but not for CPI inflation, since the inflation trend explains a larger fraction of core inflation than of CPI inflation. The advantage of the trend-cycle model with respect to the UC-SV models is most likely due to the Phillips curve which captures cyclical co-movements. This explains why the advantage is more significant at shorter horizons, where the cyclical components in the forecast are larger than at long horizons. The trend-cycle model and the BVARs have similar performance in relation to the other variables with the exception of employment one quarter ahead where both BVARs outperforms our model with a difference which is statistically significant at the 10% level.

Results seem to indicate that despite the large number of parameters and the imposition on the data of structural relationships dictated by economic theory, the model provides a stable historical decomposition in a pseudo real-time exercise and very good performance in forecasting. We consider this as evidence providing support to the claim that the model

Table 7: Relative Root Mean Squared Errors

Horizon	Variable	TC Model	MN-SOC-BVAR	PLR-BVAR	UC-SV
h=1	Real GDP	1.00	0.95	**0.94**	x
	Employment	0.94	0.76*	**0.75***	x
	Unemployment rate	0.82	0.68	**0.63**	x
	Oil price	**1.06**	1.09	1.08	x
	CPI Inflation	0.97	0.91	**0.86**	1.00***
	Core CPI Inflation	1.00	1.03	**0.97**	1.01***
	UOM: Expected inflation	1.03	1.04	**0.99**	x
	SPF: Expected CPI	**1.00**	1.06	1.06	x
h=2	Real GDP	1.02	**0.96**	0.97	x
	Employment	0.95	**0.75**	**0.75**	x
	Unemployment rate	0.80	0.72	**0.65**	x
	Oil price	**1.08**	1.18	1.19	x
	CPI Inflation	0.95	0.97	**0.92**	0.99***
	Core CPI Inflation	**0.95**	1.13	1.04	0.99***
	UOM: Expected inflation	**1.01**	1.09	1.04	x
	SPF: Expected CPI	**0.97**	1.18**	1.24*	x
h=4	Real GDP	1.04	**1.04**	1.04	x
	Employment	0.99	**0.82**	0.81	x
	Unemployment rate	**0.81**	0.84	0.75	x
	Oil price	**1.12**	1.26	1.26	x
	CPI Inflation	**0.95**	1.12	1.05	0.98**
	Core CPI Inflation	**0.89**	1.22*	1.12	0.96***
	UOM: Expected inflation	**1.11**	1.15	1.10	x
	SPF: Expected CPI	**0.91**	1.28*	1.42**	x
h=8	Real GDP	**1.11**	1.21	1.16	x
	Employment	1.07	1.01	**0.95**	x
	Unemployment rate	**0.81**	1.02***	0.85	x
	Oil price	**1.10**	1.34	1.35	x
	CPI Inflation	**0.85**	1.07	0.95	0.96*
	Core CPI Inflation	**0.83**	1.30**	1.13*	0.91
	UOM: Expected inflation	**1.02**	1.29	1.16	x
	SPF: Expected CPI	**0.86**	1.33*	1.31**	x

Note: This table shows the RMSEs relative to the random walk with drift. The MN-SOC-BVAR is a BVAR with "Minnesota" and "Sum-of-coefficients" priors and was estimated using Giannone et al. (2015). The PLR-BVAR is a BVAR with "long-run prior" as in Giannone et al. (2019). The UC-SV model was first proposed in Stock and Watson (2007). We test that the forecasts of each other model are statistically different from the trend-cycle model forecasts using Diebold and Mariano (1995) with a quadratic loss function and the modification from Harvey et al. (1997). *p <0.1, **p<0.05, ***p<0.01.

is able to capture important features of the data generating process.

Appendix H Priors for the Long-Run

In Appendix G we compare the forecast of the trend-cycle model with the forecast of a Bayesian VAR with the priors for the long run proposed in Giannone et al. (2019). Those priors require us to elicit a matrix H that captures the cointegration relationships between the variables in our information set $\{y_t, e_t, u_t, oil_t, \pi_t, \pi_t^c, F_t^{uom}\pi_{t+4}, F_t^{spf}\pi_{t+4}\}$. In the forecast exercise we adopt the following H matrix, in line with the assumptions made in the trend-cycle model:

$$H = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 \end{pmatrix} \begin{array}{l} \leftarrow \text{Idio trend in GDP} \\ \leftarrow \text{Idio trend in Employment} \\ \leftarrow \text{Idio trend in Unemployment Rate} \\ \leftarrow \text{Idio trend in Oil prices} \\ \leftarrow \text{Common trend in inflation and expectations} \\ \leftarrow \text{CPI and core inflation are cointegrated} \\ \leftarrow \text{CPI and UoM expectations are cointegrated} \\ \leftarrow \text{CPI and SPF expectations are cointegrated} \end{array} \tag{10}$$

References

Baumeister, Christiane and James D Hamilton, "Structural interpretation of vector autoregressions with incomplete identification: Revisiting the role of oil supply and demand shocks," *American Economic Review*, 2019, *109* (5), 1873–1910.

Beveridge, Stephen and Charles R Nelson, "A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the business cycle," *Journal of Monetary Economics*, 1981, *7* (2), 151–174.

Clark, Peter K, "The cyclical component of US economic activity," *The Quarterly Journal of Economics*, 1987, *102* (4), 797–814.

Cuba-Borda, Pablo, Alexander Mechanick, and Andrea Raffo, "Monitoring the World Economy : A Global Conditions Index," IFDP Notes 2018-06-15, Board of Governors of the Federal Reserve System (U.S.) June 2018.

Diebold, Francis X. and Roberto S. Mariano, "Comparing Predictive Accuracy," *Journal of Business & Economic Statistics*, 1995, *13* (3), 253–263.

Durbin, James and Siem Jan Koopman, "A simple and efficient simulation smoother for state space time series analysis," *Biometrika*, 2002, pp. 603–615.

_ and _ , *Time series analysis by state space methods*, Vol. 38, OUP Oxford, 2012.

Giannone, Domenico, Michele Lenza, and Giorgio E Primiceri, "Prior selection for vector autoregressions," *Review of Economics and Statistics*, 2015, *97* (2), 436–451.

_ , _ , and Giorgio E. Primiceri, "Priors for the Long Run," *Journal of the American Statistical Association*, 2019, *114* (526), 565–580.

Harvey, Andrew C, "Trends and cycles in macroeconomic time series," *Journal of Business & Economic Statistics*, 1985, *3* (3), 216–227.

_ , **Thomas M Trimbur, and Herman K Van Dijk**, "Trends and cycles in economic time series: A Bayesian approach," *Journal of Econometrics*, 2007, *140* (2), 618–649.

Harvey, David, Stephen Leybourne, and Paul Newbold, "Testing the equality of prediction mean squared errors," *International Journal of Forecasting*, 1997, *13* (2), 281 – 291.

Herbst, E.P. and F. Schorfheide, *Bayesian Estimation of DSGE Models* The Econometric and Tinbergen Institutes Lectures, Princeton University Press, 2015.

Jarociński, Marek, "A note on implementing the Durbin and Koopman simulation smoother," *Computational Statistics & Data Analysis*, 2015, *91*, 1–3.

Kilian, Lutz, "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, June 2009, *99* (3), 1053–69.

Koopman, Siem J and James Durbin, "Fast filtering and smoothing for multivariate state space models," *Journal of Time Series Analysis*, 2000, *21* (3), 281–296.

Stock, James H and Mark W Watson, "Why has US inflation become harder to forecast?," *Journal of Money, Credit and banking*, 2007, *39* (s1), 3–33.

Warne, Anders, "YADA Manual," *Manuscript, December*, 2008, *18*.