

Analysing inflation with semi-structural models

Thomas Hasenzagl¹, Filippo Pellegrino², Lucrezia Reichlin³, and
Giovanni Ricco⁴

¹*University of Minnesota*

²*Imperial College London*

³*London Business School, Now-Casting Economics, and CEPR*

⁴*CREST – École Polytechnique, University of Warwick, OFCE-SciencesPo, and CEPR*

Abstract

The chapter discusses semi-structural time series models for the analysis, forecast and now-cast of inflation. We define a semi-structural time series model as a multivariate structural time series model in the tradition of [Harvey \(1985\)](#) and [Harvey \(1990\)](#), where minimal economic restrictions are used to identify common and idiosyncratic trend and cyclical components of the observable data. We discuss the potential of this approach for inflation conjunctural analysis, forecasting and now-casting in comparison with more widely used models in empirical macroeconomics such as factor models and VARs.

Keywords: real-time forecasting, output gap, Phillips curve, structural time series, semi-structural models, Bayesian estimation.

JEL Classification: C11, C32, C53, E31, E32, E52.

We are grateful to the 13th International Conference on Computational and Financial Econometrics, European Seminar on Bayesian Econometrics 2021 and Bank of Finland, University of Strathclyde, and Université de Montréal seminars participants for helpful comments and suggestions. An early version of this work appeared under the title “Monitoring the Economy in Real Time: Trends and Gaps in Real Activity and Prices”.

Introduction

Inflation in advanced economies is a difficult variable to model and forecast. There are many reasons for this. The first is that it features a slowly changing trend which reflects changes in monetary policy regimes and structural forces over time. The second is that it is very sensitive to unpredictable and highly volatile fluctuations of commodity prices. The dynamics at business cycle frequency, possibly influenced by the degree of slack in economic activity (the so-called Phillips curve), is dwarfed by lower and higher frequency variation. This is why the question on the existence and size of the Phillips curve has been and continues to be controversial (see the chapter on ‘The Slope of the Phillips Curve’, in this handbook, for a review of the debate).

The literature has been traditionally split between those asserting that the best models for forecasting inflation are univariate and those who claim that measures of the slack of the economy have a role to play. Early supporting evidence for the univariate view is provided by [Atkeson and Ohanian \(2001\)](#) who argued that the random walk is the best model to forecast inflation. More recently, [Angeletos et al. \(2020\)](#) estimated that only 7% of US inflation can be explained by a business cycle shock. [Stock and Watson \(2009\)](#), in a forecasting evaluation including many models, concluded that inflation forecasts based on the Phillips curve in the traditional specification of [Gordon \(1990\)](#) do well in some sub-samples but, for a long sample including the period of the great moderation, does not improve over the univariate unobserved component model of [Stock and Watson \(2007\)](#). They also concluded that results are in general sample dependent. Recent literature confirms these results overall.¹

In sum, instability over time, reflecting changes in policy or changes in size and frequency of supply shocks, makes it difficult to establish robust conclusions about inflation at business cycle frequencies. This motivates a modelling approach that is able to both

¹See also [Mavroeidis et al. \(2014\)](#) for a recent discussion the related issue of the identification of the New Keynesian Phillips Curve.

capture the variability of the trend over time and to ‘clean’ business cycle variation from changes explained by energy prices and other supply-side factors. Multivariate structural time series models offer a promising modelling strategy in this regard. In this chapter, we explore this conjecture.

Structural time series models are well-established tools in statistics. These are models in which time series are represented by components which capture dynamics at different frequencies (classical references are [Harvey, 1985](#) and [Harvey, 1990](#)). In the original formulation, the univariate version of these models included a trend, a cycle, a seasonal component and an irregular component, each component being stochastic and mutually uncorrelated. The multivariate generalisation of these models can handle common stochastic trends and cycles.

These models have had multiple applications in forecasting since they offer a parsimonious representation of dynamic relationships. They have also been used to represent stylised facts and to handle seasonality. Indeed, the early literature (e.g. [Harvey and Jaeger, 1993](#)) shows how this approach gives a more accurate description of cyclical characteristics of economic data than methods based on detrending techniques such as differencing or band-pass filtering which typically generate spurious cycles.

In macroeconomics, multivariate structural time series models can be particularly useful to characterise the co-movements of macroeconomic and financial variables, at business cycle frequency. In the early days of business cycle analysis, [Burns and Mitchell \(1946\)](#) used heuristic methods to identify cyclical variations of the data. Later, [Sargent and Sims \(1977\)](#) suggested that the business cycle could be identified as a common factor extracted from multiple macroeconomic series. These ideas were then applied and refined by [Stock and Watson \(1989\)](#) and, for large panels of time series by [Forni and Reichlin \(1998\)](#), [Forni et al. \(2000\)](#), [Stock and Watson \(2002\)](#) and related literature.

As in factor analysis, multivariate structural time series models can be used to estimate unobserved common components from observations in a multivariate setting. When

used to extract common factors from the data, these factors are identified for different frequency ranges, distinguishing between trends and cycles, and between common and idiosyncratic components. The assumption of orthogonality amongst components, and explicit modelling of their dynamic characteristics, provide the ‘structure’ to the dynamic system. As in factor models with stochastic trends (see, for instance, [Barigozzi and Luciani, 2023](#)) and VAR with stochastic trends (see [Del Negro et al., 2017](#), and [Ascari and Fosso, 2024](#)), they capture changing low-frequency variation of the data, avoiding the trend-cycle contamination which results from estimation of VARs with a deterministic trend. However, by explicitly modelling cycles at different frequencies, multivariate structural models have the additional advantage of extracting a clean signal of the business cycle and distinguishing it from other temporary fluctuations, possibly driven by supply-side factors.

Structural time series models are represented in state space form, where the system’s state captures various unobserved components, and can be estimated both with frequentist and Bayesian methods. In the frequentist approach, the estimate of the unobservable state can be updated using a filtering procedure as new observations become available. Predictions are then made by projecting these estimated components into the future. Smoothing algorithms provide the best estimate of the state at any point within the sample period. For a Gaussian model, the likelihood function is obtained from the innovation produced by the Kalman filter and maximised with respect to unknown parameters. Alternatively, Bayesian methods involve applying priors to key parameters, such as the relative variability of the cycle and the trend, and then using standard algorithms to compute the posterior distributions of these parameters. The Bayesian approach is especially useful for assessing uncertainty in the state estimates and for managing models with a large state space.

Recently, there has been a renewed interest in these techniques in macroeconomics. Relevant references are [Jarociński and Lenza \(2018\)](#) (Euro Area output gap), [Hasenzagl](#)

[et al. \(2022a\)](#) (US trend inflation and the Phillips curve) and [Bianchi et al. \(2022\)](#) (the US Phillips curve). These are medium-scale models which exploit economic theory to inform the identifying restrictions. From a statistical standpoint, these models are multivariate structural time series models. From the economic standpoint, we define them as semi-structural since they exploit economic theory to identify the components.

The focus of this chapter is to use the semi-structural approach to describe inflation trends and cycles in conjunction with real variables and to forecast and nowcast monthly inflation. Given the observation that inflation is characterized by a variable low-frequency component and is affected by the commodity price cycle, a semi-structural time series model as we have defined it, is a promising method for extracting information on cyclical inflation and the extent of its commonality with indicators of the real economy.

Borrowing from the quarterly model in [Hasenzagl et al. \(2022a\)](#), we illustrate the methodology and how minimal assumptions from economic theory can be used to identify the components. We then extend the model to the case in which the data includes series observed at different frequencies (in our case quarterly and monthly) and have missing observations at the end of the sample reflecting the different lags at which data are released. These additional features make the model not only suited for data description and forecasting but also for now-casting. Indeed our method combines the insight of the now-casting literature (see [Giannone et al., 2008](#) for seminal contribution and [Bańbura et al., 2011](#), [Modugno, 2013](#) for early results on nowcasting inflation) with those of structural time series analysis.

For both models, the specification is coherent with the commonly accepted description of the economy in terms of trends – such as potential output, the natural rate of unemployment or NAIRU – and trend inflation –, and cycles – in particular, the output gap and its link to prices via a Phillips curve, and to labour market variables via the Okun’s law. We estimate the system by Bayesian methods as a way to handle the relatively large dimension of the state space specification.

Both the quarterly and the mixed frequency models are estimated on US data including activity indicators, labour market variables, inflation and inflation expectations variables. For both, we provide out-of-sample forecasts and trends-cycle decompositions with the aim of assessing in particular the importance of the Phillips curve in explaining the cyclical behaviour of inflation. The mixed frequency model generates, as a by-product, a monthly estimate of the output gap. For that model, we also produce an analysis of nowcasting results in relation to the flow of data releases.

As an additional exercise, based on the mixed frequency model, we estimate a constrained version that is identical to the baseline model except that the output gap is treated as an observed quantity and measured by the Congressional Budget Office's (CBO) estimate. The baseline and the constrained model provide equally plausible alternative ways to fit the data, but different characterizations of trends and cycles. The comparison of the results provides insight into the implications of the CBO view on the output gap for the trend-cycle decomposition of other variables such as employment, unemployment and different measures of inflation. This helps assess its validity through the lenses of a model.

The results reported in this chapter confirm findings in [Hasenzagl et al. \(2022a\)](#) about a sizeable Phillips curve component of the inflation cycle. Forecasting results show an improvement over univariate and multivariate benchmarks in forecasting inflation and labour market variables at a medium-term horizon between two and three years. At the now-casting horizon, on the other hand, the only variable that matters besides inflation itself is oil prices which is more timely than the other indicators. As the literature has shown, at the very short horizon, the timeliness of the data release is what matters.

The comparison between the CBO-constrained model and our baseline points to many similarities between the two specifications but also significant differences. The first model reads the US economy as having been constantly below potential since the 2001 recession, reflecting the implicit view that output potential has been unaffected by events since

then. Conversely, the unconstrained model estimates an almost symmetrical output gap fluctuating around a trend whose slope has declined since 2001. The divergent views are reflected in differences in the estimate of the slope of the Phillips curve, especially since 2001. The unconstrained model identifies a larger common cycle between inflation and real variables and implies, on average, a slightly larger cyclical component in inflation related to the output gap. Finally, a real-time evaluation of the CBO output gap nowcast during the COVID sample, shows that the revisions of the output gap in the baseline model are smaller than those in the implied monthly CBO output gap.

The paper is organised as follows. In Section 1, we introduce a toy macroeconomic model which motivates the identifying restrictions used in the statistical specification. In section two, we describe the three models used in the empirical analysis. Section 3 illustrates the components' decompositions resulting from in-sample analysis while Section 4 provides a real-time analysis of the output gap during the COVID period. The last section concludes. An Online Appendix provides details on the Bayesian estimation of the models and additional results for all of the models discussed in the paper.

1 Motivation: a stylised view of the economy

The modelling approach we describe in this chapter is motivated by a conventionally accepted stylised representation of economic variables in terms of trends and cycles. Let us review it. Output is generally described as fluctuating around a long-run trend (potential output), that is driven by demographic trends, capital accumulation and technological innovation. The trend component is often modelled as a non-stationary unit root process. Different shocks can push output above or below its potential. The fluctuations off equilibrium are defined in terms of an output gap, often modelled as an AR(p) stationary

process, that can be seen as the ‘primitive’ measure of business cycles. Such a stylised description can be formulated as

$$y_t = \tau_t^y + \hat{y}_t^{gap} = \tau_t^y + \psi_t^{gap} , \quad (1)$$

$$\psi_t^{gap} = \rho(L)\psi_{t-1}^{gap} + v_t , \quad (2)$$

$$\tau_t^y = \mu + \tau_{t-1}^y + \eta_t , \quad (3)$$

where τ_t and ψ_t^{gap} are the output potential and the output gap. The first is a unit root process with a drift μ , and the second is an autoregressive stationary process. The shocks v_t and η_t are i.i.d. innovations to the two components, while L is the lag operator.

Slack in the economy is reflected in labour market variables via the Okun’s law, with unemployment oscillating around a long-run equilibrium level (τ_t^u). This is the unemployment rate consistent with output at its potential and no inflationary pressure, and is commonly referred to as the non-accelerating inflation rate of unemployment (NAIRU):

$$u_t = \tau_t^u + \hat{u}_t^{gap} = \tau_t^u + \gamma_u \hat{y}_t^{gap} . \quad (4)$$

Prices fluctuate at business cycle frequencies around an underlying trend inflation, τ_t^π , that is anchored by the inflation target of a credible central bank and is reflected into the long-run expectations of agents. Deviations from trend inflation are either due to the transmission of cyclical pressure to prices (the Phillips curve), or to short-lived idiosyncratic disturbances, ψ_t^{epc} , possibly related to energy prices that directly enter the basket of consumption, i.e.

$$\pi_t = \lim_{h \rightarrow \infty} E_t \pi_{t+h} + \hat{\pi}_t^{gap} + \psi_t^{epc} = \tau_t^\pi + \gamma_\pi \hat{y}_t^{gap} + \psi_t^{epc} . \quad (5)$$

Such a description of the economy can be summarised by a model of idiosyncratic and common components, capturing the long-run behaviour of the variables and their business-cycle fluctuations,

$$\begin{pmatrix} y_t \\ u_t \\ \pi_t \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \gamma_u & 0 \\ \gamma_\pi & 1 \end{pmatrix} \begin{pmatrix} \psi_t^{gap} \\ \psi_t^{epc} \end{pmatrix} + \begin{pmatrix} \tau_t^y \\ \tau_t^u \\ \tau_t^\pi \end{pmatrix}. \quad (6)$$

While such a description is in line with textbooks and, possibly, the policymakers' view, it is too stylised for empirical analysis. To bring the model to the data we need to introduce lags to account for heterogeneous dynamics across variables, and idiosyncratic shocks reflecting measurement errors or wedges with the theory. Indeed, it is well known that unemployment and prices react with lags to the slack in the economy and that potentially several idiosyncratic components can distort both the formation of economic expectations and the dynamics of the variables themselves. In the next sections, we describe the empirical version of the model and the econometric methodology.

2 Semi-structural models of trends and cycles

Our empirical framework adopts and generalises the model described in the previous section to capture the joint dynamics of real activity – i.e. output, employment and unemployment rate –, nominal variables – i.e. consumer price inflation and oil prices –, and expectations – i.e. professional forecasts of inflation and output, consumers' expectations of inflation.

In this section, we provide three examples of semi-structural modelling to analyse and forecast inflation, jointly with other macro aggregates. First, we discuss the model of [Hasenzagl et al. \(2022a\)](#) that incorporates only quarterly variables to introduce the framework and illustrate our core modelling assumptions. We then introduce two models

Table 1: US data and common components

Variable name	Label	Model			Loads on			
		Q.tly	Un.	Track.	<i>BC</i>	<i>EPC</i>	<i>GDP trend</i>	<i>Trend π</i>
CBO: cycle of real GDP	gap_t^{cbo}	.	.	Q	✓			
Real GDP	y_t	Q	Q	Q	✓		✓	
SPF: expected real GDP	$F_t^y \pi_{t+12}$.	Q	Q	✓		✓	
Unemployment rate	u_t	Q	M	M	✓			
Employment	e_t	Q	M	M	✓			
WTI spot oil price	oil_t	Q	M	M		✓		
CPI	π_t	Q	M	M	✓	✓		✓
Core CPI	π_t	Q	.	.	✓	✓		✓
SPF: expected inflation	$F_t^{spf} \pi_{t+12}$	Q	Q	Q	✓	✓		✓
UoM: expected inflation	$F_t^{uom} \pi_{t+12}$	Q	M	M	✓	✓		✓

Notes: Data used in the three trend-cycle models discussed in this section: the quarterly model (Q.tly), the undisciplined model (Un.) and the tracking model (Track.). The columns under ‘Model’ show, for each model, the variables and the frequencies incorporated in each specification. All data is in levels, except for CPI which is in YoY (%). ‘UoM: expected inflation’ is the University of Michigan, 12-months ahead expected inflation. ‘SPF: expected inflation’ is the Survey of Professional Forecasters, 4-quarters ahead expected inflation rate. Data includes observations from Jan-1985 to Dec-2019.

both including a mix of monthly and quarterly variables and missing variables at the end of the sample reflecting asynchronous data releases (these characteristics made them tailored to provide a nowcast of the variables of interest, in real-time). [Table 1](#) summarises the variables considered, and the frequencies at which they enter the different models, as well as some of the key modelling choices that are discussed in the remainder of this section.

2.1 A quarterly semi-structural model

In a semi-structural model, the core modelling choices are a set of assumptions defining a number of common cycles and trends that are meant to capture structural components and their dynamics, plus several variable-specific components that absorb idiosyncratic shocks and measurement errors. These idiosyncratic components can also be seen as ‘empirical’ wedges capturing the gap between observed data and the assumed structural relationships between variables.

It is important to stress that the decision on which multivariate relationships to explicitly model, and which others to leave unmodelled, is important and has to be be

based on the scope of the model as well as on the evaluation of the relative benefits of complexity and parsimony in estimation and forecasting.

Let us describe the assumptions that underpin our quarterly model, starting from the trends.

Assumption 1 (Output potential). The output potential is the stochastic trend driving output in the long run. In the spirit of [Beveridge and Nelson \(1981\)](#), it coincides with the long-run forecast of output implied by the model.

Assumption 2 (Labor market trends). Employment and the unemployment rate have each their own trend defined as their long-run forecast. We denote them as τ_t^e and τ_t^u , respectively. τ_t^u is the estimate of the non-accelerating inflation rate of unemployment (NAIRU).

Assumption 3 (Trend inflation). Trend inflation, τ_t^π , is the common trend shared by inflation and inflation expectations. It is also the long-run model-based forecast of inflation.

The cyclical components are modelled as stationary stochastic cycles, under the following set of assumptions.

Assumption 4 (Output gap). In the spirit of [Burns and Mitchell \(1946\)](#), the output gap ψ_t^{gap} is defined as an economy-wide stationary stochastic component common to all real variables, labour market variables, inflation, and survey expectations. It informs the price gap via the Phillips curve, and the unemployment gap via the Okun’s law. Both relationships are modelled as moving averages of output gap realisations over the previous three months.²

We also consider a second common stationary component, which we call the ‘energy price component’ that captures the direct effect of energy shocks on headline inflation.

²It is worth observing that all variables, except for real GDP and the CBO’s output cycle, are connected to the output gap with a lag polynomial. This is to allow the model to nest, under parametric restrictions, the case of rational expectations, as discussed in [Hasenzagl et al. \(2022a\)](#).

This may be thought of as capturing the role of energy price disturbances as mark-up shocks.

Assumption 5 (Energy price component). The energy price component ψ_t^{epc} is a stationary stochastic common cyclical component connecting oil prices, inflation, and inflation expectations.

A number of idiosyncratic stationary components absorb various forms of noise which could distort the empirical estimates of the structural relationships.

Assumption 6 (Idiosyncratic stationary components). All variables have an idiosyncratic stationary component, $\psi_{i,t}$, which absorbs different sources of idiosyncratic dynamics such as idiosyncratic shocks, non-classic measurement error, differences in definitions, and other sources of noise.

Finally, we introduce a number of non-stationary components to capture persistent time-varying biases in survey data.

Assumption 7 (Bias in Expectations). Agents' expectations can deviate from a rational forecast due to time-varying bias – respectively $\mu_t^{spf,y}$, $\mu_t^{spf,\pi}$ for the professional forecasters' and $\mu_t^{uom,\pi}$ for consumers' expectations. The bias terms are modelled as stochastic random walk components.

Taken together these assumptions imply that the quarterly model provides a representation of the variables of interest of the form

$$\begin{pmatrix} y_t \\ u_t \\ e_t \\ oil_t \\ \pi_t \\ \pi_t^c \\ F_t^{spf} \pi_{t+12} \\ F_t^{uom} \pi_{t+12} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 0 \\ \sum_{j=0}^1 \gamma_{2,j} L^j & 0 \\ \sum_{j=0}^1 \gamma_{3,j} L^j & 0 \\ \sum_{j=0}^1 \gamma_{4,j} L^j & 1 \\ \sum_{j=0}^1 \gamma_{5,j} L^j & \sum_{j=0}^2 \delta_{5,j} L^j \\ \sum_{j=0}^1 \gamma_{6,j} L^j & \sum_{j=0}^2 \delta_{6,j} L^j \\ \sum_{j=0}^2 \gamma_{7,j} L^j & \sum_{j=0}^2 \delta_{7,j} L^j \\ \sum_{j=0}^2 \gamma_{8,j} L^j & \sum_{j=0}^2 \delta_{8,j} L^j \end{pmatrix}}_{\text{Common \& Idiosyncratic Cycles}} \begin{pmatrix} \psi_{1,t} \\ \psi_{2,t} \\ \psi_{3,t} \\ \psi_{4,t} \\ \psi_{5,t} \\ \psi_{6,t} \\ \psi_{7,t} \\ \psi_{8,t} \end{pmatrix} + \begin{pmatrix} \psi_t^{gap} \\ \psi_t^{epc} \end{pmatrix} + \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 \end{pmatrix}}_{\text{Trends \& Biases}} \begin{pmatrix} \tau^y \\ \tau_t^u \\ \tau_t^e \\ \tau_t^{oil} \\ \tau_t^\pi \\ \mu_t^{spf,\pi} \\ \mu_t^{uom,\pi} \end{pmatrix}, \quad (7)$$

where L is the lag operator.

To complete the state-space representation of the model, we specify the dynamic equations governing the evolution of the unobserved components over time following the approach of [Harvey \(1985\)](#).

Assumption 8 (State dynamics). The stationary cycles are all modelled as ARMA(2,1) stochastic processes, with coefficients restricted to produce stationary oscillations of defined periodicity. The trends are random walks. Specifically, potential output and trend employment are random walks with drift, and all the remaining trends are driftless random walks. All of the processes have mutually orthogonal stochastic innovations.

It is worth noticing that ARMA(2,1) processes display pseudo-cyclical behaviour and can be conveniently written in a VAR(1) representation as

$$\begin{aligned} \widehat{\psi}_t &= \rho \cos(\lambda) \widehat{\psi}_{t-1} + \rho \sin(\lambda) \widehat{\psi}_{t-1}^* + v_t, \\ \widehat{\psi}_t^* &= -\rho \sin(\lambda) \widehat{\psi}_{t-1} + \rho \cos(\lambda) \widehat{\psi}_{t-1}^* + v_t^*, \end{aligned} \quad (8)$$

where the parameters $0 \leq \lambda \leq \pi$ and $0 \leq \rho \leq 1$ can be interpreted, respectively, as the frequency and the damping factor on the amplitude of the cycle (the process is stationary

for $\rho < 1$). $\widehat{\psi}_t^*$ is an auxiliary cycle that supports the VAR(2) representation, and v_t and v_t^* are uncorrelated white noise disturbances (see [Harvey, 1990](#)).³ The disturbances make the cycle stochastic rather than deterministic. As discussed in [Hasenzagl et al. \(2022a\)](#), the empirical specification defined by Equations 7 and 8 reduces to the case of rational expectations, under suitable parametric restrictions.

2.2 A mixed-frequency trend-cycle framework

The mixed frequency models are closely related to the quarterly model and incorporate data at monthly and quarterly frequencies (see [Table 1](#)). In particular, we consider two specifications:

1. an *undisciplined model*, our baseline, that is as specified in [Equation 9](#) but does not include $cycle_t^{cbo}$. Hence the output gap is an unobserved component that the model has to estimate, as in the quarterly case.
2. a *tracking model* that incorporates the CBO measure of the cycle of GDP as an observed quarterly measure of the output gap, and that is reported in [Equation 9](#);⁴

³It is straightforward to show that the model can be rewritten as

$$(1 - 2\rho \cos(\lambda)L + \rho^2 L^2)\widehat{\psi}_t = (1 - \rho \cos(\lambda)L)v_t + (\rho \sin(\lambda)L)v_t^* .$$

Hence, under the restriction $\sigma_v^2 = 0$, the solution of the model is an AR(2), otherwise an ARMA(2,1). The intuition for the use of the auxiliary cycle is closely related to the standard multivariate AR(1) representation of univariate AR(p) processes.

⁴In the general form of the mixed-frequency model, we allow for biases to affect both consumers' and professionals' expectations. However, empirically only consumers' expectations exhibit persistent biases (see [Coibion and Gorodnichenko, 2015](#), for a discussion). In line with this observation, in the empirical section of this paper, we set $\mu_t^{spf,\pi}$ to zero (i.e. we assume that the professional forecasts for inflation (SPF) are 'on trend' at all times), and we only allow the constant $\mu^{spf,y}$ to account for measurement differences in the expected output trend, possibly due to measurement and aggregation issues.

$$\begin{pmatrix} cycle_t^{cbo} \\ y_t \\ F_t^{spf} y_{t+12} \\ u_t \\ e_t \\ oil_t \\ \pi_t \\ F_t^{spf} \pi_{t+12} \\ F_t^{uom} \pi_{t+12} \end{pmatrix} = \underbrace{\begin{pmatrix} \sum_{j=0}^2 L^j & 0 \\ \sum_{j=0}^2 L^j & 0 \\ \sum_{j=0}^3 \gamma_{3,j} L^j & 0 \\ \sum_{j=0}^3 \gamma_{4,j} L^j & 0 \\ \sum_{j=0}^3 \gamma_{5,j} L^j & 0 \\ 0 & 1 \\ \sum_{j=0}^3 \gamma_{7,j} L^j & \delta_7 \\ \sum_{j=0}^3 \gamma_{8,j} L^j & \delta_8 \\ \sum_{j=0}^3 \gamma_{9,j} L^j & \delta_9 \end{pmatrix}}_{\text{Common \& Idiosyncratic Cycles}} + \begin{pmatrix} \sum_{j=0}^2 L^j \psi_{1,t} \\ \sum_{j=0}^2 L^j \psi_{1,t} \\ \psi_{3,t} \\ \psi_{4,t} \\ \psi_{5,t} \\ \psi_{6,t} \\ \psi_{7,t} \\ \psi_{8,t} \\ \psi_{9,t} \end{pmatrix} + \underbrace{\begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \sum_{j=0}^2 L^j & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 3 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 \end{pmatrix}}_{\text{Trends \& Biases}} \begin{pmatrix} \tau^y \\ \mu_t^{spf,y} \\ \tau_t^u \\ \tau_t^e \\ \tau_t^{oil} \\ \tau_t^\pi \\ \mu_t^{spf,\pi} \\ \mu_t^{uom,\pi} \end{pmatrix}, \quad (9)$$

Moreover, differently from the quarterly model, these models do not include core inflation, but incorporate SPF forecasts for GDP to stabilise the output trend that may otherwise show an excess of volatility due to the quarterly nature of the observations on GDP. This implies an updated assumption on output potential:

Assumption 1' (Output potential). The output potential is the common trend between real GDP and expected real GDP.

A crucial component of these models is the aggregation of variables at monthly frequencies into quarterly-frequency indicators. In line with the nowcasting literature, we model the data by assuming that low-frequency (quarterly) series have a high-frequency (monthly) representation and estimate it at monthly frequency. The model treats quarterly indicators as monthly data with missing observations for which the model has to deliver estimates, employing a set of restrictions similar to those proposed in [Mariano and Murasawa \(2003\)](#).

We specify the aggregation procedures in the following assumptions.

Assumption 9 (Aggregation rules for the CBO cycle and real GDP). For the CBO cycle and real GDP, the quarterly data are linked to latent monthly figures (denoted with the use of a tilde) as

$$cycle_t^{cbo} = (1 + L + L^2) \widetilde{cycle}_t^{cbo},$$

$$y_t = (1 + L + L^2) \tilde{y}_t,$$

where

$$\widetilde{cycle}_t^{cbo} = \tilde{y}_t - \tau_t^y, \quad (10)$$

$$\tilde{y}_t = \psi_t^{gap} + \psi_{1,t} + \tau_t^y, \quad (11)$$

for any t . This aggregation approach is standard and used in several papers including [Giannone et al. \(2008\)](#) and [Bańbura and Modugno \(2014\)](#).

Assumption 10 (Aggregation rules for expectational data). Professionals' expectations for real GDP and inflation are aggregated in different ways. At any t we have⁵

$$F_t^{spf} y_{t+12} = F_t^{spf} \tilde{y}_{t+12} + F_t^{spf} \tilde{y}_{t+11} + F_t^{spf} \tilde{y}_{t+10},$$

since $F_t^{spf} y_{t+12}$ is quarterly and follows equivalent aggregation rules to the ones in [Assumption 9](#). In other words, we link the professional expectation $F_t^{spf} y_{t+12}$ with implied predictions for monthly real GDP figures computed with the same conditioning set (i.e., the one available at time t).

The model needs to enforce a mixed-frequency aggregation rule only for the real GDP professional expectations.⁶ Not knowing the exact prediction rule followed by professional

⁵We use $F_t x_{t+h}$ to indicate survey expectations at time t for a variable x_{t+h} to distinguish them from mathematical expectations, $E_t x_{t+h}$.

⁶SPF data for inflation are produced at monthly frequency. Therefore, we simply consider $F_t^{spf} \pi_{t+12}$ as the end-of-month one-year ahead forecast for inflation.

forecasters, we cannot recover the expectations for the latent monthly output figures in their entirety. However, we know that the persistent component of these expectations should be linked to the trends of real GDP. For simplicity, we assume them to be the same and, thus, enforce an aggregation rule according to which the trend of $F_t^{spf} y_{t+12}$ is considered as $(1 + L + L^2)\tau_{t+12}^y$. Since the trend is a random walk with drift, the trend can also be written as the sum of $3\tau_t^y$ plus a time-invariant drift. The time-invariant drift, denoted as $\mu^{spf,y} \equiv \mu_t^{spf,y}$ for every t , is estimated and so are the loadings.

2.3 Bayesian estimation

The model can be cast in a linear state-space form and estimated with Bayesian techniques, employing an Adaptive Metropolis-Within-Gibbs algorithm (details are provided in the Online Appendix). We adopt the simulation smoother of [Durbin and Koopman \(2002\)](#) along with the [Jarociński \(2015\)](#)'s modification to condition our estimates of cycles and trends on the full sample.

Data of each variable are normalised by dividing them by the standard deviation of their first differences.⁷ To deal with missing observations, we employ a Kalman filter approach (see, as a reference, the discussion in [Shumway and Stoffer, 1982](#)), and reconstruct the data on the basis of the information available at each point in time.

3 Trends and gaps in the US economy

How do semi-structural models read the inflation dynamics and the business cycle fluctuations in the US economy? We start by summarising the results from the quarterly model (for an extensive discussion, see [Hasenzagl et al., 2022a](#)), to then focus on the mixed frequency models. While we use fully revised data in this section, the following section provides a real-time appraisal of the two models' performances.

⁷As discussed in [Hasenzagl et al. \(2022a\)](#) this normalisation gives set data on a similar scale and provides better mixing in the Metropolis algorithm.

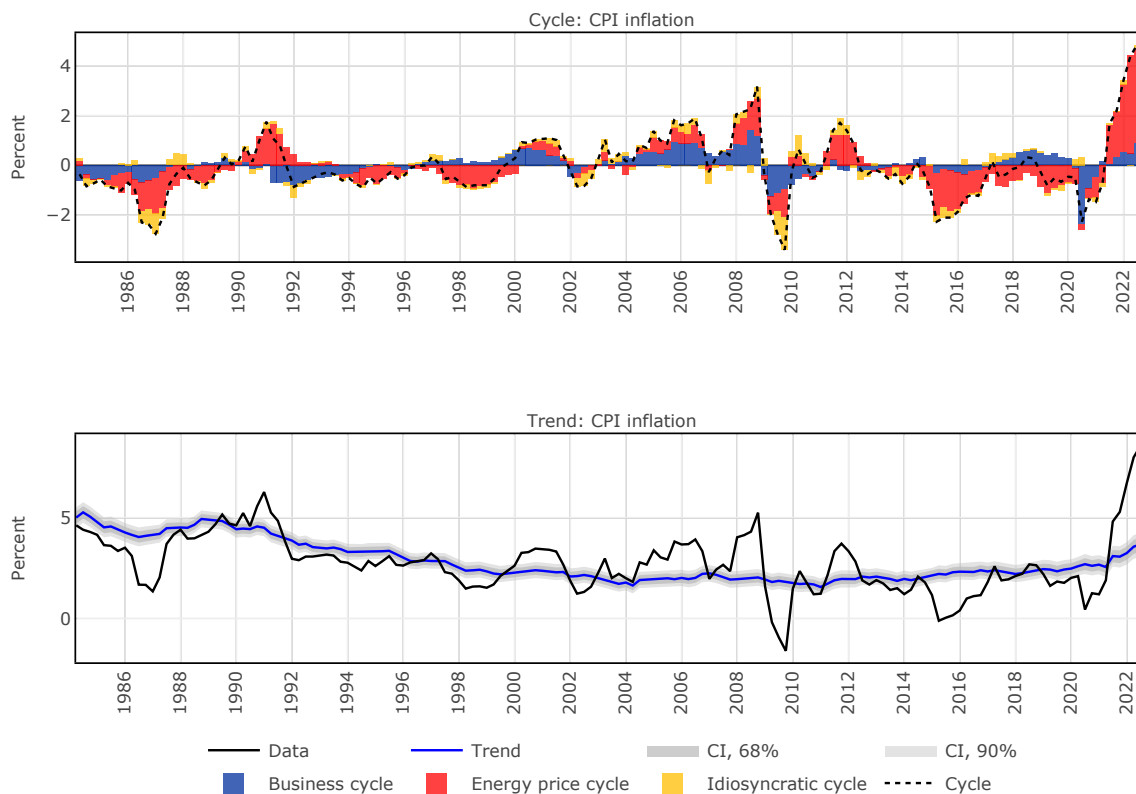


Figure 1: Top: Decomposition of the cycle of CPI inflation into common (in blue and red) and independent (in yellow) components, as estimated by the model in [Hasenzagl et al. \(2022a\)](#). Bottom: Trend of CPI inflation (in blue), with relative coverage intervals at 68% coverage (dark shade) and 90% coverage (light shade), as estimated by the model in [Hasenzagl et al. \(2022a\)](#).

3.1 Quarterly model

The quarterly model provides a reading of inflation dynamics (Figure 1) as due to three main components: (i) the inflation trend, in the lower panel, that can be seen as the long-term inflation expectations; (ii) a Phillips curve component (the blue area in the upper panel) that reflects the slack in the economy as captured by the output gap into the price dynamics; (iii) the effects of oil price disturbances (the red area) that can move inflation away from the nominal-real relationship captured by the Phillips curve.

The model is estimated on the sample from Q1-1984 to Q3-2022. Let us remark on three results. First, the inflation trend has come down from higher levels in the early eighties to become anchored at the 2% target before showing a slight increase in the post-pandemic period. Second, a well-identified, steep, and stable reduced-form Phillips

curve relationship captures a cyclical component of CPI inflation with maximum power at around eight years periodicity. Third, energy price disturbances that happen at a higher frequency than the Phillips curve fluctuations, have a large impact on CPI inflation often overpowering the Phillips curve component. In the presence of large oil price shocks, this component may dominate and cloud the signal of cyclical inflation.

It is worth observing that the Phillips curve is steeper than some of the recent estimates in the literature. This is because the model can estimate a cleaner output gap measure by separating it from the effects of energy price disturbances, which are independent of local real economic conditions. The effects of energy price disturbances can be seen as confounding factors in other studies, reducing the apparent correlation between the slack in the economy and price pressure.

3.2 Mixed-frequency models

We now dig deeper into how our semi-structural model assesses inflation and business cycle dynamics by focusing on the two mixed-frequency models. The mixed frequency models are estimated on a sample from January 1985 to December 2022. To avoid a bias on the inference of the cycle's periodicity caused by the COVID-19 pandemic we consider the year 2020 as missing when we estimate the model parameters. The undisciplined model provides a similar overall reading of the US economy as the quarterly model while presenting some key differences from the tracking model. [Figure 2](#) compares the two sets of results, on the sample from January 1985 to September 2022, and reports the cyclical components of all variables for the two models: (i) the output gap and the business cycle (blue), (ii) the energy price component (red), and a residual idiosyncratic component reflecting measurement error and unmodelled components.

The tracking model – in line with the assessment of the CBO onto which it is geared – estimates an output gap that shows significant contractions in the cyclical component of real GDP after the early 1990s recession – i.e., the dot-com bubble and the Great

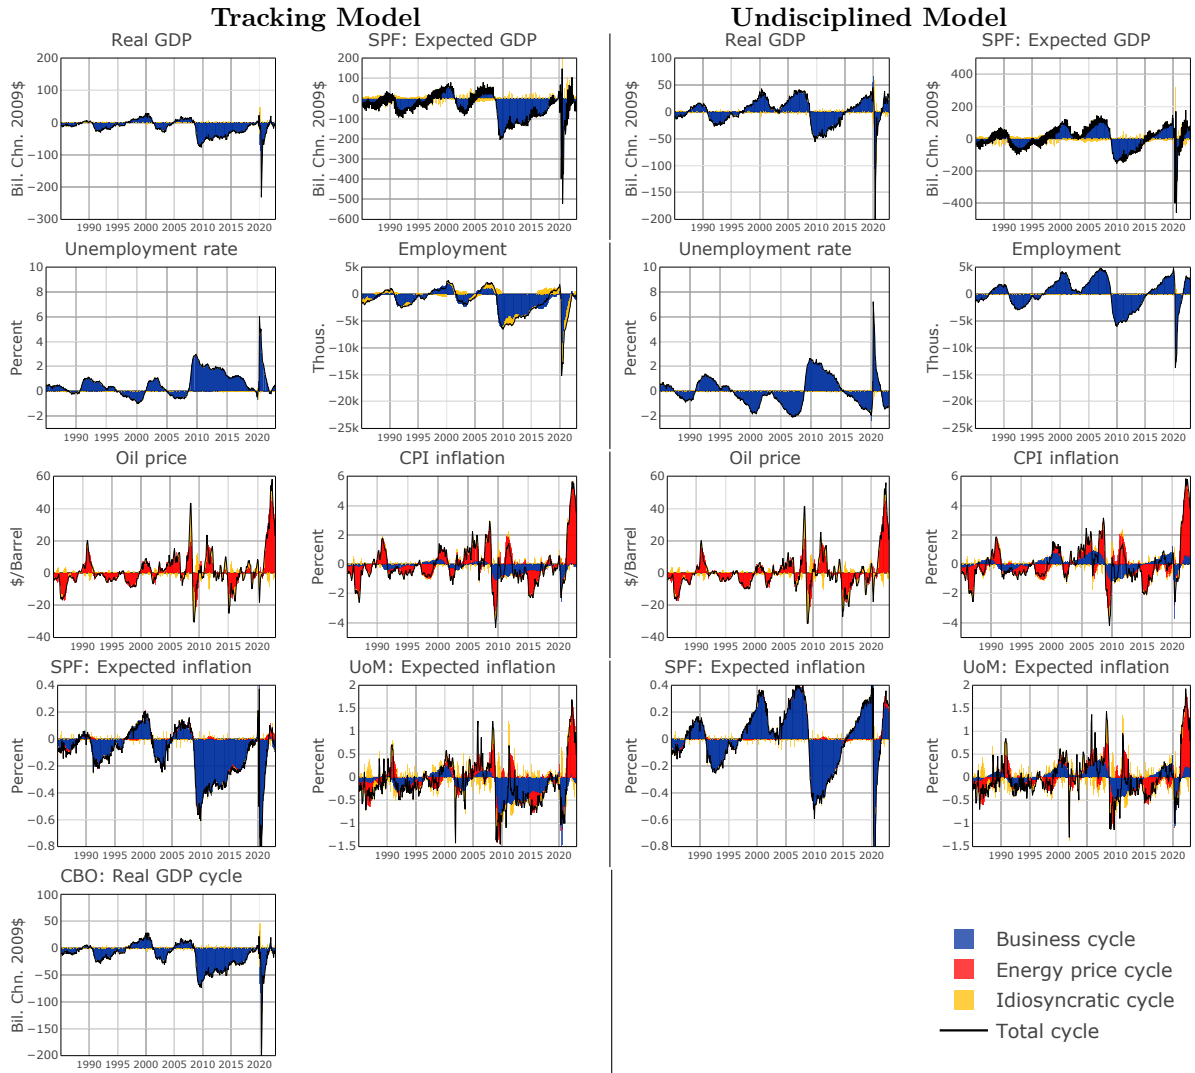


Figure 2: Historical decomposition of the stationary components of all the variables in the tracking model (left) and the undisciplined model (right).

Recession (see Figure 2). This contrasts with the assessment of the undisciplined model, which shows larger expansions in the late 1990s and early 2000s that culminated in the dot-com crisis and the Great Recession. The idiosyncratic component plays a small role in both specifications, which implies that differences across models in assessing the cyclical component reflect differences in the evaluation of potential output. These differences are then reflected in the cycles of employment and unemployment, which are linked to the output gap by Okun’s law and in the Phillips curve part of the inflation cycle. For example, the tracking model reads negative or neutral inflation pressure from the

real economy before the financial crisis while the undisciplined model identifies positive pressures.

The correlation between the unemployment and the inflation cycles is similar across models, -0.45 for the undisciplined and -0.39 for the tracking model.⁸ However, the common cycle between inflation and the real economy is masked by the highly volatile energy component unrelated to domestic business cycle fluctuations. Consequently, the overall cyclical part of inflation is similar across models.

As observed, the differences in the measure of the output gap between models are due to differences in the estimated potential, which is the gap between output and its trend. Estimates of the trends for all the variables are reported in [Figure 3](#), where they are plotted against their associated observable variables. It can be easily seen that the undisciplined model fits an output gap that fluctuates almost symmetrically around the trend, as would be the case in a standard Neoclassical or New Keynesian macroeconomic model. Conversely, in the model informed by the CBO, potential output is above GDP most of the time, and recessions appear as shortfalls against this higher level. The undisciplined model attributes a larger part of the output variance to the trend, interpreting the slowdown since 2001, especially since 2008, as a change to potential output rather than as cyclical fluctuations. Consequently, the estimate of the NAIRU in the second half of the sample is higher. Not surprisingly, trend inflation is almost identical across the two models.

⁸In comparison, the correlation of the raw unemployment and inflation data is -0.14.

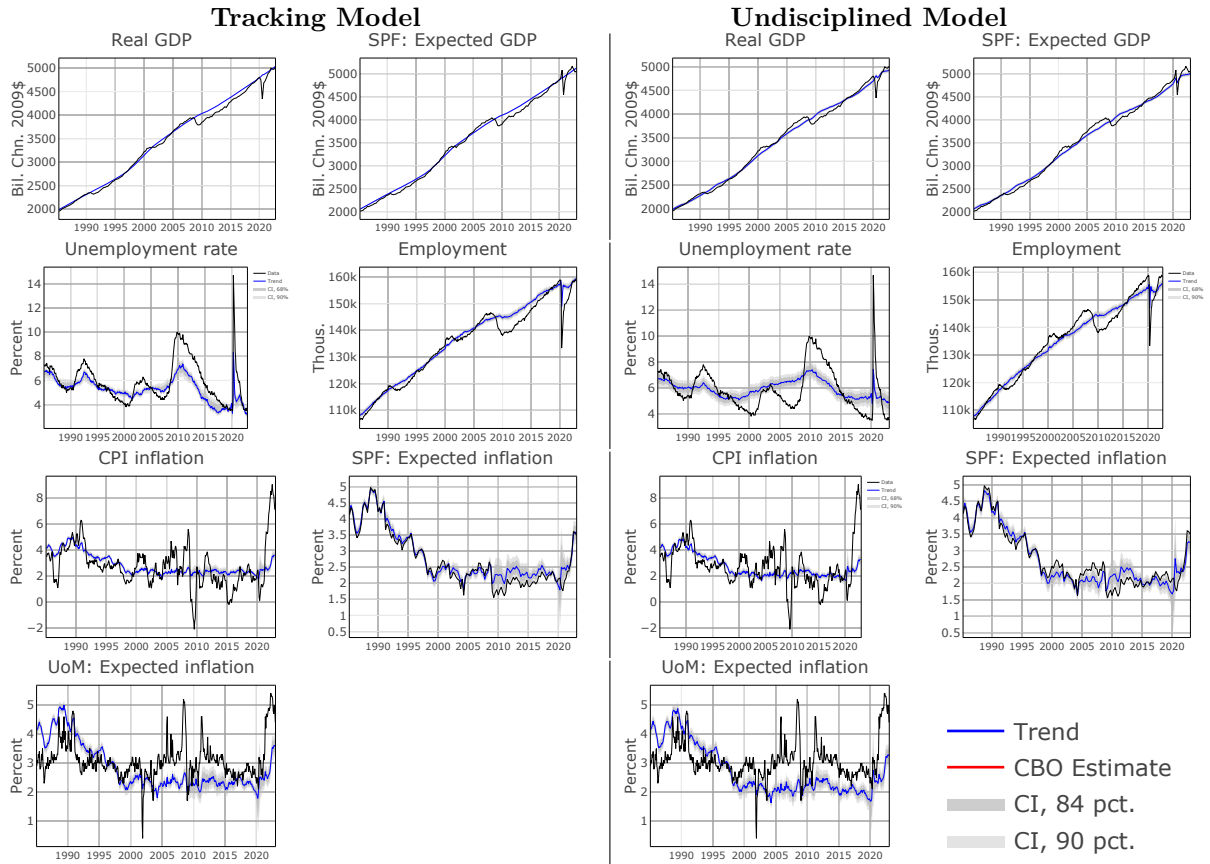


Figure 3: Trends for all the variables in the tracking model (left) and the undisciplined model (right), along with 84% 90% posterior coverage bands and The charts also report the CBO estimates for output potential and NAIRU.

4 Now-casting and forecasting

We now provide an appraisal of the forecasting performances of the semi-structural models out-of-sample and with real-time data. To this end, we first provide a pseudo out-of-the-sample appraisal of the quarterly model and we benchmark its performances against other standard models. We then assess the forecasting of the mixed-frequency models and the evolution of their now-cast over the data releases, on a fully real-time dataset.

4.1 Quarterly model: forecasting

Table 2 reports the out-of-sample performances on horizons from one quarter to two years for the quarterly model (i) on the pre-COVID sample (as reported in the online appendix of Hasenzagl et al., 2022a), and (ii) on extended sample that includes the pandemic

and more recent data, as well as for (iii) a BVAR where priors are set as in [Giannone et al. \(2015\)](#), and (iv) the univariate unobserved components IMA(1,1) with stochastic volatility (UC-SV) model proposed by [Stock and Watson \(2007\)](#) and acknowledged to be tough benchmark for inflation forecast.

Results indicate that the semi-structural model performs well at forecasting inflation, especially at the median horizon. The good performance of the model at the medium horizon, especially for inflation, is related to its ability to capture the business cycle component and a steep Philips curve. This allows the model to improve over a random walk benchmark or the UC-SV model.

Our results also show that the semi-structural model has a similar performance when we consider the sample including Covid-19 than when we exclude it. The analysis for the pre-COVID sample [Hasenzagl et al. \(2022b\)](#) showed that, at the medium horizon, the model had also a better performance than the BVAR benchmark.

The advantage of the model, however, was not as pronounced and generalized to all variables at all horizons as shown here for the longer sample. We explain this advantage by the semi-structural model's flexibility in estimating unit root trends and capturing the oil shocks by modelling explicitly an oil cycle. The poor performance of the BVAR can be explained by the difficulty of a fixed parameter model to cope with the extreme observations of the pandemic period (see [Lenza and Primiceri, 2022](#) for a discussion).

4.2 Mixed frequency: forecasting

We next compare the out-of-sample performance of the undisciplined model and the tracking model by reporting the average mean squared error of forecasts up to three years ahead, in real-time evaluation (see [Figure 4](#)). To this aim, we construct a set of real-time data vintages from [Federal Reserve Bank of St. Louis \(2021\)](#) and [Federal Reserve Bank of Philadelphia \(2021\)](#) starting on January 1, 2005, and using the prior 20 years as our pre-sample. We iterate over the real-time release calendar of the variables

Table 2: Quarterly out-of-sample evaluation

Horizon	Variable	Pre COVID- 19	Up-to- date fig- ures	Up-to- date BVAR	Up-to- date UCSV
h=1	Real GDP	1.00	1.00	1.50	-
	Employment	0.94	0.99	1.44	-
	Unemployment rate	0.82	0.98	1.51	-
	Oil price	1.06	1.05	1.12	-
	CPI Inflation	0.97	0.97	1.02	1.00
	Core CPI Inflation	1.00	0.97	1.14	1.00
	UOM: Expected inflation	1.03	1.04	1.09	-
	SPF: Expected CPI	1.00	1.02	1.30	-
h=2	Real GDP	1.02	1.00	1.77	-
	Employment	0.95	0.99	1.78	-
	Unemployment rate	0.80	0.96	1.92	-
	Oil price	1.08	1.09	1.19	-
	CPI Inflation	0.95	0.95	1.13	0.99
	Core CPI Inflation	0.95	0.93	1.27	0.99
	UOM: Expected inflation	1.01	1.02	1.21	-
	SPF: Expected CPI	0.97	1.01	1.72	-
h=4	Real GDP	1.04	0.98	1.91	-
	Employment	0.99	0.96	2.12	-
	Unemployment rate	0.81	0.89	2.27	-
	Oil price	1.12	1.11	1.21	-
	CPI Inflation	0.95	0.96	1.34	0.97
	Core CPI Inflation	0.89	0.91	1.69	0.96
	UOM: Expected inflation	1.11	1.09	1.44	-
	SPF: Expected CPI	0.91	0.95	2.11	-
h=8	Real GDP	1.11	0.99	1.52	-
	Employment	1.07	0.93	1.83	-
	Unemployment rate	0.81	0.80	1.83	-
	Oil price	1.10	1.05	1.22	-
	CPI Inflation	0.85	0.88	1.10	0.96
	Core CPI Inflation	0.83	0.92	1.80	0.94
	UOM: Expected inflation	1.02	0.99	1.28	-
	SPF: Expected CPI	0.86	0.86	1.77	-

Note: This table shows the RMSEs relative to the random walk with drift. The second column reports the results listed in [Hasenzagl et al. \(2022a\)](#), which were computed pre-COVID-19. The third column includes up-to-date results from the same model. The following columns include relevant and up-to-date benchmarks. The period from Q1 1984 to Q4 1998 is employed as the pre-sample, while the evaluation sample starts in Q1 1999 and ends either in 2018 (second column) or 2022 (third column). ‘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.

in the model and update our estimates of the trends and gaps at each new data release.

We also project the trends and gaps forward and use them to forecast the variables in

the model. To decrease the computational burden, we re-estimate all model parameters at the first release of each year, and then keep them fixed for the remainder of the year.

The undisciplined model outperforms the tracking model for output and labour market variables at the horizon beyond eight months, confirming that the model's ability to capture comovement at business cycle frequency gives it an advantage for the medium-term forecast. This suggests that the CBO output gap is not as strongly correlated with labour market variables at business cycle frequencies as the outgap model estimated by the statistical model.

The two models, however, show almost identical performance for inflation which suggests that, since the business cycle component of inflation is relatively small and dwarfed by the trend and the oil cycle, differences in output gap measures do not affect the out-of-sample forecasting performance of inflation. Those differences, however, matter for forecasting SPF expected inflation, giving an advantage to the undisciplined model beyond five months.

4.3 Mixed frequency: now-casting inflation

We complete our discussion on the forecasting ability of semi-structural models by presenting the evolution of the root mean squared forecast error (RMSFE) of the undisciplined model for the nowcast of inflation, in relation to the real-time flow of data releases.

Figure 5 reports the average RMSFE for inflation for the nowcasting, forecasting, and backcasting periods. The now-casting period is defined as the month that the specific inflation release refers to, while the forecasting period corresponds to the month preceding the nowcasting period, and the backcasting period to the month following the inflation release, prior to the publication of the inflation figure. The arrangement of the bars from left to right proceeds with the publication of the variables, where the leftmost bar represents a variable release that, on average, is published furthest in advance of the

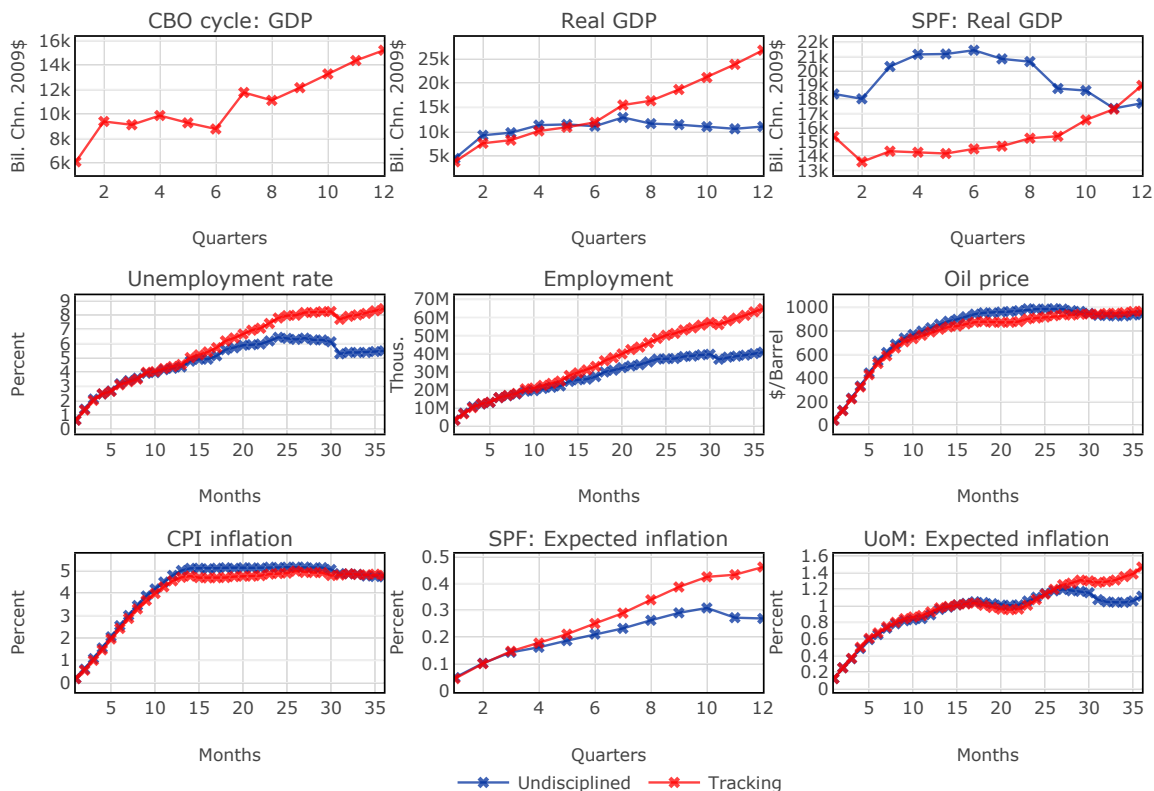


Figure 4: The chart reports the average mean squared error of the undisciplined model (in blue) and the tracking model (in red). The out-of-sample evaluation starts in January 2005 and ends in September 2020.

inflation release date. Conversely, the rightmost bar represents the variable that, on average, is released closest to the date of the inflation release.

Table 3 helps read the results. It shows the average publication lag of each series with respect to the reference period expressed in days. CPI inflation is published on the 15th of each month and refers to the previous month. At that date, we already have information on the labour market variables for the previous month and expected inflation but the only variable referring to the current month is oil. Since oil is available at a higher frequency but we aggregate it monthly, we arbitrarily attribute the release to the end of the month.

The bars show an overall improvement over the forecast horizon, with forecast errors generally declining with the new data releases. A few observations are in order. First, real variables, including those which are more timely than the inflation release such as

Table 3: Avg. publication lags vs. reference period.

Indicator	Avg. lag (days)
Real GDP	29
SPF: GDP	-33
Unemp. rate	5
Employment	5
Oil price	0
CPI infl.	15
SPF: Exp. infl.	-33
UoM: Exp. infl.	-3

Note: Table shows avg. publication lag in days post the reference period’s end. Negative lags imply pre-release forecasts.

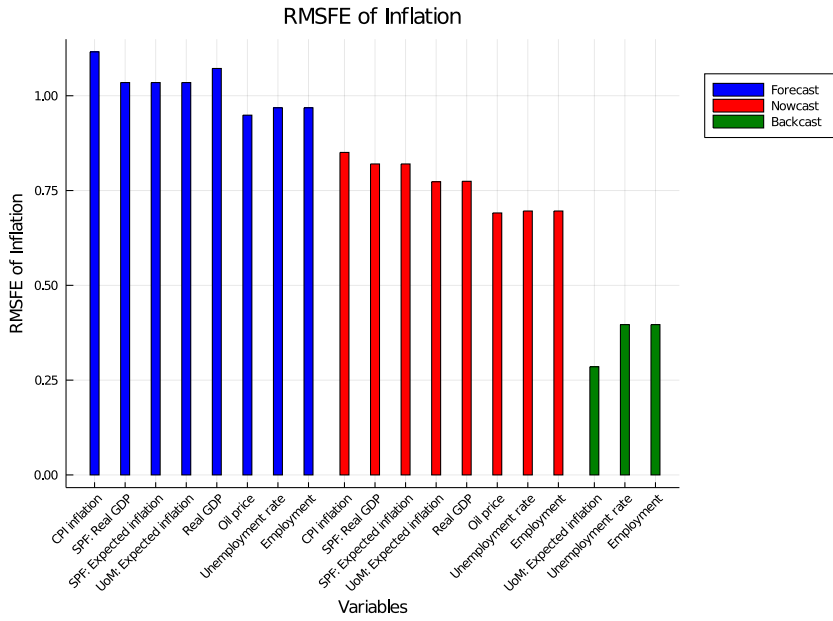


Figure 5: The chart presents the average Root Mean Square Forecast Error (RMSFE) of inflation across three distinct periods: nowcasting, forecasting, and backcasting. The nowcasting period corresponds to the month that the specific inflation release refers to, the forecasting period corresponds to the month preceding the nowcasting period, and the backcasting period corresponds to the month following the inflation release, prior to the publication of the inflation figure. The arrangement of the bars in the chart is deliberate, with the leftmost bar represents a variable release that, on average, is published furthest in advance of the inflation release date. Conversely, the rightmost bar represents the variable that, on average, is released closest to the date of the inflation release.

employment and unemployment do not help. If anything, they blur the signal in the backcasting period. The oil price release, on the other hand, has a larger impact as compared to other variables. This is not surprising since it is the most timely variable. Surprisingly, however, is that inflation expectations do not have any impact.

Since the performances of now-casting models are driven by the availability of timely indicators, these results were to be expected and suggest that using this method for now-casting GDP, which is available only quarterly, would be more promising. Timely monthly data help the now-cast of GDP but the timeliness advantage of monthly series for now-casting a monthly series such as inflation is smaller. The relevance of the release of oil prices, however, suggests that an extension to a weekly-monthly mixed frequency model might be promising.

4.4 Mixed frequency models: now-casting the output gap

We conclude this section by assessing the stability in real time of the estimates of business cycle components (Figure 6).⁹ Comparing the estimates for the two mixed frequency models in real-time allows us to better gauge the role of data releases.

For the tracking model, the sum of the business cycle and idiosyncratic components equals the CBO's output gap estimates by construction. In the pre-sample period, before 2005, the different lines reflect the instability of in-sample estimates, while from 2005 onwards, they are also affected by data revisions. In the tracking model, revisions of the output gap can be due to revisions of the CBO estimates themselves, as reported by the CBO (or the model's revisions of the CBO gap on the forecasting horizon).

We compute two statistics to understand the relative stability of output gap measures across models. First, we calculate the standard deviation of the output gap and potential output across all vintages for each reference period. We then compute our first statistics by averaging these standard deviations across all reference periods. For the undisciplined model, this statistic is 0.51 for the output gap and 6.37 for potential output compared to 0.58 and 7.88 for the tracking model. The second statistic is the average of the maximum absolute value of revisions for each reference month. For the undisciplined model, this

⁹The estimates plotted in the figure are defined as the ratio of the sum of the business cycle component (the 'true' output gap) plus and the GDP idiosyncratic cycle, over the GDP trend.

measure is 1.96 for the output gap and 15.99 for potential output, compared to 2.33 and 15.31 for the tracking model.




Figure 6: The chart reports the real-time estimates of the output gap from the undisciplined (left) and the tracking model (right). The out-of-sample evaluation starts in January 2005 and ends in September 2020.

The results suggest that, on average, the estimated output gap is more stable in the undisciplined model than in the tracking model. The same is true for potential output, although the difference across models is smaller. This leads to the conclusion that the difference in variability of the gap measures across models is due to the CBO’s larger judgmental revisions in the output gap and not simply the data revisions themselves.

The COVID-19 pandemic period also provides a good illustration of the framework’s flexibility. During that period, the estimate of potential output (but for a slight uptick) tracks developments smoothly despite the enormous size of the economic shock and its unprecedented nature (see [Figure 6](#)).

5 Concluding comments

This chapter has explored the performance of medium-scale semi-structural time series models for describing trends and cyclical components of inflations in a multivariate setting. Models in this class are designed to capture the dynamic correlations between inflation, labour market, output and expectation as captured by survey data. They aim to describe complex multivariate dynamic relationships in a parsimonious way and help

in forecasting at medium-term horizons. By featuring variable trends and cycles defined over different frequency ranges, these models are ideally suited for describing inflation dynamics and assessing the relevance of multivariate information for forecasting inflation, a topic that has been debated over the years and remains controversial.

In this approach, the cyclical and trend components of the macroeconomic series are identified using minimal economic restrictions coming from economic theory. This provides the identified structural components with a clear interpretation. From a statistical point of view, imposing a structure in the form of orthogonal common trends and cycles helps obtain an estimate of the Phillips curve component which is cleaned by low-frequency movements and higher-frequency volatility driven by oil shocks. As a signal extraction tool, this class of models presents some advantages with respect to VAR-based estimates of the Phillips curve with or without stochastic trends and to factor models with stochastic trends.

At the descriptive level, our analysis points to a steeper Phillips curve, than what has been estimated in the literature, and shows that this helps forecast inflation at the medium-term horizon. We also show that a benchmark model can easily be adapted to include mixed frequency data and missing observations and therefore used for now-casting as well as forecasting.

SUPPLEMENTARY MATERIALS

The Online Appendix provides details on the dataset, and the estimation of the models described in the manuscript, as well as additional results and robustness exercises. (pdf file)

References

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas, “Business-Cycle Anatomy,” *American Economic Review*, October 2020, *110* (10), 3030–70.

- Ascari, Guido and Luca Fosso**, “The international dimension of trend inflation,” *Journal of International Economics*, 2024, p. 103896.
- Atkeson, Andrew and Lee E. Ohanian**, “Are Phillips curves useful for forecasting inflation?,” *Quarterly Review*, 2001, (Win), 2–11.
- Bañbura, Marta and Michele Modugno**, “Maximum likelihood estimation of factor models on datasets with arbitrary pattern of missing data,” *Journal of Applied Econometrics*, 2014, 29 (1), 133–160.
- Bañbura, Marta, Domenico Giannone, and Lucrezia Reichlin**, “193 Nowcasting,” in “The Oxford Handbook of Economic Forecasting,” Oxford University Press, 07 2011.
- Barigozzi, Matteo and Matteo Luciani**, “Measuring the output gap using large datasets,” *Review of Economics and Statistics*, 2023, 105 (6), 1500–1514.
- 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.
- Bianchi, Francesco, Gianni Nicoletti, and Dongho Song**, “Inflation and Real Activity over the Business Cycle,” Working Paper, Johns Hopkins University October 2022.
- Burns, Arthur F. and Wesley C. Mitchell**, *Measuring Business Cycles* number burn46-1. In ‘NBER Books.’, National Bureau of Economic Research, Inc, July 1946.
- Coibion, Olivier and Yuriy Gorodnichenko**, “Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics*, January 2015, 7 (1), 197–232.

- Del Negro, Marco, Domenico Giannone, Marc P. Giannoni, and Andrea Tambalotti**, “Safety, Liquidity, and the Natural Rate of Interest,” *Brookings Papers on Economic Activity*, 2017, 48 (1 (Spring)), 235–316.
- Durbin, James and Siem Jan Koopman**, “A simple and efficient simulation smoother for state space time series analysis,” *Biometrika*, 2002, pp. 603–615.
- Federal Reserve Bank of Philadelphia**, “Survey of Professional Forecasters,” <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/survey-of-professional-forecasters>, 2021. Accessed Dec. 1, 2021.
- Federal Reserve Bank of St. Louis**, “Archival FRED,” <https://alfred.stlouisfed.org/>, 2021. Accessed Dec. 1, 2021.
- Forni, Mario and Lucrezia Reichlin**, “Let’s Get Real: A Factor Analytical Approach to Disaggregated Business Cycle Dynamics,” *The Review of Economic Studies*, July 1998, 65, 453–473.
- , **Marc Hallin, Marco Lippi, and Lucrezia Reichlin**, “The generalized dynamic-factor model: Identification and estimation,” *Review of Economics and Statistics*, 2000, 82 (4), 540–554.
- Giannone, Domenico, Lucrezia Reichlin, and David Small**, “Nowcasting: The real-time informational content of macroeconomic data,” *Journal of Monetary Economics*, May 2008, 55 (4), 665–676.
- , **Michele Lenza, and Giorgio E Primiceri**, “Prior selection for vector autoregressions,” *Review of Economics and Statistics*, 2015, 97 (2), 436–451.
- Gordon, Robert J.**, “U.S. Inflation, Labor’s Share, and the Natural Rate of Unemployment,” in “in” 1990.

- Harvey, A. C. and A. Jaeger**, “Detrending, stylized facts and the business cycle,” *Journal of Applied Econometrics*, 1993, 8 (3), 231–247.
- Harvey, Andrew C**, “Trends and cycles in macroeconomic time series,” *Journal of Business & Economic Statistics*, 1985, 3 (3), 216–227.
- , *Forecasting, structural time series models and the Kalman filter*, Cambridge university press, 1990.
- Hasenzagl, Thomas, Filippo Pellegrino, Lucrezia Reichlin, and Giovanni Ricco**, “A Model of the Fed’s View on Inflation,” *The Review of Economics and Statistics*, 2022, 104 (4), 686–704.
- , – , – , **and** – , “Online Appendix for ‘A Model of the Fed’s View on Inflation’,” Supplement to “The Review of Economics and Statistics” 07 2022. Available online.
- Jarociński, Marek**, “A note on implementing the Durbin and Koopman simulation smoother,” *Computational Statistics & Data Analysis*, 2015, 91, 1–3.
- Jarociński, Marek and Michele Lenza**, “An Inflation-Predicting Measure of the Output Gap in the Euro Area,” *Journal of Money, Credit and Banking*, September 2018, 50 (6), 1189–1224.
- Lenza, Michele and Giorgio E. Primiceri**, “How to estimate a vector autoregression after March 2020,” *Journal of Applied Econometrics*, 2022, 37 (4), 688–699.
- Mariano, Roberto S and Yasutomo Murasawa**, “A new coincident index of business cycles based on monthly and quarterly series,” *Journal of applied Econometrics*, 2003, 18 (4), 427–443.
- Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H. Stock**, “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” *Journal of Economic Literature*, March 2014, 52 (1), 124–88.

Modugno, Michele, “Now-casting inflation using high frequency data,” *International Journal of Forecasting*, 2013, 29 (4), 664–675.

Sargent, Thomas J. and Christopher A. Sims, “Business cycle modeling without pretending to have too much a priori economic theory,” Technical Report 1977.

Shumway, Robert H and David S Stoffer, “An approach to time series smoothing and forecasting using the EM algorithm,” *Journal of time series analysis*, 1982, 3 (4), 253–264.

Stock, James and Mark W Watson, “Phillips Curve Inflation Forecasts,” in Little J Olivei G Fuhrer J, Kodrzycki Y, ed., *Understanding Inflation and the Implications for Monetary Policy*, Cambridge: MIT Press, 2009.

Stock, James H. and Mark W. Watson, “New Indexes of Coincident and Leading Economic Indicators,” in Plivier J Blanchard and Stanley Fisher, eds., *Macroeconomics Annual*, Vol. 4, M.I.T. Press, 1989.

Stock, James H and Mark W Watson, “Forecasting using principal components from a large number of predictors,” *Journal of the American Statistical Association*, 2002, 97 (460), 1167–1179.

— **and** —, “Why has US inflation become harder to forecast?,” *Journal of Money, Credit and banking*, 2007, 39 (s1), 3–33.