

Demand and supply shocks over the business cycle

Stylianos Zlatanov*

August 2025

First online version: June 2025

Abstract

This paper examines the drivers of U.S. business cycle fluctuations using a Trend–Cycle Bayesian VAR, motivated by evidence that a single “main business cycle” shock leaves much of inflation unexplained. The analysis splits demand into monetary and non-policy components and explicitly models cost-push and oil supply shocks. Applying a generalized “Max Share” procedure with sign restrictions to the stationary cyclical components, the results indicate that demand dominates real activity and drives a large share of nominal fluctuations at horizons that include the short run. At medium-run business cycle horizons (6–32 quarters), supply forces—cost-push and oil—often become pivotal for inflation (especially when identification targets output). Crucially, allowing for these multiple shocks explains the vast majority of cyclical variation in both output and inflation, closing the gap left by one-shock analyses. Overall, by disentangling multiple demand channels and explicitly modeling oil shocks, this framework offers a more precise understanding of the U.S. business cycle.

Keywords: Business cycles, Trend–Cycle VAR, demand, supply

JEL Classification: C32, E31, E32.

*King’s Business School, King’s College London. Email: stylianos.zlatanov@kcl.ac.uk. I am indebted to Francesca Monti and George Kapetanios for their invaluable guidance and support. For helpful comments and suggestions at different stages of this project, I thank Gert Peersman, Michele Lenza, Georgios Georgiadis and Aristeidis Raftapostolos. I also thank seminar participants at the Bank of England, Ghent University Workshop on Empirical Macroeconomics and Sailing the Macro Workshop. I welcome comments, including references to related papers I may have inadvertently overlooked. All errors are my own. Please do not cite without permission.

1 Introduction

What drives the business cycle? A recent strand of literature aims to shed light on the likely cause of U.S. business cycle fluctuations from an empirical perspective. Angeletos et al. (2020) concludes that a “main business cycle” shock can explain most of the cyclical variation in output, albeit showing a puzzling disconnection between the real and nominal sides of the economy. By contrast, Bianchi et al. (2025) show that this puzzle largely disappears once the slow-moving long run trends are properly accounted for. Both studies indicate that the “main business cycle” shock has demand-like characteristics and while Bianchi et al. (2025) find that this shock explains about 60% of the cyclical variation in output, a substantial portion of inflation fluctuations remains unexplained—despite evidence that, at business cycle horizons, inflation does indeed move in tandem with real activity in line with the New Keynesian framework.

Motivated by these findings, this paper shows that once we properly distinguish different types of demand shocks (i.e., monetary and non-policy demand shocks), and incorporate explicit cost-push and oil supply disturbances, nearly the entire business cycle variation in both real activity and nominal variables can be explained within a single Trend–Cycle VAR framework.

Methodologically, I build on the approach developed by Del Negro, Giannone, et al. (2017)—and recently employed by Bianchi et al. (2025)—to extract low-frequency trends from each variable and then identify the structural drivers over the business cycle. I extend their setup in two key aspects. First, I split aggregate demand into two distinct channels: a *monetary policy* shock capturing the direct influence of central bank actions, and a *non-policy demand* shock reflecting broader demand conditions. Second, I explicitly model oil prices to isolate *oil supply* shocks, separate from *cost-push* shocks. Accounting for oil shocks is essential: historical episodes such as the oil crises of the 1970s and early 1980s illustrate how major swings in global oil supply can profoundly affect both inflation and output—yet these movements are often misattributed to generic cost pressures or omitted.

Main findings Demand-side forces dominate real activity. Over horizons that include the short run (identification based on FEV *accumulated over* 1–32 quarters), demand explains about two-thirds of fluctuations in output and unemployment; focusing on the medium run (FEV *accumulated over* 6–32 quarters), demand’s share rises further, with non-policy demand the single largest contributor. On the nominal side, demand accounts for most movements in the policy rate, again led by non-policy demand. For inflation, demand is responsible for roughly two-thirds when short-run horizons are included, whereas at medium-run horizons supply shocks—cost-push and oil—become pivotal, especially under the output-targeted identification. Oil prices are predominantly driven by oil-supply shocks. The IRFs are theory-consistent: a monetary tightening lowers activity and inflation; non-policy demand lifts both and elicits policy tightening with mild “overshooting” in output; supply shocks are stagflationary and trigger little systematic policy response. Splitting demand (policy vs. non-policy) and

supply (cost-push vs. oil) largely closes the inflation “gap” left by one-shock analyses.

Related literature This paper contributes to the literature on the real–nominal linkages over the business cycle and to other studies focused on identifying the main structural drivers of the U.S. business cycle (Angeletos et al., 2020; Bianchi et al., 2025). It reconciles the large role that Bianchi et al. (2025) attribute to demand-like forces in real activity with the empirical importance of those same forces for price-setting decisions. Moreover, adding explicit supply disturbances provides a clearer account of how nominal and real variables comove.

Beyond these direct predecessors, two parallel recent contributions, Forni et al. (2024) and Granese (2024), revisit the Blanchard and Quah (1989) framework that views business cycles as driven by supply and demand shocks. Motivated by potential informational limitations in the smaller-scale VAR of Angeletos et al. (2020), both papers employ a large-dimensional Structural Dynamic Factor model using extensive U.S. time-series data. In contrast to these predominantly frequency-domain approaches that examine one demand and one supply shock only, this paper adopts a different path within the Trend–Cycle Bayesian VAR framework of Bianchi et al. (2025) and Del Negro, Giannone, et al. (2017), leveraging sign restrictions to distinguish multiple demand from supply forces that drive business cycle fluctuations.

A second strand of related literature involves the identification of shocks using the “Max Share” approach, originally proposed by Faust (1998), later extended by Uhlig (2003), and more recently generalized by Carriero and Volpicella (2024) to allow for multiple shocks simultaneously. My identification procedure relies on the Generalized Max Share approach developed in the latter.

The remainder of the paper proceeds as follows. Section 2 introduces the extended trend–cycle model, discusses the data, and outlines the estimation procedure. Section 3 details the identification strategy. Section 4 presents the main empirical findings, and Section 5 concludes with a discussion.

2 Model

In this section, I present the econometric framework that distinguishes long-term trends from business cycle fluctuations in real and nominal variables. The approach builds on the Trend–Cycle Bayesian VAR of Del Negro, Giannone, et al. (2017), which explicitly decomposes each observable series into a nonstationary trend and a stationary cyclical component. Such a separation proves crucial: as highlighted by Bianchi et al. (2025), relying on a conventional VAR that does not properly filter out low-frequency movements risks conflating persistent trends with short-run dynamics, especially over periods featuring structural breaks.

My baseline specification follows the setup in Bianchi et al. (2025) but makes one key extension: I include oil prices among the observables in order to explicitly capture oil supply shocks. Neglecting oil market dynamics could otherwise lead to misattributions

of certain supply- or demand-driven fluctuations, since large swings in oil prices often affect both output and inflation in ways that differ from more general cost-push or demand shocks.

Formally, each observable z_t is decomposed into a trend component τ_t evolving via a unit-root process and a cyclical component c_t governed by a stationary VAR. As shown in Section 3, once these trends are extracted, I can focus on identifying multiple structural shocks over business cycle horizons directly from the cyclical residuals of the VAR.

The system of variables included in the model is as follows: Real GDP per capita growth, denoted by g_t , is modeled as

$$g_t = \tau_{g,t} + (c_{y,t} - c_{y,t-1}),$$

where $\tau_{g,t}$ represents the persistent trend component of GDP growth. The term $(c_{y,t} - c_{y,t-1})$ captures cyclical deviations around that trend, allowing the model to identify the cycle in the *level* of real GDP, $c_{y,t}$, which is the relevant variable for business cycle analysis. Similarly, the unemployment rate u_t and the one-year-ahead unemployment expectations are decomposed as

$$u_t = \tau_{u,t} + c_{u,t}, \quad u_t^{e,1y} = \tau_{u,t} + c_{u,t}^{e,1y}$$

Next, the effective federal funds rate f_t is specified as

$$f_t = (\tau_{r,t} + \tau_{\pi,t}) + c_{f,t},$$

with the trend component given by the sum of the natural rate of interest and trend inflation, and the stationary component $c_{f,t}$ capturing short-run deviations. In other words, consistent with the literature, I assume that the Fisher relation holds in the long run. The inflation rate π_t is expressed as

$$\pi_t = \tau_{\pi,t} + c_{\pi,t}$$

Inflation expectations at one and ten years ahead are modeled as

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^e, \quad \pi_t^{e,10y} = \tau_{\pi,t} + \delta c_{\pi,t}^e + \eta_{\pi,t}^{e,10y},$$

so that these measures share the same underlying inflation trend with realized inflation. They also share a cyclical component which loads with $\delta < 1$ for the longer horizon. For long run inflation expectations, we also allow for an idiosyncratic measurement error. Finally, oil price growth o_t is represented by

$$o_t = \tau_{o,t} + c_{o,t},$$

where $\tau_{o,t}$ is the trend growth rate of oil prices.

For notational convenience, the observed variables are stacked in the vector

$$z_t = (g_t, u_t, u_t^{e,1y}, f_t, \pi_t, \pi_t^{e,1y}, \pi_t^{e,10y}, o_t)'$$

and similarly define the corresponding trend and cyclical component vectors as

$$\tau_t = (\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}, \tau_{o,t})', \quad c_t = (c_{y,t}, c_{u,t}, c_{u,t}^{e,1y}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^e, \varepsilon_{c,o,t})'.$$

The measurement error is collected in $\eta_t = (\eta_{\pi,t}^{e,10y})'$ and reduced form errors are collected in

$$\varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}, \varepsilon_{\tau,o,t}\}', \quad \varepsilon_{c,t} = \{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,u,t}^{e,1y}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}, \varepsilon_{c,\pi,t}^e, \varepsilon_{c,o,t}\}'.$$

Then, the dynamics of the latent state components can be written as a system of equations as follows:

$$\begin{aligned} \tau_t &= \tau_{t-1} + \varepsilon_{\tau,t}, \\ c_t &= \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + \dots + \Phi_p c_{t-p} + \varepsilon_{c,t}, \\ \eta_t &= \varepsilon_{\eta,t}, \end{aligned} \tag{1}$$

where τ_t , c_t , and η_t denote, respectively, the trend, cyclical, and measurement error components; the innovations $\varepsilon_{\tau,t}$, $\varepsilon_{c,t}$, and $\varepsilon_{\eta,t}$ are assumed to be mutually independent with

$$\varepsilon_t = \begin{pmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \\ \varepsilon_{\eta,t} \end{pmatrix} \sim \mathcal{N}\left(0, \begin{pmatrix} \Sigma_{\tau} & 0 & 0 \\ 0 & \Sigma_c & 0 \\ 0 & 0 & \Sigma_{\eta} \end{pmatrix}\right).$$

Here, Σ_{τ} , Σ_c , and Σ_{η} are positive definite matrices.

State-space representation In order to estimate the system above, one can express it in state-space form by linking the observed vector to the latent states. Let the observed variables be stacked in the vector z_t , it can be expressed as

$$z_t = \Lambda x_t = \Lambda_{\tau} x_{\tau,t} + \Lambda_c x_{c,t} + \Lambda_{\eta} x_{\eta,t} \tag{2}$$

where $x_t = \{x_{\tau,t}, x_{c,t}, x_{\eta,t}\}'$, $x_{\tau,t} = \tau_t$, $x_{c,t} = \{c_t, c_{t-1}, \dots, c_{t-(p-1)}\}'$, $x_{\eta,t} = \{\eta_t\}'$, and p denotes the lags of the stationary cyclical components. The $n \times n_{\tau}$ matrix Λ_{τ} captures $(n - n_{\tau})$ cointegrating relationships, while $\Lambda_c = [\Lambda_{c,0}, \dots, \Lambda_{c,p-1}]$ and Λ_{η} maps the measurement error state into the observed variables.

The vector of states x_t evolves based on the following state-transition equation:

$$x_t = \Phi x_{t-1} + \mathcal{R} \varepsilon_t \tag{3}$$

The initial conditions of the trend and cyclical components are distributed as

$$x_{\tau,0} \sim \mathcal{N}(\underline{\tau}, \underline{V}_{\tau}), \quad x_{c,0} \sim \mathcal{N}(0, \underline{V}_c)$$

where $\underline{V}_{\tau}$ is an identity matrix, and $\underline{V}_c$ is the unconditional variance of c_0 , a function of the VAR coefficients $\varphi = \{\Phi_1, \dots, \Phi_p\}'$ and variance Σ_c .

2.1 Data

Eight quarterly U.S. time series are used to estimate the Trend–Cycle VAR. Growth and price changes are annualized rates; the policy rate and inflation are expressed in percentages; unemployment and expectations are in percent. In particular, the dataset includes: (i) the growth rate of real per-capita GDP, g_t ; (ii) the unemployment rate, u_t ; (iii) the median four-quarter-ahead unemployment rate forecasts, $u_t^{e,1y}$, sourced from the SPF; (iv) the effective federal funds rate, f_t , where observations at the zero lower bound are treated as missing following the approach in Del Negro, Giannone, et al. (2017)¹; (v) the inflation rate, π_t , computed as the log difference of the GDP deflator (PGDP); (vi) the median four-quarter-ahead inflation expectations for the PGDP, $\pi_t^{e,1y}$, also from the SPF; and (vii) an average measure of ten-year-ahead inflation expectations, $\pi_t^{e,10y}$, which is constructed by combining survey forecasts on ten-year-ahead CPI inflation from the SPF with those from the Blue Chip Economic Indicators survey and then adjusting for the historical discrepancy between CPI and PGDP inflation.² The period from 1955:Q1 to 1959:Q4 is designated as the pre-sample, and the estimation sample runs from 1960:Q1 to 2019:Q4.

2.2 Inference

The estimation strategy follows a Bayesian approach using a Gibbs sampler that leverages the state-space representation in (2) and (3). As the estimation of the model closely follows the baseline specification of Bianchi et al. (2025), I adopt similar priors and initial conditions. In particular, for the assumptions regarding the initial conditions and prior distributions, the approach of Del Negro, Giannone, et al. (2017) is followed. Standard priors are imposed on the covariance matrices of the trend shocks, Σ_τ , and of the cyclical shocks, Σ_c , as well as on the VAR coefficients $\varphi = \{\Phi_1, \dots, \Phi_p\}'$. Specifically, I assume

$$\begin{aligned} p(\Sigma_\tau) &= \mathcal{IW}\left(\kappa_\tau, (\kappa_\tau + n_\tau + 1) \underline{\Sigma}_\tau\right), \\ p(\Sigma_c) &= \mathcal{IW}\left(\kappa_c, (\kappa_c + n_c + 1) \underline{\Sigma}_c\right), \\ p(\phi \mid \Sigma_c) &= \mathcal{N}\left(\underline{\phi}, \Sigma_c \otimes \underline{\Omega}\right) \mathcal{I}(\phi), \end{aligned}$$

where $\phi = \text{vec}(\varphi)$, and $\mathcal{IW}(\cdot, \cdot)$ denotes the inverse Wishart distribution with mode $\underline{\Sigma}$ and appropriate degrees of freedom. $\mathcal{N}(\cdot)$ denotes the normal distribution, while the indicator function $\mathcal{I}(\phi)$ equals 1 if the VAR is non-explosive and 0 otherwise.

Priors The prior for the initial conditions of the latent trends, τ_0 , is centered at the pre-sample means of the corresponding series. In our specification, the initial trend for

¹In that case, the Kalman filter used in the state-space estimation provides an estimate of this series during the ZLB period.

²This measure of long run inflation expectations is constructed based on Del Negro and Schorfheide (2013).

real GDP growth is set to 1%, for unemployment to 5%, for the real interest rate to 0.1%, for inflation to 2.5%, and for oil prices to 3%.

To construct the prior for the covariance matrix of the trend shocks, Σ_τ , I assume these shocks are a priori uncorrelated. Accordingly, we impose a standard deviation of 1% for the expected change in the annualized trend of real GDP growth over a 40-quarter horizon, while for unemployment, the real interest rate, and inflation the standard deviation is set at 1% over a 20-quarter period. For oil prices, the expected change in the trend is assigned a standard deviation of 1% over a 10-quarter period. In addition, a tight prior is enforced by setting $\kappa_\tau = 100$.

For the shocks to the cyclical components, we similarly assume independence, thereby specifying a diagonal structure for Σ_c . Their standard deviations are calibrated to reflect the observed pre-sample variability: the cyclical component of real per-capita GDP growth is assigned a standard deviation of 5%, while that for unemployment is set to 1.1%. The shocks affecting the nominal interest rate and inflation are fixed at 0.8% and 1.5%, respectively. For the cyclical component of four-quarter-ahead unemployment rate expectations, a standard deviation of 0.9% is imposed—consistent with its lower variability relative to the observed unemployment rate. Similarly, the common cyclical component of inflation expectations is given a prior standard deviation of 1.2%, slightly below that of realized inflation. Finally, the cyclical shock for oil prices is calibrated with a standard deviation of 8%, in line with its pre-sample variance. We set $\kappa_c = n_c + 2$, where n_c denotes the number of cyclical components.

For the VAR coefficients, denoted by $\phi = \text{vec}(\varphi)$ with $\varphi = \{\Phi_1, \dots, \Phi_p\}$, a conventional Minnesota prior is employed with an overall tightness hyperparameter of 0.2. Because the cyclical components are assumed stationary, the prior for each variable’s own lag is centered at 0 rather than at 1.

The model is estimated using 50,000 draws, with the initial 25,000 discarded as burn-in. The remaining draws are then thinned by retaining every 25th draw, resulting in a final sample of 1,000 draws for the structural identification procedure described in the next section.

3 Identifying demand and supply shocks

Identification strategy In this section, I describe the identification strategy followed to disentangle four business cycle shocks: a monetary policy shock, a non-policy demand shock, a cost-push shock, and an oil supply shock. While Bianchi et al. (2025) shows that a single “demand-like” shock can already explain much of the cyclical variation of real and nominal variables, it cannot pinpoint which specific mechanisms are at play—for instance, whether a demand impulse arises from monetary policy choices or broader private/government spending, or whether rising costs reflect general inflationary pressures (cost-push) or are rooted in shifts in global oil supply. By using a Trend–Cycle framework that cleans out low-frequency fluctuations, I explain how one can apply a multi-shock identification directly to the cyclical residuals to isolate these

separate channels more precisely.

Specifically, I focus on four types of shocks, each corresponding to a distinctive channel. A monetary policy shock captures the direct influence of central bank actions on output and inflation, while a non-policy demand shock represents fluctuations in aggregate spending that do not stem from policy interventions. An adverse cost-push shock captures generalized upward pressure on prices due to increased wages or other input costs, excluding oil, whereas an adverse oil supply shock arises from shifts in global oil production or geopolitical tensions that disproportionately affect oil prices and their impact on the economy. Together, these four shocks offer a rich decomposition of cyclical dynamics that a “main business cycle” shock would otherwise obscure.

The overall goal is to identify the shocks by *maximizing their contribution to the (conditional) forecast error variance (FEV)* of the output cycle $c_{y,t}$ or, as a robustness check, the unemployment cycle $c_{u,t}$. I consider two cases: (i) horizons $h = 1, \dots, 32$ and (ii) horizons $h = 6, \dots, 32$ (excluding cycles shorter than roughly 1.5 years). This maximization of the target variable is combined with sign restrictions in order to be able to uniquely identify each shock. Importantly, no additional relative magnitude restrictions on the contribution of each shock to the FEV of each variable are imposed (e.g., I do not force the monetary shock to explain a larger share of output variance than the demand shock), so that the data determine these relative contributions in a statistical manner.

Generalized Max Share The core of this strategy is a “Max Share” identification, which finds the structural shock that explains the largest possible fraction of a target variable’s forecast error variance. The specific implementation here relies on a Generalized Max Share procedure (Carriero & Volpicella, 2024) that identifies all shocks simultaneously rather than one at a time. This approach avoids the conventional “sequential” method—which typically requires an arbitrary ordering of shocks—and thus reduces the risk of confounding or blending one shock with another.³ Also when combined with sign restrictions, in contrast to traditional sign-restricted set-identification, this procedure yields point-identified shocks.

I begin by extracting the stationary (cyclical) component from the data. Recall from (1) that the cyclical dynamics of the system are given by

$$c_t = \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + \dots + \Phi_p c_{t-p} + \varepsilon_{c,t},$$

where $\varepsilon_{c,t}$ is the vector of reduced-form residuals with covariance matrix Σ_c . The goal is to decompose $\varepsilon_{c,t}$ into $k < n_c$ structural shocks that each account for a distinct share of the (conditional) forecast error variance (FEV) of a designated target variable, accumulated over a chosen set of horizons, where n_c is the number of cyclical components.

A standard representation writes the reduced-form residuals as:

$$\varepsilon_{c,t} = \Sigma_c^{1/2} \mathbf{Q} \mathbf{u}_t, \quad \text{where} \quad \mathbf{Q}'\mathbf{Q} = \mathbf{I}.$$

³On shock confounding within the Max Share framework, see Francis and Kindberg-Hanlon (2022) and Dou et al. (2024).

Moreover, $\Sigma_c^{1/2}$ is the (lower) Cholesky factor of Σ_c , $\mathbf{u}_t$ is the vector of structural shocks, and $\mathbf{Q}$ is an orthonormal rotation matrix. The objective is to select k columns, $\mathbf{q}_1, \dots, \mathbf{q}_k$, of $\mathbf{Q}$ that best explain the target's FEV over the horizons of interest.

Following Carriero and Volpicella (2024), let $\Upsilon_h^i(\phi)$ denote the h -step-ahead FEV decomposition matrix for target variable i implied by reduced-form parameters ϕ . To target *over* horizons, define the horizon-summed objective

$$\tilde{\Upsilon}_H^i(\phi) \equiv \sum_{h \in H} \Upsilon_h^i(\phi), \quad H \in \{\{1, \dots, 32\}, \{6, \dots, 32\}\},$$

and solve

$$\mathbf{Q}_{1:k}^* = \arg \max_{\mathbf{Q}_{1:k}} \sum_{i=1}^k \mathbf{q}_i' \tilde{\Upsilon}_H^i(\phi) \mathbf{q}_i, \quad (4)$$

subject to

$$\begin{aligned} \mathbf{q}_i' \mathbf{q}_i &= 1 \quad \text{for } i = 1, \dots, k, \\ \mathbf{S}_i(\phi) \mathbf{q}_i &\geq 0 \quad \text{for all } i \in \mathcal{I}_S, \end{aligned} \quad (5)$$

where $\mathbf{S}_i(\phi)$ collects the impact sign restrictions for shock i . Targeting *at* a single horizon h^* is the special case $H = \{h^*\}$, which replaces $\tilde{\Upsilon}_H^i(\phi)$ with $\Upsilon_{h^*}^i(\phi)$.

For each posterior draw of the reduced-form parameters, I solve (4)–(5) to obtain $\mathbf{Q}_{1:k}^*$, recover the associated structural shocks and impulse responses, and then aggregate across draws to form the posterior medians and coverage intervals reported in Section 4.2.

Sign constraints The sign restrictions are deliberately minimal and imposed only on *impact* to distinguish the dynamic responses of the shocks. Table 1 summarizes the restrictions: (a) a monetary policy tightening raises the policy rate and reduces output and inflation on impact; (b) a positive non-policy demand shock increases output, inflation, and the policy rate on impact; (c) an adverse cost-push shock decreases output and increases inflation while oil prices fall on impact; and (d) an adverse oil supply shock decreases output and increases inflation, with oil prices rising on impact. We require oil prices to move oppositely across supply shocks—up on oil-supply shocks, down on adverse non-oil cost-push shocks—to isolate oil-market disturbances from generalized input-cost pressures. These restrictions are consistent with standard macroeconomic theory and recent studies (e.g., Ascari et al., 2024; Peersman, 2005).

Table 1: Sign Restrictions

Variable	Monetary	Non-policy Demand	Cost-push	Oil Supply
Output	—	+	—	—
Unemployment				
Unemployment Exp.				
Policy Rate	+	+		
Inflation	—	+	+	+
Inflation Exp.				
Oil Price			—	+

Note: Positive signs (+) indicate an increase, while negative signs (—) indicate a decrease in the respective variable on their impact response. No sign means that the response is left unrestricted.

4 Results

In this section, I present the main findings of the paper starting with the unobserved trends and cyclical components extracted from the model estimation. Thereafter, I analyze the contribution of each of the four identified shocks to the cyclical variation of output and unemployment, and how each shock propagates to the business cycle fluctuations of the real and nominal variables.

4.1 Trends and cycles

Figure 1 displays the unobserved trends and cyclical components estimated by the model over 1960–2019. In Panel (1a), the red lines represent the observed data used in the estimation, while the dark blue lines depict the posterior median of the latent trends—derived from the model’s *trend-cycle decomposition*—together with their 68% and 90% coverage intervals (blue shaded regions).

These trends capture several well-known historical patterns of the U.S. economy. Notably, throughout the 1960s and 1970s, the model infers a pronounced rise in the trend inflation component, coupled with an upward drift in the trend unemployment rate. This finding aligns with the challenges policymakers faced in that era—such as responding to productivity slowdowns and accommodating heightened government spending in the mid-1960s—factors that jointly pushed both inflation and the “natural rate” of unemployment higher. Subsequently, with the shift to a more stringent monetary policy stance (the Volcker period), the trend components for inflation and long-term inflation expectations decline substantially. Even without explicit constraints linking those two series, they move in tandem, suggesting that incorporating long-term expectations helps the model distinguish trend from cycle. During the Great Financial Crisis, the trend in unemployment rises further, aligning with other estimates of the natural rate, whereas trend inflation remains fairly stable, consistent with literature pointing to anchored expectations in the post-2000 period.

Turning to the cyclical components in Panel (1b), a clear business cycle pattern

emerges. The unemployment rate spikes during recessions and gradually subsides in expansions, reflecting the familiar rise-and-fall dynamics of labor market slack. Under a New Keynesian perspective—where inflation tends to weaken when unemployment is above its natural level—cyclical inflation moves inversely with unemployment: it falls when unemployment is high and rebounds when unemployment is low. A similar but smoother pattern appears for one-year-ahead inflation expectations, highlighting the model’s ability to separate persistent trends from the more transitory swings in nominal variables. If one also considers output (Panel (1a)), it mirrors these fluctuations, with below-trend output accompanying high unemployment and vice versa, further supporting the interpretation that the model captures standard business cycle co-movements.

4.2 Demand, supply shocks and the business cycle

In this section, we examine how different types of demand- and supply-side shocks shape U.S. business cycle dynamics. First, we turn to impulse response functions—illustrating how each shock propagates through the cyclical components of key macroeconomic variables—and then to forecast error variance decompositions, which indicate the percentage of each variable’s business cycle fluctuations that are attributable to each shock.


Impulse Responses Figure 2 displays the median impulse responses of the cyclical components for key macroeconomic variables over a five-year horizon, alongside their 68% posterior coverage intervals. These shocks are identified using the baseline specification, which targets the forecast error variance of output over horizons from 1 to 32 quarters and imposes the sign restrictions summarized in Table 1.

A contractionary monetary policy shock raises the interest rate on impact. As expected, this tightening decreases output by approximately 2% on impact, an effect that persists for about three years, while the unemployment rate increases over a similar span. In tandem, inflation exhibits a sustained decrease that lasts for about four years, with an initial decline of 0.4%. Oil prices show no statistically significant reaction, confirming that a typical monetary policy shock has limited influence on their dynamics.


The non-policy demand shock induces a strong and persistent expansion in real activity, with an impact response of 2% that aligns with the findings of Bianchi et al. (2025). In response to this demand-driven boom, the Federal Reserve systematically tightens policy for approximately three years. This policy reaction is a likely cause for the subsequent “overshooting” of output, which temporarily dips below trend before normalizing. Unemployment falls by around 0.2% on impact and returns to its baseline after about four years. Inflation and inflation expectations rise on impact by 0.5% and 0.2% respectively, with expectations responding more strongly to this shock than to any other.

Both adverse supply shocks are stagflationary. Following a cost-push shock, unemployment rises by 0.1%, and the policy rate reaction is statistically insignificant. Similarly, after an oil-supply shock, inflation jumps by 0.5%, but the policy rate shows

Figure 1: Estimated trends and cycles



(a) Trends



(b) Cycles

Notes: Panel (A) displays the observed data (red lines) used to estimate the Trend–Cycle VAR model for the period 1960–2019, together with the posterior median estimates of the latent trend components (blue lines) and their corresponding 68% and 90% posterior coverage intervals (blue shaded regions). Panel (B) shows the posterior median estimates of the latent cyclical components (blue lines) along with their 68% and 90% coverage intervals (blue shaded regions). NBER recessions are indicated by grey shaded areas.

only a brief, insignificant easing. This muted reaction is consistent with policymakers “looking through” temporary supply-driven price spikes. Although the IRFs of oil-supply and non-oil cost-push shocks are close in shape, the identification separates them along two clear dimensions. First, we impose opposite impact signs on the oil price—up under oil-supply, down under cost-push—which pins down the rotation. Second, the oil-supply shock accounts for the dominant share of the oil-price FEV, whereas the cost-push shock contributes more to inflation outside the oil block.

The importance of short-run dynamics becomes clear when the identification targets only core business cycle horizons of 6–32 quarters (Figure 3). While the responses to non-policy demand and supply shocks remain broadly consistent, the results for the monetary policy shock change markedly: the response of inflation becomes statistically insignificant across most horizons, and inflation expectations rise on impact.

This pattern should not be taken as evidence of a central-bank information effect. With quarterly data and without external instruments or high-frequency surprises, such an effect is not cleanly identifiable. Rather, down-weighting short-run horizons weakens the separation of a pure policy innovation from other forces (e.g., systematic policy responses or financial frictions), allowing the estimated “policy” disturbance to load on those components. Consistent with this view, the baseline 1–32 specification—where short-run signatures discipline the rotation—recovers the conventional declines in inflation and expectations. A cleaner separation of policy and information components would require augmenting the identification with high-frequency monetary surprises or narrative instruments, which is beyond the scope of this paper.

In contrast to the monetary shock, the impulse responses to the non-policy demand shock identified over the 6–32-quarter horizons remain broadly consistent with those from the full-horizon identification, as do the responses to the two supply shocks. The primary differences are a more muted initial output contraction and a more pronounced positive impact on inflation expectations following the two supply-side shocks.

Finally, the remaining figures 4 and 5 display the impulse responses when targeting the forecast error variance of unemployment (over horizons 1–32 and 6–32, respectively). Both qualitatively and quantitatively, these responses are very similar to those obtained when targeting output, confirming the robustness of the main results.

Variance contributions Tables 2 through 5 report the shares of cyclical variation explained by the four identified shocks. Across specifications, a clear narrative emerges: *demand is the primary driver of the business cycle*. The two demand shocks (Monetary Policy and Non-policy Demand) together account for about one-half to three-quarters of the variance in the core real-activity variables. Taken together, the four shocks explain most of the business-cycle forecast error variance in both real and nominal variables—typically 85–95% for real activity and the policy rate, with expectation series somewhat lower (around 80%).

A more nuanced picture appears across horizons. In the baseline specification (targeting the FEV of output over 1–32 quarters), which includes short-run movements,

monetary policy is a strong driver of real activity. At the same time, *non-policy demand is the single largest contributor to the policy rate and to inflation* when horizons include the short run, indicating that underlying shifts in private and public spending and preferences are central for price setting.

When I exclude high-frequency cycles and focus on the core business cycle (6–32 quarters), *non-policy demand becomes dominant for real activity* across targets. For inflation, the split depends on the identification target: in the output-targeted specification (Table 3), supply (cost-push + oil) accounts for the majority share (nearly 57%); in the unemployment-targeted specification (Table 5), demand edges out supply (about 50% vs. 40%). In all cases, oil prices are predominantly driven by the oil-supply shock, consistent with the sign restrictions and the intended separation from non-oil cost pressures.

Comparing these results to Bianchi et al. (2025) is revealing. In their trend-cycle VAR, a single “main business cycle” shock accounts for about 60% of cyclical real activity but a smaller share of inflation. By contrast, once I split demand into *monetary* versus *non-policy* components and supply into *cost-push* versus *oil* disturbances, the fraction of inflation explained by identified channels rises substantially. In essence, modeling multiple demand and supply channels closes much of the unexplained inflation gap, highlighting that several structural forces simultaneously underpin real and nominal fluctuations. The core narrative—that demand drives real activity while both demand and supply drive inflation—holds whether I target output or unemployment, demonstrating the robustness of the results.

5 Discussion

This paper set out to disentangle the drivers of U.S. business cycle fluctuations using a Trend-Cycle framework that distinguishes between multiple demand and supply shocks. The results reveal a clear hierarchy of influence: demand-side forces—composed of both monetary policy and broader non-policy components—are the primary engine of fluctuations in output and unemployment. Supply shocks, while playing a crucial role, are secondary drivers of real activity but become paramount for explaining inflation dynamics, especially over medium-term business cycle horizons.

The impulse responses reveal theoretically consistent dynamics that lend considerable credence to the identification. For instance, the response to a positive non-policy demand shock exhibits an “overshooting” pattern, where output rises for roughly two years before temporarily dipping below trend. This is not an anomaly but rather a reflection of the endogenous monetary policy response potentially captured by the model. The Federal Reserve is seen hiking interest rates persistently to cool the demand-driven boom, which eventually restrains the economy. This feedback loop is precisely what standard macroeconomic models predict and is a testament to the model’s ability to capture real-world dynamics. Similarly, the identified supply shocks are appropriately stagflationary, while the model correctly shows the central bank tending to “look


through” supply-driven price spikes—including oil—with only small and short-lived policy-rate movements.

The analysis also yields an important insight into the nature of monetary policy shocks themselves. While the baseline specification (1-32 quarters) successfully identifies a monetary policy shock with theoretically consistent effects, this result changes when the identification focuses only on medium-run horizons (6-32 quarters). Interestingly, under this alternative specification, the identified monetary shock no longer has a significant effect on inflation. This may be interpreted not as a failure of the model, but rather a finding. It suggests that the unique, identifying signature of a monetary policy shock is most evident in its short-run dynamics. By instructing the model to disregard this crucial short-run information, the procedure has more difficulty distinguishing a pure monetary shock from other disturbances. This result therefore reinforces the choice of the baseline specification and highlights that the short-run real-nominal trade-off is a key feature of monetary policy shocks in the data.

From a policy perspective, these findings carry several implications. The key challenge for effective stabilization policy is not just to identify a “demand shock” but to accurately diagnose its origin. An economic expansion driven by non-policy demand warrants a different monetary policy response than one fueled by the central bank’s own actions. The results provide a rationale for policymakers to look beyond aggregate labels to understand the specific forces at play.


As a work in progress, this paper opens several avenues for future research. The next step is to extend the framework to incorporate household inflation expectations from the Michigan Survey of Consumers. This extension would enable a direct comparison with professional forecasts and assess whether the identified shocks propagate similarly through household beliefs.

Figure 2: Impulse responses of the identified shocks targeting output over horizons from 1 to 32 quarters




Note: Impulse responses of the identified shocks targeting the conditional FEV of output over horizons from 1 to 32 quarters. The continuous lines depict the posterior median at each horizon (1 to 5 years), while the shaded regions show the 68% posterior coverage intervals of the impulse responses.

Figure 3: Impulse responses of the identified shocks targeting output over horizons from 6 to 32 quarters




Note: Impulse responses of the identified shocks targeting the conditional FEV of output over horizons from 6 to 32 quarters. The continuous lines depict the posterior median at each horizon (1 to 5 years), while the shaded regions show the 68% posterior coverage intervals of the impulse responses.

Figure 4: Impulse responses of the identified shocks targeting unemployment over horizons from 1 to 32 quarters



Note: Impulse responses of the identified shocks targeting the conditional FEV of unemployment over horizons from 1 to 32 quarters. The continuous lines depict the posterior median at each horizon (1 to 5 years), while the shaded regions show the 68% posterior coverage intervals of the impulse responses.

Figure 5: Impulse responses of the identified shocks targeting unemployment over horizons from 6 to 32 quarters



Note: Impulse responses of the identified shocks targeting the conditional FEV of unemployment over horizons from 6 to 32 quarters. The continuous lines depict the posterior median at each horizon (1 to 5 years), while the shaded regions show the 68% posterior coverage intervals of the impulse responses.

Table 2: Variance contribution of business cycle shocks (FEV *accumulated* over horizons 1–32), output-identified

Variable	Monetary	Non-policy Demand	Cost-push	Oil Supply	Total Demand	Total Supply	Total
Output	33.2 (26.1, 39.9)	26.0 (15.1, 39.6)	19.6 (14.2, 25.9)	17.4 (11.9, 23.4)	59.2	37.0	96.2
Unempl.	30.1 (21.1, 39.2)	31.7 (17.5, 48.5)	15.6 (9.2, 23.1)	12.7 (7.2, 20.5)	61.8	28.2	90.0
Unempl. exp.	24.2 (16.0, 34.2)	28.5 (13.5, 44.7)	15.9 (9.1, 25.6)	13.1 (6.9, 21.7)	52.7	29.0	81.7
Policy rate	19.1 (11.4, 28.9)	59.5 (42.0, 72.3)	5.7 (2.4, 12.2)	5.8 (2.5, 14.1)	78.5	11.5	90.1
Inflation	17.1 (5.7, 36.8)	51.0 (33.7, 62.8)	11.3 (7.8, 15.5)	13.7 (7.6, 21.3)	68.1	25.0	93.1
Inflation exp.	13.3 (6.4, 26.2)	59.1 (42.9, 70.0)	4.5 (2.1, 7.6)	4.4 (1.9, 8.9)	72.4	8.9	81.3
Oil price	3.5 (1.2, 9.7)	20.7 (11.6, 32.2)	28.7 (19.8, 37.3)	41.6 (27.5, 52.7)	24.1	70.3	94.4

Note: The figures represent variance contributions of the identified shocks with 68% coverage intervals in parentheses. 'Total Demand', 'Total Supply', and 'Total' are aggregated figures across the respective shocks. Totals below 100 reflect residual orthogonal shocks not targeted by the identification.

Table 3: Variance contribution of business cycle shocks (FEV *accumulated* over horizons 6–32), output-identified

Variable	Monetary	Non-policy Demand	Cost-push	Oil Supply	Total Demand	Total Supply	Total
Output	19.2 (11.7, 29.0)	54.2 (36.7, 67.0)	11.6 (6.3, 19.0)	9.1 (4.9, 17.4)	73.4	20.7	94.1
Unempl.	14.4 (8.3, 23.7)	62.0 (47.0, 72.9)	7.3 (3.6, 14.1)	6.1 (3.3, 12.7)	76.4	13.4	89.9
Unempl. exp.	11.5 (6.7, 19.9)	58.7 (44.4, 68.5)	6.2 (3.1, 13.6)	6.2 (3.1, 13.5)	70.2	12.5	82.6
Policy rate	19.5 (12.6, 27.4)	58.5 (48.0, 69.0)	7.8 (3.6, 12.8)	6.1 (3.0, 11.2)	78.1	13.9	92.0
Inflation	6.4 (3.6, 10.8)	30.6 (19.6, 45.0)	24.5 (17.5, 30.1)	32.4 (22.1, 40.5)	37.0	56.9	93.8
Inflation exp.	12.9 (6.9, 19.6)	40.7 (30.3, 54.1)	13.9 (8.2, 19.5)	16.8 (9.0, 24.6)	53.5	30.8	84.3
Oil price	12.9 (3.3, 28.8)	7.1 (3.0, 14.9)	18.0 (10.6, 26.8)	55.1 (38.2, 68.9)	20.0	73.2	93.1

Note: The figures represent variance contributions of the identified shocks with 68% coverage intervals in parentheses. 'Total Demand', 'Total Supply', and 'Total' are aggregated figures across the respective shocks. Totals below 100 reflect residual orthogonal shocks not targeted by the identification.

Table 4: Variance contribution of business cycle shocks (FEV *accumulated* over horizons 1–32), unemployment-identified

Variable	Monetary	Non-policy Demand	Cost-push	Oil Supply	Total Demand	Total Supply	Total
Output	29.5 (23.4, 36.0)	22.6 (14.9, 31.6)	23.6 (17.9, 28.7)	20.6 (15.2, 26.0)	52.1	44.2	96.2
Unempl.	31.1 (24.2, 38.5)	18.3 (10.0, 30.8)	23.4 (16.5, 30.5)	17.3 (11.9, 23.9)	49.3	40.8	90.1
Unempl. exp.	24.9 (17.8, 32.1)	14.5 (7.7, 26.4)	26.9 (19.2, 34.6)	17.3 (11.2, 24.6)	39.3	44.1	83.5
Policy rate	23.8 (14.9, 34.7)	37.9 (21.5, 57.4)	14.5 (6.5, 25.1)	8.8 (3.9, 17.2)	61.8	23.3	85.0
Inflation	33.5 (14.7, 52.6)	41.3 (15.5, 61.3)	8.0 (4.9, 11.6)	6.3 (3.0, 11.3)	74.9	14.3	89.2
Inflation exp.	23.7 (11.8, 38.3)	43.5 (18.5, 62.5)	5.3 (1.9, 11.9)	3.3 (1.4, 6.6)	67.2	8.7	75.9
Oil price	3.6 (1.3, 8.5)	14.8 (5.2, 29.0)	29.9 (16.5, 42.4)	37.1 (19.9, 51.7)	18.5	67.0	85.4

Note: The figures represent variance contributions of the identified shocks with 68% coverage intervals in parentheses. 'Total Demand', 'Total Supply', and 'Total' are aggregated figures across the respective shocks. Totals below 100 reflect residual orthogonal shocks not targeted by the identification.

Table 5: Variance contribution of business cycle shocks (FEV *accumulated* over horizons 6–32), unemployment-identified

Variable	Monetary	Non-policy Demand	Cost-push	Oil Supply	Total Demand	Total Supply	Total
Output	25.3 (17.6, 32.7)	39.6 (22.1, 53.5)	16.6 (10.5, 25.6)	14.7 (9.7, 21.9)	64.9	31.3	96.1
Unempl.	21.4 (13.7, 30.6)	49.9 (29.9, 63.4)	11.3 (5.7, 21.7)	8.9 (5.0, 17.0)	71.4	20.3	91.6
Unempl. exp.	16.8 (9.8, 25.4)	47.4 (27.6, 60.4)	12.0 (5.6, 24.2)	8.8 (4.4, 17.4)	64.2	20.8	84.9
Policy rate	20.2 (13.3, 27.1)	64.5 (52.4, 73.8)	5.0 (2.3, 11.0)	4.5 (2.0, 8.9)	84.7	9.4	94.1
Inflation	10.1 (4.1, 26.5)	40.1 (26.3, 55.5)	19.3 (11.3, 27.4)	20.9 (11.2, 30.5)	50.3	40.2	90.5
Inflation exp.	10.4 (5.5, 19.5)	50.9 (37.9, 62.9)	10.0 (4.8, 18.8)	10.0 (4.4, 17.8)	61.2	20.0	81.2
Oil price	2.9 (1.0, 8.5)	11.2 (5.2, 19.7)	16.0 (6.0, 27.1)	55.7 (34.3, 70.3)	14.1	71.8	85.9

Note: The figures represent variance contributions of the identified shocks with 68% coverage intervals in parentheses. 'Total Demand', 'Total Supply', and 'Total' are aggregated figures across the respective shocks. Totals below 100 reflect residual orthogonal shocks not targeted by the identification.

References

- Angeletos, G.-M., Collard, F., & Dellas, H. (2020). Business-Cycle Anatomy. *American Economic Review*, 110(10), 3030–3070.
- Ascari, G., Bonam, D., Mori, L., & Smadu, A. (2024). *Fiscal Policy and Inflation in the Euro Area*.
- Bianchi, F., Nicolò, G., & Song, D. (2025). *Inflation and Real Activity over the Business Cycle*. National Bureau of Economic Research.
- Blanchard, O. J., & Quah, D. (1989). The Dynamic Effects of Aggregate Demand and Supply Disturbances. *American Economic Review*, 79(4), 655–673.
- Carriero, A., & Volpicella, A. (2024). Max Share Identification of Multiple Shocks: An Application to Uncertainty and Financial Conditions. *Journal of Business & Economic Statistics*, 1–13.
- Del Negro, M., Giannone, D., Giannoni, M. P., & Tambalotti, A. (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity*, 2017(1), 235–316.
- Del Negro, M., & Schorfheide, F. (2013). DSGE Model-Based Forecasting. In *Handbook of Economic Forecasting* (pp. 57–140, Vol. 2). Elsevier.
- Dou, L., Ho, P., & Lubik, T. A. (2024). Max-Share Misidentification.
- Faust, J. (1998). The robustness of identified VAR conclusions about money. *Carnegie-Rochester Conference Series on Public Policy*, 49, 207–244.
- Forni, M., Gambetti, L., Granese, A., Sala, L., & Soccorsi, S. (2024). An American Macroeconomic Picture: Supply and Demand Shocks in the Frequency Domain. *American Economic Journal: Macroeconomics*.
- Francis, N., & Kindberg-Hanlon, G. (2022). Signing Out Confounding Shocks in Variance-Maximizing Identification Methods. *AEA papers and proceedings*, 112, 476–480.
- Granese, A. (2024). Two Main Business Cycle Shocks are Better than One
Accepted: 2024-04-10T06:59:43Z.
- Peersman, G. (2005). What caused the early millennium slowdown? Evidence based on vector autoregressions. *Journal of Applied Econometrics*, 20(2), 185–207.
- Uhlig, H. (2003). What moves real GNP? *Unpublished manuscript*.

A Data Sources and Transformations

The model is estimated using eight quarterly U.S. time series for the period 1960:Q1 to 2019:Q4. The period from 1955:Q1 to 1959:Q4 is used as a pre-sample for prior calibration. All data are sourced from the Federal Reserve Economic Database (FRED) at the Federal Reserve Bank of St. Louis, unless otherwise noted.

- **Real GDP per capita growth (g_t):** We use the annualized quarterly growth rate of Real Gross Domestic Product Per Capita (FRED ID: A939RX0Q048SBEA).
- **Unemployment rate (u_t):** The Civilian Unemployment Rate (FRED ID: UNRATE), which we convert from a monthly to a quarterly series by averaging.
- **Effective Federal Funds Rate (f_t):** The Effective Federal Funds Rate (FRED ID: FEDFUNDS), averaged to a quarterly frequency. Following Del Negro, Giannone, et al. (2017), we treat observations during the Zero Lower Bound period (2008:Q4–2015:Q4) as missing values.
- **Inflation (π_t):** For inflation, we use the annualized quarterly growth rate of the GDP Implicit Price Deflator (FRED ID: GDPDEF).
- **One-year-ahead unemployment expectations ($u_t^{e,1y}$):** The median one-year-ahead forecast for the unemployment rate from the Survey of Professional Forecasters (SPF).
- **One-year-ahead inflation expectations ($\pi_t^{e,1y}$):** The median one-year-ahead forecast for GDP deflator inflation from the SPF.
- **Ten-year-ahead inflation expectations ($\pi_t^{e,10y}$):** A measure of long run inflation expectations constructed by combining forecasts from the SPF and Blue Chip surveys, adjusted to be consistent with GDP deflator inflation, following the methodology of Del Negro and Schorfheide (2013).
- **Oil price growth (o_t):** The annualized quarterly growth rate of the WTI spot price (FRED ID: WTISPLC) deflated by the GDP deflator. To construct this series, we convert the monthly price data to quarterly, adjust for general inflation using the GDP deflator, and then calculate the growth rate.

B Model Estimation and State-Space Representation

B.1 Settings

The cyclical components of the model are assumed to evolve according to a VAR with two lags ($p = 2$). The model is estimated using a Gibbs sampler with 50,000 draws.

The first 25,000 draws are discarded as a burn-in period, and the remaining draws are thinned by a factor of 25, resulting in a final set of 1,000 posterior draws used for inference.

B.2 State-Space Representation

The Trend–Cycle VAR can be written in the following state-space form:

$$z_t = \Lambda x_t = \Lambda_\tau x_{\tau,t} + \Lambda_c x_{c,t} + \Lambda_\eta x_{\eta,t} \quad (6)$$

$$x_t = \Phi x_{t-1} + \mathcal{R} \varepsilon_t \quad (7)$$

where the vectors and matrices are defined as follows. The vector of $n = 8$ observable variables is:

$$z_t = (g_t, u_t, u_t^{e,1y}, f_t, \pi_t, \pi_t^{e,1y}, \pi_t^{e,10y}, o_t)'$$

The state vector x_t comprises the latent trend components $x_{\tau,t}$, the cyclical components $x_{c,t}$, and the measurement error components $x_{\eta,t}$:

$$x_t = (x'_{\tau,t}, x'_{c,t}, x'_{\eta,t})'$$

The $n_\tau = 5$ trend components are collected in τ_t :

$$x_{\tau,t} = \tau_t = (\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}, \tau_{o,t})'$$

The $n_c = 7$ cyclical components and their lags are collected in $x_{c,t}$:

$$x_{c,t} = (c'_t, c'_{t-1})', \quad \text{where} \quad c_t = (c_{y,t}, c_{u,t}, c_{u,t}^{e,1y}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^e, c_{o,t})'$$

The $n_\eta = 1$ measurement error component is collected in $x_{\eta,t}$:

$$x_{\eta,t} = \eta_t, \quad \text{where} \quad \eta_t = (\eta_{\pi,t}^{e,10y})$$

The mapping from the states to the observables in (6) is defined by the Λ matrices:

$$\Lambda_\tau = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Lambda_c = [\Lambda_{c,0} \quad \Lambda_{c,1}], \quad \Lambda_\eta$$

with

$$\Lambda_{c,0} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \delta \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \quad \Lambda_{c,1} = \begin{pmatrix} -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$\Lambda_\eta = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}$$

The state transition dynamics in (7) are governed by:

$$\Phi = \begin{pmatrix} I_{n_\tau} & 0 & 0 \\ 0 & \Phi_c & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad \mathcal{R} = \begin{pmatrix} I_{n_\tau} & 0 & 0 \\ 0 & \mathcal{R}_c & 0 \\ 0 & 0 & I_{n_\eta} \end{pmatrix}$$

where Φ_c contains the VAR coefficients $\{\Phi_1, \Phi_2\}$ and identity matrices to handle the lags, and $\mathcal{R}_c$ is an identity matrix stacked on zeros. The vector of innovations is $\varepsilon_t = (\varepsilon'_{\tau,t}, \varepsilon'_{c,t}, \varepsilon'_{\eta,t})'$, with covariance matrix Σ :

$$\text{Var}(\varepsilon_t) \equiv \Sigma = \begin{pmatrix} \Sigma_\tau & 0 & 0 \\ 0 & \Sigma_c & 0 \\ 0 & 0 & \Sigma_\eta \end{pmatrix}$$

B.3 Gibbs Sampler

We use a Gibbs sampler to estimate the model unknowns, leveraging the state-space representation. For the j -th iteration of the sampler, the algorithm proceeds as follows:

1. **Draw Latent States:** Conditional on the parameter draws from the previous iteration (Θ^j), draw the full time series of the state vector, $x_{1:T}^{j+1}$. This is achieved using the Carter and Kohn (1994) simulation smoother, which consists of a forward pass (the Kalman filter) and a backward pass (the smoothing recursion).
 - The Kalman filter recursively computes the filtered states $x_{t|t}$ and their covariance matrices $P_{t|t}$. The forecasting and updating equations are standard:

$$\begin{aligned} x_{t|t-1} &= \Phi x_{t-1|t-1} \\ P_{t|t-1} &= \Phi P_{t-1|t-1} \Phi' + \mathcal{R} \Sigma \mathcal{R}' \\ x_{t|t} &= x_{t|t-1} + P_{t|t-1} \Lambda' (\Lambda P_{t|t-1} \Lambda')^{-1} (z_t - \Lambda x_{t|t-1}) \\ P_{t|t} &= P_{t|t-1} - P_{t|t-1} \Lambda' (\Lambda P_{t|t-1} \Lambda')^{-1} \Lambda P_{t|t-1} \end{aligned}$$

- The backward recursion then draws the smoothed states x_t^{j+1} for $t = T - 1, \dots, 1$ from $N(x_{t|t+1}, P_{t|t+1})$, where:

$$\begin{aligned} x_{t|t+1} &= x_{t|t} + P_{t|t} \Phi' P_{t+1|t}^{-1} (x_{t+1}^{j+1} - \Phi x_{t|t}) \\ P_{t|t+1} &= P_{t|t} - P_{t|t} \Phi' P_{t+1|t}^{-1} \Phi P_{t|t} \end{aligned}$$

2. **Draw Model Parameters:** Conditional on the newly drawn history of the states $x_{1:T}^{j+1}$, draw the new parameter values Θ^{j+1} .

- **VAR coefficients φ :** Given the smoothed cyclical components c_t , the VAR can be written as a system of linear regressions. With a Normal-inverse-Wishart prior, the posterior for the coefficients $\varphi = \text{vec}(\{\Phi_1, \Phi_2\})$ is conditionally Normal. We draw from this posterior and retain the draw if it satisfies the stationarity condition for the cyclical VAR.
- **Covariance matrices $\Sigma_\tau, \Sigma_c, \Sigma_\eta$:** Given the smoothed states, we can compute the corresponding innovations $\varepsilon_{\tau,t}, \varepsilon_{c,t}, \varepsilon_{\eta,t}$. With conjugate inverse-Wishart priors, the posterior distributions for the covariance matrices are also inverse-Wishart, from which we can draw directly.
- **Parameter δ :** The loading parameter δ in the measurement equation for ten-year inflation expectations is drawn from its conditional posterior distribution, which is also Normal under a Normal prior.

The algorithm iterates these two main blocks until the distribution of the draws converges.