NBER WORKING PAPER SERIES

INFLATION AND REAL ACTIVITY OVER THE BUSINESS CYCLE

Francesco Bianchi
Giovanni Nicolò
Dongho Song

Working Paper 31075
http://www.nber.org/papers/w31075

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2023, Revised March 2025

We would like to thank our discussants Domenico Giannone, Marta Ba bura, and Vasco Cúrdia, as well as Gianni Amisano, Guido Ascari, Dario Caldara, Thomas Drechsel, Filippo Ferroni, Francesco Furlanetto, Etienne Gagnon, Ed Herbst, Margaret Jacobson, Ben Johannsen, Leonardo Melosi, Emi Nakamura, Anna Orlik, Juan Rubio-Ramírez, Luca Sala, Frank Schorfheide, Mark Watson, and all seminar participants at George Washington University, Georgetown, the 2022 Central Bank Research Association Annual Meeting, the 2022 mid-year meeting of the NBER EFSF Workgroup on Methods and Applications for DSGE Models, the Inflation: Drivers and Dynamics 2023 Conference, the International Association for Applied Econometrics 2023 Conference, the Computing in Economics and Finance 2023 Conference, and the 2024 Annual Meeting of the Society for Economic Dynamics for useful comments and suggestions. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Board, the Federal Reserve System, or the National Bureau of Economic Research.

Inflation and Real Activity over the Business Cycle
Francesco Bianchi, Giovanni Nicolò, and Dongho Song
NBER Working Paper No. 31075
March 2023, Revised March 2025
JEL No. C32, E31, E32

ABSTRACT

We study the relation between inflation and real activity over the business cycle. We employ a Trend-Cycle VAR to control for low-frequency movements in inflation, unemployment, and growth that are pervasive in the post-WWII period. We show that cyclical fluctuations of inflation are related to cyclical movements in real activity and unemployment, in line with what is implied by the New Keynesian framework. We then discuss the reasons for which our results relying on a Trend-Cycle VAR differ from the findings of previous studies based on VAR analysis. We explain empirically and theoretically how to reconcile these differences.

Francesco Bianchi
Johns Hopkins University
Department of Economics
3400 N. Charles Street
544E Wyman Bldg.
Baltimore, Maryland 21218
and CEPR
and also NBER
francesco.bianchi@jhu.edu

Giovanni Nicolò
Federal Reserve Board
20th St & Constitution Avenue NW
Washington, DC 20551
giovanni.nicolo@frb.gov

Dongho Song
Johns Hopkins University
Carey Business School
100 International Drive
Baltimore, MD 21202
dongho.song@jhu.edu

1 Introduction

From a traditional Keynesian Phillips-curve viewpoint, there is an established link between inflation and unemployment. However, during the Great Moderation, the empirical evidence supporting a connection between real economic activity and inflation diminished (Stock and Watson, 2020). This shift prompted economists to reconsider the core principles of New Keynesian models. Recent contributions offer explanations for the weakening of this empirical relationship (Del Negro et al., 2020, McLeay and Tenreyro, 2020) or modifications to the New Keynesian model to improve its empirical fit (Gust et al., 2022, Farmer and Nicolò, 2018, 2019). Other authors abandon the New Keynesian framework and develop business-cycle theories that abstract from inflation (Beaudry et al., 2020; Basu et al., 2024) or in which shocks driving the business-cycle are reminiscent of demand shocks but have no inflationary effects (Beaudry and Portier, 2013; Angeletos et al., 2018).

An important empirical justification for these alternative theoretical frameworks is provided by Angeletos et al. (2020). In their seminal contribution, the authors follow an extensive literature that uses vector autoregressions (VARs) as a "model free," but still structural approach to the data. Using a VAR applied to post-WWII US data, they identify a "business-cycle" shock that accounts for the largest share of variation in GDP or unemployment at business-cycle frequencies. This shock explains most of the cyclical fluctuations in various measures of real activity but has little impact on the business-cycle variation of inflation. The authors conclude that their findings contradict the New Keynesian framework, which posits a strong link between inflation and economic slack.

As we argue next, the approach of using a VAR to identify shocks in the frequency domain has some limitations when the goal is to assess the business-cycle relationship between real and nominal variables over the US post-WWII period. The main reason is that a standard fixed-coefficient VAR might be unable to correctly disentangle business-cycle and low-frequency movements in those variables over a relatively short period of time that features structural breaks (Clarida et al., 2000; Sims and Zha, 2006; Bianchi, 2013; Bianchi and Ilut, 2017). In a VAR, a single set of parameters and shocks need to accommodate the variation at all frequencies observed over a relatively short period. As a result, a procedure that uses the estimated parameters to identify variation at business-cycle frequency might be biased. The problem is particularly severe if one of the variables of interest shows significant variation at low frequency, as it is the case with inflation. Ultimately, the identified shock might fail to capture a business-cycle relationship between the two variables even when such a relation is in fact in the data. To remedy this limitation of the VAR for the specific question at hand, we adopt a more flexible model that explicitly extracts business-cycle movements in the variables of interest. Specifically, we argue that a Trend-Cycle VAR (TC-VAR) is better suited to analyze the business-cycle relation between inflation and real activity.

We start by presenting simple, but insightful, evidence that serves as motivation for our analysis. We consider a measure of inflation—the GDP deflator—and two measures of real economic activity—real GDP per capita and unemployment—over the period between 1960 and 2019. Using a bandpass filter, we extract movements in those measures at frequencies between 6 and 32 quarters—labeled "business-cycle frequencies"—and between 8 and 30 years—labeled "medium-cycle frequencies." At business-cycle frequencies, the correlation between inflation and real GDP per capita (unemployment) is positive (negative) and roughly equal to about 0.2 (negative 0.4). The correlations become larger (in absolute value) when considering the relationship between current inflation and lagged measures of real economic activity, peaking at about 0.45 (negative 0.45) when considering real GDP per capita (unemployment) lagged by four (two) quarters. In addition, over the medium cycle, these correlations can be nearly 50% larger than at business-cycle frequencies. This evidence is puzzling in light of the existing literature because it suggests that inflation is related to real activity at business-cycle frequencies, at least to some extent.

Motivated by this analysis, we adopt a more rigorous empirical framework and estimate a multivariate Trend-Cycle VAR building on the work of Watson (1986), Stock and Watson (1988, 2007), Villani (2009) and Del Negro et al. (2017). We consider the sample between 1960 and 2019 using seven time series. The first four time series are commonly used in previous studies: GDP growth, unemployment, the federal funds rate (FFR), and inflation. We then include three additional variables. First, to better capture low-frequency movements in inflation, we add ten-year-ahead inflation expectations. Agents' long-term inflation expectations are informative about the *current* level of trend inflation, even if not necessarily good predictors of future inflation. Second, we use one-year-ahead inflation expectations as a variable that should respond to business cycle variation in inflation and be less affected by transitory shocks. Third, we include one-year-ahead expectations of unemployment to inform the estimates of the latent trend of unemployment.

Given that a TC-VAR already separates trends from cycles, we identify the shock that maximizes the variation of the latent cyclical component of unemployment, and we study its contribution to the volatility of all cyclical components. A series of important results emerge from our analysis. First, under our baseline specification, the shock targeting unemployment explains about 70% of the unemployment cycle and close to 40% of the inflation cycle. This large share suggests that it is important to account for the low-frequency movements in real and nominal variables when studying their cyclical relationship. Second, the unemployment-identified shock explains about 55% of the inflation expectations cycle. This result provides further support for the notion that business-cycle movements in inflation are in fact related to real activity, given that expected inflation is obviously related to actual inflation, but less affected by high-frequency fluctuations due to transitory shocks. In line with this reasoning, when we focus on frequencies that correspond to cyclical fluctuations with duration of at least 1.5 years, the results become stronger. In this case, the shock identified targeting unemployment explains a higher percentage of the business-cycle volatility of inflation (44%) and inflation expectations (58%) compared to when all frequencies are considered. Finally, the results are very similar if we use GDP instead of unemployment to identify the real-activity shock.

A TC-VAR has four clear advantages relative to a VAR when identifying shocks in the frequency domain (Uhlig, 2003, Giannone et al., 2019b). First, the inference exercise automatically separates trends and cycles. Second, cyclical variation is controlled by a different set of parameters with respect to low-frequency variation. Third, we do not need to take a stance on the typical length of the business cycle. This is important in light of the fact that expansions have become progressively longer in the sample under consideration. Finally, by allowing for changes in trend growth, trend inflation, and long-run unemployment, a TC-VAR accommodates the notion that what matters for cyclical movements in inflation is the output or unemployment gap, while at the same time nesting a typical VAR specification if the data do not feature significant variation in the trend component.

A standard VAR cannot easily recover the business-cycle relationship between nominal and real variables even if we use the same data. To illustrate this point, we estimate a VAR with two lags, following Angeletos et al. (2020). We explore different priors, starting with flat priors, followed by optimized Minnesota priors, and finally optimized Minnesota priors combined with long-run restrictions *à la* Giannone et al. (2019a). For each specification, we identify a max-share shock that targets unemployment at business-cycle frequencies. As in Angeletos et al. (2020), the contribution of the shock to the variability of inflation at the same frequencies is very low, ranging from about 13% when using a Minnesota prior to about 17% when also assuming long-run priors.

We then lay out theoretical arguments for why the two methodological approaches reach such different conclusions. We demonstrate that a fixed-coefficient VAR estimated over a period of time that presents low-frequency variation is misspecified, if the goal is to assess the commovement at business-cycle frequency. The misspecification problem associated with the use of a VAR to describe a data generating process characterized by both low- and high-frequency movements cannot be resolved. An econometrician would need infinite data to reconstruct the VAR representation of a TC-VAR. Even in that case, the reduced-form innovations that she would recover would map into the innovations affecting the latent persistent and stationary components *as well as* the estimation error associated with the latent components. In reality, these issues are exacerbated by the fact that the VAR parameters estimated over a finite sample would be distorted because a single set of parameters needs to account for both trend and cycle fluctuations.

We conclude the paper by providing an illustration of these issues with a simple model of unemployment and inflation. We generate Monte Carlo simulations of the two series using a TC-VAR. We then show that if the simulated data contain low-frequency variation in any of the two series, the VAR fails to recover a relation between the two variables even when the inflation cycle is *exclusively* driven by the unemployment cycle. In contrast, if the model does not feature any relation between output and inflation, a TC-VAR would correctly uncover the absence of commovement.

Our analysis connects to Sargent and Sims (1977) that shows that macro data can be well explained by two factors, one of which captures low-frequency movements in nominal variables. The factor-analysis literature has repeatedly confirmed this key insight (Giannone et al., 2006; Watson, 2004; Stock and Watson, 2011; Forni et al., 2025). Related branches of the literature use unobserved component models to measure long-run inflation (Stock and Watson, 2007; Mertens, 2016) and long-run interest rates (Laubach and Williams, 2003, 2015; Lubik and Matthes, 2015; Del Negro et al., 2017; Del Negro et al., 2019; Holston et al., 2017; Lewis and Vazquez-Grande, 2019; and Johannsen and Mertens, 2021).

Our results are consistent with the findings of Hazell et al. (2022). These authors show that a stable relation between real activity and inflation can be recovered from the data when controlling for long-term inflation expectations. Their evidence is based on a cross-sectional analysis across US states, while we take a time-series approach. Ascari and Sbordone (2014) emphasize the importance of controlling for trend inflation when analyzing the conduct of monetary policy. Our findings also relate to the work of Hall and Kudlyak (2024) who suggest that the flat Phillips curve is an illusion caused by assuming that natural unemployment has little or no movement over the business cycle. Ascari and Fosso (2024) use a methodology similar to the one adopted in this paper to study the role of imported intermediate goods in explaining the lack of sensitivity of inflation to a business-cycle shock in the post-Millennial period. In line with our results, they find a stronger link between inflation and real activity compared to Angeletos et al. (2020). Our paper provides an explanation to reconcile these differences.

Our findings align with the evidence presented by Stock and Watson (2010) for the US and Smets (2010) for the Euro Area, indicating that the relationship between inflation and unemployment is more pronounced during recessions, when the cyclical component is likely more significant. Accounting for a relevant inflation-output relationship can improve inflation forecasts as shown in Stock and Watson (2008) for the US and Smets (2010) and Giannone et al. (2014) for the Euro Area. Similar results hold also in data-rich environments (Bańbura et al., 2015; Crump et al., 2025).

The literature has proposed various approaches to address the long-standing problem of dealing with trends in the data (Elliott, 1998). Using structural VARs, Fernald (2007) argues that the introduction of trend breaks in the level of labor productivity yields the robust empirical finding that hours worked fall on impact following a technology improvement. Relatedly, Bergholt et al. (2024) discuss how the poor identification of the deterministic component of VARs can translate into imprecise estimates of historical shock decompositions. Other recent approaches relay on the Beveridge-Nelson decomposition for the study of macroeconomic dynamics (Berger et al., 2023; Morley et al., 2024).

2 Motivating evidence

In this section, we provide motivating evidence for our subsequent empirical analysis. We aim to show that even a simple detrending exercise suggests that inflation and real activity are related over the business cycle once we control for their trends. The goal of this simple exercise is to set the stage for the more rigorous multivariate time series analysis that we pursue in the rest of the paper.

We consider inflation—measured as the log difference in the GDP deflator—and two measures of real economic activity—the log real GDP per capita and unemployment—over the period 1955:Q1-2019:Q4. Using a bandpass filter (Christiano and Fitzgerald, 2003), we extract the corresponding filtered time series over two frequency bands: the business cycle—defined as fluctuations between 6 and 32 quarters—and the medium cycle—defined as fluctuations between 8 and 30 years.

Panel (A) of Figure 1 plots business-cycle fluctuations for inflation and real GDP per

Figure 1: Inflation and real economic activity at business- and medium-cycle frequencies

(A) Business-cycle frequencies (6-32 quarters)

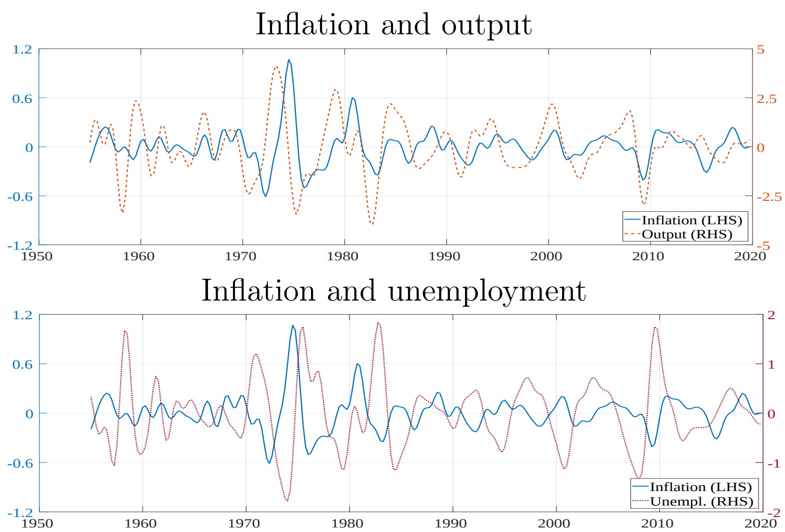

(B) Medium-cycle frequencies (8-30 years)

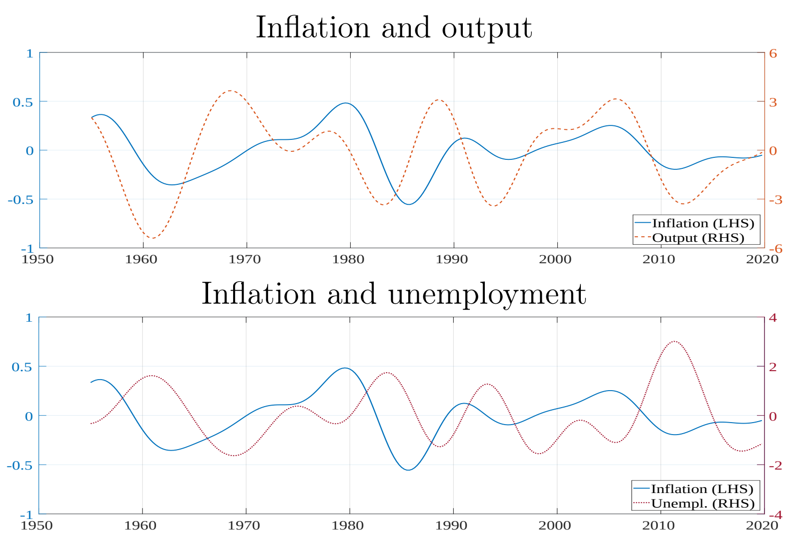

Notes: Inflation is defined as the log difference in the GDP deflator. For the two measures of real economic activity, we consider the log of real GDP per capita and unemployment. Data sample is from 1955:Q1 to 2019:Q4. Using the bandpass filter proposed by Christiano and Fitzgerald (2003), we extract the corresponding filtered time series over two frequency bands: the business cycle—defined as the period between 6 and 32 quarters—and the medium cycle—defined as the period between 8 and 30 years.

capita as well as for inflation and unemployment. Panel (B) of Figure 1 plots the corresponding medium-cycle fluctuations. The plots offer two main empirical facts. First, inflation appears correlated with both measures of real economic activity at both business-cycle and medium-cycle frequencies. As expected, inflation is positively correlated with real GDP per capita and negatively correlated with unemployment. Second, movements in both measures of real economic activity lead changes in the inflation dynamics at business-cycle and medium-cycle frequencies. High (low) levels of real GDP per capita (unemployment) are associated with subsequent high levels of inflation, and vice versa.

To formalize the notion that cyclical fluctuations in inflation commove with real activity, Table 1 reports the correlations between current filtered inflation and current and lagged filtered real GDP per capita and unemployment at time $(t-j)$ for $j=\{0,2,6,8\}$. We consider both business-cycle and medium-cycle frequencies. Over the business cycle, the positive (negative) correlation of inflation peaks with real GDP per capita (unemployment) lagged by four (two) quarters at about 0.47 (negative 0.44). Over the medium cycle, the correlation of inflation with real GDP per capita (unemployment) lagged by eight quarters peaks at 0.64 (negative 0.54). The results are very similar if instead of the band-pass filter, we employ the filtering approach advocated by Hamilton (2018).

The empirical evidence presented in this section motivates us to adopt a dynamic, multivariate framework that allows to study the relationship between inflation and real economic activity over the business cycle while controlling for low-frequency variation in those variables. We discuss the adopted framework in the next section.

3 The Trend-Cycle VAR

In this section, we present the TC-VAR used to model the joint dynamics of GDP, unemployment, the FFR, and inflation, as well as three expectations measures: the one-year-ahead unemployment expectations and the one- and ten-year-ahead inflation expectations.

Our baseline specification has four trends and six cycles. Unemployment u_t evolves as

$$u_t = \tau_{u,t} + c_{u,t}, \tag{1}$$

where $\tau_{u,t}$ and $c_{u,t}$ are the trend and cyclical components, respectively.

In line with the typical approach employed in structural macroeconomic models, we assume that real GDP per capita follows a process of the form $Y_t = \bar{Y}_t \exp(c_{y,t})$, where $c_{y,t}$ represents the cyclical movements of real GDP around a stochastic trend $\bar{Y}_t$. By definition,

Table 1: Correlations of inflation with lagged measures of real economic activity

Business-cycle frequencies (6-32 quarters)

	$j=8$	$j=6$	$j=4$	$j=2$	$j=0$
Output	0.15	0.36	0.47	0.42	0.22
Unemployment	0.09	-0.13	-0.34	-0.44	-0.36

Medium-cycle frequencies (8-30 years)

	$j=8$	$j=6$	$j=4$	$j=2$	$j=0$
Output	0.64	0.61	0.54	0.45	0.33
Unemployment	-0.54	-0.50	-0.44	-0.36	-0.25

Notes: Inflation is defined as the log difference in the GDP deflator. For the two measures of real economic activity, we consider the log level of real GDP per capita and unemployment. Data sample is from 1955:Q1 to 2019:Q4. Using the bandpass filter proposed by Christiano and Fitzgerald (2003), we extract the corresponding filtered time series over two frequency bands: the business cycle—defined as the period between 6 and 32 quarters—and the medium cycle—defined as the period between 8 and 30 years. We provide the correlations between current (filtered) inflation and current and lagged (filtered) levels of real GDP per capita and unemployment at time $(t-j)$ for $j = \{0, 2, 6, 8\}$.

real GDP per capita growth, $g_t \equiv \ln(Y_t/Y_{t-1})$, can then be expressed as

$$g_t = \tau_{g,t} + (c_{y,t} - c_{y,t-1}), \tag{2}$$

where $\tau_{g,t} \equiv \ln\left(\bar{Y}_t/\bar{Y}_{t-1}\right)$ is the trend component of GDP growth. It is worth emphasizing that we model the cycle in real GDP, as opposed to GDP growth, because what matters for inflation and unemployment dynamics is the output gap, not output growth.

One-year-ahead unemployment expectations share a common trend with realized unemployment, while also following a separate cyclical component

$$u_t^{e,1y} = \tau_{u,t} + c_{u,t}^e. \tag{3}$$

Assuming that the Fisher relation holds in the long run, the FFR evolves as

$$f_t = (\tau_{r,t} + \tau_{\pi,t}) + c_{f,t}, \tag{4}$$

where $\tau_{r,t}$ and $\tau_{\pi,t}$ are the trends of the real interest rate and inflation, respectively. Inflation and the one- and ten-year-ahead inflation expectations are decomposed as

$$\pi_t = \tau_{\pi,t} + c_{\pi,t}, \tag{5.1}$$
$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^e, \tag{5.2}$$
$$\pi_t^{e,10y} = \tau_{\pi,t} + \delta c_{\pi,t}^e + \eta_{\pi,t}^{e,10y}, \tag{5.3}$$

thus sharing a common trend $\tau_{\pi,t}$. We assume that the cyclical component for expected inflation, $c_{\pi,t}^e$, is shared across the one- and ten-year-ahead inflation expectation surveys. Because the ten-year inflation expectation $\pi_t^{e,10y}$ is fairly stable over time, we estimate its loading with the belief that it is less than one: $\delta < 1$. This parameterization is consistent with the definitions of one-year-ahead and ten-year-ahead inflation expectations that measure expected average inflation over the respective horizons. We allow for idiosyncratic errors in the ten-year-ahead inflation expectations, $\eta_{\pi,t}^{e,10y}$.

For ease of exposition, we collect observables and state variables in vectors

$$z_t = \left\{g_t, u_t, u_t^{e,1y}, f_t, \pi_t, \pi_t^{e,1y}, \pi_t^{e,10y}\right\}', \quad \tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}', \quad c_t = \left\{c_{y,t}, c_{u,t}, c_{u,t}^e, c_{f,t}, c_{\pi,t}, c_{\pi,t}^e\right\}',$$
$$\eta_t = \left\{\eta_{\pi,t}^{e,10y}\right\}, \quad \varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}\}', \quad \varepsilon_{c,t} = \left\{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,u,t}^e, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}, \varepsilon_{c,\pi,t}^e\right\}'.$$

The dynamics of the trend τ_t and cyclical component c_t are given as

$$\tau_t = \tau_{t-1} + \varepsilon_{\tau,t}, \tag{6.1}$$
$$c_t = \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + ... + \Phi_p c_{t-p} + \underbrace{\Phi_\tau(L_\tau^{-1}\varepsilon_{\tau,t}) + \tilde{\varepsilon}_{c,t}}_{\varepsilon_{c,t}}. \tag{6.2}$$

We allow for correlation in the innovations of the trend $\varepsilon_{\tau,t}$ and cycle $\varepsilon_{c,t}$ components:

$$\varepsilon_t = \begin{bmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \Sigma_\tau & L_\tau \Phi_\tau' \\ \Phi_\tau L_\tau' & \Phi_\tau \Phi_\tau' + \Sigma_c \end{bmatrix} \right), \tag{7}$$

where the matrices Σ_τ and Σ_c are conforming positive definite matrices associated with $\varepsilon_{\tau,t}$ and $\tilde{\varepsilon}_{c,t}$, respectively, which are orthogonal to each other; and L_τ is a lower triangular matrix resulting from the Cholesky decomposition of Σ_τ. We use the matrix Φ_τ to parameterize the correlation between cycle and trend innovations. This parameterization does not affect the identification strategy based on the max-share variance described below, but it facilitates inference. If we set $\Phi_\tau = \mathbf{0}$, we obtain a model in which trend and cycle innovations are uncorrelated, the typical setup studied in the literature. Morley et al. (2003) show that a necessary condition for the equivalence between the univariate trend-cycle representation with correlated shocks and an ARIMA process is that $p \geq 2$. This condition can be generalized to the multivariate case and it holds for the specification considered below.

The model can be recast into state-space form. Let n be the number of observables which can be decomposed into n_τ trends and n_c cycles, where $0 < n_\tau \le n$ and $0 < n_c \le n$.

Measurement equation. The measurement equation can be expressed as

$$z_t = \Lambda_x x_t + \Lambda_\eta \eta_t = \Lambda_\tau x_{\tau,t} + \Lambda_c x_{c,t} + \Lambda_\eta \eta_t, \tag{8}$$

where $x_t = \left\{x'_{\tau,t}, x'_{c,t}\right\}'$, $x_{\tau,t} = \tau_t$, $x_{c,t} = \left\{c'_t, c'_{t-1}, \ldots, c'_{t-(p-1)}\right\}'$, $\eta_t = \eta^{e,10y}_{\pi,t}$. The conformable matrices Λ_τ, Λ_c, and Λ_η map the trend, cycles, and observation error into the observables. The matrix Λ_c is defined as $\Lambda_c = [\Lambda_{c,0}, \ldots, \Lambda_{c,p-1}]$. Finally, $\Lambda_x = [\Lambda_\tau, \Lambda_c]$.

State-transition equation. The transition equation for the state x_t is given by

$$x_t = \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \tag{9}$$

or equivalently,

$$\begin{bmatrix} x_{\tau,t} \\ x_{c,t} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \Phi_c \end{bmatrix} \begin{bmatrix} x_{\tau,t-1} \\ x_{c,t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathcal{R}_c \end{bmatrix} \begin{bmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \end{bmatrix},$$

where

$$\Phi_c = \begin{bmatrix} \Phi_1 & \Phi_2 & \ldots & \Phi_p \\ \mathbf{I} & \mathbf{0} & \ldots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \vdots \\ \mathbf{0} & \ldots & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \mathcal{R}_c = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}.$$

The initial conditions are distributed as $x_{\tau,0} \sim \mathcal{N}\left(\underline{\tau}, \underline{V_\tau}\right)$ and $x_{c,0} \sim \mathcal{N}(0, \underline{V_c})$ where $\underline{V_\tau}$ is an identity matrix and $\underline{V_c}$ is the unconditional variances of $x_{c,0}$ based on (7).

4 Inference

We describe the data, the priors, and the methodology employed in our empirical analysis.

4.1 Data

We estimate the TC-VAR with tho lags ($p = 2$) using the following seven quarterly time series which are expressed at annualized rates: i) the growth rate of real GDP per capita, g_t; ii) unemployment, u_t; iii) the median four-quarter-ahead SPF unemployment expectations, $u_t^{e,1y}$; iv) the FFR f_t by treating observations at the zero lower bound as missing following Del Negro et al. (2017); v) inflation π_t, measured as the log difference in GDP deflator (PGDP); vi) the median four-quarter-ahead average PGDP inflation expectations, $\pi_t^{e,1y}$, from the SPF; vii) a measure of average ten-year-ahead inflation expectations, $\pi_t^{e,10y}$, which, following Del Negro and Schorfheide (2013), we construct by combining survey expectations on average ten-year-ahead CPI inflation from the SPF and Blue Chip Economic Indicators survey, and adjusting it for the historical difference between CPI and PGDP inflation. We use the period between 1955:Q1 and 1959:Q4 as pre-sample and estimate the TC-VAR over the period from 1960:Q1 to 2019:Q4. Appendix A provides the definitions, sources, and transformations for the data series.

4.2 Priors and initial conditions

For the initial conditions and priors, we mainly follow Del Negro et al. (2017). We consider standard priors for covariance matrices Σ_τ and Σ_c and for the VAR coefficients

$$p\left(\Sigma_\tau\right) = \mathcal{IW}\left(\kappa_\tau, (\kappa_\tau + n_\tau + 1)\underline{\Sigma_\tau}\right), \tag{10.1}$$

$$p\left(\Sigma_c\right) = \mathcal{IW}\left(\kappa_c, (\kappa_c + n_c + 1)\underline{\Sigma_c}\right), \tag{10.2}$$

$$p\left(\phi|\Sigma_c\right) = \mathcal{N}\left(\underline{\phi}, \Sigma_c \otimes \underline{\Omega}\right)\mathbf{I}\left(\phi\right), \tag{10.3}$$

where $\phi = vec\left[\Phi_1, \ldots, \Phi_p, \Phi_\tau\right]$ collects the vectorized matrices of VAR coefficients, $\mathcal{IW} = (\kappa, (\kappa + n + 1)\underline{\Sigma})$ corresponds to the inverse Wishart distribution with mode $\underline{\Sigma}$ and k degrees of freedom, and $\mathcal{I}(\phi)$ is an indicator function equal to 1 if the VAR in (9) is stationary.

The prior means for the initial conditions of the trends $\underline{\tau}$ are centered on the pre-sample means of the respective variables. Specifically, the annualized trends for real GDP growth and unemployment are set to 1% and 5%, respectively, while those for the real interest rate and inflation are centered at 0.1% and 2.5%.

We center the prior for the covariance matrix of the shocks to the trends Σ_τ on a diagonal matrix. The priors are such that the standard deviation of the expected change in annualized trend real GDP growth and unemployment is 1% over 10 years. For all the remaining variables, we assume a 1% standard deviation for the expected change in their trends over 5 years. As in Del Negro et al. (2020), we assume a tight prior by setting κ_τ to 100.

The prior for the covariance matrix of the shocks to the cyclical components Σ_c is also centered on a diagonal matrix. We calibrate the standard deviation of the shocks affecting the stationary components of annualized real GDP per capita and unemployment to 5% and 1.1%, reflecting their pre-sample standard deviations. The standard deviation of the shocks affecting the cycles of the nominal interest rate and inflation are also set to their pre-sample standard deviations of 0.8% and 1.5% respectively. We also need to specify the priors on the standard deviations of the cyclical component of the one-year-ahead unemployment expectations and of the common cyclical component of the two inflation expectation measures. As these surveys are unavailable for the pre-sample period, we set the standard deviations of unemployment and inflation expectations to 1.1% and 1.5%, respectively, matching the standard deviations of realized unemployment and inflation. As in Kadiyala and Karlsson (1997) and Giannone et al. (2015), we set $\kappa_c = n_c + 2$.

For the prior of the VAR coefficients ϕ in (10.3), we assume a conventional Minnesota prior, in line with Giannone et al. (2015). Because the cyclical component in (6.2) is assumed to be stationary, we center the prior for each variable's own lag on 0, rather than 1. Recall that we have augmented $\phi = vec\,[\Phi_1, \ldots, \Phi_p, \Phi_\tau]$ to include Φ_τ, the vector controlling the correlation between innovations to the trend and cycle components. Thus, the Minnesota prior also implies a prior centered on the case of no-correlation between cycle and trend innovations, a natural *a-priori* hypothesis. Three elements comprise $\underline{\Omega}$ in (10.3): the hyperparameters governing the overall tightness and lag decay of the Minnesota prior, set to 0.2 and 2, respectively, and the prior standard deviation for Φ_τ. For this last parameter, we choose a value that implies that *a-priori* the correlation in the innovations falls within ± 0.1 with 68% probability. Our baseline prior allows for correlation while limiting parameter uncertainty and overfitting. Looser priors produce stronger but more uncertain results (see the Online Appendix).

We employ a Gibbs sampling algorithm to draw from the posterior of the TC-VAR parameters. Details on the Bayesian algorithm and robustness exercises with respect to the priors are reported in the Online Appendix.

4.3 Identifying shocks that drive business-cycle fluctuations

To identify the business-cycle shock, we adopt the max-share identification strategy proposed by Faust (1998) and Uhlig (2003) and implemented in the frequency domain by Giannone et al. (2019b), Angeletos et al. (2020), and Basu et al. (2024), among others. The approach identifies a shock, or a combination of shocks, by finding the linear combination of the reduced-form residuals that maximizes its contribution to the volatility of a particular variable over a particular frequency band.

The TC-VAR delivers a decomposition between trends and cycles. Given that the cyclical components are already cleaned of low-frequency movements, we do not need to remove the low-frequency variation by using spectral analysis. Instead, we look for the combination of reduced-form shocks that explains the largest possible share of unemployment or output *cycles*, without having to take a stance on which frequencies correspond to the business cycle. In our baseline analysis, we ask how much the unemployment-identified or output-identified shock contributes to the volatility of the cyclical component of the other variables, with a special focus on inflation and inflation expectations. As a robustness check, we also ask if the results are sensitive to further removing high-frequency movements in the cycles. In this second case, we ask how much the unemployment-identified shock contributes to the volatility of the cyclical component of the other variables at frequencies that correspond to fluctuations with duration of at least 1.5 years. Thus, in this second methodology we take into account that the cyclical component of the variables could present some residual high-frequency movements that are not related to the business cycle.

Formally, in our TC-VAR, the vector of states x_t evolves as in (9), where the innovations of the trend $\varepsilon_{\tau,t}$ and cycle $\varepsilon_{c,t}$ components are correlated as described in (7). Therefore, the dynamics of the state vector can be expressed as

$$x_t = \Psi(L)\varepsilon_t \tag{11}$$

where $\Psi(L) = (\mathbf{I} - \Phi L)^{-1}\mathcal{R}$ is an infinite matrix polynomial in the lag operator L. To extract the dynamics of the cyclical components, we specify the mapping

$$c_t = M_z x_t = M_z \Psi(L)\varepsilon_t, \tag{12}$$

where $M_z \equiv [\mathbf{0}_{n_c\times n_\tau} \quad \mathbf{I}_{n_c\times n_c}]$. The residuals ε_t can then be written as

$$\varepsilon_t = SPu_t, \tag{13}$$

where S is the lower-triangular Cholesky decomposition of the covariance matrix associated with the innovations ε_t, the matrix P is an orthonormal matrix such that $P' = P^{-1}$, and u_t are i.i.d. "structural" shocks such that $E(u_t u_t') = \mathbf{I}$.

Combining equations (12) and (13), the cyclical components evolve as

$$c_t = [M_z \Delta(L)]\, Pu_t, \tag{14}$$

where $\Delta(L) = \Psi(L)S$ is an infinite matrix polynomial. Equation (14) expresses the cycles as a linear combination of the structural shocks while accounting for the possibility that shocks to the trends and cycles are correlated. This representation is used to implement the max-share identification strategy. The objective is to identify a shock by maximizing its contribution to the volatility of a particular variable over a given frequency band. In our model, the effect of any structural shock $j \in \{1, \ldots, dim(\varepsilon_t)\}$ on any cyclical component $k \in \{1, \ldots, n_c\}$ at any horizon $l \in \{0, 1, \ldots\}$ corresponds to $M_z^{[k]}\Delta_l\,\mathbf{p}$, where $M_z^{[k]}$ selects the $k-$th row of the matrix Δ_l and $\mathbf{p}$ denotes the $j-$th column of the orthonormal matrix P. Then, the responses of the $k-$th cyclical component to the $j-$th structural shock correspond to the sequence $\{M_z^{[k]}\Delta_l\,\mathbf{p}\}_{l=0}^{\infty}$. This implies that its contribution to the spectral density of the cyclical component over the frequency band $\omega \in [\underline{\omega}, \bar{\omega}]$ is

$$\overline{\mathbf{p}}\ \Xi(k, \underline{\omega}, \bar{\omega})\ \mathbf{p}, \tag{15}$$

where

$$\Xi(k, \underline{\omega}, \bar{\omega}) \equiv \int_{\omega \in [\underline{\omega}, \bar{\omega}]} \overline{M_z^{[k]}\Delta\left(e^{-i\omega}\right)} M_z^{[k]}\Delta\left(e^{-i\omega}\right)$$

is the covariance matrix of the overall volatility of the $k-$th cyclical component over the specified frequency band and for any vector α, we denote by $\overline{\alpha}$ its transpose.

Equation (15) can be written for any structural shock. The shock is therefore identified by maximizing its contribution to the volatility of the chosen cyclical component. This identification strategy corresponds to finding the vector $\mathbf{p}$ that maximizes (15) subject to the constraint $\overline{\mathbf{p}}\,\mathbf{p} = 1$. In the literature, it is well known that such vector is the eigenvector associated to the largest eigenvalue of the matrix $\Xi(k, \underline{\omega}, \bar{\omega})$ that can be obtained by the estimation of the TC-VAR.

In the context of our TC-VAR, the estimation of our model extracts cycles that are already cleaned of frequencies other than business cycles. So, when targeting a cyclical component for the shock identification in (15), we consider all frequencies or remove high frequencies. Instead, in the case of a VAR, the model does not automatically isolate the cycle component. We therefore follow Angeletos et al. (2020) and identify the combination of shocks that explain the largest share of variability at cycles between 6 and 32 quarters.

5 Results

In this section, we discuss the main results of the paper. We first present the decomposition of the variables in trends and cycles. We then analyze the extent to which the unemployment-identified shock accounts for the cyclical fluctuations in inflation and inflation expectations, and how it influences both real and nominal variables.

5.1 Estimated latent trends and cycles

Panel (A) of Figure 2 plots the data (orange lines) over the 1960-2019 period as well as the posterior median of their latent trends (blue lines) and the corresponding 90-percent posterior-coverage intervals (shaded blue area). Panel (B) of Figure 2 plots the posterior

Figure 2: Data, trends and cycles

(A) Trends

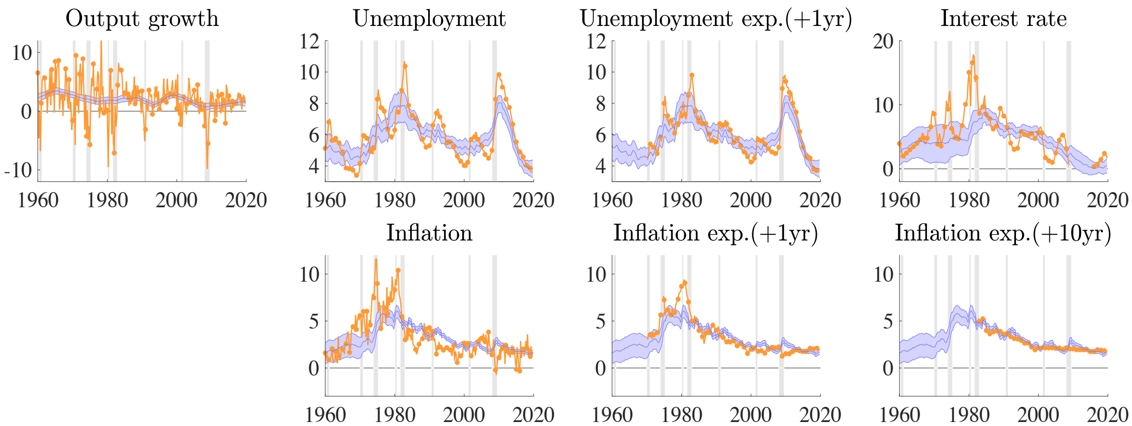

(B) Cycles

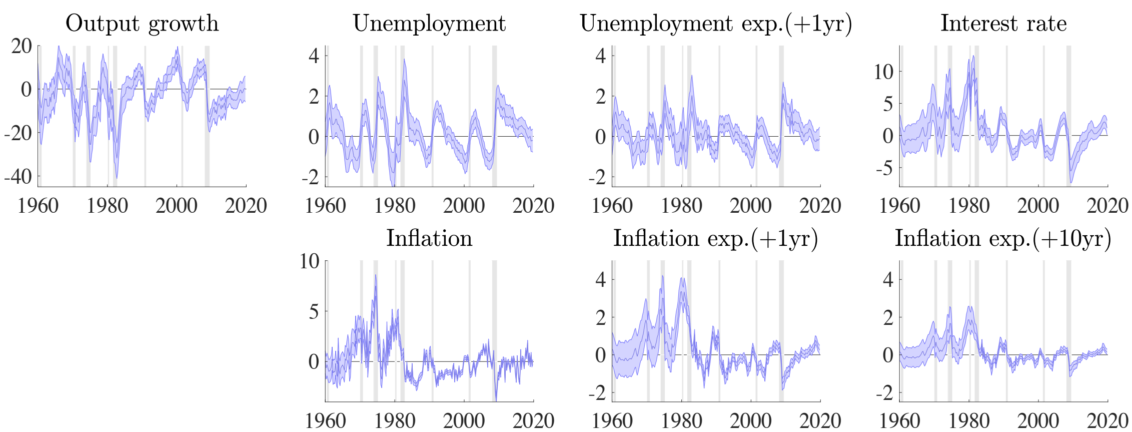

Notes: The figure plots the data (orange lines) used for the estimation of the TC-VAR over 1960-2019 period as well as the posterior median of their latent trends (blue lines) in panel (A) and latent cycles (blue lines) in panel (B) and the corresponding 90-percent posterior-coverage intervals (shaded blue areas). NBER recessions are denoted by shaded grey areas.

median of the latent cycles (blue lines) and the corresponding 90-percent posterior-coverage intervals (shaded blue area).

The results confirm some stylized facts about the US economy that are commonly accepted. In the 1960s and 1970s the US economy experienced an increase in trend inflation. This was possibly caused by the attempt of policymakers to counteract a break in productivity that manifested itself with an increase in natural unemployment or to partially accommodate the inflationary pressure resulting from a large increase in spending that occurred starting from the mid-1960s (Bianchi et al., 2023). These stylized facts are captured by an increase in the trend components of inflation and unemployment, and a slowdown in trend growth. Based on the median, trend growth moved from a peak of 3.4% in 1965:Q4, to a minimum of 1.7% at the end of the 1970s, while the trend component of unemployment rose from 4.8% in 1965:Q4 to a peak of 7.9% in 1981:Q1, a change in line with the results presented for natural unemployment in Crump et al. (2019). At the same time, trend inflation moved from a minimum of 2.4% in 1965:Q4 to a peak of 6.0% in 1981:Q1.

The appointment of Volcker marked a change in the conduct of monetary policy. Trend inflation declined, and so did long-term inflation expectations. The inflation trend and long-term inflation expectations largely coincide, with the corresponding cycle displaying relatively small fluctuations. Thus, including long-term inflation expectations helps to separate trend and cycle fluctuations. During the same years, the trend component of unemployment also declined steadily, while trend growth experienced an acceleration in the 1990s, consistent with the narrative associated with the productivity improvements brought forward by the Internet revolution. Finally, trend unemployment rose during the Great Financial Crisis to levels consistent with estimates of natural unemployment reported by Hall and Kudlyak (2024) and Crump et al. (2019) and based on the New Keynesian model of Galí et al. (2011). During that period, trend inflation stayed relatively stable due to the anchoring of long-term inflation expectations.

The behavior of the cycles reported in Panel (B) suggests a pattern consistent with a popular view of how the economy behaves over the business cycle. Unemployment increases during recessions and smoothly declines over time as the economy recovers. Based on a cursory look at the cycles, inflation seems to behave as the New Keynesian framework would suggest: Declining during a recession, when unemployment is high, and increasing during an expansion, when unemployment is low. This is especially visible when focusing on inflation expectations at the one-year horizon: its cyclical component behaves very much like inflation, but it is smoother. In what follows, we formalize this hypothesis.

5.2 Inflation and unemployment over the business cycle

We now move to formally study the relation between the real economy and inflation over the business cycle. We use the estimated TC-VAR to identify the unemployment-cycle shock using the method described in Subsection 4.3. Specifically, the shock is identified by maximizing its contribution to the volatility of the cyclical component of unemployment. As explained above, we consider two cases. In the first case, the shock is chosen to maximize

Table 2: Variance contributions of unemployment shock

All frequencies ($0-\infty$ quarters)

Unemployment	Output	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
71.6	57.9	68.8	63.9	39.4	54.7
[60.9,84.8]	[47.2,71.6]	[56.5,81.1]	[40.1,84.3]	[16.9,66.9]	[28.3,81.4]

All-but-short-run frequencies ($6-\infty$ quarters)

Unemployment	Output	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
72.4	57.9	72.4	64.9	44.4	57.8
[61.5,86.2]	[46.5,72.4]	[59.1,86.0]	[40.0,85.8]	[19.2,72.7]	[30.2,84.9]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment cycle. We consider two cases. In the first case, the shock is chosen to maximize the fraction of the volatility over all the frequencies of the cycle, while in the second case we exclude frequencies that imply cycles less than 1.5 years. We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of the cycle of all variables over the corresponding frequencies.

the fraction of the volatility over all the frequencies of the cycle, while in the second case we exclude frequencies that imply cycles less than 1.5 years.

The top panel of Table 2 reports the median and the 68% posterior-coverage interval for the contribution of the identified shock to the variance of the *cycle* of all the other variables. In the second panel, we repeat the exercise by excluding frequencies that imply cycles less than 1.5 years. Not surprisingly, the shock can explain a large share of the fluctuations of the unemployment cycle. However, the shock can also explain a sizable fraction of the cyclical component of inflation. In the baseline scenario, the unemployment-identified shock can explain nearly 40% of the inflation cycle. When excluding cycles shorter than 1.5 years, the unemployment-identified shock explains more than 44% of inflation cyclical variability. These are large shares when considering that the unemployment shock explains around 72%, rather than the entirety, of unemployment cyclical fluctuations.

The results are even stronger when focusing on the cyclical component of inflation expectations: approximately 55% in the baseline scenario—or 58%, in the alternative scenario—of the business-cycle variability of inflation expectations is explained by the unemployment-identified shock. Given that the cycle of inflation expectations appears to be a smoother version of the cycle of realized inflation, this result corroborates the finding that inflation moves in a way consistent with the New Keynesian framework over the business cycle.

Table 3: Variance contributions of GDP business-cycle shock

All frequencies ($0 - \infty$ quarters)

Unemployment	Output	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
61.9	66.0	58.2	55.3	49.0	58.3
[45.1,77.9]	[57.5,76.7]	[41.0,74.3]	[11.8,91.7]	[12.2,78.4]	[12.9,89.9]

All-but-short-run frequencies ($6 - \infty$ quarters)

Unemployment	Output	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
61.3	66.9	59.8	58.9	54.9	64.1
[44.3,79.1]	[57.6,77.2]	[42.2,78.0]	[13.2,92.6]	[13.9,83.0]	[15.7,93.1]

Notes: The shock is identified by maximizing its contribution to the volatility of the cyclical component of real GDP (in loglevels). We consider two cases. In the first case, the shock is chosen to maximize the fraction of the volatility over all the frequencies of the cycle, while in the second case we exclude frequencies that imply cycles less than 1.5 years. We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of the cycle of all variables over the corresponding frequencies.

When comparing the results of the baseline and alternative cases, we find that the contribution for inflation and inflation expectations goes visibly up when removing the short cycles. This implies that there is likely to be some residual high-frequency variation in these variables that it is not related to the business cycle. In addition, the shock explains about 64% of the volatility of the nominal interest rate cycle over both frequency bands. Combined with the previous two findings, this result implies that the shock also explains the volatility of the cyclical component of the real interest rate, defined as the difference between the FFR and expected inflation.

Finally, the shock explains a large portion of the volatility of the GDP cycle over all frequencies and also when excluding cycles shorter than 6 quarters. The shock also explains a share of the volatility of the cyclical component of expected unemployment similar to the corresponding share for realized unemployment. These findings support the evidence that the identified shock is the main driver of the business-cycle.

The results are similar when using GDP to identify the business-cycle shock. Table 3 reports the median contribution—and the corresponding 68-percent posterior-coverage interval—of the shock identified targeting the cycle of real GDP to the variance of the *cycles* of all the variables. As before, we consider two cases. In the first case, we identify

Figure 3: Impulse responses to unemployment shock

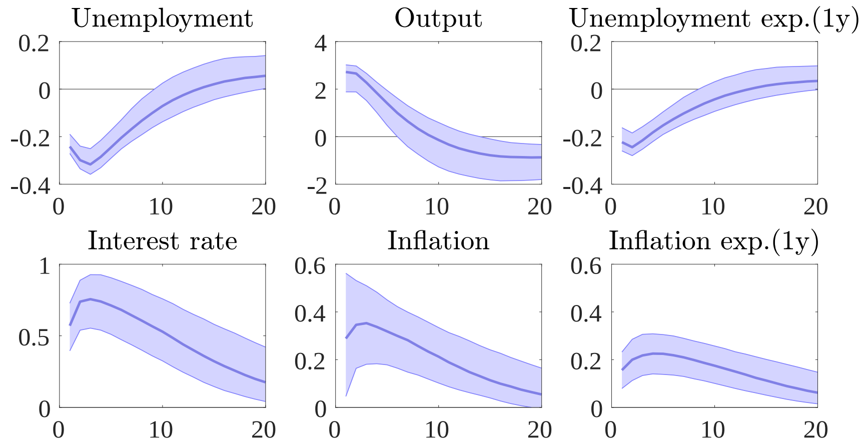

Notes: The figure shows the response to the unemployment-identified shock of the cyclical component of all the variables over a period of 20 quarters. The figure plots the posterior median (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area).

the shock and compute its contributions based on all frequencies of the cycles. In the second case, we consider frequencies that imply fluctuations of at least 1.5 years. As in the case of unemployment, the shock can explain a large share of the cyclical fluctuations of real GDP. In line with the results for the unemployment-identified shock in Table 2, the shock also explains a sizable fraction of the cyclical component of realized and expected inflation as well as realized and expected unemployment. Additionally, the contribution of the output-identified shock for the variability of the realized and expected inflation increases noticeably when movements at frequencies shorter than 1.5 years are excluded. Given that the results are similar across the two specifications, in the rest of the paper we focus on the unemployment-identified shock.

Figure 3 plots the median response—and corresponding 68-percent posterior-coverage intervals—of each cyclical component to the unemployment-identified shock over a period of 20 quarters. The resulting interpretation of the unemployment-identified shock is in line with a demand shock in a canonical New-Keynesian model. Unemployment decreases about 0.2% percent on impact and subsequently returns to its initial level after about 3 years. The response of one-year-ahead unemployment expectations is quantitatively and qualitatively similar to that of unemployment. Considering real GDP, the shock causes an increase by nearly 3% on impact, and its effect gradually vanishes after about 2 years. When considering the nominal variables, the shock leads to a contemporaneous increase of about 0.3% and 0.2% in the cyclical components of realized and expected inflation, respectively. Both variables slowly decline back to their initial levels over about 5 years. Finally, the nominal interest rate peaks at nearly 0.7% after about a year and gradually returns to its initial value thereafter.

To summarize, the responses of the real side of the economy are consistent with the findings of Angeletos et al. (2020) who point to the presence of a main shock driving the fluctuations of real economic activity over the business cycle. However, differently from their findings, the unemployment shock that we identify has significant effects also on the nominal side of the economy. In the Online Appendix, we consider a series of robustness checks that show how these results are robust to different model specifications.

6 VARs and the link between inflation and real activity

In this section, we show that the use of a standard VAR, as opposed to a TC-VAR, can lead to very different conclusions about the link between real activity and inflation over the business cycle. We check whether this discrepancy disappears when imposing alternative priors or when considering different variables. We find that while long-run priors help, the results are still very different from our baseline analysis based on the TC-VAR.

We proceed to estimate a VAR with two lags. Given that a VAR does not automatically separate trends from cycles, we follow Angeletos et al. (2020) and identify the business-cycle unemployment shock in the frequency domain using the method described in Subsection 4.3. Specifically, we look for the combination of shocks that explains the largest share of unemployment fluctuations for cycles between 1.5 and 8 years. We then compute the contribution of the shock to the volatility of inflation over the same frequencies.

To understand the role of different priors, we consider three specifications. In the first case, we use a flat prior on all VAR parameters. In the second case, we follow Angeletos et al. (2020) and use an optimized Minnesota prior in which hyperparameters are optimized for shrinkage. In the third specification, we combine a long-run prior *à la* Giannone et al. (2019a) with a Minnesota prior and jointly optimize them. The cointegrating relationships that we assume in this case are in line with those of Giannone et al. (2019a) and are described in the Online Appendix. We allow for separate shrinkage for each active row of the matrix that captures the cointegrating relationship among the variables.

To understand the role of different datasets, we consider two different sets of variables. In the first case, we use exactly the same variables used in the TC-VAR, while in the second case we use the same variables used in Angeletos et al. (2020). The most noticeable difference is that the dataset used in this paper also includes expectations. Consistently

Table 4: Shocks targeting u_t: VAR analysis

	Baseline data		Angeletos et al. (2020) data	
	Unemployment	Inflation	Unemployment	Inflation
Flat prior	72.0 $[62.7, 81.0]$	13.6 $[6.4, 24.3]$	71.2 $[63.7, 78.8]$	10.4 $[4.4, 19.1]$
Optimized Minnesota prior	72.9 $[64.6, 80.7]$	12.8 $[6.8, 21.6]$	73.0 $[66.5, 79.4]$	7.7 $[3.5, 13.3]$
Optimize Minnesota + Long-run prior	78.1 $[69.7, 85.0]$	16.8 $[8.4, 26.4]$	91.9 $[88.7, 94.3]$	10.3 $[4.9, 16.8]$

Notes: We identify the max-share unemployment shock using a VAR based on frequencies corresponding to cycles between 6 and 32 quarters. Optimizing hyperparameters for the long-run prior involves estimating the degree of shrinkage for each active row of H individually. The first two columns use the same data and estimation sample as the baseline TC-VAR. In this case, the VAR is estimated in state-space form to handle missing observations, such as survey expectations largely absent in the early part of the estimation sample and nominal interest rates during the zero lower bound period. The third and fourth columns report results using the same variables of Angeletos et al. (2020).

with the approach used for the TC-VAR, when using the baseline dataset, the VAR is estimated in state-space form to handle missing observations, such as survey expectations in the early part of the sample and the FFR during the zero lower bound period.

The first two columns of Table 4 report the contributions of the max-share shock for the cyclical fluctuations of unemployment and inflation when using the TC-VAR data series. The last two columns report results using the same series used in Angeletos et al. (2020). Across the different rows, we assess how the results change based on the priors. The results indicate that when using a VAR, the contribution of the business-cycle shock to the cyclical fluctuations of inflation is greatly reduced with respect to what obtained with a TC-VAR. Comparing the first two columns, with the third and fourth columns, we notice that when using the same series of Angeletos et al. (2020), the business-cycle shock has a negligible impact on inflation. The contribution is larger when using the baseline dataset that includes

Figure 4: Impulse responses to unemployment shock

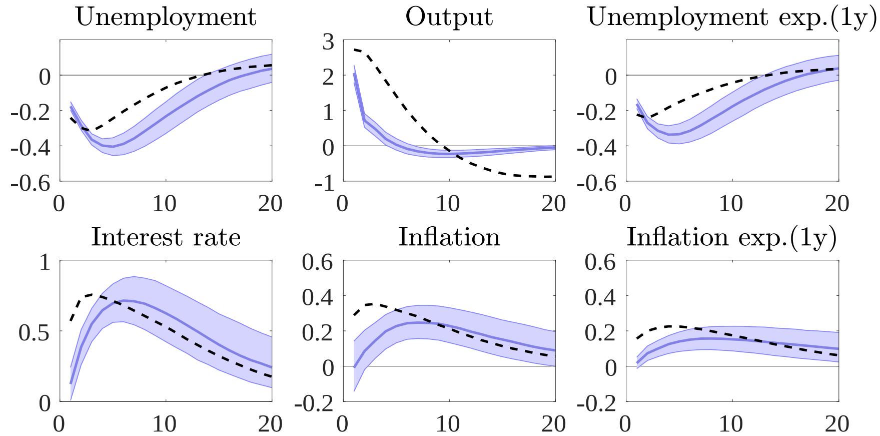

Notes: The figure shows the response of all variables to the max-share unemployment shock identified using a VAR and over frequencies corresponding to cycles between 6 and 32 quarters. The VAR is estimated using the same data and estimation sample as the baseline TC-VAR, an optimized Minnesota priors as in Angeletos et al. (2020), and relying on a state-space form to handle missing observations, such as survey expectations largely absent in the early part of the estimation sample and nominal interest rates during the zero lower bound period. The figure plots the posterior median (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). The figure also reports the posterior median contributions (dashed black lines) for the baseline TC-VAR and plotted in Figure 3.

expectations, but still lower than when using the TC-VAR.

For both datasets, the results modestly improve when combining the Minnesota priors with long-run priors, but the contribution to inflation cycles is still around a third of what obtained with the TC-VAR. The contribution of the business-cycle shocks to the cyclical fluctuations of inflation is visibly reduced under the Optimized Minnesota priors used in Angeletos et al. (2020). Figure 4 reports impulse responses for this last case (solid lines). In line with the results shown in Table 4, a larger response in unemployment corresponds to a smaller response in inflation compared to the TC-VAR (dashed lines).

In summary, a TC-VAR and a VAR yield significantly different outcomes when examining whether a shock that largely influences cyclical fluctuations in unemployment also contributes substantially to cyclical volatility in inflation. As previously discussed, the TC-VAR has logical advantages for this type of analysis because it is designed to isolate cyclical fluctuations. Nonetheless, it would be beneficial to demonstrate that the TC-VAR has good properties for the data at hand beyond the specific research question addressed in this paper. With this goal in mind, we assess whether the TC-VAR provides better forecasts for unemployment, inflation, and GDP growth than the VAR.

We recursively estimate the TC-VAR and different VAR specifications varying the sample

Figure 5: Relative log-determinant differentials across models

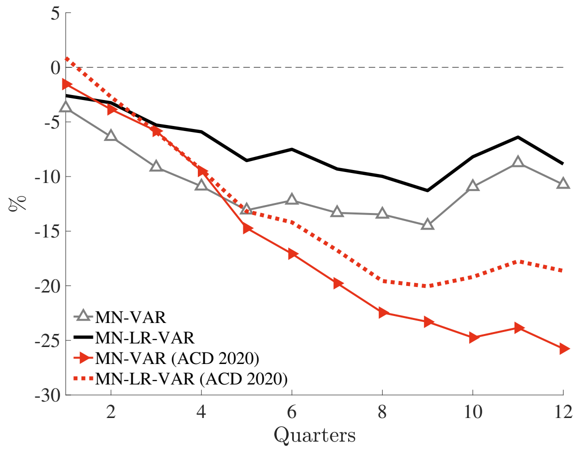

Notes: The figure reports the log-determinant differentials of different models with respect to the baseline TC-VAR. A negative value favors the baseline TC-VAR over alternative models. Forecasts for inflation, unemployment, and GDP growth are produced using recursive sample from 1960:Q1-2000:Q2 to 1960:Q1-2016:Q4. We consider VARs with Minnesota (MN) and Minnesota and Long Run (MN-LR) priors. We report results for the baseline dataset and an alternative dataset based on Angeletos et al. (2020).

from 1960:Q1-2000:Q2 to 1960:Q1-2016:Q4. We then compute out-of-sample forecasts up to 12 quarters ahead. To summarize the multivariate forecast performance, we use the scaled log-determinant differential (LDD) of the forecast error covariance matrix (Doan et al., 1984; Schorfheide and Song, 2015):

$$\text{LDD} \equiv \frac{100}{2N_v}\left[\ln\left(\left|\frac{1}{\overline{T}-\underline{T}}\sum_{t=\underline{T}}^{\overline{T}}\zeta_t\zeta_t'\right|\right)^{\text{TC-VAR}} - \ln\left(\left|\frac{1}{\overline{T}-\underline{T}}\sum_{t=\underline{T}}^{\overline{T}}\zeta_t\zeta_t'\right|\right)^{\text{Alt.}}\right], \tag{16}$$

where ζ_t represents forecast errors, the factor 2 converts mean-squared errors into root mean square errors, N_v averages across variables in ζ_t, and the factor 100 converts the differential into percentages. The metric (16) compares the baseline TC-VAR with alternative models with $\underline{T} = 2000{:}Q2$, $\overline{T} = 2016{:}Q4$, and $N_v = 3$. For a given horizon, a negative value of the LDD implies that the overall forecasting power of the corresponding model is inferior to the forecasting power of the TC-VAR.

Figure 5 reports the results. All different VAR specifications produce forecasts that are overall inferior to the TC-VAR forecasts. Furthermore, the ranking in the forecasting performance of the different VAR specifications aligns with their ability to recover a significant relation between the unemployment and inflation cycles. The forecasting performance of the VAR is particularly weak when using the same data used in Angeletos et al. (2020).

The forecasting performance improves when using the dataset that includes expectations. However, in all cases, the TC-VAR returns better forecasts at all horizons. The best VAR specification still implies a 10% loss in forecasting power with respect to the TC-VAR.

7 Reconciling VAR and TC-VAR Results

In this section, we first provide theoretical arguments that explain the large discrepancy between the TC-VAR and VAR results. We then consider a series of Monte Carlo simulations to elucidate the theoretical arguments.

7.1 Trend-Cycle models and their VAR(∞) representation

In what follows, we show that a fixed-coefficient VAR estimated over a period of time that presents low-frequency variation is misspecified, if the goal is trying to assess the commovement at business-cycle frequency. The misspecification problem associated with the use of a VAR to describe a data generating process characterized by both low- and high-frequency movements cannot be easily resolved.[1] This is because the VAR parameters need to account at the same time for the low-frequency and business-cycle frequency variation observed in the data with a finite number of observations. Even if an econometrician could correctly reconstruct the VAR representation of the TC-VAR, the parameter estimates of the misspecified model would confound low-frequency movements associated with the trend with those at business-cycle frequencies related to the cycle. Moreover, the reduced-form innovations would capture not only the innovations to the latent persistent and stationary components, but also the estimation errors associated with the latent components.

Our goal is to map the state-space representation introduced in Section 3 into a VAR. In doing so, we follow the approach proposed in Fernández-Villaverde et al. (2007). We provide the key steps, while we leave the details for the Online Appendix. For convenience, we report here the state-space representation of our baseline TC-VAR model:

$$z_t = \Lambda_x x_t, \tag{17.1}$$

$$x_t = \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \tag{17.2}$$

where $\varepsilon_t = Qw_t$, $E(w_t w_t') = \mathbf{I}$, and $E(\varepsilon_t \varepsilon_t') = \Sigma$. For all the specifications considered

[1] Watson (1986) discusses the equivalence between an unobserved component model and its autoregressive, integrated, moving average (ARIMA) representation, thus pointing to the misspecification problem characterizing an AR representation of a TC-AR model.

in our analysis, the overall number of shocks of the TC-VAR is strictly larger than the number of observables. Equivalently, $dim\,(w_t) = (n_\tau + n_c) > n = dim\,(z_t)$. As a result, the *'poor man's invertibility condition'* proposed in Fernández-Villaverde et al. (2007) cannot be tested because it requires the number of shocks and observables to coincide. We therefore seek to find the *'innovation representation'* of (17).

Because the innovation representation results from the application of the Kalman filter to the state-space representation, we first ensure the suitability of the filter for our purpose and more specifically its asymptotic stability and convergence. Clearly, these properties of the filter depend on the properties of (17) and should not be taken for granted in our setup: In the transition equation (17.2), the cyclical components are assumed to be stationary, while trends follow unit-root processes. We follow Anderson and Moore (1979) who suggest to verify two conditions: i) the pair $(\Phi, \mathcal{R}Q)$ is stabilizable; ii) the pair (Φ', Λ_x') is detectable—or equivalently, the pair (Φ', Λ_x') is stabilizable. Both conditions are satisfied for each draw from the posterior of the TC-VAR presented in Section 3.

Having verified the suitability of the Kalman filter, we derive the innovation representation. We first express (17) as

$$x_{t+1} = Ax_t + Bw_{t+1}, \tag{18.1}$$

$$z_{t+1} = Cx_t + Dw_{t+1}, \tag{18.2}$$

where $A = \Phi$, $B = \mathcal{R}Q$, $C = \Lambda_x A$, $D = \Lambda_x B$ and $E\,(w_t w_t') = \mathbf{I}$. Defining the linear projection of x_t on $z^t \equiv \{z_j\}_{j=1}^{t}$ as $\hat{x}_t \equiv E\,(x_t|z^t)$, the one-step-ahead error associated with the forecast of z_{t+1} as $\hat{D}_{t+1}\nu_{t+1}$, and the term updating the filtered state for the next period $\hat{x}_{t+1}$ as $\hat{B}_{t+1}\nu_{t+1}$, the application of the Kalman filter to (18) delivers the innovation representation

$$\hat{x}_{t+1} = A\hat{x}_t + \hat{B}_{t+1}\nu_{t+1}, \tag{19.1}$$

$$z_{t+1} = C\hat{x}_t + \hat{D}_{t+1}\nu_{t+1}, \tag{19.2}$$

where $x_0 \sim (\hat{x}_0, \Omega_0)$ and $\nu_t \sim (\mathbf{0}, \mathbf{I})$. Under this representation, the number of shocks and observables coincide. Because the innovation ν_t is fundamental for z_t by definition, it is uncorrelated with z_{t-s} and ultimately ν_{t-s} for any $s \geq 0$.

The innovation representation shows that, with a finite sample $\{z_t\}_{t=1}^{T}$, $T < \infty$, it is *not* possible to derive a VAR representation because the matrices $\hat{B}_t$ and $\hat{D}_t$ depend on time t. As a result, we consider the limit case for T approaching infinity. Because the asymptotic stability and convergence of the Kalman filter hold, the matrices $\hat{B}_t$ and $\hat{D}_t$ also converge to their time-invariant counterparts $\hat{B}$ and $\hat{D}$. We then derive the VAR(∞) representation

$$z_{t+1} \;=\; \sum_{s=0}^{\infty} C\left(A - \hat{B}\hat{D}^{-1}C\right)^{s}\hat{B}\hat{D}^{-1}z_{t-s} + \hat{D}\nu_{t+1}. \tag{20}$$

Equation (20) leads us to three significant conclusions. First, the state-space representation of the TC-VAR in (17) maps into the infinite-order VAR representation in (20) under the assumption that infinite data are available. As a result, a finite-order VAR with finite data cannot capture the dynamics described by the decomposition of the observables z_t into trends and cycles. Second, even if infinite data were available, equation (20) clarifies that estimates of the autoregressive parameters associated with the VAR(∞) representation confound movements of z_{t+1} that are driven by both the trend and cycle. Equivalently, the VAR(∞) representation cannot disentangle movements of z_{t+1} at low frequencies from those at cyclical frequencies. Finally, as shown in Fernández-Villaverde et al. (2007), the innovations associated with the VAR(∞) representation, $\hat{D}\nu_{t+1}$, capture not only the shocks to the latent trends and cycles, Dw_{t+1}, but also the error associated with the estimate of those latent components, $C\left(x_t - \hat{x}_t\right)$. In the Online Appendix, we provide an analytical example based on Stock and Watson (2007) to further illustrate these points.

These results are not meant to establish the unconditional superiority of a TC-VAR over a VAR. Over the past four decades, economists have used VARs as extremely flexible econometric models capable of uncovering a variety of enlightening empirical results. However, for the specific question of assessing the strength of the relation between inflation and real activity over the business cycle, a TC-VAR appears to be a more effective tool.

7.2 Monte Carlo simulations of a bivariate TC-VAR

To illustrate the implications of the theoretical results of Subsection 7.1, let us assume that the data generating process for unemployment and inflation, $z_t = \{u_t, \pi_t\}'$, is described by the measurement equation $z_t = \tau_t + c_t$ and the transition equations:

$$\tau_t \;=\; \tau_{t-1} + \varepsilon_{\tau,t}, \tag{21.1}$$

$$c_t \;=\; \Phi_1 c_{t-1} + \varepsilon_{c,t}, \tag{21.2}$$

where $\tau_t = \{\tau_{u,t}, \tau_{\pi,t}\}'$, $c_t = \{c_{u,t}, c_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,u,t}, \varepsilon_{\tau,\pi,t}\}'$ and $\varepsilon_{c,t} = \{\varepsilon_{c,u,t}, \varepsilon_{c,\pi,t}\}'$. In this example, we assume that (21.2) is

$$\begin{bmatrix} c_{u,t} \\ c_{\pi,t} \end{bmatrix} = \begin{bmatrix} \rho_{uu} & 0 \\ -(1-\rho_{\pi\pi})\kappa & \rho_{\pi\pi} \end{bmatrix} \begin{bmatrix} c_{u,t-1} \\ c_{\pi,t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_{c,u,t} \\ \varepsilon_{c,\pi,t} \end{bmatrix},$$

implying that, while the cyclical component of unemployment only depends on its lag, the cyclical component of inflation depends on its lag as well as on the lagged cyclical component of unemployment. We assume that the shocks are *iid*:

$$\varepsilon_t = \begin{bmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \Sigma_\tau & \mathbf{0} \\ \mathbf{0} & \Sigma_c \end{bmatrix} \right), \quad \Sigma_\tau = \begin{bmatrix} \sigma_{\tau,u} & \mathbf{0} \\ \mathbf{0} & \sigma_{\tau,\pi} \end{bmatrix}, \quad \Sigma_c = \begin{bmatrix} \sigma_{c,u} & \mathbf{0} \\ \mathbf{0} & \sigma_{c,\pi} \end{bmatrix}.$$

Within this framework, we consider four cases of interest. In the first case, unemployment does not feature low-frequency variation and it is only driven by its business-cycle movements, while inflation also features changes in the trend. We consider different degrees of low-frequency variation for inflation, while always maintaining the assumption that its cycle is only driven by the unemployment cycle. Specifically, for unemployment, we set the autoregressive parameter ρ_{uu} to 0.95, normalize the standard deviation of the shock to the cycle $\sigma_{c,u}$ to 1, and turn off the shocks to its trend ($\sigma_{\tau,u}$=0). For inflation, we assume $\rho_{\pi\pi} = 0$, $\kappa = 1$, and $\sigma_{c,\pi} = 0$. We then consider three values for the standard deviation of the shock to the inflation trend ($\sigma_{\tau,\pi}$). Relative to the standard deviation of the unemployment-cycle shock, the standard deviation of the inflation trend shock is one order of magnitude smaller ($\sigma_{\tau,\pi} = 0.1$), equivalent ($\sigma_{\tau,\pi}$=1) or twice as large ($\sigma_{\tau,\pi}$=2).

For each calibration, we produce long Monte Carlo simulations that we use to fit a VAR, identify the shock targeting unemployment at business-cycle frequencies, and compute the median contribution of the identified shock to the variance of both series over the same frequencies.[2] Table 5 reports the median and 68% intervals of the median contributions of the identified shock. As expected, the identified shock fully explains the cyclical movements of unemployment regardless of the chosen calibration. However, the explanatory power of the unemployment-identified shock for inflation varies with the extent of low-frequency fluctuations in inflation. The higher the low-frequency variation, the lower the degree to which the unemployment-identified shock explains business-cycle movements in inflation. This conclusion holds even with long data samples and in presence of strong assumptions about the cyclical relationship between unemployment and inflation. In line with our

[2]We generate 500 Monte Carlo simulation of 50,000 observations of which we keep the last 1,000. We choose the lags with the lowest Bayesian Information Criterion (BIC), use a Minnesota prior as in Angeletos et al. (2020), and keep the last 1,000 of 50,000 draws to identify the shock.

Table 5: Variance contribution of unemployment shock (data simulated with $\sigma_{\tau,u} = 0$)

$\sigma_{\tau,\pi}$	Unemployment	Inflation
0.1	100.0	98.8
	$[100.0, 100.0]$	$[98.7, 98.9]$
1	100.0	17.9
	$[100.0, 100.0]$	$[15.3, 20.4]$
2	100.0	4.5
	$[100.0, 100.0]$	$[3.3, 6.1]$

Notes: The shock is identified by maximizing its contribution to the volatility of unemployment over business-cycle frequencies (6-32 quarters). We report the median and the corresponding 68-percent interval of the median contributions of the shock to the variance of all variables over the same frequencies. To simulate the data, we use the following calibrations. For unemployment, we set $\rho_{uu} = 0.95$, $\sigma_{c,u} = 1$ and $\sigma_{\tau,u} = 0$. For inflation, we set $\rho_{\pi\pi} = 0$, $\kappa = 1$, and $\sigma_{c,\pi} = 0$ and $\sigma_{\tau,\pi} = \{0.1, 1, 2\}$.

results above, the identification of the shock at business-cycle frequencies does not succeed in extracting the cyclical relationship between unemployment and inflation because the fixed-coefficient VAR fails to separate cycle and trend innovations in inflation.

The results of the other three cases are briefly presented here, while we discuss the details in the Online Appendix. In the second exercise, we introduce low-frequency movements in unemployment, while inflation only follows the unemployment cycle. In the third case, the underlying true data generating process features trends in both inflation and unemployment. The results for these two cases confirm the importance of controlling for low-frequency movements in both inflation and unemployment to appropriately extract their business-cycle relationship. In the fourth and last case, unemployment and inflation cycles are assumed to be unrelated, and we ask whether the TC-VAR can correctly recover the truth. To this end, we simulate a model with independent persistent processes for inflation and unemployment and then fit the TC-VAR on the simulated data. We find that a bivariate TC-VAR correctly recovers the disconnect between the two simulated series. Thus, the fact that the model has the flexibility of separating trends from cycles does not mean that it would automatically try to recover commovement between the two cycles.

8 Conclusions

In this paper, we adopt a TC-VAR to study the relation between inflation and real activity over the business cycle. A TC-VAR has the virtue of removing low-frequency movements in inflation and real activity that can contaminate inference when using a fixed-coefficient

VAR. We show that at business-cycle frequencies, fluctuations of inflation are in fact related to movements in real activity. We explain why evidence based on VARs can be misleading. We see three directions for future research: first, to investigate the drivers of low-frequency movements in the macroeconomy; second, to allow for the possibility of multiple shocks at business-cycle frequencies to separate demand-driven and supply-driven fluctuations; third, to apply the same methods to study the cyclical behavior of other key macroeconomic variables that feature trends, such as the employment-to-population ratio (Fukui et al., 2023) and the share of employment in middle-skilled jobs (Jaimovich and Siu, 2020).

The literature has proposed alternative ways to control for changes in the statistical properties of macroeconomic variables. A popular approach involves the use of models with smoothly time-varying parameters (Cogley et al., 2010; Galí and Gambetti, 2009; Debortoli et al., 2020). We opt for a time-invariant TC-VAR because it is well-suited to separating trends and cycles and conducting a max-share analysis of cyclical fluctuations. A natural extension would allow for time-varying volatility (Stock and Watson, 2007; Galí and Gambetti, 2009, Debortoli et al., 2020; Johannsen and Mertens, 2021). Such an extension would make the application of the max-share approach more subtle, as the covariance structure would change over time. The max-share could be computed under an "anticipated utility assumption" in which future parameter instability is disregarded. Alternatively, parameter instability could be modeled with regime changes (Sims and Zha, 2006; Bianchi, 2013; Bianchi and Ilut, 2017). In this case, the spectrum could be calculated taking into account the possibility of regime changes (Bianchi, 2016). We regard this as a promising direction for future research.

References

Anderson, B. D. O. and Moore, J. B. (1979). *Optimal Filtering.* Prentice-Hall, Inc., New Jersey.

Angeletos, G.-M., Collard, F., and Dellas, H. (2018). Quantifying Confidence. *Econometrica*, 86(5):1689–1726.

Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business-Cycle Anatomy. *American Economic Review*, 110(10):3030–3070.

Ascari, G. and Fosso, L. (2024). The International Dimension of Trend Inflation. *Journal of International Economics*, 148(103896).

Ascari, G. and Sbordone, A. M. (2014). The Macroeconomics of Trend Inflation. *Journal of Economic Literature*, 52(3):679–739.

Bańbura, M., Giannone, D., and Lenza, M. (2015). Conditional forecasts and scenario analysis with vector autoregressions for large cross-sections. *International Journal of Forecasting*, 31(3):739–756.

Basu, S., Candian, G., Chahrour, R., and Valchev, R. (2024). Risky Business Cycles. NBER Working Paper No. 28693.

Beaudry, P., Galizia, D., and Portier, F. (2020). Putting the Cycle Back into Business Cycle Analysis. *American Economic Review*, 110(1):1–47.

Beaudry, P. and Portier, F. (2013). Understanding Noninflationary Demand-Driven Business Cycles. *NBER Macroeconomics Annual*, 28(1):69–130.

Berger, T., Morley, J., and Wong, B. (2023). Nowcasting the output gap. *Journal of Econometrics*, 232(1):18–34.

Bergholt, D., Canova, F., Furlanetto, F., Maffei-Faccioli, N., and Ulvedal, P. (2024). What Drives the Recent Surge in Inflation? The Historical Decomposition Roller Coaster. Norges Bank Working Paper 7/2024.

Bianchi, F. (2013). Regime Switches, Agents' Beliefs, and Post-Wold War II U.S. Macroeconomic Dynamics. *Review of Economic Studies*, 80(2):463–490.

Bianchi, F. (2016). Methods for Measuring Expectations and Uncertainty in Markov-Switching Models. *Journal of Econometrics*, 190(1):79–99.

Bianchi, F., Faccini, R., and Melosi, L. (2023). A fiscal theory of persistent inflation. *The Quarterly Journal of Economics*, 138(4):2127–2179.

Bianchi, F. and Ilut, C. (2017). Monetary/Fiscal Policy Mix and Agents' Beliefs. *Review of Economic Dynamics*, 26:113–139.

Christiano, L. and Fitzgerald, T. (2003). The Band Pass Filter. *International Economic Review*, 44(2):435–465.

Clarida, R., Galí, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180.

Cogley, T., Primiceri, G. E., and Sargent, T. J. (2010). Inflation-Gap Persistence in the U.S. *American Economic Journal: Macroeconomics*, 2(1):43–69.

Crump, R. K., Eusepi, S., Giannone, D., Qian, E., and Sbordone, A. M. (2025). A Large Bayesian VAR of the US Economy. *International Journal of Central Banking*. Forthcoming.

Crump, R. K., Eusepi, S., Giannoni, M., and Şahin, A. (2019). A Unified Approach to Measuring u^*. *Brookings Papers on Economic Activity*, Spring 2019.

Debortoli, D., Galí, J., and Gambetti, L. (2020). On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint. In Eichenbaum, M., Hurst, E., and Parker, J., editors, *NBER Macroeconomics Annual 2020*, volume 34, pages 141–170. MIT Press, Cambridge, MA.

Del Negro, M., Giannone, D., Giannoni, M., and Tambalotti, A. (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity*, Spring:235–294.

Del Negro, M., Giannone, D., Giannoni, M., and Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, 118:248–262.

Del Negro, M., Lenza, M., Primiceri, G. E., and Tambalotti, A. (2020). What's up with the Phillips Curve? *Brookings Papers on Economic Activity*, Spring:301–357.

Del Negro, M. and Schorfheide, F. (2013). DSGE Model-Based Forecasting. In Elliot, G. and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 2, pages 57–140. Elsevier, New York.

Doan, T., Litterman, R., and Sims, C. (1984). Forecasting and conditional projection using realistic prior distributions. *Econometric reviews*, 3(1):1–100.

Elliott, G. (1998). On the Robustness of Cointegration Methods When Regressors Almost Have Unit Roots. *Econometrica*, 66(1):149–158.

Farmer, R. E. A. and Nicolò, G. (2018). Keynesian Economics without the Phillips Curve. *Journal of Economic Dynamics and Control*, 89:137–150.

Farmer, R. E. A. and Nicolò, G. (2019). Some International Evidence for Keynesian Economics without the Phillips Curve. *The Manchester School*, 89(S1):1–22.

Faust, J. (1998). The robustness of identified var conclusions about money. In *Carnegie-Rochester conference series on public policy*, volume 49, pages 207–244. Elsevier.

Fernald, J. G. (2007). Trend breaks, long-run restrictions, and contractionary technology improvements. *Journal of Monetary Economics*, 54:2467–2485.

Fernández-Villaverde, J., Rubio-Ramírez, J., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review*, 97(3):1021–1026.

Forni, M., Gambetti, L., Granese, A., Sala, L., and Soccorsi, S. (2025). An American Macroeconomic Picture: Supply and Demand Shocks in the Frequency Domain. *American Economic Journal: Macroeconomics*. Forthcoming.

Fukui, M., Nakamura, E., and Steinsson, J. (2023). Women, Wealth Effects, and Slow Recoveries. *American Economic Journal: Macroeconomics*, 15(1):269–313.

Galí, J. and Gambetti, L. (2009). On the Sources of the Great Moderation. *American Economic Journal: Macroeconomics*, 1(1):26–57.

Galí, J., Smets, F., and Wouters, R. (2011). Unemployment in an Estimated New Keynesian Model. In *NBER Macroeconomics Annual*, volume 26, pages 329–360. University of Chicago Press.

Giannone, D., Lenza, M., Momferatou, D., and Onorante, L. (2014). Short-term inflation projections: A Bayesian vector autoregressive approach. *International Journal of Forecasting*, 30(3):635–644.

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics*, 97(2):436–451.

Giannone, D., Lenza, M., and Primiceri, G. E. (2019a). Priors for the Long Run. *Journal of the American Statistical Association*, 114(526):565–580.

Giannone, D., Lenza, M., and Reichlin, L. (2019b). Money, Credit, Monetary Policy, and the Business Cycle in the Euro Area: What Has Changed Since the Crisis? *International Journal of Central Banking*, 15(5):137–173.

Giannone, D., Reichlin, L., and Sala, L. (2006). VARs, common factors and the empirical validation of equilibrium business cycle models. *Journal of Econometrics*, 132(1):257–279.

Gust, C., Herbst, E., and López-Salido, D. (2022). Short-Term Planning, Monetary Policy, and Macroeconomic Persistence. *American Economic Journal: Macroeconomics*, 14(4):174–209.

Hall, R. E. and Kudlyak, M. (2024). The Active Role of the Natural Rate of Unemployment. NBER Working Paper No. 31848.

Hamilton, J. D. (2018). Why You Should Never Use the Hodrick-Prescott Filter. *Review of Economics and Statistics*, 100(5):831–843.

Hazell, J., Herreño, J., Nakamura, E., and Steinsson, J. (2022). The Slope of the Phillips Curve: Evidence from U.S. States. *The Quarterly Journal of Economics*, 137(3):1299–1344.

Holston, K., Laubach, T., and Williams, J. C. (2017). Measuring the Natural Rate of Interest: International Trends and Determinants. *Journal of International Economics*, 108(1):S59–S75.

Jaimovich, N. and Siu, H. (2020). Job Polarization and Jobless Recoveries. *Review of Economics and Statistics*, 102(1):129–147.

Johannsen, B. and Mertens, E. (2021). A Time Series Model of Interest Rates With the Effective Lower Bound. *Journal of Money Credit and Banking*, 53(5):1005–1046.

Kadiyala, R. K. and Karlsson, S. (1997). Numerical Methods for Estimation and Inference in Bayesian VAR-models. *Journal of Applied Econometrics*, 12(2):99–132.

Laubach, T. and Williams, J. C. (2003). Measuring the Natural Rate of Interest. *Review of Economics and Statistics*, 85(4):1063–1070.

Laubach, T. and Williams, J. C. (2015). Measuring the Natural Rate of Interest Redux. Working Paper 15, Hutchins Center on Fiscal & Monetary Policy at Brookings.

Lewis, K. F. and Vazquez-Grande, F. (2019). Measuring the Natural Rate of Interest: A Note on Transitory Shocks. *Journal of Applied Econometrics*, 34(3):425–436.

Lubik, T. A. and Matthes, C. (2015). Calculating the Natural Rate of Interest: A Comparison of Two Alternative Approaches. Economic Brief 15-10, Federal Reserve Bank of Richmond.

McLeay, M. and Tenreyro, S. (2020). Optimal Inflation and the Identification of the Phillips Curve. In Eichenbaum, M., Hurst, E., and Parker, J., editors, *NBER Macroeconomics Annual 2020*, volume 34, pages 199–255. MIT Press, Cambridge, MA.

Mertens, E. (2016). Measuring the Level and Uncertainty of Trend Inflation. *Review of Economics and Statistics*, 98(5):950–967.

Morley, J. C., Nelson, C. R., and Zivot, E. (2003). Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different? *Review of Economics and Statistics*, 85(2):235–243.

Morley, J. C., Tran, T. D., and Wong, B. (2024). A Simple Correction for Misspecification in Trend-Cycle Decompositions with an Application to Estimating r^*. *Journal of Business & Economic Statistics*, 42(2):665–680.

Sargent, T. and Sims, C. (1977). Business cycle modeling without pretending to have too much a priori economic theory. Federal Reserve Bank of Minneapolis, Working Papers no. 55.

Schorfheide, F. and Song, D. (2015). Real-time forecasting with a mixed-frequency var. *Journal of Business & Economic Statistics*, 33(3):366–380.

Sims, C. A. and Zha, T. (2006). Were There Regime Switches in U.S. Monetary Policy? *American Economic Review*, 96(1):54–81.

Smets, F. (2010). Commentary: Modeling Inflation After the Crisis. *Macroeconomic Policy: Post-Crisis and Risks Ahead, Federal Reserve Bank of Kansas City Symposium*. Jackson Hole, Wyoming, August 26-28.

Stock, J. H. and Watson, M. W. (1988). Testing for Common Trends. *Journal of the American Statistical Association*, 83(404):1097–1107.

Stock, J. H. and Watson, M. W. (2007). Why Has Inflation Become Harder to Forecast? *Journal of Money, Credit and Banking*, 39(S1):3–33.

Stock, J. H. and Watson, M. W. (2008). Phillips Curve Inflation Forecasts. *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective, Federal Reserve Bank of Boston Conference*, June 9-11.

Stock, J. H. and Watson, M. W. (2010). Modeling Inflation After the Crisis. *Macroeconomic Policy: Post-Crisis and Risks Ahead, Federal Reserve Bank of Kansas City Symposium*. Jackson Hole, Wyoming, August 26-28.

Stock, J. H. and Watson, M. W. (2011). Dynamic Factor Models. In Clements, M. P. and Hendry, D. F., editors, *Oxford Handbook on Economic Forecasting*, pages 35–60. Oxford University Press, Oxford.

Stock, J. H. and Watson, M. W. (2020). Slack and cyclically sensitive inflation. *Journal of Money, Credit and Banking*, 52(S2):393–428.

Uhlig, H. (2003). What moves real GNP? Unpublished.

Villani, M. (2009). Steady-state Priors for Vector Autoregressions. *Journal of Applied Econometrics*, 24(4):630–650.

Watson, M. (1986). Univariate Detrending Methods with Stochastic Trends. *Journal of Monetary Economics*, 18(1):49–75.

Watson, M. (2004). Comment on “Monetary Policy in Real Time,” by Domenico Giannone, Lucrezia Reichlin, and Luca Sala. *NBER Macroeconomics Annual*, 19:216–221.

A Data

The following data series are from the Federal Reserve Economic Database (FRED) maintained by the Federal Reserve Bank of St. Louis:

- Real GDP per capita, A939RX0Q048SBEA, quarterly frequency. We transform the series by taking quarterly growth rates at annual rate and express these rates in percentages.
- Unemployment rate, UNRATE, monthly frequency. We transform the series by taking quarterly averages.
- Inflation, GDPDEF, quarterly frequency. We transform the series for the GDP price index by taking quarterly growth rates at annual rate and express these rates in percentages.
- Effective federal funds rate, FEDFUNDS, monthly frequency. Because the series is already expressed at annual rate, we take quarterly averages.

The following data series are available from the Real-Time Data Research Center maintained by the Federal Reserve Bank of Philadelphia:[3]

- One-year-ahead inflation expectations, INFPGDP1YR, quarterly frequency. The series corresponds to the median forecast for one-year-ahead annual average inflation measured by the GDP price index. The series starts in 1970:Q2.
- Ten-year-ahead inflation expectations. We follow Del Negro and Schorfheide (2013) to construct this time series. Specifically, we combine longer-run inflation expectations from the SPF and the Blue Chip Economic Indicators survey. We use the ten-year-ahead Consumer Price Index (CPI) inflation expectations from the Blue Chip survey—from 1979:Q4 to 1991:Q3 and available twice a year—and those from the SPF (INFCPI10YR)—available each quarter starting from 1991:Q4. To combine the measures, we subtract from the ten-year-ahead CPI inflation expectations the historical average difference between CPI and GDP annualized inflation over the estimation period.

[3] The data may contain missing observations. More details are available at the webpage: https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/inflation-forecasts.

- One-year-ahead unemployment expectations, UNEMP6 , quarterly frequency. The series corresponds to the median forecast for one-year-ahead unemployment. The series starts in 1968:Q4.