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 July 2026

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 University, 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 July 2026
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
and CEPR
and also NBER
francesco.bianchi@jhu.edu

Giovanni Nicolò
Federal Reserve Board of Governors
giovanni.nicolo@frb.gov

Dongho Song
Johns Hopkins University
Carey Business School
dongho.song@jhu.edu

1 Introduction

This paper addresses an important conundrum. Benchmark New Keynesian (NK) models link movements in inflation to various measures of economic slack, such as the output or unemployment gap. Yet a recent time-series literature argues that inflation appears largely disconnected from real economic activity (Angeletos et al., 2020). In other words, although the NK Phillips curve remains central to contemporary macroeconomic analysis of monetary policy and the interaction between inflation and real activity, researchers have become increasingly skeptical of this relationship. We argue that this apparent disconnect reflects neglected movements in trend inflation and that, once these are accounted for, the empirical relationship between inflation and economic slack reemerges.

Three features of macroeconomic data complicate the analysis of the cyclical relationship between inflation and real activity. First, macroeconomic series exhibit substantial low-frequency fluctuations (Sargent and Sims, 1977). This feature is especially pronounced for inflation but is also evident in unemployment and output growth. Second, the duration of business-cycle fluctuations is uncertain and may differ across variables. Traditionally, business-cycle fluctuations have been associated with cycles lasting between 6 and 32 quarters. Beaudry et al. (2020) argue that this benchmark may be outdated, as it is rooted in pre–World War II (WWII) data. Using postwar evidence, they advocate focusing on longer cycles, with horizons on the order of 8 to 10 years. Third, the relatively short sample of macroeconomic data poses econometric challenges, especially when estimating models with many parameters. As a result, researchers typically employ VARs with only a few lags or relatively tight priors to avoid overfitting. We argue that these common modeling choices may affect a model's ability to recover the business-cycle relationships of interest.

To address these challenges, we employ a Trend-Cycle (TC) VAR (Watson, 1986; Stock and Watson, 1988, 2007; Villani, 2009; Del Negro et al., 2017). The TC-VAR provides a parametric structure that separately models trend and cyclical fluctuations. Relative to a conventional VAR with many lags, it is more tightly parameterized while still offering substantial flexibility. A distinct set of parameters, separate from those that determine trend dynamics, governs how the cyclical components evolve, shaping the cyclical relationship between inflation and real activity. These features allow us to adopt a max-share identification strategy to isolate the shock that explains the largest share of variation in the unemployment cycle without taking a stance on the duration of the business cycle. We show that this shock accounts for a substantial fraction of the variance of the cyclical components of both inflation and expected inflation.

We estimate the TC-VAR over the 1960-2019 period using seven time series. The first four are standard in the literature: GDP growth, unemployment, the federal funds rate, and inflation. We then augment the system with three additional variables. First, to capture movements in trend inflation more accurately, we include ten-year-ahead inflation expectations. While long-term expectations may not forecast future inflation precisely, they are informative about the *current* level of trend inflation. Second, we incorporate one-year-ahead inflation expectations, which should respond to business-cycle movements in inflation while being less sensitive to purely transitory shocks. Finally, we add one-year-ahead unemployment expectations, which provide additional information about the latent unemployment trend.

We identify the shock that maximizes the variance of the latent cyclical component of unemployment and examine its contribution to the volatility of all cyclical components in the system. In our baseline specification, the unemployment-identified shock explains approximately 72% of the unemployment cycle and nearly 41% of the inflation cycle. This finding points to a strong relation between the drivers of real activity and inflation over the business cycle. The shock also explains roughly 57% of the cyclical variation in inflation expectations. This result further supports the view that business-cycle movements in inflation are closely related to real activity, given that expected inflation is linked to actual inflation but less affected by high-frequency transitory disturbances.

Our results differ sharply from the seminal contribution by Angeletos et al. (2020). Using a VAR applied to post-WWII US data, they identify a "business-cycle" shock in the *frequency domain* that accounts for the largest share of variation in GDP or unemployment at frequencies corresponding to cycles between 6 and 32 quarters (Uhlig, 2003, Giannone et al., 2019b). They find that this shock explains most of the cyclical fluctuations in various measures of real activity but has little impact on the variation of inflation at the same frequencies. The authors explain that while the unemployment-identified shock may effectively reflect a combination of structural shocks, their findings nonetheless point to a substantive disconnect between unemployment and inflation dynamics. Thus, even if an NK Phillips curve were to exist, its importance would be dwarfed by the shocks driving unemployment fluctuations.

We first verify that we can recover results similar to Angeletos et al. (2020) if we follow their VAR approach. We then highlight two key differences between our approach and the one employed by Angeletos et al. (2020). First, we use a TC-VAR as opposed to a

VAR. Second, while both approaches rely on a max-share identification strategy, we work directly with the cycles of the target variables, as opposed to targeting variation of the overall variables at specific frequencies.

We investigate the role of these differences. We first show that the unemployment shock identified with our baseline approach explains significantly more inflation variation at 6 to 32 quarters than the VAR-identified shock. Consistent with this finding, while the shocks are positively correlated (0.64), when fed through the corresponding models they imply quite different cyclical fluctuations, especially around key events such as the Great Inflation and the aftermath of the Great Recession. Finally, even when the shocks are identified in the frequency domain under both models, we obtain a stronger relation between unemployment and inflation using the TC-VAR than using a VAR. The difference between the two models is larger when considering longer cycles as advocated by Beaudry et al. (2020).

These results underscore two main insights. First, the parametric model used to fit the data materially shapes the outcomes, even when conditioning on a given frequency band. Second, the choice of frequencies itself matters. Together, these points validate our focus on the variability of cycles extracted by the TC-VAR rather than on frequency-domain methods. The TC-VAR cleanly separates trend and cycle dynamics without requiring assumptions about which frequencies constitute the business cycle.

In the final section, we formalize these ideas. Building on Fernández-Villaverde et al. (2007), we show that a TC-VAR can be represented as a VAR with infinitely many lags. This observation has important implications. First, using a finite-lag VAR, as is common in the literature, implies an inherent misspecification if the data are generated by a TC-VAR. As a result, identification in the frequency domain based on a finite-order VAR can produce results that differ markedly from those obtained with a TC-VAR. Second, even when the VAR includes a very large number of lags, it tends to understate the contribution of the max-share shock because it cannot cleanly separate trend shocks from cyclical shocks. Third, variance decompositions are highly sensitive to the choice of frequency band, implying that quantitative conclusions depend critically on how the business cycle is defined. Our baseline approach avoids these issues by working directly with the cyclical representation.

While our baseline approach advances the understanding of the inflation-unemployment relationship by controlling for movements in trend inflation, it does not speak to the underlying sources of changes in trend inflation itself. This is because the TC-VAR provides a parsimonious and informative representation of the data, but it is ultimately silent on the origins of trend-inflation fluctuations. Although these movements are sizable and merit deeper investigation in future work, our results show that the Phillips-curve-type tradeoff remains relevant for policymakers. At the same time, our findings point to the need for models that incorporate meaningful variation in trend inflation and can account for the sources and implications of such low-frequency movements.

Our paper contributes to several literatures in macroeconomics and time-series econometrics. From a traditional Keynesian Phillips curve perspective, 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 NK 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 NK model to improve its empirical fit (Gust et al., 2022, Farmer and Nicolò, 2018, 2019). Other authors abandon the NK 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).

Our analysis connects to Sargent and Sims (1977) which shows that macro data can be well explained by two factors, one of which captures low-frequency movements in nominal variables, while the other accounts for business cycle variation. 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 that reconciles 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 also hold in data-rich environments (Bańbura et al., 2015; Crump et al., 2025).

The literature has proposed alternative 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 rely on the Beveridge-Nelson decomposition for the study of macroeconomic dynamics (Berger et al., 2023; Morley et al., 2024).

We opt for a time-invariant TC-VAR as it is well-suited for decomposing trends and cycles *and* conducting a max-share analysis of cyclical fluctuations. In principle, models with smoothly time-varying parameters could also be used to capture low-frequency variation (Cogley et al., 2010; Galí and Gambetti, 2009; Debortoli et al., 2020). Furthermore, our model could be extended to include time-varying volatility (Stock and Watson, 2007; Galí and Gambetti, 2009; Debortoli et al., 2020; Johannsen and Mertens, 2021). However, in both cases, the max-share approach becomes more nuanced because the covariance structure evolves over time. One could compute the max-share under an "anticipated utility" assumption, disregarding parameter instability. Alternatively, instability could be modeled via regime changes (Sims and Zha, 2006; Bianchi, 2013; Bianchi and Ilut, 2017), allowing for a max-share approach that properly accounts for parameter changes (Bianchi, 2016). We regard this as a promising direction for future research.

2 A Trend-Cycle VAR

In this section, we present the TC-VAR used to model the joint dynamics of GDP, unemployment, the federal funds rate (FFR), and inflation, as well as three expectations measures: one-year-ahead unemployment expectations and one- and ten-year-ahead inflation expectations. The model features four latent trend components and six latent cyclical components, collected in the vectors τ_t and c_t. The economic interpretation of these latent components is determined by how they load into the observed variables through the measurement equations. We first describe the dynamics of the latent components, then explain the measurement mapping variable by variable, and finally introduce the state-space representation of the TC-VAR.

2.1 State-transition dynamics

In our model, τ_t captures persistent low-frequency movements, while c_t captures stationary cycles. The joint dynamics of the trend and cyclical components are:

$$\tau_t = \mu + \varphi\tau_{t-1} + \varepsilon_{\tau,t}, \tag{1.1}$$

$$c_t = \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + \cdots + \Phi_p c_{t-p} + \underbrace{\Phi_\tau(L_\tau^{-1}\varepsilon_{\tau,t}) + \tilde{\varepsilon}_{c,t}}_{\varepsilon_{c,t}}, \tag{1.2}$$

where μ is a column vector and φ is a diagonal matrix. We restrict all elements of μ to zero and the diagonal elements φ to 1 except for the trend of real GDP per capita growth. This implies that all trends follow a random walk (as typically assumed in the literature), except for trend-growth, which is instead modeled as an AR(1) process.

The innovations to trend and cycles, $\varepsilon_{\tau,t}$ and $\varepsilon_{c,t}$, are allowed to be correlated:

$$\varepsilon_t = \begin{bmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \begin{bmatrix} \Sigma_\tau & L_\tau\Phi_\tau' \\ \Phi_\tau L_\tau' & \Phi_\tau\Phi_\tau' + \Sigma_c \end{bmatrix} \right). \tag{2}$$

The matrices Σ_τ and Σ_c are conforming positive definite covariance matrices associated with $\varepsilon_{\tau,t}$ and $\tilde{\varepsilon}_{c,t}$, respectively, which are orthogonal to each other. The matrix L_τ is the lower triangular factor from the Cholesky decomposition of Σ_τ. We use the matrix Φ_τ to parameterize the correlation between trend and cycle innovations. This parameterization does not affect the identification strategy based on the max-share variance described below, but it facilitates inference. Setting $\Phi_\tau = \mathbf{0}$ yields the standard specification with orthogonal trend and cycle shocks. As shown by Morley et al. (2003), a necessary condition for equivalence between a univariate trend–cycle representation with correlated shocks and an ARIMA process is that $p \geq 2$. This condition generalizes to the multivariate case and holds for the specification considered here.

2.2 Measurement equations

Unemployment and inflation, whose dynamics are central to our analysis, are modeled as the sum of their own trends and cyclical components:

$$u_t = \tau_{u,t} + c_{u,t}, \tag{3}$$

$$\pi_t = \tau_{\pi,t} + c_{\pi,t}. \tag{4}$$

For real GDP per capita, output is assumed to evolve as $Y_t = \bar{Y}_t \exp(c_{y,t})$, where the cyclical component captures deviations of output from a stochastic trend. We model the cycle in the level of real GDP rather than in GDP growth, because inflation and unemployment dynamics are governed by the output gap rather than output growth. By definition, real GDP per capita growth, $g_t \equiv \ln(Y_t/Y_{t-1})$, then follows as

$$g_t = \tau_{g,t} + (c_{y,t} - c_{y,t-1}). \tag{5}$$

The long-run Fisher relation implies that the federal funds rate loads on the sum of the trend real interest rate and trend inflation, together with a cyclical component:

$$f_t = (\tau_{r,t} + \tau_{\pi,t}) + c_{f,t}. \tag{6}$$

In addition to the four core macroeconomic observables, survey-based expectations provide additional information on the latent trend components. One-year-ahead unemployment expectations share the same trend component as realized unemployment in (3) but load on a distinct cyclical component:

$$u_t^{e,1y} = \tau_{u,t} + c_{u,t}^{e}. \tag{7}$$

Inflation expectations share the same inflation trend as realized inflation in (4):

$$\begin{aligned} \pi_t^{e,1y} &= \tau_{\pi,t} + c_{\pi,t}^{e}, \\ \pi_t^{e,10y} &= \tau_{\pi,t} + \delta c_{\pi,t}^{e} + \eta_{\pi,t}^{e,10y}. \end{aligned} \tag{8}$$

The one-year-ahead expectation features its own cyclical component. The ten-year-ahead expectation also loads on this cyclical component, but with a loading $\delta < 1$, as long-term expectations are naturally less affected by cyclical movements. We add measurement error on long-term expectations to avoid singularity in the measurement system.

2.3 State-space representation

The observed, trend, and cycle vectors are, $z_t = \{g_t, u_t, u_t^{e,1y}, f_t, \pi_t, \pi_t^{e,1y}, \pi_t^{e,10y}\}'$, $\tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}'$, $c_t = \{c_{y,t}, c_{u,t}, c_{u,t}^{e}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^{e}\}'$, respectively. The state vector stacks the trend and cycle states. The trend state is $x_{\tau,t} = \tau_t$, while the cycle state stacks current and lagged cycles, $x_{c,t} = \{c_t', c_{t-1}', \ldots, c_{t-(p-1)}'\}'$. The full state vector is $x_t = \{x_{\tau,t}', x_{c,t}'\}'$, and $\eta_t = \{\eta_{\pi,t}^{e,10y}\}$ denotes the idiosyncratic measurement error.

The measurement equation takes the form

$$z_t = \Lambda_x x_t + \Lambda_\eta \eta_t. \tag{9}$$

Each row of Λ_x corresponds to one of the equations (3)–(8), which specify how the observables load on the latent trends and cycles. The matrix Λ_η loads measurement error in the long-run inflation expectation, which is the only source of measurement error in the model.

The transition equation for the state x_t is given by

$$x_t = \mathcal{C} + \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \quad \varepsilon_t \sim N(\mathbf{0}, \Sigma_\varepsilon) \tag{10}$$

or equivalently,

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

where

$$\Phi_c = \begin{bmatrix} \Phi_1 & \Phi_2 & \dots & \Phi_p \\ \mathbf{I} & \mathbf{0} & \dots & \mathbf{0} \\ \mathbf{0} & \ddots & \ddots & \vdots \\ \mathbf{0} & \dots & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad \mathcal{R}_c = \begin{bmatrix} \mathbf{I} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix}, \quad \Sigma_\varepsilon = \begin{bmatrix} \Sigma_\tau & L_\tau \Phi_\tau' \\ \Phi_\tau L_\tau' & \Phi_\tau \Phi_\tau' + \Sigma_c \end{bmatrix}.$$

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 $x_{c,0}$.

3 Inference

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

3.1 Data

We estimate the TC-VAR with two lags ($p = 2$) and conduct robustness checks with four lags ($p = 4$). We use seven quarterly time series 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, treating observations at the zero lower bound as in 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.

3.2 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

$$p(\Sigma_\tau) = \mathcal{IW}\left(\kappa_\tau, (\kappa_\tau + n_\tau + 1)\underline{\Sigma_\tau}\right), \tag{11.1}$$

$$p(\Sigma_c) = \mathcal{IW}\left(\kappa_c, (\kappa_c + n_c + 1)\underline{\Sigma_c}\right), \tag{11.2}$$

$$p(\phi|\Sigma_c) = \mathcal{N}\left(\underline{\phi}, \Sigma_c \otimes \underline{\Omega}\right) \mathbf{I}(\phi), \tag{11.3}$$

where $\phi = vec[\Phi_1, \ldots, \Phi_p, \Phi_\tau]$ 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 (10) 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 (11.3), we assume a conventional Minnesota prior, in line with Giannone et al. (2015). Because the cyclical component in (1.2) 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 (11.3): 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. Qualitatively similar but less precise results obtain under looser priors.

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

3.3 Identifying shocks that drive business-cycle fluctuations

We first describe our baseline max-share identification strategy (Faust, 1998, Uhlig, 2003). Given that the TC-VAR explicitly isolates cyclical fluctuations, our baseline approach applies the max-share criterion directly to the unemployment cycle. We then discuss an alternative identification strategy based on the volatility of the observed variables at specific frequencies. This second approach will be useful when relating our results to the existing literature that employs analogous frequency-domain methods within a VAR framework (Giannone et al., 2019b, Angeletos et al., 2020, Basu et al., 2024).

Baseline identification Given the state-space representation in (9) and (10), let M_c be a selector matrix that extracts the cyclical component from x_t. Under stationarity,

$$c_t \equiv M_c x_t = M_c(\mathbf{I} - \Phi L)^{-1}\mathcal{R}\,\varepsilon_t, \tag{12}$$

where L is the lag operator. Let e_s select the sth element of c_t, and define $c_{s,t} \equiv e_s' c_t$. In our application, e_s selects unemployment. Using (12), the scalar cycle can be written as

$$c_{s,t} \equiv e_s' c_t = \sum_{j=0}^{\infty} a_j\,\varepsilon_{t-j}, \qquad a_j \equiv e_s' M_c \Phi^j \mathcal{R}.$$

We identify a linear combination $q'\varepsilon_t$, normalized so that $\mathrm{Var}(q'\varepsilon_t) = q'\Sigma_\varepsilon q = 1$, that explains the largest possible share of the unconditional variance of $c_{s,t}$.[1] Let $c_{s,t}(q)$ denote the component of $c_{s,t}$ driven by this shock,[2] i.e.,

$$c_{s,t}(q) = \sum_{j=0}^{\infty} \frac{a_j \Sigma_\varepsilon q}{q'\Sigma_\varepsilon q} \left(q'\varepsilon_{t-j}\right) = \sum_{j=0}^{\infty} a_j \Sigma_\varepsilon q \left(q'\varepsilon_{t-j}\right).$$

We define q_s as any solution to

$$q_s \in \mathrm{argmax}_{q'\Sigma_\varepsilon q=1} \; \mathrm{Var}\big(c_{s,t}(q)\big) \quad \text{where} \quad \mathrm{Var}\big(c_{s,t}(q)\big) = q'\Big(\Sigma_\varepsilon \sum_{j=0}^{\infty} a_j' a_j \Sigma_\varepsilon\Big) q. \tag{13}$$

The maximizer q_s corresponds to the eigenvector associated with the largest eigenvalue of the quadratic form pinning down $\mathrm{Var}(c_{s,t}(q))$ in (13).

Alternative identification We now describe an alternative max-share identification that targets the volatility of the observed variable at specific frequencies.

Let $f_{z_s}^{(q)}(\omega)$ denote the spectral density of the component of the scalar variable $z_{s,t}$ driven by the linear combination $q'\varepsilon_t$, constructed using the measurement-equation selector M_x:

$$f_{z_s}^{(q)}(\omega) = \frac{1}{2\pi} H_s(e^{-i\omega}) \, \Sigma_\varepsilon^{(q)} \, H_s(e^{-i\omega})', \quad H_s(e^{-i\omega}) \equiv e_s' M_x (\mathbf{I} - \Phi e^{-i\omega})^{-1} \mathcal{R}, \quad \Sigma_\varepsilon^{(q)} \equiv \frac{\Sigma_\varepsilon q q' \Sigma_\varepsilon}{q'\Sigma_\varepsilon q}.$$

The frequency-domain max-share identification is defined as the solution to

$$q_s \in \mathrm{argmax}_{q'\Sigma_\varepsilon q=1} \int_{\mathcal{B}} f_{z_s}^{(q)}(\omega)\, d\omega, \tag{14}$$

which measures the share of the variance of $z_{s,t}$ attributable to the linear combination $q'\varepsilon_t$ over the frequency band $\mathcal{B}$. Angeletos et al. (2020) restrict $\mathcal{B}$ to cycles of 6 to 32 quarters. We will also consider longer frequency bands as advocated by Beaudry et al. (2020).

[1] Equivalently, one may express the linear combination in an orthonormalized shock representation. In that representation, the weights are given by a unit-norm vector Q (where $Q'Q = 1$). The corresponding weights on ε_t in our representation are then $q = (L_\varepsilon^{-1})'Q$, where $\Sigma_\varepsilon = L_\varepsilon L_\varepsilon'$.

[2] This derivation is based on the population linear projection $\mathbb{E}(\varepsilon_t \mid q'\varepsilon_t) = \mathrm{Cov}(\varepsilon_t, q'\varepsilon_t)/\mathrm{Var}(q'\varepsilon_t)\,(q'\varepsilon_t)$.

4 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.

4.1 Estimated latent trends and cycles

Panel (A) of Figure 1 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 1 plots the posterior 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

Figure 1: Data, trends and cycles

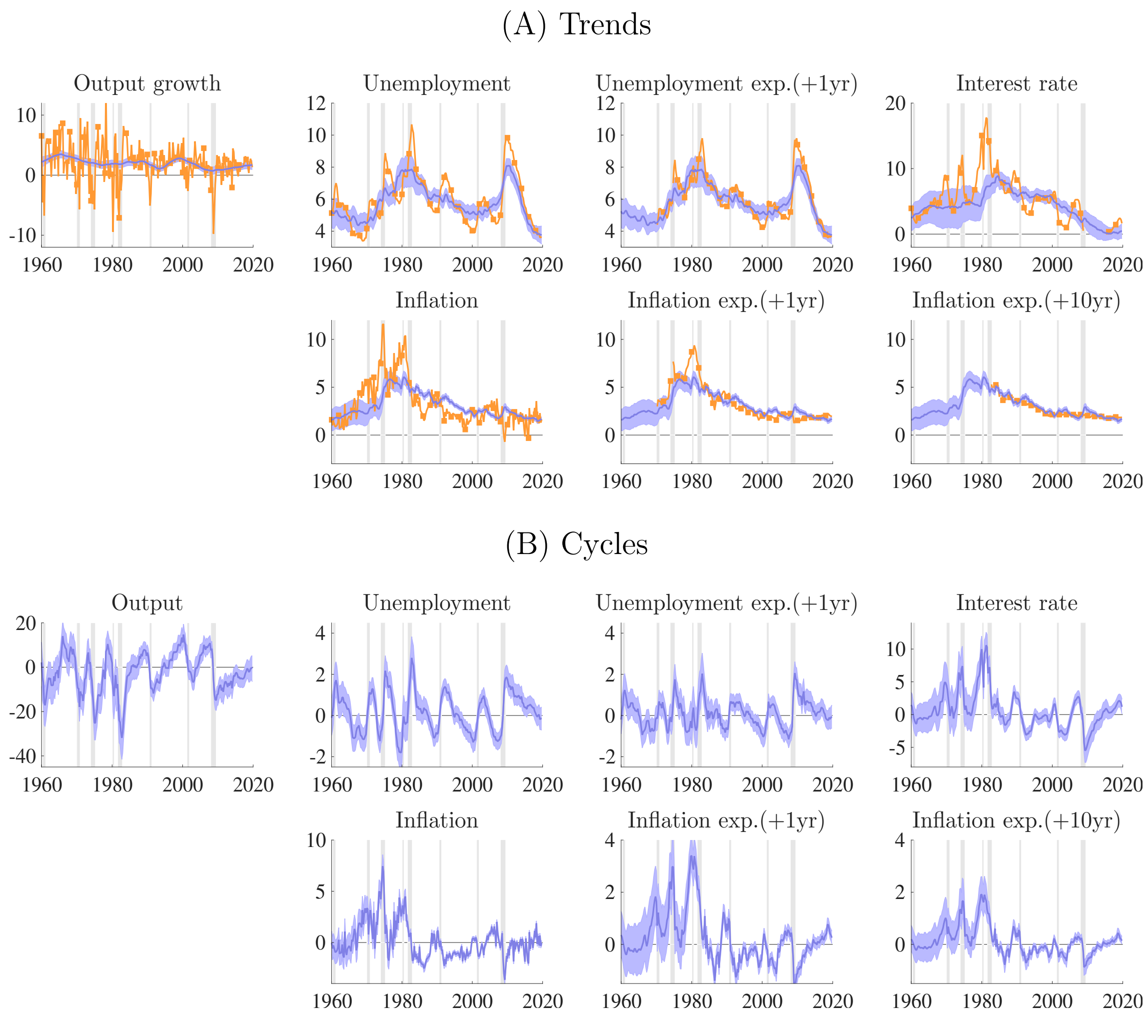

Notes: The figure plots the data (orange lines) used for the estimation of the TC-VAR over 1960-2019 period as well as the posterior median of their latent trends (blue lines) in panel (A) and latent cycles (blue lines) in panel (B) and the corresponding 90-percent posterior-coverage intervals (shaded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

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

Table 1: Variance contributions of the unemployment shock

Unemployment	Output	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
72.4	58.6	69.0	66.1	40.6	57.1
[61.8,85.8]	[46.9,71.0]	[57.2,81.2]	[44.5,85.3]	[19.1,66.6]	[31.6,80.8]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment cycle. 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.

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.

4.2 Inflation and unemployment over the business cycle

We now 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 3.3. Specifically, the shock is identified by maximizing its contribution to the volatility of the cyclical component of unemployment. The Online Appendix shows that results are similar if we use GDP to identify the business-cycle shock.

Table 1 reports the median and the 68% posterior-coverage interval for the contribution of the identified shock to the variance of the *cycle* of all variables. Not surprisingly, the unemployment-identified shock explains a large share of unemployment cyclical fluctuations (about 72%). Importantly, it also accounts for a sizable fraction of cyclical inflation, explaining nearly 41% of the inflation cycle. These shares are economically large, particularly given that the shock does not fully account for unemployment cyclical fluctuations.

The results are even stronger when focusing on the cyclical component of inflation expectations: approximately 57% 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 inflation cycle, this finding corroborates the view that inflation dynamics are strongly related to the real business cycle.

The shock also explains a substantial share of fluctuations in nominal interest rates, accounting for about 66% of the volatility of the nominal interest rate cycle. Taken together with the results for inflation and inflation expectations, this implies that the shock explains

Figure 2: Impulse responses to the unemployment shock

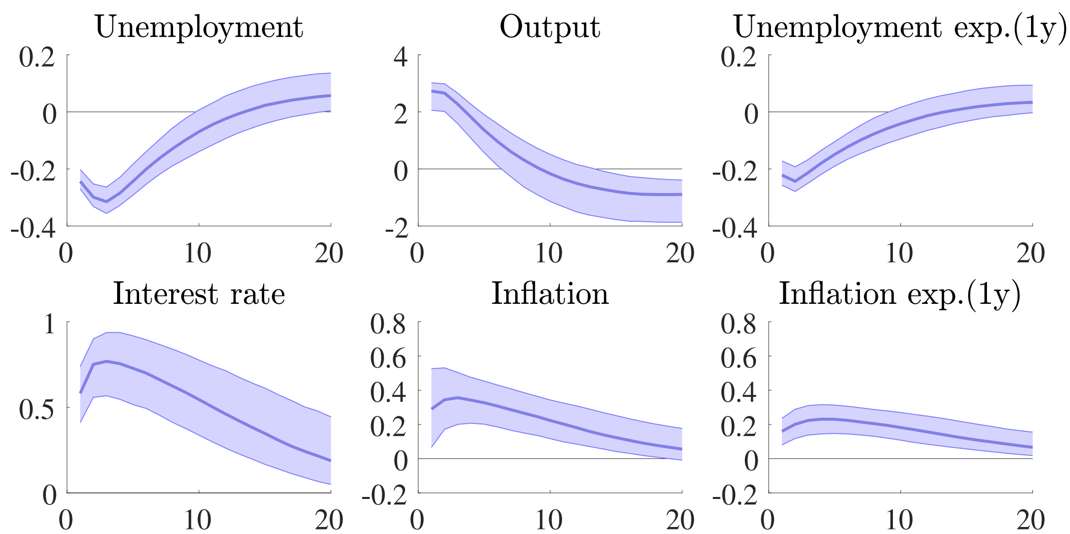

Notes: The figure shows the response to the unemployment-identified shock of the cyclical component of all the variables over a period of 20 quarters. The figure plots the posterior median (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area).

a large portion of the volatility of the cyclical component of the real interest rate, defined as the difference between the federal funds rate and expected inflation.

Finally, the shock explains a large share of the volatility of the GDP cycle. It also explains a share of the volatility of the cyclical component of expected unemployment that is similar to that for realized unemployment. Taken together, these findings support the conclusion that the identified shock is a central driver of business-cycle fluctuations.

Figure 2 plots the impulse responses of each cyclical component to the unemployment-identified shock over a period of 20 quarters. The resulting interpretation of the 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

Figure 3: Implied series from the unemployment shock

(A) Comparison with observed variables

(B) Comparison with cyclical components

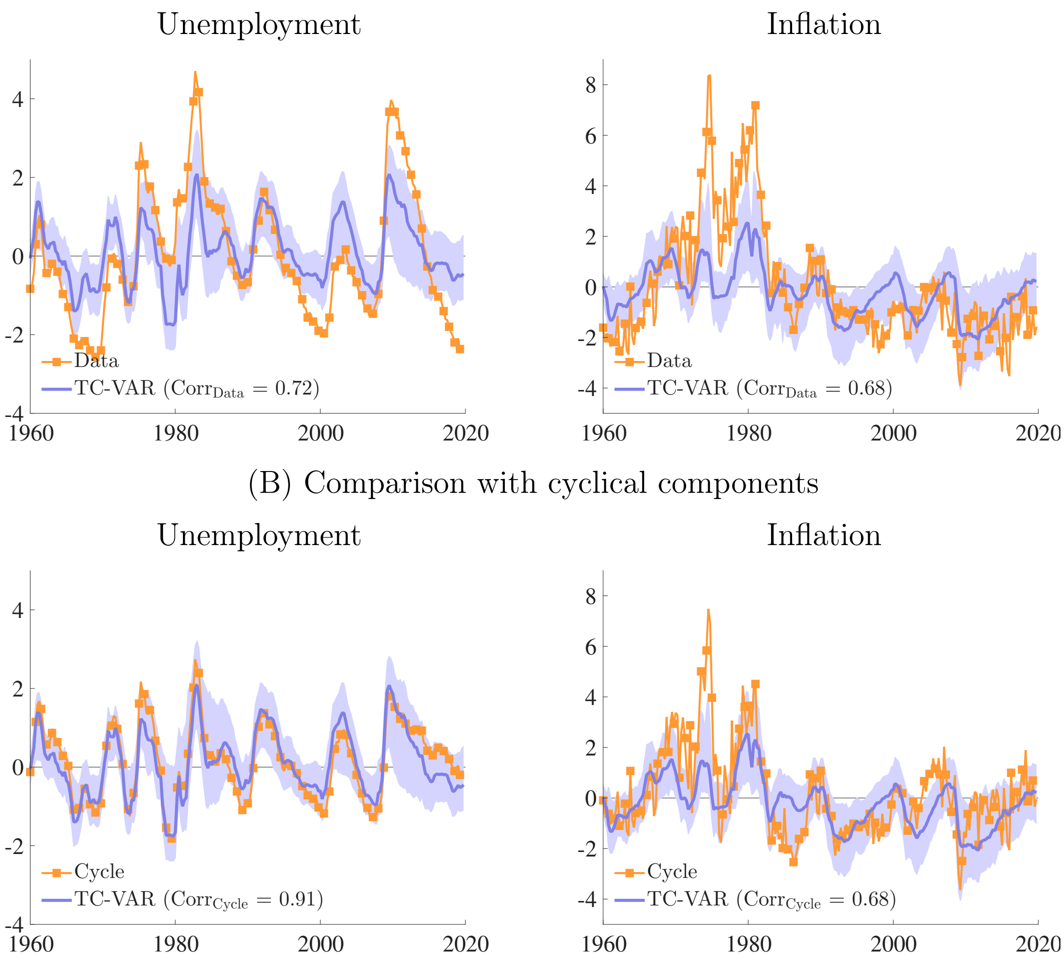

Notes: We plot the implied unemployment and inflation series obtained by isolating the max-share unemployment shock in the TC-VAR, identified by maximizing its contribution to cyclical volatility. Results are based on all posterior draws, with bands denoting 90% posterior intervals and solid lines indicating posterior medians. For ease of comparison, we overlay the demeaned data in panel (A) and the posterior median cyclical components in panel (B).

returns to its initial value thereafter.

Figure 3 displays the implied unemployment and inflation series and cycles obtained by isolating the max-share unemployment shock. These counterfactual series are generated by feeding the TC-VAR model only with the identified max-share unemployment shock. The series can be interpreted as the amount of *historical* variation accounted for by the shock. We compare these counterfactual series with the actual data series (top panel) and with the corresponding cycles (bottom panel). Since the shock is constructed to explain the largest share of cyclical unemployment variation, it is not surprising that it accounts for most of the historical cyclical variation in this series. The correlation with the unemployment cycle is 0.91 and the correlation with observed unemployment is also high (0.72).

Noticeably, the shock also accounts for a large share of the historical variation in inflation. Its correlations with actual inflation and with the inflation cycle are both high, at 0.68. The only notable divergence for the inflation cycle occurs during the major oil shock of 1974, when inflation spiked due to the direct impact of the energy shock. These results complement the evidence presented in Table 1. Recall that those results are based on parametric inference for the TC-VAR model and do not directly use the filtered shocks employed in this figure. This provides compelling evidence of the validity and empirical strength of our baseline identification strategy.

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. The results are essentially unchanged when we allow for $p = 4$ rather than $p = 2$ lags in the VAR characterizing the cycles, or when we assume that shocks to trend and cycles are uncorrelated. The results remain qualitatively unchanged, though somewhat weaker in magnitude, when expectations are removed. However, the model without expectations also exhibits worse out-of-sample forecasting performance, as highlighted in Section 5.3. We interpret this as evidence that expectations contain valuable information for sorting trend and cycle fluctuations.

5 Connecting to the existing literature

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 then consider two alternative approaches that allow for a more direct comparison between the VAR and the TC-VAR.

Table 2: Variance contribution of unemployment-identified shock: VAR analysis

	Baseline data		Angeletos et al. (2020) data	
	Unemployment	Inflation	Unemployment	Inflation
Flat prior	72.0	13.6	71.2	10.4
	$[62.7, 81.0]$	$[6.4, 24.3]$	$[63.7, 78.8]$	$[4.4, 19.1]$
Optimized Minnesota prior	72.9	12.7	73.0	7.7
	$[64.6, 80.7]$	$[6.8, 21.6]$	$[66.5, 79.4]$	$[3.5, 13.3]$
Optimized Minnesota	78.1	16.8	91.9	10.3
+ Long-run prior	$[69.7, 85.0]$	$[8.4, 26.4]$	$[88.7, 94.3]$	$[4.9, 16.8]$

*N*otes: 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).

5.1 VAR evidence

We proceed to estimate a VAR with two lags as in Angeletos et al. (2020). 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 3.3. Specifically, we look for the combination of shocks that explains the largest share of unemployment fluctuations for cycles between 6 and 32 quarters. 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 employ exactly the same variables used in the TC-VAR, while in the second case we use the same variables as Angeletos et al. (2020). The most noticeable difference is that our dataset also includes expectations. Consistently 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 2 report the contributions of the max-share shock for the variability 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 expectations, but still significantly 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).

5.2 Comparing VAR and TC-VAR results

The previous subsection has shown that a researcher would reach very different conclusions regarding the business cycle relationship between inflation and unemployment if they used the VAR approach with identification in the frequency domain rather than our baseline approach, which employs a TC-VAR to separately model trends and cycles and identifies shocks using the covariance structure of the cycles. In this subsection, we consider a series of variations of our baseline approach to understand what drives the difference in results.

We begin by noting that the TC-VAR does not require the researcher to take a stance on which frequencies correspond to the business cycle. As a result, the two approaches differ in two important dimensions. First, in our baseline approach, the importance of the unemployment shock for inflation business cycle fluctuations is based on its contribution to the inflation *cycle*. In contrast, the VAR approach relies on the contribution of the identified shock to the variability of *overall* inflation at certain frequencies. Second, the shock itself is identified differently. Under the TC-VAR, the shock is identified based on its contribution to the unemployment *cycle*, whereas under the VAR, the shock is identified based on its contribution to the volatility of *overall* unemployment at certain frequencies. In what follows, we examine the role of these features of the two methodologies in accounting for the differences in results.

Do the two approaches recover the same shocks? A first possibility is that the results between the two approaches differ because in our baseline exercise the importance of the shock is extracted using the unemployment cycle, while in the other case it is extracted based on the volatility of overall unemployment at certain frequencies. To assess this hypothesis, we first check whether the two identification strategies recover similar shocks. In Figure 4, we report the time series of the unemployment cycle shock identified in the TC-VAR by targeting cyclical unemployment (our baseline) and its VAR counterpart identified by targeting observed unemployment over the 6-32 quarter frequency range. The two shocks commove with a correlation of 0.64, and the discrepancies are scattered rather than concentrated in particular episodes. Thus, to a large extent the shocks that are recovered with the two methods are quite similar.

However, when the shocks are fed through the corresponding models we observe important differences. To show this, we construct two counterfactual simulations. The first one inputs the max-share shock identified with the VAR through the VAR model, while the second one is identical to what reported in Figure 3 and obtained by feeding the TC-VAR with the shock identified using our baseline approach. These results are reported in the second and third plots of Figure 4. Both models imply a strong correlation between unemployment and the counterfactual series. However, the correlation of inflation with its TC-VAR counterfactual is more than twice as large as with the VAR counterfactual (.68 vs .29). Moreover, for both unemployment and inflation, the counterfactual series exhibit substantial discrepancies during key historical episodes. For inflation, differences are particularly pronounced during the two oil shocks of the 1970s, while the Great Recession and its aftermath show a large divergence for unemployment.

While these historical decompositions do not directly enter the baseline results presented in Table 1, they provide additional evidence that the parametric model used to analyze the data makes a material difference for our understanding of the link between unemployment and inflation. Shocks that appear similar might lead to very different cyclical fluctuations once fed through different models.

Figure 4: Shocks and implied series under baseline and VAR identification strategy

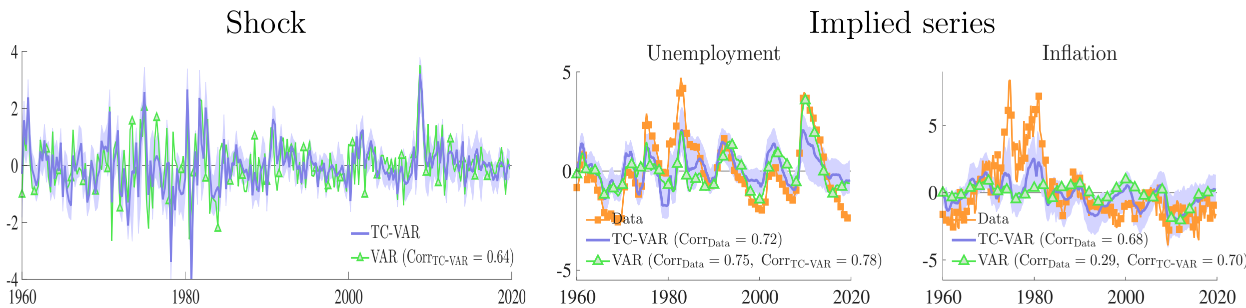

Notes: TC-VAR targets cyclical unemployment in identifying the unemployment shock, while the VAR targets observed unemployment over the 6-32 quarter band. In the first panel, we show the identified shocks, while in the second and third panels we report the implied unemployment and inflation series. Bands denote 90% posterior intervals, solid lines indicate posterior medians, and demeaned data are overlaid.

Aligning the target We then ask whether the different conclusions arise from differences in how the importance of the identified shock for inflation is evaluated. In this exercise, we keep the identification used in our baseline approach unchanged but align the statistic used to measure the share of inflation volatility. Specifically, we consider the spectrum of the entire TC-VAR and ask how much of the variance of overall inflation at certain frequencies is explained by the shock that maximizes the volatility of the unemployment cycle. The answer to this question depends on the recovered TC-VAR parameters. The shock sequence and the associated counterfactual series presented above are clearly related to these parameters, but they do not directly affect the results presented here.

Table 3 reports the results. For comparison, Panel (A) contains our baseline results, while Panel (B) considers three frequency bands: $6-32$ quarters, $16-50$ quarters, and $32-50$ quarters. The first frequency band corresponds to what used in Angeletos et al. (2020). The last one roughly aligns with what advocate by Beaudry et al. (2020), while the second one bridges the two by allowing for longer cycles, while removing very short ones.

Two important results emerge. First, the unemployment cycle shock identified with our baseline approach still explains a sizable portion of the variance of observed inflation at frequencies between 6 and 32 quarters. This shares is still larger than what obtained with a VAR over the same frequencies. Second, the results get very close to what obtained in our baseline analysis when considering longer cycles. The shock can explain around 34% and 39% of inflation volatility when considering cycles of $16-50$ and $32-50$ quarters duration, compared to around 41% under our baseline case.

Summarizing, the difference in results cannot be simply explained by the way the im-

Table 3: Variance contributions of unemployment-identified shock

(A) Contribution to cycles

Baseline

Unemployment	Output (cycle)	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
72.4	58.6	69.0	66.1	40.6	57.1
[61.8,85.8]	[46.9,71.0]	[57.2,81.2]	[44.5,85.3]	[19.1,66.6]	[31.6,80.8]

(B) Contribution to observables

Frequencies $6 - 32$ quarters

Unemployment	Output (level)	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
75.1	72.2	58.5	50.9	21.6	27.1
[60.6,79.8]	[51.5,83.0]	[45.7,65.6]	[32.0,68.2]	[9.4,42.3]	[12.7,45.3]

Frequencies $16 - 50$ quarters

Unemployment	Output (level)	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
74.2	77.8	60.1	56.0	33.7	36.3
[61.8,80.1]	[58.1,85.8]	[48.4,67.9]	[35.9,75.2]	[14.8,55.7]	[18.6,56.2]

Frequencies $32 - 50$ quarters

Unemployment	Output (level)	Unemployment exp.(1y)	Interest rate	Inflation	Inflation exp.(1y)
71.5	78.6	58.2	58.5	39.2	39.2
[58.9,78.9]	[59.6,86.1]	[46.6,67.1]	[37.9,78.0]	[18.0,61.9]	[20.2,59.5]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment cycle. The contribution of the identified shock to the variance of the cycle of all variables is evaluated (panel A). The contribution of the identified shock to the variance of all observables is evaluated over frequencies that imply cycles between 6 and 32 quarters, 16 and 50 quarters or 32 and 50 quarters (panel B). We report the median contribution and the corresponding 68-percent posterior-coverage intervals.

portance of the unemployment cycle shock for inflation volatility is measured. Even under a dogmatic view of what frequencies correspond to the business cycle, the two methods deliver different results and the gap opens up when longer frequencies are considered. As explained in the Introduction, one advantage of our approach is that a researcher does not need to take a stance on which frequencies correspond to the business cycle.

Aligning identification strategy and target We now go one step further to make the comparison between the VAR and TC-VAR even more direct. We use the TC-VAR to identify the unemployment shock in the frequency domain for overall unemployment and then assess its contribution to inflation volatility in the frequency domain. In other words, we replicate what the literature has done with the VAR, but using a TC-VAR. While we believe that our baseline approach fully exploits the parametric properties of the TC-VAR, this variation helps to further understand the gap between our findings and the seminal contribution by Angeletos et al. (2020).

We start by comparing the impulse responses for the two models based on the frequency bands employed by Angeletos et al. (2020), $6 - 32$ quarters, and an alternative range that removes high-frequency fluctuations, $16 - 50$ quarters. Results for the $32 - 50$ quarters case are similar and available upon request. Figure 5 reports the results: The shaded areas correspond to the 68% credible sets for the TC-VAR, while the dashed lines correspond to the 68% credible sets for the VAR. Despite using the same approach to identify the shocks, the impulse responses appear quite different. The impulse responses for inflation and interest rates present a significantly larger jump under the TC-VAR than under the VAR. Under the VAR, the response of inflation is more sluggish and builds over time. For output, instead, the initial response is quite similar, but differences open up over time. The TC-VAR implies a relatively quick reversal, while in the VAR the variable keeps increasing for a while. Finally, we observe some differences in the response of unemployment itself, with the VAR implying a more prolonged decline.

The last step consists of assessing whether these differences translate in a different contribution of the unemployment shock to inflation volatility at different frequencies. We report the results in Table 4. Panel (A) contains the results based on the TC-VAR, while Panel (B) reports the results based on the VAR. In each case, columns report results keeping fixed the frequency used to identify the shocks, while rows vary the frequencies used to assess the contribution to inflation. On the main diagonal we have the standard approach, in which the frequency bands coincide for the identification of the shock and the assessment of its contribution to inflation volatility.

Three main results emerge. First, across all frequency combinations, the TC-VAR delivers a larger contribution to inflation volatility, despite using the same data and a frequency

Figure 5: Impulse responses to unemployment shock identified in the frequency domain

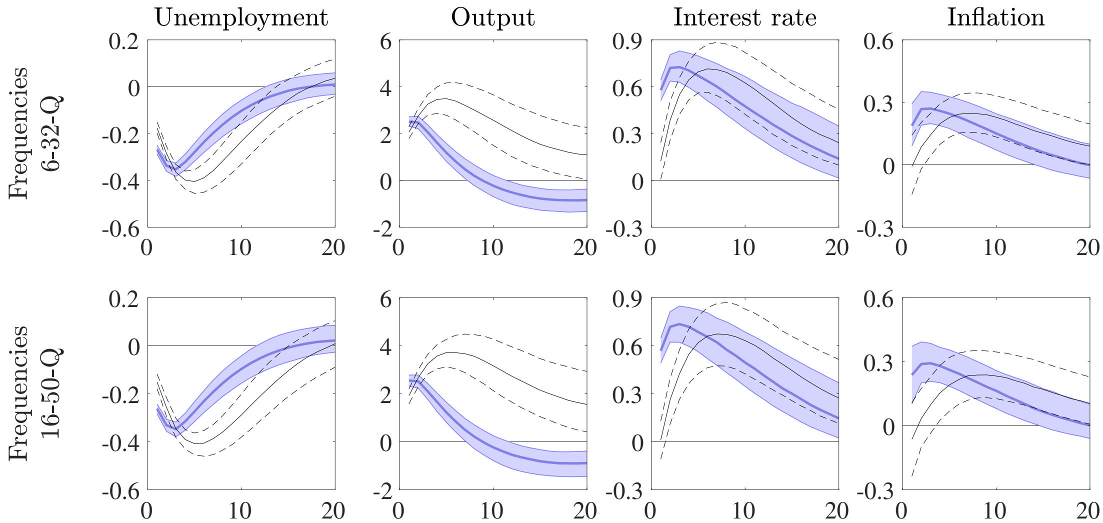

Notes: The figure shows the response of unemployment, output (log-level), interest rate and inflation to the unemployment shock identified in the frequency domain using a VAR and a TC-VAR. The shock is identified using either the baseline TC-VAR model or a VAR model and over frequencies that imply cycles between 6 and 32 quarters or 16 and 50 quarters. The figure plots the posterior median (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area) estimated using the baseline TC-VAR. 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). The figure plots the posterior median (black lines) and the corresponding 68-percent posterior-coverage intervals (dashed black lines) estimated using the VAR.

domain approach. Second, both the TC-VAR and VAR deliver larger contributions when longer cycles are considered. Specifically, if the unemployment shock is identified based on cycles of $6-32$ quarters, the contribution to inflation variability is 16.5% for the TC-VAR and 12.7% for the VAR. If the exercise is repeated using cycles of $32-50$ quarters, the contribution to inflation variability is 36.0% for the TC-VAR and 22.8% for the VAR. Third, the choice of frequency implies more drastic changes in the results for the TC-VAR than for the VAR.

These results suggest two important considerations. First, the parametric model fit to the data matters for the results, even when conditioning on a certain frequency. Second, the choice of frequencies matters for a given model. We believe that these two considerations validate the choice of focusing on the variability of cycles as implied by the TC-VAR, as opposed to using a frequency domain approach. This choice exploits the ability of the TC-VAR to separate the dynamics of trends and cycles and does not require to take a stance on which frequencies correspond to the business cycle.

Table 4: Contribution to inflation volatility of the unemployment shock identified in the frequency domain

(A) Based on a TC-VAR

Frequencies of contribution on inflation (quarters)	Frequencies of identification on unemployment (quarters)		
	$6-32$	$16-50$	$32-50$
$6-32$	16.5	19.5	20.7
	[9.5,24.9]	[10.5,30.4]	[10.2,35.2]
$16-50$	25.4	29.3	31.3
	[14.9,37.0]	[16.8,43.5]	[16.4,47.6]
$32-50$	29.7	33.2	36.0
	[17.8,42.2]	[20.1,49.0]	[20.0,53.1]

(B) Based on a VAR

$6-32$	12.7	13.8	13.8
	[6.8,21.6]	[7.3,22.8]	[6.9,23.7]
$16-50$	19.4	20.2	20.1
	[10.4,31.5]	[10.6,33.4]	[9.7,33.7]
$32-50$	22.1	23.0	22.8
	[11.4,35.9]	[11.0,37.8]	[10.3,38.8]

Notes: The table reports the contribution to the variance of inflation of the max-share unemployment shock identified using either the baseline TC-VAR model (panel A) or a VAR model (panel B). The shock is identified by maximizing its contribution to the volatility of unemployment. The frequency bands over which the shock is identified correspond to cycles between 6 and 32 quarters, 16 and 50 quarters, or 32 and 50 quarters. The contribution of the shock to the volatility of inflation is evaluated over those three frequency bands. 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. We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of inflation over the corresponding frequencies.

5.3 Forecasting performance of VAR and TC-VAR

We have shown above that 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 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):

$$\text{LDD} \equiv \frac{100}{2N_v}\left[\ln\left(\left|\frac{1}{\overline{T}-\underline{T}}\sum_{t=\underline{T}}^{\overline{T}}\zeta_t\zeta_t'\right|\right)^{\text{TC-VAR}} - \ln\left(\left|\frac{1}{\overline{T}-\underline{T}}\sum_{t=\underline{T}}^{\overline{T}}\zeta_t\zeta_t'\right|\right)^{\text{Alt.}}\right], \tag{15}$$

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 (15) compares the baseline TC-VAR with alternative models with $\underline{T} = 2000\text{:Q2}$, $\overline{T} = 2016\text{:Q4}$, 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 6 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. Lastly, we find that removing survey expectations from the TC-VAR model leads to a substantial deterioration in out-of-sample forecast performance, which provides support for their inclusion in the baseline estimation.

6 Reconciling TC-VAR and VAR Results

In the previous section, we showed that our baseline approach based on a TC-VAR leads to conclusions about the relationship between unemployment and inflation over the business cycle that differ markedly from those obtained using a max-share approach based on

Figure 6: Relative log-determinant differentials across models

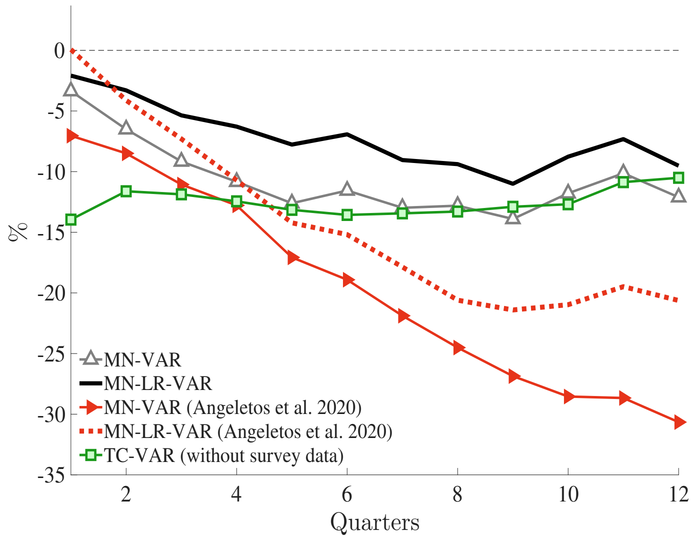

Notes: The figure reports the log-determinant differentials of different models with respect to the baseline TC-VAR. A negative value favors the baseline TC-VAR over alternative models. Forecasts for inflation, unemployment, and GDP growth are produced using recursive sample from 1960:Q1-2000:Q2 to 1960:Q1-2016:Q4. 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). We also examine a TC-VAR specification that excludes survey expectations.

a VAR. We also showed that these differences persist even when conducting the analysis in the frequency domain of the observables, rather than relying on the trend-cycle decomposition implied by the TC-VAR. Together, these findings suggest that the parametric representation of the data plays an important role in shaping the resulting conclusions.

This insight is perhaps not entirely surprising, given that the max-share approach ultimately relies on model parameters to reconstruct the systematic comovement between variables. When the data-generating process exhibits pronounced low-frequency movements alongside business-cycle fluctuations, a standard VAR imposes no structure that separates forces operating at different frequencies. By contrast, the TC-VAR explicitly models trend and cyclical components, which is especially relevant for variables such as inflation that display substantial low-frequency variation.

In this section, we draw on the work of Fernández-Villaverde et al. (2007) to formalize these ideas. We first derive the VAR(∞) representation of the TC-VAR. We then explore the implications of this mapping to address two fundamental questions. First, how does the mapping between the two models affect a max-share analysis of the type commonly used in the literature? Second, what are the consequences of working in the frequency domain as opposed to explicitly separating trend and cycle dynamics?

For ease of exposition, we restate the state-space representation introduced in Section 2:

$$z_t = \Lambda_x x_t, \qquad x_t = \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \qquad \varepsilon_t = L_\varepsilon w_t, \qquad E(w_t w_t') = \mathbf{I}. \tag{16}$$

We abstract away from measurement error in this analysis. The system in (16) can be cast into a forecast-error representation of the form

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + \hat{B}\nu_{t+1}, \\ z_{t+1} &= C\hat{x}_t + \hat{D}\nu_{t+1}, \end{aligned} \tag{17}$$

where $A = \Phi$ and $C = \Lambda_x \Phi$, $\hat{x}_t$ denotes the filtered state, and $\nu_t \sim (\mathbf{0}, \mathbf{I})$ denotes the one-step-ahead forecast error of z_t. In this representation, the matrices $\hat{B}$ and $\hat{D}$ govern the response of the filtered state and observables to forecast-error innovations. By contrast, the matrices $B = \mathcal{R}L_\varepsilon$ and $D = \Lambda_x \mathcal{R} L_\varepsilon$ characterize the response to the underlying structural shocks w_t in the population state-space system; see Fernández-Villaverde et al. (2007) for details. Iterating (17) backward yields the VAR(∞) representation

$$z_{t+1} = \sum_{s=0}^{\infty} C\left(A - \hat{B}\hat{D}^{-1}C\right)^s \hat{B}\hat{D}^{-1} z_{t-s} + \hat{D}\nu_{t+1}. \tag{18}$$

Equation (18) highlights two implications that are central for our analysis. First, when the data-generating process features both trend and cycle components, finite-order autoregressive representations can be misleading, as noted by Watson (1986). Second, even abstracting from lag truncation, the reduced-form innovations $\hat{D}\nu_{t+1}$ do not isolate structural shocks to trends and cycles, Dw_{t+1}, but instead combine them with forecast errors arising from the estimation of latent components, as emphasized by Fernández-Villaverde et al. (2007). Taken together, these features can have a material impact on variance decompositions and, in particular, on the contribution of max-share identified shocks, even before accounting for finite-sample considerations.

We consider a deliberately simple data-generating process to illustrate these points numerically. The TC-VAR features two observables, unemployment and inflation. Unemployment is assumed to be driven solely by a cyclical component and does not exhibit low-frequency variation, whereas inflation is driven by both trend and cyclical components. The cyclical components are jointly modeled as a VAR of order two, which is the minimal specification that allows for distinct persistence and richer cyclical dynamics across the two components. Its state-space representation is given by (16), with $x_t = [\tau_{\pi,t}, c_{u,t}, c_{\pi,t}, c_{u,t-1}, c_{\pi,t-1}]$, $z_t = [u_t, \pi_t]$, and the corresponding coefficient matrices:

$$\Lambda_x = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \end{bmatrix}, \tag{19}$$

$$\Phi = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & \Phi_{1,uu} & \Phi_{1,u\pi} & \Phi_{2,uu} & \Phi_{2,u\pi} \\ 0 & \Phi_{1,\pi u} & \Phi_{1,\pi\pi} & \Phi_{2,\pi u} & \Phi_{2,\pi\pi} \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{bmatrix}, \quad \mathcal{R} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \quad L_\varepsilon = \begin{bmatrix} \sigma_{\tau,\pi} & 0 & 0 \\ 0 & \sigma_{c,u} & 0 \\ 0 & 0 & \sigma_{c,\pi} \end{bmatrix}.$$

The simulation parameter values in (19) are reported in the legend to Figure 7. We chose the parameters to consider a case in which the unemployment shock explains a similar fraction of inflation and unemployment cycles: 83% and 78%, respectively, for a ratio between the two of around 1.06. We verified that if we were to simulate the model repeatedly and apply the baseline identification strategy, we would correctly recover this relation (82%, [69%, 89%], and 78%, [75%, 81%]), consistent with the fact that such strategy is aligned with the true data generating process. In what follows, we therefore focus on what would happen if the analysis were conducted in the frequency domain and if a VAR were used instead of the correctly specified TC-VAR.

Figure 7 reports the variance *at certain frequencies* explained by the identified max-share unemployment shock under three distinct identification strategies. First, we consider the baseline TC-VAR identification (solid line), which targets the variance of the latent unemployment cycle ($c_{u,t}$). While identification is based on the overall covariance structure of the cycle, in the figure we report the contribution of the shock to the variance of the observables at certain frequencies. This case serves as benchmark because the identification strategy is aligned with the data generating process. Second, we use the TC-VAR to identify the unemployment shock over specified frequency bands for observed unemployment (solid line with triangles). Finally, we apply an analogous frequency-based identification strategy using a VAR with p lags, with p varying from 1 to 40 (solid line with squares). In summary, the two TC-VAR specifications differ in whether the max-share shock is identified in the frequency domain or not, while the latter two specifications differ only in the underlying model used, holding identification in the frequency domain fixed.

Figure 7: Variance explained by the unemployment-identified shock

(A) Based on population coefficients

Frequency bands: 6 to 32 quarters | Frequency bands: 16 to 50 quarters

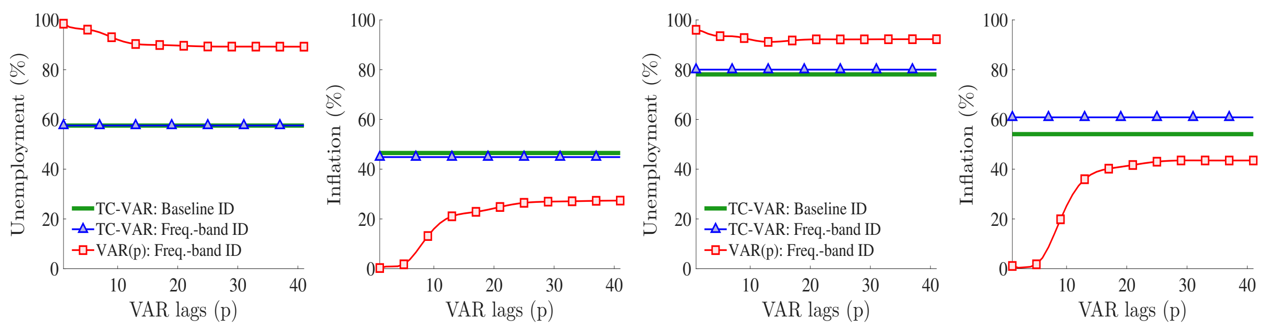

(B) Based on finite-sample estimates

Frequency bands: 6 to 32 quarters | Frequency bands: 16 to 50 quarters

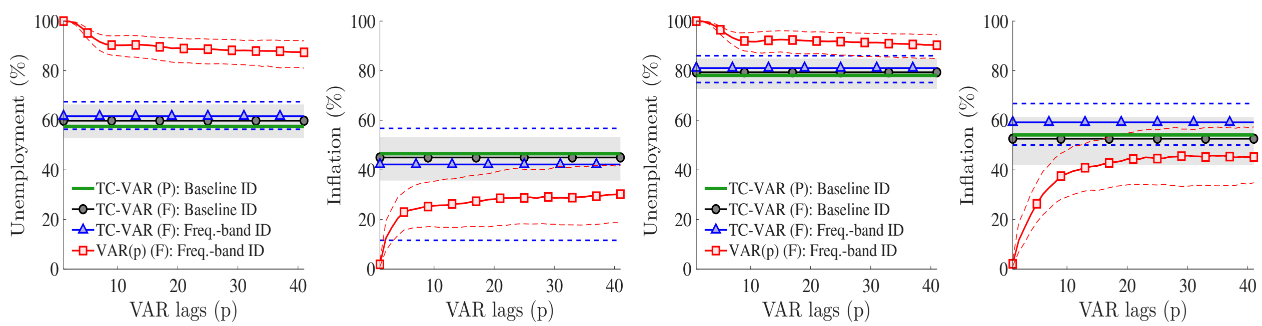

Notes: For panel (A), the corresponding coefficient matrices in (16) are calibrated as follows

$$\begin{bmatrix} \Phi_{1,uu} & \Phi_{1,u\pi} \\ \Phi_{1,\pi u} & \Phi_{1,\pi\pi} \end{bmatrix} = \begin{bmatrix} 1.6229 & -0.1122 \\ -0.1122 & 1.8112 \end{bmatrix}, \quad \begin{bmatrix} \Phi_{2,uu} & \Phi_{2,u\pi} \\ \Phi_{2,\pi u} & \Phi_{2,\pi\pi} \end{bmatrix} = \begin{bmatrix} -0.7546 & 0.0689 \\ 0.0689 & -0.8704 \end{bmatrix},$$

$\sigma_{\tau,\pi} = 1.0$, $\sigma_{c,u} = 0.2$, $\sigma_{c,\pi} = 0.2$. For panel (B), we simulate datasets of length $T = 1{,}000$ over $N = 500$ Monte Carlo simulation based on the calibrated parameter values reported in panel (A), and estimate the TC-VAR and VAR(p) models using MLE and OLS, respectively. Results are summarized by median estimates and 68% intervals.

In Panel (A), we report results based on knowledge of the parameters of the TC-VAR and its implied VAR representation. For the VAR, we show how the results vary as the VAR is truncated to $p = \{1, \ldots, 40\}$ lags. In Panel (B), we instead report results based on a large number of simulations in which we estimate the two models. For each panel, we consider two possible frequency bands: 6 to 32 and 16 to 50 quarters.

We begin with Panel (A), which is based on population coefficients. For frequencies between 16 and 32 quarters, the TC-VAR with identification in the frequency domain delivers results that are virtually identical to those obtained using the baseline identification strategy. In both cases, the TC-VAR implies that the unemployment-identified shock explains similar shares of unemployment and inflation volatility. Both shares are smaller than the ones implied by examining the variances of the corresponding cycles (78% and 83%, respectively, as reported above). By contrast, the VAR-based approach substantially overestimates the importance of the max-share shock for unemployment, while markedly underestimating its contribution to inflation. This discrepancy persists even as the lag length increases and is particularly severe when the VAR representation is truncated to a few lags. The max-share unemployment shock identified using a VAR(2) explains around 97% of the variation in unemployment, but less than 1% of the variation in inflation.

When the analysis is repeated for lower-frequency bands of 16 to 50 quarters, a VAR truncated to two lags still fails to capture the contribution of the shock to inflation volatility. VARs with several lags attribute greater importance to the unemployment shock, explaining around 44% of inflation volatility, but the corresponding ratio relative to unemployment, about 0.5, remains well below the TC-VAR benchmark under the baseline identification (0.7). By contrast, the TC-VAR with frequency-domain identification performs well, with the identified shock explaining 80% of unemployment volatility and 60% of inflation volatility, only slightly above the baseline value for inflation (around 54%).

Panel (B) reports variance decompositions based on finite-sample estimates. The insights from the population analysis carry over, with one difference: in finite samples, the estimated VAR(p) coefficients do not coincide with the coefficients implied by the corresponding p-lag truncation of (18), since the former are chosen to match finite-sample moments up to lag p. Accordingly, the results improve slightly for a VAR with only a few lags. However, even accounting for this better approximation, a VAR(2) implies that the identified unemployment shock explains around 99.85% (99.87%) of unemployment volatility and only 12.50% (11.72%) of inflation volatility at the 6 to 32 (16 to 50) quarters frequencies, yielding a ratio of 0.13 (0.12). This leaves a large gap relative to the results obtained when fitting a TC-VAR, that still recovers a strong relation between the two variables even when identifying shocks in the frequency domain.

Overall, two key features emerge from this example. First, variance decompositions are highly sensitive to the choice of frequency band, implying that quantitative conclusions depend critically on how the business cycle is defined. This sensitivity persists even when using a TC-VAR. In this respect, our baseline approach circumvents the problem by working directly with the cyclical representation. Second, even with very long lag orders, a VAR(p) can systematically understate the contribution of the max-share shock relative to the TC-VAR because it cannot correctly separate shocks to trend and shocks to cycle, a point that echoes the results presented in Fernández-Villaverde et al. (2007). This issue is especially salient given that standard lag choices in quarterly VAR applications are typically small, often in the single digits (e.g., two lags in Angeletos et al., 2020). In this case, the VAR is a very poor approximation of the TC-VAR. Although, in principle, more lags could be included, doing so with a finite amount of data raises practical challenges due to the proliferation of parameters. The TC-VAR has the advantage of being a more tightly parameterized model that maps into an infinite order VAR with cross-equation restrictions.

7 Conclusions

In this paper, we adopt a TC-VAR to study the relationship between inflation and real activity over the business cycle. A TC-VAR has the advantage of explicitly controlling for low-frequency variation that can otherwise contaminate inference about cyclical relationships. We show that the shock explaining the largest share of the unemployment cycle also accounts for a large fraction of the inflation cycle. We then investigate why our results differ from those obtained using a VAR with frequency-domain identification. Two features play an important role. First, a VAR with a small number of lags is typically a poor approximation to a TC-VAR, and even an infinite-order VAR may fail to separate trend and cycle shocks. Second, the results are highly sensitive to the choice of frequency bands. Our baseline approach has the advantage of explicitly modeling trends and cycles while avoiding assumptions about the relevant frequencies.

We outline three directions for future research. First, it would be valuable to investigate the drivers of low-frequency movements in inflation and unemployment, as the current model is silent on the forces shaping these trends. Second, our methodology could be applied to study the cyclical behavior of other key macroeconomic variables that exhibit pronounced 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). Finally, macroeconomists should work toward developing micro-founded models that allow for changes in trends.

References

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.

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.

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.