

Inflation and Real Activity over the Business Cycle

Francesco Bianchi*
Johns Hopkins University
NBER and CEPR

Giovanni Nicolò
Federal Reserve Board

Dongho Song
Johns Hopkins University
Carey Business School

October 26, 2022

Abstract

We use a Trend-Cycle VAR model to study the relation between inflation and the real economy over the business cycle. The Trend-Cycle VAR model allows us to control for low-frequency movements in inflation and real activity. We show that at business-cycle frequencies, fluctuations of inflation are tightly related to movements in real activity, in line with what is implied by the New Keynesian framework. We explain why our results differ from previous studies that argue inflation is unrelated to real activity at business-cycle frequencies.

*Bianchi: Department of Economics, Wyman Park Building, 3100 Wyman Park Drive, Baltimore, MD 21211 (francesco.bianchi@jhu.edu). Corresponding author: Nicolò, Federal Reserve Board, 20th Street and Constitution Avenue N.W., Washington, DC 20551 (giovanni.nicolo@frb.gov). Song: Johns Hopkins Carey Business School, 100 International Drive, Baltimore, MD 21202 (dongho.song@jhu.edu).

We are grateful to Domenico Giannone who provided comments that led to improvements in our work. 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 or the Federal Reserve System.

1 Introduction

A proper understanding of the empirical relationship between real and nominal macroeconomic variables is central to our ability to discern across different macroeconomic models and to validate their theoretical foundations. In turn, these models provide a rigorous framework to describe how shocks transmit through the economy and to design policies meant to deliver a more stable macroeconomic environment.

A large New Keynesian literature relies on empirical evidence suggesting that, in response to demand shocks, various measures of real economic activity are related to inflation at business-cycle frequencies. In New Keynesian models, this relationship is explained in light of nominal rigidities, leading to a structural price Phillips curve. In this theoretical framework, a positive demand shock generates a positive gap of the measures of real economic activity relative to their flexible-price counterparts and ultimately leads to inflationary pressures.

However, the increasing evidence of a weaker relationship between real economic activity and inflation since the 1990s has led some economists to rethink the empirical foundations of New Keynesian models. This apparent weak relation between inflation and real activity has led to different, but not necessarily mutually exclusive responses. A series of recent contributions focuses on identifying explanations for the weakening of this empirical relationship ([Primiceri, 2013](#), [Del Negro et al., 2020](#) and references therein) or on modifying features of the New Keynesian model to improve its empirical fit ([Gust et al., 2022](#)). By contrast, other authors have interpreted this evidence as a reason to abandon the New Keynesian apparatus and have proposed different frameworks in which non-inflationary demand shocks are the main drivers of business-cycle fluctuations ([Beaudry and Portier, 2013](#); [Angeletos et al., 2018](#)). Alternatively, other studies develop business-cycle theories that entirely abstract from discussing implications for inflation ([Beaudry et al., 2020](#); [Basu et al., 2021](#)).

An important empirical justification for these alternative theoretical frameworks is offered by [Angeletos et al. \(2020\)](#). Looking at the U.S. post-war period, [Angeletos et al. \(2020\)](#) find that business-cycle fluctuations are explained by a real activity shock which is largely disconnected from movements in inflation. In their clever empirical analysis, the authors follow an extensive empirical literature that uses vector autoregressions (VARs) as a “model free,” but still structural way to look at the data. The authors use the VAR to identify a “business-cycle” shock (or a combination of shocks) that explains the largest possible share of real activity variation at business-cycle frequency. They then show that while this single shock can explain most of real activity at business-cycle frequency, it explains close to nothing of the business-cycle variation of inflation. They conclude that their results are at odds with the premise of the NK framework that demand shocks lead to business cycle fluctuations because of nominal rigidities that prevent an immediate adjustment in prices.

As we argue next, the approach of using a VAR to identify shocks is conventional in the literature, but has relevant implications when focusing on the specific question that [Angeletos et al. \(2020\)](#) were after, the business-cycle relationship between real and nominal variables. The main reason is that a standard VAR might be unable to correctly disentangle business-cycle and low-frequency movements in those variables over a relatively short period of time. This is because in a VAR, a single set of parameters and reduced form shocks need to accommodate the variation at all frequencies observed over a relatively short period. As a result, a procedure that uses the estimated parameters and reduced form shocks to identify variation at business-cycle frequency might be biased. The problem is particularly severe if one of the variables of interest shows significant variation at low frequency, as it is the case with inflation. Ultimately, the identified shock might fail to capture a business-cycle relationship between the two variables even when such a relation is in fact in the data. To remedy this limitation of the VAR for the specific question at hand, we adopt a more flexible model that explicitly extracts business-cycle movements in the variables of interest. Specifically, we argue that a Trend-Cycle VAR (TC-VAR) model is better suited to analyze the relation between inflation and real activity at business-cycle frequencies.¹

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

Motivated by this analysis, we adopt a more rigorous empirical framework and estimate a multivariate Trend-Cycle VAR model building on the work of [Watson \(1986\)](#), [Stock and Watson \(1988, 2007\)](#), [Villani \(2009\)](#) and, more recently, [Del Negro et al. \(2017\)](#). We consider the sample between 1960 and 2019 using six time series. The first four time series are commonly used in previous studies: i) the growth rate of real GDP per capita; ii) the unemployment rate; iii) the effective federal funds rate; and iv) inflation. Importantly, we include two additional time series. First, to

¹We do not mean that the approach that we propose in the paper is always superior to using a VAR. In many cases, the goal is not to isolate the effects of a shock at a particular frequency.

better capture low-frequency movements in inflation, we add a measure of average 10-year-ahead inflation expectations by combining information from the Survey of Professional Forecasters (SPF) and Blue Chip (BC) Economic Indicators survey as in [Del Negro and Schorfheide \(2013\)](#). Second, we use the median 1-year-ahead average inflation expectations from the SPF as a measure of inflation that abstracts from the high-frequency movements observed in the time series for inflation, but it is still likely to be related to business-cycle movements in inflation.

In our baseline specification, we assume that the trend of the real per-capita GDP growth, unemployment rate, and real interest rate are uncorrelated with each other and trend inflation. By contrast, we consider a common trend that captures low-frequency movements in realized inflation and 1- and 10-year-ahead inflation expectations. Each of the three inflation measures is driven by a cyclical component. We also allow for a (negative) moving average component in inflation to control for high-frequency movements in inflation. The moving average component allows to purge movements in inflation from observation errors and transitory movements in prices. Given that our model already separates trend- from cyclical variation, we identify the shock that maximizes the variation of the latent cyclical component of the unemployment rate and study its contribution to the volatility of the cyclical components of the remaining variables. We also consider an alternative specification in which we remove frequencies that correspond to fluctuations with duration of less than 1.5 years. This allows us to clean the results from residual high-frequency movements that are arguably unrelated to the business cycle.

This approach has four clear advantages with respect to identifying shocks in the frequency domain based on a VAR.² First, the inference exercise automatically separates the trends from the cycles. Second, cyclical variation is controlled by a different set of parameters with respect to low-frequency variation. Third, we do not need to take a stance on the typical length of the business cycle. This is important in light of the fact that expansions have become progressively longer in the sample under consideration. Finally, by allowing for changes in trend growth, trend inflation, and variation in long-run unemployment, the model accommodates the notion that what matters for cyclical movements in inflation is the output (or unemployment) gap.

A series of important results emerge from this exercise. First, under our baseline specification, the shock targeting the unemployment rate explains about 45% of the volatility of inflation over the 1960-2019 period. This result suggests that it is important to account for the low-frequency movements in real and nominal variables when studying their relationship at business-cycle frequencies. Second, the shock targeting the unemployment rate explains nearly 60% of the business-cycle volatility of both inflation expectation measures. This result provides further support for the notion that business-cycle movements in inflation are in fact related to the business cycle, given that expected inflation is obviously related to actual inflation, but less affected by high-frequency

²The approach of identifying shocks in the frequency domain starting from a VAR was pioneered by [Uhlig \(2003\)](#). The approach has lately adopted by many others, including [Giannone et al. \(2019a\)](#), [Angeletos et al. \(2020\)](#) and [Basu et al. \(2021\)](#).

fluctuations. Furthermore, the result is in line with the NK framework, in which agents' inflation expectations depend on the state of the economy. In line with this reasoning, when we focus on frequencies that correspond to fluctuations *in the cyclical component* with duration of at least 1.5 years, the results become stronger. In this case, the shock identified targeting the rate of unemployment explains a higher percentage of the business-cycle volatility of all inflation measures compared to when all the frequencies are considered: 50% for realized inflation and 65% for both inflation expectation measures. Once again, this result is in line with the evidence presented in our motivating evidence exercise, suggesting a strong relation between inflation and real activity over the business-cycle.

All previous results are robust to two alternative specifications. Rather than allowing for an idiosyncratic cyclical component for realized inflation and the two measures of inflation expectations, the first alternative specification admits a cyclical component in realized inflation and a common cycle for the two measures of inflation expectation. The second alternative specification instead allows a common cycle between realized inflation and the one-year-ahead inflation expectations. The 10-year-ahead inflation expectations are instead driven only by the common inflation trend in addition to being subject to measurement errors. For both alternative specifications, the main result applies: the shock identified targeting the cyclical component of the unemployment rate explains large portions of the variability of the alternative cyclical inflation measures considered.

In addition, we assess the importance for our results of modeling time-varying latent trends. If we model the latent trends as constant, rather than time-varying, we find evidence of a disconnect between nominal and real variables, thus recovering the results of a disconnect between real and nominal variables commonly found when adopting a standard VAR model (Giannone et al., 2019b; Angeletos et al., 2020). Crucially, this finding shows that the TC-VAR model is flexible enough to deliver results in line with the existing literature, but the data do not support this evidence, given that they present important trends that can confound the business-cycle analysis.

We then verify that a standard VAR model cannot easily capture the uncovered business-cycle relationship between nominal and real variables even if we choose priors and sample periods to account for low-frequency movements in the data. Given our focus on the U.S. economy, we borrow the framework of Angeletos et al. (2020) who estimate a VAR model with a Minnesota prior on U.S. data. Extending their dataset by two years to the end of 2019, we estimate their baseline VAR model with a Minnesota prior as well as combining it with long-run priors *à la* Giannone et al. (2019a). For each specification, we follow their approach and identify a shock that targets the unemployment rate at business-cycle frequencies. The contribution of the shock to the variability of inflation at the same frequencies only marginally increases from about 8% when using a Minnesota prior to about 11% when also assuming long-run priors. Both estimates are significantly lower than our baseline estimates between 45% and 50%.

Better results emerge when we focus on the Great Moderation period, a sample that presents

less low-frequency variation in inflation. As before, we use the dataset of [Angeletos et al. \(2020\)](#). However, we estimate a VAR model over the period between 1984 and 2008 and assuming alternative priors. Specifically, we allow for combinations of both Minnesota and long-run priors for which either we set the corresponding hyperparameters in a dogmatic way or we optimize them. The largest contributions of the unemployment-identified shock results when the hyperparameters of the long-run priors are either optimized or set in a dogmatic manner, regardless of whether we also assume a Minnesota prior. However, those estimates are still between 30% and 45% lower than under our baseline specification, respectively. For a more direct comparison to our results, we also consider our dataset over the period of the Great Moderation. We estimate a VAR model using a Minnesota prior alone or combining it with long-run priors. The contribution of the unemployment-identified shock ranges between 13% and 20%, thus substantially lower than our baseline estimate.

Finally, we lay out theoretical arguments for why our results differ from recent contributions that challenge the validity of the New Keynesian paradigm. We demonstrate that a fixed-coefficient VAR estimated over a period of time that presents structural changes is misspecified, if the goal is trying to assess the comovement at business-cycle frequency. The misspecification problem associated with the use of a VAR model to describe a data generating process characterized by both low- and high-frequency movements cannot be resolved. An econometrician would need infinite data to reconstruct the VAR representation of a TC-VAR model. Even in that case, the reduced-form innovations that she would recover would map into the innovations affecting the latent persistent and stationary components *as well as* the error associated with the estimates of the latent components. In reality, these issues are exacerbated by the fact that the VAR parameters estimated over a finite sample would be distorted because a single set of parameters need to account for both trend and cycle fluctuations.

We provide an illustration of these issues with a simple example. We generate a Monte Carlo simulation of a TC-VAR model of unemployment rate and inflation under alternative calibrations that reflect different degrees to which the cyclical component of inflation depends on the cyclical component of unemployment rate. Even for a long simulation in which *realized* inflation is unrealistically assumed to depend exclusively on the *cyclical* component of the unemployment rate, the estimation of a VAR leads to an unemployment-identified shock explaining up to only 9% of the variability of inflation over the business cycle.

Overall, our findings have implications for both the New Keynesian literature and the recent and growing literature that proposes alternative explanations for the sources of business-cycle fluctuations. For the former, the results support the evidence of a relationship between real and nominal variables over the business cycle, while highlighting the importance of accounting for both low-frequency movements in real and nominal variables and the observed high-frequency movements in inflation to properly quantify that relationship. For the latter, the results suggest that the

alternative explanations for the drivers of business cycles should also propose transmission mechanisms which are consistent with the empirical evidence on the connection between movements in inflation and measures of real economic activity.³

Our results connect to [Sargent and Sims \(1977\)](#) that shows that two dynamic factors could explain a large fraction of the variance of a series of important macroeconomic variables, including output, employment, and prices. Two factors are in fact necessary to control for the low-frequency movements in nominal variables. The subsequent factor-analysis literature has repeatedly confirmed this key insight ([Giannone et al., 2006](#); [Watson, 2004](#); [Stock and Watson, 2011](#)). A related literature uses unobserved component models to estimate measures of 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 \(2021\)](#) only to cite a few).

[Ascari and Fosso \(2021\)](#) use a TC-VAR to study the role of imported intermediate goods in explaining the lack of sensitivity of inflation to business-cycle movements in the post-Millennial period. These authors focus on the changes in the importance of the business-cycle shock for inflation. They find that the contribution of the shock decreases from about 60% for the period starting in 1960 and ending in 1984 to nearly 30% for the period between 1985 and 2019. In line with our results, these estimates are about four and six times larger than the counterparts of 17% and 5% that [Angeletos et al. \(2020\)](#) find for the pre- and post-Volcker periods using a VAR and interpret as evidence of a disconnect. Our paper provides an explanation to reconcile these differences. Our results are also consistent with [Justiniano et al. \(2013\)](#). These authors use an estimated NK DSGE model to show that most inflation fluctuations are demand driven and commove with the output gap. Our results are obtained without imposing at the onset the NK framework, something that we find valuable if the goal is to provide external validation for such framework.

2 Motivating Empirical Evidence

In this section, we provide motivating evidence for our subsequent empirical analysis. We aim to show that inflation and real activity are related over the business cycle once we control for their trends. In particular, we underscore the importance of adopting a unified empirical framework that distinguishes between business-cycle and low-frequency movements in these variables. Moreover, given that our goal is to study the relationship between inflation and real economic activity over the business cycle, we also want to control for higher-frequency movements. These are especially important for the dynamics of inflation.

³For an alternative theory that integrates Keynesian economics with general equilibrium theory without relying on nominal rigidities, refer to [Farmer and Platonov \(2019\)](#) and [Farmer and Nicolò \(2018, 2019\)](#).

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

Business-cycle frequencies (6-32 quarters)					
	$j = 8$	$j = 6$	$j = 4$	$j = 2$	$j = 0$
Output	0.15	0.36	0.47	0.42	0.22
Unemployment rate	0.09	-0.13	-0.34	-0.44	-0.36
Medium-cycle frequencies (8-30 years)					
	$j = 8$	$j = 6$	$j = 4$	$j = 2$	$j = 0$
Output	0.64	0.61	0.54	0.45	0.33
Unemployment rate	-0.54	-0.50	-0.44	-0.36	-0.25

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

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

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


To formalize the notion that cyclical fluctuations in inflation appear to comove with real activity, Table 1 reports the correlations between current (filtered) inflation and current and lagged (filtered) levels of real per-capita GDP and unemployment rate at time $(t - j)$ for $j = \{0, 2, 6, 8\}$. We consider

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

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



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



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

both business-cycle and medium-cycle frequencies. Over the business cycle, the positive (negative) correlation of inflation peaks with real per-capita GDP (unemployment rate) lagged by four (two) quarters at about 0.45 (negative 0.45). Over the medium cycle, the correlation of inflation with real per-capita GDP (unemployment rate) lagged by eight quarters peaks at about 0.65 (negative 0.55). These correlations are larger (in absolute value) than the counterparts with real per-capita GDP and the rate of unemployment lagged by four quarters over the business cycle.

The empirical evidence presented in this section motivates us to adopt a dynamic, multivariate framework that allows to study the relationship between inflation and real economic activity over the business cycle and without the analysis being contaminated by movements at high and low frequencies. We discuss the adopted framework in the next section.

3 The Trend-Cycle VAR Model

In this section, we present the TC-VAR used to model the joint dynamics of the growth rate of real GDP per capita, the unemployment rate, the federal funds rate, the inflation rate, and the 1- and 10-year-ahead inflation expectations.

3.1 Baseline specification

Our baseline specification allows for four trend and six cycles. Real per-capita GDP growth g_t and the unemployment rate u_t evolve according to

$$\begin{aligned} g_t &= \tau_{g,t} + (c_{y,t} - c_{y,t-1}), \\ u_t &= \tau_{u,t} + c_{u,t}, \end{aligned}$$

where we assume that the trend of the unemployment rate $\tau_{u,t}$ and the growth rate of real, per-capita GDP $\tau_{g,t}$ are unrelated. It is worth emphasizing that we model the cycle in real GDP, as opposed to GDP growth. This is consistent with typical macroeconomic models in which output moves around a stochastic trend.

We assume that the Fisher equation holds in the long run. This implies that the federal funds rate can be expressed as

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

where $\tau_{r,t}$ and $\tau_{\pi,t}$ correspond to the trend of real interest rate and inflation respectively, and $\tau_{f,t} = \tau_{r,t} + \tau_{\pi,t}$. Finally, realized inflation and the 1- and 10-year-ahead inflation expectations are

decomposed as

$$\pi_t = \tau_{\pi,t} + c_{\pi,t} + (\eta_{\pi,t} - \eta_{\pi,t-1}), \quad (1.1)$$

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^{e,1y}, \quad (1.2)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + c_{\pi,t}^{e,10y}. \quad (1.3)$$

thus sharing a common trend $\tau_{\pi,t}$. We allow for an iid measurement error in the log-level of the GDP deflator, $\eta_{\pi,t}$, which implies that, after taking log difference, realized inflation features a negative moving average measurement error component. In the baseline, inflation and inflation expectations feature their own cyclical components. The VAR component below allows these cyclical components to comove over time, but we do not impose such a restriction here.

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

$$y_t = \left\{ g_t, u_t, f_t, \pi_t, \pi_t^{e,1y}, \pi_t^{e,10y} \right\}', \quad \tau_t = \left\{ \tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t} \right\}', \quad c_t = \left\{ c_{y,t}, c_{u,t}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^{e,1y}, c_{\pi,t}^{e,10y} \right\}',$$

$$\eta_t = \eta_{\pi,t}, \quad \varepsilon_{\tau,t} = \left\{ \varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t} \right\}', \quad \varepsilon_{c,t} = \left\{ \varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}, \varepsilon_{c,\pi,t}^{e,1y}, \varepsilon_{c,\pi,t}^{e,10y} \right\}'.$$

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

$$\tau_t = \tau_{t-1} + \varepsilon_{\tau,t}, \quad (2.1)$$

$$c_t = \Phi_1 c_{t-1} + \Phi_2 c_{t-2} + \dots + \Phi_p c_{t-p} + \varepsilon_{c,t}, \quad (2.2)$$

and the measurement errors follow

$$\eta_t = \varepsilon_{\eta,t}.$$

We assume that the shocks are independent and identically distributed as

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

where the Σ s are conforming positive definite matrices, $\Sigma = E(QQ')$, and $\mathcal{N}(\cdot, \cdot)$ denotes the multivariate Gaussian distribution.

3.2 Alternative specifications

For robustness, we also consider two alternative specifications for the 1- and 10-year-ahead inflation expectations, while leaving unchanged the decomposition assumed for the other variables. In the first alternative specification, we consider the following alternative decomposition for the two

inflation expectations surveys

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^e, \quad (3.1)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + \delta c_{\pi,t}^e + \eta_{\pi,t}^{e,10y}. \quad (3.2)$$

The 1- and 10-year-ahead inflation expectation surveys are assumed to share a common cyclical component, $c_{\pi,t}^e$. Because the 10-year inflation expectations $\pi_t^{e,10y}$ are by definition related to short-term inflation expectations but less volatile over time, we assume that its loading is less than one $\delta < 1$. Lastly, we allow for idiosyncratic errors for the 10-year-ahead inflation expectations. Under our baseline specification, we collect observables and state variables in vectors (6) where $\tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}\}'$, $c_t = \{c_{y,t}, c_{u,t}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^e\}'$, $\eta_t = \{\eta_{\pi,t}, \eta_{\pi,t}^{e,10y}\}'$, and $\varepsilon_{c,t} = \{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}, \varepsilon_{c,\pi,t}^e\}'$.

In the second alternative specification, we assume that the 1- and 10-year-ahead inflation expectation measures are decomposed as

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t} + \eta_{\pi,t}^{e,1y}, \quad (4.1)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + \eta_{\pi,t}^{e,10y}, \quad (4.2)$$

thus assuming that there is one common cyclical component $c_{\pi,t}$ for inflation which is shared across realized inflation in (1.1) and the 1-year-ahead inflation expectation. Because the 10-year inflation expectation $\pi_t^{e,10y}$ is fairly stable over time, we assume that it proxys for the trend component but does not include any cyclical component. Under this alternative specification, we keep the same assumptions on the standard deviation of the shocks to the trends as under the baseline. In this case, the vectors of observables and state variables are $\tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}\}'$, $c_t = \{c_{y,t}, c_{u,t}, c_{f,t}, c_{\pi,t}\}'$, $\eta_t = \{\eta_{\pi,t}, \eta_{\pi,t}^{e,1y}, \eta_{\pi,t}^{e,10y}\}'$, and $\varepsilon_{c,t} = \{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}\}'$.

3.3 State-space representation

Our baseline and alternative specifications can be cast into a state-space representation.⁴ Using generic notation, let us begin with n observables which can be decomposed into n_τ trends and n_c cycles, where $0 < n_\tau \leq n$ and $0 < n_c \leq n$.

Measurement equation. Allowing for n_η measurement errors η_t , observables can be expressed as

$$y_t = \Lambda x_t = \Lambda_\tau x_{\tau,t} + \Lambda_c x_{c,t} + \Lambda_\eta x_{\eta,t}, \quad (5)$$

where $x_t = \{x_{\tau,t}, x_{c,t}, x_{\eta,t}\}'$, $x_{\tau,t} = \tau_t$, $x_{c,t} = \{c_t, c_{t-1}, \dots, c_{t-(p-1)}\}'$, $x_{\eta,t} = \{\eta_t, \eta_{t-1}\}'$ and p

⁴Appendix A provides details on the construction of the matrices in (5) and (6) for our baseline specification described in Subsection 3.1.

denotes the number of lags used to model the stationary cyclical component. The $n \times n_\tau$ matrix Λ_τ captures $(n - n_\tau)$ cointegrating relationships, while $\Lambda_c = [\Lambda_{c,0}, \dots, \Lambda_{c,p-1}]$ and $\Lambda_\eta = [\Lambda_{\eta,0}, \Lambda_{\eta,1}]$.

State-transition equation. The vector of states x_t evolves as

$$x_t = \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \quad (6)$$

or equivalently,

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

where

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

The initial conditions are distributed as

$$\tau_0 \sim \mathcal{N}(\underline{\tau}, \underline{V}_\tau), \quad c_0 \sim \mathcal{N}(0, \underline{V}_c), \quad (7)$$

where \underline{V}_τ is an identity matrix, and \underline{V}_c is the unconditional variance of c_0 consistent with (6) and thus a function of the VAR coefficients $\varphi = \{\Phi_1, \dots, \Phi_p\}'$ and variance Σ_c .

4 Inference

We now describe the data and priors used in our analysis and the methodology used to assess the strength of the relation between inflation and real activity over the business cycle.

4.1 Data

We estimate the TC-VAR model using the following six quarterly time series which are expressed at annualized rates: i) the growth rate of real, per-capita GDP g_t ; ii) the unemployment rate u_t ; iii) the effective federal funds rate (EFFR) f_t by treating observations at the zero lower bound (ZLB) as missing following [Del Negro et al. \(2017\)](#); iv) the inflation rate π_t , measured as the log difference in GDP deflator (PGDP); v) the median four-quarter-ahead average PGDP inflation expectations, $\pi_t^{e,1y}$, from the SPF; vi) a measure of 10-year-ahead inflation expectations, $\pi_t^{e,10y}$, which, following [Del Negro and Schorfheide \(2013\)](#), we construct by combining survey expectations

on average 10-year-ahead CPI inflation from the SPF and BC 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 model over the period from 1960:Q1 to 2019:Q4. Finally, in Appendix B, we report details about the Bayesian inference, including settings and the Gibbs sampler adopted to estimate the TC-VAR model.

4.2 Priors and initial conditions

For our assumptions about initial conditions and prior distributions, we mainly follow the approach of Del Negro et al. (2017). We consider standard priors for covariance matrices Σ_τ and Σ_c and for the VAR coefficients $\varphi = \{\Phi_1, \dots, \Phi_p\}'$

$$p(\Sigma_\tau) = \mathcal{IW}(\kappa_\tau(\kappa_\tau + n_\tau + 1)\underline{\Sigma}_\tau), \quad (8.1)$$

$$p(\Sigma_c) = \mathcal{IW}(\kappa_c(\kappa_c + n_c + 1)\underline{\Sigma}_c), \quad (8.2)$$

$$p(\phi|\Sigma_c) = \mathcal{N}(\underline{\phi}, \Sigma_c \otimes \underline{\Omega}) \mathcal{I}(\phi), \quad (8.3)$$

where $\phi = \text{vec}(\varphi)$, $\underline{\phi} = \text{vec}(\underline{\varphi})$, $\mathcal{IW} = (\kappa(\kappa + n + 1)\underline{\Sigma})$ corresponds to the inverse Wishart distribution with mode $\underline{\Sigma}$ and k degrees of freedom, and $\mathcal{I}(\phi)$ is an indicator function that equals 0 if the VAR in (6) is explosive and 1 otherwise.

We center the prior distribution for the initial conditions of the trends τ_0 in (7) to the pre-sample mean of the corresponding variables. For the annualized trend real, per-capita GDP growth and the unemployment rate, the prior for their initial condition is set to 1% and 5%. For the trend real interest rate and inflation at annualized rate, the priors for their initial condition are centered at 0.1% and 2.5%. Therefore, the vector of initial conditions for the trend components is $\underline{\tau} = \{1, 5, 0.1, 2.5\}$.

To specify the prior for the covariance matrix of the shocks to the trends Σ_τ in (8.1), we assume that those shocks are *a priori* uncorrelated. We then set the standard deviation for the expected change in the annualized trend of real, per-capita GDP growth to 1% over a time period of 40 years. For all the remaining variables, we assume a one-percent standard deviation for the expected change in their trends over 20 years. These assumptions imply that the prior covariance matrix of the shocks to the trends is diagonal with the following elements on the main diagonal $\text{diag}(\underline{\Sigma}_\tau) = [1/40, 1/20, 1/20, 1/20]$. As in Del Negro et al. (2020), we assume a tight prior by setting κ_τ to 100.

The shocks to the cyclical components Σ_c in (8.2) are also assumed to be uncorrelated *a priori*. We calibrate the standard deviation of the shocks affecting the stationary component of the (annualized) real, per-capita GDP growth and the unemployment rate to 5% and 1.1%, reflecting their pre-sample standard deviations. The standard deviation of the shocks affecting the cycle 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 assumptions for the priors on the standard deviation of the cyclical components of the two measures of inflation expectations. Because these surveys are not available for the pre-sample period, we assume those standard deviations to be smaller than the pre-sample counterpart of 1.5% for realized inflation. For the one-year-ahead inflation expectations, we set the prior for the standard deviation of the cyclical component to 1.2% and assume it to be 0.3% for the ten-year-ahead inflation expectations. As a result of these assumptions, the prior covariance matrix of the shocks to the cycles is diagonal and such that the values on the main diagonal approximately correspond to $\text{diag}(\underline{\Sigma}_c) = [25, 1.3, 0.6, 2.2, 1.4, 0.1]$. 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 (8.3), we assume a conventional Minnesota prior with hyperparameter for the overall tightness equal to 0.2 in line with [Giannone et al. \(2015\)](#). Because the cyclical component in (2.2) is assumed to be stationary, we then center the prior for each variable's own lag to 0, rather than 1, as in [Del Negro et al. \(2017\)](#).

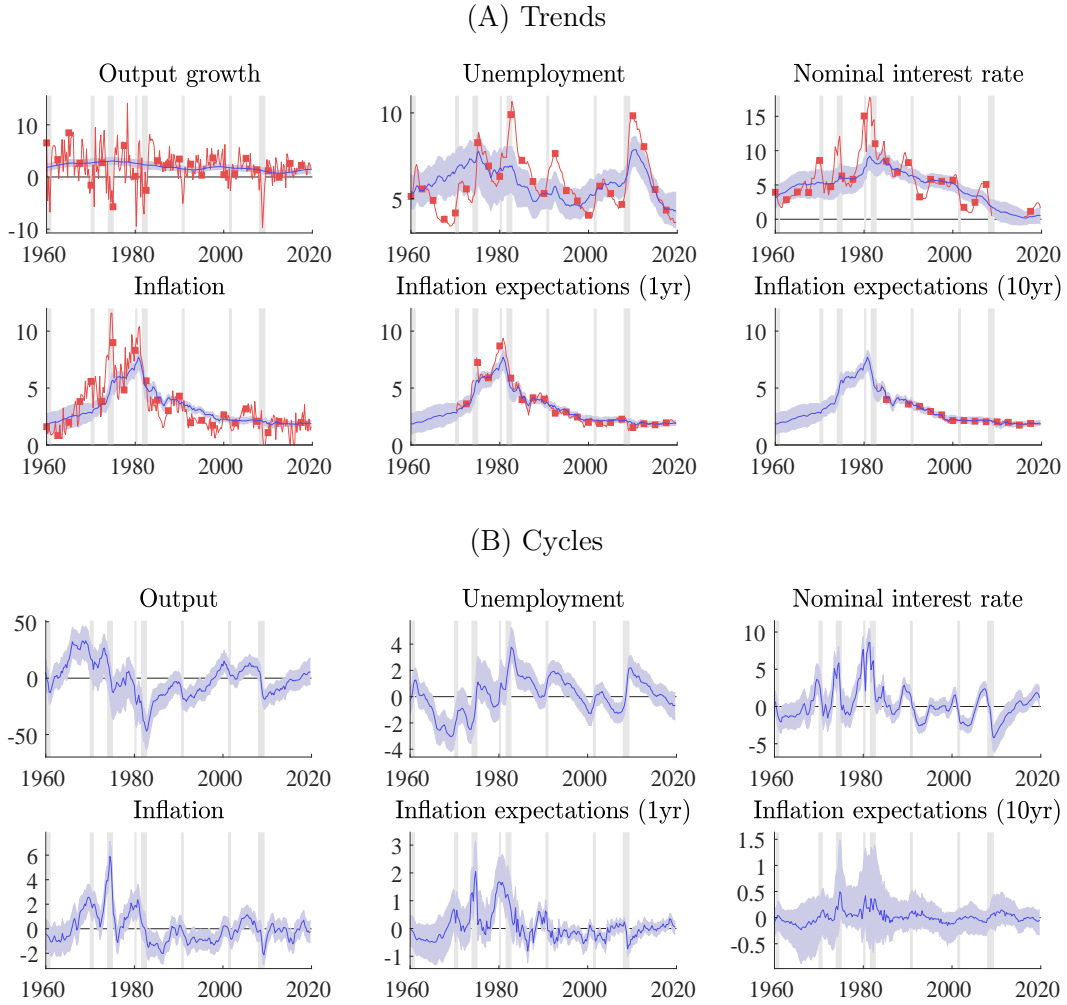
4.3 Identifying shocks that drive business-cycle fluctuations

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

5 Results

In this section, we present the main results of the paper. We first present the decomposition of the variables in trends and cycles. We then proceed to analyze how much of the cyclical variation of inflation and inflation expectations can be explained by the unemployment-identified shocks. We conclude with two robustness checks.

Figure 2: Data, trends and cycles



Notes: The figure plots the data (red lines) used for the estimation of the TC-VAR model 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 68-percent posterior-coverage intervals (shaded blue areas). NBER recessions are denoted by shaded grey areas.

5.1 Estimated latent trends and cycles

Panel (A) of Figure 2 plots the data (red lines) used for the estimation of the VAR with common trends over the 1960-2019 period as well as the posterior median of their latent trends (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). Panel (B) of Figure 2 plots the posterior median of the latent cycles (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area) in addition to the data (red lines).

The results confirm some stylized facts about the US economy that are commonly accepted.

First, in the 1960s and 1970s the US economy experienced an increase in trend inflation. This was possibly caused by the attempt of the Federal Reserve to counteract a break in productivity that manifested itself with an increase in the natural rate of unemployment. These two stylized facts are captured by an increase in the trend components of inflation and unemployment rate during those years. The appointment of Volcker marked a change in the conduct of monetary policy. Trend inflation declined, and so did the long-term inflation expectations. Note that even if we do not impose any restriction on the mapping from the trend component of inflation to long-term inflation expectations, the two variables largely coincide. Thus, including long-term inflation expectations helps in separating trend and cycle fluctuations.

The behavior of the cycles is reported in panel (B) of Figure 2. From this figure, a clear pattern emerges, consistent with our understanding of how the economy behaves over the business cycle. The unemployment rate increases during recessions and smoothly declines over time as the economy recovers. What is interesting is that, visually, inflation seems to behave as the New Keynesian framework would suggest: Declining during a recession, when the unemployment rate is high, and increasing during an expansion when the unemployment rate is low. This is particularly visible when focusing on inflation expectations at the one-year horizon. The cyclical component of this variable behaves very much like inflation, but it is smoother.

5.2 Inflation and unemployment over the business cycle

We now move to formally study the relation between the real economy and inflation over the business cycle. We use the estimated TC-VAR to identify the unemployment cycle shock using the method described in subsection 4.3. Specifically, the shock is identified by maximizing its contribution to the volatility of the cyclical component of the unemployment rate. As explained above, we consider two cases. In the first case, the shock is chosen to maximize the fraction of the volatility over all the frequencies of the cycle, while in the second case we exclude frequencies that imply cycles less than 1.5 years.

Table 2 reports the median contribution—and the corresponding 68-percent posterior-coverage interval—of the identified shock to the variance of the *cycle* of all the other variables over two frequency bands: a band that considers all the frequencies of the cycle—as for the construction of the shock—and the alternative band specified over all-but-short-run frequencies, thus starting from 6 quarters onward—as for the construction of the shock. By construction, the shock can explain a large share of the fluctuations of the unemployment-rate cycle. However, the shock can also explain a sizable fraction of the cyclical component of realized inflation. In the baseline scenario, the unemployment-identified shock can explain around 46% of the inflation cycle. When excluding cycles shorter than 1.5 years, the unemployment-identified shock explains about 50% of inflation variability. The result is even more striking when looking at the cyclical component of the one-year-ahead inflation expectations: approximately 58%, in the baseline scenario, or 64%, in

Table 2: Variance contributions of unemployment shock

All frequencies ($0 - \infty$ quarters)					
Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)	Inflation Exp.(10y)
72.7	81.1	58.9	46.3	57.7	56.6
[62.5,83.2]	[70.0,89.4]	[21.4,85.2]	[31.9,63.3]	[28.7,81.8]	[19.9,84.3]
All-but-short-run frequencies ($6 - \infty$ quarters)					
Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)	Inflation Exp.(10y)
73.6	81.8	58.5	51.5	64.3	65.8
[63.1,83.9]	[71.1,89.9]	[21.5,85.2]	[35.9,67.8]	[34.5,87.1]	[24.8,88.3]

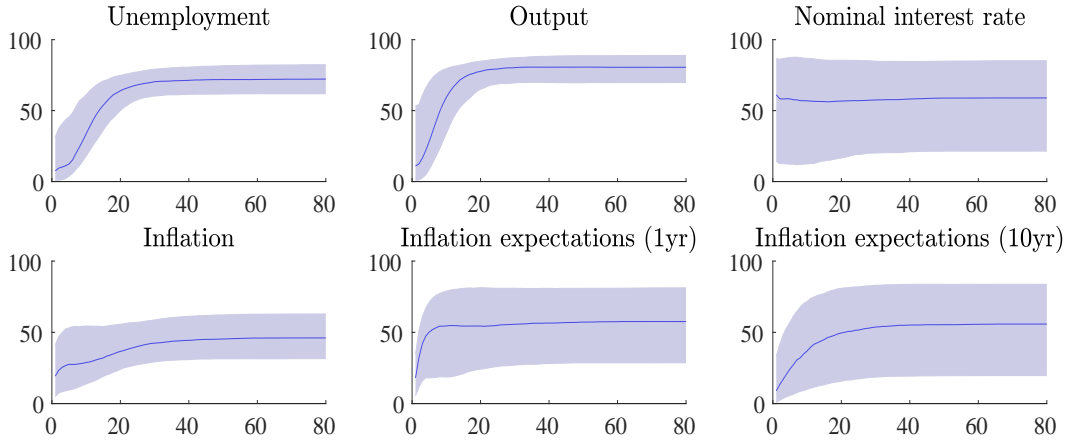
Notes: The shock is identified by maximizing its contribution to the volatility of the cyclical component of the unemployment rate. We consider two cases. In the first case, the shock is chosen to maximize the fraction of the volatility over all the frequencies of the cycle, while in the second case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

the alternative scenario, of the business-cycle variability of one-year-ahead inflation expectations is explained by the unemployment-identified shock. Given that the cycle of the one-year-ahead inflation expectations appears to be a smoother version of the cycle of realized inflation, this result corroborates the finding that inflation over the business cycle moves in a way consistent with the New Keynesian framework.

When comparing the baseline results with the alternative specification in which we remove cycles shorter than 1.5 years, we find that the contribution for inflation and inflation expectations goes visibly up when removing the short cycles. This implies that there is likely to be some residual high frequency variation in these variables that it is not related to the business cycle. The shock explains about 58% of the volatility of cyclical nominal interest rate over both frequency bands. Combined with the previous two findings, this result implies that the shock also explains the volatility of the cyclical component of the real interest rate, defined as the FFR minus expected inflation. Finally, the identified shock explains a large portion of the volatility of cyclical real GDP (in loglevels) over all frequencies and also when excluding frequencies associated with the first 6 quarters of the cycles. This finding supports the evidence of a main driver of U.S. business-cycle fluctuations as found by [Angeletos et al. \(2020\)](#).

In Figure 3, we show the contribution of the unemployment-identified shock to the forecast error variance (FEV) of the cyclical component of all the variables at different time horizons. The explanatory power of the shock reaches about 70% and 80% of the cyclical movements in unemployment and output after the first five years and about 60% of the movements in the nominal

Figure 3: Forecast error variances



Notes: The figure plots the posterior median (blue lines) and the corresponding 68-percent posterior-coverage intervals (shaded blue area).

interest rate at any horizon. Moreover, the shock explains a large portion of the movements in inflation, reaching about 40% after five years. The result for inflation and inflation expectations is consistent with the motivating evidence of Section 2 that shows that inflation cycles lag output cycles. This percentage rises above 50% after the first year for the one-year-ahead inflation expectations and after about five years for the ten-year-ahead inflation expectations. These results are in line with the shock’s contributions reported in Table 2 and clearly point to the large explanatory power of the unemployment-identified shock for the business-cycle movements in all inflation measures.

5.3 Robustness checks

In this subsection, we present results for the two alternative specifications described in Subsection 3.2. As discussed next, the results of our baseline specification are robust to both specifications.⁵ To save space, we only discuss the variance contribution of the unemployment shock on inflation.

In the first alternative specification, the one- and ten-year-ahead inflation expectations are decomposed as described in equation (3). Namely, we assume that the one- and ten-year-ahead inflation expectation surveys share a common cyclical component. Over the business-cycle frequencies, the identified shock explains nearly 50% of the cyclical component of realized inflation and nearly 60% of the cyclical component of inflation that is assumed to be shared between the two inflation-expectation measures. Over the frequencies that exclude the first 1.5 years, the shock explains 50% and just above 60% respectively. The second alternative decomposition defined in

⁵In Appendix C, we present details on the alternative specifications, describe the assumptions about initial conditions and priors used for the model estimation, and report the corresponding results.

Table 3: Variance contributions of unemployment shock: When trends are constant

Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)	Inflation Exp.(10y)
93.9	34.0	38.3	8.0	10.4	3.8
[90.5,96.6]	[27.0,40.9]	[28.7,46.8]	[3.8,13.4]	[5.6,16.4]	[1.2,8.4]

Notes: The shock is identified by maximizing its contribution to the volatility of the cyclical component of the unemployment rate over business-cycle frequencies (6-32 quarters). 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 the other variables over the same frequencies.

(4) is based on the assumption that realized inflation and one-year-ahead inflation expectations share a common cycle, while the measure of ten-year-ahead inflation expectations does not. We find that the shock explains nearly 40% of the volatility of the common cyclical component of inflation. When excluding movements in cyclical components at frequencies higher than 1.5 years, the results point to an explanatory power of nearly 45% of the identified shock for the common cyclical component of inflation.

We conclude that our results are robust to additional restrictions on the trend-cycle decomposition. We favor our baseline specification because more flexible.

6 Importance of Time-varying Trends

Our previous results highlight the importance of properly capturing low-frequency movements in real and nominal variables. By treating the latent trends as time-varying objects, our model is able to assess the relationship between those variables at business-cycle frequencies. To further support our claim, we show that, by considering a constrained specification of our model in which the trends are constant, the model reproduces the findings of [Angeletos et al. \(2020\)](#) who argue for a disconnect between real activity and inflation over the business cycle.

Specifically, we leave unchanged all other assumptions about priors and initial conditions and simply set the standard deviation of shocks to the trends to zero. Table 3 reports the contributions of the identified shock to the volatility of the cyclical components of the unemployment rate as well as the remaining variables over the same frequencies.

In this specification which does not allow for low-frequency movements in the variables of interest, the shock only explains 8% of the volatility of the realized inflation cycle and about 10% and 3% of the volatility of the one- and ten-year-ahead inflation expectations cycle. The estimate of the shock’s contribution to the volatility of inflation is well within the range of estimates provided in [Angeletos et al. \(2020\)](#). The similarity of our findings under this scenario with those of [Angeletos et al. \(2020\)](#) is not surprising. Once the latent trends are assumed to be constant, our

model collapses to a VAR model, with the only distinction that output is modelled in deviations from a linear trend and the other variables are demeaned.

7 Revisiting VAR models

Given the empirical evidence provided in the two previous sections, our goal is now to demonstrate that the adoption of a standard VAR, as opposed to a TC-VAR, induces very different conclusions about the disconnect between real activity and inflation over the business cycle. Results are more similar when considering subsamples over which trends are less relevant, but we are never able to recover the same results obtained with the TC-VAR. We proceed in two ways. In Subsection 7.1, we consider the data and specification proposed by [Angeletos et al. \(2020\)](#) and explore alternative sample periods and assumptions about the priors. The evidence strongly points to a disconnect between real activity and inflation over the business cycle. In Subsection 7.2, we bring our attention back to the data used in Section 5. Even when using the same data, the adoption of a VAR model leads to the conclusion of a disconnect between inflation and real activity over the business cycle. This is true even after exploring alternative assumptions about the priors.

7.1 Using a VAR to assess the link between inflation and real activity

Alternative priors. In this subsection, we revisit the results obtained by [Angeletos et al. \(2020\)](#) by allowing for different priors and data periods. We start by extending the estimation sample of [Angeletos et al. \(2020\)](#) by two years so that the periods match those of our baseline analysis, i.e., 1955:Q1 to 2019:Q4. As detailed in Section I of their paper, the data consist of quarterly observations on the following macroeconomic variables: the unemployment rate (u_t); the real, per-capita levels of GDP (Y_t), investment (I_t), consumption (C_t); hours worked per person (h_t); the level of utilization-adjusted total factor productivity (TFP_t); the labor share ($w_t h_t / Y_t$); the inflation rate (π_t), as measured by the rate of change in the GDP deflator; and the nominal interest rate (R_t), as measured by the federal funds rate.⁶

As in their baseline, we proceed to estimate a VAR model with two lags using Bayesian methods. However, we consider two sets of priors. In one case, we follow their approach and use a Minnesota prior (“Minnesota”). In the other case, we combine a long-run prior *à la* [Giannone et al. \(2019a\)](#) with the Minnesota prior (“Minnesota+Long-run”). The cointegrating relationships that we assume among the variables in the long run are in line with those of [Giannone et al. \(2019a\)](#) and are described in Appendix D.1. We allow for separate shrinkage for each active row of the matrix that captures the cointegrating relationship among the variables. In the estimation, we consider the same number of draws and pre-sample years as in our baseline analysis. Hyperparameters for

⁶We drop labor productivity in the VAR model because it can be measured by the ratio between output and hours worked.

Table 4: Variance contributions of unemployment shock: 1955-2019

	Unemployment	Output	Investment	Consumption	Hours
Minnesota	91.9 [88.3, 94.5]	71.4 [66.9, 77.5]	74.4 [69.4, 79.5]	47.3 [37.9, 54.4]	75.1 [70.7, 78.8]
Minnesota + Long-run	92.4 [88.7, 94.7]	72.8 [68.6, 76.1]	74.6 [70.7, 77.6]	49.7 [44.2, 58.4]	75.9 [73.1, 79.4]
	TFP	Labor share	Inflation	FFR	
Minnesota	19.2 [11.3, 28.1]	15.1 [9.8, 22.7]	7.6 [3.9, 17.9]	36.3 [26.5, 45.4]	
Minnesota + Long-run	18.5 [10.2, 24.4]	14.9 [11.0, 21.6]	11.5 [5.9, 17.8]	38.2 [32.2, 45.9]	

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies.

priors are optimized following [Giannone et al. \(2019a\)](#). For each estimation, we then identify the shock targeting the unemployment rate at business-cycle frequencies using the method described in section 4.3. Table 4 reports the contribution of that shock for each of the variables of the VAR model over the same frequencies.

The results in Table 4 indicate that, even when combining the Minnesota and long-run priors, we still recover the evidence of a disconnect between real and nominal variables. While the identified shock explains nearly 8% of the volatility of inflation over the business-cycle when only using Minnesota priors, its contribution slightly increases to about 11% when those priors are combined with long-run priors *à la* [Giannone et al. \(2019a\)](#). These results are in line with those obtained from a vast variety of specifications that [Angeletos et al. \(2020\)](#) consider including vector error correction models.

Alternative samples. We now consider an estimation sample that is least impacted by time-varying trends. Specifically, our sample now starts in 1984:Q1 and ends in 2008:Q4. We estimate the VAR model with two lags and use the two alternative assumptions about the priors: only a Minnesota prior or combining that prior with the long-run priors. Table 5 reports the contribution of the shock identified targeting the unemployment rate for the volatility of the other variables at business-cycle frequencies. As expected, the results in Table 5 show that restricting the sample to the period of the Great Moderation improves the contribution of the identified shock on inflation. However, such contribution is overall small (about 16%) regardless of priors.

Dogmatic priors. As a final check, we explore if these results can be changed if priors are not

Table 5: Variance contributions of unemployment shock: 1984-2008

	Unemployment	Output	Investment	Consumption	Hours
Minnesota	92.8 [85.5, 95.6]	54.3 [43.2, 67.7]	55.3 [44.8, 67.9]	34.9 [21.0, 49.5]	69.4 [61.3, 79.5]
Minnesota + Long-run	92.8 [87.7, 95.3]	52.7 [42.2, 66.6]	55.9 [45.8, 66.1]	33.1 [21.3, 53.4]	69.6 [61.7, 78.1]
	TFP	Labor share	Inflation	FFR	
Minnesota	9.4 [3.5, 22.1]	11.6 [3.7, 21.3]	15.8 [7.3, 39.3]	63.2 [56.9, 73.9]	
Minnesota + Long-run	9.6 [3.4, 21.5]	9.9 [3.1, 17.5]	16.3 [7.7, 29.1]	63.5 [54.9, 72.3]	

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies.

optimized and instead are imposed dogmatically. The objective is to understand if one can still use VAR models with the help of priors. We consider several different combinations of hyperparameters such that the two Minnesota and long-run priors can be either not imposed, optimized, or set dogmatically. For this exercise, it is important to understand that we are imposing same degree of shrinkage for each active row of the matrix that captures the cointegrating relationship among the variables.

Table 6 reports the contribution of unemployment-identified shock on inflation at business-cycle frequencies. The results are straightforward to understand. When only dogmatic Minnesota priors are imposed (see lower left corner of Table 6), the inflation variance contribution of the unemployment shock is nearly zero. It is not exactly zero because the shocks are correlated even when the individual series follows a random-walk process. However, the opposite case of dogmatic long-run priors (see upper right corner of Table 6) leads to about 33%.⁷ While the numbers are certainly lower than those reported in Table 2, our findings strongly suggest that it is important to consider long-run relationship amongst variables and not properly accounting those features in the data could lead to imprecise conclusions about how real activity and inflation are connected over the business cycle.

⁷The fact that the number 26.2% in the middle of Table 6 does not coincide with 16.3% in Table 5 is because we are imposing same degree of shrinkage for each active row of the matrix that captures the cointegrating relationship among the variables in Table 6 but not in Table 5 which allows for separate shrinkage.

Table 6: Inflation variance contribution of unemployment shock: 1984-2008

	No long-run	Optimized long-run	Dogmatic long-run
No Minnesota	25.2 [9.2, 42.6]	26.2 [13.3, 52.4]	33.1 [11.0, 49.0]
Optimized Minnesota	20.1 [11.0, 31.5]	26.2 [13.3, 52.4]	33.1 [11.0, 49.0]
Dogmatic Minnesota	2.0 [0.4, 5.6]	26.2 [13.3, 52.4]	–

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies. For long-run priors, we are imposing same degree of shrinkage for each active row of the matrix that captures the cointegrating relationship among the variables.

7.2 Estimating our data with a VAR model based on alternative priors

As a final exercise, we bring our attention back to the data used in Section 5 and described in Section 4.1. We focus on the period of the Great Moderation between 1984 and 2008 and examine the extent to which the estimated VAR model can reproduce empirical evidence provided by our baseline TC-VAR model.⁸ We estimate the model using only a Minnesota prior (“Minnesota”) or combining it with long-run priors (“Minnesota+Long-run”). The cointegrating relationships that we impose across variables are reported in Appendix D.2. For both specifications, we identify the shock targeting the unemployment rate at business-cycle frequencies and report in Table 7 its contribution for other variables. In sum, while the long-run priors are better at delineating long-run from short-run dynamics, the evidence suggests one needs to go beyond the VAR dynamics to examine the (dis)connect between real activity and inflation over the business cycle.

8 Explaining VARs Through the Lens of Trend-cycle-VARs

In this section, we provide theoretical arguments that motivate the adoption of a TC-VAR model, rather than a standard VAR model, for the results presented in Section 5. In Subsection 8.1, we show that a fixed-coefficient VAR model estimated over a period of time that presents structural changes is misspecified. The misspecification problem associated with the use of a VAR model to describe a data generating process characterized by both low- and high-frequency move-

⁸Relative to the baseline specification, we drop the ten-year-ahead inflation expectations as it is not straightforward to impose long-run relationships across two inflation expectations series.

Table 7: Variance contributions of unemployment shock: 1984-2008

	Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)
Minnesota	88.3 [78.2, 92.1]	61.5 [48.0, 74.0]	70.7 [61.7, 78.9]	13.4 [6.0, 27.3]	28.2 [13.5, 42.2]
Minnesota + Long-run	88.4 [82.3, 93.2]	65.7 [53.5, 73.2]	74.5 [60.1, 80.5]	20.1 [7.0, 29.1]	28.6 [12.6, 47.1]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies.

ments cannot be easily resolved.⁹ Even if an econometrician could correctly reconstruct the VAR representation of the TC-VAR model, the parameter estimates of the misspecified model would confound low-frequency movements associated with the trend with those at business-cycle frequencies related to the cycle. Moreover, the reduced-form innovations that she would recover would not only map into the innovations affecting the latent persistent and stationary components. By contrast, the reduced-form innovations would also capture the error associated with the estimates of the latent components. Of course, in reality, these issues would be exacerbated by the fact that the VAR parameters estimated over a finite sample would be distorted because a single set of parameters would need to account for both trend and cycle fluctuations.

To build the intuition about the underlying issues, in Subsection 8.2, we generate data from a Monte Carlo simulation of a bivariate TC-VAR model and show that a VAR model does not succeed in capturing the assumed cyclical relationship between the two variables, even with a long data sample. Moreover, in Appendix E.3, we consider a univariate trend-cycle autoregressive (TC-AR) model of U.S. inflation based on [Stock and Watson \(2007\)](#) and derive analytically its infinite-order autoregressive (AR) representation.

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

In Subsection 3.3, we provide the state-space representation of our model in (5) and (6) and that we report below in (9)

$$y_t = \Lambda x_t, \tag{9.1}$$

$$x_t = \Phi x_{t-1} + \mathcal{R}\varepsilon_t, \tag{9.2}$$

⁹[Watson \(1986\)](#) discusses the equivalence between an unobserved component model and its autoregressive, integrated, moving average (ARIMA) representation, thus pointing to the misspecification problem characterizing a VAR representation of a TC-VAR model.

where $\varepsilon_t = Qw_t$ and $E(w_t w_t') = \mathcal{I}$. Our goal is to map the state-space representation in (9) into a VAR model. Importantly, for all the considered specifications, the overall number of shocks of the TC-VAR model is strictly larger than the number of observables. Equivalently, $\dim(w_t) = (n_\tau + n_c + n_\eta) > n = \dim(y_t)$. As a result, the ‘*poor man’s invertibility condition*’ proposed in [Fernandez-Villaverde et al. \(2007\)](#) cannot be tested because it requires the number of shocks and observables to coincide. We therefore seek to find the ‘*innovations representation*’ of (9).

Because the innovations representation results from the application of the Kalman filter to the state-space representation, we first ensure the suitability of the filter for our purpose and more specifically its asymptotic stability and convergence. Clearly, these properties of the filter depend on the properties of (9) and should not be taken for granted in our setup: In the transition equation (9.2), the cyclical components are assumed to be stationary, while trends follow unit-root processes. However, we follow [Anderson and Moore \(1979\)](#) who suggest to verify two conditions: i) the pair $(\Phi, \mathcal{R}Q)$ is stabilizable; ii) the pair (Φ, Λ') is detectable—or equivalently, the pair (Φ', Λ') is stabilizable.¹⁰ For all the specifications of the TC-VAR model that we consider in Section 3, both conditions are satisfied for each draw obtained from the model estimation.

Having verified the suitability of the Kalman filter for our purpose, we follow [Fernandez-Villaverde et al. \(2007\)](#) to derive the innovations representation. Specifically, we write the representation in (9) as

$$x_{t+1} = Ax_t + Bw_{t+1}, \quad (10.1)$$

$$y_{t+1} = Cx_t + Dw_{t+1}, \quad (10.2)$$

where $A = \Phi$, $B = \mathcal{R}Q$, $C = \Lambda A$, $D = \Lambda B$ and $E(w_t w_t') = \mathcal{I}$. Defining the linear projection of x_t on $y^t \equiv \{y_j\}_{j=1}^t$ as $\hat{x}_t \equiv E(x_t | y^t)$, the one-step-ahead error associated with the forecast of y_{t+1} as $\hat{D}_{t+1}\nu_{t+1}$ and the term updating the estimator of the state for the next period \hat{x}_{t+1} as $\hat{B}_{t+1}\nu_{t+1}$, the application of the Kalman filter to (10) delivers the innovations representation

$$\hat{x}_{t+1} = A\hat{x}_t + \hat{B}_{t+1}\nu_{t+1}, \quad (11.1)$$

$$y_{t+1} = C\hat{x}_t + \hat{D}_{t+1}\nu_{t+1}, \quad (11.2)$$

where $x_0 \sim (\hat{x}_0, \Sigma_0)$, the covariance matrix Σ_0 is positive semi-definite, and ν_t is a vector of mean-zero, normal and uncorrelated white-noise innovations such that $E(\nu_t \nu_t') = \mathcal{I}$. Notably, under this representation, the number of shocks and observables coincide. Also, because the innovation ν_t is fundamental for y_t by definition, it is thus uncorrelated with y_{t-s} and ultimately ν_{t-s} for any

¹⁰We provide the definition of stabilizability in Appendix E.1. For further details, refer to [Burrige and Wallis \(1983\)](#), [Sargent \(1988\)](#), [Anderson et al. \(1996\)](#) and [Hansen and Sargent \(2008, 2013\)](#) among others. We refer the reader to pp. 76-82 and Appendix C in [Anderson and Moore \(1979\)](#) for the details on the conditions mentioned here and the definition of stabilizability.

$s \geq 0$.

Importantly, the innovations representation in (11) shows that, with a finite sample $\{y_t\}_{t=1}^T$ where $T < \infty$, it is *not* possible to derive a VAR representation because the matrices \hat{B}_t and \hat{D}_t depend on time t . As a result, we consider the limit case for T approaching infinity. Because the asymptotic stability and convergence of the Kalman filter hold, the matrices \hat{B}_t and \hat{D}_t also converge to their time-invariant counterparts \hat{B} and \hat{D} .¹¹ Therefore, we can derive the infinite-order VAR representation. Solving (11.2) for ν_{t+1} and combining with (11.1), we obtain

$$\left[\mathcal{I} - \left(A - \hat{B}\hat{D}^{-1}C \right) L \right] \hat{x}_{t+1} = \hat{B}\hat{D}^{-1}y_{t+1} \quad (12)$$

where L denotes the lag operator. Importantly, the asymptotic properties of the Kalman filter also guarantee that all the eigenvalues of the matrix $\left(A - \hat{B}\hat{D}^{-1}C \right)$ are all strictly inside the unit circle in modulus (Anderson and Moore, 1979). Therefore, we solve equation (12) for \hat{x}_{t+1} and plug the solution in the time-invariant version of equation (11.2) to obtain the following VAR(∞) representation

$$\begin{aligned} y_{t+1} &= C \left[\mathcal{I} - \left(A - \hat{B}\hat{D}^{-1}C \right) L \right]^{-1} \hat{B}\hat{D}^{-1}y_t + \hat{D}\nu_{t+1} \\ &= \sum_{s=0}^{\infty} C \left(A - \hat{B}\hat{D}^{-1}C \right)^s \hat{B}\hat{D}^{-1}y_{t-s} + \hat{D}\nu_{t+1}, \end{aligned} \quad (13)$$

where the inverted matrix in square brackets in the first equation is a square summable polynomial in L . Equation (13) shows three important conclusions. First, the state-space representation of the TC-VAR model in (9) maps into the infinite-order VAR representation in (13) under the assumption that infinite data are available. As a result, a finite-order VAR with finite data cannot capture the dynamics described by the decomposition of the observables y_t into trends and cycles. Second, even if infinite data were available, equation (13) clarifies that estimates of the autoregressive parameters associated with the VAR(∞) representation confound movements of y_{t+1} that are driven by both the trend and cycle. Equivalently, the VAR(∞) representation cannot disentangle movements of y_{t+1} at low frequencies from those at cyclical frequencies. Finally, as shown in Fernandez-Villaverde et al. (2007), the innovations associated with the VAR(∞) representation, $\hat{D}\nu_{t+1}$, capture not only the shocks to the latent trends and cycles, Dw_{t+1} , but also the error associated with the estimate of those latent components, $C(x_t - \hat{x}_t)$.

Finally, in Appendix E.3, we provide a simple analytical example based on an unobserved components model of U.S. inflation by Stock and Watson (2007) to show the intuition for these results.

With these results, we are not arguing for the unconditional superiority of a TC-VAR over a

¹¹Appendix E.2 provides the details on the equations for matrices \hat{B} and \hat{D} .

VAR. Over the past four decades, economists have used VARs as extremely flexible econometric models capable of uncovering enlightening empirical results. Our point is that for the specific question of assess the strength of the relation between inflation and real activity over the business cycle, a TC-VAR appears to be a more effective tool.

8.2 A Monte Carlo simulation of a bivariate TC-VAR model

To provide an intuitive example, let us assume that the data generating process for the unemployment rate and inflation rate, $y_t = \{u_t, \pi_t\}'$, is described by the measurement equation $y_t = \tau_t + c_t$, where the dynamics of trend τ_t and cyclical component c_t can be modeled as in (2) and that we report below in (14)

$$\tau_t = \tau_{t-1} + \varepsilon_{\tau,t}, \quad (14.1)$$

$$c_t = \Phi_1 c_{t-1} + \varepsilon_{c,t}, \quad (14.2)$$

where $\tau_t = \{\tau_{u,t}, \tau_{\pi,t}\}'$, $c_t = \{c_{u,t}, c_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,u,t}, \varepsilon_{\tau,\pi,t}\}'$ and $\varepsilon_{c,t} = \{\varepsilon_{c,u,t}, \varepsilon_{c,\pi,t}\}'$. In this example, we assume that (14.2) is

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

implying that, while the cyclical component of the unemployment rate only depends on its lag, the cyclical component of inflation depends on its own lag as well as on the lagged cyclical component of the unemployment rate. We also assume that the shocks are independent and identically distributed as

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

To generate a Monte Carlo simulation, we consider a parameterization that would help the misspecified VAR model to capture a relationship between the two series over the business cycle. For the unemployment rate, we set ρ_{uu} to 0.9, thus assuming a persistent process for the cyclical component of the unemployment rate. Additionally, we standardize the gaussian innovations to the trend and cycle of the unemployment rate, that is $\sigma_{\tau,u} = \sigma_{c,u} = 1$. For the inflation rate, we assume that its process is fully explained by the (lagged) cyclical component of the unemployment rate. This assumption is evidently unrealistic but is chosen to ensure that the innovations to the *cyclical component of the unemployment rate* are the only source of fluctuations for *realized inflation*. Implementing this assumption requires setting the innovations to both the trend and cyclical components of inflation to zero, that is $\sigma_{\tau,\pi} = \sigma_{c,\pi} = 0$. Additionally, we set to zero

Table 8: Variance contribution of unemployment shock

Unemployment	Inflation
96.5	9.4
[94.9, 97.8]	[7.4, 11.7]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies.

the autoregressive parameter $\rho_{\pi\pi}$. These assumptions imply that the simulated inflation rate is described by the process¹²

$$\pi_t = -c_{u,t-1}. \quad (15)$$

Using this calibration, we produce a long Monte Carlo simulation and then estimate a VAR model choosing the lags to deliver the lowest Bayesian Information Criterion (BIC) and using a Minnesota prior as in [Angeletos et al. \(2020\)](#).¹³ We then identify the shock targeting the simulated unemployment rate at business-cycle frequencies.

Table 8 reports the contribution of the identified shock for both series over business-cycle frequencies. Even with a long data sample and in presence of strong assumptions about the cyclical relationship between the unemployment rate and inflation, the results in the table lead to conclude that the shock driving the unemployment rate does not explain the inflation rate. Intuitively, the identification of the shock at business-cycle frequencies does not succeed in extracting the cyclical relationship because of the misspecification of the standard VAR model that fails to separate cycle and trend innovations in unemployment. It goes without saying, that if we were to complicate the model further, for example by introducing a trend in inflation, the results would get even worse.

9 Conclusions

In recent years, a series of papers have called into question the validity of the New Keynesian framework. One important argument is the observation that inflation seems to be unresponsive to business-cycle movements in real activity. In this paper, we used a trend-cycle VAR model to study the relation between inflation and the real economy over the business cycle. The trend-cycle VAR model has the virtue to remove low-frequency movements in inflation and real activity that can contaminate inference about the VAR parameters and innovations. We show that at business-

¹²In Appendix E.4, we also consider an alternative calibration by setting $\rho_{\pi\pi} = 0.5$ and $\sigma_{\tau,\pi} = 1$, therefore allowing the inflation rate to also depend on its lagged cyclical component and the innovations to its trend component. The results show that the contribution of the unemployment-rate shock to inflation at business-cycle frequencies is about 70% lower than the corresponding contribution reported in Table 8.

¹³We generate a Monte Carlo simulation of 50,000 observations and keep the last 1,000 as simulated data. To estimate the VAR, we set 50,000 draws and keep the last 1,000 to identify the shock as in [Angeletos et al. \(2020\)](#).

cycle frequencies, fluctuations of inflation are tightly related to movements in real activity, in line with what implied by the New Keynesian framework.

References

- Anderson, B. D. O. and Moore, J. B. (1979). *Optimal Filtering*. Prentice-Hall, Inc., New Jersey.
- Anderson, E. W., Hansen, L. P., McGrattan, E. R., , and Sargent, T. J. (1996). Mechanics of Forming and Estimating Dynamic Linear Economies. In Amman, H. M., Kendrick, D. A., and Rust, J., editors, *Handbook of Computational economics*, volume 1, pages 171–252. North-Holland.
- 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. (2021). The Inflation Rate Disconnect Puzzle: On the International Component of Trend Inflation and the Flattening of the Phillips Curve. Norges Bank Research, Working Paper no. 17.
- Basu, S., Candian, G., Chahrour, R., and Valchev, R. (2021). Risky Business Cycles. NBER WP 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*, 28:69–130.
- Burridge, P. and Wallis, K. F. (1983). Signal Extraction in Nonstationary Series. University of Warwick working paper no. 234.
- Carter, C. K. and Kohn, R. (1994). On Gibbs Sampling for State Space Models. *Biometrika*, 81(3):541–553.
- Christiano, L. and Fitzgerald, T. (2003). The Band Pass Filter. *International Economic Review*, 44(2):435–65.
- Del Negro, M., Giannone, D., Giannoni, M., and Tambalotti, A. (2017). Safety, Liquidity, and the Natural Rate of Interest. *Brookings Papers on Economic Activity*, 48:235–94.
- Del Negro, M., Giannone, D., Giannoni, M., and Tambalotti, A. (2019). Global trends in interest rates. *Journal of International Economics*, 118:248–62.
- Del Negro, M., Lenza, M., Primiceri, G. E., and Tambalotti, A. (2020). What’s up with the Phillips Curve? NBER Working Papers 27003, National Bureau of Economic Research, Inc.
- 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.

- Durbin, J. and Koopman, S. J. (2001). *Time Series Analysis by State Space Methods*. Oxford University Press, second edition.
- 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*, pages 1– 22.
- Farmer, R. E. A. and Platonov, K. (2019). Animal Spirits in a Monetary Model. *European Economic Review*, 115:60–77.
- Fernandez-Villaverde, J., Rubio-Ramirez, J., Sargent, T. J., and Watson, M. W. (2007). ABCs (and Ds) of Understanding VARs. *American Economic Review*, 97(3):1021–1026.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics*, 97(2):436–51.
- Giannone, D., Lenza, M., and Primiceri, G. E. (2019a). Priors for the Long Run. *Journal of the American Statistical Association*, 114(526):565–80.
- 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):137–73.
- Giannone, D., Reichlin, L., and Sala, L. (2006). VARs, common factors and the empirical validation of equilibrium business cycle models. *Journal of Econometrics*, 132:257–79.
- 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.
- Hansen, L. P. and Sargent, T. J. (2008). *Robustness*. Princeton University Press, Princeton, NJ.
- Hansen, L. P. and Sargent, T. J. (2013). *Recursive Models of Dynamic Linear Economies*. Princeton University Press, Princeton, NJ.
- Holston, K., Laubach, T., and Williams, J. C. (2017). Measuring the Natural Rate of Interest: International Trends and Determinants. *Journal of International Economics*, 108:59–75.
- 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):1005–46.
- Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2013). Is There a Trade-Off between Inflation and Output Stabilization? *American Economic Journal: Macroeconomics*, 5(2):1–31.
- 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*, 85:1063–70.

- 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*, 34:425–36.
- 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.
- Primiceri, G. E. (2013). Comment of “Understanding Noninflationary Demand-Driven Business Cycles,” by Beaudry and Portier. *NBER Macroeconomics Annual*, 28:131–143.
- 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.
- Sargent, T. J. (1988). Linear Optimal Control, Filtering, and Rational Expectations. Federal Reserve Bank of Minneapolis, Working Paper no. 224.
- 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(1):3–33.
- 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.
- Uhlig, H. (2003). What moves real GNP? Unpublished.
- Villani, M. (2009). Steady-state Priors for Vector Autoregressions. *Journal of Applied Econometrics*, 24:630–50.
- Watson, M. (1986). Univariate Detrending Methods with Stochastic Trends. *Journal of Monetary Economics*, 18:49–75.
- Watson, M. (2004). Comment on “Monetary Policy in Real Time,” by Domenico Giannone, Lucrezia Reichlin, and Luca Sala. *NBER Macroeconomics Annual*, 19:216–21.

A The State-space Representation

In Subsection 3.3, we discussed the state-state representation in (5) and (6) that we report below in (16)

$$y_t = \Lambda_\tau x_{\tau,t} + \Lambda_c x_{c,t} + \Lambda_\eta x_{\eta,t}, \quad (16.1)$$

$$x_t = \Phi x_{t-1} + \mathcal{R} \varepsilon_t, \quad (16.2)$$

where $\Lambda = [\Lambda_\tau, \Lambda_c, \Lambda_\eta]$, $\Lambda_c = [\Lambda_{c,0}, \dots, \Lambda_{c,p-1}]$ and $\Lambda_\eta = [\Lambda_{\eta,0}, \Lambda_{\eta,1}]$. For our baseline specification in 3.1, these matrices are constructed as

$$\Phi = \begin{bmatrix} \mathbf{I}_{4 \times 4} & \mathbf{0}_{4 \times 6} & \mathbf{0}_{4 \times 6} & \mathbf{0}_{4 \times 1} & \mathbf{0}_{4 \times 1} \\ \mathbf{0}_{6 \times 4} & \Phi_1 & \Phi_2 & \mathbf{0}_{6 \times 1} & \mathbf{0}_{6 \times 1} \\ \mathbf{0}_{6 \times 4} & \mathbf{I}_{6 \times 6} & \mathbf{0}_{6 \times 6} & \mathbf{0}_{6 \times 1} & \mathbf{0}_{6 \times 1} \\ \mathbf{0}_{1 \times 4} & \mathbf{0}_{1 \times 6} & \mathbf{0}_{1 \times 6} & 0 & 0 \\ \mathbf{0}_{1 \times 4} & \mathbf{0}_{1 \times 6} & \mathbf{0}_{1 \times 6} & 1 & 0 \end{bmatrix}, \mathcal{R} = \begin{bmatrix} \mathbf{I}_{4 \times 4} & \mathbf{0}_{4 \times 6} & \mathbf{0}_{4 \times 1} \\ \mathbf{0}_{6 \times 4} & \mathbf{I}_{6 \times 6} & \mathbf{0}_{6 \times 1} \\ \mathbf{0}_{6 \times 4} & \mathbf{0}_{6 \times 6} & \mathbf{0}_{6 \times 1} \\ \mathbf{0}_{1 \times 4} & \mathbf{0}_{1 \times 6} & 1 \\ \mathbf{0}_{1 \times 4} & \mathbf{0}_{1 \times 6} & 0 \end{bmatrix} \quad (17)$$

$$\varepsilon_t = \begin{bmatrix} \varepsilon_{\tau,t} \\ \varepsilon_{c,t} \\ \varepsilon_{\eta,t} \end{bmatrix}, \Sigma = \begin{bmatrix} \Sigma_\tau & \mathbf{0}_{4 \times 6} & \mathbf{0}_{4 \times 1} \\ \mathbf{0}_{6 \times 4} & \Sigma_c & \mathbf{0}_{6 \times 1} \\ \mathbf{0}_{1 \times 4} & \mathbf{0}_{1 \times 6} & \Sigma_\eta \end{bmatrix}, \Lambda = \begin{bmatrix} \Lambda_{\tau,0} & \Lambda_{c,0} & \Lambda_{c,1} & \Lambda_{\eta,0} & \Lambda_{\eta,1} \end{bmatrix}, \Lambda_{c,0} = \mathbf{I}_{6 \times 6},$$

$$\Lambda_{\tau,0} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \Lambda_{c,1} = \begin{bmatrix} -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \Lambda_{\eta,0} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \Lambda_{\eta,1} = -\Lambda_{\eta,0}.$$

B Estimation Details

B.1 Settings

For all the considered specifications, we assume that the cyclical components evolve according to a VAR model with two lags ($p = 2$) following the baseline approach of Angeletos et al. (2020). For each estimation, we adopt the Gibbs sampler described in Appendix B.2. We use 50,000 draws to estimate the TC-VAR model. We then discard the first 25,000 draws and keep one in every 25 draws, thus leaving 1,000 draws, to use for the identification of the shocks.

B.2 Gibbs Sampler

We assume that Σ_η is known. We collect parameters that need to be estimated in

$$\Theta_\tau = \{\Sigma_\tau\}, \quad \Theta_c = \{\Phi_1, \Phi_2, \Sigma_c\}.$$

We use the Gibbs sampler to estimate the model unknowns. We rely on the state-space representation of (5) and (6). For the j th iteration,

- Run Kalman smoother to generate $\tau_{1:T}^j$ and $c_{0:T}^j$ conditional on $\Theta_\tau^j, \Theta_c^j$: This is explained in Section B.2.1.
- Obtain posterior estimates of $\Theta_\tau^{j+1}, \Theta_c^{j+1}$ from the MNIW conditional on $\tau_{1:T}^j$ and $c_{0:T}^j$: This is explained in Section B.2.2.

B.2.1 Kalman smoother

We rely on the state-space representation in equations (5) and (6). Conditional on the j th draw of $\Theta_\tau^j, \Theta_c^j$, we apply the standard Kalman filter as described in Durbin and Koopman (2001). Suppose that the distribution of

$$s_{t-1} | \{y_{1:t-1}, \Theta_\tau^j, \Theta_c^j\} \sim N(s_{t-1|t-1}, P_{t-1|t-1}).$$

Then, the Kalman filter forecasting and updating equations take the form

$$\begin{aligned} s_{t|t-1} &= \Gamma s_{t-1|t-1} \\ P_{t|t-1} &= \Gamma P_{t-1|t-1} \Gamma' + \Omega \Sigma \Omega' \\ s_{t|t} &= s_{t|t-1} + (\Lambda P_{t|t-1})' (\Lambda P_{t|t-1} \Lambda')^{-1} (y_t - \Lambda s_{t|t-1}) \\ P_{t|t} &= P_{t|t-1} - (\Lambda P_{t|t-1})' (\Lambda P_{t|t-1} \Lambda')^{-1} (\Lambda P_{t|t-1}). \end{aligned}$$

In turn,

$$s_t | \{y_{1:t}, \Theta_\tau^j, \Theta_c^j\} \sim N(s_{t|t}, P_{t|t}).$$

Next, the backward smoothing algorithm developed by Carter and Kohn (1994) is applied to recursively generate draws from the distributions $s_t | (S_{t+1:T}, Y_{1:T}, \Theta_\tau^j, \Theta_c^j)$ for $t = T-1, T-2, \dots, 1$. The last element of the Kalman filter recursion provides the initialization for the simulation smoother:

$$\begin{aligned} s_{t|t+1} &= s_{t|t} + P_{t|t} \Gamma' P_{t+1|t}^{-1} (s_{t+1} - \Gamma s_{t|t}) \\ P_{t|t+1} &= P_{t|t} - P_{t|t} \Gamma' P_{t+1|t}^{-1} \Gamma P_{t|t} \\ s_t^j &\sim N(s_{t|t+1}, P_{t|t+1}), \quad t = T-1, T-2, \dots, 1. \end{aligned} \tag{18}$$

In sum, we obtain smoothed estimates of $\tau_{1:T}^j$ and $c_{0:T}^j$.

B.2.2 Posterior draw

We treat the smoothed estimates of $\tau_{1:T}^j$ and $c_{0:T}^j$ as data points. The objective is to draw Θ_τ^{j+1} and Θ_c^{j+1} .

VAR coefficients. For ease of exposition, we omit the superscript j below. For $t \in \{2, \dots, T\}$, we express the VAR as

$$c'_t = \underbrace{\begin{bmatrix} c'_{t-1} & c'_{t-2} \end{bmatrix}}_{w'_t} \underbrace{\begin{bmatrix} \Phi'_1 \\ \Phi'_2 \end{bmatrix}}_{\beta} + \epsilon'_{c,t}, \quad \epsilon_{c,t} \sim N(\mathbf{0}, \Sigma_c). \quad (19)$$

Define $X = [c_2, \dots, c_T]'$, $W = [w_2, \dots, w_T]'$, and $\epsilon_c = [\epsilon_{c,2}, \dots, \epsilon_{c,T}]'$ conditional on the initial observations. If the prior distributions for β and Σ_c are

$$\beta | \Sigma \sim MN(\underline{\beta}, \Sigma \otimes (\underline{V}_\beta \xi)), \quad \Sigma_c \sim IW(\underline{\Psi}, \underline{d}), \quad (20)$$

then because of the conjugacy the posterior distributions can be expressed as

$$\beta | \Sigma \sim MN(\bar{\beta}, \Sigma \otimes \bar{V}_\beta), \quad \Sigma_c \sim IW(\bar{\Psi}, \bar{d}) \quad (21)$$

where

$$\begin{aligned} \bar{\beta} &= (W'W + (\underline{V}_\beta \xi)^{-1})^{-1} (W'X + (\underline{V}_\beta \xi)^{-1} \underline{\beta}), \\ \bar{V}_\beta &= (W'W + (\underline{V}_\beta \xi)^{-1})^{-1}, \\ \bar{\Psi} &= (X - W\bar{\beta})'(X - W\bar{\beta}) + (\bar{\beta} - \underline{\beta})'(\underline{V}_\beta \xi)^{-1}(\bar{\beta} - \underline{\beta}) + \underline{\Psi}, \\ \bar{d} &= T - 2 + \underline{d}. \end{aligned} \quad (22)$$

We follow the exposition in [Giannone et al. \(2015\)](#) in which ξ is a scalar parameter controlling the tightness of the prior information in (20). For instance, prior becomes more informative when $\xi \rightarrow 0$. In contrast, when $\xi = \infty$, then it is easy to see that $\bar{\beta} = \hat{\beta}$, i.e., an OLS estimate.

In sum, we draw β^{j+1} and Σ^{j+1} from (21). Hence, we obtain $\Theta_c^{j+1} = \{\Phi_1^{j+1}, \Phi_2^{j+1}, \Sigma_c^{j+1}\}$.

Trend component variances. Conditional on $\tau_{1:T}^j$, the objective is to draw $\Theta_\tau^{j+1} = \{\Sigma_\tau^{j+1}\}$. For ease of exposition, we omit the superscript j below. Define $X = [\tau_2, \dots, \tau_T]'$ and $W = [\tau_1, \dots, \tau_{T-1}]'$. Similarly as before we draw from

$$\Sigma_\tau \sim IW(\bar{\Psi}, \bar{d}), \quad \bar{\Psi} = (X - W)'(X - W) + \underline{\Psi}, \quad \bar{d} = T - 1 + \underline{d}. \quad (23)$$

C Robustness Checks

In this appendix, we investigate the robustness of our results to different specifications of the TC-VAR model. We consider two robustness analyses that relate to alternative specifications for the one- and ten-year-ahead inflation expectations surveys.

C.1 Two inflation cycles

Specification and priors. In the baseline specification, we decompose the inflation rate and the one- and ten-year-ahead inflation expectation measures as in (1) and that we report below in (24)

$$\pi_t = \tau_{\pi,t} + c_{\pi,t} + (\eta_{\pi,t} - \eta_{\pi,t-1}), \quad (24.1)$$

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^e, \quad (24.2)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + \delta c_{\pi,t}^e + \eta_{\pi,t}^{e,10y}. \quad (24.3)$$

While we leave unchanged the decomposition for the inflation rate in (24.1), we now consider the following alternative specification for the two inflation expectations surveys

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t}^e, \quad (25.1)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + \delta c_{\pi,t}^e + \eta_{\pi,t}^{e,10y}. \quad (25.2)$$

The one- and ten-year-ahead inflation expectation surveys are assumed to share a common cyclical component, $c_{\pi,t}^e$. Because the ten-year inflation expectations $\pi_t^{e,10y}$ are fairly stable over time, we assume that its loading is less than one $\delta < 1$. In particular, we estimate δ assuming a normal prior distribution centered at 0 and with variance 0.1. Lastly, we allow for idiosyncratic errors for the ten-year-ahead inflation expectations. Under our baseline specification, the joint dynamics of trend and cyclical components are described in (6) where $\tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}\}'$, $c_t = \{c_{y,t}, c_{u,t}, c_{f,t}, c_{\pi,t}, c_{\pi,t}^e\}'$ and $\varepsilon_{c,t} = \{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}, \varepsilon_{c,\pi,t}^e\}'$.

Initial conditions and priors. Under this alternative specification, we modify the assumptions on the prior for the standard deviation of the shocks to the trends. In particular, we set the standard deviation for the expected change in the annualized trend of real GDP growth and unemployment rate to 1% over a time period of 100 and 50 years respectively. For the annualized real interest rate and inflation, we assume a one-percent standard deviation for the expected change in their trends over 25 years. These assumptions imply that the prior covariance matrix of the shocks to the trends is diagonal with the following elements on the main diagonal $diag(\underline{\Sigma}_{\tau}) = [1/100, 1/50, 1/25, 1/25]$.

Finally, because pre-sample data is not available for either the one- or ten-year-ahead inflation

expectations, we set the prior for the standard deviation of the shocks affecting the common cyclical component of the two inflation expectations survey to 1.5%, thus coinciding with the corresponding standard deviation for the inflation rate. As a result of these assumptions, the prior covariance matrix of the shocks to the cycles is diagonal and such that the value on the main diagonal approximately correspond to $diag(\underline{\Sigma}_c) = [25, 1.3, 0.6, 2.2, 2.2]$.

Results. We estimate the TC-VAR model in (5) and (6) over the 1960-2019 period. We then identify the shock targeting $c_{u,t}$ and report in Table 9 its contributions to the volatility of the cyclical components of the other variables. Even in this case, the results are in line with those of the baseline specification.

Table 9: Two inflation cycles: Variance contribution of unemployment shock

All frequencies ($0 - \infty$ quarters)				
Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)
67.4	74.6	52.5	47.3	58.5
[58.5,78.3]	[54.7,85.3]	[15.2,86.4]	[24.2,67.2]	[18.6,82.8]
All-but-short-run frequencies ($6 - \infty$ quarters)				
Unemployment	Output	FFR	Inflation	Inflation Exp.(1y)
68.0	75.4	52.6	50.3	62.1
[59.2,78.9]	[55.8,85.9]	[15.5,87.0]	[25.9,70.6]	[20.1,86.6]

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

C.2 One common inflation cycle

Specification and priors. In addition to the decomposition for realized inflation in (24.1), this specification assumes the following alternative decomposition for the one- and ten-year-ahead inflation expectation measures

$$\pi_t^{e,1y} = \tau_{\pi,t} + c_{\pi,t} + \eta_{\pi,t}^{e,1y}, \quad (26.1)$$

$$\pi_t^{e,10y} = \tau_{\pi,t} + \eta_{\pi,t}^{e,10y}. \quad (26.2)$$

The specification in (26) assumes that there is one common cyclical component $c_{\pi,t}$ for inflation which is shared across realized inflation π_t in (24.1) and the one-year-ahead inflation ex-

pectation $\pi_t^{e,1y}$ in (26.1). Because the ten-year inflation expectation $\pi_t^{e,10y}$ is fairly stable over time, we assume that it proxys for the trend component but does not include any cyclical component. In this case, the joint dynamics of trend and cyclical components are described in (6) where $\tau_t = \{\tau_{g,t}, \tau_{u,t}, \tau_{r,t}, \tau_{\pi,t}\}'$, $\varepsilon_{\tau,t} = \{\varepsilon_{\tau,g,t}, \varepsilon_{\tau,u,t}, \varepsilon_{\tau,r,t}, \varepsilon_{\tau,\pi,t}\}'$, $c_t = \{c_{y,t}, c_{u,t}, c_{f,t}, c_{\pi,t}\}'$, $\varepsilon_{c,t} = \{\varepsilon_{c,y,t}, \varepsilon_{c,u,t}, \varepsilon_{c,f,t}, \varepsilon_{c,\pi,t}\}'$.

Under this alternative specification, we keep the same assumptions on the standard deviation of the shocks to the trends as under the baseline, that is $diag(\underline{\Sigma}_{\tau}) = [1/40, 1/20, 1/20, 1/20]$. We also assume that the prior for the standard deviation of the shock to the common cyclical component of inflation coincides with the pre-sample standard deviation of realized inflation. Therefore, $diag(\underline{\Sigma}_c) = [25, 1.3, 0.6, 2.2]$.

Table 10: One common inflation cycle: Variance contribution of unemployment shock rate

All frequencies (0 – ∞ quarters)			
Unemployment	Output	FFR	Inflation
74.5	81.8	43.9	38.8
[64.4,85.0]	[71.1,90.1]	[19.4,73.3]	[25.1,54.2]
All-but-short-run frequencies (6 – ∞ quarters)			
Unemployment	Output	FFR	Inflation
75.7	82.7	43.7	43.9
[65.4,85.9]	[71.9,90.8]	[19.7,73.0]	[28.6,60.0]

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

Results. Estimating the proposed alternative specification of the TC-VAR model over the 1960-2019 period, we identify the shock targeting $c_{u,t}$ and report in Table 10 its contributions to the volatility of the cyclical components of the other variables. The results show that the findings under the baseline specification carry over. In fact, the shock explains nearly 40% of the volatility of the common cyclical component of inflation. When excluding movements in cyclical components at frequencies higher than 1.5 years, the results point to a somewhat stronger explanatory power of the identified shock for the common cyclical component of inflation.

D Details on Long-run Priors

D.1 Long-run priors for [Angeletos et al. \(2020\)](#)

When estimating the VAR model using the long-run priors, we consider the arbitrary ordering of the observables $y_t = \{Y_t, I_t, C_t, u_t, h_t, TFP_t, w_t h_t / Y_t, \pi_t, R_t\}'$. We assume that the following matrix H captures the cointegrating relationships in the long run

$$H = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 & 1 \end{bmatrix}.$$

D.2 Long-run priors for VAR model in [Subsection 7.2](#)

When estimating the VAR model using the long-run priors, we consider the arbitrary ordering of the observables $y_t = \{g_t, u_t, f_t, \pi_t, \pi_t^{e,1y}\}'$. We assume that the following matrix H captures the cointegrating relationships in the long run

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & -1 & 1 \end{bmatrix}. \tag{27}$$

E Supplementary Material to Section 8

E.1 Definition of stabilizability

Definition 1 *The pair (A, B) is stabilizable if any of the following conditions holds:*

- *There exists no left eigenvector of A associated with an eigenvalue having nonnegative real part that is orthogonal to the columns of B ;*

$$\begin{cases} \nu^* A = \lambda \nu & (\text{Re}[\lambda(A)] \geq 0) \\ \nu^* B = 0 \end{cases} \Rightarrow \nu = 0.$$

- *rank $[\lambda I - A \quad B] = \dim(A)$ for all $\text{Re}[\lambda(A)] \geq 0$.*

E.2 Matrices for time-invariant innovations representation

As shown in [Fernandez-Villaverde et al. \(2007\)](#), the time-invariant matrices \hat{B} and \hat{D} satisfy the following equations:

$$\Sigma = A \Sigma A' + B B' - (A \Sigma C' + B D') (C \Sigma C' + D D')^{-1} (A \Sigma C' + B D')', \quad (28.1)$$

$$K = (A \Sigma C' + B D') (C \Sigma C' + D D')^{-1}, \quad (28.2)$$

$$\hat{D} \hat{D}' = D D' + C \Sigma C', \quad (28.3)$$

$$\hat{B} = K \hat{D}. \quad (28.4)$$

E.3 The AR representation of a univariate TC-AR model

In this appendix, we provide an intuitive, analytical example that considers the unobserved components model used by [Stock and Watson \(2007\)](#) to test whether the U.S. inflation process experienced a structural change since the beginning of the Great Moderation. Inflation is described by the following state-space representation

$$\pi_t = \Lambda_\tau \tau_{\pi,t} + \Lambda_\eta \varepsilon_{\eta,\pi,t}, \quad (29.1)$$

$$\tau_{\pi,t} = \Phi_\tau \tau_{\pi,t-1} + \mathcal{R} \varepsilon_{\tau,\pi,t}, \quad (29.2)$$

where

$$\Phi_\tau = 1, \quad \mathcal{R} = 1, \quad \Lambda_\tau = 1, \quad \Lambda_\eta = 1,$$

and $\varepsilon_{\tau,\pi,t} = Q w_{\tau,\pi,t}$, $Q = \sigma_\tau$, $\varepsilon_{\eta,\pi,t} = \sigma_\eta w_{\eta,\pi,t}$ such that $E(w_t w_t') = \mathcal{I}$ where $w_t = \{w_{\tau,\pi,t}, w_{\eta,\pi,t}\}'$. Defining $y_t = \pi_t$ and $x_t = \tau_t$, the representation in (29) coincides with (5) and (6) where, following the notation in [Anderson and Moore \(1979\)](#), we appended the measurement error $\varepsilon_{\eta,\pi,t}$ directly

in (5) as opposed to redundantly defining it as $\eta_{\pi,t}$ in the transition equation (6). Thus, we verify two conditions: i) the pair $(\Phi_\tau, \mathcal{R}Q)$ is stabilizable; ii) the pair $(\Phi'_\tau, \Lambda'_\tau)$ is stabilizable. Because Φ_τ has only one (unit-root) eigenvalue, then

$$\text{rank} [I - \Phi_\tau \quad \mathcal{R}Q] = \text{rank} [0 \quad \sigma_\tau] = 1,$$

and

$$\text{rank} [I - \Phi'_\tau \quad \Lambda'_\tau] = \text{rank} [0 \quad 1] = 1.$$

Therefore, the asymptotic properties of the Kalman filter hold, and we can derive the AR(∞) representation of (29). In particular, we write (29) as in (10) where

$$A = 1, \quad B = \begin{bmatrix} \sigma_\tau & 0 \end{bmatrix}, \quad C = 1, \quad D = \begin{bmatrix} \sigma_\tau & \sigma_\eta \end{bmatrix},$$

assuming that the shock to the trend is $\varepsilon_{\tau,\pi,t} = \sigma_\tau w_{\tau,\pi,t}$ and the measurement error is $\eta_{\pi,t} = \sigma_\eta w_{\eta,\pi,t}$ and $E(w_t w_t') = \mathcal{I}$ where $w_t = \{w_{\tau,\pi,t}, w_{\eta,\pi,t}\}'$. Defining $\sigma_{\hat{\tau}} = \Sigma$ as the standard deviation of the error associated with the estimate of the trend, $(\tau_{\pi,t} - \hat{\tau}_{\pi,t})$, we can use the equation (28) in Appendix E.2 to derive the time-invariant matrices \hat{B} and \hat{D} as

$$\sigma_{\hat{\tau}}^2 = \frac{1}{2} \left(-\sigma_\tau^2 + \sqrt{\sigma_\tau^4 + 4\sigma_\tau^2 \sigma_\eta^2} \right) > 0, \quad (30.1)$$

$$K = 1 - \delta, \quad (30.2)$$

$$\hat{D}\hat{D}' = (\sigma_{\hat{\tau}}^2 + \sigma_\tau^2 + \sigma_\eta^2), \quad (30.3)$$

$$\hat{B} = (1 - \delta) (\sigma_{\hat{\tau}}^2 + \sigma_\tau^2 + \sigma_\eta^2)^{1/2}, \quad (30.4)$$

where $\delta = \sigma_\eta^2 / (\sigma_{\hat{\tau}}^2 + \sigma_\tau^2 + \sigma_\eta^2) < 1$. Finally, using (13) and (30), we map the state-space representation in (29) maps into the infinite-order autoregression, AR(∞),

$$\begin{aligned} \pi_{t+1} &= \sum_{s=0}^{\infty} C \left(A - \hat{B}\hat{D}^{-1}C \right)^s \hat{B}\hat{D}^{-1}y_{t-s} + \hat{D}\nu_{t+1} \\ &= \sum_{s=0}^{\infty} C (A - KC)^s K\pi_{t-s} + \hat{D}\nu_{t+1}, \\ &= (1 - \delta) \sum_{s=0}^{\infty} \delta^s \pi_{t-s} + (\sigma_{\hat{\tau}}^2 + \sigma_\tau^2 + \sigma_\eta^2)^{1/2} \nu_{t+1}, \end{aligned} \quad (31)$$

where $\nu_t \sim \mathcal{N}(0, 1)$. Equation (31) shows that, even with infinite data, the estimation of the AR(∞) representation leads to parameter estimates and VAR residuals that confound the standard deviation of both the measurement error σ_η and the innovations to the trend σ_τ with the standard deviation $\sigma_{\hat{\tau}}$ resulting from the error associated with the estimate of the trend, $(\tau_{\pi,t} - \hat{\tau}_{\pi,t})$.

E.4 Bivariate example: An alternative calibration

In this appendix, we consider an alternative parameterization of the model in Subsection 8.2. While we leave unchanged the parameterization for the unemployment rate process, we set ρ_π to 0.5, thus assigning an equal loading on cyclical inflation of its lag and the cyclical unemployment rate even though with opposite signs. We further assume that $\Sigma_\tau = \mathcal{I}$, $\sigma_{c,u} = 1$ and $\sigma_{c,\pi} = 0$, where the latter assumption ensures that the inference about the contribution of the unemployment-identified shock is not affected by the shocks to the cyclical component of inflation. Finally, we follow the same steps described for the previous specification and, after producing a new data simulation, we estimate a bivariate VAR model and subsequently identify the shock targeting the unemployment rate at business-cycle frequencies.

Table 11: Bivariate model: Variance contribution of unemployment shock

Unemployment	Inflation
99.8	3.1
[99.6, 99.9]	[1.8, 4.6]

Notes: The shock is identified by maximizing its contribution to the volatility of the unemployment rate over business-cycle frequencies (6-32 quarters). We report the median contribution and the corresponding 68-percent posterior-coverage interval of the identified shock to the variance of all the other variables over the same frequencies.

Table 11 reports the associated contribution of the shock for both series at those frequencies. As expected, the results show that the contribution of the unemployment-rate shock to inflation at business-cycle frequencies is lower—by about about 70%—than the shock contribution reported in Table 8. Intuitively, under this alternative calibration, the inflation rate depends on its lagged cyclical component and the innovations to its trend component in addition to the movements in the lagged cyclical component of the unemployment rate. Consequently, the estimated bivariate VAR confounds all these effects, implying a lower contribution of the unemployment-identified shock on inflation at business-cycle frequencies. To conclude, even if the long simulation was generated under the assumption that their cyclical components were evidently related, the identified shock does not capture this feature of the simulated data.