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

August 22, 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 implied by the New Keynesian framework. We explain why our results differ from previous contributions that conclude that inflation is unrelated to real activity at business-cycle frequencies.

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

The 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 (recently, see Del Negro et al., 2020). 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. Importantly, monetary policy is non-neutral in the short run and can be designed to optimally respond to such shocks. At the same time, monetary policy is assumed to be neutral in the long run. Thus, while monetary policy can be used to stabilize the economy, it cannot permanently affect the level of real activity. Any attempt of pushing the economy beyond potential would result in higher inflation and no benefit on real activity.

However, the increasing evidence of a weaker relationship between real economic activity and inflation since the 1990s questions the empirical foundations of New Keynesian models. As a result, a series of recent contributions have focused on identifying explanations for the weakening of this empirical relationship or on modifying structural features of the New Keynesian framework to improve the empirical fit of New Keynesian models. By contrast, a recent literature interprets the recent evidence as a reason to abandon the New Keynesian way of thinking altogether and proposes alternative explanations for the sources of business-cycle fluctuations. For example, Angeletos et al. (2018) show that, in presence of uncertainty about the beliefs of others, agents can generate waves of optimism and pessimism about the short-term economic outlook leading to demand-driven fluctuations with no inflationary effects. Alternatively, while abstracting from the implications for inflation, Beaudry et al. (2020) propose a model that, by means of real-financial linkages and an accumulated capital stock, generates endogenous propagation mechanisms characterized by boom-bust cycles driving business-cycle movements. The empirical foundation of these alternative theoretical frameworks is offered by Angeletos et al. (2020). They find that business-cycle fluctuations are explained by a demand-side shock which is disconnected from movements in inflation over the entire post-war period.

In this paper, we provide empirical evidence seeking to resolve the debate on the empirical foundations of these two strands of the literature and ultimately answer the question: Does

inflation comove with measures of real economic activity over the business cycle? In other words, is there evidence that inflation behaves as suggested by the New Keynesian framework at business-cycle frequencies? To this purpose, we adopt a dynamic, multivariate approach in line with several previous studies. However, we argue that two features are key to properly understand the empirical relationship between inflation and measures of real economic activity. First, we have to distinguish business-cycle and low-frequency movements of those variables within a unified empirical framework. Second, we need to disentangle business-cycle fluctuations in inflation from its high-frequency movements. Equivalently, measures of inflation and real activity *over the business cycle* have to be explicitly purged of *both* low- and high-frequency movements.

We start by presenting a simple, but insightful, analysis that motivates the adoption of the two key features just described. 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". We find that, 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 0.20 (negative 0.35). 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.

Motivated by this analysis, we adopt a more rigorous empirical framework and estimate a multivariate Trend-Cycle VAR (TC-VAR) model, building on the work of Stock and Watson (1988) 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, we add the median 10-year-ahead average inflation expectations from the Survey of Professional Forecasters (SPF) to better capture low-frequency movements in inflation. 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.

In our baseline specification, we consider a common trend that captures low-frequency movements in inflation and 1- and 10-year-ahead inflation expectations. By contrast, we assume that the trend of the real per-capita GDP growth, unemployment rate and real interest rate are uncorrelated. To identify shocks driving business-cycle fluctuations, we adapt the

method proposed by Angeletos et al. (2020) for a VAR to our TC-VAR model. In particular, 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.

The key difference with respect to a VAR is that the TC-VAR model is specified in a way to separate cycle and low-frequency fluctuations. Thus, a different set of parameters control the variability of the variables at different frequencies. We consider two specifications. In the first one, we consider all frequencies of the cycle component. In the second one, we further refine the results by focusing on the business-cycle frequencies of the cycle component. This allows to clean the results from the high-frequency movements of inflation that are arguably unrelated to the business cycle.

A series of important results emerge from this exercise. First, over all the frequencies of the cyclical components, the shock targeting the unemployment rate explains nearly 30% of the volatility of inflation over the 1960-2019 period. This result suggests that the low-frequency movements in the considered real and nominal variables need to be explicitly accounted for when studying their relationship at business-cycle frequencies. Second, the shock targeting the unemployment rate explains roughly 55% of the business-cycle volatility of 1-year-ahead inflation expectations. Therefore, the adoption of a measure of inflation which abstracts from its high-frequency movements points to evidence of a noticeably stronger connection between inflation and unemployment rate than found when using a measure of inflation that is instead subject to changes at high frequencies.

Third, when focusing on the business-cycle frequencies of the cycle component, the results become even stronger. The shock identified targeting the rate of unemployment explains a higher percentage of the business-cycle volatility of both inflation and 1-year-ahead inflation expectations compared to when all the frequencies are considered: 35% and nearly 70%, respectively. This result in line with the evidence presented in our motivating evidence exercise, suggesting a strong relation between inflation and real activity at business-cycle frequencies.

When estimating the model over the period between 1960 and 2008, the explanatory power of the shock targeting the unemployment rate for the business-cycle volatility of inflation and 1-year-ahead inflation expectations is stronger than over the period between 1960 and 2019. This can be interpreted as evidence of a recent weakening of the relationship between real and nominal variables. However, this change could also be due to the particular policies implemented at the zero lower bound (Bianchi and Melosi, 2017). Fifth, the previous four results are robust to alternative specifications that allow for the estimation of the initial conditions of the trend components of the VAR and also for the estimation of a cointegrating relationship between the trend growth rate of real per-capita GDP and the trend of both the unemployment rate and the real interest rate.

In addition, we assess the importance for our results of modeling time-varying latent trends or including data on short- and longer-term inflation expectations. 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 Angeletos et al. (2020) who adopt a standard VAR model. If we exclude the data on short- and longer-term inflation expectations, the main result that inflation and the unemployment rate are related over the business cycle persists, but the resulting latent trend of inflation displays a substantially higher volatility and shows a delay in capturing the Volcker disinflation.

We then explain why our results differ from recent contributions that challenge the validity of the New Keynesian paradigm. We show that a fixed coefficient VAR 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 movements cannot be resolved. Even if an econometrician could correctly reconstruct the VAR representation of the TC-VAR, the reduced-form innovations that she would recover would not map into the innovations affecting the latent persistent and stationary components. By contrast, the reduced-form innovations would correspond to a linear combination of the shocks to the latent components. 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.

Overall, these 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.

2 Motivating Empirical Evidence

In this section, we provide motivating evidence for the main thesis proposed in this paper: Inflation and real activity are in fact related over the business cycle once we control for trends in the variables of interest. In particular, we seek to 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.

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

Correlations of $\tilde{\pi}_t$ at business-cycle frequencies $(6-32$ quarters)

	$j=8$	$j=6$	$j=4$	$j=2$	$j=0$
$\tilde{Y}_{t-j}$	0.15	0.35	0.46	0.42	0.22
$\tilde{u}_{t-j}$	0.09	-0.13	-0.34	-0.43	-0.36

Correlations of $\tilde{\pi}_t$ at medium-cycle frequencies $(8-30$ years)

	$j=8$	$j=6$	$j=4$	$j=2$	$j=0$
$\tilde{Y}_{t-j}$	0.66	0.64	0.59	0.52	0.42
$\tilde{u}_{t-j}$	-0.59	-0.56	-0.51	-0.44	-0.33

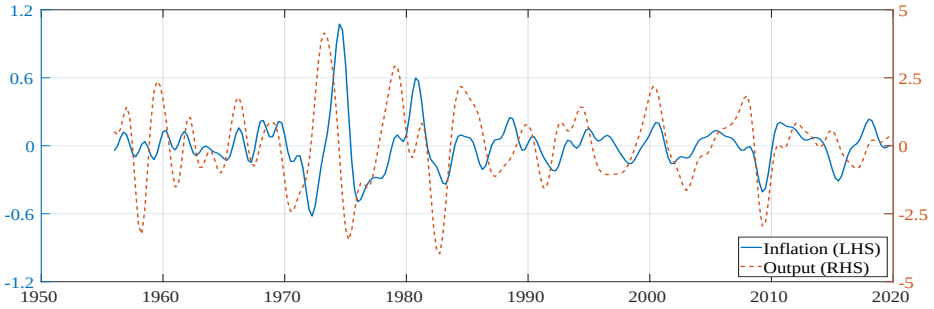
We consider the inflation rate—measured as the log difference in the GDP deflator, π_t—and two measures of real economic activity—the log level of real GDP per capita, Y_t, and the unemployment rate, u_t—over the period starting from 1956:Q1 to 2019:Q4. Using a bandpass filter, we extract the corresponding 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.

Figure 1 plots the filtered time series at business-cycle frequencies for inflation, $\tilde{\pi}_t$, and real per-capita GDP, $\tilde{Y}_t$—panel (a)—and for inflation and the rate of unemployment, $\tilde{u}_t$—panel (b). Figure 2 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 further investigate the role of measures of real economic activity as leading indicators for inflation, Table 1 reports the correlations between current (filtered) inflation , $\tilde{\pi}_t$, and current and lagged (filtered) levels of real per-capita GDP and unemployment rate, $\left\{\tilde{Y}_{t-j}, \tilde{u}_{t-j}\right\}$ for $j=\{0,2,6,8\}$, at 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 0.60). These correlations are larger (in absolute value) than the counterparts with real per-capita GDP and the rate of unemployment lagged by

Figure 1: Business-cycle frequencies (6-32 quarters)

(a) Inflation and output



(b) Inflation and unemployment rate

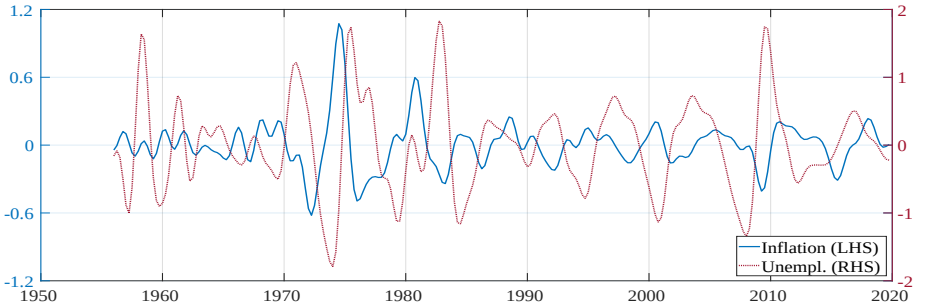
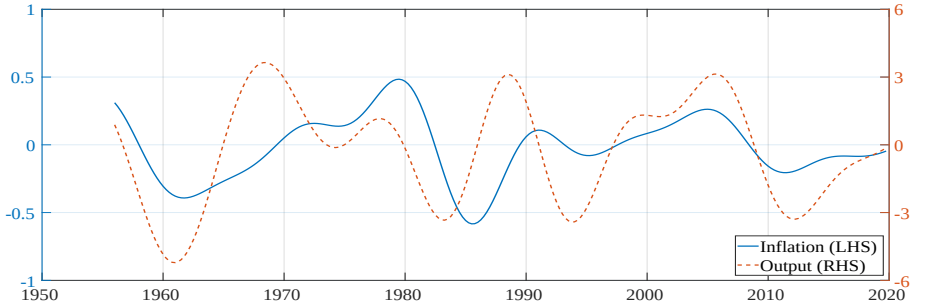
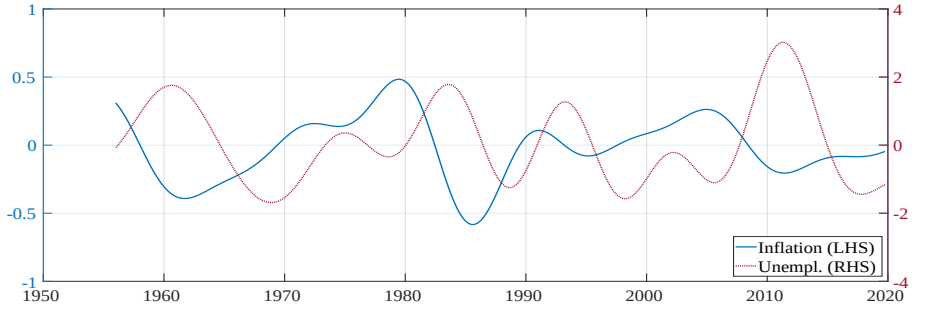


Figure 2: Medium-cycle frequencies (8-30 years)



(a) Inflation and output



(b) Inflation and unemployment rate

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

3 The Trend-Cycle VAR Model

The representation of the TC-VAR model follows the work of Stock and Watson (1988), and more recently Del Negro et al. (2017) and Johannsen and Mertens (2021). The measurement equation can be represented as

$$y_t = \Lambda s_t, \tag{1}$$

where y_t is a vector of observables, and $s_t = [\tau_t, c_t, c_{t-1}, u_t, u_{t-1}]'$ is a vector of states including vectors of trends τ_t, cycles c_t and measurement errors u_t.
The trend component evolves according to a unit-root process

$$\tau_t = \tau_{t-1} + \varepsilon_{\tau,t}, \tag{2}$$

and the cyclical component follows a stationary VAR process

$$\Phi(L)c_t = \varepsilon_{c,t}, \tag{3}$$

where $\Phi(L) = I - \sum_{l=1}^{p} \Phi_l L^l$. The resulting transition equation can be expressed as

$$s_t = \Gamma s_{t-1} + \Omega \varepsilon_t, \tag{4}$$

where the shocks are independent and identically distributed according to

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

the Σs are conforming positive definite matrices, and $\mathcal{N}(\cdot, \cdot)$ denotes the multivariate Gaussian distribution. The initial conditions are distributed according to

$$\tau_0 \sim \mathcal{N}\left(\underline{\tau_0}, \underline{V_0}\right)$$

$$c_{0:-p+1} \sim \mathcal{N}\left(0, V\left(\Phi, \Sigma_c\right)\right),$$

where $\underline{V_0}$ is an identity matrix and $V(\Phi, \Sigma_\varepsilon)$ is the unconditional variance of $c_{0:-p+1}$ implied by equation (3).

As explained in more detail in Section 5, we consider various specifications of the VAR with common trends. Each of those specification is characterized by a difference set of assumptions about four elements: i) the cointegrating relationships between the trend components of the model; ii) the parameters $\underline{\tau_0}$ used to center the distribution for the initial conditions of the trend components, τ_0; iii) the variance-covariance matrix adopted to specify the prior distribution for the corresponding matrix of the shocks to the trends, Σ_τ; and iv) the variance-covariance matrix which defines the prior distribution for the corresponding matrix of the shocks to the trends, Σ_c.

4 Inference

4.1 Data

We estimate the TC-VAR model using the following six quarterly time series:[1] i) the growth rate of real GDP per capita, g_t, measured as difference in log real GDP (Y_t); ii) the unemployment rate, u_t; iii) the effective federal funds rate (EFFR), R_t;[2] iv) the inflation rate, π_t, measured as the log difference in GDP deflator (PGDP); v) the median 4-quarter-ahead average PGDP inflation expectations, $\pi_t^{e,4}$, from the Survey of Professional Forecasters (SPF); vi) the median 40-quarter-ahead average CPI inflation expectations adjusted to account for difference between CPI inflation and PGDP inflation over the estimation sample, $\pi_t^{e,40}$, also taken from the SPF.

We use the period between 1956:Q1 and 1959:Q4 as pre-sample and subsequently estimate the VAR with common trends over two alternative sample periods: the period from 1960:Q1 to 2019:Q4 and the period from 1960:Q1 to 2008:Q3, thus excluding the period starting with the onset of the Great Financial Crisis (GFC).

In the following, we present the results assuming that the cycle component evolves according to a VAR with two lags $p = 2$. The results with four lags are robust and only slightly weaker. For each estimation of the VAR with common trends, we adopt the Gibbs sampler described in Appendix A. In particular, we use 100,000 draws, discard the first 50,000 and finally consider one in every 50 draws, thus leaving 1,000 draws to use for the identification of the shocks using the method of Angeletos et al. (2020) described next.

[1] With the exception of the unemployment rate, all variables are expressed at quarterly rates.

[2] Following Del Negro et al. (2017), we treat observations of the FFR at the zero lower bound (ZLB) as missing.

4.2 Identification

To identify the shocks, we adapt the methodology proposed by Angeletos et al. (2020) to the latent cycles evolving according to equation (3). Their approach is to identify a shock by finding the linear combination of the reduced-form residuals that maximizes its contribution to the volatility of a particular variable over a particular frequency band. To compare the results in our framework to those in Angeletos et al. (2020), we account for the two following considerations.

First, we estimate the TC-VAR model using the Gibbs sampler described in Appendix A. Second, we apply their method to the latent cycles in (3). Given that the cycle component is already cleaned of movements at frequencies other than business cycles, we do not need to remove the low frequencies. Thus, we consider two specifications. The shock targeting the unemployment rate is always identified over all the frequencies of the cycle component. However, we then consider the two frequency bands when studying the explanatory power of that shock for the volatility of the remaining variables. In the first one, we ask how much the unemployment-identified shock contributes to the volatility of the cycle component of the other variables at all frequencies. In the second case, we ask how much the unemployment-identified shock contributes to the volatility of the cycle 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 cycle component of other variables such as inflation present some high-frequency movements that are not related to the business cycle.

5 Results

This section presents the main results.

5.1 Baseline model

In the baseline model, the growth rate of real GDP and the unemployment rate evolve according to

$$\begin{aligned} g_t &= \bar{g}_t + \tilde{g}_t = \bar{g}_t + \left(\tilde{Y}_t - \tilde{Y}_{t-1}\right) \\ u_t &= \bar{u}_t + \tilde{u}_t, \end{aligned} \tag{5}$$

where we assume that the trend of the real GDP growth rate and of the unemployment rate are unrelated. Moreover, assuming that the Fisher equation holds in the longer run, the federal funds rate is expressed as

$$R_t = \bar{r}_t + \bar{\pi}_t + \tilde{R}_t, \tag{6}$$

where $\bar{r}_t$ and $\bar{\pi}_t$ correspond to trend real interest rate and trend inflation, respectively. Finally, the inflation rate and the expectations of inflation over the next 4 and 40 quarters are described by equation 5.1 below

$$\begin{aligned} \pi_t &= \bar{\pi}_t + \tilde{\pi}_t + (u_{\pi,t} - u_{\pi,t-1}) \\ \pi_t^{e,4} &= \bar{\pi}_t + \tilde{\pi}_t^{e,4} \\ \pi_t^{e,40} &= \bar{\pi}_t + \tilde{\pi}_t^{e,40}, \end{aligned}$$

and thus share the same trend inflation $\bar{\pi}$. In addition, we allow for a multiplicative measurement error for the level of the PGDP, which implies that, after taking log difference, the PGDP inflation rate features a negative moving average measurement error component. Thus, we can define the vector of observables $y_t \equiv \left\{g_t, u_t, R_t, \pi_t, \pi_t^{e,4}, \pi_t^{e,40}\right\}'$, the trends $\tau_t \equiv \{\bar{g}_t, \bar{u}_t, \bar{r}_t, \bar{\pi}_t\}'$, and the cycles $c_t \equiv \left\{\tilde{g}_t, \tilde{u}_t, \tilde{R}_t, \tilde{\pi}_t, \tilde{\pi}_t^{e,4}, \tilde{\pi}_t^{e,40}\right\}'$.

We also calibrate the initial conditions $\underline{\tau_0}$ for the trends τ_t. We set the initial condition for the trend of the unemployment rate, $\underline{u_0}$, the quarterly growth rate of real GDP, $\underline{g_0}$, and quarterly inflation, $\underline{\pi_0}$, to 5%, 0.2%, and 0.6% respectively, in line with their pre-sample averages. The initial condition for the trend real interest rate, $\underline{r_0}$, is set to 0.5%, reflecting the estimates of Del Negro et al. (2017) for the 1960s.

To specify the prior for the variance-covariance matrix of the shocks to the trends, we follow Del Negro et al. (2017) and assume that the matrix $\underline{\Sigma_\tau}$ is diagonal. Dividing each element of the diagonal by 400, we equivalently make assumptions on the standard deviation of the change in each trend expected *over 100 years*. In particular, we set the standard deviation for that expected change in the trend of annual growth rate of real GDP, annual real interest rate and annual inflation to 2.6%, 2.8% and 3.4%, respectively. Additionally, we set the standard deviation for that expected change in the trend unemployment rate to 0.85%.

We follow Del Negro et al. (2017) to specify the prior for the variance-covariance matrix of the shocks to the cycles Σ_c. We assume that the matrix $\underline{\Sigma_c}$ is diagonal, and we calibrate the standard deviation of the shocks affecting the stationary component of the quarterly growth rate of real GDP and quarterly inflation to 1.3% and 0.45% respectively, reflecting pre-sample standard deviations. The standard deviation of the shocks affecting the cycle of the unemployment rate and nominal interest rate are set to 1.5% and 0.7% respectively. These standard deviations are only slightly higher than their pre-sample standard deviations of 1.2% and 0.2% respectively to better capture the cyclical fluctuations in the unemployment rate and nominal interest rate observed during the post-war period. Because pre-sample data is not available for the quarterly 4- and 40-quarter inflation expectations, we set the standard deviation of the shocks affecting their stationary components to 0.4% and 0.3% respectively, thus lower that the corresponding standard deviation for inflation.

Results for baseline model. Figure 4 plots the data (solid black line) used for the estima-

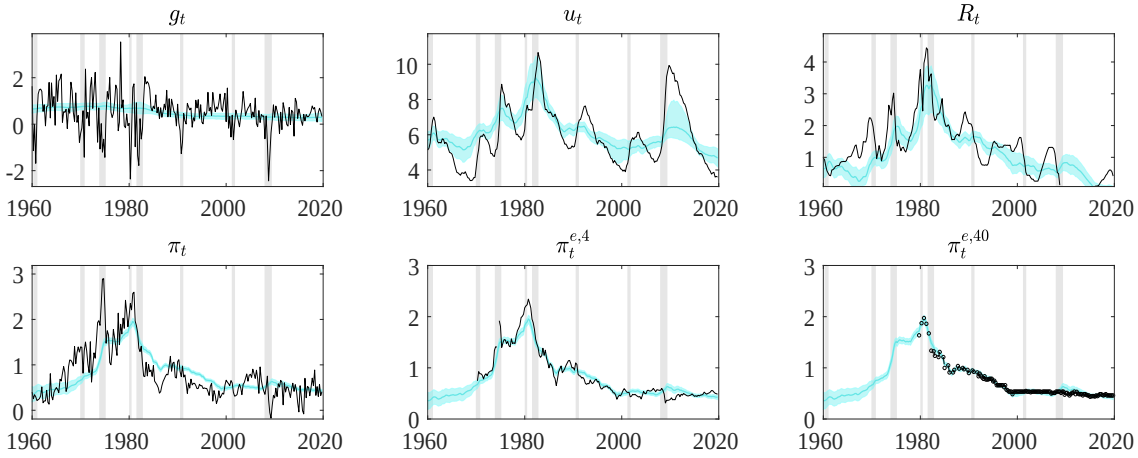
tion of the VAR with common trends over 1960-2019 period as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). Figure 5 plots the posterior median of the latent cycles (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). 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.

Given the focus of this paper, the behavior of the cycle components is reported in Figure 5. 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 cycle component of this variable behaves very much like inflation, but it is smoother.

To formalize the visual impression, after estimating the VAR model with common trends, we identify the shock that targets the cycle component of the unemployment rate, $\tilde{u}_t$, using the method of Angeletos et al. (2020). In particular, the shock is identified by maximizing its contribution to the volatility of the cycle 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.

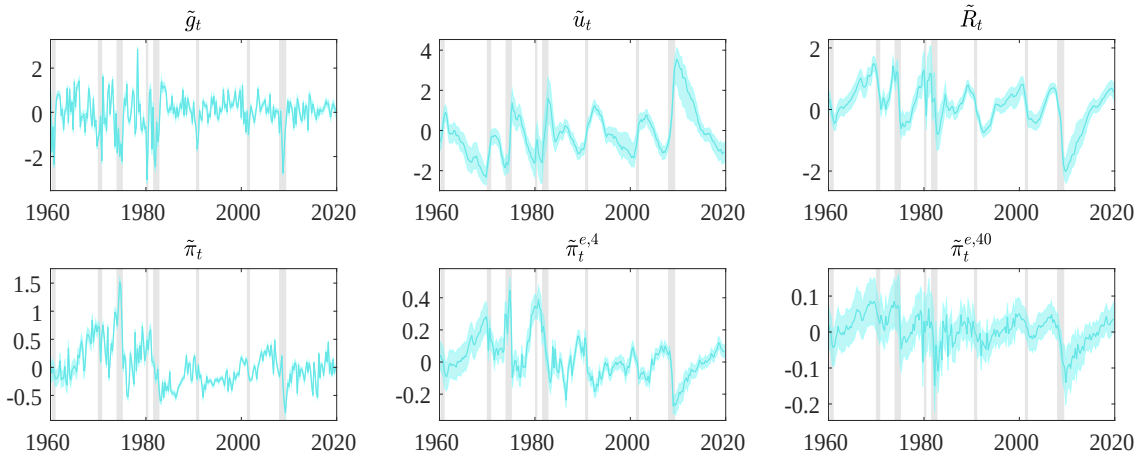
Table 2 reports the median contribution—and the corresponding 68-percent posterior-coverage interval—to the variance of the cycle of all the other variables over two frequency bands: a band that considers all the frequencies of the cycle—as for the construction of the shock—and the alternative band specified over all-but-short-run frequencies, thus starting from 6 quarters onward. By construction, the shock can explain the bulk of the fluctuations of the unemployment-rate cycle. However, the shock can also explain a sizable fraction of the cycle component of inflation. When excluding cycles shorter than 1.5 years, the unemployment-

Figure 4: Baseline model: Data and trends



The figure plots the data (solid black line) used for the estimation of the TC-VAR model over 1960-2019 period as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Figure 5: Baseline model: Cycles



The figure plots the posterior median of the latent cycles (solid blue line) of the TC-VAR model estimated over 1960-2019 period along with the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Table 2: Baseline model: Shock targeting $\tilde{u}_t$

All frequencies ($0 - \infty$ quarters)

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
95.8	9.6	28.6	93.7	55.5	39.3
[93.8,97.2]	[0.7,35.7]	[14.9,44.6]	[90.0,95.7]	[31.7,73.5]	[14.1,61.5]

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

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
97.7	9.6	35.2	96.1	68.8	60.0
[96.5,98.6]	[0.7,35.7]	[19.3,52.8]	[94.1,97.5]	[43.0,84.9]	[27.5,79.8]

identified shock explains more than 36% of inflation variability. The result is even more striking when looking at the cyclical component of one-year ahead inflation expectations: 68.8% of the business cycle variability is explained by the unemployment-identified shock. Given that the one-year-inflation expectations appear to be a smoother version of the inflation cycle component, this result corroborates the finding that inflation over the business cycle moves in a way consistent with the New Keynesian framework.

Four main conclusions can be drawn from the results reported in Table 2. First, over all the frequencies of the stationary components, the shock targeting $\tilde{u}_t$ explains about 30%, 55% and 40% of the volatility of $\tilde{\pi}_t$, $\tilde{\pi}_t^{e,4}$ and $\tilde{\pi}_t^{e,40}$, respectively. Second, over all-but-short-run frequencies, the identified shock explains more of the volatility of $\tilde{\pi}_t$, $\tilde{\pi}_t^{e,4}$ and $\tilde{\pi}_t^{e,40}$ than over all frequencies. Third, the shock explains roughly the entirety of the volatility of $\tilde{R}_t$ over both frequency bands. Combined with the previous two findings, this result likely implies that the shock also explains the volatility of the cyclical component of the real interest rate. Finally, the identified shock explains little of the volatility of $\tilde{Y}_t$ over all frequencies and also when excluding frequencies associated with the first 6 quarters of the cycles.

We also estimate the VAR model with common trends over the 1960-2008 period. We find that the dynamics of the latent trends and cycles are similar to those obtained from the model estimation over the 1960-2019 period and we then construct the identified shock targeting $\tilde{u}_t$.

Table 3 reports its median contribution—and the corresponding 68-percent posterior-coverage interval—to the variance of the cycle of all the other variables over the two considered frequency bands.

Compared to the results reported in Table 2, we find that the shock identified over the 1960-2008 period explains a larger percentage of the volatility of both $\tilde{\pi}_t$ and $\tilde{\pi}_t^{e,4}$ than their counterparts over the 1960-2019 period. Moreover, the exclusion of the period starting from the onset of the GFC somewhat restores the explanatory power of the identified shock for the volatility of $\tilde{Y}_t$.

Table 3: Baseline model: Shock targeting $\tilde{u}_t$ (sample: 1960-2008)

All frequencies ($0-\infty$ quarters)

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
93.4	26.5	42.9	88.6	60.1	36.0
[91.0,95.3]	[7.6,48.0]	[28.6,58.2]	[84.2,92.0]	[45.7,72.3]	[14.6,56.7]

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

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
96.5	26.5	51.4	92.6	73.9	53.7
[94.6,97.7]	[7.6,48.0]	[35.2,66.3]	[89.2,95.0]	[60.5,84.0]	[24.7,73.5]

5.2 Two Key Features

In this subsection, we show that two features are key to properly capture low-frequency movements in inflation and ultimately assess the relationship between real and nominal variables: the treatment of the latent trends as time-varying objects (Subsection 5.2.1) and the use of data on inflation expectations to capture low-frequency movements in inflation (Subsection 5.2.2).

5.2.1 Constant latent trends

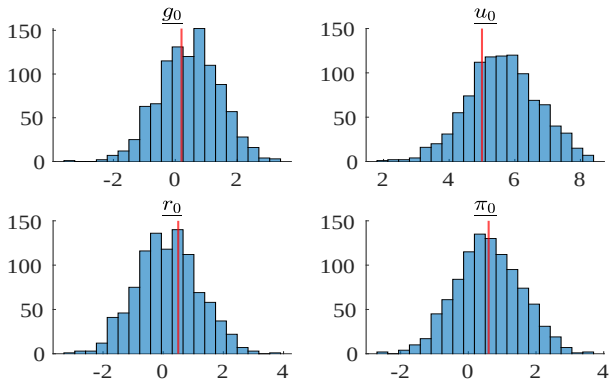
Differently from the previous specifications, we assume that the trends are constant by setting the standard deviation of all the shocks to the trends to zero. To allow some flexibility for the model to match the data, we leave unrestricted the initial conditions $\underline{\tau_0}$ as in Subsection 6.1 and set the scale parameter such that the acceptance rate is 38%.

Figure 7 plots the histogram of the initial conditions of the trends and, as for the previous specifications, the initial values are within the range of estimates obtained for these parameters. The 68-percent posterior-coverage intervals associated with these estimates correspond to the shaded blue areas in Figure 8, thus representing the constant latent trends of the estimated model. The resulting latent cycles are plotted in Figure 9 and clearly capture the difference between the data and the constant trends, both plotted in Figure 8.

Using the estimated coefficients and residuals from the VAR model describing the latent cycles, we identify the shock targeting the cyclical component of the unemployment rate over the two frequency bands considered for the previous specifications. Table 4 reports the contributions of the identified shock to the volatility of the cyclical components of the unemployment rate as well as the remaining variables.

In this specification in which we do not allow for low-frequency movements in the variables of interest, the identified shock explains about 10% of the volatility of not only inflation, but also short- and longer-term inflation expectations. The estimate of the shock's contribution to the volatility of inflation is roughly in line with the one provided in Angeletos et al. (2020).

Figure 7: Histogram of initial conditions $\underline{y_0}$



The figure plots the histogram of the initial conditions for the growth rate of real per-capita GDP $\underline{g_0}$ (top-left panel), unemployment rate $\underline{u_0}$ (top-right panel), real interest rate $\underline{r_0}$ (bottom-left panel) and inflation rate $\underline{\pi_0}$ (bottom-right panel) obtained from the estimation of the TC-VAR model with constant trends and unrestricted initial conditions over 1960-2019 period.

Table 4: Constant trends and unrestricted initial conditions: Shock targeting $\tilde{u}_t$

All frequencies ($0 - \infty$ quarters)

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
72.2	11.4	11.4	26.9	11.9	12.8
[61.4,83.4]	[1.4,37.7]	[2.7,54.0]	[10.6,58.1]	[2.1,58.7]	[3.1,60.6]

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

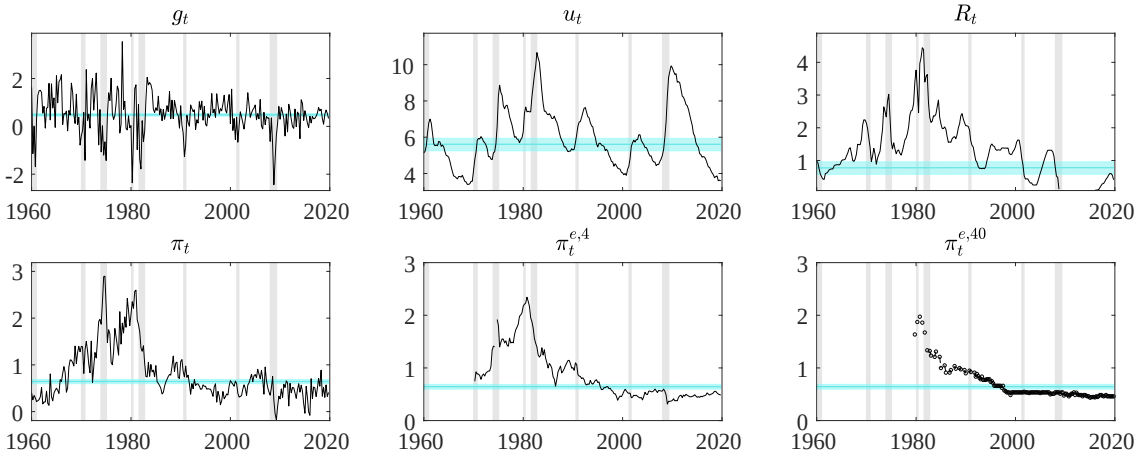
$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
72.8	11.4	11.6	27.2	11.9	12.9
[61.8,84.3]	[1.4,37.7]	[2.4,55.4]	[10.8,58.6]	[2.1,59.2]	[3.0,61.0]

Importantly, the results show that the assumption of a constant latent trend is not innocuous and can erroneously point to the absence of a relationship between real and nominal variables at business-cycle frequency. The similarity of our findings with those of Angeletos et al. (2020) is not surprising. Once the latent trends are assumed to be constant, our model collapses to a fixed-efficient VAR model similar to the one adopted in their paper with the exception that our setup also assumes that the constant trend of inflation is the same for the corresponding trend of short- and longer-run inflation expectations.

5.2.2 Inflation expectations

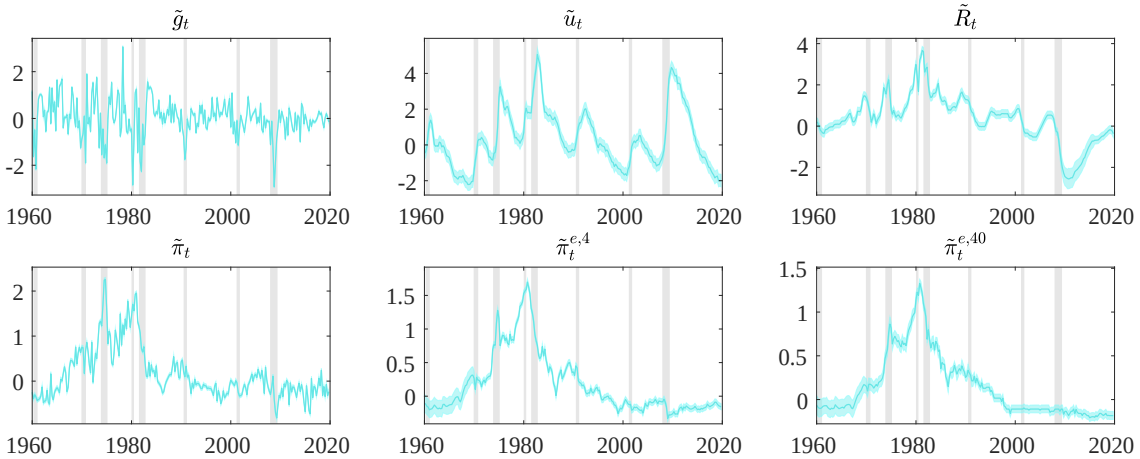
Having established the importance of modeling time-varying, rather than constant, trends, we show that the inclusion of data on inflation expectations is also important to appropriately capture low-frequency movements in inflation. To this purpose, we estimate the baseline

Figure 8: Constant trends and unrestricted initial conditions: Data and trends



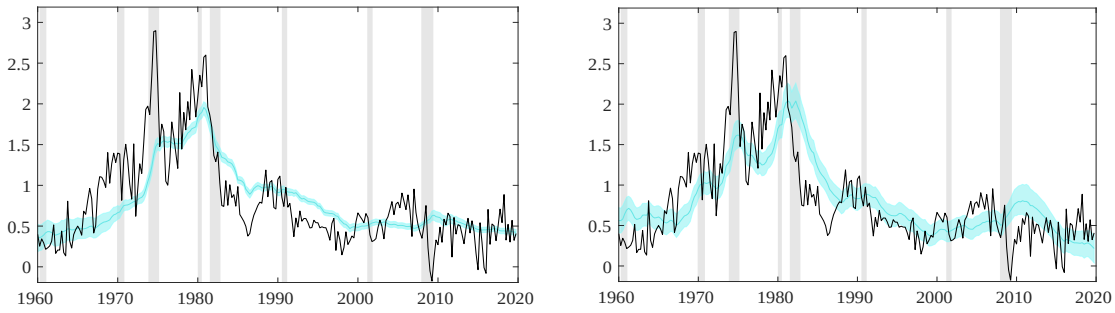
The figure plots the data (solid black line) used for the estimation of the TC-VAR model over 1960-2019 period as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Figure 9: Constant trends and unrestricted initial conditions: Cycles



The figure plots the posterior median of the latent cycles (solid blue line) of the TC-VAR model estimated over 1960-2019 period along with the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Figure 11: Baseline model: Inflation data and trend



(a) With inflation expectations

(b) Without inflation expectations

For both panels plot the data for inflation (solid black line) as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area) after estimating the baseline TC-VAR model is estimated over 1960-2019 period. However, the model is estimated either with data on inflation expectation—panel (a)—or without—panel (b). NBER recessions are denoted by shaded grey areas.

TC-VAR model as in Section 5, but without using data on short- and long-run inflation expectations. Figure 11 reports the data on inflation and its corresponding latent trend for the cases in which the model is estimating with data on inflation expectations—panel (a)—or without—panel (b). Note that the the plot in panel (a) is equivalent to the corresponding panel for inflation in Figure 4 and discussed in Section 5.

We find that the main result of a strong connection between inflation and measures of real economic activity holds if we estimate the model without data on inflation expectations. However, the comparison of the trend components of inflation when the model is estimated with or without data on inflation expectations highlights three important differences. First, the exclusion of data on inflation expectations implies that trend inflation captures not only low-frequency movements in inflation, but also higher-frequency changes. This result is supported by the noticeably higher volatility particularly for the period prior to the Volcker disinflation. Second, when excluding data on inflation expectations, the reversal in the upward trend in inflation induced by the Volcker disinflation takes place only with a noticeable lag. In this case, the estimation of the model cannot rely on the information associated with the change in inflation expectations induced by the switch to a macroeconomic regime characterized by an aggressive monetary-policy stance and fiscal backing (Bianchi and Ilut, 2017). Third, as indicated by the wider 68-percent posterior-coverage intervals, the estimates of the latent inflation trend are subject to a substantially higher degree of uncertainty if data on inflation expectations are not used to estimate the model. As a result, these differences point to the adoption of such data to properly capture the low-frequency movements in inflation.

6 Robustness

In this subsection, we show that our results are robust to different specifications of the TC-VAR model. We consider two robustness checks. In the first one, we allow for the estimation of initial conditions of trends, $\underline{\tau_0}$. In the second one, we allow for the estimation of the initial conditions of trends, $\underline{\tau_0}$, and two cointegrating relationships—between $\bar{g}_t$ and $\bar{u}_t$ and between $\bar{g}_t$ and $\bar{r}_t$.

6.1 Unrestricted initial conditions

Unrestricted initial conditions. In this subsection, rather than calibrating the initial conditions $\underline{\tau_0}$ for the trends τ_t as in the baseline specification in Subsection 5.1, we estimate them jointly with the TC-VAR model in equations (2) and (3).[3]

To this end, we introduce a random walk Metropolis-Hastings step in the Gibbs sampler used for the estimation of the model. Defining the starting value $\underline{\tau_0}^{(0)}$ with the initial conditions $\underline{\tau_0}$ chosen for the baseline specification, we make a draw ϑ for each step j of the Gibbs sampler from the following proposal distribution:

$$z\left(\vartheta|\underline{\tau_0}^{(s-1)}\right) = \mathcal{N}\left(\underline{\tau_0}^{(s-1)}, c\,\mathbf{I}\right) \tag{7}$$

where $\mathbf{I}$ is an identity matrix, and we set c to ensure an acceptance rate between 30% and 40% depending on the data sample and lags used. We accept the jump from $\underline{\tau_0}^{(s-1)}$ to ϑ $\left(\underline{\tau_0}^{(s)} = \vartheta\right)$ with probability $\min\left\{1, r_j\left(\underline{\tau_0}^{(s-1)}, \vartheta|Y\right)\right\}$ or otherwise reject the proposed draw and set $\underline{y_0}^{(s)} = \underline{y_0}^{(s-1)}$, where

$$r_j\left(\underline{y_0}^{(s-1)}, \vartheta|Y\right) = \frac{\mathcal{L}\left(\vartheta|Y\right)}{\mathcal{L}\left(\underline{y_0}^{(s-1)}|Y\right)}. \tag{8}$$

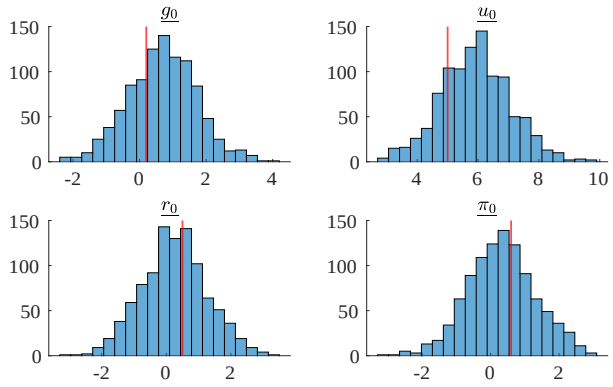
Results. We estimate the TC-VAR model over the 1960-2019 period and assuming two lags of the cyclical component in equation (3).[4] Figure 12 plots the histogram of the initial conditions $\underline{\tau_0}$ for the latent trends as well as the values used for the calibration of those initial conditions under the baseline specification (solid red lines). The figure shows that the calibrated values fall well within the range of estimates obtained when jointly estimating the initial conditions and the VAR model with trends.

Figure 13 and 14 report the data and the latent trends and cycles obtained from the estimation of this alternative specification. The latent trends and cycles obtained under this alternative specification are very similar to the ones obtained under the baseline specification.

[3] Compared with the baseline specification, we keep unchanged the assumptions related to the cointegration relationships, $\underline{\Sigma_c}$ and the standard deviations of the change in the trend of annual growth rate of real GDP, annual real interest rate and annual inflation expected over the next 100 years are also unchanged. We increase the standard deviation for that expected change in the trend unemployment rate from 0.85% to 0.9%.

[4] The results are roughly unchanged when we assume four lags of the VAR in equation (3).

Figure 12: Histogram of initial conditions $\underline{\tau_0}$



The figure plots the histogram of the initial conditions for the growth rate of real per-capita GDP $\underline{g_0}$ (top-left panel), unemployment rate $\underline{u_0}$ (top-right panel), real interest rate $\underline{r_0}$ (bottom-left panel) and inflation rate $\underline{\pi_0}$ (bottom-right panel) obtained from the estimation of the TC-VAR model with unrestricted initial conditions over 1960-2019 period.

Table 5: Unrestricted initial conditions: Shock targeting $\tilde{u}_t$

All frequencies $(0 - \infty$ quarters)

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
96.0	12.2	31.3	94.1	59.0	40.7
[93.9,97.4]	[1.3,37.0]	[17.3,48.1]	[91.3,96.2]	[38.8,76.3]	[17.4,62.3]

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

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
97.9	12.2	38.1	96.4	72.0	63.1
[96.5,98.7]	[1.3,37.0]	[21.9,56.0]	[94.3,97.8]	[51.7,86.6]	[32.9,80.6]

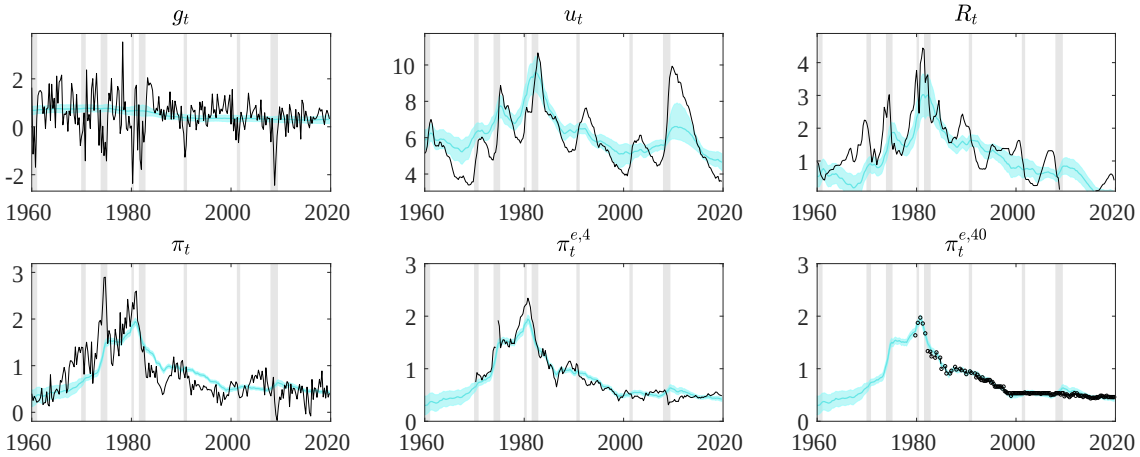
Using the model estimated over the 1960-2019 period, we identify the shock targeting $\tilde{u}_t$ and report in Table 5 its contributions to the volatility of the cyclical components of the other variables. The results show that the findings under the baseline specification carry over to the specification that estimates the initial conditions of the trends. In fact, the results are even stronger under this alternative specification.[5]

6.2 Unrestricted initial conditions and cointegrating relationships

Allowing for cointegration in trends. Relative to the specification in Subsection 6.1, the specification in this subsection relies on a different set of assumptions describing the cointegrating relationships between the trend of the real GDP growth, $\bar{g}_t$, and the trend of both the unemployment rate, $\bar{u}_t$, and the real interest rate, $\bar{r}_t$. In particular, we assume that

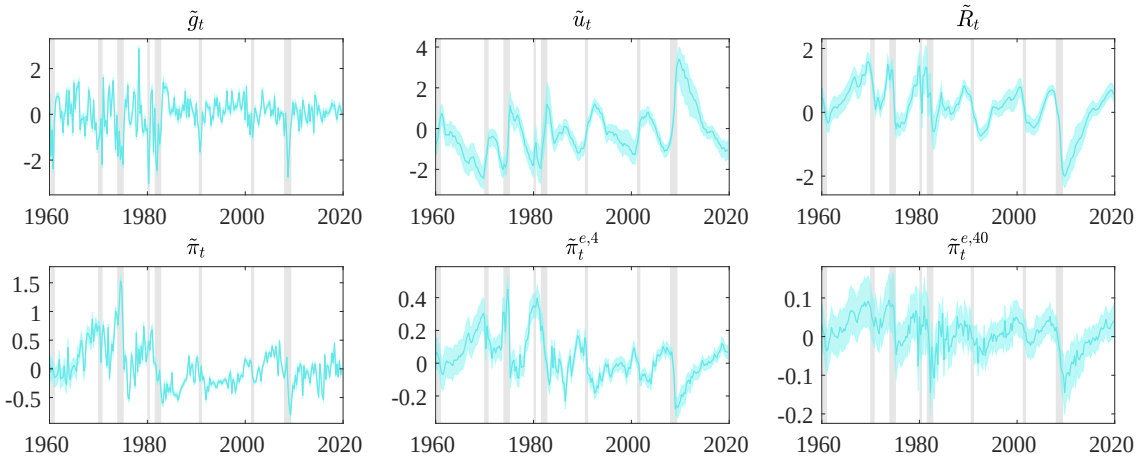
[5]We draw similar conclusions when estimating the model under the assumption of unrestricted initial conditions over the 1960-2008 period.

Figure 13: Unrestricted initial conditions: Data and trends



The figure plots the data (solid black line) used for the estimation of the TC-VAR model and the initial conditions $\underline{\tau_0}$ over 1960-2019 period as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Figure 14: Unrestricted initial conditions: Cycles



The figure plots the posterior median of the latent cycles (solid blue line) of the TC-VAR model estimated over 1960-2019 period along with the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

$$\bar{u}_t \quad = \quad \gamma_u \bar{g}_t + \bar{\delta}_{u,t}, \tag{9}$$

where $g_t = \bar{g}_t + \tilde{g}_t$ as in the baseline specification, and the parameter γ_u captures the relationship between the trend of real GDP growth and of the rate of unemployment. Because the trend of the unemployment rate may also be explained by other factors, we introduce the term $\bar{\delta}_{u,t}$ as a residual driver of the movements in $\bar{u}_t$.

Additionally, along with the baseline assumption in (6) that the Fisher equation holds in the longer run, we consider the possibility of the following cointegrating relationship

$$\bar{r}_t \quad = \quad \gamma_r \bar{g}_t + \bar{\delta}_{r,t}, \tag{10}$$

where $\bar{\delta}_{r,t}$ captures factors driving movements in the trend of the real interest rate such as the convenience yield as highlighted by Del Negro et al. (2017). We do not modify the assumptions about the cointegrating relationships of the other variables. Under this specification, we choose to estimate the parameters γ_u and γ_r jointly with the initial conditions, $\left\{\underline{g_0}, \underline{\delta_{u,0}}, \underline{\delta_{r,0}}, \underline{\pi_0}\right\}$. As for the alternative specification in Subsection 6.1, we use a random walk Metropolis-Hastings algorithm that we introduce as an additional step of the Gibbs sampler.

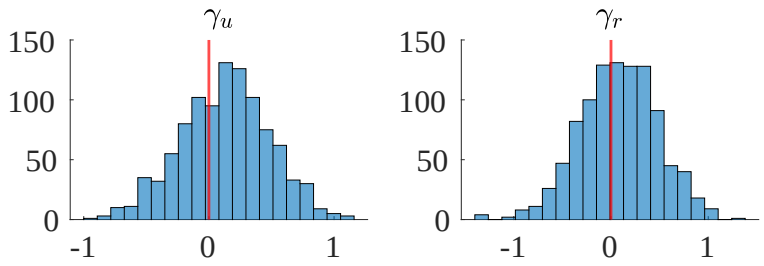
For the parameter γ_u, we set the starting value $\gamma_u^{(0)}$ to 0, thus coinciding with the value used for the calibration of the baseline specification. Similarly, we set the starting value $\gamma_r^{(0)}$ to 0. A derivation of the relationship in (10) from a conventional intertemporal IS would suggest a positive value of γ_r—corresponding to 1 for the case of a logarithmic utility function or to the intertemporal elasticity of substitution for the case of a utility function with a constant relative risk aversion. However, Lunsford and West (2019) study long-run correlations for the real interest rate and, in contrast to standard theory, find mixed results for the correlation with trend GDP growth depending on the considered sample period and the adopted method. Moreover, we assume that the proposal distribution for both parameters is standard normal and set the scale parameter to 0.3. For the initial conditions $\left\{\underline{\delta_{u,0}}, \underline{\delta_{r,0}}\right\}$, we use the same starting values and proposal distribution as for $\left\{\underline{u_0}, \underline{r_0}\right\}$ in Subsection 6.1 and we set c to 0.4 in equation (8).

Under the alternative specification proposed in this subsection, we modify the assumptions on the standard deviation of the shocks to the trends. Relative to the previous case, we keep unchanged the assumptions about the shocks to the trends with the exception of a decrease of the standard deviation for the change in the trend of annual growth rate of real GDP expected over the next 100 years from 2.6% to 1.6% because changes in this trend also affect movements in the trend of the unemployment rate and real interest rate.

As a result of these assumptions, the acceptance rate of the random walk Metropolis-Hastings algorithm is about 35%.

Results. As for the previous specifications, we report the results of estimating the TC-VAR

Figure 16: Histogram of γ_u and γ_r



The figure plots the histogram of the parameters that capture the relationship between the trend of real GDP growth and either the rate of unemployment γ_u (left panel) or the real interest rate γ_r (right panel) obtained from the estimation of the TC-VAR model with unrestricted initial conditions and cointegrating relationships over 1960-2019 period.

model over the 1960-2019 period and assuming two lags of the cyclical component in equation (3).[6]

Figure 16 plots the histogram of the parameters γ_u (left panel) and γ_r (right panel) as well as the corresponding values used for their calibration under the baseline specification (solid red lines). Moreover, Figure 17 plots the histogram of the initial conditions $\underline{\tau_0}$ for the latent trends as well as the corresponding values used for the calibration of those initial conditions under the baseline specification (solid red lines).

The figures shows that the calibrated values fall well within the range of estimates obtained when jointly estimating the parameters γ_u and γ_r, the initial conditions $\underline{\tau_0}$ and the VAR model with common trends.

Using the model estimated over the 1960-2019 period, we identify the shock targeting $\tilde{u}_t$ and report in Table 6 its contributions to the volatility of the cyclical components of the other variables.[7] Even in this case, the results are in line with those of the previous specifications considered.[8]

7 Conclusions

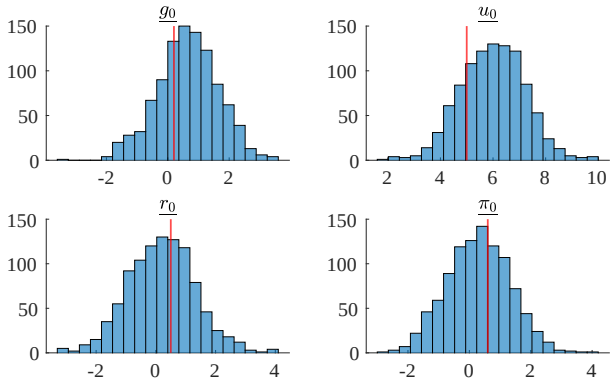
In recent years, a series of papers have called into question the validity of the New Keynesian framework. One important argument is the observation that inflation seems to be unresponsive to business cycle movements in the real activity. In this paper, we used a trend-cycle VAR model to study the relation between inflation and the real economy over the business

[6] The results for the 1960-2008 period are similar to those under the baseline specification.

[7] Appendix B reports the plots of the data as well as the latent trends and cycles obtained from the estimation of this specification and confirms the similarity of their dynamics to those found for the previous specifications.

[8] We draw similar conclusions when estimating the model under the alternative specification 2 over the 1960-2008 period and two lags.

Figure 17: Histogram of initial conditions $\underline{\tau_0}$



The figure plots the histogram of the initial conditions for the growth rate of real per-capita GDP $\underline{g_0}$ (top-left panel), unemployment rate $\underline{u_0}$ (top-right panel), real interest rate $\underline{r_0}$ (bottom-left panel) and inflation rate $\underline{\pi_0}$ (bottom-right panel) obtained from the estimation of the TC-VAR model with unrestricted initial conditions and cointegrating relationships over 1960-2019 period.

Table 6: Unrestricted initial conditions and cointegrating relationships: Shock targeting $\tilde{u}_t$

All frequencies ($0 - \infty$ quarters)

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
96.2	11.3	29.4	95.0	56.4	35.2
[94.3,97.5]	[1.1,36.3]	[16.9,45.2]	[92.5,96.8]	[36.5,72.0]	[12.9,60.7]

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

$\tilde{u}_t$	$\tilde{Y}_t$	$\tilde{\pi}_t$	$\tilde{R}_t$	$\tilde{\pi}_t^{e,4}$	$\tilde{\pi}_t^{e,40}$
98.0	11.3	36.2	96.9	69.8	56.3
[96.9,98.8]	[1.1,36.3]	[21.5,53.3]	[95.0,98.1]	[48.4,83.5]	[25.0,79.2]

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

References

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

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

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

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

Bianchi, F. and Melosi, L. (2017). Escaping the Great Recession. *American Economic Review*, 107(4):1030–58.

Carter, C. K. and Kohn, R. (1994). On Gibbs Sampling for State Space Models. *Biometrika*, 81(3):541–553.

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

Del Negro, M., Lenza, M., Primiceri, G. E., and Tambalotti, A. (2020). What's up with the Phillips Curve? NBER Working Papers 27003, National Bureay of Economic Research, Inc.

Durbin, J. and Koopman, S. J. (2001). *Time Series Analysis by State Space Methods*. Oxford University Press, second edition.

Giannone, D., Lenza, M., and Primiceri, G. E. (2015). Prior Selection for Vector Autoregressions. *Review of Economics and Statistics*, 97(2):436–51.

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.

Lunsford, K. G. and West, K. (2019). Some Evidence on Secular Drivers of US Safe Real Rates. *American Economic Journal: Macroeconomics*, 11(4):113–139.

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

A Gibbs Sampler

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

$$\Theta_\tau = \{\Sigma_\tau\}, \quad \Theta_c = \{\Phi_1, \Phi_2, \Sigma_c\}. \tag{11}$$

We use the Gibbs sampler to estimate the model unknowns. We rely on the state-space representation of (1) and (4). For the jth 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 A.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 A.2.

A.1 Kalman smoother

We rely on the state-space representation in equations (1) and (4). Conditional on the jth 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

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

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 $t = T-1, T-2, \ldots, 1$. The last element of the Kalman filter recursion provides the initialization for the simulation smoother:

$$\begin{aligned}
s_{t|t+1} &= s_{t|t} + P_{t|t}\Gamma' P_{t+1|t}^{-1}\left(s_{t+1} - \Gamma s_{t|t}\right) \\
P_{t|t+1} &= P_{t|t} - P_{t|t}\Gamma' P_{t+1|t}^{-1}\Gamma P_{t|t} \\
s_t^j &\sim N(s_{t|t+1}, P_{t|t+1}), \quad t = T-1, T-2, ..., 1.
\end{aligned} \tag{12}$$

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

A.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}.

A.2.1 VAR coefficients

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

$$c_t' = \underbrace{\begin{bmatrix} c_{t-1}' & c_{t-2}' \end{bmatrix}}_{w_t'} \underbrace{\begin{bmatrix} \Phi_1' \\ \Phi_2' \end{bmatrix}}_{\beta} + \epsilon_{c,t}', \quad \epsilon_{c,t} \sim N(\mathbf{0}, \Sigma_c). \tag{13}$$

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

$$\beta|\Sigma \sim MN\big(\underline{\beta}, \Sigma \otimes (\underline{V}_\beta \xi)\big), \quad \Sigma_c \sim IW(\underline{\Psi}, \underline{d}), \tag{14}$$

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

$$\beta|\Sigma \sim MN\big(\overline{\beta}, \Sigma \otimes \overline{V}_\beta\big), \quad \Sigma_c \sim IW(\overline{\Psi}, \overline{d}) \tag{15}$$

where

$$\begin{aligned} \overline{\beta} &= \big(W'W + (\underline{V}_\beta \xi)^{-1}\big)^{-1}\big(W'X + (\underline{V}_\beta \xi)^{-1}\underline{\beta}\big), \\ \overline{V}_\beta &= \big(W'W + (\underline{V}_\beta \xi)^{-1}\big)^{-1}, \\ \overline{\Psi} &= (X - W\overline{\beta})'(X - W\overline{\beta}) + (\overline{\beta} - \underline{\beta})'(\underline{V}_\beta \xi)^{-1}(\overline{\beta} - \underline{\beta}) + \underline{\Psi}, \\ \overline{d} &= T - 2 + \underline{d}. \end{aligned} \tag{16}$$

We follow the exposition in Giannone et al. (2015) in which ξ is a scalar parameter controlling the tightness of the prior information in (14). For instance, prior becomes more informative when $\xi \to 0$. In contrast, when $\xi = \infty$, then it is easy to see that $\overline{\beta} = \hat{\beta}$, i.e., an OLS estimate.

In sum, we draw β^{j+1} and Σ^{j+1} from (15). Hence, we obtain $\Theta_c^{j+1} = \{\Phi_1^{j+1}, \Phi_2^{j+1}, \Sigma_c^{j+1}\}$.

A.2.2 Trend component variances

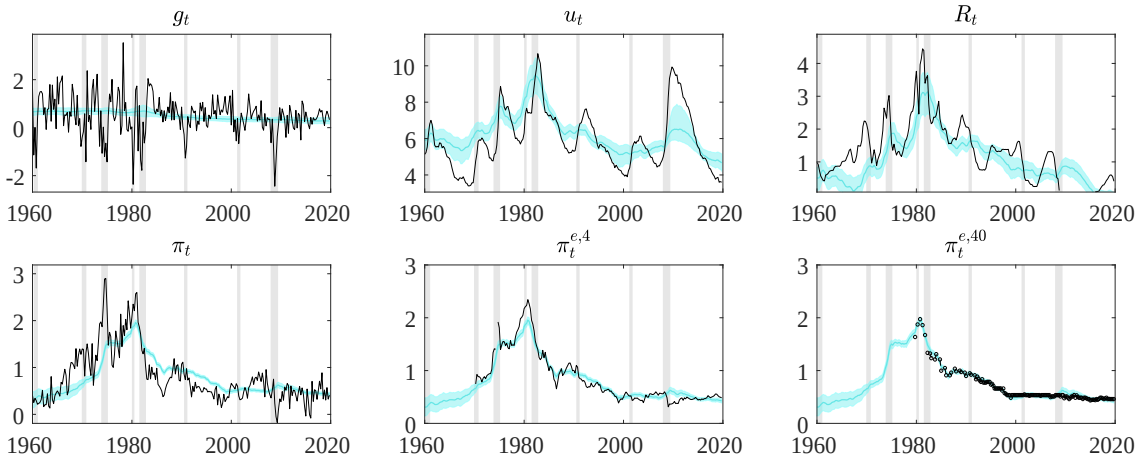
Conditional on $\tau_{1:T}^{j}$, the objective is to draw $\Theta_{\tau}^{j+1} = \{\Sigma_{\tau}^{j+1}\}$. For ease of exposition, we omit the superscript j below. Define $X = [\tau_2, ..., \tau_T]'$ and $W = [\tau_1, ..., \tau_{T-1}]'$. Similarly as before we draw from

$$\Sigma_{\tau} \sim IW(\overline{\Psi}, \overline{d}), \quad \overline{\Psi} = (X - W)'(X - W) + \underline{\Psi}, \quad \overline{d} = T - 1 + \underline{d}. \tag{17}$$

B TC-VAR model with unrestricted initial conditions and cointegrating relationships

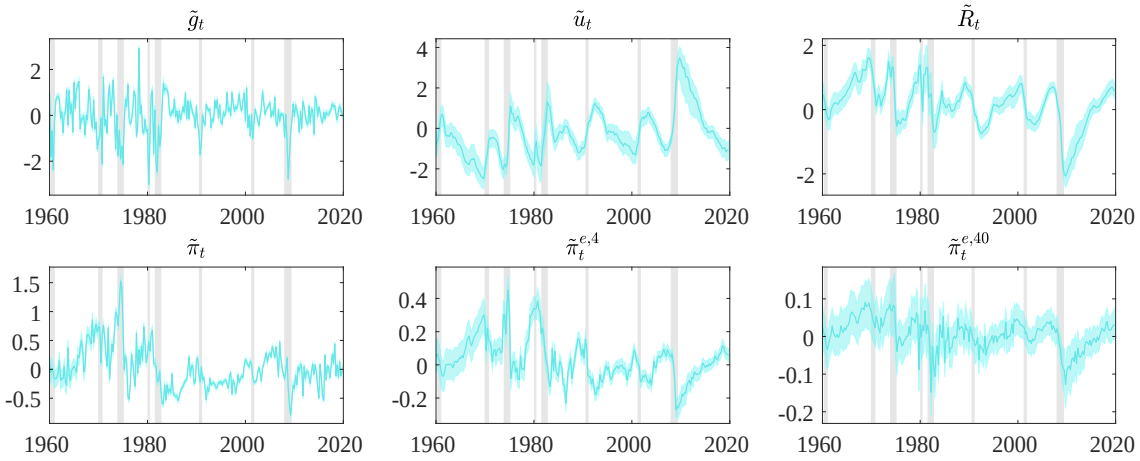
Figure 18 reports the data and the latent trends obtained from the estimation of this specification, and Figure 19 shows the associated cycles. The comparison of the latent trends and cycles obtained under this specification with those under the baseline specification in Section 5 and the specification in Subsection 6.2 confirms the similarity of their dynamics.

Figure 18: Unrestricted initial conditions and cointegrating relationships: Data and trends



The figure plots the data (solid black line) used for the estimation of the TC-VAR model over 1960-2019 period as well as the posterior median of their latent trends (solid blue line) and the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.

Figure 19: Unrestricted initial conditions and cointegrating relationships: Cycles



The figure plots the posterior median of the latent cycles (solid blue line) of the TC-VAR model estimated over 1960-2019 period along with the corresponding 68-percent posterior-coverage intervals (shaded blue area). NBER recessions are denoted by shaded grey areas.