

CAHIER D'ÉTUDES WORKING PAPER

N° 201

MEASURING BUSINESS CYCLES USING VARs

PATRICK FEVE

ALBAN MOURA

OCTOBER 2025



BANQUE CENTRALE DU LUXEMBOURG
EUROSYSTEME

MEASURING BUSINESS CYCLES USING VARS

PATRICK FÈVE AND ALBAN MOURA

ABSTRACT. We propose to measure business cycles using vector autoregressions (VARs). Our method builds on two insights: VARs automatically decompose the data into stable and unstable components, and variance-based shock identification can extract meaningful cycles from the stable part. This method has appealing properties: (1) it isolates a well-defined component associated with typical fluctuations; (2) it ensures stationarity by construction; (3) it targets movements at business-cycle frequencies; and (4) it is backward-looking, ensuring that cycles at each date only depend on current and past shocks. Since most existing filters lack one or more of these features, our method offers a valuable alternative. In an empirical application, we show that the two shocks with the largest cyclical impact effectively capture postwar U.S. business cycles and we find a tighter link between real activity and inflation than previously recognized. We compare our method with standard alternatives and document the plausibility and robustness of our results.

JEL Codes: C32, E32.

Keywords: business cycles, detrending, filtering, shocks, vector autoregressions.

August 2025. Patrick Fève: Toulouse School of Economics (TSE) and Université Toulouse 1 - Capitole, 1 esplanade de l'université, 31000 Toulouse, France (patrick.feve@tse-fr.eu). Alban Moura: Banque centrale du Luxembourg (BCL), Département Économie et Recherche, 2 boulevard Royal, L-2983 Luxembourg, Luxembourg (alban.moura@bcl.lu). This paper has been produced in the context of the partnership agreement between the BCL and TSE. We thank Fabrice Collard, Paolo Guarda, Olivier Pierrard, and BCL colleagues for helpful comments. Patrick Fève acknowledges funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EURE-0010, and under the project ANR-23-CE26-0011-01. This paper should not be reported as representing the views of the BCL or the Eurosystem. The views expressed are those of the authors and may not be shared by other research staff or policymakers in the BCL or the Eurosystem.

RÉSUMÉ NON TECHNIQUE

L'analyse des cycles économiques est essentielle pour les banques centrales, qui stabilisent l'inflation en ajustant leur politique monétaire en fonction des fluctuations de l'activité. Pourtant, il n'existe pas de consensus sur la manière de mesurer ces cycles. Le principal défi tient à la multiplicité des méthodes disponibles, chacune produisant des résultats potentiellement divergents. Cette hétérogénéité rend difficile l'obtention d'une lecture cohérente des cycles économiques.

Ce travail propose une nouvelle méthode de mesure du cycle, fondée sur un modèle VAR décrivant conjointement la dynamique de variables macroéconomiques. La démarche s'articule en deux temps : d'abord elle isole et élimine les composantes instables des données, puis elle identifie les chocs dont les effets ont une durée proche de celle conventionnellement attribuée aux cycles économiques. Notre mesure du cycle est définie par les contributions de ces chocs aux fluctuations passées des variables observées.

La méthode présente quatre atouts principaux. Premièrement, elle ancre la décomposition tendance-cycle dans les propriétés statistiques des données. Deuxièmement, les cycles obtenus reviennent régulièrement vers une moyenne constante, ce qui permet de capter des écarts transitoires par rapport aux tendances de long terme. Troisièmement, elle cible les fluctuations dont la durée est typique des cycles économiques, tout en préservant la structure temporelle des données. Quatrièmement, elle met en évidence des fluctuations typiques, facilitant ainsi l'identification d'épisodes atypiques.

Appliquée aux données des États-Unis, la méthode montre qu'il suffit de deux chocs pour rendre compte des principales caractéristiques des cycles économiques. Les cycles estimés coïncident étroitement avec la chronologie officielle des récessions et expansions, et reconnaissent la Grande Récession comme la plus marquée. En outre, les résultats suggèrent que le lien entre inflation et activité réelle est plus étroit que généralement admis. Enfin, les tests de robustesse confirment la fiabilité de notre méthode.

1. INTRODUCTION

Economists disagree on the measurement of business cycles, the alternating periods of expansion and recession observed in broad measures of economic activity. The problem lies in the infinite number of possible transformations of any given series that can be applied to extract a cycle.

The literature has explored various approaches, each reflecting a different perspective on the nature of trends and cycles in the data. Some studies assume deterministic time trends; others focus on growth rates instead of the (transformed) level of the series. Popular statistical methods impose restrictions on the smoothness of the trend component, as in the HP filter (Hodrick and Prescott, 1981, 1997), or on the frequency range where the power of the cyclical component is concentrated, as in bandpass filters (Baxter and King, 1999; Christiano and Fitzgerald, 2003). Other approaches rely on parametric time-series models, either unobserved-components models (UCMs; Harvey, 1985; Watson, 1986; Harvey and Jaeger, 1993; Harvey, Trimbur, and Van Dijk, 2007) or models that exploit the relationship between trends and medium- to long-term forecasts (Beveridge and Nelson, 1981; Hamilton, 2018).

This heterogeneity would not matter if different methods yielded a consistent picture of business cycles. However, Canova (1998, 1999) shows that they do not: both quantitatively and qualitatively, the cyclical properties of U.S. data vary significantly across detrending methods. Moreover, as discussed below, existing approaches have substantial limitations.

In this paper, we propose a novel method with attractive properties. Our analysis begins with a vector autoregression (VAR) characterizing the joint dynamics of key macroeconomic variables. The method proceeds in two steps. In the *detrending* step, we decompose the VAR into stable and unstable (unit-root-like) components based on the eigenvalues of its companion form, following Casals, Jerez, and Sotoca (2002), and retain only the stable component. In the *filtering* step, we rotate the VAR residuals to extract orthogonal shocks ranked by their contribution to selected variables at business-cycle frequencies, as in Angeletos, Collard, and Dellas (2020); by construction, the top-ranked shocks define the cyclical behavior of the target variables. Our main contribution is to show that the historical contributions of these primary business-cycle shocks yield cyclical estimates with desirable properties.

We highlight four key features of our method.

First, the VAR should reasonably approximate the data-generating process (DGP), ensuring that the cycles reflect genuine properties of the data. In contrast, simple time trends, bandpass filters, and the HP filter often produce cycles misaligned with the DGP (Nelson and Kang, 1981; Cogley and Nason, 1995; Hamilton, 2018), and even fully specified UCMs may yield spurious results (Nelson, 1988). An unrestricted VAR in levels also avoids difficult modeling choices regarding unit roots and cointegration, as these properties emerge

automatically in long samples (Sims, Stock, and Watson, 1990; Hamilton, 1994); deterministic trends can likewise be accommodated easily. In addition, a VAR greatly simplifies specification and estimation relative to multivariate UCMs, and our Bayesian implementation provides straightforward measures of statistical uncertainty, which are often lacking in alternative approaches.

The use of VARs to distinguish between permanent and transitory dynamics has a long tradition (e.g., Cochrane and Sbordone, 1988; Blanchard and Quah, 1989; King, Plosser, Stock, and Watson, 1991; Cochrane, 1994; Coibion, Gorodnichenko, and Ulate, 2018). We build on the key insight that multivariate analysis can reveal stationary components that remain undetected in univariate settings, but we adopt a more flexible specification that does not impose exact unit roots or cointegration on the data. The resulting trend-cycle decomposition removes random-walk components from integrated variables while preserving the transitory dynamics captured by the classic Beveridge and Nelson (1981) filter. In addition, the multivariate structure ensures cross-variable coherence, as all cyclical dynamics originate from a common set of shocks.

Second, our cycles are stationary by construction, an essential property for variables intended to capture transitory deviations from long-run dynamics. This marks an improvement over alternatives like the HP filter, which may fail to ensure stationarity in practice (Phillips and Jin, 2021). Our eigenvalue-based detrending also improves on the treatment of unit roots in Angeletos, Collard, and Dellas (2020), who address instability by truncating the posterior distribution, discarding draws that imply explosive behavior. Instead, our strategy retains the full posterior and accommodates exact unit roots. More importantly, we will show that it produces more credible cyclical estimates than the truncation approach, which fails to eliminate trend components in our empirical application.

Third, unlike approaches such as the Beveridge-Nelson decomposition or the Hamilton filter, our cycles conform to the well-established view that business-cycle fluctuations occur at specific periodicities (Burns and Mitchell, 1946; Stock and Watson, 1999). Crucially, our method targets these frequencies while producing backward-looking estimates that depend only on current and past shocks, thereby preserving the temporal structure of the data. This sets it apart from standard two-sided filters like the bandpass or HP filters, which achieve frequency selection through non-causal smoothing that relies on future observations, an approach that can introduce boundary distortions and generate spurious cyclical patterns (Cogley and Nason, 1995; Hamilton, 2018). Our estimates also avoid the distortions introduced by the Hamilton filter, which, by defining cycles as forecast errors and mechanically delaying trend adjustments, often produces implausible decompositions, such as rising trends during recessions and falling trends at the onset of recoveries (Moura, 2024).

Fourth, our method captures ‘*business cycles as usual*,’ that is, the typical pattern of expansions and recessions caused by the shocks that determine the cyclical behavior of macroeconomic variables over the full sample. This focus reveals the number of relevant shocks and yields summary business-cycle statistics that avoid conflating distinct disturbances across variables and time, providing a valuable input for structural modeling that standard filters cannot offer. It also establishes a benchmark for historical comparison: by defining typical cycles, our method helps identify and interpret deviations from regular patterns, offering new insights into the nature of specific episodes.

To summarize, the VAR cycles we propose isolate a stationary and well-defined component of the data, aligned with the conventional notion of business cycles and preserving the information structure. These properties are widely regarded as desirable for a cyclical estimate. Since most existing alternatives lack one or more of these features, our method represents a valuable addition to the empirical macroeconomist’s toolbox.

Of course, our method has limitations. In particular, it assumes that the VAR provides a good approximation of the DGP, an assumption open to debate. For example, Harvey and Jaeger (1993) argue that autoregressive models offer a poor representation of cyclical dynamics relative to UCMs, while Ravenna (2007) and Chari, Kehoe, and McGrattan (2008) show that VARs may approximate dynamic stochastic general-equilibrium models poorly when there are hidden endogenous state variables. Relatedly, Lippi and Reichlin (1993, 1994) and Fernández-Villaverde, Rubio-Ramírez, Sargent, and Watson (2007) highlight that VARs may be non-fundamental, meaning that residuals need not correspond to true innovations, raising the possibility that our VAR cycles could reflect spurious shocks. Finally, the assumption of constant coefficients may also be contestable (McConnell and Perez-Quiros, 2000; Primiceri, 2005). Ultimately, the adequacy of the VAR is an empirical matter that can be assessed through specification tests, while non-fundamentality concerns can be addressed by expanding the information set with macroeconomic factors to capture unobserved states (Forni and Gambetti, 2014; Beaudry, Fève, Guay, and Portier, 2019).

Our focus on typical cycles may also limit applications aimed at analyzing the specific features of individual historical episodes. Still, our typical cycle should remain a useful benchmark to understand how particular periods diverged from standard patterns. Another potential issue is that our multivariate estimates depend on the set of included variables. While this could be problematic if results were highly sensitive to VAR specification, we find that the cyclical estimates remain robust when the analysis includes core U.S. macroeconomic variables.

Indeed, in the second part of the paper, we apply our method to postwar U.S. business cycles. We estimate a BVAR including key macroeconomic variables — gross domestic product, consumption, investment, hours worked, the unemployment rate, the labor share of income, inflation, the nominal interest rate, total factor productivity, and credit spreads

— and identify shocks based on their contribution within the 6-to-32-quarter frequency band. We find that two shocks suffice to capture the main features of business cycles and document a close alignment between the resulting VAR cycles and the NBER chronology of expansions and recessions. The estimated cycles reveal that inflation is strongly procyclical conditional on the two shocks, indicating a tight link between price pressures and real activity consistent with inflationary demand shocks. Compared to standard alternatives, our method is the only one that identifies the Great Recession as deeper than previous contractions and the subsequent recovery as slower than earlier ones, in line with common wisdom (Hall, 2011; Galí, Smets, and Wouters, 2012). Finally, we establish the robustness of our results by varying the VAR order, the treatment of low-frequency movements, the set of included variables, and the sample period. We also discuss the performance of our method during the unusual COVID recession.

We organize the paper as follows. Section 2 presents our VAR method and discusses its theoretical properties. Section 3 applies the method to U.S. business cycles, detailing the VAR specification, the resulting cycles, and the selection of the number of retained shocks. Section 4 documents the robustness of our results and extends the analysis to the exceptional COVID episode. Section 5 concludes.

2. MEASURE

This section outlines the construction of our VAR cycles, which proceeds in two steps. First, we decompose the data into stable and unstable components to eliminate unit roots. Second, we extract meaningful cycles from the stable component by identifying shocks with a large cyclical footprint and discarding the rest.

2.1. VAR representation. Our starting point is the $n \times 1$ vector $\mathbf{x}_t$, which stacks the macroeconomic time series of interest. We assume that a finite-order VAR model provides a reasonable approximation of the DGP, so that $\mathbf{x}_t$ evolves according to

$$\mathbf{A}(L) \mathbf{x}_t = \mathbf{u}_t, \quad (1)$$

where $\mathbf{A}(L) = \mathbf{I}_n - \sum_{j=1}^p \mathbf{A}_j L^j$ is a matrix polynomial in the lag operator L , p is the lag order, and $\mathbf{u}_t$ is an $n \times 1$ vector of residuals satisfying $E(\mathbf{u}_t \mathbf{u}_t') = \Sigma$, with Σ a positive-definite $n \times n$ matrix. Without loss of generality, we omit deterministic terms from the presentation. These can be included in estimation without affecting the analysis.

We specify the VAR in (log) levels for two key reasons. First, this aligns with our goal of extracting cycles from the raw data: including GDP in levels yields GDP cycles, whereas including GDP in growth rates would produce cycles in the growth rate, which are less natural objects for business-cycle analysis. An exception may apply to price variables, which can enter in levels when focusing on price cycles, or in growth rates when focusing on inflation cycles. Second, the levels specification offers greater robustness to alternative trend patterns,

whether deterministic (if included in the VAR), stochastic, or involving cointegration (Sims, Stock, and Watson, 1990; Hamilton, 1994).¹ This robustness is particularly valuable for mitigating specification concerns.²

2.2. Detrending step. Potential deterministic trends are removed at the estimation stage. The detrending step decomposes the VAR into stable and unstable (unit-root-like) components, eliminating the latter. This ensures that the estimated cycles, derived from the stable component, are stationary by construction.

This step is essential when $\mathbf{x}_t$ includes variables such as aggregate quantities or relative prices, which may inherit stochastic trends from permanent technology shocks (King, Plosser, Stock, and Watson, 1991; Fisher, 2006). More generally, macroeconomic time series tend to be highly persistent (Nelson and Plosser, 1982), and even supposedly stationary series, such as per-capita hours worked, inflation, and nominal interest rates, display pronounced low-frequency fluctuations (Francis and Ramey, 2009; Uribe, 2022). The detrending step eliminates these persistent movements, which lack the mean-reverting properties of genuine cycles.

One shortcut, used by Angeletos, Collard, and Dellas (2020), estimates the VAR with Bayesian methods and discards unstable posterior draws. This approach has three drawbacks. From a theoretical standpoint, *ad hoc* truncations of the posterior are inconsistent with Bayesian principles, which encode strong views about parameters in the prior. Discarding unstable draws is particularly troubling in this context, as Angeletos, Collard, and Dellas (2020) employ a standard Minnesota prior centered on an (unstable) multivariate random walk. From an empirical perspective, the approach lacks generality. In our empirical application, up to 95% of posterior draws are unstable when using prior distributions that emphasize unit roots and cointegration, as in Sims and Zha (1998) and Giannone, Lenza, and Primiceri (2019); discarding most of the posterior and analyzing only a narrow subset is clearly problematic. The issue also arises in a frequentist setting if the point estimate lies in the unstable region. Finally, and perhaps more importantly, if the data contain at least one stochastic trend, focusing on stable draws misclassifies the permanent component as transitory, thus failing to properly detrend the series. As we will show, this can lead to less credible estimates of the cyclical component.

¹In the presence of cointegration, the VAR admits an error-correction representation (Engle and Granger, 1987). We prefer the more general VAR in levels because it is difficult to determine which linear combinations of the variables are stationary.

²As noted in the Introduction, several authors have questioned the empirical usefulness of VARs. In our view, any parametric model necessarily entails strong assumptions, and the generality and flexibility of VARs are notable advantages. Economists with strong prior beliefs in favor of alternative models may still find our approach of interest, as it can be applied directly to any state-space representation, including linear(-ized) DSGE models.

To overcome these limitations, we propose an alternative strategy to eliminate unstable components, drawing on the eigenvalue decomposition of state-space models proposed by Casals, Jerez, and Sotoca (2002). Specifically, we first express the VAR in its companion representation,

$$\mathbf{X}_t = \mathbf{F}\mathbf{X}_{t-1} + \mathbf{G}\mathbf{u}_t, \quad \mathbf{x}_t = \mathbf{H}\mathbf{X}_t, \quad (2)$$

where

$$\underbrace{\mathbf{X}_t}_{np \times 1} = \begin{bmatrix} \mathbf{x}_t \\ \dots \\ \mathbf{x}_{t-p+1} \end{bmatrix}, \quad \underbrace{\mathbf{F}}_{np \times np} = \begin{bmatrix} \mathbf{A}_1 & \dots & \mathbf{A}_{p-1} & \mathbf{A}_p \\ & \mathbf{I}_{n(p-1)} & & \mathbf{0}_{n(p-1) \times n} \end{bmatrix}, \quad \underbrace{\mathbf{G}}_{np \times n} = \begin{bmatrix} \mathbf{I}_n \\ \mathbf{0}_{n(p-1) \times n} \end{bmatrix}, \quad \underbrace{\mathbf{H}}_{n \times np} = \mathbf{G}',$$

with $\mathbf{I}_z$ the identity matrix of size z .

Next, we compute the eigenvalues of $\mathbf{F}$ to assess the stability of $\mathbf{x}_t$. If all np eigenvalues have modulus strictly below one, the process is stationary and no unstable root needs to be removed. In this case, we proceed directly to Section 2.3.

If some eigenvalues have modulus equal to or greater than one, we decompose $\mathbf{F}$ to separate the stable and unstable components of $\mathbf{x}_t$. Suppose there are n_u unstable eigenvalues and $n_s = np - n_u$ stable ones. Assuming for simplicity that $\mathbf{F}$ is diagonalizable over the reals, there exists an invertible real matrix $\mathbf{U}$ such that

$$\mathbf{T} = \mathbf{U}^{-1}\mathbf{F}\mathbf{U} = \begin{bmatrix} \mathbf{T}_u & \mathbf{0}_{n_u \times n_s} \\ \mathbf{0}_{n_s \times n_u} & \mathbf{T}_s \end{bmatrix},$$

where $\mathbf{T}_u$ is an $n_u \times n_u$ diagonal matrix of unstable eigenvalues and $\mathbf{T}_s$ is an $n_s \times n_s$ diagonal matrix of stable eigenvalues. This decomposition is unique up to the ordering of eigenvalues within each block, which does not affect the results.³

Defining $\mathbf{Y}_t = \mathbf{U}^{-1}\mathbf{X}_t$, $\mathbf{W} = \mathbf{U}^{-1}\mathbf{G}$, and $\mathbf{J} = \mathbf{H}\mathbf{U}$, and partitioning the matrices and vectors according to the blocks in $\mathbf{T}$, we obtain the diagonal representation,

$$\begin{bmatrix} \mathbf{Y}_{u,t} \\ \mathbf{Y}_{s,t} \end{bmatrix} = \begin{bmatrix} \mathbf{T}_u & \mathbf{0}_{n_u \times n_s} \\ \mathbf{0}_{n_s \times n_u} & \mathbf{T}_s \end{bmatrix} \begin{bmatrix} \mathbf{Y}_{u,t-1} \\ \mathbf{Y}_{s,t-1} \end{bmatrix} + \begin{bmatrix} \mathbf{W}_u \\ \mathbf{W}_s \end{bmatrix} \mathbf{u}_t, \quad \mathbf{x}_t = \begin{bmatrix} \mathbf{J}_u & \mathbf{J}_s \end{bmatrix} \begin{bmatrix} \mathbf{Y}_{u,t} \\ \mathbf{Y}_{s,t} \end{bmatrix}. \quad (3)$$

This system decomposes $\mathbf{x}_t$ into a non-stationary component, $\mathbf{J}_u \mathbf{Y}_{u,t}$, and a stationary component, $\mathbf{J}_s \mathbf{Y}_{s,t}$. Our detrending step removes the non-stationary component and retains the stationary one, denoted $\mathbf{x}_{s,t}$:

$$\mathbf{Y}_{s,t} = \mathbf{T}_s \mathbf{Y}_{s,t-1} + \mathbf{W}_s \mathbf{u}_t, \quad \mathbf{x}_{s,t} = \mathbf{J}_s \mathbf{Y}_{s,t}. \quad (4)$$

Crucially, when at least one eigenvalue is stable ($n_s > 0$), the resulting stable component $\mathbf{x}_{s,t}$ remains n -dimensional, matching the dimension of the original vector $\mathbf{x}_t$. This confirms that our method removes permanent components, not variables. Conversely, in the unlikely case

³If $\mathbf{F}$ is defective (not diagonalizable over the reals), the analysis carries through by replacing standard diagonalization with the real block diagonalization described in Golub and Van Loan (1996, Theorem 7.6.3). For robustness, our code implements this more general case.

where all eigenvalues are unstable, it would correctly conclude that no stable component exists.

This detrending step builds on the common-trends representation from the cointegration literature, which decomposes an n -dimensional vector of $I(1)$ variables into n_u common random walks and $n - n_u$ stationary combinations (Stock and Watson, 1988; Cochrane and Sbordone, 1988). When the unstable eigenvalues of $\mathbf{F}$ correspond to exact unit roots, our procedure recovers the standard common-trends representation and yields a multivariate extension of the Beveridge and Nelson (1981) decomposition, eliminating the permanent component $\mathbf{J}_u \mathbf{Y}_{u,t}$ and retaining the stationary component $\mathbf{J}_s \mathbf{Y}_{s,t}$. Our detrending step allows greater flexibility: instead of imposing exact unit roots or requiring *a priori* knowledge of cointegration relationships, we let the data determine persistence and classify the components through estimated eigenvalues.

The well-known small-sample downward bias in autoregressive coefficient estimates motivates setting a stability threshold below unity. Following Casals, Jerez, and Sotoca (2002), we classify eigenvalues with modulus above $\bar{\rho} = 0.99$ as unstable. This cutoff reflects a trade-off: it buffers against small-sample bias that can make unit roots appear stable, while avoiding the excessive removal of persistent yet genuinely stationary dynamics. It also sharpens the focus on shorter-term cycles by filtering out very low-frequency movements. As shown in our empirical analysis, the cyclical estimates are robust to reasonable variations in $\bar{\rho}$.

2.3. Filtering step. The detrending step produces a genuine stationary component. One option is to stop here and interpret $\mathbf{x}_{s,t}$ as a multivariate cycle, as in Casals, Jerez, and Sotoca (2002). However, not all stationary dynamics correspond to conventional business cycles: some may be too short-lived, others too persistent, and some may display patterns inconsistent with the typical comovement of macroeconomic variables. To refine our estimate, we introduce an additional filtering step that enhances its cyclical properties.

Specifically, following Faust (1998), Uhlig (2004), and Angeletos, Collard, and Dellas (2020), we determine the orthogonal rotation of the residual vector $\mathbf{u}_t$ in equation (4) that ranks shocks by their contributions to the variance of selected variables at business-cycle frequencies. The top-ranked shocks under this criterion shape the cyclical properties of the target variables and thus capture the main drivers of macroeconomic fluctuations, linking statistical prominence to economic relevance. Retaining these shocks while discarding the others provides an effective filter, isolating meaningful cyclical components and attenuating movements unrelated to aggregate fluctuations. Our multivariate method is essential here, as univariate autoregressions with a single residual cannot recover distinct shocks.

Formally, we express $\mathbf{u}_t$ as a linear transformation of an $n \times 1$ vector of orthogonal shocks $\boldsymbol{\epsilon}_t$,

$$\mathbf{u}_t = \mathbf{S}\boldsymbol{\epsilon}_t,$$

with $\mathbf{S}$ an invertible $n \times n$ matrix and $E(\boldsymbol{\epsilon}_t \boldsymbol{\epsilon}_t') = \mathbf{I}_n$.⁴ This mapping imposes only that $\mathbf{S}\mathbf{S}' = \boldsymbol{\Sigma}$, which holds for infinitely many matrices. Any valid choice must satisfy $\mathbf{S} = \boldsymbol{\Sigma}^{1/2}\mathbf{Q}$, with $\boldsymbol{\Sigma}^{1/2}$ the lower Cholesky factor of $\boldsymbol{\Sigma}$ and $\mathbf{Q}$ an $n \times n$ orthonormal matrix. Substituting this into equation (4), we obtain the Wold representation of $\mathbf{x}_{s,t}$,

$$\mathbf{x}_{s,t} = \mathbf{J}_s (\mathbf{I}_{n_s} - \mathbf{T}_s L)^{-1} \mathbf{W}_s \boldsymbol{\Sigma}^{1/2} \mathbf{Q} \boldsymbol{\epsilon}_t = \boldsymbol{\Gamma}(L) \mathbf{Q} \boldsymbol{\epsilon}_t, \quad (5)$$

with $\boldsymbol{\Gamma}(L) = \mathbf{J}_s (\mathbf{I}_{n_s} - \mathbf{T}_s L)^{-1} \mathbf{W}_s \boldsymbol{\Sigma}^{1/2}$. The stability of $\mathbf{T}_s$ is crucial for ensuring the validity of this representation, highlighting the importance of the detrending step.

2.3.1. Single target variable. To set the stage, we begin with the case in which the ranking criterion relies on a single target variable. The objective is to rank the shocks by their contributions to the variance of the j th element of $\mathbf{x}_{s,t}$ over a specified frequency band $[\underline{\omega}, \bar{\omega}]$.

The population spectrum of $\mathbf{x}_{s,t}$, derived from equation (5), satisfies

$$\mathbf{S}_{\mathbf{x}_s}(\omega) = (2\pi)^{-1} \boldsymbol{\Gamma}(e^{-i\omega}) \boldsymbol{\Gamma}(e^{i\omega})',$$

where i is the imaginary number and ω is a real scalar. Using the identity $\mathbf{I}_n = \mathbf{Q} (\sum_{k=1}^n \mathbf{E}_k) \mathbf{Q}'$, where $\mathbf{E}_k = \mathbf{e}_k \mathbf{e}_k'$ and $\mathbf{e}_k$ is the n -dimensional column vector with a one in the k th entry and zeros elsewhere, we decompose the spectrum into orthogonal contributions from the n shocks in $\boldsymbol{\epsilon}_t$,

$$\mathbf{S}_{\mathbf{x}_s}(\omega) = (2\pi)^{-1} \left[\sum_{k=1}^n \boldsymbol{\Gamma}(e^{-i\omega}) \mathbf{Q} \mathbf{E}_k \mathbf{Q}' \boldsymbol{\Gamma}(e^{i\omega})' \right].$$

The contribution of the k th shock in $\boldsymbol{\epsilon}_t$ to the spectrum of the j th variable in $\mathbf{x}_{s,t}$ follows as

$$(2\pi)^{-1} \mathbf{e}_j' [\boldsymbol{\Gamma}(e^{-i\omega}) \mathbf{Q} \mathbf{E}_k \mathbf{Q}' \boldsymbol{\Gamma}(e^{i\omega})'] \mathbf{e}_j = \mathbf{q}_k' [(2\pi)^{-1} \boldsymbol{\Gamma}(e^{i\omega})' \mathbf{E}_j \boldsymbol{\Gamma}(e^{-i\omega})] \mathbf{q}_k,$$

where $\mathbf{q}_k$ is the k th column vector of $\mathbf{Q}$. Integrating this expression over the frequency band $[\underline{\omega}, \bar{\omega}]$ gives the contribution of the k th shock to the variance of the j th variable in $\mathbf{x}_{s,t}$, denoted

$$\gamma_{k,j}(\underline{\omega}, \bar{\omega}) = \mathbf{q}_k' \left[(2\pi)^{-1} \int_{\underline{\omega}}^{\bar{\omega}} \boldsymbol{\Gamma}(e^{i\omega})' \mathbf{E}_j \boldsymbol{\Gamma}(e^{-i\omega}) d\omega \right] \mathbf{q}_k = \mathbf{q}_k' \boldsymbol{\Theta}_j(\underline{\omega}, \bar{\omega}) \mathbf{q}_k,$$

with

$$\boldsymbol{\Theta}_j(\underline{\omega}, \bar{\omega}) = \left[(2\pi)^{-1} \int_{\underline{\omega}}^{\bar{\omega}} \boldsymbol{\Gamma}(e^{i\omega})' \mathbf{E}_j \boldsymbol{\Gamma}(e^{-i\omega}) d\omega \right].$$

Since $\boldsymbol{\Theta}_j(\underline{\omega}, \bar{\omega})$ is real-valued, $\gamma_{k,j}(\underline{\omega}, \bar{\omega})$ forms a well-defined symmetric quadratic form in $\mathbf{q}_k$.

Ranking the shocks by their variance contributions requires finding n orthogonal unit-length vectors, $\mathbf{q}_1, \dots, \mathbf{q}_n$, such that $\gamma_{1,j}(\underline{\omega}, \bar{\omega}) \geq \gamma_{2,j}(\underline{\omega}, \bar{\omega}) \geq \dots \geq \gamma_{n,j}(\underline{\omega}, \bar{\omega})$. This reduces

⁴Although we refer to the entries of $\boldsymbol{\epsilon}_t$ as ‘shocks,’ these are statistical rotations of the VAR residuals rather than structural economic disturbances.

to a standard eigenvalue problem, implying that $\mathbf{q}_k$ are the eigenvectors of $\Theta_j(\underline{\omega}, \bar{\omega})$ ordered by their corresponding eigenvalues in descending order.⁵

2.3.2. Several target variables. We extend the analysis to cases in which the ranking procedure relies on multiple variables. Multivariate restrictions are particularly appealing because they require shocks to explain the joint dynamics of several series, thereby inducing comovement at target frequencies. This aligns with Lucas's (1977) view that short-term macroeconomic comovement is a defining feature of business cycles. Moreover, a long literature argues that fluctuations across key macroeconomic aggregates are largely driven by a small number of common forces (Sargent and Sims, 1977; Stock and Watson, 2016). This strengthens the case for multivariate targeting: if business cycles reflect such shared drivers, then requiring shocks to account for comovement should sharpen identification.

In the absence of a standard multivariate variance measure, we adopt a simple and transparent benchmark: the average variance of the individual target variables.⁶ Target selection should be guided by the researcher's interest. For instance, studies of standard business cycles would typically target output, investment, consumption, hours worked, and unemployment, while analyses of financial cycles might instead target credit aggregates, interest-rate spreads, and asset prices.

Let $\mathcal{T}$ denote the subset of target variables in $\mathbf{x}_{s,t}$, with $n_{\mathcal{T}}$ as its cardinal. Building on earlier results, the impact of the k th shock in $\boldsymbol{\epsilon}_t$ on the average variance of the variables in $\mathcal{T}$ is

$$\bar{\gamma}_{k\mathcal{T}}(\underline{\omega}, \bar{\omega}) = n_{\mathcal{T}}^{-1} \sum_{j \in \mathcal{T}} \gamma_{kj}(\underline{\omega}, \bar{\omega}).$$

Equivalently,

$$\bar{\gamma}_{k\mathcal{T}}(\underline{\omega}, \bar{\omega}) = \mathbf{q}'_k \mathbf{V}(\underline{\omega}, \bar{\omega}, \mathcal{T}) \mathbf{q}_k, \quad \text{with } \mathbf{V}(\underline{\omega}, \bar{\omega}, \mathcal{T}) = n_{\mathcal{T}}^{-1} \sum_{j \in \mathcal{T}} \Theta_j(\underline{\omega}, \bar{\omega}). \quad (6)$$

⁵In practice, computing $\Theta_j(\underline{\omega}, \bar{\omega})$ is necessary to solve the maximization problem. A simple quadrature rule is effective. Define

$$\Theta_j(\underline{\omega}, \bar{\omega}) = (2\pi)^{-1} \left[\int_{\underline{\omega}}^{\bar{\omega}} \mathbf{G}_j(\omega) d\omega \right], \quad \text{with } \mathbf{G}_j(\omega) = \boldsymbol{\Gamma}(e^{i\omega})' \mathbf{E}_j \boldsymbol{\Gamma}(e^{-i\omega}).$$

For N large, an accurate approximation is

$$\hat{\Theta}_j = \frac{\bar{\omega} - \underline{\omega}}{2\pi N} \left[\frac{\mathbf{G}_j(\underline{\omega}) + \mathbf{G}_j(\bar{\omega})}{2} + \sum_{z=1}^{N-1} \mathbf{G}_j \left(\underline{\omega} + \frac{z(\bar{\omega} - \underline{\omega})}{N} \right) \right].$$

⁶We also experimented with a weighted target that scales each variable's contribution by its unconditional variance over the relevant frequency band. This normalization gives less weight to more volatile variables, which may or may not be desirable depending on context. We will show that results are identical across weighted and unweighted targets in our application to U.S. business cycles, confirming the robustness of our identification to this choice.

We solve the ranking problem as before, ordering shocks by their variance contributions using the eigenvalues and eigenvectors of $\mathbf{V}(\underline{\omega}, \bar{\omega}, \mathcal{T})$.

2.4. VAR cycles and interpretation. Our cyclical estimate is the historical contribution of the first shock(s) identified in the filtering step to the stationary component $\mathbf{x}_{s,t}$. The construction follows standard methods (Kilian and Lutkepohl, 2017, Chapter 4).

In population, the contribution of the first K shocks to $\mathbf{x}_{s,t}$ is

$$\mathbf{x}_{s,t}(\mathbf{q}_1, \dots, \mathbf{q}_K) = \sum_{j=0}^{\infty} \sum_{k=1}^K \mathbf{\Gamma}_j \mathbf{q}_k \epsilon_{k,t-j}.$$

Given the estimated VAR, the detrending and filtering steps provide estimates of the $\mathbf{\Gamma}$'s and $\mathbf{q}$'s. The shocks can also be estimated from the VAR residuals:

$$\hat{\epsilon}_{k,t} = \hat{\mathbf{q}}_k' \hat{\Sigma}^{-1/2} \hat{\mathbf{u}}_t, \quad k = 1, \dots, K, \quad t = 1, \dots, T,$$

where hats denote estimates and T is the sample size. Thus, the estimated contribution of the shocks to historical movements in $\mathbf{x}_{s,t}$ follows as

$$\hat{\mathbf{x}}_{s,t}(\hat{\mathbf{q}}_1, \dots, \hat{\mathbf{q}}_K) = \sum_{j=0}^{t-1} \sum_{k=1}^K \hat{\mathbf{\Gamma}}_j \hat{\mathbf{q}}_k \hat{\epsilon}_{k,t-j}, \quad t = 1, \dots, T. \quad (7)$$

Equation (7) defines our VAR cycles. Comparing it with its population counterpart reveals that it omits the effects of pre-sample shocks, which cannot be estimated. Since our procedure recovers shocks that generate stable dynamics, the estimated coefficients $\hat{\mathbf{\Gamma}}_j \hat{\mathbf{q}}_k$ decline as we move further in the past, reducing the influence of previous shocks as they become more distant in time. The choice of an eigenvalue stability threshold $\bar{\rho}$ slightly below one reinforces this decay. Therefore, apart from a small number of initial observations, we expect our estimate to closely approximate its population counterpart.

How should these cycles be interpreted? From an econometric standpoint, the answer is straightforward: they represent a well-defined stationary component of the VAR, driven by orthogonal rotations of the residuals with the largest impact at business-cycle frequencies. More intuitively, our VAR cycles capture the common fluctuations across variables most closely associated with the timing, persistence, and comovement patterns characteristic of business cycles. They isolate movements that are both temporary (stationary) and cyclical, while excluding permanent trends and limiting non-cyclical variation.

The economic interpretation is more challenging. Our filtering step is deliberately agnostic about the nature of the shocks, which are identified based on their impact at business-cycle frequencies rather than through economically motivated restrictions, as in structural VARs. As emphasized by Dieppe, Francis, and Kindberg-Hanlon (2021) and Francis and Kindberg-Hanlon (2022), max-variance identification typically recovers combinations of underlying economic disturbances, weighted by their relevance for cyclical dynamics, rather than true

structural shocks. Consequently, attempts to assign structural labels to these shocks, for instance via impulse-response analysis, are fragile. The difficulty is compounded when retaining two shocks or more, as any additional orthogonal rotation alters the interpretation of individual shocks without affecting the implied VAR cycles.⁷

2.5. Summary. To summarize, our proposed method for measuring business cycles involves the following steps:

- (1) Select from the vector $\mathbf{x}_t$ the subset $\mathcal{T}$ of target variables that defines business cycles.
- (2) Estimate the VAR model in equation (1), including deterministic terms (intercept, time trends) if necessary.
- (3) Detrending step: compute matrix $\mathbf{F}$ in the companion form (2), set the eigenvalue stability threshold $\bar{\rho}$, determine the number of unstable eigenvalues, and diagonalize the system to isolate the stationary component $\mathbf{x}_{s,t}$ as in equation (3).
- (4) Filtering step: construct $\mathbf{V}(\underline{\omega}, \bar{\omega}, \mathcal{T})$ using equation (6) for a specified frequency band, then solve the eigenvalue problem to obtain the impulse vectors.
- (5) Select a number K of shocks to retain and construct the VAR cycle $\hat{\mathbf{x}}_{s,t}$ using equation (7).

Each step offers flexibility. We must select the variables for the analysis and the subset of target variables (step 1); choose the lag order, deterministic component, and estimation approach (step 2); set the stability threshold (step 3); specify the target frequency band (step 4); and determine how many shocks to retain in the final estimate (step 5).

General principles can guide these choices:

- (1) The VAR should include variables central to the research question, e.g., output, consumption, investment, hours worked, and unemployment for standard business cycles. Additional variables can be included either to extract their own cyclical component or to help address non-fundamentality. Target series should feature pronounced fluctuations to sharpen cyclical identification.
- (2) Standard practices apply for lag selection and the specification of deterministic components: with quarterly data, $p = 4$ lags and at least a constant term are typically appropriate. Bayesian estimation is especially attractive, as it offers automatic shrinkage for large systems and quantifies uncertainty through posterior confidence bands.
- (3) A stability threshold of $\bar{\rho} = 0.99$ balances the risk of excluding genuine cyclical dynamics against the risk of retaining near-unit roots.

⁷A robust economic interpretation requires a reference structural model. The detrending and filtering steps can then be applied either to the model's state-space form or to its population-implied VAR approximation, allowing the cycles to be expressed in terms of structural shocks and frictions and enabling formal comparisons between model and VAR dynamics. Given our focus on measurement, we deliberately refrain from such structural interpretation in this paper.

- (4) The frequency range $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$ spans fluctuations between 6 quarters and 8 years, in line with standard definitions of business cycles (Burns and Mitchell, 1946; Stock and Watson, 1999). Other applications may call for different ranges.
- (5) Conservative choices for the number of retained shocks, like $K \leq 2$, strike a balance between capturing relevant dynamics and avoiding excess noise or weakening the link to business cycles. Historical and variance decompositions formalize this trade-off, much as scree plots are used in principal component analysis (PCA) and dynamic factor models.⁸

Finally, it is essential to assess robustness to alternative variable selections, VAR and prior specifications, stability thresholds, number of retained shocks, and sample periods.

The steps outlined above also highlight key aspects that differentiate our VAR method from existing cycle-extraction techniques.

Flexibility and robustness. Our method fully leverages the flexibility of VARs while addressing key limitations of parametric alternatives. Unlike the Beveridge-Nelson decomposition, which imposes exact unit roots, we assess persistence based on estimated eigenvalues. Our VAR-based method should also be more robust to specification error than UCMs, which require specifying the dynamics of multiple latent variables and are prone to identification issues.⁹ Compared to univariate techniques, our method exploits comovement across series, allowing it to more effectively separate permanent and transitory components (Cochrane, 1994; Coibion, Gorodnichenko, and Ulate, 2018). It also provides a consistent characterization of fluctuations, ensuring cross-variable coherence as VAR cycles reflect shared dynamics stripped of idiosyncratic movements.

Stationarity. Our VAR cycles are stationary by construction, addressing a fundamental limitation of alternatives like the HP filter, which may yield nonstationary cycles in realistic samples (Phillips and Jin, 2021). Stationarity is an essential property of meaningful cycles, which represent temporary deviations from long-run trends. Crucially, our method will only identify a stationary cycle if one is present in the VAR: if the data follows a multivariate random walk, all variation should be attributed to trends, and no cycle would be detected. In contrast, bandpass, HP, and Hamilton filters extract cycles even from pure random walks, generating spurious stationary dynamics that are absent from the DGP.

⁸The analogy between our method and factor analysis rests on the shared goal of reducing the stochastic dimension of the data while preserving relevant information (Sargent and Sims, 1977; Stock and Watson, 2016). However, there are two key differences. Theoretically, factors and shocks are distinct concepts: a (static) factor affects all observables proportionally in each period; a shock triggers different dynamic responses across variables. Empirically, unlike our VAR method, PCA cannot handle non-stationary data and dynamic factor models usually focus on stationary dynamics (Stock and Watson, 2016, Section 2.1.4).

⁹See Harvey (1985), Watson (1986), and Harvey and Jaeger (1993) for univariate UCMs, and Harvey, Trimbur, and Van Dijk (2007) for a multivariate extension. For discussions of identification and specification challenges, see Morley, Nelson, and Zivot (2003), Kiley (2020), and Buncic, Pagan, and Robinson (2023).

Focus on cyclical frequencies. Our procedure explicitly targets the frequency range commonly associated with business cycles, unlike alternatives such as the Beveridge-Nelson decomposition and the Hamilton filter, which are not frequency-specific and may extract stationary components that lack cyclical patterns. For example, the Beveridge-Nelson cycle for U.S. GDP is neither volatile nor persistent and fails to align with NBER chronology (Kamber, Morley, and Wong, 2018), while the Hamilton filter amplifies fluctuations longer than typical business-cycle durations (Schüler, 2024). In contrast, we will show that our VAR cycles exhibit plausible cyclical behavior in U.S. data.

Preservation of timing and dynamics. Our VAR cycles are backward-looking, depending only on current and past shocks, and thereby preserve the timing and propagation mechanisms estimated from the data. In contrast, two-sided filters like bandpass and HP rely on non-causal smoothing, incorporating future information unavailable in real time and potentially generating spurious patterns (Cogley and Nason, 1995; Hamilton, 2018). The Hamilton filter also distorts timing: defining cycles as two-year-ahead forecast errors mechanically delays trend adjustments and often yields implausible decompositions, with trends rising during recessions and falling at the start of the next recovery (Moura, 2024).

Business cycles as usual. By tracing the effects of shocks identified over the full sample, our VAR cycles measure ‘business cycles as usual,’ i.e., typical patterns observed over extended periods. This focus offers three key advantages. First, constructing cycles from well-defined shocks allows our method to retain information on the stochastic dimension of the cyclical component, a crucial input for developing structural models that univariate methods cannot provide. Second, focusing on typical patterns driven by well-defined shocks is a better basis for producing summary business-cycle statistics, while more flexible approaches may blend different disturbances without distinction, complicating interpretation. Third, establishing what constitutes the ‘usual’ cycle enables systematic comparisons, revealing how specific episodes diverged from typical patterns and offering new insights into the nature of each period.

Estimation uncertainty. A Bayesian implementation of our method yields credible intervals around VAR cycles through posterior simulation. In contrast, most alternative filters provide point estimates without associated measures of uncertainty. While UCMs can quantify uncertainty, they are more complex and remain vulnerable to the specification and identification issues mentioned earlier.

3. MEASURING U.S. BUSINESS CYCLES

This section applies our method to study postwar U.S. business cycles. We outline the baseline VAR specification, present our cyclical estimates, and compare them with common alternatives. Throughout, we focus on pre-COVID ‘business cycles as usual,’ which our method is designed to capture.

3.1. Specification. Our baseline VAR uses quarterly data on ten key macroeconomic variables: the logarithms of real per-capita GDP (y), consumption (c), and investment (i); the log of per-capita hours worked (h); the unemployment rate (u); the labor share (ls); the federal funds rate (r); the log-difference of the GDP deflator (π); the log of utilization-adjusted TFP (tfp); and the spread between BAA-rated corporate-bond yields and 10-year U.S. government-bond yields (spr). Appendix A provides details on definitions and sources. All series except the federal funds rate, TFP, and the spread are seasonally adjusted.

This selection reflects several considerations. The VAR includes core national-accounts aggregates that are central to canonical business-cycle models (Smets and Wouters, 2007; Justiniano, Primiceri, and Tambalotti, 2011), along with key labor market indicators that capture employment dynamics relevant for NBER recession dating. Including both price and quantity variables enables a joint analysis of real and nominal developments. TFP and credit spreads are added to account for productivity and financial conditions, two potential business-cycle drivers. All variables enter in levels.

The estimation sample runs from 1955Q1, the earliest available date for the federal funds rate, to 2019Q4. We construct the cyclical components over this same period. Consistent with our focus on typical business cycles, we end the sample before the onset of the COVID-19 pandemic to exclude the extreme outliers associated with that episode, which we analyze separately below.

Our baseline VAR includes $p = 4$ lags and no deterministic time trend. We estimate the model using Bayesian methods, adopting an independent Normal - inverse Wishart prior of the Minnesota type.¹⁰ Posterior distributions are obtained via Gibbs sampling with 3,000 draws, retaining the last 1,000 for inference. To define the cyclical components, we eliminate eigenvalues with modulus above $\bar{\rho} = 0.99$ and target the average variance of GDP, consumption, investment, hours worked, and unemployment, the five variables used by Angeletos, Collard, and Dellas (2020) to characterize the primary business-cycle shock. We focus on the 6-to-32-quarter frequency band and retain $K = 2$ shocks, with Section 3.4 justifying this choice of K .

Section 4 assesses the sensitivity of our results to alternative specifications. We examine smaller VARs that exclude some variables, larger VARs that incorporate additional series, variations in prior distributions and lag order, and the inclusion of deterministic time trends. We also evaluate the stability of the cycles across subsamples. These robustness checks

¹⁰Expressing the VAR as $\mathbf{x}_t = \mathbf{c} + \sum_{j=1}^p \mathbf{A}_j \mathbf{x}_{t-j} + \mathbf{u}_t$, this prior assigns a mean of zero to the elements of $\mathbf{c}$ and $\mathbf{A}_j$ for $j = 2, \dots, p$, as well as to the off-diagonal elements of $\mathbf{A}_1$, while the diagonal elements of $\mathbf{A}_1$ have a prior mean of one. The prior variances are $(\gamma_1/k^{\gamma_3})^2$ for the i, i th entry of $\mathbf{A}_k$, $[\sigma_i \gamma_1 \gamma_2 / (\sigma_j k^{\gamma_3})]^2$ for the i, j th entry of $\mathbf{A}_k$, $i \neq j$, and $(\sigma_i \gamma_4)^2$ for the i th entry of $\mathbf{c}$, where σ_i is the residual standard deviation from a univariate AR(4) model for variable i . The γ 's control the tightness of the prior and we set $\gamma_1 = 0.2$, $\gamma_2 = 1$, $\gamma_3 = 2$, and $\gamma_4 = 100,000$. The prior on Σ , the residual variance matrix, is centered at $\mathbf{I}_n$ with $n + 1$ degrees of freedom.

confirm the reliability of our method. Finally, we extend the sample through late 2022 to include the COVID period and analyze how the pandemic-related recession differs from earlier downturns.

3.2. VAR cycles. Figure 1 presents our VAR cycles. Table 1 reports the corresponding business-cycle statistics, including the volatility of the cycles, their persistence, and their comovement with GDP cycles.

Visually, the estimated cyclical components exhibit no discernible trend and clear mean reversion, confirming that our method effectively removes non-stationary components.

The cycles in GDP, investment, consumption, hours worked, and unemployment display plausible magnitudes, with pre-2008 recessions featuring cyclical losses of approximately 3-4% for GDP, 1-2% for consumption, 10-15% for investment, 2-3% for hours worked, and a 1-2 percentage-point increase in unemployment. The standard deviations in Table 1 align with the typical view that GDP cycles are more volatile than those in consumption, less volatile than those in investment, and comparable to the volatility of cycles in hours worked. Interpreting volatility of cyclical components is more challenging for other variables due to the lack of established benchmarks. One approach, pursued below, is to compare the VAR cycles with some common alternatives.


The charts also underscore the exceptional severity of the Great Recession, with cyclical losses roughly twice as large as in earlier downturns. This aligns with direct evidence on recession size: according to Christiano (2017), during the 2007-09 recession, GDP fell by 7.2% (versus a 4.4% postwar recession average), consumption by 5.4% (2.1%), investment by 33.5% (17.8%), and hours worked by 8.7% (3.2%). That our ‘business-cycles-as-usual’ estimates capture this event well supports earlier findings by Stock and Watson (2012), who argued that the Great Recession was driven by shocks similar to those in earlier downturns, though of greater magnitude. However, we show below that other detrending methods often fail to reflect this severity.

More broadly, our cycles match the NBER business-cycle chronology closely, with every recession in the sample coinciding with falling activity and rising unemployment. This provides initial support that our two-shock estimates effectively capture the main features of postwar U.S. fluctuations.

The cyclical components exhibit strong persistence, with first-order autocorrelations in Table 1 ranging from 0.87 for inflation to 0.97 for hours worked, unemployment, and TFP. This persistence is consistent with standard views on business cycles and reflects our procedure’s emphasis on cycles lasting several quarters to multiple years.

Investment and hours worked are strongly procyclical, since their correlations with GDP cycles are above 0.95, while unemployment is strongly countercyclical (correlation of -0.97). Consumption is also procyclical but less tightly aligned (GDP correlation of 0.82).

FIGURE 1. VAR cycles — Baseline specification



Notes. Solid lines depict posterior median estimates, with shaded bands indicating 68% credible intervals. Vertical bands mark NBER recessions. See Table 1 notes for data and VAR details.

TABLE 1. Business-cycle statistics — Baseline specification

	Standard deviation	First-order autocorrelation	Correlation with GDP cycle
Cyclical component of			
gross domestic product	1.81	0.94	1.00
investment	7.23	0.94	0.96
consumption	0.86	0.88	0.82
hours worked	1.90	0.97	0.97
unemployment	1.23	0.97	-0.97
labor share	1.12	0.92	0.28
nominal interest rate	1.50	0.93	0.79
inflation rate	0.73	0.87	0.66
total factor productivity	0.64	0.97	-0.30
credit spread	0.46	0.89	-0.77

Notes. Statistics are based on posterior median estimates from the 1955-2019 sample. Cyclical components are expressed in percent for GDP, investment, consumption, hours worked, and TFP; in percentage points for unemployment and the labor share; and in annualized percentage points for the nominal interest rate, inflation, and credit spread. The VAR includes $p = 4$ lags. Cycles are driven by the $K = 2$ shocks that contribute most to the average variance of GDP, consumption, investment, hours worked, and unemployment within the frequency band $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, after eliminating eigenvalues larger than $\bar{\rho} = 0.99$.

Turning to other variables, the labor share typically peaks during recessions, drops sharply at the start of the subsequent expansion, and then gradually rises until the next downturn. This pattern, which is well documented in the literature on markups (Rotemberg and Woodford, 1999; Nekarda and Ramey, 2020), reflects the relative stability of labor income compared to the much more procyclical capital income. It results in a low correlation of 0.28 with the cyclical component of GDP. In contrast, the nominal interest rate is strongly procyclical (correlation of nearly 0.80), reflecting the Federal Reserve’s sharp rate cuts during recessions.

Inflation is also procyclical, with correlations of 0.66 with GDP and -0.58 with unemployment. This stands in stark contrast to Angeletos, Collard, and Dellas (2020), who emphasize a disconnect between inflation and the business cycle. Why do we reach such a different conclusion? The divergence does not stem from our decision to retain $K = 2$ shocks. Since Angeletos, Collard, and Dellas focus on the primary business-cycle shock, their results effectively correspond to a $K = 1$ specification. To match this setup, we compute single-shock VAR cycles and find an even stronger connection between inflation and real activity: correlations rise to 0.81 with GDP and -0.97 with unemployment.

Instead, the difference stems from the methodology used to assess comovement: correlations between cyclical components provide a more robust measure than impulse responses

or variance decompositions. As discussed in Section 3.4, our VAR also attributes only a modest share of the variance of inflation at business-cycle frequencies to the primary shock, but this does not diminish the relevance of our finding. Even if most inflation fluctuations originate from shocks that are orthogonal to real activity, such as markup shocks, the strong procyclicality of inflation conditional on the main business-cycle shocks provides important information to understand aggregate dynamics. This pattern becomes even more economically significant when retaining two shocks, as in our baseline cycles, which increases the contribution to the variance of inflation at cyclical frequencies to 30% while preserving its procyclical nature.


Overall, our VAR cycles support the New Keynesian view that procyclical marginal costs generate upward price pressure during booms and disinflation during recessions. This Phillips-curve mechanism aligns with recent empirical evidence for the U.S., e.g., McLeay and Tenreyro (2019), Stock and Watson (2019), Hazell, Herreno, Nakamura, and Steinsson (2022), and Bianchi, Nicoló, and Song (2023). It emerges as an important stylized fact, suggesting that inflationary demand shocks play a greater role in business cycles than implied by Angeletos, Collard, and Dellas (2020).

The bottom panels in Figure 1 show that TFP and spreads are both countercyclical, as their cyclical components are negatively correlated with the cycle in GDP. This is expected for credit spreads, which tend to rise in recessions and spiked sharply during the financial crisis.

The countercyclicality of utilization-adjusted TFP is more surprising. Setting aside potential measurement error in Fernald’s (2012) series, our VAR cycles provide clear empirical evidence that the primary forces behind U.S. business cycles are largely orthogonal to cyclical TFP movements. While this does not rule out a role for permanent TFP changes, we verify that neither of our two business-cycle shocks induces permanent movements in TFP by tracing their effects on the unit-root component eliminated in the detrending step. This confirms the disconnect and suggests that TFP fluctuations are not a primary driver of business cycles. Given the atheoretical nature of our measurement approach, this finding complements and reinforces the literature questioning the role of technology shocks in short-term dynamics (Gali, 1999; Francis and Ramey, 2005; Kimball, Fernald, and Basu, 2006; Ramey, 2016).

Finally, Figure 1 plots the cyclical components within posterior confidence bands. This intuitive measure of uncertainty represents a clear advantage over most detrending and filtering approaches, which typically provide estimates without confidence measures (with UCMs as the main exception). Sampling uncertainty varies remarkably across series. For key macroeconomic variables such as GDP, investment, hours worked, and unemployment, posterior bands are relatively narrow. This provides strong confidence that our estimated cycles are precise, an important result given the heavily parametrized nature of our VAR

FIGURE 2. Comparison with CBO output gap



Notes. The dashed blue line depicts the posterior median VAR GDP cycle, with shaded bands indicating a 68% credible interval. The solid red line shows the CBO output gap. See Table 1 notes for data and VAR details.

framework. That our method delivers sharp inference for these central aggregates reinforces its empirical relevance.

Instead, estimated cycles for variables such as consumption and TFP display much wider confidence bands. These are the most persistent series in our system, and our method attributes a larger share of their variation to the trend, resulting in smaller and more uncertain estimates of the cyclical component for these variables, with broader posterior confidence intervals. Nonetheless, all estimated cycles, except for TFP, display statistically significant fluctuations, with peaks and troughs clearly distinct from zero.

3.3. Comparison to common measures of the cycle. We now compare our VAR cycles to several common alternatives, both to validate their credibility and to underscore key differences that distinguish our method.

Figure 2 compares our GDP cycle to the output gap published by the Congressional Budget Office (CBO), a widely used cyclical indicator known for its alignment with the NBER chronology and frequent use in business-cycle analysis (Kamber, Morley, and Wong, 2018, 2025). Visually, the two measures track each other closely, particularly after 1985. This is confirmed by formal correlations: 0.77 over the full 1955-2019 sample and 0.91 over 1985-2019. While the CBO gap may not be an ideal measure of output fluctuations, its close alignment with our VAR cycle, despite not being used in the estimation, provides strong external validation of our method and adds to the empirical credibility of our estimates.¹¹

¹¹It is also instructive to examine the discrepancies between the CBO gap and our VAR cycle, notably the stronger expansion in the early 1960s and the deeper recessions in 1980 and 1981-82 implied by the CBO. Two (not mutually exclusive) explanations arise: our method may filter out idiosyncratic shocks with limited relevance for broader postwar dynamics, or it may assign part of these fluctuations to the trend. Evidence favors the latter: even when we remove the filtering step and retain *all* transitory GDP movements, the VAR cycle remains distinct from the CBO gap. This finding aligns with Coibion, Gorodnichenko, and Ulate

An additional strength of our method lies in its multivariate nature and its focus on well-defined shocks. By tracing the effects of these shocks across all variables, we obtain a coherent representation of business cycles. This perspective is especially valuable for macroeconomic modeling: in particular, our results imply that a successful business-cycle model should reproduce the cycles in Figure 1 and the summary statistics in Table 1 using only two shocks. By contrast, standalone measures like the CBO gap or cycles from univariate filters offer little guidance on the number of driving shocks, and cannot assess whether there are distinct shocks acting on different variables or episodes.

These insights extend to comparisons with standard alternatives, including the HP filter with smoothing parameter $\lambda = 1,600$, the Christiano and Fitzgerald (2003) bandpass filter targeting 6-32 quarters, and the Hamilton filter with $p = 4$ and $h = 8$.¹² Table 2 summarizes the properties of these alternative measures of the cycle and reports their correlations with our VAR cycles. Figure 3 provides visual comparisons for GDP, unemployment, and inflation; Online Appendix I reports full results.

For key real variables (GDP, investment, consumption, hours worked, unemployment), the VAR cycles exhibit slightly greater volatility than those from the HP and bandpass filters, indicating that the VAR method assigns larger fluctuations to the cycle. Correlations between the VAR and HP/bandpass cycles are modest, ranging from 0.15 to 0.8 at most. Both these observations confirm that the cyclical components isolated by our VAR method are distinct from those estimated by standard two-sided filters.

The top- and middle-left panels of Figure 3 compare VAR and HP cycles for GDP and unemployment. (Bandpass cycles closely resemble HP cycles; see Figure 9 in Online Appendix I.) Before 1990, VAR and HP cycles align closely, offering similar perspectives on 1970s-1980s recessions in particular and suggesting that our multivariate parametric estimates capture the episodes about as well as the flexible univariate HP filter. Discrepancies emerge in the 1990s and 2000s, as VAR cycles show a stronger expansion in the 1990s and, crucially, a deeper 2008-2009 recession. This likely reflects the reliance of two-sided filters on future data: deeper downturns like the Great Recession prompt larger trend adjustments driven by future observations, mechanically reducing the depth of estimated cyclical

(2018), who argue that the CBO measure of potential output underreacts to permanent shocks, distorting the inferred gap. Visually, Figure 2 reveals a clear downward trend in the CBO gap over 1960-1985, which is absent from our VAR cycle.


¹²The HP and bandpass filters are well known. The Hamilton filter defines the trend as the expected value of the series at date t , based on its behavior up to $t - h$, using a linear regression on a constant, the value h periods ago, and $p - 1$ additional lags; the cycle is given by the residual. We also considered the univariate Beveridge-Nelson decomposition, but the resulting cycles were small in amplitude, weakly persistent, and poorly aligned with the NBER chronology; see Kamber, Morley, and Wong (2018) for a discussion.

TABLE 2. Comparison with alternative cyclical estimates

	VAR				HP				BP				HF			
	σ	ρ	ρ_y	σ	ρ	ρ_y	ρ_{var}	σ	ρ	ρ_y	ρ_{var}	σ	ρ	ρ_y	ρ_{var}	σ
Cyclical component of																
gross domestic product	1.83	0.94	1.00	1.47	0.86	1.00	0.68	1.39	0.92	1.00	0.52	3.21	0.90	1.00	0.65	
investment	7.33	0.94	0.96	5.48	0.85	0.93	0.80	5.20	0.92	0.95	0.62	11.03	0.89	0.91	0.74	
consumption	0.87	0.89	0.82	0.84	0.84	0.80	0.50	0.79	0.92	0.85	0.39	2.02	0.88	0.84	0.35	
hours worked	1.93	0.97	0.97	1.32	0.89	0.86	0.72	1.26	0.93	0.88	0.52	2.52	0.90	0.82	0.74	
unemployment	1.25	0.97	0.97	0.74	0.91	0.86	0.72	0.68	0.92	0.88	0.50	1.35	0.91	0.80	0.78	
labor share	1.13	0.93	0.28	1.15	0.63	-0.21	0.55	0.95	0.89	-0.15	0.43	2.19	0.84	-0.09	0.57	
nominal interest rate	1.52	0.93	0.79	1.46	0.83	0.42	0.77	1.34	0.93	0.40	0.62	2.68	0.91	0.31	0.73	
inflation rate	0.74	0.87	0.66	1.05	0.46	0.17	0.52	0.86	0.91	0.21	0.49	1.76	0.78	0.14	0.53	
total factor productivity	0.65	0.97	-0.30	0.94	0.66	-0.07	0.25	0.83	0.89	-0.16	0.17	2.02	0.88	0.07	0.26	
credit spread	0.46	0.89	-0.77	0.47	0.78	-0.60	0.71	0.44	0.90	-0.62	0.60	0.71	0.88	-0.57	0.71	

Notes. ‘VAR’ refers to posterior median VAR cycles, ‘HP’ to HP cycles ($\lambda = 1,600$), ‘BP’ to bandpass cycles (6-32 quarters), and ‘HF’ to Hamilton cycles ($p = 4$, $h = 8$). For all cycles, σ is the standard deviation, ρ the first-order autocorrelation, and ρ_y the contemporaneous correlation with GDP cycles. For HP, BP, and HF cycles, ρ_{var} is the contemporaneous correlation with the corresponding VAR cycle. The sample accounts for the loss of 11 initial observations in the Hamilton filter. See Table 1 notes for data and VAR details.

FIGURE 3. Selected comparison with alternative cyclical estimates



Notes. Dashed blue lines depict posterior median VAR cycles, with shaded bands indicating 68% credible intervals. Solid red lines show HP cycles ($\lambda = 1,600$) in the left column, and Hamilton cycles ($p = 4, h = 8$) in the right column. See Table 1 notes for data and VAR details.

troughs.¹³ In contrast, our VAR method separates transitory from permanent components based on component persistence and appears to recover the magnitude of the Great Recession more accurately, aligning with the CBO output gap.

VAR cycles also reflect the slow recovery in GDP and unemployment following the financial crisis better than HP cycles, which quickly revert toward zero and understate the persistence of slack. This matters for macroeconomic analysis, as the protracted weakness in activity and labor markets was a defining feature of the post-crisis period, shaping both policy responses and subsequent research (Hall, 2011; Galí, Smets, and Wouters, 2012). Accurately representing this drawn-out recovery is essential to understand the full extent of cyclical damage associated with the Great Recession.

¹³Consistent with this view, the one-sided HP filter, which ignores future data, implies a GDP decline roughly 2 percentage points larger than the two-sided filter during the 2008-09 recession, closer to our VAR cycle and the CBO output gap.

VAR cycles for the labor share, nominal interest rate, and credit spreads are as volatile as those from HP and bandpass filters. VAR cycles for inflation and TFP are about one-third less volatile (see the lower-left panel of Figure 3 for inflation), suggesting that these variables are largely driven by shocks unrelated to the two driving our VAR cycles. Univariate filters, which do not condition on shocks, cannot make this comparison across variables: they extract cyclical components separately for each variable, potentially attributing fluctuations to different underlying disturbances. This limits their usefulness for joint economic interpretation. For example, the HP and bandpass filters find weak comovement between output and inflation cycles, with correlations near 0.20, while our VAR cycles yield a much stronger correlation of 0.66. This suggests that by filtering out idiosyncratic inflation movements our VAR method reveals a tighter link to real activity.


Finally, the Hamilton filter produces cyclical components that are markedly more volatile, as shown in Table 2 and the right column of Figure 3. The top panel, focusing on GDP, reveals that the filter implies larger cyclical losses during the 1958 and 1973-75 recessions than during the Great Recession, despite the fact that the latter involved output losses exceeding 5%, a full percentage point more than in the earlier downturns. In addition, following recessions Hamilton cycles feature sudden recoveries, far more pronounced than in other estimates. As noted earlier and emphasized in Moura (2024), these exaggerated rebounds lack a meaningful economic interpretation, as they result from the delayed trend adjustments inherent in the Hamilton filter.

Overall, the comparisons confirm that different cycle-extraction strategies can yield markedly different results. This highlights the need for an approach that exploits data properties and enhances economic interpretability, precisely the strength our VAR method.


3.4. Rationale for retaining $K = 2$ shocks. Having established the empirical credibility of our VAR cycles, we now justify our decision to retain $K = 2$ shocks. The key tradeoff is familiar: too few shocks risk missing relevant dynamics, while too many dilute the connection to business cycles and increase estimation noise. We explore this tradeoff using historical and variance decompositions.

Figures 4 and 5 compare cycles with $K = 1$ or $K = 3$ shocks (in blue) to our baseline $K = 2$ estimates. One-shock cycles closely track our benchmark after 1990, especially during the Great Recession, reflecting the prominence of that episode in identifying the dominant shock. However, one-shock cycles understate fluctuations in the 1960s-1980s, missing the depth of recessions and strength of expansions, especially in 1973-75 and 1981-82. In contrast, as shown in Figure 3, two-shock cycles better match both the magnitude and timing of these episodes, performing about as well as the HP filter for GDP and unemployment and providing a more accurate representation of pre-1990 dynamics.

This finding suggests a shift in cyclical drivers over time. By capturing distinct episodes across the full sample, the $K = 2$ specification delivers a more comprehensive account


FIGURE 4. VAR cycles — Retaining $K = 1$ shock

Notes. Solid blue lines depict posterior median VAR cycles driven by the $K = 1$ shock that contributes most to the average variance of GDP, consumption, investment, hours worked, and unemployment within the frequency band $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, after eliminating eigenvalues larger than $\bar{\rho} = 0.99$. Shaded bands indicate associated 68% credible intervals. Dashed black lines show our reference VAR cycles driven by $K = 2$ shocks. See Table 1 notes for data and VAR details.

FIGURE 5. VAR cycles — Retaining $K = 3$ shocks

Notes. Solid blue lines depict posterior median VAR cycles driven by the $K = 3$ shocks that contribute most to the average variance of GDP, consumption, investment, hours worked, and unemployment within the frequency band $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, after eliminating eigenvalues larger than $\bar{\rho} = 0.99$. Shaded bands indicate associated 68% credible intervals. Dashed black lines show our reference VAR cycles driven by $K = 2$ shocks. See Table 1 notes for data and VAR details.

FIGURE 6. Cyclical variance contributions



Notes. The x -axis rank shocks by their contributions to the average variance of GDP, consumption, investment, hours worked, and unemployment within the frequency band $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, after eliminating eigenvalues larger than $\bar{\rho} = 0.99$. The y -axis shows the cumulative contribution to individual variances. Solid lines depict posterior median estimates, with shaded bands indicating 68% credible intervals. See Table 1 notes for data and VAR details.

of U.S. business cycles. This marks a key departure from Angeletos, Collard, and Dellas (2020), whose single-shock focus overlooks the evolving nature of aggregate fluctuations and the important role of the second shock before the 1990s.

Figure 5 shows that adding a third shock does little to enhance the cyclical representation of core variables like GDP, investment, hours worked, and unemployment. Instead, it mostly captures idiosyncratic movements in consumption, inflation, and TFP, precisely the type of fluctuations our method aims to filter out. The three-shock specification also yields wider uncertainty bands without meaningful gains in economic interpretation, further supporting our choice of $K = 2$ shocks as the most informative and parsimonious characterization of U.S. business cycles.

The population variance decompositions in Figure 6 reinforce this conclusion. The first shock accounts for the bulk of cyclical variance in real variables (60% for GDP, 80% for investment, 65% for unemployment), while the second shock raises the explained share above 90% for most of these variables. These are statistically significant gains, with a clear and common ‘elbow’ in the variance profiles indicating a natural breakpoint at $K = 2$.¹⁴ The second shock also contributes meaningfully to the cyclical behavior of nominal rates and inflation, strengthening real-nominal linkages. By contrast, the third shock adds little to core macro variables and primarily reflects variation orthogonal to business cycles.

These decompositions also clarify which variables systematically comove with the cycle and which are shaped by unrelated shocks. The two retained shocks account for most cyclical variation in core real variables, as well as a substantial share of nominal-rate and credit-spread fluctuations. In contrast, they explain only about 30% of cyclical variation in inflation and little of TFP, suggesting that these series are largely driven by disturbances unrelated to business cycles. Crucially, our VAR cycles isolate the portion of each series that responds to the main business-cycle shocks, filtering out idiosyncratic components and retaining only the relevant comovement. Such quantitative insights sharpen interpretation and usefully inform modeling, while also distinguishing our method from traditional filters, which blend together all sources of variation and cannot identify whether different variables respond to common shocks.

4. ROBUSTNESS AND COVID CYCLES

This final section establishes the robustness of our VAR method across alternative specifications and extends the analysis to the atypical COVID episode.

¹⁴See Jolliffe (2002) and Stock and Watson (2016) for a discussion of elbow tests in PCA and dynamic factor models. Our logic is similar, although we cannot formally test one value of K against another.

4.1. Robustness. Unlike standard cycle-extraction approaches, which are typically univariate and nonparametric, our VAR method is multivariate and parametric.¹⁵ This allows us to exploit cross-variable dynamics to identify trends and cycles, but also raises concerns about robustness. Here, we document that our VAR cycles remain stable across a wide range of alternative implementations, underscoring our method’s reliability.

For each variable, Table 3 reports the correlations between our baseline VAR cycles and those obtained under alternative specifications. All estimates use $K = 2$ shocks, identified (unless noted otherwise) by targeting the average variance of GDP, consumption, investment, hours worked, and unemployment over the frequency band $[\underline{\omega}, \bar{\omega}] = [2\pi/32, 2\pi/6]$, after eliminating eigenvalues larger than $\bar{\rho} = 0.99$. Correlations are computed over the baseline 1955-2019 sample, adjusted when necessary. For completeness, Figures 11-28 in Online Appendix I display the cycles from each alternative specification alongside the baseline estimates.

Rows 1-2 assess sensitivity to the lag order, using $p = 2$ and $p = 8$ instead of the baseline $p = 4$.

Rows 3-4 evaluate the identification of business-cycle shocks. Given the weak cyclical behavior of consumption, Row 3 excludes it from the set of target variables. Row 4 replaces the baseline unweighted variance target by a weighted version that scales each variable’s contribution by its unconditional variance within the relevant frequency band.

Rows 5-8 compare the baseline cycles to those from alternative VAR systems. Rows 5-7 expand the information set by including additional variables: Row 5 adds standard cyclical indicators — the CBO output gap, manufacturing capacity utilization, and macroeconomic uncertainty from Jurado, Ludvigson, and Ng (2015) — to sharpen the focus on business cycles; Row 6 incorporates the first two macroeconomic factors from McCracken and Ng (2020) to address potential non-fundamentalness; Row 7 adds government consumption, exports, and imports. Instead, Row 8 restricts the system to a smaller core of real variables (GDP, investment, consumption, hours worked, and unemployment), omitting direct information on financial and nominal dynamics.

Rows 9-14 examine the role of slow-moving trends. Row 9 replaces our detrending step with the simpler approach of Angeletos, Collard, and Dellas (2020), discarding unstable posterior draws. Row 10 introduces a deterministic, variable-specific cubic trend, which is flexible enough to approximate piecewise linear trends. Rows 11-12 report sensitivity to the stability threshold $\bar{\rho}$ for classifying eigenvalues in the detrending step, setting it to 0.95 and 1.0 instead of the baseline 0.99. Rows 13-14 add prior components that shrink the VAR toward unit roots and/or cointegration: Row 13 uses the long-run prior of Giannone, Lenza, and Primiceri

¹⁵BP and HP filters are canonical univariate, nonparametric approaches. The Beveridge-Nelson and Hamilton filters are parametric but univariate. Multivariate UCMs are the closest to our VAR method in terms of dimensionality and parametrization.

TABLE 3. Robustness

	Correlations with baseline cycles									
	<i>y</i>	<i>i</i>	<i>c</i>	<i>h</i>	<i>u</i>	<i>ls</i>	<i>r</i>	π	<i>tfp</i>	<i>spr</i>
<i>VAR order</i>										
1- VAR(2)	1.00	1.00	0.98	1.00	1.00	1.00	0.99	0.99	0.98	1.00
2- VAR(8)	1.00	1.00	1.00	1.00	1.00	0.99	1.00	1.00	0.99	1.00
<i>Identification strategy</i>										
3- No consumption	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00
4- Weighted target	0.99	0.99	0.95	0.98	0.98	0.98	0.98	0.96	0.96	0.98
<i>VAR system</i>										
5- Cyclical	0.85	0.90	0.73	0.91	0.90	0.93	0.84	0.61	0.73	0.93
6- Factors	0.98	0.98	0.95	0.97	0.98	0.96	0.98	0.95	0.93	0.98
7- Large	0.90	0.94	0.82	0.90	0.90	0.92	0.90	0.84	0.87	0.95
8- Small	0.88	0.96	0.20	0.88	0.91	--	--	--	--	--
<i>Trends</i>										
9- ACD	0.75	0.94	0.36	0.98	1.00	0.97	0.97	0.98	0.75	0.99
10- Deterministic	0.86	0.89	0.66	0.92	0.90	0.92	0.90	0.86	0.83	0.93
11- Cutoff 0.95	1.00	1.00	0.99	0.99	0.99	0.99	0.97	0.93	0.80	1.00
12- Cutoff 1.0	0.85	0.97	0.55	0.99	1.00	0.99	0.98	0.99	0.82	1.00
13- Long-run prior	0.95	0.95	0.88	0.95	0.96	0.77	0.90	0.87	0.45	0.97
14- Sims-Zha prior	0.94	0.93	0.71	0.94	0.96	0.65	0.85	0.83	0.67	0.96
<i>Subsample stability</i>										
15- 1955-1983	0.90	0.93	0.65	0.95	0.95	0.69	0.93	0.77	0.56	0.84
16- 1984-2019	0.78	0.94	0.19	0.92	0.94	0.48	0.86	0.45	0.30	0.91

Notes. Entries are correlations between baseline VAR cycles and alternative specifications (posterior medians, $K = 2$ shocks). Rows 1-2: Lag order. Rows 3-4: Shock identification (exclusion of consumption from targets; weighted variance target). Rows 5-8: VAR system ('cyclical' adds output gap, capacity utilization, uncertainty; 'factor' adds macro factors; 'large' adds government consumption, exports, imports; 'small' restricts to GDP, investment, consumption, hours worked, unemployment). Rows 9-14: Trend specification ('ACD' excludes unstable draws; cubic deterministic trend; alternative stability cutoffs $\bar{\rho}$; long-run prior; Sims-Zha prior). Rows 15-16: Subsample stability. Correlations computed over 1955-2019, adjusted for initial observations in Rows 1-2, for data availability in Rows 5-6, and for sample splits in Rows 15-16.

(2019), tilting estimates toward cointegration among GDP, consumption, investment, and TFP on the one hand, and between nominal rates and inflation on the other; Row 14 follows Sims and Zha (1998), combining the sum-of-coefficients prior (favoring unit roots) and the initial-observations prior (favoring cointegration).¹⁶ Finally, Rows 15-16 assess temporal stability by comparing full-sample cycles to those estimated separately over the 1955-1983 and 1984-2019 subsamples.¹⁷

The correlations reported in Table 3 confirm the robustness of the VAR method. The lag order (Rows 1-2, Figures 11-12), the exclusion of consumption from the target variables (Row 3, Figure 13), and the use of a weighted target (Row 4, Figure 14) have minimal impact on the results. Thus, consumption does not contribute independent information about business cycles, and using a simple average variance target does not distort inference about the dominant shocks.

Modifying the information set by adding or removing variables affects the estimates, though correlations with the baseline cycles remain high throughout. Augmenting the system with off-the-shelf cyclical indicators lowers the depth of the estimated trough during the Great Recession, likely because neither the output gap nor utilization declined more in 2008-09 than in earlier recessions (Row 5, Figures 15-16). Instead, including the McCracken and Ng (2020) factors or additional GDP components has little impact, suggesting that the baseline specification is not subject to non-fundamentality (Rows 6-7, Figures 17-20). Interestingly, restricting attention to a smaller real system has little effect on the cycles for GDP, investment, hours worked, and unemployment, suggesting that nominal variables are not essential for identifying business cycles in key macro series (Row 8, Figure 21). However, the consumption cycle differs markedly, reflecting the challenge of separating trend and cycle for this near-unit-root series — a difficulty already apparent in our baseline estimates from the wide uncertainty bands.

The VAR cycles are also sensitive to the treatment of trends, though the impact remains moderate. The Angeletos, Collard, and Dellas (2020) strategy, which simply discards unstable posterior draws, yields cycles that are highly correlated with our baseline estimates for most variables, but not for GDP, consumption, and TFP (Row 9, Figure 22). For these three series, the alternative detrending approach produces much wider and noisier cycles, featuring implausible patterns such as a persistent 5% positive cyclical component in GDP throughout the 1990s and 2000s, or an steadily growing consumption ‘cycle’ that remains strongly positive during the 2008–09 recession. These anomalies suggest that discarding unstable draws is not effective in removing the permanent component: by selecting stable

¹⁶Following Giannone, Lenza, and Primiceri (2019), we adopt a hierarchical Bayesian approach in these exercises, treating the hyperparameters defining prior tightness as random variables.

¹⁷We adopt the same break date as Stock and Watson (2016), and our findings are robust to alternative break points in the early 1980s. See McConnell and Perez-Quiros (2000) and Stock and Watson (2003) for empirical evidence on the structural shifts in U.S. business cycles during this period.

parameter vectors from a model estimated on trending data, the approach tends to understate the persistence of unit roots and thus misclassifies permanent movements as cyclical. This confirms the superiority of our detrending step from an empirical perspective.

Introducing a deterministic trend in the VAR has only a moderate effect on the estimated cycles: correlations with the baseline remain close to or above 0.85 in most cases, indicating that our method is robust to trend specification (Row 10, Figure 23). Similarly, lowering the eigenvalue stability threshold $\bar{\rho}$ to 0.95 has little impact, while raising it 1.0 produces more persistent GDP and consumption cycles, with implausible consumption dynamics during the 1990s and 2000s (Rows 11-12, Figures 24-25). These findings support the use of a stability threshold below unity to guard against misclassifying near-unit-root behavior as part of the stable component.¹⁸ Prior distributions that more strongly tilt the VAR toward unit roots and/or cointegration have limited effects overall, except for the typically sensitive consumption and TFP cycles (Rows 13-14, Figures 26-27). Crucially, these priors yield posterior distributions in which nearly 95% of draws are unstable; discarding them to focus on a narrow, unrepresentative subset of stable draws risks distorting inference, emphasizing the advantage of our more general detrending step.

Finally, applying the method to subsamples split in the early 1980s yields cyclical estimates largely consistent with the baseline, with only consumption and TFP cycles displaying notable deviations in the second subsample (Rows 15–16, Figure 28). In particular, the consumption cycle rises sharply during the Great Recession, indicating that the method attributes the decline to the trend rather than the cycle, likely a symptom of the difficulty in separating the two components in shorter samples. Aside from this discrepancy, the estimated cycles remain closely aligned with the full-sample results, as reflected in the generally strong correlations reported in Table 3, and continue to convey a consistent picture of past business cycles. Achieving such alignment despite significant shifts in cyclical dynamics across the two periods and the highly parametric nature of VARs is a striking result.

Taken together, these results confirm the robustness of our method to a wide range of modeling choices. Despite some variation for consumption and TFP, the broader picture of business-cycle dynamics remains remarkably stable across specifications. In addition, highlighting the uncertainty surrounding consumption and TFP cycles is an advantage of our method, as most alternatives would produce point estimates without revealing that separating trend from cycle is more difficult for these series than for the others.

¹⁸We experimented with stability thresholds ranging from 0.80 to 1.0. For all variables except TFP, correlations with the baseline cycles obtained with $\bar{\rho} = 0.99$ remain above 0.90 for thresholds between 0.92 and 0.98; see Figure 29 in Online Appendix II. Lower thresholds lead to a clear deterioration: with $\bar{\rho} = 0.88$, most correlations fall near zero as the method then attributes most persistent fluctuations to the trend. These findings suggest that, in business-cycle applications, the stability threshold should generally lie in the 0.95–0.99 range to avoid misclassification of persistent cyclical movements.

4.2. COVID cycles. While our VAR method is designed to capture typical business-cycle patterns, examining its performance during the COVID episode yields valuable insights into both the robustness of historical cyclical relationships and the dimensions along which this unprecedented shock diverged from typical postwar fluctuations.

Figure 7 presents the VAR cycles extended to include the COVID period. Following recent econometric practice, e.g., Schorfheide and Song (2021) and Lenza and Primiceri (2022), we hold the VAR parameters fixed at their values estimated from the 1955–2019 sample, assuming that pandemic-related outliers contain no information about standard macroeconomic dynamics. As a result, the pre-COVID cycles are identical to those in Figure 1, while post-2020 cycles incorporate the shocks inferred from the residuals associated with the COVID episode. For comparison, the figure also reports HP and Hamilton cycles.


For several key macroeconomic variables — GDP, consumption, the labor share, TFP, and credit spreads — our VAR cycles capture the COVID recession reasonably well. The implied dynamics align closely with fluctuations in the raw data and in alternative estimates of cyclical components, suggesting that core aspects of the contraction followed patterns consistent with earlier postwar downturns, despite the episode’s unique origin. That a model trained exclusively on pre-2020 data can track these variables highlights the robustness of certain cyclical relationships during economic downturns.

More revealingly, the VAR cycles diverge from actual data for specific variables in ways that are economically interpretable. The method understates the steep fall in hours worked and the surge in unemployment, missing the unprecedented speed and scale of labor-market disruption. Instead, it overpredicts the decline in the nominal interest rate, which quickly hit the effective lower bound — a nonlinearity that cannot be accommodated within the VAR’s linear structure. It also substantially overstates the drop in inflation, which proved more resilient than historical patterns would suggest.

These discrepancies can serve as diagnostic signals to pinpoint how the COVID recession differed from historical norms. Because our method reflects typical cyclical regularities, it highlights the precise dimensions along which the pandemic was exceptional. For instance, given the magnitude of the output contraction, historical patterns would have implied smaller labor-market losses. The divergence between these counterfactual predictions and actual outcomes likely reveals the role of abrupt policy interventions, including lockdowns and unique constraints on economic activity. More broadly, this illustrates a key strength of our method: by providing a benchmark grounded in typical dynamics, it enables researchers to identify when and how current developments deviate from standard cyclical behavior.

Figure 7 yields two further insights. First, the VAR cycles revert to zero relatively quickly after the COVID recession, capturing the transient nature of the disruption. By 2022, most cycles had returned near baseline, in line with two-sided HP cycles that benefit from hindsight, indicating that the pandemic left little lasting cyclical imprint. Second, the charts

FIGURE 7. Cyclical estimates during COVID



Notes. Solid blue lines depict posterior median VAR estimates, with shaded bands indicating 68% credible intervals. Dashed red and black lines show HP and Hamilton cycles. Vertical bands mark NBER recessions. See Table 1 notes.

clearly illustrate the artificial base effects introduced by the Hamilton filter, which generate exaggerated spikes in the estimated cycles exactly two years after the COVID downturn.

An alternative strategy for handling exceptional episodes would be to augment the VAR with episode-specific dummy variables, as proposed by Ferroni, Fisher, and Melosi (2024) in the context of DSGE models. This approach would absorb unusual shocks, such as those associated with the COVID recession, into the dummy terms, while preserving the estimated autoregressive structure for standard dynamics. As a result, much of the pandemic episode would be attributed to these dummies, leaving little of it to the underlying VAR. While this would offer a cleaner separation between regular fluctuations and exceptional disturbances, it would come at the cost of discarding the informative deviations discussed above. We view our approach as complementary: by retaining the full residual structure, it allows for a diagnostic reading of atypical episodes through the lens of historical business-cycle patterns. Exploring the potential of such hybrid strategies is an interesting avenue for future work.

5. CONCLUSION

We propose a new method to measure business cycles using VARs. This involves two steps: first, detrending to remove nonstationary components; second, filtering to extract meaningful cycles from the stable component. The resulting cycles isolate a well-defined VAR component that corresponds to typical fluctuations, are stationary and backward-looking by construction, and reflect the idea that business cycles occur at specific periodicities. In an empirical application, we show that two shocks suffice to capture postwar U.S. business cycles and reveal a tighter link between inflation and real activity than previously recognized. We also compare our method to other common filters, assess the plausibility and robustness of our results, and analyze the COVID episode in detail.

We conjecture that our method could be fruitfully applied in other contexts. Beyond straightforward extensions to other countries, a promising application is the measurement of the international business cycle, defined as the historical contribution of the shock(s) with the largest impact on national indicators within a given group of countries. By preserving conditional comovements, VAR cycles may shed new light on longstanding debates, such as the correlation between consumption and real exchange rates (Backus and Smith, 1993), or the relationship between GDP and trade (Frankel and Rose, 1998). Another relevant application is the measurement of longer-term economic and financial cycles, capturing cyclical movements in asset prices and credit conditions, as well as their interaction with real activity (Comin and Gertler, 2006; Jordà, Schularick, and Taylor, 2017). Finally, on the policy side, our VAR method could help estimate the impact of cyclical shocks on government deficits, disentangling structural imbalances from automatic stabilizers (Fedelino, Horton, and Ivanova, 2009).

REFERENCES

- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2020): “Business-Cycle Anatomy,” *American Economic Review*, 110(10), 3030–3070.
- BACKUS, D. K., AND G. W. SMITH (1993): “Consumption and Real Exchange Rates in Dynamic Economies with Non-Traded Goods,” *Journal of International Economics*, 35(3-4), 297–316.
- BAXTER, M., AND R. G. KING (1999): “Measuring Business Cycles: Approximate Band-Pass Filters For Economic Time Series,” *The Review of Economics and Statistics*, 81(4), 575–593.
- BEAUDRY, P., P. FÈVE, A. GUAY, AND F. PORTIER (2019): “When is Nonfundamentalness in SVARs a Real Problem?,” *Review of Economic Dynamics*, 34, 221–243.
- BEVERIDGE, S., AND C. R. NELSON (1981): “A New Approach to Decomposition of Economic Time Series into Permanent and Transitory Components with Particular Attention to Measurement of the ‘Business Cycle’,” *Journal of Monetary Economics*, 7(2), 151–174.
- BIANCHI, F., G. NICOLÓ, AND D. SONG (2023): “Inflation and Real Activity over the Business Cycle,” NBER Working Papers 31075, National Bureau of Economic Research, Inc.
- BLANCHARD, O., AND D. QUAH (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, 79(4), 655–73.
- BUNCIC, D., A. PAGAN, AND T. ROBINSON (2023): “Recovering Stars in Macroeconomics,” CAMA Working Papers 2023-43, Centre for Applied Macroeconomic Analysis, Crawford School of Public Policy, The Australian National University.
- BURNS, A. F., AND W. C. MITCHELL (1946): *Measuring Business Cycles*. National Bureau of Economic Research, Inc (Studies in Business Cycles).
- CANOVA, F. (1998): “Detrending and Business Cycle Facts,” *Journal of Monetary Economics*, 41(3), 475–512.
- (1999): “Does Detrending Matter for the Determination of the Reference Cycle and the Selection of Turning Points?,” *Economic Journal*, 109(452), 126–150.
- CASALS, J., M. JEREZ, AND S. SOTOCA (2002): “An Exact Multivariate Model-Based Structural Decomposition,” *Journal of the American Statistical Association*, 97, 553–564.
- CHARI, V., P. J. KEHOE, AND E. R. MCGRATTAN (2008): “Are Structural VARs with Long-Run Restrictions Useful in Developing Business Cycle Theory?,” *Journal of Monetary Economics*, 55(8), 1337–1352.
- CHRISTIANO, L. J. (2017): “The Great Recession: A Macroeconomic Earthquake,” Economic Policy Paper 17-1, Federal Reserve Bank of Minneapolis.
- CHRISTIANO, L. J., AND T. J. FITZGERALD (2003): “The Band Pass Filter,” *International Economic Review*, 44(2), 435–465.

- COCHRANE, J. H. (1994): “Permanent and Transitory Components of GNP and Stock Prices,” *The Quarterly Journal of Economics*, 109(1), 241–265.
- COCHRANE, J. H., AND A. M. SBORDONE (1988): “Multivariate Estimates of the Permanent Components of GNP and Stock Prices,” *Journal of Economic Dynamics and Control*, 12(2-3), 255–296.
- COGLEY, T., AND J. M. NASON (1995): “Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series: Implications for Business Cycle Research,” *Journal of Economic Dynamics and Control*, 19(1-2), 253–278.
- COIBION, O., Y. GORODNICHENKO, AND M. ULATE (2018): “The Cyclical Sensitivity in Estimates of Potential Output,” *Brookings Papers on Economic Activity*, 49(2 (Fall)), 343–441.
- COMIN, D., AND M. GERTLER (2006): “Medium-Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- DIEPPE, A., N. FRANCIS, AND G. KINDBERG-HANLON (2021): “The Identification of Dominant Macroeconomic Drivers: Coping with Confounding Shocks,” Working Paper Series 2534, European Central Bank.
- EDGE, R. M., AND R. S. GURKAYNAK (2010): “How Useful Are Estimated DSGE Model Forecasts for Central Bankers?,” *Brookings Papers on Economic Activity*, 41(2 (Fall)), 209–259.
- ENGLE, R. F., AND C. W. J. GRANGER (1987): “Co-integration and Error Correction: Representation, Estimation, and Testing,” *Econometrica*, 55(2), 251–276.
- FAUST, J. (1998): “The Robustness of Identified VAR Conclusions about Money,” *Carnegie-Rochester Conference Series on Public Policy*, 49(1), 207–244.
- FEDELINO, A., M. A. HORTON, AND A. IVANOVA (2009): “Computing Cyclically-Adjusted Balances and Automatic Stabilizers,” IMF Technical Notes and Manuals 2009/005, International Monetary Fund.
- FERNALD, J. G. (2012): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- FERNÁNDEZ-VILLAYERDE, J., J. F. RUBIO-RAMÍREZ, T. J. SARGENT, AND M. W. WATSON (2007): “ABCs (and Ds) of Understanding VARs,” *American Economic Review*, 97(3), 1021–1026.
- FERRONI, F., J. D. FISHER, AND L. MELOSI (2024): “Unusual Shocks in our Usual Models,” *Journal of Monetary Economics*, 147(C).
- FISHER, J. D. M. (2006): “The Dynamic Effects of Neutral and Investment-Specific Technology Shocks,” *Journal of Political Economy*, 114(3), 413–451.
- FORNI, M., AND L. GAMBETTI (2014): “Sufficient Information in Structural VARs,” *Journal of Monetary Economics*, 66(C), 124–136.

- FRANCIS, N., AND G. KINDBERG-HANLON (2022): “Signing Out Confounding Shocks in Variance-Maximizing Identification Methods,” *AEA Papers and Proceedings*, 112, 476–480.
- FRANCIS, N., AND V. A. RAMEY (2005): “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,” *Journal of Monetary Economics*, 52(8), 1379–1399.
- (2009): “Measures of per Capita Hours and Their Implications for the Technology-Hours Debate,” *Journal of Money, Credit and Banking*, 41(6), 1071–1097.
- FRANKEL, J. A., AND A. K. ROSE (1998): “The Endogeneity of the Optimum Currency Area Criteria,” *Economic Journal*, 108(449), 1009–1025.
- GALI, J. (1999): “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?,” *American Economic Review*, 89(1), 249–271.
- GALÍ, J., F. SMETS, AND R. WOUTERS (2012): “Slow Recoveries: A Structural Interpretation,” *Journal of Money, Credit and Banking*, 44, 9–30.
- GIANNONE, D., M. LENZA, AND G. E. PRIMICERI (2019): “Priors for the Long Run,” *Journal of the American Statistical Association*, 114(526), 565–580.
- GOLUB, G. H., AND C. F. VAN LOAN (1996): *Matrix Computations*. The Johns Hopkins University Press, third edn.
- HALL, R. E. (2011): “The Long Slump,” *American Economic Review*, 101(2), 431–469.
- HAMILTON, J. D. (1994): *Time Series Analysis*. Princeton University Press.
- HAMILTON, J. D. (2018): “Why You Should Never Use the Hodrick-Prescott Filter,” *The Review of Economics and Statistics*, 100(5), 831–843.
- HARVEY, A. C. (1985): “Trends and Cycles in Macroeconomic Time Series,” *Journal of Business & Economic Statistics*, 3(3), 216–227.
- HARVEY, A. C., AND A. JAEGER (1993): “Detrending, Stylized Facts and the Business Cycle,” *Journal of Applied Econometrics*, 8(3), 231–247.
- HARVEY, A. C., T. M. TRIMBUR, AND H. K. VAN DIJK (2007): “Trends and Cycles in Economic Time Series: A Bayesian Approach,” *Journal of Econometrics*, 140(2), 618–649.
- HAZELL, J., J. HERRENO, E. NAKAMURA, AND J. STEINSSON (2022): “The Slope of the Phillips Curve: Evidence from U.S. States,” *The Quarterly Journal of Economics*, 137(3), 1299–1344.
- HODRICK, R. J., AND E. PRESCOTT (1981): “Post-War U.S. Business Cycles: An Empirical Investigation,” Discussion Papers 451, Northwestern University, Center for Mathematical Studies in Economics and Management Science.
- HODRICK, R. J., AND E. PRESCOTT (1997): “Postwar U.S. Business Cycles: An Empirical Investigation,” *Journal of Money, Credit and Banking*, 29(1), 1–16.
- JOLLIFFE, I. (2002): *Principal Components Analysis, 2nd Edition*, Springer Series in Statistics. Springer.

- JORDÀ, O., M. SCHULARICK, AND A. M. TAYLOR (2017): “Macrofinancial History and the New Business Cycle Facts,” *NBER Macroeconomics Annual*, 31(1), 213–263.
- JURADO, K., S. C. LUDVIGSON, AND S. NG (2015): “Measuring Uncertainty,” *American Economic Review*, 105(3), 1177–1216.
- JUSTINIANO, A., G. PRIMICERI, AND A. TAMBALOTTI (2011): “Investment Shocks and the Relative Price of Investment,” *Review of Economic Dynamics*, 14(1), 101–121.
- KAMBER, G., J. MORLEY, AND B. WONG (2018): “Intuitive and Reliable Estimates of the Output Gap from a Beveridge-Nelson Filter,” *The Review of Economics and Statistics*, 100(3), 550–566.
- (2025): “Trend-Cycle Decomposition in the Presence of Large Shocks,” *Journal of Economic Dynamics and Control*, 173(C).
- KILEY, M. T. (2020): “What Can the Data Tell Us about the Equilibrium Real Interest Rate?,” *International Journal of Central Banking*, 16(3), 181–209.
- KILIAN, L., AND H. LUTKEPOHL (2017): *Structural Vector Autoregressive Analysis*, Themes in Modern Econometrics. Cambridge University Press.
- KIMBALL, M. S., J. G. FERNALD, AND S. BASU (2006): “Are Technology Improvements Contractionary?,” *American Economic Review*, 96(5), 1418–1448.
- KING, R. G., C. I. PLOSSER, J. H. STOCK, AND M. W. WATSON (1991): “Stochastic Trends and Economic Fluctuations,” *American Economic Review*, 81(4), 819–840.
- LENZA, M., AND G. E. PRIMICERI (2022): “How to Estimate a Vector Autoregression after March 2020,” *Journal of Applied Econometrics*, 37(4), 688–699.
- LIPPI, M., AND L. REICHLIN (1993): “The Dynamic Effects of Aggregate Demand and Supply Disturbances: Comment,” *American Economic Review*, 83(3), 644–52.
- (1994): “VAR Analysis, Nonfundamental Representations, Blaschke Matrices,” *Journal of Econometrics*, 63(1), 307–325.
- LUCAS, R. E. (1977): “Understanding Business Cycles,” *Carnegie-Rochester Conference Series on Public Policy*, 5(1), 7–29.
- MCCONNELL, M. M., AND G. PEREZ-QUIROS (2000): “Output Fluctuations in the United States: What Has Changed since the Early 1980’s?,” *American Economic Review*, 90(5), 1464–1476.
- MCCRACKEN, M. W., AND S. NG (2020): “FRED-QD: A Quarterly Database for Macroeconomic Research,” Working Papers 2020-005, Federal Reserve Bank of St. Louis.
- MCLEAY, M., AND S. TENREYRO (2019): “Optimal Inflation and the Identification of the Phillips Curve,” in *NBER Macroeconomics Annual 2019, volume 34*, NBER Chapters, pp. 199–255. National Bureau of Economic Research, Inc.
- MORLEY, J. C., C. R. NELSON, AND E. ZIVOT (2003): “Why Are the Beveridge-Nelson and Unobserved-Components Decompositions of GDP So Different?,” *The Review of Economics and Statistics*, 85(2), 235–243.

- MOURA, A. (2024): “Why You Should Never Use the Hodrick-Prescott Filter. A Comment on Hamilton (The Review of Economics and Statistics, 2018),” *Journal of Comments and Replications in Economics*, 3(2024-1), 1–17.
- NEKARDA, C. J., AND V. A. RAMEY (2020): “The Cyclical Behavior of the Price-Cost Markup,” *Journal of Money, Credit and Banking*, 52(S2), 319–353.
- NELSON, C. R. (1988): “Spurious Trend and Cycle in the State Space Decomposition of a Time Series with a Unit Root,” *Journal of Economic Dynamics and Control*, 12(2-3), 475–488.
- NELSON, C. R., AND H. KANG (1981): “Spurious Periodicity in Inappropriately Detrended Time Series,” *Econometrica*, 49(3), 741–751.
- NELSON, C. R., AND C. I. PLOSSER (1982): “Trends and Random Walks in Macroeconomic Time Series : Some Evidence and Implications,” *Journal of Monetary Economics*, 10(2), 139–162.
- PHILLIPS, P. C. B., AND S. JIN (2021): “Business Cycles, Trend Elimination, and the HP Filter,” *International Economic Review*, 62(2), 469–520.
- PRIMICERI, G. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *The Review of Economic Studies*, 72(3), 821–852.
- RAMEY, V. (2016): “Macroeconomic Shocks and Their Propagation,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and H. Uhlig, vol. 2 of *Handbook of Macroeconomics*, chap. 2, pp. 71–162. Elsevier.
- RAVENNA, F. (2007): “Vector autoregressions and reduced form representations of DSGE models,” *Journal of Monetary Economics*, 54(7), 2048–2064.
- ROTEMBERG, J. J., AND M. WOODFORD (1999): “The Cyclical Behavior of Prices and Costs,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 16, pp. 1051–1135. Elsevier.
- SARGENT, T. J., AND C. A. SIMS (1977): “Business Cycle Modeling without Pretending to Have Too Much A Priori Economic Theory,” Working Papers 55, Federal Reserve Bank of Minneapolis.
- SCHORFHEIDE, F., AND D. SONG (2021): “Real-Time Forecasting with a (Standard) Mixed-Frequency VAR During a Pandemic,” NBER Working Papers 29535, National Bureau of Economic Research, Inc.
- SCHÜLER, Y. (2024): “Filtering Economic Time Series: On the Cyclical Properties of Hamilton’s Regression Filter and the Hodrick-Prescott Filter,” *Review of Economic Dynamics*, 54.
- SIMS, C. A., J. H. STOCK, AND M. W. WATSON (1990): “Inference in Linear Time Series Models with Some Unit Roots,” *Econometrica*, 58(1), 113–144.
- SIMS, C. A., AND T. ZHA (1998): “Bayesian Methods for Dynamic Multivariate Models,” *International Economic Review*, 39(4), 949–968.

- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- STOCK, J. H., AND M. W. WATSON (1988): “Testing for Common Trends,” *Journal of the American Statistical Association*, 83, 1097–1107.
- (1999): “Business Cycle Fluctuations in US Macroeconomic Time Series,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 1, pp. 3–64. Elsevier.
- (2003): “Has the Business Cycle Changed and Why?,” in *NBER Macroeconomics Annual 2002, Volume 17*, NBER Chapters, pp. 159–230. National Bureau of Economic Research, Inc.
- (2012): “Disentangling the Channels of the 2007-09 Recession,” *Brookings Papers on Economic Activity*, 43(1), 81–156.
- (2016): “Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and H. Uhlig, vol. 2 of *Handbook of Macroeconomics*, chap. 8, pp. 415–525. Elsevier.
- (2019): “Slack and Cyclically Sensitive Inflation,” NBER Working Papers 25987, National Bureau of Economic Research, Inc.
- UHLIG, H. (2004): “What Moves GNP?,” Econometric Society 2004 North American Winter Meetings 636, Econometric Society.
- URIBE, M. (2022): “The Neo-Fisher Effect: Econometric Evidence from Empirical and Optimizing Models,” *American Economic Journal: Macroeconomics*, 14(3), 133–162.
- WATSON, M. W. (1986): “Univariate Detrending Methods with Stochastic Trends,” *Journal of Monetary Economics*, 18(1), 49–75.

APPENDIX A. DATA

Most series are sourced from the Federal Reserve Economic Database (FRED). Exceptions include TFP from Fernald (2012), macroeconomic uncertainty from Jurado, Ludvigson, and Ng (2015), and macroeconomic factors from McCracken and Ng (2020). Tables 4 and 5 document the original series and the transformations applied prior to inclusion in the VARs.

TABLE 4. Data sources

Data	Mnemonic	Symbol
<i>GDP and subcomponents</i>		
Real GDP per capita	A939RX0Q048SBEA	Y
GDP: Implicit price deflator	GDPDEF	P
GDP share: Consumption expenditures on nondurable goods	DNDGRE1Q156NBEA	CNDY
GDP share: Consumption expenditures on services	DSERRE1Q156NBEA	CSY
GDP share: Gross private domestic investment	A006RE1Q156NBEA	IY
GDP share: Expenditures on durable goods	DDURRE1Q156NBEA	DGY
GDP share: Government consumption and investment	A822RE1Q156NBEA	GY
GDP share: Exports	B020RE1Q156NBEA	XY
GDP share: Imports	B021RE1Q156NBEA	MY
<i>Labor market</i>		
Nonfarm business sector: average weekly hours	PRS85006023	AH
Civilian non-institutional population	CNP16OV	POP
Civilian employment level	CE16OV	EMP
Civilian unemployment rate	UNRATE	UN
Nonfarm business sector: labor share	PRS85006173	LS
<i>Interest rates</i>		
Effective federal funds rate	FEDFUNDS	FF
Moody's Baa corporate bond spread relative to 10-year treasuries	BAA10YM	SPR
<i>Productivity</i>		
Fernald's utilization-adjusted TFP (log difference)	—	DTFP
<i>Cyclical measures and controls</i>		
Capacity utilization (manufacturing)	CAPUTLB00004SQ	UT
Jurado et al.'s macroeconomic uncertainty (one-year-ahead)	—	MU
McCracken and Ng's macroeconomic factors	—	F1, F2
Congressional Budget Office output gap	GDPC1_GDPPOT	OG

Notes. Mnemonics refer to FRED series identifiers. Symbols represent the variables in Table 5.

TABLE 5. Variables entering the VARs

Variable	Definition
GDP	$y = 100 \times \log(Y)$
Inflation rate	$\pi = 400 \times \log[P/P(-1)]$
Consumption	$c = 100 \times \log[(CNDY + CSY) \times Y]$
Investment	$i = 100 \times \log[(IY + DGY) \times Y]$
Government expenditures	$g = 100 \times \log(GY \times Y)$
Exports	$x = 100 \times \log(XY \times Y)$
Imports	$m = 100 \times \log(MY \times Y)$
Hours worked	$h = 100 \times \log(AH \times EMP/\overline{POP})$
Unemployment rate	$u = UN$
Labor share	$ls = LS$
Interest rate	$r = FF$
Credit spread	$spr = SPR$
TFP	$tfp = \log[\text{cumulative sum}(DTFP/4)]$
Uncertainty	$mu = MU$
Macroeconomic factors	$f_1, f_2 = F1, F2$
Output gap	$og = OG$


Notes. As recommended by Edge and Gurkaynak (2010), we apply a very gradual HP filter to smooth population when computing per-capita hours worked, eliminating artificial spikes from Census revisions. We set the smoothing parameter to 100,000 and retain the trend, denoted $\overline{POP}$.

ONLINE APPENDIX
“MEASURING BUSINESS CYCLES USING VARs”

ONLINE APPENDIX I. COMPARISON TO ALTERNATIVES


Figures 8 to 10 compare our baseline two-shock VAR cycles with standard alternatives — the HP filter with $\lambda = 1,600$, the Christiano and Fitzgerald (2003) bandpass filter targeting the 6-32-quarter frequency band, and the Hamilton filter with $p = 4$ and $h = 8$. As in Section 3, we focus on pre-COVID cycles. Summary business-cycle statistics are reported in Table 1.

FIGURE 8. Cyclical estimates — VAR vs. HP filter




Notes. Dashed blue lines depict posterior median VAR cycles, with shaded bands indicating a 68% credible interval. Solid red lines show HP cycles ($\lambda = 1,600$). See Table 1 notes.

FIGURE 9. Cyclical estimates — VAR vs. bandpass filter



Notes. Dashed blue lines depict posterior median VAR cycles, with shaded bands indicating a 68% credible interval. Solid red lines show bandpass cycles (6-32 quarters). See Table 1 notes.

FIGURE 10. Cyclical estimates — VAR vs. Hamilton filter




Notes. Dashed blue lines depict posterior median VAR cycles, with shaded bands indicating a 68% credible interval. Solid red lines show Hamilton cycles ($p = 4, h = 8$). See Table 1 notes.

ONLINE APPENDIX II. ROBUSTNESS


Figures 11 to 28 compare our baseline VAR cycles with those obtained under a wide range of alternative specifications, including variations in lag order, shock identification, system composition, trend treatment, and sample split, as detailed in Section 4.1. As in the main text, we focus on pre-COVID cycles.

FIGURE 11. Robustness — VAR(2) vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles from a BVAR(2) model, while dashed black lines depict baseline cycles.

FIGURE 12. Robustness — VAR(8) vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles from a BVAR(8) model, while dashed black lines depict baseline cycles.

FIGURE 13. Robustness — Consumption excluded from targets vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when consumption is excluded from target variables, while dashed black lines depict baseline cycles.

FIGURE 14. Robustness — Weighted variance target vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when using a weighted variance target, while dashed black lines depict baseline cycles.

FIGURE 15. Robustness — Cyclical VAR vs. baseline (1/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with the output gap, capacity utilization, and macroeconomic uncertainty, while dashed black lines depict baseline cycles.

FIGURE 16. Robustness — Cyclical VAR vs. baseline (2/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with the output gap, capacity utilization, and macroeconomic uncertainty, while dashed black lines depict baseline cycles.

FIGURE 17. Robustness — Factor-augmented VAR vs. baseline (1/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with the first two McCracken and Ng (2020) macroeconomic factors, while dashed black lines depict baseline cycles.

FIGURE 18. Robustness — Factor-augmented VAR vs. baseline (2/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with the first two McCracken and Ng (2020) macroeconomic factors, while dashed black lines depict baseline cycles.

FIGURE 19. Robustness — Large VAR vs. baseline (1/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with government consumption, exports, and imports, while dashed black lines depict baseline cycles.

FIGURE 20. Robustness — Large VAR vs. baseline (2/2)




Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with government consumption, exports, and imports, while dashed black lines depict baseline cycles.

FIGURE 21. Robustness — Small VAR vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when restricting the VAR to GDP, investment, consumption, hours worked, and unemployment, while dashed black lines depict baseline cycles.

FIGURE 22. Robustness — Discarding unstable draws vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when replacing the detrending step by the discarding of unstable posterior draws as in Angeletos, Collard, and Dellas (2020), while dashed black lines depict baseline cycles.


FIGURE 23. Robustness — Deterministic time trend vs. baseline



Notes. See Figure 1 notes. Blue lines show cycles when augmenting the VAR with a variable-specific cubic time trend, while dashed black lines depict baseline cycles.


FIGURE 24. Robustness — Eigenvalue stability threshold $\bar{\rho} = 0.95$ vs. baseline

Notes. See Figure 1 notes. Blue lines show cycles when using the eigenvalue stability threshold $\bar{\rho} = 0.95$, while dashed black lines depict baseline cycles.

FIGURE 25. Robustness — Eigenvalue stability threshold $\bar{\rho} = 1.00$ vs. baseline


Notes. See Figure 1 notes. Blue lines show cycles when using the eigenvalue stability threshold $\bar{\rho} = 1.00$, while dashed black lines depict baseline cycles.

FIGURE 26. Robustness — Long-run prior vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when estimating the VAR with the long-run prior of Giannone, Lenza, and Primiceri (2019), while dashed black lines depict baseline cycles.

FIGURE 27. Robustness — Sims-Zha prior vs. baseline




Notes. See Figure 1 notes. Blue lines show cycles when estimating the VAR with the Sims and Zha (1998) dummy prior, while dashed black lines depict baseline cycles.

FIGURE 28. Robustness — Split sample vs. baseline



Notes. See Figure 1 notes. Blue and red lines show cycles from BVAR(2) models estimated on the 1955-1983 and 1984-2019 subsamples, while dashed black lines depict baseline cycles.

FIGURE 29. Robustness — Correlations with baseline as function of stability threshold $\bar{\rho}$



Notes. In each panel, the x -axis reports the eigenvalue stability threshold $\bar{\rho}$ and the y -axis shows the correlation between the implied VAR cycles and the baseline estimates using $\bar{\rho} = 0.99$. See Table 1 notes for details on the VAR.



BANQUE CENTRALE DU LUXEMBOURG

EUROSYSTEME

2, boulevard Royal
L-2983 Luxembourg

Tél.: +352 4774-1
Fax: +352 4774 4910

www.bcl.lu • info@bcl.lu