



Principal component analysis for a mix of stationary and nonstationary variables[★]

James D. Hamilton^{a,*}, Xinwei Ma^a, Jin Xi^b

^a University of California at San Diego, United States

^b Center for Forecasting Science, Academy of Mathematics and Systems Science, Chinese Academy of Sciences, China

ARTICLE INFO

Keywords:

Principal components

Nonstationary

Economic activity indexes

ABSTRACT

This paper develops a procedure for uncovering the common cyclical factors that drive a mix of stationary and nonstationary variables. The method does not require knowing which variables are nonstationary or the nature of the nonstationarity. An application to the FRED-MD macroeconomic dataset demonstrates that the approach offers similar benefits to those of traditional principal component analysis with some added advantages.

1. Introduction

Principal component analysis (PCA) has become a key tool for building dynamic models of vector time series with a large cross-sectional dimension. A typical application first subtracts the sample mean from each variable and divides the demeaned variables by their sample standard deviations. PCA finds linear combinations of the standardized variables that have maximum sample variance. These principal components are then used to build dynamic models for each individual series. For surveys of PCA and its usefulness in economics see [Bai and Ng \(2008\)](#) and [Stock and Watson \(2016\)](#).

One difficulty with PCA is that many of the time series encountered in economics and finance are nonstationary. For a nonstationary variable, the population mean is undefined and the sample standard deviation diverges to infinity as the number of time-series observations gets large. [Onatski and Wang \(2021\)](#) detailed some of the problems that can arise from trying to apply PCA to nonstationary data. The typical solution to this problem is for researchers to examine each series individually by hand to determine the transformation that needs to be applied to that series before calculating principal components of the set of variables.

This approach has two shortcomings. First, while for some variables it may be fairly clear what transformation is necessary to achieve stationarity, for others it is far from obvious. For example, interest rates have a strong downward trend since 1980. Should yields be treated as stationary? If nonstationary, should we take their first differences or deviations from a time trend before performing PCA? Some authors work with first differences of inflation and the unemployment rate while others leave those variables as is. Many decisions like these have to be made before applying PCA to large data sets. In this paper, we propose an automatic procedure that allows the researcher to avoid these judgment calls.

A second fundamental problem is the appropriateness of the methodology itself. Suppose we somehow overcame the first problem and knew for certain the true nature of the trend in each individual series. Suppose for illustration we knew correctly that the first

[★] An earlier version of this paper was circulated under the title, “Principal Component Analysis for Nonstationary Time Series.” We thank Serena Ng, Mark Watson, Ken West, and anonymous referees and editors for helpful comments. Xi gratefully acknowledges financial support from the National Natural Science Foundation of China (No. 72503234). Authors are listed in alphabetical order; all contributed equally to this work.

^{*} Corresponding author.

E-mail addresses: jhamilton@ucsd.edu (J.D. Hamilton), x1ma@ucsd.edu (X. Ma), xijin@amss.ac.cn (J. Xi).

variable y_{1t} is a stationary $AR(1)$ process with autoregressive coefficient $\rho = 0.99$ while the second variable y_{2t} is a random walk. The currently prescribed procedure would instruct the researcher to use the first variable as is and the second variable in first differences. But while the levels of y_{1t} and y_{2t} exhibit very similar properties, the level of y_{1t} and first difference of y_{2t} are radically different. Should we expect that there is some linear combination of y_{1t} and Δy_{2t} that can summarize the common economic drivers behind the two variables? If differencing is the appropriate transformation for a random walk, it seems we should be using some similar transformation for an $AR(1)$ process whose root is close but not quite equal to unity.

This paper proposes an approach to PCA that solves these problems. We first note as in [Hamilton \(2018\)](#) that it is possible to use OLS to estimate an h -period-ahead forecast of the level of any variable as a linear function of its own lags without knowing the nature of the nonstationarity. Moreover, the errors from these linear forecasts are stationary for a broad class of underlying nonstationary processes. Our proposal is to use PCA to identify common factors behind the forecast errors. Specifically, we estimate an OLS regression of y_{it} on $(1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$, where $i = 1, \dots, N$ represents the index of the variable, h is the forecasting horizon, and p is the number of lags used for the forecast. We then calculate principal components of the residuals. This approach solves both of the problems identified above. The procedure is fully automatic and treats every variable in the same way. The transformation of y_{it} is a continuous function of the estimated autoregressive coefficients and implies no discontinuity in the way that persistent stationary series are treated as the largest autoregressive root tends toward unity.

[Hamilton \(2018\)](#) suggested that a two-year-ahead forecast error could be interpreted as the stationary cyclical component of the series. Our empirical application follows this suggestion and uses $h = 24$ for monthly data to uncover the common cyclical factors behind a range of economic and financial indicators. We demonstrate that using $h = 24$ also offers a practical solution to the problem of outliers which has posed a severe challenge to using PCA on macroeconomic observations in 2020.

[Bai and Ng \(2004\)](#) and [Barigozzi et al. \(2021\)](#) developed methods that are suitable when the observed data are an unknown mix of $I(0)$ and $I(1)$ variables. Their approaches start by applying PCA to first-differences of all of the variables whether stationary or not. [Bai and Ng \(2004\)](#) estimated factors for the original data by accumulating factors estimated from the differenced data. [Barigozzi et al. \(2021\)](#) estimated factors by applying the loadings estimated from differenced data to the detrended levels of the original data. [Barigozzi et al. \(2021\)](#) showed that their procedure can estimate the set of factors that account for common stochastic trends in the original data. In contrast to these approaches, the goal of our analysis is to uncover the cyclical components that are driving both the stationary and nonstationary variables.

Our results are also useful for estimation of dynamic factor models. One common approach to estimating dynamic factor models is to apply a transformation that is believed to make the individual series stationary, calculate principal components of the transformed data, and then fit a vector autoregression to the estimated factors ([Stock and Watson, 2002, 2016](#); [Bai and Ng, 2002, 2006](#)). In contrast, our approach allows the researcher to implement the initial transformations without the need for subjective judgments about each individual series.

[Section 2](#) discusses the use of forecasting regressions to isolate a stationary component of a possibly nonstationary time series. [Section 3](#) describes how PCA could be used to uncover the common factors behind the forecast errors if we somehow could know the true process followed by each variable. [Section 4](#) analyzes the case when PCA is applied to the residuals from estimated OLS regressions and establishes consistency of the method. Simulations in [Section 5](#) report results for a variety of possible data-generating processes. We find that using two-year-ahead forecast errors works well when we have 50 years of data and that $h = 1$ performs quite well even in much shorter samples. [Section 6](#) illustrates the promise of our approach in an empirical analysis using the FRED-MD large macroeconomics data set. Proofs of theorems and additional simulations and empirical results are provided in an online supplement.

2. Isolating a stationary component from a nonstationary series

Let y_{it} denote the level of a possibly nonstationary variable at time t . Collect its p most recent values as of date $t - h$ along with a constant term in a vector $z_{i,t-h} = (1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$. The population linear projection of y_{it} on $z_{i,t-h}$ is given by $\mathbb{P}(y_{it}|z_{i,t-h}) = \alpha'_{i0} z_{i,t-h}$ where the coefficient α_{i0} is defined as the vector that minimizes the expected squared error of a forecast of y_{it} based on a linear function of $z_{i,t-h}$:

$$\alpha_{i0} = \arg \min_{\alpha} E(y_{it} - \alpha' z_{i,t-h})^2. \quad (1)$$

For example, if y_{it} is covariance stationary with $E(z_{i,t-h} z'_{i,t-h})$ nonsingular, the population linear projection coefficient is given by

$$\alpha_{i0} = [E(z_{i,t-h} z'_{i,t-h})]^{-1} E(z_{i,t-h} y_{it}). \quad (2)$$

If y_{it} is ergodic for second moments, the population parameter α_{i0} can be consistently estimated by an OLS regression of y_{it} on $z_{i,t-h}$:

$$\hat{\alpha}_{i0} = \left[\sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \right]^{-1} \left[\sum_{t=1}^T z_{i,t-h} y_{it} \right] = \left[T^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \right]^{-1} \left[T^{-1} \sum_{t=1}^T z_{i,t-h} y_{it} \right] \xrightarrow{p} \alpha_{i0}.$$

Note that there is no assumption in the definition (1) nor in its stationary solution (2) that the process is linear or that it can be characterized by an ARMA representation of known order. There may be better forecasts of y_{it} that could be obtained using a nonlinear function or more lags like $y_{i,t-h-p}, y_{i,t-h-p-1}, \dots$. But there is an optimal forecast within the class of linear forecasts that use only p lags. For a covariance-stationary process, the optimal forecast within that class is characterized by (2).

A solution to (1) also exists if y_{it} is nonstationary but its first difference Δy_{it} is covariance stationary. This can be seen from the accounting identity that holds for any time series:

$$y_{it} = y_{i,t-h} + \Delta y_{i,t-h+1} + \Delta y_{i,t-h+2} + \dots + \Delta y_{it}. \quad (3)$$

For an $I(1)$ process,¹

$$\mathbb{P}(y_{it}|z_{i,t-h}) = y_{i,t-h} + \mathbb{P}(\Delta y_{i,t-h+1}|z_{i,t-h}^\dagger) + \mathbb{P}(\Delta y_{i,t-h+2}|z_{i,t-h}^\dagger) + \dots + \mathbb{P}(\Delta y_{it}|z_{i,t-h}^\dagger) \quad (4)$$

for $z_{i,t-h} = (1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$, $z_{i,t-h}^\dagger = (1, \Delta y_{i,t-h}, \Delta y_{i,t-h-1}, \dots, \Delta y_{i,t-h-p+2})'$ and $\mathbb{P}(\Delta y_{i,t-h}|z_{i,t-h}^\dagger) = [E(\Delta y_{i,t-h} z_{i,t-h}^{\dagger'})] [E(z_{i,t-h}^\dagger z_{i,t-h}^{\dagger'})]^{-1} z_{i,t-h}^\dagger$. Moreover, the population linear projection coefficient α_{i0} can again be consistently estimated by an OLS regression of the level of y_{it} on $(1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$. The intuition for this is that an OLS levels regression minimizes the sample analog to the population minimization in (1):

$$\hat{\alpha}_i = \arg \min_{\alpha} T^{-1} \sum_{t=1}^T (y_{it} - \alpha' z_{i,t-h})^2. \quad (5)$$

If the OLS estimate $\hat{\alpha}_i$ picks up the unit root, the value of (5) converges to a finite number as T gets large. If it does not, the value of (5) would diverge to infinity. Thus in large samples, OLS estimation is dominated by the incentive inherent in the objective function of OLS to remove any nonstationary features of the data. More formally, the fitted values from the regression (5) are numerically identical to the fitted values from an OLS regression of y_{it} on $(y_{i,t-h}, z_{i,t-h}^{\dagger'})$. In the latter regression, the OLS coefficient on the lagged level $y_{i,t-h}$ will converge in probability to one whenever the dependent variable is an $I(1)$ process. In fact, the OLS estimate of this parameter is superconsistent, converging at rate T rather than at rate $\sqrt{T}$, and the asymptotic distribution of the other coefficients is identical to what would result if we regressed Δy_{it} on $z_{i,t-h}^\dagger$; see Hamilton (2018).

Hamilton (2018) showed that these results generalize to any nonstationary process that is covariance stationary around a deterministic polynomial in time t provided that the order of the polynomial is no greater than the number of lags p , or if the process is $I(d)$ with the d th difference of y_{it} covariance stationary for some d no greater than p . One does not need to know the order of the polynomial or the value of d in order to consistently estimate the population linear projection coefficient α_{i0} using a regression of the level of y_{it} on a constant and its p most recent levels as of date $t-h$.

Hamilton (2018) also showed that the residual from the population linear projection, $c_{it} = y_{it} - \alpha_{i0}' z_{i,t-h}$, is covariance stationary for any of the above nonstationary processes. For example, for an $I(1)$ process, we see from (3) and (4) that

$$\begin{aligned} c_{it} &= y_{it} - \mathbb{P}(y_{it}|z_{i,t-h}) \\ &= [\Delta y_{i,t-h+1} - \mathbb{P}(\Delta y_{i,t-h+1}|z_{i,t-h}^\dagger)] + [\Delta y_{i,t-h+2} - \mathbb{P}(\Delta y_{i,t-h+2}|z_{i,t-h}^\dagger)] + \\ &\quad \dots + [\Delta y_{it} - \mathbb{P}(\Delta y_{it}|z_{i,t-h}^\dagger)] \end{aligned}$$

is covariance stationary for any finite h . It is thus possible to isolate a stationary component of a large range of nonstationary processes from the residuals of an OLS regression of y_{it} on $(1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$ without knowing whether the variable y_{it} is stationary or nonstationary.²

It is instructive to compare this approach with other common ways of thinking about the trend in a nonstationary series. If the i th observed variable is a deterministic function of time plus a zero-mean stationary ergodic process, $y_{it} = \delta_i^{\det}(t) + c_{it}^{\det}$, we typically could describe the deterministic time trend as the limit of a forecast that would have been made in the arbitrarily distant past:

$$\delta_i^{\det}(t) = \lim_{h \rightarrow \infty} \lim_{p \rightarrow \infty} E(y_{it} | y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1}).$$

By contrast, if the first difference of the i th variable is a zero-mean stationary process, Beveridge and Nelson (1981) suggested that we think of the trend as the forecast of the variable in the arbitrarily distant future. They decomposed $y_{it} = \delta_{it}^{BN} + c_{it}^{BN}$ where

$$\delta_{it}^{BN} = \lim_{h \rightarrow \infty} \lim_{p \rightarrow \infty} E(y_{i,t+h} | y_{it}, y_{i,t-1}, \dots, y_{i,t-p+1}).$$

While these concepts of trend have some appeal, they have the significant practical drawback that both are based on the properties of forecasts at infinite horizons. They thus depend on conjectures of what happens at infinity, conjectures that are impossible to verify on the basis of a finite sample of observations.

¹ Note that the space spanned by $z_{i,t-h}$ is the same as the space spanned by $(z_{i,t-h}^\dagger, y_{i,t-h})$ which is bigger than the space spanned by $z_{i,t-h}^\dagger$. However, the coefficient on $y_{i,t-h}$ in the projection $\mathbb{P}(\Delta y_{it}|z_{i,t-h}^\dagger, y_{i,t-h})$ is zero, because multiplying the $I(1)$ variable $y_{i,t-h}$ by any nonzero value would produce a forecast of y_{it} that has infinite MSE. For this reason $\mathbb{P}(\Delta y_{it}|z_{i,t-h}) = \mathbb{P}(\Delta y_{it}|z_{i,t-h}^\dagger)$.

² Specifically, if either: (i) y_{it} is stationary around a deterministic polynomial function of time of order $d_i \leq p$ satisfying

$$T^{-1/2} \sum_{s=1}^{[Tr]} (y_{it} - \delta_{i0} - \delta_{i1}t - \delta_{i2}t^2 - \dots - \delta_{id_i}t^{d_i}) \xrightarrow{d} \omega_i W_i(r)$$

where $[Tr]$ denotes the largest integer no greater than Tr , $W_i(r)$ denotes standard Brownian motion, and $\xrightarrow{d}$ denotes convergence in distribution; or alternatively if (ii) d_i differences of y_{it} are stationary for some $d_i \leq p$ satisfying

$$T^{-1/2} \sum_{s=1}^{[Tr]} (\Delta^{d_i} y_{it} - \mu_i) \xrightarrow{d} \omega_i W_i(r);$$

then Hamilton (2018) showed that c_{it} is stationary and that the estimated coefficient $\hat{\alpha}_i$ from an OLS regression gives a consistent estimate of the population parameter α_{i0} .

By contrast, forecasts at a two-year horizon are something that can be reasonably investigated without auxiliary assumptions in samples of typical size. If we choose h to correspond to a two-year horizon, [Hamilton \(2018\)](#) argued that the decomposition

$$y_{it} = \mathbb{P}(y_{it}|1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1}) + c_{it} \quad (6)$$

can be viewed as a way to implement traditional approaches to trend-cycle decomposition that is practical and robust. Moreover, the primary reason we would go wrong in making a two-year-ahead forecast of most economic time series is due to unforeseen cyclical changes. For example, the value of y_{it} will be significantly below $\mathbb{P}(y_{it}|z_{i,t-h})$ if the economy goes into a recession after $t-h$ and significantly above the forecast if recovery from a downturn is more robust than expected. For this reason, [Hamilton \(2018\)](#) proposed to use h corresponding to a two-year horizon as a way to isolate the stationary cyclical component of the i th variable, for example, $h = 24$ for monthly data.

Another important benefit of setting $h = 24$ is that it offers a solution to the huge practical problem of dealing with outliers. To see why this is the case, suppose for illustration that the i th variable is characterized by a random walk: $y_{it} = y_{i,t-1} + \varepsilon_{it}$. For a monthly random walk, the residual from the two-year population linear projection is given by $c_{it} = \sum_{s=0}^{23} \varepsilon_{i,t-s}$. From the Central Limit Theorem, c_{it} typically exhibits much less kurtosis than the one-period-ahead forecast error ε_{it} .³ We find in our empirical application in [Section 6](#) that this feature is extremely helpful in dealing with outliers such as those seen during the COVID-19 pandemic in 2020.

Notwithstanding, these benefits do not come without a cost. The larger the value of h , the more serial correlation there will be in c_{it} , and the more prone PCA can be in small samples to the spurious factor problem identified by [Onatski and Wang \(2021\)](#). Our simulations in [Section 5](#) suggest that if we have 50 years of data, using two-year-ahead forecast errors works pretty well. For shorter samples, the spurious factor problem can be more serious, and our recommendation for those applications is to look for common factors in the one-year-ahead ($h = 12$) or one-month-ahead ($h = 1$) forecasts. In the latter case, additional correction for outliers is likely needed.

To summarize, our procedure is to estimate the same regression for every variable y_{it} , regardless of whether we think it is stationary and without making any conjecture about the nature of any nonstationarity. To allow for persistent seasonal components in y_{it} , we recommend choosing p to be the number of observations in a year. Our procedure estimates the following regression by OLS for every variable,

$$y_{it} = k_i + \alpha_{i1}y_{i,t-h} + \alpha_{i2}y_{i,t-h-1} + \dots + \alpha_{ip}y_{i,t-h-p+1} + c_{it}, \quad (7)$$

with $h = 8$ and $p = 4$ for quarterly data and $h = 24$ and $p = 12$ for monthly data. We will refer to the residual from the estimated regression $\hat{c}_{it}$ as the OLS residual for variable y_{it} and the residual from the population linear projection c_{it} as the true cyclical component. The value $\hat{c}_{it}$ is a consistent estimate of c_{it} , and the true value c_{it} is stationary as long as any nonstationarity in y_{it} is characterized by either a polynomial time trend of order d_i or an $I(d_i)$ process with $d_i \leq p$. Our procedure is to perform PCA on the regression residuals $\{\hat{c}_{1t}, \dots, \hat{c}_{Nt}\}$.

One practical decision is whether a nonlinear transformation of the raw data is necessary for $\Delta^{d_i}y_{it}$ to be stationary for some d_i . If taking first differences of the log is the correct way to produce a stationary series, then taking the change in the level would not produce a stationary series. We recommend using logs for variables like output or prices which are usually described in terms of growth rates. For such variables we use the log of the level of the variable, $y_{it} = \log Y_{it}$ as the variable in the regression (7). For variables like interest rates or the unemployment rate that are already quoted in percentage terms, we use the raw data $y_{it} = Y_{it}$ in the regression.

The true cyclical component c_{it} has mean zero and is stationary for a wide range of processes. However, the population value of α_{i0} is not known but must be estimated by regression. In [Section 3](#) we characterize our assumptions about the factor structure that we hypothesize describes the true values of c_{it} , and use standard results to establish that these population factors could be consistently estimated if the true values of c_{it} were observed without error. [Section 4](#) considers the case when we do not know the value of d_i for each series, do not know whether it is stationary or characterized by a deterministic time trend or an $I(d_i)$ process, and the c_{it} are not observed. In that section we analyze the consequences of performing PCA on the estimated OLS residuals $\hat{c}_{it}$.

³ For example, if $c_{it}^{(h)} = \varepsilon_{it} + \dots + \varepsilon_{i,t-h+1}$ and ε_{it} is i.i.d. with $E(\varepsilon_{it}) = 0$, $E(\varepsilon_{it}^2) = \sigma^2$, and $E(\varepsilon_{it}^4) = \kappa$, then $E[(c_{it}^{(h)})^2] = h\sigma^2$ and $E[(c_{it}^{(h)})^4] = h\kappa + 3h(h-1)\sigma^4$. The excess kurtosis of $c_{it}^{(h)}$ is then

$$\frac{E[(c_{it}^{(h)})^4]}{\{E[(c_{it}^{(h)})^2]\}^2} - 3 = \frac{h\kappa + 3h(h-1)\sigma^4}{h^2\sigma^4} - 3 = \frac{[\kappa/\sigma^4] - 3}{h}.$$

If ε_{it} has positive excess kurtosis ($\kappa/\sigma^4 > 3$), then the excess kurtosis of $c_{it}^{(h)}$ declines monotonically to zero as h increases. For a numerical illustration, suppose ε_{it} comes from a mixture of two Normals, where $\varepsilon_{it} \sim N(0, \sigma_1^2)$ with probability p and $N(0, \sigma_2^2)$ with probability $1-p$. Then $c_{it}^{(h)} = \sum_{s=0}^{h-1} \varepsilon_{i,t-s}$ is a mixture of $h+1$ Normals, distributed

$$N(0, k\sigma_1^2 + (h-k)\sigma_2^2) \text{ with probability } \binom{h}{k} p^k (1-p)^{h-k}$$

and variance $E[c_{it}^{(h)}]^2 = V^{(h)} = h[p\sigma_1^2 + (1-p)\sigma_2^2]$. For $\sigma_1 = 1$, $\sigma_2 = 5$, and $p = .9$, the probability that $c_{it}^{(h)}$ exceeds four of its standard deviations, that is, $\text{Prob}(|c_{it}^{(h)}| > 4\sqrt{V^{(h)}})$, is 1.4% for $h = 1$, 0.1% for $h = 12$, and 0.005% for $h = 24$.

3. Principal component analysis when the true cyclical component is observed

In the previous section we defined the true cyclical component c_{it} to be the residual from a population linear projection of y_{it} on $(1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$, and noted that c_{it} is stationary for a broad class of possible processes. In this section we provide sufficient conditions under which the true cyclical components for a collection of N different variables would have a factor structure that could be consistently estimated using PCA if we observed the true value of c_{it} for each variable. The set-up and results in this section closely follow [Stock and Watson \(2002\)](#).

3.1. Assumed factor structure of the true cyclical components

Collect the true cyclical components for the N different series at time t in an $(N \times 1)$ vector $C_t = (c_{1t}, \dots, c_{Nt})'$. We postulate that these are characterized by a factor structure of the form

$$C_t = \underset{(N \times 1)}{\Lambda} \underset{(N \times r)}{F_t} + \underset{(N \times 1)}{e_t}. \quad (8)$$

The number of latent factors r is much less than the number of variables N , but the r factors are assumed to account for most of the variance of C_t in a sense made formal below. Since the factors are unobserved, $C_t = \Lambda H^{-1} H F_t + e_t$ would imply the identical observable model as (8). Thus some normalizations are necessary in order to talk about consistently estimating the j th factor f_{jt} . In empirical estimation, practitioners typically resolve this ambiguity by estimating the j th column of Λ by the eigenvector associated with the j th largest eigenvalue of $T^{-1} \sum_{t=1}^T C_t C_t'$. Note that such a procedure implies a normalization in which the columns of Λ are orthogonal to each other and the elements of F_t are uncorrelated with each other and ordered by the size of their variance. We follow [Stock and Watson \(2002\)](#) in how to characterize these conventions as the cross-section dimension N and time-series dimension T get large.⁴

Assumption 1 (factor structure).

1. $(\Lambda' \Lambda / N) \rightarrow I_r$.
2. $E[F_t F_t'] = \Omega_{FF}$, where Ω_{FF} is a diagonal matrix with $\omega_{ii} > \omega_{jj} > 0$ for $i < j$.
3. $|\lambda_{ij}| \leq \bar{\lambda} < \infty$.
4. $T^{-1} \sum_t F_t F_t' \xrightarrow{p} \Omega_{FF}$.

In addition to implementing the property that eigenvectors of a symmetric matrix are orthogonal, [Assumption 1\(I\)](#) requires that each factor makes a nonnegligible contribution to the average variance of c_{it} across i . That is, if we were to imagine adding more variables (increasing N) with $\lambda_{ij} = 0$ for all i greater than some fixed N_0 , then [Assumption 1\(I\)](#) could not hold. Likewise [Assumption 1\(2\)](#) and (4) require that each factor continues to matter as the number of time-series observations T grows. These conditions are consistent with serial dependence of the factors, but rely on the fact that C_t is stationary.

Let γ denote an $(N \times 1)$ vector and $\Gamma = \{\gamma : \gamma' \gamma / N = 1\}$. Note that if γ were the j th column of Λ , the scalar $\gamma' \Lambda F_t / N$ would converge to f_{jt} and $(N^2 T)^{-1} \sum_{t=1}^T \gamma' \Lambda F_t F_t' \Lambda' \gamma \xrightarrow{p} \omega_{jj}$. The assumption that the idiosyncratic elements e_t do not have a factor structure requires that there is no value of γ for which the analogous operation applied to e_t would lead to anything other than zero: $\sup_{\gamma \in \Gamma} (N^2 T)^{-1} \sum_{t=1}^T \gamma' e_t e_t' \gamma \xrightarrow{p} 0$. [Stock and Watson \(2002\)](#) used the following assumptions to guarantee the absence of a factor structure in e_t .

Assumption 2 (moments of the errors).

- i $\lim_{N \rightarrow \infty} \sup_t \sum_{s=-\infty}^{\infty} |E[e_t' e_{t+s} / N]| < \infty$.
- ii $\lim_{N \rightarrow \infty} \sup_t N^{-1} \sum_{i=1}^N \sum_{j=1}^N |E[e_{it} e_{jt}]| < \infty$, where e_{it} denotes the i th element of e_t .
- iii $\lim_{N \rightarrow \infty} \sup_{t,s} N^{-1} \sum_{i=1}^N \sum_{j=1}^N |cov[e_{is} e_{it}, e_{js} e_{jt}]| < \infty$.

Some might be concerned that we have simply postulated that the true cyclical components are characterized by [Assumptions 1](#) and [2](#). But something very similar is done in traditional applications that assume conditions like these characterize specified stationary transformations of the original data. Indeed, insofar as the cyclical components have a common primitive definition in terms of h -period-ahead forecast errors, we would argue that these assumptions are easier to defend in our application than in many others.

3.2. Consequences of applying PCA to the true cyclical components

Recall that the $(N \times 1)$ vector of true cyclical components C_t is stationary and has population mean zero. If C_t was observed directly, its estimated sample variance matrix would be $S = T^{-1} \sum_{t=1}^T C_t C_t'$ and a linear combination $\gamma' C_t$ for any $(N \times 1)$ vector γ would have sample variance $\gamma' S \gamma$. If C_t was observed, the first estimated principal component (denoted $\tilde{f}_{1t} = N^{-1} \tilde{\lambda}_1' C_t$) would be defined as the linear combination that has maximum sample variance subject to a normalization condition such as $\gamma \in \Gamma = \{\gamma : \gamma' \gamma / N = 1\}$:

$$\tilde{\lambda}_1 = \arg \sup_{\gamma \in \Gamma} \tilde{R}(\gamma) \quad (9)$$

⁴ See [Bai and Ng \(2013\)](#) and [Stock and Watson \(2016\)](#) for discussion of alternative normalizations.

$$\tilde{R}(\gamma) = (N^2 T)^{-1} \gamma' \sum_{t=1}^T C_t C_t' \gamma. \quad (10)$$

Note we are normalizing $\tilde{\lambda}_1' \tilde{\lambda}_1 / N = 1$ as we did asymptotically for the columns of Λ in [Assumption 1](#)(1). We also divide the sample variance of $\gamma' C_t$ by N^2 in anticipation of the result that $\tilde{R}(\tilde{\lambda}_1)$ will converge to a fixed constant as N and T grow. The solution to (9) is obtained by setting $\tilde{\lambda}_1$ proportional to the eigenvector of $S = T^{-1} \sum_{t=1}^T C_t C_t'$ associated with the largest eigenvalue. For example, if we calculated eigenvectors of this matrix using code that normalizes eigenvectors to have unit length and orders eigenvalues by decreasing size, $\tilde{\lambda}_1$ would be $\sqrt{N}$ times the first eigenvector. The largest eigenvalue of S is equal to $T^{-1} \sum_{t=1}^T \tilde{f}_{1t}^2$, the sample variance of the first principal component. The j th principal component $N^{-1} \tilde{\lambda}_j' C_t$ is found by maximizing $\tilde{R}(\gamma)$ subject to the constraint that γ is orthogonal to $\tilde{\lambda}_1, \dots, \tilde{\lambda}_{j-1}$. The solution for $\tilde{\lambda}_j$ is proportional to the eigenvector of S associated with the j th largest eigenvalue.

Alternatively, if we observed the true factors F_t and loadings Λ , we could calculate the component of the variance of $\gamma' C_t$ that is attributable to the r factors alone:

$$R^*(\gamma) = (N^2 T)^{-1} \gamma' \sum_{t=1}^T \Lambda F_t F_t' \Lambda' \gamma. \quad (11)$$

[Stock and Watson \(2002\)](#) showed that under [Assumptions 1](#) and [2](#), the maximum value for (10) (which is given by the largest eigenvalue of S) and the supremum of (11) over all $\gamma \in \Gamma$ converge in probability to the same number ω_{11} , which is the population variance of the first factor, and that $\tilde{\lambda}_j' C_t / N$ gives a consistent estimate of f_{jt} up to a sign. If we were to estimate $k > r$ principal components, the first r would consistently estimate f_{jt} up to a sign normalization and the last $k - r$ would asymptotically have zero variance. Their results are a special case of [Theorem 1](#) developed in the following section.

4. Principal component analysis when the cyclical component must be estimated

In this section we assume that we do not observe the true cyclical component c_{it} of series i but have an estimate $\hat{c}_{it} = c_{it} + \hat{v}_{it}$. Let $\hat{C}_t = (\hat{c}_{1t}, \dots, \hat{c}_{Nt})'$ and $\hat{V}_t = (\hat{v}_{1t}, \dots, \hat{v}_{Nt})'$. We investigate the properties of principal components calculated from the estimated cyclical components:

$$\hat{f}_{jt} = N^{-1} \hat{\lambda}_j' \hat{C}_t \quad (12)$$

$$\hat{\lambda}_j = \arg \sup_{\{\gamma \in \Gamma, \gamma' \hat{\lambda}_1 = \dots = \gamma' \hat{\lambda}_{j-1} = 0\}} \hat{R}(\gamma)$$

$$\hat{R}(\gamma) = (N^2 T)^{-1} \gamma' \sum_{t=1}^T \hat{C}_t \hat{C}_t' \gamma.$$

We first state high-level sufficient conditions under which PCA applied to the estimated cyclical components $\hat{C}_t$ gives consistent estimates of the true factors F_t . Let $\hat{v}_{it} = \hat{c}_{it} - c_{it}$ denote the difference between the estimated and true cyclical component of series i at date t . The conditions require N to grow with T so as to ensure that $N^{-1} \sum_{i=1}^N \hat{v}_{it}^2$ and $T^{-1} \sum_{t=1}^T \hat{v}_{it}^2$ are $o_p(1)$.

Assumption 3 (high-level conditions on $\hat{v}_{it}$). The cross-section dimension N grows with the number of time series observations T according to a function $N(T)$ such that for all $\delta, \epsilon > 0$ there exists a $T_0(\delta, \epsilon)$ such that for all $T > T_0(\delta, \epsilon)$:

$$Prob \left\{ \frac{1}{N(T)} \sum_{i=1}^{N(T)} \hat{v}_{it}^2 > \delta \right\} < \epsilon \quad (i)$$

$$Prob \left\{ \max_{1 \leq i \leq N(T)} \frac{1}{T} \sum_{t=1}^T \hat{v}_{it}^2 > \delta \right\} < \epsilon. \quad (ii)$$

The following result establishes that if the error in estimating the cyclical component satisfies [Assumption 3](#), then the results $(\hat{f}_{jt}, \hat{\lambda}_j, \hat{R}(\gamma))$ of applying PCA to the estimated cyclical components $\hat{C}_t$ give consistent estimates of the magnitudes that characterize the true cyclical components C_t . For a proof see online Appendix A.

Theorem 1. Suppose that C_t and e_t in [Eq. \(8\)](#) satisfy [Assumptions 1](#) and [2](#). Let $\hat{C}_t = C_t + \hat{V}_t$ for $\hat{V}_t = (\hat{v}_{1t}, \dots, \hat{v}_{Nt})'$ where $\hat{v}_{it}$ satisfy [Assumption 3](#). Let $\hat{f}_{1t}, \dots, \hat{f}_{kt}$ denote the first k estimated principal components of $\hat{C}_t$ ($\hat{f}_{jt} = \hat{\lambda}_j' \hat{C}_t / N$) with $k \geq r$ and let $\hat{F}_t = (\hat{f}_{1t}, \dots, \hat{f}_{rt})'$ ($r \times 1$)

and $\hat{\Lambda}_{(N \times r)} = \begin{bmatrix} \hat{\lambda}_1 & \dots & \hat{\lambda}_r \end{bmatrix}$. Then

1. $T^{-1} \sum_{t=1}^T \hat{f}_{jt}^2 = \hat{R}(\hat{\lambda}_j) \xrightarrow{p} \omega_{jj}$ for $j = 1, \dots, r$;
2. $T^{-1} \sum_{t=1}^T \hat{f}_{jt}^2 \xrightarrow{p} 0$ for $j = r + 1, \dots, k$;
3. $\hat{\Lambda}' \Lambda / N \xrightarrow{p} I_r$ where $\hat{\Lambda}$ is a diagonal matrix whose row j column j element is $+1$ if $\hat{\lambda}_j' \lambda_j > 0$ and -1 if $\hat{\lambda}_j' \lambda_j < 0$;
4. $\hat{\Lambda}' \hat{F}_t - F_t \xrightarrow{p} 0$.

Before presenting formal sufficient conditions for verifying [Assumption 3](#), we first discuss the intuition for why we might expect it to hold. For $z_{i,t-h} = (1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$, the true cyclical component c_{it} is the residual from a population linear projection of y_{it} on $z_{i,t-h}$ and $\hat{c}_{it}$ is the residual from the corresponding estimated regression:

$$c_{it} = y_{it} - \alpha'_{i0} z_{i,t-h}$$

$$\hat{c}_{it} = y_{it} - \hat{\alpha}'_i z_{i,t-h}$$

$$\hat{v}_{it} = \hat{c}_{it} - c_{it} = (\alpha_{i0} - \hat{\alpha}_i)' z_{i,t-h}$$

$$(\alpha_{i0} - \hat{\alpha}_i) = - \left(\sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \right)^{-1} \left(\sum_{t=1}^T z_{i,t-h} c_{it} \right) \quad (13)$$

$$\hat{v}_{it}^2 = (\alpha_{i0} - \hat{\alpha}_i)' z_{i,t-h} z'_{i,t-h} (\alpha_{i0} - \hat{\alpha}_i) \quad (14)$$

$$\sum_{t=1}^T \hat{v}_{it}^2 = (\alpha_{i0} - \hat{\alpha}_i)' \left(\sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \right) (\alpha_{i0} - \hat{\alpha}_i). \quad (15)$$

Expression (15) will be recognized as the OLS Wald statistic for testing the true null hypothesis $H_0 : \alpha_i = \alpha_{i0}$ multiplied by $\hat{\sigma}_i^2 = (T-k)^{-1} \sum_{t=1}^T \hat{c}_{it}^2$, the average squared regression residual. As we now sketch, this statistic would be expected to be $O_p(1)$ for a wide class of stationary and nonstationary processes. This would lead us to expect that the individual terms $\hat{v}_{it}^2$ should be $o_p(1)$, with some uniformity conditions across i then guaranteeing [Assumption 3\(i\)](#). Likewise if we divide (15) by T we should again get an $o_p(1)$ random variable, as required by [Assumption 3\(ii\)](#).

We now explore the intuition for why we would typically expect (15) to be $O_p(1)$. Note that the expression can be written

$$\begin{aligned} \sum_{t=1}^T \hat{v}_{it}^2 &= \left(\sum_{t=1}^T c_{it} z'_{i,t-h} \right) \left(\sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \right)^{-1} \left(\sum_{t=1}^T z_{i,t-h} c_{it} \right) \\ &= \left(\sum_{t=1}^T c_{it} z'_{i,t-h} \Upsilon_{iT}^{-1} \right) \left(\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \Upsilon_{iT}^{-1} \right)^{-1} \left(\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} c_{it} \right) \end{aligned} \quad (16)$$

where Υ_{iT}^{-1} could be any sequence of nonsingular matrices. Note that the researcher does not need to know the value of Υ_{iT} . This matrix is just a device to calculate the asymptotic properties of the left side of (16) (which does not depend on Υ_{iT}) under various possible forms of nonstationarity for z_{it} .

If y_{it} were stationary, we would use $\Upsilon_{iT}^{-1} = T^{-1/2} I_{p+1}$ to find the asymptotic distribution of $\sum_{t=1}^T \hat{v}_{it}^2$. In this case, $z_{i,t-h} c_{it}$ is a stationary random variable which by the definition of the true cyclical component c_{it} has population mean zero. In this case, a central limit theorem could be used to establish that $\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} c_{it} = T^{-1/2} \sum_{t=1}^T z_{i,t-h} c_{it}$ converges to a Normal distribution. Likewise in the stationary case,

$$\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \Upsilon_{iT}^{-1} = T^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \xrightarrow{p} E(z_{i,t-h} z'_{i,t-h}).$$

For this reason, $\sum_{t=1}^T \hat{v}_{it}^2$ would typically be $O_p(1)$ when y_{it} is stationary.

Consider next the case when $d_i = 1$ and Δy_{it} (denoted below by u_{it}) is a zero-mean stationary process. Taking a regression with $p = 3$ lags for illustration, an OLS regression of y_{it} on $z_{i,t-h} = (1, y_{i,t-h}, y_{i,t-h-1}, y_{i,t-h-2})'$ has the identical fitted values as an OLS regression of y_{it} on $R_i^{-1} z_{i,t-h} = (u_{i,t-h}, u_{i,t-h-1}, 1, y_{i,t-h})'$. In this case to calculate the asymptotic distribution of (16) we would take

$$\begin{aligned} \Upsilon_{iT}^{-1} &= \begin{bmatrix} T^{-1/2} I_3 & 0 \\ 0 & T^{-1} \end{bmatrix} R_i^{-1} \\ \Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \Upsilon_{iT}^{-1} &= \begin{bmatrix} T^{-1} \sum_{t=1}^T u_{i,t-h}^2 & T^{-1} \sum_{t=1}^T u_{i,t-h} u_{i,t-h-1} & T^{-1} \sum_{t=1}^T u_{i,t-h} & T^{-3/2} \sum_{t=1}^T u_{i,t-h} y_{i,t-h} \\ T^{-1} \sum_{t=1}^T u_{i,t-h-1} u_{i,t-h} & T^{-1} \sum_{t=1}^T u_{i,t-h-1}^2 & T^{-1} \sum_{t=1}^T u_{i,t-h-1} & T^{-3/2} \sum_{t=1}^T u_{i,t-h-1} y_{i,t-h} \\ T^{-1} \sum_{t=1}^T u_{i,t-h} & T^{-1} \sum_{t=1}^T u_{i,t-h-1} & 1 & T^{-3/2} \sum_{t=1}^T y_{i,t-h} \\ T^{-3/2} \sum_{t=1}^T y_{i,t-h} u_{i,t-h} & T^{-3/2} \sum_{t=1}^T y_{i,t-h} u_{i,t-h-1} & T^{-3/2} \sum_{t=1}^T y_{i,t-h} & T^{-2} \sum_{t=1}^T y_{i,t-h}^2 \end{bmatrix}. \end{aligned}$$

Under standard unit-root asymptotics (e.g., [Hamilton, 1994](#), p. 506), we expect $\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} z'_{i,t-h} \Upsilon_{iT}^{-1} \xrightarrow{d} Q_i$ with

$$Q_i = \begin{bmatrix} \gamma_{i0} & \gamma_{i1} & 0 & 0 \\ \gamma_{i1} & \gamma_{i0} & 0 & 0 \\ 0 & 0 & 1 & \omega_i \int_0^1 W_i(r) dr \\ 0 & 0 & \omega_i \int_0^1 W_i(r) dr & \omega_i^2 \int_0^1 [W_i(r)]^2 dr \end{bmatrix} \quad (17)$$

for $\gamma_{ij} = E(u_{it} u_{i,t-j})$, $\omega_i^2 = \sum_{j=-\infty}^{\infty} \gamma_{ij}$, and $W_i(r)$ standard Brownian motion. Similar tools can be used to establish that $\Upsilon_{iT}^{-1} \sum_{t=1}^T z_{i,t-h} c_{it} \xrightarrow{d} q_i$ and $\sum_{t=1}^T \hat{v}_{it}^2 \xrightarrow{d} q_i' Q_i^{-1} q_i \sim O_p(1)$.

If instead $E(\Delta y_{it}) = \mu_i \neq 0$, the level y_{it} would be dominated asymptotically by a deterministic time trend $\mu_i t$.⁵ In this case (and for general $p \geq 2$) we would use $R_i^{-1} z_{i,t-h} = (\Delta y_{i,t-h} - \mu_i, \Delta y_{i,t-h-1} - \mu_i, \dots, \Delta y_{i,t-h+2} - \mu_i, 1, y_{i,t-h})'$ and

$$Y_{iT}^{-1} = \begin{bmatrix} T^{-1/2} I_p & 0 \\ 0 & T^{-3/2} \end{bmatrix} R_i^{-1}$$

to establish that $\sum_{t=1}^T \hat{v}_{it}^2 \sim O_p(1)$. General cases for $d_i \leq p$ are examined in Hamilton (2018). Note that we do not need to know the value of d_i or Y_{iT} to estimate any magnitudes – in every case we are talking about a regression of y_{it} on the lagged levels $z_{i,t-h} = (1, y_{i,t-h}, y_{i,t-h-1}, \dots, y_{i,t-h-p+1})'$. Instead, Y_{iT} is just a device to analyze the properties if the levels regression is applied to a variety of different stationary and nonstationary processes.

We are now in a position to state formally sufficient conditions under which Assumption 3 would be satisfied, for which we let $\|X\|$ denote the Frobenius/Euclidean norm (the square root of the sums of squares of the elements of a vector or matrix X).⁶

Assumption 4 (sufficient conditions for Assumption 3). For each i , there exist some nonsingular matrix Y_{iT} and a possibly random matrix Q_i , such that

1. For some positive constant c_1 ,

$$\max_{1 \leq i \leq N} \max_{1 \leq t \leq T} E(\|\sqrt{T} Y_{iT}^{-1} z_{i,t-h}\|^2) \leq c_1 \quad \text{and} \quad \max_{1 \leq i \leq N} E(\|\sum_{t=1}^T Y_{iT}^{-1} z_{i,t-h} c_{it}\|^2) \leq c_1.$$

2. For some positive constants c_2, c_3, c_4 , and c_5 ,

$$\max_{1 \leq i \leq N} \text{Prob}(\rho_{\min}(Q_i) < \delta) \leq c_2 \exp\left\{-c_3 \left(\frac{1}{\delta}\right)^{1/c_4}\right\} \quad \text{for all } \delta \leq c_5,$$

where $\rho_{\min}$ denotes the minimum eigenvalue. In addition, $N \log^{2c_4}(N)/T \rightarrow 0$.

3. The cross-section dimension N grows with the number of time series observations T according to a function $N(T)$, such that

$$\text{Prob}\left(\max_{1 \leq i \leq N} \left\|\sum_{t=1}^T Y_{iT}^{-1} z_{i,t-h} z'_{i,t-h} Y_{iT}^{-1} - Q_i\right\| > c_6(N)\right) \rightarrow 0$$

for some positive sequence $c_6(N) = o(\log^{-c_4}(N))$.

Condition Assumption 4(1) is quite mild, as it only requires finite moments on the transformed regressors and the cyclical components after suitable normalization. For example, when Δy_{it} is a zero-mean stationary process, the first inequality requires Δy_{it} and $y_{it}/\sqrt{T}$ to have finite variance. Conditions Assumption 4(2-3) are key to our theoretical developments. Specifically, they posit that the matrix that gets inverted in (16) can be uniformly approximated by a possibly random limit Q_i , and that the minimum eigenvalue of the limiting matrix exhibits an exponential-type tail at zero.

Theorem 2. Assumption 4 implies Assumption 3.

Next we relate Assumption 4 to conditions for general linear processes.

Assumption 5 (linear process). There exist positive constants c_7, c_{11} , such that for each $i \in \{1, 2, \dots, N\}$ there exist a μ_i and an integer d_i for which $u_{it} = (1 - L)^{d_i} y_{it} - \mu_i = \sum_{l=0}^{\infty} \psi_{il} \eta_{i,t-l}$, and

1. $\max_{1 \leq i \leq N} |\psi_{it}| \leq c_7 t^{-c_8}$ with $c_8 > 2$.
2. $\min_{1 \leq i \leq N} |\sum_{l=0}^{\infty} \psi_{il}| \geq c_9 > 0$.
3. η_{it} is an i.i.d. process with zero mean, unit variance, and $\max_{1 \leq i \leq N} E(\eta_{it}^4) \leq c_{10}$.
4. Let $\tilde{Q}_i$ be the covariance matrix of $(u_{it}, u_{i,t-1}, \dots, u_{i,t-p+d_i+1})$. Then $\rho_{\min}(\tilde{Q}_i) \geq c_{11} > 0$.

Conditions Assumption 5(1) and (2) would be satisfied if u_{it} is any zero-mean stationary $AR(g)$ process for some g ,

$$u_{it} = \phi_{i1} u_{i,t-1} + \dots + \phi_{ig} u_{i,t-g} + \eta_{it} \quad \text{with } \|z\| > 1 \text{ for any } z \text{ satisfying } |1 - \phi_{i1} z - \dots - \phi_{ig} z^g| = 0.$$

Alternatively, Assumption 5(1) and (2) would also be satisfied for any $MA(g)$ process that has a representation that is not arbitrarily close to being noninvertible,

$$u_{it} = \eta_{it} + \theta_{i1} \eta_{i,t-1} + \dots + \theta_{ig} \eta_{i,t-g} \quad \text{with } |1 + \theta_{i1} + \dots + \theta_{ig}| > c_9.$$

Condition Assumption 5(2) defines the necessary degree of differencing d_i . For example, if $y_{it} = \psi_i(L) \eta_{it}$ with $\psi_{i\ell}$ satisfying $\psi_i(1) > c_9$, then the $MA(\infty)$ coefficients $\tilde{\psi}_{i\ell}$ of the first difference $(1 - L)y_{it} = (1 - L)\psi_i(L)\eta_{it} = \tilde{\psi}_i(L)\eta_{it}$ would not satisfy (2) because $\tilde{\psi}_i(1) = (1 - 1)\psi_i(1) = 0$.

⁵ That is, for any nonnegative integer v ,

$$T^{-(v+1)} \sum_{t=1}^T y_{it}^v - T^{-(v+1)} \sum_{t=1}^T (\mu_i t)^v \xrightarrow{p} 0$$

as $T \rightarrow \infty$ and $T^{-(v+1)} \sum_{t=1}^T (\mu_i t)^v \rightarrow \mu_i^v / (v + 1)$.

⁶ An earlier version of this paper avoided the requirement that T grows faster than N by assuming that the true cyclical component c_{it} is uniformly bounded.

Theorem 3. Suppose that $p \geq 1$ and [Assumption 5](#) is satisfied for every $i \in \{1, 2, \dots, N\}$ for some $d_i \in \{0, 1\}$, and that the true cyclical components c_{it} satisfy [Assumptions 1 and 2](#). If $N \log^{12}(N)/T \rightarrow 0$, then results (i)–(iv) of [Theorem 1](#) hold.

We now briefly outline our proof technique and main technical contribution. It is helpful to compare the current setting with the case where N is fixed and T grows. The fixed- N case is straightforward, as one can establish either convergence in probability to a nonrandom Q_i (under weak stationarity) or convergence in distribution to a random Q_i (as functionals of Brownian motion). This aligns with the reasoning provided earlier (see [\(17\)](#) and the related discussion). However, when N is allowed to grow, the theoretical analysis requires uniform control of the estimation error across $i = 1, \dots, N$, as implied by [Assumption 3](#). The core of our proof is to decompose the quantity of interest into a sum of i.i.d. components plus negligible remainder terms. We then apply the coupling method of [Komlós et al. \(1976\)](#) to establish a distributional approximation that is uniformly valid for $i = 1, \dots, N$.

While our proof of [Theorem 3](#) only covers the case when $d_i \leq 1$, we conjecture that a similar result could be obtained for any $d_i \leq p$. Extending the proof to this case requires verifying the three conditions in [Assumption 4](#). [Assumption 4\(1\)](#) requires only finite moments of the regressors and the cyclical components after suitable transformation and rescaling by the appropriate Y_{iT} . This is true for processes satisfying [Assumption 5](#) with $d_i = 2$, although the calculations demonstrating this would be more tedious than the $d_i = 1$ case we review. It is also possible to establish [Assumption 4\(3\)](#) for $d_i \geq 2$. The key step in our proof builds on the coupling result of [Komlós et al. \(1976\)](#), which provides a uniform (in both $1 \leq i \leq N$ and $1 \leq t \leq T$) strong approximation of a random walk by sums of independent Gaussian variables. This result can also be applied to provide a uniform (in $1 \leq i \leq N$) approximation of the Q_{iT} matrix using d_i -fold integrals of Brownian motion. As another extension, we show in [Appendix A.4](#) that the strong approximation result also holds for series that are local to unity.

The main technical challenge in extending [Theorem 3](#) to the case when $d_i \geq 2$ and the local-to-unity scenario is verifying [Assumption 4\(2\)](#). When $d_i \geq 2$ and $\mu_i = 0$, the limiting random matrix Q_i involves n -fold integrals of Brownian motion. Based on the literature of small ball probabilities for n -fold integrals of Brownian motion ([Chen and Li, 2003](#); [Gao et al., 2003](#)), we conjecture that [Assumption 4\(2\)](#) holds with $c_4 = 2 \max_{1 \leq i \leq N} d_i$, though we were unable to locate a concrete demonstration of this in the literature. Regarding the local-to-unity case, existing results on small ball probabilities for Ornstein-Uhlenbeck processes ([Ai, 2016](#)) suggest that [Assumption 4\(2\)](#) holds with $c_4 = 2$. Thus, although our current [Theorem 3](#) applies only to stationary and $I(1)$ processes, we expect that it extends to series with any order of integration and those that are local to unity, with the only change that the exponent on the log of N may differ.

Finally, we note that no additional proof is necessary to extend [Theorem 3](#) to the case when $2 \leq d_i \leq p$ if the number of series for which $d_i \geq 2$ does not increase with N or T .

5. Results from simulations

In this section we report results from applying our method in a variety of different settings. In these simulations we take the number of cross-section variables to be $N = 100$ and vary the number of time-series observations T from 100 to 1000. For each generated sample, we calculate the $(N \times N)$ correlation matrices of: (1) the raw data; (2) the OLS regression residuals for that data set; and (3) the true cyclical components for that data set implied by the particular process that was used to generate that sample. For each of these correlation matrices, let $\hat{\xi}_j$ denote the j th largest eigenvalue of the correlation matrix. The fraction of the variance of the sample explained by the j th principal component is

$$R_j^2 = \hat{\xi}_j / N. \quad (18)$$

We also calculated the number of factors that would be selected for that sample based on the IC_{p_2} criterion of [Bai and Ng \(2002\)](#) recommended by [Stock and Watson \(2016, p. 436\)](#),

$$r^* = \arg \left\{ \min_{r \in \{0, 1, \dots, r_0\}} \log \left(1 - \frac{\sum_{j=0}^r \hat{\xi}_j}{N} \right) + r \left(\frac{N+T}{NT} \right) \log(\min\{N, T\}) \right\}, \quad (19)$$

with $\hat{\xi}_0$ defined to be 0. We took $r_0 = 10$ and for each of the different cases generated 1000 different samples.

5.1. Mix of unrelated stationary and nonstationary variables

For our first example, half the variables are random walks and the other half are white noise,

$$y_{it} = \begin{cases} y_{it-1} + \varepsilon_{it} = \varepsilon_{it} + \varepsilon_{i,t-1} + \dots + \varepsilon_{i1} & \text{for } i = 1, 2, \dots, N/2 \\ \varepsilon_{it} & \text{for } i = (N/2) + 1, \dots, N \end{cases}$$

for $t = 1, \dots, T$. The innovations $\varepsilon_{it} \sim N(0, 1)$ are independent across all i and t . Thus each of the N variables is completely independent of the others and there is no factor structure in the true data-generating process.

The first two columns of [Table 1](#) report the results from calculating principal components of the raw data. In a sample of $T = 100$ observations, the first three principal components seem to account for 38% of the variance of the full set of $N = 100$ variables. This result is entirely spurious, and illustrates the cautions raised by [Onatski and Wang \(2021\)](#) about using PCA when some of the variables are nonstationary. The criterion (19) would always lead us incorrectly to conclude that there is more than one factor in a sample of size $T = 100$. This problem in the apparent number of factors gets even worse when the sample size increases. The latter is the expected result, since (19) is in the class of criteria for which [Onatski and Wang \(2021, p. 602\)](#) demonstrated that the number of factors selected diverges as N and T go to infinity.

Table 1
Mixture of independent random walks and white noise.

j	Raw data		c_t ($h=24$)		$\hat{c}_t$ ($h=24$)		$\hat{c}_t$ ($h=12$)		$\hat{c}_t$ ($h=8$)		$\hat{c}_t$ ($h=1$)	
	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$T = 100$												
0	—	0	—	0	—	1	—	13	—	96	—	100
1	22.6	0	16.0	0	11.9	26	10.3	48	8.4	4	4.1	0
2	9.9	67	12.0	6	9.4	54	8.5	35	7.1	0	3.8	0
3	5.6	33	8.3	94	7.3	19	6.9	4	6.1	0	3.6	0
$T = 200$												
0	—	0	—	0	—	0	—	78	—	100	—	100
1	22.7	0	9.4	0	10.2	0	6.6	22	5.2	0	2.9	0
2	9.4	13	7.8	3	8.2	8	5.7	0	4.6	0	2.7	0
3	5.3	87	6.6	97	6.5	92	4.9	0	4.2	0	2.6	0
$T = 400$												
0	—	0	—	23	—	3	—	100	—	100	—	100
1	22.5	0	5.9	51	6.5	34	4.3	0	3.5	0	2.2	0
2	9.4	0	5.2	22	5.5	47	3.8	0	3.2	0	2.1	0
3	5.1	100	4.5	4	4.7	16	3.5	0	3.0	0	2.1	0
$T = 600$												
0	—	0	—	99	—	79	—	100	—	100	—	100
1	22.4	0	4.6	1	5.0	21	3.4	0	2.9	0	1.9	0
2	9.3	0	4.1	0	4.4	0	3.1	0	2.7	0	1.9	0
3	5.2	100	3.7	0	3.9	0	2.9	0	2.5	0	1.8	0
$T = 800$												
0	—	0	—	100	—	100	—	100	—	100	—	100
1	22.7	0	4.0	0	4.2	0	3.0	0	2.6	0	1.8	0
2	9.3	0	3.6	0	3.7	0	2.7	0	2.4	0	1.7	0
3	5.0	100	3.3	0	3.4	0	2.6	0	2.3	0	1.7	0
$T = 1000$												
0	—	0	—	100	—	100	—	100	—	100	—	100
1	22.6	0	3.5	0	3.7	0	2.7	0	2.4	0	1.7	0
2	9.3	1	3.2	0	3.3	0	2.5	0	2.2	0	1.7	0
3	5.1	99	3.0	0	3.0	0	2.4	0	2.1	0	1.6	0

Notes to Table 1. R^2 indicates the percentage of total variance accounted for by the j th principal component for $j = 1, 2$ or 3 . r^* indicates the percentage of samples for which the criterion (19) selects the number of factors to be $j = 0, 1, 2$, or ≥ 3 . In every case, the true number of factors is $r = 0$ and the cross-section dimension is $N = 100$.

The next two columns of Table 1 report what the results would be if we somehow knew the true cyclical component of each variable. For this case, we applied PCA to a sample of $T - h$ observations for which the i th observed variable for $t = h + 1, h + 2, \dots, T$ is given by

$$c_{it} = \begin{cases} y_{it} - y_{i,t-h} = \varepsilon_{it} + \varepsilon_{i,t-1} + \dots + \varepsilon_{i,t-h+1} & \text{for } i = 1, 2, \dots, N/2 \\ \varepsilon_{it} & \text{for } i = (N/2) + 1, \dots, N \end{cases} \quad (20)$$

Note that the cyclical components in (20) can be serially correlated, but this autocorrelation vanishes for observations separated by more than h periods. The cyclical components c_{it} thus satisfy by construction the conditions under which Bai and Ng (2002) demonstrated that (19) would asymptotically select the correct number of factors. We find in our simulations that (19) does indeed correctly conclude there is no factor structure for these data sets provided the number of time-series observations is 600 or larger. For smaller T it is less reliable. The reason is that there is a small-sample version of the Onatski and Wang (2021) spurious factor problem that arises from the serial correlation of some of the variables that is induced by the definition of the true cyclical component in Eq. (20). If T is large enough, this problem goes away, but for smaller T it can make a difference.

Columns (5) and (6) examine the case where the analysis is based on the residuals from running an OLS regression on the raw data Y_{it} for all variables $i = 1, \dots, N$ without making any judgments about which variables are stationary and which are not. For large samples, the results are similar to those that we would obtain if we somehow knew the exact correct transformation to use for every variable.

The small-sample problem in columns (3)–(6) results from the serial correlation that is a consequence defining the cyclical component to be the error from a 24-period-ahead forecast. Columns (7) and (8) report results if we instead were to look for common factors in the 12-period-ahead forecast errors. This typically would reach the correct conclusion even in a sample of only $T = 200$ observations. Columns (9) and (10) consider 8-period-ahead forecast errors, such as our suggested cyclical calculation would use for quarterly data. The results indicate that if we have more than 50 years of data ($T = 600$ for monthly data or $T = 200$ for quarterly data), conducting PCA on the two-year-ahead OLS forecast residuals ($h = 24$ for monthly data or $h = 8$ for quarterly data) is

reasonably reliable. With less than 50 years of data, some researchers might want to use a smaller value for h or place less reliance on (19) as a criterion for selecting the number of factors.

The last two columns of Table 1 examine looking for common factors in the one-period-ahead forecast errors. In these simulations, this reaches the correct conclusion 100% of the time that there are zero factors in these data sets even for a sample of $T = 100$ observations. Thus our proposed method appears to be quite reliable if the interest is in identifying common factors behind one-period-ahead forecast errors. However, one-period-ahead forecast errors are more sensitive to outliers. This is an important consideration, as will be demonstrated in our analysis of actual data in Section 6.

5.2. Mix of unrelated stationary variables with differing persistence

In our second example, for $i = 1, 2, \dots, N/2$ the variables are generated by a stationary but persistent AR(1) process:

$$y_{i1} \sim N(0, 1/(1 - \rho^2))$$

$$y_{it} = \rho y_{i,t-1} + \varepsilon_{it} \quad \text{for } t = 2, 3, \dots, T. \quad (21)$$

The remaining $N/2$ variables are white noise ($y_{it} = \varepsilon_{it}$ for $i = (N/2) + 1, \dots, N$). Our example uses $\rho = 0.99$, so all the variables are stationary but half of them are highly persistent. The innovations $\varepsilon_{it} \sim N(0, 1)$ are independent across all i and t , so there is no factor structure in the true data-generating processes.

Columns (1) and (2) of Table 2 report the results from applying PCA to the raw data. Note that for this example, the raw data themselves satisfy the (Bai and Ng, 2002) conditions for asymptotic validity of PCA. Nevertheless, even in a sample of size $T = 1000$, the first principal component alone appears to explain a third of the data, and the criterion in (19) would always conclude incorrectly that there is at least one factor. This is a small-sample manifestation of the Onatski and Wang (2021) spurious factor phenomenon. In a sufficiently large sample, this problem would go away. But $T = 1000$ is not large enough for persistence characterized by $\rho = 0.99$.

In columns (3)–(4) we apply PCA to the residuals from a 24-period-ahead forecasting regression. Again we estimated the same regression for all variables, whether persistent or not. And again using regression residuals solves the problem pretty reliably in samples larger than $T = 600$. If we look for a factor structure in the one-period-ahead regression residuals as in columns (5)–(6), the problem is solved 100% of the time even in a sample of $T = 100$.

5.3. Results for other data-generating processes

We also report simulations for a number of other data-generating processes. For details of these and other simulations see online Appendix C. We first considered a dataset in which half the variables are cointegrated with a single common factor and the other are independent white noise. Columns (7) and (8) of Table 2 reproduce the well-known result that PCA on raw nonstationary cointegrated data can correctly pick out the single common factor; for more discussion see Harris (1997) and Onatski and Wang (2018). We find that this result is preserved when all N series are prefiltered as proposed here, as seen in columns (9)–(10) of Table 2.

We also considered an example where the common factor is stationary while the idiosyncratic components are a mix of stationary and nonstationary processes. When PCA is applied directly to the observed data y_{it} , the familiar (Onatski and Wang, 2021) result is observed in column (11) of Table 2: higher-order factors spuriously appear to explain a large amount of the variance of the data. When PCA is applied instead to the $h = 24$ -period-ahead regression forecast residuals $\hat{\varepsilon}_{it}$, the contribution of higher-order factors $\hat{F}_{jt}$ for $j \geq 1$ is substantially lower, though the selection criterion (19) would still typically incorrectly conclude that $r > 2$ for $T \leq 600$. The correct conclusion ($r = 1$) would be reached 86% of the time when $T = 1000$. Moreover, we found that the average correlation between the true realization of F_t for a particular simulation and the first estimated factor $\hat{F}_{1t}$ of the OLS regression residuals from that simulation is 0.98 for $T \geq 600$. In other words, the first principal component of the OLS regression residuals accurately uncovers the true single common feature of these data.

Table C1 in online Appendix C examines local-to-unity and fractionally integrated processes. The results in these simulations turn out to be in between the unit-root example in Table 1 and the stationary-but-persistent example in columns (1)–(6) of Table 2. Table C2 compares our approach with the PANIC method of Bai and Ng (2004). They proposed to take first-differences of all the original data, apply PCA to the changes, and then accumulate the resulting principal components. We find PANIC performs very similarly to our approach for these examples when we set $h = 1$.

Overall, these results confirm the asymptotic theory that applying PCA to OLS regression residuals is a promising approach to handling both nonstationarity and stationary persistence of unknown form provided that the time-series dimension T is reasonably large.

Table 2
Monte Carlo results in other settings.

<i>j</i>	$\rho = 0.99$						Cointegrated				Mixed			
	Raw data		$\hat{c}_i (h=24)$		$\hat{c}_i (h=1)$		Raw data		$\hat{c}_i (h=24)$		Raw data		$\hat{c}_i (h=24)$	
	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*	R^2	r^*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
<i>T</i> = 100														
0	—	0	—	0	—	100	—	0	—	0	—	0	—	0
1	43.0	98	15.8	3	4.1	0	46.1	100	40.3	100	27.5	0	23.8	0
2	3.7	2	10.6	97	3.8	0	2.8	0	3.4	0	18.1	0	11.2	1
3	2.8	0	5.5	0	3.6	0	2.5	0	3.1	0	8.5	100	8.7	99
<i>T</i> = 200														
0	—	0	—	0	—	100	—	0	—	0	—	0	—	0
1	42.5	83	10.8	0	2.9	0	48.1	100	45.7	100	25.1	0	22.6	0
2	3.3	17	8.7	1	2.7	0	2.1	0	2.3	0	18.0	0	9.8	0
3	2.3	0	7.0	99	2.6	0	2.0	0	2.2	0	8.7	100	7.9	100
<i>T</i> = 400														
0	—	0	—	5	—	100	—	0	—	0	—	0	—	0
1	40.5	52	6.3	33	2.2	0	48.8	100	46.5	100	23.8	0	21.1	0
2	3.4	47	5.4	45	2.1	0	1.8	0	1.8	0	17.9	0	6.4	0
3	2.2	1	4.8	18	2.1	0	1.7	0	1.7	0	8.8	100	5.4	100
<i>T</i> = 600														
0	—	0	—	91	—	100	—	0	—	0	—	0	—	0
1	38.6	38	4.9	9	1.9	0	49.3	100	47.0	100	23.3	0	20.8	0
2	3.5	58	4.3	0	1.9	0	1.6	0	1.6	0	17.9	0	4.9	4
3	2.3	5	3.9	0	1.8	0	1.5	0	1.6	0	8.9	100	4.3	96
<i>T</i> = 800														
0	—	0	—	100	—	100	—	0	—	0	—	0	—	0
1	37.0	33	4.1	0	1.8	0	49.4	100	47.0	100	23.2	0	20.5	33
2	3.5	60	3.7	0	1.7	0	1.5	0	1.5	0	17.7	0	4.1	43
3	2.4	8	3.3	0	1.7	0	1.5	0	1.5	0	9.2	100	3.7	24
<i>T</i> = 1000														
0	—	0	—	100	—	100	—	0	—	0	—	0	—	0
1	35.5	33	3.6	0	1.7	0	49.5	100	47.1	100	22.8	0	20.4	86
2	3.5	59	3.3	0	1.7	0	1.5	0	1.5	0	17.6	0	3.7	14
3	2.5	8	3.0	0	1.6	0	1.4	0	1.4	0	9.2	100	3.3	0

Notes to Table 2. R^2 indicates the percentage of total variance accounted for by the j th principal component for $j = 1, 2$ or 3. r^* indicates the percentage of samples for which the criterion (19) selects the number of factors to be $j = 0, 1, 2$, or ≥ 3 . In columns (1)–(6), the true number of factors is $r = 0$. In columns (7)–(14), the true number of factors is $r = 1$. In every case, the cross-section dimension is $N = 100$.

6. Characterizing a large macroeconomic data set

The use of large macroeconomic data sets was pioneered by [Stock and Watson \(1999\)](#). Their goal was to use the information of 168 different macroeconomic variables to produce better forecasts of inflation. They found that the first principal component of indicators of real economic activity produced the best inflation forecasts over the period 1959:1 to 1997:9. Their findings led to the development of the Chicago Fed National Activity Index (CFNAI).

[McCracken and Ng \(2016\)](#) developed the FRED-MD database whose 2015:4 vintage covered 134 macroeconomic variables. These include monthly measures in eight broad categories: (1) output and income; (2) labor market; (3) housing; (4) consumption, orders, and inventories; (5) money and credit; (6) interest and exchange rates; (7) prices; and (8) stock market. This data set offers benefits of continuity and continuous updating and is the basis for the analysis in this paper.

In previous applications of PCA to large macroeconomic data sets, each of the variables needed to be transformed using a detrending method that was selected individually for each series. For details of how this has been done for the CFNAI see [Federal Reserve Bank of Chicago \(2021\)](#) and for FRED-MD see the data appendix to [McCracken and Ng \(2016\)](#). [Fig. 1](#) illustrates these transformations for three important macroeconomic indicators. The first column plots the raw data, while the second column plots the data as transformed by [McCracken and Ng \(2016\)](#) using the same data set as in their original paper. Everyone agrees that industrial production (row 1) is nonstationary, and all previous researchers have used first differences of the log of industrial production shown in panel (1,2). While there is little doubt that this is a valid way to generate a stationary series for this variable, monthly growth rates of industrial production exhibit a lot of high-frequency fluctuations around the dominant cyclical patterns. For the unemployment rate (row 2), it is less clear whether the series should be regarded as stationary. [McCracken and Ng \(2016\)](#) used first differences of unemployment, which behave quite differently from the level. The purchasing managers composite index from the Institute of




Fig. 1. Level, transformed value, and cyclical component of industrial production, unemployment, and PMI Composite, 1962:3 to 2014:12.

Supply Management (row 3) appears to be stationary, and McCracken and Ng (2016) entered this series directly into PCA without any transformation.

The third column of Fig. 1 plots the cyclical components of industrial production, unemployment, and PMI as estimated by the residuals of the OLS regression (7) with $h = 24$ and $p = 12$.⁷ PMI is almost impossible to predict two years in advance, and our cyclical component is almost identical to the original series. Thus both our method and the traditional approach use this variable essentially as is. There is some but not much predictability of the unemployment rate at the two-year horizon, so for this variable our transformation much more closely resembles the original series than it does the first-difference transformation. For industrial production, our approach takes out the broad trend while retaining the essential cyclical behavior observed in the raw data. The three variables in the third column, unlike those in the second column, all share a common characterization of what is happening over the business cycle. Consistent with a long tradition in business cycle research, when plotted this way PMI appears as a leading indicator, industrial production as a coincident indicator, and unemployment as a coincident or lagging indicator, with all three clearly following the same cycle.

The top panel of Fig. 2 plots the first principal component of the transformed series arrived at by McCracken and Ng (2016).⁸ This inherits some of the high-frequency fluctuations seen in the (1,2) and (2,2) panels of Fig. 1. Indeed, McCracken and Ng (2016) regarded this series as too volatile to reliably identify business cycles and turning points, and instead plotted in their Fig. 3 the accumulation of this series. The CFNAI (shown in panel 2 of Fig. 2) is very similar to the first principal component of the FRED-MD macro data set.

The first principal component of the estimated cyclical components of the variables in the data set is plotted in the bottom panel of Fig. 2. Unlike the CFNAI, this provides a very clean summary of historical business cycles. There is another interesting difference between the third panel and the first two. The NBER defines the business-cycle trough (the end point of the shaded regions) as the low point in the *level* of overall economic activity. For example, in the first month of a new expansion, the unemployment rate is still

⁷ For those series that McCracken and Ng (2016) transformed using logs, first differences of logs, or second differences of logs (their transformations 4–6), we simply took the log of the variable before performing the regression. Thus for example the series plotted in the upper left panel of Fig. 1 is 100 times the natural logarithm of the industrial production index. For those series that they used as is, as first differences, or second differences (their transformations 1–3), we simply used the variable as is. They employed a special transformation (7) for nonborrowed reserves. One would have expected to take logs of a variable like this, but the variable took on negative values in 2008. For this series their transformation was $y_{it} = \Delta(x_t/x_{t-1} - 1.0)$ and we used $y_{it} = x_t/x_{t-1}$.

⁸ We generated this figure using the exact data and code posted at <https://research.stlouisfed.org/econ/mccracken/fred-databases/>. Note that we have multiplied the series by -1 in order to give it the property that the factor declines in recessions, and that their Fig. 3 plots the accumulations ($s_{1t} = \sum_{j=1}^t \tilde{f}_{1t}$) whereas our graph shows $\tilde{f}_{1t}$ itself.




Fig. 2. First PC of FRED-MD variables as transformed by [McCracken and Ng \(2016\)](#), the Chicago fed national activity index, and first PC of cyclical components of FRED-MD variables, 1962:3 to 2014:12.

very high, but it has started to come down. Our series in the bottom panel of [Fig. 2](#) captures this feature very well, reaching a trough at exactly the point identified by the NBER Business Cycle Dating Committee. By contrast, the low point in the series plotted in the first two panels typically comes more towards the middle of the recession. This is because the *rate of decline* of real output (a common raw input in the variables as usually transformed) starts to ease well before the recession has ended. Both in terms of the cleanness of the series and its timing, we would suggest that our approach offers a better characterization of the state of the U.S. business cycle over this sample period.

6.1. Choice of h

Our method identifies common components in h -period-ahead forecast errors. Any value of $h \geq 1$ is sufficient to produce a stationary series. Which value of h to use depends on which common components the researcher is interested in.

McCracken and Ng applied the ([Bry and Boschan, 1971](#)) algorithm for dating business-cycle turning points to the first principal component that resulted from their analysis. They found that the resulting series was in agreement with NBER recession dates only 65% of the time, meaning that 65% of NBER recession periods fell within the peak-to-trough phases of the series. We were able to reproduce this result by adapting the code in [Stock and Watson \(2014\)](#).⁹ We found a similar result (67% agreement) when Bry and Boschan is applied to the CFNAI.

The top panel of [Fig. 3](#) plots the percent agreement with NBER recessions that results from our method for every value of h between 1 and 36 months. The common component of $h = 1$ - or 2-month-ahead forecast errors has less correspondence with NBER recession dates than either the McCracken and Ng series or the CFNAI. We attribute this to the fact that replacing stationary series like the PMI in the bottom row of [Fig. 1](#) with a 1-month-ahead forecast error removes what we would normally think of as the business-cycle indicator provided by the series. By contrast, for every $h \geq 3$, our approach has a closer correspondence to NBER recessions than either McCracken-Ng or the CFNAI. The maximal agreement (99%) is obtained using $h = 25$. Choosing $h = 24$ is very similar (98%).

McCracken and Ng noted that agreement of the original $\hat{f}_t^{MN}$ with NBER expansion dates was even weaker. We found that only 51% of the expansion dates based on $\hat{f}_t^{MN}$ were also designated as expansion by NBER and still only 61% agreement using $\hat{F}_t^{MN}$. Our series with $h = 24$ does a little better than either of these with 65% agreement.

⁹ As discussed in [Stock and Watson \(2014, f.n. 2\)](#), the first step in applying the Bry and Boschan algorithm is to take a centered moving average of the series for which turning points are to be assigned, which in this case is a moving average of the first principal component. Specifying $nma = 12$ seems to come closest to reproducing the numbers reported in [McCracken and Ng \(2016, p.583\)](#). Our numbers, however, were not in all instances identical. This appears to be in part because McCracken and Ng were reporting results for a data set that has some slightly different series from those in 2015-04.csv. In the case of the percent agreement with NBER recession dates, our calculation (65%) exactly reproduces theirs. Other numbers reported here in the text are based on our reproduction using the 2015-04.csv data, which sometimes differ from those reported in [McCracken and Ng \(2016, p.583\)](#).




Fig. 3. Agreement with NBER-dated recessions and number of series exhibiting outliers as a function of forecast horizon h , 1967:2–2014:12. Notes to Fig. 3. Horizontal axis: Forecast horizon h . Top panel: Percent of NBER-dated recession periods that fall within Bry-Boschan-designated downturns. Bottom panel: Number of series exhibiting outliers. Horizontal lines denote results for the first principal component found by McCracken and Ng (2016).

6.2. Outliers

Previous users of large macro data sets devoted a lot of attention to outliers and implemented procedures to mitigate their influence. Prior to the COVID-19 recession of 2020, the CFNAI discarded observations that were more than six times the interquartile range, as did Stock and Watson (1999) in some of their analysis. McCracken and Ng (2016) discarded observations that were more than ten times the interquartile range. This criterion identifies 79 different observations on 22 different variables as outliers in the 1960:3 to 2014:12 data set.

We noted in footnote 3 that as a result of the Central Limit Theorem, we should expect fewer outliers when we calculate principal components of h -period-ahead forecasts for larger values of h . To calculate outliers for our method, we constructed forecast errors from leave-one-out regressions and designated values exceeding ten times the interquartile range as outliers.¹⁰ The number of series exhibiting outliers is plotted as a function of h in the bottom panel of Fig. 3. Using $h = 1$ results in 31 series exhibiting outliers, substantially more than in calculating the (McCracken and Ng, 2016) principal components. Many of the additional outliers come from interest rate spreads. McCracken and Ng (2016) used these as is without filtering, which turns out to be close to our definition of the cyclical component of these variables. By contrast, using $h = 1$ -month-ahead forecast errors highlights some of the unusual behavior of interest rates during the Volcker monetary contraction in 1980–81. The number of outliers steadily decreases as h is increased. Only two series exhibit outliers for $h \geq 21$.¹¹

Our recommended procedure is to use two-year-ahead regression residuals and make no corrections for outliers. The series that we have plotted in the bottom panel of Fig. 2 is the unadjusted first principal component of the full set of OLS residuals $\hat{c}_{it}$.

Outliers are an even bigger issue when data for 2020 are included. For the 2024:12 vintage of FRED-MD, the McCracken-Ng procedure would identify 40 of the 126 variables as all being outliers in the single month of 2020:4. Despite dropping all of these 40 observations, the first principal component calculated using their algorithm shows an enormous decline in this month. Indeed, in order to include the 2020 observations in the top panel of Fig. 4, the scale must be so large that it makes all the previous cyclical fluctuations barely noticeable. The CFNAI modified its procedure for dealing with anomalous observations to handle these observations. Even so the CFNAI still displays an unprecedented drop in 2020, as seen in the second panel.

¹⁰ That is, we calculated $\hat{c}_{it} = y_{it} - \hat{\alpha}_{i,t-h}$ with $\hat{\alpha}_{i,t-h} = \left(\sum_{s=1, s \neq t}^T \bar{z}_{i,s-h} \bar{z}_{i,s-h}' \right)^{-1} \left(\sum_{s=1, s \neq t}^T \bar{z}_{i,s-h} y_{is} \right)$ estimated separately for each i and t and then divided $\hat{c}_{i,t}$ by its observed interquartile range.

¹¹ These are total bank reserves and nonborrowed reserves in the Fed's response to the financial crisis.




Fig. 4. First PC of FRED-MD variables as transformed by McCracken and Ng (2016), the Chicago Fed National Activity Index, and first and second PC of cyclical components of FRED-MD variables, 1962:3 to 2024:9.

By contrast, only two variables are identified as outliers for 2020:4 for purposes of our approach, these being new claims for unemployment insurance and the number unemployed for less than 5 weeks. The result of applying our procedure to the FRED-MD database available as of December of 2024 with no corrections for outliers is displayed in the third panel of Fig. 4. Note that, unlike the top two panels, our series describes the downturn in 2020 on a comparable scale as earlier recessions, although our series indicates that the speed of the downturn was unprecedented, as was the growth in the first two months of the recovery. Also in contrast to the first two panels, our series indicates (correctly, in our view) that the economy did not fully recover from the COVID-19 shock until September of 2021. The sharp spike up in our series in April of 2022 reflects the fact that most macro variables were substantially higher in April 2022 than one would have predicted based on observations in April 2020. This conclusion is also consistent with the aggressive actions of policy makers in the spring of 2020. Our series further indicates that economic activity remained unusually strong through the fall of 2023.

Another issue comes from discontinued, newly added, or missing variables. McCracken and Ng (2016) adapted the (Stock and Watson, 2002) algorithm for unbalanced panels, though they found in their original data set that the results are essentially identical if one simply drops variables as needed to create a balanced panel. For our application, we have simply calculated principal components of $\hat{e}_{it}$ on a balanced panel, though there is no obstacle to applying the (Stock and Watson, 2002) algorithm to an unbalanced panel of $\hat{e}_{it}$.

6.3. Uses of macroeconomic cyclical factors

A key use of PCA is to summarize the statistical information in a large cross section of indicators. The movement in variable i that is captured by the j th factor alone is given by $\hat{\lambda}_{ij}\hat{f}_{jt}$. Since $\hat{e}_{it}$ is normalized to have unit variance, the fraction of the variance of the stationary component of variable i that is explained by the j th macro factor is given by $\hat{\lambda}_{ij}^2 T^{-1} \sum_{t=1}^T \hat{f}_{jt}^2$.

Table B1 in the online appendix reports the R^2 explained by the first macroeconomic cyclical factor for each of the 119 variables used in our analysis of the 2023 vintage data set. The first factor alone accounts for almost 2/3 of the variance of typical indicators of real output or income and more than half of the variance of typical indicators of labor-market conditions. The first cyclical factor is far less successful at describing financial indicators and nominal prices. It is interesting that when we add the second cyclical factor,

the R^2 for the median price indicator rises to 66%. The first factor thus seems mainly to capture real economic conditions and the second characterizes nominal prices and interest rates.

The fourth panel in Fig. 4 plots the second cyclical factor. This by construction is orthogonal to the first, and often continues to fall even as the recovery in real economic activity is beginning. This is consistent with the view that nominal variables may respond sluggishly to business-cycle developments. It describes events beginning in 2022 as a third big U.S. inflation wave, though less dramatic than the big inflations of 1973–74 and 1979–81.

We next explore the use of the cyclical factors in forecasting. Stock and Watson (1999) demonstrated that the first principal component of a large data set of real macroeconomic variables could be very helpful for forecasting inflation. Their finding gave rise to the Chicago Fed National Activity Index (CFNAI), a PCA-based indicator that is still widely used today. We compare the usefulness for forecasting of the CFNAI (denoted $\hat{f}_t^{CF}$), the first or second principal component calculated using the algorithm and data set of McCracken and Ng (2016) (denoted $\hat{f}_t^{MN1}$ and $\hat{f}_t^{MN2}$, respectively), or the first or second principal component of the forecasting residuals (denoted $\hat{f}_t^{HMX1}$ or $\hat{f}_t^{HMX2}$) calculated from the FRED-MD data set using Eq. (12). Our approach to comparing different forecasts is similar to that used by Stock and Watson (1999) and McCracken and Ng (2016).

A particular model m uses a set of variables $\mathbf{x}_t^m$ that are observed at date t to try to forecast the value of a variable of interest y_{t+h}^h that will not be observed until $t+h$:

$$y_{t+h}^h = \pi^m \mathbf{x}_t^m + u_{t+h}^{m,h}. \quad (22)$$

We estimated the value of π^m by OLS regression on the subsample $t = T_0, T_0 + 1, \dots, T_1 - h - 1$ and used these coefficients to forecast $y_{T_1}^h$. We then augmented the sample by one observation, estimating the regression over $t = T_0, T_0 + 1, \dots, T_1 - h$ and using those coefficients to forecast $y_{T_1+1}^h$. We repeated this for an evaluation period T_1 to T_2 and calculated the average squared forecast error over this evaluation period. We performed this analysis using three different evaluation periods. The first evaluation period is specified by $T_1 = 1970:1$ to $T_2 = 1996:12$, which was the evaluation period in the original study by Stock and Watson (1999). The second evaluation period is $T_1 = 1997:1$ to $T_2 = 2014:12$, which corresponds to the new data used by McCracken and Ng (2016) that were not available to Stock and Watson (1999). The third evaluation period is $T_1 = 2015:1$ to $T_2 = 2024:9$, which is the new data available since publication of McCracken and Ng (2016). The models we considered were a pure autoregressive model, $\mathbf{x}_t^{AR} = (1, y_t^1, y_{t-1}^1, \dots, y_{t-5}^1)'$, and models that add to the autoregressive model six lags of one of the principal components estimates. For example, $\mathbf{x}_t^{CF} = (\mathbf{x}_t^{AR}, \hat{f}_t^{CF}, \hat{f}_{t-1}^{CF}, \dots, \hat{f}_{t-5}^{CF})'$. This differs a little from the forecast evaluations performed by Stock and Watson (1999) and McCracken and Ng (2016) in that these authors used BIC to select different lag lengths for the autoregressive and principal components and for each subsample, whereas we set the lag length to six for every evaluation. Also, since $\hat{f}_t^{CF}$ is only available beginning in 1967:3, we used $T_0 = 1967:9$ as the first date for estimation of all models. In every case, for each T we re-estimated the coefficients $\hat{\pi}^m$ for the forecasting regression (22) using an expanding data set ending h periods before the variable being forecast.

A separate question is the data set used to estimate the factors $\hat{f}_t$ themselves. Insofar as the factors are only identified up to sign, the meaning of coefficients multiplying $\hat{f}_t$ could change across expanding samples. We formed data sets to calculate the factors $\hat{f}_t$ for each of the three methods as follows. (i) For the explanatory variable $\hat{f}_t^{MN}$, for the first two evaluation samples we followed McCracken and Ng (2016) in calculating the factor $\hat{f}_t^{MN}$ using the full historical vintage of the FRED-MD database available as of 2015:4.¹² For the third evaluation sample, we estimated $\hat{f}_t^{MN}$ using the full database available as of 2024:12. (ii) For the explanatory variable $\hat{f}_t^{CF}$ we used the value of the CFNAI as it was reported in 2024 for all historical dates.¹³ (iii) For the explanatory variable $\hat{f}_t^{HMX}$, we need to specify both (a) the data set used to estimate the coefficients $\hat{\alpha}_i$ for the detrending regression (7) and (b) the data set used to calculate principal components of the forecast errors $\hat{\epsilon}_{it} = y_{it} - \hat{\alpha}'_i z_{i,t-h}$. For (a), for all three evaluation samples, we estimated $\hat{\alpha}_i$ from the 2015:4 vintage data set. For (b), for the first two evaluation samples, we calculated principal components of $\hat{\epsilon}_{it}$ using the 2015:4 vintage data set. For the last evaluation sample, we calculated principal components of $\hat{\epsilon}_{it}$ using the 2024:12 data set. Thus our forecasting exercise for the third subsample is not affected by the potential concern that the detrending regression used information that was not available at the time the forecast was made. This is because the coefficients for the detrending regression were all known as of 2015. We also repeated the exercise using instead for (b) trend coefficients estimated using all the data through 2024. This alternative specification produced similar results (not reported here).

For our first set of evaluations we set y_{t+h}^h to be the average inflation rate between month t and $t+h$, quoted at an annual rate,

$$y_{t+h}^h = (1200/h) \log(CPI_{t+h}/CPI_t),$$

where CPI_t denotes the level of the consumer price index in month t .¹⁴ The column labeled AR in Table 3 reports the simulated out-of-sample mean squared error of a purely autoregressive model for each of the three evaluation samples and for forecast horizons of $h = 1, 6$, or 12 months.¹⁵ The first panel reproduces the finding of Stock and Watson (1999) that an index like CFNAI significantly improves forecasts for longer horizons over the 1970–1996 period. The alternative measures $\hat{f}_t^{MN1}$ or $\hat{f}_t^{HMX1}$ offer similar improvements. The

¹² Our series for $\hat{f}_t^{MN1}$ for this subsample is almost (but not quite) identical to the series analyzed by McCracken and Ng (2016). We have been unable to identify the source of the small discrepancies.

¹³ We downloaded $\hat{f}_t^{CF}$ on January 13, 2025 from the FRED database at <https://\protect\penalty~z@fred.stlouisfed.org/series/CFNAI>.

¹⁴ Here again our evaluation design differs slightly from that in Stock and Watson (1999) and McCracken and Ng (2016) in that those authors took the object of interest to be to forecast the *change* in the inflation rate as a function of lagged changes:

$$y_{t+h}^h = (1200/h) \log(CPI_{t+h}/CPI_t) - 1200 \log(CPI_t/CPI_{t-1}).$$

¹⁵ Stock and Watson (1999) reported results for $h = 12$ months whereas McCracken and Ng (2016) reported results for $h = 1, 6$, and 12.

Table 3
Mean squared forecast errors for different models.

sample	horizon	Consumer Price Index					
		AR	CF	MN1	HMX1	MN2	HMX2
1970–1996	h = 1	7.91	1.00	0.99	1.03	0.98	0.90
	h = 6	4.26	0.81	0.77	0.80	0.95	0.88
	h = 12	5.32	0.70	0.62	0.74	1.02	1.33
1997–2014	h = 1	12.26	1.03	1.04	1.02	0.97	1.09
	h = 6	6.08	1.23	1.23	1.23	0.96	1.11
	h = 12	4.21	1.22	1.22	1.28	0.95	1.17
2015–2024	h = 1	8.07	1.51	1.27	1.14	0.94	1.04
	h = 6	3.70	1.95	1.48	1.76	1.03	1.01
	h = 12	3.64	1.71	1.30	1.49	1.09	0.96

sample	horizon	Industrial Production					
		AR	CF	MN1	HMX1	MN2	HMX2
1970–1996	h = 1	76.73	0.97	0.94	0.96	1.01	1.02
	h = 6	38.66	0.92	0.93	0.83	0.71	0.79
	h = 12	27.19	1.01	1.06	1.21	0.49	0.87
1997–2014	h = 1	58.90	0.85	0.83	0.98	1.04	1.00
	h = 6	22.61	0.93	0.94	1.05	1.24	1.12
	h = 12	20.11	0.96	1.01	1.06	1.26	1.11
2015–2024	h = 1	506.67	1.73	0.93	1.10	1.02	1.00
	h = 6	87.37	2.29	1.16	1.03	1.05	0.94
	h = 12	38.10	2.31	1.14	0.85	1.03	0.93

Notes to [Table 3](#). AR column reports simulated out-of-sample mean squared forecast error for purely autoregressive model evaluated over three different out-of-sample periods. CF column reports the relative MSE when lags of the Chicago Fed National Activity Index are added to the autoregression, with a value less than one indicating the variable is useful for forecasting. MN1 column reports the relative MSE when lags of the first principal component calculated using the procedures in [McCracken and Ng \(2016\)](#) are used in place of the CFNAI. HMX1 reports relative MSE when lags of the first principal component of the estimated cyclical components are used in place of CFNAI. MN2 and HMX2 report results when the second principal component is used instead of the first.

indexes offer little or no improvement for one-month-ahead forecasts over this period, but again are similar to each other. All three indexes are outperformed by simple autoregressive forecasts over either of the later two evaluation periods. The observation that inflation has become much harder to forecast in data since 1996 has been reported by a number of other researchers, including [Atkeson and Ohanian \(2001\)](#), [Fisher et al. \(2002\)](#), [Stock and Watson \(2007\)](#), and [Stock and Watson \(2009\)](#). Interestingly, the second cyclical factor $\hat{f}_t^{HMX2}$ does better than any of the other four indexes at forecasting inflation at the one-month horizon for the 1970–1996 sample and at the 6–12 month horizons for the 2015–2024 sample.

[Table 3](#) also reports forecasts of industrial production, setting

$$y_{t+h}^h = (1200/h) \log(IP_{t+h}/IP_t),$$

for IP_t the level of the industrial production index in month t . The indexes $\hat{f}_t^{CF}$, $\hat{f}_t^{MN1}$, and $\hat{f}_t^{HMX1}$ all help forecast industrial production over near horizons in the first two evaluation periods. The CFNAI does particularly poorly at forecasting either inflation or industrial production at any horizon for 2015–2024. Both $\hat{f}_t^{HMX1}$ and $\hat{f}_t^{HMX2}$ do significantly better than CFNAI in every case over this period.

We conclude that our approach offers similar benefits to conventional PCA when evaluated in terms of simulated out-of-sample forecasts, and does much better than measures like the CFNAI for recent data. We share the conclusion of the earlier literature that the usefulness of any principal-component-based measure for purposes of forecasting depends on the variable, evaluation period, and horizon of the forecast.

7. Conclusion

Calculating principal components of medium-horizon forecast errors is a viable approach to identifying the common cyclical factors that drive a large collection of potentially nonstationary economic indicators. This avoids the need to decide how to detrend each individual series and is much more promising than approaches such as the Chicago Fed National Activity Index for handling data that include the large outliers of 2020.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Supplementary material

Supplementary material associated with this article can be found in the online version at [10.1016/j.jeconom.2026.106317](https://doi.org/10.1016/j.jeconom.2026.106317).

References

- Ai, X., 2016. A note on Karhunen–Loeve expansions for the demeaned stationary Ornstein–Uhlenbeck process. *Stat. Probab. Lett.* 117, 113–117.
- Atkeson, A., Ohanian, L.E., 2001. Are Phillips curves useful for forecasting inflation? *Fed. Reserve Bank Minneap. Q. Rev.* 25 (1), 2–11.
- Bai, J., Ng, S., 2002. Determining the number of factors in approximate factor models. *Econometrica* 70 (1), 191–221.
- Bai, J., Ng, S., 2004. A PANIC attack on unit roots and cointegration. *Econometrica* 72 (4), 1127–1177.
- Bai, J., Ng, S., 2006. Confidence intervals for diffusion index forecasts and inference for factor-augmented regressions. *Econometrica* 74 (4), 1133–1150.
- Bai, J., Ng, S., 2008. *Large Dimensional Factor Analysis*. Now Publishers Inc.
- Bai, J., Ng, S., 2013. Principal components estimation and identification of static factors. *J. Econom.* 176 (1), 18–29.
- Barigozzi, M., Lippi, M., Luciani, M., 2021. Large-dimensional dynamic factor models: estimation of impulse–response functions with I(1) cointegrated factors. *J. Econom.* 221 (2), 455–482.
- Beveridge, S., Nelson, C.R., 1981. A new approach to decomposition of economic time series into permanent and transitory components with particular attention to measurement of the “Business cycle”. *J. Monet. Econ.* 7 (2), 151–174.
- Bry, G., Boschan, C., 1971. Cyclical analysis of time series: selected procedures and computer programs. NBER Technical Paper 20, National Bureau of Economic Research, New York.
- Chen, X., Li, W.V., 2003. Quadratic functionals and small ball probabilities for the m -fold integrated brownian motion. *Ann. Probab.* 31 (2), 1052–1077.
- Federal Reserve Bank of Chicago, 2021. Background on the Chicago Fed national activity index. <https://www.chicagofed.org/-/media/publications/cfnai/background/cfnai-technical-report-pdf.pdf>
- Fisher, J. D.M., Liu, C.T., Zhou, R., 2002. When can we forecast inflation? *Econ. Perspect. Fed. Reserve Bank Chic.* 26 (1), 32–44.
- Gao, F., Hannig, J., Torcaso, F., 2003. Integrated brownian motions and exact l_2 -small balls. *Ann. Probab.* 31 (3), 1320–1337.
- Hamilton, J.D., 1994. *Time Series Analysis*. Princeton University Press.
- Hamilton, J.D., 2018. Why you should never use the Hodrick–Prescott filter. *Rev. Econ. Stat.* 100 (5), 831–843.
- Harris, D., 1997. Principal components analysis of cointegrated time series. *Econom. Theory* 13 (4), 529–557.
- Komlós, J., Major, P., Tusnády, G., 1976. An approximation of partial sums of independent RV's, and the sample DF. II. *Z. für Wahrscheinlichkeitstheorie verwandte Geb.* 34, 33–58.
- McCracken, M.W., Ng, S., 2016. FRED-MD: a monthly database for macroeconomic research. *J. Bus. Econ. Stat.* 34 (4), 574–589.
- Onatski, A., Wang, C., 2018. Alternative asymptotics for cointegration tests in large VARs. *Econometrica* 86 (4), 1465–1478.
- Onatski, A., Wang, C., 2021. Spurious factor analysis. *Econometrica* 89 (2), 592–614.
- Stock, J.H., Watson, M.W., 1999. Forecasting inflation. *J. Monet. Econ.* 44 (2), 293–335.
- Stock, J.H., Watson, M.W., 2002. Forecasting using principal components from a large number of predictors. *J. Am. Stat. Assoc.* 97 (460), 1167–1179.
- Stock, J.H., Watson, M.W., 2007. Why has US inflation become harder to forecast? *J. Money Credit Bank.* 39, 3–33.
- Stock, J.H., Watson, M.W., 2009. Phillips curve inflation forecasts. in: Fuhrer, J., Kodrzyck, Y.K., Little, J.S., Olivei, G.P. (Eds.), *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, MIT Press, Cambridge, MA, pp. 99–202. <https://doi.org/10.7551/mitpress/9780262013635.003>.
- Stock, J.H., Watson, M.W., 2014. Estimating turning points using large data sets. *J. Econom.* 178, 368–381.
- Stock, J.H., Watson, M.W., 2016. Dynamic factor models, factor-augmented vector autoregressions, and structural vector autoregressions in macroeconomics. In: *Handbook of Macroeconomics*. Elsevier. Vol. 2, pp. 415–525.