Gutenberg School of Management and Economics

& Research Unit “Interdisciplinary Public Policy”

Discussion Paper Series

Filtering And Inference

Luca A. Pennacchio

31st August 2026

Discussion paper number 2607

Johannes Gutenberg University Mainz
Gutenberg School of Management and Economics
Jakob-Welder-Weg 9
55128 Mainz
Germany
https://wiwi.uni-mainz.de/

Contact details

Luca A. Pennacchio
Department of Law and Economics
Johannes Gutenberg University
Jakob-Welder-Weg 4
55128 Mainz
Germany
luca.pennacchio@uni-mainz.de

Filtering And Inference

Luca A. Pennacchio

GSEFM Goethe University Frankfurt am Main,

Johannes Gutenberg University Mainz

luca.pennacchio@uni-mainz.de

31. August 2026

Abstract

Many empirical macroeconomic questions rely on filtering non-stationary data to extract a stationary, business-cycle-like component. Common practice is to apply a linear filter with a standard passband, although economically relevant cycles may lie outside this passband. The resulting signal-extraction error likely affects subsequent regression estimates. This paper proposes a Continuous Wavelet Transform-informed filter that uses the CWT scalogram to identify statistically significant passbands in raw, non-stationary data relative to a researcher-specified null model. These passbands are supplied to a flexible Butterworth bandpass filter to extract a denoised, stationary cycle. Simulations show that the method improves signal extraction for periodic cycles and recovers statistically informative variation in stochastic-cycle settings missed by commonly used BK, HP, and Hamilton filters, while performing comparably or better in terms of correlation and RMSE. Applications demonstrate its use for cycle extraction, seasonal adjustment, and frequency-dependent regression analysis.

Keywords: Business cycles, Signal extraction, Continuous Wavelet Transform, Bandpass filtering, Spectral analysis

1 Introduction

A central problem in empirical macroeconomics is to extract from a non-stationary time series a stationary cyclical component that captures economically meaningful fluctuations. Common approaches include detrending, first differencing, Hodrick-Prescott (HP) filtering, Hamilton (H) filtering, and bandpass filtering. A fundamental issue, however, is which fluctuations should be regarded as economically meaningful. Following Burns and Mitchell (1946), business-cycle fluctuations are commonly associated with periodicities between 6 and 32 quarters. This fixed range has subsequently become a standard benchmark for measuring and comparing business-cycle properties across countries; see, for example, Uribe and Schmitt-Grohé (2017). More importantly, cyclical components isolated over this fixed frequency band are routinely employed in structural estimation, from standard regression specifications to DSGE models. But there is no guarantee that this conventional frequency band captures the economically meaningful variation in the data. Its principal advantage is comparability: standard practice defines the business cycle ex ante as fluctuations occurring between $1/32$ and $1/6$ cycles per quarter, implicitly treating variation outside this range as irrelevant. Yet it is difficult to justify the assumption that economically meaningful cyclical fluctuations across countries, variables, and historical periods invariably lie within the same fixed band. If, for example, GDP in a particular country exhibits a recurring component with a period of 36 quarters, should that variation be discarded simply to preserve comparability? This points to a potential trade-off between comparability and the preservation of relevant cyclical information. In practice, however, the trade-off may be largely illusory. Comparability is of limited value if it is achieved by excluding economically meaningful fluctuations while retaining variation that is little more than filtered noise. The shortcomings of a one-size-fits-all frequency band become particularly consequential when the resulting cyclical component is subsequently used in regression analysis or structural estimation. Then, OLS estimates are usually inconsistent due to a signal-extraction error (Watson (1986), Wooldridge (2002)).

This paper argues against applying a one-size-fits-all range of periodicities in filtering, proposing instead that only statistically significant periodicities be used to reduce potential signal-extraction error. For stationary time series, this problem can be addressed by testing for statistically significant periodicities in the periodogram. For non-stationary time series, periodogram-based testing first requires transforming the series so that it can be treated as stationary. However, the transformation itself may distort the periodogram. For example, suppose a researcher incorrectly attributes the non-stationarity of a time series to a unit root. First differencing may then introduce high-frequency artifacts while attenuating low-frequency components, thereby altering which periodicities appear statistically significant. More generally, without knowledge of the underlying data-generating process (DGP), it is difficult to determine how a particular transformation distorts the spectral content of the series. For observed data that may exhibit complex spectral structure and multiple forms of non-stationarity, it is therefore inadvisable to select a transformation ad hoc and subsequently conduct statistical inference on the resulting periodogram. Such an approach risks confounding genuine spectral features with artifacts induced by the transformation itself.

This paper argues that the continuous wavelet transform (CWT) (Grossmann and Morlet (1984)) should be used instead to explore the non-stationary time series. Not because the CWT scalogram[1] generally provides a clearer picture than the periodogram, but because it provides an additional time dimension. By providing a high-resolution time-frequency representation, the CWT makes it possible to identify and interpret frequency components whose importance possibly varies over time, without generally requiring the series to be transformed beforehand. The conventional use of the CWT is to reconstruct a signal by applying the inverse wavelet transform over selected scales. This paper instead proposes using the CWT as an exploratory device to identify statistically significant passbands, which are then supplied to a bandpass filter of choice. A Butterworth-based bandpass filter is particularly well suited to this purpose because its transition-band slope can be adjusted flexibly, permitting the construction of relatively narrow passbands. Relative to direct reconstruction through the inverse CWT, the proposed procedure is more closely aligned with the conventional use of filtering in macroeconomics and therefore provides a more tractable framework for extracting a stationary component from a non-stationary series. The resulting series is denoised, stationary and restricted to statistically significant frequency content. Simulation results demonstrate substantial reductions in signal-extraction error when the underlying cycle is periodic. For stochastic AR(2) cycles, the proposed procedure extracts a cyclical component that retains statistically relevant information about the true model cycle that is not captured by the BK, H, or HP filters.

The remainder of the paper is organized as follows. Section 2 defines the signal-extraction error and demonstrates that it can arise even under ideal conditions. Section 3 reviews the

[1] The CWT magnitude scalogram can be viewed as a time-localized extension of the periodogram, showing how spectral magnitude varies across both time and frequency.

existing continuous wavelet transform (CWT) methodology, with particular emphasis on the CWT scalogram; readers already familiar with the CWT scalogram may skip Section 3.2. Section 4 introduces a novel CWT-informed filtering procedure and compares its performance with alternative filters using the simulated processes of Hodrick (2020), for which the underlying cycles are known. The paper concludes with several empirical applications.

2 Signal-Extraction Error

2.1 Definition

Using a filtered cycle—an estimated explanatory variable—in a regression model can lead to inconsistent OLS estimates (Watson (1986), Wooldridge (2002) and Fuller (1995)). Model-based approaches can, of course, account explicitly for the fact that the explanatory variable is estimated rather than observed. Such approaches, however, require the researcher to specify a model for the latent cyclical component and its relationship to the observed data, thereby imposing potentially strong restrictions on the data-generating process. This paper therefore does not pursue model-based solutions, but instead focuses on linear filtering as it is commonly applied in empirical macroeconomics.

To illustrate the problem in the context of linear filtering, let c_t denote the true but unobserved cyclical component. Because c_t is latent, the researcher instead works with a filtered estimate, $\tilde{c}_t$. The true cycle can therefore be decomposed into the extracted component and a signal-extraction error:

$$c_t = \tilde{c}_t + a_t, \tag{1}$$

where a_t is the signal-extraction error; the part of the true cycle that a filter failed to extract due to a one-size-fits-all passband. Suppose the true cycle is the sum of two periodic components:

$$c_t = s_{1t} + s_{2t}. \tag{2}$$

Further suppose that s_{1t} repeats every 16 quarters and s_{2t} every 36 quarters. If the one-size-fits-all passband of 6 to 32 quarters per cycle is applied, the filter passes only s_{1t}, such that

$$\tilde{c}_t = s_{1t}, \qquad a_t = s_{2t}. \tag{3}$$

Suppose a researcher is interested in the true model:

$$y_t = \beta(s_{1t} + s_{2t}) + \varepsilon_t, \tag{4}$$

but, given the selected passband, the regression becomes

$$y_t = \beta s_{1t} + u_t, \tag{5}$$

with

$$u_t = \beta s_{2t} + \varepsilon_t. \tag{6}$$

The probability limit of the OLS estimator based on s_{1t} becomes

$$\operatorname{plim} \hat{\beta}_{OLS} = \beta + \frac{\beta \operatorname{Cov}(s_{1t}, s_{2t}) + \operatorname{Cov}(s_{1t}, \varepsilon_t)}{\operatorname{Var}(s_{1t})}. \tag{7}$$

Hence, (assuming $\operatorname{Cov}(s_{1t}, \varepsilon_t) = 0$) OLS is inconsistent whenever

$$\operatorname{Cov}(s_{1t}, s_{2t}) \neq 0. \tag{8}$$

In practice, the exact orthogonality condition is often too strong. Even if two components have different frequencies, they need not be orthogonal in finite samples. The problem can also arise through other regressors. Suppose the estimated regression includes an additional variable z_t. Since the regression disturbance contains the missed component s_{2t}, OLS also requires

$$\operatorname{Cov}(z_t, u_t) = 0. \tag{9}$$

Thus, even if the passed component s_{1t} is orthogonal to the missed component s_{2t}, OLS can still be inconsistent if another included regressor contains information about the missed part of the cycle. The same issue appears in dynamic regressions. If lags of the extracted cycle are included, then the estimated model may contain regressors such as $s_{1,t-j}$. Since the disturbance contains the missed component s_{2t}, consistency requires

$$\operatorname{Cov}(s_{1,t-j}, s_{2t}) = 0 \tag{10}$$

for all included lags j. This condition may fail even when the contemporaneous covariance between s_{1t} and s_{2t} is zero. There is also a related misspecification issue. If the true relationship depends differently on s_{1t} and s_{2t}, then the model with only s_{1t} is not merely affected by signal-extraction error; it is estimating the wrong object. The only way to eliminate the endogeneity issue is to use the optimal linear projection of the true cycle on the full relevant information set. The optimal linear projection guarantees that the extracted cycle is orthogonal to the signal-extraction error. However, a filter with a one-size-fits-all passband does not generally have this property. Therefore, unless the filter coincides with the optimal linear projection, the extracted cycle may remain correlated with the missed component, with other regressors, or with lagged included cycles, leading to inconsistent OLS estimates.

These considerations make the quality of cycle extraction consequential not only for the measurement of business-cycle facts, but also for subsequent econometric inference. When a stationary cyclical component is extracted for use in regressions or structural estimation, the objective should therefore be to retain as much of the relevant cyclical signal as possible while excluding noise and other components that do not belong to the cycle of interest. A poorly chosen passband can both discard relevant cyclical information and admit unwanted variation, thereby distorting measured business-cycle properties and, once the extracted cycle enters an estimated model, potentially inducing endogeneity. The choice of passband is therefore not merely a descriptive filtering decision: it can directly affect the consistency and interpretation of subsequent parameter estimates. Accordingly, passband selection should be guided by statistical evidence on the frequency content of the observed series, rather than by a uniformly imposed a priori range.

2.2 Numerical example

The following example demonstrates that filters with a priori fixed passbands can generate substantial signal-extraction error by admitting different amounts of noise into the extracted cyclical component. Suppose you simulate the following arbitrary model:

$$w_t = 0.02t + \sin\left(\frac{2\pi t}{P}\right) + n_t, \quad n_t \sim \mathcal{N}(0,1), \quad t = 1, 2, \ldots, 1000, \tag{11}$$

where the permanent component and the cycle are modeled as a simple linear trend and a sine function with periodicity $P = 16$, respectively. The Baxter and King (1999), Hamilton (2018) and Hodrick and Prescott (1997) filters are applied to (11) to extract the simulated

4-year cycle. Figure 1 shows the autocorrelation functions[2] (ACFs) of the estimated cycles

Figure 1: ACFs across filters

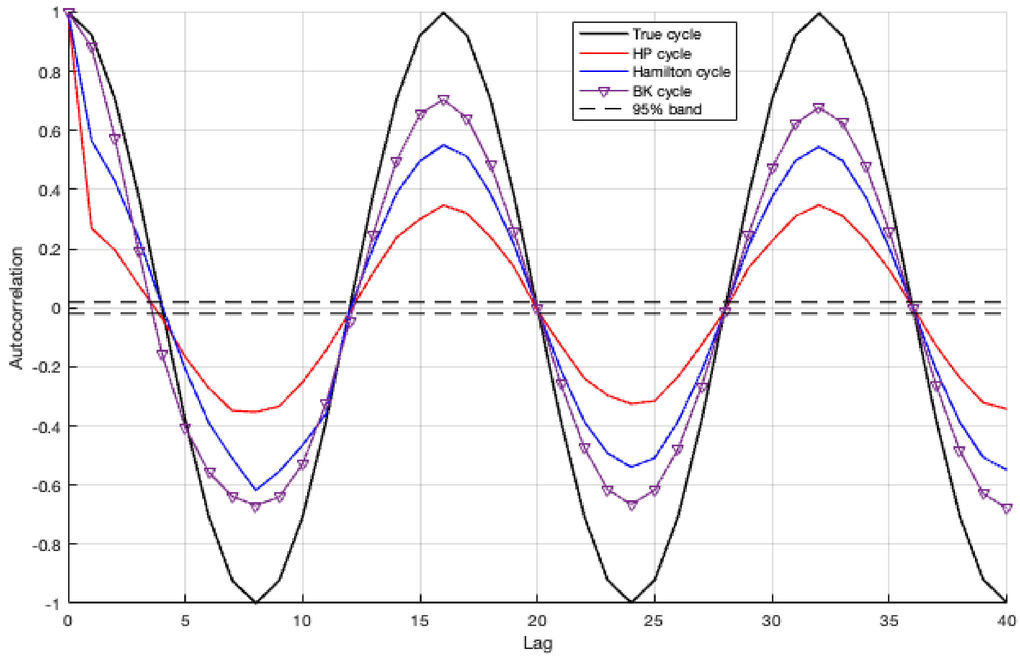

Source: Own computations. ACFs of the estimated cycles across filters are shown and compared to the ACF of the true cycle. The standard settings are the HP filter with smoothing parameter $\lambda = 1600$, the BK filter with lower and upper cutoff periods of 6 and 32 quarters and lag length 12, and the Hamilton filter with lead length $h = 8$ and lag length $p = 4$. The sample size T is chosen to be large to mitigate finite-sample bias.

across filters comparing them to the true cycle for one example realization. The persistence of each estimated cycle is altered by the extent to which filtered measurement noise is included. The visible bias is the lowest for the BK-filtered cycle since high frequency components are left out and the highest for the HP-filtered cycle since much more filtered noise is soaked up. However, the bias for the BK-filtered cycle remains considerable. Two metrics describing the signal-extraction error are calculated: a standard root mean squared error (RMSE) and the sample Pearson correlation coefficient between the extracted cycle and the true cycle, where $\hat{c}_{j,t}$ denotes the cycle estimated by filter $j \in \{\mathrm{BK}, \mathrm{H}, \mathrm{HP}\}$, $c_{j,t}$[3] denotes the true cycle, and T_j is the number of valid observations used in the comparison:

$$\mathrm{RMSE}_j = \sqrt{\frac{1}{T_j} \sum_{t=1}^{T_j} (\hat{c}_{j,t} - c_{j,t})^2} \tag{12}$$

[2]The continuity of the ACFs is artificial. The data points have been connected by direct interpolation to improve visual comparability.

[3]The BK and H filter produce 'NaNs' at the beginning and the end of a series. Therefore, the true cycle has been trimmed as to be perfectly aligned with the extracted cycles. This is why the true cycle also has a subscript j.

$$\mathrm{Corr}_j = \frac{\sum_{t=1}^{T_j} \left(\hat{c}_{j,t} - \bar{\hat{c}}_j\right)\left(c_{j,t} - \bar{c}_j\right)}{\sqrt{\sum_{t=1}^{T_j} \left(\hat{c}_{j,t} - \bar{\hat{c}}_j\right)^2 \sum_{t=1}^{T_j} \left(c_{j,t} - \bar{c}_j\right)^2}}. \tag{13}$$

Equation (11) is simulated 5000 times, the three filters are applied to each realization and the average value of the RMSE and the correlation statistic are shown in Table 1. For this specific simulation, the BK filter performs best in describing the true cycle as its estimated cycle has the lowest RMSE and the highest correlation with the true cycle. The estimated cycle from the H filter shows the highest RMSE but has a much higher correlation with the true cycle compared to the cycle estimated by the HP filter. The HP filter outperforms the H filter regarding the RMSE but not regarding the correlation with the true cycle. Importantly, these issues should not be interpreted as shortcomings of the filters themselves. The key issue is not knowing the true data-generating process (DGP). Even if the permanent component could be reliably separated from the cycle, which is difficult in itself, it is generally unknown how much filtered noise is mixed into the cycle. While excluding fluctuations with periods shorter than 6 quarters may be justified in many applications, the rationale for retaining all variation between 6 and 32 quarters is much less clear. In this example, only a narrow passband centered around 16 quarters per cycle would recover the underlying cycle with little bias. As Section 4 will show, statistical testing can identify this range directly, eliminating the need to impose a one-size-fits-all passband a priori.

Table 1: True cycle versus BK-, H-, HP-filtered cycle

Filter	RMSE		Corr.	
BK $(6-32, 12)$	0.507	(0.00031)	0.826	(0.00023)
Hamilton $(h=8, p=4)$	1.299	(0.00054)	0.689	(0.00025)
HP $(\lambda = 1600)$	0.965	(0.00032)	0.579	(0.00027)

Source: Own computations. The table reports the RMSE and the sample Pearson correlation coefficient between the extracted cycle and the true simulated cycle for the BK, H, and HP filters. The columns report Monte Carlo averages over 5000 realizations. Monte Carlo standard errors are shown in parentheses. In each realization, the observed series consists of a linear trend, a sinusoidal cycle with period 16 quarters, and Gaussian white noise, and the sample size is chosen to be $T = 1000$ to mitigate finite-sample and endpoint bias.

3 Continuous Wavelet Transform

3.1 Interpreting frequency content of a non-stationary time series

In spectral analysis, the frequency content of a time series is usually examined through the spectrum, which is derived from the Fourier transform. The spectrum shows how the variance of a stationary time series is allocated over frequencies. With observed data, the spectrum must be estimated – for example, using a smoothed periodogram. Two practical issues arise: First, the conventional interpretation of a smoothed periodogram relies on the assumption that the time series is stationary. Any transformation to render a non-stationary time series stationary can introduce artifacts into the periodogram. Without knowing the DGP, it is generally unclear which transformation causes which artifacts. Especially, for observed data with complicated spectral features. To avoid introducing artifacts through transformation, the researcher must identify the source of the non stationarity with high precision – a difficult task when working with observed data. Second, the periodogram does not contain a time dimension. Therefore, it is generally unknown if a peak in the periodogram reflects a genuine cycle or a singular event characterized by extreme fluctuations.

The CWT allows for analyzing the frequency content of non-stationary time series (Grossmann and Morlet (1984)). The general idea is to compare the time series with shifted and scaled versions of a chosen prototype or basis function, sometimes called the mother wavelet, and measure how much the time series resembles that function locally. This prototype function can be derived, e.g., from the Gaussian probability density function (PDF). Wavelets can also be step-like or real and complex-valued. Figure 2 is taken from Percival and Walden (2000) to illustrate basic wavelet shapes. The first wavelet is the classic Haar wavelet which is useful to describe sharp jumps in a time series. The second and third (the latter also called Mexican hat) wavelet are normalized versions of the first and second derivative of the Gaussian PDF, respectively. Both are more useful when changes in a time series are more smooth. The key requirements are that the wavelet, regardless of its specific shape, be a localized, oscillatory function with zero mean and finite energy (Percival and Walden (2000)). More formally, the CWT coefficient for a specific scale and time is the inner product of a time series with a shifted and scaled version of the mother wavelet. For a time series $x(t)$, the CWT coefficient at scale

Figure 2: Wavelet types

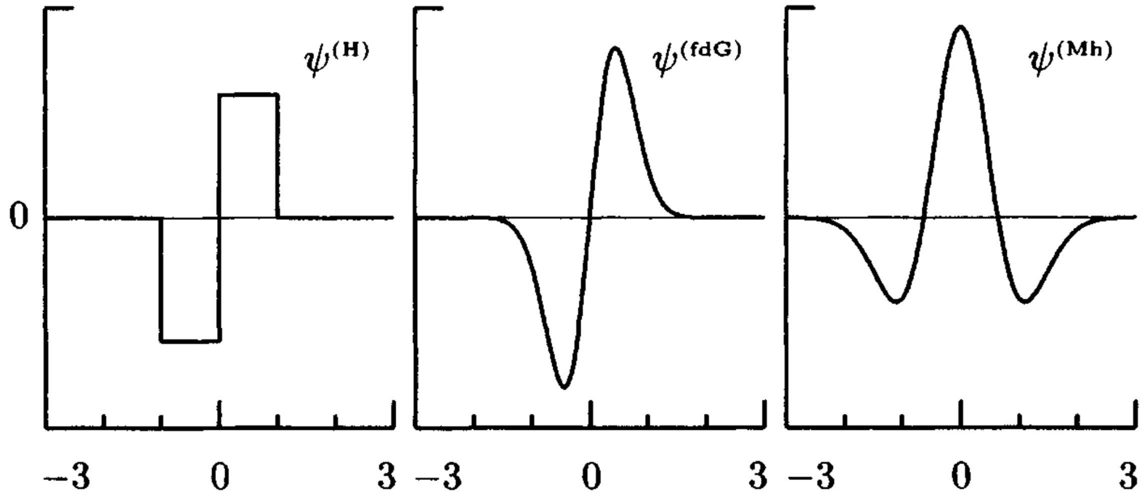

Source: Percival and Walden (2000), p. 3.

$j > 0$ and time translation k is

$$W_x(j,k) = \int_{-\infty}^{\infty} x(t)\, \psi^*_{j,k}(t)\, dt, \tag{14}$$

where the shifted and scaled wavelet is defined as

$$\psi_{j,k}(t) = \frac{1}{\sqrt{j}} \psi\left(\frac{t-k}{j}\right). \tag{15}$$

Here, j controls the scale of the wavelet, k controls its location in time, and $(\cdot)^*$ denotes complex conjugation (Ramsey (2002)). The resulting coefficients $W_x(j,k)$ are large if the wavelet matches the time series locally well. For a high scale, the wavelet is dilated or strechted in the time domain. Conversely, at a low scale, the wavelet is compressed in the time domain. Essentially, the wavelet is constantly reshaped and compared to the time series across scale and time (see Daubechies and Paul (1988) and Grossmann and Morlet (1984) for a rigorous exposition; Percival and Walden (2000), Ramsey (2002) and Crowley (2007) for an introduction).

The conventional approach to signal extraction from wavelet coefficients is to apply the inverse wavelet transform over a selected set of scales. This requires retaining scales deemed relevant while excluding those associated with unwanted variation. For example, reconstructing the signal while omitting the largest scales removes very low-frequency components. In the economics literature many applications use the discrete wavelet transform (DWT) instead of the CWT, because the reconstruction of a signal via the inverse transform is easy to implement and requires low computational effort (Aguiar-Conraria and Soares (2014)). While the theoretical CWT uses all possible scales and points in time and therefore includes a lot of redundant information, the DWT samples a wavelet transform on a dyadic scale-time grid which reduces redundancies significantly (see also Lo Cascio (2007)). This paper focuses on the CWT because its redundant representation is advantageous when the objective is to obtain as detailed a picture as possible of the time-varying frequency content of the series. This paper does not reconstruct the time series directly from selected CWT scales via the inverse transform, as doing so would simply shift the central problem to deciding which scales contain meaningful variation and which should be excluded. Instead, the CWT scalogram is used as an exploratory tool to identify the statistically informative frequency content of the series and its evolution over time (discussed in Section 4).

Readers familiar with the CWT scalogram may skip Section 3.2.

3.2 CWT scalogram

A standard way to visualize the CWT coefficients $W_x(j, k)$ is the magnitude scalogram. Figure 3 shows the magnitude scalogram of one realization of the simulated process from equation (11). The vertical axis measures a normalized frequency in cycles per sample based on the scales. In this example, samples refer to quarters. The horizontal axis represents the time span of the time series. The wavelet coefficients are represented in a heat map where bright yellow colors indicate strong local resemblance of the scaled and shifted mother wavelet with the time series. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts. As the mother wavelet gets stretched in time with higher scales, edge effects reach further inward, so the reliable region becomes smaller. Therefore, the high magnitude regions in the bottom left and right corners are mere artifacts and should not be interpreted as true resemblance of the scaled wavelet with the time series. The horizontal signal that is perfectly consistent over time appears at 0.0625 cycles per sample, which corresponds exactly to $1/0.0625 = 16$ quarters per cycle – the true simulated signal from equation (11). For this simplistic simulation, both the scalogram and a periodogram will reveal the true cycle. However, the scalogram offers the additional time dimension which is missing in the periodogram. When studying observed non-stationary data, the extra time dimension is very valuable as one can see how the frequency content changes

Figure 3: Magnitude Scalogram

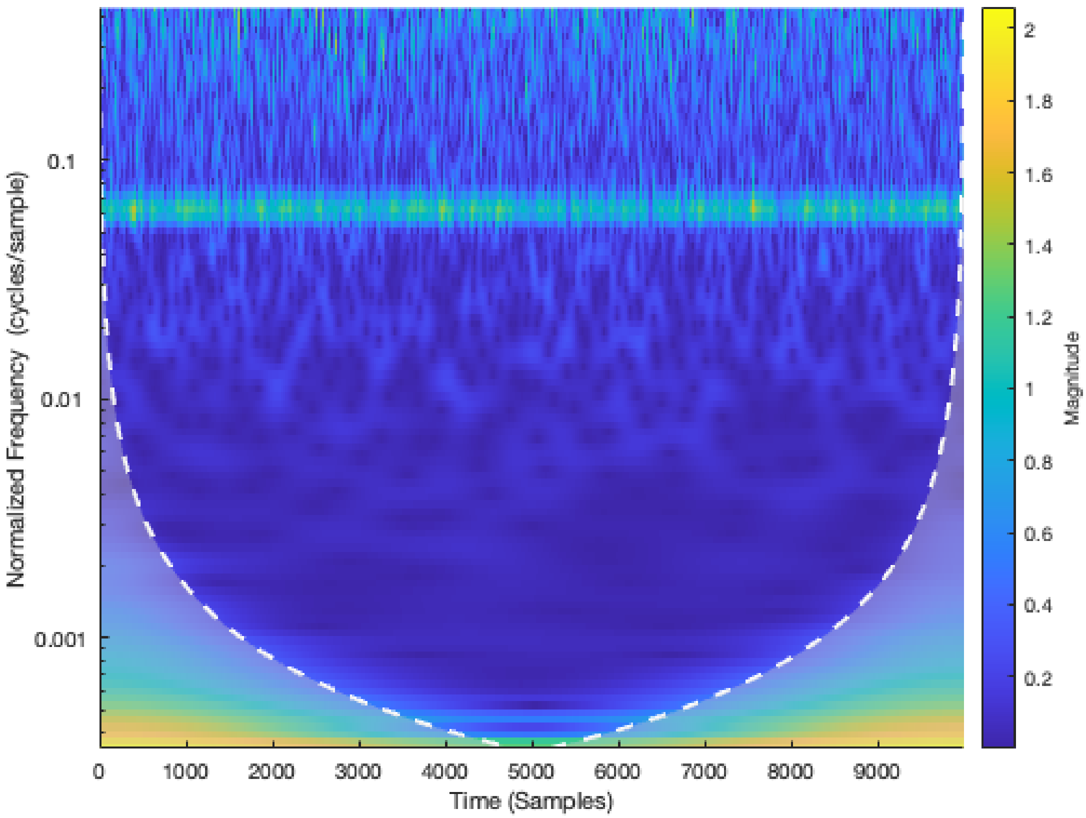

Source: Own computations. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

over time; structural breaks and singular events become visible. Additional important features that become visible are the very low frequency and the noise component. In the low scales (high frequency), the noise is visible by alternating stripes of high and low magnitudes. At intermediate frequencies, the noise appears as a leopard-like pattern that extends further into the lower-frequency range. The permanent component is visible in the very high scales (low frequency) with time consistent, thin horizontal lines. Note that this illustration shows an ideal case where a sine wave can be perfectly disentangled from the very large simulated dataset.

To illustrate a more common setting, eight arbitrary processes have been simulated for $T = 300$ and $\varepsilon_t \sim \mathcal{N}(0, 1)$:

$$\text{White Noise:} \qquad y_t = \varepsilon_t, \tag{16}$$

$$\text{Random Walk with Drift:} \qquad y_t = 0.2 + y_{t-1} + \varepsilon_t, \tag{17}$$

$$\text{AR(1):} \qquad y_t = 0.8 y_{t-1} + \varepsilon_t, \tag{18}$$

$$\text{AR(2):} \qquad y_t = 0.5y_{t-1} + 0.2y_{t-2} + \varepsilon_t, \tag{19}$$

$$\text{MA(1):} \qquad y_t = \varepsilon_t + 0.5\varepsilon_{t-1}, \tag{20}$$

$$\text{MA(2):} \qquad y_t = \varepsilon_t + 0.5\varepsilon_{t-1} + 0.2\varepsilon_{t-2}, \tag{21}$$

$$\text{ARMA(1,1):} \qquad y_t = 0.6y_{t-1} + \varepsilon_t + 0.5\varepsilon_{t-1}, \tag{22}$$

$$\text{ARIMA(2,1,2):} \quad (1 - 0.5L + 0.2L^2)(1 - L)y_t = (1 + 0.4L + 0.2L^2)\varepsilon_t, \tag{23}$$

When comparing scalograms across different processes, it is important to recognize that CWT coefficient patterns are not unique to a given DGP. As Figure 4 illustrates, similar time-frequency features can arise across distinct processes. For example, the MA(2) realization in panel 6 cannot, in general, be reliably distinguished from the ARMA(1,1) realization in panel 7 on the basis of the scalogram alone. No particular time-frequency pattern can therefore be attributed uniquely to a specific process with certainty. At most, certain recurring features may be suggestive of broader spectral characteristics; for instance, a scalogram dominated by alternating, leopard-like magnitude patterns is more likely to reflect noise-like structure than a persistent cyclical signal. Comparing panel 1 with panel 5 and 6 also suggests that additional

Figure 4: Common time series scalograms 1 out of 5000

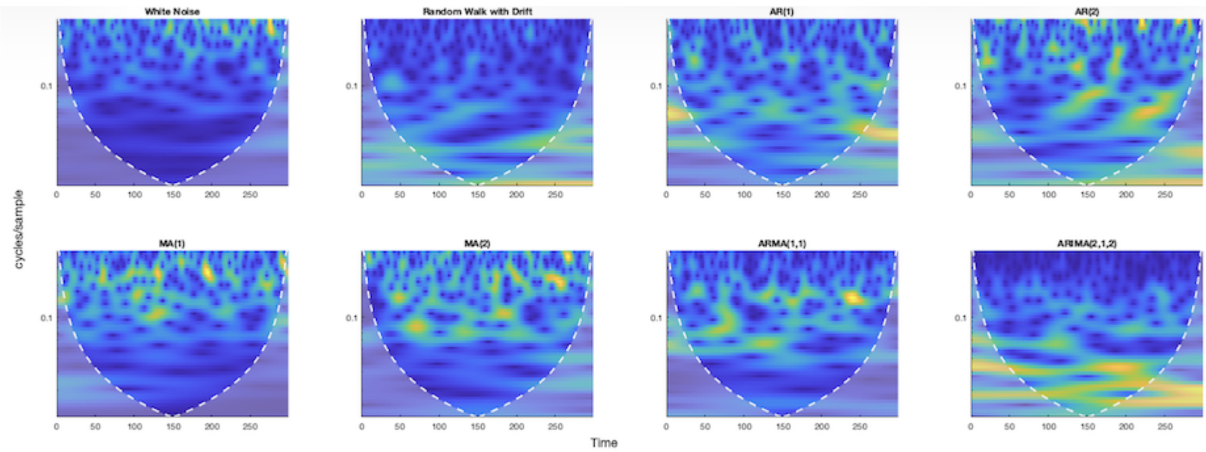

Source: Own computations. 8 processes - one per subplot - are simulated 5000 times and this figure shows one realization per process. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

serial dependence increases the magnitude of this noise signature. The next extreme case is the random walk with drift in panel 2. Here, as expected, magnitude is higher in the very low frequencies. Panel 3 shows that the AR(1) realization looks similar to panel 2 but with slightly lower magnitude in the lower frequencies and higher magnitude in the middle frequencies. Finally, the mixed processes in panel 7 and 8 show that, generally, the dominating feature will be visible in the scalogram. To see this, compare the low frequency content between panel 7 and 8. Figure 4 displays only a single realization of each process. To assess whether the observed patterns are representative, Figures 7–11 in the Appendix show realizations drawn randomly without replacement from a pool of $5,000$ simulations for each of the eight processes.[4] The broad patterns identified in Figure 4 remain evident across these additional realizations.

Figure 5 illustrates this from a different perspective. A process realization is fixed and then the parameters are increased. For realizations of the white-noise and MA(1) processes, varying the innovation variance or the parameter θ does not materially alter the scalogram. The dominant magnitude stays in the high frequencies and magnitude increases with additional

Figure 5: Common time series scalograms for different parameters

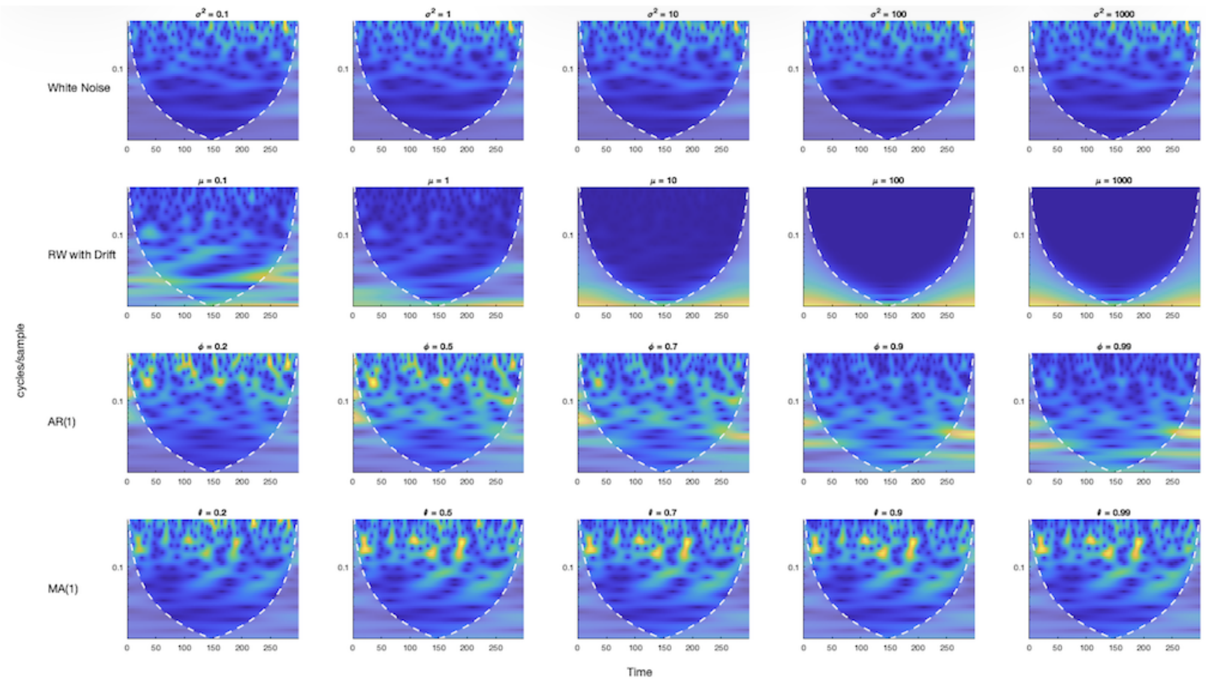

Source: Own computations. 4 process realizations - one per row - are fixed and parameters given in the subplot titles are changed in each column. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

serial dependence. For the realizations based on the random walk with drift and the AR(1) process, varying the parameters alter the scalogram significantly. Increasing the drift parameter

[4] Displaying scalograms for 50 randomly selected realizations per process is more informative and tractable than constructing a single summary scalogram across all 5,000 simulations. The former preserves the underlying time-frequency structure, whereas aggregation tends to obscure these features and produce a cluttered representation.

across the second row shows how the scalogram becomes dominated by the strong trend, which appears as thin, horizontally persistent lines at very low frequencies. In the third row, the magnitude moves from high to low frequencies.

As noted above, CWT magnitude patterns do not uniquely identify the underlying DGP. A scalogram should therefore not be interpreted as a direct classifier of stochastic processes. Nevertheless, the location and persistence of high-magnitude regions across time and frequency can provide useful diagnostic information. Three features are particularly suggestive: (i) an alternating, leopard-like pattern of high- and low-magnitude regions may be indicative of noise-dominated variation; (ii) when the leopard-like pattern extends into intermediate frequencies with elevated magnitude, it may be consistent with short-memory serial dependence, such as that generated by a MA process; and (iii) sustained high magnitude at low frequencies may signal strong persistence. These features should be viewed as diagnostic clues rather than unique signatures of specific processes.

4 A Novel CWT-Informed Filter

4.1 CWT-informed Butterworth filter

The basis for reconstructing a signal from the CWT is the coefficient matrix defined in equation (14), whose rows correspond to scales and whose columns correspond to points in time. Regardless of whether reconstruction is based on the DWT, CWT, or maximal-overlap DWT (MODWT[5]), the researcher must ultimately decide which scales contain relevant signal and which should be discarded. This distinction is generally not known a priori. Among these transforms, the CWT is particularly well suited to informing this choice because it provides a highly resolved representation of how coefficient magnitude varies jointly across time and frequency. A natural way to identify potentially relevant frequency content is therefore to examine where coefficient magnitude is concentrated and how persistent these concentrations are over time. Building on this information, I propose an algorithm for systematically identifying statistically significant passbands from the CWT coefficient matrix:

Let x_t, $t = 1, \ldots, T$, denote the observed time series, and let its CWT be given by the coefficients in (14). Each scale j is associated with a frequency f_j. The goal is to identify

[5]The MODWT provides a compromise between the DWT and CWT by relaxing the rigid dyadic sampling scheme of the DWT. As a result, more coefficients are retained and redundancy is reduced less aggressively. See Section 4 of Crowley (2007) for further details.

frequency bands for which the wavelet transform displays both large magnitude and consistency over time. For each scale j, define the maximum absolute wavelet magnitude as

$$M_x(f_j) = \max_{1 \le k \le T} |W_x(j,k)| \,. \tag{24}$$

$M_x(f_j)$ is a magnitude score. It measures the largest absolute wavelet coefficient observed at frequency f_j. A high value of $M_x(f_j)$ indicates that the series contains a strong component at that frequency at least at some point in time. A large value of $M_x(f_j)$ may, however, reflect an isolated event rather than a persistent cyclical component. Therefore, for each scale j, also define the average absolute wavelet magnitude as

$$C_x(f_j) = \frac{1}{T} \sum_{k=1}^{T} |W_x(j,k)| \,. \tag{25}$$

$C_x(f_j)$ is a consistency score. It measures whether the magnitude at frequency f_j is consistently high over time rather than concentrated in only a few time periods.

To assess whether the magnitude and consistency scores are statistically significant, they are compared with scores generated under a background-noise null model, following the approach of Torrence and Compo (1998). This empirical testing framework has two advantages: the null model is explicit, and it can be adapted to the application at hand. A deliberately conservative benchmark, for example, is white noise with a standard deviation equal to that of the detrended input series. Importantly, the purpose of the null model is not to estimate the cycle of interest, but to generate a plausible level of random background variation against which the observed spectral structure can be evaluated. The method therefore does not depend on any particular null specification; the conservative white-noise model is merely a starting point. Researchers can instead incorporate prior information or examine robustness across a range of alternative null models, depending on the empirical setting.

Let M denote the number of Monte Carlo simulations, and generate null surrogates

$$e_t^{(m)} \sim \mathcal{N}(0, \sigma_x^2), \qquad t = 1, \ldots, T, \qquad m = 1, \ldots, M, \tag{26}$$

where σ_x is, e.g., the standard deviation of the detrended input series. For each surrogate $e_t^{(m)}$, compute the CWT coefficients $W_e^{(m)}(j,k)$. The corresponding simulated magnitude and consistency scores are

$$M_e^{(m)}(f_j) = \max_{1 \le k \le T} \left| W_e^{(m)}(j,k) \right| \tag{27}$$

and

$$C_e^{(m)}(f_j) = \frac{1}{T} \sum_{k=1}^{T} \left| W_e^{(m)}(j,k) \right|. \tag{28}$$

For a chosen significance level α, define the simulated critical values at frequency f_j as

$$q_M^{1-\alpha}(f_j) = Q_{1-\alpha}\left(\left\{ M_e^{(m)}(f_j) \right\}_{m=1}^{M} \right) \tag{29}$$

and

$$q_C^{1-\alpha}(f_j) = Q_{1-\alpha}\left(\left\{ C_e^{(m)}(f_j) \right\}_{m=1}^{M} \right), \tag{30}$$

where $Q_{1-\alpha}(\cdot)$ denotes the empirical $(1-\alpha)$-quantile. The magnitude score is statistically significant at frequency f_j if

$$M_x(f_j) > q_M^{1-\alpha}(f_j), \tag{31}$$

and the consistency score is statistically significant at frequency f_j if

$$C_x(f_j) > q_C^{1-\alpha}(f_j). \tag{32}$$

A frequency band is retained only if both conditions hold simultaneously:

$$\mathcal{F}_\alpha = \left\{ f_j \in \mathcal{F} : M_x(f_j) > q_M^{1-\alpha}(f_j) \ \wedge \ C_x(f_j) > q_C^{1-\alpha}(f_j) \right\}. \tag{33}$$

The set $\mathcal{F}_\alpha$ contains frequencies at which the original series exhibits wavelet magnitudes that are both large and consistent relative to the chosen null model. Thus, only frequency bands for which both the magnitude score and the consistency score are statistically significant over the same passband are retained. These passbands are then supplied to a bandpass filter.

I propose the usage of a bandpass filter derived from a Butterworth (1930) (BW) filter. There are several advantages in using a BW filter. For one, Gómez (1999) showed that the HP filter actually is a BW filter. So, the gain function of the BW filter can be designed to meet specific requirements. Specifically, the first parameter in the BW filter decides at which frequency the gain is exactly $\frac{1}{2}$ and the second parameter (the degree of the BW filter) changes the slope of the gain function, increasing the sharpness of the filter where necessary (Gómez

(2001)). This is advantageous for narrow passbands. A common problem encountered while filtering a time series is phase shift. Meaning that important features of the original series appear shifted in time in the filtered series (see for example Schüler (2024)). This is a major concern once a filtered series is used in a regression model. A BW bandpass filter can easily be implemented with the forward-backward filtering method. Here, a chosen offline filter is applied forwards, then backwards on the reversed filtered series and the output is reversed again to produce the final filtered series (Gustafsson (2002)). The forward and backward operation cancel phase shift. By applying the proposed method, the resulting filtered series is a denoised, stationary series with statistically significant frequency content against a null model.

4.2 Simulation results

4.2.1 Deterministic trend and periodic cycle

To illustrate how the CWT-informed filter described in the previous Section 4.1 can help reducing the signal-extraction error, Table 1 from Section 2.2 is reprinted with six new rows including the metrics for the CWT-informed filter, the inverse CWT (iCWT) reconstruction and the BK filter with the statistically significant passbands. For the simulated series (11), the algorithm indicates a significant passband at a 5% level between 0.052 and 0.075 cycles per sample with a center frequency (geometric average) of 0.0624 cycles per sample which is very close to the true frequency of 0.0625 cycles per sample (16 quarters per cycle). This passband is now supplied to the filters in rows 4 - 6 to extract the implied cycle from the non-stationary series (11). The difference between the cycle implied by the CWT-informed filter and the true cycle, the signal-extraction error, decreased sharply compared to the filters in rows 1 - 3. Also, there is a significant improvement in the correlation metric. The results for the iCWT reconstruction are virtually identical. This finding is line with Yogo (2008). However, as rows 7 and 8 illustrate, this may not always be the case. On a 0.01%-significant passband, the signal-extraction error for the cycle implied by the CWT-informed filter can be reduced again while the correlation metric increases to 0.962. For the iCWT on the same 0.01%-significant passband the signal-extraction error increases by 40%. The BK filter, as any bandpass filter, can be modified to pass the same significant passbands (see rows 6 and 9). Both the signal-extraction error and the correlation improve compared to the standard settings in row 1. Compared to the CWT-informed filter and the iCWT reconstruction, the

Table 2: True cycle versus implied cycles

Filter	RMSE		Corr.	
BK $(6-32, 12)$	0.507	(0.00031)	0.826	(0.00023)
Hamilton $(h=8, p=4)$	1.299	(0.00054)	0.689	(0.00025)
HP $(\lambda=1600)$	0.965	(0.00032)	0.579	(0.00027)
CWT-informed (5%-significant passband)	0.213	(0.00029)	0.957	(0.00012)
iCWT (5%-significant passband)	0.219	(0.00026)	0.956	(0.00011)
BK (5%-significant passband)	0.396	(0.00021)	0.920	(0.00016)
CWT-informed (0.01%-significant passband)	0.197	(0.00025)	0.962	(0.00010)
iCWT (0.01%-significant passband)	0.308	(0.00025)	0.959	(0.00011)
BK (0.01%-significant passband)	0.501	(0.00014)	0.920	(0.00016)

Source: Own computations. The table reports in rows 1 - 3 the RMSE and the sample Pearson correlation coefficient between the extracted cycle and the true simulated cycle for the BK, H, and HP filters (identical to Table 1). Rows 4 - 6 show the same metrics for the implied cycles of the CWT-informed filter, the inverse CWT reconstruction and for the BK filter for a 5% significant passband (0.052 to 0.075 cycles per sample). Rows 7 - 9 show the same as rows 4 - 6 just for a 0.01% significant passband (0.054 to 0.067 cycles per sample). The columns report Monte Carlo averages over 5000 realizations. Monte Carlo standard errors are shown in parentheses. In each realization, the observed series consists of a linear trend, a sinusoidal cycle with period 16 quarters, and Gaussian white noise, and the sample size is chosen to be $T = 1000$ to mitigate finite-sample and endpoint bias.

BK filter with the identical passband induces almost double the signal-extraction error and is slightly less correlated with the true cycle. The signal-extraction error actually gets worse when using the 0.01%-significant passband in the BK filter. This highlights the fact that in principle many filters can be modified to achieve good results. What is necessary, however, is to first identify the significant passband. This is what the CWT-informed filter can achieve. One reason the CWT-informed filter can outperform the BK filter, even when applied to the same passband, is that its steeper gain response allows a narrower range of frequencies to be isolated with less spectral distortion. This is particularly useful because a steeper gain response implies a narrower transition band[6], thereby reducing the range of frequencies that are only partially attenuated. For applications where periodic oscillations are suspected and where measurement noise is present, the CWT-informed filter outperforms the BK-, H- and HP filter. The performance is almost identical to the iCWT reconstruction for the 5%-significant passband, but not for the 0.01%-significant passband. Here, the CWT-informed filter has a lower signal-extraction error and a higher correlation with the true cycle.

Next the CWT-informed filter is tested for processes which have stochastic trend and stochastic cycles in a Hodrick (2020) style.

[6]The transition band is the frequency interval over which the filter gain moves between the passband and stopband.

4.2.2 Stochastic trends and cycles

In his response to Hamilton (2018), Hodrick (2020) compared the cyclical components produced by the HP and BK filters with those produced by the H filter across several DGPs. Hodrick (2020)'s findings are that the H filter performs better for simple time series models like a random walk. For more complicated DGPs, the HP and BK filter outperform the H filter. To make this comparison, Hodrick (2020) compares a set of metrics from which I use a reduced set containing the standard deviation of the estimated cycle, the correlation between the true and the estimated cycle, the RMSE between the true and the estimated cycle, the slope coefficient of the regression

$$c_t - \widehat{c}_{j,t} = \alpha_j + \delta_j \left(\widehat{c}_{\mathrm{CWT},t} - \widehat{c}_{j,t}\right) + \varepsilon_{j,t}, \tag{34}$$

and the adjusted R^2 of this regression. The slope coefficient δ_j indicates whether or not the cycle implied by the CWT-informed filter can explain the true cycle c_t in the presence of filter $j \in \{iCWT, H, HP, BK\}$. When the slope coefficient is close to 1, the cycle implied by the CWT-informed filter explains the true cycle with no additional explanatory power from filter j. Then, the cycle implied by the CWT-informed filter contains statistically important information not captured by the cycle produced by filter j. Conversely, if the slope coefficient is close to 0, the CWT-informed filter contributes little to explaining the true cycle. Hodrick (2020) tests the filters for six DGPs, but only the last four have a defined cycle. The following filter comparison focuses on the DGPs where both trend and cycle are defined objects. Each observed series y_t is decomposed into a trend component g_t and a cyclical component c_t. Innovations in Hodrick (2020)'s cycles are part of the true cycle. The following 4 DGPs are identical to Hodrick (2020) to ensure comparability. All simulations have the same sample size ($T = 289$), sample mean and sample standard deviation as the rate of growth of quarterly US real GDP between 1947 and 2019. Innovations in trend and cycle are uncorrelated (for further details see Hodrick (2020)).

The first DGP combines a stochastic trend modeled as ARIMA(1,1,0), where the change in trend is serially correlated and the cyclical component is modeled as ARIMA(2,0,0):

$$\begin{aligned} \Delta g_t &= 0.07786 + 0.9\Delta g_{t-1} + \nu_t, \quad \nu_t \sim \mathcal{N}(0, 0.002), \\ c_t &= 1.25 c_{t-1} - 0.45 c_{t-2} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, 0.6385). \end{aligned} \tag{35}$$

In the second DGP, the change in trend has a conditional mean. The conditional mean is modeled as a random walk. The cyclical component is modeled as an ARIMA(2,0,0):

$$\begin{aligned} \Delta g_t &= d_{t-1} + 0.545 w_t, \\ d_t &= d_{t-1} + 0.021 u_t, \\ c_t &= 1.51 c_{t-1} - 0.565 c_{t-2} + 0.603 v_t, \end{aligned} \tag{36}$$

where w_t, u_t, and v_t are independent standard normal innovations.

The change in trend in the third DGP includes a changing unconditional mean and an ARIMA(2,0,0) cycle:

$$\Delta g_t = \mu_t + 0.9 \Delta g_{t-1} + \nu_t, \quad \nu_t \sim \mathcal{N}(0, 0.002),$$

with

$$\mu_t = \begin{cases} 0.1, & t \leq \lfloor T/3 \rfloor, \\ 0.07786, & \lfloor T/3 \rfloor < t \leq \lfloor 2T/3 \rfloor, \\ 0.0554, & t > \lfloor 2T/3 \rfloor, \end{cases} \tag{37}$$

$$c_t = 1.25 c_{t-1} - 0.45 c_{t-2} + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, 0.6385).$$

The fourth DGP is the same as the third one only that the intercept in trend growth is allowed to change smoothly over time:

$$\begin{aligned} \Delta g_t &= 0.07786 + \frac{0.04428(145 - t)}{289} + 0.9 \Delta g_{t-1} + \nu_t, && \nu_t \sim \mathcal{N}(0, 0.002), \\ c_t &= 1.25 c_{t-1} - 0.45 c_{t-2} + \epsilon_t, && \epsilon_t \sim \mathcal{N}(0, 0.6385). \end{aligned} \tag{38}$$

All processes are initialized at zero. For each process, the filters are applied to y_t to extract an estimate of the true cycle. Columns 3 - 5 of each table are replications of the corresponding tables in Hodrick (2020) (see pages 47 - 50 therein). To facilitate replication, the simulations reported in Table 3 are based on 100 Monte Carlo repetitions and $1,000$ white-noise surrogates for the CWT-informed filter. Table 6 in the Appendix confirms the results with $5,000$ Monte Carlo repetitions and $10,000$ white-noise surrogates. The latter computations were performed on the MOGON KI high-performance computing cluster at Johannes Gutenberg University Mainz.

The AR(2) structure imposed by Hodrick (2020) on the cyclical component provides a useful example of how the null model can be adapted from a deliberately conservative benchmark to one that incorporates prior information about cyclical dynamics. More generally, the null model can represent any process against which the researcher wishes to assess the observed spectral structure. In the present simulation exercise, the null is constructed under the assumption that the cycle follows a general ARMA(p, q) process. The input series is first detrended, an ARMA(p, q) specification is selected using the Bayesian information criterion (BIC), and the standard deviation of the fitted model's residuals is used to calibrate the innovations of the white-noise null model. This makes explicit which spectral features are deemed significant relative to a clearly specified benchmark. In this setting, the ARMA residuals can provide a good estimate of the remaining random variation, allowing the underlying AR(2) structure to emerge as significant spectral structure rather than treating all variation in the detrended series as noise. The ARMA assumption imposes relatively little structure on the cyclical component. A sufficiently general ARMA(p, q) class can accommodate a wide range of stationary dynamics, making it a flexible benchmark rather than a tightly specified model of the cycle. When little or no prior information is available, the researcher can instead begin with a deliberately conservative null model and assess the robustness of the resulting significant passbands across alternative null specifications. It is worth emphasizing that the method does not depend on any particular class of null models. Its advantage is instead that the choice of significant passbands is made relative to an explicit and transparent benchmark, rather than being determined by an ad hoc choice of frequencies.

Table 3, panel 1 shows the results for the first DGP (constant unconditional mean change in trend modeled as ARIMA(1,1,0) and an ARIMA(2,0,0) cycle). The correlation of the cycle implied by the CWT-informed filter with the true cycle is almost identical and the RMSE comparable to the cycles implied by the HP and BK filters. The slope coefficients indicate that the cycle implied by the CWT-informed filter is meaningful in the presence of the other filters and contains important information in explaining the true cycle. The lowest slope coefficient of 0.41 is found for the cycle implied by the HP filter (rejection rate of $H_0 : \delta_{HP} = 0$ at a 5% level is 0.84). To put this into perspective: Hodrick (2020)'s regression tested the HP and BK filter against the H filter. The H-filtered cycle took the place of the CWT-informed filtered cycle in regression (34). The slope coefficients in Hodrick (2020) on the same DGP are 0.087 (adjusted R^2: 0.041) for the HP-filtered cycle and 0.15 (adjusted R^2: 0.082) for the

Table 3: Results across the four DGPs

DGP 1	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.67	2.08	2.48	1.44	1.37
Correlation	0.87	0.70	0.73	0.88	0.86
RMSE	0.91	1.52	1.71	0.85	0.93
Slope coefficient $\widehat{\delta}_j$		0.95	0.89	0.41	0.63
Adjusted R^2		0.67	0.74	0.18	0.22
Rejection rate (5%)		1.00	1.00	0.84	0.96

DGP 2	CWT-informed	iCWT	H	HP	BK
Standard deviation	2.37	2.44	3.42	1.60	1.53
Correlation	0.75	0.72	0.51	0.64	0.67
RMSE	1.89	1.99	3.15	2.12	2.09
Slope coefficient $\widehat{\delta}_j$		0.86	0.92	0.72	0.71
Adjusted R^2		0.14	0.65	0.27	0.24
Rejection rate (5%)		0.68	1.00	0.95	0.92

DGP 3	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.68	2.04	2.57	1.44	1.38
Correlation	0.85	0.71	0.70	0.88	0.85
RMSE	0.97	1.48	1.84	0.88	0.96
Slope coefficient $\widehat{\delta}_j$		0.94	0.91	0.37	0.56
Adjusted R^2		0.60	0.75	0.16	0.19
Rejection rate (5%)		1.00	1.00	0.79	0.88

DGP 4	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.67	2.01	2.49	1.44	1.38
Correlation	0.86	0.72	0.73	0.88	0.86
RMSE	0.93	1.43	1.73	0.85	0.93
Slope coefficient $\widehat{\delta}_j$		0.92	0.90	0.39	0.61
Adjusted R^2		0.61	0.74	0.16	0.20
Rejection rate (5%)		1.00	1.00	0.81	0.95

Source: Own computations with 100 Monte Carlo repetitions and $1,000$ white-noise surrogates to facilitate replication. The table reports six metrics across five filters. The first three rows report the standard deviation of the estimated cycle, its correlation with the true cycle c_t, and the corresponding RMSE. The final three rows report the slope coefficient, adjusted R^2, and rejection rate of $H_0 : \delta_j = 0$ at the 5% significance level in regression (34) for filter j.

BK-filtered cycle. Hodrick concludes that these results show that the H filter has very little to no additional explanatory power in the presence of the HP and BK filter. This is not the case for the CWT-informed filter as the slope coefficients and the adjusted R^2 (together with the rejection rates) are large. This implies that the cycle extracted by the CWT-informed filter retains large explanatory power in the presence of the cycles produced by filter j.

Panel 2 shows the results for the second DGP (conditional mean as a random walk in the change of trend and an ARIMA(2,0,0) cycle). The correlation is highest and the RMSE the lowest for the CWT-informed filtered cycle. The results for the iCWT are comparable. The H-, HP- and BK-filtered cycles show significantly less correlation and a higher RMSE with the true cycle. This is also reflected by very high slope coefficients.

Panel 3 presents the results for the third DGP (changing unconditional mean change in trend and an ARIMA(2,0,0) cycle). The correlation and RMSE of the CWT-informed filtered cycle is, again, comparable to the HP- and BK-filtered cycle, but outperforms the iCWT and H-filtered cycle. The slope coefficients indicate that also here the CWT-informed filtered cycle includes important information in the presence of the other filters. Panel 4 shows the results for the fourth DGP (slowly changing unconditional mean change in trend and an ARIMA(2,0,0) cycle). A similar picture emerges as in panel 3.

The CWT-informed filter generally performs comparably to the HP and BK filters in terms of correlation and RMSE, while outperforming both the H filter and the iCWT-based reconstruction. The main exception is the second DGP, in which the conditional mean follows a random walk; here, the CWT-informed filter outperforms all competing methods. More importantly, however, there is substantial difference in informational content. Across all DGPs, the slope coefficients associated with the CWT-informed cycle are large relative to those reported in Hodrick (2020). This indicates that, even when its RMSE and correlation are similar to those of the HP or BK filter, the CWT-informed cycle retains considerable explanatory power. In other words, the CWT-informed estimate is not merely an alternative reconstruction with comparable fit; it captures statistically informative variation that is missed by the competing filters. These results also provide evidence that the method can successfully extract an informative cycle when the null model is adapted to the task at hand.

4.3 Applications

Figure 6 shows a graph that visualizes how the algorithm identifies significant passbands. The input series is the raw and not seasonally adjusted German quarterly real GDP from 1991.1 to 2025.4. In the top panel the magnitude scalogram is shown with the cone of influence marked

Figure 6: Significant passbands for German real GDP

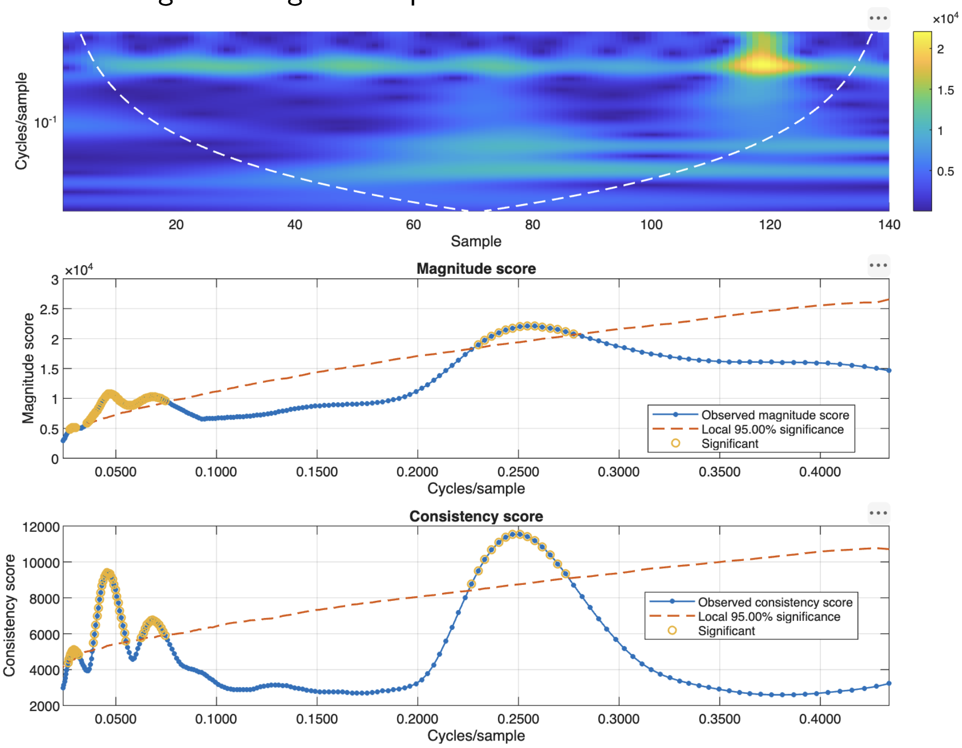

Source: Own computations. This graph shows the CWT scalogram with a normalized frequency of cycles per sample on the vertical axis and time on the horizontal axis. The heat map indicates strong local resemblance in yellow with the time series and the scaled and shifted wavelet. The cone of influence is marked by the white dashed line. Panel 2 shows the magnitude score in blue and the 95th percentile of what background noise would have generated in red. Yellow circles mark the significant passbands. Panel 3 shows the same for the consistency score.

by the white dashed line. The second panel shows the magnitude score and the third panel the consistency score in blue connected dots. The red dashed line marks the 95th percentile of what the white noise background would have generated for each frequency. The yellow circles indicate the significant passbands. Because both scores are statistically significant over the same frequency range, only these passbands are retained for cycle extraction, rather than imposing a one-size-fits-all band a priori. The scalogram suggests three periodic cycles and one seasonal cycle at 0.25 cycles per quarter or 1 cycle per year in German real GDP. This yields a first application of the CWT-informed filter: raw, non-seasonally adjusted data can be used directly, without first imposing transformations to induce stationarity or relying on seasonal adjustments produced by statistical agencies. A passband centered on the annual frequency isolates the seasonal component; conversely, excluding that frequency range yields a seasonally adjusted series directly from the observed data.

A second application emerges from the three distinct peaks at intermediate frequencies in Figure 6, which appear to characterize the cyclical component of German real GDP. Table 4 reports the corresponding extracted passbands, ordered from the lowest to the highest p-value that remains statistically significant and expressed in years per cycle. For Germany, the

Table 4: German, US and UK real GDP

Germany real GDP			US real GDP			UK real GDP		
Passband		Center periodicity	Passband		Center periodicity	Passband		Center periodicity
4.54	6.42	5.40	7.21	10.49	8.70	7.42	9.35	8.33
0.91	1.09	1.00	5.17	6.90	5.97	0.88	1.14	1.00
3.35	3.99	3.66	0.94	1.07	1.00	4.67	6.15	5.36
8.21	9.35	8.76	3.40	3.93	3.66	3.16	4.10	3.60
						2.66	2.86	2.76

Source: Own computations. This table shows the statistically significant passbands and center periodicities as a geometric average in years per cycle ranked by statistical significance (the first passband with the lowest p-value) for different countries. The input series are German and US quarterly real GDP, not seasonally adjusted, from 1991.1 to 2025.4 and UK quarterly real GDP, not seasonally adjusted, from 1995.1 to 2025.4. For the German input series, the null model is white noise with a standard deviation of the detrended input series and for the US and UK input series, the null model is white noise with a standard deviation of the residuals of a general ARMA model selected by BIC.

passband with the lowest p-value spans approximately 4.5 to 6.4 years per cycle and is centered at 5.4 years. Two additional significant cycles are centered at roughly 3.7 and 8.8 years, with the longest significant periodicity extending to about 9.4 years. A standard BK filter would exclude this lowest-frequency component. The limitation of course does not arise from the BK filter per se, but from imposing a fixed passband without first establishing which periodicities are statistically informative. Strikingly, the United States and the United Kingdom exhibit very similar, in some cases nearly identical, center periodicities. In all three economies, the annual seasonal cycle is present, alongside cycles of approximately $5-6$ and 8.5 years. Moreover, each country displays a significant cycle centered at about 3.6 years, suggesting a degree of synchronization in cyclical dynamics across these economies.

A third application concerns regression analysis. Suppose a researcher is interested in the co-movement between unemployment and inflation, but allows this relationship to vary across periodicities. The significant passbands reported in Table 5 indicate that the two series share common frequency ranges. In particular, the upper bound of the first significant passband is identical for unemployment and inflation, suggesting that part of their co-movement may be concentrated within a similar passband. The last three columns show that the relationship

Table 5: Euro Area Unemployment and Inflation

Euro Area Unemployment			Euro Area Inflation			$\pi_t = \alpha + \beta u_t + \epsilon_t$		
Passband		Center periodicity	Passband		Center periodicity	$\hat{\beta}$	SE	adj. R^2
1.96	8.08	3.98	1.15	8.08	3.05	-0.91^{***}	(0.07)	0.32
0.76	1.27	0.98	0.59	0.85	0.71	-0.01	(0.03)	-0.00
1.43	1.49	1.46						

Source: Own computations. This table shows the statistically significant passbands and center periodicities as a geometric average in years per cycle ranked by statistical significance (the first passband with the lowest p-value) for the Euro Area unemployment rate and HICP year-on-year growth rate from 1993.1 to 2019.12. For both input series, the null model is white noise with a standard deviation of the residuals of a general ARMA model selected by BIC. In the third super column, slope coefficients, standard errors and adjusted R^2 of regressing inflation on unemployment are shown.

can differ substantially depending on the periodicity. According to this purely illustrative regression, there seems to be no relationship between cycles of very high frequency but a strong negative relationship in middle to low frequencies. The CWT-informed filter has the potential to uncover more such relationships as a function of frequency.

5 Conclusion

Many questions in empirical macroeconomics rely on extracting a stationary, business-cycle-type cyclical component. Established filtering methods are commonly used to extract a cyclical component that repeats within a range of 6 to 32 quarters per cycle. This paper argues that instead of using a one-size-fits-all range of periodicities, only statistically significant periodicities should be used. Especially, whenever the estimated cycle is used in regression models, since a signal-extraction error likely yields inconsistent OLS estimates. This paper proposes to use the CWT scalogram as an exploratory tool for analyzing non-stationary data, since it avoids the need for transformations that could otherwise distort the periodogram. The proposed method is based on an algorithm that identifies statistically significant passbands against a null model that can be adapted according to the researcher's needs. Significant passbands are then supplied to a bandpass filter. Here, a Butterworth bandpass filter is recommended due to the ability to sharpen the gain function to improve performance on narrow passbands. The extracted cycle is a denoised, stationary component with statistically significant frequency content. The CWT-informed filter outperforms commonly used filters for periodic cycles. For Hodrick (2020)'s DGPs with stochastic trends and cycles, the proposed filter extracts statistically important information about the true model cycle that is missed by the cycles implied by the BK, HP, and H filter. In terms of RMSE and correlation the CWT-informed filter achieves similar results for stochastic trends and cycles. When the trend includes a random walk, the CWT-informed filter is able to outperform commonly used filters. The applications demonstrate that the CWT-informed filter can be used to extract statistically significant cycles, seasonally adjust data, and uncover frequency-dependent relationships between variables.

ABSTRACT: Many empirical macroeconomic questions rely on filtering non-stationary data to extract a stationary, business-cycle-like component. Common practice is to apply a linear filter with a standard passband, although economically relevant cycles may lie outside this passband. The resulting signal-extraction error likely affects subsequent regression estimates. This paper proposes a Continuous Wavelet Transform-informed filter that uses the CWT scalogram to identify statistically significant passbands in raw, non-stationary data relative to a researcher-specified null model. These passbands are supplied to a flexible Butterworth bandpass filter to extract a denoised, stationary cycle. Simulations show that the method improves signal extraction for periodic cycles and recovers statistically informative variation in stochastic-cycle settings missed by commonly used BK, HP, and Hamilton filters, while performing comparably or better in terms of correlation and RMSE. Applications demonstrate its use for cycle extraction, seasonal adjustment, and frequency-dependent regression analysis.

References

Aguiar-Conraria, L., & Soares, M. J. (2014). The continuous wavelet transform: Moving beyond uni-and bivariate analysis. *Journal of economic surveys*, *28*(2), 344–375.

Baxter, M., & King, R. G. (1999). Measuring business cycles: approximate band-pass filters for economic time series. *Review of economics and statistics*, *81*(4), 575–593.

Burns, A. F., & Mitchell, W. C. (1946). *Measuring business cycles* (No. burn46-1). National bureau of economic research.

Butterworth, S. (1930). On the theory of filter amplifiers. *Wireless Engineer*, *7*(6), 536–541.

Crowley, P. M. (2007). A guide to wavelets for economists. *Journal of Economic Surveys*, *21*(2), 207–267.

Daubechies, I., & Paul, T. (1988). Time-frequency localisation operators-a geometric phase space approach: Ii. the use of dilations. *Inverse problems*, *4*(3), 661–680.

Fuller, W. A. (1995). *Introduction to statistical time series*. John Wiley & Sons.

Gómez, V. (1999). Three equivalent methods for filtering finite nonstationary time series. *Journal of Business & Economic Statistics*, *17*(1), 109–116.

Gómez, V. (2001). The use of butterworth filters for trend and cycle estimation in economic time series. *Journal of Business & Economic Statistics*, *19*(3), 365–373.

Grossmann, A., & Morlet, J. (1984). Decomposition of hardy functions into square integrable wavelets of constant shape. *SIAM journal on mathematical analysis*, *15*(4), 723–736.

Gustafsson, F. (2002). Determining the initial states in forward-backward filtering. *IEEE Transactions on signal processing*, *44*(4), 988–992.

Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *Review of Economics and Statistics*, *100*(5), 831–843.

Hodrick, R. J. (2020). *An exploration of trend-cycle decomposition methodologies in simulated data* (Working Paper No. 26750). National Bureau of Economic Research.

Hodrick, R. J., & Prescott, E. C. (1997). Postwar us business cycles: an empirical investigation. *Journal of Money, credit, and Banking*, 1–16.

Lo Cascio, I. (2007). Wavelet analysis and denoising: New tools for economists. *Queen Mary University of London, Working paper No. 200*.

Percival, D. B., & Walden, A. T. (2000). *Wavelet methods for time series analysis* (Vol. 4). Cambridge university press.

Ramsey, J. B. (2002). Wavelets in economics and finance: Past and future. *NYU Working Paper No. S-MF-02-02*.

Schüler, Y. S. (2024). Filtering economic time series: On the cyclical properties of hamilton's regression filter and the hodrick-prescott filter. *Review of Economic Dynamics*, *54*, 101237.

Torrence, C., & Compo, G. P. (1998). A practical guide to wavelet analysis. *Bulletin of the American Meteorological society*, *79*(1), 61–78.

Uribe, M., & Schmitt-Grohé, S. (2017). *Open economy macroeconomics*. Princeton University Press.

Watson, M. W. (1986). Univariate detrending methods with stochastic trends. *Journal of monetary economics*, *18*(1), 49–75.

Wooldridge, J. M. (2002). *Econometric analysis of cross section and panel data* (Vol. 108) (No. 2). MIT press,.

Yogo, M. (2008). Measuring business cycles: A wavelet analysis of economic time series. *Economics Letters*, *100*(2), 208–212.

Appendix

Additional Figures

Figure 7: Common time series scalograms 10 out of 5000

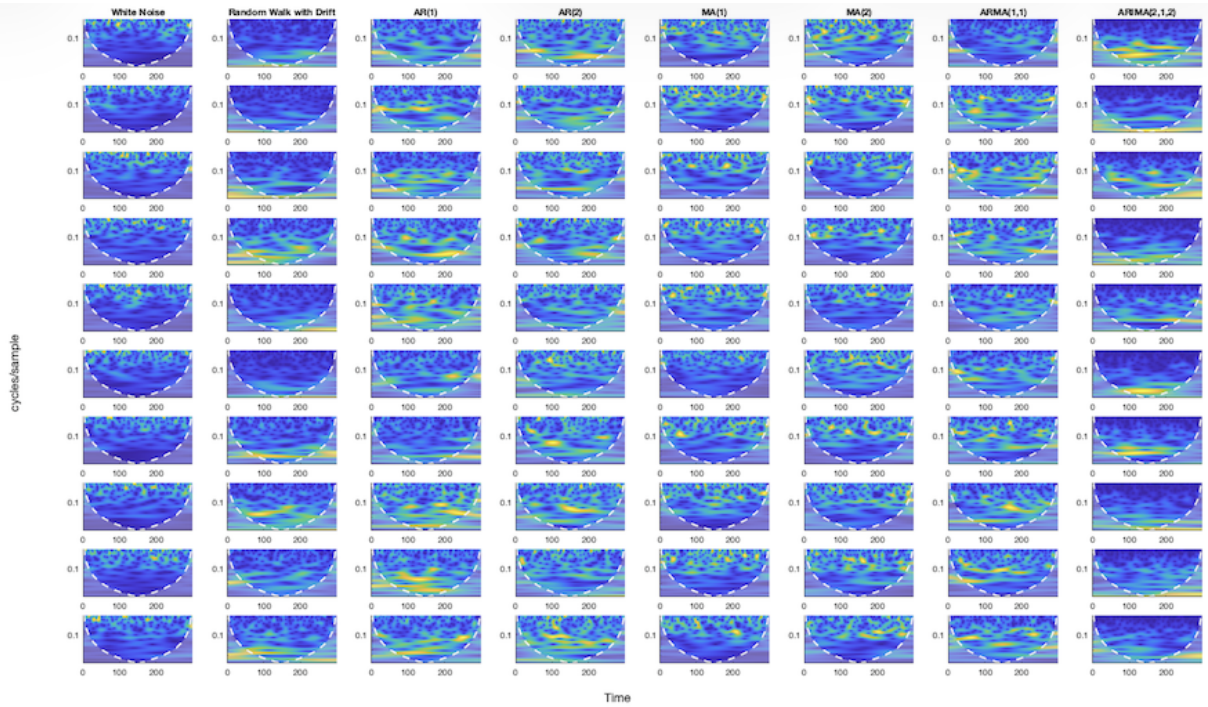

Source: Own computations. 8 processes - one per column - are simulated 5000 times and this figure shows 10 realization per process randomly selected from the simulated pool without replacement. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

Figure 8: Common time series scalograms 10 out of 4990

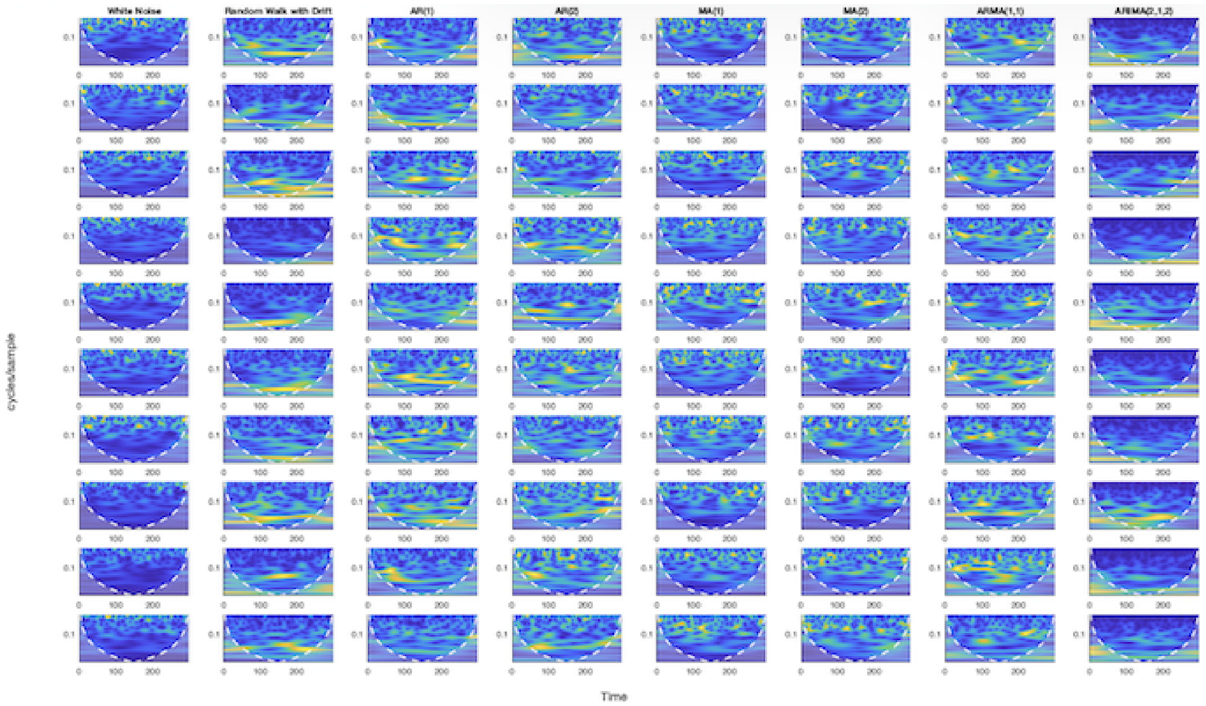

Figure 9: Common time series scalograms 10 out of 4980

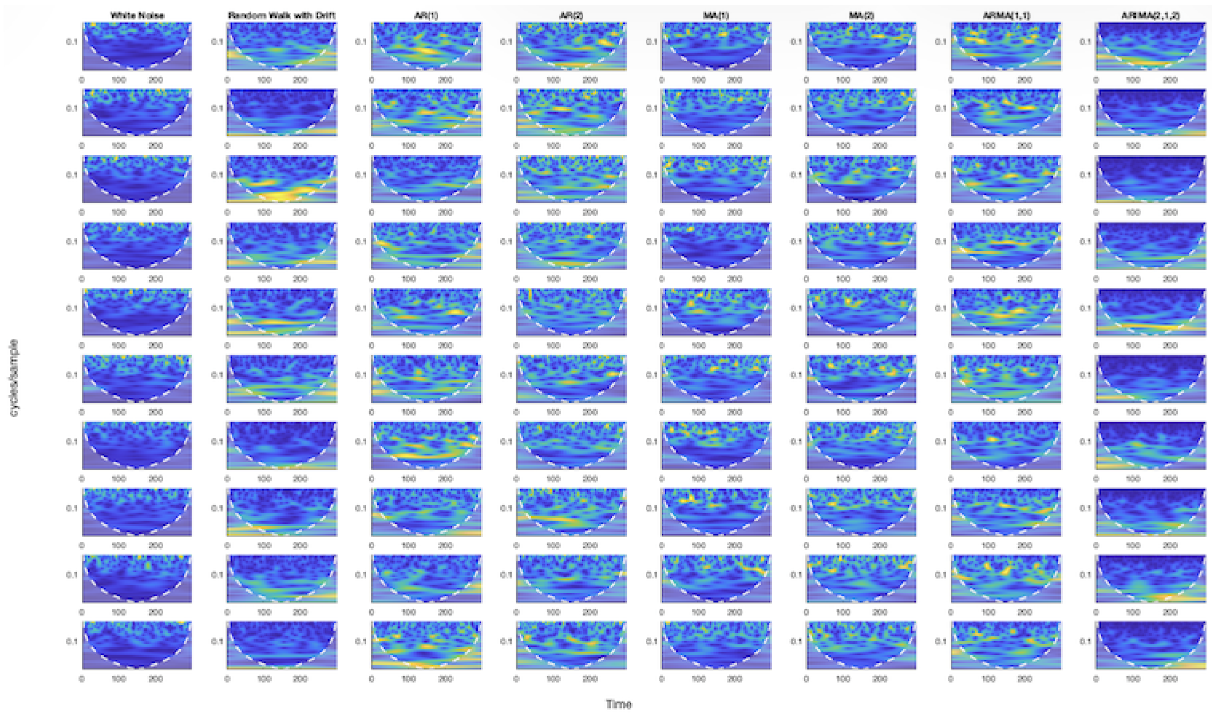

Source: Own computations. 8 processes - one per column - are simulated 5000 times and this figure shows 10 realization per process randomly selected from the simulated pool without replacement. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

Figure 10: Common time series scalograms 10 out of 4970

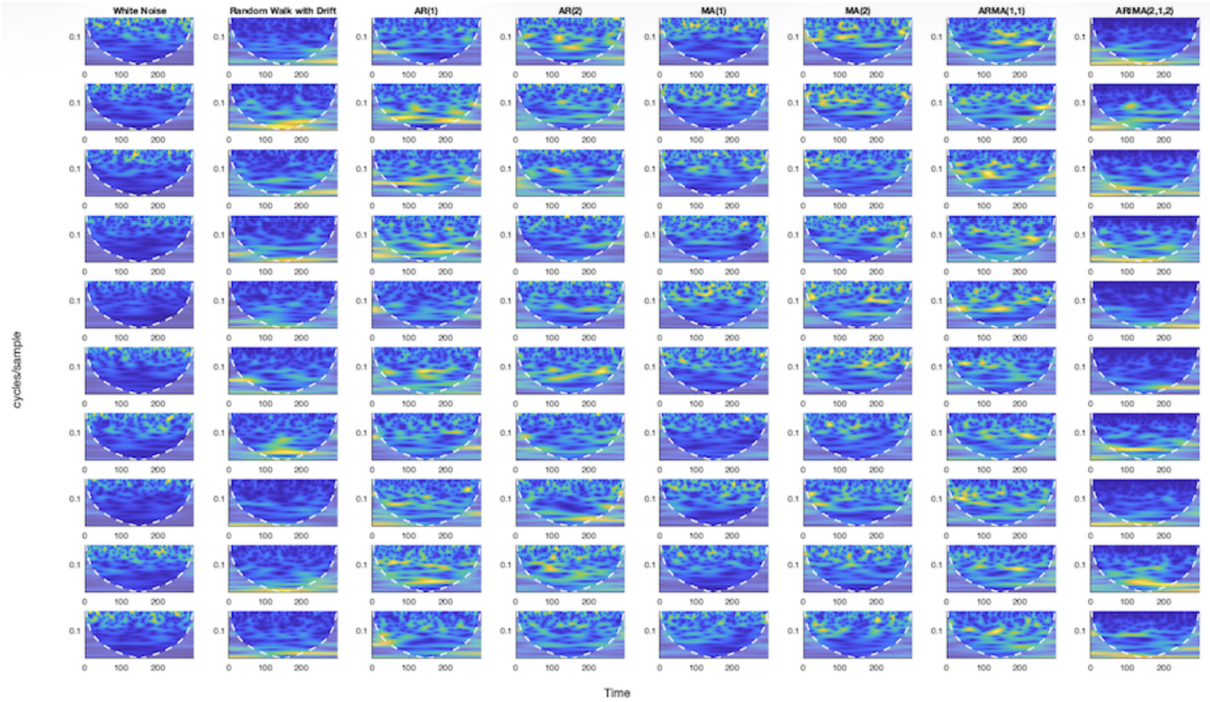

Figure 11: Common time series scalograms 10 out of 4960

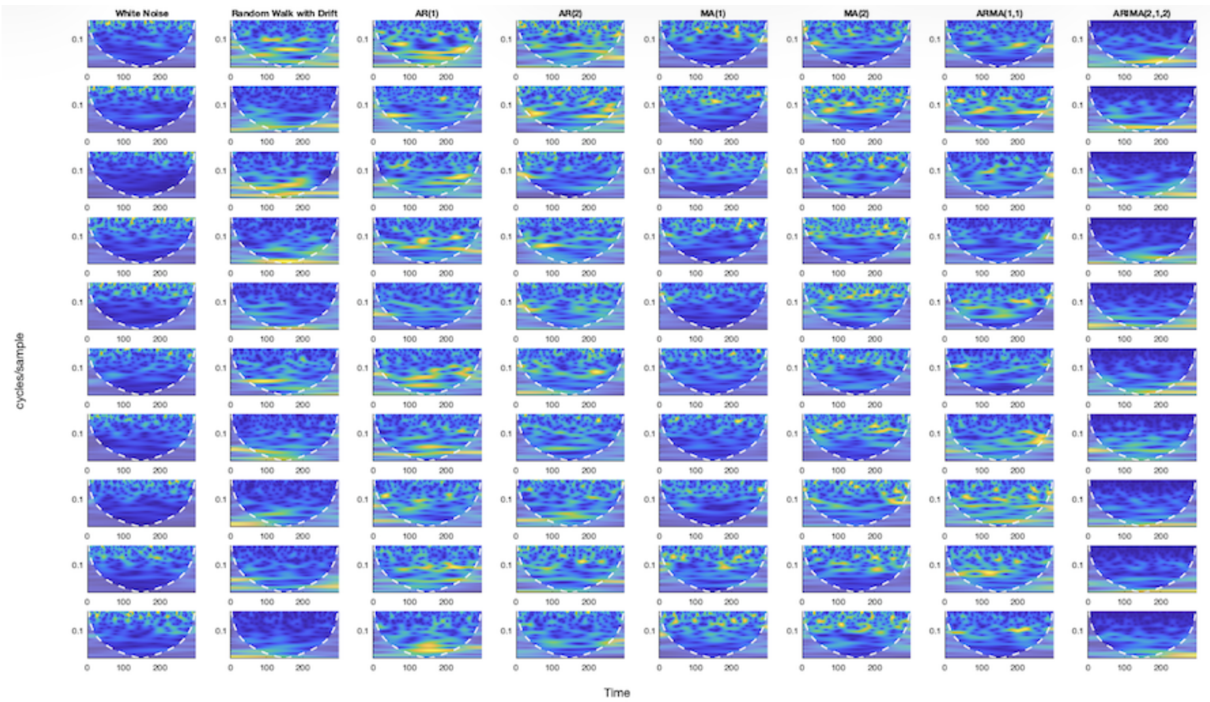

Source: Own computations. 8 processes - one per column - are simulated 5000 times and this figure shows 10 realization per process randomly selected from the simulated pool without replacement. CWT coefficients in absolute values are shown in the heat map, where bright yellow indicates strong local resemblance of the scaled and shifted mother wavelet (Morse) with the time series. The vertical axis displays a normalized frequency which is defined as cycles per sample (quarters). The horizontal axis displays time in quarters. The white dashed line and the shaded region is called cone of influence. This area is potentially affected by edge-effect artifacts and should therefore not be interpreted.

Additional Tables

Table 6: Results across the four DGPs on MOGON KI

DGP 1	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.66	2.06	2.46	1.43	1.37
Correlation	0.86	0.70	0.73	0.88	0.86
RMSE	0.91	1.51	1.69	0.85	0.94
Slope coefficient $\widehat{\delta}_j$		0.96	0.89	0.41	0.63
Adjusted R^2		0.66	0.74	0.18	0.22
Rejection rate (5%)		1.00	1.00	0.84	0.92

DGP 2	CWT-informed	iCWT	H	HP	BK
Standard deviation	2.39	2.48	3.44	1.63	1.56
Correlation	0.75	0.73	0.52	0.65	0.68
RMSE	1.92	1.98	3.15	2.13	2.09
Slope coefficient $\widehat{\delta}_j$		0.77	0.92	0.71	0.71
Adjusted R^2		0.13	0.65	0.26	0.24
Rejection rate (5%)		0.65	1.00	0.94	0.91

DGP 3	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.67	2.02	2.54	1.43	1.37
Correlation	0.85	0.71	0.70	0.88	0.85
RMSE	0.95	1.47	1.81	0.86	0.94
Slope coefficient $\widehat{\delta}_j$		0.93	0.90	0.37	0.57
Adjusted R^2		0.61	0.75	0.16	0.19
Rejection rate (5%)		1.00	1.00	0.80	0.89

DGP 4	CWT-informed	iCWT	H	HP	BK
Standard deviation	1.66	1.99	2.47	1.43	1.36
Correlation	0.86	0.72	0.73	0.88	0.86
RMSE	0.93	1.43	1.71	0.85	0.93
Slope coefficient $\widehat{\delta}_j$		0.92	0.89	0.39	0.60
Adjusted R^2		0.60	0.73	0.17	0.21
Rejection rate (5%)		1.00	1.00	0.82	0.90

Source: Computations were performed on the MOGON KI high-performance computing cluster at Johannes Gutenberg University Mainz using $5,000$ Monte Carlo repetitions and $10,000$ white-noise surrogates. The table reports six metrics across five filters. The first three rows report the standard deviation of the estimated cycle, its correlation with the true cycle c_t, and the corresponding RMSE. The final three rows report the slope coefficient, adjusted R^2, and rejection rate of $H_0 : \delta_j = 0$ at the 5% significance level in regression (34) for filter j.